Edmund G. Brown Jr.
Governor

HIGH-PERFORMANCE HIGH-TECH BUILDINGS

PIER FINAL PROJECT REPORT

Prepared For:
California Energy Commission
Public Interest Energy Research Program

Prepared By:
Lawrence Berkeley National Laboratory

March 2011
CEC-500-2009-090

Prepared By:
Lawrence Berkeley National Laboratory
William Tschudi
Berkeley, CA 94720
Commission Contract No.: 500-02-004
Commission Work Authorization No: MR-016

Prepared For:
Public Interest Energy Research (PIER)
California Energy Commission

Paul Roggensack
Contract Manager

Michael Lozano
Program Area Lead
Industrial, Agricultural, Water

Virginia Lew
Office Manager
ENERGY EFFICIENCY RESEARCH OFFICE

Laurie ten Hope
Deputy Director
ENERGY RESEARCH AND DEVELOPMENT DIVISION

Melissa Jones
Executive Director

DISCLAIMER

This report was prepared as the result of work sponsored by the California Energy Commission. It does not necessarily represent the views of the Energy Commission, its employees or the State of California. The Energy Commission, the State of California, its employees, contractors and subcontractors make no warrant, express or implied, and assume no legal liability for the information in this report; nor does any party represent that the uses of this information will not infringe upon privately owned rights. This report has not been approved or disapproved by the California Energy Commission nor has the California Energy Commission passed upon the accuracy or adequacy of the information in this report.

Disclaimer

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor Lawrence Berkeley National Laboratory, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, or the Regents of the University of California.

Ernest Orlando Lawrence Berkeley National Laboratory is an equal opportunity employer.

Acknowledgements

Special thanks are extended to the Industry Partners and project participants for their support of this project.

Principal Lawrence Berkeley National Laboratory (LBNL) Investigator:

- Bill Tschudi

Additional LBNL Investigators:

- Dale Sartor
- Tengfang Xu
- Evan Mills
- David Faulkner
- Geoffrey Bell
- Steve Greenberg
- Bruce Nordman

Subcontractors:

- Peter Rumsey, Kim Traber, and Larry Chu: Rumsey Engineers, Inc.
- Jay Madden and Michael Tatoyan:EYP Mission Critical Facilities
- My Ton and Chris Calwell: Ecos Consulting
- Arshad Mansoor and Brian Fortenbury: EPRI
- Tom Smith: Exposure Control Technologies

Industry Partners and Other Collaborators:

- 365 Main
- Rockwell International
- Abgenix
- Johnson & Johnson
- Hewlett Packard
- Cisco
- Jazz Semiconductor
- Balazs Analytical Services
- Equinix
- International Sematech
- Earthlink
- Cornell University

- State University of New York at Binghamton
- Georgia Institute of Technology
- IBM
- Pacific Gas and Electric Company
- Intel Corporation
- Sun Microsystems
- Additional firms acknowledged in DC Power Demonstration Report
- Critical Facilities Roundtable

Please cite this report as follows:

Lawrence Berkeley National Laboratory. 2011. *High-Performance High-Tech Buildings.* California Energy Commission, Public Interest Energy Research. CEC-500-2009-090.

Preface

The California Energy Commission's Public Interest Energy Research (PIER) Program supports public interest energy research and development that will help improve the quality of life in California by bringing environmentally safe, affordable, and reliable energy services and products to the marketplace.

The PIER Program conducts public interest research, development, and demonstration (RD&D) projects to benefit California.

The PIER Program strives to conduct the most promising public interest energy research by partnering with RD&D entities, including individuals, businesses, utilities, and public or private research institutions.

- PIER funding efforts are focused on the following RD&D program areas:
- Buildings End-Use Energy Efficiency
- Energy Innovations Small Grants
- Energy-Related Environmental Research
- Energy Systems Integration
- Environmentally Preferred Advanced Generation
- Industrial/Agricultural/Water End-Use Energy Efficiency
- Renewable Energy Technologies
- Transportation

High-Performance High-Tech Buildings is the final report for the Lawrence Berkeley National Laboratory, UC Research Agreement project (Contract Number 500-02-004, Work Authorization Number MR-016) conducted by Lawrence Berkeley National Laboratory and it subcontractors. The information from this project contributes to PIER's Industrial, Water, and Agriculture Program.

For more information about the PIER Program, please visit the Energy Commission's website at www.energy.ca.gov/research/ or contact the Energy Commission at 916-654-4878.

Table of Contents

Preface ... iii
Abstract ... vii
Executive Summary ... 1
1.0 Introduction ... 7
1.1. Background and Overview ... 7
2.0 Project Objectives ... 9
2.1. Cleanrooms and Laboratories ... 9
2.2. Berkeley Fume Hood Development ... 10
2.3. Data Centers ... 10
2.4. Demonstration Projects ... 11
2.5. Technology Transfer ... 11
3.0 Project Approach ... 13
3.1. Cleanrooms and Laboratories ... 13
3.2. Berkeley Fume Hood Development ... 15
3.3. Data Centers ... 16
3.4. Demonstration Projects ... 18
3.5. Technology Transfer ... 18
4.0 Project Outcomes ... 21
4.1. Cleanrooms and Laboratories ... 21
4.2. Berkeley Fume Hood Development ... 26
4.3. Data Centers ... 27
4.4. Demonstration Projects ... 31
4.5. Technology Transfer ... 31
5.0 Conclusions and Recommendations ... 33
5.1. Conclusions ... 33
5.2. Recommendations ... 35
5.3. Benefits to California ... 36
6.0 References ... 37
7.0 Glossary ... 39

Attachment I. Cleanroom/Laboratory Benchmarking

Attachment II. Development of Fan-Filter Test Procedure

Attachment III. Investigate Demand Controlled Ventilation

Attachment IV. Investigate Mini-Environments

Attachment V. Side-by-Side Fume Hood Testing

Attachment VI. Data Center Benchmarking

Attachment VII. Improve Uninterruptible Power Supplies

Attachment VIII. Investigate Computer Performance Metrics

Attachment IX. Energy-Efficient Power Supplies

Attachment X. Demonstrate Demand Controlled Ventilation

Attachment XI. Demonstrate Standardized Testing of Fan-Filter Units

Attachment XII. Demonstrate DC Powering of IT Equipment

Attachment XIII. Energy Efficient Scientific Computing

Attachment XIV. Technology Transfer Activities

List of Figures

Figure 1. Recirculation fan energy savings during setback at Facility L 21

Figure 2. Recirculation fan energy savings during setback at Facility K 22

Figure 3. Cleanroom recirculation air system efficiency 23

Figure 4. Particle counts and fan speed 25

Figure 5. Ratio of IT power to total data center power 28

Figure 6. UPS system efficiencies 29

Figure 7. Performance of power supplies 30

Abstract

This report summarizes research, demonstration, and market transformation activities directed at the high-tech buildings market. Included are elements focused on laboratory, cleanroom, and data center facilities. This work spans 2003 through 2007 and supports the Public Interest Energy Research (PIER) Program research roadmaps for these building types. The project resulted in a self-benchmark protocol for data centers, a standard method of testing and reporting performance of fan-filter units, testing of the Berkeley Fan Hood, which is a fume hood designed to remove any toxic gas from a laboratory, that showed equivalent or better containment with 50 percent less airflow, identification of a wide variation in energy use among uninterruptible power supplies, a demonstration of a 20 to 30 percent energy savings if direct current power is distributed directly into information technology equipment, and other research that will contribute to energy savings in these sectors. This report is organized with stand-alone appendices that provide more detail for specific topics. The research project benefits the high-tech industry and data centers by making them more efficient and safer.

Keywords: Cleanroom, data center, UPS, uninterruptible power supplies, energy efficient high tech buildings, laboratory, fume hood, demand controlled ventilation, particle counts, mini-environments, energy benchmarking, fan-filter performance, standby generation, power supplies, DC power, power conversion

Executive Summary

Introduction

Prior research and investigations by Lawrence Berkeley National Laboratory focused on energy intensive buildings that contain laboratories, cleanrooms, and data centers. Lawrence Berkeley National Laboratory studies and data provided by the California Energy Commission and California public utilities concluded that energy use in the industries and institutions that rely on these facilities represented a large and growing portion of the state's energy consumption. These facilities are characterized as having large, continuous electrical demand as well as significant gas demand in many cases, and far higher energy intensities (that is, energy per unit of floor area) than conventional buildings.

Preliminary investigations into energy efficiency improvement opportunities conducted within this project involved multiple sponsoring organizations[1] and industry partners, which led to a better understanding of how they presently operate, and identified research that would be needed to achieve greater efficiency in the future. The Public Interest Energy Research "energy research roadmaps" for these building types highlighted numerous areas ripe for public interest research support, where dramatic energy efficiency gains were possible. The research activities that ultimately were included in this project were prioritized by industry leaders through workshops and public utility-hosted meetings. This project was designed to address specific high-priority opportunities identified in the roadmaps for each of the building types. A variety of research, development, demonstration, and deployment activities were included.

Two project advisory committees—one focused on laboratories and cleanrooms and the other on data centers—were assembled and provided valuable project review and input. These committees consisted of industry leaders representing various stakeholders: public utilities, facility operators, design firms, manufacturers, industry associations, and other public interest organizations.

Purpose

This project was to advance energy efficiency research related to California's high-tech buildings and develop ongoing market transformation activities to communicate findings to stakeholders.

Project Objectives

The objectives for the various elements of this project were as follows:

---

1 Previous sponsoring organizations included California Institute for Energy and Environment, the California Energy Commission's Public Interest Energy Research Program, Pacific Gas and Electric, San Diego Gas & Electric, Northwest Energy Efficiency Alliance, New York State Energy Research and Development Authority, U.S. Department of Energy, and the U.S. Environmental Protection Agency.

- Benchmark laboratories, cleanrooms, and data centers to provide insight into opportunities for efficiency improvement.
- Develop a self-benchmarking protocol for data centers.
- Identify best practices for cleanrooms and data centers.
- Develop case studies.
- Develop a standard method of testing and reporting energy performance of fan-filter units commonly used in cleanrooms, and test a large sampling of units.
- Research and demonstrate a control strategy for cleanrooms that controls airflow based upon cleanliness levels in the cleanroom.
- Investigate the energy implications of the use of mini-environments (self-contained clean spaces) in cleanrooms.
- Develop test methods to validate that the Berkeley Fume Hood developed in prior Public Interest Energy Research projects achieved containment performance equal to or better than standard fume hoods.
- Research energy efficiency of commercially available uninterruptible power supply systems commonly used in data centers.
- Research energy issues associated with standby generators maintaining their readiness.
- Research energy efficiency of power supplies in information technology equipment.
- Investigate options for performance metrics applicable to information technology equipment.
- Identify demonstration projects to showcase emerging or underused technologies.
- Implement a variety of outreach outlets for technology transfer.

Project Outcomes

The results of this project include many high-value findings described within this report. Highlights include:

- As a result of the benchmarking of laboratories, cleanrooms, and data centers, it was obvious that energy performance varied tremendously, and by examining the better performing facilities many opportunities for efficiency improvement were identified. Many efficiency recommendations were relayed to the participating industry partners and presented at industry meetings. Best practices revealed through the case studies were summarized and posted on Lawrence Berkeley National Laboratory's high-tech buildings website along with case study reports and benchmark summaries. In addition, a self-benchmarking protocol for data centers was developed to assist industry in obtaining benchmark data. A similar resource for laboratory buildings is available through the U.S. Environmental Protection Agency/U.S. Department of Energy Laboratories for the 21st Century website.

- Industry experts developed and reviewed a standard method of testing and reporting performance of fan-filter units commonly used in cleanrooms and mini-environments. This procedure enables end users to compare energy use in a standard way when making purchasing decisions. Use of the procedure was demonstrated in the testing of 17 units. The results illustrated the wide range of performance that exists and they can be used to establish baselines for use in utility incentive programs. Extensive targeted print and electronic publications, as well as in-person presentations, have made industry aware of this resource.
- A technology that controls cleanroom airflow based upon real-time monitoring of particles (contamination levels) in the room was developed and demonstrated in a pilot study and in two operating industrial cleanrooms. This technology enables cleanroom operators to reduce air flow based upon actual cleanliness conditions, resulting in large energy savings. This task successfully demonstrated the viability and cost-effectiveness of this technology.
- Energy implications for the use of mini-environments were studied. Two case studies provided insight into the efficiency of the mini-environments themselves, as well as the energy impact of their use within cleanrooms. Findings suggest opportunities for large energy savings in both areas.
- The high-performance Berkeley Fume Hood, which exhausts 50 percent less conditioned air than a conventional fume hood, was evaluated in side-by-side testing with a conventional fume hood. To accomplish this, industry consensus was obtained to define new static and dynamic tracer gas tests to verify containment capability. The testing successfully demonstrated that the Berkeley hood containment performance equaled or exceeded that of a conventional hood while providing large energy savings. The results of this testing were successfully used as technical justification in California Occupational Safety and Health Administration (Cal/OSHA) variance hearings wherein Cal/OSHA accepted the use of the hoods with minimal supplemental requirements. Unfortunately, Cal/OSHA would not revise its mandated face-velocity requirement, so this variance procedure would have to be followed for any installation of the hood. As a result PIER management decided to terminate the planned industrial demonstrations and further work on this technology. The rationale for terminating further work was that unless the California standard was changed, there would not be widespread adoption in California, and there are alternative technologies available to achieve savings.
- Uninterruptible power supplies, used extensively in data centers, cleanrooms, and hospitals, account for a constant energy loss due to conversions between alternating current and direct current. A test protocol was developed and used to test the major commercially available uninterruptible power supply systems. Test results showed a wide variation in energy loss among these competing systems, suggesting that use of the more efficient systems in data centers could save up to 20 percent of the electricity powering information technology equipment plus approximately equal savings due to the reduction in heating, ventilating, and air-conditioning energy use for not having to remove the heat generated in the power conversions. In addition, since uninterruptible

power supply systems are often deployed in redundant configurations, the losses can be magnified. Various redundancy schemes were studied to determine the relative efficiency of commonly used configurations. This information can be used to establish utility incentives to encourage use of more efficient uninterruptible power supply units and their deployment in redundant system configurations. Finally, draft labeling criteria were developed that could be used in California or ENERGY STAR® programs to provide information for those that specify and purchase these systems.

- Energy required to keep standby generation systems ready to function was measured. This data reveal that standby generation consumes more energy than it will ever generate and that, statewide, these losses could be significant.
- Various industry experts were consulted to begin developing consensus on metrics that could be used to evaluate computing performance and associated energy use. These types of metrics will be useful in evaluating energy use together with computing capability. A U.S. Environmental Protection Agency-sponsored conference in January 2006, attended by more than 250 data center professionals, recommended that the industry focus on developing a consensus for such metrics. Subsequently, a performance metric procedure that considers computational performance and energy performance was developed by an industry group for one of the common computational benchmarks.
- Electrical power conversions within information technology equipment accounts for a large percentage of its power consumption. A calculator was developed and posted on Lawrence Berkeley National Laboratory's website to quantify energy savings that could result if more efficient components were used. Researchers also measured the energy performance of the majority of the most popular class of AC-DC power supplies used in low-end servers. Similar to uninterruptible power supply systems, a wide variation in performance was observed, suggesting that encouraging use of highly efficient power supplies could immediately save an additional 15 percent or more at the equipment level coupled with a similar saving for facility energy use. This work helped to raise awareness of the inefficiencies in power supplies, and over the life of this contract, significant improvement in power supply efficiency became evident.
- Potential demonstration projects were evaluated and proposed. Five of these were selected to highlight the savings opportunity for industry:
    - Direct DC powering of a rack of computers
    - DC powering at the facility level
    - Alternative cooling (subsequently cancelled)
    - Fan-filter testing
    - Demand controlled filtration
    - Data center design process
- Numerous outlets were employed to transfer technology and findings from this project to industry stakeholders. Highlights of this activity include:

- More than 100 meetings with industry and other stakeholders.
- Extensive Web-based publishing of results and data (http://hightech.lbl.gov).
- Video documentary of the DC Powering demonstration.
- Articles in 33 trade publications, 10 professional journals, and best-practice guides produced by project researchers.
- Dozens of interviews and references to the project by reporters in the trade media plus coverage in 14 mainstream "popular" media outlets.
- Periodic newsletters titled *News from the Hood* and *High-Tech News* (each distributed to about 400 people).
- Active participation in 21 industry technical committees or standard-setting groups.
- Provided numerous professional training sessions and curriculum development through the California utilities, community colleges, and other entities.
- Collaboration with Industry Associations:
    - American Society of Heating, Refrigerating and Air-Conditioning Engineers
    - 7x24 Exchange
    - Sematech
    - Uptime Institute
    - Silicon Valley Leadership Group
    - Critical Facilities Roundtable
    - Laboratories for the 21st Century (Labs 21)
    - Green Grid.

Conclusions

High-tech buildings are an integral part of California's industries and institutions. Since energy intensity in these facilities is large, their energy use represents a significant operating cost for a large number of California businesses and institutions. Research has shown that there is no single strategy that will improve energy performance; rather, there are opportunities at many levels to improve efficiency. The specific research areas of this project each illustrate that the focus areas identified in the Public Interest Energy Research roadmaps do have the potential to curb, and eventually reduce, the energy growth associated with high-tech buildings.

Recommendations

As a result of this project and other industry actions, there is an increased awareness of the potential for improvement in high-tech buildings' energy performance, yet there is reluctance to change or embrace new technology for a number of reasons. Fear of upsetting production, regulatory barriers, reliability or safety concerns, or misinformation often stand in the way of major breakthroughs for these buildings. Public Interest Energy Research projects have had a dramatic impact by:

- Illustrating ranges of possible improvement.
- Documenting best practices.
- Providing research that industry would not accomplish on its own.
- Demonstrating new or underused technologies.
- Training and dissemination of information.

The authors recommend that research roadmaps continue to be followed and updated periodically as technologies change. Demonstration projects should be encouraged and supported under the Public Interest Energy Research Program. Specific recommendations for each of the topic areas in this project are provided in the report.

Benefits to California

High-tech industries are extremely important to the California economy. Many high-tech firms are headquartered in California and represent most major sectors of the economy, such as biotech, semiconductor, aerospace, and health care. Energy use and intensity is large in these industries and continues to grow as these companies flourish. This project has provided insight into the energy efficiency opportunity in cleanrooms, laboratories, and data center facilities. The various technologies explored show that there is potential for large energy savings using solutions available today. Energy savings of 20 to 40 percent have been shown to be possible. This project has also helped to raise the awareness of industry and has garnered the interest of other public interest programs. California is now seen as leading efficiency efforts in the high-tech sectors.

1.0 Introduction

1.1. Background and Overview

California has long been known for its innovative and robust high-tech industries, as well as a research focus that spawned growth in the industries and institutions that rely on laboratory and cleanroom facilities. In addition, California's businesses and institutions operate numerous computing facilities—data centers—with seemingly endless growth in size, complexity, and energy intensity. The Public Interest Energy Research (PIER) Industrial program, recognizing the importance of California's high-tech industries and institutions and the increasing reliance on computing systems, previously developed "research roadmaps," which were meant to guide and prioritize public interest research related to high-tech buildings. These were developed by Lawrence Berkeley National Laboratory (LBNL) utilizing industry input, and they provide a comprehensive summary of research needs relating to laboratories and cleanrooms (Tschudi et al. 2002) and data centers (Tschudi et al. 2003a). Prior LBNL research confirmed that the high-tech building "market" was characterized as large and growing in terms of total size of market and in energy intensity. Further research and case studies confirmed that energy efficiency measures, if developed, could readily be applied resulting in 40% or more in energy savings. Since high-tech buildings typically operate continuously with high energy intensity, energy reduction often translates into very impressive life cycle cost savings while lowering utility peak demand.

This project was designed to advance the research agenda for high-tech buildings by focusing on high-priority research on selected topics (as determined by industry feedback) in each of the three facility types. To help guide the current research activities, two project advisory committees (PACs) were formed, one focused on cleanrooms and laboratories and the other on data centers. These PACs were made up of many leading firms, public interest organizations, and industry experts. In addition, a comprehensive tech-transfer and market transformation activity was included to enable the results to reach the targeted industries and other stakeholders.

This report is organized so that each task is addressed in summary fashion in the body of the report. Detailed stand-alone reports and supplemental materials were prepared and are attached as appendices to this overall project report. These collectively constitute the deliverables for the project.

2.0 Project Objectives

This project included objectives in the following areas:

2.1. Cleanrooms and Laboratories

The objective of this task was to improve energy efficiency of cleanroom and laboratory facilities by executing selected high-priority activities from the PIER high-tech buildings roadmap as described below:

Benchmark laboratory and cleanroom facilities and document best practices

The objective of this task was to solicit industry partners who operate laboratory or cleanroom facilities to participate in in-depth case studies of their facilities. This involved measurement of the energy performance of building systems in order to develop a more robust set of benchmark data building upon previous benchmark results. This also was to include, for the first time, performance information on standby generation—a typical feature of most high-tech buildings. Site reports were to be developed for each facility and were to include benchmark data along with observations of areas for potential efficiency improvement. Using this data combined with prior benchmark data, the technologies or techniques that contributed to better performance were to be identified in "best practice" summaries. This information along with the benchmark data was to be used to begin market transformation through interaction with industry professionals.

In particular, cleanroom energy-intensive air systems (recirculation, make-up, and/or exhaust) were targeted. The findings were to be reviewed against current industry-accepted practices such as use of air-change ranges recommended by the Institute of Environmental Sciences and Technology (IEST). The objective was to provide insight into whether the existing practices established by rule-of-thumb were appropriate and whether lower airflow—with resulting energy savings—could be used to achieve acceptable production.

This task was notable in that it entailed a far more detailed level of benchmarking than is common, which, in turn, enabled more precise opportunity identification.

- **Advance energy-efficient filtration**

  This task's objective was to develop a standard method of testing and reporting performance of fan-filter units commonly used in cleanrooms and other enclosures by collaborating with the Industrial Technology Research Institute (ITRI), of Taiwan, and the Air Movement and Control Association (AMCA, the U.S. fan manufacturer's association) or other industry organizations.

- **Investigate improvements in efficiency of heating, ventilating, and airconditioning in high-tech buildings**

  The objective of this task was to research promising technologies that could contribute to efficiency savings in cleanroom heating, ventilating, and air-conditioning (HVAC) systems. Researchers conducted investigations into energy efficiency opportunities using demand-controlled filtration (i.e., use of real-time particle monitoring to control airflow in

cleanrooms) and in the use of mini-environments commonly used in many high-tech industries.

2.2. Berkeley Fume Hood Development

The objective of this task was to continue the development of the Berkeley Fume Hood by overcoming institutional barriers and demonstrating energy savings by using the fume hood in industrial settings. This objective was to be accomplished through the following tasks:

- **Overcoming institutional barriers**

  This task's objective was to devise and provide technical justification acceptable to the California Division of Occupational Safety and Health (CAL/OSHA) to support granting of variances authorizing the use of the Berkeley Hood in industrial settings.

- **Conduct side-by-side tests with conventional hoods**

  In this task, the objective was to demonstrate that the Berkeley Hood's safety and energy performance was equivalent to or exceeded that of a conventional fume hood when subject to identical static and dynamic tests.

- **Perform industrial field demonstrations**

  The primary objective of the industrial demonstrations task was to demonstrate that the Berkeley Hood performed safely in real-life industrial settings while achieving significant energy savings.

2.3. Data Centers

Similar to the cleanrooms and laboratories tasks, the objective of the data center tasks was to perform high priority research as identified in the PIER "High-Performance Data Centers energy research roadmap" to develop solutions that could lead to 40% or more energy efficiency improvement. The activities selected are detailed below:

- **Investigate data center efficiency opportunities and perform energy benchmarking**

  The objective of this task was to solicit industry partners to participate in case studies and to measure energy performance of data center and associated standby generation systems. Detailed energy benchmark data for data centers was to be collected in order to develop a more robust set of benchmark data when combined with prior benchmark results. For standby generation, the goal was to obtain initial assessments. Site reports were to be developed for each facility and were to include benchmark data along with observations of areas for potential efficiency improvement. A further objective was to review this data along with prior benchmark data to determine the technologies or techniques that contributed to better performance. These “best practices” were to be summarized for use in market transformation activities. In addition, a self-benchmarking protocol was to be developed with the objective of encouraging data center professionals to obtain their own benchmarks. Ultimately the objective of these activities was to identify energy efficiency measures, encourage industry action, and reduce energy consumption in data centers by 10%–20%.

- **Improve uninterruptible power supplies**

  This task's objective was to research and test energy performance of existing uninterruptible power supply (UPS) systems commonly used in data center environments with a possible longer-term objective of establishing labeling criteria. Possible labeling options were to be explored. In addition, various redundancy strategies were to be investigated with the objective of determining their relative energy efficiencies.

- **Performance metrics**

  The objective of this task was to investigate metrics that could be used to evaluate computing performance along with energy consumption. This task replaced a task to monitor and evaluate new or emerging rack systems used in data centers.

- **Energy-efficient power supplies**

  The objective of this task was to investigate the energy performance of selected power supplies currently used in servers and investigate the energy efficiency opportunity to improving their performance.

2.4. Demonstration Projects

The objective of this task was to first identify demonstration projects for consideration. Once the candidate demonstration projects were selected by PIER's Industrial Program, the objective was to organize and perform the demonstrations in conjunction with industry partners. New or underutilized technologies and practices were selected to be demonstrated in commercial settings. The demonstrations were authorized in a second phase of the project. Demonstrations were expected to highlight promising products, technologies, or approaches; or better application of existing technologies. The objective of the demonstrations was to encourage wider adoption of these technologies to a broad audience of high-tech building design and operations professionals. As a result, a substantial technology transfer component was included to disseminate the benefits of the demonstrations to larger targeted audiences.

2.5. Technology Transfer

The objective of the technology transfer activities was to convey the knowledge gained, the experimental results, and the lessons learned to key decision makers in the various industries and institutions involved with high-tech buildings. Targeted audiences included facility designers and operators, manufacturers and suppliers, and policy makers and public interest organizations.

3.0 Project Approach

3.1. Cleanrooms and Laboratories

Several project tasks focused on energy use in cleanroom and laboratory facilities. These tasks included the following activities:

- **Benchmark laboratory and cleanroom facilities and document best practices**

  The approach for the energy benchmarking portion of this task involved soliciting benchmarking sites through contacts within individual companies, industry associations, California public utilities, and project advisors. Initial contacts focused on Southern California since previous benchmarking efforts were focused on Northern California facilities. At each site, typically, one or more meetings were held to explain the benchmarking process and its benefits to the host site. Once an industry partner agreed to participate, a meeting was held along with LBNL's subcontractor (Rumsey Engineers) to describe the system's evaluation and monitoring that would take place and request a limited amount of design information that would be used to help determine monitoring points. Utility customer representatives for the sites were invited to participate in both the kick-off meeting and the final meeting where results were presented and discussed. The subcontractor then arranged to be on site for several days to collect measurements and obtain existing information from building management systems. Since cleanroom facilities' energy use is typically not dominated by weather conditions, a relatively short monitoring period was planned, usually three to four days. Then, over a period of several weeks, the team analyzed the data and developed a site-specific report with measured results and energy efficiency recommendations, which were based upon observations while on site. In a final meeting with the host site, the report was reviewed and explained. The host site was given the opportunity to correct any factual errors in the report, and then the anonymous reports were posted on LBNL's website: http://hightech.lbl.gov/cr-benchmarking-results.html.

  The benchmarking focused mainly on cleanroom air systems and standby generation systems; however, other data were obtained if readily available. All benchmark results were entered into the Access™ database developed through a prior PIER project. Summary graphs were then prepared in order to focus on the energy performance for selected systems or components. Results were compared to typical industry practice for various parameters. By reviewing all of the benchmark data and recommendations from the current case studies combined with information contained in case studies obtained from various sources, many areas where better energy performance was achieved became evident. From these, a series of best-practice summaries were prepared.

  In all, 11 separate cleanrooms at four sites were benchmarked. Two of the sites were in Southern California, and two were in the San Francisco Bay Area. In addition, a laboratory facility in Southern California was benchmarked.

- **Advance energy-efficient filtration: Standard testing for fan-filter units**

  This task involved development of a standard method of testing and reporting performance of fan-filter units (FFUs), which are very common ventilation system components in cleanrooms. In current practice, there is a lack of consistent measurement and reporting of FFU energy performance, making it difficult for those specifying them to make informed choices among manufacturer's products. Initially, there was interest within the fan manufacturers association, the Air Movement and Control Association (AMCA), to work with LBNL to develop a standard test method. However, for their continued involvement, a large number of FFU manufacturers (members) were needed, and they would need to agree to pay for testing. Since there were not enough manufacturers that were AMCA members, AMCA consequently decided to reduce their involvement in the procedure development yet maintained high interest in collaborating in its development. With LBNL's urging, IEST then agreed to take the lead as the "standards body" by developing an IEST "eecommended practice" (RP), which would include the test procedure. The IEST RP also plans to address other non-energy considerations such as noise and vibration.

  During the project, LBNL developed several drafts of the FFU testing procedure and collaborated with numerous industry professionals who provided input to the draft procedures. LBNL also collaborated with the Industrial Technology Research Institute (ITRI) in Taiwan. In this collaboration, ITRI sponsored a trial application of the test procedure in its AMCA-certified test facility. LBNL participated in the test and gained additional insight that was useful in the development of the test procedure. For this test, an FFU was donated by a manufacturer, and its energy performance was obtained through use of the draft test procedure.

  Discussions with Pacific Gas and Electric (PG&E) also led to interest in performing a trial application of the fan-filter test procedure at the PG&E test facility in San Ramon, California. Consequently, PG&E sponsored a test at this facility where performance of the same unit that was tested in Taiwan was again tested. This test arrangement used a somewhat different configuration. PG&E, in collaboration with LBNL, performed the testing and obtained a second set of test data.

- **Investigate improvements in efficiency of heating, ventilating, and air conditioning in high-tech buildings: Demand controlled filtration and use of mini-environments**

  Demand Controlled Filtration – The approach for the demand controlled filtration task was to first perform a pilot study in an operating LBNL cleanroom. In this study, the technology of sensing particles in the room and directly controlling cleanroom airflow based upon particle concentrations was attempted using various particle counters. The study focused on determining variations in particle size and quantities that could be used for controlling the airflow and achieving desired cleanliness levels. Based upon the pilot study, a demonstration of the technology in an operating industrial cleanroom was proposed and implemented in the second phase.

Mini-environments – The approach to investigating energy efficiency issues when mini-environments are utilized was to locate and work with an industry partner who manufactures mini-environments. The manufacturer provided access to a mini-environment for assessment of its performance. Airflow and electrical fan energy were measured throughout its full operating range. Although beyond the original scope of this task, an additional study or "macro" study was able to be accomplished. This involved working with one of the benchmarking site partners to study the aggregate effects of using many mini-environments within a larger cleanroom.

3.2. Berkeley Fume Hood Development

The approach in the continuing development of the Berkeley Fume Hood involved actions designed to overcome the institutional barrier impeding its acceptance in California, performing side-by side testing with a conventional fume hood in order to document the relative performance of each, and demonstrating containment (safety) and energy performance by using the fume hood in industrial settings.

- **Overcoming institutional barriers**

  The primary institutional barrier to acceptance of the Berkeley Hood was the CAL/OSHA requirement to maintain 100 feet per minute (ft/min.) of airflow through the face of the hood opening. The Berkeley Hood provides containment through a different approach, so a variance to the requirement was necessary since face velocities were, by design, far lower than the CAL/OSHA threshold. The approach to overcoming this barrier was to negotiate acceptance criteria with the regulatory body, CAL/OSHA, and then obtain and provide that information through variance requests. Numerous meetings and communications with CAL/OSHA attempted to define technical criteria that would provide assurance that the Berkeley Hood would provide containment that was as safe (or safer) than a conventional fume hood. Standard ASHRAE tests, along with non-standard tracer gas tests, were proposed, and although CAL/OSHA would not agree to formal acceptance criteria, the variance procedure provided the best opportunity for eventual acceptance.

- **Conduct side-by-side tests with conventional hoods**

  To support the variance process, test protocols for static and dynamic side-by-side testing of the Berkeley Hood compared to a conventional hood were prepared, and a series of tests were then conducted. A highly respected third-party industry expert, Tom Smith of Exposure Control Technologies, performed and/or witnessed the tests and provided an independent assessment of the results. The test results were provided to CAL/OSHA through the variance process to enable demonstration of the hood in industrial settings.

- **Perform industrial field demonstrations**

  Industry partners interested in demonstrating the Berkeley Hood were solicited, and a California fume hood manufacturer, Genie Scientific, was selected to manufacture the Berkeley Hoods. Variance applications were prepared for two of the industry partners:

the National Food Service Laboratory and Chevron Texaco. Several hearings were held with Cal/OSHA in order to attempt to develop mutually agreeable criteria that would allow them to issue variances for this technology.

3.3. Data Centers

Similar to the cleanrooms and laboratories tasks, the approach for the data center tasks involved performing high priority research as identified in the PIER research roadmap entitled "High-Performance Data Centers." The approach to each task is detailed below:

- **Investigate data center efficiency opportunities and perform energy benchmarking**

  The approach for the energy benchmarking portion of this task involved soliciting data center sites for benchmarking through contacts within individual companies, industry associations, California public utilities, and project advisors. The initial focus was on obtaining benchmarking partners from Southern California, since previous data center benchmarking was primarily conducted in Northern California. Typically, one or more meetings were held to explain the benchmarking process and its benefits to the host site. Once an industry partner agreed to participate, a meeting was held with the host site and LBNL's subcontractor, EYP Mission Critical Facilities. The discussion included a description of the systems that would be monitored and how the monitoring would take place. A limited amount of design information that would be used to help determine monitoring points was requested. The subcontractor then arranged to be on site for several days to collect measurements and/or obtain existing information from building management systems. Since data center energy use is typically not dominated by weather conditions unless outside air economizers are used, a relatively short monitoring period was used—usually three to four days. Then over a several-week period, the team analyzed the data and developed site specific reports. These reports provided the data collected by measurement or other means, along with energy efficiency recommendations based upon observations while on site. Finally, the report was reviewed and explained with the host site. The host site was given the opportunity to correct any facts in the report, and then the anonymous reports were posted on LBNL's website: http://hightech.lbl.gov/benchmarking_dc.html.

  Utility representatives were invited to attend meetings where benchmark summary results were presented.

  The benchmarking determined energy intensity and end use within the center and focused on the efficiency of key systems and standby generation. All benchmark results were entered into spreadsheets, and summary graphs including prior benchmarks were prepared in order to focus on the energy performance for selected systems and components.

  In a synergistic effort, PG&E began developing data center design guides during this project. The LBNL team collaborated with PG&E's contractor, Rumsey Engineers, in the development of "best practice" design guides. By reviewing all of the benchmark data and recommendations from the current case studies, combined with the information

contained in case studies from a prior PIER project and various other sources, the team identified key areas where improved design and operation led to better energy performance (best practice). Due to the timing of this effort, the team simply developed the design guides rather than preparing separate summary documents.

In all, seven data center spaces at five sites were benchmarked during this phase. Two of the sites were in Southern California, and three were in the San Francisco Bay Area.

Finally, a self-benchmarking protocol was developed based upon the metrics and methods used in the PIER benchmarking task. This protocol was submitted to a number of industry partners for review and comment. Comments were incorporated, and the protocol was posted on LBNL's website for use by data center professionals.

- **Improve uninterruptible power supplies**

  Ecos Consulting and Electric Power Research Institute (EPRI) were selected as subcontractors to investigate uninterruptible power supply (UPS) system efficiency. First, in order to determine the energy performance of existing systems, a test protocol was developed. Then all of the major manufacturers of larger UPS systems used in data centers were surveyed and test data for individual systems was obtained. The test results were then plotted for comparison purposes. Redundancy strategies (e.g., 2N, N+1, 2N+1, etc.) have an effect on overall energy consumption, and the more common configurations were studied to determine their relative efficiency.

  A set of draft labeling criteria was developed based upon a similar labeling criteria produced in Switzerland. The labeling criteria could be used in programs such as ENERGY STAR® or utility incentive programs.

- **Performance metrics**

  This investigation sought to identify metrics that could be used to compare the computing and energy performance of like kinds of information technology (IT) equipment, i.e., comparisons during different operational states. To accomplish this, various manufacturers were contacted, and available benchmark metrics were evaluated. During the project, the U.S. Environmental Protection Agency's (EPA's) ENERGY STAR program became interested in developing a metric and convened an industry workshop to seek agreement on applicable approaches. Also, the U.S. EPA has endorsed the concept of developing performance metrics and is planning to develop criteria for ENERGY STAR rating of servers and will utilize a performance metrics approach.

- **Energy efficient power supplies**

  In this task LBNL worked with subcontractors Ecos Consulting and EPRI to investigate the market for power supplies commonly used in servers. The most prevalent "form factor" (typically used in low-end servers) was selected for more detailed energy performance evaluations. A testing protocol was developed, and many individual power supplies were tested to determine their relative performance. Energy

performance over a full range of loading was measured to see how energy efficiency varied as the power supply was loaded. This information was summarized graphically to readily determine the range of efficiencies in the market.

3.4. Demonstration Projects

This task focused on identifying possible technologies or strategies that would be appropriate to demonstrate to industry through public interest support. Initially, selection criteria were developed that included the following key criteria:

- Use of new or underutilized technologies, strategies, and/or concepts applicable to any of the target building types. The demonstrations were to focus on technologies that were ready to be implemented, i.e., no research would be involved.
- Industry participation in demonstrations ideally involving key industry leaders in California and reaching a large number of California companies involving a cross-section of the state. Typical cost sharing of 25%–75% was expected.
- Technologies that have the potential for significant energy savings if implemented on a wide scale.
- Additional benefits such as cross-cutting applicability in other building types, or other non-energy benefits such as improved worker safety.
- Technologies that industry would not otherwise undertake or readily adopt.
- A reasonable probability of success.
- Cofunded demonstrations receiving higher priority.
- Priority to those demonstrations involving multiple facets of the industry (e.g., building owners, designers, equipment suppliers).

Dozens of potential demonstrations were suggested by project advisors, industry partners, or other researchers. Of these, 15 were presented as possible demonstrations, and PIER selected the following 5 demonstrations to be included in the project:

- DC powering a rack of servers
- Alternative cooling of a rack of servers (subsequently cancelled)
- Demand-controlled filtration in an industrial cleanroom
- Fan-filter testing of 15 or more units
- Data center design process

3.5. Technology Transfer

The approach to technology transfer activities was to use as many channels of communication as possible, with efforts tailored for distinct target audiences, in order to reach the largest number of stakeholders. This included the following:

- An extensive website where detailed technical information was presented and continuously updated and maintained

- Periodic newsletters distributed to large audiences
- Trade publications
- Video documentaries
- TV interviews
- Journal articles
- Workshops with industry
- Collaboration with industry associations and professional societies
- Interim reports of findings
- Best practice summaries
- Utility workshops/training
- Individual requests for information
- Interaction with PIER
- Industry press releases

Project findings for each task are summarized below and detailed task reports are provided as appendices to the report.

4.0 Project Outcomes

4.1. Cleanrooms and Laboratories

- **Benchmark laboratory and cleanroom facilities and document best practices**

  The project successfully arranged for laboratory and cleanroom facilities to be studied. In total, 11 cleanrooms at four sites and one laboratory were benchmarked. Individual case study reports for each facility are attached as Attachment I. For each site, energy benchmarks were obtained through direct measurement or from building management systems where direct measurement was not practical. In some cases design information was used in lieu of direct measurement. For example, if entry into the cleanroom to obtain airflow measurement was not allowed, then design or test and balance were obtained. The case studies each contain recommendations of potential energy efficiency improvement areas that were observed during the course of the benchmarking. Key findings from these case studies include:

  - Two of the cleanroom sites were reducing airflow in their cleanrooms at night and weekends when no workers were in the rooms. In one case, 70% energy reduction was evident. See Figure 1.

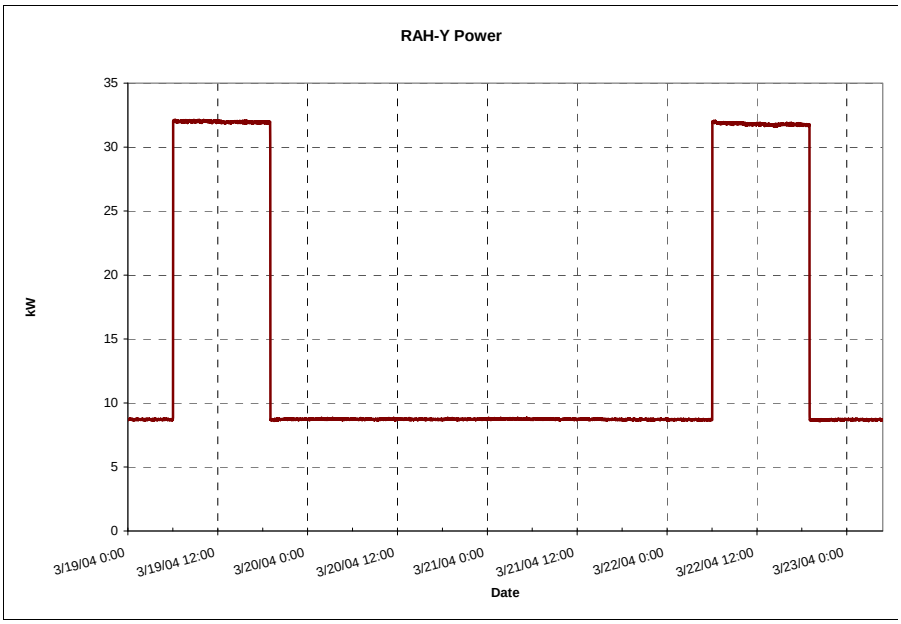

Figure 1. Recirculation fan energy savings during setback at Facility L

Source: Lawrence Berkeley National Laboratory

In the other case, set-back was not occurring as planned due to a controls problem. This highlights the need for continuous commissioning. Once the controls problem was identified and corrected, this site showed a similar large energy reduction of approximately 75%, as illustrated in Figure 2, resulting in estimated yearly savings of approximately $138,000.

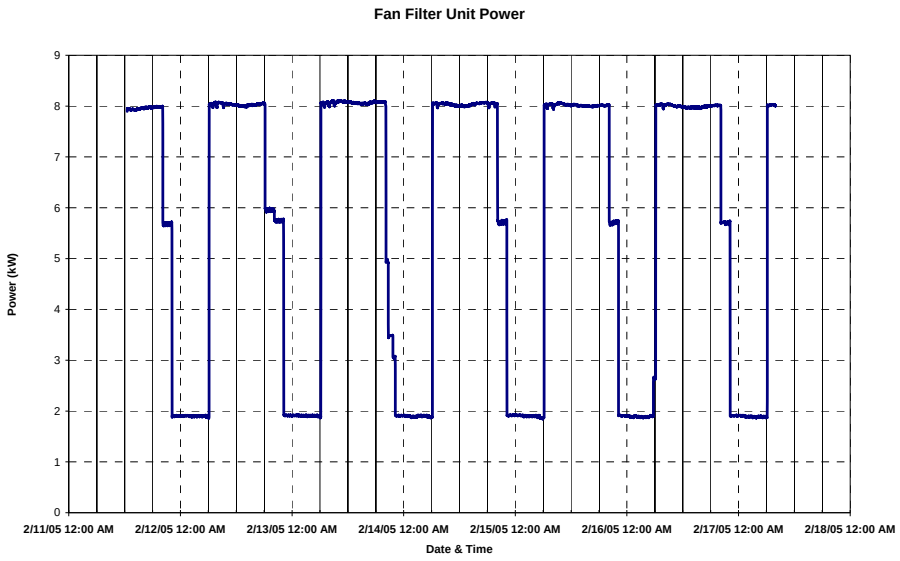

Figure 2. Recirculation fan energy savings during setback at Facility K

Source: Lawrence Berkeley National Laboratory

As an indicator of the energy efficiency opportunity for these building types, many observations of efficiency opportunity were presented for each of the sites. Recommendations included no/low-cost measures as well as more capital-intensive improvements. Recommendations suggested that lower air-change rates (recirculation) in cleanrooms should be possible.

It was found that many systems could benefit from better humidity control.

Benchmark results were recorded in the Access™ database developed during the prior PIER benchmarking project. Summary graphs of various metrics were prepared and posted on LBNL's website: http://hightech.lbl and are also included in Attachment I.

Figure 3 illustrates benchmark results for air recirculation systems in cleanrooms.

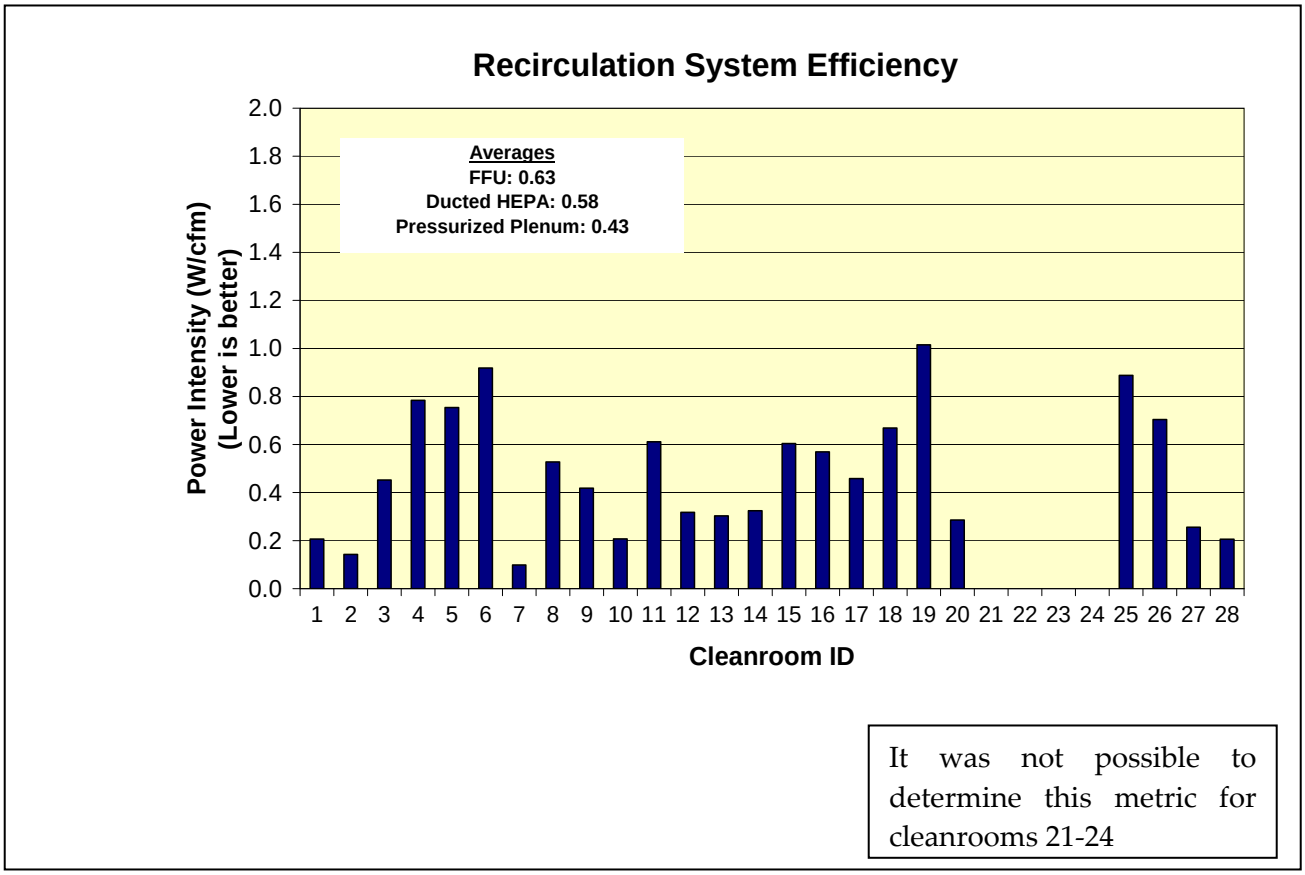

Figure 3. Cleanroom recirculation air system efficiency

Source: Lawrence Berkeley National Laboratory

A review of current and prior case studies was performed to identify the "best practices" that helped in achieving higher energy efficiency in cleanrooms. Based upon this review, best practice topics were identified, and summaries were prepared that focused on the following measures:

- Low pressure drop design
- Recirculation air-change rate optimization
- Demand-controlled filtration
- Fan-filter efficiency
- Use of mini-environments
- Exhaust optimization
- "Right-sizing" systems
- Process systems vacuum pump optimization
- Variable-speed pumping
- Variable-speed chillers
- Use of free cooling
- Dual-temperature cooling loops

- Cooling tower and condenser water optimization
- Control of chilled water systems
- Air-recirculation system type
- Exhaust systems

Through collaboration with PG&E, the best practice summaries were used to develop design guides and were finalized by Rumsey Engineers for PG&E. Lawrence Berkeley National Laboratory provided review and input to the design guides. The best practice summaries are available here: http://hightech.lbl.gov/cleanrooms_bpg.html and contained in Attachment I of this report.

- **Advance energy-efficient filtration: Standard testing for fan-filter units**

  Fan-filter units are commonly used in cleanrooms for both new construction and retrofits of existing spaces in virtually all types of industries and institutions that use cleanrooms. Lower capital cost, ease of construction, and flexibility are frequently cited reasons for their selection. Without a standard test procedure, however, an owner or designer wishing to compare competing units based upon their energy performance had no common basis for comparison. In this task, first a standard test method was developed in collaboration with the Industrial Technology Research Institute (ITRI), in Taiwan, members of the Institute of Environmental Sciences and Technology (IEST), and other interested industry partners. IEST is developing a "recommended practice" for fan-filter units that will address many facets of its design other than energy use, and it has committed to including the test procedure in the recommended practice document.

  The proposed standard test procedure is attached in Attachment II.

  A trial run of the procedure was performed in conjunction with ITRI at their AMCA certified test facility in Taiwan. A second test was performed in conjunction with PG&E at its San Ramon, California, test facility. The two test configurations were significantly different, as was the electrical power supply. As a result, the tests were not able to obtain repeatable results. Further testing in a controlled test setup was recommended, and subsequently a task to demonstrate testing of multiple units was included in the second phase of the project. (See demonstrations.)

  A report of the fan-filter test procedure development and its use in the initial two tests (LBNL Report: LBNL-57727) is attached in Attachment II.

- **Improvements in efficiency of heating, ventilating, and air conditioning in high-tech buildings: Demand-controlled filtration and use of mini-environments**

  Demand-controlled filtration

  The pilot study conducted in LBNL's cleanroom demonstrated the technology of sensing particles in the room and directly controlling cleanroom airflow based upon particle concentrations. Various particle counters were used to study the impact of particle sizes

and concentrations. This study concluded that the technology to control airflow based upon real-time particle counts is viable; however, additional study is needed to examine the relationships between air change rates and particle counts. Figure 4 illustrates typical data that were collected and analyzed during the pilot study. This data provides insight into the relationship between contamination (particle counts) and air change rates. Also, a better understanding is needed concerning the ability to control by counting various sizes of particles.

A report of the pilot study is provided in Attachment III.

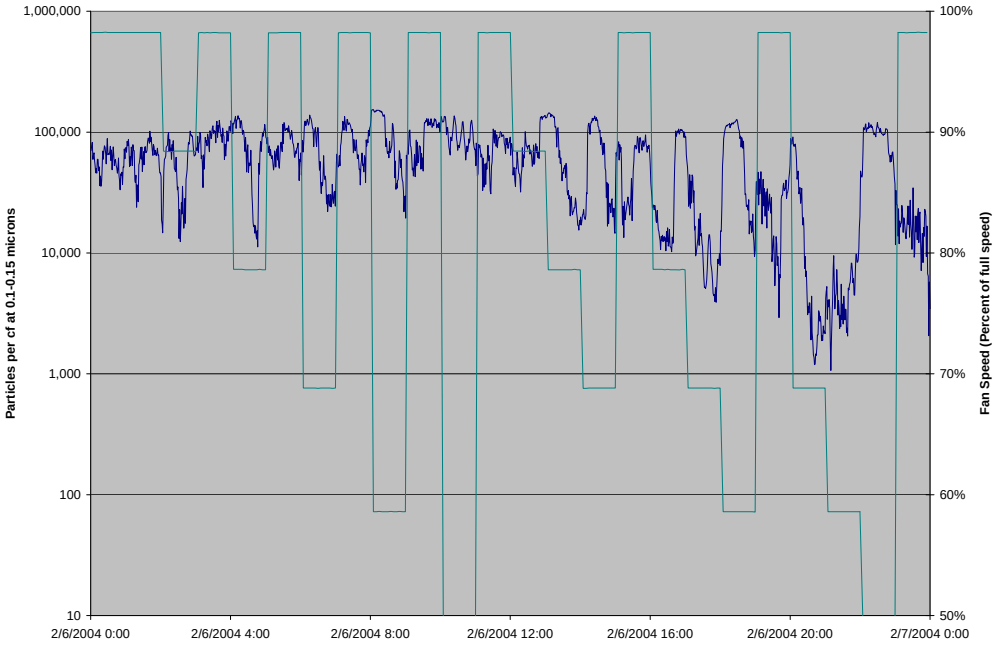

Figure 4. Particle counts and fan speed

Source: Lawrence Berkeley National Laboratory

Based upon the pilot study, a demonstration of the technology in an operating industrial cleanroom was proposed and subsequently added to the project. See Section 4.4 – Demonstrations for a description of the demonstration activities.

Mini-environments

The mini-environment study included investigation into the energy efficiency opportunity of mini-environment devices (micro level) and an investigation into the energy implications when mini-environments are deployed throughout cleanrooms (macro level). The micro-level investigation evaluated the energy performance of the fan-filter unit, which is used to provide filtration through high-efficiency particulate air (HEPA) filters. This study evaluated considerations such as leakage from the mini-environment, exhaust opening, fan speed control, air change rates, and more.

The study investigating energy efficiency issues in mini-environments was performed with an industry partner who manufactures mini-environments. The manufacturer

provided access to a mini-environment for assessment of its performance. Airflow and electrical fan energy were measured throughout its full operating range. Although beyond the original scope of this task, an additional study or "macro" study was able to be accomplished. This involved working with one of the benchmarking sites to study the aggregate effects of using many mini-environments within a larger cleanroom. Reports of both studies are included in Attachment IV.

4.2. Berkeley Fume Hood Development

This task was designed to overcome the institutional barrier presented by the requirement of 100 ft./min face velocity mandated by CAL/OSHA. This work led to a better understanding of the Berkeley Fume Hood performance.

- **Overcoming institutional barriers**

  The primary institutional barrier to acceptance of the Berkeley Hood was the CAL/OSHA requirement to maintain 100 ft/min. of airflow through the face of the hood opening. Evaluation criteria were negotiated with Cal/OSHA through a number of meetings and information exchanges. While LBNL attempted to get agreement on acceptance criteria, there was no agreed-upon level of performance that would satisfy Cal/OSHA either before or after the performance tests. Side-by-side tests were agreed upon in order to compare the performance of the Berkeley Hood with a standard hood. Although the results of the side-by-side testing provided sufficient justification for Cal/OSHA to issue a variance, it became evident that the only way that Cal/OSHA would accept use of the Berkeley Hood would be through individual variances rather than a blanket acceptance. At this point, PIER management and LBNL decided to terminate further work in this area.

- **Conduct side-by-side tests with conventional hoods**

  A series of tests were conducted using the protocols for static and dynamic tests for both a conventional fume hood and the Berkeley Hood. These tests involved standard industry tests but also included more rigorous dynamic tests that are not typically performed for standard hoods. The testing was observed/conducted by a third-party fume hood expert, Tom Smith of Exposure Control Technologies, who provided an independent assessment of the results. The test results showed that the Berkeley Hood performed as well or better than a standard fume hood, and these results were acceptable to Cal/OSHA to grant a variance for use of the hood in one location.

- **Perform industrial field demonstrations**

  Industry partners interested in demonstrating the Berkeley Hood were identified. The National Food Service Laboratory, Dublin, California; Chevron Texaco, Richmond, California; and Genentech expressed interest in participating in the demonstrations. A California fume hood manufacturer, Genie Scientific, was selected to manufacture the Berkeley Hoods, and LBNL procured three hoods for demonstration—although actual manufacture was deferred pending variance

approvals. Variance applications were prepared for the National Food Service Laboratory and Chevron Texaco. Several hearings were held with Cal/OSHA, and a variance was granted for the first installation at the National Food Service Laboratory.

However, since it became evident that Cal/OSHA would not grant a blanket variance for this technology, the fume hood work was terminated by PIER, so consequently the demonstrations were cancelled.

Various reports containing the testing protocol and test results are attached as Attachment V.

4.3. Data Centers

- **Investigate data center efficiency opportunities and perform energy benchmarking**

  Data centers were identified for benchmarking through interaction with industry associations (e.g., Critical Facilities Roundtable), subcontractors, and public utilities. Initial efforts focused on obtaining benchmark sites in Southern California, since previous benchmarking primarily involved Northern California. As a result of these efforts, two Southern California companies were identified having three distinct data center spaces. In addition two firms in Northern California were recruited, accounting for an additional three data center spaces. The centers included Web hosting, Internet service provider, and network equipment manufacturers.

  For each site, energy use was determined through direct measurement, use of building management data, or other available means. Case study reports were prepared to summarize the energy end-use breakdown in the centers and to provide recommendations based upon observations. The case studies were anonymous reports to protect the confidentiality of the host site. These reports are included in Attachment VI and available through LBNL website: http://hightech.lbl.gov/benchmarking_dc.html. The report, *Data Centers and Energy Use – Let's Look at the Data* (Tschudi et. al. 2003b), discussed much of the benchmark data.

  The case study findings from this study and prior benchmarking revealed a large variation in performance in virtually every aspect of data center operation. The observed differences suggested that applying better design and operational practices could lead to greatly improved efficiency. Figure 5 illustrates the wide range of performance exhibited by all of the centers benchmarked by LBNL to date. This metric is the ratio of IT power to the total power delivered to the data center and its supporting systems. It gives an indication of the amount of electrical power that is actually powering the IT equipment.

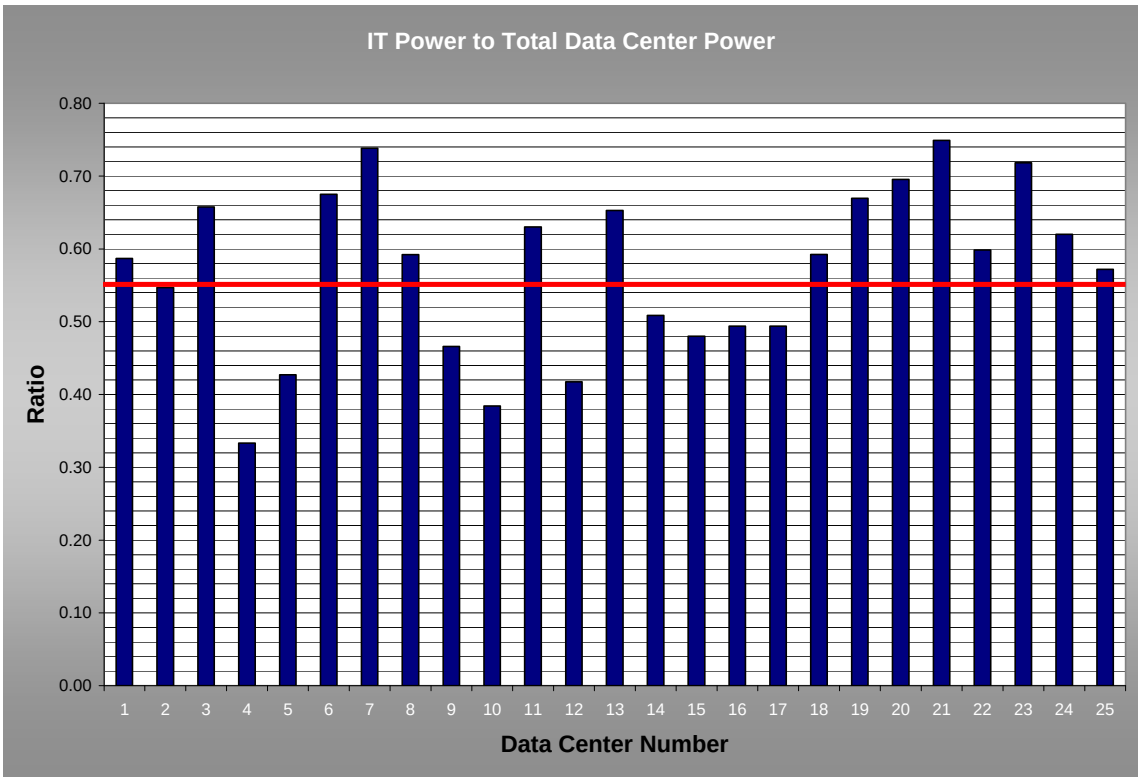

Figure 5. Ratio of IT power to total data center power

Source: Lawrence Berkeley National Laboratory

By examining the design of key systems in centers that had higher ratios, best practices were identified. Design guidelines for selected systems were developed in collaboration with PG&E and Rumsey Engineers. The guidelines include the following topics:

1. Air management (guidelines for improving air flow)
2. Air economizers
3. Centralized air handling
4. Cooling plant optimization
5. Direct liquid cooling
6. Free cooling via water-side economizers
7. Humidification controls alternatives
8. Power supplies
9. Self-generation
10. Uninterruptible power supply systems

A self-benchmarking protocol was developed to provide guidance for end users to obtain the same benchmark data for their data centers. This protocol outlines metrics and the procedure to obtain benchmarks similar to the data collected by LBNL. Use of this protocol will enable comparison of like metrics among the benchmarked facilities. The protocol is included in Attachment VI and is available through the LBNL website: http://hightech.lbl.gov/datacenters.

- **Improve uninterruptible power supplies**

  Ecos Consulting and EPRI measured UPS system performance throughout the operating range for essentially all commercially available UPS systems. The results revealed that there is a large variation in performance, as shown in Figure 6. A 30% variance in performance was found at low utilization levels, and a 10% variance was found for fully loaded equipment.

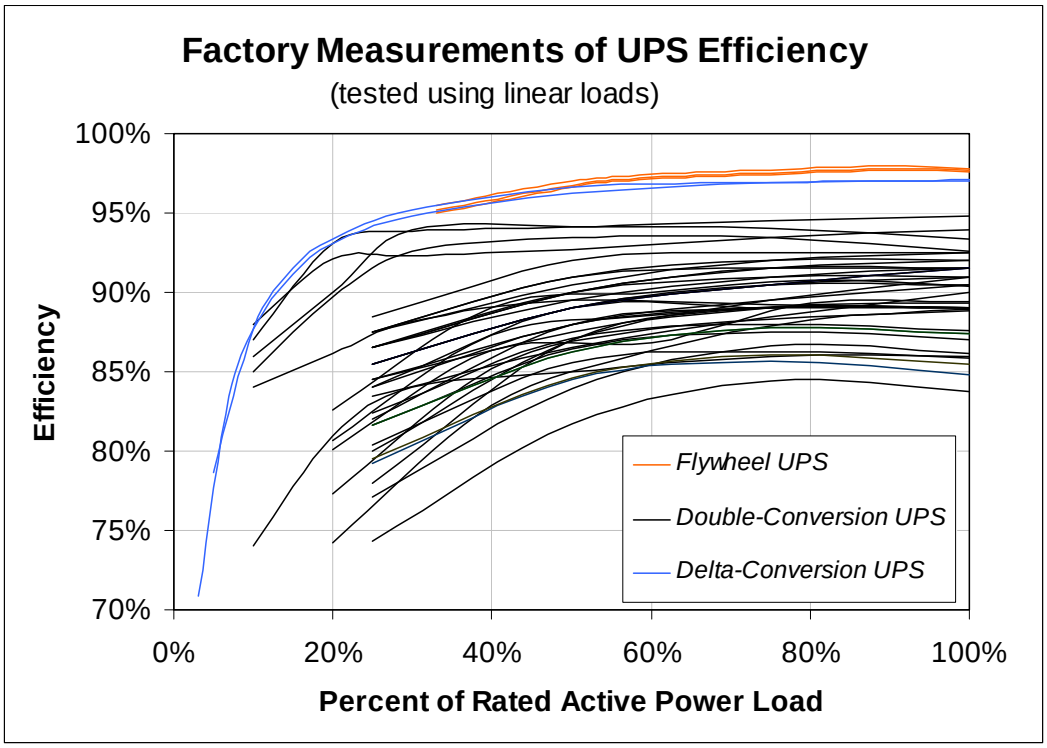

Figure 6. UPS system efficiencies
Source: Lawrence Berkeley National Laboratory

Redundancy configurations also influenced the ultimate efficiency of the electrical distribution systems.

A draft labeling criteria was developed that could be used in programs such as ENERGY STAR or utility incentive programs.

The full report of the UPS task is in Attachment VII.

- **Performance metrics**

  In this task, the concept of coupling computational work with energy used to produce it was advanced. By working with various industry groups, the metrics appropriate to various computational workloads began to be identified. Different computational work loads require different metrics to evaluate performance. For example, scientific computing has a much different metric than, say, Web hosting. LBNL collaborated with the Standard Performance Evaluation Corporation (SPEC) to develop the first performance metric protocol. The research team continued collaboration with the ENERGY STAR program as it sought to develop performance metrics for a broader range of applications. The work in this task is further discussed in Attachment VIII.

- **Energy-efficient power supplies**

  In this task subcontractors Ecos Consulting and EPRI investigated the market for power supplies used in low-end servers. Measurements of the most prevalent "form factor" were performed. A testing protocol was developed and then used to evaluate many individual power supplies to determine their relative performance. Energy performance over a full range of loading was measured to see how energy efficiency varied as the power supply was loaded. This information showed that, similar to UPS systems, there was a wide range of performance, as shown in Figure 7. A factor-of-two variance in performance was found at low utilization levels, and a 30% variance was found for fully loaded equipment.

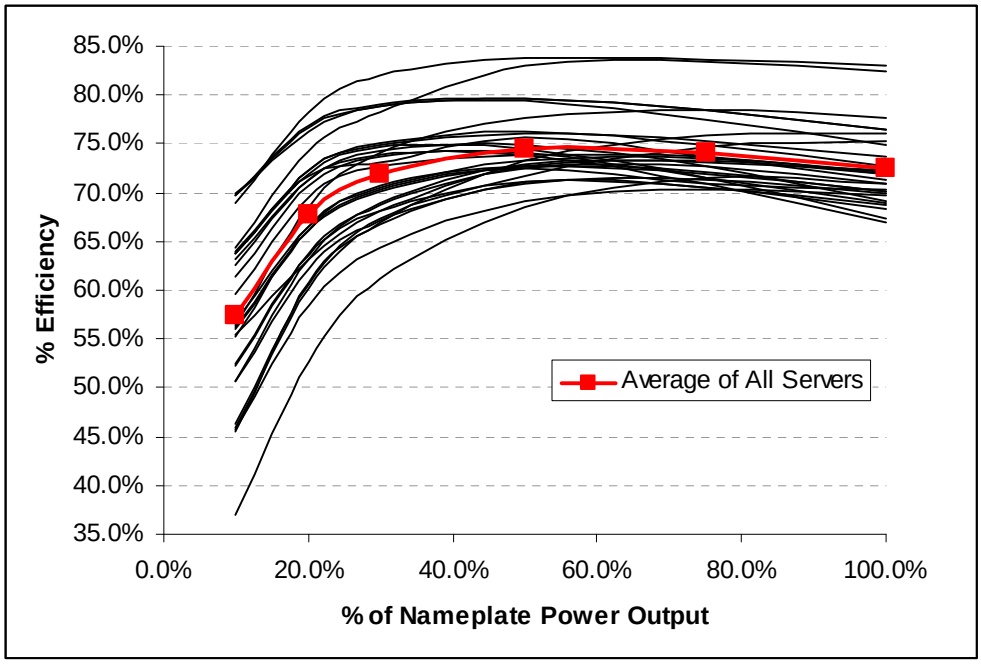

Figure 7. Performance of power supplies

Source: Lawrence Berkeley National Laboratory

The report of this task is included in Attachment IX.

4.4. Demonstration Projects

The following demonstrations were conducted:

- Demand-controlled filtration (See Attachment X.)
- Fan-filter testing (See Attachment XI.)
- DC powering of servers (See Attachment XII.)
- Data center design process (See Attachment XIII.)

In each demonstration, industry partners participated by providing equipment, services, and/or use of their facility for the demonstration. Highlights from the demonstrations are as follows:

- Demand-controlled filtration – The host facility thought that cleanroom airflow was being controlled by a timer, but in actuality no setback of airflow was occurring. Having this corrected saved the host facility considerable energy. Use of the demand control strategy was also shown to be effective in controlling cleanroom airflow.
- Fan-Filter Testing – The test procedure developed in an earlier phase of the project was successfully used to test 17 fan-filter units donated by manufacturers. A wide range of performance was observed.
- DC powering of servers – Over 25 firms collaborated on this demonstration by donating equipment, services, and a demonstration site. Intel and Sun Microsystems modified servers to directly accept 380-volt DC. This side-by-side demonstration illustrated that commercially available equipment is available today for a DC system. Energy savings of over 25% compared with typical systems is possible.
- Data center design process – The design process for the design of a new supercomputer facility illustrated the differences in understanding between IT and facility professionals. Through an understanding of the different focus of these organizations, it is possible to bridge the gaps and develop energy-efficient solutions.

4.5. Technology Transfer

Attachment XIV includes a summary of the extensive technology transfer activities included in this task. The technology transfer activities enabled the project results to be provided to a large number of California, national, and international stakeholders. The PIER Program is recognized as leading efficiency efforts in the high-tech sector, as evidenced by the synergistic activities begun by the New York State Energy Research and Development Authority (NYSERDA), U.S. EPA, the U.S. Department of Energy (DOE), and public utilities since this project began. Reporters from the professional and popular press routinely seek information from members of the project team. Technology transfer activities were critical to informing the industry and other public goods efforts concerning the research and demonstrations undertaken through the PIER project.

5.0 Conclusions and Recommendations

5.1. Conclusions

The following conclusions can be drawn from this project:

- Benchmarking of laboratories, cleanrooms, and data centers provides insight into areas for efficiency improvement, and motivates end users to take action. Benchmarking helped to identify best practices and areas where further efforts should be targeted. A wide variation in performance illustrates that some facilities can achieve energy-efficient operation while supporting the critical nature of the operation.
- The self-benchmarking protocol developed for data centers can assist industry in obtaining its own benchmark data. A similar resource for laboratory buildings is available through the EPA/DOE Laboratories for the 21st Century website. Obtaining similar benchmark data will enable comparison to others and comparison of similar facilities within corporations.
- The standard method of testing and reporting performance of fan-filter units developed in this project will enable an "apples to apples" comparison of their energy performance, where before there was no consistent method. This procedure will enable end users to compare energy use when making purchasing decisions. This procedure could also be used in incentive programs.
- Demand-controlled filtration is a viable technology to control cleanroom airflow through real-time monitoring of contamination levels linked to air-flow controls. This technology will enable cleanroom operators to reduce air flow based upon actual cleanroom cleanliness conditions, resulting in large energy savings.
- Energy efficiency improvement opportunities are suggested for mini-environments. In addition, use of mini-environments within cleanrooms can reduce the energy required for the room itself, resulting in large energy savings.
- Side-by-side testing of the high-performance Berkeley Fume Hood compared to a conventional hood illustrated equivalent or better containment with 50% less airflow The institutional barrier introduced by the Cal/OSHA requirement of a 100 ft.-per-minute face velocity, however, would require individual variances for each installation. Midway through the project it appeared that Cal/OSHA would not provide a blanket approval of the Berkeley Hood for use in California. At this point, PIER and LBNL decided to terminate further demonstrations of the technology. The Berkeley Hood has since been licensed and will be marketed outside of California.
- Uninterruptible power supply (UPS) test results showed a wide variation in energy loss, suggesting that use of the more efficient systems in data centers could save 20% or more of the electricity powering IT equipment, in addition to savings from the reduction in HVAC energy use. Benchmarking results indicate that approximately an equal amount of HVAC energy would be saved. In addition, awareness of the energy implications of various redundancy configurations can lead to improved overall system efficiency. This information can be used to establish utility incentives to encourage use of more efficient

UPS units and provide incentives for more efficient redundancy configurations. The draft labeling criteria developed in this project could also be used in California or ENERGY STAR programs to provide information for those that specify and purchase these systems.

- Standby generation consumes more energy than it will ever generate. Further study of options to current standby generation schemes (e.g., use of on-site generation with utility backup) may lead to improved efficiency in this area. Efficiency opportunities also exist within the current standby generator systems.
- Performance metrics considering computational ability coupled with energy performance can be effective metrics to compare performance of IT equipment. Different metrics would need to be developed for use with various applications. For example, Web hosting performance would be quite different than scientific computing.
- Power supplies within IT equipment account for a large percentage of the IT equipment power consumption. Similar to UPS systems, a wide variation in performance was observed, suggesting that encouraging use of highly efficient power supplies could immediately save a significant amount of energy both for the direct loss and for the cooling required to remove the heat produced.
- The demonstration of direct DC powering of IT equipment illustrated that 10%–20% energy savings are possible if DC power is distributed directly into IT equipment. The demonstration showed that 380-volt DC power can be distributed using Underwriters Laboratories (UL)-rated equipment that is commercially available today. Sun Microsystems and Intel Corporation demonstrated that power supplies in IT equipment can be easily modified to accept 380-volt DC. Over 25 firms participated in the demonstration, and this was observed by several hundred industry professionals. The industry is interested in this technology; however, further efforts are needed to standardize distribution voltages and connectors, address safety concerns, and generally overcome current paradigms. Through very wide coverage in industry and mainstream media, this demonstration has already had an impact in raising awareness of the efficiency of the power conversion processes, and it will likely lead to improvements in AC distribution schemes as well.
- The demonstration of the fan-filter test procedure validated the procedure and illustrated that repeatable results can be obtained. Further experimentation also suggested that the procedure is applicable to “push” or “pull” configurations and that the test rig could be shortened in order to make its implementation more attractive to the industry. An analysis of the results of the testing led to a method to graphically show the interrelation of pressure, flow, and energy use.
- The demonstration of demand controlled filtration illustrated how this control strategy can be deployed in industry cleanrooms using existing commercially available equipment. Cleanroom airflow was directly controlled through use of particle counters communicating with cleanroom fan systems. Alternative airflow reduction strategies such as use of occupancy sensors were also shown to be effective.

- The demonstration of the data center design process showed a process where industry experts could interact as a group to brainstorm energy efficiency solutions. The "charrette" process enabled a group of industry experts to explore best practice and cutting-edge solutions without the constraints of a typical compartmentalized design process. The charrette identified energy efficiency strategies and concepts that could be explored in more detail later in the design process.
- Technology transfer activities enabled the project to reach large numbers of stakeholders through various channels for each of the high-tech facility types. Frequent outreach activities were successful in informing a large component of California's high-tech industries and the industries around the world. California is viewed as leading in developing high–tech energy efficiency solutions. PIER is now recognized as the first and leading supporter of research in this domain.

5.2. Recommendations

Research opportunities identified in prior PIER research roadmaps should be reviewed and updated to reflect the rapid pace of innovation in this field. Important opportunities should be pursued. The recommendations from this project are in line with the original roadmaps' recommendations, and include:

- Existing benchmarking data should be made more readily available to industry stakeholders, e.g., through an action-oriented interactive web-based tool and such as that currently being developed by the PIER Buildings program.
- Additional benchmarking of high-tech buildings should be performed to assist in:
    - Motivating industry to take action.
    - Comparing like facilities.
    - Benchmarking performance over time.
    - Identifying trends in industry.
    - Recognizing/labeling to encourage high performance.
    - Defining incentive program baselines.
    - Identifying best practices.
- The standard method of testing and reporting performance of fan-filter units developed in this project should be adopted by standards organizations. Currently IEST is planning to include the procedure in a "recommended practice" (RP). The PIER Program should consider continuing interaction with IEST to ensure that the provisions of the procedure are adopted. Adoption by other organizations such as AMCA should also be encouraged.
- Market transformation activities to promote use of demand-controlled filtration should be pursued for the various sectors that operate cleanrooms (e.g., biotech). This technology should also be explored for use in healthcare facilities.

- Best practices should be publicized through demonstrations and case studies for training and in order to promote wider adoption.
- New and emerging data center cooling solutions should be evaluated to determine their energy implications. Strategies such as demonstrations or incentives should be explored to encourage use of solutions that offer the most efficient operation.
- The same efficiency issues identified through the evaluation of UPS and standby generation systems in data centers exist in a large number of other facilities that provide for backup power (e.g., hospitals). Outreach targeting a broader cross-section of other market sectors would expand the energy savings potential.
- Performance metrics for servers to enable energy and computational comparisons should continue to be developed with collaboration with industry.
- Power supply efficiency improvements should be encouraged through programs similar to the 80+ program.
- Power distribution efficiency within data centers and other high-tech facilities should continue to be a major focus for research, best practice identification, and outreach. Efficiencies of various power distribution designs utilizing AC or DC power should be studied. Collaborations with industry organizations such as Green Grid, IEEE, and others should be established to help identify energy-efficient schemes and evaluate performance of their power conversion devices. Specific activities related to overcoming barriers identified in the demonstration of the use of 380-volt DC power are (1) gaining consensus on distribution voltage, (2) gaining consensus and standardizing design of connectors, and (3) developing pilot installations in order to accelerate adoption.

5.3. Benefits to California

High-tech industries are extremely important to the California economy. Many high-tech firms are headquartered in California and represent most major sectors of the economy, such as biotech, semiconductor, aerospace, and health care. Energy use and intensity are large in these industries and continue to grow as these companies flourish. This project has provided insight into the energy efficiency opportunity in cleanrooms, laboratories, and data center facilities. The various technologies explored show that there is potential for large energy savings using solutions available today. Energy savings of 20%–40% have been shown to be possible. This project has also helped to raise the awareness of industry and has garnered the interest of other public interest programs. California is now seen as leading efficiency efforts in the high-tech sectors.

6.0 References

Bell, G., and D. Sartor. 2004. Side-by-Side Fume Hood Testing: ASHRAE 110 Containment Report—Comparison of a Conventional and a Berkeley Fume Hood. LBID-2560.

Bell, G., and D. Sartor. 2004. Side-by-Side Fume Hood Testing: Human-as-Mannequin Report—Comparison of a Conventional and a Berkeley Fume Hood. LBID-2561.

Ecos Consulting and EPRI Solutions. 2005. High Performance Buildings: Data Centers—Uninterruptible Power Supplies (UPS).

Ecos Consulting and EPRI Solutions. 2005. High-Performance Buildings: Data Centers—Server Power Supplies.

Exposure Control Technologies, Inc. 2005. 'Side by Side' Fume Hood Tests Using ASHRAE 110 and 'Human As Mannequin' (HAM) Test Methodology to Compare Performance of a Jamestown Conventional Fume Hood and a LBNL High Performance Hood. Final Report (March 21).

Faulkner, D., D. DiBartolomeo, and D. Wang. 2007. Demand Controlled Filtration in an Industrial Cleanroom. Lawrence Berkeley National Laboratory Report No. 63420.

Greenberg, S., W. Tschudi, and J. Weale. 2006. Self-Benchmarking Guide for Data Center Energy Performance (Version 1.0). Lawrence Berkeley National Laboratory.

Greenberg, S., E. Mills, B. Tschudi, P. Rumsey, and B. Myatt. 2006. Best Practices for Data Centers: Results from Benchmarking 22 Data Centers. *Proceedings of the 2006 ACEEE Summer Study on Energy Efficiency in Building.*

Jeng, Ming-Shan, T. Xu, and Chao-Ho Lan. 2004. Toward Green Systems for Cleanrooms: Energy Efficient Fan Filter Units. *Proceedings of SEMI Technology Symposium: Innovations in Semiconductor Manufacturing*, SEMICON West 2004. San Francisco (July). pp 73–77. LBNL-55039.

Koomey, G. K., O. Sezgen, and R. Steinmetz. 2005. Estimating Total Power Used by Data Centers in California.

Mills, E. 2007. "Eco-Solutions." *Forbes.* April 23.

Mills, E., and D. Sartor. 2005. "Energy Use and Savings Potential for Laboratory Fume Hoods." LBNL-55400. *Energy* 30:1859–1864.

Mills, E., G. Bell, D. Sartor, A. Chen, D. Avery, M. Siminovitch, S. Greenberg, G. Marton, A. de Almeida, and L. E. Lock. 1996. Energy Efficiency in California Laboratory-Type Facilities. Prepared for the California Institute for Energy Efficiency. LBNL-39061.

Rocky Mountain Institute (RMI). 2003. Final Report on Energy Efficient Data Centers - A Rocky Mountain Institute Design Charrette. Hayes Mansion Conference Center, San Jose, California.

Rumsey Engineers and Lawrence Berkeley National Laboratory. 2006. "High Performance Cleanrooms: A Design Guideline Sourcebook." Part of *Energy Design Resources* collection.

Rumsey Engineers and Lawrence Berkeley National Laboratory. 2006. "High Performance Data Centers: A Design Guideline Sourcebook." Part of the *Energy Design Resources* collection.

Shamshoian, G., M. Blazek, P. Naughton, R. Seese, E. Mills, and W. Tschudi. 2005. "High-Tech Means High-Efficiency." *Forbes*. November 18.

Shamshoian, G., M. Blazek, P. Naughton, R. S. Seese, E. Mills, and W. Tschudi. 2005. High-Tech Means High-Efficiency: The Business Case for Energy Management in High-Tech Industries. Lawrence Berkeley National Laboratory Report.

Tschudi, W. 2007. "DC Power for Data Centres." *Zerodowntime*. March 21

Tschudi, W. F., and P. Rumsey. 2004. "Using Benchmarking to Identify Energy Efficiency Opportunities in Cleanrooms: The Labs21 Approach" *ASHRAE Symposium*.

Tschudi, W. T., D. Faulkner, and A. Hebert. 2005. "Energy Efficiency Strategies for Cleanrooms without Compromising Environmental Conditions." *ASHRAE Transactions*, DE-05-9-1 10pp.

Tschudi, W., P. Rumsey, E. Mills, T. Xu. 2005. "Measuring and Managing Energy Use in Cleanrooms." *HPAC Engineering*, LBNL/PUB-946.

Tschudi, W., T. Xu, D. Sartor, and J. Stein. 2003a. *High-Performance Data Centers: A Research Roadmap.* Lawrence Berkeley National Laboratory Report No. 53483.

Tschudi, W. F., P. Sreedharan, T. Xu, D. Coup, and P. Roggensack. 2003b. "Data Centers and Energy Use - Let's Look at the Data." ACEEE Conference Paper

Tschudi, W., D. Sartor, E. Mills, and T. Xu. 2002. *High-Performance Laboratories and Cleanrooms: A Technology Roadmap.* Lawrence Berkeley National Laboratory Report No. 50599.

U.S. Environmental Protection Agency. 2007. Report to Congress on Server and Data Center Energy Efficiency: Public Law 109-431. Washington, D.C.: U.S. Environmental Protection Agency, ENERGY STAR Program. August 2.

Xu, T. 2007. "Characterization of Minienvironments in a Clean Room: Design Characteristics and Environmental Performance." *Building and Environment*, 42:2993–3000

Xu, T., C-H. Lan, and M-S. Jeng. 2007. "Performance of Large Fan-filter Units for Cleanroom Applications." *Building and Environment*, 42:2299–2304.

7.0 Glossary

AMCA	Air Movement and Control Association (fan manufacturers association)
ASHRAE	American Society of Heating, Refrigeration, and Air Conditioning Engineers
BTU	British thermal unit
CAL/OSHA	California Occupational Safety and Health Authority
Cfm	cubic feet per minute, a measure of ventilation rate
Commission	California Energy Commission
CPR	Critical Program Review
DOE	United States Department of Energy
GW	Gigawatt ($1x10^9$ Watt)
HEPA	High efficiency particulate air
HVAC	Heating, ventilation, and air-conditioning system
IEST	Institute of Environmental Standards and Technology
IT	Information Technology
ITRI	Industrial Technology Research Institute, Taiwan
KW	Kilowatt (1000 Watts)
LBNL	Lawrence Berkeley National Laboratory
MW	Megawatt ($1x10^6$ Watt)
PAC	Project advisory committee
PIER	Public Interest Energy Research Program (Commission)
RD&D	Research, development, and demonstration
RP	Recommended practice
T	Temperature, measured in degrees Celsius or Fahrenheit
TC	Technical committee
TWh	Terawatt hour ($1x10^9$ kWh)

List of Attachments

Attachment Number	Description	Deliverables
I	Cleanroom/laboratory benchmarking	Case study reports Summary benchmark graphs Best practices
II	Development of fan-filter test procedure	Fan-filter test procedure
III	Investigation of demand controlled ventilation	Demand controlled ventilation case study
IV	Investigation of mini-environments	Mini-environment case study
V	Side-by-side fume hood testing	Summary of side-by-side testing
VI	Data center benchmarking	Case study reports Best practices (Design Guidelines) Self-benchmarking protocol
VII	Improvement of uninterruptible power supplies	UPS report
VIII	Investigation of computer performance metrics	Performance metrics report
IX	Energy efficient power supplies	Power Supply report
X	Demonstration of demand controlled ventilation	Demand controlled ventilation report
XI	Demonstration of standardized testing of fan-filter units	Fan-filter testing report
XII	Demonstration of DC powering of IT equipment	DC power demonstration report
XIII	Energy efficient scientific computing	Scientific computing summary report
XIV	Technology transfer activities	Technology transfer report

Attachment I: Cleanroom/Laboratory Benchmarking

CLEANROOM ENERGY BENCHMARKING SITE REPORT

FACILITY K

NORTHERN CALIFORNIA

JANUARY 2005

SPONSORED BY:

LAWRENCE BERKELEY NATIONAL LABORATORY

PREPARED BY:
RUMSEY ENGINEERS, INC
99 LINDEN STREET
OAKLAND, CA 94607
WWW.RUMSEYENGINEERS.COM

Copyright (c) 2004 Lawrence Berkeley National Laboratory. All rights reserved.

Acknowledgements

Special thanks to the entire team for their generous assistance and cooperation throughout the benchmarking process.

Project Team

Lawrence Berkeley National Laboratory	**Rumsey Engineers, Inc**
Bill Tschudi	Peter Rumsey
Dale Sartor	Achai Broner
	Larry Chu

Reproduction or distribution of the whole, or any part of the contents of this document without written permission of LBNL is prohibited. The document was prepared by LBNL for the exclusive use of its employees and its contractors. Neither LBNL nor any of its employees makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any data, information, method, product or process disclosed in this document, or represents that its use will not infringe any privately-owned rights, including but not, limited to, patents, trademarks, or copyrights.

Table of Contents

I. Introduction 1

II. Review of Site Characteristics 1

A. Cleanrooms 1

B. Lab Area 2

III. Site Energy Use Characteristics 3

A. Cleanroom & Lab Area Power Consumption 3

IV. System Performance Metrics 5

V. Site Observations Regarding Energy Efficiency 9

APPENDICES

A. Trended Data Graphs

B. Data Collection and Accuracy Notes

C. Measurement Methodology

I. INTRODUCTION

As part of the California Energy Commission Public Interest Energy Research (PIER) project, energy use at Facility K was monitored November 15 to November 24, 2004. The project is administered by LBNL (Lawrence Berkeley National Laboratory) and funded through the California Energy Commission.

This site report reviews the data collected by the monitoring team and presents a set of performance metrics as well as a complete set of trended data points for energy end uses for equipment supporting and located in the cleanrooms. In addition, the energy end uses for equipment supporting and located in a non-clean lab area were monitored. Energy metrics were established that allow facility owners to evaluate their energy efficiency performance and identify opportunities for improvements that reduce their overall operating costs.

First, the report reviews the site characteristics, noting design features of the cleanrooms and lab area monitored. Second, the energy use for the cleanrooms and lab area is broken down into components. Third, performance metrics recorded through the project are presented. Finally, key energy efficiency observations for the facility are noted. This is not intended to be a full energy audit, merely observations from the site measurement team. The data collected, trended graphs and methodology documentation are included among the appendices.

II. REVIEW OF SITE CHARACTERISTICS

A. Site

Facility K, built in 2001, is a single-story facility located in Northern California, which provides services to the electronics industry. The facility consists primarily of office spaces, two lab areas, and cleanrooms. The facility is a total of 22,600 square feet (sf). The total square footage of the cleanrooms is 1,821 sf. One lab area is 4,645 sf and the other lab area is 1,684 sf. The 4,645 sf lab area was monitored in this study. The remainder of the building (14,450 sf) is made up of office spaces, conference rooms, a lunch area, a shipping and receiving area, bathrooms, and other miscellaneous rooms.

The cleanrooms and one of the lab areas were selected for this study. The air systems, process equipment, and lighting were monitored. The environmental systems serving the cleanrooms and lab area run 8,760 hours a year in order to maintain conditions and provide a safe environment.

B. Cleanrooms

The cleanrooms are made up of one ISO class 4 (class 10) area, two ISO class 5 (class 100) areas, two ISO class 7 (class 10,000) areas, and a non-classified gowning room. Airflow is cascaded from the highest cleanliness to the lowest. The following table summarizes the areas of each cleanroom. The primary area is the certified rated area of the cleanroom. The secondary area is the area where the air is returned to the air handling units.

Table 1. Cleanroom Areas

Cleanroom	ISO Class	Primary Area (sf)	Secondary Area (sf)	Total Area (sf)
136 – Gowning	Not rated	99	18	117
137	7	608	52	660
138 – Inter-gowning	5	72	0	72
139	4	170	12	182
140	7	367	54	421
141	5	334	35	369
Total	N/A	1,650	171	1,821

The cleanrooms are conditioned by a rooftop package unit, AC-5, which provides both makeup air and recirculation air. The air handler supplies air to an interstitial space and then to fan filtered HEPA units (FFUs) located in the ceiling, which provide the recirculation air to the cleanroom. The air is returned via a low-sidewall return. There are four fume hoods located in the cleanrooms with each having its own exhaust fan. There is also a general exhaust fan for heat removal, which is directly connected to various equipment located in the cleanrooms.

In addition to measuring the air handling systems for these areas, measurements were taken on the power consumption of the process equipment located in the cleanrooms as well as the lighting power consumption.

C. Lab Area

The lab area is conditioned by two identical rooftop package units, AC-2 and AC-3, which both provide makeup air and recirculation air. For the purpose of this study, the package units are considered as makeup air handlers. Both air handlers are ducted to attempt an even supply of air into the lab area. There are five fume hoods located in the lab area with each having its own exhaust fan. There is also a general exhaust fan for heat removal, which is directly connected by ducts to the various equipment located in the cleanrooms. Although the facility is unoccupied at night, the systems run 24 hours a day since some experiments require more than a day to complete.

Measurements were taken on the power consumption of the process equipment and lighting in addition to the power use of the air handling systems.

III. SITE ENERGY USE CHARACTERISTICS

A. Cleanroom & Lab Area Power Consumption

The energy consumption attributed to the cleanroom is shown in the table and pie chart below. This information will help for comparisons of different cleanroom air handling system efficiencies.

Table 2. Cleanroom Power Consumption

Description	Average Load (kW)	Average Efficiency
AIR HANDLING		
Makeup Fan	6.5	1,443 cfm/kW
Recirculation Fans [1]	8.3	3,503 cfm/kW
EXHAUST FANS [2]		
EF-6 (Fume hood)	0.50	-
EF-7 (Fume hood)	0.68	-
EF-8 (Fume hood)	0.60	-
EF-9 (Fume hood)	0.49	-
EF-10 (General exhaust)	2.4	-
PROCESS POWER	18.7	11.3 W/sf
LIGHTS	3.2	2.0 W/sf
TOTAL	41.4	-

1. *Air handler, AC-5 provides both makeup and recirculation air; the fan power attributed to the recirculation air was 3.2 kW. The FFU power consumption was measured at 5.1 kW.*
2. *An efficiency for the exhaust fans could not be determined since airflow data was not available.*

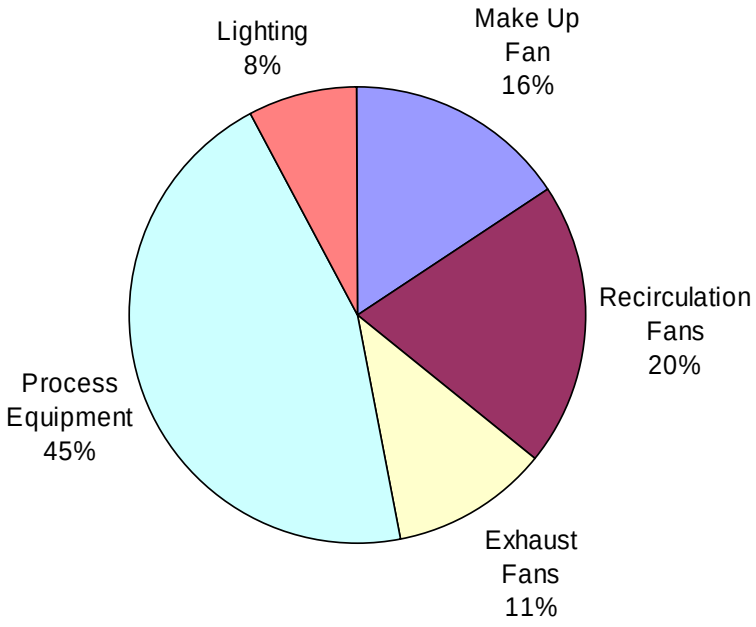

Table 3. Lab Area Power Consumption

Description	Average Load (kW)	Average Efficiency
AIR HANDLING		
AC-2 Supply Air Fan [1]	1.91	3,139 cfm/kW
AC-3 Supply Air Fan [1]	2.02	2,976 cfm/kW
EXHAUST FANS [2]		
EF-3 (Fume hood)	0.97	-
EF-4 (Fume hood)	0.89	-
EF-5 (Fume hood)	0.44	-
EF-11 (General exhaust)	1.5	-
EF-12 (Fume hood)	0.43	-
PROCESS POWER [3]	7.5	1.6 W/sf
LIGHTS	2.4	0.52 W/sf
TOTAL	18.0	-

1. *The air handlers provides both ventilation and recirculation air.*
2. *An efficiency for the exhaust fans could not be determined since airflow data was not available.*
3. *Spot measurements were taken only for the process equipment.*

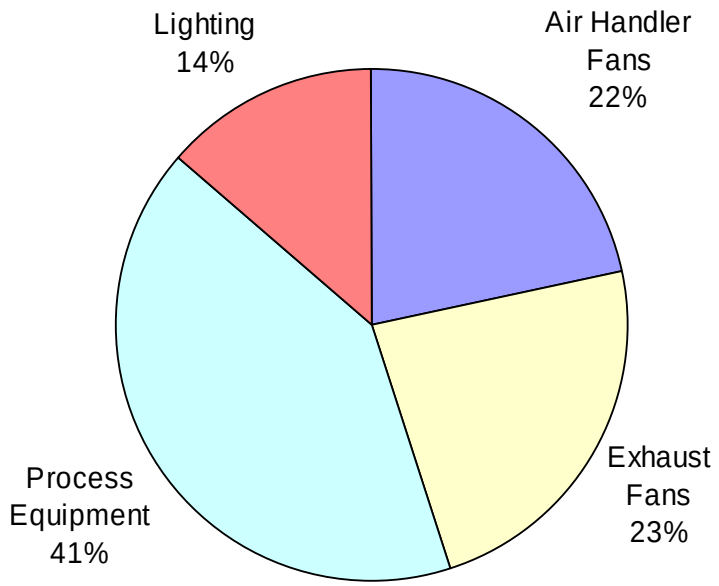

IV. SYSTEM PERFORMANCE METRICS

Metrics are ratios of important performance parameters that can characterize the effectiveness of a system or component. In order to gage the efficiency of the entire building system design and operation, this project tracks key metrics at different system levels. These metrics can be used to compare designs or determine areas with the most potential for improvement via retrofit or replacement.

For Facility K, the cleanroom HVAC components operate at a nearly constant level throughout the year. Therefore, these metrics are based on spot measurements, and various logged data to ensure accuracy. All of the metrics involving area are based on the primary cleanroom area, which is the area that passes certification, unless otherwise noted.

Table 4. Cleanroom Metrics

Description		
MAKEUP AIR HANDLING		
Makeup Airflow	cfm	9,445
Makeup Air Fan Power	kW	6.5
Makeup Air Efficiency	cfm/kW	1,443
Makeup Air	cfm/sf	8.5
Makeup Air Fan Power Density [1]	W/sf	5.9
RECIRCULATION AIR HANDLING		
Recirculation Airflow	cfm	28,908
Recirculation Air Fan Power	kW	8.3
Recirculation Air Efficiency	cfm/kW	3,503
Average Filter Face Velocity	fpm	57
Recirculation Air Fan Power Density [1]	W/sf	5.0

1. *Calculated as total kW load divided by the primary area of the cleanroom.*

Table 5. Cleanroom Airflow Rates

Description		***Room 139 ISO Class 4***	***Room 138 ISO Class 5***	***Room 141 ISO Class 5***	***Room 137 ISO Class 7***	***Room 140 ISO Class 7***	***Room 136 (unrated)***
Recirculation Airflow	cfm	8,712	2,712	14,027	3,744	3,733	375
Air Change Rate	ACH	342	251	280	41	69	25
Recirculation Air	cfm/sf	51	38	42	6	10	4
Average Filter Face Velocity	fpm	69	45	63	61	36	49

Figures 1 and 2 on the following pages compare the recirculation and makeup air handling systems benchmarked in this study. The fan filter unit recirculation air handling system for this facility is more efficient than what was benchmarked at most other facilities. This is due to the low operating ranges of face velocities for the fan filter units. The operating face velocities of the FFUs ranged between 36 to 69 fpm. The FFU face velocity measured in Facilities B.1 and B.2 was both 76 fpm, and 61 fpm in Facility E. The average FFU face velocity for this facility was 57 fpm. Lower face velocities result in lower

pressure drop (or resistance) of the air moving across the filter. This equates to lower fan power consumption to circulate the air through the unit.

Figures 3 and 4 compare the air change rates of the one ISO class 4 (class 10) and the two ISO class 5 (class 100) cleanrooms to other cleanrooms of similar class at the different facilities benchmarked. The ISO class 4 cleanroom air change rate is the lowest of the six facilities as shown in Figure 3. This shows that cleanliness in a cleanroom can still be maintained at lower amounts of airflow. In Figure 4, the cleanrooms at this facility are operating close to the average air change rate of all of the cleanrooms. An opportunity for energy savings is to turn down of the recirculation airflow for these two cleanrooms (also see *Site Observations Regarding Energy Efficiency*).

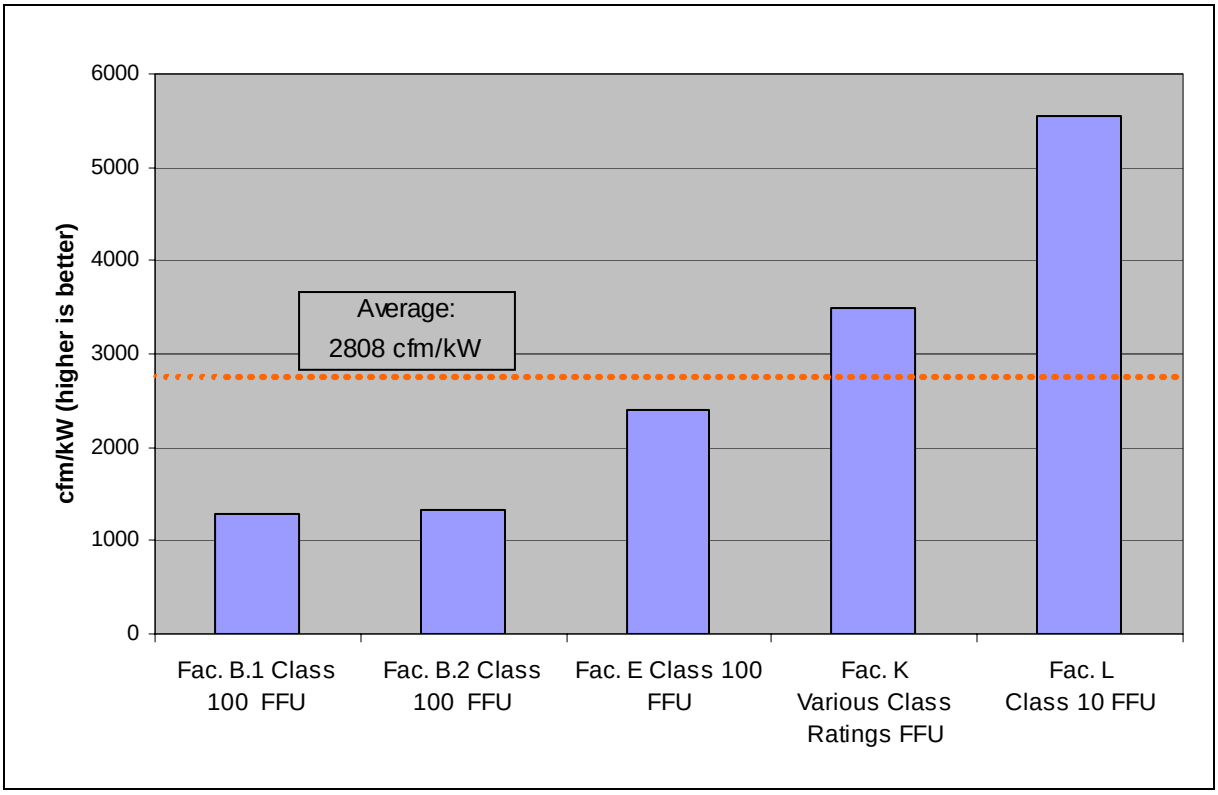

Figure 1. Fan Filter Unit Recirculation Air Handling Comparison

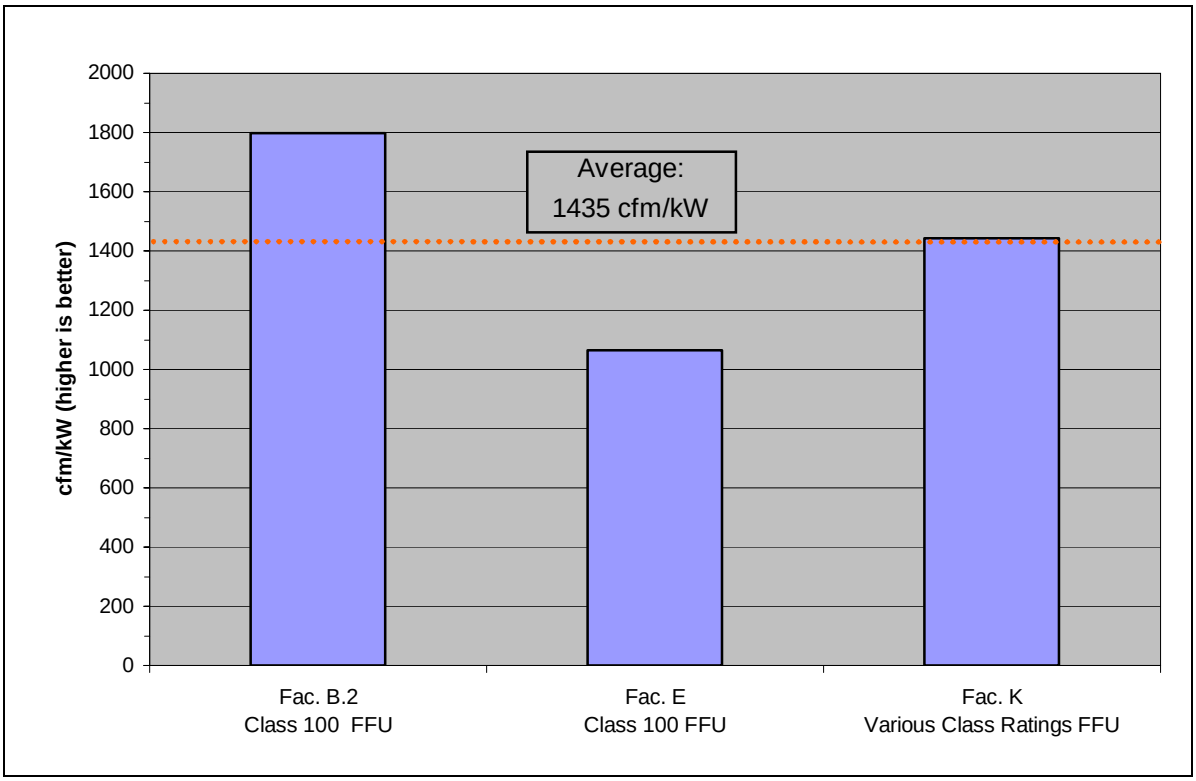

Figure 2. Makeup Air Handling Comparison

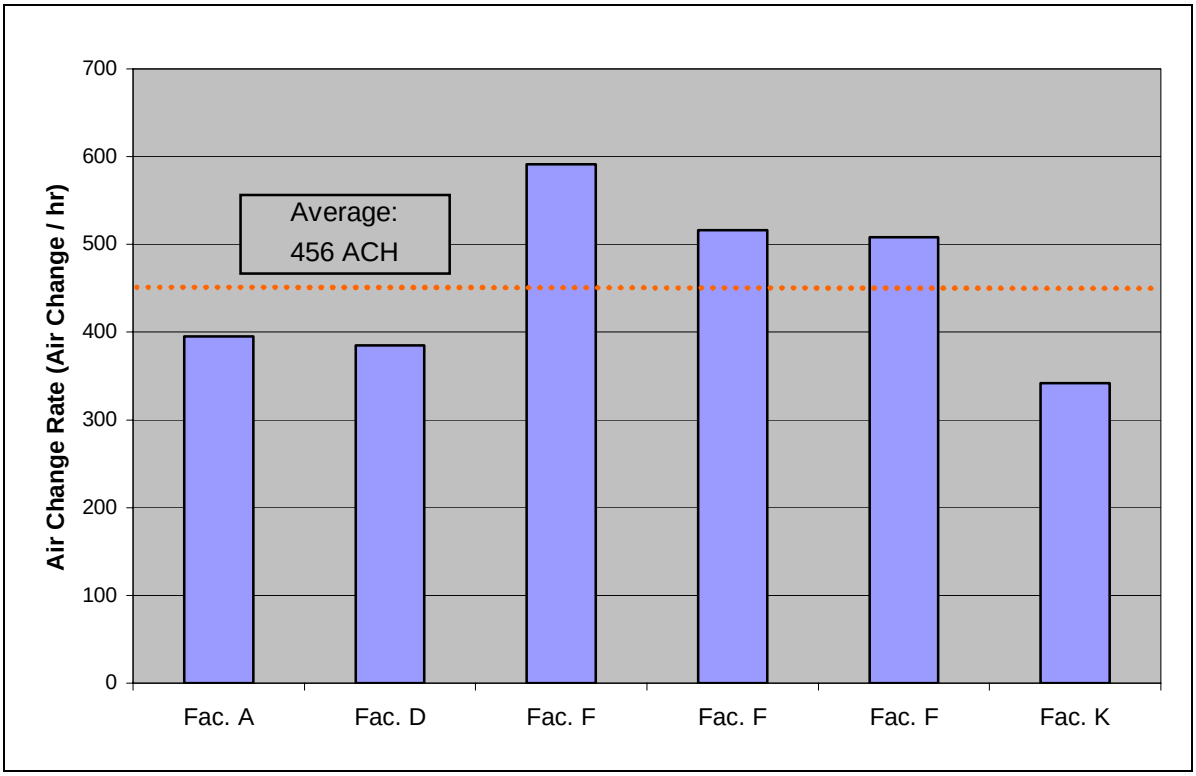

Figure 3. ISO Class 4 (class 10) Cleanroom Air Change Rate Comparison

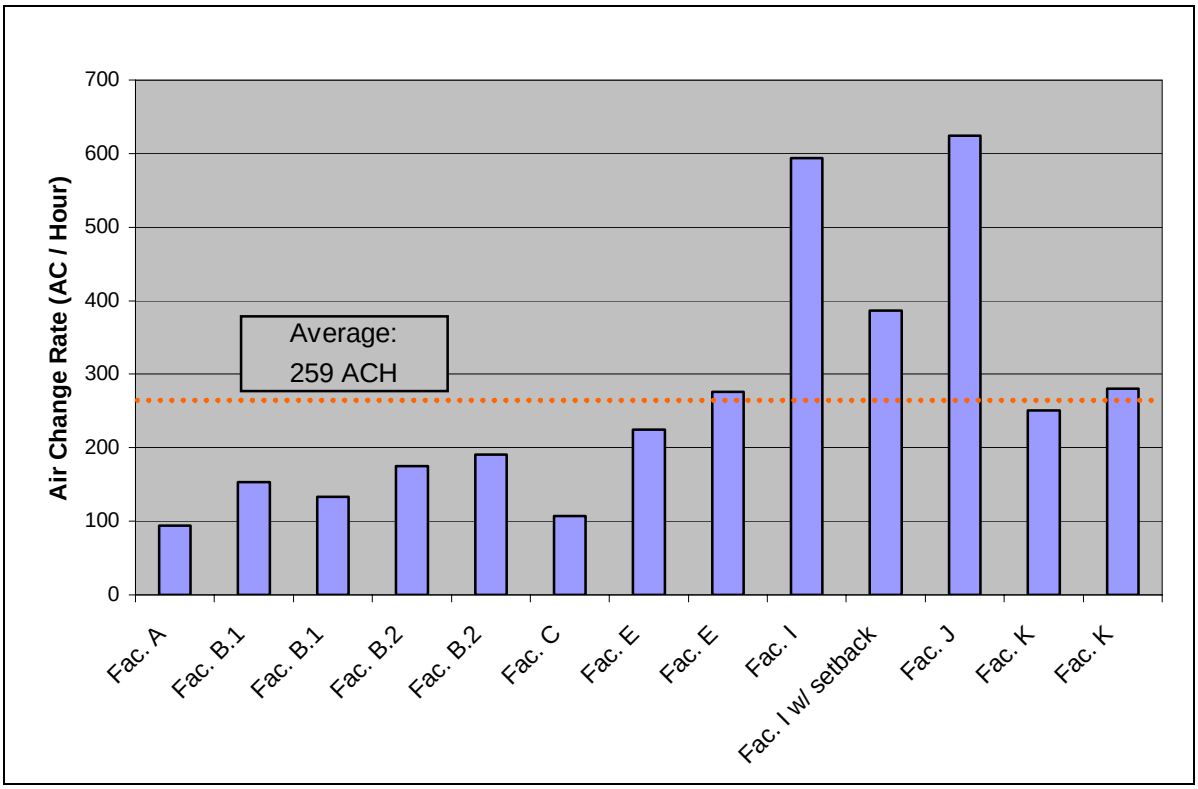

Figure 4. ISO Class 5 (class 100) Cleanroom Air Change Rate Comparison

VI. SITE OBSERVATIONS REGARDING ENERGY EFFICIENCY

Fume Hood Face Velocity

The five fume hoods in the lab area are set for a face velocity of 100 fpm with a sash height of 24". Standard laboratory practice recommends a face velocity of 100 fpm at a sash height of 18". By reducing the face area, a lower amount of airflow will need to be exhausted by each exhaust fan. The reduced airflow can be achieved by adjusting the speeds of the fan motors by changing the sheaves. This will result in fan power savings as well as cooling and heating energy savings, since less of the conditioned air will be exhausted from the space.

The cleanroom fume hoods did not indicate a face velocity. However, the face velocity should be checked and set to 100 fpm at a sash height of 18".

Vacuum Pump Exhaust

In the lab area are two vacuum pumps, which are each loosely connected to the general exhaust fan by 6" flexible ducts. The exhaust port on the vacuum pump is about 1" in diameter. The 6" flexible duct ports should be closed off with a damper, so a smaller tube can be connected to each vacuum pump and exhausted into a close by fume hood. As a result, a lower amount of air will be removed from the space, which will result in fan energy and cooling/heating energy savings.

Vacuum pumps have dramatically improved energy efficiency recently. Efficiency of new models compared to those in use should be investigated to determine pay back time. This may be eligible for a utility incentive as well.

HEPA Recirculation Hood

The HEPA recirculation hood should be turned off when not in use. A staff member indicated that the recirculation hood runs 24 hours a day. A simple modification will need to be made to the hood, which would be to add a toggle switch so the unit can be turned on when in use.

Various Process Equipment Tools

There were a couple of large process equipment tools located in the lab and cleanroom areas, which were off during the study. However, a constant volume of exhaust was being drawn from the space through the tool. An interlocked damper in the exhaust duct, which opens and closes when the machine turns on and off will lower the amount of exhaust fan energy use and the conditioned air exhausted from the space. Further energy savings can be achieved by controlling the speed of the exhaust fan with a variable frequency drive (VFD) based on a static pressure sensor located in the exhaust duct.

Demand-Controlled Filtration

LBNL is currently investigating the feasibility to implement demand-controlled filtration in the facility. Based on particle counts in the ISO class 5 (class 100) cleanrooms, the data showed that there were less than 10 particles per cubic foot of air taken at the 0.3 micron level. The cleanrooms are essentially operating at achieving ISO class 4 (class 10), which is cleaner than required. Currently, the recirculation air handling system is setback during unoccupied hours, however further setback should be possible. Demand controlled filtration (i.e. controlling based upon particle counts) should verify that additional setback is possible. By actively monitoring the particle counts in the cleanrooms and adjusting the

airflow provided by the recirculation air handling system to retain cleanliness, a significant amount of energy can be saved.

Gowning Room Door Filters

The grilles in the doors located between the cleanroom gowning room and lab area space are unnecessary. This is unusual design. Removing the filters and covering up the void spaces should result in a reduction in airflow in the gowning room.

Lab Area Temperature Setpoint

The room temperature cooling setpoint is 68°F and seems low. During the warmer months a setpoint between 74 - 78°F is recommended by ASHRAE. By raising the setpoint temperature, more energy will be saved. However, if the processes in the area require 68°F then this measure will not be possible.

Fume Hood

The fume hood currently has a sign indicating that the fume hood is not working and that it should not be used. However, the fume hood exhaust fan, EF-12 is running. The exhaust fan should be turned off until the fume hood is repaired.

CLEANROOM ENERGY BENCHMARKING SITE REPORT

FACILITY J

NORTHERN CALIFORNIA

AUGUST 2004

SPONSORED BY:

LAWRENCE BERKELEY NATIONAL LABORATORY

PREPARED BY:

99 LINDEN STREET
OAKLAND, CA 94607
WWW.RUMSEYENGINEERS.COM

Copyright (c) 2004 Lawrence Berkeley National Laboratory. All rights reserved.

Acknowledgements

Special thanks to the entire team for their generous assistance and cooperation throughout the benchmarking process.

Project Team

Lawrence Berkeley National Laboratory	**Rumsey Engineers, Inc**
Bill Tschudi	Peter Rumsey
Dale Sartor	Achai Broner
	Larry Chu
	John Weale

Reproduction or distribution of the whole, or any part of the contents of this document without written permission of LBNL is prohibited. The document was prepared by LBNL for the exclusive use of its employees and its contractors. Neither LBNL nor any of its employees makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any data, information, method, product or process disclosed in this document, or represents that its use will not infringe any privately-owned rights, including but not, limited to, patents, trademarks, or copyrights.

Table of Contents

I. Executive Summary 1

II. Introduction 1

III. Review of Site Characteristics 2

- A. Site 2
- B. Fill Suite Cleanroom Design 3
- C. Fill Suite Hall/Component Staging Cleanroom Design 3
- D. Parts Preparation/Loading/Equipment Wash Cleanroom Design 4

IV. Site Energy Use Characteristics 6

- A. Cleanroom Power Consumption 6
- B. Electrical System Power Consumption 7

V. System Performance Metrics 7

VI. Site Observations Regarding Energy Efficiency 11

APPENDICES

A. Data Reports

B. Trended Data Graphs

C. Data Collection and Accuracy Notes

D. Measurement Methodology

I. EXECUTIVE SUMMARY

As part of the California Energy Commission Public Interest Energy Research (PIER) project, energy use at Facility J was monitored May 3 to May 12, 2004. Facility J, built in approximately 2002, is a facility that houses primarily office spaces, cleanrooms and non-clean supports areas.

This site report reviews the data collected by the monitoring team and presents a set of performance metrics as well as a complete set of trended data points for energy end uses for equipment supporting and located in the cleanrooms. Some of the most important metrics are summarized below in Table 1.

Table 1. Cleanroom Metric Results for Facility J

Metric Name	Metric Value
Class 100 Recirculation Fan Efficiency	2,184 cfm/kW
Class 100 Make Up Fan Efficiency [1]	1,492 cfm/kW
Class 1,000 Recirculation Fan Efficiency	1,496 cfm/kW
Class 1,000 Make Up Fan Efficiency [1]	1,492 cfm/kW
Class 100,000 Recirculation Fan Efficiency	983 cfm/kW

1. *Make up air for the class 100 and class 1,000 cleanrooms are supplied by the same air handler, therefore, the metrics are identical.*

The metrics for the HVAC systems at Facility J show that there are opportunities for energy efficiency improvement. The monitoring team observed a number of opportunities for potential energy savings at the facility. A summary of these observations follows and a more detailed discussion can be found in Section VI "Site Observations Regarding Energy Efficiency".

II. INTRODUCTION

Energy metrics were established that allow cleanroom owners to evaluate their energy efficiency performance and identify opportunities for improvements that reduce their overall operating costs. The project is administered by LBNL and funded through the California Energy Commission.

This Site Report summarizes the data collected and presents energy performance metrics with which the facility can evaluate the performance of its cleanrooms. First, the report reviews the site characteristics, noting design features of the mechanical plant and the cleanrooms monitored. Second, the energy use for the cleanrooms and mechanical equipment is broken down into components. Third, performance metrics recorded through the project are presented. Finally, key energy efficiency observations for the facility are noted. This is not intended to be a full energy audit, merely observations from the site measurement team. The data collected, trended graphs and methodology documentation are included among the appendices.

III. REVIEW OF SITE CHARACTERISTICS

A. Site

Facility J, located in Northern California, is a two-story 100,000 sf building that is approximately two years old. The building houses primarily cleanrooms, office spaces, and non-clean support areas, such as hallways and mechanical equipment rooms.

A 2,000 kW diesel generator provides backup power in the case of a utility failure. Backup power is supplied to the makeup and recirculation air handlers serving the cleanrooms, both chilled water plants, exhaust fans and lighting for life safety reasons, various process tools, and the fire alarm and telephone systems. Two uninterruptible power supply (UPS) systems are also utilized to provide emergency power. One UPS rated at 1000 kVA provides backup power to various process tools, life safety systems and automated controls systems. The second UPS rated at 15 kVA/10.5 kW provides emergency power to the makeup air handler serving the Class 100 area to keep the room pressurized to mitigate contamination during a utility power failure.

The environmental systems serving the cleanrooms run 8,760 hours a year in order to maintain conditions. The cleanroom spaces are conditioned by sixteen air handlers served by a chilled water plant, steam boiler plant, and steam humidifiers. A VAV (variable air volume) system consisting of one large rooftop air handler serves the office areas. The HVAC chilled water plant consists of two water-cooled chillers to provide chilled water to the make up air handlers (MUAH) and recirculation air handlers (RCU). The HVAC chilled water plant also serves process equipment; i.e. CDA compressor cooling, steam condensate coolers, clean steam generator, etc. Hot water is generated through a heat exchanger supplied with steam by the boiler plant. The hot water is distributed to the coils in the make up air handlers, and the VAV boxes serving the office areas. An additional water-cooled chiller plant is utilized for cooling of additional process manufacturing tools.

The cleanrooms chosen for monitoring are the Fill Suite Cleanroom (Class 100 - 885 sf), the Fill Suite Hall/Component Staging Cleanrooms (Class 1,000 - 740 sf), and the Parts Preparation/Loading/ Equipment Wash Cleanrooms (Class 100,000 - 1,805 sf).

Chilled water is produced by two water-cooled chillers connected by a common header. Normally, one chiller runs while the other stages on as necessary to accommodate the loads during warmer days. The chilled water system employs a primary-secondary loop pumping system. There are three constant speed pumps on the primary loop, and three variable-speed driven pumps on the secondary loop; two primary and two secondary pumps normally operate with one each on emergency backup. During the monitoring period, chilled water was supplied at 42.6 ± 0.5°F.

Process cooling water is supplied by two water-cooled chillers connected by a primary-secondary loop. There are three constant-speed pumps on the primary loop, and three variable-speed driven pumps on the secondary loop; two pumps on each loop normally run with the third on emergency backup. Process cooling water is used to cool the process equipment located in the cleanrooms.

There are two steam boilers, with one on emergency backup, used to generate hot water via a heat exchanger for use in the hot water coils of the make up air handlers. Hot water is also used for the VAV reheat coils serving the office and miscellaneous spaces. Hot water is distributed by one hot water pump with an additional pump serving as an emergency backup. A clean steam plant is also utilized for generating humidification for the air handlers, and for cleaning of the process equipment. There is a separate steam boiler system used to generate clean steam to serve the process tools and the humidifiers in some of the make up air handlers. Various cleanrooms are equipped with exhaust fans without scrubbers

.

B. Fill Suite Cleanroom Design

The Fill Suite Cleanroom as measured in this report is a total of 780 sf, including both primary and secondary (air return) areas. The return area consists of 120 sf. The cleanroom utilizes a pressurized plenum and is rated at class 100. HEPA ceiling coverage in the cleanroom is 100% as mandated by the FDA. The Fill Suite cleanroom has low sidewall returns.

The Fill Suite Cleanroom is served by one make up air handler (AH-5340), and two recirculation air handlers (AH-5380 and AH-5381). The make up air unit delivers its air to the intake plenums of the RCUs. The make up air handler is served with chilled water, heating hot water and clean steam for humidification. The RCUs are served with chilled water only.

Make Up Air Handler, AH-5340

The environmental specifications for the cleanrooms are 65.5°F ± 5.5°F (18°C ± 3°C) and 45% ± 15% relative humidity. During the monitoring period, the measured temperature for the Fill Suite Cleanroom was 66°F ± 0.4°F, and the average measured relative humidity was 48% with a fluctuation of ± 6%.

C. Fill Suite Hall/Component Staging Cleanroom Design

The Fill Suite Hall/Component Staging Cleanrooms are class 1,000 cleanrooms. The Fill Suite Hall and Component Staging areas have local protection with 100% HEPA filter coverage; these areas are served by a pressurized plenum. The gowning room and de-gowning areas have 20% and 28% HEPA coverage, respectively; these areas utilize ducted HEPAs. The total cleanroom area, including both primary and secondary areas, is 870 sf. A wall physically separates the two cleanrooms. The cleanrooms are served by a recirculation air handler (AH-5360) and make up air handler, AH-5340, which also provides makeup air to the Fill Suite Cleanroom mentioned above. The cleanroom return air is directed through low sidewall returns around the perimeter of each of the cleanrooms. The air handling unit is served with chilled water and

Air Handler, AH-5360

heating hot water. and steam for humidification.

Table 2. Class 1,000 Cleanroom Areas

Cleanroom	Primary Area (sf)	Secondary Area (sf)	Total Area (sf)
Fill Suite Hall	265	35	300
Component Staging	105	15	120
Gown	240	20	260
De-Gown	85	10	95
Common Area	175	0	175
Total	870	80	950

The design specifications for the cleanroom air conditions are 68°F ± 9°F (20°C ± 5°C) and 45% ± 15% relative humidity. During the monitoring period, the measured temperature was 66°F ± 2°F, and the measured relative humidity was 51% ± 8%. The cleanroom temperature and humidity sensors may need to be calibrated.

D. Parts Preparation/Loading/Equipment Wash Cleanroom Design

The Parts Preparation/Loading/Equipment Wash Cleanrooms are class 100,000 pressurized plenum cleanrooms. The filter coverage varies throughout the rooms; areas of the Parts Preparation and Loading Cleanrooms each have local protection where HEPA coverage is 100% and air is supplied to a pressurized plenum. The average HEPA filter coverage for the Pats Preparation and Loading Cleanrooms is 39% and 44%, respectively. The air is ducted to HEPA filters for the areas without local protection. The Equipment Wash room has 11% HEPA coverage. The total cleanroom area, including both primary and secondary areas, is 1,420 sf. Physical walls of separation divide the rooms in this cleanroom area. The cleanrooms are served by a single air handler (AH-5320), which provides both makeup air and recirculation air. The cleanroom return air is directed through low sidewall returns around the perimeter of each of the cleanrooms. The air handling unit is served with chilled water and heating hot water. A steam humidifier will be installed for this unit in the future.

Table 3. Class 100,000 Cleanroom Areas

Cleanroom	Primary Area (sf)	Secondary Area (sf)	Total Area (sf)
Part Preparation	700	85	785
Loading	435	30	465
Equipment Wash	285	20	305
Total	1,420	135	1,555

The design specifications for the cleanroom air conditions are 72.5°F ± 13.5°F (22.5°C ± 7.5°C) without any relative humidity requirements. During the monitoring period, the measured temperature was 66°F ± 2°F, and the measured relative humidity was 51% ± 8%.

Table 4. Summary of Measured Cleanroom Air Handling Parameters

Description		***Fill Suite***	***Fill Suite Hall/Component Staging***	***Parts Preparation/Loading /Equipment Wash [2]***
		Class 100	*Class 1,000*	*Class 100,000*
Primary Area	sf	660	870	1,420
Ceiling Height	ft	9	9	9
Total Make Up Air [1]	cfm	956	3,967	8,695
Total Make Up Fan Power [3]	kW	3.3		20.1
Total Recirculation Air [1]	cfm	61,732	20,866	11,038
Total Recirculation Fan Power	kW	28.3	13.9	20.1
Room Air Changes per Hour	ACH	531	160	52
HEPA Filter Ceiling Coverage	%	100	varies, see description above	varies, see description above
Average Ceiling Filter Velocity [4]	fpm	115	52	26

1. *Make Up and Recirculation Air is the air delivered, based on the balance report data.*
2. *Make Up and Recirculation air for the Parts Preparation/Loading/Equipment Wash Cleanrooms are delivered by the same unit, therefore the power consumption is identical.*
3. *Make Up air is delivered to the Fill Suite Cleanroom and the Fill Suite Hall/Component Staging Cleanrooms by the same air handler.*
4. *Filter velocity based on average filter flow and 6.8 sf (85%) effective filter area.*

IV. SITE ENERGY USE CHARACTERISTICS

A. Cleanroom Power Consumption

The energy consumption attributed to the cleanroom air handling systems, exhaust fans, process tools, and lighting are reported in Tables 5, 6 and 7. This breakdown of energy use by equipment helps identify the major loads.

Table 5. Fill Suite Cleanroom Power Consumption Breakdown

Description	Average Load (kW)	Average Efficiency
AIR HANDLING		
Make Up Fans	3.3	1,492 cfm/kW
Recirculation Fans	28.3	2,184 cfm/kW
EXHAUST FANS [1]	N/A	N/A
PROCESS POWER [2]	-	-
LIGHTS [3]	0.5	0.75 W/sf
TOTAL	32.1	

1. *Cleanroom does not utilize exhaust fans.*
2. *Process tools were not being used during monitoring period since cleanroom was not fully operational.*
3. *Lighting power calculated based on lighting power density of a representative cleanroom.*

Table 6. Fill Suite Hall/Component Staging Cleanroom Power Consumption Breakdown

Description	Average Load (kW)	Average Efficiency (cfm/kW)
AIR HANDLING		
Make Up Fans [1]	3.3	1,492 cfm/kW
Recirculation Fans	13.9	1,496 cfm/kW
EXHAUST FANS	1.8	3,403 cfm/kW
PROCESS POWER [2]	-	-
LIGHTS [3]	0.65	0.75 W/sf
TOTAL	19.7	

1. *The make up air handler also serves the Fill Suite Cleanroom, therefore, the values are identical to the table above.*
2. *Process tools were not being used during monitoring period since cleanroom was not fully operational.*
3. *Lighting power calculated based on lighting power density of a representative cleanroom.*

Table 7. Parts Preparation/Loading/Equipment Wash Cleanroom Power Consumption Breakdown

Description	Average Load (kW)	Average Efficiency (cfm/kW)
AIR HANDLING		
Make Up/ Recirculation Fans [1]	20.1	983 cfm/kW
EXHAUST FANS [2]	N/A	N/A
PROCESS POWER [3]	-	-
LIGHTS [4]	1.1	0.75 W/sf
TOTAL	21.2	

1. *Air handing unit provides both make up and recirculation air.*
2. *Cleanroom does not utilize exhaust fans.*
3. *Process tools were not being used during monitoring period since cleanroom was not fully operational.*
4. *Lighting power calculated based on lighting power density of a representative cleanroom.*

B. Electrical System Power Consumption

The table below shows the power consumption of the emergency generator standby power loss. The emergency generator constantly draws power to maintain the batteries and a specific temperature of the diesel engines that drive the generator. The generator has a capacity of 2,000 kW.

Table 8. Electrical System Power Consumption

Description	Average Load
ELECTRICAL SYSTEM	
Emergency Generator Standby Power	5.3 kW

V. SYSTEM PERFORMANCE METRICS

Metrics are ratios of important performance parameters that can characterize the effectiveness of a system or component. In order to gage the efficiency of the entire building system design and operation, this project tracks key metrics at different system levels. These metrics can be used to compare designs or determine areas with the most potential for improvement via retrofit or replacement.

For Facility J, the cleanroom HVAC components operate at a nearly constant level throughout the year. Therefore, these metrics are based on spot measurements. All of the metrics involving area are based on the primary cleanroom area, which is the area that passes certification, unless otherwise noted.

The Fill Suite Cleanroom is class 100. The pressurized plenum recirculation air handling efficiency was 2,184 cfm/kW. The recirculation air handling efficiency is poor, when compared to other pressurized plenum cleanrooms of various class ratings with an average efficiency of 5,152 cfm/kW. The cleanroom air change rate for this cleanroom is relatively high for this class of cleanroom. While the air change rate

is high within the set of Class 100 facilities benchmarked, it does fall within the range recommended by other sources (see http://cr.pennnet.com/Articles/Article_Display.cfm?Section=Archives&Subsection=Display&ARTICLE_ID=165797 for a copy of CleanRooms Magazine article discussing this issue). The benchmarking project has found that significantly lower air change rates are commonly used to provide the same class cleanroom environment, and with the good gowning protocol observe it is expected that a lower rate could be used without problem. A reduction to a level of 250 - 300 ACH should be investigated to further reduce energy usage.

Table 9. Cleanroom Metrics

Description		***Fill Suite Cleanroom***	***Fill Suite Hall/Component Staging Cleanrooms***	***Parts Preparation/Loading /Equipment Wash Cleanrooms*** [3]
		Class 100	*Class 1,000*	*Class 100,000*
MUAH Efficiency	cfm/kW	1,492		983
Make Up Air	cfm/sf	1.4	4.6	6.1
Make Up Fan Power Density [1]	W/sf	2.2		14.1
Recirculation Air Handler Efficiency	cfm/kW	2,184	1,496	983
Recirculation Air	cfm/sf	79.7	32.1	6.1
Recirculation Air Changes per Hour	ACH	624	160	52
Recirculation Fan Power Density [1]	W/sf	42.8	16.0	14.1
Lighting Power Density	W/sf	0.75	0.75	0.75
Process Tools Power Density [2]	W/sf	not measured	not measured	not measured

1. *Calculated as total kW load divided by the primary area of the cleanroom.*
2. *Process tool power load was not measured since the facility was not in production.*
3. *Air handlers for the Parts Preparation/Loading/ Equipment Wash Cleanrooms provide both make up and recirculation air, therefore the data is identical for the make up air handler and recirculation air handler.*

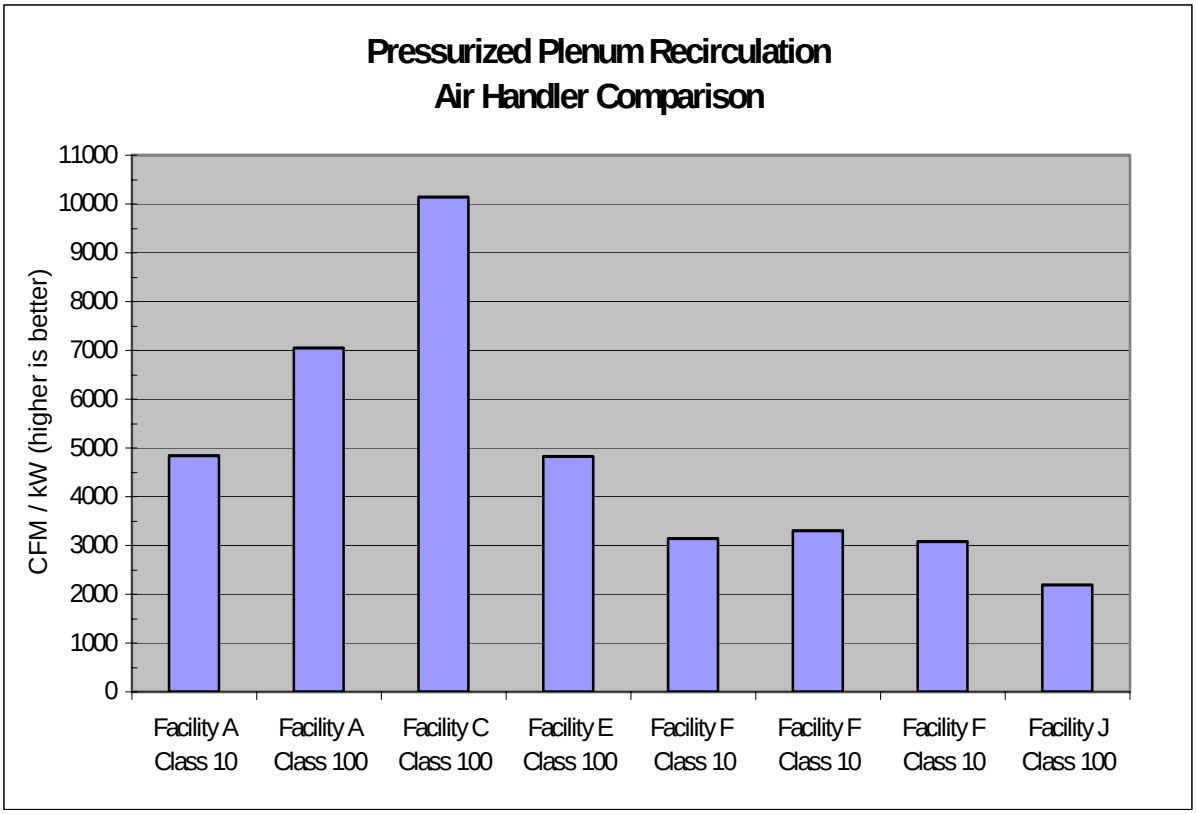

Figure 1. Pressurized Plenum Recirculation Air Handler Efficiencies of Measured Facilities (Various Class Ratings)

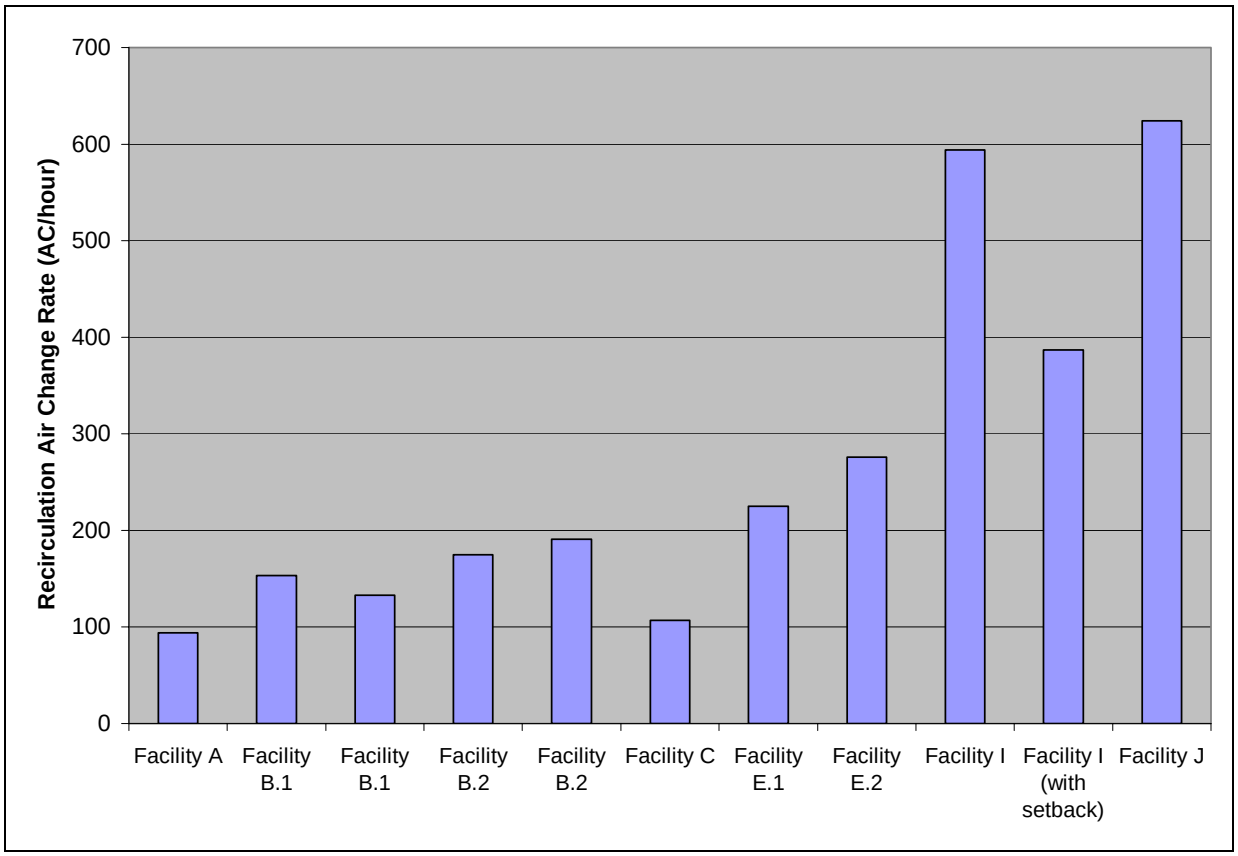

Figure 2. Recirculation Air Change Rate of Measured Class 100 Cleanrooms

VI. SITE OBSERVATIONS REGARDING ENERGY EFFICIENCY

Raise Chilled Water Supply Temperature

The 36°F low temperature loop could be reset to a higher temperature. Low temperature chilled water is inherently more energy intensive to produce due to the larger temperature delta, or 'lift,' the compressor is required to move heat through. On centrifugal compressors based chillers, an increase of one degree in the chilled water supply temperature improves the efficiency of the chiller by 1 to 2%. The chiller manufacturer should be consulted before radical modification of the supply temperature, but a change of a few degrees should allow for an improvement in efficiency without risk.

Condenser Water Temperature Reset

The chillers in the main plant may also benefit from a reduction in lift in two different ways. The first would be to ensure the condenser water reset is working properly and verifying with the factory that it cannot be taken lower. A reduction in the condenser water temperature of a couple degrees will achieve meaningful energy savings, however chiller stability can be compromised if it is taken too low, hence the need to verify with the chiller manufacturer before modifying this point.

Chilled Water Temperature Reset

The other way to improve chiller performance is to implement a chilled water temperature reset. When dehumidification is *not* required, there is typically a significant surplus of coil capacity. This allows the chilled water temperature to be increased, improving the chiller efficiency. There are a number of chilled water resets that have been shown to work. The simplest is to set the CHW temperature based on the outside air temperature, with the lowest temperature chilled water setpoint achieved during the highest outside air temperatures and during dehumidification. A more sophisticated approach is to poll all the chilled water valves on the loop for their position. If no valve is at 90% or more open, then maximum cooling is not required and the chilled water temperature can be increased. When a valve position exceeds 90%, it indicates a space is calling for additional cooling and the chilled water temperature can be reset down. The savings from a chilled water reset are maximized when the periods of dehumidification are minimized.

Operate Cooling Towers in Parallel

The lowest energy operation of a cooling tower is when heat is rejected without the use of any fan power, with air driven strictly by natural convective forces. To achieve this type of operation, the amount of surface area being wetted should be maximized. As many towers as possible should be operated simultaneously. The number of towers in operation is limited typically by the minimum flow that allows the tower nozzles to sufficiently achieve even tower media wetting.

Relative Humidity Control

The humidification setpoint is more restrictive than the specification requires. Loosening the humidification setpoint can reduce the hours of humidification/dehumidification in this very mild climate. Controlling the humidity levels in the supply air is very energy intensive; the fewer hours of control, the lower the overall energy cost.

The current humidification control algorithm appears to have stability problems. Since humidity control is so energy intensive, control instabilities can consume significant amounts of energy during unnecessary overshoot. The stability of the space conditions can also suffer. It is recommended that the pychrometric parameter 'Dewpoint Temperature' be used instead of RH for supply air humidity control loops. Control to RH tends to be prone to control problems due to the volatility of the RH reading – the RH of air changes quickly in response to both temperature changes and changes in the quantity of water in the air. The dewpoint changes only in response to changes in the quantity of water in the air, providing a more

direct measurement point for humidification loops. The space should still be controlled to RH, since RH is typically the parameter of interest in interactions with cleanroom activities and workers.

Steam-in-place Heat Recovery

The SIP, steam-in-place, process for cleaning tanks may offer some energy saving opportunities. During the SIP process, the tank is increased in temperature to a point higher than the hot water loop using clean steam. The tank is then cooled using compressed air, at an energy cost of about 200 watts per cfm, until it is cool enough to use chilled water in the jacket for final cooling. This process takes significant time, and also uses energy intensive compressed air just to reject heat to the room, where it then has to be removed through the chiller system. There is an opportunity to use the *hot water* loop for the initial stages of cooling. Using the hot water loop would reduce the use of compressed air as well as actually recover some of the tank's heat into the hot water loop for use elsewhere in the facility. The amount of heat recovered could be optimized by lowering the hot water loop temperature as much as is permissible.

Uninterruptible Power Supply (UPS) System Efficiency

Most uninterruptible power supply systems operate more efficiently when they are more heavily loaded. The efficiency of a battery UPS tends to decrease significantly when it is loaded at 30% or lower. See the figure below. Properly sizing a UPS system will minimize the energy losses. In addition, adding non-critical circuits to the UPS to increase its load may also be a strategy to decrease the overall power consumption of the facility.

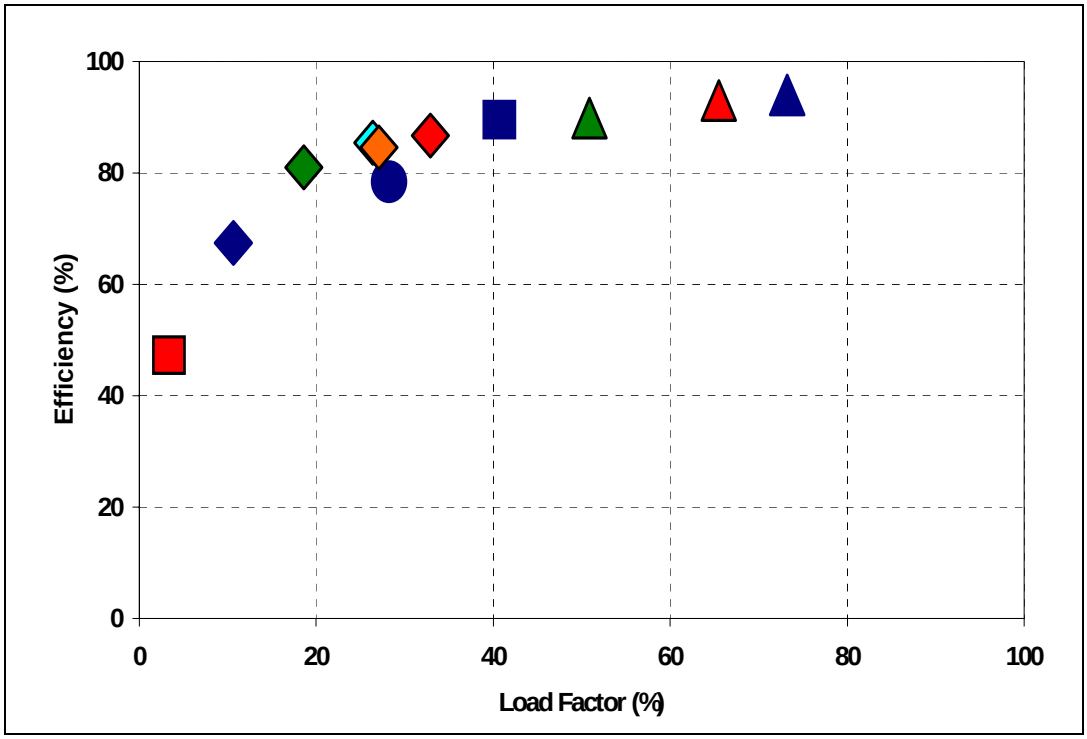

Figure 3. UPS Efficiency

CLEANROOM ENERGY BENCHMARKING AND EFFICIENCY TECHNOLOGY TRANSFER ASSISTANCE SITE REPORT

FACILITY I

SOUTHERN CALIFORNIA

JUNE 2004

SPONSORED BY:

LAWRENCE BERKELEY NATIONAL LABORATORY

PREPARED BY:
RUMSEY ENGINEERS, INC
99 LINDEN STREET
OAKLAND, CA 94607
WWW.RUMSEYENGINEERS.COM

Copyright (c) 2004 Lawrence Berkeley National Laboratory. All rights reserved.

Acknowledgements

Special thanks to the entire facilities staff for their generous assistance and cooperation throughout the benchmarking process.

Project Team

Lawrence Berkeley National Laboratory	**Rumsey Engineers, Inc**
Bill Tschudi	Peter Rumsey
Dale Sartor	Achai Broner
	Larry Chu
	John Weale

Reproduction or distribution of the whole, or any part of the contents of this document without written permission of LBNL is prohibited. The document was prepared by LBNL for the exclusive use of its employees and its contractors. Neither LBNL nor any of its employees makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any data, information, method, product or process disclosed in this document, or represents that its use will not infringe any privately-owned rights, including but not, limited to, patents, trademarks, or copyrights.

Table of Contents

I. Executive Summary 1

II. Introduction 2

III. Review of Site Characteristics 2

- A. Site 2
- B. MBE/Metal/Junction Formation Cleanroom Design 4
- C. Hybridization/Micro-Optics/ONM Array Production Assembly/Bonding Cleanroom Design 5

IV. Site Energy Use Characteristics 7

- A. Site Energy Use 7
- B. Cleanroom Power Consumption 7
- C. Major Mechanical and Electrical Systems Power Consumption 9
- D. Recirculation System Setback 9

V. System Performance Metrics 12

VI. Site Observations Regarding Energy Efficiency 14

APPENDICES

A. Data Reports

B. Trended Data Graphs

C. Data Collection and Accuracy Notes

D. Measurement Methodology

E. Cleanrooms Magazine Article – Air Change Rates

I. EXECUTIVE SUMMARY

As part of the California Energy Commission PIER study, energy use at Facility I was monitored March 15 to March 23, 2004. Facility I, built approximately two years ago, is a facility that houses primarily cleanrooms, office spaces, and non-clean laboratory and manufacturing areas.

This site report reviews the data collected by the monitoring team and presents a set of performance metrics as well as a complete set of trended data points for the end uses of energy for equipment supporting and located in the cleanrooms. Some of the most important metrics are summarized below in Tables 1 and 2.

Table 1. Cleanroom Metric Results for Facility I

Metric Name	Metric Value
Class 100 Recirculation Fan Efficiency	1,655 cfm/kW
Class 100 Make Up Fan Efficiency	995 cfm/kW
Class 10,000 Recirculation Fan Efficiency	1,756 cfm/kW
Class 10,000 Make Up Fan Efficiency	1,000 cfm/kW

Table 2. Chilled Water System Metric Results for Facility I

Metric Name	Metric Value
Chiller Efficiency	0.96 kW/ton
Chilled Water Pumps Efficiency	0.06 kW/ton
Total Chilled Water System Efficiency	1.06 kW/ton
Process Chiller Efficiency	1.52 KW/ton

These efficiency numbers are averages of 1-minute samples. Data was taken for the chiller over a period of 6 days (March 18 through March 23, 2004) when all the equipment was running and monitored simultaneously. Data for the process chiller was taken over a period of 3 days (March 18 through March 20, 2004). See Appendix B for charts of the trended data.

The metrics for the HVAC systems at Facility I show that there are opportunities for energy efficiency. The values in Table 2 indicate that the chillers have poor efficiency. However, air-cooled chillers are inherently not very energy efficient.

The monitoring team observed a number of opportunities for potential energy savings at the facility. A summary of these observations follows and a more detailed discussion can be found in Section VI "Site Observations Regarding Energy Efficiency".

II. INTRODUCTION

Energy metrics were established that allow cleanroom owners to evaluate their energy efficiency performance and identify opportunities for improvements that reduce their overall operating costs. The project is administered by LBNL and funded through the California Energy Commission.

With this report, Facility I is receiving the energy monitoring data collected at its facilities as a service provided by LBNL to participants in this project. This Site Report summarizes the data collected and presents energy performance metrics with which the facility can evaluate the performance of its cleanrooms. First, the report reviews the site characteristics, noting design features of the mechanical plant and the cleanrooms monitored. Second, the energy use for the cleanrooms and major mechanical equipment is broken down into major components. Third, performance metrics recorded through the project are presented. Finally, key energy efficiency observations for the facility will be noted. The data collected, trended graphs and methodology documentation are included among the appendices.

III. REVIEW OF SITE CHARACTERISTICS

A. Site

Facility I, located in Southern California, is a two-story 69,310 sf building that is approximately two years old and situated on a 151,200 sf lot. The building houses primarily cleanrooms, office spaces, and non-clean laboratory and manufacturing areas. The building is broken down as follows: cleanroom areas account for 15,430 sf (22%); office areas occupy 20,800 sf (30%); non-clean laboratory and manufacturing areas are 12,340 sf (18%); the second floor equipment area also referred to as the mezzanine, and the corridors/miscellaneous areas account for 20,750 sf (30%).

A 1000 kVA/800 kW diesel generator provides backup power in the case of a utility failure. Backup power is supplied to the recirculation air handlers, exhaust fans and lighting for life safety reasons, various process tools, and the fire alarm and telephone systems. Uninterruptible power supply (UPS) systems, the largest rated for 120 kW, are also utilized to provide emergency power to additional process tools located in the cleanrooms.

The majority of employees work from 8AM to 5PM, Monday through Friday, although the environmental systems serving the cleanrooms run 8,760 hours a year in order to maintain conditions. However, during the non-working periods, the make up and recirculation air handlers serving the cleanrooms are set to provide less airflow, and thus run at a lower fan speed. Since people, being the main source of contaminants in a cleanroom, are not present, less air is required to maintain cleanliness. As a result, fan energy and cooling/heating energy are lowered during the non-working hours.

The cleanroom spaces are conditioned by a chilled water plant, hot water boiler plant, and electric humidifiers. A VAV (variable air volume) system consisting of two large rooftop air handlers serves the non-clean laboratory areas. The office areas are conditioned by rooftop package units. Chilled water is provided by three air-cooled chillers to the make up air handlers (MUAH). Hot water is provided by three hot water boilers to the hot water coils in the make up air handlers, the reheat coils in the return ducts of the recirculation air units, and to the VAV boxes serving the laboratory areas. A separate air-cooled chiller is utilized to supply process cooling for the cleanroom and non-clean manufacturing/laboratories' process tools. Deionized (DI) water, process vacuum, and compressed air also are generated for use in the cleanrooms.

The cleanrooms chosen for monitoring are the MBE/Metal/Junction Formation Cleanrooms (3,660 sf), and the Hybridization/Micro-Optics/ONM Array Production/Bonding Assembly Cleanrooms (4,310 sf).

The MBE/Metal/Junction Formation Cleanrooms are class 100. Although there are walls of separation between these cleanrooms, they were grouped together for this study since they are served by the same make up and recirculation air handling units (RCU). The Hybridization/Micro-Optics/ONM Array Production Assembly/Bonding Cleanrooms are the only class 10,000 cleanrooms in this facility. The remainder of the cleanrooms are class 100 and account for 7,460 sf of this building.

Chiller

Chilled water is produced by three 170 ton air-cooled chillers connected by a common header. Normally, two chillers run with one on emergency backup. The chilled water system employs a primary only loop pumping system. There are three variable-speed driven pumps; two normally operate with one on emergency backup. The pumps supply the chilled water to the make up air handlers only. During the monitoring period, primary chilled water was supplied at 40.8 ± 2.7°F.

Over the chiller monitoring period from March 18, 2004 through March 23, 2004 the outside air temperature ranged from 53°F to 72°F (see Appendix B for trended data). During that time two chillers operated at a combined average load of 201 tons with an overall range from 73 to 328 tons.

There are three 1,530 MBH boilers, with one on emergency backup, used to generate hot water for use in the hot water coils of the make up air handlers. Hot water is also used for reheat coils located in the ductwork of the recirculation air handlers, and in the VAV reheat coils serving the laboratory and manufacturing spaces. Hot water is distributed by one hot water pump with an additional pump serving as an emergency backup.

Process cooling water is supplied by a 50 ton air-cooled chiller coupled to two variable-speed driven pumps and a storage tank. One pump normally runs with the other on emergency backup to cool the process tools located in the cleanrooms, and in the manufacturing and laboratory areas. The storage tank provides backup cooling in the event of an emergency. When there is a power failure, chilled water in the tank continues to be distributed by the pump connected to the backup generator to the tools without the operation of the process cooling chiller. Over the process cooling chiller monitoring period from March 18 through March 23, 2004, the process cooling water supply temperature was 62.9 ± 2.1°F.

Process Cooling Water Pumps

Compressed air and DI water are produced at the central location for use in the cleanrooms. There are several small, individual process vacuum pumps dispersed throughout the cleanrooms. The cleanrooms are equipped with exhaust fans; no air scrubbers are fitted. Corrosives, solvents and non-contaminated air are exhausted from various cleanrooms.

B. MBE/Metal/Junction Formation Cleanroom Design

The MBE/Metal/Junction Formation Cleanrooms as measured in this report is a total of 3,660 sf, including both primary and secondary (air return) areas. See the table below for a breakdown of the areas. The cleanrooms utilize ducted HEPAs and are rated at class 100. HEPA ceiling coverage in the cleanrooms is 100%. The MBE cleanroom has low sidewall returns. The Metal cleanroom utilizes a combination of a low sidewall return and service chase for the return air. The Junction Formation cleanroom is a raised floor cleanroom with a service chase for return air. The service chases are class 100,000 rated and house the tools used to support the production in the Metal and Junction Formation Cleanrooms.

Table 3. Class 100 Cleanroom Areas

Cleanroom	Primary Area (sf)	Secondary Area (sf)	Total Area (sf)
MBE	565	150	715
Metal	655	695	1,350
Junction Formation	1,085	510	1,595
Total	2,305	1,355	3,660

The MBE/Metal/Junction Formation Cleanrooms are served by one make up air handler (AHU-2), and two recirculation air handlers (RAH-3 and RAH-4). The make up unit delivers its air to the intake plenums of the RCUs. The make up air handler is served with chilled water and heating hot water. In the return ducts of the recirculation air handlers are hot water coils and electric humidifiers. The electric humidifiers are supplied with deionized water. There are 2 general exhaust fans that serve the cleanroom; they exhaust directly to the outdoors without a scrubber.

Recirculation Air Handler Unit, RAH-3

During non-working hours in the cleanrooms, the make up and recirculation air handler fans are turned down to supply a lower amount of airflow to a point where cleanliness continues to be maintained. Since people being the major source of contaminants in a cleanroom are not present, this allows for lower airflows in the cleanroom, thus lower fan and heating/cooling energy are consumed. During the monitoring period, the two RCUs were operating at a combined power of 124.1 kW during normal working hours. During the turndown mode, the RCUs were consuming about 72% less power (34.2 kW) than in the normal operating mode. The MUAH airflow was not turned down during the monitoring period due to a failed VFD.

The condition specifications for the cleanrooms are 69°F ± 2°F and 45% ± 10% relative humidity. During the monitoring period, the measured temperature for the MBE Cleanroom was 66°F ± 2°F, and the average measured relative humidity was 48% with a fluctuation of ± 6%; the measured temperature for the Junction Formation Cleanroom was 67°F ± 1°F, and the average measured relative humidity was 48% with a fluctuation of ± 4%. The cleanroom temperature and humidity sensors may need to be calibrated.

C. Hybridization/Micro-Optics/ONM Array Production Assembly/Bonding Cleanroom Design

Make Up Air Handler, AHU-3

The Hybridization/MicroOptics/ ONM Array Production Assembly/ Bonding Cleanrooms are class 10,000 ducted HEPA cleanrooms with 25% filter coverage. The total cleanroom area, including both primary and secondary areas, is 4,310 sf. Most of the investigated cleanrooms are open to one another, except for the Bonding Cleanroom which has its own physical walls of separation. Although, this cleanroom does share a common low sidewall return air space with the other cleanrooms. The class 10,000 cleanrooms are served by a single make up air handler (AHU-3) and one recirculation air handler (RAH-6). The make up unit delivers its air to the intake plenum of the RCU. The cleanroom return air is directed through low sidewall returns around the perimeter of each of the cleanrooms.

Table 4. Class 10,000 Cleanroom Areas

Cleanroom	Primary Area (sf)	Secondary Area (sf)	Total Area (sf)
Hybridization	1,500	190	1,690
Micro-Optics	1,120	145	1,265
ONM Array Production Assembly	645	120	765
Bonding	280	70	350
Cleanroom Entry	205	35	240
Total	3,750	560	4,310

The make up air handler is served with chilled water, and heating hot water. In the return ducts of the recirculation air handler are hot water coils and electric humidifiers. The electric humidifiers are supplied with deionized water.

During non-working hours in the cleanrooms, the make up and recirculation air handler fans are turned down to supply a lower amount of airflow to a point where cleanliness continues to be maintained. Since people being the major source of contaminants in a cleanroom are not present, this allows for lower airflows in the cleanroom, thus lower fan and heating/cooling energy are consumed. During the monitoring period, the RCU and the MUAH were operating at a 22.9 kW and 13.7 kW, respectively during normal working hours. During the turndown mode, the RCU was consuming about 62% less power (8.7 kW) than in the normal operating mode. The MUAH was consuming 81% less power (2.6 kW) during the turndown mode.

The design specifications for the cleanroom air conditions are 69°F ± 2°F and 45% ± 15% relative humidity. During the monitoring period, the measured temperature was 66°F ± 2°F, and the measured relative humidity was 51% ± 8%. The cleanroom temperature and humidity sensors may need to be calibrated.

Table 5. Summary of Measured Cleanroom Air Handling Parameters

Description		***MBE/Metal/ Junction Formation***	***Hybridization/Micro-Optics/ONM Array Production Assembly/ Bonding***
		Class 100	*Class 10,000*
Primary Area	sf	2,305	2,570
Ceiling Height	ft	9	9
Total Make Up Air [1]	cfm	40,000	13,750
Total Make Up Fan Power [2]	kW	40.2	13.7
Total Recirculation Air [1]	cfm	205,330	40,260
Total Recirculation Fan Power [2]	kW	124.1	22.9
Room Air Changes per Hour	ACH	594	72
HEPA Filter Ceiling Coverage	%	100	14
Average Ceiling Filter Velocity [3]	fpm	105	90

1. *Make Up and Recirculation Air is the air delivered, based on the balance report data.*
2. *Make Up and Recirculation fan power reported for cleanrooms in normal operating mode.*
3. *Filter velocity based on average filter flow and 6.8 sf (85%) effective filter area.*

IV. SITE ENERGY USE CHARACTERISTICS

A. Site Energy Use

Facility I paid over $1 million in calendar year 2003 for energy use; Table 6 gives a breakdown. Table 7 calculates two key values to use in comparing the facility to other facilities with similar operations. Facility I pays average rates of $0.121/kWh for electricity and $0.64/therm for natural gas.

Table 6. Annual Energy Use

Annual Electricity Usage (MWh/yr)	Annual Electricity Cost ($/yr)	Annual Natural Gas Usage (therms/yr)	Annual Natural Gas Cost ($/yr)	Annual Total Cost ($/yr)
8,260	999,900	119,900	77,000	1,076,900

Source: Power bills for the year 2003.

Table 7. Annual Energy Utilization Intensity (EUI) and Energy Cost per Square Foot

Area (sf)	Energy Utilization Intensity (kWh/sf·yr)	Annual Energy Cost per Square Foot ($/sf·yr)
69,310	170	15.5

Energy from natural gas has been converted to kWh for the EUI calculation.

B. Cleanroom Power Consumption

The energy consumption attributed to the cleanroom air handling systems, exhaust fans, process tools, and lighting are reported in Tables 8 and 9. This breakdown of energy use by equipment helps identify the major loads.

This cleanroom has a relatively low process load and typically high HVAC power requirements. The HVAC power usage actually exceeds the process power usage by a factor of 6 for the class 100 cleanrooms, and by a factor of 3 for the class 10,000 cleanrooms.

Table 8. MBE/Metal/Junction Formation Cleanroom Power Consumption Breakdown

Description	Average Load (kW)	Average Efficiency
AIR HANDLING [1]		
Make Up Fan	40.2	995 cfm/kW
Recirculation Fans	124.1	1,655 cfm/kW
EXHAUST FANS		
EF-8A & 8B (Strobic Air)	23.3	888 cfm/kW
EF-13	1.8	1,944 cfm/kW
EF-15	0.4	1,579 cfm/kW
PROCESS	20.3	N/A
LIGHTS	4.2	N/A

1. *Make Up and Recirculation Fan power reported for cleanrooms in normal operating mode.*

Table 9. Hybridization/Micro-Optics/ONM Array Production Assembly/Bonding Cleanroom Power Consumption Breakdown

Description	Average Load (kW)	Average Efficiency (cfm/kW)
AIR HANDLING [1]		
Make Up Fan	13.7	1,000 cfm/kW
Recirculation Fan	22.9	1,756 cfm/kW
EXHAUST FANS		
EF-18	0.47	1,702 cfm/kW
PROCESS	6.9	N/A
LIGHTS	6.1	N/A

1. *Make Up and Recirculation fan power reported for cleanrooms in normal operating mode.*

C. Major Mechanical and Electrical Systems Power Consumption

The table below shows the power consumption of the chiller system for space cooling, process cooling chiller system, the emergency generator standby, and the UPS. The 1000kVA/800 kW emergency generator constantly draws power to maintain the oil temperature in the motors that drive the generator. The UPS has a capacity of 120 kW and was loaded at 36%.

Table 10. Energy Use by Major Components

Description	Average Load (kW)	Average Efficiency
COOLING SYSTEMS		
Chillers	195.5	0.96 kW/ton
Chilled Water Pumps	12.4	0.06 kW/ton
Process Cooling Chiller	21.8	1.52 kW/ton
Process Cooling Water Pump	7.8	-
ELECTRICAL SYSTEMS		
Emergency Generator Standby Power	1.73	N/A
UPS		
Input Power	43.3	86%
Output Power	37.3	

D. Recirculation Air System Setback

The recirculation air handling system uses an innovative reset to save power. When the cleanroom is not in use, there are many fewer sources of particles in the space. As evidenced by the facilities high gowning protocols, the site designer and operators are well aware that the human operators are the primary source of containments in the cleanroom space. The design intent of recirculation airflow is to continuously sweep particles from the space and remove them via the ceiling HEPA filters. With the recognition of people in the spaces as the primary source of particles, it is natural to connect the recirculation airflow quantities with the number of people, particle sources, in the space. Facility I has made this connection and implemented an off-shift airflow setback with excellent results. The following figure shows the fan power of a recirculation unit for a Friday through Monday period.

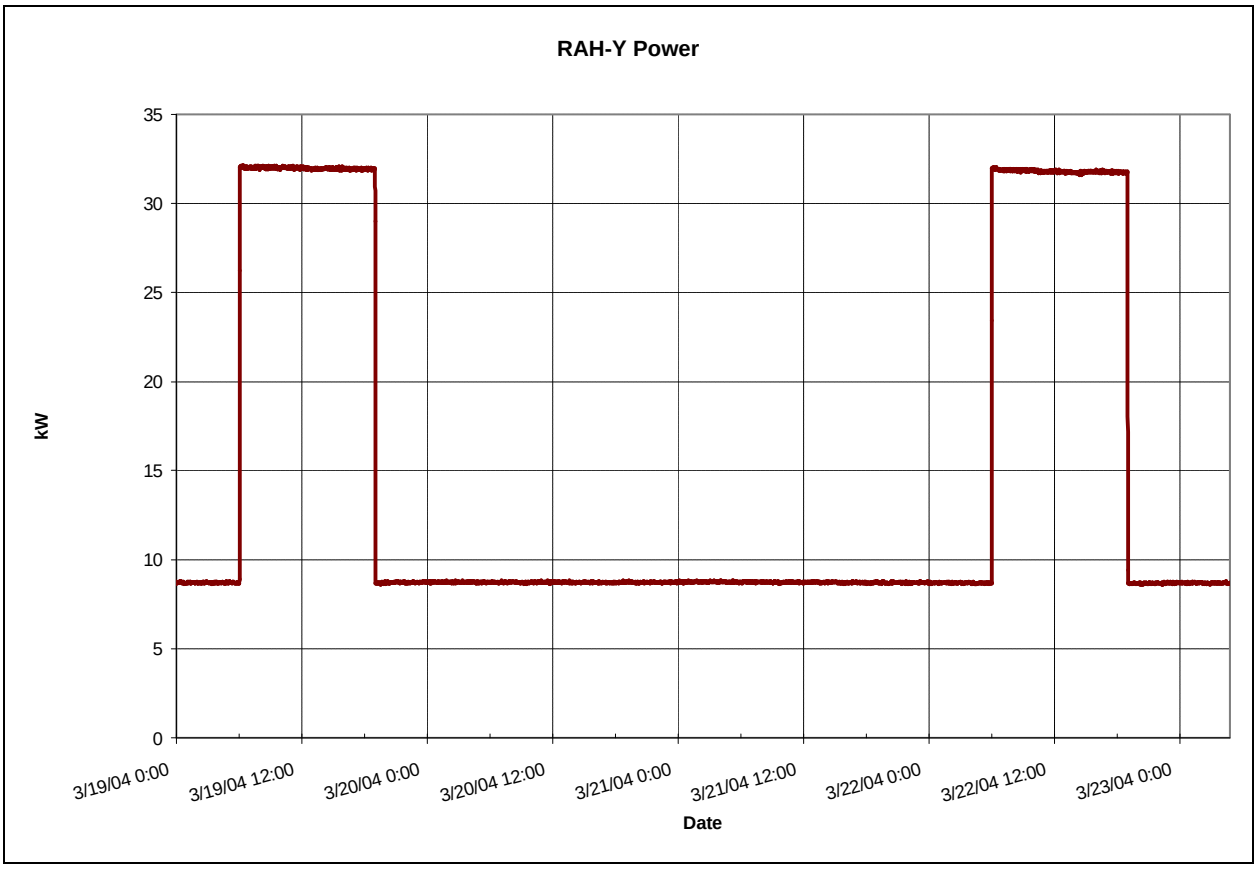

Figure 1. Recirculation Air Handler Power (One Minute True RMS Power Log)

While the airflow is setback less than 40% from the full occupied airflow, the actual measured power savings exceed 70%. Note that the power savings are measured values and do not follow the fan power 'laws' exactly. The power reduction results in significant real world energy savings, as summarized in the table below. This very simple energy efficiency measure is achieving significant savings with no impact on the facility operation.

Table 11. Recirculation Air System – Setback Metrics and Savings

Description		***MBE/Metal/ Junction Formation Cleanrooms***	***Hybridization/Micro-Optics/ONM Array Production Assembly/ Bonding Cleanrooms***
		Class 100	*Class 10,000*
Recirculation Air Power, Setback	kW	34.2	8.7
Recirculation Air Handler Volume, Setback [1]	cfm	128,000	28,000
Recirculation Air Setback Efficiency	cfm/kW	3,740	3,200
Recirculation Air Changes per Hour, Setback	ACH	371	50
Sitewide Savings			
Recirculation Air Annual Energy Savings [2]	kWh	1,250,000	
Recirculation Air Annual Cost Savings [2]	$/yr	138,000	

1. *Estimated using fan laws to scale flow and measured power data.*
2. *Extrapolated out to full site to include 3 identical but unmeasured recirculation units, $0.11/kWh average assumed.*

V. SYSTEM PERFORMANCE METRICS

Metrics are ratios of important performance parameters that can characterize the effectiveness of a system or component. In order to gage the efficiency of the entire building system design and operation, this project tracks key metrics at different system levels. These metrics can be used to compare designs or determine areas with the most potential for improvement via retrofit or replacement.

MBE/Metal/Junction Formation Cleanrooms & Hybridization/Micro-Optics/ONM Array Production Assembly/Bonding Cleanrooms

For Facility I, the cleanroom HVAC components operate at a nearly constant level throughout the year. Therefore, these metrics are based on spot measurements. All of the metrics involving area are based on the primary cleanroom area, which is the area that passes certification, unless otherwise noted.

The MBE/Metal/Junction Formation Cleanrooms are class 100. The ducted HEPA recirculation air handling efficiency was 1,655 cfm/kW. The recirculation air handling efficiency is average, when compared to other class 100, ducted HEPA cleanroom facilities with efficiencies ranging from 1,090 – 2,210 cfm/kW.

The class 10,000 cleanrooms also utilize a ducted HEPA design. The recirculation air handler efficiency is slightly better when compared to another class 10,000 ducted HEPA cleanroom facility with an efficiency of 1,635 cfm/sf. However, the RCU efficiency at this site is average when compared to the recirculation air handler efficiency between all tested facilities of various class ratings (see Figure 2 below).

The make up air handler efficiency for the class 100 and class 10,000 cleanrooms was 995 cfm/kW and 1,000 cfm/kW, respectively. The efficiency of the two make up air handlers is average when compared to the other tested facilities of various class ratings ranging from class 10 to 10,000. Make up air handler efficiencies at these tested facilities ranged from 537 to 1,797 cfm/kW.

Table 12. Cleanroom Metrics

Description		***MBE/Metal/ Junction Formation Cleanrooms***	***Hybridization/Micro-Optics/ONM Array Production Assembly/ Bonding Cleanrooms***
		Class 100	*Class 10,000*
MUAH Efficiency	cfm/kW	995	1,000
Make Up Air	cfm/sf	17.4	3.7
Make Up Fan Power Density [1]	W/sf	17.4	3.7
Recirculation Air Handler Efficiency	cfm/kW	1,655	1,756
Recirculation Air	cfm/sf	89.1	10.7
Recirculation Air Changes per Hour	ACH	594	72
Recirculation Fan Power Density [1]	W/sf	53.9	6.1
Lighting Power Density [2]	W/sf	1.2	1.6
Process Tools Power Density [2]	W/sf	5.9	1.8

1. *Calculated as total kW load divided by the primary area of the cleanroom.*
2. *Calculated as total kW load divided by the combined area of the cleanroom and the support room (secondary area) that contains the lighting and process tools.*

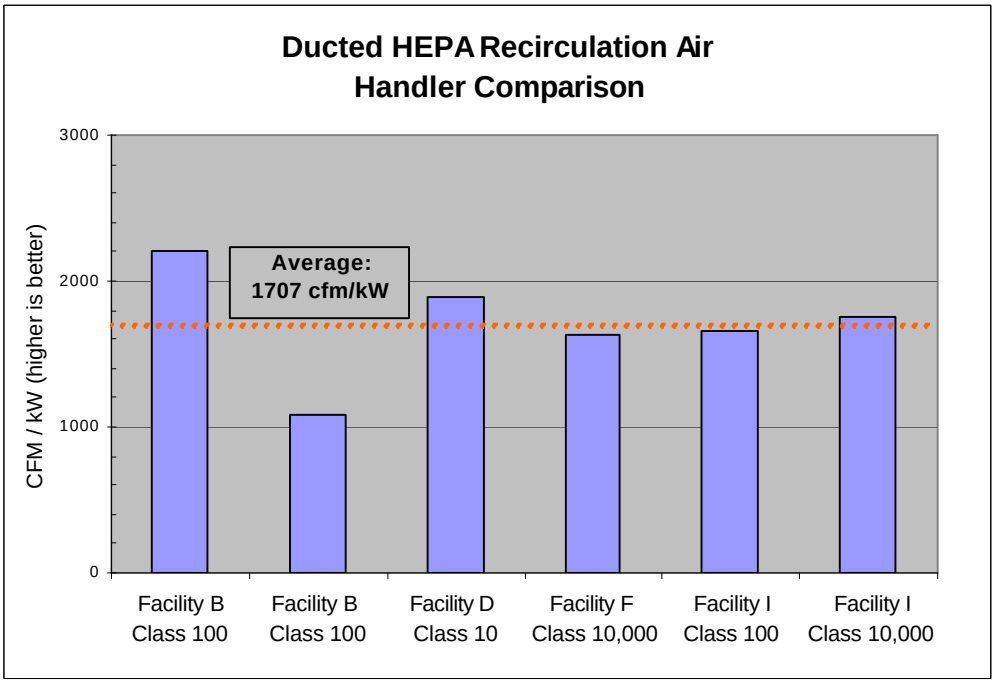

Figure 2. Ducted HEPA Recirculation Air Handler Efficiencies of Measured Facilities (Various Class Ratings)

Chilled Water Systems

Metrics of "kW/ton" are based on the total average equipment power and the average operating tonnage of the chilled water plant. These figures are useful for making comparisons between facilities, but more substantial information is expressed in the metric plots in Appendix B that reflect "kW/ton" performance at a sampling frequency of one minute over the course of a week. This type of information can be used to diagnose operational problems as well as evaluate the overall design performance.

Table 13. Chiller Efficiencies

Component	***Metric***
Chillers	0.96 kW/ton
Chilled Water Pumps	0.06 kW/ton
Chilled Water System	1.03 kW/ton
Cooling Load Density [1]	76.8 sf/ton
Process Cooling Chiller	1.52 kW/ton

1. *Cooling Load Density is the total conditioned area of the building served by the central plant, divided by the average plant tonnage.*

VI. SITE OBSERVATIONS REGARDING ENERGY EFFICIENCY

There are a number of potential areas for energy savings at Facility I. This section includes a general description of the most significant opportunities observed by the monitoring team.

Target the Humidity Control Sequences for Energy Reduction

During the measurement period it was observed that the air handling units were cycling frequently between humidification and dehumidification. This is a very inefficient mode of operation. During dehumidification, the air must be sub cooled, that is cooled below the sensible temperature at which it is to be delivered, resulting in the need for a reheat coil to raise the air to the supply temperature. The resulting energy use is two fold – extra energy is used to cool the air enough to dehumidify it, and then additional extra energy is used to increase the air after the dehumidification process is complete.

The majority of cleanroom spaces are controlled to much tighter humidity requirements (43% +/- 2%) than originally specified (up to 50% +/- 10%). The frequent, and energy intensive, cycling between dehumidification and humidification indicates the control loops would benefit from tuning to reduce the continuous oscillation. The control is made significantly more difficult by the current tight humidity control setpoints. Dehumidification requires that the air be "sub-cooled," meaning the air is cooled to a point that heating energy must be used to reheat it – a waste of both cooling and heating energy if the dehumidification is not actually required to meet the space conditions.

The chilled water humidity control setpoints should be modified to match the control ranges called out in the design documentation (drawing M5.8). Wider humidity control bands will significantly reduce the dehumidification and reheat energy requirements.

Optimize Chiller Operation with a Chilled Water Reset

Once dehumidification is brought under control, a chilled water reset can be implemented to decrease chiller energy use. The chiller will operate more efficiently when supplying a higher temperature chilled water. A higher temperature of chilled water, about 46°F rather than 42°F, should be capable of meeting the facility's load when the system is not dehumidifying. For internal load dominated plants such as this, it is recommended that the chilled water reset be based upon valve position rather than outdoor air temperature. This works by polling all the chilled water valves for there percent open value. If all the chilled water valves are less than 90% open, it is a direct indication that less cooling is required to meet the space loads. The chilled water supply temperature is then slowly increased until a least one valve is 90% open. When a valve position exceeds 90%, the chilled water temperature setpoint is reduced to meet the additional demand. Also, when dehumidification is called for the chilled water temperature must be set down to around 42°F to provide the subcooling required.

Implement Free Cooling for Process Chilled Water

Currently a 50 ton air-cooled chiller is used to supply a process chilled water loop. Due to the nature of the chiller and the low load on the chiller, the process chilled water chiller is operating at very poor efficiency – with the majority of operation in the range of 1.5 kW/ton to 2.5 kW/ton (Title 24 minimum performance is 1.25 kW/ton for an air cooled chiller and 0.8 kW/ton for a water cooled chiller).

The process chilled water supply is maintained at a temperature of about 63°F. This temperature could be

easily supplied for a significant portion of the year via an evaporative cooling tower at an energy cost of less than 0.2 kW/ton. Make up water for the tower could be supplied for the evaporative cooling tower from the RO reject water stream, which would provide from 2,500 – 4000 gallons per day (dependent on RO water usage). Use of RO reject water for cooling tower makeup has been successfully adopted at other cleanroom sites in arid climates, such as Phoenix, AZ, and is highly recommended. Use of free cooling would save from 100,000kWh to 200,000 kWh per year.

Reduce Cleanroom Air Change Rate

The cleanroom design air change rates are relatively high for this class of cleanroom, in some cases over 500 ACH for class 100 space. While the air change rate is high within the set of Class 100 facilities benchmarked, it does fall within the range recommended by other sources (see http://cr.pennnet.com/Articles/Article_Display.cfm?Section=Archives&Subsection=Display&ARTICLE_ID=165797 or Appendix for a copy of CleanRooms Magazine article discussing this issue). The benchmarking project has found that significantly lower air change rates are commonly used to provide the same class cleanroom environment, and with the good gowning protocol observe it is expected that a lower rate could be used without problem. This has been recognized by site personnel who have already reduced the air change rates somewhat. A reduction to a level of 250 - 300 ACH should be investigated to further reduce energy usage. The setback is currently operating at an air change rate equivalent to the normal air change rate of many operation Class 100 cleanrooms; it too offers an opportunity for reduction.

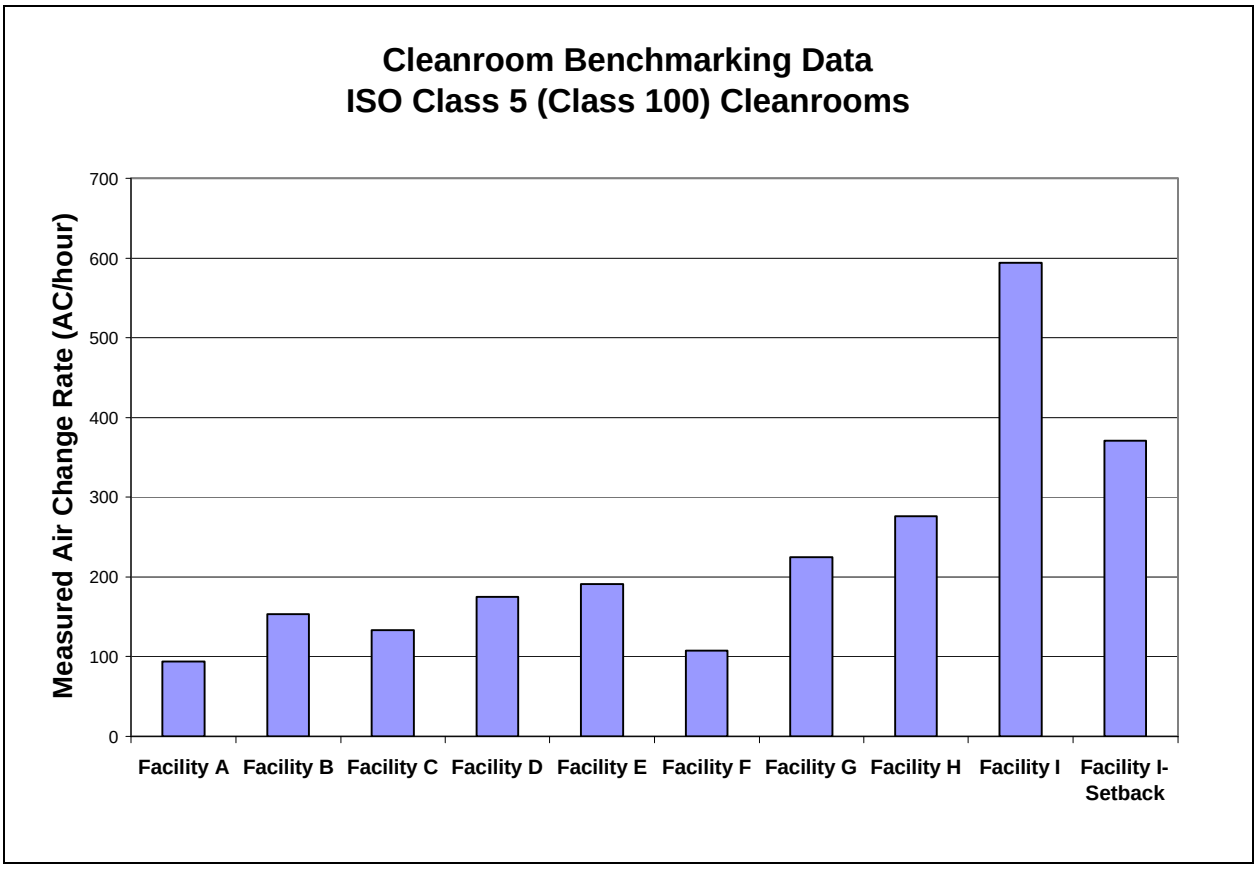

Unoccupied Recirculation Flow Setback

The current recirculation air handler control is very energy efficient, with nighttime recirculation volumes being reduced. For RAH-4, A reduction in fan speed of about 30% at night, when the space is unoccupied, yields a measured power consumption reduction of 72% – recirculation air handler power drops from 64 kW to less than 18 kW. When all the recirculation units are considered, this is obviously a significant savings. The magnitude of savings is indicative of the power reductions that occur from small reductions in the airflow volume. Implementing a more active and aggressive control scheme for the fan power setback would increase the power savings further.

One active control methodology is to install occupancy sensor based lighting controls in the cleanroom space. By means of current transducers, monitor the lighting circuits to detect when the cleanroom lights are off. Since the cleanroom spaces have little to no natural lighting, the controls can assume that when the lights are off the space is unoccupied and recirculation rates can be reduced. As demonstrated by the current setback operation, a small setback is all that is required to achieve a large energy savings. A setback of 30% achieves the majority of energy savings and presents little threat of space hygiene problems even if the system is setback during an occupied period inadvertently.

Appendix A

Data Reports

Site Data

Customer Name: ***Facility I***

Site Name: ***Building I***

Customer

Address:	anon.	Contact:	anon.
City:	anon.	Contact Phone:	anon.
Zip:	anon.	First Contact:	anon.

Site

Address:	anon.	Contact:	anon.
City:	anon.	Contact Phone:	anon.
Service Territory	SCE		

Total Facility Area	69310	sf	Year Built:	2002
			Annual Hours Use:	8760
Total Cleanroom Area (Class 3 thru 8 only):	15430	sf	Customer Owned	☐
Class 3 - 4 (1 - 10) Area	-99	sf	Corporate Payback	-77
Class 5 (100) Area:	11120	sf	Self Evaluation	-77
Class 6 (1000) Area:	-99	sf	Industry Type:	Other
Class 7 (10,000) Area	4310	sf		
Class 8 (100,000) Area	-77	sf		
Mini-environment Area	-77	sf		
Support Clean Area	-77	sf		
Support NonClean Area	-77	sf		
Sub Cleanroom Area	-77	sf		

Utility Billing

Annual Electric Use:	8260000	kWh/yr
Annual Electic Cost:	999900	$/yr
Average Electric Rate	0.121	$/kWh
Annual Fuel Use:	11900	Therms/yr
Annual Fuel Cost:	77000	$/yr
Peak Power:	1706	kW
Avg Power Factor	-77	

Billing Notes

Cleanroom Data

Customer Name: ***Facility I***
Site Name: ***Building I***
Cleanroom Name: ***MBE/Metal/Junction Formation - Class 100***

Class:	Class 100	HEPA Filter Coverage	100	%	Annual Hours Use:	8760	
ISO Class:	ISO Class 5	Filter Effective Area	85	%	Heat Recovery	☐	
Fan System Type	Recirculation AHU	Ceiling Filter Efficiency	-77		Raised Floor:	☑	
					Exhaust Fan(s):	☑	
Building Area:	69,310 sf						
Primary Cleanroom Area	2,305 sf	Primary Cleanroom Ceiling Heigh	9 ft				
Secondary Cleanroom Area	1,355 sf	Secondary Cleanroom Ceiling Heigh	-77 ft				

Power

Monitoring Start Date: 3/15/2004
Monitoring End Date: 3/23/2004

Design				Measured				
	Unit	Value	Source	Value	Max	Min	Source	Accuracy (+/-)
Lighting:	kW	-77		4.2	-77	-77	Rumsey Eng	5%
Process:	kW	-77		20.3	-77	-77	Rumsey Eng	5%
Other:	kW	-77		-77	-77	-77		

Space Conditions

Design					Measured				
	Unit	Value	+/-	Source	Value	Max	Min	Source	Accuracy (+/-)
Temperature:	F	69	2	Submittals	66.8	67.7	65.9	Rumsey Eng	5%
Humidity:	%RH	45	10	Submittals	47.6	51.2	43.9	Rumsey Eng	10%
Ceiling Velocity	fpm	-77	-77		105	-77	-77	Calculation	20%
Pressurization:	in wg	-77	-77		0.14	-77	-77	Rumsey Eng	10%

Cleanroom Description

This cleanroom is served by one makeup air handler and two recirculation air handlers.

MUAH Data

Customer Name: ***Facility I***
Site Name: ***Building I***
Cleanroom Name: ***MBE/Metal/Junction Formation - Class 100***

MUAH Name: ***AHU-2***

Number of Units: 1

Design

	Unit	Value	Source
Air Flow:	cfm	40000	Drawings
MUAH Fan Power:	kW	-77	
Efficiency:	cfm/kW	-99	
Total Air Flow:	cfm	40000	Drawings
Total Power:	kW	-99.0	
VFD Speed:	Hz	-77	
Supply Air Temperature	°F	-77	
Supply Setpoint	°F	47	Drawings
Supply RH:	%	-77	
RH Setpoint:	%	-77	
Fan Pressure Rise:	in	-77	
Filter Pressure Drop:	in	-77	
Sensible Cooling Load	tons	-77	
Coil Face Velocity	fpm	-77	

Measured

Average	Max	Min	Source	Accuracy (+/-)
995			Calculation	20%
40.2			Rumsey Eng	5%
25				
995				
40.2				
60			Rumsey Eng	5%
-77	-77	-77		
-77				
-77	-77	-77		
-77				
-77				
-77				
-77	-77	-77		
-77				

MUAH Description

Fan running at full speed due to burned out variable frequency drive.

RCU Data

Customer Name: ***Facility I***
Site Name: ***Building I***
Cleanroom Name: ***MBE/Metal/Junction Formation - Class 100***

RCU Name: ***RAH-3 & RAH-4***

Number of Units: 2

Design

	Unit	Value	Source
Air Flow:	cfm	26000	Drawings
RCU Fan Power:	kW	-77	
Efficiency:	cfm/kW	-99	
VFD Speed:	Hz	-77	
Supply Air Temperature	°F	-77	
Return Air Temperature	°F	-77	
Supply Setpoint	°F	-77	
RH Setpoint:	%	-77	
Supply RH:	%	-77	
Fan Pressure Rise:	in w.g.	-77	
Filter Pressure Drop:	in w.g.	-77	
Sensible Cooling Load	tons	-77	
Coil Face Velocity	fpm	-77	

Measured

Average	Max	Min	Source	Accuracy (+/-)
205330			Balance Report	20%
124.1			Rumsey Eng	5%
1655				
48			Rumsey Eng	5%
66	68	61	Rumsey Eng	5%
67	70	65	Rumsey Eng	2%
-77				
43			EMS/BMS	
-77	-77	-77		
2.2			Rumsey Eng	10%
0.7			Rumsey Eng	10%
-77				
-77				

RCU Description

Exhaust Data

Customer Name: ***Facility I***
Site Name: ***Building I***
Cleanroom Name: ***MBE/Metal/Junction Formation - Class 100***

Exhaust Name: ***EF-8A & 8B***

Number of Units: 2

Design

	Unit	Value	Source
Exhaust Flow:	cfm	20700	Drawings
Fan Power:	kW	-77	
Efficiency:	cfm/kW	-99	
Total Exhaust Flow	cfm	41400	
Total Power:	kW	-99.0	
Fan Pressure Rise:	in	3	Drawings

Measured

Average	Source	Accuracy (+/-)
20700	Calculation	20%
11.6	Rumsey Eng	5%
1784		
41400		
23.2		
-77		0%

Exhaust Description

Exhaust Data

Customer Name: ***Facility I***
Site Name: ***Building I***
Cleanroom Name: ***MBE/Metal/Junction Formation - Class 100***

Exhaust Name: ***EF-13***

Number of Units: 1

Design

	Unit	Value	Source
Exhaust Flow:	cfm	3500	Drawings
Fan Power:	kW	-77	
Efficiency:	cfm/kW	-99	
Total Exhaust Flow	cfm	3500	
Total Power:	kW	-99.0	
Fan Pressure Rise:	in	3	Drawings

Measured

Average	Source	Accuracy (+/-)
3500	Calculation	20%
1.8	Rumsey Eng	5%
1944		
3500		
1.8		
-77		0%

Exhaust Description

Exhaust Data

Customer Name: ***Facility I***
Site Name: ***Building I***
Cleanroom Name: ***MBE/Metal/Junction Formation - Class 100***

Exhaust Name: ***EF-15***

Number of Units: 1

Design

	Unit	Value	Source
Exhaust Flow:	cfm	600	Drawings
Fan Power:	kW	-77	
Efficiency:	cfm/kW	-99	
Total Exhaust Flow	cfm	600	
Total Power:	kW	-99.0	
Fan Pressure Rise:	in	0.5	Drawings

Measured

Average	Source	Accuracy (+/-)
600	Calculation	20%
0.38	Rumsey Eng	5%
1579		
600		
0.4		
-77		0%

Exhaust Description

Cleanroom Data

Customer Name: ***Facility I***
Site Name: ***Building I***
Cleanroom Name: ***Hybridization/Micro-Optics/ONM/Bonding - Cls 10K***

Class:	Class 10,000	HEPA Filter Coverage	14	%	Annual Hours Use:	8760	
ISO Class:	ISO Class 7	Filter Effective Area	85	%	Heat Recovery	☐	
Fan System Type	Recirculation AHU	Ceiling Filter Efficiency	-77		Raised Floor:	☐	
					Exhaust Fan(s):	☐	
Building Area:	69,310 sf						
Primary Cleanroom Area	3,750 sf	Primary Cleanroom Ceiling Heigh	9 ft				
Secondary Cleanroom Area	560 sf	Secondary Cleanroom Ceiling Heigh	-77 ft				

Power

Monitoring Start Date: 3/15/2004
Monitoring End Date: 3/23/2004

	Design			Measured				
	Unit	Value	Source	Value	Max	Min	Source	Accuracy (+/-)
Lighting:	kW	-77		6.1	-77	-77	Rumsey Eng	5%
Process:	kW	-77		6.9	-77	-77	Rumsey Eng	5%
Other:	kW	-77		-77	-77	-77		

Space Conditions

	Design				Measured				
	Unit	Value	+/-	Source	Value	Max	Min	Source	Accuracy (+/-)
Temperature:	F	69	2	Submittals	65.8	68.1	64	Rumsey Eng	5%
Humidity:	%RH	45	15	Submittals	50.8	58.6	44.4	Rumsey Eng	10%
Ceiling Velocity	fpm	-77	-77		90	-77	-77	Calculation	20%
Pressurization:	in wg	-77	-77		0.1	-77	-77	Rumsey Eng	20%

Cleanroom Description

This cleanroom is served by one makeup air handler and one recirculation air handler.

MUAH Data

Customer Name: ***Facility I***
Site Name: ***Building I***
Cleanroom Name: ***Hybridization/Micro-Optics/ONM/Bonding - Cls 10K***

MUAH Name: ***AHU-3***

Number of Units: 1

Design

	Unit	Value	Source
Air Flow:	cfm	15000	Drawings
MUAH Fan Power:	kW	-77	
Efficiency:	cfm/kW	-99	
Total Air Flow:	cfm	15000	Drawings
Total Power:	kW	-99.0	
VFD Speed:	Hz	-77	
Supply Air Temperature	°F	-77	
Supply Setpoint	°F	47	Drawings
Supply RH:	%	-77	
RH Setpoint:	%	-77	
Fan Pressure Rise:	in	-77	
Filter Pressure Drop:	in	-77	
Sensible Cooling Load	tons	-77	
Coil Face Velocity	fpm	-77	

Measured

Average	Max	Min	Source	Accuracy (+/-)
1000			Calculation	20%
13.7			Rumsey Eng	5%
73				
1000				
13.7				
55			Rumsey Eng	5%
-77	-77	-77		
-77				
-77	-77	-77		
-77				
-77				
-77				
-77	-77	-77		
-77				

MUAH Description

RCU Data

Customer Name: ***Facility I***
Site Name: ***Building I***
Cleanroom Name: ***Hybridization/Micro-Optics/ONM/Bonding - Cls 10K***

RCU Name: ***RAH-6***

Number of Units: 1

Design

	Unit	Value	Source
Air Flow:	cfm	65000	Drawings
RCU Fan Power:	kW	-77	
Efficiency:	cfm/kW	-99	
VFD Speed:	Hz	-77	
Supply Air Temperature	°F	-77	
Return Air Temperature	°F	-77	
Supply Setpoint	°F	-77	
RH Setpoint:	%	-77	
Supply RH:	%	-77	
Fan Pressure Rise:	in w.g.	-77	
Filter Pressure Drop:	in w.g.	-77	
Sensible Cooling Load	tons	-77	
Coil Face Velocity	fpm	-77	

Measured

Average	Max	Min	Source	Accuracy (+/-)
40260			Balance Report	20%
22.9			Rumsey Eng	5%
1758				
42			Rumsey Eng	5%
66	69	62	Rumsey Eng	5%
67	69	65	Rumsey Eng	5%
-77				
-77				
-77	-77	-77		
1.8			Rumsey Eng	10%
0.4			Rumsey Eng	10%
-77				
-77				

RCU Description

Exhaust Data

Customer Name: ***Facility I***
Site Name: ***Building I***
Cleanroom Name: ***Hybridization/Micro-Optics/ONM/Bonding - Cls 10K***

Exhaust Name: ***EF-18***

Number of Units: 1

Design

	Unit	Value	Source
Exhaust Flow:	cfm	800	Drawings
Fan Power:	kW	-77	
Efficiency:	cfm/kW	-99	
Total Exhaust Flow	cfm	800	
Total Power:	kW	-99.0	
Fan Pressure Rise:	in	0.5	Drawings

Measured

Average	Source	Accuracy (+/-)
800	Calculation	20%
0.47	Rumsey Eng	5%
1702		
800		
0.5		
-77		0%

Exhaust Description

Chiller Data

Customer Name: ***Facility I***
Site Name: ***Building I***

Chiller Name: ***Chilled Water Plant***

End Use: Combined
Nominal Tons: 170
Total Number of Chillers of this type, including Standb 3
Number of Standby Chillers of this type 1
Monitoring Start Date: 3/18/2004
Monitoring End Date: 3/23/2004

Design

	Unit	Value	Source
Total Power:	kW	-77	
Cooling Supplied	Tons	340	Drawings
Efficiency:	kW/Ton	-99.00	

Measured

Average	Max	Min	Source	Accuracy (+/-)
195.5	358.8	93.2	Rumsey Eng	0.05
201	328	108	Calculation	0.2
0.97				

Chiller Description

Chillers 1 and 3 were running during monitoring period.

Chiller Name: ***Process Chiller***

End Use: Process
Nominal Tons: 50
Total Number of Chillers of this type, including Standb 1
Number of Standby Chillers of this type 0
Monitoring Start Date: 3/18/2004
Monitoring End Date: 3/20/2004

Design

	Unit	Value	Source
Total Power:	kW	-77	
Cooling Supplied	Tons	50	Drawings
Efficiency:	kW/Ton	-99.00	

Measured

Average	Max	Min	Source	Accuracy (+/-)
21.8	28.8	14.2	Rumsey Eng	0.05
17.8	48.4	4.4	Calculation	0.2
1.22				

Chiller Description

Chilled Water Pump Loop Data

Customer Name: ***Facility I***
Site Name: ***Building I***

Pump Loop Name: **CHW Pumps**
Pump Loop Type: **Primary Chilled Water**

Number of Pumps:	3	Monitoring Start Date	3/15/2004
Number Used as Backup	1	Monitoring End Date	3/23/2004

Design

	Unit	Value	Source
Supply Temp:	°F	40	Drawings
Return Temp:	°F	55	Drawings
Total Flow:	gpm	-77	
Cooling Tons:	Tons	-99	
Total Power:	kW	-77	
Efficiency:	kW/Ton	-99.000	
Head:	ft	65	Drawings

Measured

Average	Max	Min	Source	Accuracy (+/-)
40.8	43.8	38.3	Rumsey Eng	2%
48.4	54.2	43.3	Rumsey Eng	2%
641	759	435	Rumsey Eng	20%
203				
12.4	17.5	8.2	Rumsey Eng	5%
0.061				
56.9	-77	-77	Rumsey Eng	5%

Chilled Water Pump Loop Description

Chilled water (CHW) pumps 1 and 3 were running during monitoring period.

Chilled Water Pump Loop Data

Customer Name: ***Facility I***
Site Name: ***Building I***

Pump Loop Name: **PCHW Pumps**
Pump Loop Type: **Process Chilled Water**

Number of Pumps: 2
Number Used as Backup 1

Monitoring Start Date 3/18/2004
Monitoring End Date 3/20/2004

Design

	Unit	Value	Source
Supply Temp:	°F	65	Drawings
Return Temp:	°F	75	Drawings
Total Flow:	gpm	120	Drawings
Cooling Tons:	Tons	50	
Total Power:	kW	-77	
Efficiency:	kW/Ton	-99.000	
Head:	ft	95	Drawings

Measured

Average	Max	Min	Source	Accuracy (+/-)
62.9	64.9	60.6	Rumsey Eng	2%
66.6	70.6	63.9	Rumsey Eng	2%
123	140	107	Rumsey Eng	20%
19				
7.8	8.5	7.4	Rumsey Eng	2%
0.411				
90	-77	-77	Rumsey Eng	5%

Chilled Water Pump Loop Description

PCHW pump 2 was running during monitoring period.

Appendix B

Trended Data Graphs

Outside Air Conditions

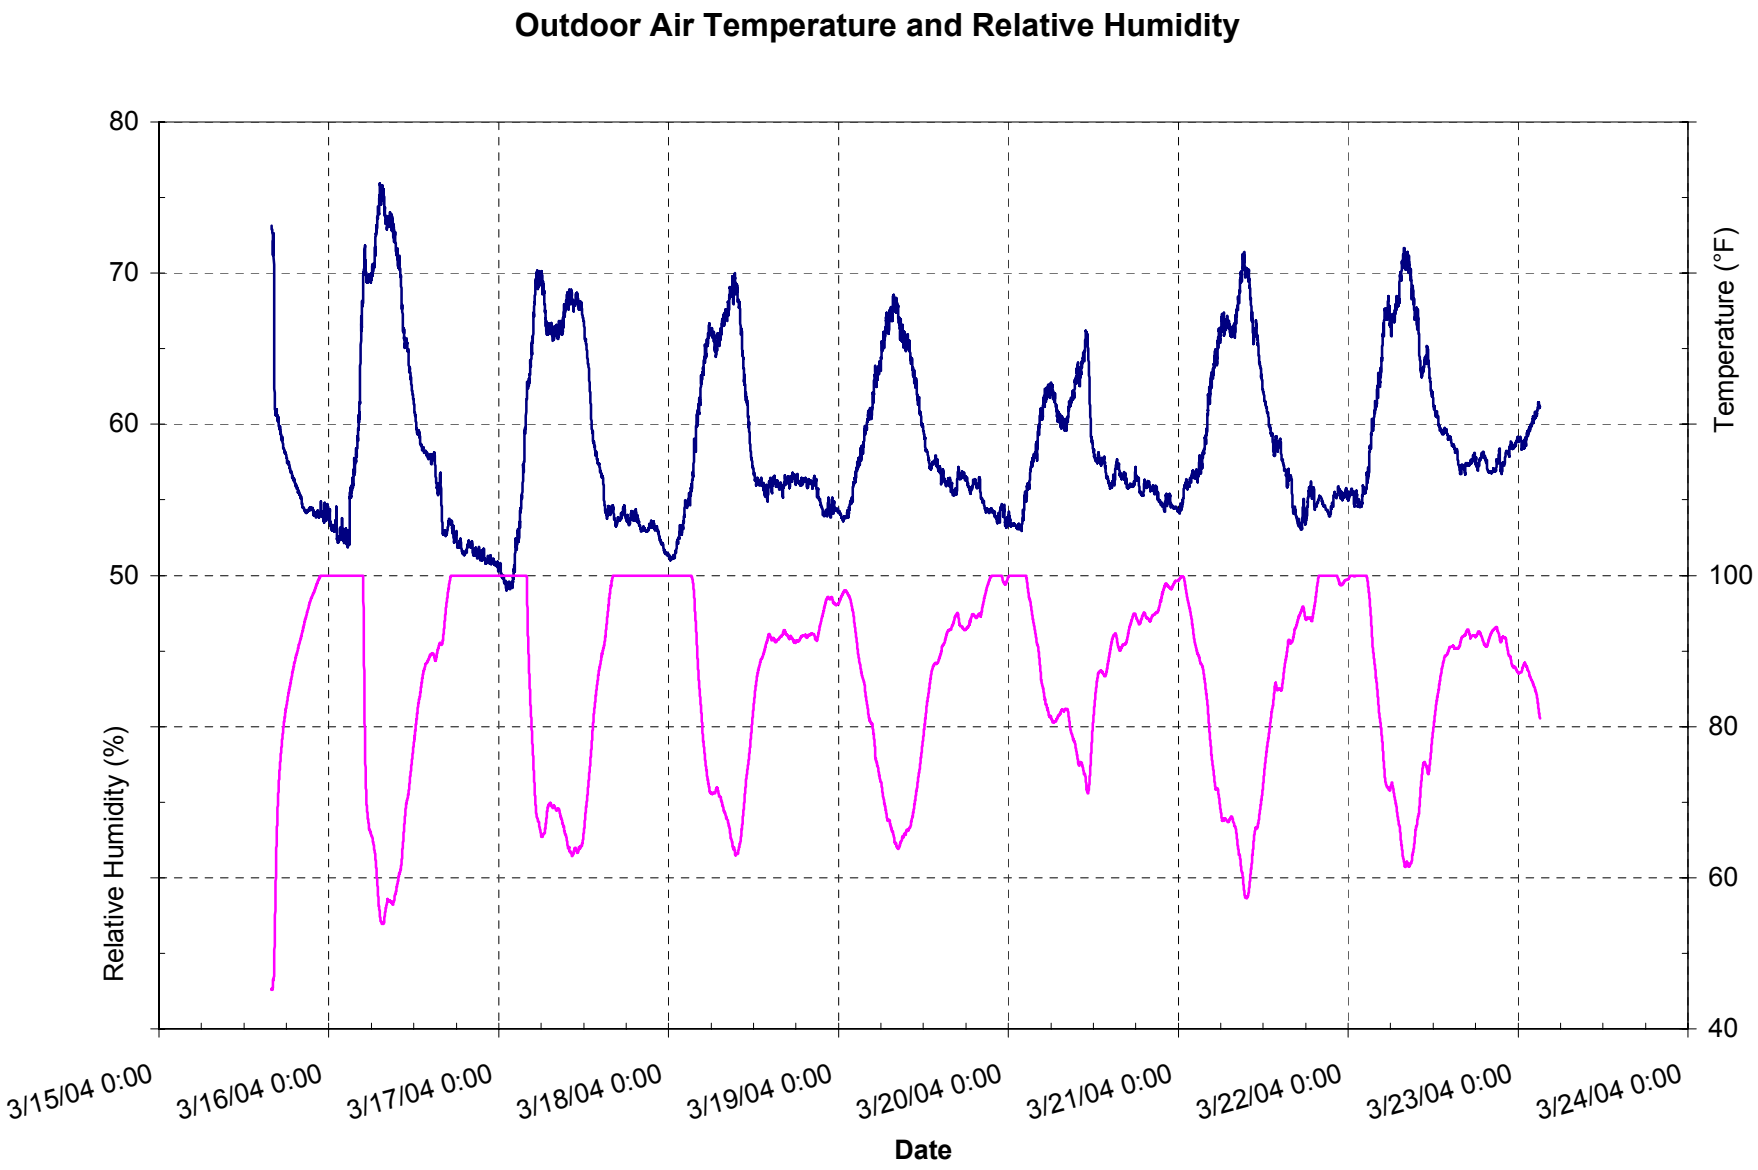
Outdoor Air Temperature and Relative Humidity
80
70
60
50
Relative Humidity (%)
100
80
60
40
Temperature (°F)
3/15/04 0:00
3/16/04 0:00
3/17/04 0:00
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
3/23/04 0:00
3/24/04 0:00
Date

MBE/Metal/Junction Formation Cleanrooms

Class 100

MBE Class 100
Temperature and Relative Humidity

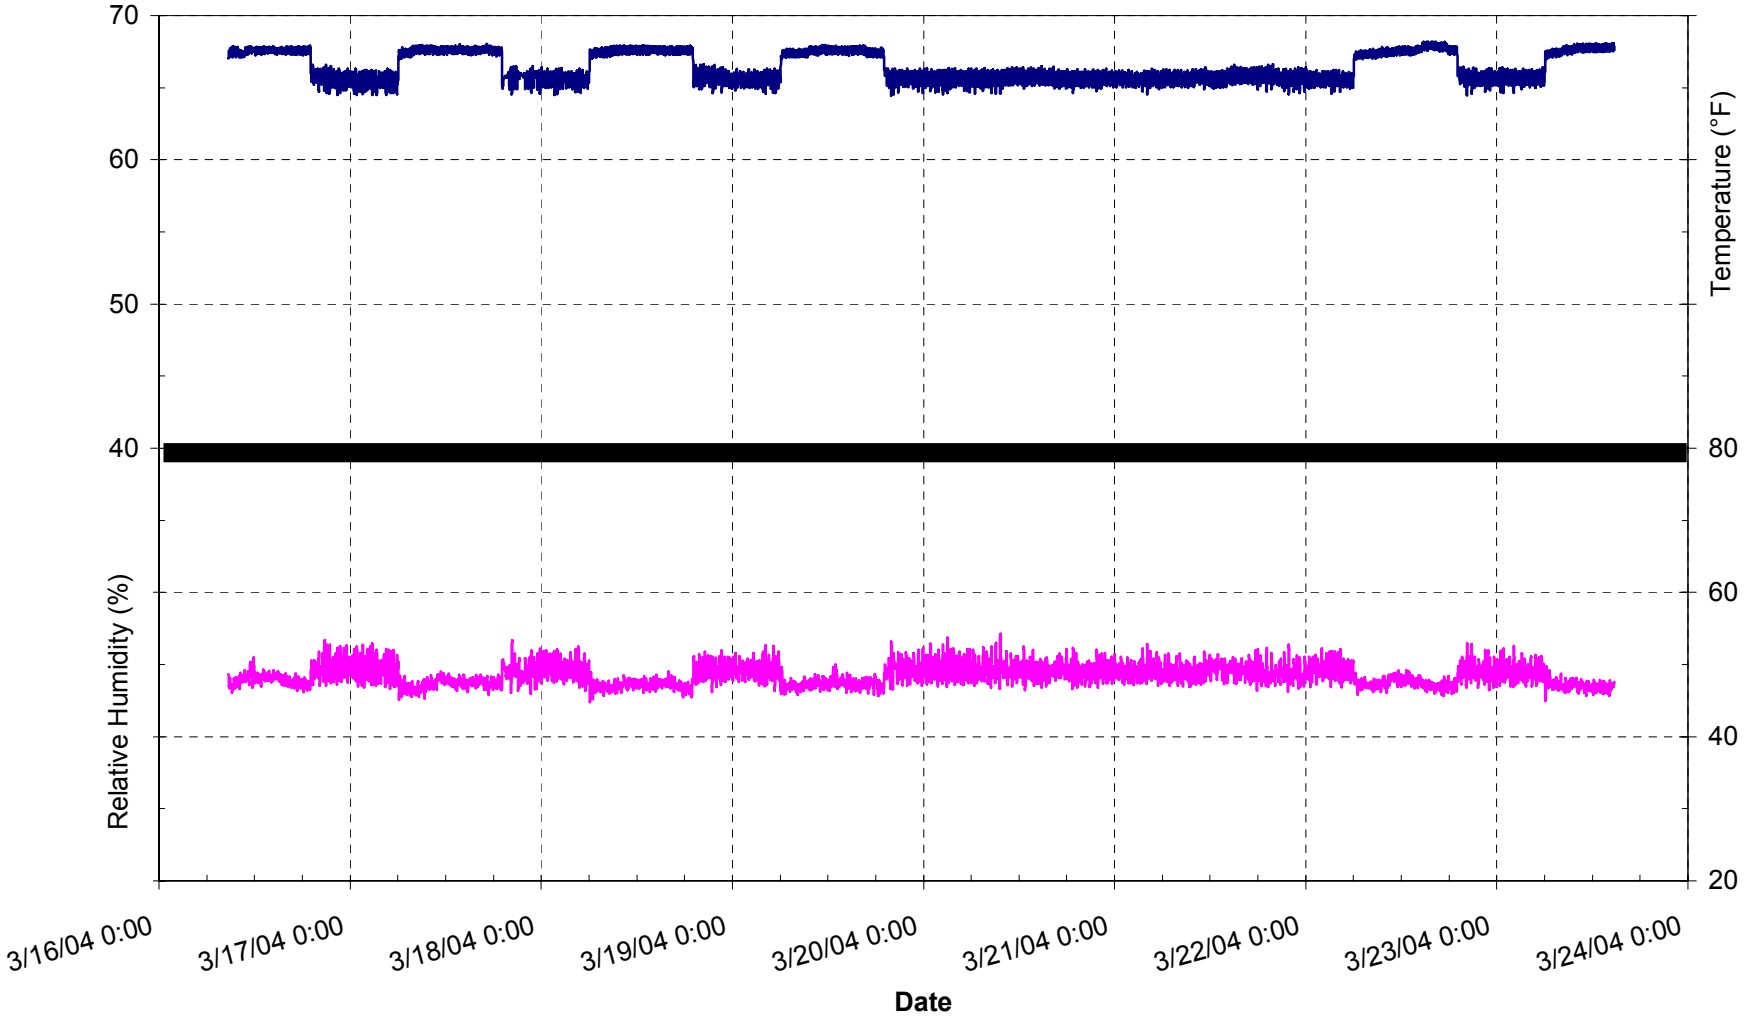

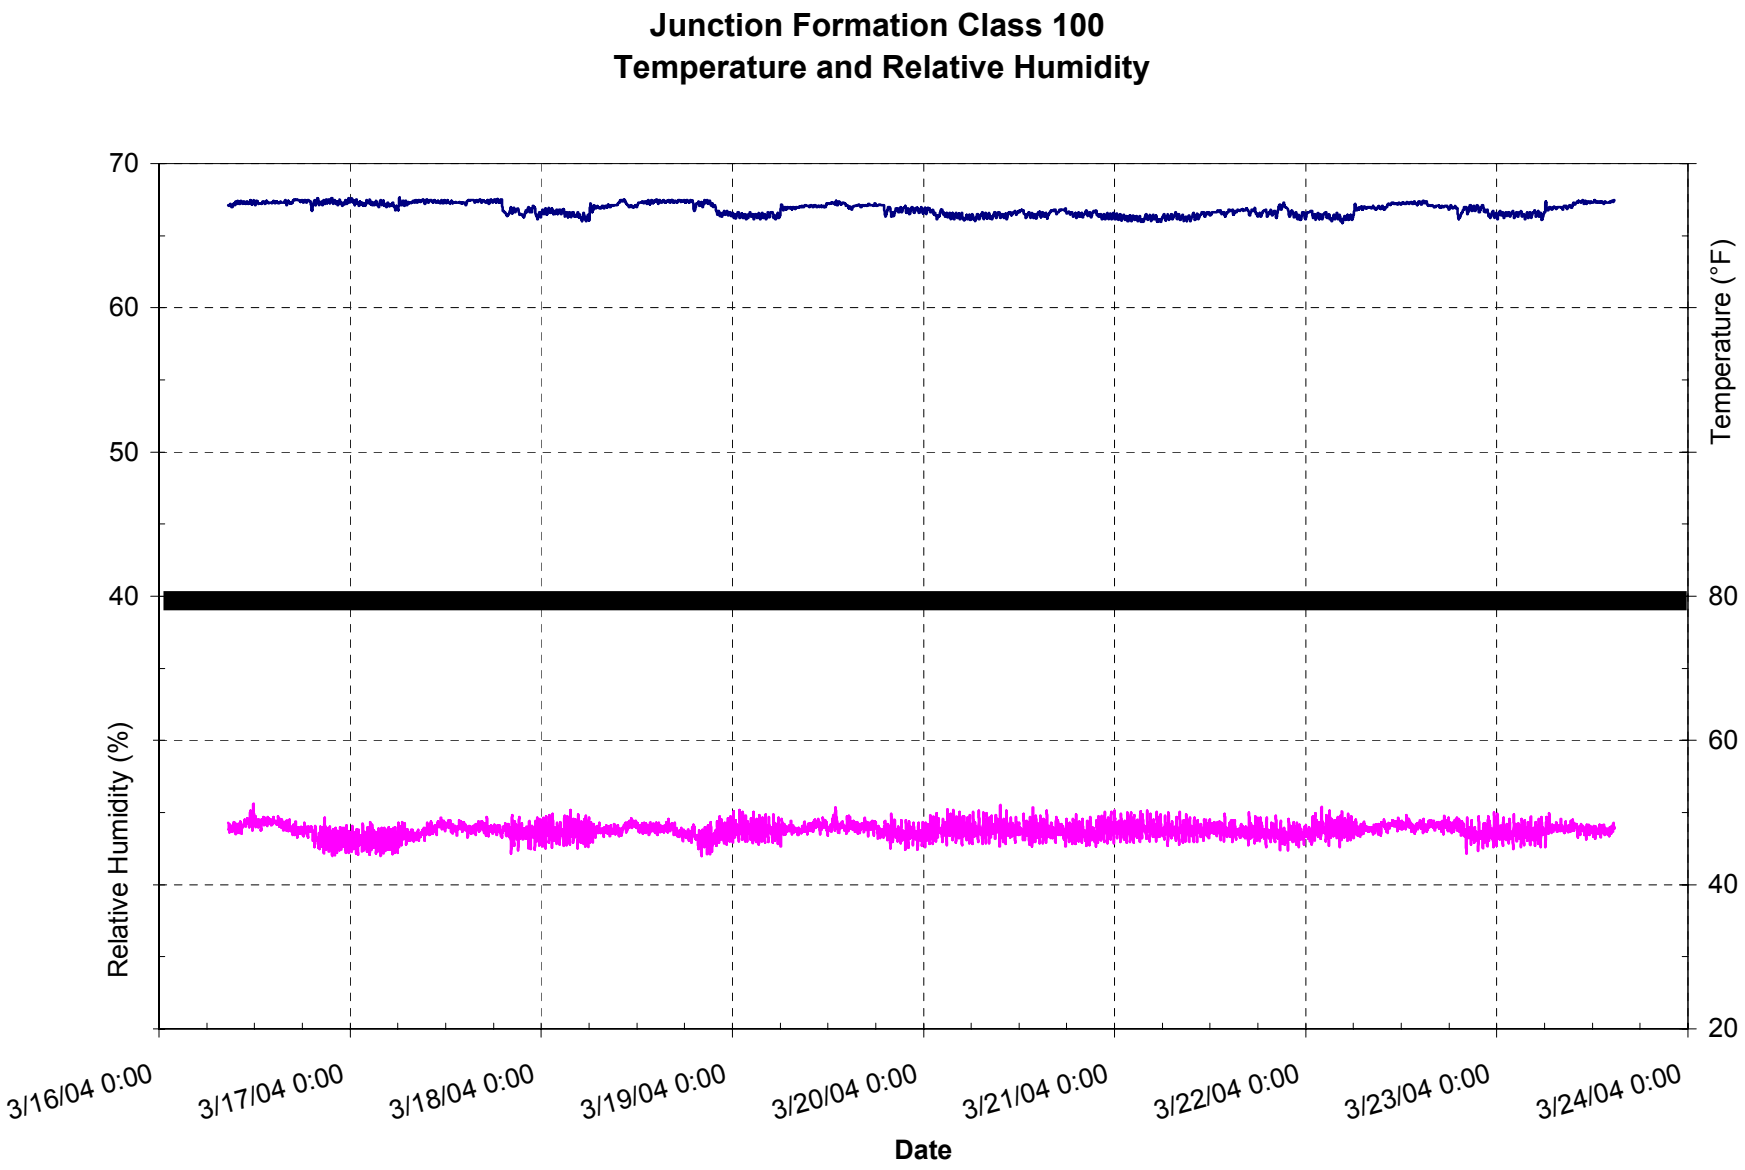
Junction Formation Class 100
Temperature and Relative Humidity
70
60
50
40
Relative Humidity (%)
80
60
40
20
Temperature (°F)
3/16/04 0:00
3/17/04 0:00
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
3/23/04 0:00
3/24/04 0:00
Date

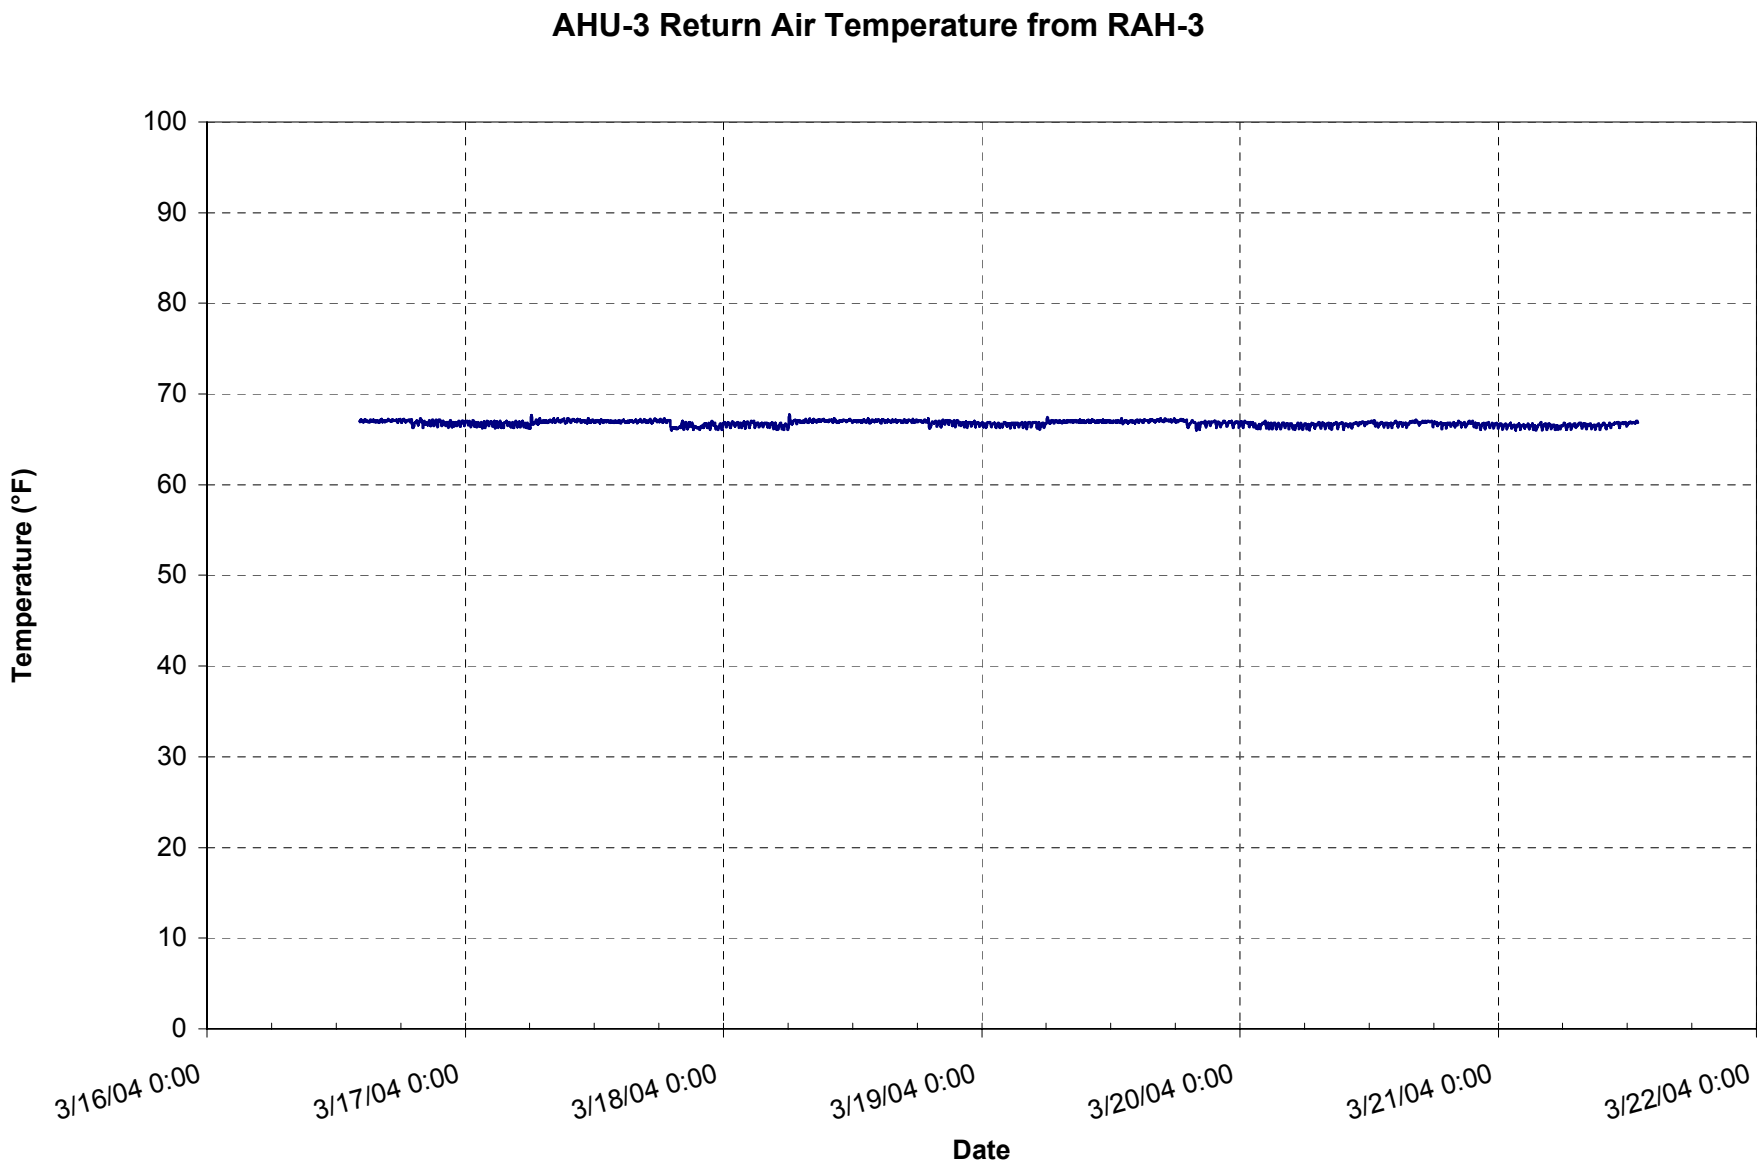
AHU-3 Return Air Temperature from RAH-3
Temperature (°F)
100
90
80
70
60
50
40
30
20
10
0
3/16/04 0:00
3/17/04 0:00
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
Date

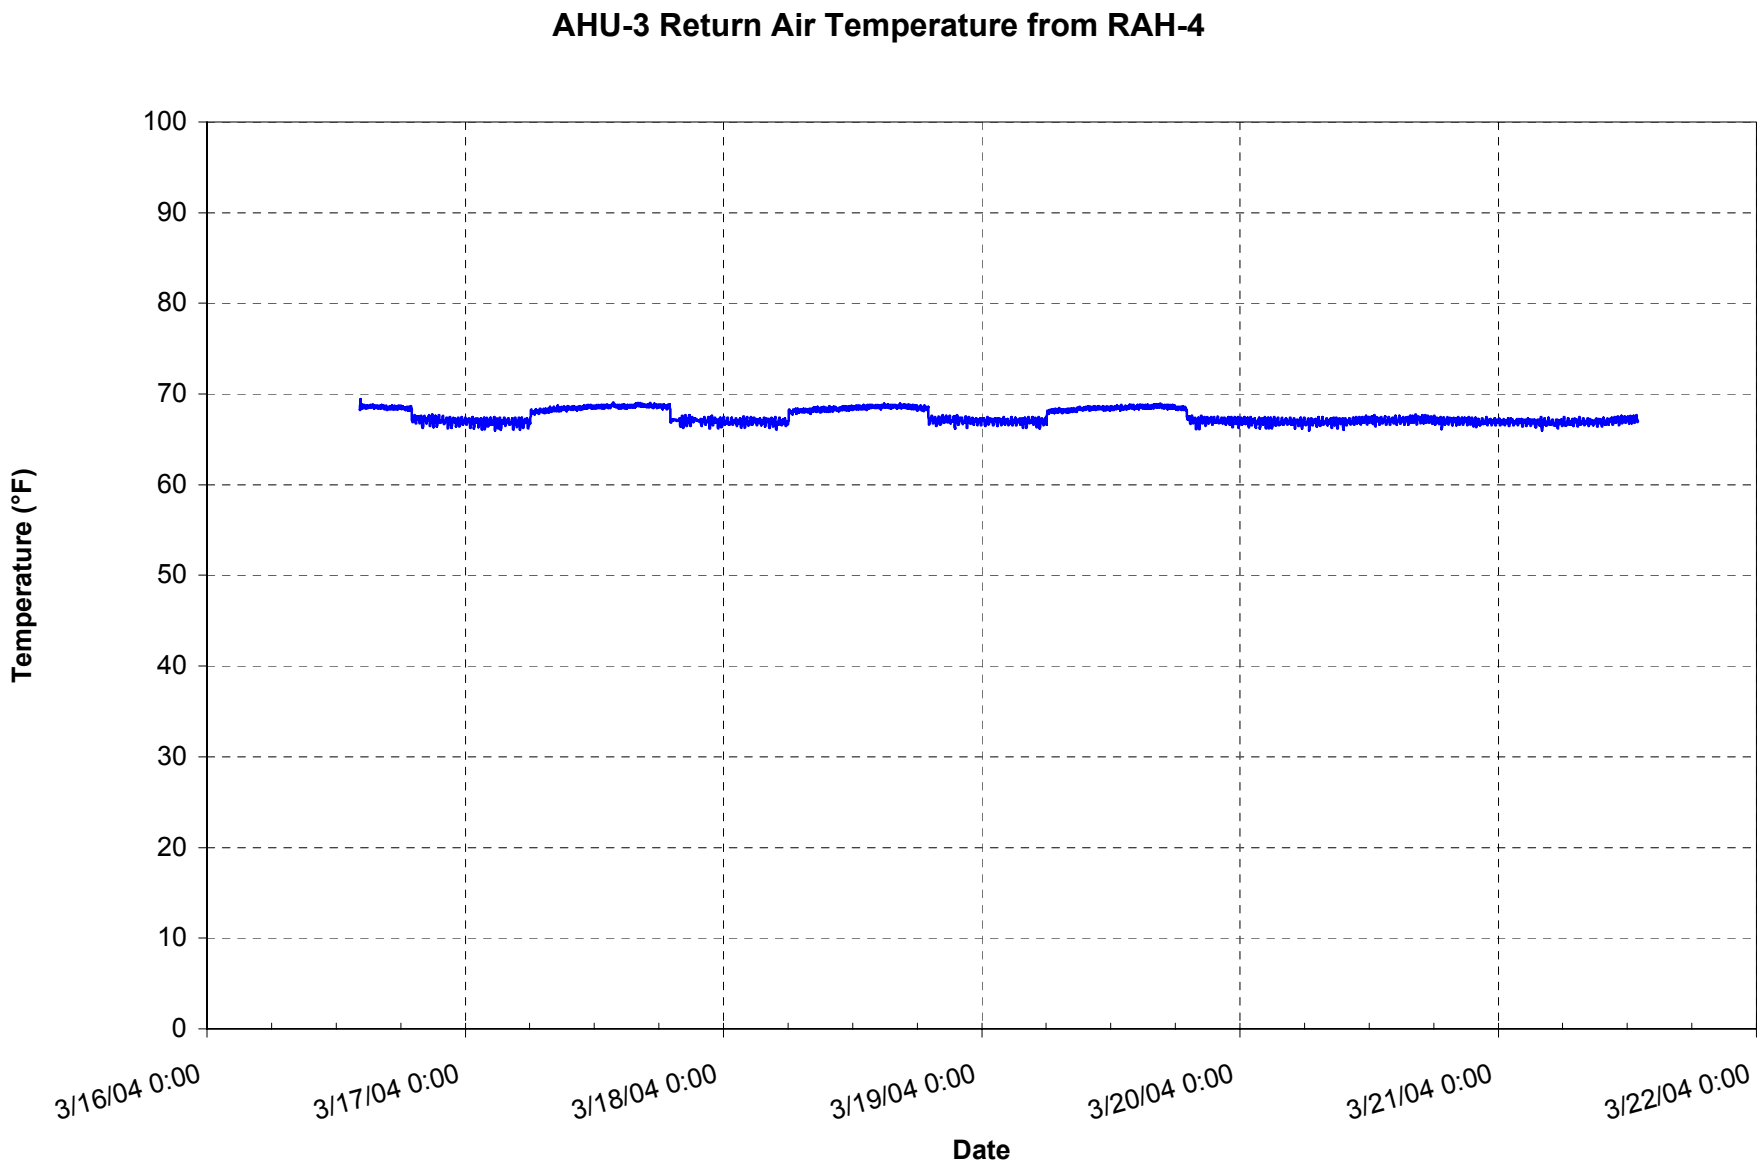
AHU-3 Return Air Temperature from RAH-4
Temperature (°F)
100
90
80
70
60
50
40
30
20
10
0
3/16/04 0:00
3/17/04 0:00
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
Date

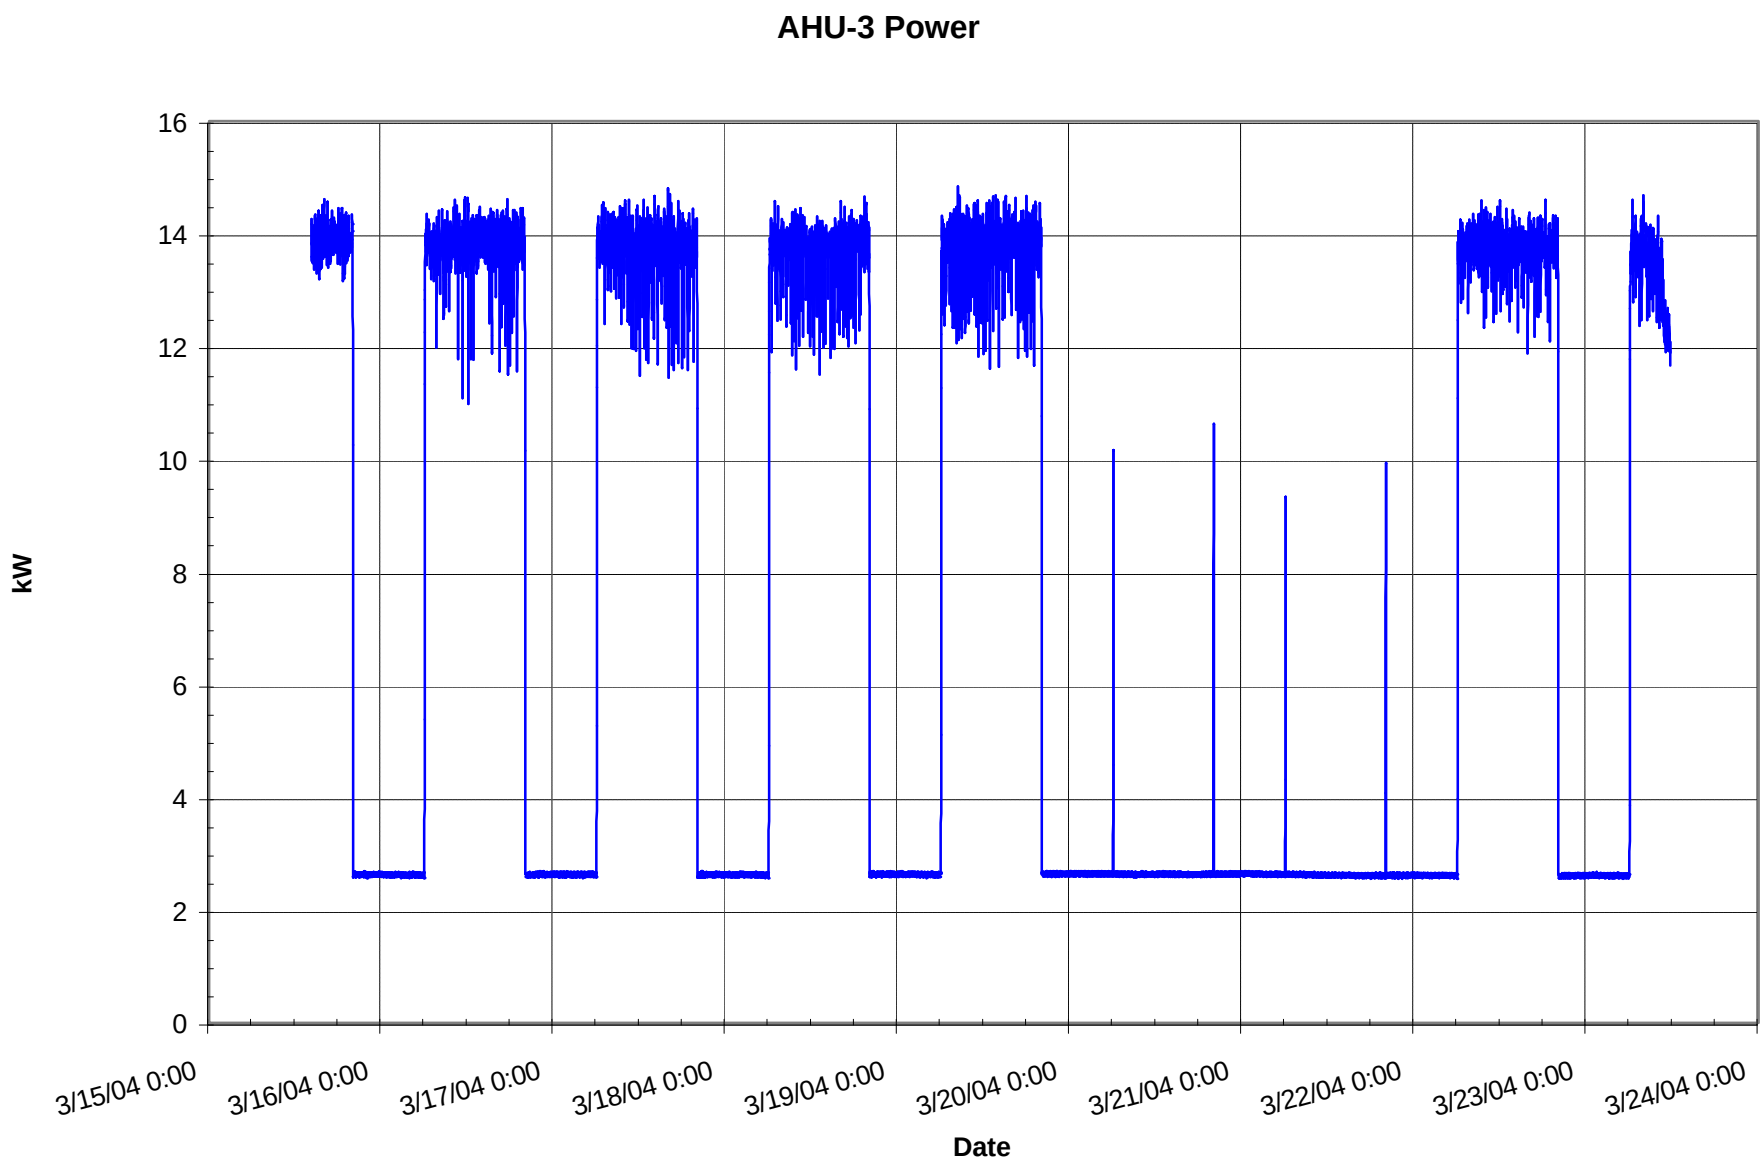
AHU-3 Power
kW
16
14
12
10
8
6
4
2
0
3/15/04 0:00
3/16/04 0:00
3/17/04 0:00
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
3/23/04 0:00
3/24/04 0:00
Date

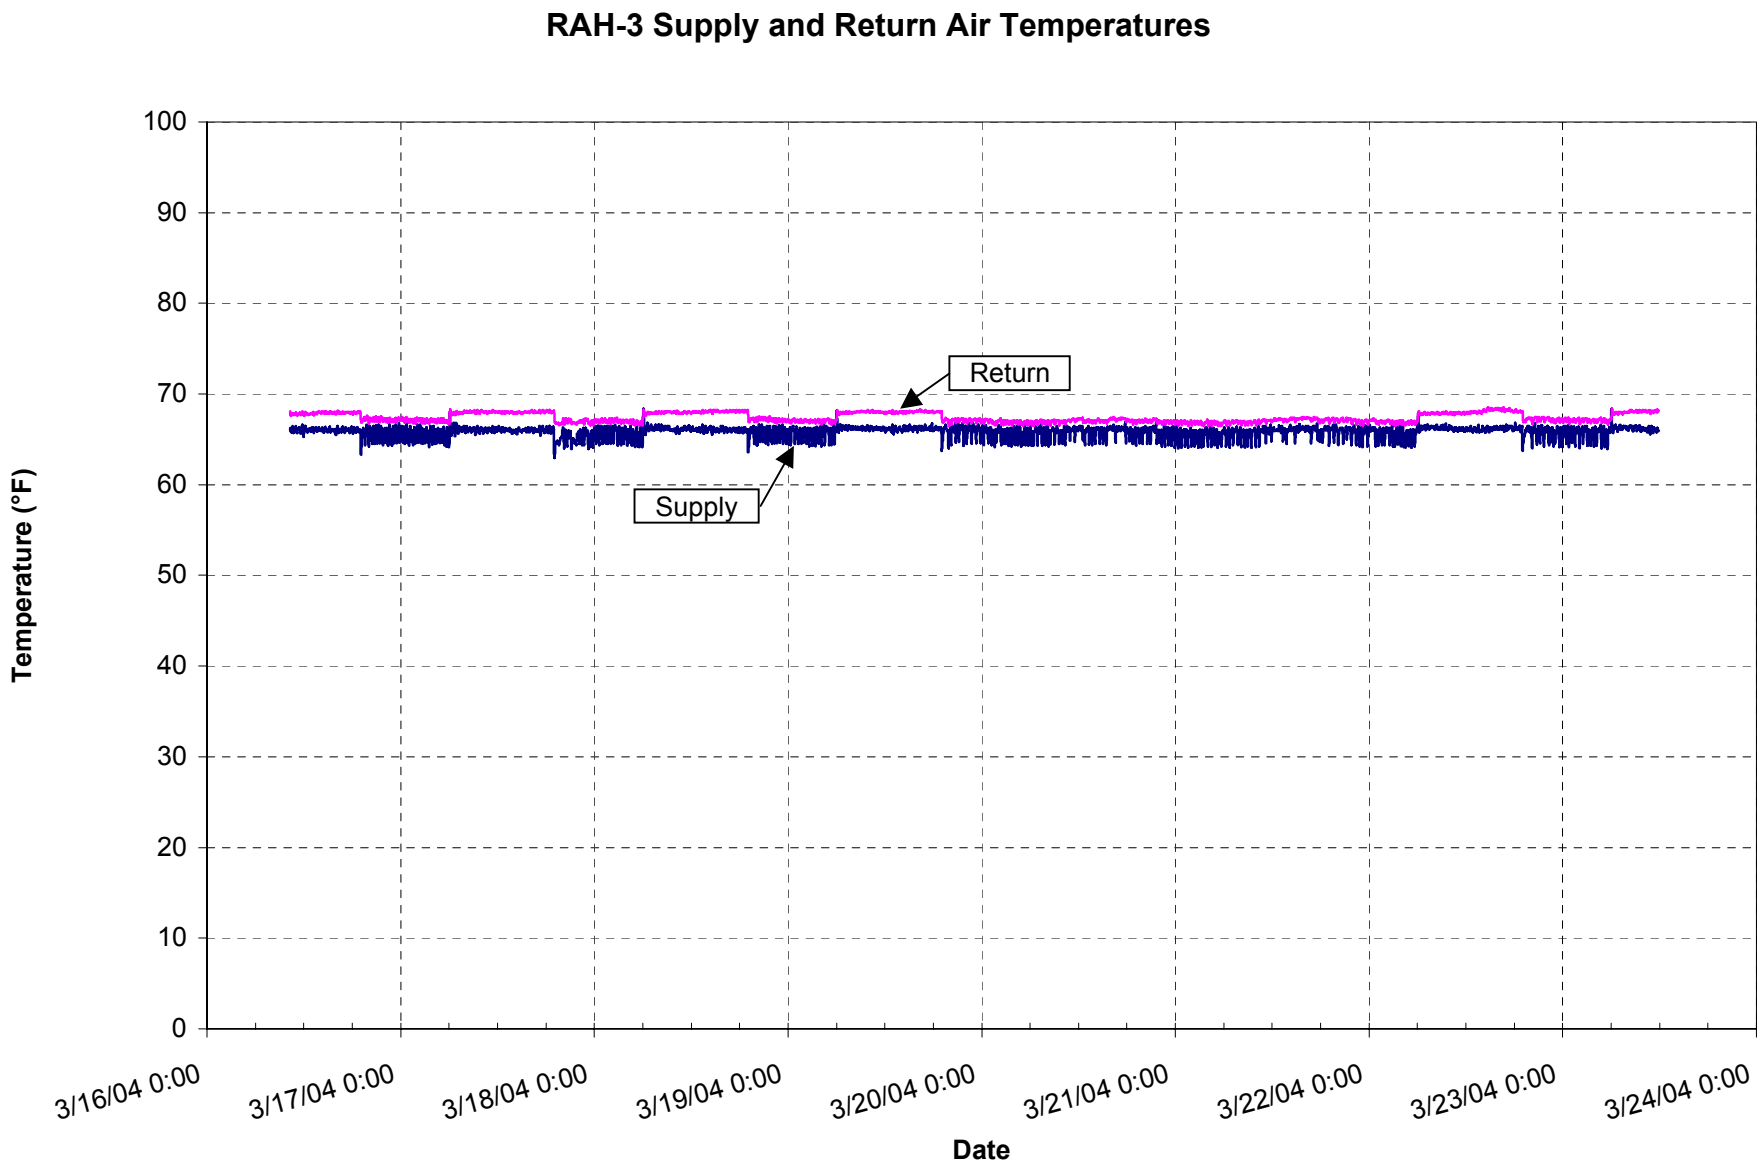
RAH-3 Supply and Return Air Temperatures
Temperature (°F)
100
90
80
70
60
50
40
30
20
10
0
Return
Supply
3/16/04 0:00
3/17/04 0:00
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
3/23/04 0:00
3/24/04 0:00
Date

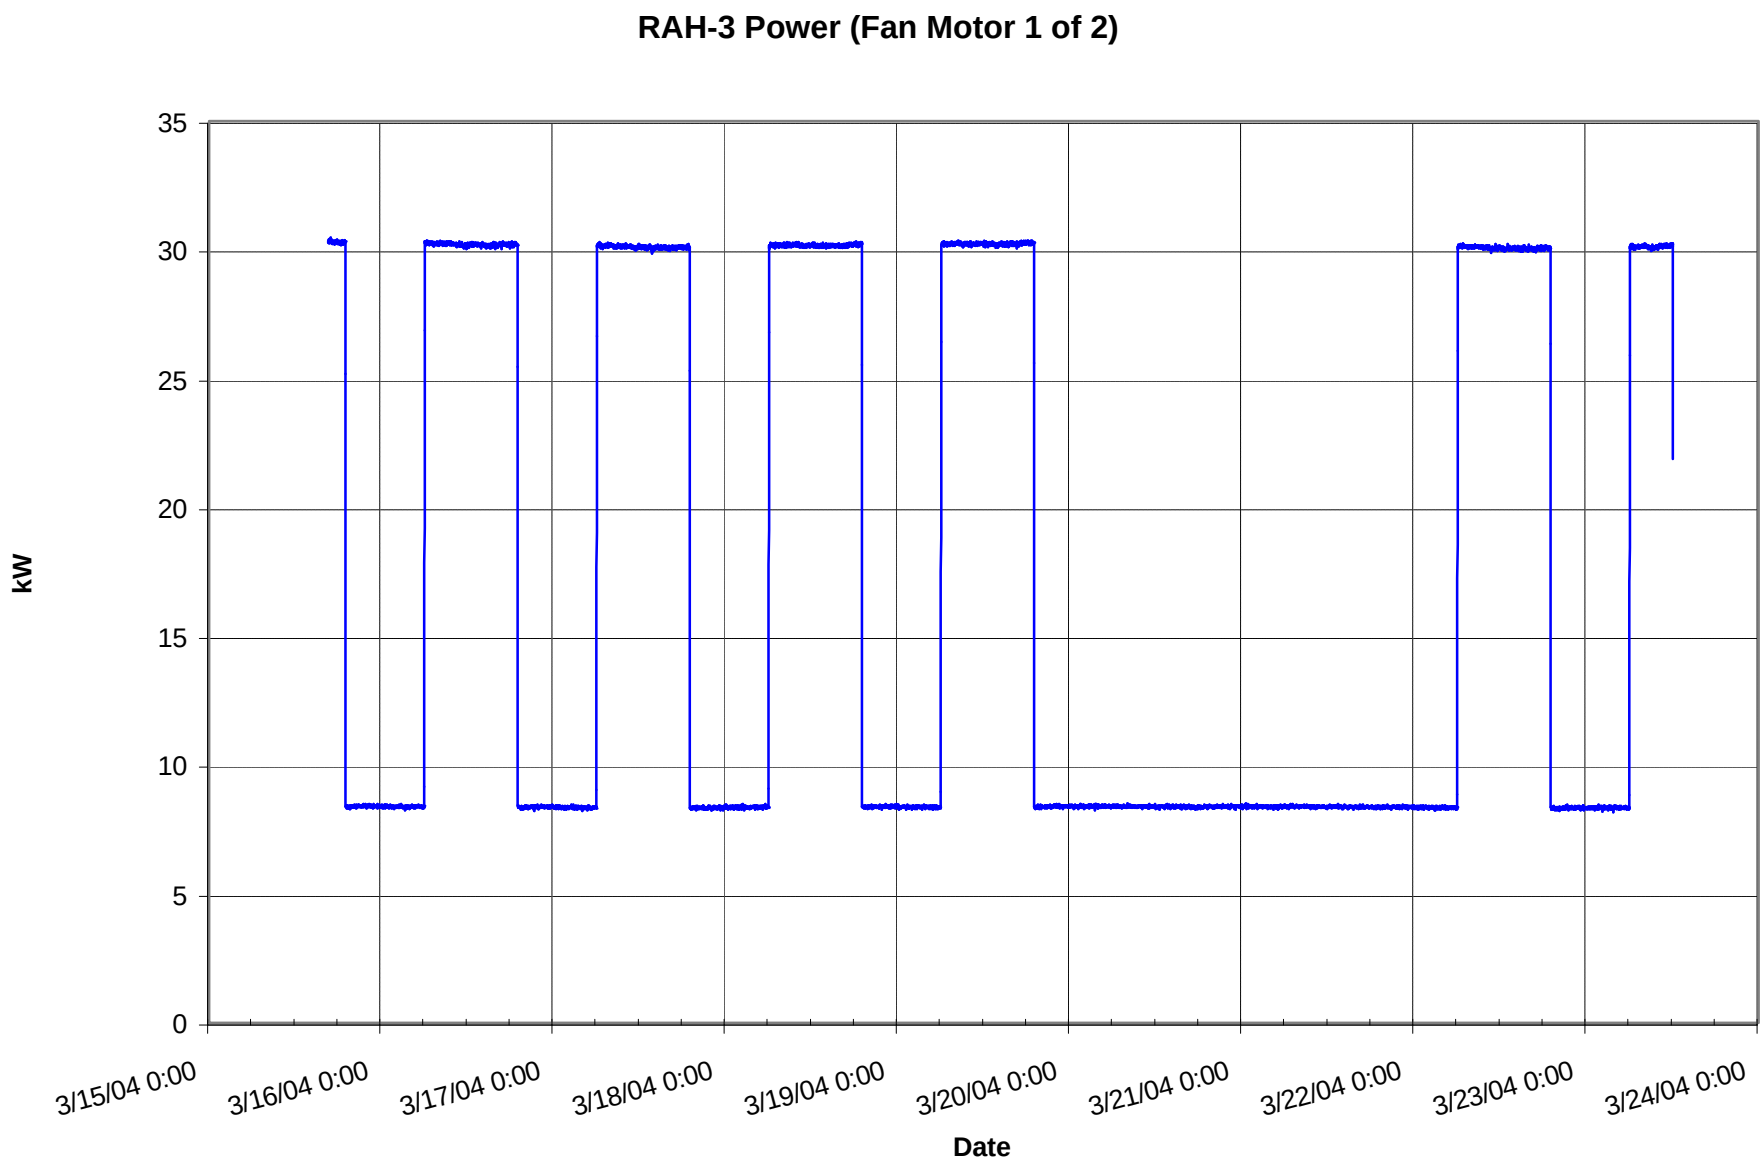
RAH-3 Power (Fan Motor 1 of 2)
kW
35
30
25
20
15
10
5
0
3/15/04 0:00
3/16/04 0:00
3/17/04 0:00
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
3/23/04 0:00
3/24/04 0:00
Date

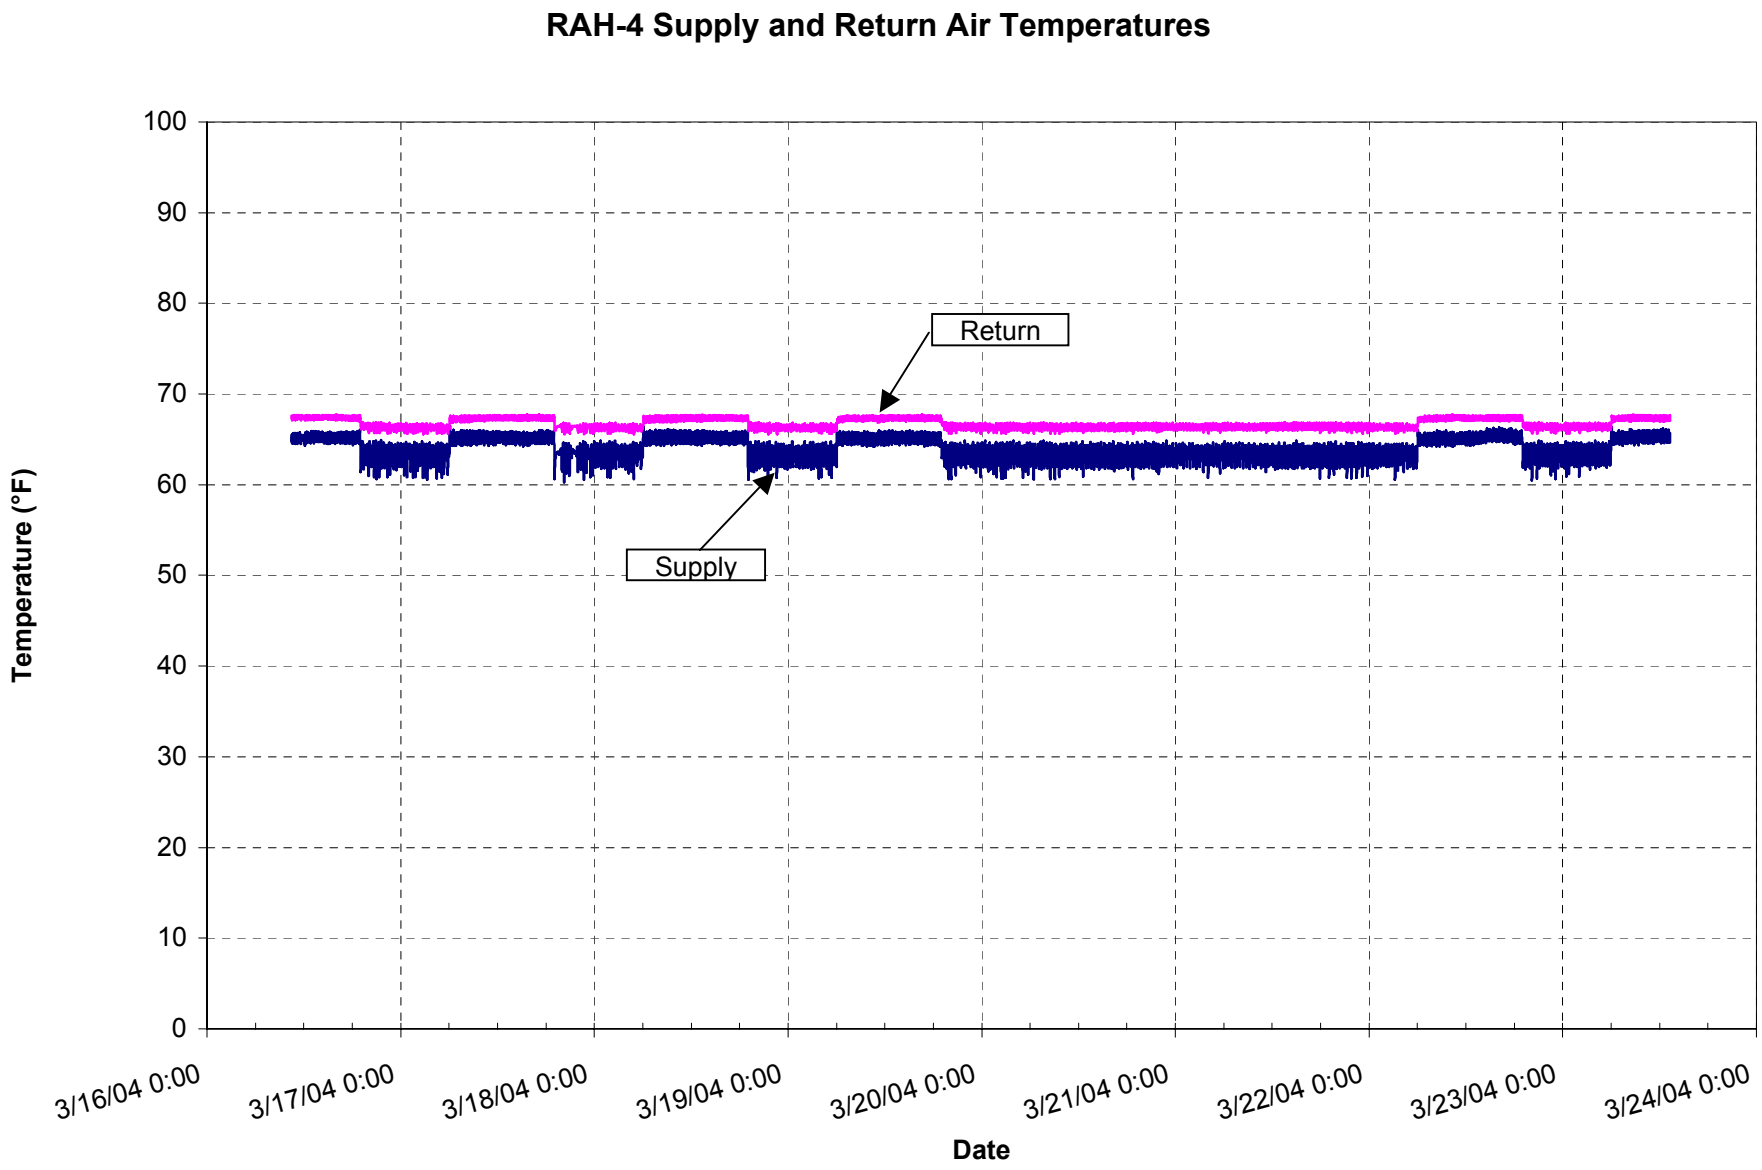
RAH-4 Supply and Return Air Temperatures
Temperature (°F)
100
90
80
70
60
50
40
30
20
10
0
Return
Supply
3/16/04 0:00
3/17/04 0:00
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
3/23/04 0:00
3/24/04 0:00
Date

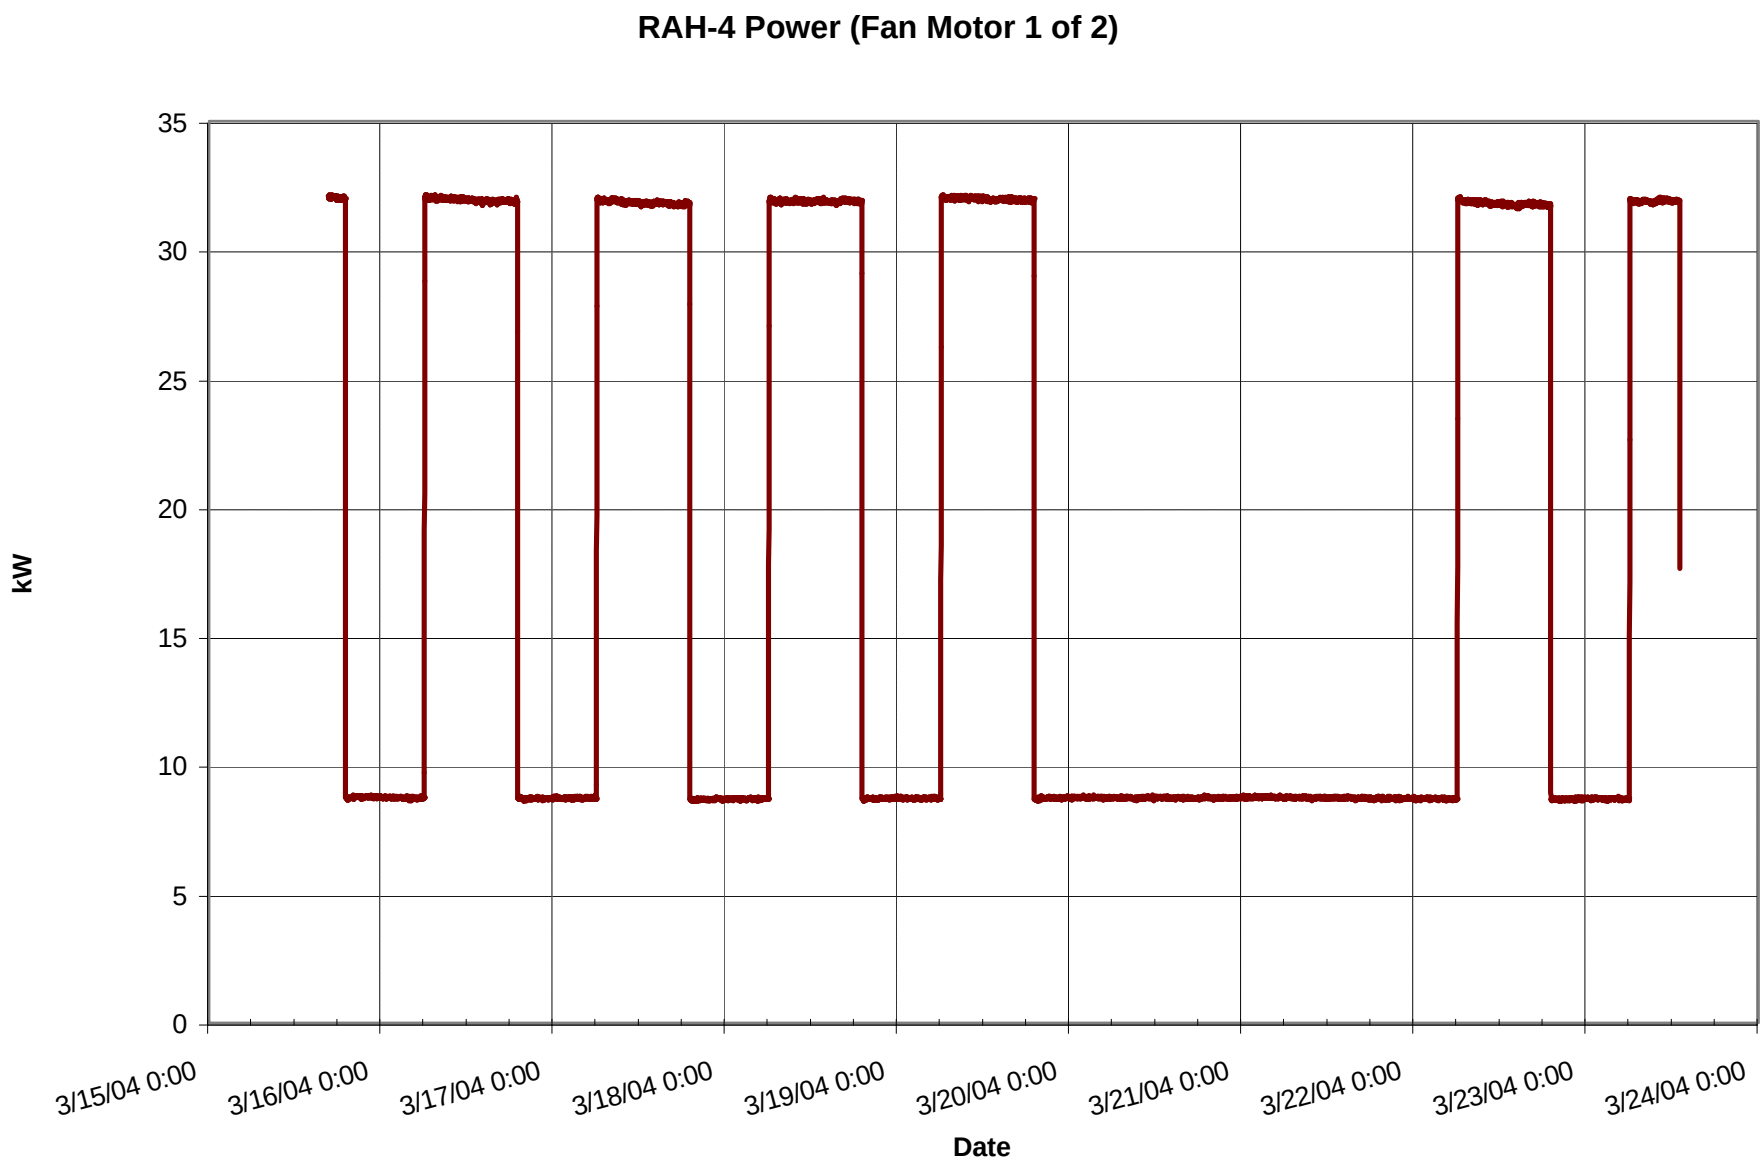
RAH-4 Power (Fan Motor 1 of 2)
kW
35
30
25
20
15
10
5
0
3/15/04 0:00
3/16/04 0:00
3/17/04 0:00
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
3/23/04 0:00
3/24/04 0:00
Date

Hybridization/Micro-Optics/ ONM Array Production Assembly/Bonding Cleanrooms

Class 10,000

Hybridlization Class 10,000
Temperature and Relative Humidity

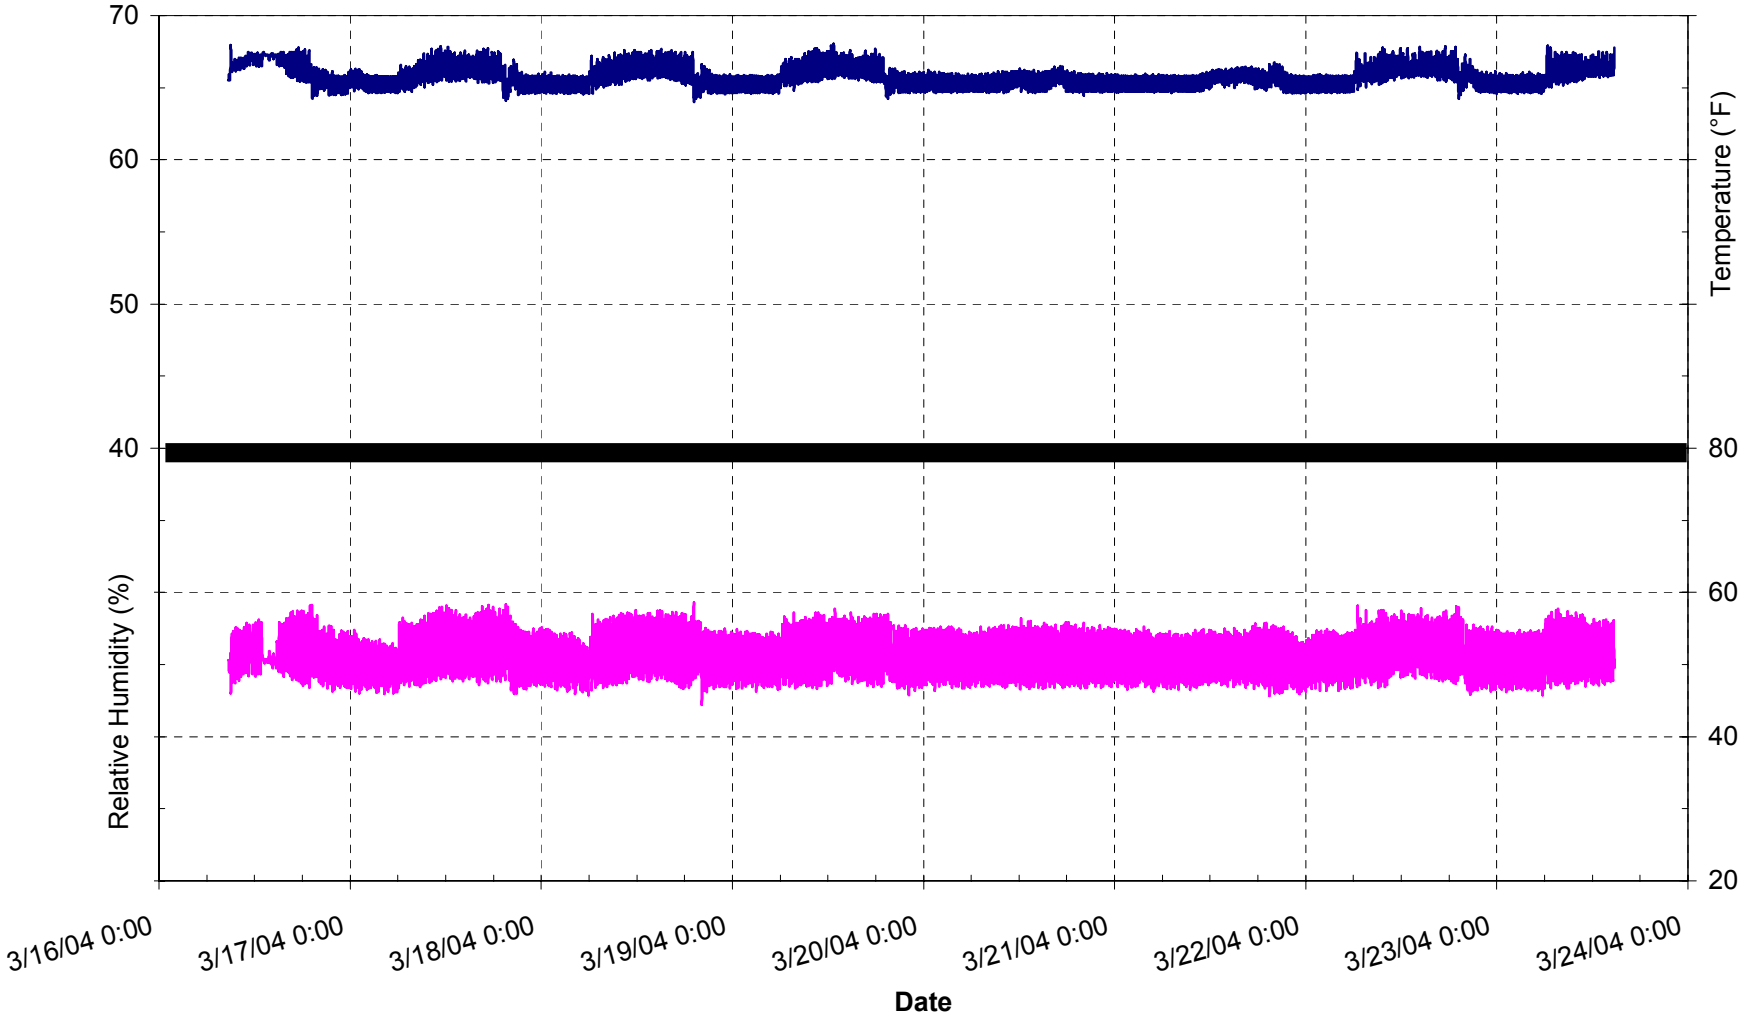

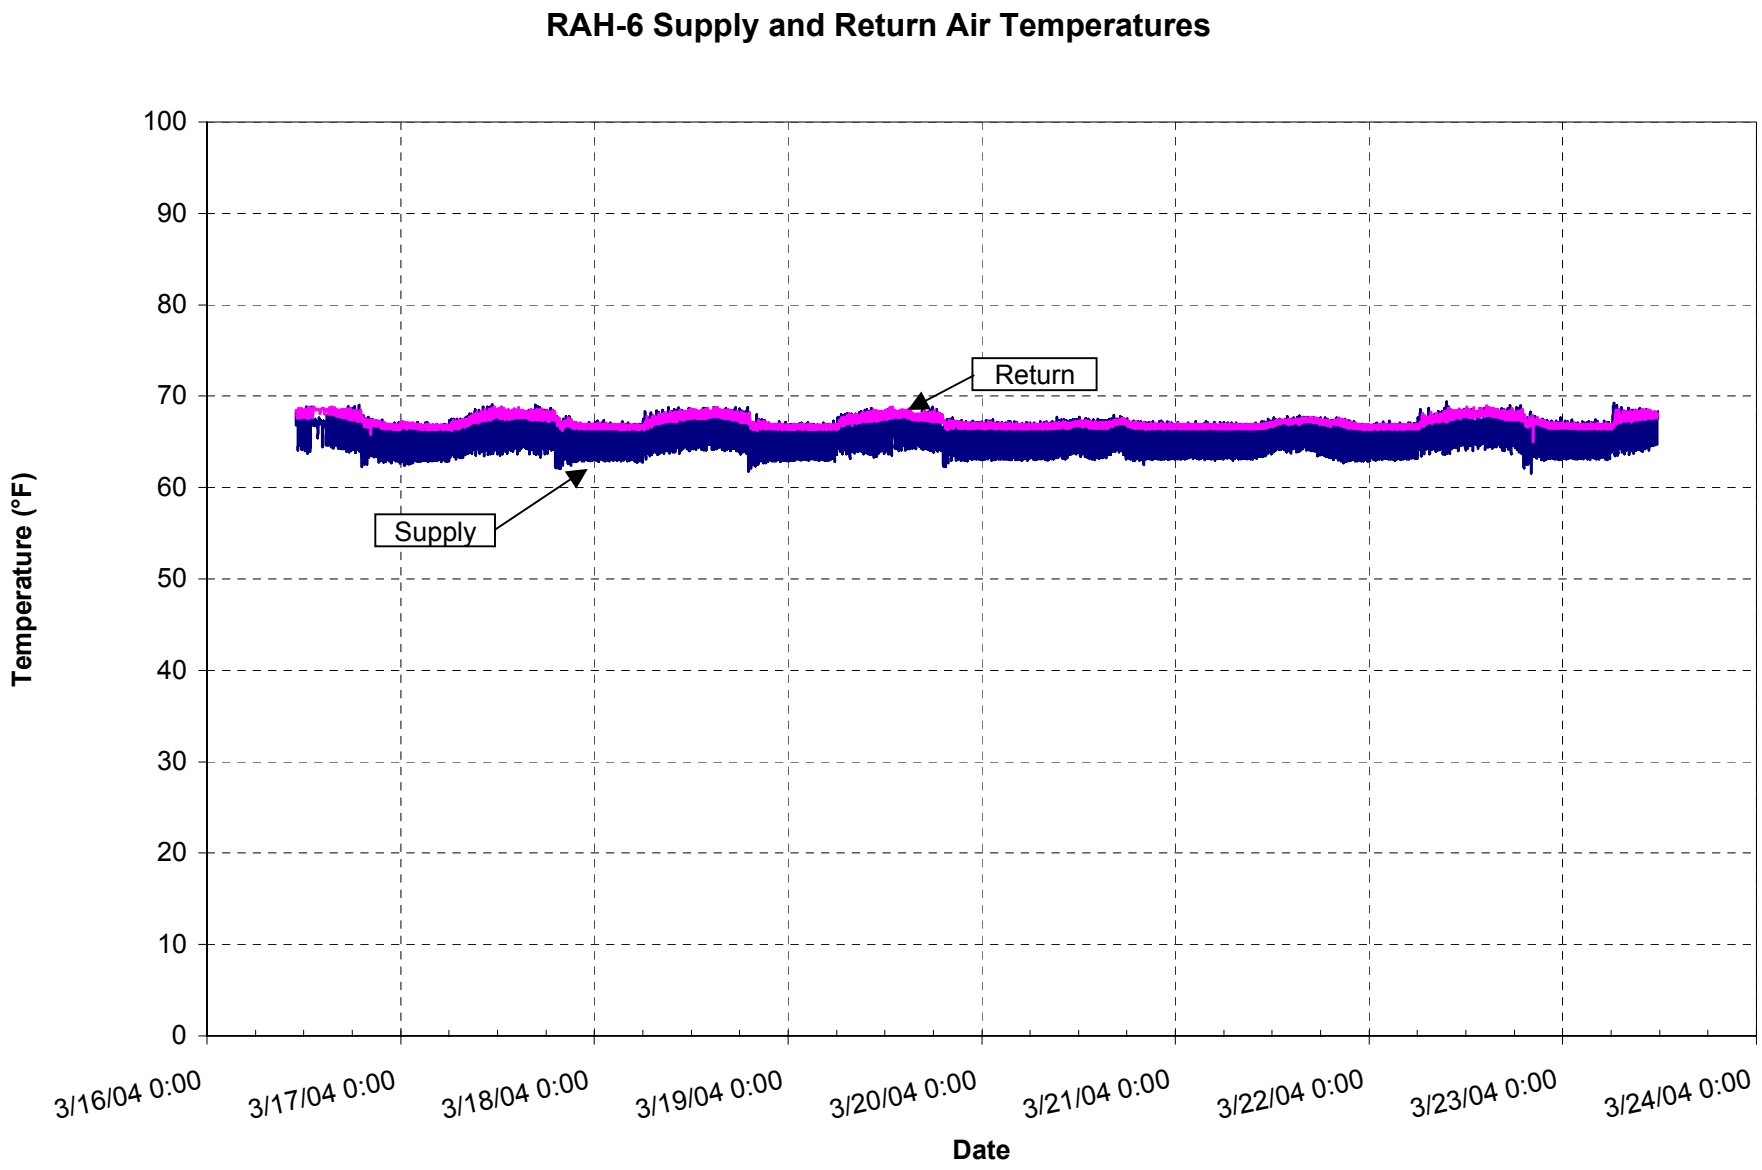
RAH-6 Supply and Return Air Temperatures
Temperature (°F)
100
90
80
70
60
50
40
30
20
10
0
Return
Supply
3/16/04 0:00
3/17/04 0:00
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
3/23/04 0:00
3/24/04 0:00
Date

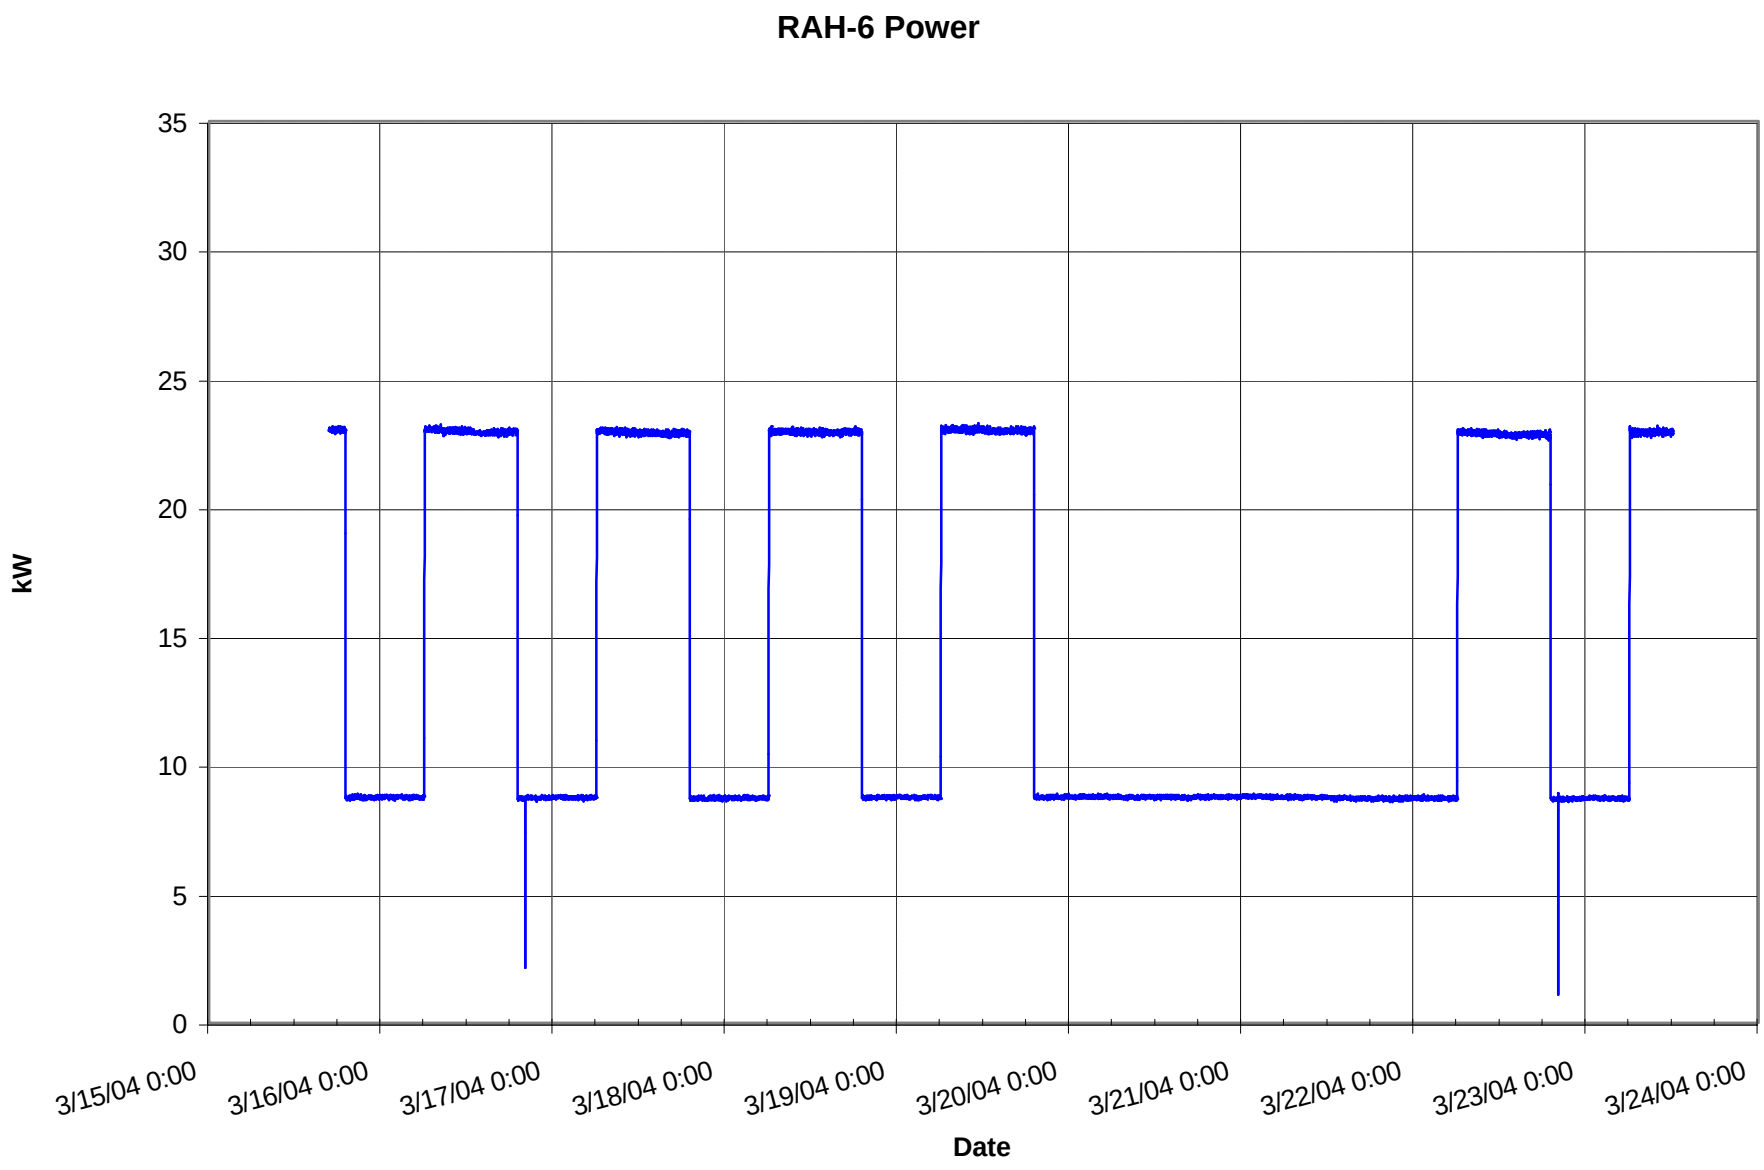
RAH-6 Power
35
30
25
20
15
10
5
0
kW
3/15/04 0:00
3/16/04 0:00
3/17/04 0:00
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
3/23/04 0:00
3/24/04 0:00
Date

Electrical Systems

UPS & Backup Generator

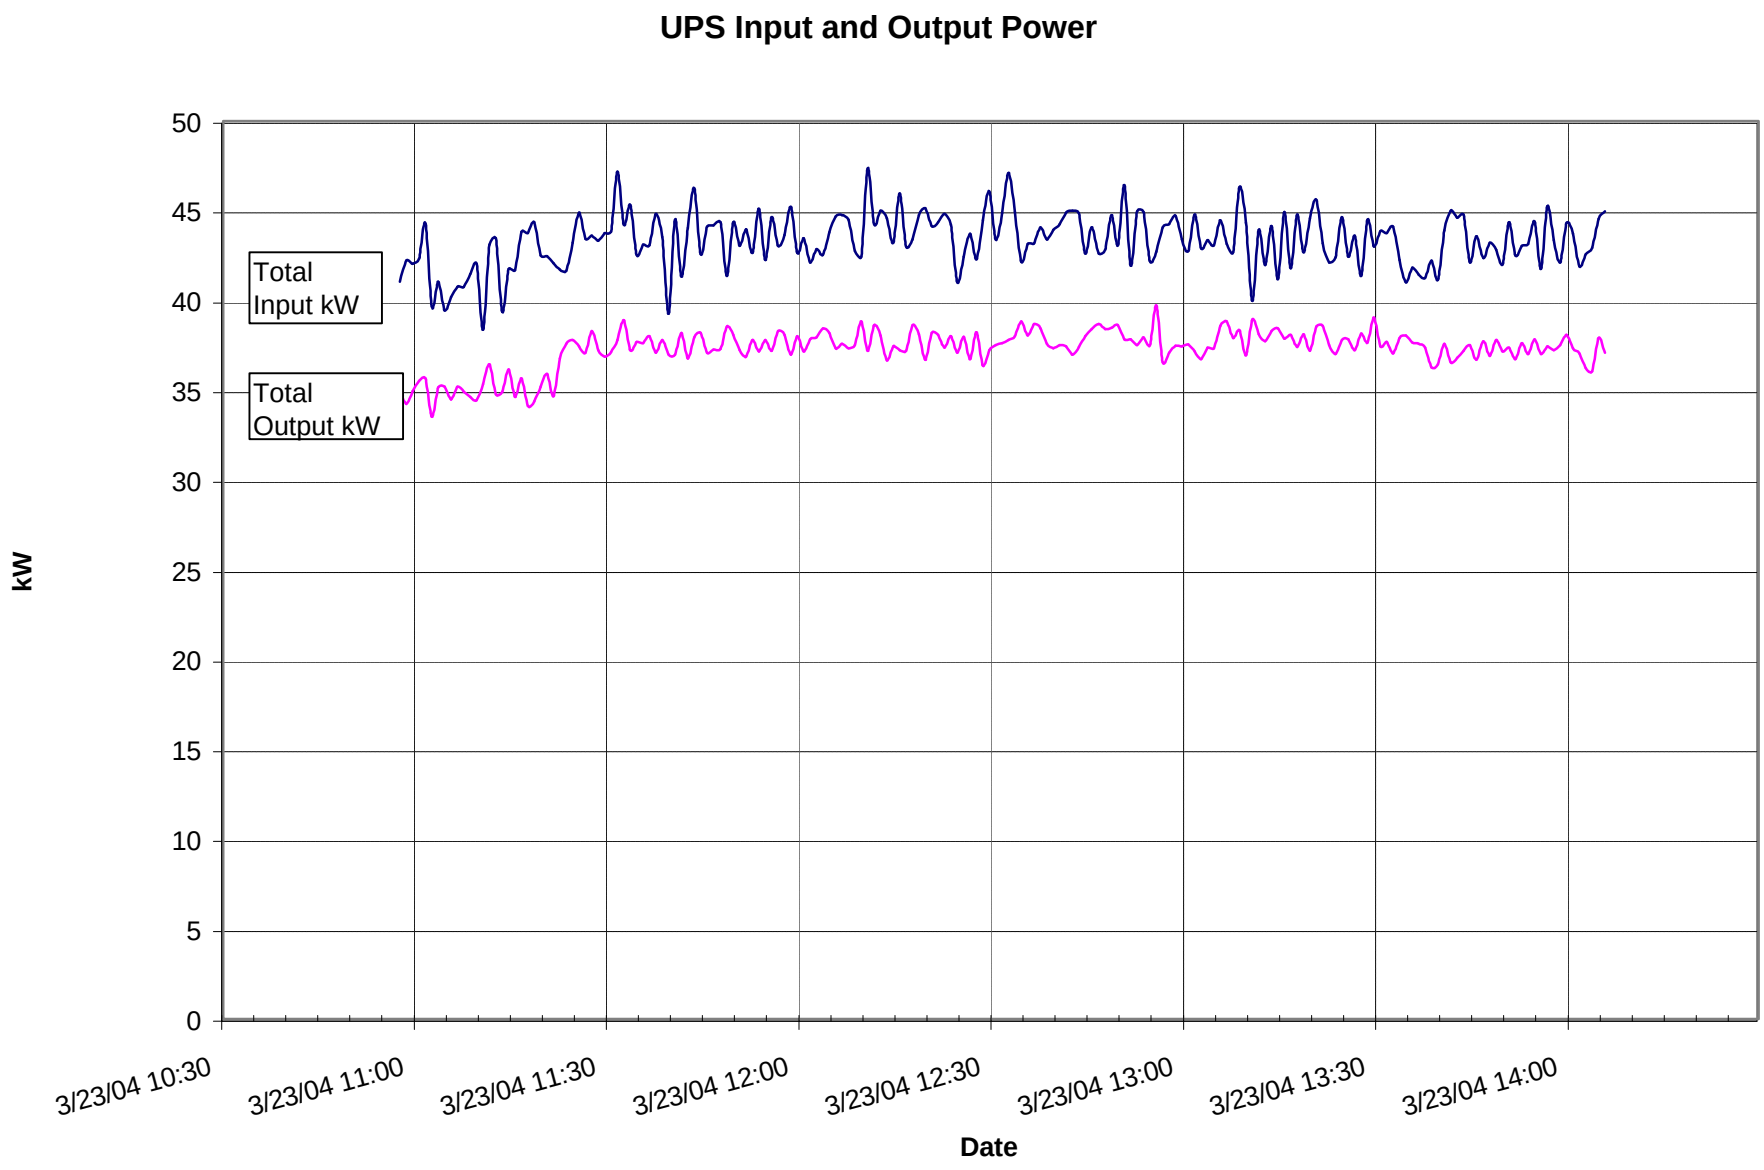
UPS Input and Output Power
Total Input kW
Total Output kW
kW
50
45
40
35
30
25
20
15
10
5
0
3/23/04 10:30
3/23/04 11:00
3/23/04 11:30
3/23/04 12:00
3/23/04 12:30
3/23/04 13:00
3/23/04 13:30
3/23/04 14:00
Date

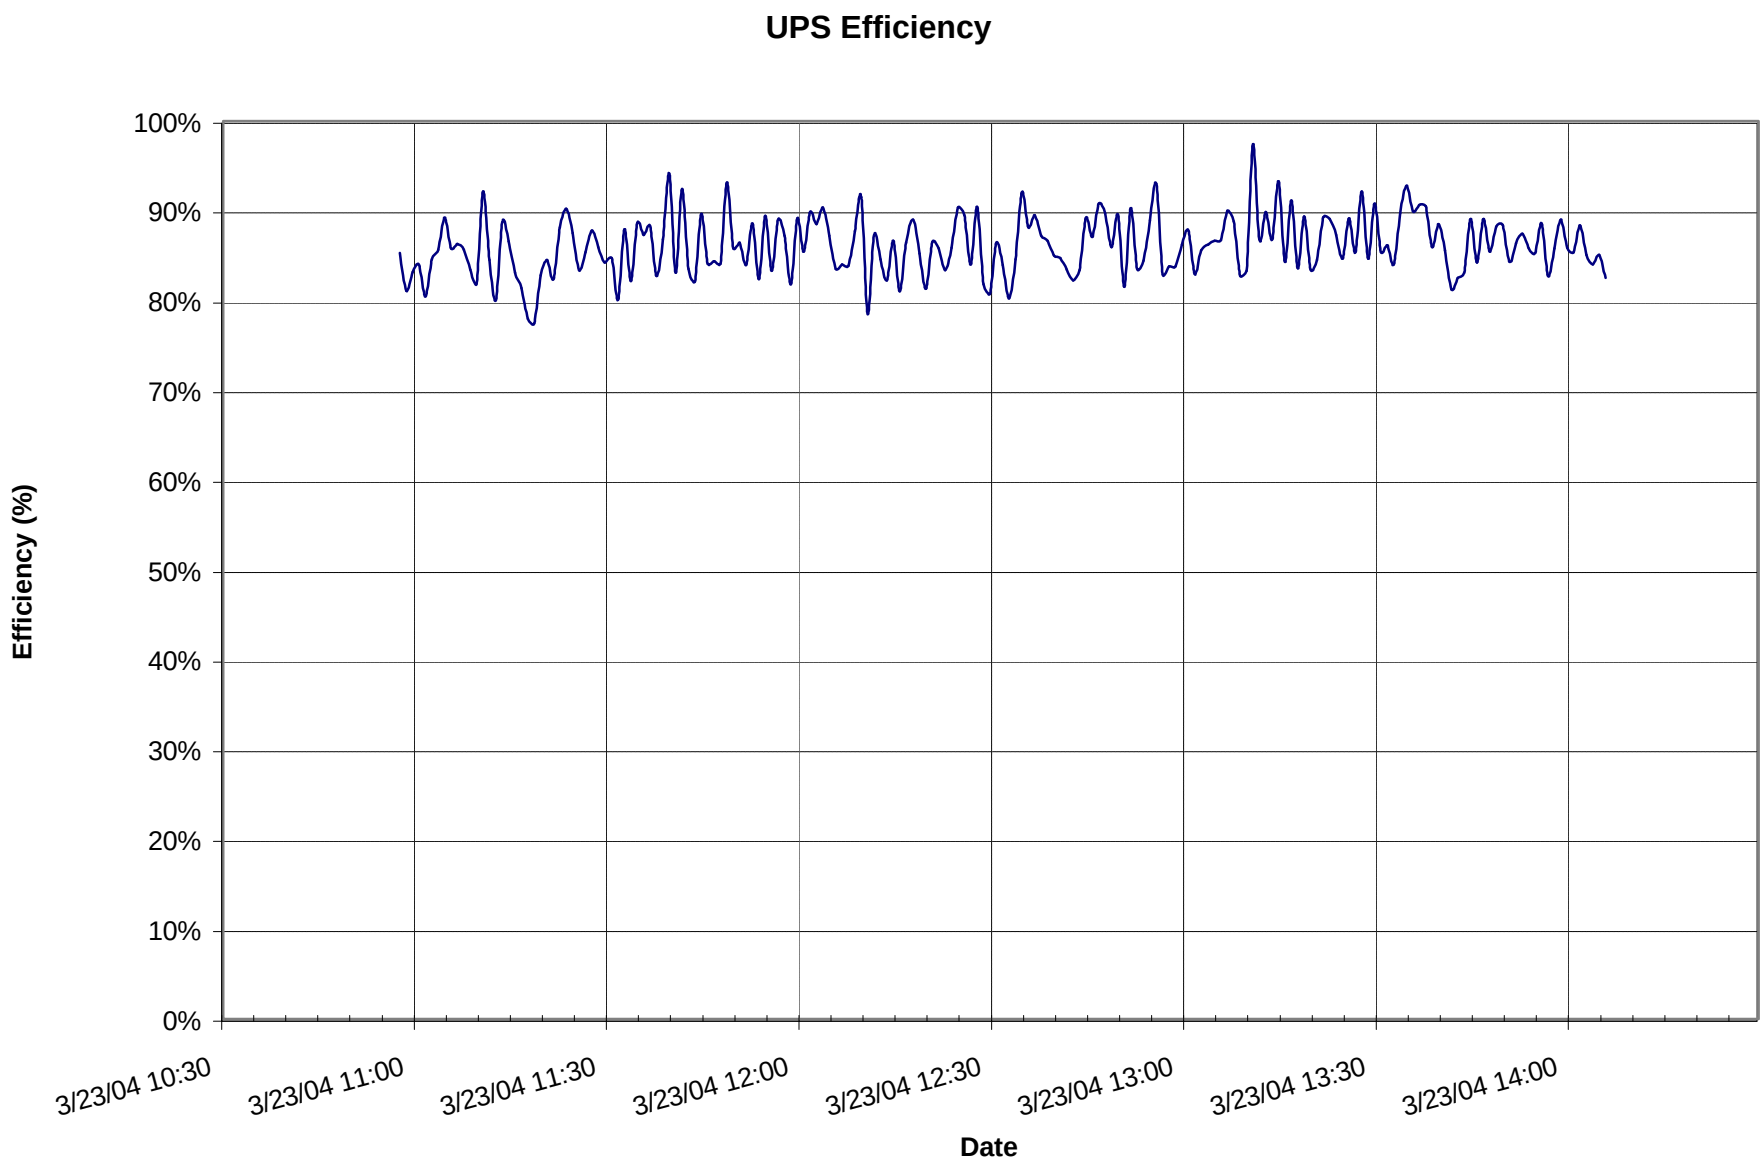
UPS Efficiency
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
Efficiency (%)
3/23/04 10:30
3/23/04 11:00
3/23/04 11:30
3/23/04 12:00
3/23/04 12:30
3/23/04 13:00
3/23/04 13:30
3/23/04 14:00
Date

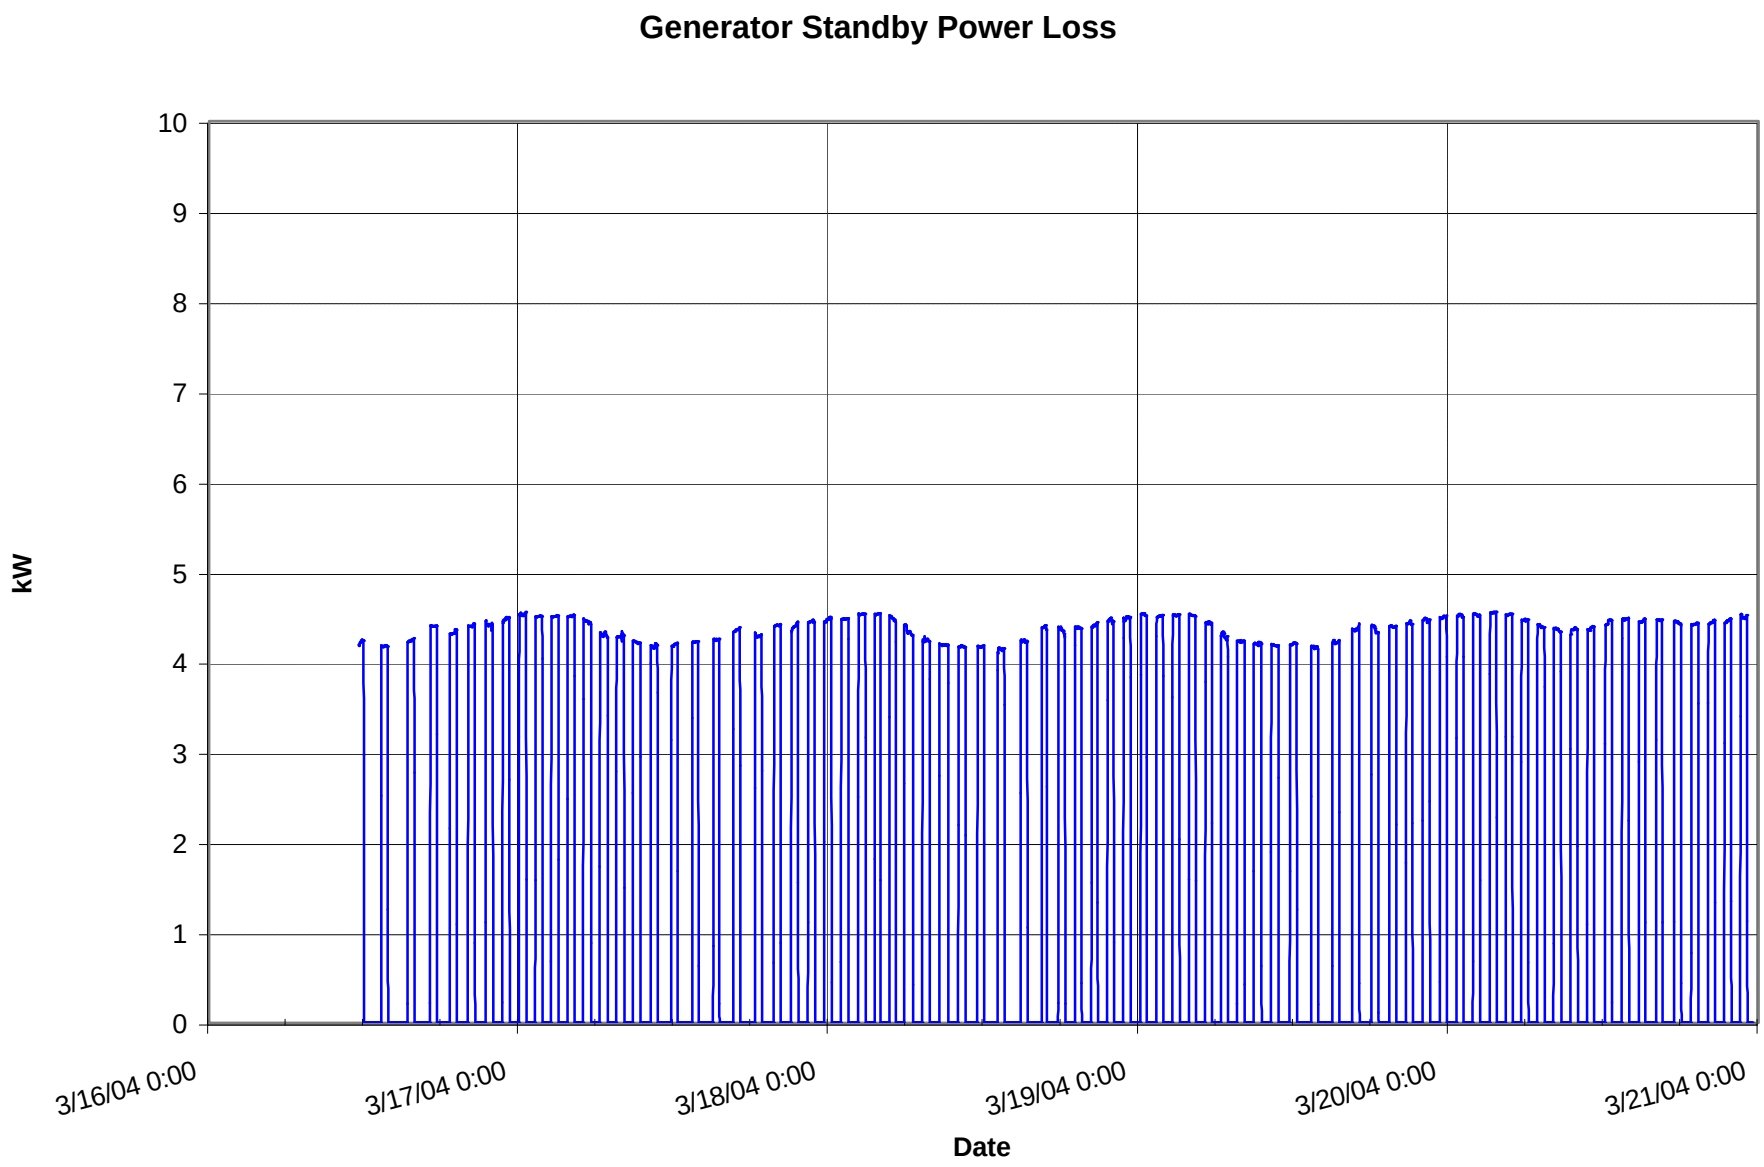
Generator Standby Power Loss
kW
10
9
8
7
6
5
4
3
2
1
0
3/16/04 0:00
3/17/04 0:00
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
Date

Chilled Water System

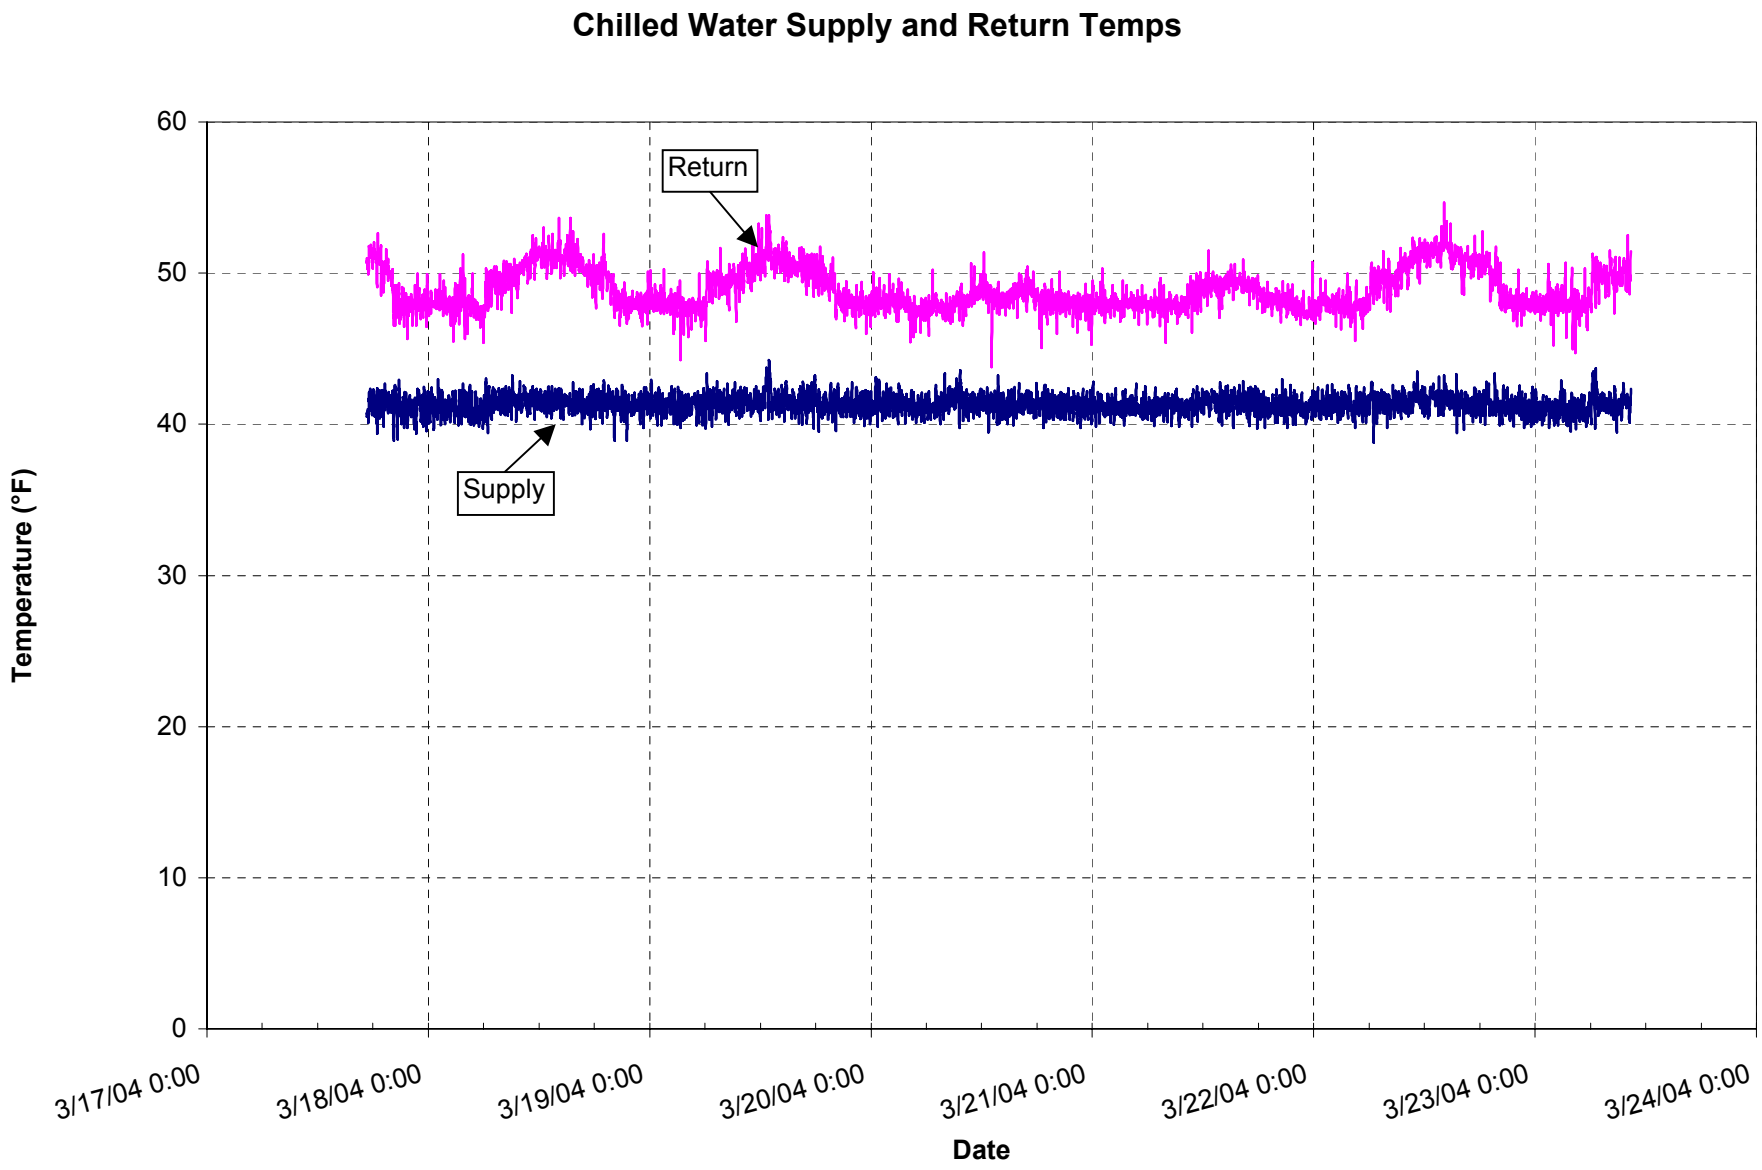
Chilled Water Supply and Return Temps
Return
Supply
Temperature (°F)
60
50
40
30
20
10
0
3/17/04 0:00
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
3/23/04 0:00
3/24/04 0:00
Date

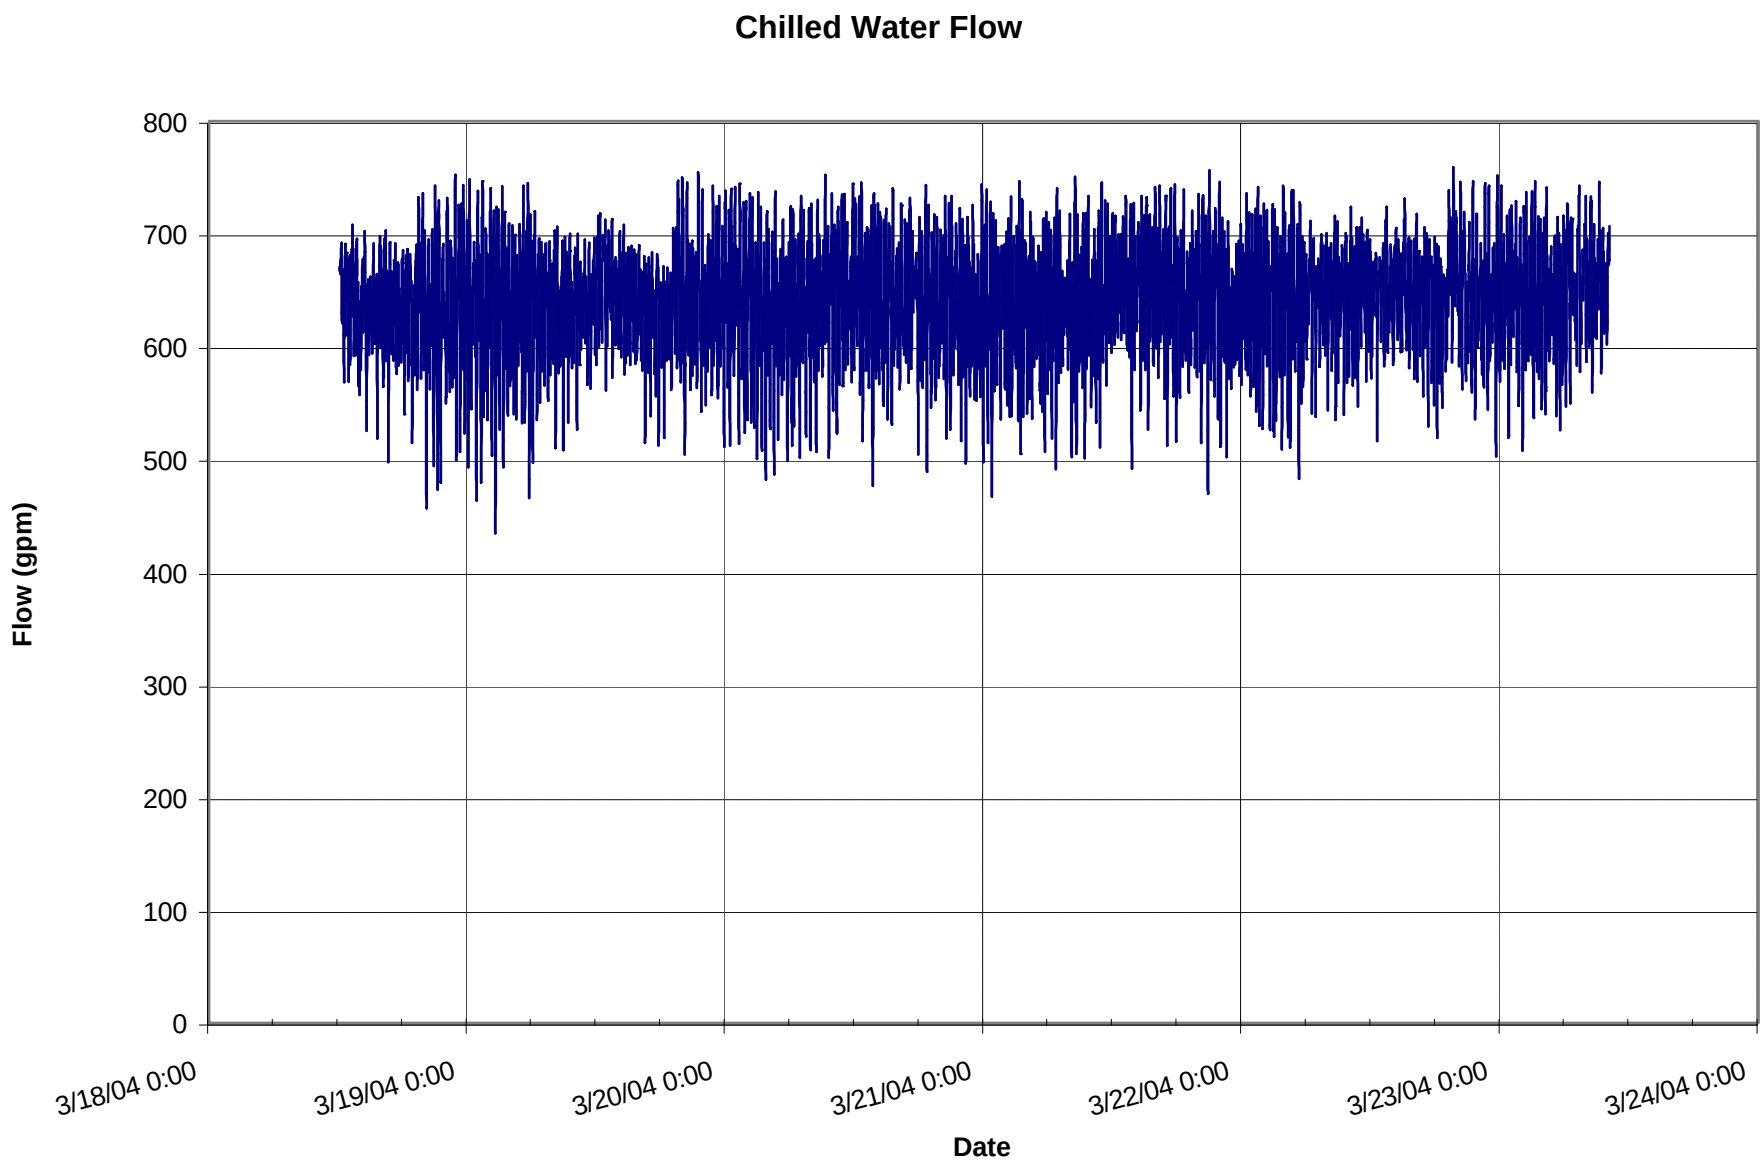
Chilled Water Flow
Flow (gpm)
800
700
600
500
400
300
200
100
0
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
3/23/04 0:00
3/24/04 0:00
Date

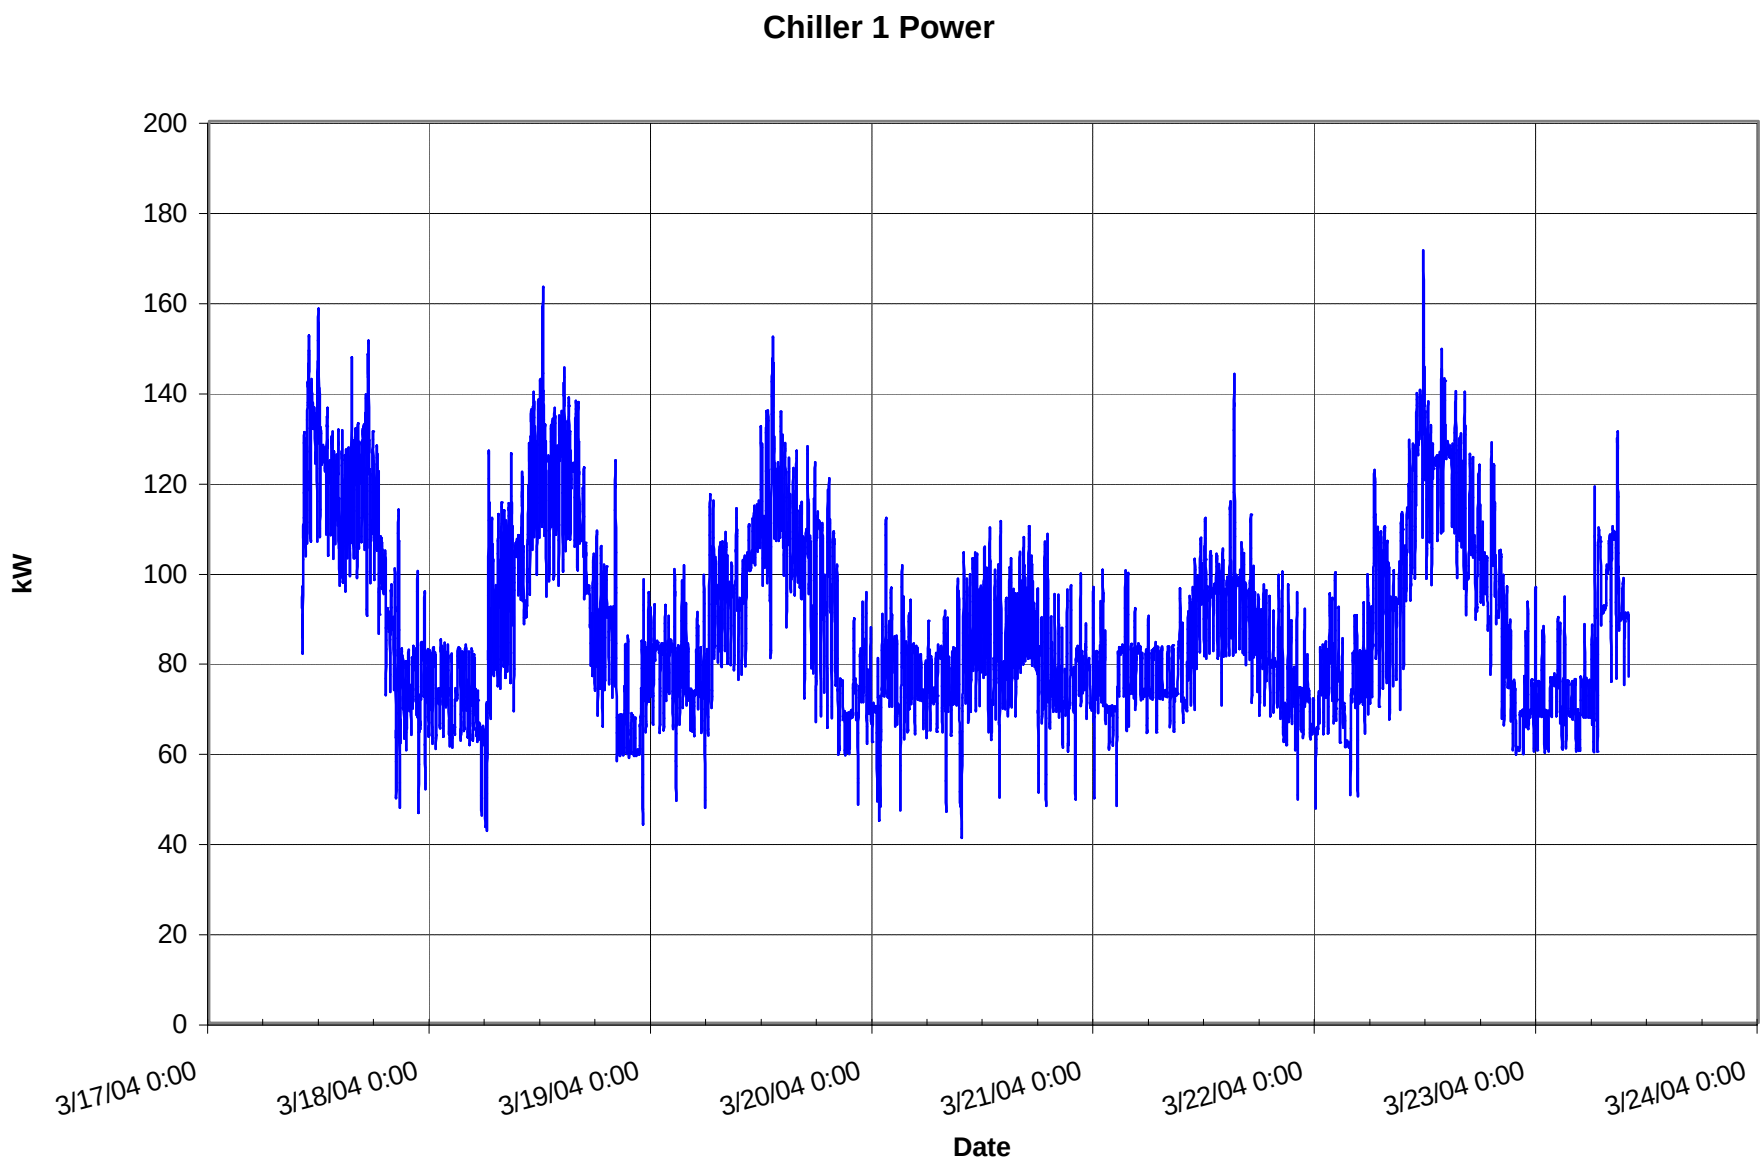
Chiller 1 Power
kW
200
180
160
140
120
100
80
60
40
20
0
3/17/04 0:00
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
3/23/04 0:00
3/24/04 0:00
Date

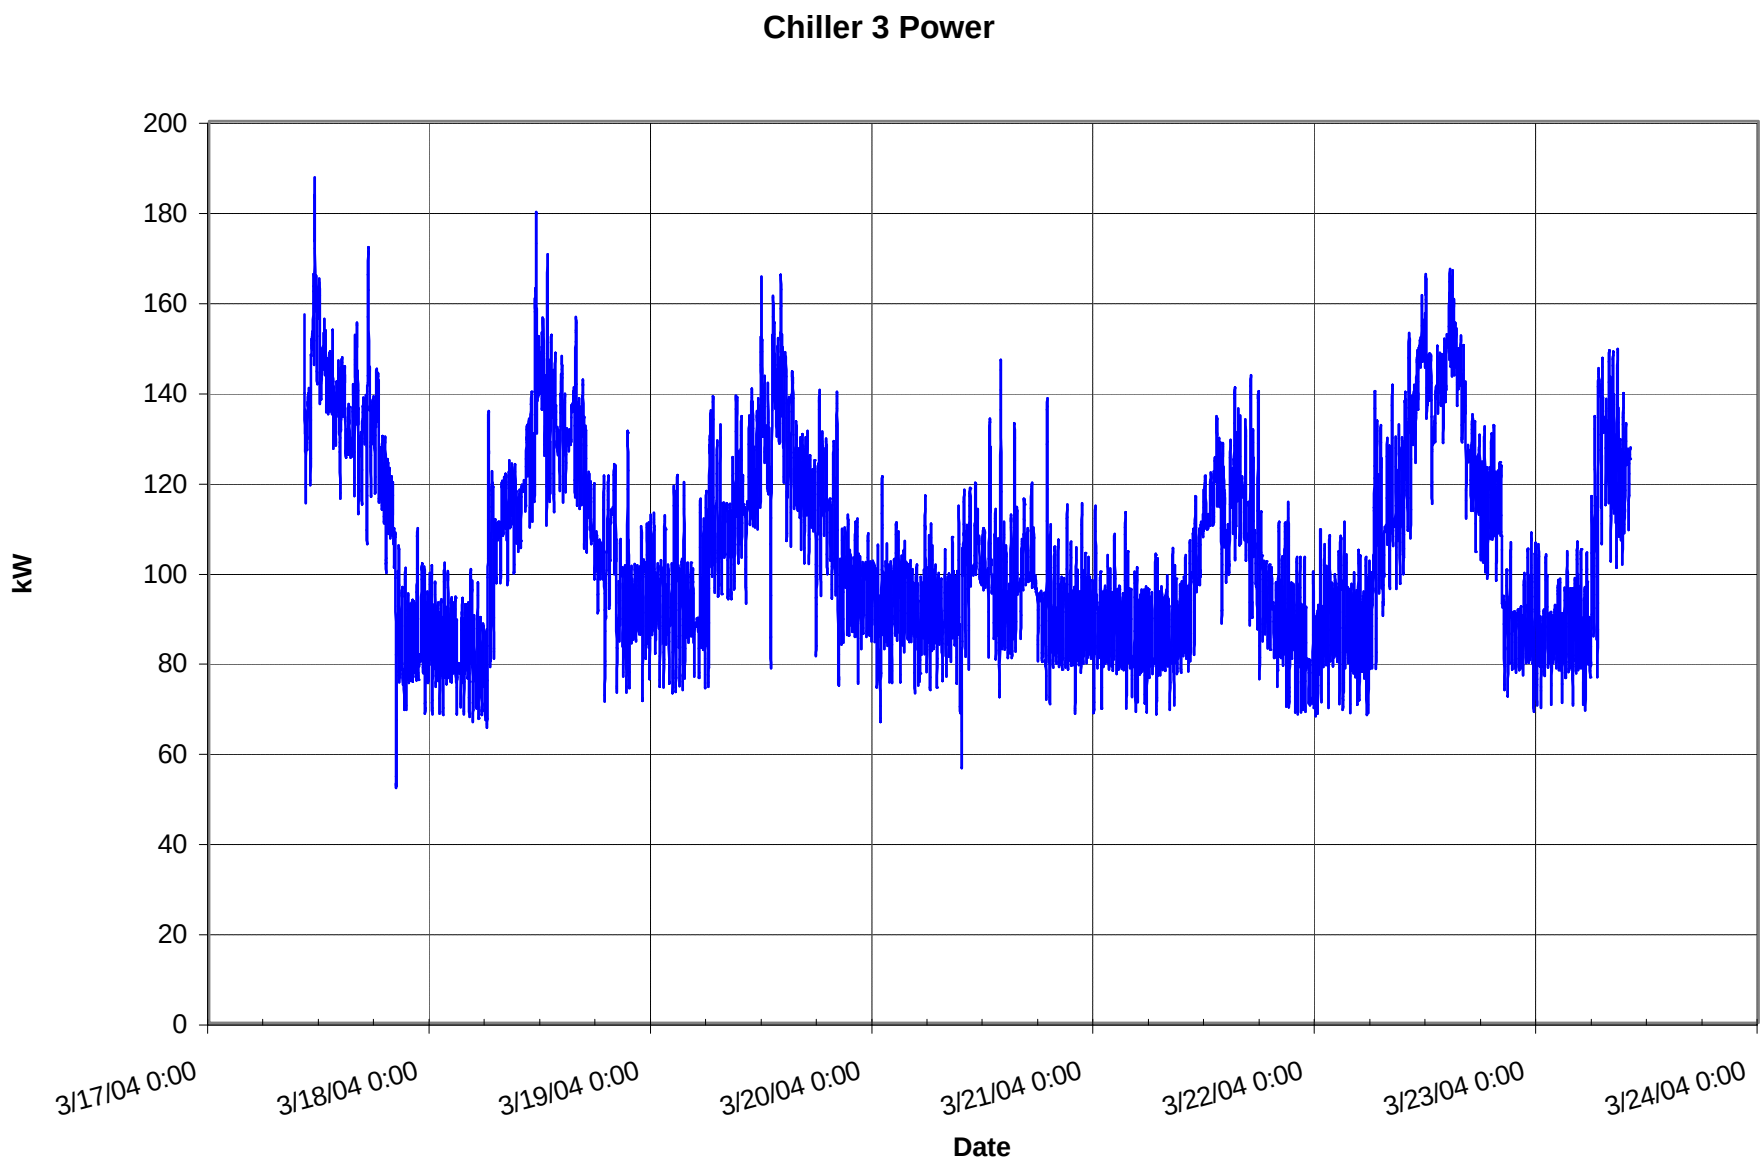
Chiller 3 Power
kW
200
180
160
140
120
100
80
60
40
20
0
3/17/04 0:00
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
3/23/04 0:00
3/24/04 0:00
Date

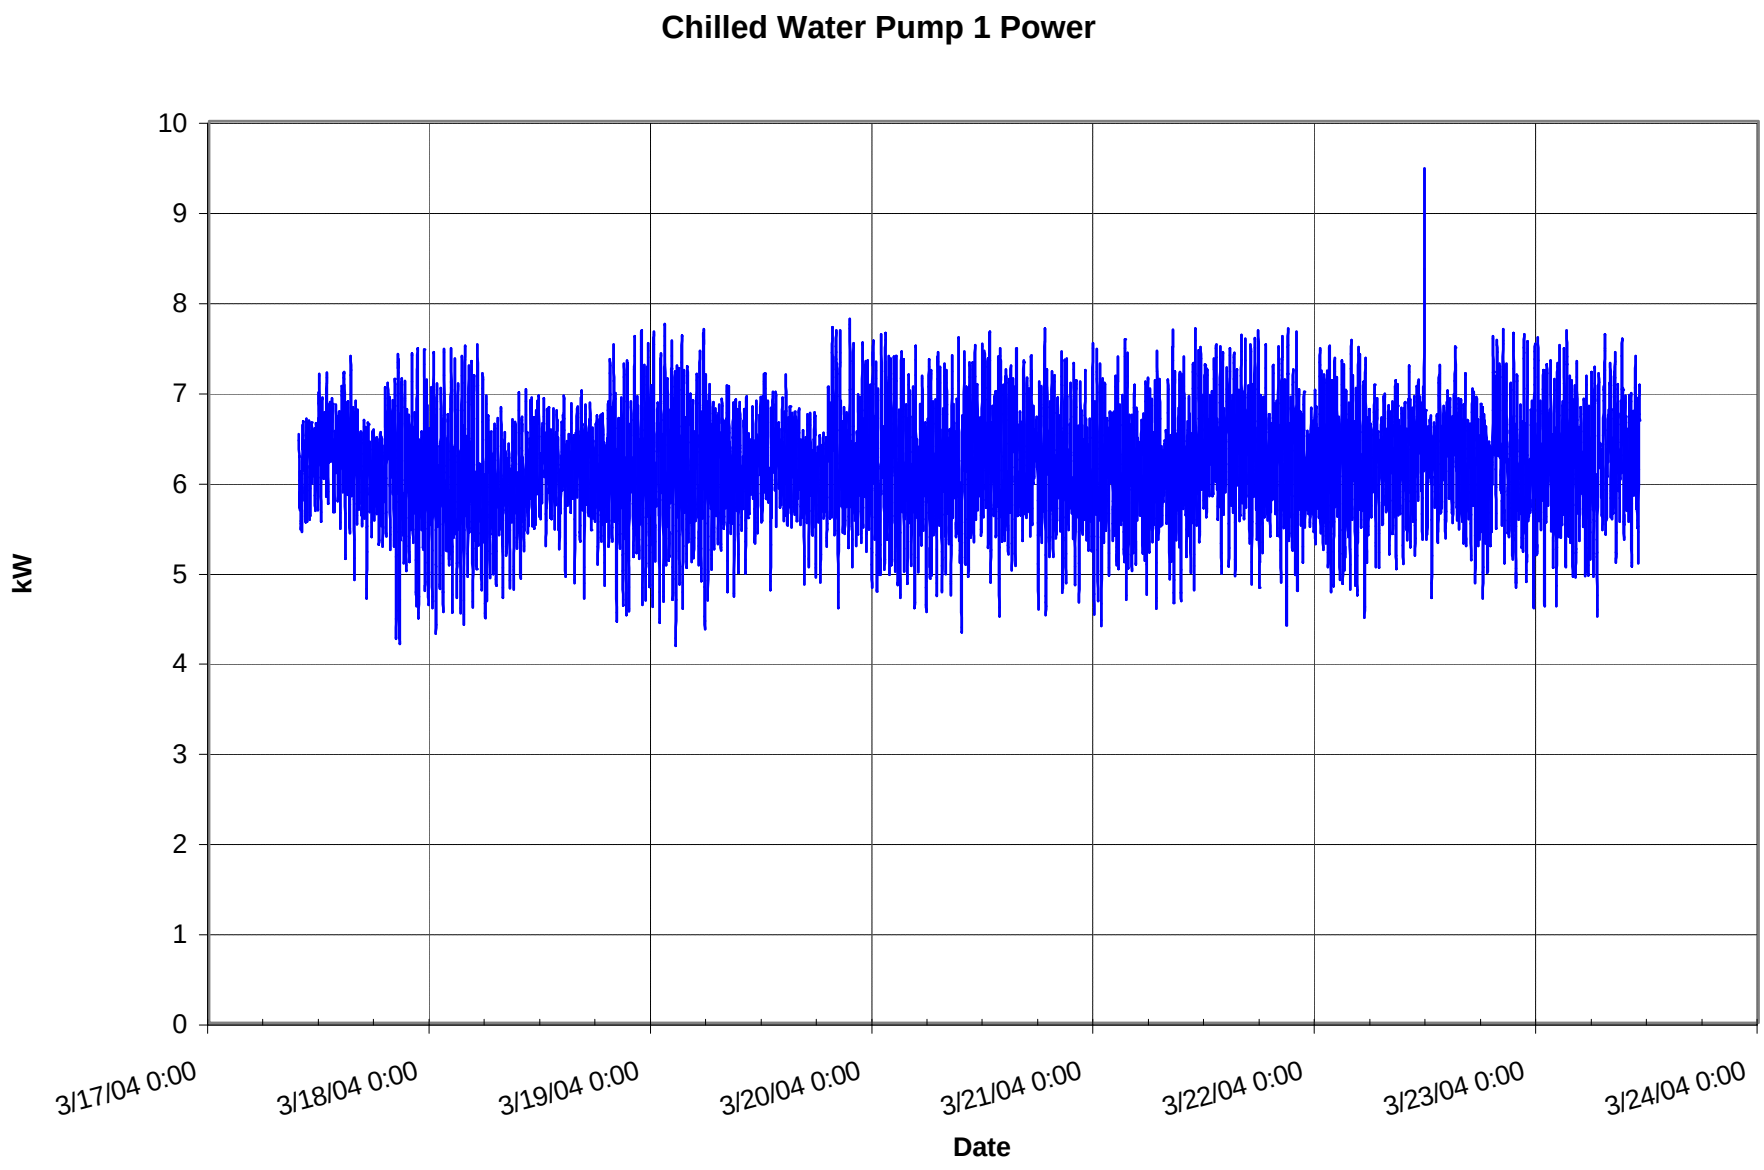
Chilled Water Pump 1 Power
kW
10
9
8
7
6
5
4
3
2
1
0
3/17/04 0:00
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
3/23/04 0:00
3/24/04 0:00
Date

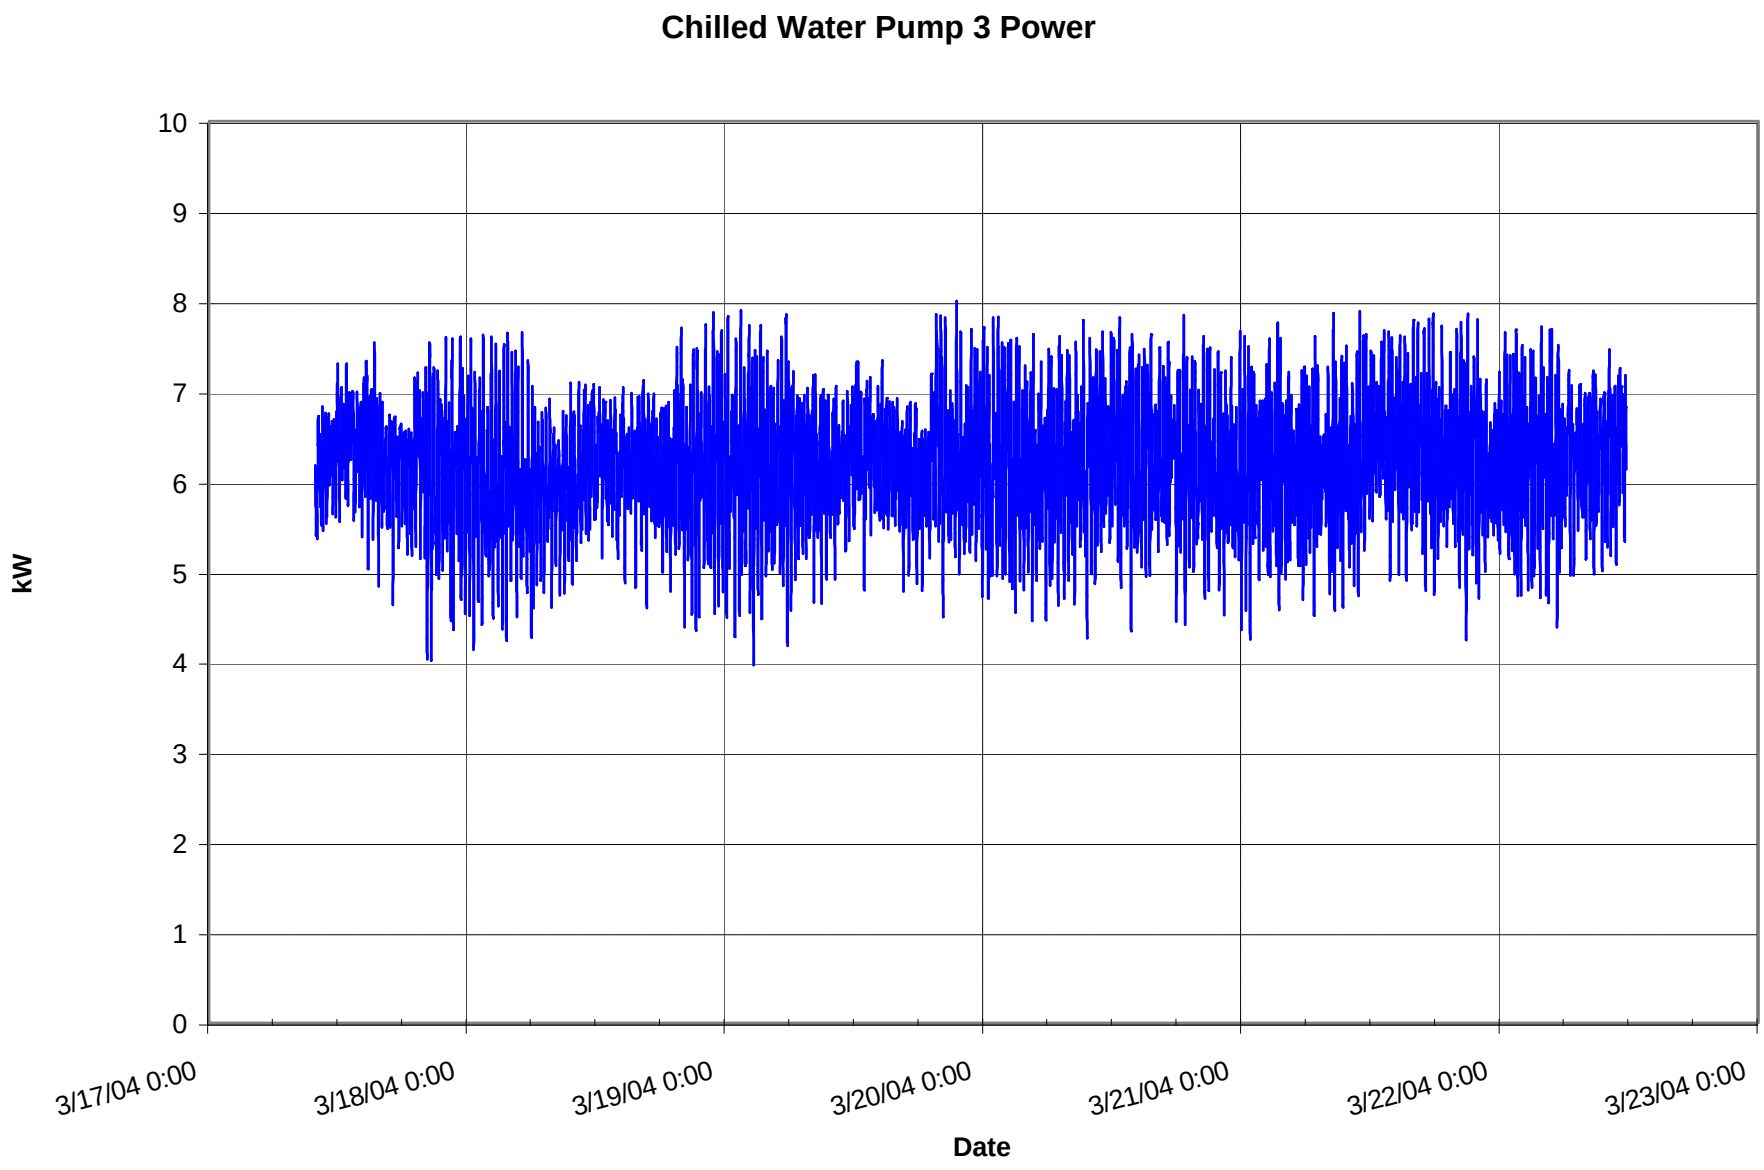
Chilled Water Pump 3 Power
kW
10
9
8
7
6
5
4
3
2
1
0
3/17/04 0:00
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
3/23/04 0:00
Date

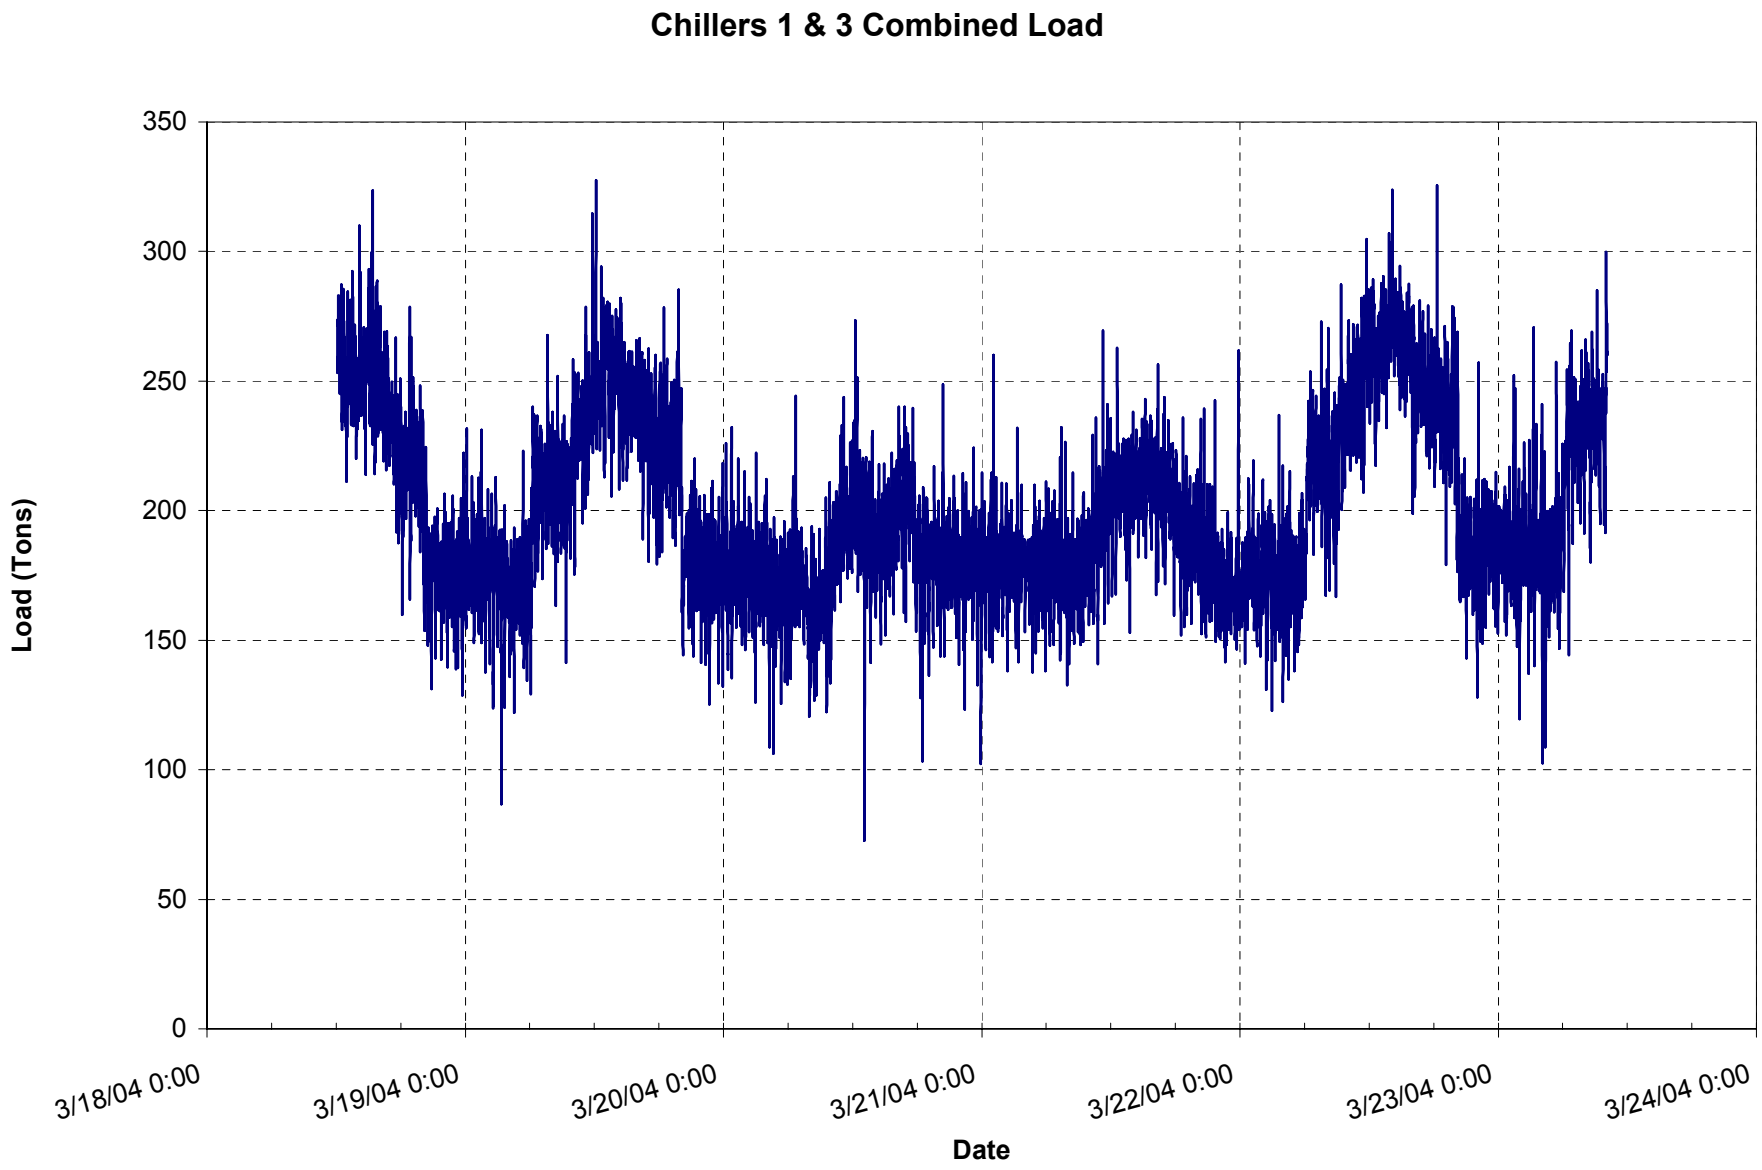
Chillers 1 & 3 Combined Load
Load (Tons)
350
300
250
200
150
100
50
0
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
3/23/04 0:00
3/24/04 0:00
Date

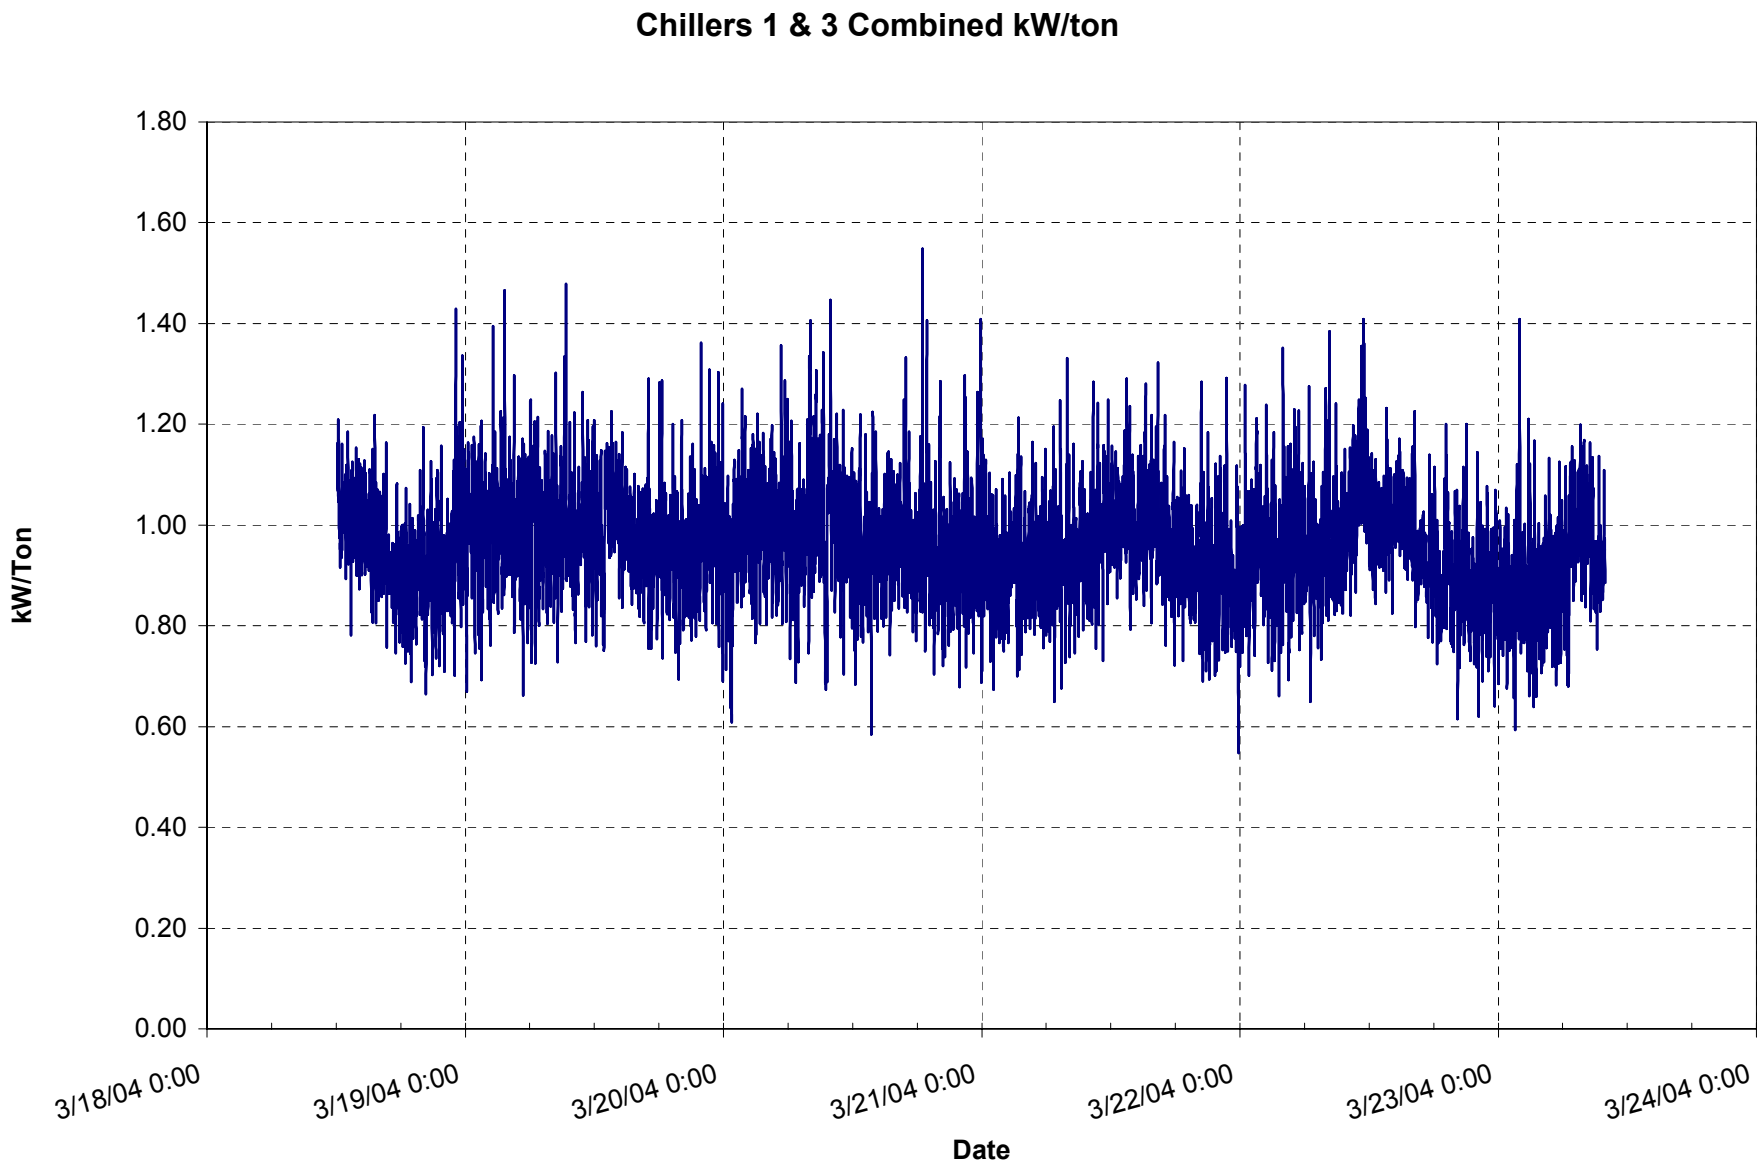
Chillers 1 & 3 Combined kW/ton
kW/Ton
1.80
1.60
1.40
1.20
1.00
0.80
0.60
0.40
0.20
0.00
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
3/23/04 0:00
3/24/04 0:00
Date

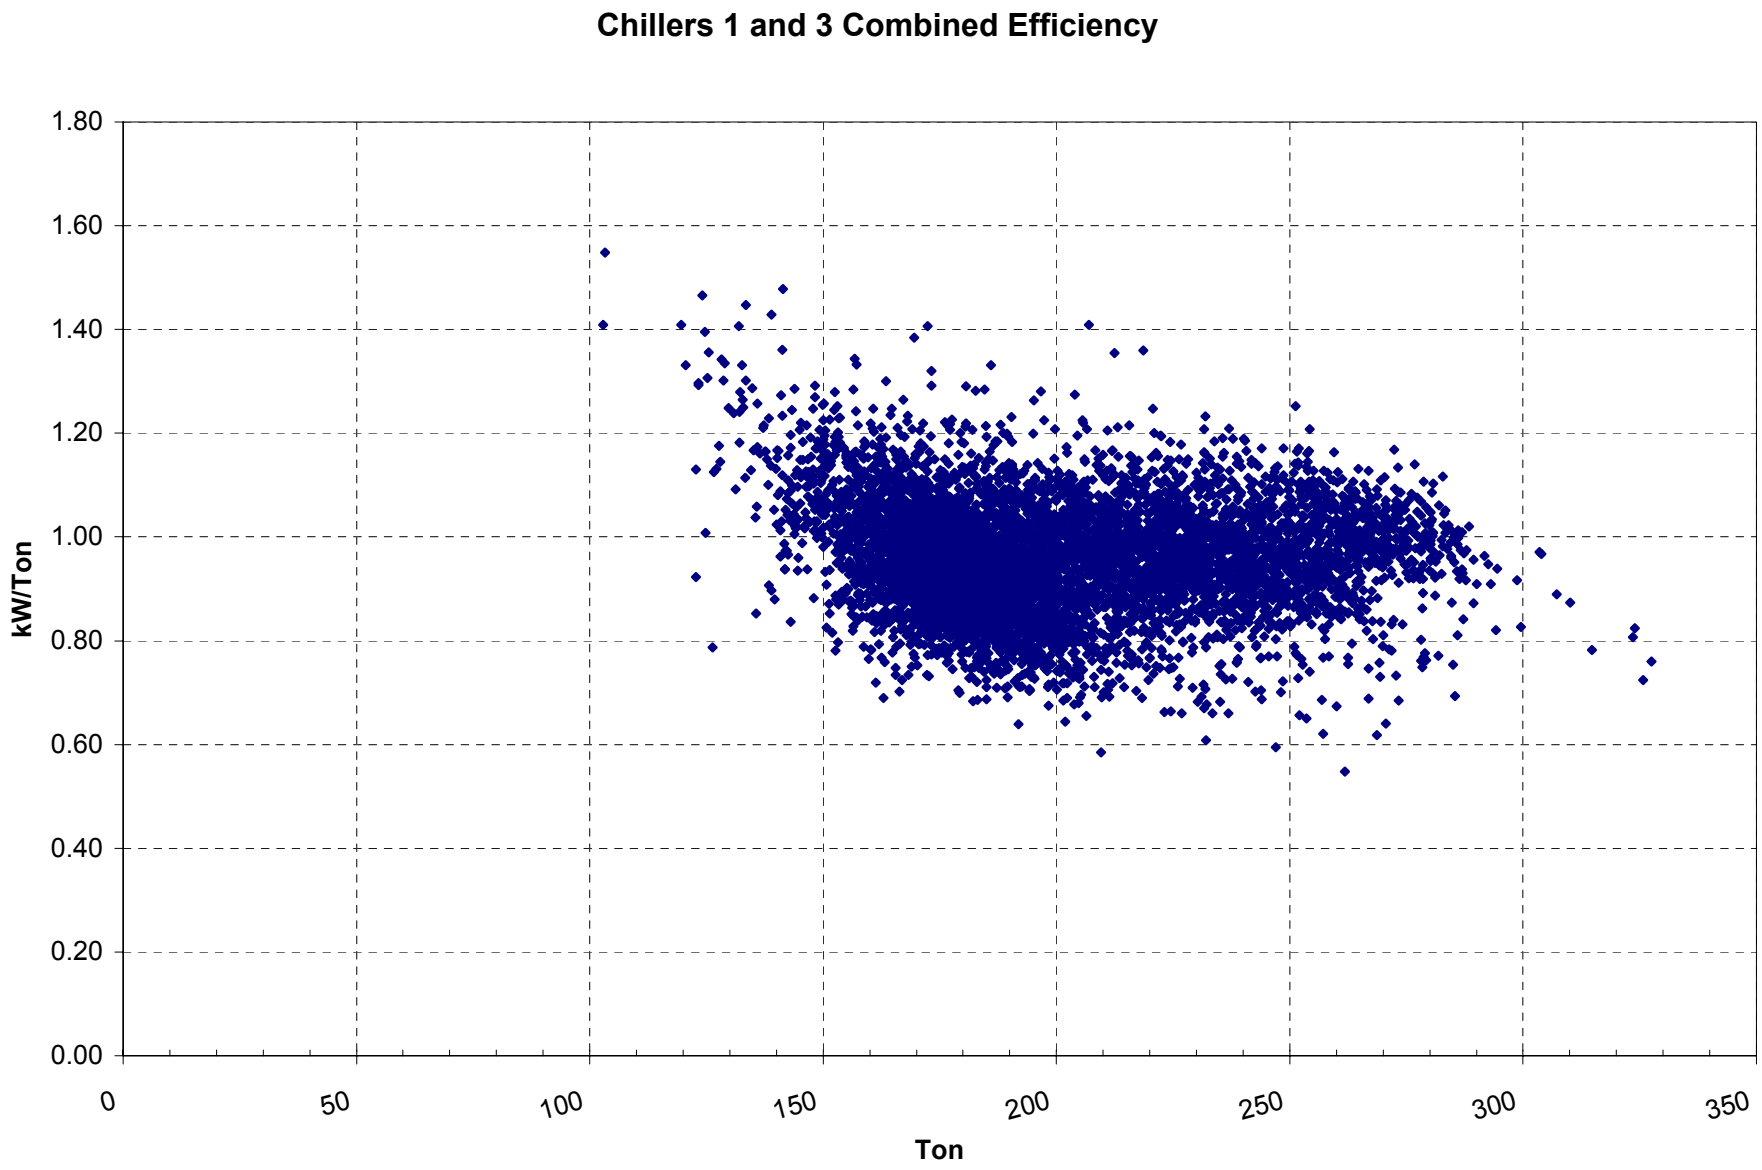
Chillers 1 and 3 Combined Efficiency
kW/Ton
1.80
1.60
1.40
1.20
1.00
0.80
0.60
0.40
0.20
0.00
0
50
100
150
200
250
300
350
Ton

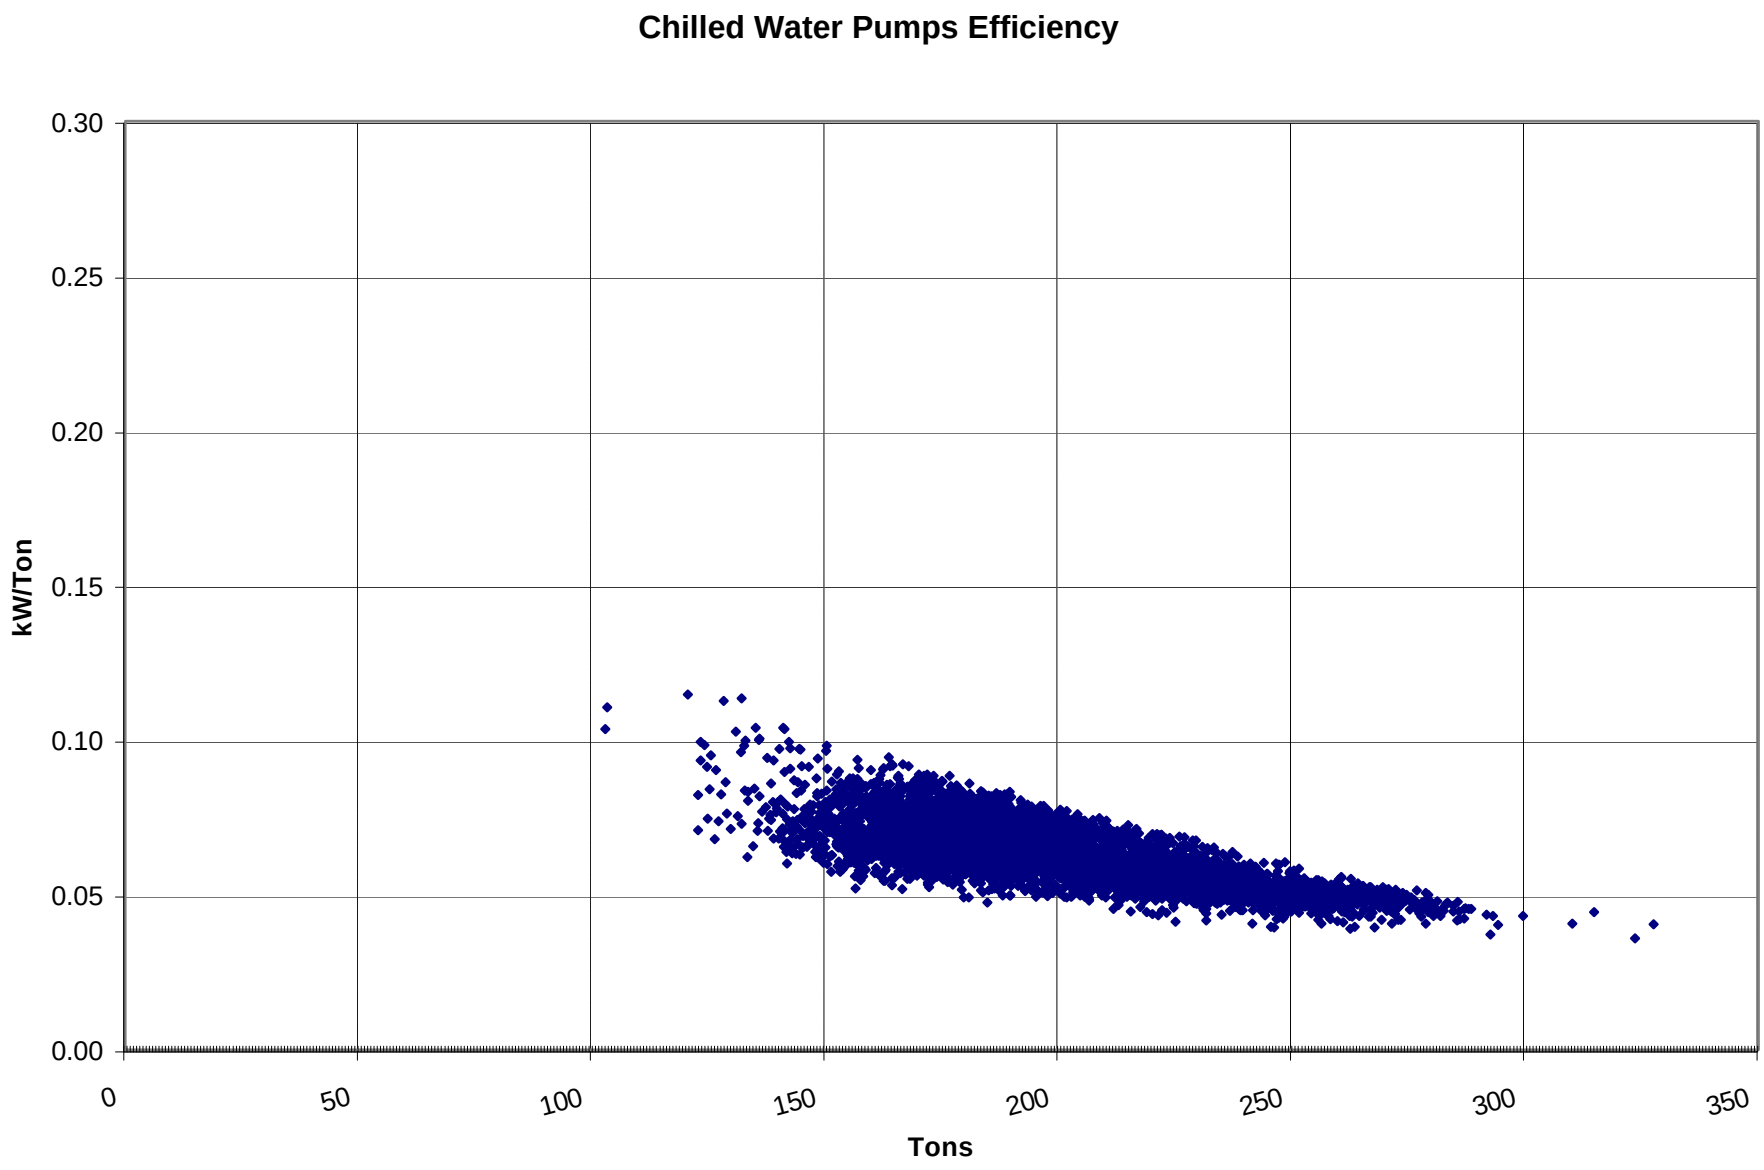
Chilled Water Pumps Efficiency
0.30
0.25
0.20
0.15
0.10
0.05
0.00
kW/Ton
0
50
100
150
200
250
300
350
Tons

Process Chilled Water System

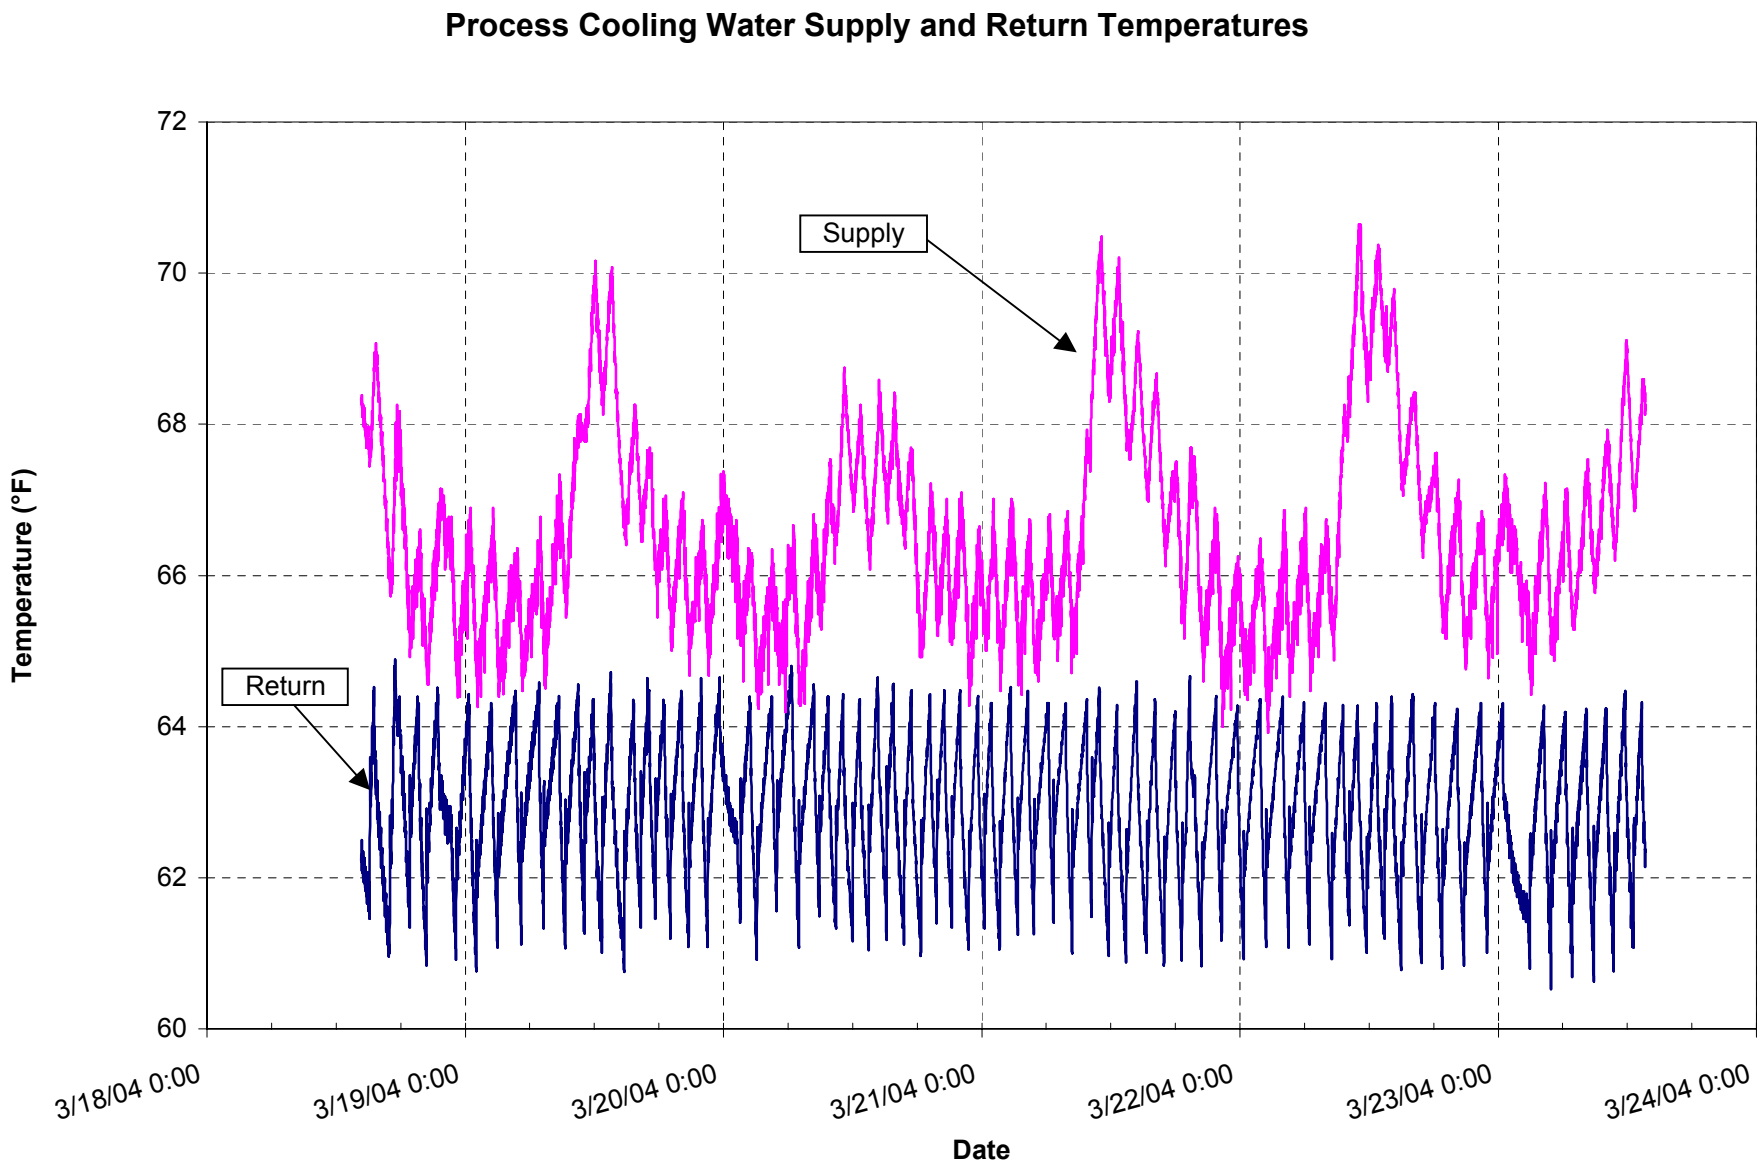
Process Cooling Water Supply and Return Temperatures
Supply
Return
Temperature (°F)
72
70
68
66
64
62
60
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
3/23/04 0:00
3/24/04 0:00
Date

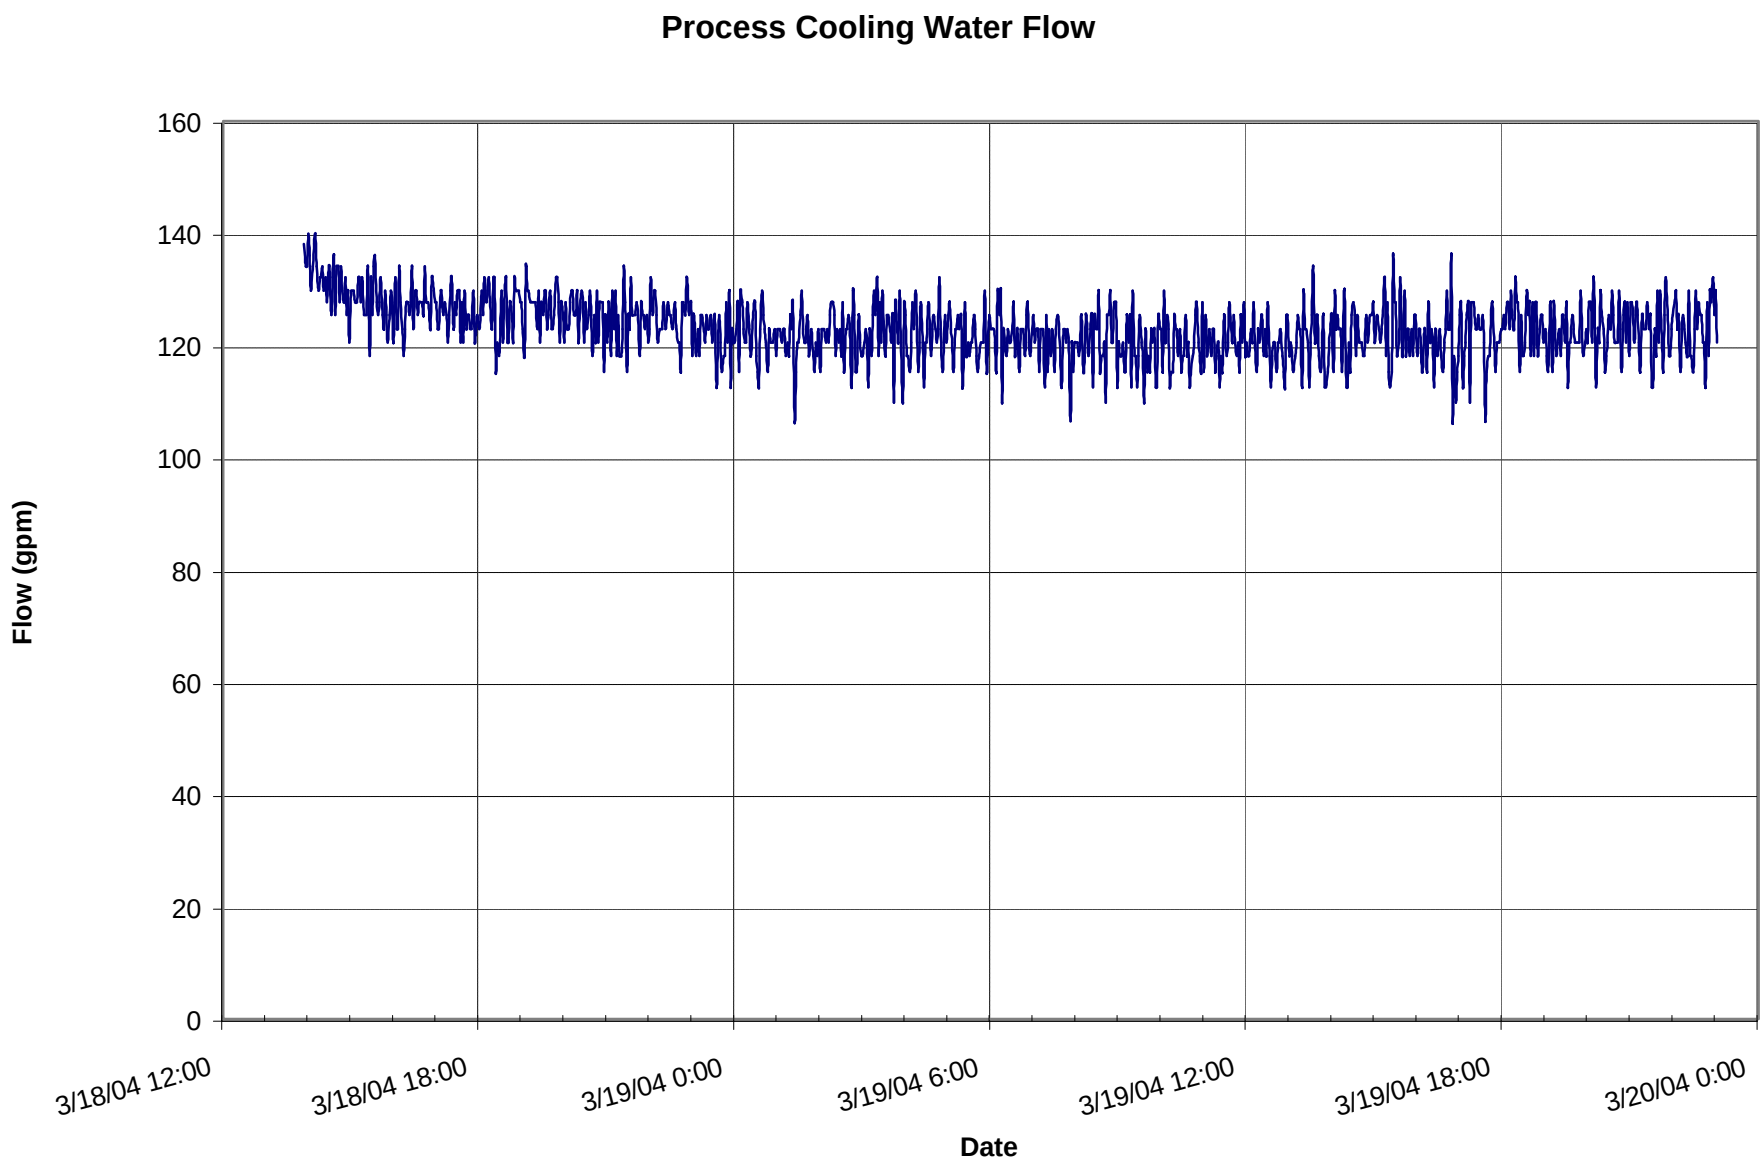
Process Cooling Water Flow
Flow (gpm)
160
140
120
100
80
60
40
20
0
3/18/04 12:00
3/18/04 18:00
3/19/04 0:00
3/19/04 6:00
3/19/04 12:00
3/19/04 18:00
3/20/04 0:00
Date

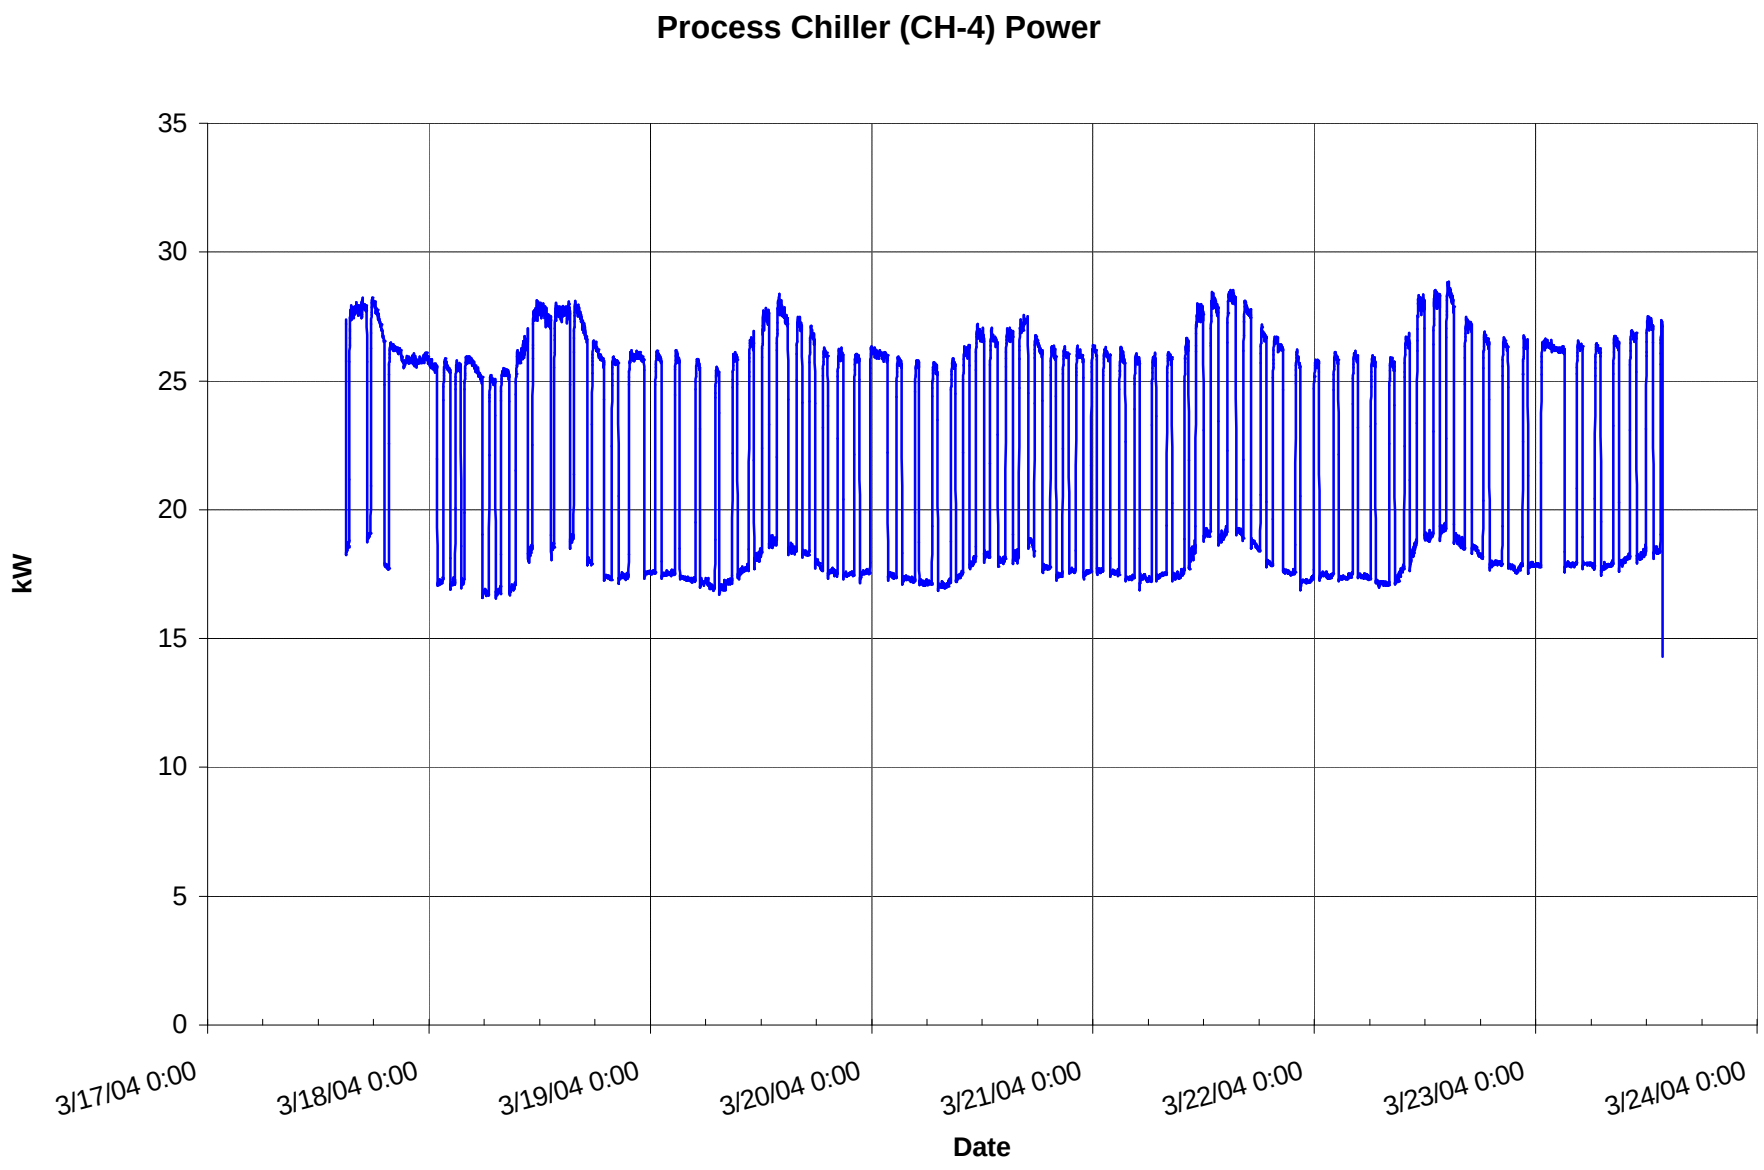
Process Chiller (CH-4) Power
35
30
25
20
15
10
5
0
kW
3/17/04 0:00
3/18/04 0:00
3/19/04 0:00
3/20/04 0:00
3/21/04 0:00
3/22/04 0:00
3/23/04 0:00
3/24/04 0:00
Date

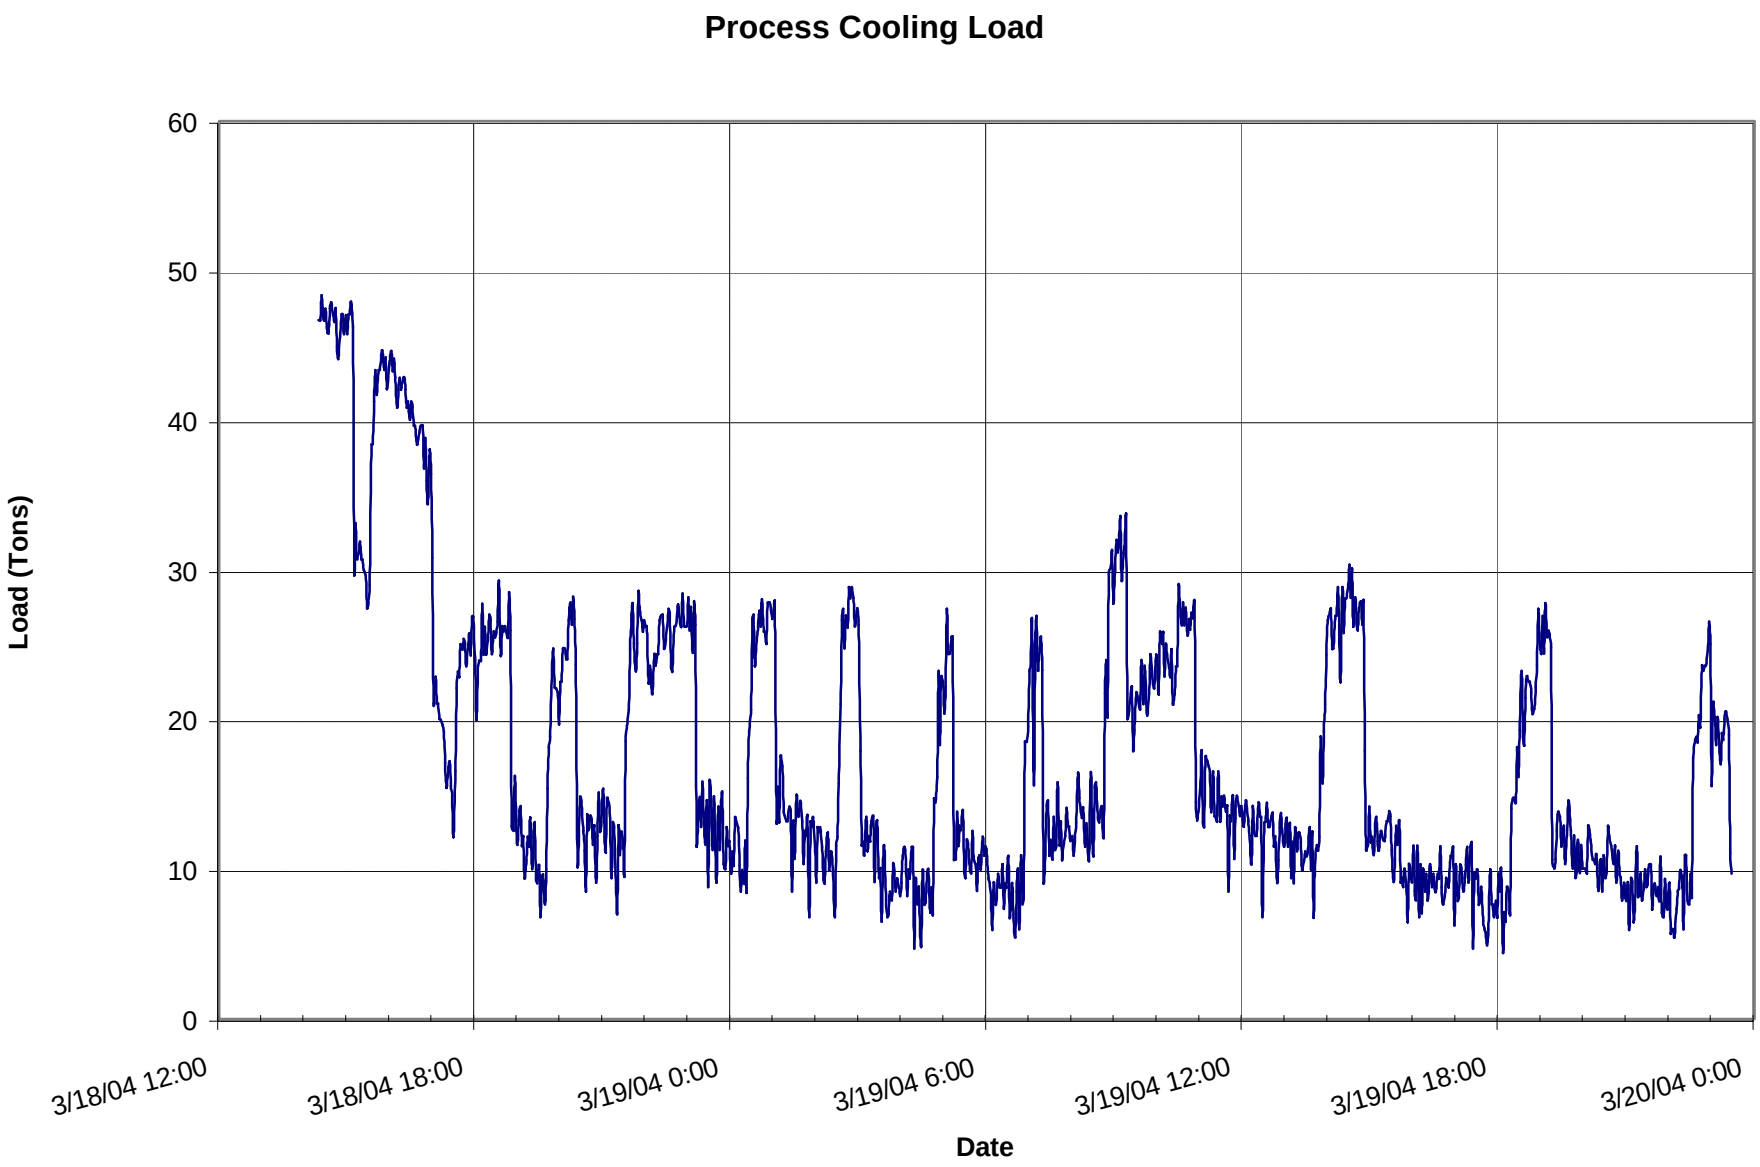
Process Cooling Load
Load (Tons)
60
50
40
30
20
10
0
3/18/04 12:00
3/18/04 18:00
3/19/04 0:00
3/19/04 6:00
3/19/04 12:00
3/19/04 18:00
3/20/04 0:00
Date

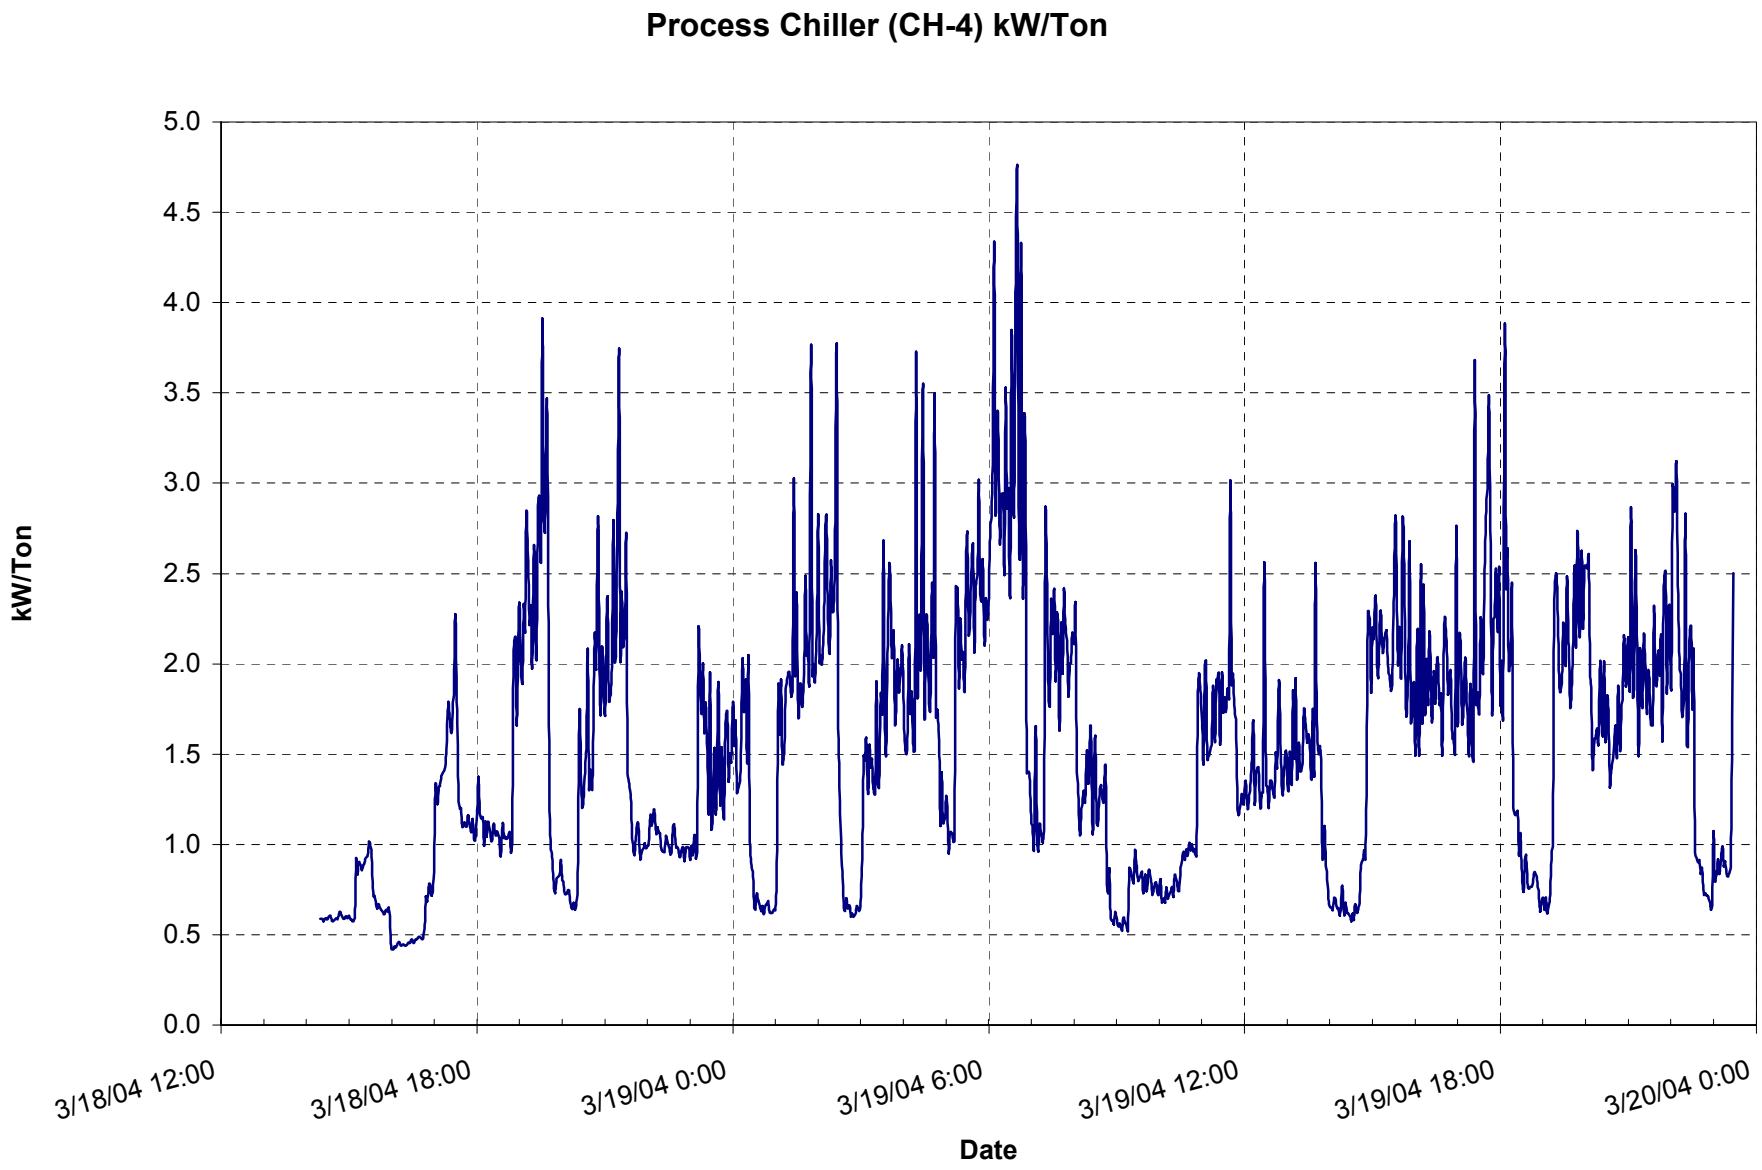
Process Chiller (CH-4) kW/Ton
kW/Ton
5.0
4.5
4.0
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
3/18/04 12:00
3/18/04 18:00
3/19/04 0:00
3/19/04 6:00
3/19/04 12:00
3/19/04 18:00
3/20/04 0:00
Date

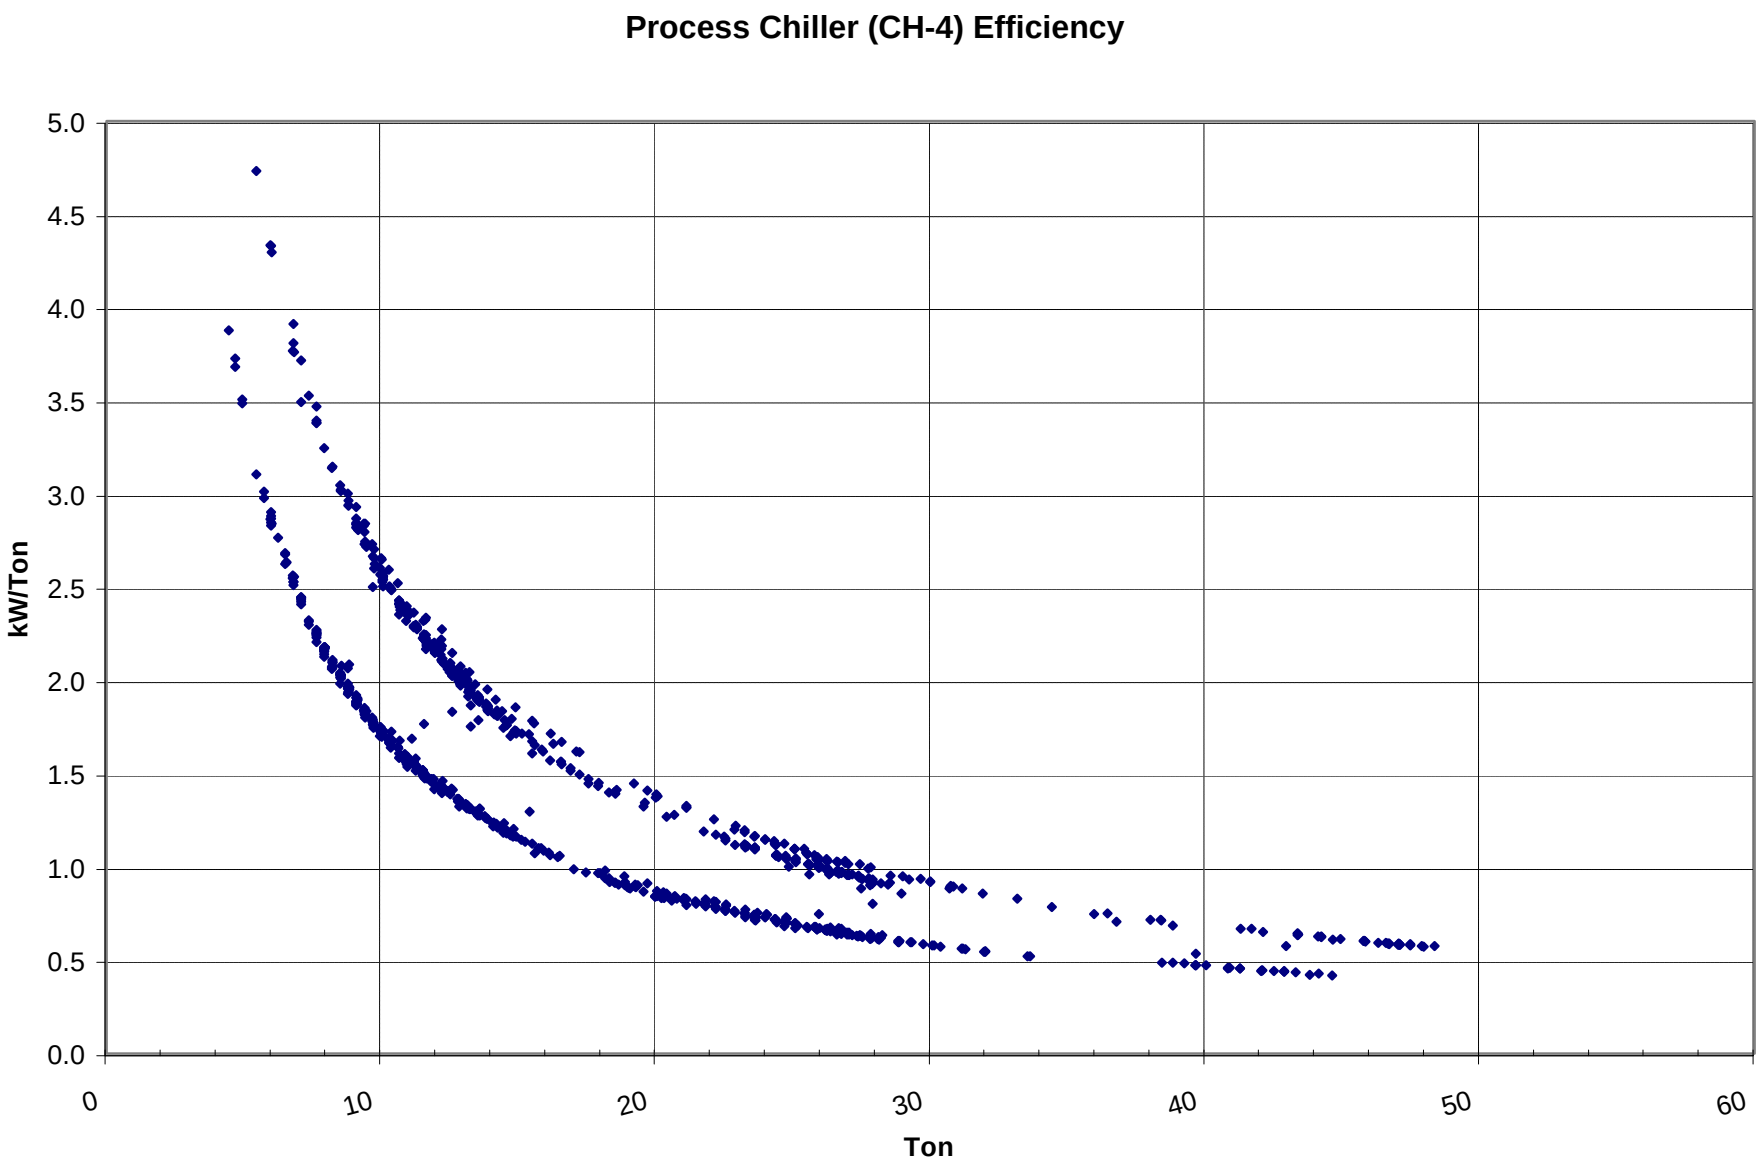
Process Chiller (CH-4) Efficiency
kW/Ton
5.0
4.5
4.0
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
0
10
20
30
40
50
60
Ton

Appendix C

Data Collection and Accuracy Notes

Data Collection and Accuracy Notes

The following notes describe specific measurements and what assumptions were made in preparing calculated results.

Cleanroom Data

Lighting Power	Lighting loads measured directly at their respective breakers.
Process Power	Process tool loads measured directly at their respective breakers.
Average Filter Velocity	Filter velocity based on airflow divided by the assumed effective HEPA filter area of 6.8 square feet.
Primary Cleanroom Area	The area certified at a rated class level. This is taken from the drawings.
Secondary Cleanroom Area	Air return area.

Recirculation Air

Air Flow	All air flow measurements were provided by certification report.
Fan Power	Fan power was measured directly
RCU Efficiency	Number of cubic feet of recirculation air delivered to a given cleanroom, divided by the total kW of the units providing the recirculation air.

Make-Up Air Handling Unit (MUAH)

Air Flow	
Fan Power	Fan power was measured directly.
MUAH Efficiency	Number of cubic feet of air delivered per kW.

Chiller

Total Power	Chiller loads measured directly at their respective breakers.
Total Cooling Supplied	Determined by direct flow and temperature measurements of supply and return chilled water at the common header of the three chillers. Accuracy of this calculation (20%) reflects the uncertainty of the chilled water flow.
Cooling Tons	Standard engineering calculation based on temperature and flow.
Efficiency Metric	Amount of chiller power (kW) per ton of cooling supplied by the chilled water plant.

Pumps

Total Power	Pump loads measured directly at their respective breakers.
Efficiency Metric	Amount of pumping power (kw) per ton of cooling supplied by the chilled water plant.

Utility Billing

Annual Electric Use	Use for the entire site is taken from the most recent complete calendar year (2003) bills as provided.
Annual Electric Cost	Same as above.
Annual Natural Gas Use	Same as above.
Annual Natural Gas Cost	Same as above.

Appendix D

Measurement Methodology

Measurement Methodology

Data collection measurements were made according to the following procedures:

Ultrasonic Flowmeter – Controlotron

- Equipment: Controlotron Uniflow 1010
- Identify straight pipe run.
- Remove pipe insulation
- Clean pipe surface.
- Program flow meter with pipe characteristics.
- Measure wall thickness using thickness gauge.
- Verify meter setup.
- Plug meter into AC outlet.
- Secure transducer rails on pipe.
- Clean transducer surfaces.
- Apply gel to transducers.
- Secure transducers on rails.
- Read real time data. Verify flow conditions.
- Log data.

Water Temperature – Pete's Plug

- Equipment: Pace Scientific 4-channel pocket loggers model XR440, and 4" 30kOhm thermistors.
- Attach thermistor temperature sensor to pocket logger channel block.
- Setup pocket logger using product software.
- Verify channels set to correct sensor type and operation.
- Upload setup to pocket logger to launch logging.
- Insert thermistor into Pete's plug.
- Secure pocket logger to pipe.
- Read real time data. Verify setup and actual conditions.
- Log data.

Air Temperature/Relative Humidity – AHUs and Cleanrooms

- Equipment: Pace Scientific 4-channel pocket loggers model XR440, 4" 30kOhm thermistors, or Temperature/Relative Humidity sensors.
- Attach temperature/relative humidity sensor to pocket logger channel block.
- Setup pocket logger using product software.
- Verify channels set to correct sensor type and operation.
- Set RH linear scale specified by the sensor.
- Upload setup into pocket logger to launch logging.
- Read real time data. Verify setup and actual conditions.
- Log data.

Power Trend – Elite Logger

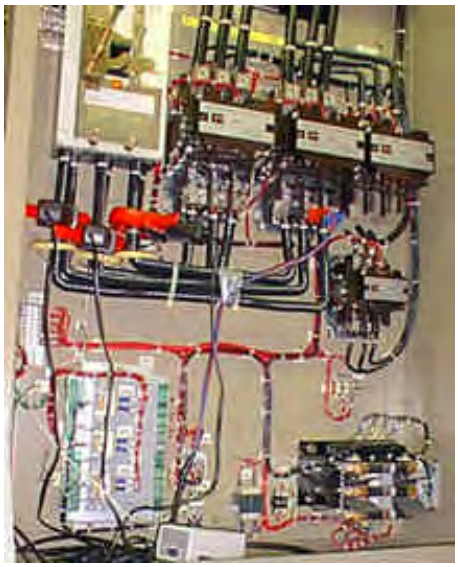

- Equipment: Elite Logger and ELOG 97c software.
- Select current transducers (CTs) appropriate for the measurement and panel space constraints.
- Attach current transducers (CTs) to Elite logger channel block.
- Plug Elite logger into AC supply.
- Setup Elite logger using product software.
- Electrician installation of voltage sensors in the electrical panel.
- Electrician installation of CTs in electrical panel for the specified load to be measured.
- Read real time data.
- Verify balanced current as well as appropriate, balanced voltage readings.
- Secure panel door and attach caution tape and warning notice if panel cannot be locked shut.
- Log data.

Power Spot Measurement – Power Sight

- Equipment: Power Sight PS 3000
- Plug Power Sight into AC supply if necessary.
- Connect current transducers (CTs) and voltage sensors to Power Sight.
- Electrician installation of voltage sensors in the electrical panel.
- Electrician installation of CTs in electrical panel for the specified load to be measured.
- Verify balanced current as well as appropriate, balanced voltage readings.
- Read and record the real time power reading for spot measurement.
- Log data for selected measurements.

Appendix E

CleanRooms Magazine: “An examination of ACRs: An opportunity to reduce energy and construction costs”

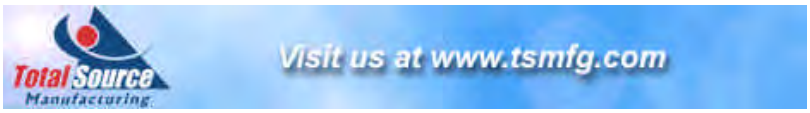

An examination of ACRs: An opportunity to reduce energy and construction costs

Now that cost-cutting has become paramount, it's time to discuss putting the growing list of energy-saving recommendations into practice

By Peter Rumsey PE, CEM

There are several conflicting sets of recommendations on what is the best airflow for cleanrooms. Recent articles in Cleanrooms magazine have explored the different ways of measuring or describing air flows and have discussed the upcoming Institute of Environmental Sciences and Technology (IEST; Rolling Meadows, Ill.) recommended changes; however, few industry observers have examined actual practices and the foreseeable impact on construction and energy costs.[1,2]

A recent benchmarking project conducted by Pacific Gas and Electric Company (San Francisco) and Lawrence Berkeley National Laboratory (Berkeley, Calif.) that measured air change rates in several cleanrooms verified that there is no consistent design strategy for air change rate, even for cleanrooms of the same cleanliness classification. Air change rates per hour (ACRs) are crucial for cleanroom designers because they have a significant impact on fan sizing and energy use.

Using best-practice ACRs can result in clean-filtered air, lower construction costs and reduced energy costs—a win-win situation for cleanroom owners.

Current design recommendations

Today, designers and cleanroom operators have a variety of sources to choose from when looking for an ACR recommendation. There is no clear consensus on what is an optimum ACR, and many of the established guidelines are outdated.

A recent article in Cleanrooms magazine pointed out that many of the recommended ACRs are based on relatively low-efficiency filters that were prevalent 10 years ago.3 For example, today's widely-used 99.99 percent efficient filters are three times more effective at filtering out 0.3 micron particles than the 99.97 percent filters that were common 10 years ago. Ultra-low penetration air (ULPA) filters are even more efficient than those of a decade ago.

When Rumsey Engineers (Oakland, Calif.) conducted a review of recommended cleanroom ACRs, it found that there is no agreement on a correct rate. Most sources suggest a range of rates. These ranges tend to be wide and do not provide clear guidance to designers who need to use a set ACR value to specify fan sizes. Figure 1 shows the result of our comparative review of recommended ACRs.

Air changes affect energy and construction costs

ACRs are the single largest factor in cleanroom fan sizing, building configuration and energy costs. As shown in Figure 1, recommended rates can vary from 250 to more than 700 air changes per hour for an ISO Class 5 cleanroom.

Click here to enlarge image

Figure 1: Recommended air change rates (ACRs) for ISO Class 5 cleanrooms

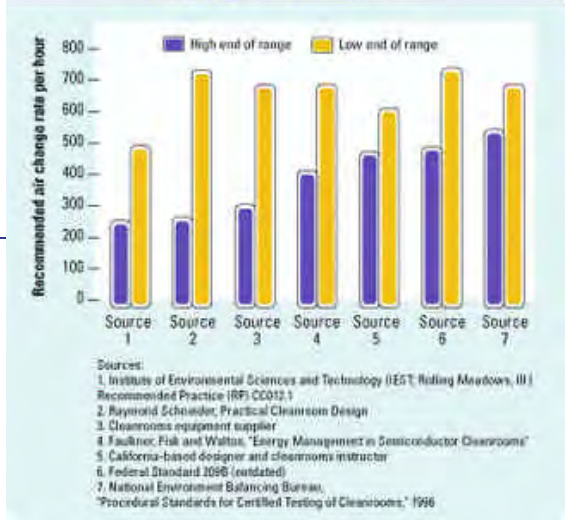

After gathering the results from its comparative review of recommended ACRs, Rumsey Engineers discovered that recommended rates for ISO Class 5 cleanrooms vary widely from source to source.

The high end of that range is almost three times the rate at the low end, yet the impact of this difference on fan sizing and motor horsepower is radically greater. According to the fan affinity laws, the power difference is close to the cube of the flow or air change rate difference. For example, a 50 percent reduction in flow will result in up to a factor of eight, or 87.5 percent reduction in fan power. Due to filter dynamics, the cube law does not apply exactly and, typically, the reduction is between a cube and a square relationship.

Even relatively conservative reductions of 10 percent to 20 percent in ACR provide significant benefits. A 20 percent decrease in ACR will enable close to a 50 percent reduction in fan size, with reduction calculation: 1 - 80%3. The energy savings opportunities are comparable to the potential fan size reductions.

While energy costs are not high on the priority list during the design and construction of cleanrooms, capital costs or construction costs are always important. Not so long ago, electronics, pharmaceutical and biotechnology companies did not need to worry much about construction costs. Currently, however, any designer would be irresponsible if construction costs or energy costs were ignored.

Click here to enlarge image

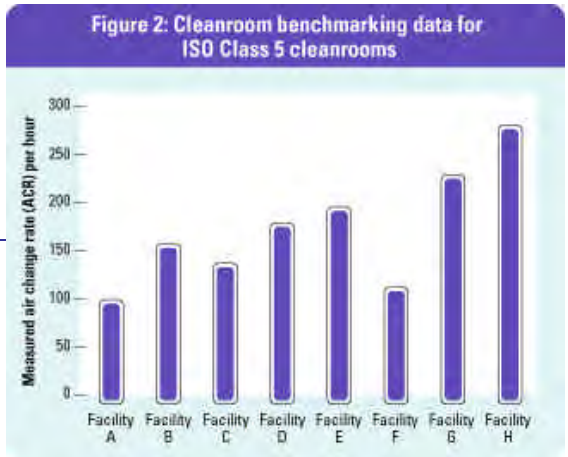

According to the results of the Pacific Gas and Electric Company and Lawrence Berkeley National Laboratory benchmarking data, cleanroom operators can use ACRs that are lower than what is recommended practice for ISO Class 5 facilities.

It's a common assumption that making a cleanroom more efficient will drive up construction costs, which is often impossible in today's tight-fisted climate. However, well-planned ACR reductions can reduce both construction and energy costs. This is that elusive goal, a true win-win situation, which decreases the amount of work the mechanical system has to perform and offers high leverage for downsizing equipment.

Current practice

Pacific Gas and Electric Company and Lawrence Berkeley National Laboratory recently conducted a cleanroom energy benchmarking study.4 A variety of systems and practices were measured, including air change rates in eight ISO Class 5 cleanrooms. The results were surprising.

While the recommended design ranges for ACRs are from 250 to 700 air changes per hour, the actual operating ACRs ranged from 90 to 250 (see Figure 2). All of these cleanrooms were certified and performing at ISO Class 5 conditions. This shows that cleanroom operators can use ACRs that are far lower than what is recommended practice without compromising either production or cleanliness requirements.

This is often done to lower energy costs. However, these facilities did not take advantage of the fan sizing reduction opportunities during construction. As a result, most of the fan systems were operating at very low variable speed drive speeds.

What others have found

Air cleanliness is a critical component of any cleanroom, far outweighing energy saving priorities. Designers and operators need evidence from others who have tried similar strategies in order to address the perceived risks of lowering air change rates.

Fortunately, a growing body of data, case studies and research are available that document success. In a recent study by International Sematech (Austin, Texas), no noticeable increase of particle generation was found when air change rates were lowered by 20 percent in ISO Class 4 cleanrooms.5 A recent study at the Massachusetts Institute of Technology (MIT; Cambridge, Mass.) found that in a raised-floor-type cleanroom "with a small decrease in air velocity, such facilities will decrease particle deposition and maintain air unidirectionality."6

Other successes have been noted by cleanroom operators at Intel (Santa Clara, Calif.) and Sandia National Laboratories (Albuquerque, N.M.). Michael Dever, Intel's Oregon site utility manager, reported that an Intel project aiming to reduce both air change rates and ceiling HEPA velocities succeeded in achieving a 20 percent fan energy savings goal at a low cost of implementation. Sandia National Laboratories has also successfully reduced air change rates in their state-of-the-art ISO Class 4 and 5 cleanrooms. This is especially significant because Sandia pioneered laminar flow cleanrooms in the early 1960s.

Conclusions and recommendations

There is no doubt that more clarification and justification of optimal and safe air change rates are required. From the Pacific Gas and Electric Company and Lawrence Berkeley National Laboratory benchmarking data, it is clear that air change rates can be lower than what is currently recommended by several sources.

The benchmarking data suggests that an ISO Class 5 facility should be designed with an air change rate of around 200 air changes per hour. A conservative upper limit should be about 300, significantly lower than the high range of 700 indicated by some sources.

Facility designers and operators tend to err on the side of conservatism in their efforts to provide high reliability cleanroom support. More independent research on optimized air change rates based on contemporary filter efficiencies needs to be conducted to reduce the perceived risks of modifying standard practices.

Biotechnology and pharmaceutical cleanrooms are currently designed to meet current good manufacturing practices (cGMPs) that require high air change rates. These ACRs should be re-examined as part of upcoming revisions to the cGMP. In addition, IEST recommended practices updates should include lower ACR guidelines.

Using better air change rate practices will allow designers to offer lower construction costs as well as reduced energy costs while maintaining the high level of air cleanliness that is required in cleanroom facilities.III

***Peter Rumsey PE**, CEM, president of Rumsey Engineers (Oakland, Calif.), specializes in cleanroom design and other critical applications. He has over 20 years of experience internationally in commercial, governmental and scientific projects. Rumsey can be reached at prumsey@rumseyengineers.com.*

References

- 1.Fitzpatrick, Mike and Goldstein, Ken, "Cleanroom airflow measurement: Velocity, air changes per hour or percent filter coverage?" Cleanrooms magazine, May 2002.
- 2.Fitzpatrick, Mike and Goldstein, Ken, "Cleanroom airflows Part II: The messy details," Cleanrooms magazine, July 2002.
- 3.Jaisinghani, Raj, "New ways of thinking about air handling," Cleanrooms magazine, January 2001.
- 4.http://ateam.lbl.gov/cleanroom/benchmarking/index.htm.
- 5.Huang, Tom, "Tool and Fab Energy Reduction," Spring 2000 Northwest Microelectronics Workshop, Northwest Energy Efficiency Alliance.
- 6.Vazquez, Maribel and Glicksman, Leon, "On the Study of Altering Air Velocities in Operational Cleanrooms," 1999 International Conference on Advanced Technologies and Practices for Contamination Control.

Click here to enlarge image

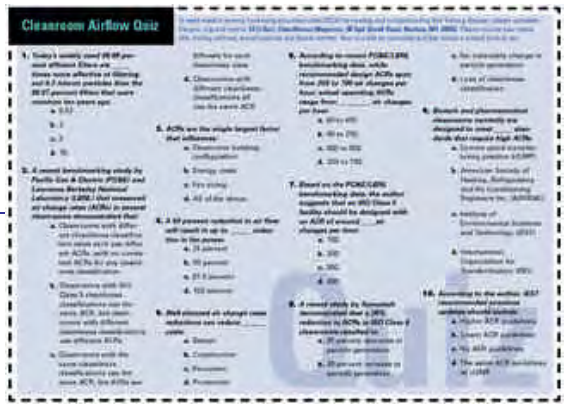

CleanRooms January, 2003
Author(s) : Peter Rumsey PE

Interested in a subscription to CleanRooms Magazine?
Click here to subscribe!

Copyright © 2004 - PennWell Corporation. All rights reserved.

CLEANROOM ENERGY BENCHMARKING AND EFFICIENCY TECHNOLOGY TRANSFER ASSISTANCE SITE REPORT

FACILITY K

SOUTHERN CALIFORNIA

AUGUST 2004

SPONSORED BY:

LAWRENCE BERKELEY NATIONAL LABORATORY

PREPARED BY:
RUMSEY ENGINEERS, INC
99 LINDEN STREET
OAKLAND, CA 94607
WWW.RUMSEYENGINEERS.COM

Copyright (c) 2004 Lawrence Berkeley National Laboratory. All rights reserved.

Acknowledgements

Special thanks to the entire team for their generous assistance and cooperation throughout the benchmarking process.

Project Team

Lawrence Berkeley National Laboratory	**Rumsey Engineers, Inc**
Bill Tschudi	Peter Rumsey
Dale Sartor	Achai Broner
	Larry Chu
	John Weale

Reproduction or distribution of the whole, or any part of the contents of this document without written permission of LBNL is prohibited. The document was prepared by LBNL for the exclusive use of its employees and its contractors. Neither LBNL nor any of its employees makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any data, information, method, product or process disclosed in this document, or represents that its use will not infringe any privately-owned rights, including but not, limited to, patents, trademarks, or copyrights.

Table of Contents

I. Introduction 1

II. Review of Site Characteristics 1

A. Site 1

B. "North Wing" Chemistry and Biology Laboratories 2

C. "East Wing" Chemistry and Biology Laboratories 3

III. Site Energy Use Characteristics 4

A. Laboratory Power Consumption 4

B. Electrical System Power Consumption 6

IV. System Performance Metrics 6

V. Site Observations Regarding Energy Efficiency 9

APPENDICES

A. Trended Data Graphs

B. Data Collection and Accuracy Notes

C. Measurement Methodology

I. INTRODUCTION

As part of the California Energy Commission PIER study, energy use at Facility K was monitored June 21 to June 25, 2004. The project is administered by LBNL (Lawrence Berkeley National Laboratory) and funded through the California Energy Commission. Facility K is a building that houses primarily laboratory spaces, office spaces, common areas, and support areas.

This site report reviews the data collected by the monitoring team and presents a set of performance metrics as well as a complete set of trended data points for the end uses of energy for equipment supporting and located in cleanrooms or laboratories. The chilled water plant at this facility was not measured in detail since several chilled water plants have been benchmarked as part of this study at prior sites. Energy metrics were established that allow cleanroom and laboratory owners to evaluate their energy efficiency performance and identify opportunities for improvements that reduce their overall operating costs.

With this report, Facility K is receiving the energy monitoring data collected along with a few energy efficiency improvement observations. This Site Report summarizes the data collected and presents energy performance metrics with which will allow the facility to compare selected systems' performance to others. A database of other lab energy performance can be found at the Laboratories for the 21st Century website (http://www.dc.lbl.gov/Labs21/Labs21intro.php).

First, this report reviews the site characteristics, noting design features of the mechanical plant and the cleanrooms or laboratories monitored. Second, the energy use for the cleanrooms or laboratories and major mechanical equipment is broken down into major components. Third, performance metrics recorded through the project are presented. Finally, key energy efficiency observations for the facility will be noted. The data collected, trended graphs and methodology documentation are included among the appendices.

II. REVIEW OF SITE CHARACTERISTICS

A. Site

Facility K, located in Southern California, is a 113,000 square foot, three-story building that is approximately five years old. The building houses primarily laboratory spaces, office spaces, and common areas, such as hallways and restrooms. The 1st and 2nd floors have a nearly identical layout; the building has two wings, with one referred to as the "North Wing" and the other as the "East Wing". Each wing has both a Chemistry and Biology Laboratory. The basement consists of additional laboratory and common areas.

This site has on-site power production created by a natural gas cogeneration system. The waste heat from this system is recovered and utilized for an absorption chiller and heating hot water. A 1,875 kVA/1,500 kW diesel generator provides backup power in the case of a utility failure. A 100 kVA uninterruptible power supply (UPS) system is also utilized to provide backup power to the computer servers.

Absorption Chiller

The environmental systems serving the laboratories run 8,760 hours a year in order to maintain conditions and provide a safe lab environment. The 1st and 2nd floor laboratory and office spaces are conditioned by two large air handlers served by a chilled water plant, and a hot water boiler plant. The air handling systems serving these areas are both VAV (variable air volume) systems providing once-through, 100% outside air. Exhaust is provided to these areas also. The basement laboratory spaces are conditioned also by an air handler providing once-through, 100% outside air. This air handler is also served by the chilled water plant, and the hot water boiler plant mentioned above.
The common areas on all three floors are served by three traditional air handlers, which provide both make up and recirculation air. The common areas are made up of the hallways, lobby, restrooms, and some office areas. These air handlers are also served with chilled and hot water from the plants mentioned above. The air handlers serving these spaces are AH-3, AH-4, AH-6 and AH-7.

The spaces chosen for monitoring are the Chemistry and Biology Laboratories located on the 1st and 2nd floors. The air handler that serves the Chemistry and Biology Laboratories in the "North Wing" is AH-1/2 rated for a nominal air delivery of 111,500 cfm (cubic feet per minute). An exhaust system, consisting of six identical fans connected to a common plenum serves the "North Wing". The air handler that serves the Chemistry and Biology Laboratories in the "East Wing" is AH-5/6 rated for a nominal air delivery of 101,000 cfm. Exhaust is provided to the "East Wing" by four exhaust fans connected to a common plenum.

The chilled water plant consists of: two 600 ton, variable speed driven centrifugal chillers; two induced-draft cooling towers; three primary and two secondary pumps. The system utilizes a constant flow primary-variable flow secondary pumping system. Currently there is a new building under construction, which shares its chilled water plant with this building. The absorption chiller, which is part of the new building, is being used for the existing building to take advantage of the waste heat being generated by the cogeneration system.

B. "North Wing" Chemistry and Biology Laboratories

The "North Wing" laboratories account for 33,100 sf or 29% of the total building area. The Chemistry Labs are 12,450 sf and the Biology Labs are 20,650 sf. The lab is made up of actual lab areas where the experiments are conducted and office areas adjacent to the labs; this is what defines the laboratories. There is no physical wall of separation between the labs and the office areas. The air is delivered to the office areas spaces via a ducted VAV system, and cascaded to the lab areas. The labs operate 24 hours a day, although they are not continually occupied. The Chemistry Laboratories consist of primarily ten modules (five on each floor). The modules are very similar to one another in that they are approximately the same square footage and contain similar equipment. Each module contains four bench mounted fume hoods, which have occupancy sensors and sash position sensors. These sensors work in conjunction with the VAV dampers to minimize the amount of air exhausted and supplied to the areas to save fan, cooling, and heating energy. The Biology

Air Handler AH-1/2

Laboratories are made up of primarily eight modules (four on each floor). The modules are very similar in square footage and equipment.

The "North Wing" is served by one make up air handler (AH-1/2) with a nominal air delivery rate of 111,500 cfm. The make up air unit delivers air into the laboratories and office spaces via ceiling diffusers. The air handler consists of two sections; each section consists of a variable speed driven fan, cooling coil and hot water coil. The two sections are joined by a common supply plenum.

C. "East Wing" Chemistry and Biology Laboratories

The "East Wing" laboratories account for 23,000 sf or 20% of the total building area. The Chemistry Labs are 9,400 sf and the Biology Labs are 13,600 sf. The laboratories also are made up of office spaces similar to the "North Wing" Laboratories. The air is delivered to the office areas spaces via a ducted VAV system, and cascaded to the lab areas. The labs operate 24 hours a day, although they are not continually occupied. The Chemistry Laboratories consists of primarily eight modules (four on each floor). The modules are very similar to one another as described above. Each module contains four bench mounted fume hoods, which have occupancy sensors and sash position sensors similar to the "North Wing" Labs mentioned above. The Biology Laboratories are made up of primarily eight modules (four on each floor). The modules are very similar in square footage and equipment.

Air Handler AH-5/6

The "East Wing" is served by one make up air handler (AH-5/6) rated for a nominal air delivery rate of 101,000 cfm. The make up air unit delivers air into the laboratories and office areas via ceiling diffusers. The air handler consists of two sections, similar to AH-1/2, which are joined by a common supply plenum.

III. SITE ENERGY USE CHARACTERISTICS

A. Laboratory Power Consumption

The power consumption attributed to the laboratory air handling systems, exhaust fans, process tools, and lighting are reported in Tables 1, 2 and 3. These values were directly measured. Note that the cooling and heating energy are not included in the tables and chart. This breakdown of energy use by equipment helps identify the major loads. Table 1 is also shown graphically in the chart below.

Table 1. Total Laboratory Power Consumption [1]

Description	Average Load (kW)
AIR HANDLING	
Make Up Air Fans	88.7
Exhaust Fans	195.6
PROCESS POWER	
Laboratory Areas	137.2
Laboratory Office Areas	62.9
LIGHTS	70.0
TOTAL [1]	**554.4**

1. The total power consumption does not include the power and gas consumption of the chillers and boilers cooling and heating energy produced.

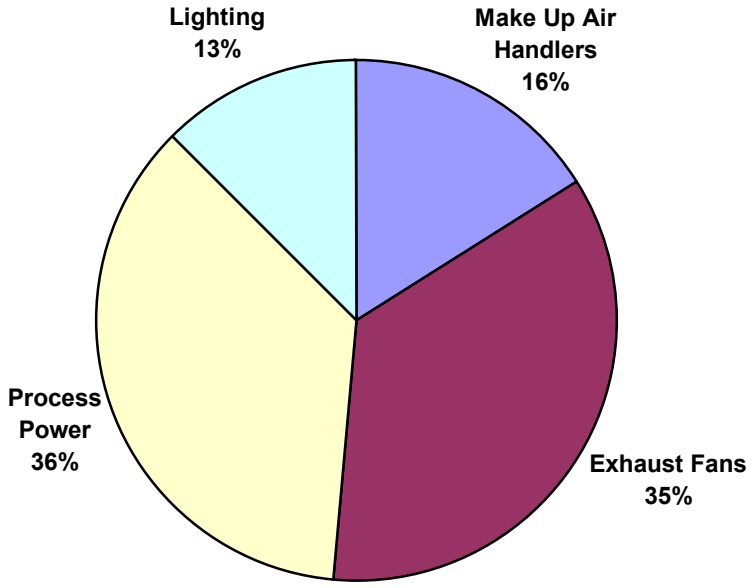

Figure 1. Total Laboratory Power Consumption

Table 2. "North Wing" Chemistry and Biology Labs Power Consumption Breakdown

Description	Average Load (kW)	Metric
AIR HANDLING		
Make Up Air Fans	46.5	0.66 W/cfm
Exhaust Fans	112.3	[3]
PROCESS POWER [1]	-	-
LIGHTS [2]		
Chemistry Labs + Offices	7.7	0.62 W/sf
Biology Labs + Offices	34.1	1.7 W/sf

1. *Process power was not determined for the individual wing.*
2. *Lighting power calculated based on lighting power density of a representative laboratory module*
3. *The metric was not calculated since an operating airflow was not determined.*

Table 3. "East Wing" Chemistry and Biology Labs Power Consumption Breakdown

Description	Average Load (kW)	Metric
AIR HANDLING		
Make Up Air Fans	42.2	0.60 W/cfm
Exhaust Fans	83.3	[3]
PROCESS POWER [1]	-	-
LIGHTS [2]		
Chemistry Labs + Offices	5.8	0.62 W/sf
Biology Labs + Offices	22.5	1.7 W/sf

1. *Process power was not determined for the individual wing.*
2. *Lighting power calculated based on lighting power density of a representative laboratory module*
3. *The metric was not calculated since an operating airflow was not determined.*

B. Electrical System Power Consumption

The table below shows the power consumption of the emergency generator standby power loss. The 1,500 kW emergency generator constantly draws power to maintain the batteries and a specific temperature of the diesel engines that drive the generator.

Table 4. Electrical System Power Consumption

Description	Average Load
ELECTRICAL SYSTEM	
Emergency Generator Standby Power	2.3 kW

IV. SYSTEM PERFORMANCE METRICS

Metrics are ratios of important performance parameters that can characterize the effectiveness of a system or component. In order to gage the efficiency of the entire building system design and operation, this project tracks key metrics at different system levels. These metrics can be used to compare designs or determine areas with the most potential for improvement via retrofit or replacement. The data used for comparison was obtained from the Laboratories for the 21st Century website (http://www.dc.lbl.gov/Labs21/Labs21intro.php).

Table 5. Laboratory Metrics

Description	*"North Wing" Laboratories*	*"East Wing" Laboratories*	*Total for All Laboratories*
VENTILATION			
cfm	70,060	70,700	140,760
Fan kW	46.5	42.2	88.7
W/cfm (measured)	0.66	0.60	0.63
cfm/sf-lab (measured)	2.1	3.1	2.5
LIGHTING [1]			
W/sf-lab (measured)	1.1	1.1	1.1
PROCESS [1]			
W/sf-lab (measured)	10.9	10.9	10.9

1. Power densities calculated based on power measurements of a representative laboratory module.

The figures on the following pages compare the facility's airflow with other facilities documented in the Labs for the 21st Century online database. The facilities used for comparison are located in different parts of the nation, which essentially are representative of various climate zones (i.e. cool-humid, cold-dry, cold-humid, warm-marine, hot-humid, etc.). Note that the data shown for this facility was actually based on a spot measurement, which closely represents an average value rather than a peak value. The ventilation supply cfm/sf for Facility K is higher than a few of the compared facilities. There are a number of facilities that are actually being maintained at lower flow rates, thus a reduction in airflow to the lab spaces at this facility should be considered for energy savings (see "Site Observations Regarding Energy Efficiency" section).

Figure 2. Laboratory Ventilation Peak W/cfm

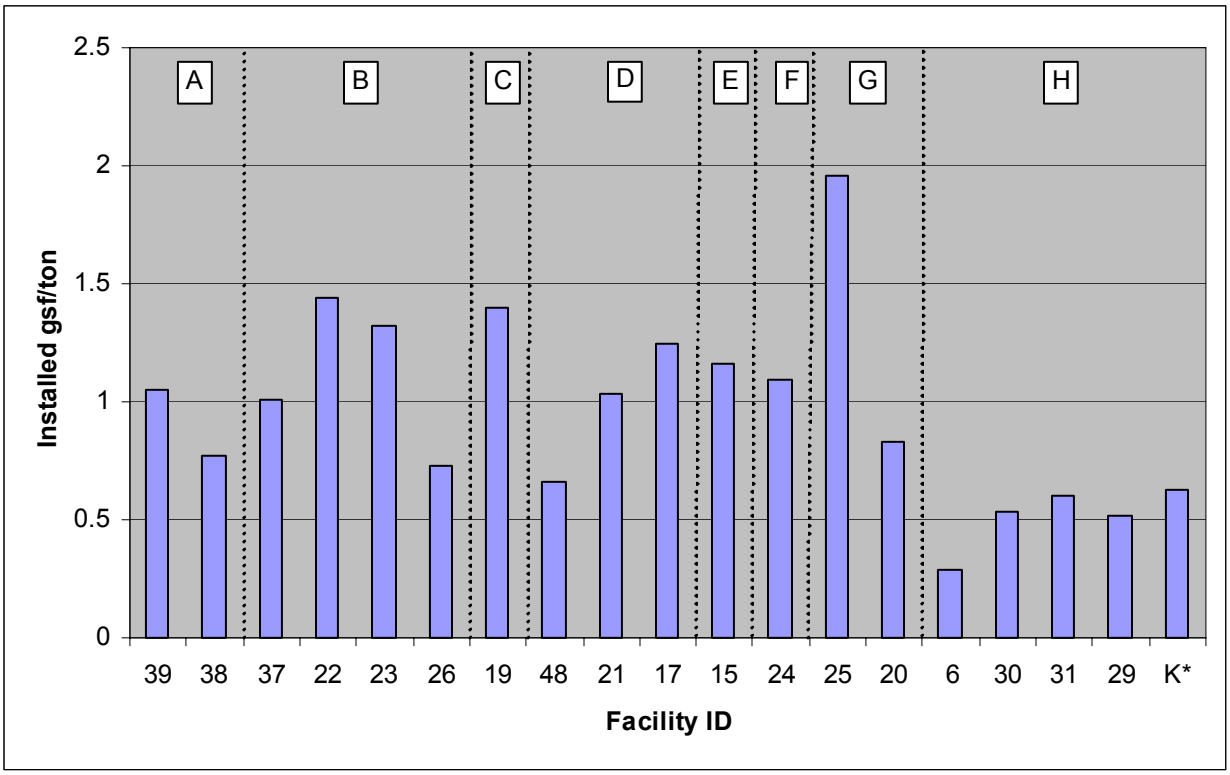

** Data for Facility K was based on a spot measurement, which closely represents an average value rather than a peak value.*

Figure 2 Legend – Climate Type, Representative City

A Cool-Dry (Boise, ID)	**E** Mixed-Marine (Salem, OR)
B Cool-Humid (Chicago, IL)	**F** Warm-Dry (El Paso, TX)
C Mixed-Dry (Albuquerque, NM)	**G** Warm-Humid (Memphis, TN)
D Mixed-Humid (Baltimore, MD)	**H** Warm-Marine (San Francisco, CA)

Figure 3. Laboratory Ventilation Peak Supply cfm/sf

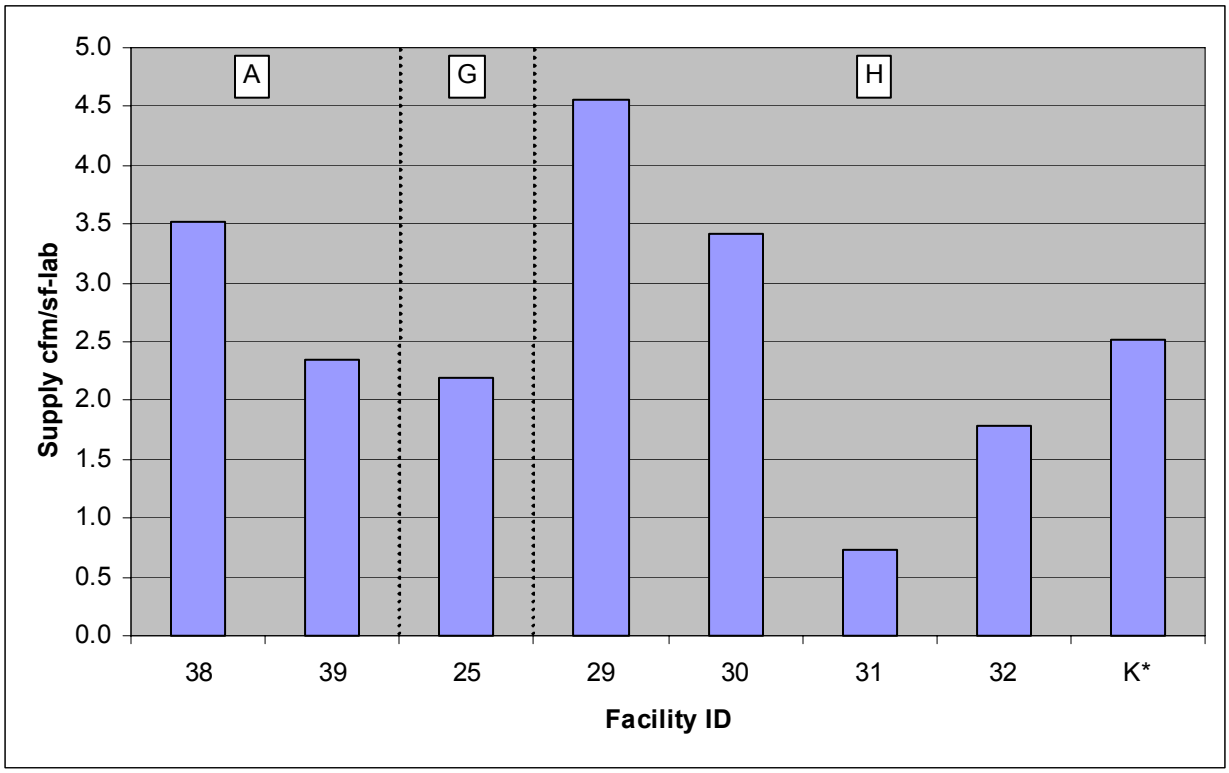

** Data for Facility K was based on a spot measurement, which closely represents an average value rather than a peak value.*

Figure 3 Legend – Climate Type, Representative City

A Cool-Dry (Boise, ID)
G Warm-Humid (Memphis, TN)
H Warm-Marine (San Francisco, CA)

Table 6. Facility Metrics

Description	*Metric*
COOLING	
gsf/ton (installed) [1]	58.8
W/gsf (installed)	9.7

1. gsf – gross square feet

The figure below shows the facility's installed cooling capacity based on the total facility area (gsf/ton). This chart is used to compare the installed cooling capacity of this facility to other laboratory facilities.

Figure 4. Installed Cooling – Gross Square Feet per Ton

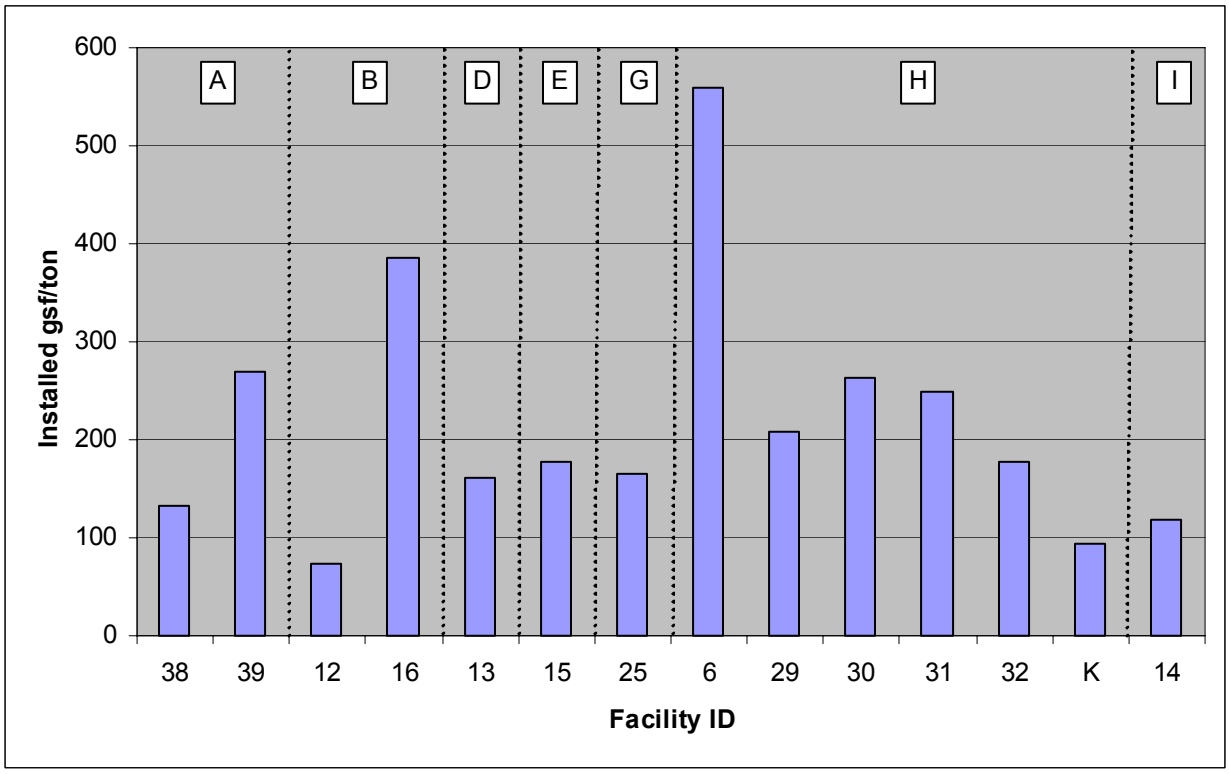

Figure 4 Legend – Climate Type, Representative City

A Cool-Dry (Boise, ID)
B Cool-Humid (Chicago, IL)
D Mixed-Humid (Baltimore, MD)
E Mixed-Marine (Salem, OR)
G Warm-Humid (Memphis, TN)
H Warm-Marine (San Francisco, CA)
I Cold-Humid (Burlington, VT)

VI. SITE OBSERVATIONS REGARDING ENERGY EFFICIENCY

Many notable energy efficiency features have already been implemented into the systems at this facility. Some of the key features are:

- Cogeneration / on-site power production
- Recovery of waste heat from cogeneration system utilized for absorption chiller and reheat in the new building
- The chilled water loop in the new building is coupled to the chilled water loop in this building, allowing the use of the absorption chiller and the cogeneration system's waste heat, and also allowing a chiller to run at higher loads. The overall chilled water plant efficiency is improved since chillers running at low partial loads are inherently inefficient.
- Cooling towers operating in parallel
- Variable air volume supply and exhaust air flow in laboratories based on fume hood sash position
- Protocol for closing fume hoods to minimum position during unoccupied hours
- Low pressure drop pre-filters on make up air handlers
- Parallel operation of make up air handler fans

- Variable speed driven centrifugal chillers
- Manifolded exhaust fans staged on and off based on maintaining a plenum pressure setpoint
- Lighting controlled via occupancy sensors
- The air handlers in the new building were sized based on measurements made in this building. The make up air handlers for the original building were sized for an 80% diversity factor versus 60% diversity for the new building.

However, there are a few additional opportunities that can be implemented to improve the overall energy efficiency at this site.

Chilled Water Temperature Reset

Chiller performance can be improved by implementing a chilled water temperature reset. When dehumidification is *not* required, there is typically a significant surplus of coil capacity. The 41 to 42°F low temperature loop could be reset to a higher temperature. Low temperature chilled water is inherently more energy intensive to produce due to the larger temperature delta, or 'lift,' the compressor is required to move heat through. On centrifugal compressors based chillers, an increase of one degree in the chilled water supply temperature improves the efficiency of the chiller by 1 to 2%. The chiller manufacturer should be consulted before radical modification of the supply temperature, but a change of a few degrees should allow for an improvement in efficiency without risk.

There are a number of chilled water resets that have been shown to work. The simplest is to set the CHW temperature based on the outside air temperature, with the lowest temperature chilled water setpoint achieved during the highest outside air temperatures and during dehumidification. A more sophisticated approach is to poll all the chilled water valves on the loop for their position. If no valve is at 90% or more open, then maximum cooling is not required and the chilled water temperature can be increased. When a valve position exceeds 90%, it indicates a space is calling for additional cooling and the chilled water temperature can be reset down. The savings from a chilled water reset are maximized when the periods of dehumidification are minimized.

Make Up Air Handler Supply Air Temperature Reset

Another strategy to save energy is to implement a supply air temperature reset on the make up air. Typically the supply air temperature delivered by the air handlers is based on the lowest room temperature or lowest average temperature (of a group of rooms) that is required. A study should be done to determine the actual requirements in each of the spaces. Additionally, reheat energy will be saved if the supply air temperature of the air handler is raised. By implementing this measure along with a chilled water temperature reset and differential pressure setpoint reset, additional energy savings will be realized.

Condenser Water Temperature Reset

The centrifugal chillers may also benefit from a reduction in lift. A reduction in the condenser water temperature of a couple degrees will achieve meaningful energy savings, however chiller stability can be compromised if it is taken too low, hence the need to verify with the chiller manufacturer before modifying this point. A bypass and a dedicated recirculation pump on the condenser water loop feeding the absorption chiller should be installed so that the warmer condenser water temperature required by the chiller can be maintained.

Add Additional Plates to the Cogeneration System Heat Exchangers (HX-9 & HX-10)

The plate and frame heat exchangers for transferring the heat generated by the cogeneration system to the cooling tower will result in pumping energy savings if more plates are added to lower the pressure drop through the device. Additionally, adding more plates will increase the heat transfer effectiveness and may result in lower flow rates for both the condenser water and engine water loops.

Chilled Water Differential Pressure Setpoint Reset

The secondary pumps on the chilled water supply loop are controlled via VFDs and a set pressure differential setpoint. An actively controlled differential setpoint can be utilized to save pumping energy. The differential setpoint can be reset based on the cooling valve positions. When the valves are partially closed, the differential setpoint can be reduced since lower water flow is required for each coil, hence the lower pressure drop required through the loop. As the valves begin to open, the differential setpoint should be increased so that adequate flow is supplied to the coils. Typically the most energy savings are realized when the coils are oversized.

Air Change Rate Nighttime Setback

Energy savings are currently achieved in the Laboratories by lowering the air change rates based on the fume hood sash positions. Setback of the airflow for the Biology Labs during unoccupied hours should be investigated to further increase energy savings. During occupied hours, the Biology Labs operate at approximately 29 ACH (air changes per hour). OSHA recommends 4 to 12 ACH for normally adequate laboratory ventilation (*source: OSHA 29 Code of Federal Regulations*). The National Institutes of Health recommends a minimum of 6 ACH. A case study at the Fred Hutchinson Cancer Research Center in Seattle has shown that a reduction in its minimum air change rate from 10 ACH to 6 ACH (occupied) and 4 ACH (unoccupied) to be successful. By reducing the air change rate during unoccupied hours from 10 to 12 ACH, significant energy savings can be achieved while safety is maintained.

Emergency Generator Standby Power Loss Reduction

The 1,500 kW emergency generator constantly draws power to maintain the batteries and a specific temperature of the diesel engines that drive the generator. An average of 2.3 kW is being consumed to prime the generator for a power outage. Claims have been made that the engine(s) of a generator can be maintained at lower temperatures without introducing any physical damage when started. The manufacturer of the generator should be contacted to explore energy reduction options.

LBNL's Low-Flow Fume Hood

LBNL has developed a low-flow fume hood, which significantly requires a lower amount of airflow as compared to a standard fume hood. Although, the airflow through the low-flow fume hood is lower than normal, a safe working environment protecting the health of the users is still maintained. Lower airflow rates translate into lower fan energy, and lower cooling and heating energy requirements.

Place Refrigerators and Freezers in Separate Rooms

Refrigerators and freezers give off a substantial amount of heat to the surrounding spaces in which they are located. By placing refrigerators and freezers in a separate room, then the cooling requirements for the space may be reduced. In addition, by grouping refrigerators into a common room, the temperature requirements for that room can be relaxed (i.e. increase the room temperature setpoint) for energy savings depending on what is tolerable to the users.

Appendix A

Trended Data Graphs

Outside Air Conditions

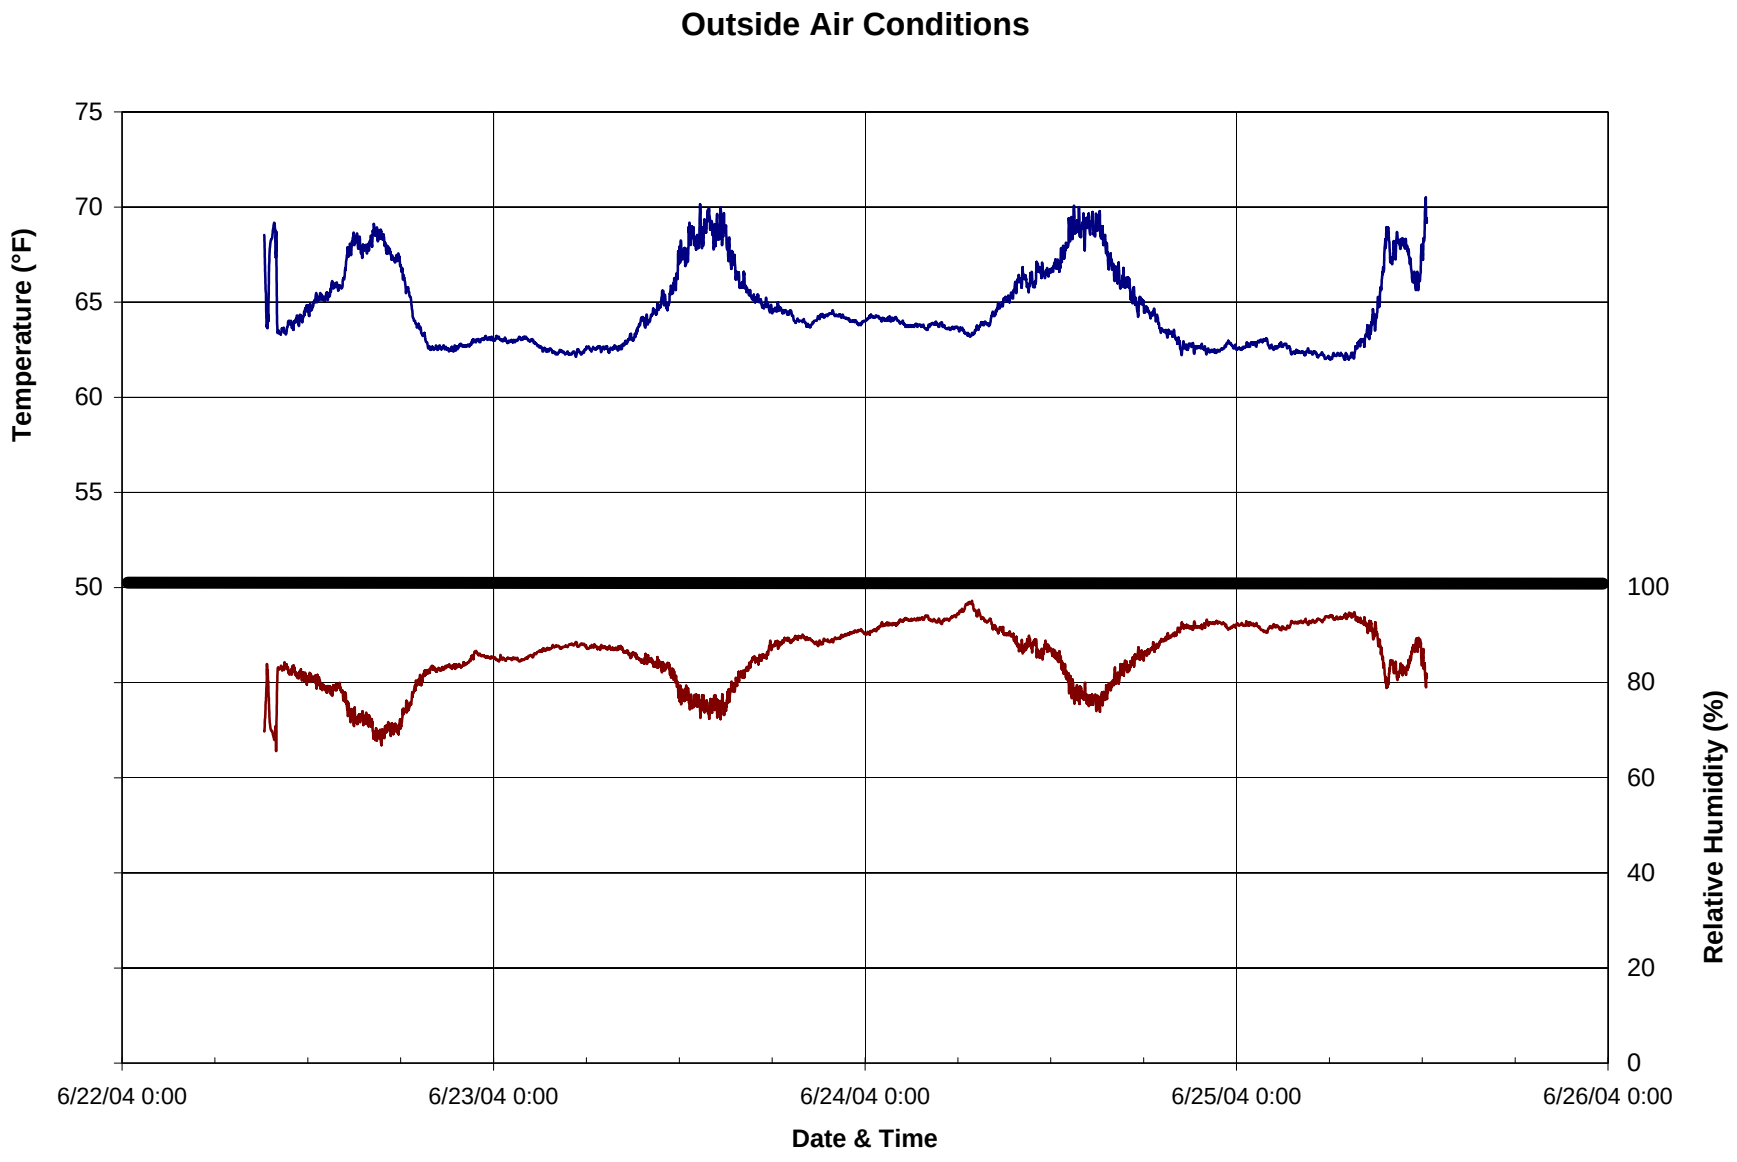
Outside Air Conditions
Temperature (°F)
75
70
65
60
55
50
100
80
60
40
20
0
Relative Humidity (%)
6/22/04 0:00
6/23/04 0:00
6/24/04 0:00
6/25/04 0:00
6/26/04 0:00
Date & Time

“North Wing” Chemistry & Biology Laboratories

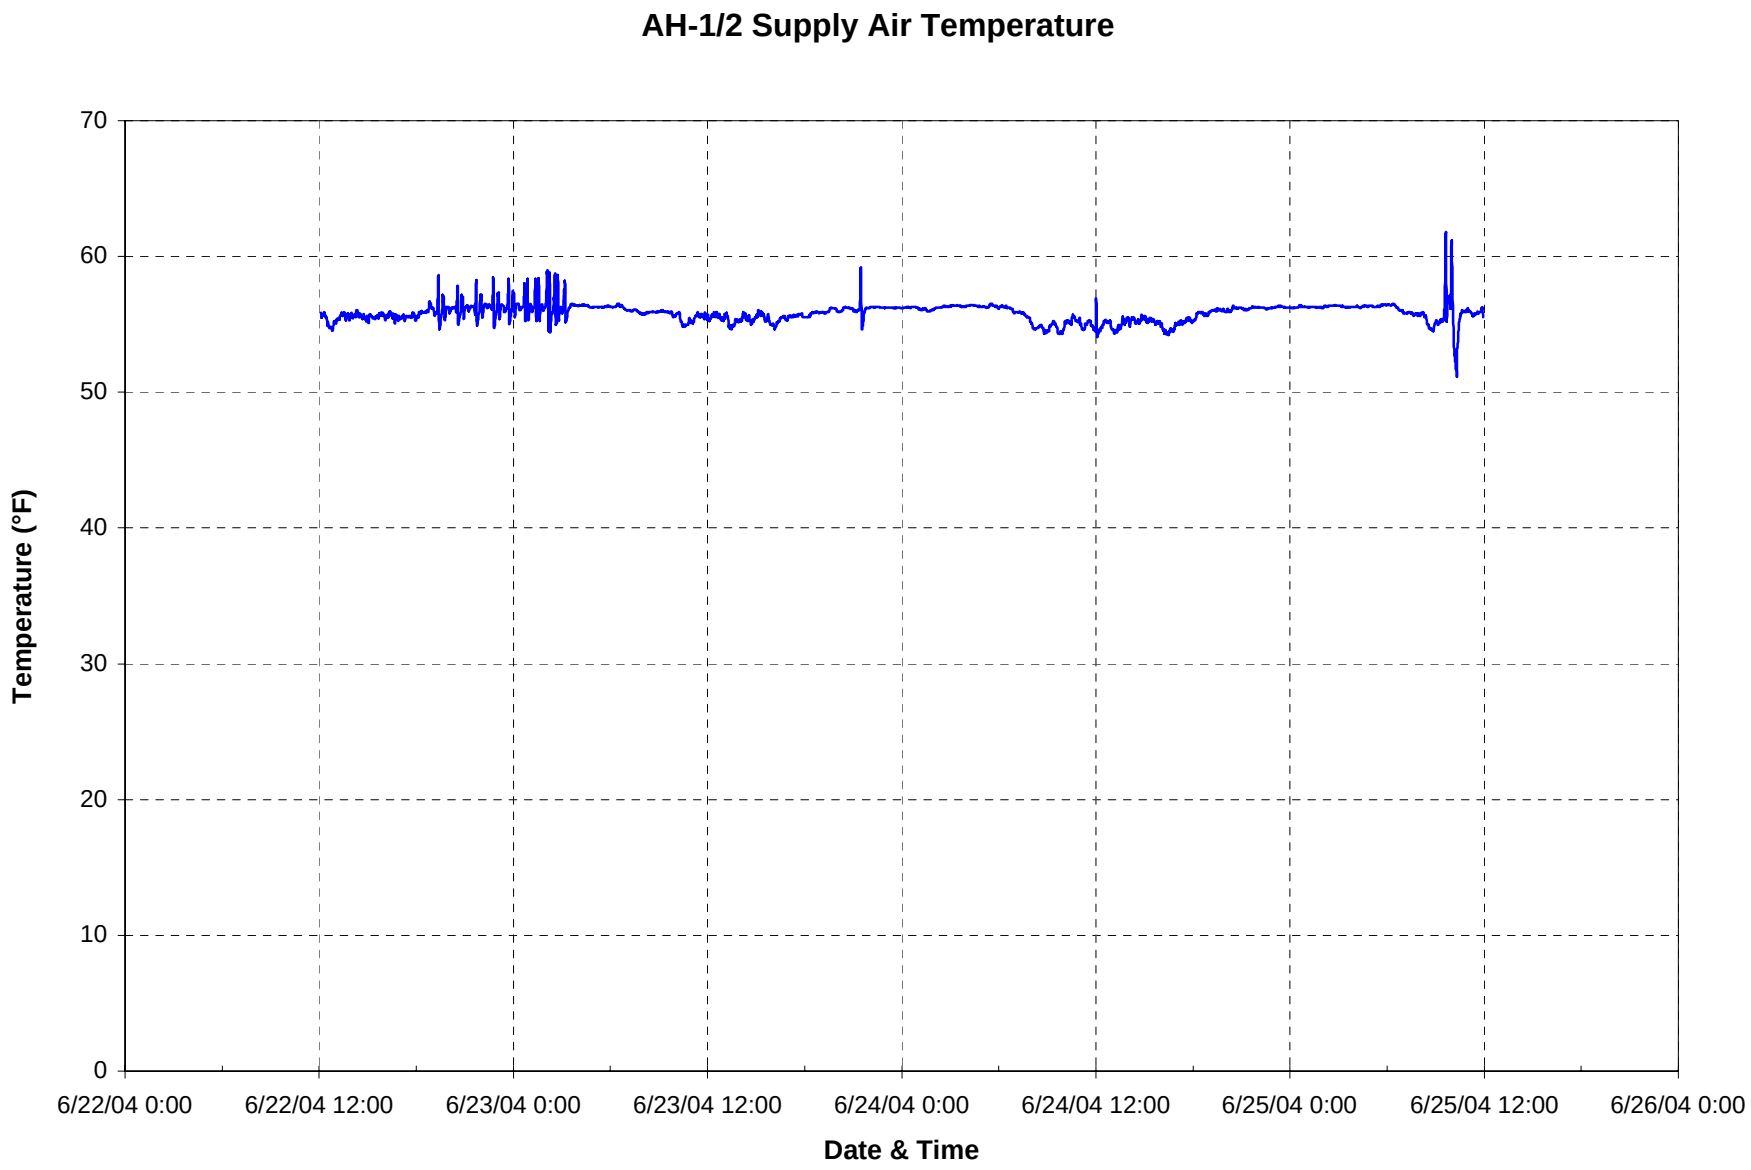
AH-1/2 Supply Air Temperature
Temperature (°F)
70
60
50
40
30
20
10
0
6/22/04 0:00
6/22/04 12:00
6/23/04 0:00
6/23/04 12:00
6/24/04 0:00
6/24/04 12:00
6/25/04 0:00
6/25/04 12:00
6/26/04 0:00
Date & Time

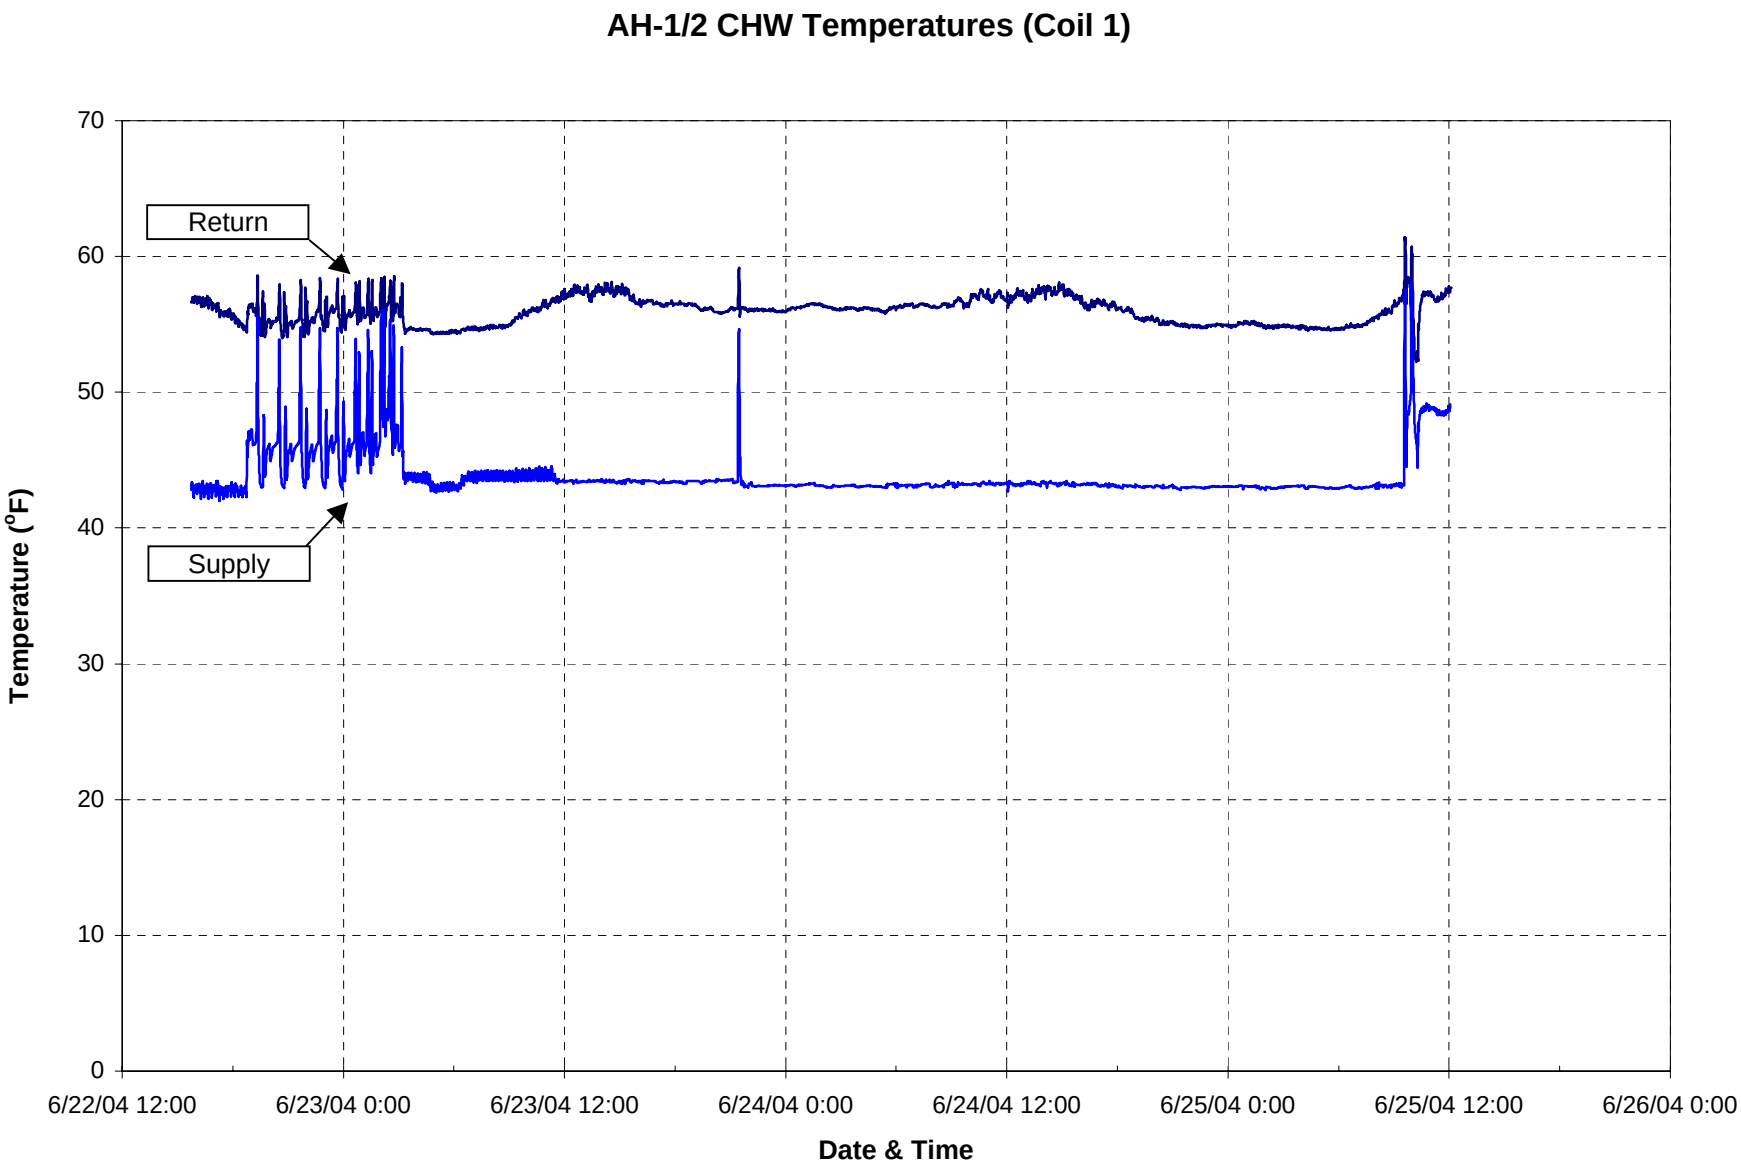
AH-1/2 CHW Temperatures (Coil 1)
Return
Supply
Temperature (°F)
70
60
50
40
30
20
10
0
6/22/04 12:00
6/23/04 0:00
6/23/04 12:00
6/24/04 0:00
6/24/04 12:00
6/25/04 0:00
6/25/04 12:00
6/26/04 0:00
Date & Time

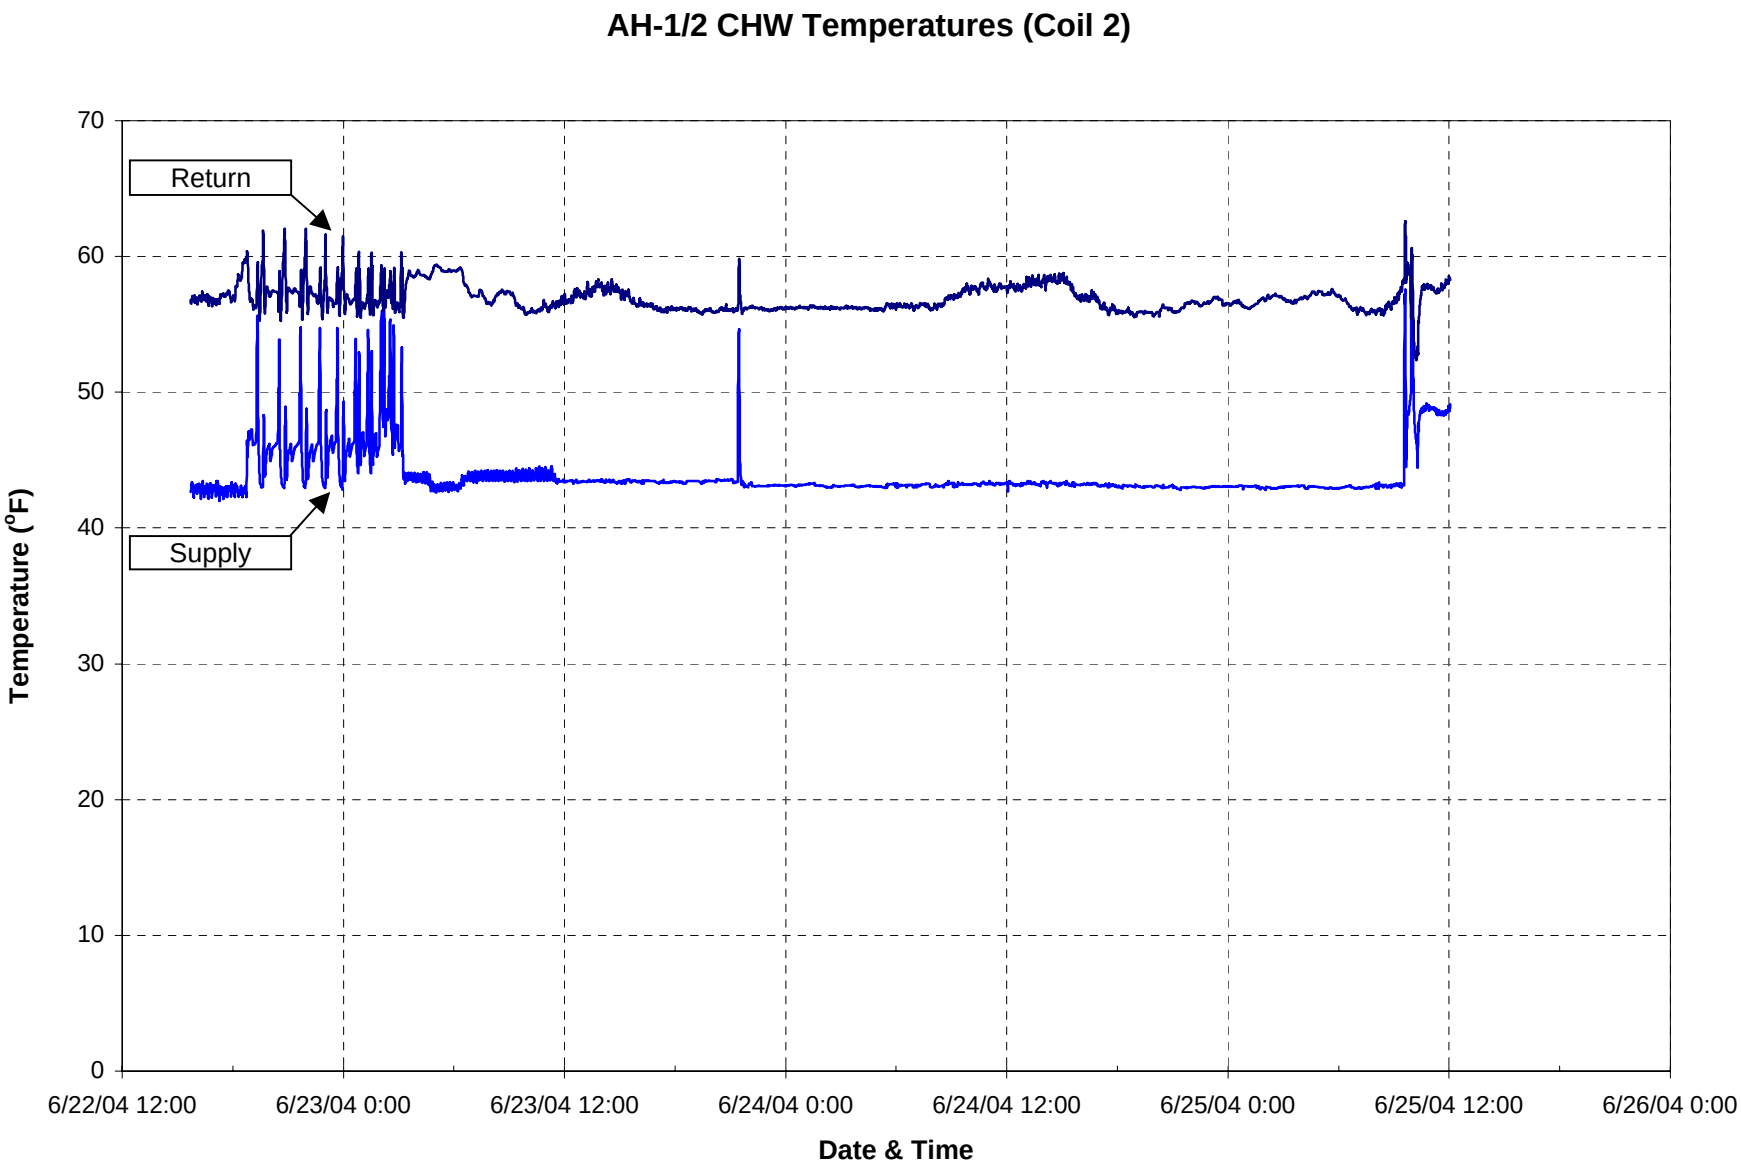
AH-1/2 CHW Temperatures (Coil 2)
Return
Supply
Temperature (°F)
70
60
50
40
30
20
10
0
6/22/04 12:00
6/23/04 0:00
6/23/04 12:00
6/24/04 0:00
6/24/04 12:00
6/25/04 0:00
6/25/04 12:00
6/26/04 0:00
Date & Time

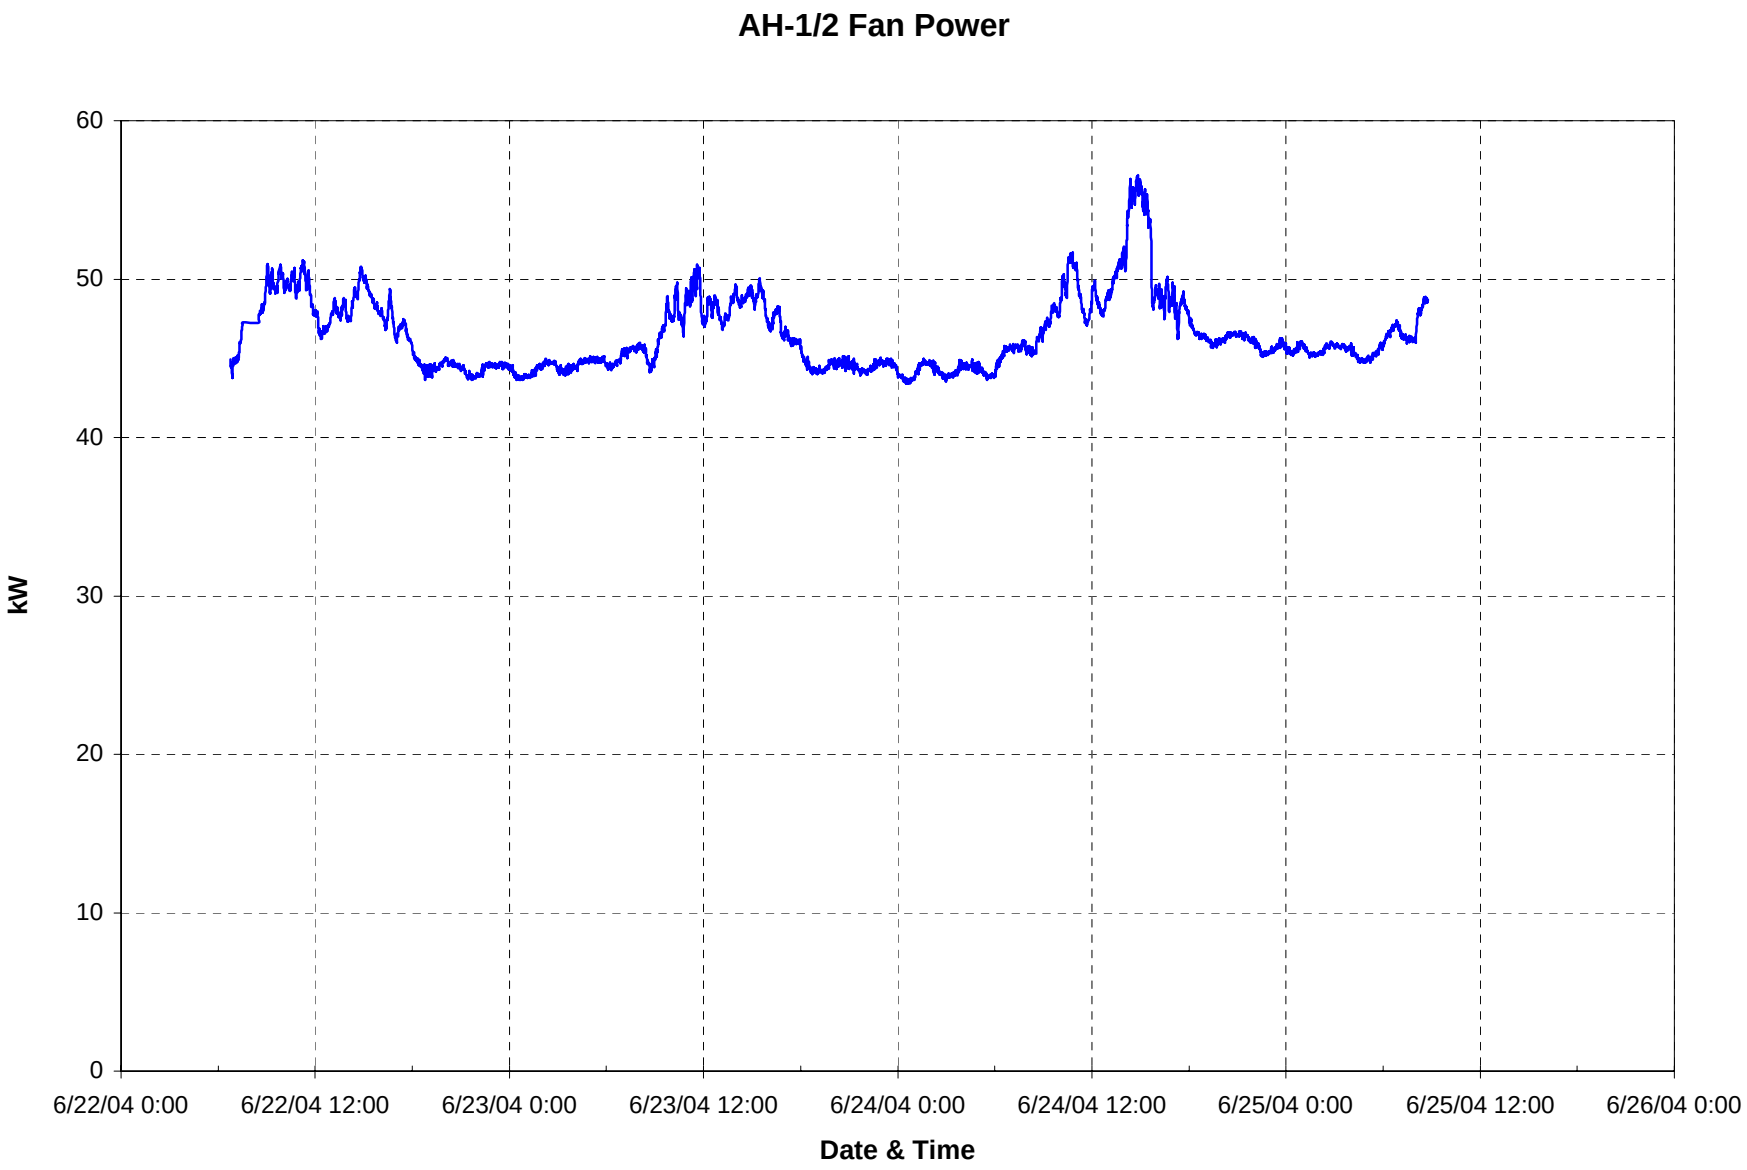
AH-1/2 Fan Power
60
50
40
30
20
10
0
kW
6/22/04 0:00
6/22/04 12:00
6/23/04 0:00
6/23/04 12:00
6/24/04 0:00
6/24/04 12:00
6/25/04 0:00
6/25/04 12:00
6/26/04 0:00
Date & Time

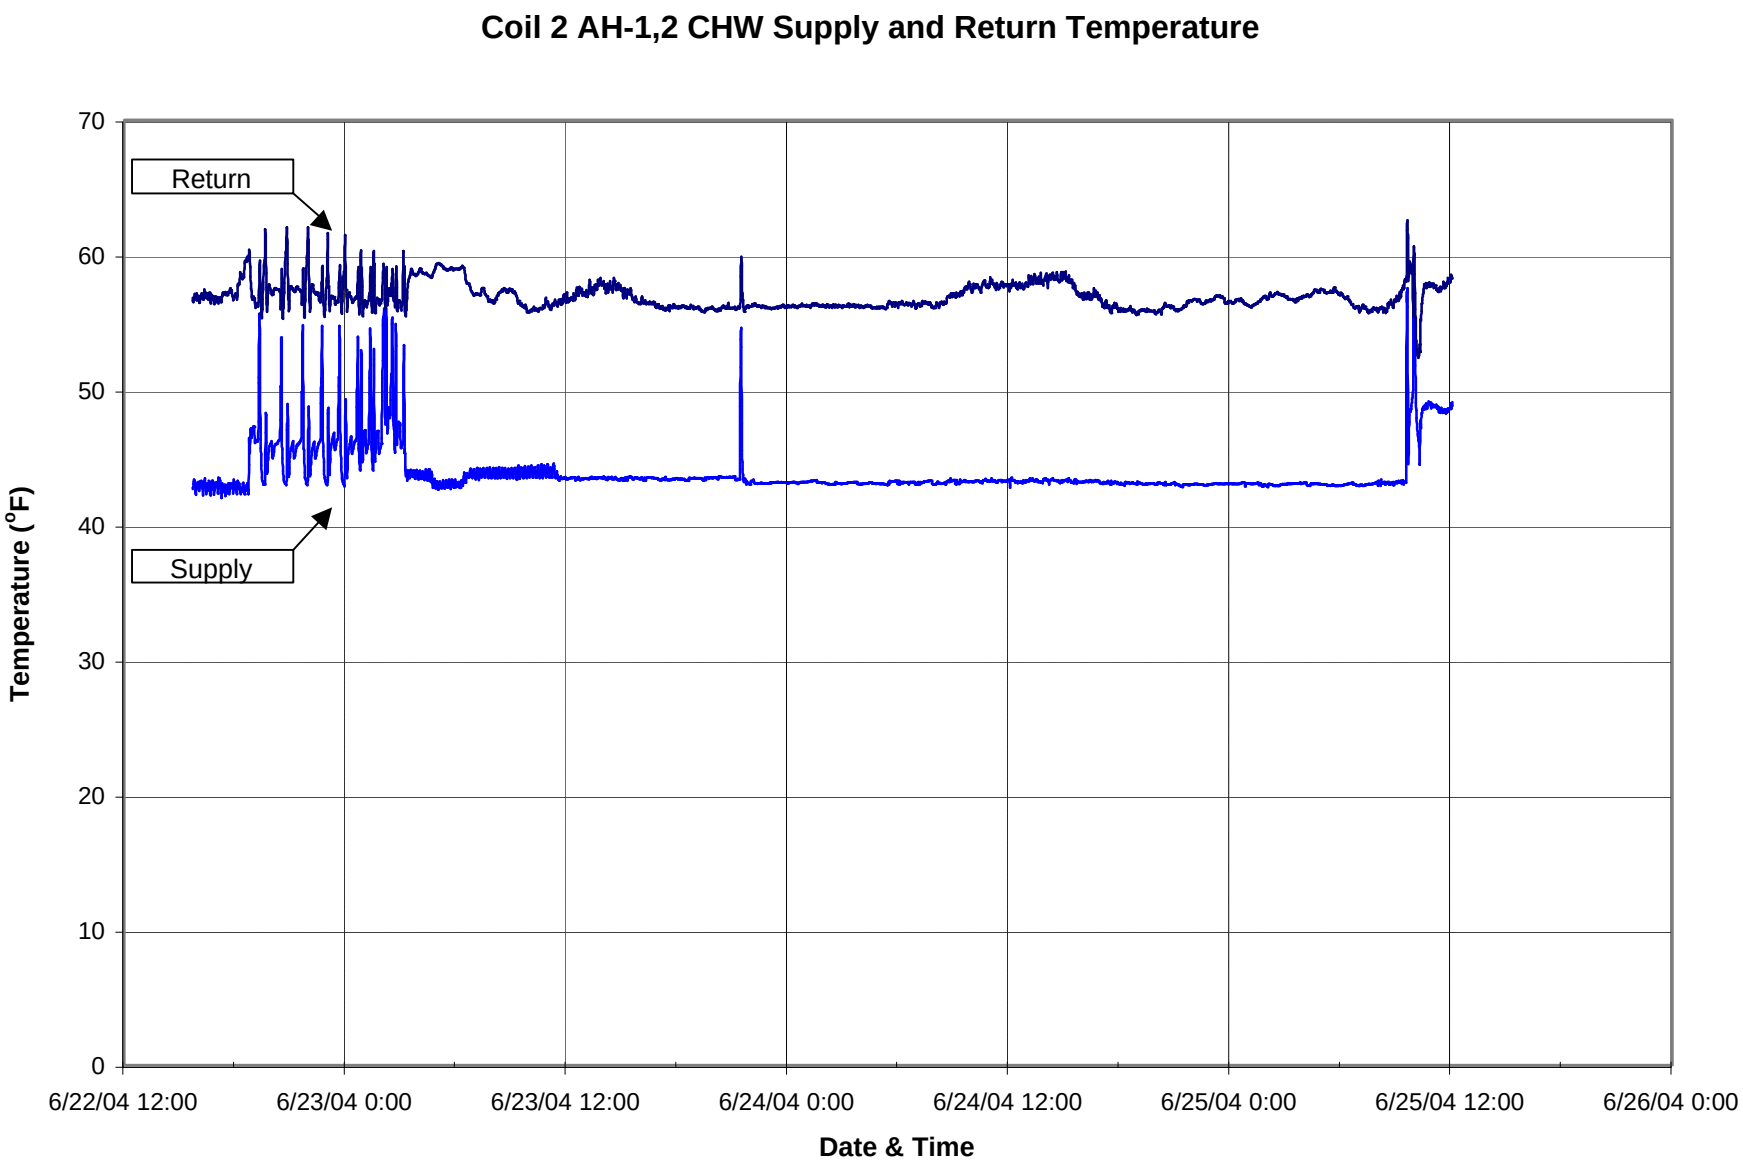
Coil 2 AH-1,2 CHW Supply and Return Temperature
Return
Supply
Temperature (°F)
70
60
50
40
30
20
10
0
6/22/04 12:00
6/23/04 0:00
6/23/04 12:00
6/24/04 0:00
6/24/04 12:00
6/25/04 0:00
6/25/04 12:00
6/26/04 0:00
Date & Time

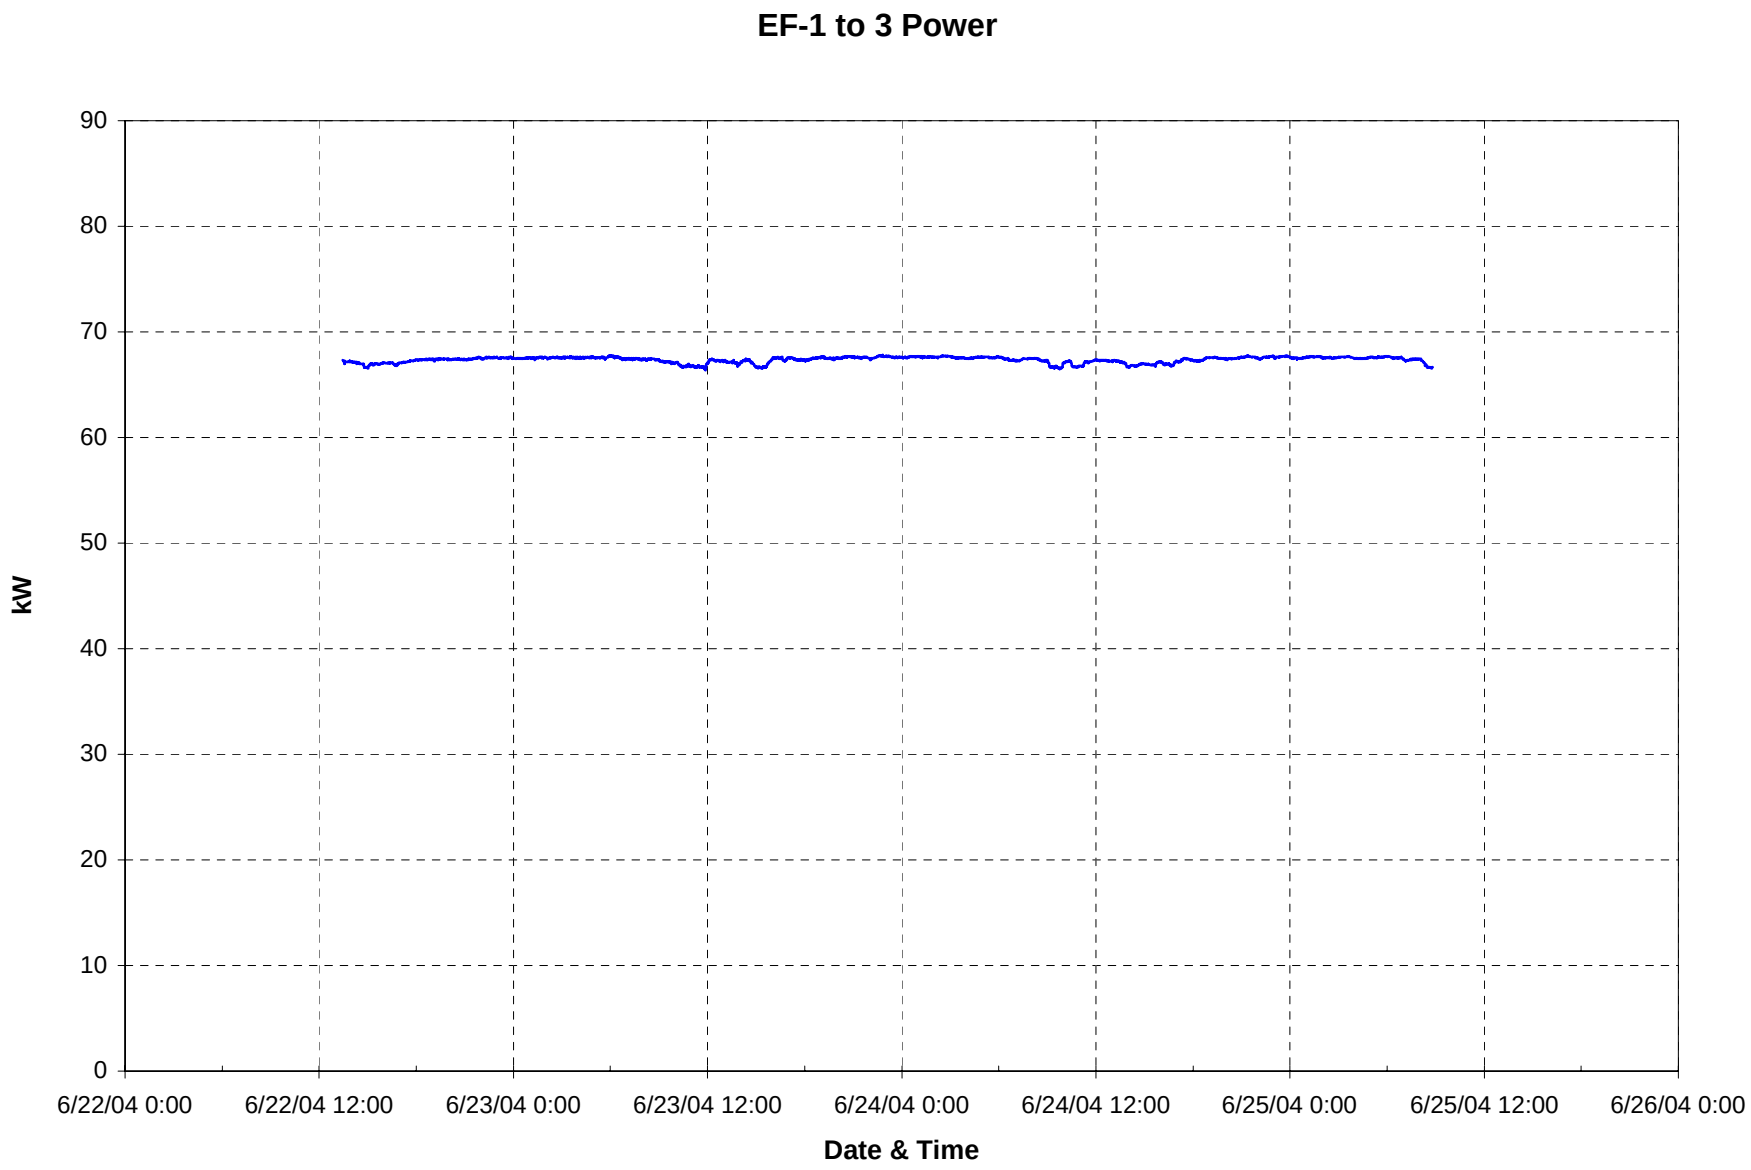
EF-1 to 3 Power
90
80
70
60
50
40
30
20
10
0
kW
6/22/04 0:00
6/22/04 12:00
6/23/04 0:00
6/23/04 12:00
6/24/04 0:00
6/24/04 12:00
6/25/04 0:00
6/25/04 12:00
6/26/04 0:00
Date & Time

“East Wing” Chemistry & Biology Laboratories

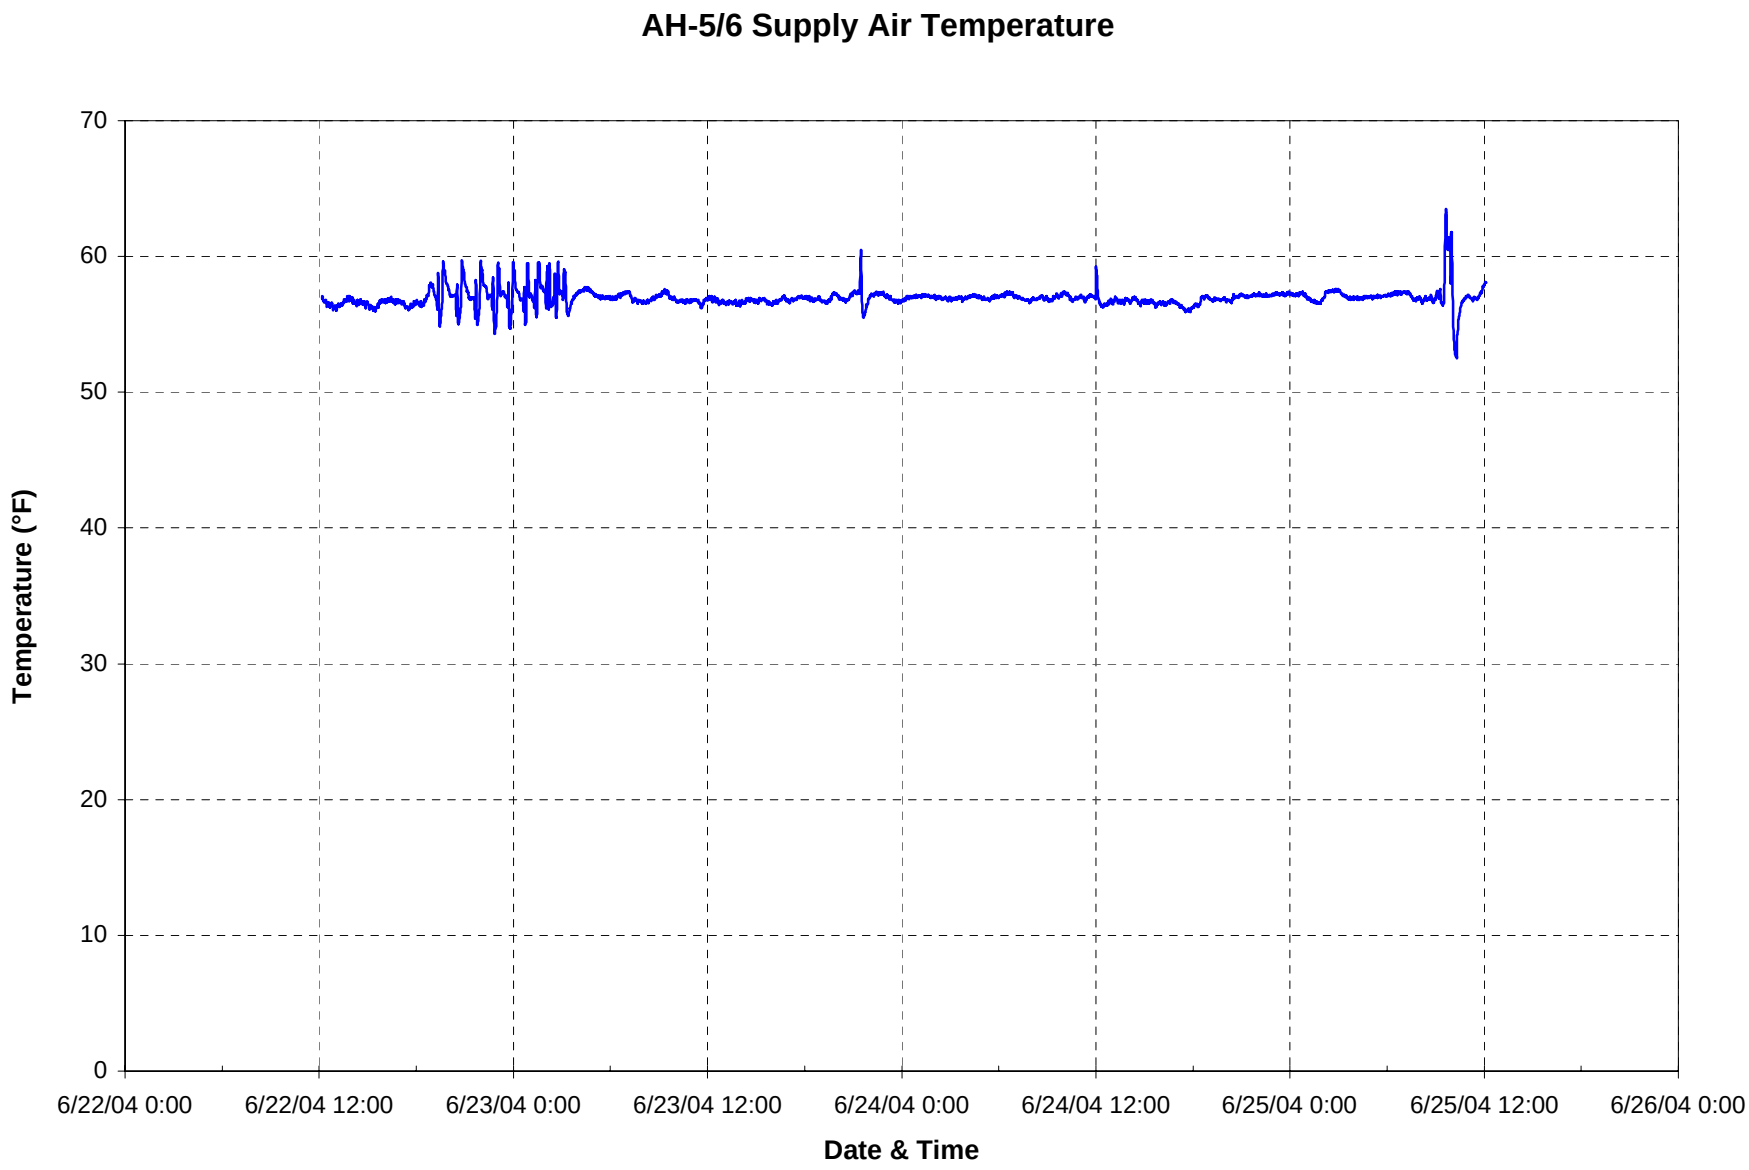
AH-5/6 Supply Air Temperature
Temperature (°F)
70
60
50
40
30
20
10
0
6/22/04 0:00
6/22/04 12:00
6/23/04 0:00
6/23/04 12:00
6/24/04 0:00
6/24/04 12:00
6/25/04 0:00
6/25/04 12:00
6/26/04 0:00
Date & Time

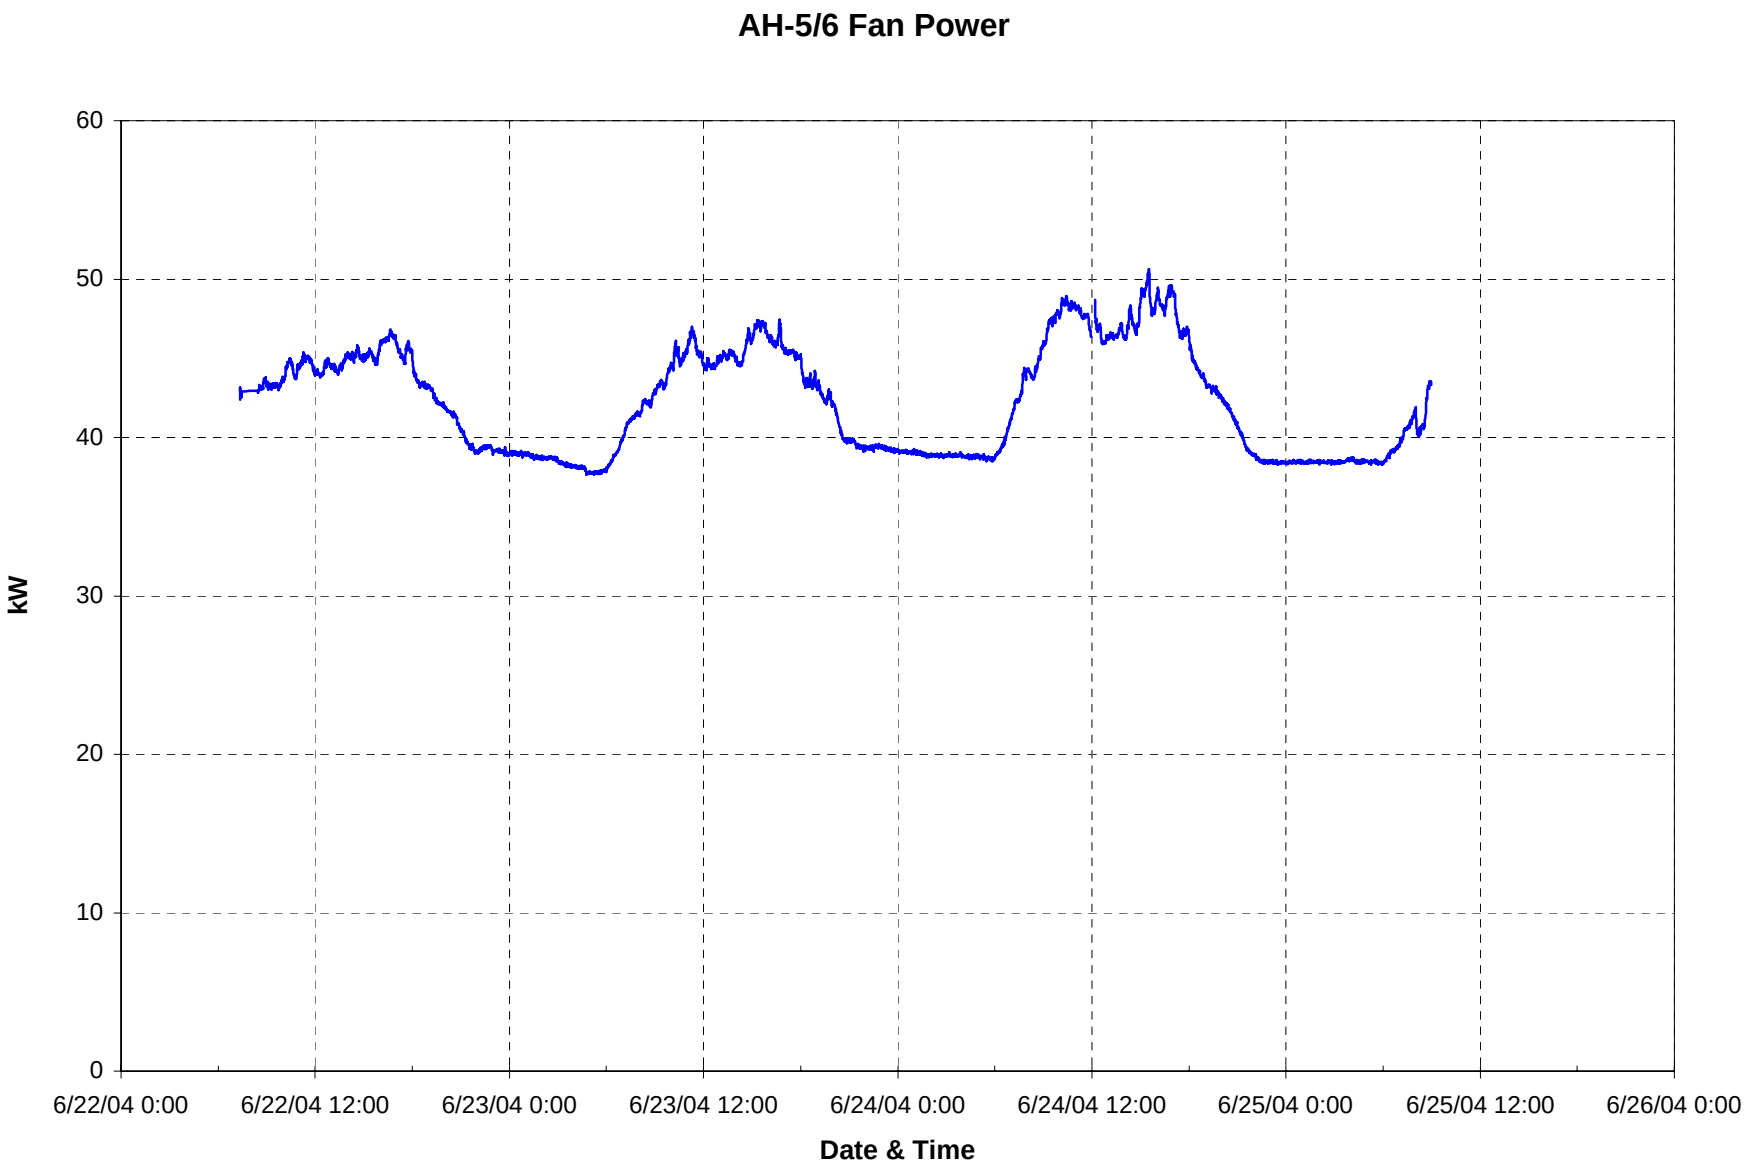
AH-5/6 Fan Power
60
50
40
30
20
10
0
kW
6/22/04 0:00
6/22/04 12:00
6/23/04 0:00
6/23/04 12:00
6/24/04 0:00
6/24/04 12:00
6/25/04 0:00
6/25/04 12:00
6/26/04 0:00
Date & Time

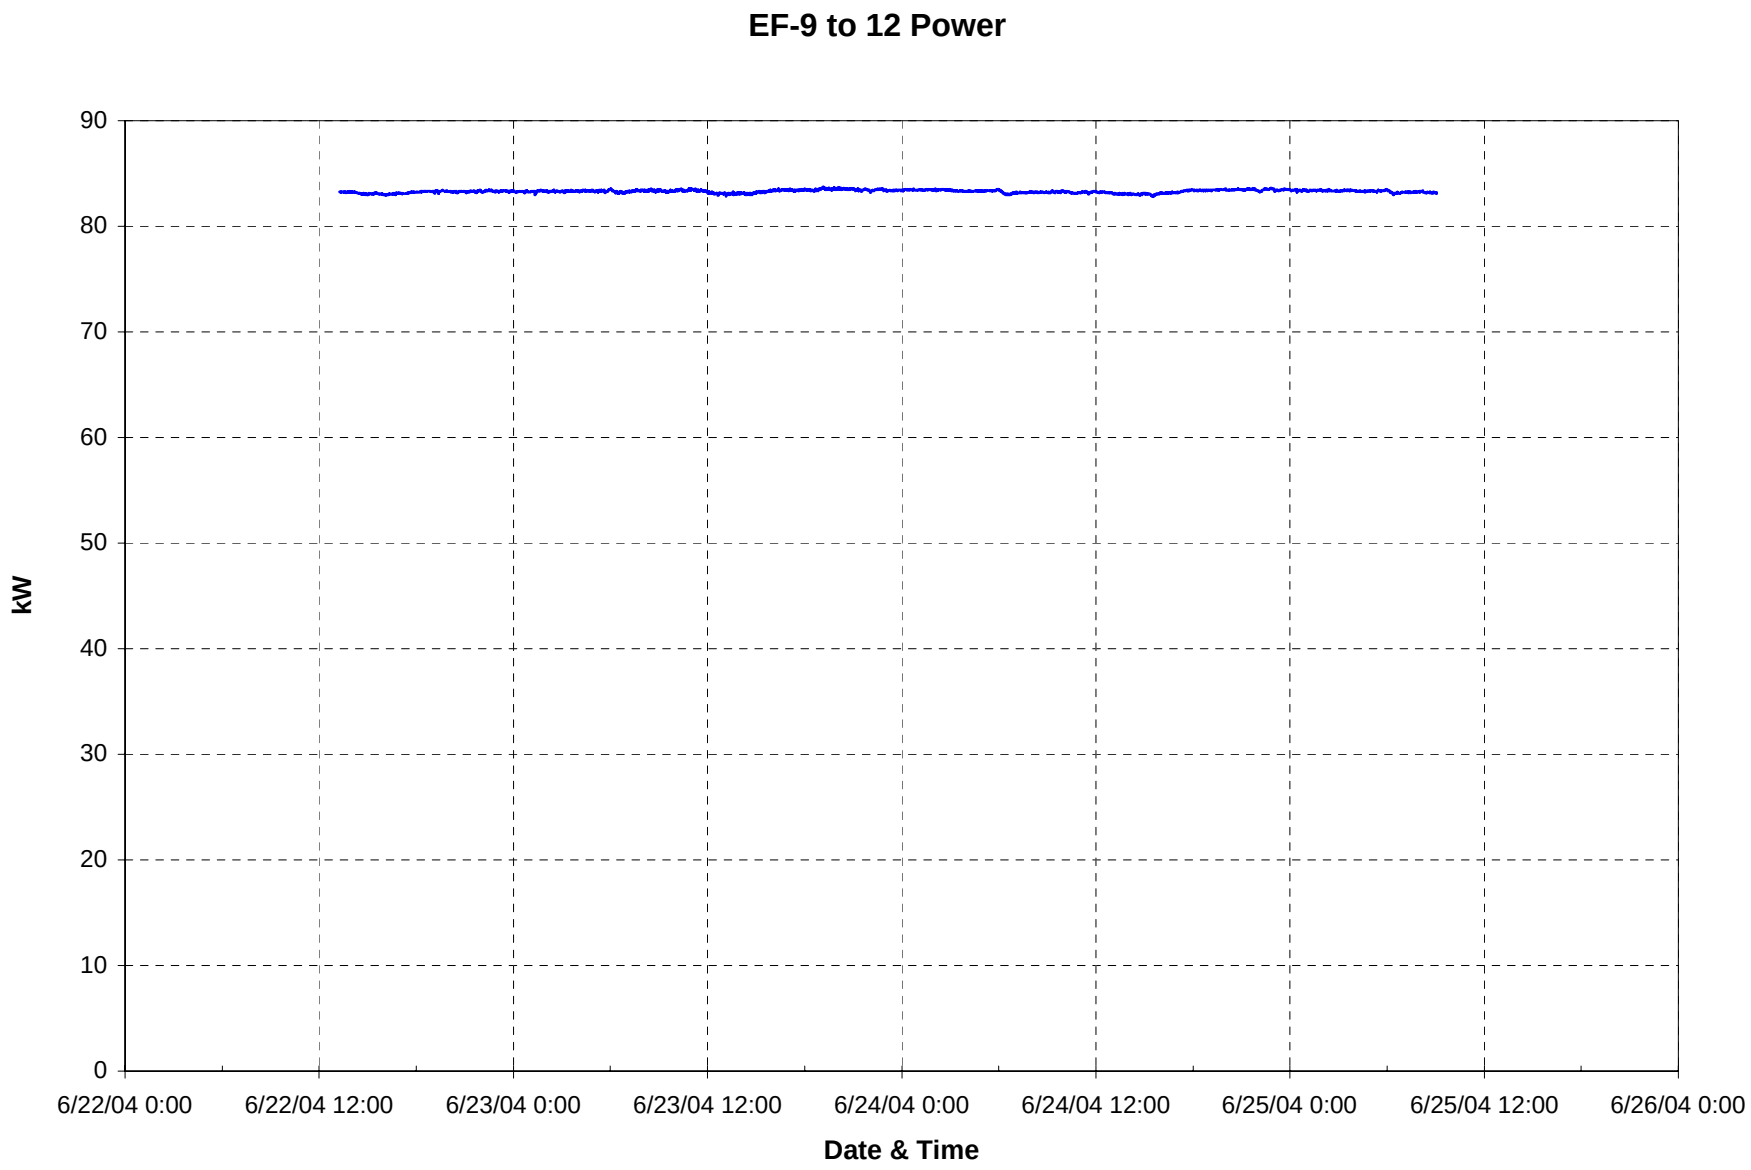
EF-9 to 12 Power
90
80
70
60
50
40
30
20
10
0
kW
6/22/04 0:00
6/22/04 12:00
6/23/04 0:00
6/23/04 12:00
6/24/04 0:00
6/24/04 12:00
6/25/04 0:00
6/25/04 12:00
6/26/04 0:00
Date & Time

Electrical Systems

Backup Generator

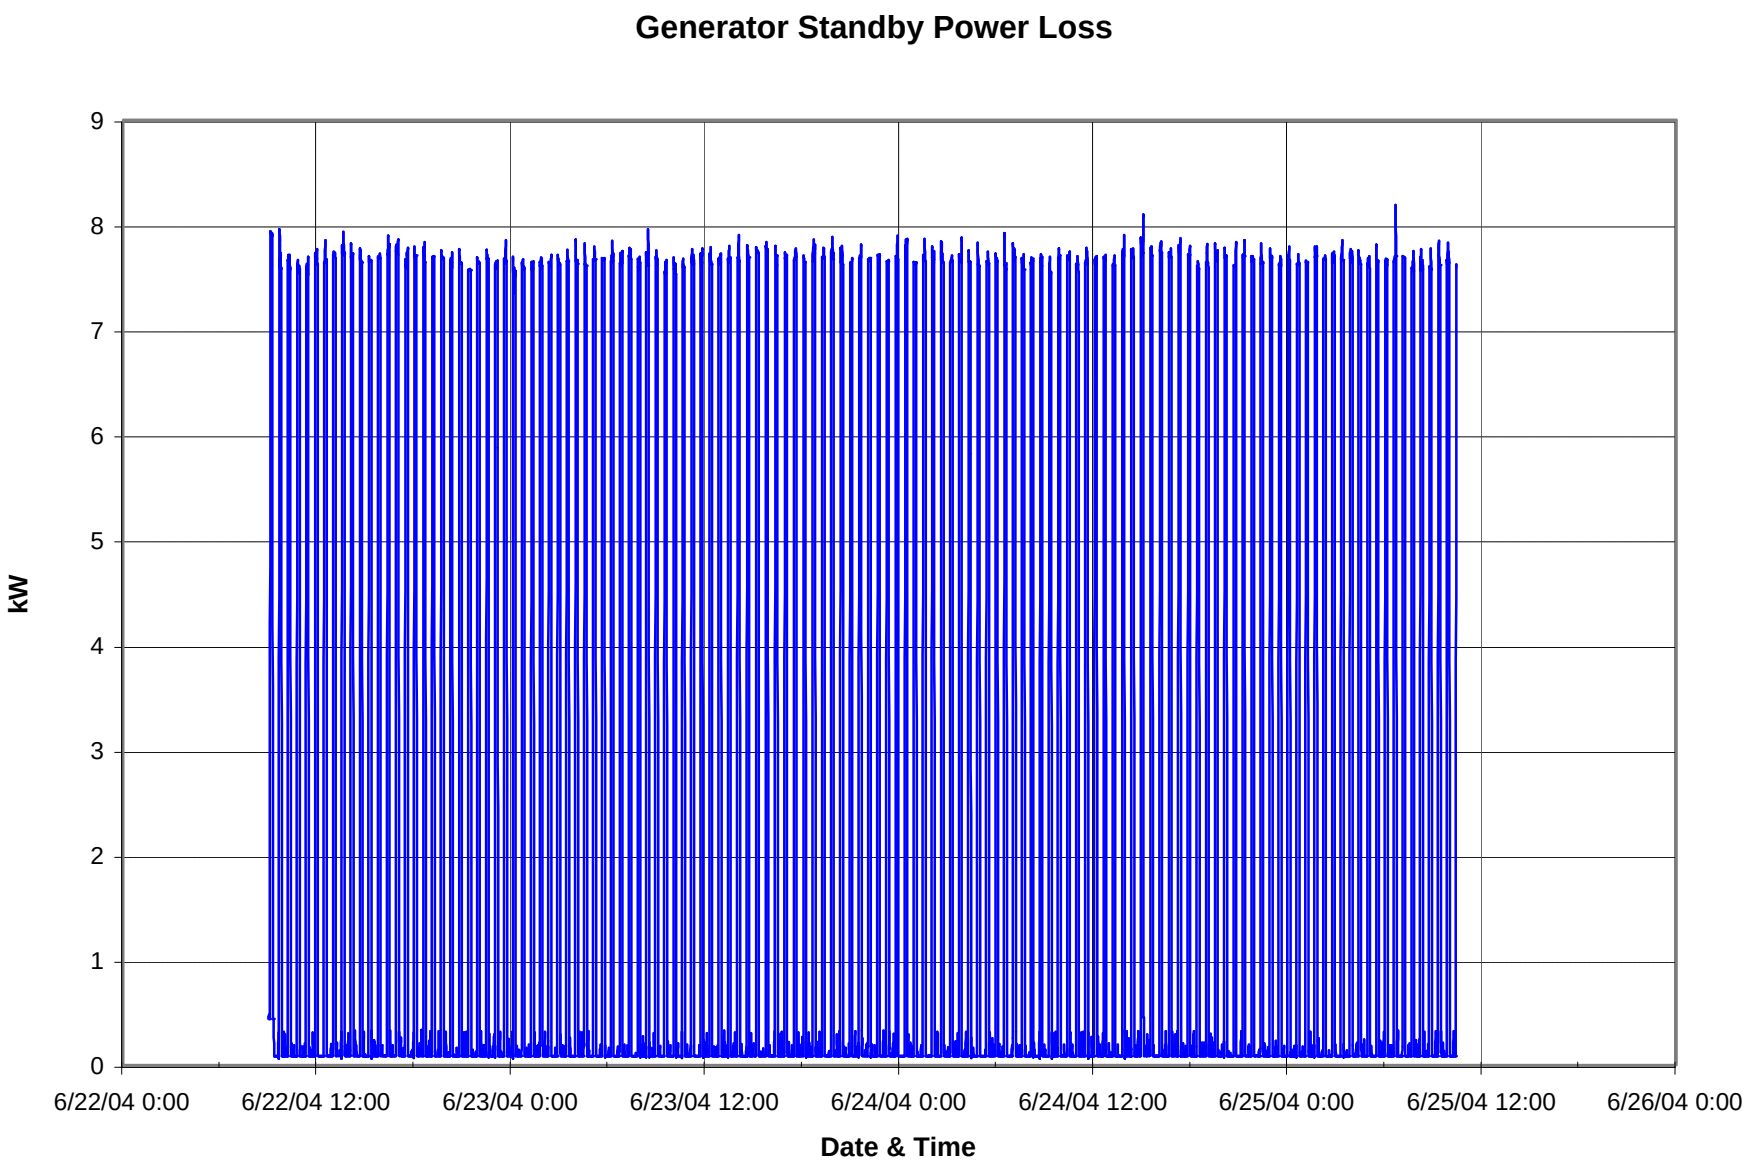
Generator Standby Power Loss
kW
9
8
7
6
5
4
3
2
1
0
6/22/04 0:00
6/22/04 12:00
6/23/04 0:00
6/23/04 12:00
6/24/04 0:00
6/24/04 12:00
6/25/04 0:00
6/25/04 12:00
6/26/04 0:00
Date & Time

Miscellaneous Electrical Measurements

Lighting Panel 1HA - Circuit 20

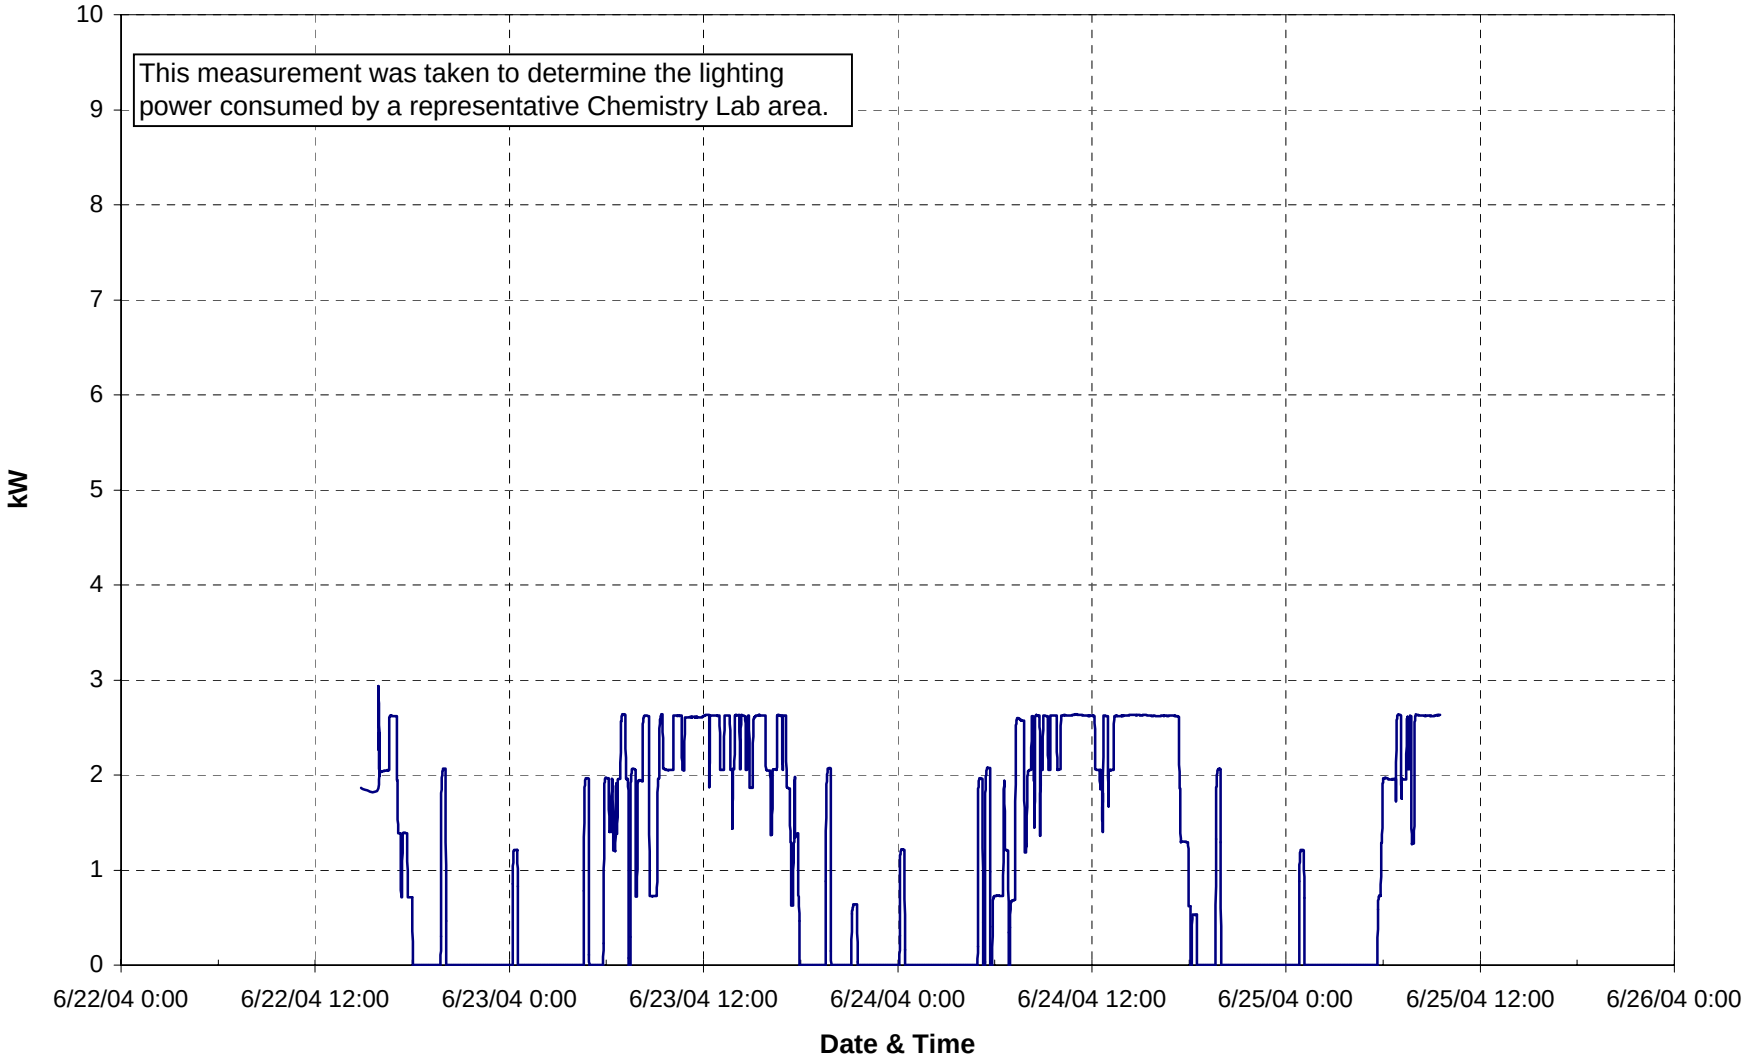

Lighting Panel 1HA - Circuits 23, 27

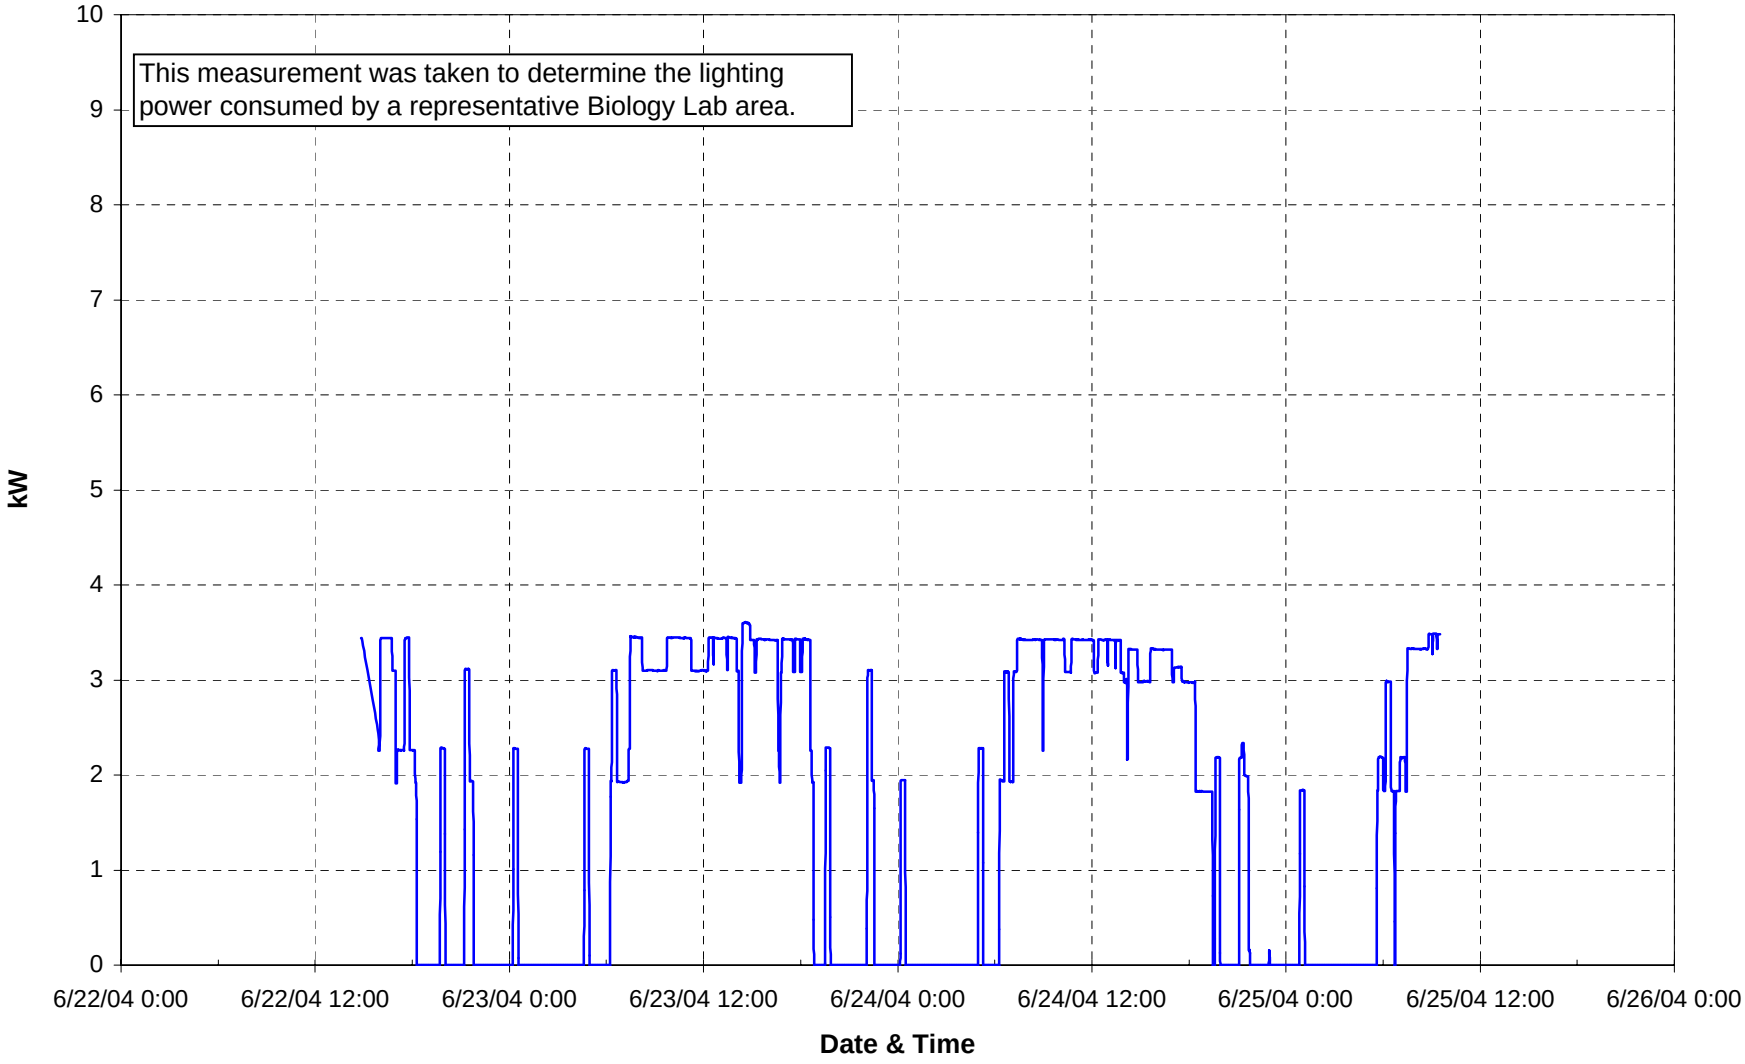

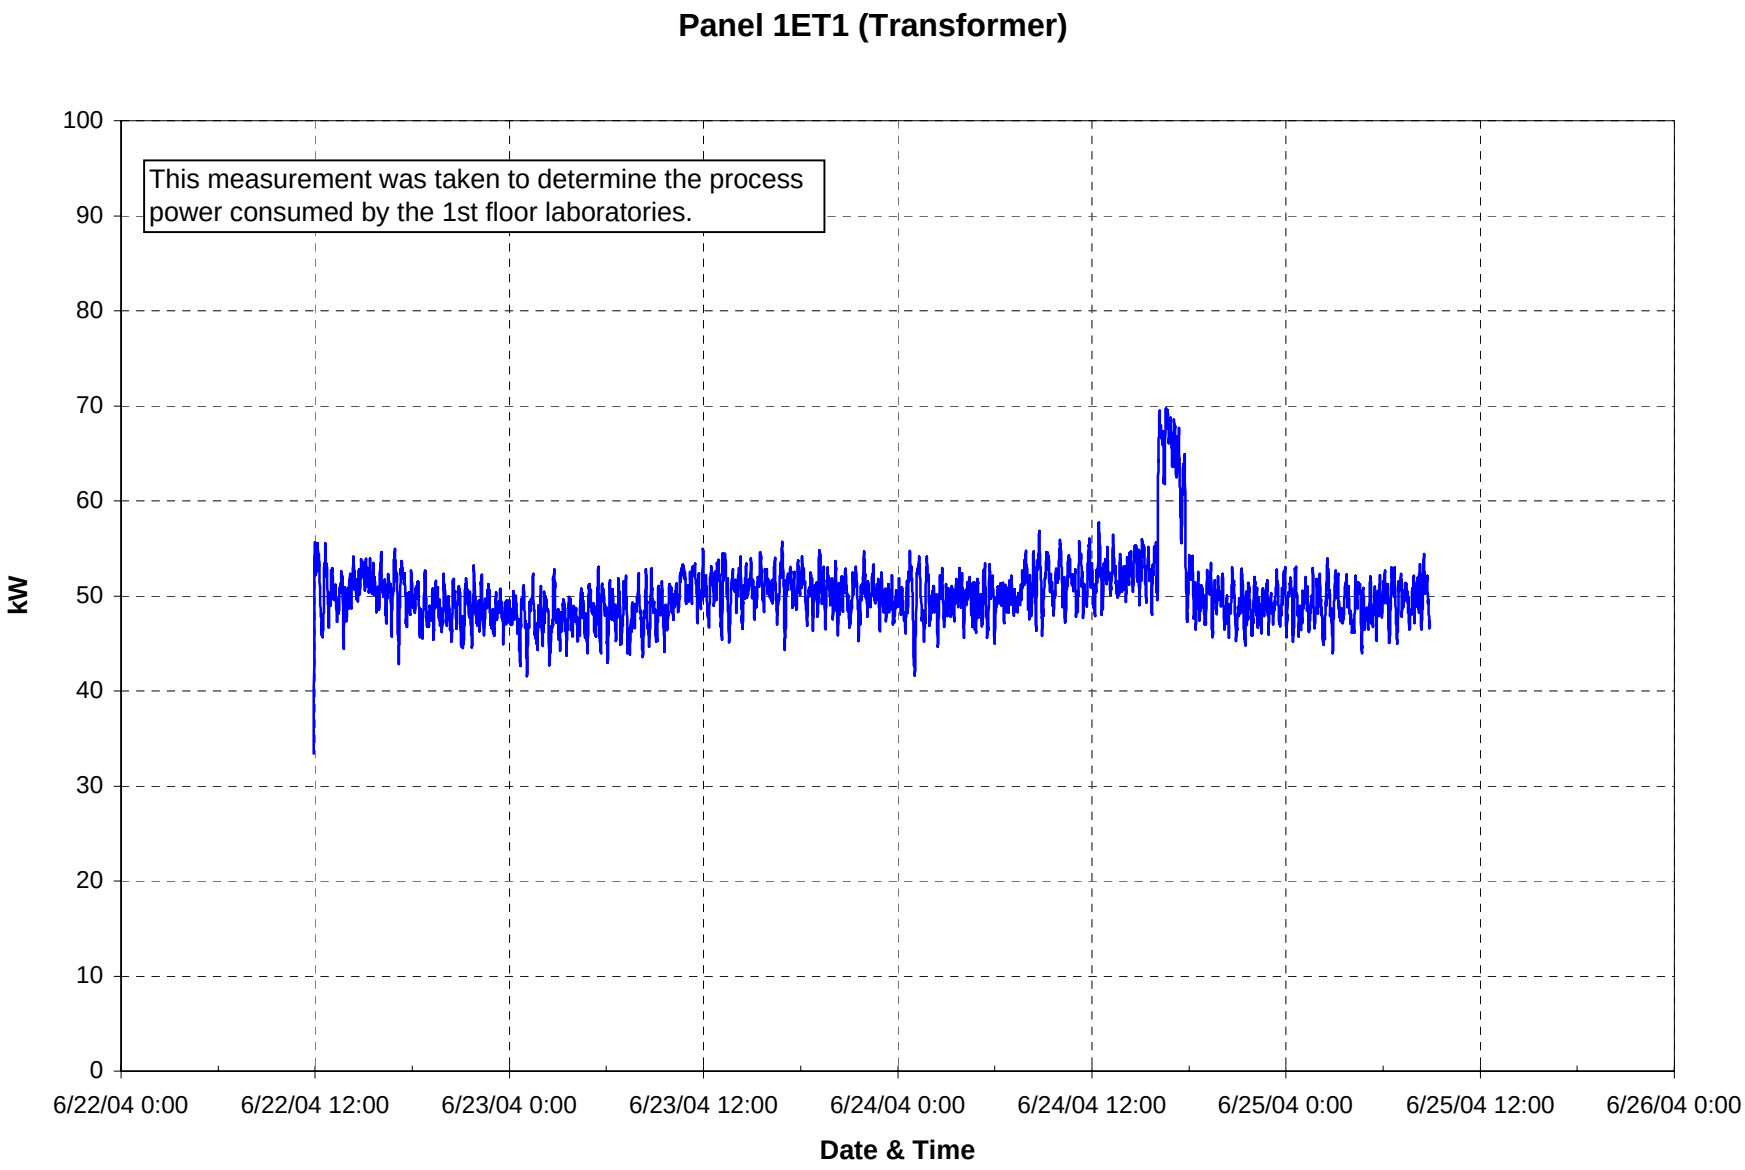
Panel 1ET1 (Transformer)
This measurement was taken to determine the process power consumed by the 1st floor laboratories.
kW
100
90
80
70
60
50
40
30
20
10
0
6/22/04 0:00
6/22/04 12:00
6/23/04 0:00
6/23/04 12:00
6/24/04 0:00
6/24/04 12:00
6/25/04 0:00
6/25/04 12:00
6/26/04 0:00
Date & Time

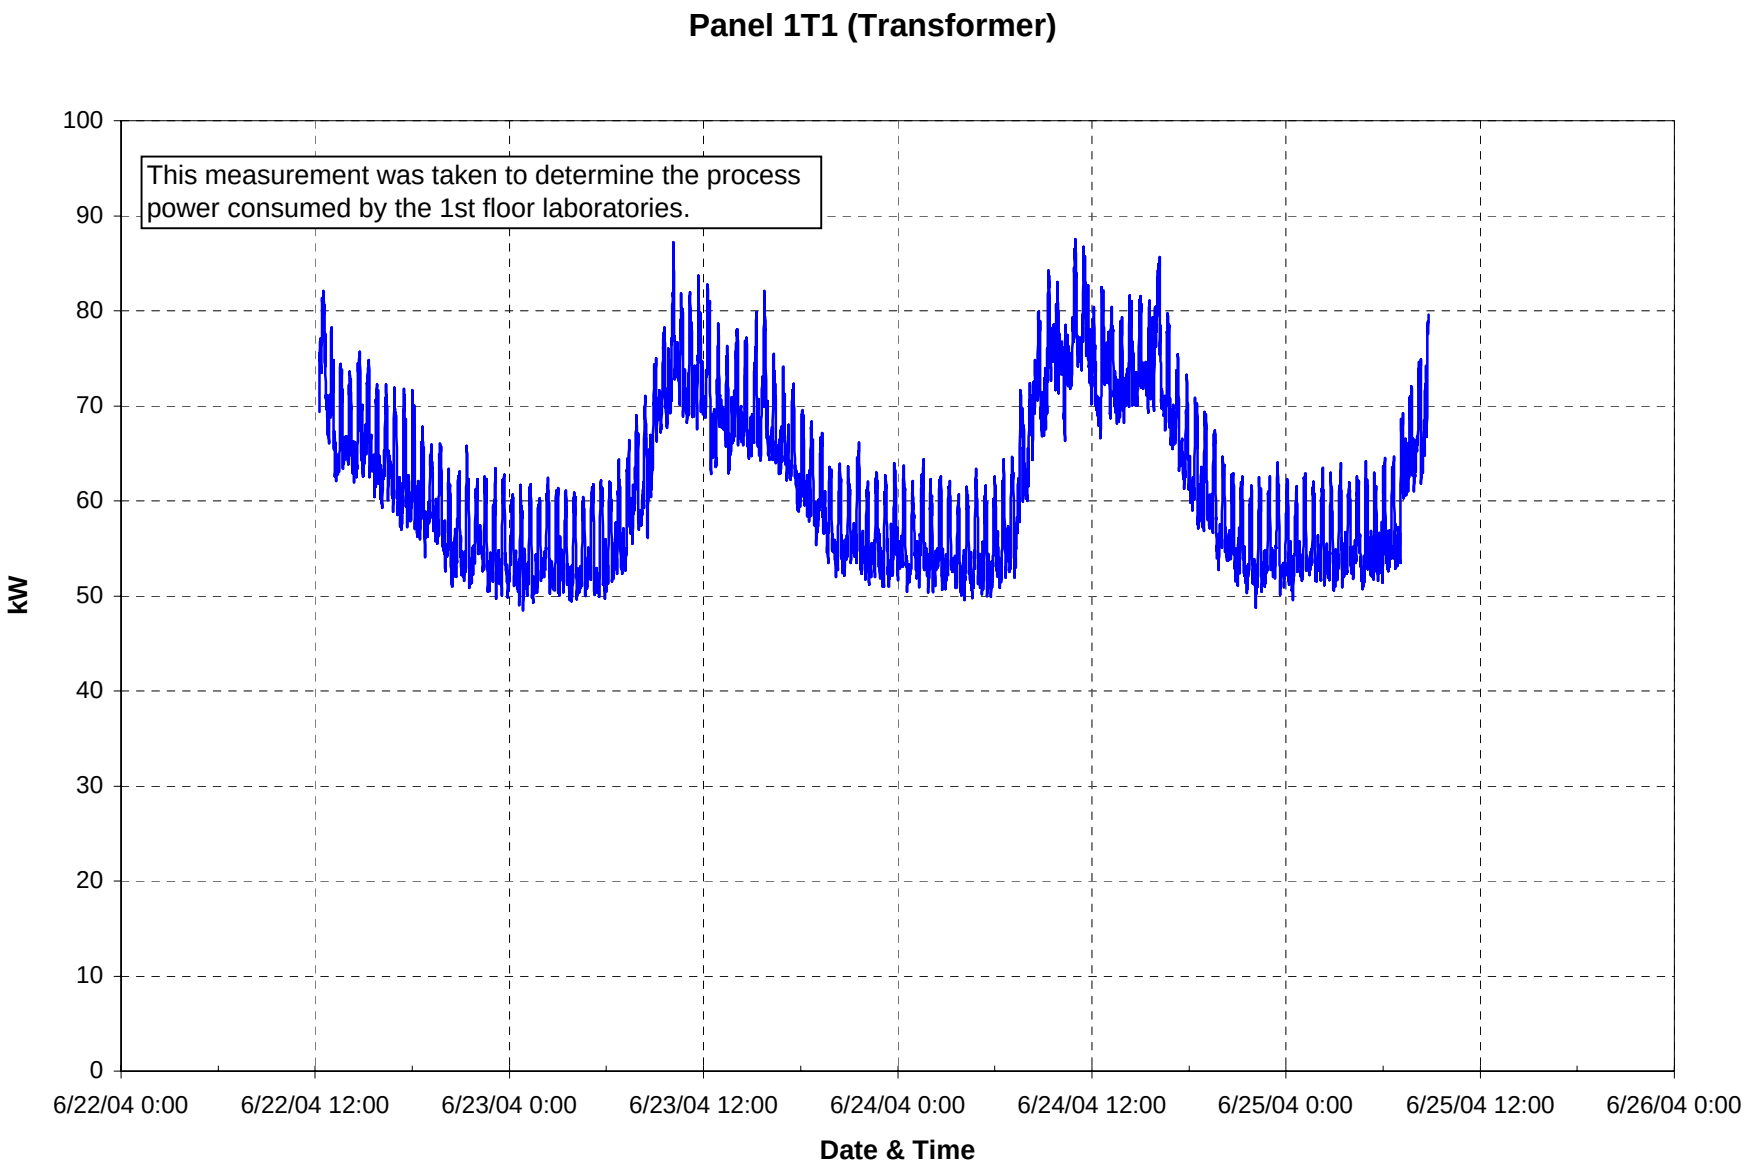
Panel 1T1 (Transformer)
This measurement was taken to determine the process power consumed by the 1st floor laboratories.
kW
100
90
80
70
60
50
40
30
20
10
0
6/22/04 0:00
6/22/04 12:00
6/23/04 0:00
6/23/04 12:00
6/24/04 0:00
6/24/04 12:00
6/25/04 0:00
6/25/04 12:00
6/26/04 0:00
Date & Time

Panel 1LF - Circuits 2, 4, 6

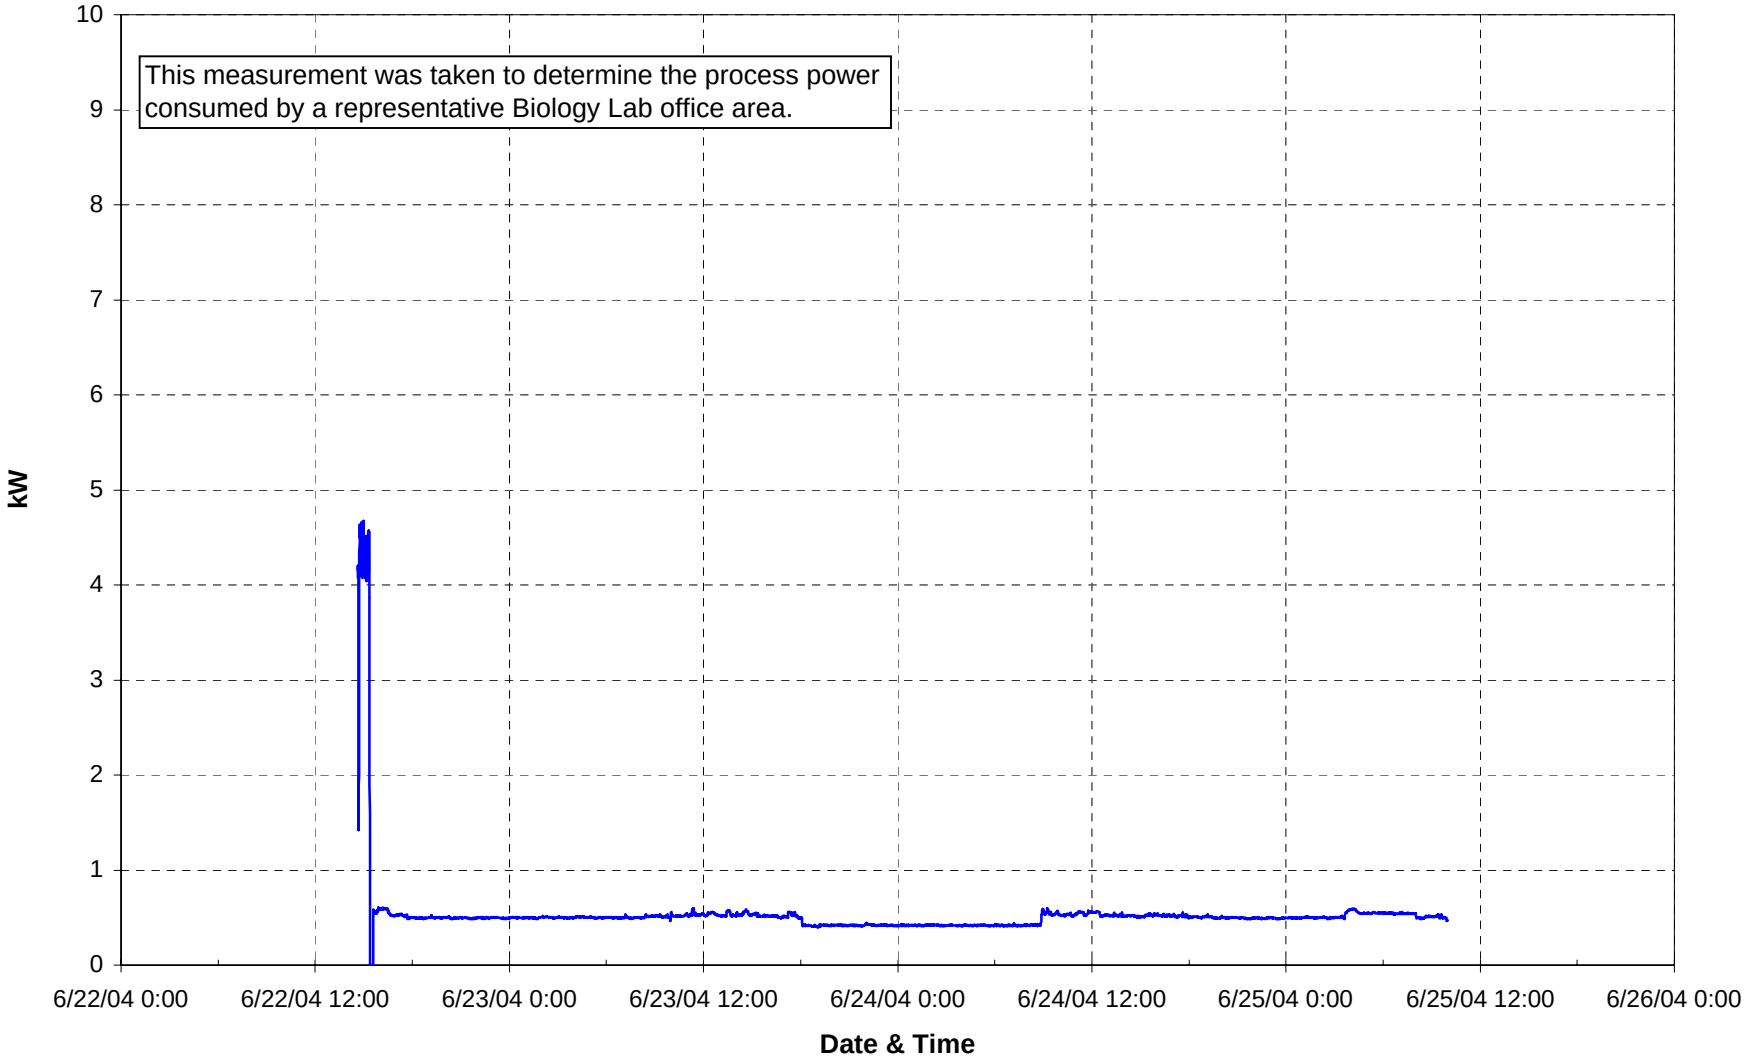

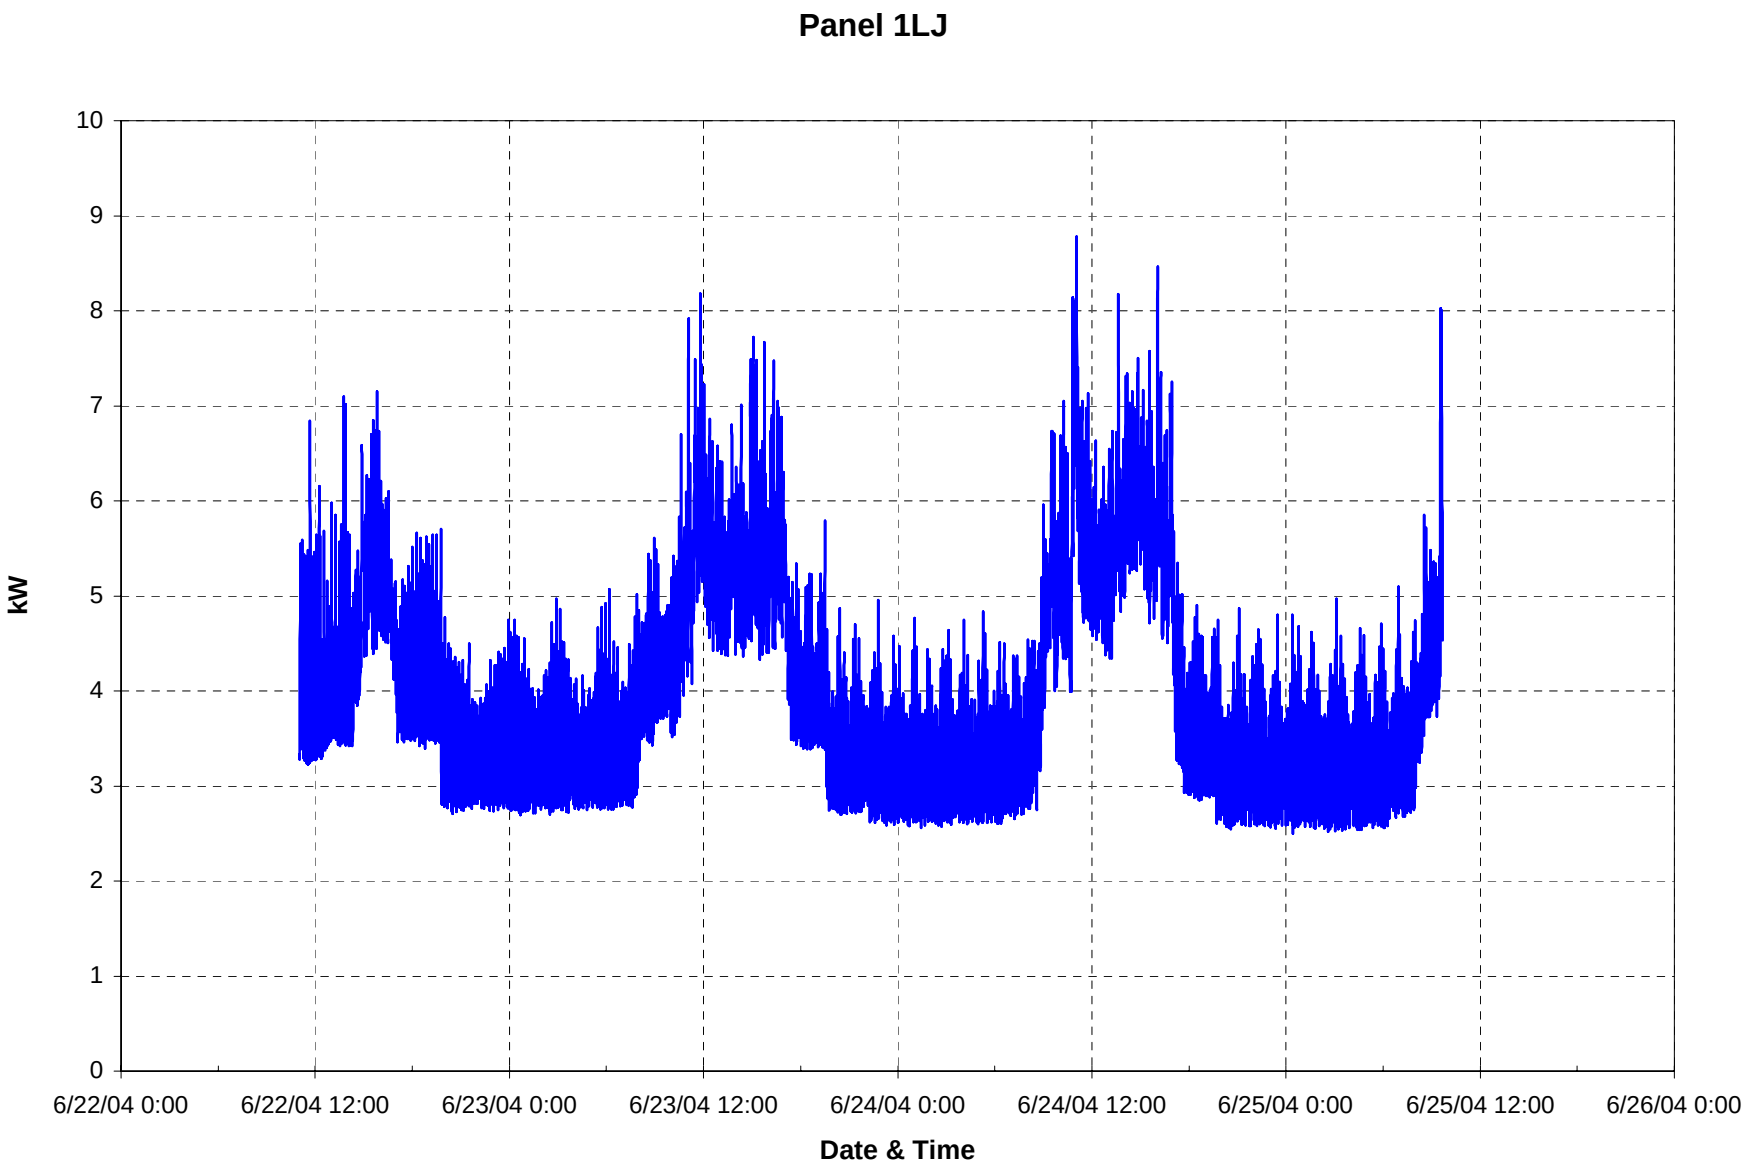
Panel 1LJ
kW
10
9
8
7
6
5
4
3
2
1
0
6/22/04 0:00
6/22/04 12:00
6/23/04 0:00
6/23/04 12:00
6/24/04 0:00
6/24/04 12:00
6/25/04 0:00
6/25/04 12:00
6/26/04 0:00
Date & Time

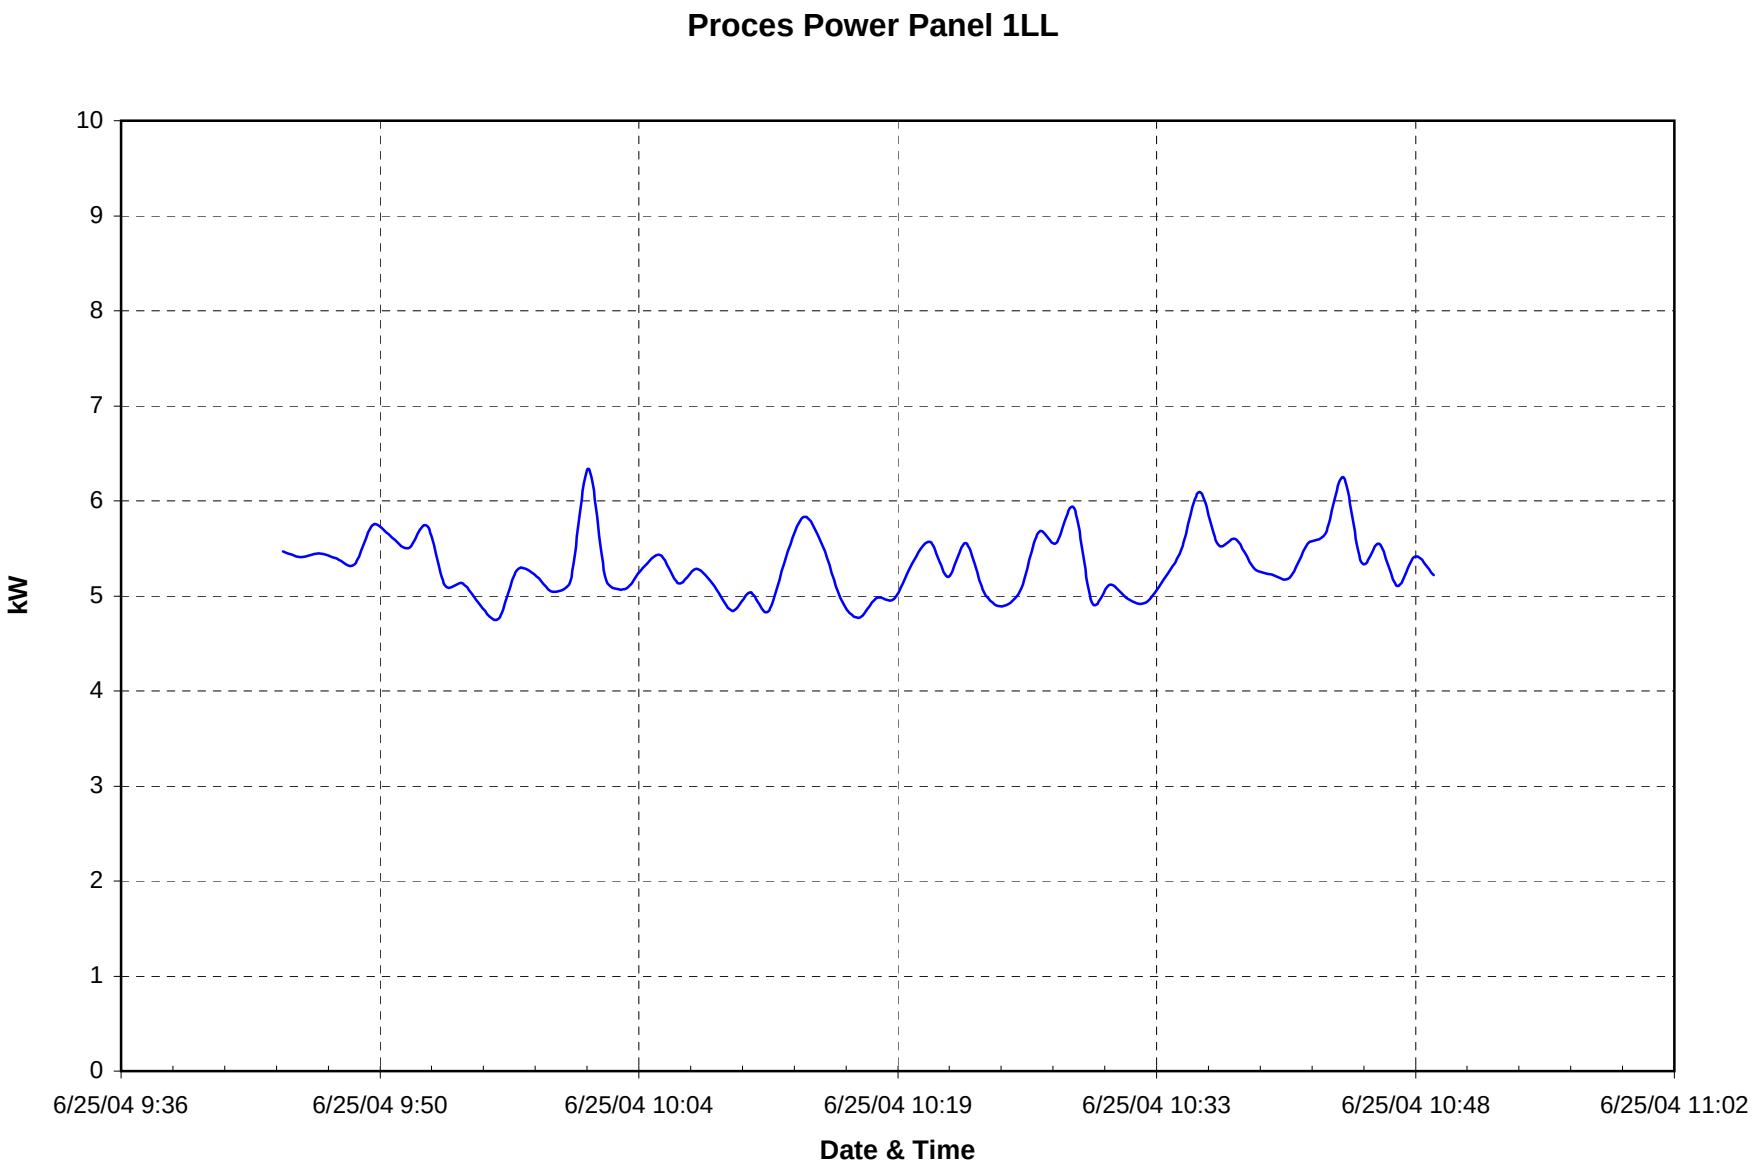
Proces Power Panel 1LL
kW
10
9
8
7
6
5
4
3
2
1
0
6/25/04 9:36
6/25/04 9:50
6/25/04 10:04
6/25/04 10:19
6/25/04 10:33
6/25/04 10:48
6/25/04 11:02
Date & Time

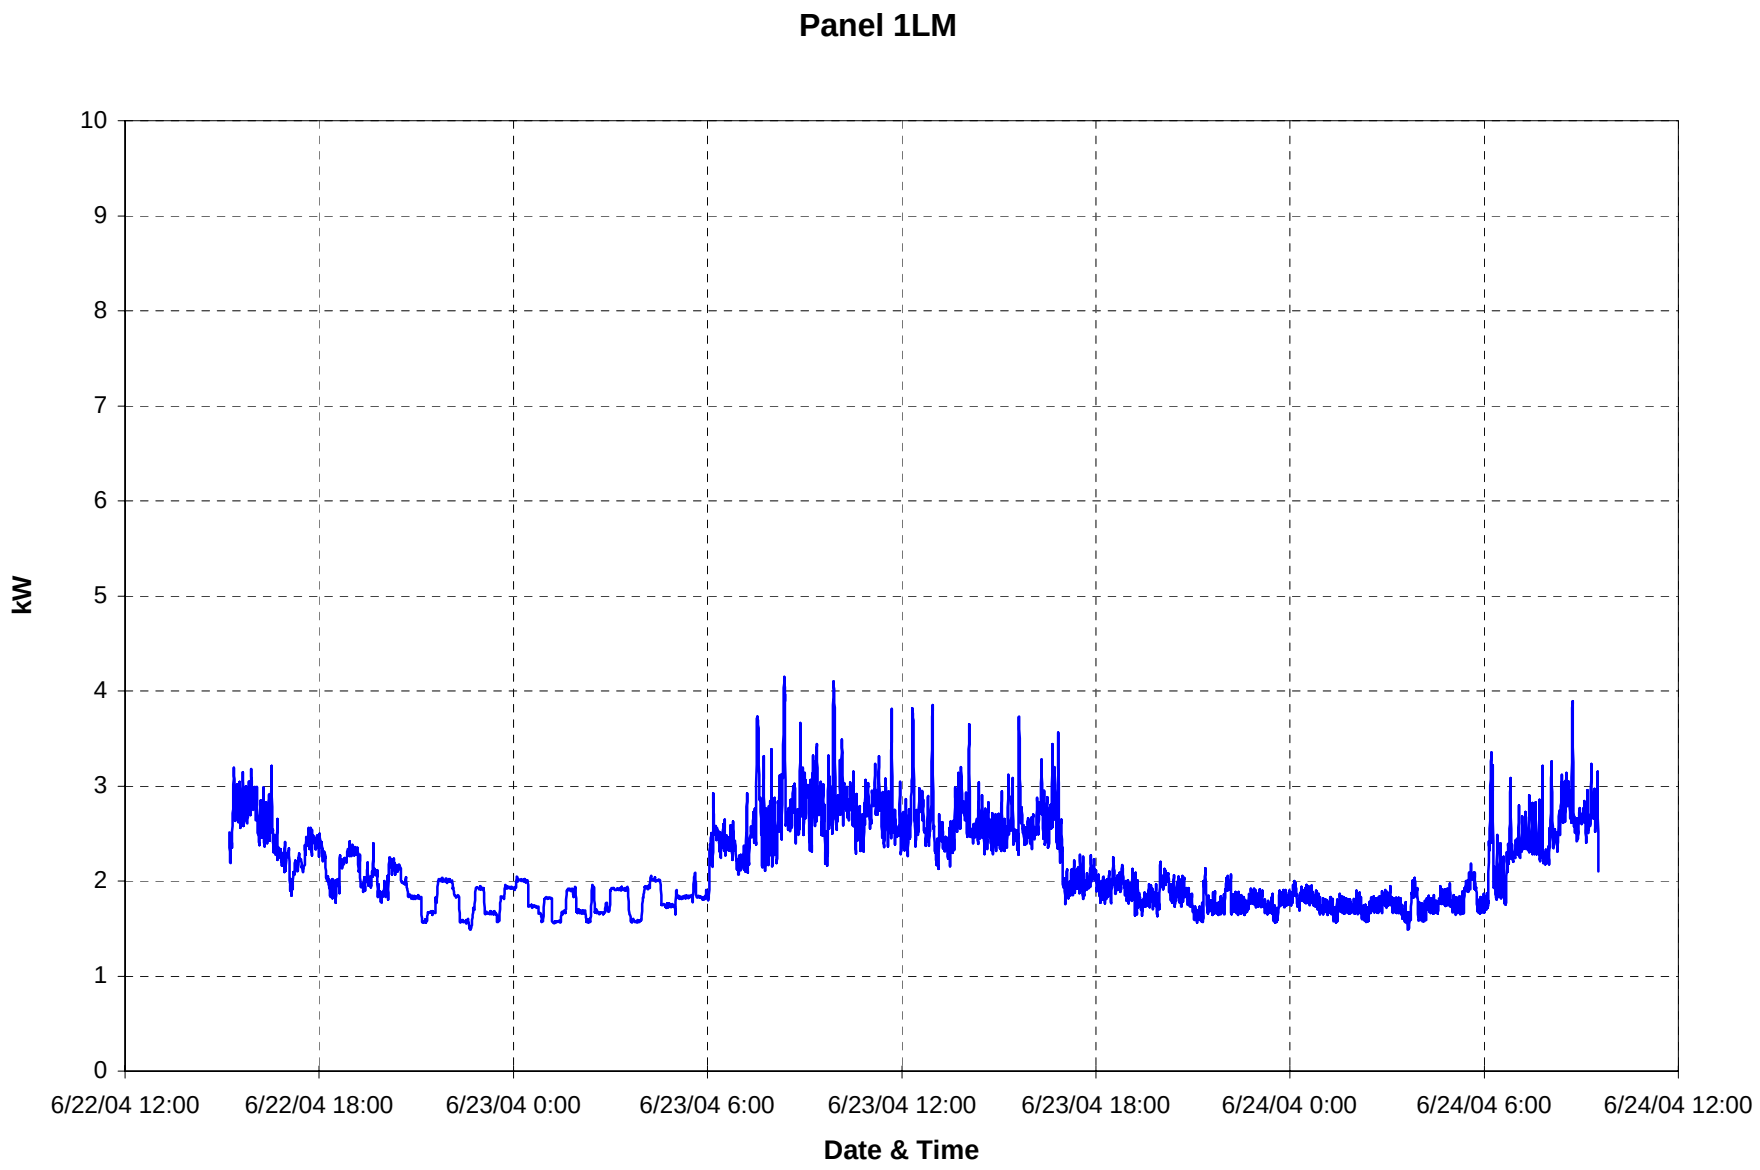
Panel 1LM
kW
10
9
8
7
6
5
4
3
2
1
0
6/22/04 12:00
6/22/04 18:00
6/23/04 0:00
6/23/04 6:00
6/23/04 12:00
6/23/04 18:00
6/24/04 0:00
6/24/04 6:00
6/24/04 12:00
Date & Time

Panel 1LM - Circuits 8, 10, 12

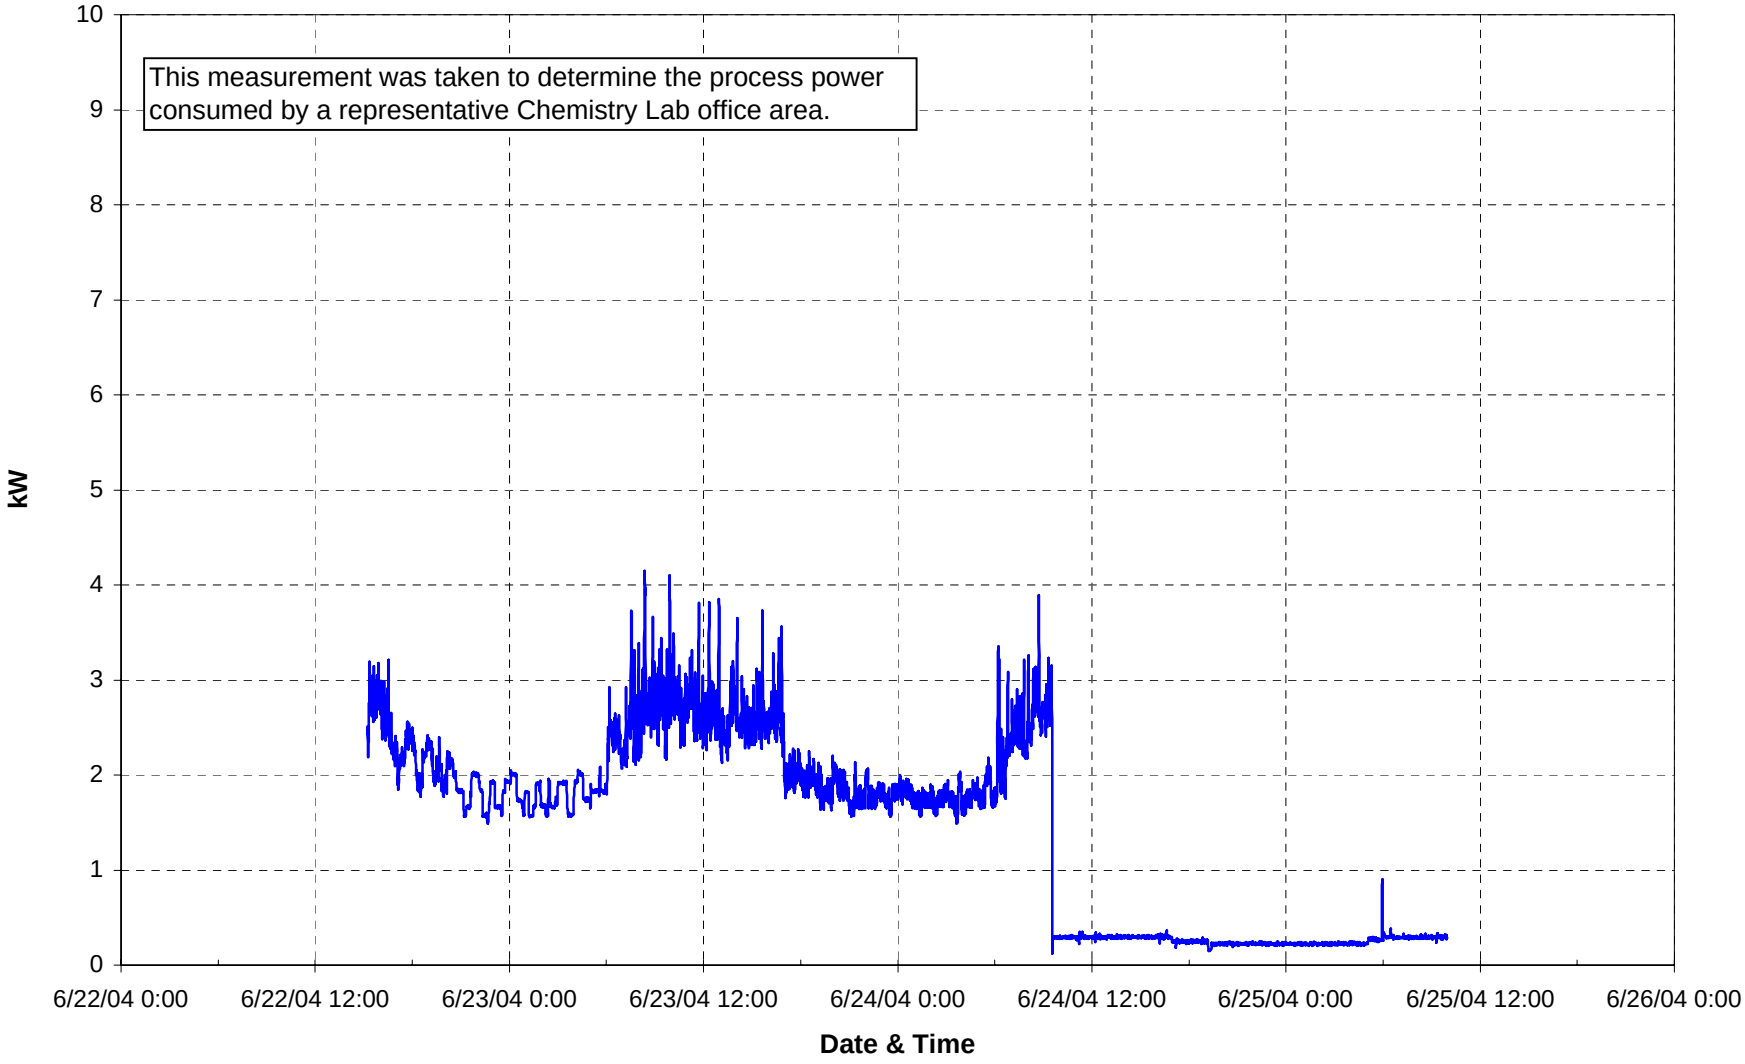

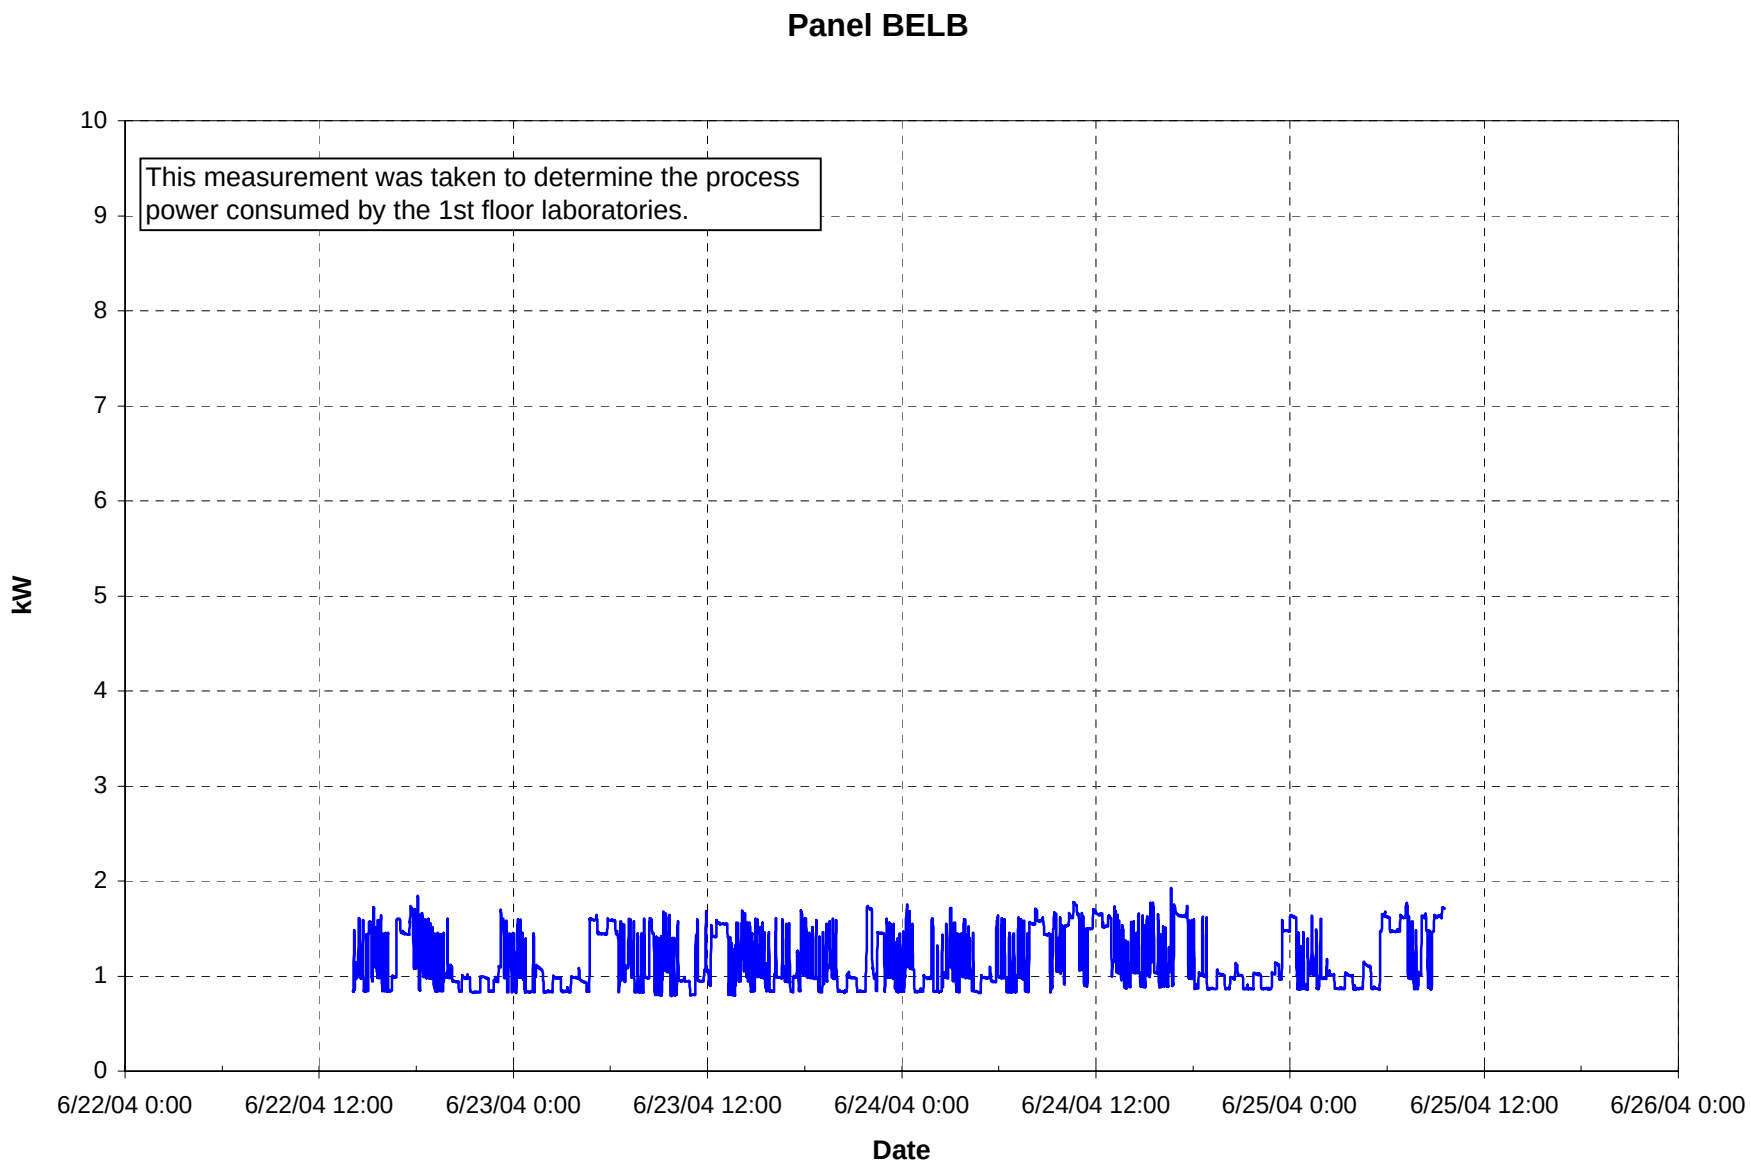
Panel BELB
This measurement was taken to determine the process power consumed by the 1st floor laboratories.
kW
10
9
8
7
6
5
4
3
2
1
0
6/22/04 0:00
6/22/04 12:00
6/23/04 0:00
6/23/04 12:00
6/24/04 0:00
6/24/04 12:00
6/25/04 0:00
6/25/04 12:00
6/26/04 0:00
Date

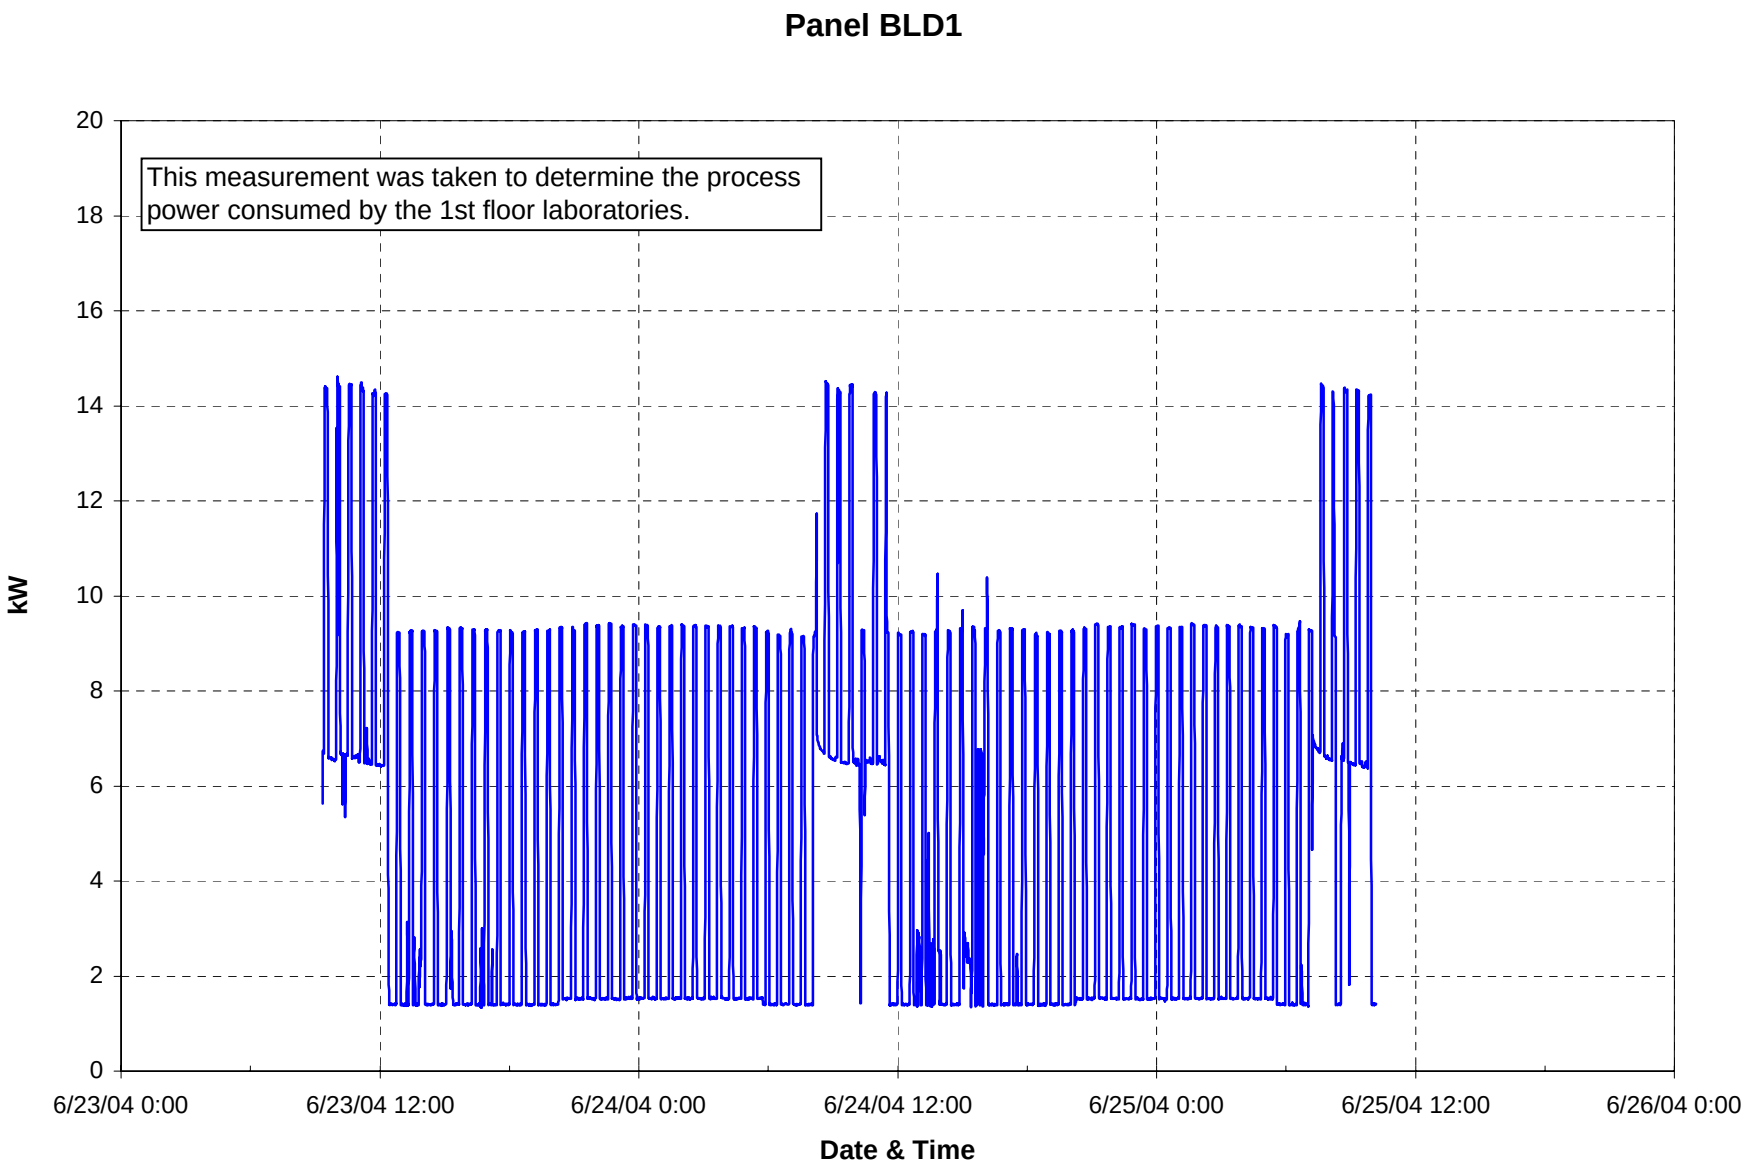
Panel BLD1
This measurement was taken to determine the process power consumed by the 1st floor laboratories.
kW
20
18
16
14
12
10
8
6
4
2
0
6/23/04 0:00
6/23/04 12:00
6/24/04 0:00
6/24/04 12:00
6/25/04 0:00
6/25/04 12:00
6/26/04 0:00
Date & Time

Absorption Chiller Measurements

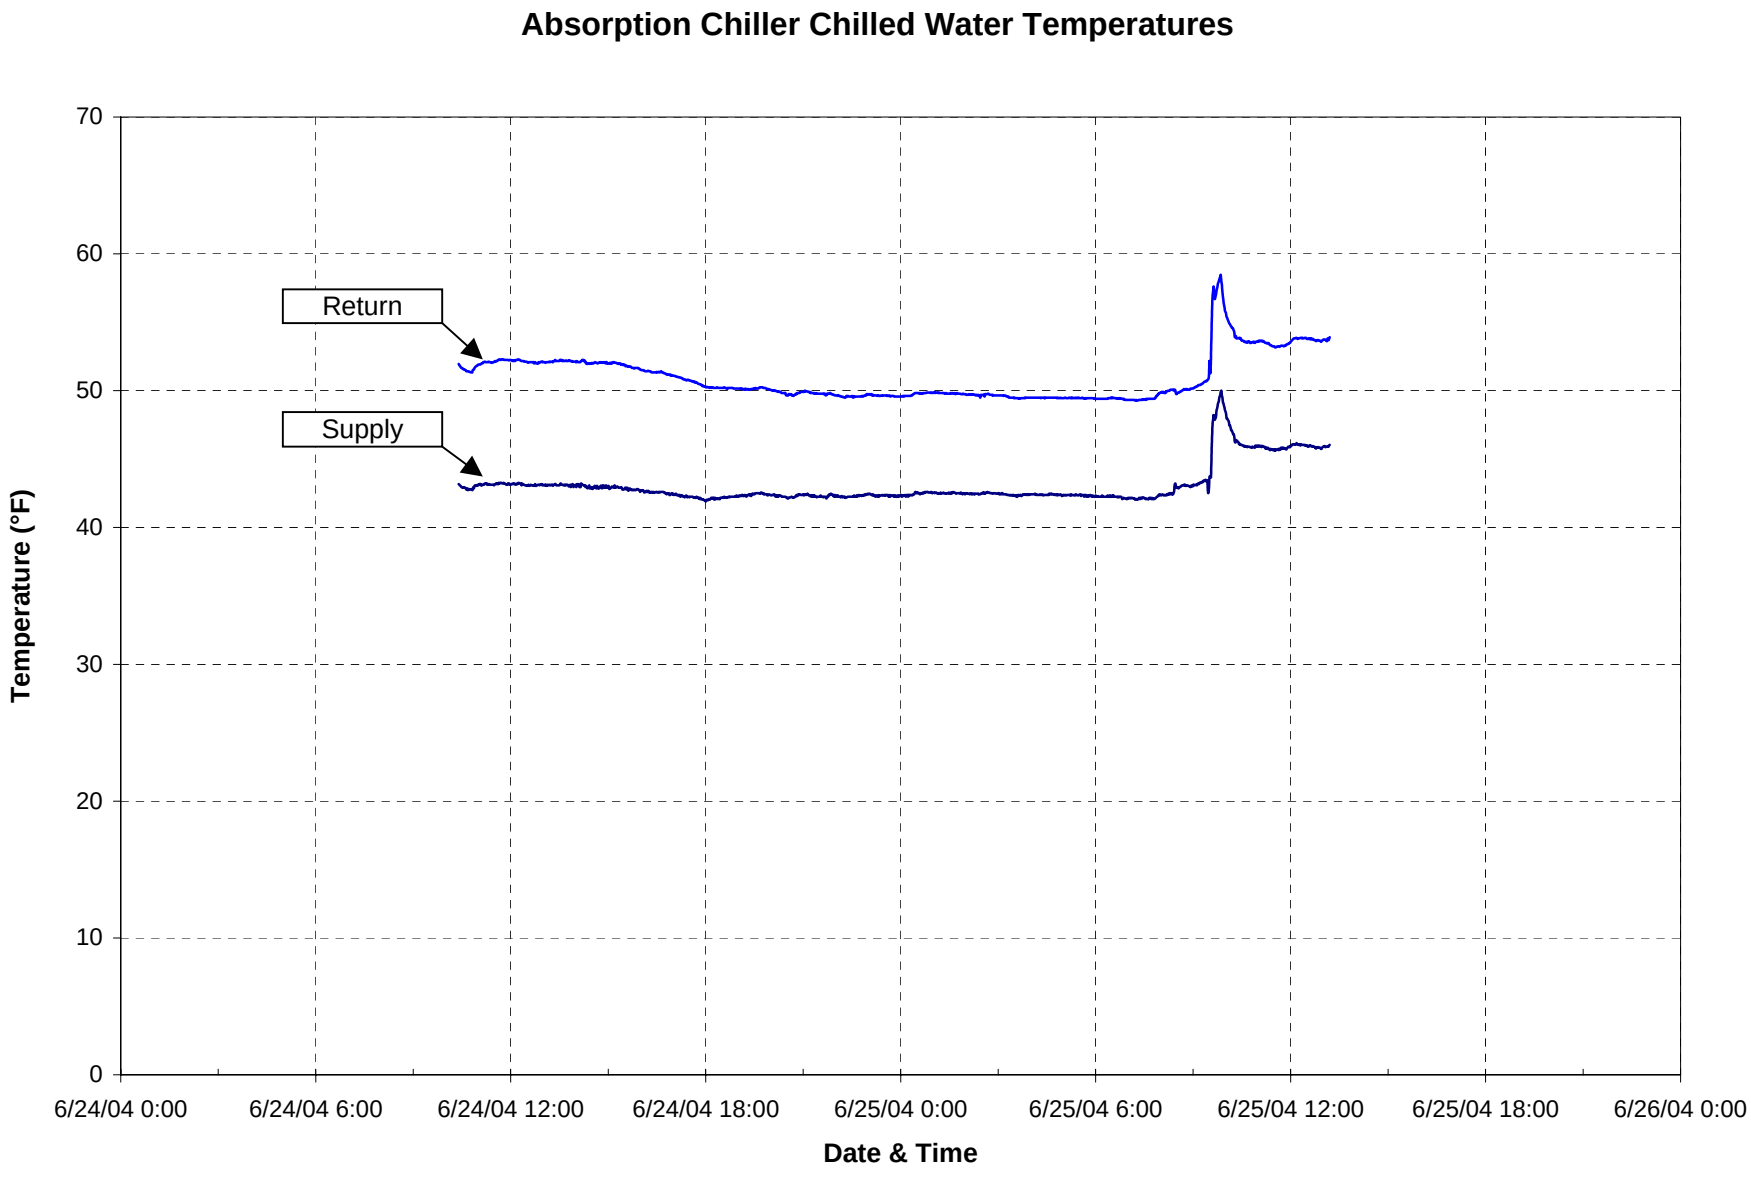
Absorption Chiller Chilled Water Temperatures
Return
Supply
Temperature (°F)
70
60
50
40
30
20
10
0
6/24/04 0:00
6/24/04 6:00
6/24/04 12:00
6/24/04 18:00
6/25/04 0:00
6/25/04 6:00
6/25/04 12:00
6/25/04 18:00
6/26/04 0:00
Date & Time

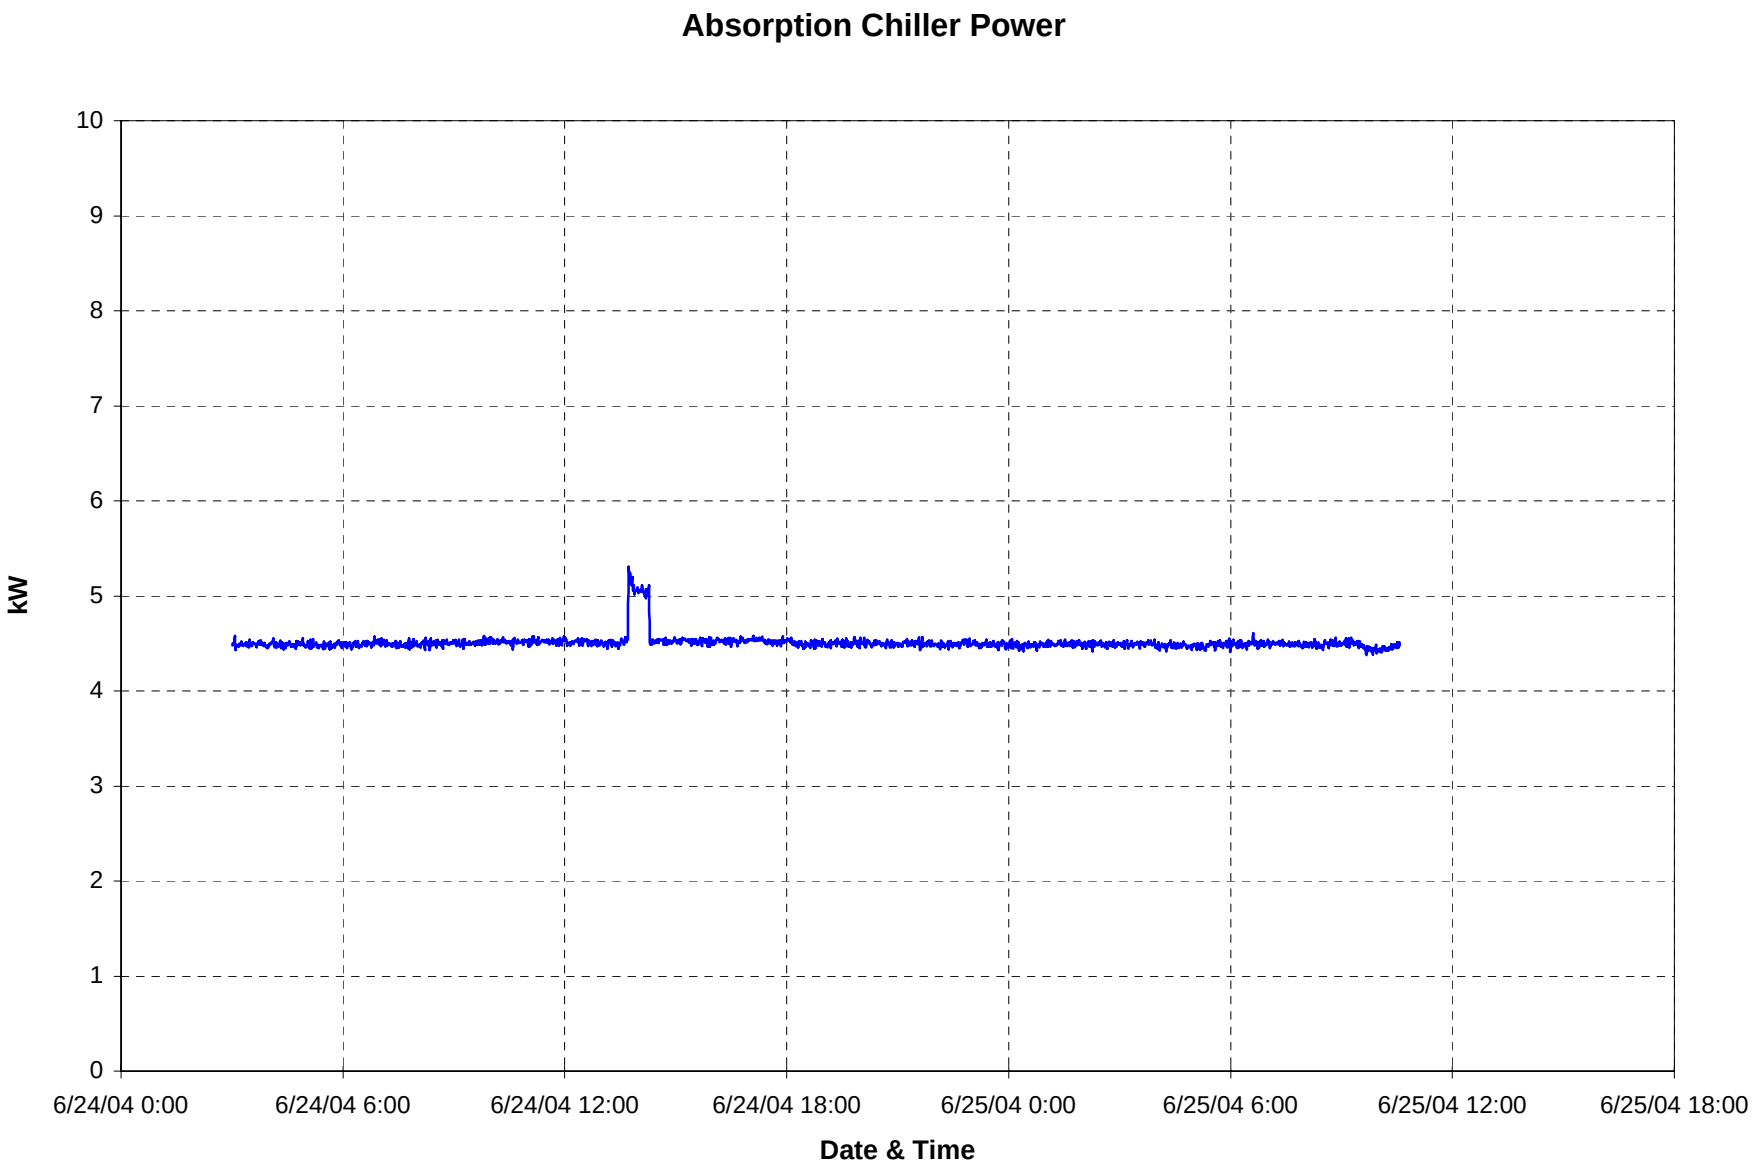
Absorption Chiller Power
kW
10
9
8
7
6
5
4
3
2
1
0
6/24/04 0:00
6/24/04 6:00
6/24/04 12:00
6/24/04 18:00
6/25/04 0:00
6/25/04 6:00
6/25/04 12:00
6/25/04 18:00
Date & Time

Appendix B

Data Collection and Accuracy Notes

Data Collection and Accuracy Notes

The following notes describe specific measurements and what assumptions were made in preparing calculated results.

Laboratory Data

Lighting Power	Lighting loads measured directly at their respective breakers for a representative laboratory area. The lighting power was then extrapolated based on the Watts/sf of the representative laboratory.
Process Power	Process tool loads measured directly at their respective breakers for a representative laboratory area. The process power was then extrapolated based on the Watts/sf of the representative laboratory.

Air Handler

Air Flow	All air flow measurements were provided by certification report.
Fan Power	Fan power was measured directly at the breaker.

Appendix C

Measurement Methodology

Measurement Methodology

Data collection measurements were made according to the following procedures:

Water Temperature – Pete's Plug

- Equipment: Pace Scientific 4-channel pocket loggers model XR440, and 4" 30kOhm thermistors.
- Attach thermistor temperature sensor to pocket logger channel block.
- Setup pocket logger using product software.
- Verify channels set to correct sensor type and operation.
- Upload setup to pocket logger to launch logging.
- Insert thermistor into Pete's plug.
- Secure pocket logger to pipe.
- Read real time data. Verify setup and actual conditions.
- Log data.

Air Temperature/Relative Humidity – Air Handler

- Equipment: Pace Scientific 4-channel pocket loggers model XR440, 4" 30kOhm thermistors, or Temperature/Relative Humidity sensors.
- Attach temperature/relative humidity sensor to pocket logger channel block.
- Setup pocket logger using product software.
- Verify channels set to correct sensor type and operation.
- Set RH linear scale specified by the sensor.
- Upload setup into pocket logger to launch logging.
- Read real time data. Verify setup and actual conditions.
- Log data.

Power Trend – Elite Logger

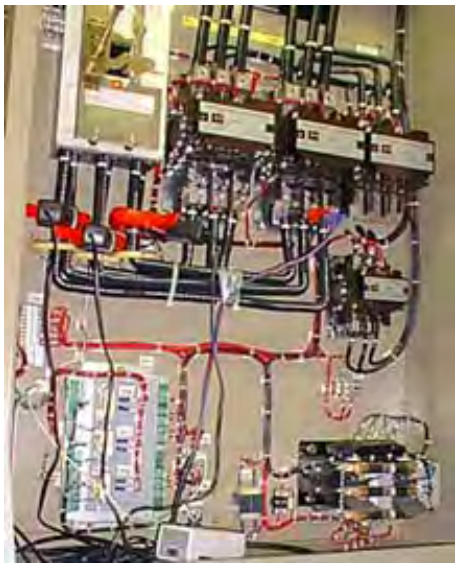

- Equipment: Elite Logger and ELOG 97c software.
- Select current transducers (CTs) appropriate for the measurement and panel space constraints.
- Attach current transducers (CTs) to Elite logger channel block.
- Plug Elite logger into AC supply.
- Setup Elite logger using product software.
- Electrician installation of voltage sensors in the electrical panel.
- Electrician installation of CTs in electrical panel for the specified load to be measured.
- Read real time data.
- Verify balanced current as well as appropriate, balanced voltage readings.
- Secure panel door and attach caution tape and warning notice if panel cannot be locked shut.
- Log data.

Power Spot Measurement – Power Sight

- Equipment: Power Sight PS 3000
- Plug Power Sight into AC supply if necessary.
- Connect current transducers (CTs) and voltage sensors to Power Sight.
- Electrician installation of voltage sensors in the electrical panel.
- Electrician installation of CTs in electrical panel for the specified load to be measured.
- Verify balanced current as well as appropriate, balanced voltage readings.
- Read and record the real time power reading for spot measurement.
- Log data for selected measurements.

CLEANROOM ENERGY BENCHMARKING AND EFFICIENCY TECHNOLOGY TRANSFER ASSISTANCE SITE REPORT

FACILITY L

SOUTHERN CALIFORNIA

JUNE 2005

SPONSORED BY:

LAWRENCE BERKELEY NATIONAL LABORATORY

PREPARED BY:
RUMSEY ENGINEERS, INC
99 LINDEN STREET
OAKLAND, CA 94607
WWW.RUMSEYENGINEERS.COM

Copyright (c) 2004 Lawrence Berkeley National Laboratory. All rights reserved.

Acknowledgements

Special thanks to the entire team for their generous assistance and cooperation throughout the benchmarking process.

Project Team

Lawrence Berkeley National Laboratory	**Rumsey Engineers, Inc**
Bill Tschudi	Peter Rumsey
Dale Sartor	Achai Broner
	Larry Chu

Reproduction or distribution of the whole, or any part of the contents of this document without written permission of LBNL is prohibited. The document was prepared by LBNL for the exclusive use of its employees and its contractors. Neither LBNL nor any of its employees makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any data, information, method, product or process disclosed in this document, or represents that its use will not infringe any privately-owned rights, including but not, limited to, patents, trademarks, or copyrights.

Table of Contents

I. Introduction 1

II. Review of Site Characteristics 1

III. Site Energy Use Characteristics 2

A. Cleanroom Power Consumption 2

B. Electrical System Power Consumption 4

C. Process Equipment Vacuum Pumps 4

IV. System Performance Metrics 5

V. Site Observations Regarding Energy Efficiency 9

APPENDICES

A. Trended Data Graphs

B. Data Collection and Accuracy Notes

C. Measurement Methodology

I. INTRODUCTION

As part of the California Energy Commission Public Interest Energy Research (PIER) project, energy use at Facility L was monitored December 6 to December 10, 2004. The project is administered by LBNL (Lawrence Berkeley National Laboratory) and funded through the California Energy Commission.

This site report reviews the data collected by the monitoring team and presents a set of performance metrics as well as a complete set of trended data points for energy end uses for air handling equipment supporting the cleanrooms. The chilled water plant, process power density, or lighting power density at this facility was not measured since all these metrics have been benchmarked as part of this study at prior sites. Energy metrics were established that allow cleanroom owners to evaluate their energy efficiency performance and identify opportunities for improvements that reduce their overall operating costs.

This Site Report summarizes the data collected and presents energy performance metrics with which the facility can evaluate the performance of its cleanrooms. First, the report reviews the site characteristics, noting design features of the mechanical plant and the cleanrooms monitored. Second, the energy use for the cleanrooms and mechanical equipment is broken down into components. Third, performance metrics recorded through the project are presented. Finally, key energy efficiency observations for the facility are noted. This is not intended to be a full energy audit, merely observations from the site measurement team. The data collected, trended graphs and methodology documentation are included among the appendices.

II. REVIEW OF SITE CHARACTERISTICS

A. Site

Facility L is a two-story semi-conductor manufacturing plant located in Southern California. The building was built in the 1960's and has gone through many different renovations since then. The first floor of the facility is primarily office and support spaces such as stock rooms and restrooms, while the second floor houses all of the cleanroom areas.

There are two 1375 kVA/1100 kW diesel generators onsite for emergency power in case of a utility failure. In addition to the generators, there are multiple uninterruptible power supplies (UPS) of various sizes supplying backup power for the cleanrooms.

The environmental systems serving the cleanrooms run 8,760 hours a year in order to maintain conditions and provide a safe cleanroom environment. The cleanrooms are conditioned by makeup air handlers (MUAH) working in conjunction with recirculation air handlers (RCU) served by a central plant. Due to the many renovations at Facility L, many different cleanroom systems are present. The different types of cleanroom systems include fan filtered HEPA units, pressurized plenums, and ducted HEPAs.

Three different areas with varying systems have been chosen, in order to compare the various air handling system efficiencies. All cleanroom areas are classified as ISO class 4 (class 10). The first area, Fab A Diffusion, has one MUAH with fan filtered HEPA units (FFUs). The second area, Fab B Diffusion, is conditioned by one fan-wall MUAH serving two fan-wall RCU's with ducted HEPA filters. The third area, Fab C Photo, has one MUAH serving two groups of RCU's, one with a pressurized plenum and another with ducted HEPA filters. All these areas operate 24 hours a day, 365 days a year and house an array of different equipment.

Fab A, which is 1,905 sf, is served by one makeup air handler, AH-A1, which supplies fan filtered HEPA units. The makeup air handler in this area provides ventilation air as well as recirculation air. The recirculation air supplies sensible cooling to the cleanroom.

Fab B, which is 1,650 sf, was recently equipped with three new fan wall air handling units, MUAH AH-B1, and RCU's AH-B2 and B3. An air handler with a fan wall is like a typical air handler, except that it employs multiple smaller fans working in conjunction, as opposed to one large fan. It is believed that these units have superior performance, as compared to a standard configuration air handler, because of lower system pressure drop and superior cooling coil performance.

Fab C, which is 4,065 sf, is served by one makeup air handler, AH-C1, and two groups of recirculation units. The first group of RCU's, AH-C2, C3, C4 and C5, serves ducted HEPA filters, which cover the other half of the cleanroom. The second group of RCU's, AH-C6, C7 and C8, serves a pressurized plenum, which approximately covers half of the cleanroom.

In addition to measuring the air systems for these zones, measurements were taken on various mini-environments and the vacuum pump power for one of the cleanroom tools. Mini-environments are special enclosures, within the cleanroom, that provide increased filtration and air circulation over critical tools and areas. The vacuum pumps for the cleanroom tool were measured to observe the power draw of the pumps as the tool is being used.

III. SITE ENERGY USE CHARACTERISTICS

A. Cleanroom Power Consumption

The energy consumption attributed to the cleanroom air handling systems. This information will help for comparisons of different cleanroom air handling system efficiencies.

Table 1. Fab A Diffusion – Fan Filtered HEPA Cleanroom

Description	Average Load (kW)	Efficiency (cfm/kW)
MAKEUP FAN [1]		
AH-A1	13.0	564
RECIRCULATION FANS [2]		
Fan Filter HEPA Units	38.5	1,127

1. *Makeup air handler provides both ventilation and recirculation air, but is considered as a makeup air handler only for this study.*

Table 2. Fab B Diffusion – Fan Wall, Ducted HEPA Cleanroom

Description	Average Load (kW)	Efficiency (cfm/kW)
MAKEUP FAN		
AH-B1	6.6	1,307
RECIRCULATION FANS		
AH-B2	26.9	1,180
AH-B3	17.1	1,798
Overall	44.0	1,421

Table 3. Fab C Photo – Ducted HEPA/Pressurized Plenum Cleanroom

Description	Average Load (kW)	Efficiency (cfm/kW)
MAKEUP FAN		
AH-C1	11.3	1,324
RECIRCULATION FANS		
Recirculation Fans w/ Ducted HEPA Filters		
AH-C2	6.4	4,672
AH-C3	6.2	3,305
AH-C4	7.7	3,371
AH-C5	4.1	4,673
Overall	24.4	3,915
Recirculation Fans w/ Pressurized Plenum		
AH-C6	4.8	4,715
AH-C7	4.8	4,986
AH-C8	3.5	4,929
Overall	13.1	4,871

Table 4. Summary of Measured Cleanroom Air Handling Parameters

Description		***Fab A Fan Filtered HEPA Cleanroom***	***Fab B Ducted HEPA Cleanroom***	***Fab C Ducted HEPA/ Pressurized Plenum Cleanroom***
		ISO Class 4	*ISO Class 4*	*ISO Class 4*
Primary Area	sf	880	890	3,680
Secondary Area	sf	1,025	760	385
Ceiling Height	ft	10	10	10
Total Makeup Air [1]	cfm	7,340	8,630	14,960
Total Makeup Fan Power	kW	13.0	6.6	11.3
Total Recirculation Air [1]	cfm	43,452	62,485	159,369
Total Recirculation Fan Power	kW	38.5	44.0	37.5
HEPA Filter Ceiling Coverage	%	100	100	100
Average Ceiling Filter Velocity [2]	fpm	58	83	51

1. *Makeup and Recirculation Air is the air delivered, based on the balance report data.*
2. *Filter velocity based on 6.8 sf (85%) effective filter area.*

B. Electrical System Power Consumption

The tables below show the power consumption of the two 1,375 kVA / 1,100 kW emergency generator's standby power loss. The two emergency generators constantly draw power to maintain the batteries and a specific temperature of the diesel engines that drive the generator.

Table 5. Electrical System Power Consumption

Description	Average Load
ELECTRICAL SYSTEM	
Emergency Generator Standby Power (2 generators)	7.8 kW

C. Process Equipment Vacuum Pumps

The vacuum pump power for one of the cleanroom process equipment tools was measured over a 17-hour period, while the tool was in production. The table below shows the average power consumption data for each vacuum pump associated with the process tool. Following is a graph for one of the vacuum pump's power consumption, which its load shape is representative of all of the other pumps. Regardless of the process equipment tool sitting idle or operating, each vacuum pump drew a constant load.

Table 6. Cleanroom Process Equipment Tool Vacuum Pump Power Consumption

Description	Average Load (kW)
VACUUM PUMP	
Chamber A	2.5
Chamber B	2.6
Chamber C	3.4
Chamber D	3.4
Transfer	1.5
Load/Lock	0.96
TOTAL	14.4

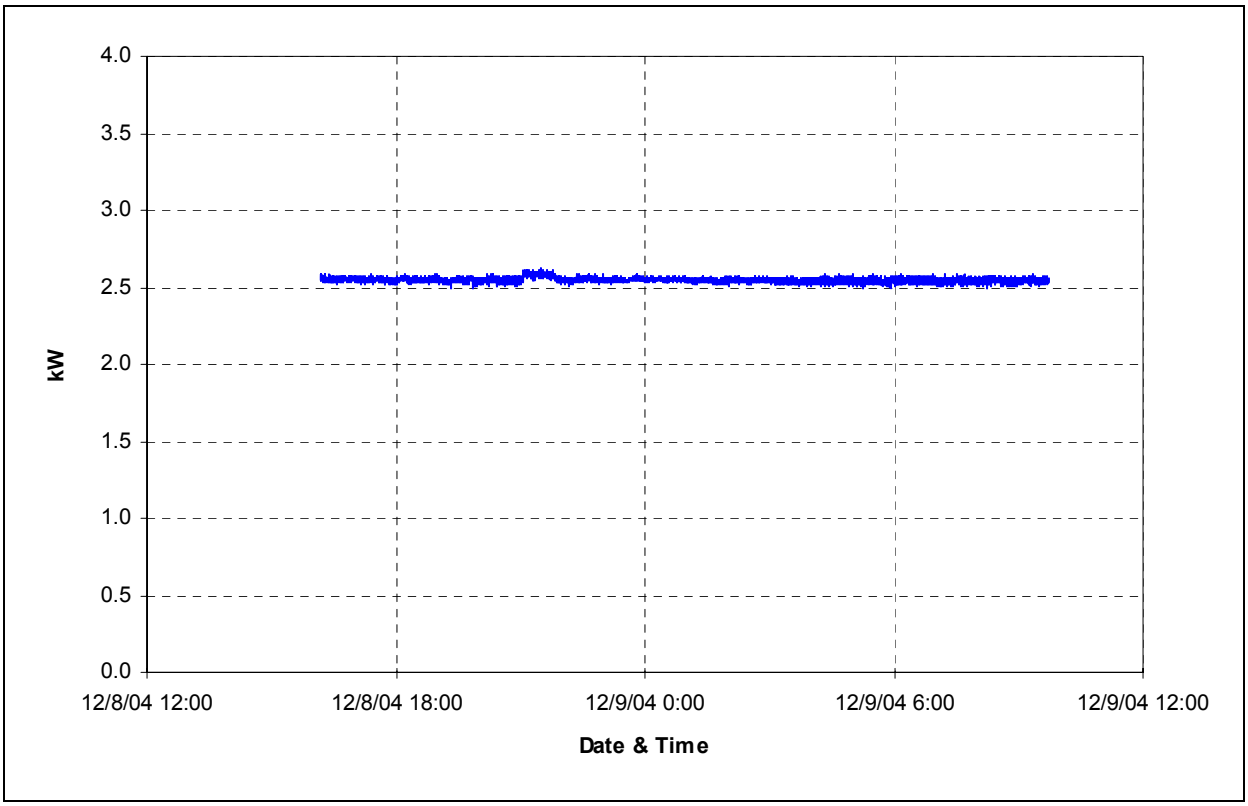

Figure 1. Cleanroom Process Tool Chamber A Vacuum Pump Power

IV. SYSTEM PERFORMANCE METRICS

Metrics are ratios of important performance parameters that can characterize the effectiveness of a system or component. In order to gage the efficiency of the entire building system design and operation, this project tracks key metrics at different system levels. These metrics can be used to compare designs or determine areas with the most potential for improvement via retrofit or replacement.

For Facility L, the cleanroom HVAC components operate at a nearly constant level throughout the year. Therefore, these metrics are based on spot measurements, and various logged data to ensure accuracy. All of the metrics involving area are based on the primary cleanroom area, which is the area that passes certification, unless otherwise noted.

Three different types of recirculation air handling systems were evaluated in this study. The recirculation air handling efficiency for the ducted HEPA system has been found to be the best performer among the other ducted HEPA cleanrooms in the entire study (see Figure 3). The recirculation air handling efficiency has an efficiency of 3,915 cfm/kW. Table 7 below shows the measured fan power density for the ducted HEPA area of Fab C to be 7 times lower than the designed fan power density. Apparently, the air handlers are grossly oversized and are operating at a much lower pressure drop than what they were designed for, hence the outstanding air handling efficiency.

Table 7. Cleanroom Metrics

Description		***Fab A Fan Filtered HEPA Cleanroom***	***Fab B Ducted HEPA Cleanroom***	***Fab C Ducted HEPA Cleanroom***	***Fab C Pressurized Plenum Cleanroom***
		ISO Class 4	*ISO Class 4*	*ISO Class 4*	*ISO Class 4*
MUAH Efficiency	cfm/kW	564	1,307	1,324	
Makeup Air	cfm/sf	8.3	9.7	4.1	
Makeup Fan Power Density [1]	W/sf	14.8	7.4	3.1	
Recirculation Air Handler Airflow	cfm	43,452	62,485	95,406	63,963
Designed Recirculation Air Handler Airflow	cfm	N/A	77,000	216,000	131,000
Recirculation Air Handler Efficiency	cfm/kW	1,127	1,421	3,915	4,871
Designed Recirculation Air Handler Efficiency	cfm/kW	N/A	1,139	1,224	1,079
Recirculation Air	cfm/sf	49	70	43	43
Recirculation Air Changes per Hour	ACH	296	421	260	259
Designed Recirculation Air Changes per Hour	ACH	N/A	577	589	591
Recirculation Fan Power Density	W/sf	20	49	11	9
Designed Recirculation Fan Power Density	W/sf	N/A	76	80	82

1. Calculated as total kW load divided by the primary area of the cleanroom.

N/A – Data not available.

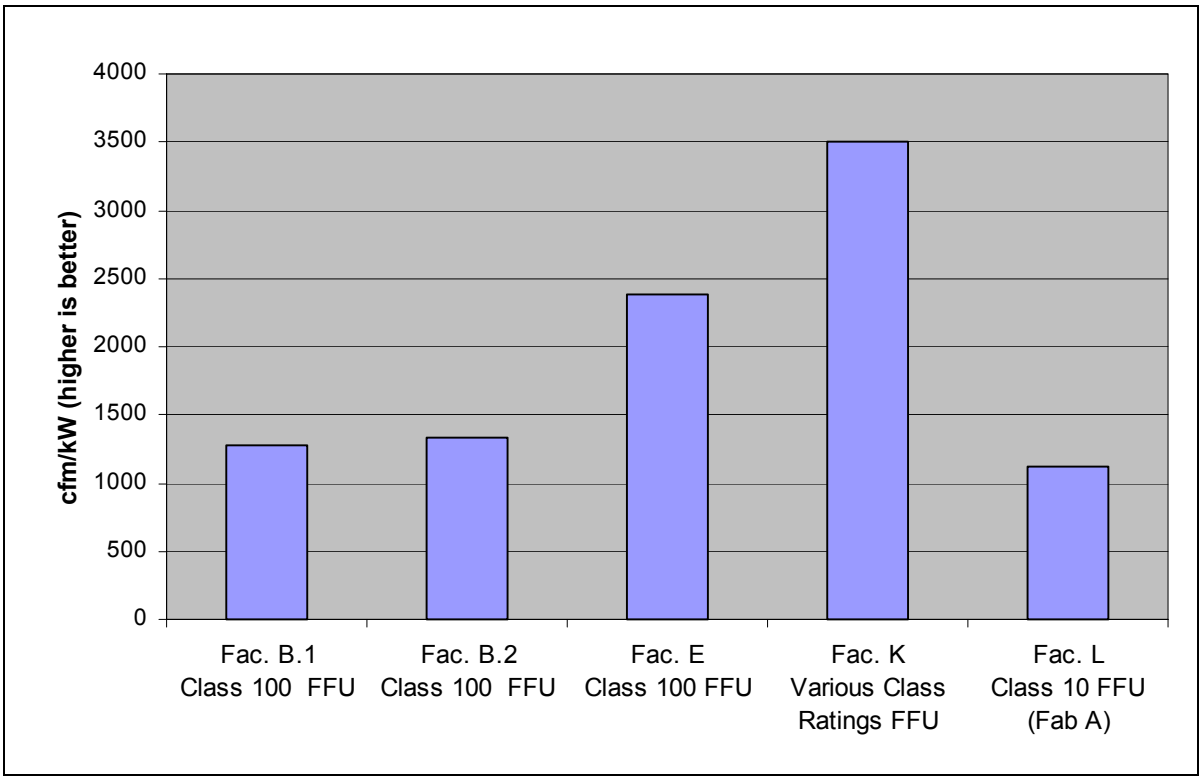

Figure 2. Fan Filter Unit Recirculation Air Handler Efficiencies of Measured Facilities (Various Class Ratings)

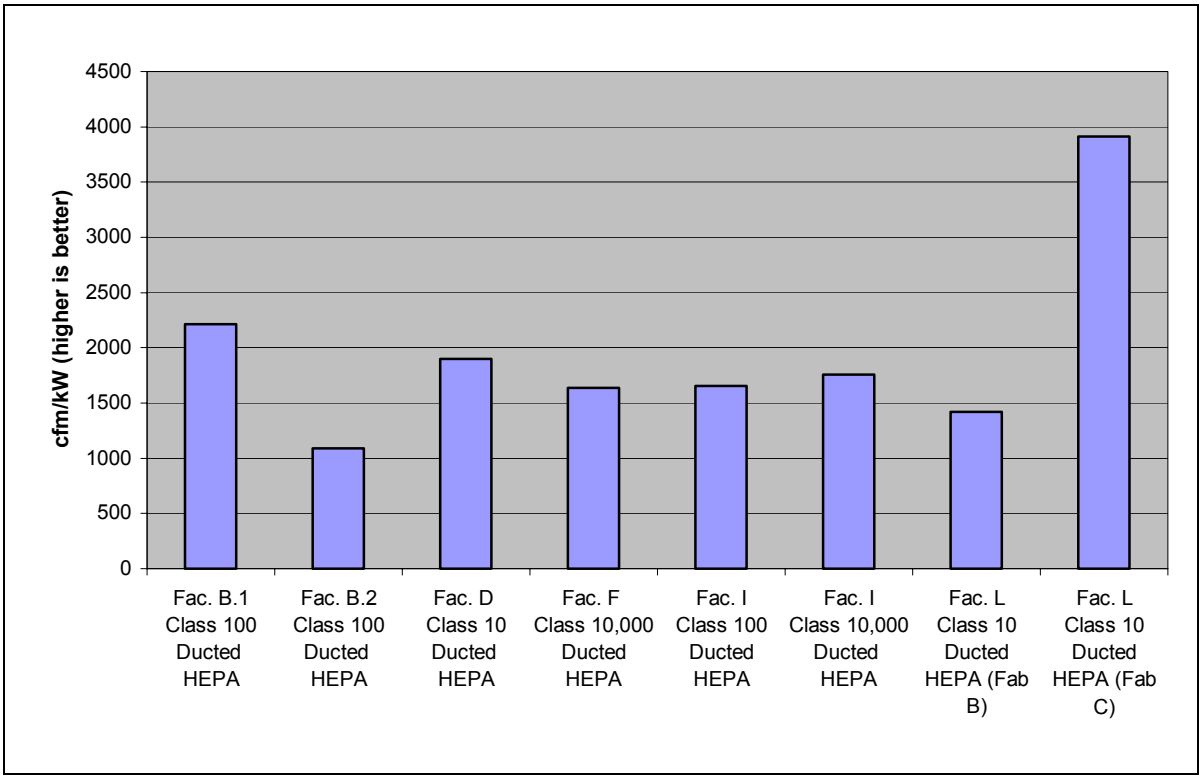

Figure 3. Ducted HEPA Recirculation Air Handler Efficiencies of Measured Facilities (Various Class Ratings)

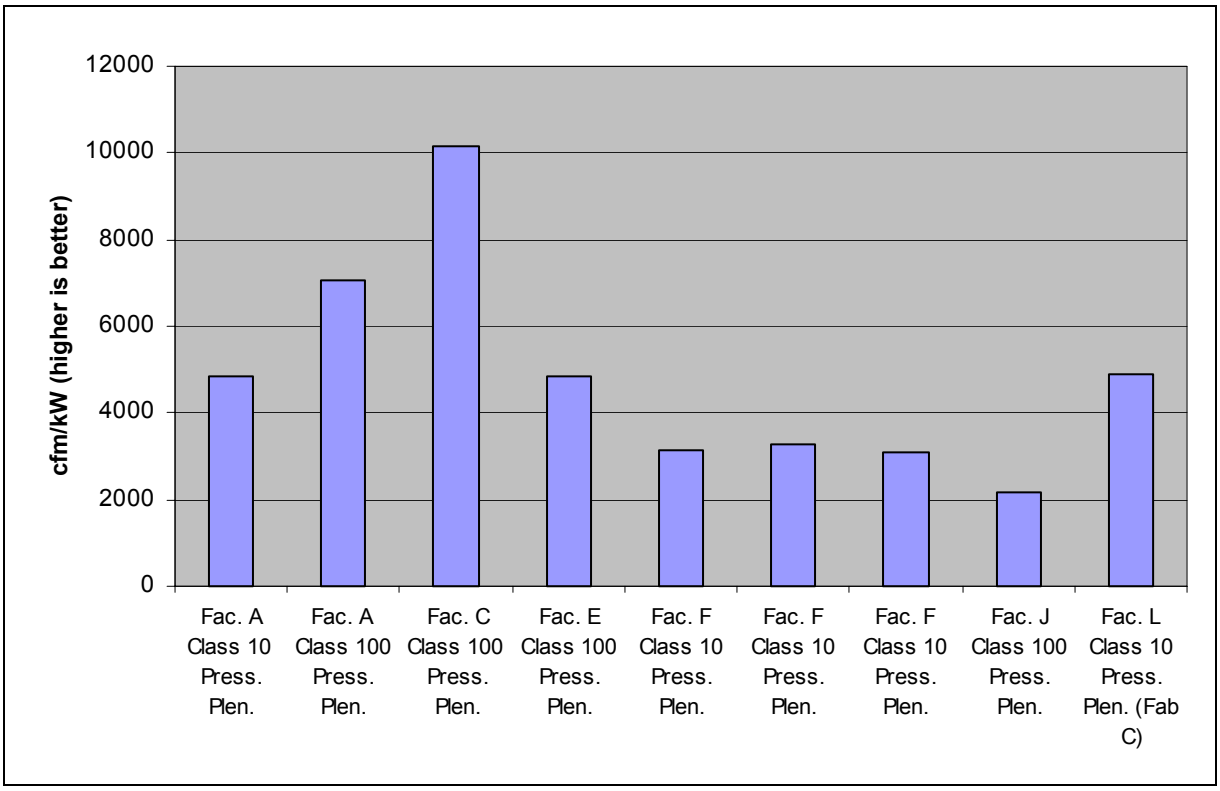

Figure 4. Pressurized Plenum Recirculation Air Handler Efficiencies of Measured Facilities (Various Class Ratings)

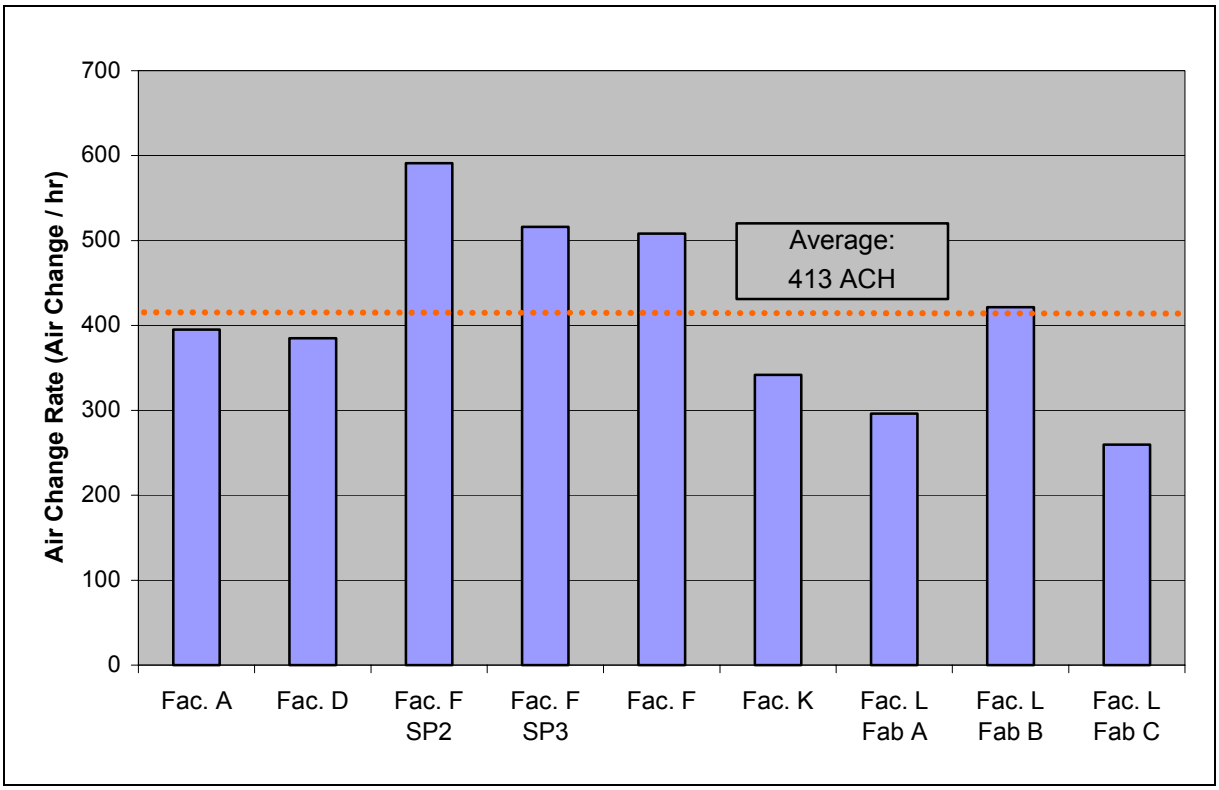

Figure 5. Air Change Rates of Measured ISO Class 4 (Class 10) Facilities

V. SITE OBSERVATIONS REGARDING ENERGY EFFICIENCY

Some notable energy efficiency features have already been implemented into the systems at this facility. Some of the key features are:

- Corrosive exhaust air scrubber retrofitted with lower pressure drop fill.
- Replacement of ceiling HEPA filters with lower pressure drop filters
- Filter velocity reduction in cleanrooms by 10-15%.

However, there are a few additional opportunities that can be implemented to improve the overall energy efficiency at this site.

Reduce Air Change Rates

The air change rate in Fab C is high within the set of ISO class 4 (class 10) facilities benchmarked. The benchmarking project has found that significantly lower air change rates are commonly used to provide the same class cleanroom environment, and with the good gowning protocol it is expected that a lower rate could be used without problem. A particle count should be conducted in each cleanroom to verify if a reduction in air change rates is feasible. Either the air handler fan can be set back to a constant speed, or implement demand-controlled filtration as a more aggressive approach.

Demand-controlled filtration is currently being investigated at the facility. By actively monitoring the particle counts in the cleanrooms and adjusting the airflow provided by the recirculation air handling system to retain cleanliness, a significant amount of energy can be saved.

Vacuum Pump Optimization

The vacuum pumps serving the process equipment tool were found to be drawing a constant load weather the tool was idle or in production. Current research is being done in this area. Investigating variable speed vacuum pumps, which reduce its load when the process tool is idle is an example of such. Opportunities to reduce the power consumption of vacuum pumps should be considered as research permits.

Replace Portable Chillers with Chiller Plant Water to Cool Vacuum Pumps

Many of the vacuum pumps each had their own portable chiller, which are inherently inefficient. The feasibility of utilizing the chilled water plant to provide cooling to the vacuum pumps serving the process equipment should be investigated. A heat exchanger would be required to supply water at a warmer temperature as required by the vacuum pumps. Supply water to one of the vacuum pumps was measured to be 65°F. A cooling tower could also provide this temperature water on cooler days.

Emergency Generator Standby Power Loss Reduction

The 1,100 kW emergency generators constantly draw power to maintain the batteries and a specific temperature of the diesel engines that drive the generator. An average of 3.9 kW per generator is being consumed to prime each generator for a power outage. Claims have been made that the engine(s) of a generator can be maintained at lower temperatures without introducing any physical damage or problems when started. The manufacturer of the generator should be contacted to explore energy reduction options.

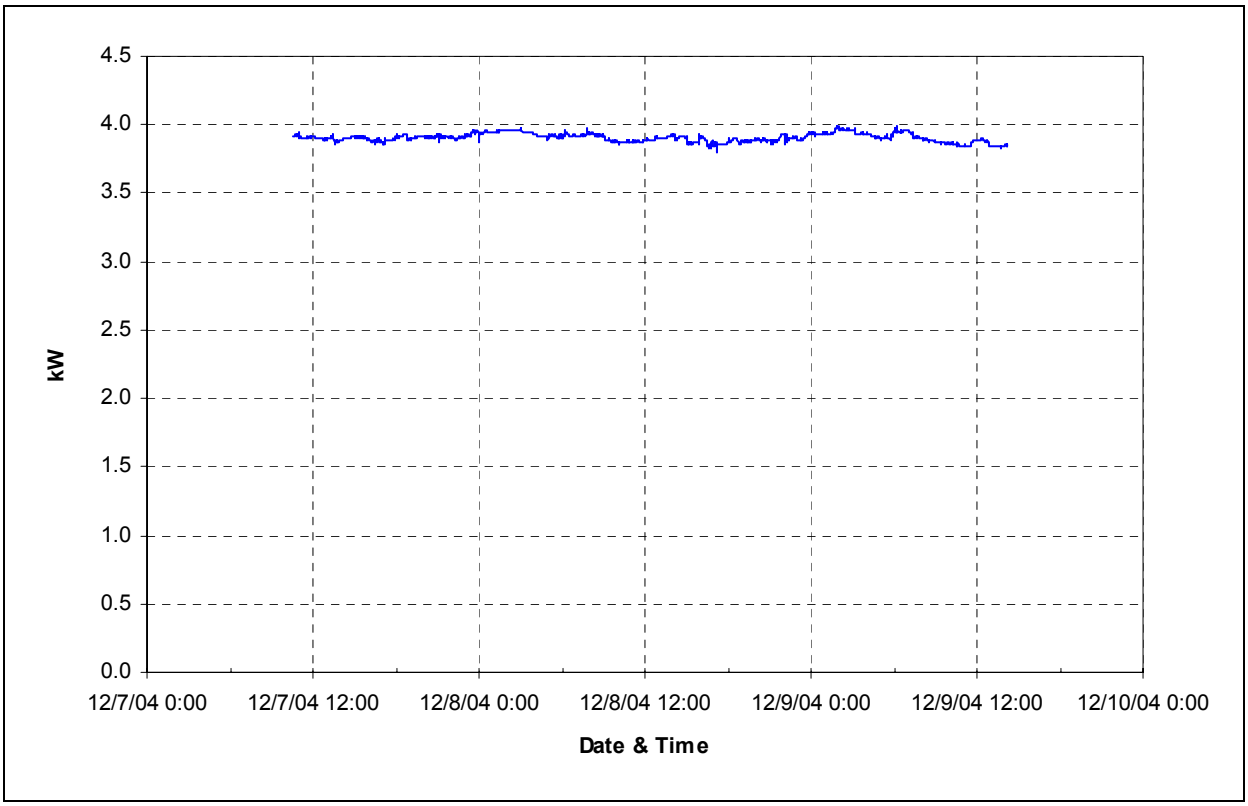

Figure 6. Generator 1 Standby Power Consumption

Condenser Water Temperature Reset

The centrifugal chillers may benefit from a reduction in lift. A reduction in the condenser water temperature of a couple degrees will achieve meaningful energy savings, however chiller stability can be compromised if it is taken too low, hence the need to verify with the chiller manufacturer before modifying this point.

Appendix A

Trended Data Graphs

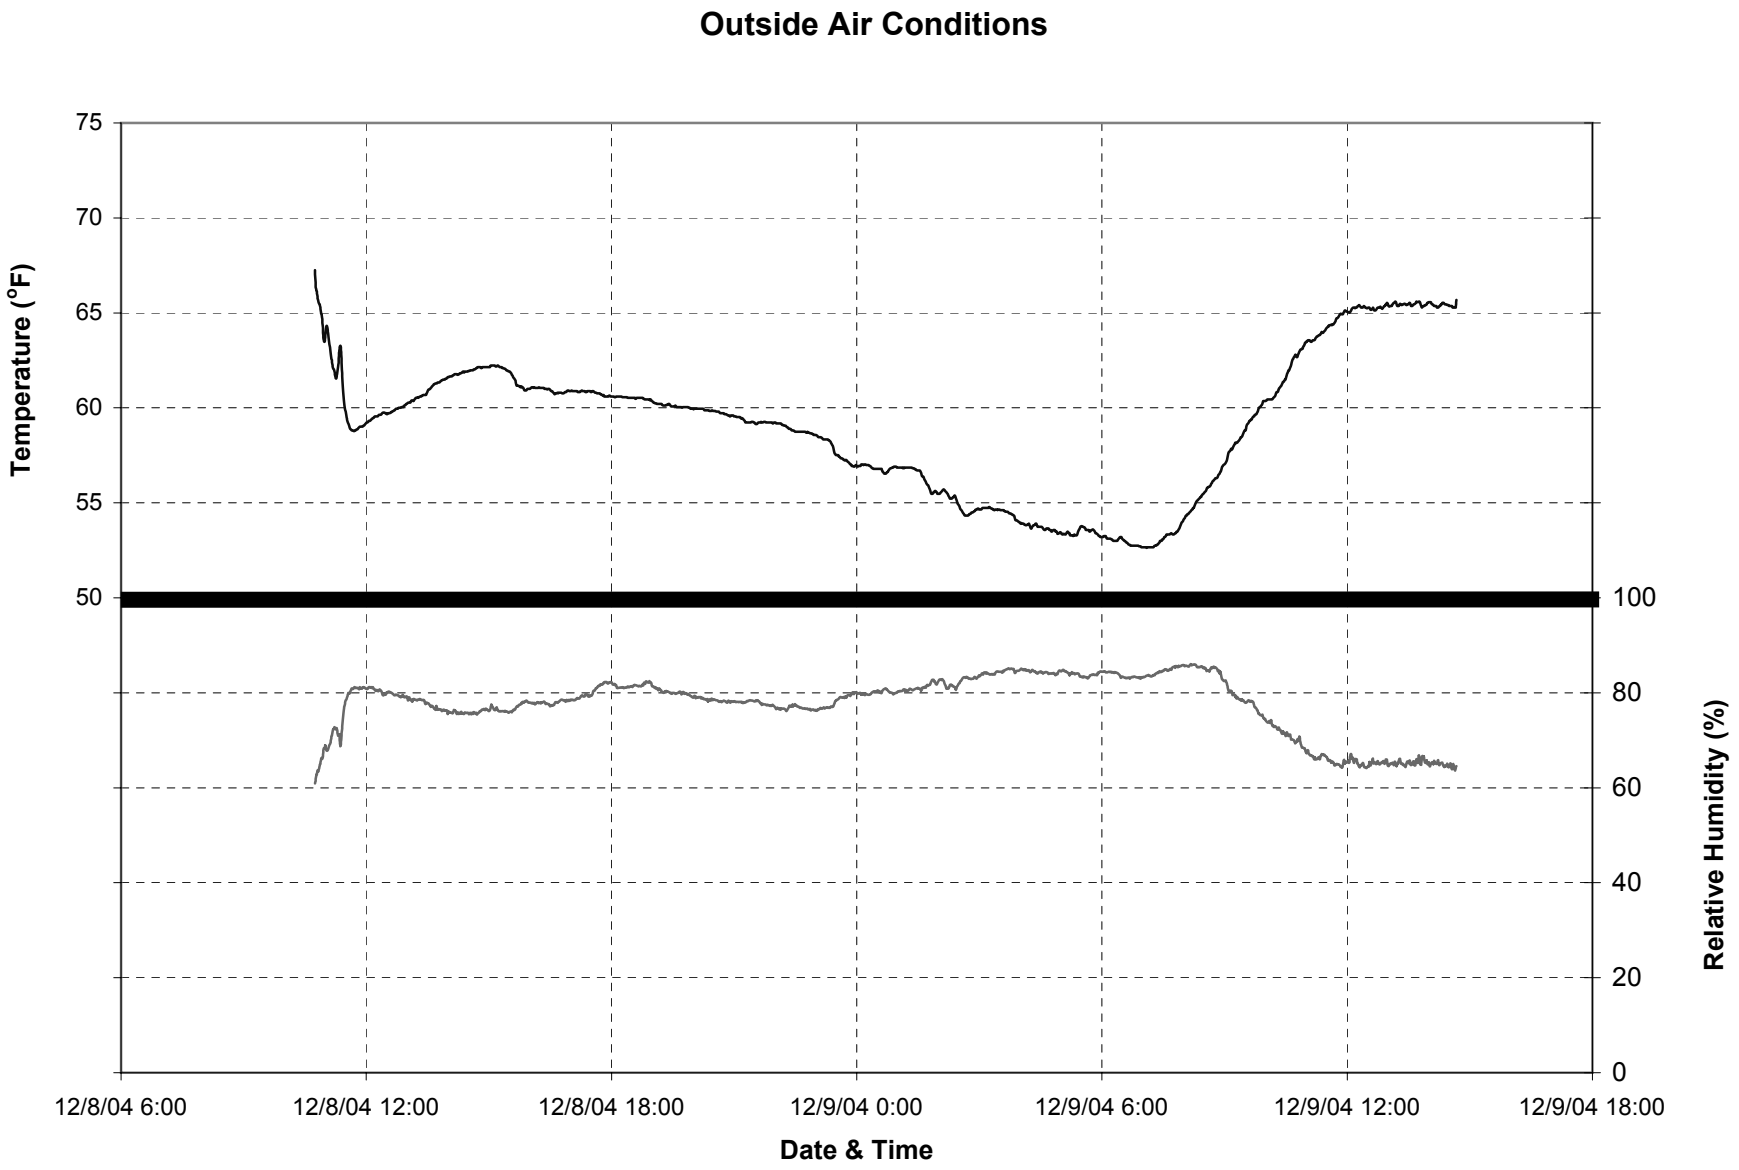
Outside Air Conditions
Temperature (°F)
75
70
65
60
55
50
100
80
60
40
20
0
Relative Humidity (%)
12/8/04 6:00
12/8/04 12:00
12/8/04 18:00
12/9/04 0:00
12/9/04 6:00
12/9/04 12:00
12/9/04 18:00
Date & Time

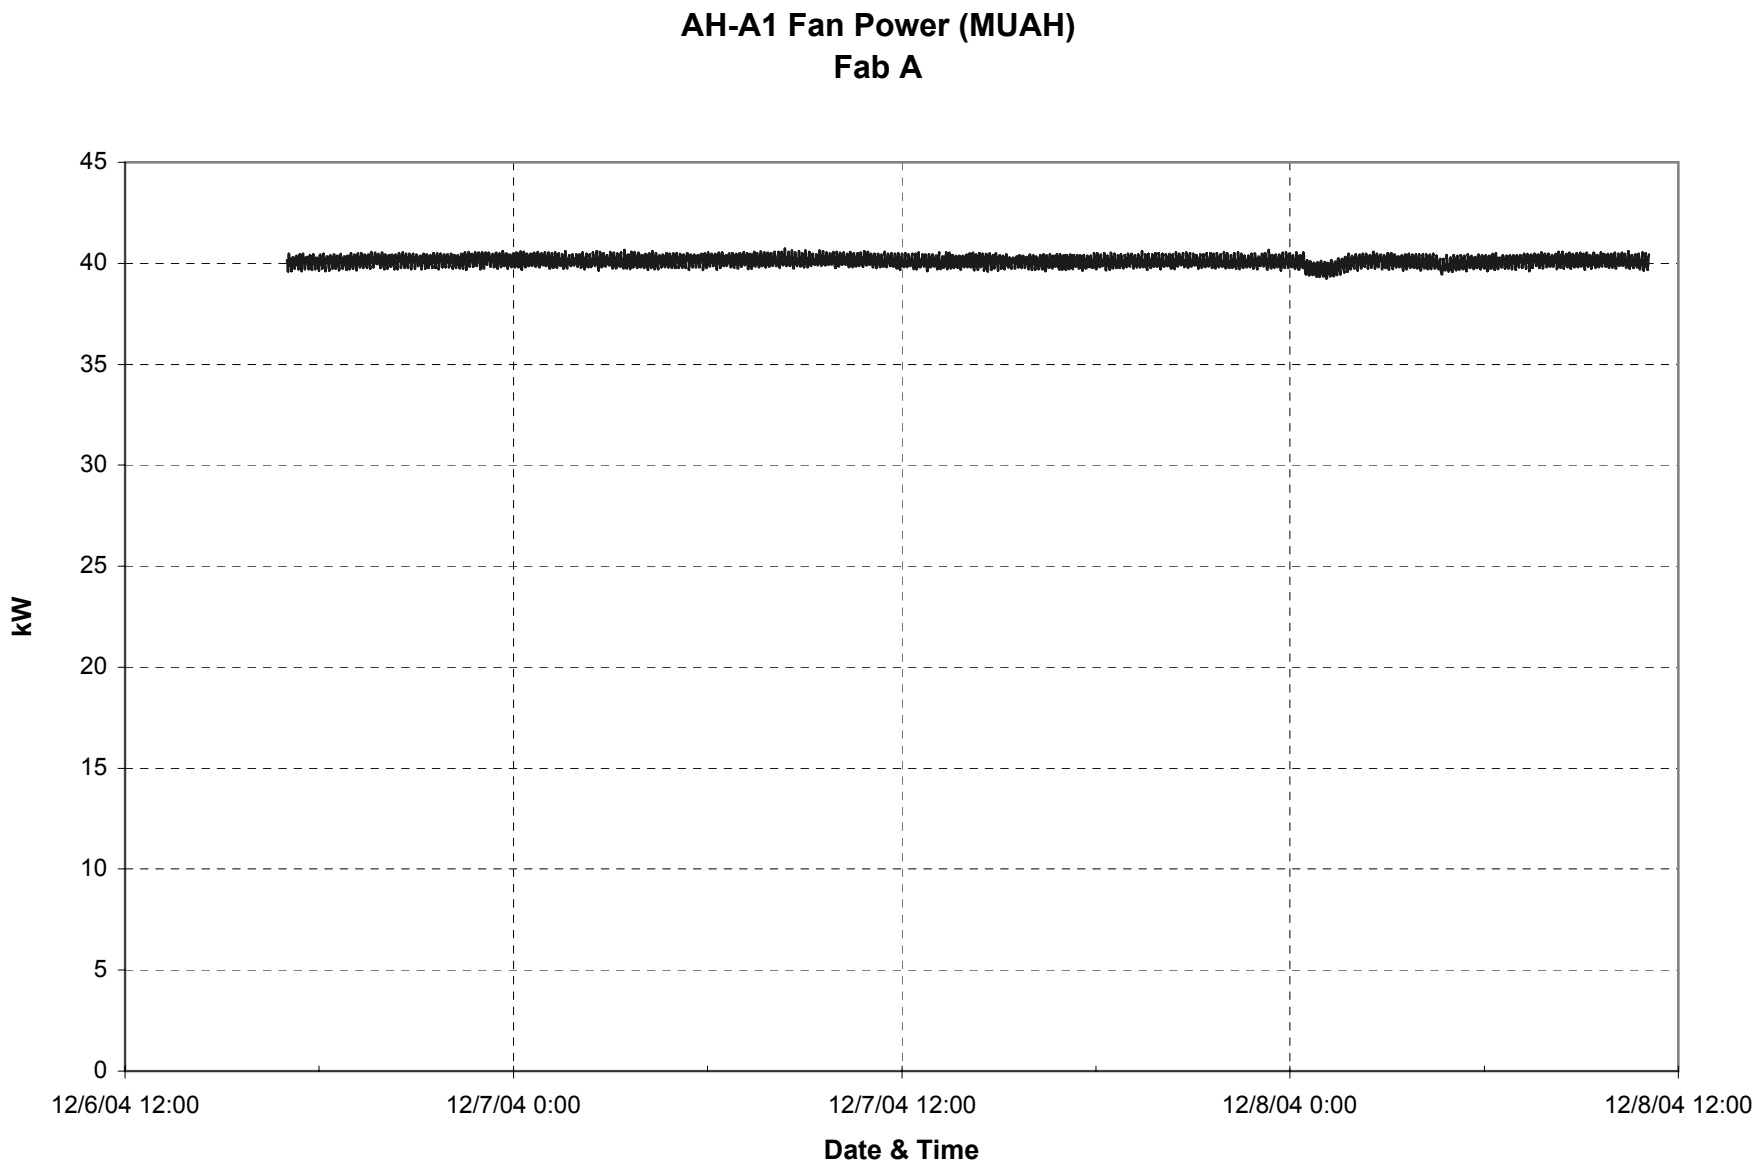
AH-A1 Fan Power (MUAH)
Fab A
45
40
35
30
25
20
15
10
5
0
kW
12/6/04 12:00
12/7/04 0:00
12/7/04 12:00
12/8/04 0:00
12/8/04 12:00
Date & Time

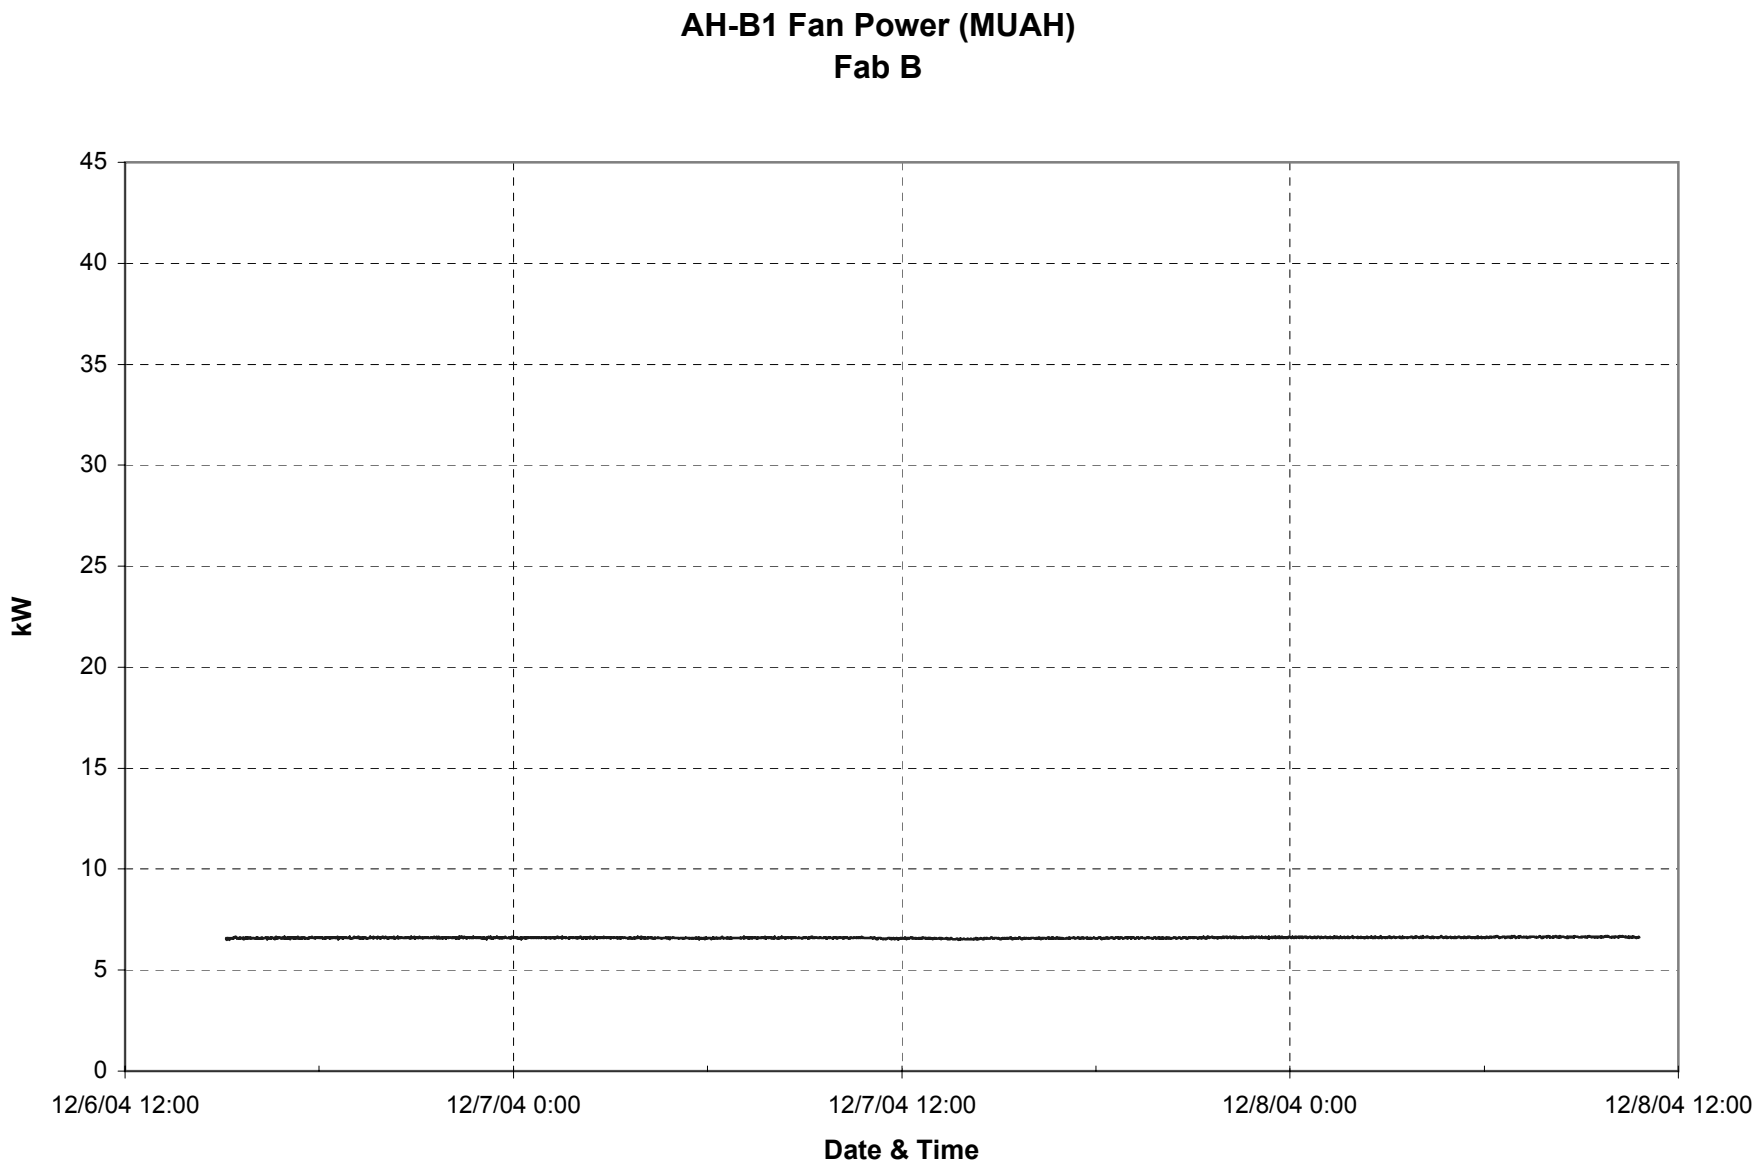
AH-B1 Fan Power (MUAH)
Fab B
45
40
35
30
25
20
15
10
5
0
kW
12/6/04 12:00
12/7/04 0:00
12/7/04 12:00
12/8/04 0:00
12/8/04 12:00
Date & Time

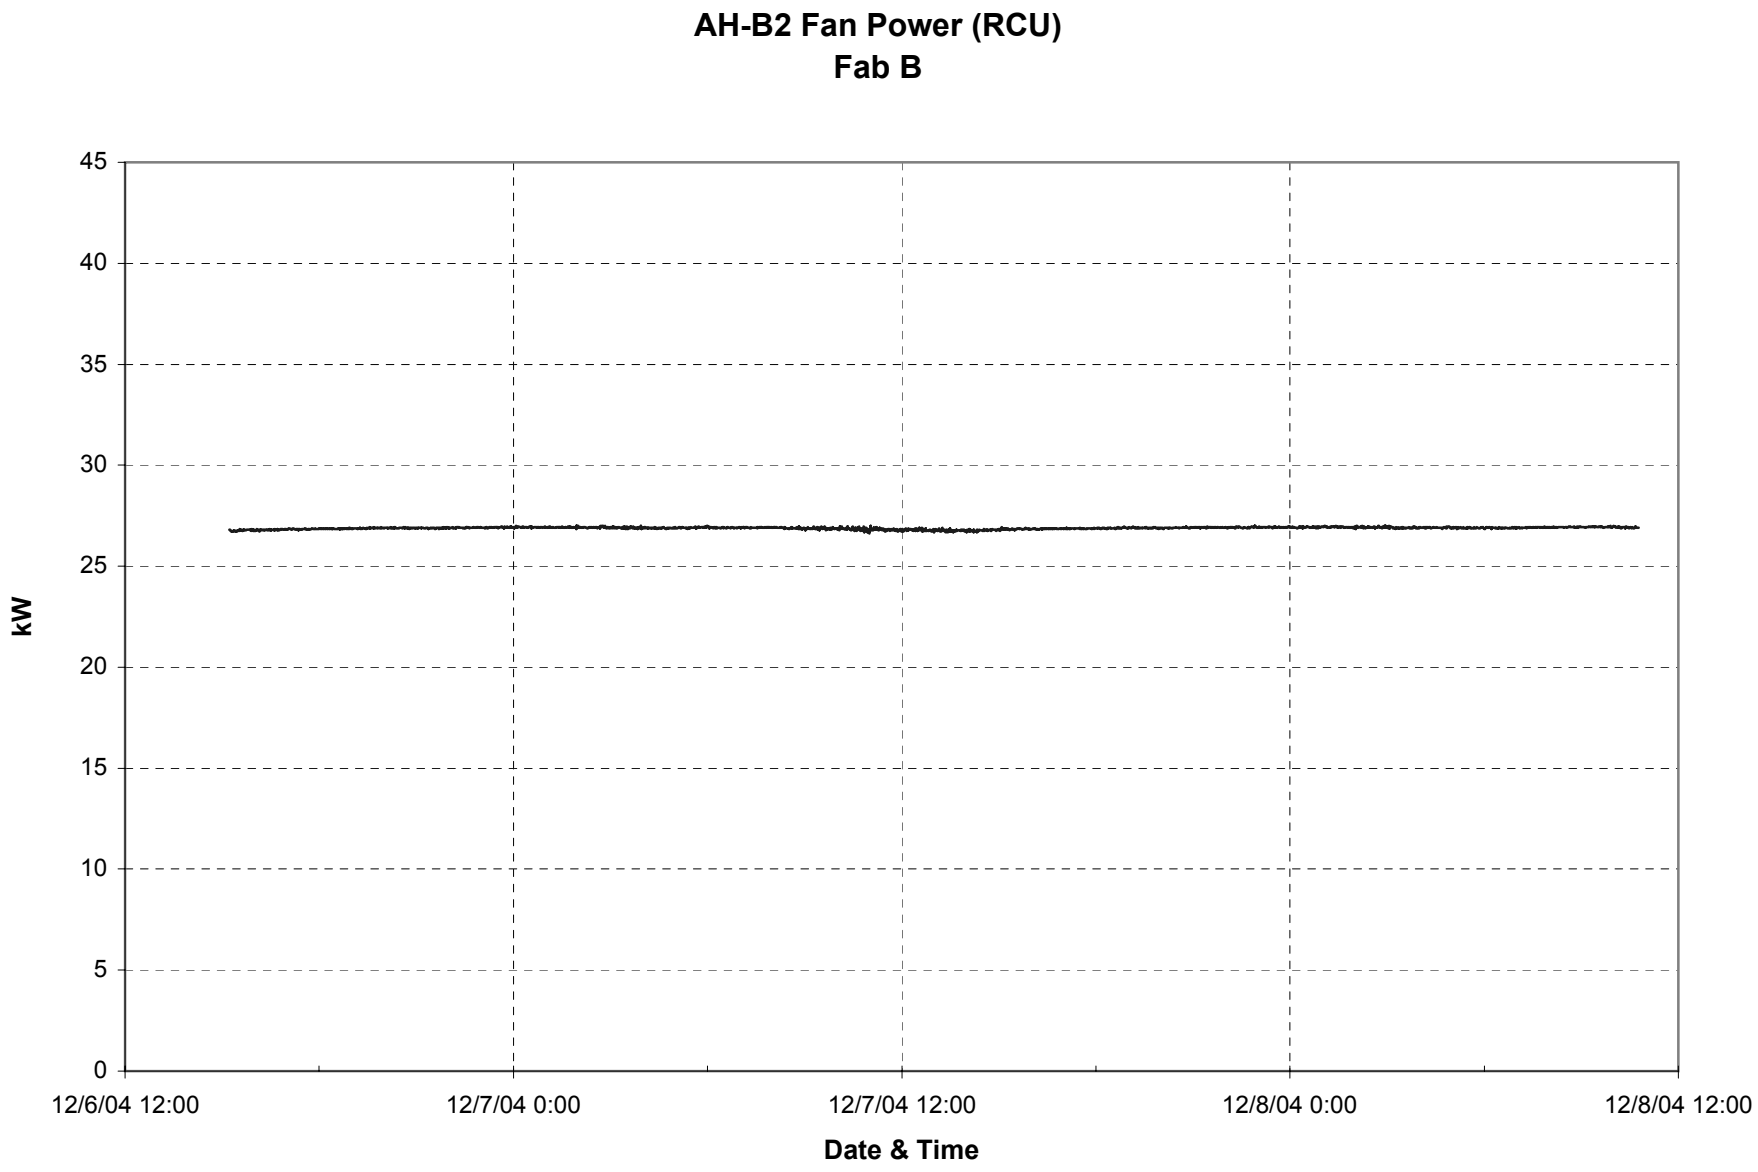
AH-B2 Fan Power (RCU)
Fab B
45
40
35
30
25
20
15
10
5
0
kW
12/6/04 12:00
12/7/04 0:00
12/7/04 12:00
12/8/04 0:00
12/8/04 12:00
Date & Time

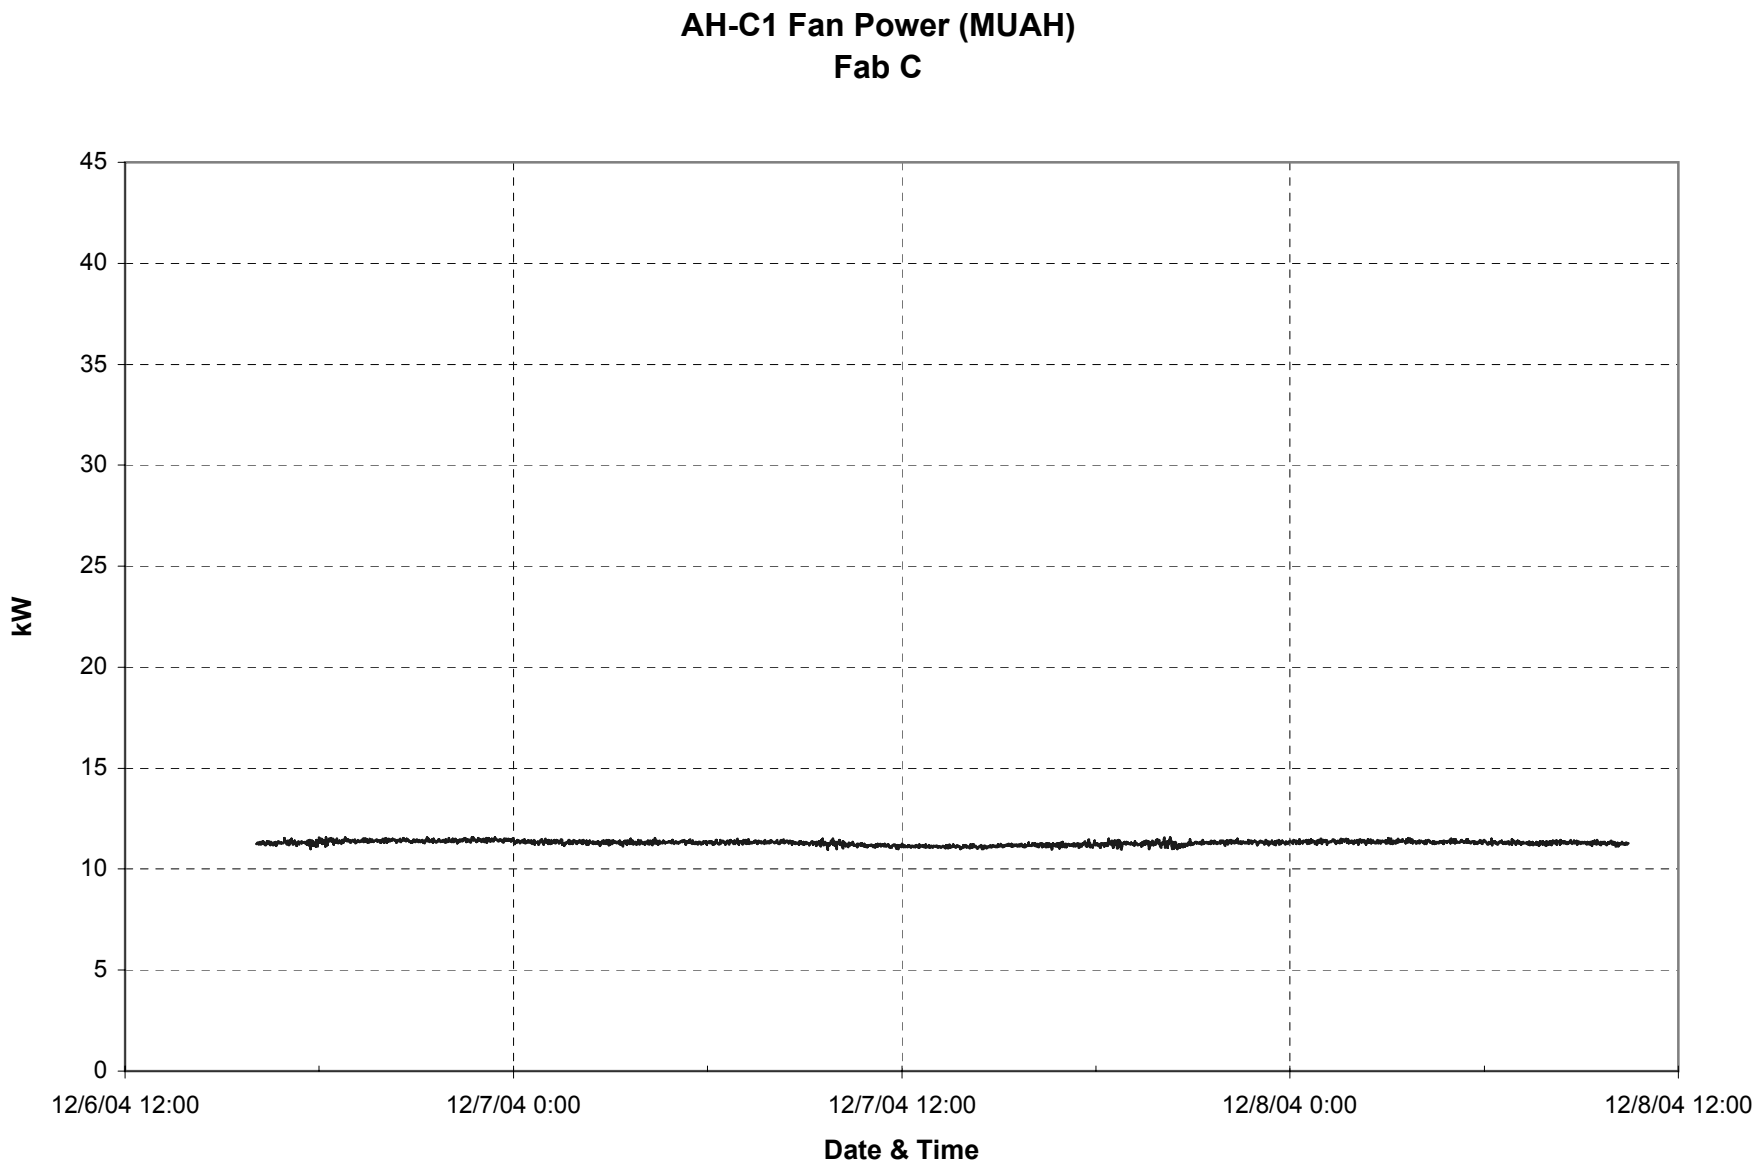
AH-C1 Fan Power (MUAH)
Fab C
45
40
35
30
25
20
15
10
5
0
kW
12/6/04 12:00
12/7/04 0:00
12/7/04 12:00
12/8/04 0:00
12/8/04 12:00
Date & Time

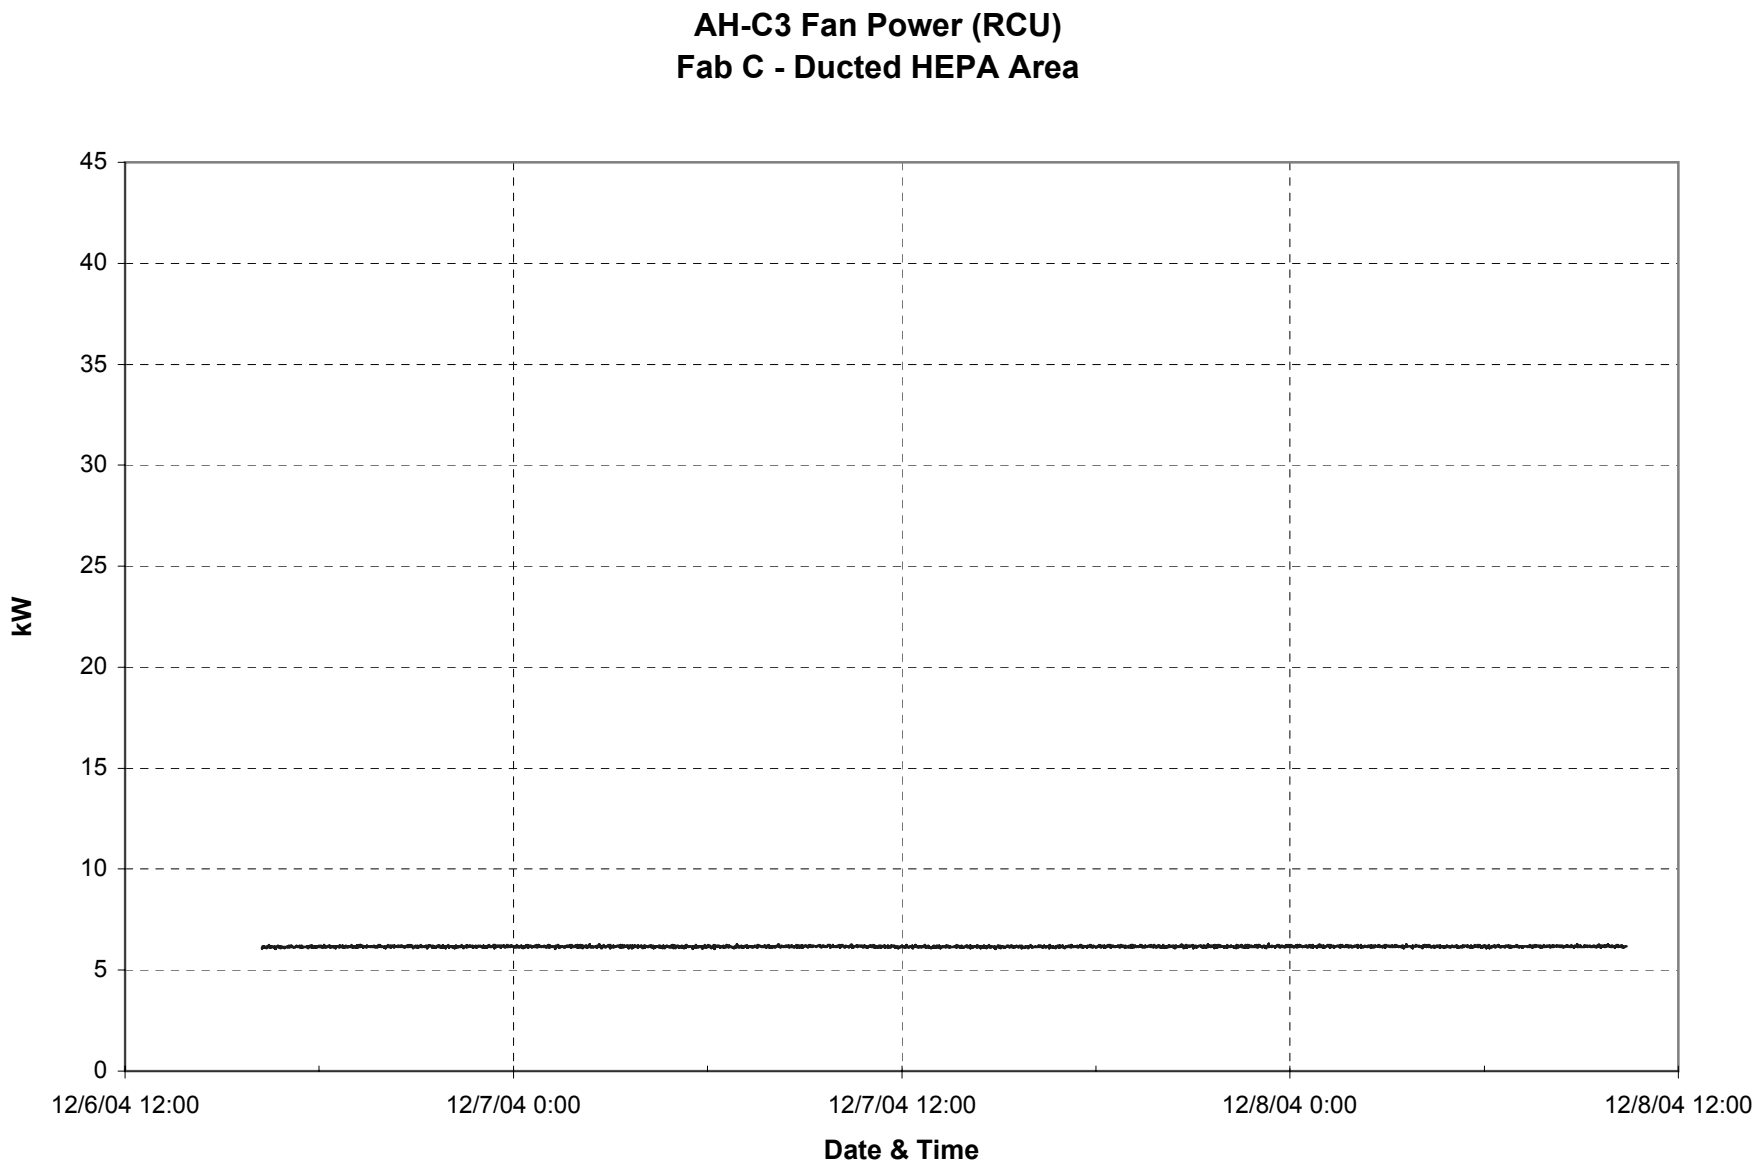
AH-C3 Fan Power (RCU)
Fab C - Ducted HEPA Area
kW
45
40
35
30
25
20
15
10
5
0
12/6/04 12:00
12/7/04 0:00
12/7/04 12:00
12/8/04 0:00
12/8/04 12:00
Date & Time

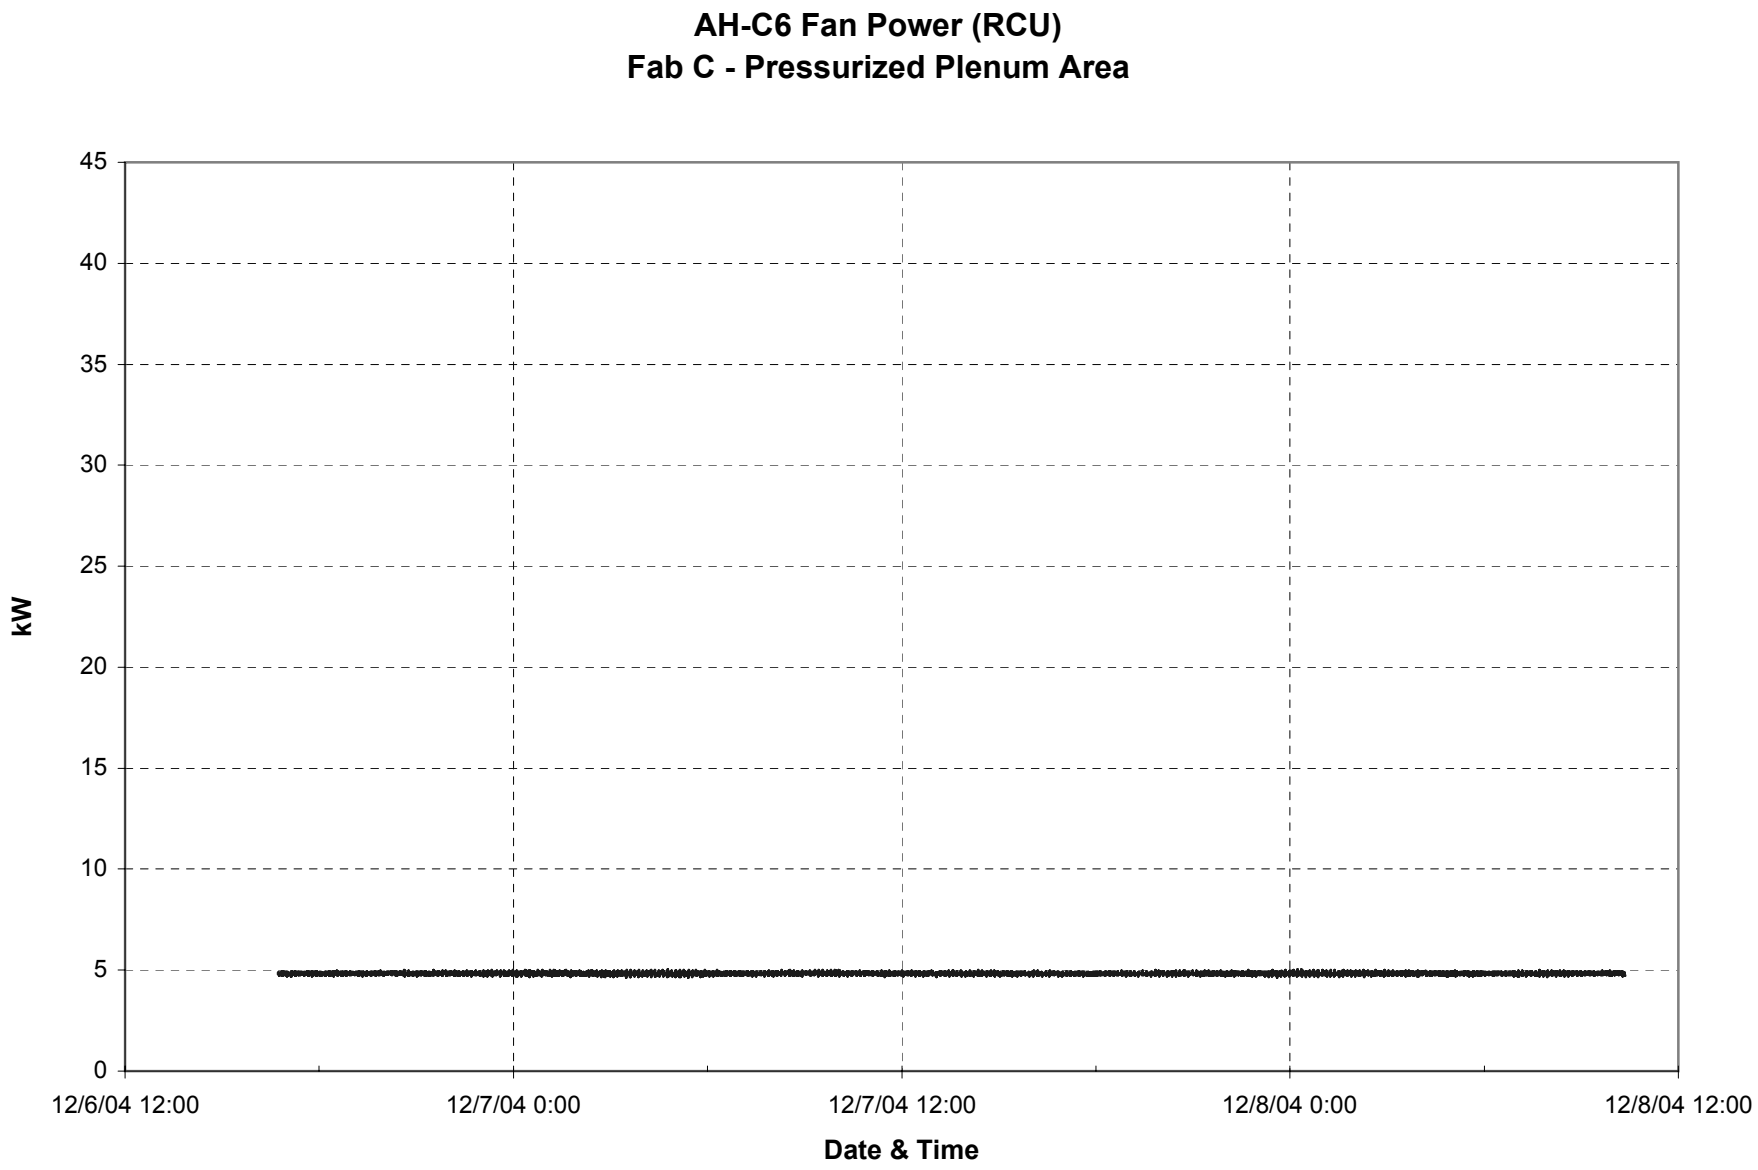
AH-C6 Fan Power (RCU)
Fab C - Pressurized Plenum Area
kW
45
40
35
30
25
20
15
10
5
0
12/6/04 12:00
12/7/04 0:00
12/7/04 12:00
12/8/04 0:00
12/8/04 12:00
Date & Time

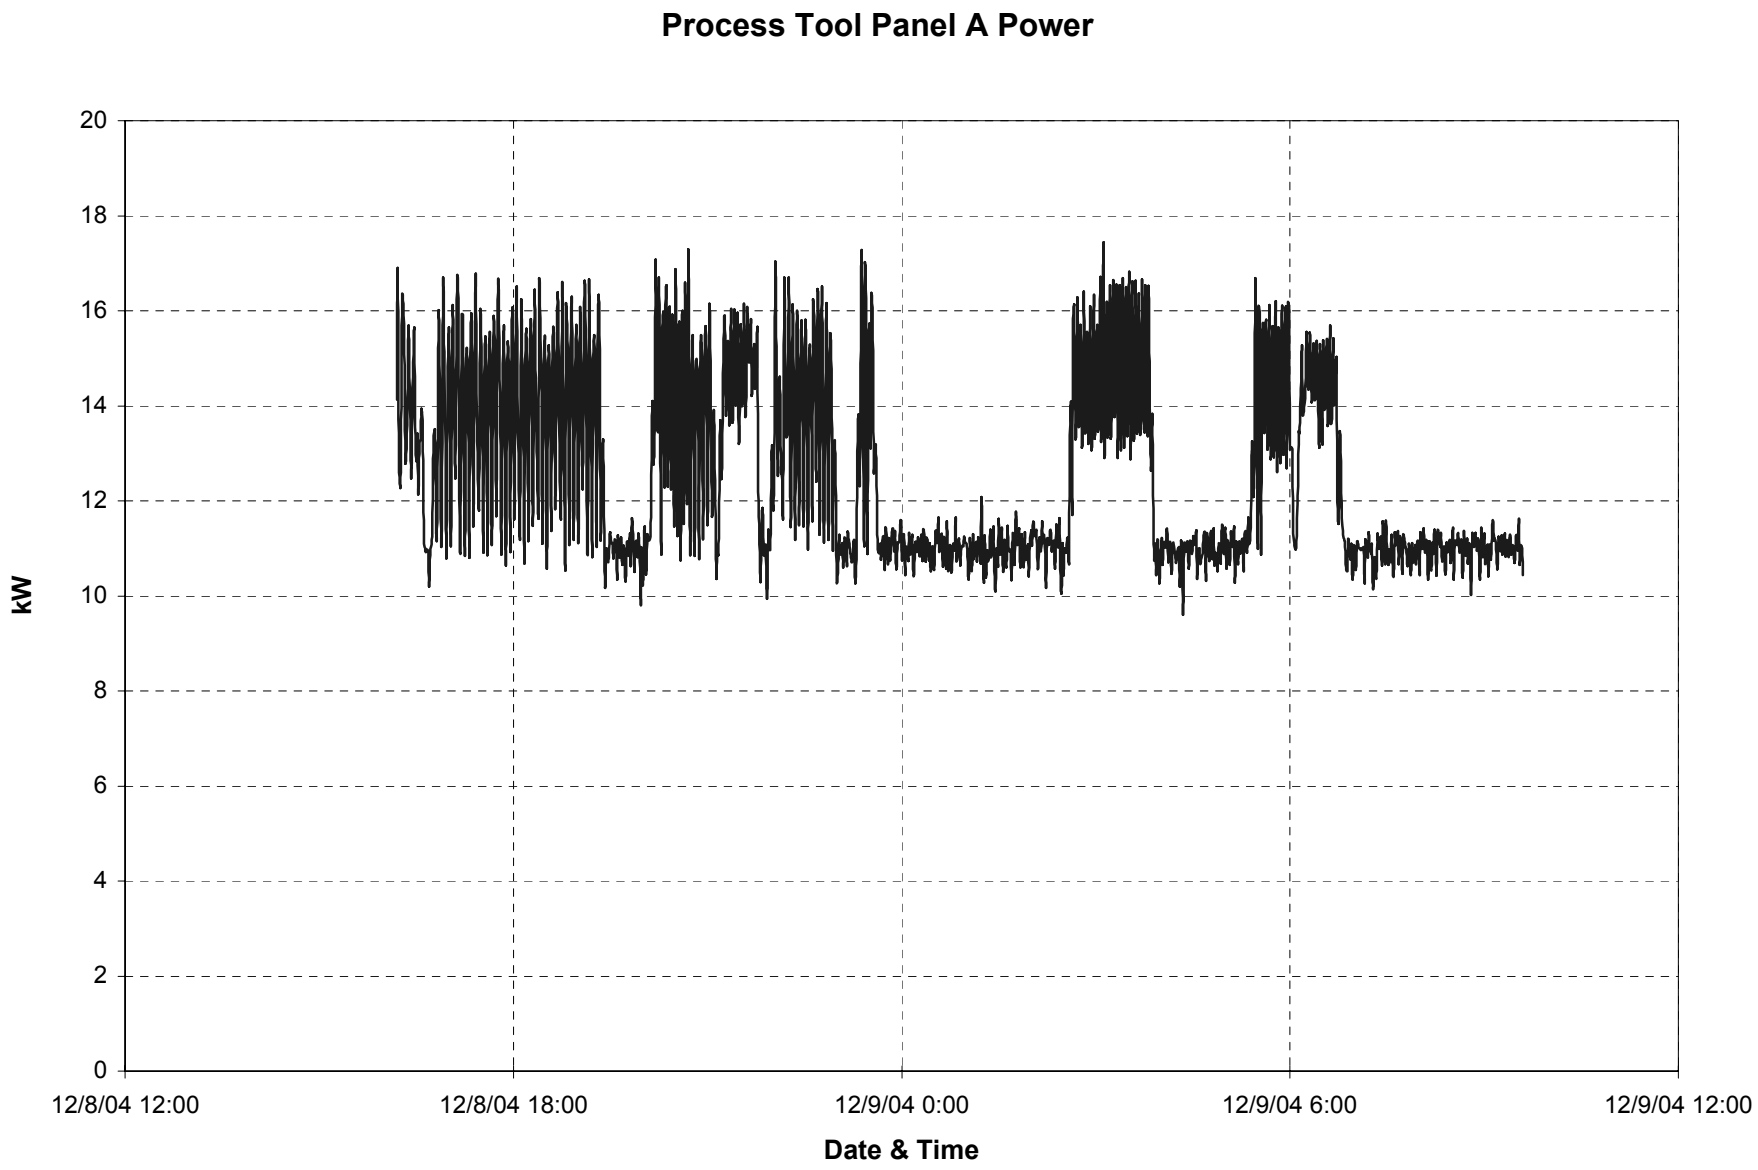
Process Tool Panel A Power
20
18
16
14
12
10
8
6
4
2
0
kW
12/8/04 12:00
12/8/04 18:00
12/9/04 0:00
12/9/04 6:00
12/9/04 12:00
Date & Time

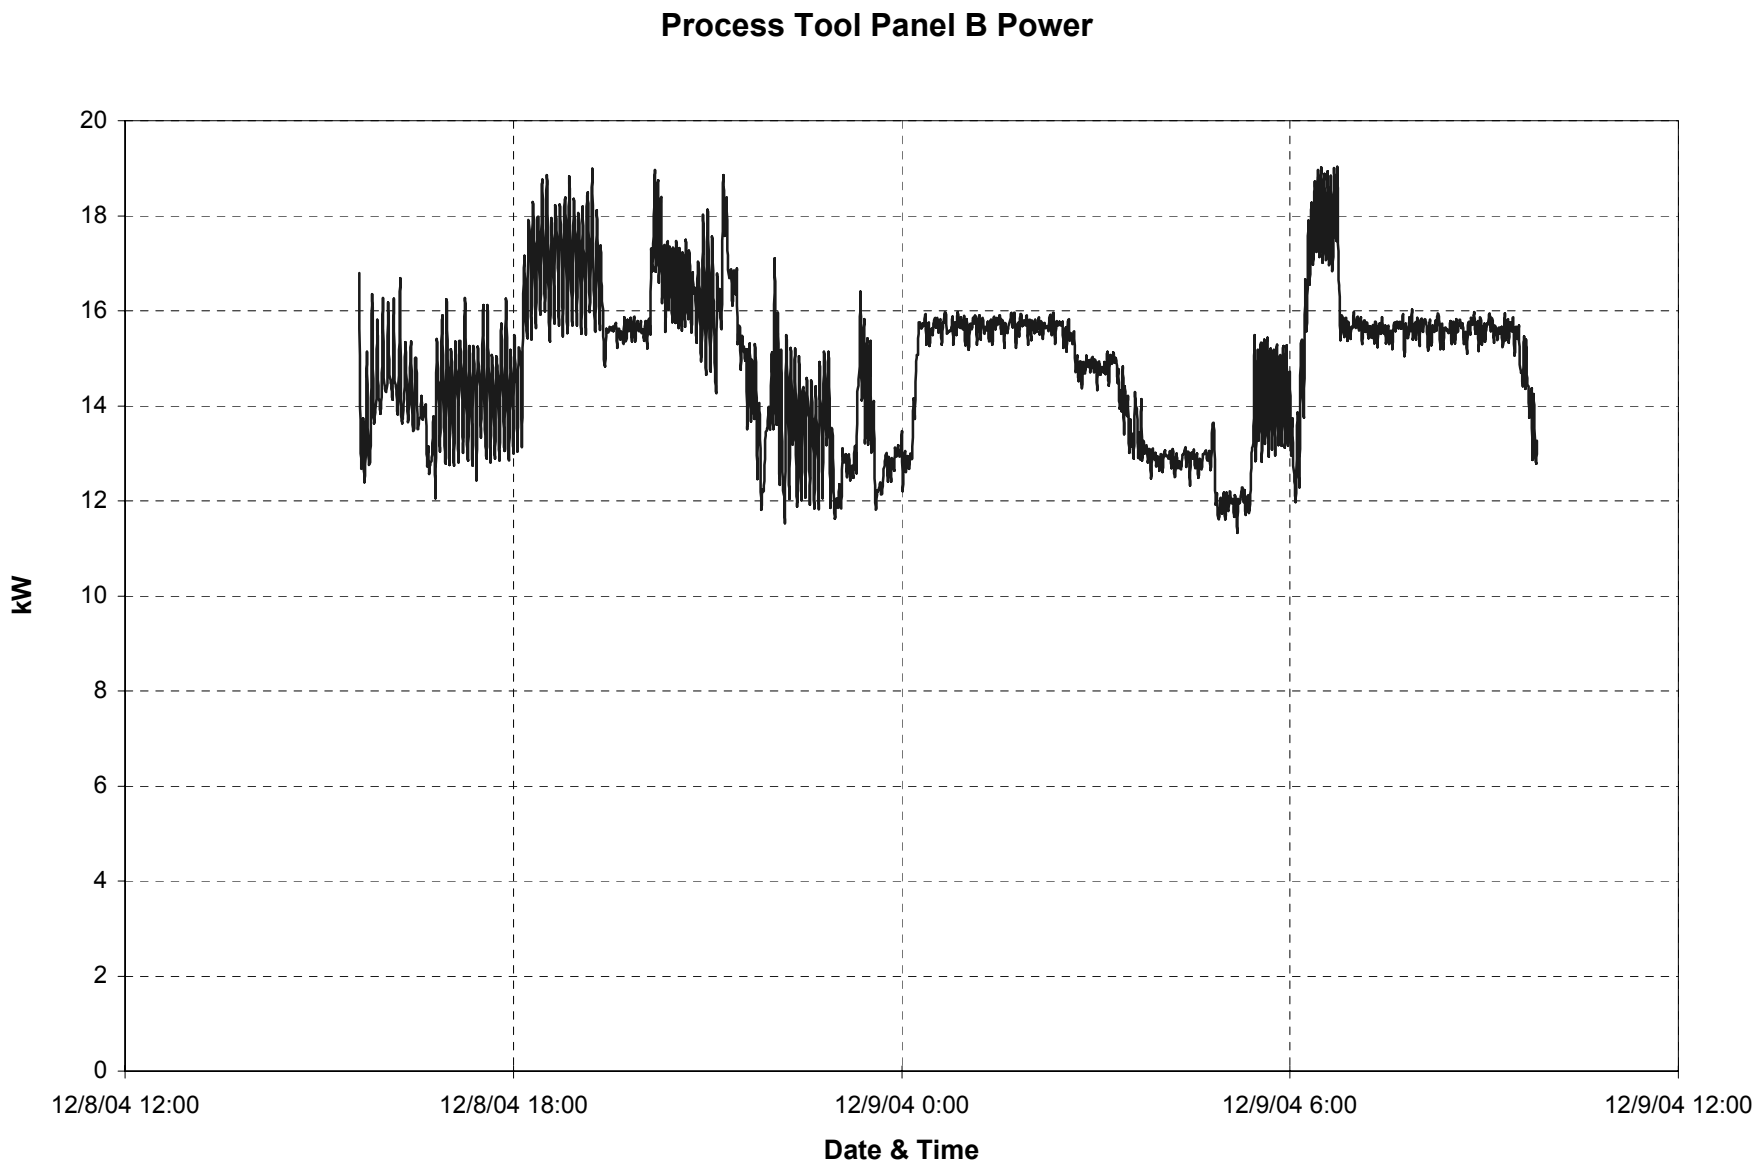
Process Tool Panel B Power
20
18
16
14
12
10
8
6
4
2
0
kW
12/8/04 12:00
12/8/04 18:00
12/9/04 0:00
12/9/04 6:00
12/9/04 12:00
Date & Time

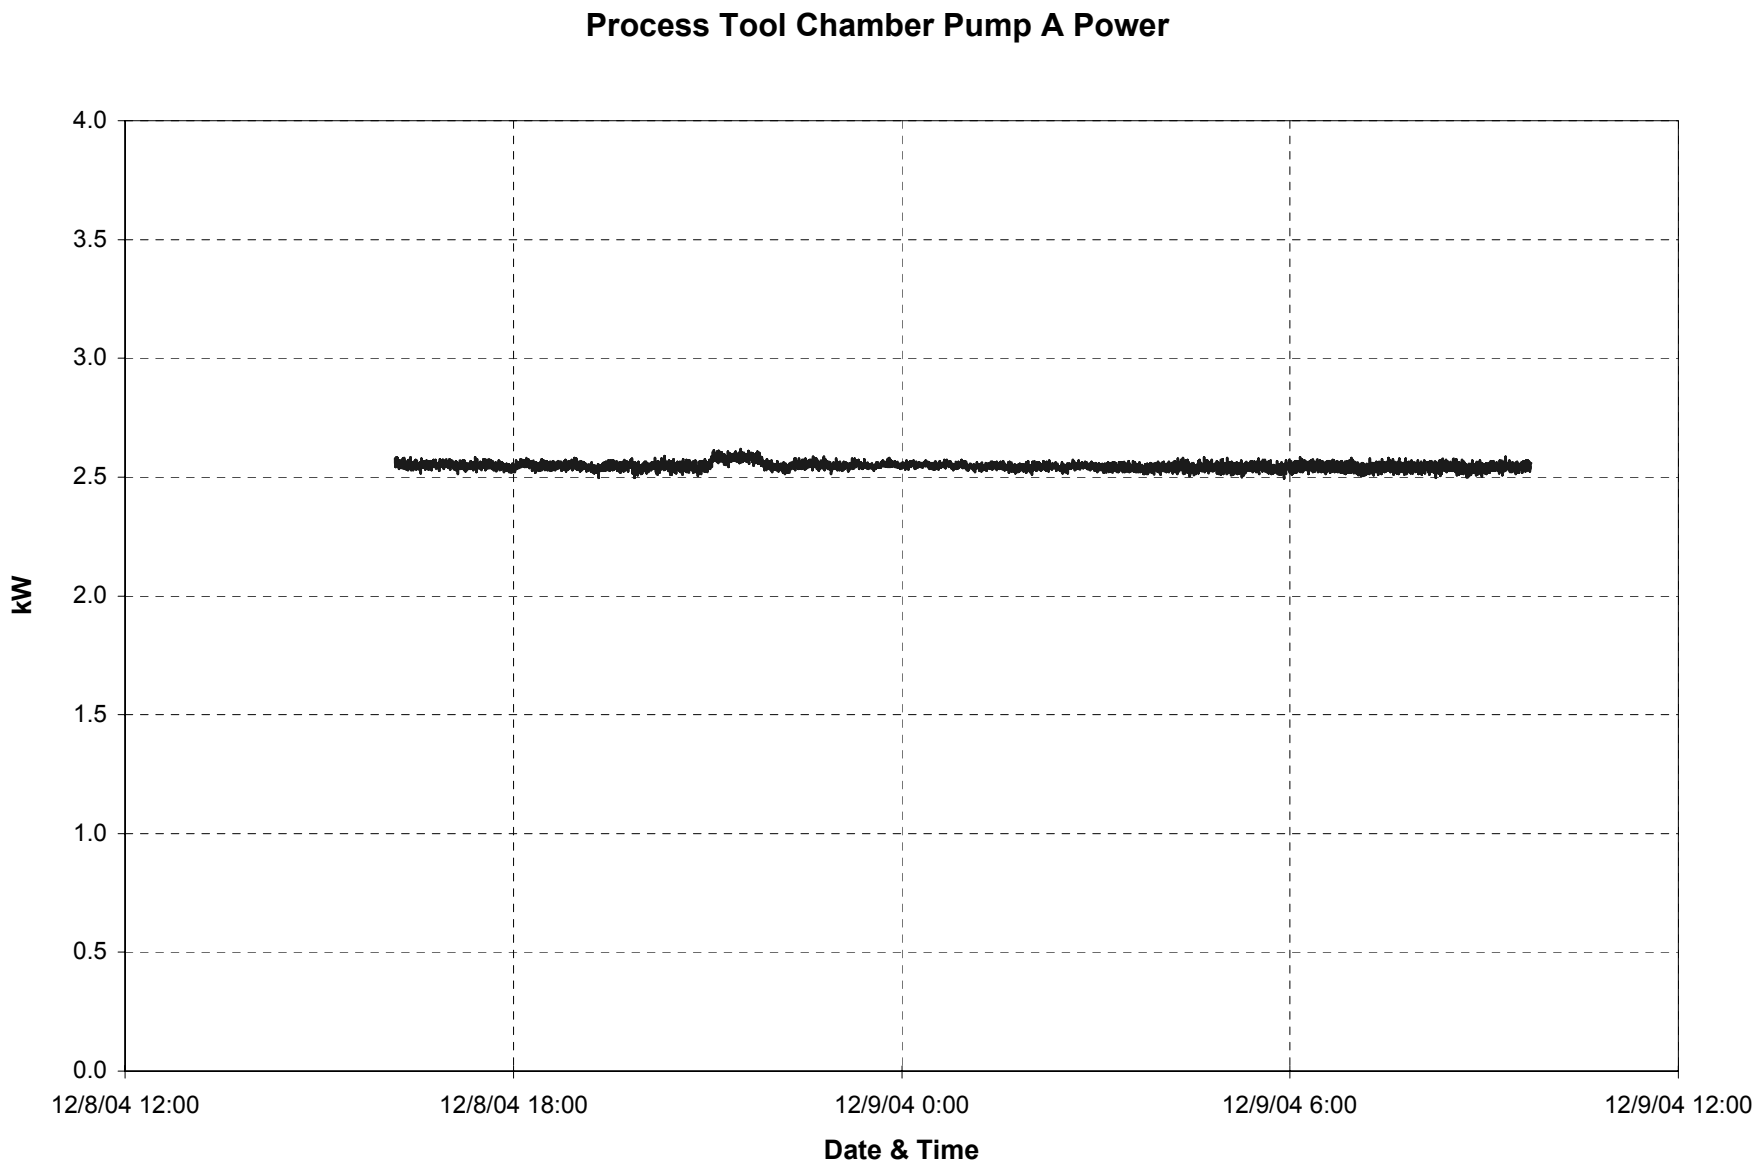
Process Tool Chamber Pump A Power
4.0
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
kW
12/8/04 12:00
12/8/04 18:00
12/9/04 0:00
12/9/04 6:00
12/9/04 12:00
Date & Time

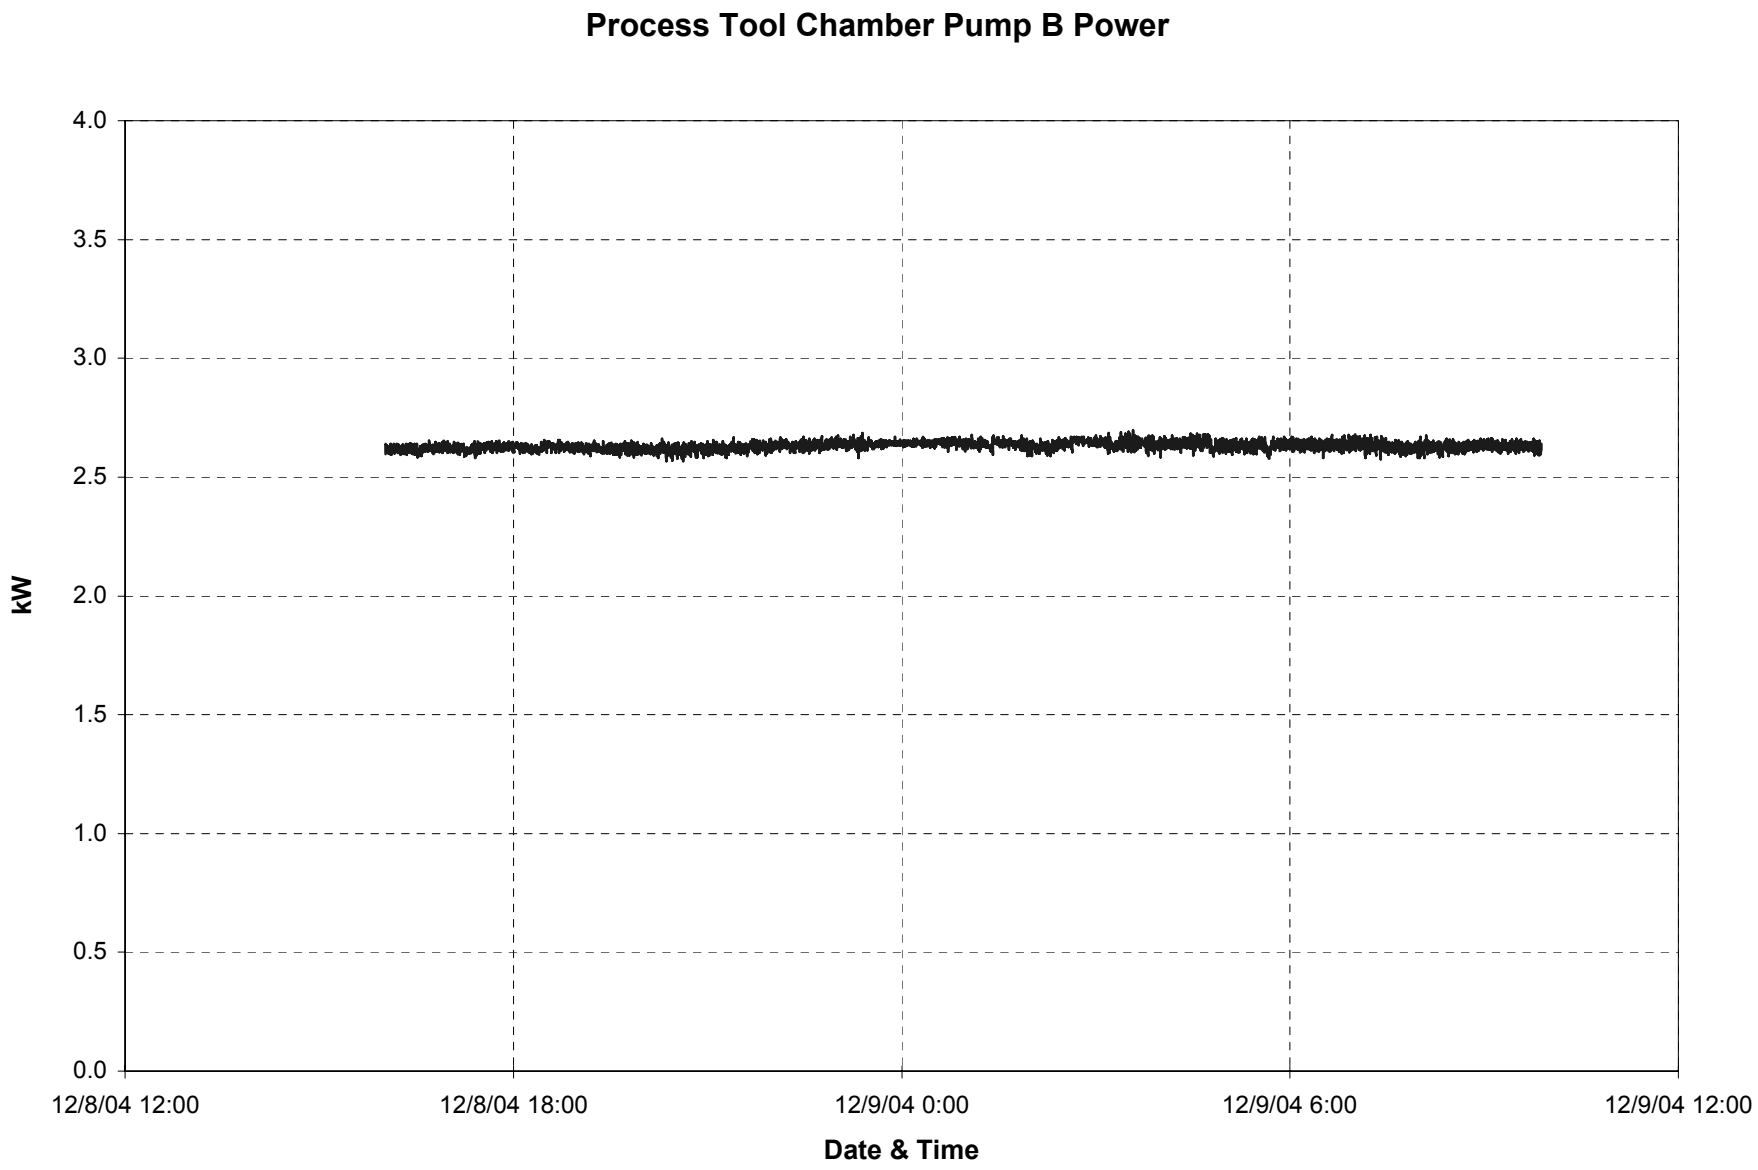
Process Tool Chamber Pump B Power
4.0
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
kW
12/8/04 12:00
12/8/04 18:00
12/9/04 0:00
12/9/04 6:00
12/9/04 12:00
Date & Time

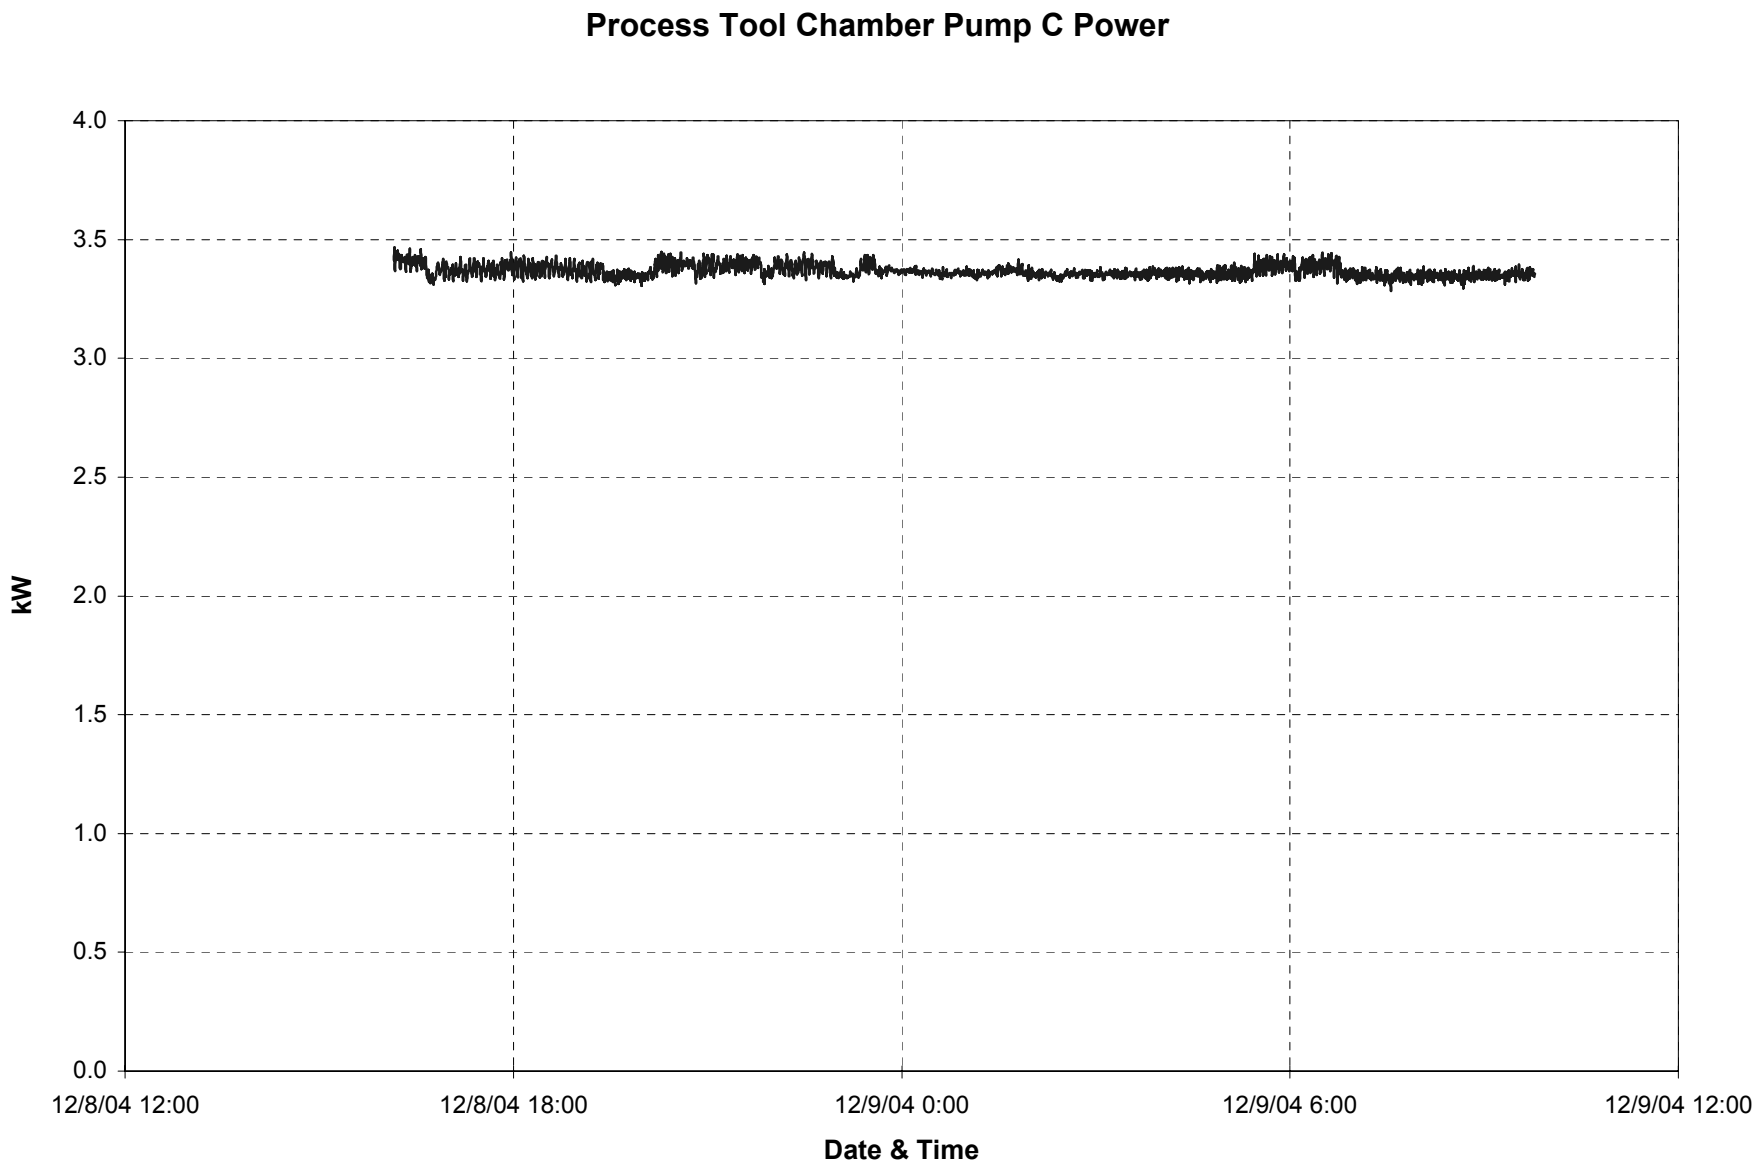
Process Tool Chamber Pump C Power
4.0
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
kW
12/8/04 12:00
12/8/04 18:00
12/9/04 0:00
12/9/04 6:00
12/9/04 12:00
Date & Time

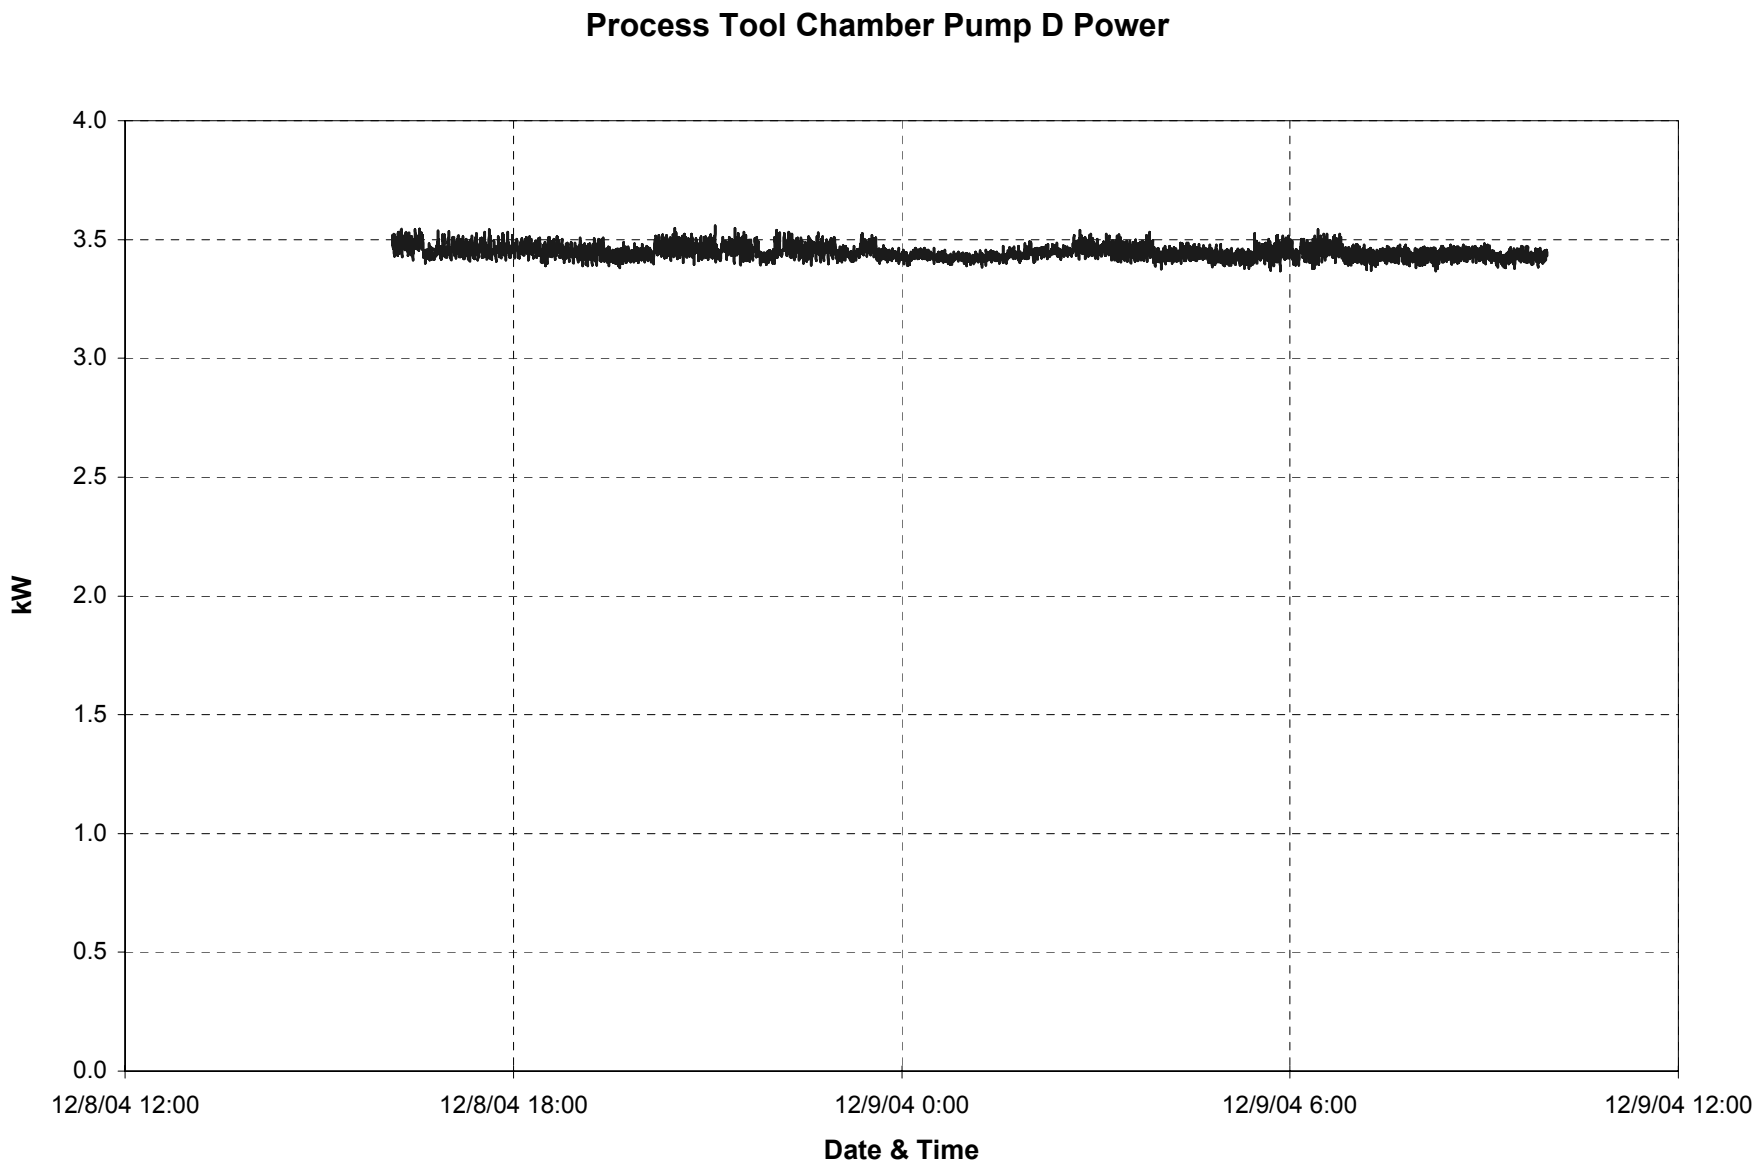
Process Tool Chamber Pump D Power
4.0
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
kW
12/8/04 12:00
12/8/04 18:00
12/9/04 0:00
12/9/04 6:00
12/9/04 12:00
Date & Time

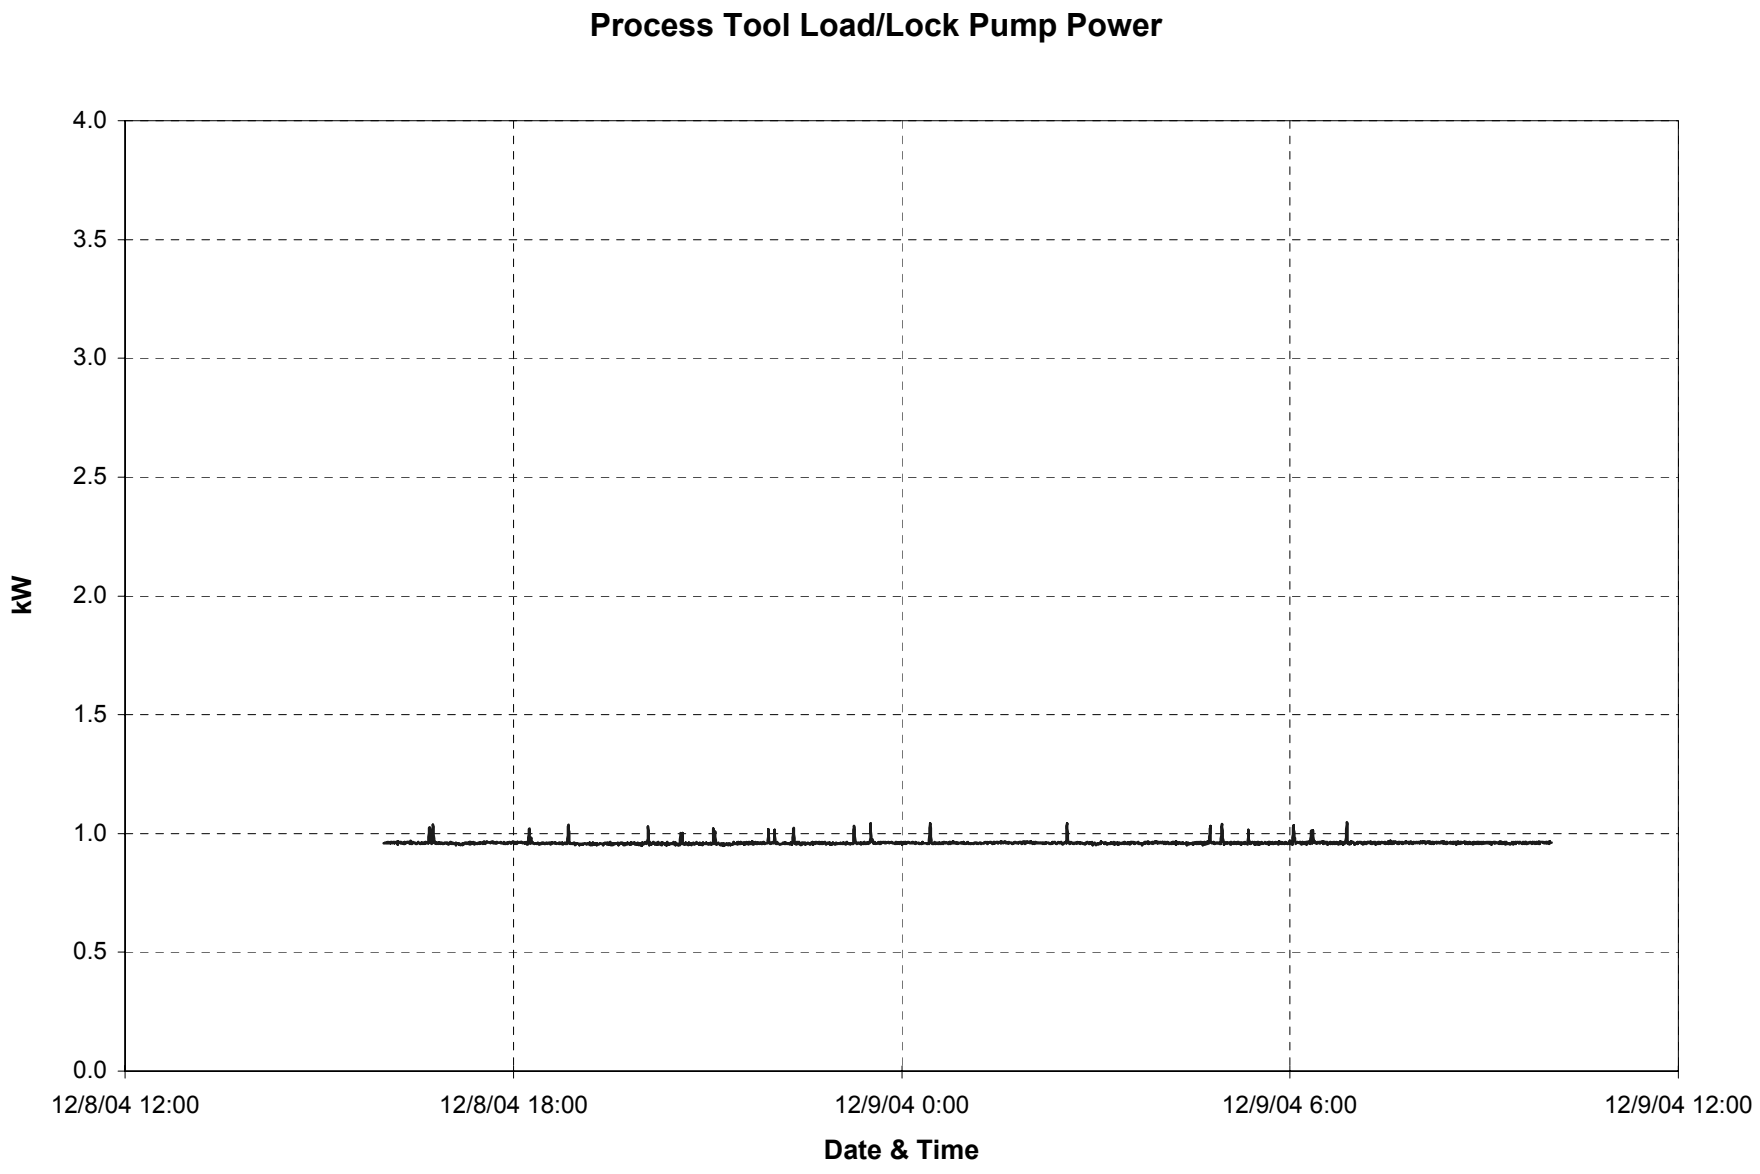
Process Tool Load/Lock Pump Power
4.0
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
kW
12/8/04 12:00
12/8/04 18:00
12/9/04 0:00
12/9/04 6:00
12/9/04 12:00
Date & Time

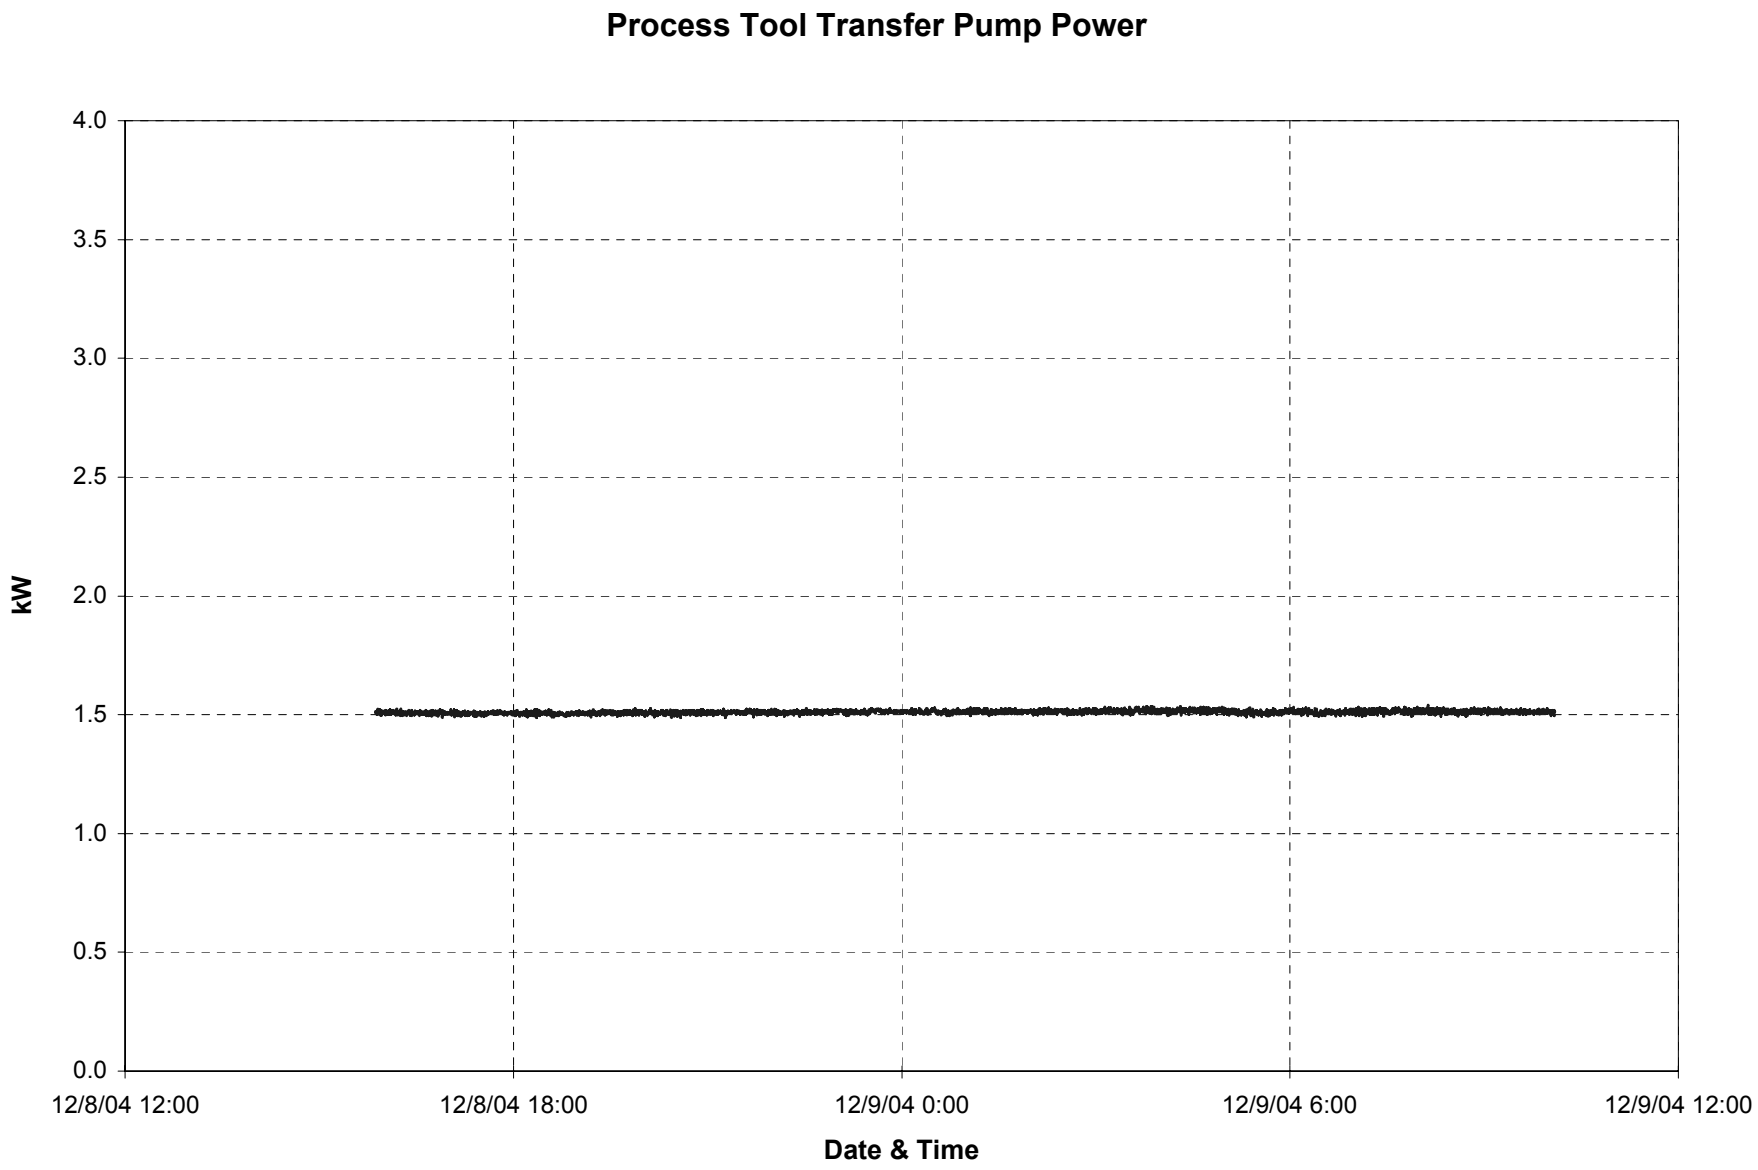
Process Tool Transfer Pump Power
4.0
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
kW
12/8/04 12:00
12/8/04 18:00
12/9/04 0:00
12/9/04 6:00
12/9/04 12:00
Date & Time

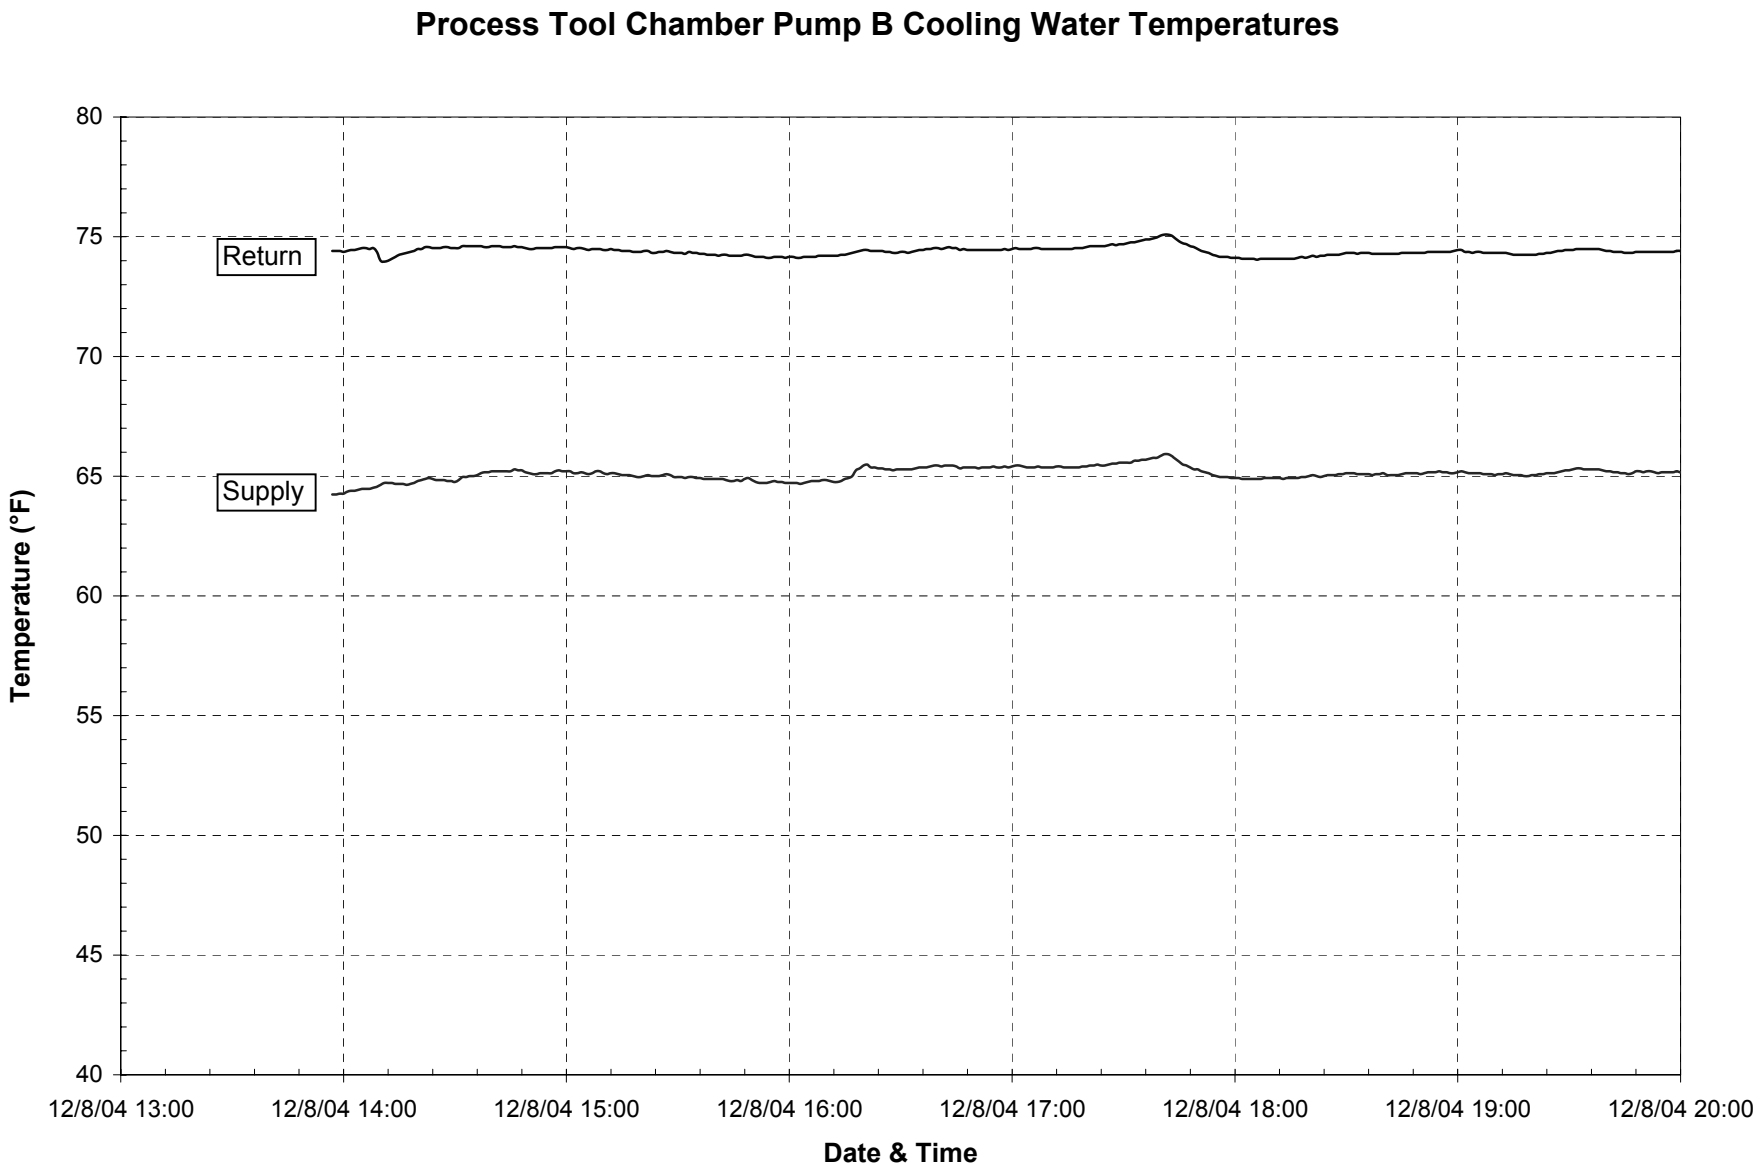
Process Tool Chamber Pump B Cooling Water Temperatures
Return
Supply
Temperature (°F)
80
75
70
65
60
55
50
45
40
12/8/04 13:00
12/8/04 14:00
12/8/04 15:00
12/8/04 16:00
12/8/04 17:00
12/8/04 18:00
12/8/04 19:00
12/8/04 20:00
Date & Time

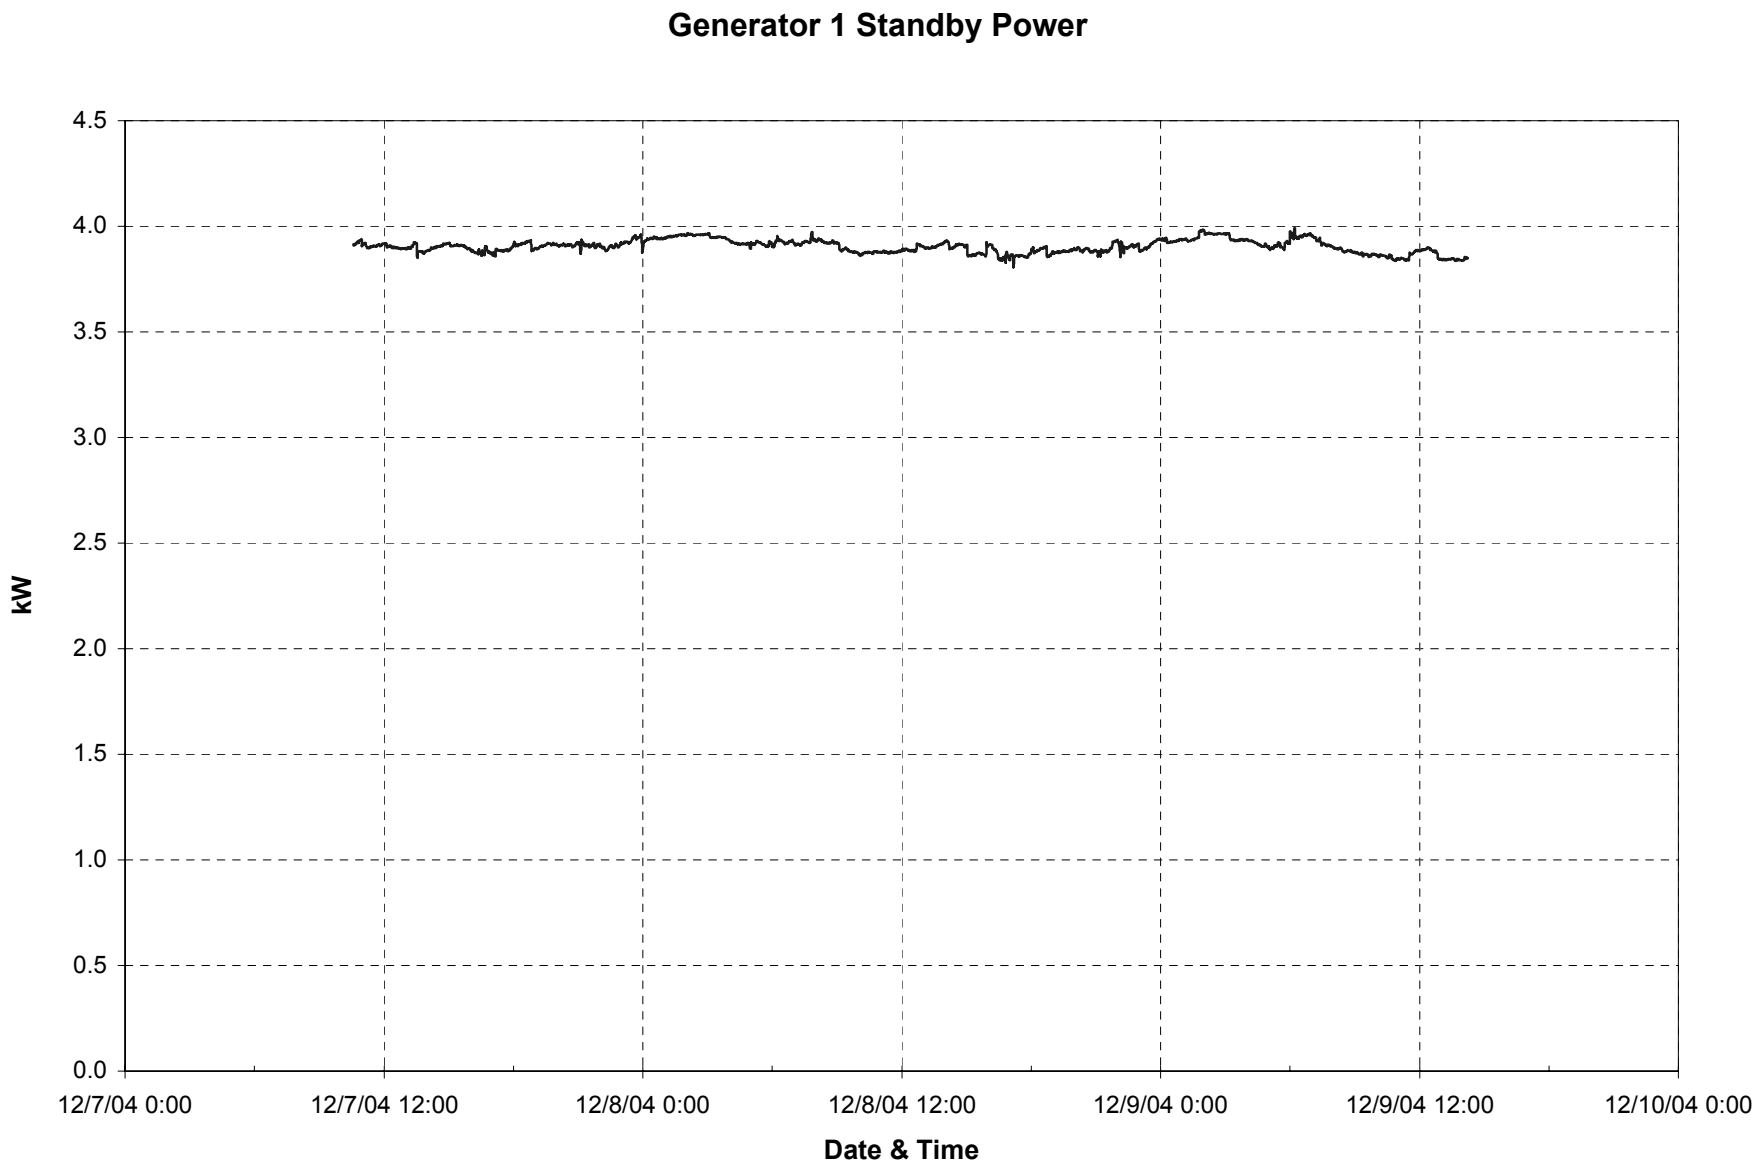
Generator 1 Standby Power
4.5
4.0
3.5
3.0
2.5
2.0
1.5
1.0
0.5
0.0
kW
12/7/04 0:00
12/7/04 12:00
12/8/04 0:00
12/8/04 12:00
12/9/04 0:00
12/9/04 12:00
12/10/04 0:00
Date & Time

Appendix B

Data Collection and Accuracy Notes

Data Collection and Accuracy Notes

The following notes describe specific measurements and what assumptions were made in preparing calculated results.

Cleanroom Data

Average Filter Velocity	Filter velocity based on airflow divided by the assumed effective HEPA filter area of 6.8 square feet.
Primary Cleanroom Area	The area certified at a rated class level. This is taken from the drawings.
Secondary Cleanroom Area	Air return area.

Recirculation Air (RCU)

Air Flow	All air flow measurements were provided by certification report.
Fan Power	Fan power was measured directly
RCU Efficiency	Number of cubic feet of recirculation air delivered to a given cleanroom, divided by the total kW of the units providing the recirculation air.

Make-Up Air Handling Unit (MAU)

Air Flow	All air flow measurements were provided by certification report.
Fan Power	Fan power was measured directly.
MUAH Efficiency	Number of cubic feet of air delivered per kW.

Mini Environment

Average Filter Velocity	Measured and supplied by client.
Air Flow	Air flow based on filter velocity multiplied by effective area. Effective area of HEPA filters assumed to be 85% of total area.
Fan Power	Fan power was measured directly with spot measurement.

Appendix C

Measurement Methodology

Measurement Methodology

Data collection measurements were made according to the following procedures:

Air Temperature/Relative Humidity

- Equipment: Pace Scientific 4-channel pocket loggers model XR440, 4" 30kOhm thermistors, or Temperature/Relative Humidity sensors.
- Attach temperature/relative humidity sensor to pocket logger channel block.
- Setup pocket logger using product software.
- Verify channels set to correct sensor type and operation.
- Set RH linear scale specified by the sensor.
- Upload setup into pocket logger to launch logging.
- Read real time data. Verify setup and actual conditions.
- Log data.

Power Trend – Elite Logger

- Equipment: Elite Logger and ELOG 97c software.
- Select current transducers (CTs) appropriate for the measurement and panel space constraints.
- Attach current transducers (CTs) to Elite logger channel block.
- Plug Elite logger into AC supply.
- Setup Elite logger using product software.
- Electrician installation of voltage sensors in the electrical panel.
- Electrician installation of CTs in electrical panel for the specified load to be measured.
- Read real time data.
- Verify balanced current as well as appropriate, balanced voltage readings.
- Secure panel door and attach caution tape and warning notice if panel cannot be locked shut.
- Log data.

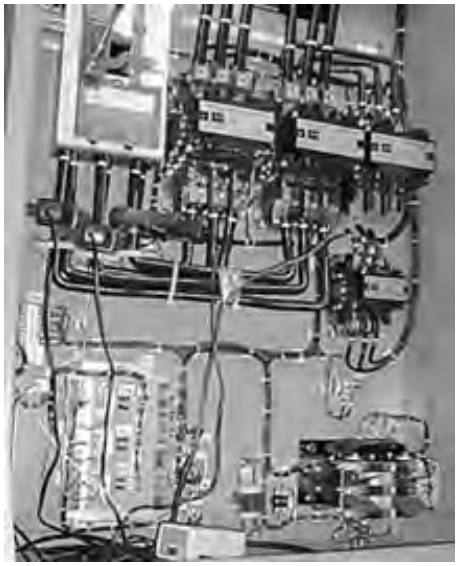

Power Spot Measurement – Power Sight

- Equipment: Power Sight PS 3000
- Plug Power Sight into AC supply if necessary.
- Connect current transducers (CTs) and voltage sensors to Power Sight.
- Electrician installation of voltage sensors in the electrical panel.
- Electrician installation of CTs in electrical panel for the specified load to be measured.
- Verify balanced current as well as appropriate, balanced voltage readings.
- Read and record the real time power reading for spot measurement.
- Log data for selected measurements.

Attachment II: Development of Fan-filter Test Procedure

Advanced Energy-Efficient Filtration:

Fan Filter Unit

Final Report (LBNL-57727)

To

California Energy Commission

2005

Tengfang Xu

Contents

1. Background 3

2. Project Objectives 4

3. Project Activities and Major Outcomes 4

4. Project Deliverables 5

5. Summary of Phase-one Demonstration Project 6

5.1 Objective 6

5.2 Demonstration Test Setup 6

5.2.1 Test Facility at PG&E 7

5.2.2 Test Facility at ITRI 8

5.3 Results 8

5.3.1 FFU Electric Power and Airflow Speed 8

5.3.2 FFU Pressure Rise and Airflow Speed 9

5.3.3 FFU Total Pressure Efficiency and Airflow Speed 10

5.3.4 FFU Energy Performance Index (EPI) and Airflow Speed 11

5.4 Conclusions and Recommendations 12

6. Acknowledgements 13

7. References 14

Advanced Energy-Efficient Filtration: Fan Filter Unit

Tengfang Xu

1. Background

With the expectation that the development of a standard method or procedure for testing and reporting performance of fan-filter units (FFUs) will continue in the industry, LBNL proposed to provide technical assistance in the testing method development in this project.

At the early stage of this project, LBNL investigated the levels of interests, expertise, and resources of a list of standard-setting organizations, which potentially would be interested in developing a national, voluntary standard for testing the performance of FFUs. LBNL communicated with the officials or members in the following organizations: Air Movement and Control Association International (AMCA), American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE), American Society for Testing and Material International (ASTM International), Institute of Environmental Sciences and Technology (IEST), and other relevant entities. Based upon the evaluation of the levels of interests, availability of expertise and resources, we identified the following key parties and players:

- Industrial Technology Research Institute (ITRI), Taiwan.
- The Air Movement and Control Association International (AMCA).
- Institute of Environmental Sciences and Technology (IEST).
- FFU suppliers and users.

Upon further discussion with the interested parties, and stakeholders that include the members of the Project Advisory Committee and the California Energy Commission, LBNL decided to interact closely with these key players and to develop a draft standard method for testing and reporting energy performance of fan-filter units.

In November 2003, IEST created and formalized a Working Group (WG) consisting of approximately twenty or more experts associated with the industries, standard-setting entities including AMCA, and the government. The IEST WG started to develop the scope of an industrial Recommended Practice (RP) document on the standard method for testing FFU performance. The WG, which is titled "Testing Fan Filter Unit," has since met twice a year and established the collaboration and made voluntary effort by the WG members to develop a Recommended Practice (RP) document for testing FFUs in laboratory setting.

Establishing a testing method to consistently report the energy performance of fan-filter units (FFU's) will help cleanroom owners and designers to make informed choices that consider energy efficiency among other important considerations. Being able to providing comparative

performance information to owners and designers can facilitate the selection of more energy efficient FFU models. Energy efficiency practice can be encouraged through utility incentive programs, which can be made possible by using the performance data reported from standardized testing, or by establishing minimal performance rating criteria.

2. Project Objectives

The objective of this project is to provide assistance in development of a standard test procedure for fan-filter units, which are gaining popularity for use in California cleanrooms. In particular, LBNL carried out collaboration with various stakeholders in the industry and took a lead in developing a draft standard method for testing the energy performance of fan-filter units, and provided assistance to California public utility companies by testing the draft method in PG&E's testing facility. Through testing more units in the future with a robust standard method, baseline performance information can be developed for use in possible energy incentive programs.

3. Project Activities and Major Outcomes

In the course of this project, LBNL collaborated with the ITRI of Taiwan, AMCA, IEST, PG&E, suppliers and users of fan-filter units. The major activities and outcomes are summarized in the following:

- Developed a draft standard method for testing and reporting performance of fan-filter unit. The draft outline and the first version of the standard test procedure were developed and issued for reviews by the Project Advisory Committee (PAC) members. A second draft was made available on LBNL website for public reviews from April 2004 to June 2005 [1]. Additionally, the second draft was sent to a wide range of stakeholders for additional reviews. Based upon input from the industry and interactions with IEST Working Group 36 "Testing Fan Filter Unit," a final draft of the standard test method was prepared along with in this final report for delivery to the California Energy Commission [2]. The developmental output from the project is well received by the stakeholders and is being integrated into the document on Recommended Practice by a national standard development organization, i.e., IEST.
- Completed demonstration tests at PG&E's testing facility, San Ramon, California. Partnering with ITRI and PG&E, LBNL carried out research and the initial demonstration test on testing the draft standard method at PG&E's facility. This initial demonstration project tested the usability of the draft standard method, and preliminarily demonstrated that testing energy performance of fan-filter units based upon a consistent method was necessary and viable. The use of this draft method has produced useful results that could be adopted to establish a baseline line criteria for energy incentive program in the future.
- In addition, in collaboration with ITRI, LBNL conducted and completed another demonstration test on testing the draft standard method at ITRI's facility, Taiwan. The same fan-filter unit was tested at both facilities. The purpose of the extra effort was to start investigating some of the technical issues affecting the draft standard method, although they are not within the scope of the existing CEC project. This additional research and

demonstration, in the meanwhile, has identified a number of technical issues that merit further investigations. Being able to address these issues would enhance the robustness of the draft standard test method developed in this phase so that its adoption by various users could be maximized.

- Overall, synergistic activities in the duration of this project have generated interests and raised energy-efficiency awareness among fan-filter professionals. LBNL's interactions with the industry, e.g., IEST, SEMI, Sematech International, and utility companies, have enhanced industry liaisons (suppliers and end-users). Through continued interactions with the stakeholders (e.g., involvement with the IEST working group and utility companies), energy-efficiency awareness can be continuously enhanced among the professionals and current practices in California cleanroom industries can be expected to improve.

4. Project Deliverables

Project deliverables include the following:

- Completed the draft standard test procedure for fan-filter units. The final draft of standard test procedure for fan-filter units was developed, and was made available in the LBNL web page. The development has stimulated significant interests from the stakeholders including users, suppliers, and various professional organizations.
- Completed the phase-one demonstration project – to test the standard test procedure. As contracted, LBNL has completed the phase-one demonstration project at PG&E's testing facility. In addition, in collaboration with ITRI, LBNL also completed an extra research and demonstration project at ITRI's facility. A summary of the phase-one demonstration project – to test the standard test procedure, is included in this report.
- PG&E has suggested that once an energy performance test standard is available, demand side management programs would like to include high efficiency fan filter units in possible rebate programs. In this project, LBNL provided technical assistance through various activities, including poster, meetings, and the phase-one demonstration at PG&E's facility in the course of developing a baseline for its possible use in establishing a rebate program.

In addition, LBNL has produced the following material resulting from synergistic activities and collaboration with others [3, 4, 5, 6, 7]. The material has been distributed among and well received by the stakeholders and the industry.

- An award-winning journal article on the fan-filter unit work, as a result of collaborating with ITRI in testing FFUs [3].
- Several articles were published in IEST, SEMICON, and Controlled Environments Magazine (formerly A2C2), in collaboration with ITRI in this phase of project. The articles addressed findings from testing the FFUs based upon ITRI's data and focused on energy performance of FFUs to inform and educate cleanroom professionals. The publications and associated presentations were well received by the industry [4, 5].

- A poster and various presentations were prepared and distributed to California Public Utilities [7]. The concept of performance testing method was introduced to the industry, including three utility companies in California, SEMATECH International, and IEST [3, 4, 5, 6, 7].

5. Summary of Phase-one Demonstration Project

5.1 Objective

The main objective of the demonstration is to test the draft standard method being developed at Lawrence Berkeley National Laboratory [1], in PG&E's facility. In collaboration with ITRI, LBNL conducted additional research and demonstration at the ITRI facility, Taiwan.

This report includes: 1) demonstration of testing the draft standard test method for aerodynamic and energy performance in two laboratory-testing facilities, and 2) comparing the tests and suggesting areas for future investigations.

5.2 Demonstration Test Setup

The FFU used in the demonstration project consisted of a small fan, and a HEPA filter, and a sealed box sized to fit into standard cleanroom ceiling grids 2 ft by 4 ft as shown in Figure 1.

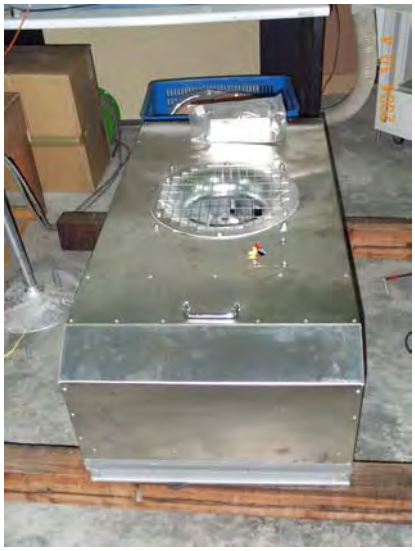

Figure 1 Fan-filter Unit Used in the Tests

The FFU modulation device was integrated with the FFU for setting fan-wheel rotation speeds (Rotation Per Minute, RPM). The fan-wheel motor in the FFU was set at different RPMs during the testing to obtain performance data. At each RPM setting, the FFU was tested at various pressure-rise across the unit, which was modulated by adjusting damper positions. The ambient conditions and the airflow conditions were recorded and were used for the air density conversion to the equivalent standard condition (i.e., 1 atm, 20°C). We assume that the airflow was isothermal, although a small fraction of heat was generated from fan motors, which was

transferred to the airflow. The reported performance data were based upon the standard air condition, i.e., with the air density of 1.20 kg/m^3, in order to directly compare the energy and aerodynamic performance.

Through interactions with collaborators, we determined the following parameters to include in the testing and reporting:

- Unit's airflow rates (or actual airflow speeds).
- Pressures rise across the unit.
- Total electric power demand.

5.2.1 Test Facility at PG&E

The FFU tested was mounted vertically on the inlet end of an air chamber for airflow measurement. The chamber contains a multiple-nozzle bank for recording airflow rates through the tested unit. Figure 2 shows the floor plan of facility setup at PG&E. The airflow through the FFU was supplied and pre-filtered from the atmosphere upstream of the FFUs. A booster fan and a damper were installed at the chamber outlet to modulate air pressures across the FFU so that the airflow rates were controlled.

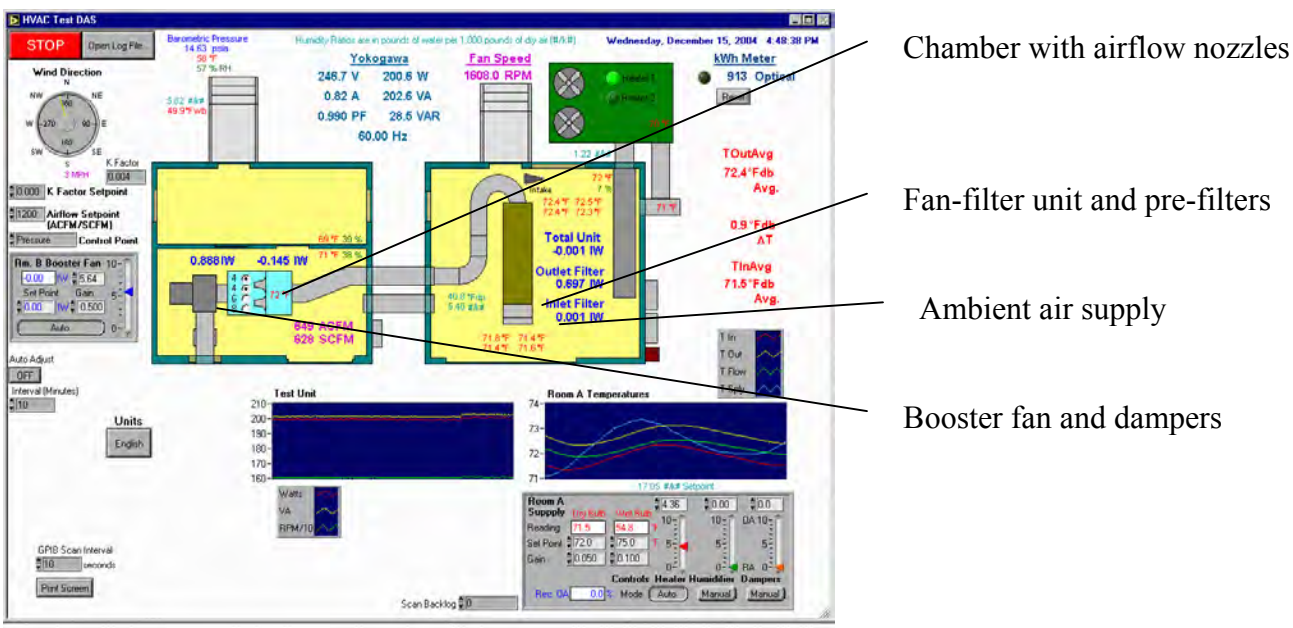

Figure 2 Laboratory measurement layout at the PG&E facility

5.2.2 Test Facility at ITRI

The FFU tested was mounted vertically on the exit end of the air chamber (

Figure 3). The chamber contains a multiple-nozzle bank for recording airflow rates through the tested unit.

Figure 3 Laboratory measurement layout at the ITRI facility

Figure 3 shows that the air from the downstream of the FFU was discharged to the atmosphere from the filter face. A booster fan and a damper were installed at the chamber inlet to modulate the airflow across the FFU. Exit airflow of the FFU discharged into a room with the air at atmospheric conditions.

5.3 Results

5.3.1 FFU Electric Power and Airflow Speed

Figure 4 shows the measured FFU electric power demand as it correlated with the airflow speeds at FFU exit at two fan rotation speeds, i.e., 1630 RPM and 1520 RPM, respectively. From the results obtained from PG&E's facility, it's clear that the electric power demand increased with the increase in airflow speeds within the operating range of the FFU for both RPMs. In addition, the electric power demand was higher with the higher RPM at any given airflow speed within the FFU's the operating range.

On the other hand, measurement results of electric power demand at PG&E's facility were slightly different from those obtained from ITRI's facility. Because the actual power supplies at the two facilities were different as a result from facility constraints, the measured data obtained from both facilities was not sufficient to address the issue on the repeatability. A power supply with 240-Volt and 60-Hertz was available and used in PG&E's facility, while a power supply with 220-Volt and 60-Hertz was available and used in ITRI's facility. The constraints of the

power supplies offered at the two test facilities and the difference in the actual RPMs of the fan wheels probably contributed to the observed difference in the electric power demand for the same unit, which would have been expected at the same airflow rate.

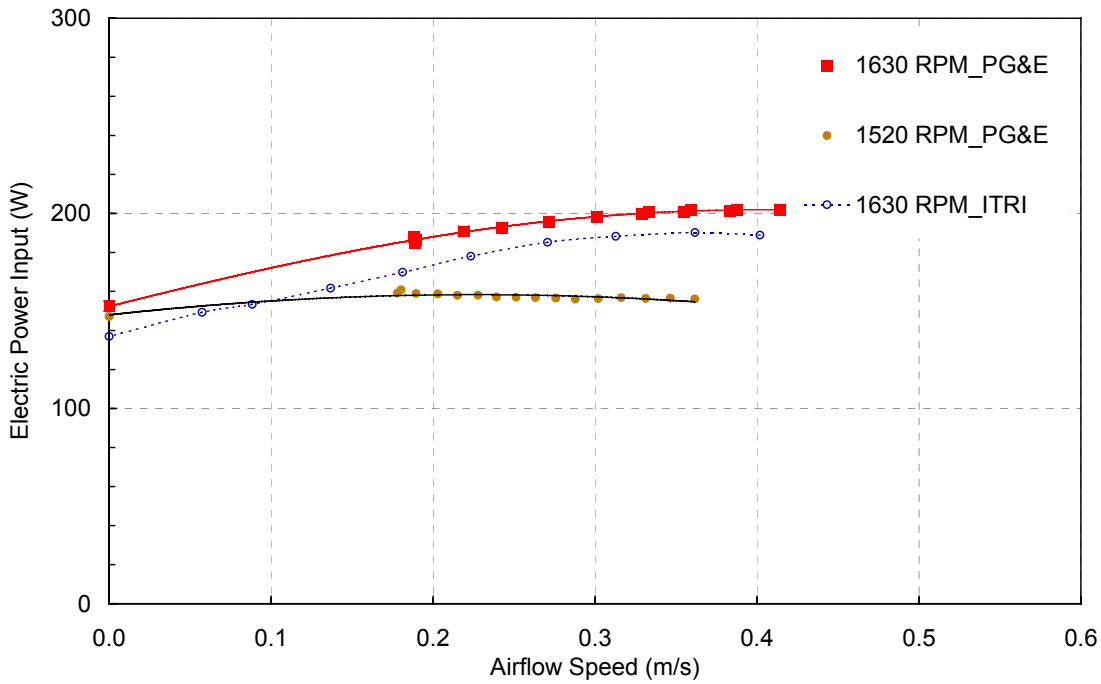

Figure 4 Total Electric Power Demand vs. Airflow Speed at FFU Exit

5.3.2 FFU Pressure Rise and Airflow Speed

Figure 5 shows the measured FFU pressure rise versus airflow speeds at FFU exit at two fan rotation speeds, i.e., 1630 RPM and 1520 RPM, respectively. From the results obtained from PG&E's facility, it's clear that the FFU pressure rise decreased with the increase in airflow speeds within the operating range of the FFU for both RPM setting. In addition, the FFU pressure rise across the unit was higher with the higher RPM at any given airflow speed. For a given airflow speed, the pressure rise across the FFU was maximized by operating the unit at its highest RPMs.

On the other hand, pressure rise across the FFU at PG&E's facility was slightly different from that obtained from the ITRI's facility at any given airflow speed. The significance of the observed difference and the possible influence from the constraints in equipment layout and power supplies is a complex issue and should become a subject for future investigations.

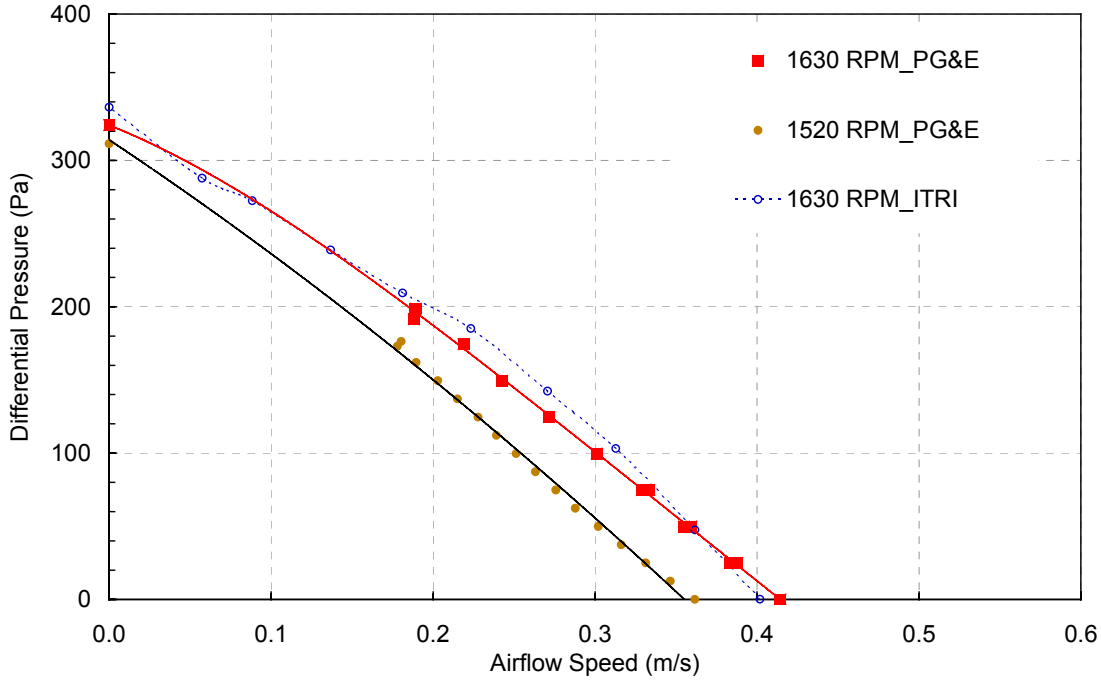

Figure 5 FFU Pressure Rise vs. Airflow Speed at FFU Exit

5.3.3 FFU Total Pressure Efficiency and Airflow Speed

The total FFU power efficiency (E_t) is defined as the actual airflow dynamic power divided by the total electric power input to the FFU unit. The total FFU power efficiency includes electrical efficiency and mechanical efficiency of the whole FFU unit and it takes into account fan motors, transformers, etc.

$$E_t = P_t Q / W \qquad \text{Equation [1]}$$

where
P_t is the FFU pressure rise (Pa)
Q the airflow rate (m^3/s), and
W is the total electric power input to FFU (W).

Figure 6 shows the performance curves of the FFU in terms of total pressure efficiency as it changed with the airflow speeds at the FFU exit for two different RPMs. The total pressure efficiency of the FFU varied considerably at various rotational speeds and airflow speeds. For a given airflow speed, the total pressure efficiency may be maximized by operating the unit at its highest RPM in general.

Additionally, the calculated values of total pressure efficiency of the unit based upon the ITRI's facility tended to be slightly higher than those obtained from PG&E's facility. This was most likely due to the lower electric power supplied (Watt) to the FFU unit tested at the ITRI facility

at a similar operating condition, even though a slightly higher pressure-rise across the FFU was developed for certain range of airflow rates (speed). Because the unit was operating by different power supplies, a same output (total pressure efficiency) at the same airflow speed should not be expected in the comparison. Adding to the possible variations in fan-wheel rotation speeds, the effect of test rig configurations on the measured results could not be sufficiently determined from the data generated from the facilities in this study. The effect of test rig configurations will need to be determined through additional investigations.

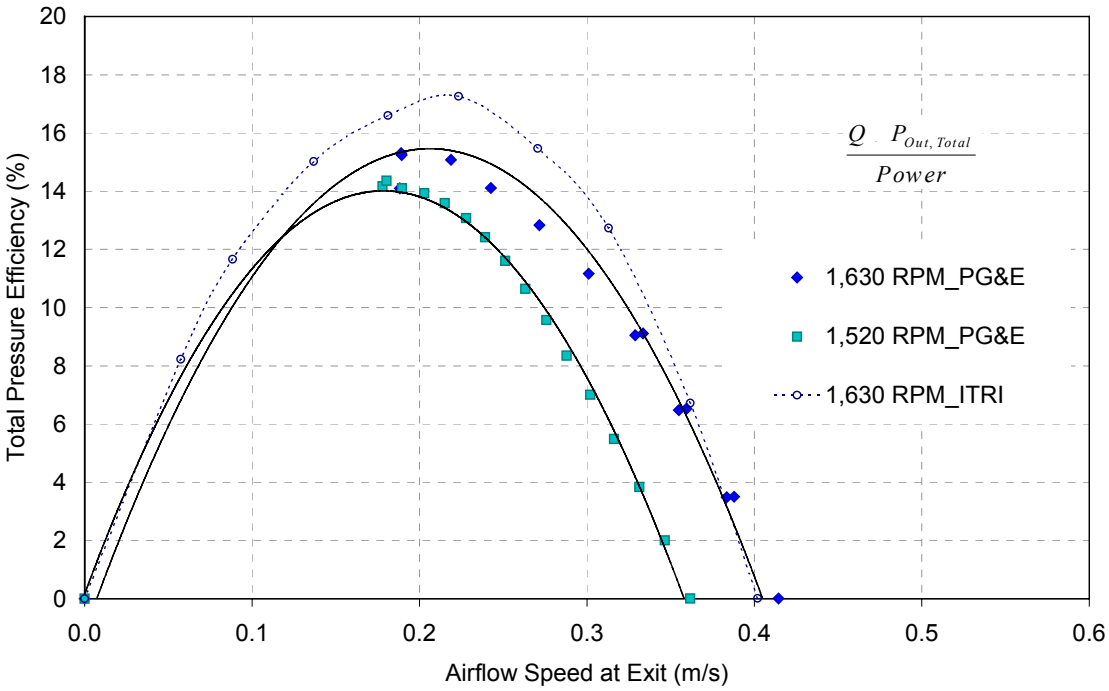

Figure 6 Total Pressure Efficiency vs. Airflow Speed at FFU Exit

5.3.4 FFU Energy Performance Index (EPI) and Airflow Speed

Figure 7 shows the measured Energy Performance Index (EPI, FFU electric power demand per airflow rate) under a range of the airflow speeds at unit's exit. The EPI values decreased with the increase in airflow speeds. In this regard, the FFU operated more efficiently in delivering the airflow at a higher airflow speed. This exhibited a similar trend found for other FFUs tested by ITRI in a previous study.

On the other hand, this indicates that the fan inside the FFU with wider damper opening would not have to work as hard as it would have to, compared to other cases in which the system-resistance increases (i.e., a narrower damper opening). In general, for a given airflow rate (speed), the EPI value may be reduced by operating the unit at a lower RPM.

Similar to the trend observed in total pressure efficiency, the calculated EPI values of the unit based upon the ITRI's facility tended to be slightly lower than those obtained from PG&E's facility. This was most likely due to lower electric power supplied (Watt) to the FFU unit tested at the ITRI facility at a similar operating condition. Because the unit was operating by different power supplies, a same EPI value at the same airflow speed would not be expected when comparing the results between the two facilities, however.

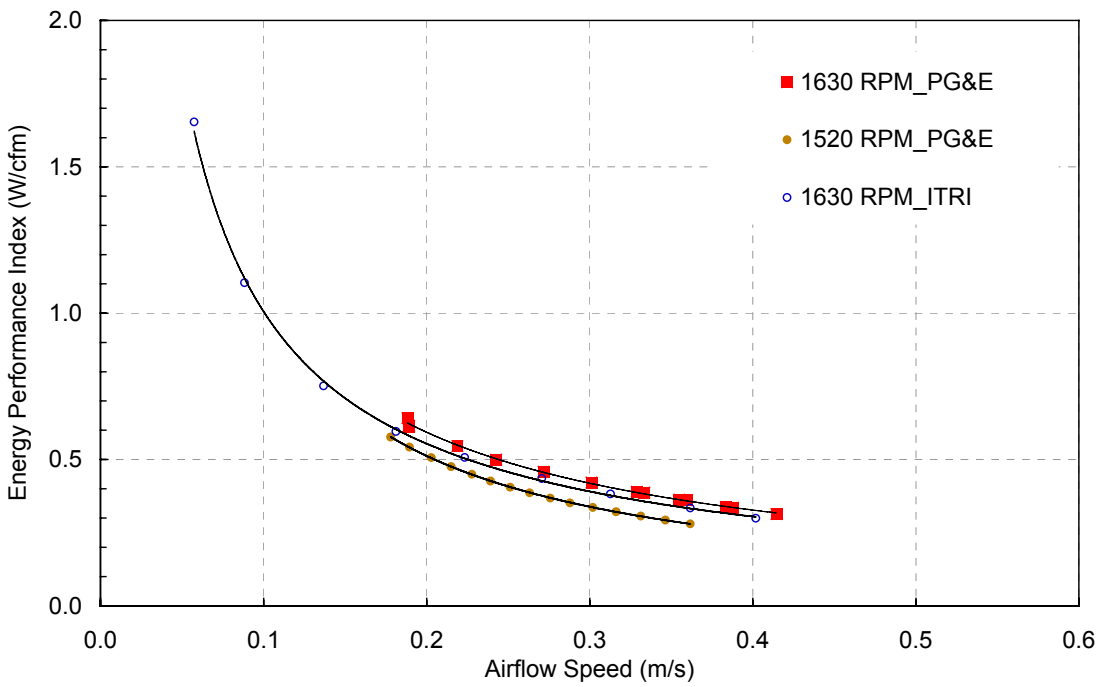

Figure 7 Energy Performance Index vs. Airflow Speed at FFU Exit

5.4 Conclusions and Recommendations

Laboratory testing of FFU energy performance can provide useful data for suppliers and end users to understand the performance of FFU products. Based upon this draft standard test method, it's feasible for users and suppliers to obtain energy performance in a consistent way. LBNL has completed the phase-one demonstration project at PG&E's testing facility.

In addition, in collaboration with ITRI, LBNL also completed an extra research and demonstration project at ITRI's facility. The additional demonstration was intended to start initial examination of the variance such as different configurations of test rigs and device. In both demonstrations, the same fan-filter unit was tested using the draft standard test method developed.

Demonstration testing at both facilities followed the options described in draft standard method in order to consistently collect the same performance parameters and to report the results using the same metrics. The measured parameters and calculated metrics included pressure-rise across the unit, energy performance index, and total pressure efficiency, under various operating conditions. Comparisons of the measured metrics calculated based upon the tests within each facility were made, and they demonstrated clear correlation between energy performance, airflow rates, and pressure rise for the unit under different conditions. The tests demonstrated that users were able to follow the protocols developed in the draft standard method and obtain comparable results.

Because of the constraints at PG&E and ITRI facilities (e.g., different power supplies, limitation in testing space and rig configuration), the experimental data obtained from both of the facilities, however, did not appear to be sufficient to address the issue on repeatability between the two facilities. The effect of test rig configurations and device will need to be determined through additional investigations.

A second-phase of demonstration is necessary to understand technical issues associated with configurations of testing rig and device, to refine the draft standard method, and to test more units to establish the baseline performance for energy-incentive program. Through the additional research and demonstration tests in collaboration with ITRI, a number of technical issues are initially identified and are recommended for future investigations. The following enlists some of the important recommendations for future investigations associated with the test method, test set up, and device:

- Examine the significance of difference between FFUs to be place in horizontal position vs. vertical position
- Examine the significance of difference between FFUs to be placed at downstream vs. upstream of the airflows
- Examine the significance of difference between FFUs with various power supplies (voltage, frequency, etc.)
- Examine the significance of difference between FFUs using various speed-controllers (constant-speed, variable speed drive, etc.)
- Examine the significance of difference between FFUs
- Examine the sensitivity of measured results in airflow rates and pressure-rise using various flow measurement methods and measurement locations
- Investigate the uniformity of airflows from a FFU
- Investigate the performance effect of various FFU design, filter type, etc.

6. Acknowledgements

The project is funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program (http://www.energy.ca.gov/). This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of

Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. The author would like to extend thanks to Robert Davis and Jim Yang of PG&E, and Dr. Chao-Ho Lan of ITRI for providing assistance and comments in the research and demonstration.

7. References

[1] Xu, T. 2004. Laboratory Method of Testing Energy Performance of Fan-Filter Units, Version 1.2 (2004, public reviewed). Lawrence Berkeley National Laboratory Report, LBNL-54626, Berkeley, California.

[2] Xu, T. 2005. Laboratory Method of Testing Energy Performance of Fan-Filter Units, Version 1.3 (2005, refined upon additional reviews). Lawrence Berkeley National Laboratory Report, LBNL-54626, Berkeley, California.

[3] Xu, T. and M. Jeng. 2004. Laboratory Evaluation of Fan-filter Units' Aerodynamic and Energy Performance. Journal of the IEST, Volume 47: 116-120. The Institute of Environmental Sciences and Technology (IEST), 5005 Newport Drive, Rolling Meadows, Illinois, October 2004. Lawrence Berkeley National Laboratory Report, LBNL-54250, Berkeley, Calif.

[4] Gale, S. 2004. FFUs: Setting a course for energy efficiency, Special Report, Cleanroom Magazine, September 2004, http://cr.pennnet.com/Articles/Article_Display.cfm?Section=Articles&Subsection=Display&Article_ID=211810

[5] Xu, T. 2004. Introducing a Standard Testing Method for FFUs, Controlled Environments (Formerly A2C2), http://www.cemag.us/articles.asp?pid=469, September 2004.

[6] M. Jeng, T. Xu, C. Lan. 2004. Toward Green Systems for Cleanrooms: Energy Efficient Fan-Filter Units, Proceedings of SEMI Technical Symposium: Innovations in Semiconductor Manufacturing (STS: ISM), SEMICONÆ West 2004, San Francisco, CA, Semiconductor Equipment and Materials International, July 2004. LBNL 55039.

[7] Xu, T. 2004. Developing a Standard Method For Evaluating Fan-Filter Unit Performance – Applications in Utility Incentive Programs. Poster distributed to three major California utilities, California Energy Commission, and other stakeholders.

Attachment III: Investigation of Demand Controlled Ventilation

LBNL – 63420

Demand Controlled Filtration in an Industrial Cleanroom

David Faulkner, Dennis DiBartolomeo, Duo Wang

Environmental Energy Technologies Division
Indoor Environment Department
Lawrence Berkeley National Laboratories
Berkeley, CA 94720

September 2007

This work is governed under U.S. Department of Energy Contract No. DE-AC03-76SF00098 with the Regents of the University of California and funded by the California Energy Commission's Public Interest Energy Research (PIER) Program through agreement with the Regents of the University of California administered by the California Institute for Energy and Environment (CIEE) (under Memorandum Agreement NO. MR-0312).

Demand Controlled Filtration in an Industrial Cleanroom

David Faulkner, Dennis DiBartolomeo, Duo Wang

Abstract

In an industrial cleanroom, significant energy savings were realized by implementing two types of demand controlled filtration (DCF) strategies, one based on particle counts and one on occupancy. With each strategy the speed of the recirculation fan filter units was reduced to save energy. When the control was based on particle counts, the energy use was 60% of the baseline configuration of continuous fan operation. With simple occupancy sensors, the energy usage was 63% of the baseline configuration. During the testing of DCF, no complaints were registered by the operator of the cleanroom concerning processes and products being affected by the DCF implementation.

Introduction

Demand controlled filtration (DCF) is a method of saving energy in cleanrooms by decreasing the recirculation fan speed when demand for filtration is low. A particle counter or an occupancy sensor can measure the demand for filtration. Another method of saving energy in cleanrooms is the use of programmable timers that decrease recirculation fan speeds when occupants are not expected, such as nights and weekends.

DCF should not be confused with demand controlled ventilation, which is the control of outside air entering a building based on the demand due to occupancy and/or processes occurring in the building.

In a past study, we found that industrial cleanroom users were hesitant to implement DCF, often because of concern that large changes in the recirculation fan speeds would produce a burst of particles discharged into the cleanroom. We have not seen this phenomenon in our previous or current study. Another concern is that for large manufacturing facilities, the potential loss of product outweighs any potential energy savings. There are many small research or industrial facilities that are not as product-oriented that can benefit from the potential energy savings from implementing DCF.

In the past few years there has been gathering interest in saving energy in cleanrooms and some businesses are experimenting with DCF. There is a large Southern California manufacturer with Class 100 (ISO Class 5) and 10,000 (ISO Class 7) cleanrooms covering about 15,000 ft^2. During non-working hours, the make-up and recirculation air handler fans at this facility are turned down to recirculate less air. A reduction in fan speed of about 30% at night, when the space is unoccupied, reduces the measured recirculation air handler power consumption from 32 kW to less than 9kW – a reduction of 75%.

Another large manufacturing facility in Southern California is in the process of implementing active, real-time DCF using particle counters. The facility manager has received rebates from the local utility for implementing DCF. The particle counters have been purchased and preliminary testing has been done. The next steps are installing the particle counters and implementing the control routines.

Finally, a large East Coast university recently modernized a building with cleanrooms. In the process they installed particle counters in every cleanroom area and plan to implement DCF.

Energy savings potential

The potential for energy savings is tremendous as fan energy has an approximately cubic relationship with fan speed. Thus, even a small reduction in fan speed can translate into large energy savings. Cleanrooms employing DCF can achieve lasting savings by implementing active control of the recirculation fans based on inputs from either a particle counter or an occupancy sensor.

Also, cleanroom operators may be able to participate in demand response programs offered by utilities. During high electricity demand days, such as hot summer days, many utilities offer incentives to their customers, such as rebates or price reductions, to voluntarily reduce electricity demand at their facilities. Typically these programs ask customers to reduce electricity consumption for periods of 1-6 hours. As long as the cleanroom fans are equipped with variable frequency drives (VFD), then the recirculation fan speeds can be adjusted lower during critical demand days. A 10-20% reduction in recirculation fan speed can translate to a reduction in energy use of over 40%.

Methods

Facility

Experiments were conducted to evaluate air quality in an industrial cleanroom in the San Francisco Bay Area while DCF was implemented. The cleanroom that was studied was one of a suite of four adjacent cleanrooms, each with different classifications from Class 10 (ISO Class 4) to Class 10,000 (ISO Class 7). The suite of cleanrooms was designed so that the inner Class 10 (ISO Class 4) cleanroom (Group-5) is at a higher pressure than the adjacent Class 100 (ISO Class 5) cleanroom (Group-6), which in turn is at higher pressure than the adjacent Class 1,000 (ISO Class 6) cleanroom (Group-3), etc., see Figure 1. There are open grilles for air to flow from the Group-5 cleanroom to the Group-6 cleanroom to the Group-3 cleanroom and finally to the Group-2 cleanroom. The cleanroom in which the experiments were conducted was the Group-2 area and was a Class 10,000 (ISO Class 7).

The floor area of the entire suite is about 1600 ft^2 and the floor area of the Group-2, the experimental cleanroom, was 600 ft^2. The layout of Group-2 is shown in Figure 2. Each cleanroom had air recirculated by ceiling-mounted 2 ft x 4 ft fan filter units with HEPA filters. The Group-2 cleanroom had 16 fan filter units covering about 20% of the ceiling area. There were 4 hoods in the middle of the cleanroom as well as 2-3 flow hoods at each end of the cleanroom. Airflow through these hoods were not controlled as a part of our study.

Outside make-up air was controlled by a separate ventilation system and was not controlled in our study. Environmental conditions, temperature and humidity were controlled by a separate control and also were not part of our study.

Before our study, a laptop computer using a proprietary software program controlled the cleanroom fan filter units. The software allowed control of the fan filter units individually or in groups. This software was used to lower the fan speeds to 50% at night (22:00 to 6:00) and on weekends. This software and corresponding laptop would occasionally fail and the fan filter units would run at the speed of the last valid command; the fan filter units are designed to always run at the last valid command. Thus, occasionally the fan filter units for all of the cleanrooms would run for several days at 50% speed before the problem was discovered.

After implementing the night and weekend energy reduction, particle concentrations at 4 locations in the cleanroom suite were recorded continuously for background data. The particles concentrations were measured with MetOne 237B particle counters (Hach Ultra, Grants Pass, OR) capable of counting particles in 6 size bins from 0.3 to 5 microns at a sample flow rate of 0.1 cfm. Some particle counters were moved to different locations within the cleanroom suite during the background measurements. Based on this data and the desires of the facility manager, the experiments were performed in the Group-2 cleanroom which was the Class 10,000 (ISO Class 7) cleanroom.

After background particle measurements were recorded, a particle counter was placed in the study cleanroom (Group-2) near the center of the room on a countertop about 3 ft above the floor. To implement DCF, the particle counter was connected to a laptop computer running an off-the-shelf software program, LabVIEW Version 5.1 (National Instruments). Control routines in LabVIEW had been developed to record particle counts, and use the particle count information to decide whether to increase or decrease fan speeds. The particle counter sent counts to the laptop every 30 seconds and the particle counts in the 0.3 and 0.5 micron size bins were compared to established upper limits. The fan speed was either increased or decreased until either the maximum or minimum fan speed was reached. A record of the fan speed, particle counts and other parameters were recorded every 60 seconds in LabVIEW.

After implementing DCF with the particle counter, we implemented DCF with control based on occupancy. The strategy was simple; if there were occupants detected, then the fan filter units were set to the programmed maximum speed, and after 30 minutes of no occupancy detected, the fans were set to 50% of programmed maximum speed.

There were 6 wireless infrared battery-operated occupancy sensors placed in the cleanroom. The 6 sensors were mounted on the walls so that each work area and traffic lane was covered by at least one occupancy sensor. A wireless signal was sent from the occupancy sensors to the laptop and the above strategy was implemented.

The data from the particle counter was converted to an equivalent ISO Class using the following equation (ISO 14644-1: 1999(E)):

$$C_n = 10^N \left(\frac{0.1}{D}\right)^{2.08} \qquad (1)$$

where,

C_n is the maximum permitted concentration (in particles per cubic meter of air) of airborne particles that are equal to or larger than the considered particle size. Value is rounded to the nearest whole number, using no more than three significant figures.

N is the ISO classification number, which shall not exceed a value of 9. Intermediate ISO classification numbers may be specified; with 0.1 the smallest permitted increment of N.

D is the considered particle size, in micrometers, which was 0.3 µm in this study.

0.1 is a constant, with a dimension of micrometers.

With algebraic manipulation equation (1) can be rearranged to solve for N, the ISO classification number:

$$N = \log_{10} C_n + 2.08(1 + \log_{10} D) \qquad (2)$$

Results

The data presented in Figures 3-6, show the ISO Class for the particle size of D = 0.3 µm when the Group-2 cleanroom was using DCF controlled by particle counts and by occupancy sensors. The concentration, C_n, is calculated from the counts measured by the particle counter.

During these experiments, the fan filter units were running at about 77% of maximum speed. Thus, in the following discussion, full speed means 77% of maximum speed and half speed means about 38% of maximum speed. Also, in all plots, the greatest fan speed will be 77%.

Figure 3 shows fan speed and particle counts in the Group-2 cleanroom (converted to ISO Class) versus time for a weekday while the fan speed was controlled by particle counts. From the graph, it is clear that the fans are actively being controlled by the particle counts most of the work day from about 6:00 to 17:00. The ISO Class stays around 3.5 to 4 during the working hours. At night the particle counts increase, despite there being no occupants in the cleanroom. The reason for this increase in particle counts while there were no occupants has not been explained by the building manager/cleanroom operator. Pressure measurements made between cleanrooms showed that at night, some of the pressure differences between cleanrooms reduced in magnitude and some actually reversed direction. It was thought that this was possibly due to changes in the whole building HVAC system. The building manager did not have details concerning nighttime setback conditions in the whole building HVAC system. Each morning, as the whole building HVAC system reset and the cleanroom fan filter units increased, the cleanroom, returned to “clean” conditions very quickly.

Figure 4 shows the fan speed and particle count data for a weekday while the cleanroom had DCF control based on occupancy sensors. The fan filter units were either at full speed (77% of maximum) while occupants were detected or at “half” speed (~38% of

maximum) with no occupants detected in the last 30 minutes. After noon, it can be seen that no occupants were detected on 2 or 3 occasions. In response, the fan filter units reduced to half speed and the ISO Class (based on particle counts) increased rapidly from about 3 to 5. But each time occupants were detected, the particle counts decreased quickly.

The rapid response in particle counts to the change in fan filter speed is seen in Figure 5, which shows similar data as in Figure 4, but over a shorter time span. The ISO Class changes from about 5 to about 3 in approximately 3-4 minutes as the fan speed increased from half speed to full speed.

Finally, Figure 6 shows the fan speed and particle count for a weekend day with DCF control based on occupancy sensors. As soon as an occupant was detected around 9:00, the fan speed increased to full speed and the particle counts decreased by about a factor of 300 (ISO Class 6 to Class 3.5). Around 10:10, an absence of occupants was detected and fan speeds were set to "half" speed and the particle counts increased. About 10 minutes later, an occupant was detected and the fan speeds increased to "full" speed and particle counts again decrease, as indicated by the ISO Class.

Energy values are shown in Table 1, for a typical week. There are 5 different configurations that are compared. Two of the configurations are used as baselines for comparison with various energy saving configurations. The baselines are 24/7, which means the fans are running all the time, and the other baseline is the night and weekend setback, in which the fans were reduced at night and on the weekends. Comparing the two baselines, the night and weekend setback saves about 28% in energy as compared to the fans running constantly.

The use of the particle counter to implement DCF saved more than the occupancy sensors, but not by much, 40% vs 36% as compared to 24/7 as a baseline and 17% vs. 12% as compared to the night and weekend setback as a baseline. Thus, using a few occupancy sensors at $30 each is less expensive than a particle counter that may cost $3,000 or more.

Table 1. Energy use (kW-h per week) and percent savings due to demand controlled filtration.

Configuration	**kW-h per week**	**Compared to 24/7**	**Compared to Night & Weekend Setback**
24/7	151		
Night & Weekend Setback	109	28%	
Occupancy Sensor	96	36%	12%
Particle Counter	91	40%	17%
"Half" Speed	70	54%	36%

Conclusions

The data collected indicate two major conclusions. First, DCF was implemented without disruption to the processes in the cleanroom. This implementation was relatively easy and inexpensive, less than $10,000 in equipment. Second, recovery of particle concentration from a "dirty" condition to a "clean" condition occurred in a matter of minutes. Thus, those cleanroom operators in industry need not be concerned about the removal of particles taking days or hours after a reduction in the fan speeds.

The energy savings realized were about 40% with either method as compared to running the cleanroom fans constantly, 24/7. Savings were about 15% with either method as compared to night-time and weekend setback.

Acknowledgements
Bill Tschudi and Mark Sippola of Lawrence Berkeley National Lab, Larry Chu and Peter Rumsey of Rumsey Engineers, and Reza Kazemi of Environmental Filtration, Inc.

References

Faulkner, D., Fisk, W.J., and Walton, J.T. (1996) Energy Savings in Cleanrooms from Demand-Controlled Filtration", Journal of the Institute of Environmental Sciences 39(2): 21-27. LBNL-38869

International Standard, Cleanrooms and associated controlled environments – Part 1: Classification of air cleanliness, ISO 14644-1: 1999(E)

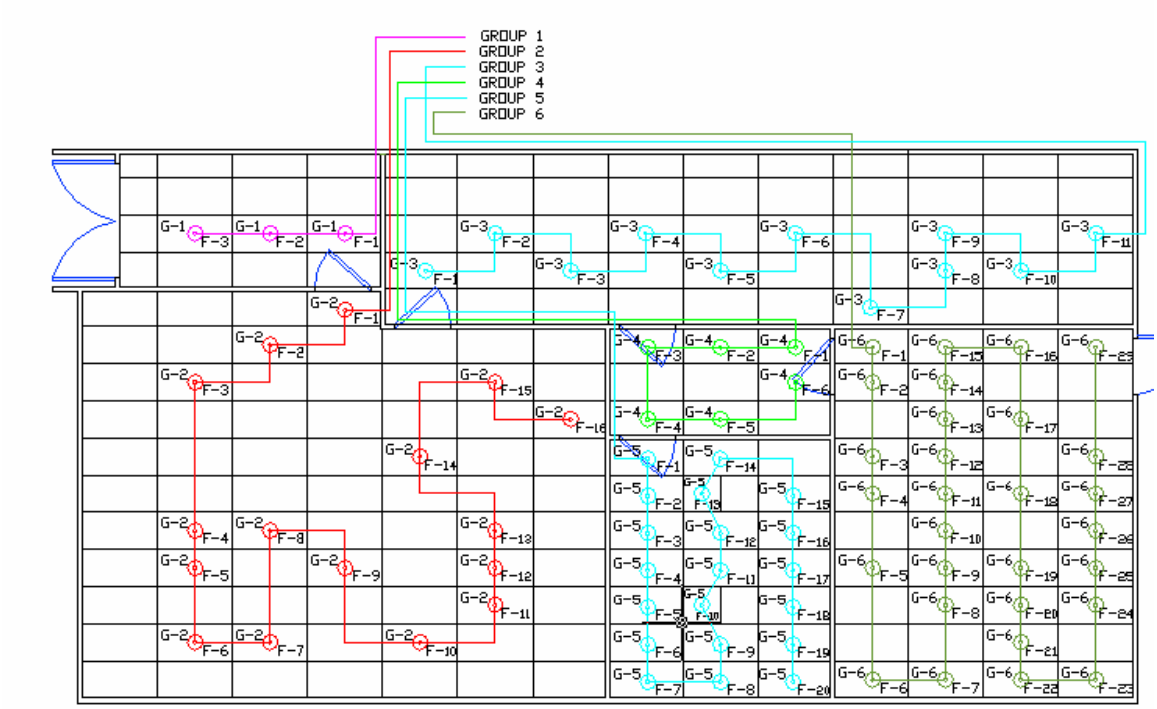

Figure 1. Layout of the suite of cleanrooms in the studied industrial facility. G-5 (Group-5) is an ISO Class 4 cleanroom. G-2 (Group-2) is an ISO Class 7 cleanroom. G-1 (Group-1) and G-4 (Group-4) are gowning rooms. G-3 (Group-3) is an ISO Class 6 cleanroom and G-6 (Group-6) is an ISO Class 5 cleanroom. Fan filter units are labeled as F-1, F-2, etc. Each rectangle is 2 x 4 ft.

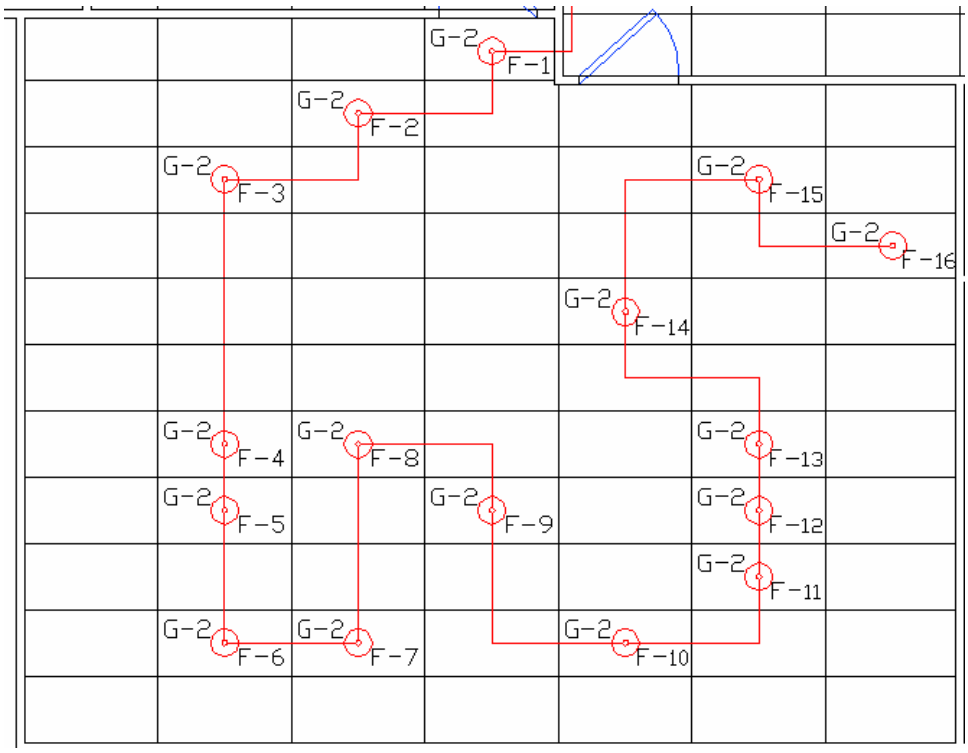

Figure 2. Layout of fan filter units in study cleanroom, Group-2. There are 16 fan filter units, shown as F-1 to F-16. Each rectangle is 2 x 4 ft.

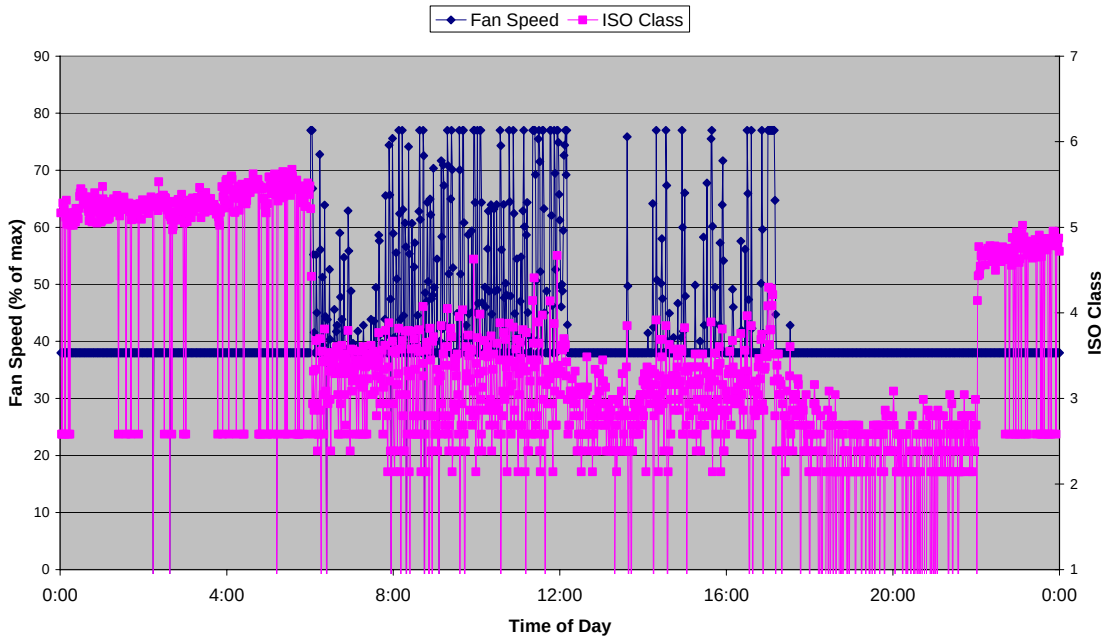

Figure 3. Fan speed and particle counts (converted to ISO Class) versus time for a weekday with DCF control based on particle counts. The cleanroom classification is based on the particle concentration at 0.3 micrometers.

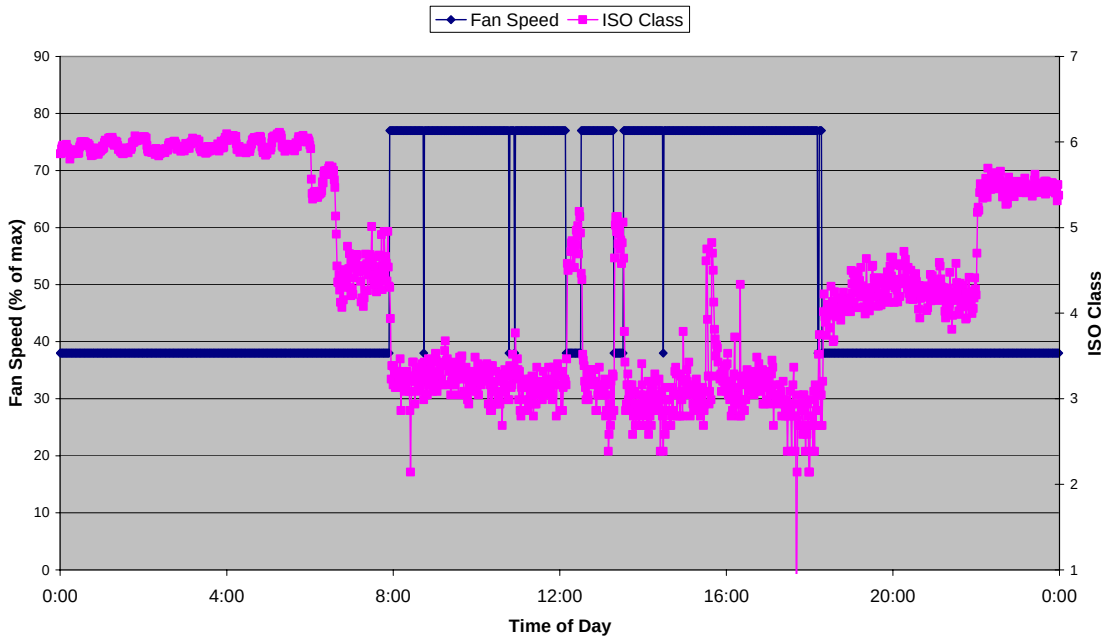

Figure 4. Fan speed and particle counts (converted to ISO Class) versus time for a weekday with DCF control based on occupancy sensors. The cleanroom classification is based on the particle concentration at 0.3 micrometers.

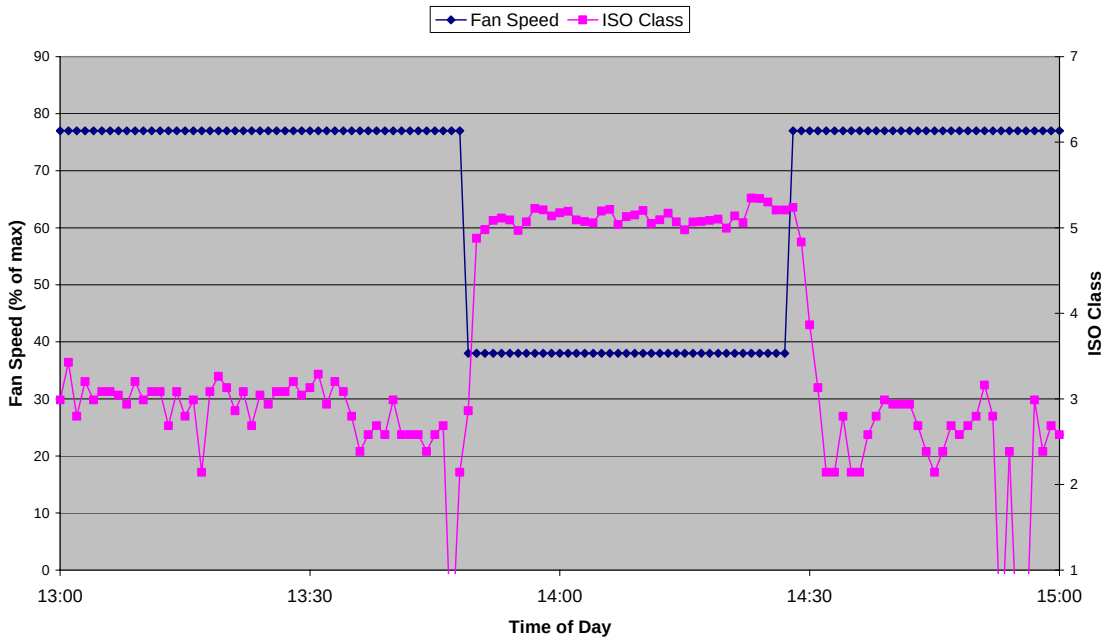

Figure 5. Fan speed and particle counts (converted to ISO Class) versus time for a weekday with DCF control. There is a rapid return to “clean” conditions within 3-4 minutes (data recorded every 1 minute) after the fans increase from a low speed to a high speed. The cleanroom classification is based on the particle concentration at 0.3 micrometers.

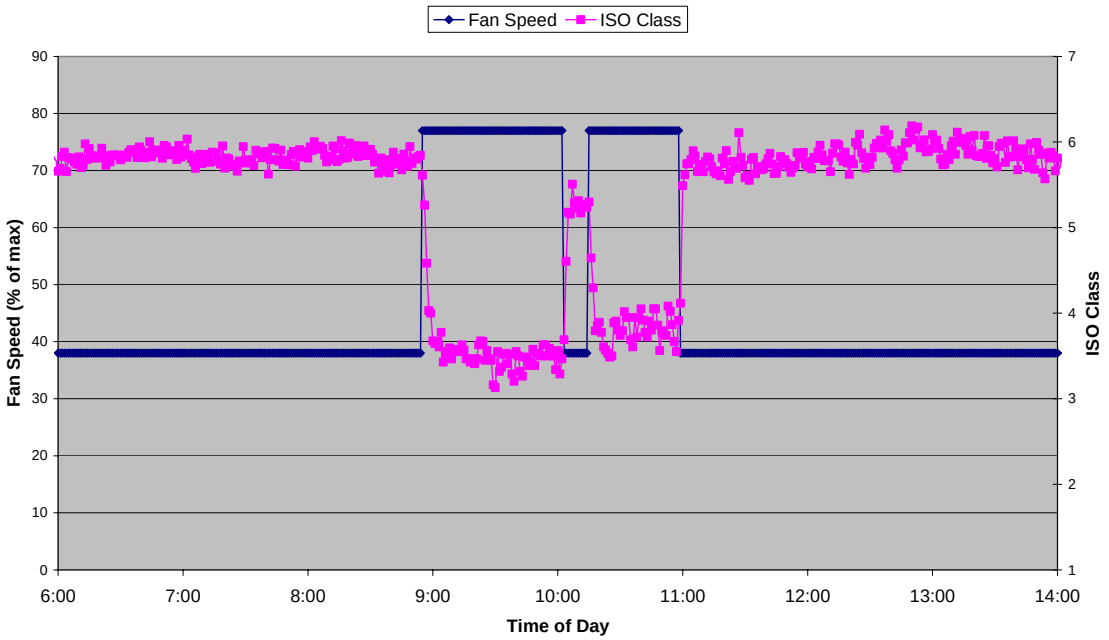

Figure 6. Fan speed and particle counts (converted to ISO Class) for a weekend day with DCF control based on occupancy sensors. Changes in particle concentrations, as expressed by the ISO Class, are very quick. The cleanroom classification is based on the particle concentration at 0.3 micrometers.

Attachment IV: Investigation of Mini-environments

Lawrence Berkeley National Laboratory Report LBNL-56312

Opportunities for Sustainable Design and Operation of Clean Spaces: A Case Study on Minienvironment System Performance

Final Report, LBNL-56312

Tengfang Xu

Friday, May 20, 2005

1. Background

In order to identify and pursue energy efficiency opportunities associated with cleanrooms, it is necessary to understand the design and operation of cleanroom systems for specific contamination control requirements. With the industrial trend toward more stringent cleanliness class and tightening clean spaces, it is vital to understand the design of minienvironment and the operational performance of its systems. A good understanding of such system performance would help to identify opportunities in efficient energy end-use and wise allocation of resources associated with processes or productions that require minienvironments and cleanrooms. This report summarizes a case study on energy performance of a common minienvironment used in semiconductor industry, and discusses the opportunities in saving energy, in particular, the opportunities in achieving efficient operation and design that entails applications of minienvironments.

2. Introduction

A minienvironment is a localized environment created by an enclosure to isolate a product or process from the surrounding environment [1,2]. Minienvironments, often termed "Separative Devices," have been gaining popularity to provide effective isolation for critical contamination control. The purpose of using minienvironments is either to protect contamination-sensitive products or processes by isolating them from the ambient environment and workers, or to protect workers or their environment from exposures to hazardous contaminants by isolating the products or processes, or both. Serving similar purposes, a minienvironment carries a variety of names in different industries and applications, in which materials or processes must be protected from the surrounding environment, or operators must be protected from the activities within the minienvironment. For instance, they are called "gloveboxes" in research laboratories and the defense industry; they are also labeled as "isolators," "separative enclosures or devices," or "barriers" in the microelectronics, chemical, and pharmaceutical industries; in addition, they are commonly named "safety cabinets" in biomedical and healthcare industry. Minienvironments can often introduce filtered air through HEPA or ULPA filters at a high airflow speed (e.g., 90 fpm) in order to achieve the desired pressure difference or unidirectional airflows to maintain specific levels of cleanliness and contamination control [3]. Depending on the actual height of minienvironment spaces, air change rates of the supplied air can be much higher than the air change rates of recirculation air in common cleanrooms that are designed to achieve similar cleanliness classification.

3. Minienvironment Design and Operation

3.1 Minienvironment Design

Minienvironments are designed to have an isolation enclosure, typically operating within a conventional clean space, to provide well-controlled environmental conditions where it is needed for specific processes or activities. Minienvironments are normally smaller clean environments than conventional cleanrooms and may be contained within cleanrooms. The use of minienvironments can provide several orders of magnitude improvement in particulate cleanliness levels. In the semiconductor industry, more and more minienvironments have, in fact, become a requirement as semiconductor device feature sizes are getting smaller, and the shift from manual product movement to fully automated wafer movement are becoming common practice - considering benefits of minienvironment Fab designs over conventional Fab designs. A minienvironment can produce very clean environmental conditions or "mini-atmospheres", and present physical barriers to protect workers from hazardous materials, or vice versa, to protect process from contamination from the occupants. A common situation is that personnel manipulate tools, processes, and products inside the minienvironment with access devices, which can be as diverse as manual glove systems or automatic robotics systems for handling or transferring products. Overall, the advantages in using minienvironments are to allow better contamination control and process integration, to allow cleanliness-class upgrade required for certain process, to improve safety, and to potentially save costs. In the meanwhile, energy intensity may be shifted from the conventional cleanroom systems to the minienvironments that enclose the specific process. On the other hand, in using minienvironments, energy intensity may be shifted from conventional cleanroom systems to the minienvironments that enclose the specific process. On the other hand, the use of minienvironments requires careful and integrated planning that should consider safety, production (yield), ergonomic requirements, and overall production efficiency. This project addresses the energy performance of such devices and overall impact on energy usage of applying minienvironments in cleanrooms.

3.2 Minienvironment Operation

Minienvironments typically introduce large quantities of filtered air through HEPA or ULPA filters (e.g., at an airflow speed of 90 fpm) in order to achieve the desired pressure difference or unidirectional airflows to maintain required levels of cleanliness and contamination control. Depending on the "height" of minienvironment spaces, the air change rates of the supplied air can be much higher than the air change rates of recirculation air in common cleanrooms that are designed to achieve similar cleanliness classification.

Anecdotal industry experience indicates that in some situations, the design and operation of the overall cleanroom might well remain largely unchanged, and that minienvironments (or isolated spaces) are simply adding another set of air movement and air conditioning, requiring more energy to operate. On the other hand, the potential for energy savings could be achieved if the cleanliness levels are correspondingly relaxed in the surrounding cleanroom space. Although there are papers and guidelines addressing minienvironments' design, construction, and operation [4,5,6,7,8,9,10], and yields and production associated with deploying minienvironments, there is nonetheless virtually no data available to quantify the energy efficiency of minienvironment systems. To understand actual energy implications of a

minienvironment system, it is necessary to investigate energy performance of a typical minienvironment, and understand its potential effects on energy end-use.

4. Case Study

4.1 Purpose and Scope

The main purposes for this case study were to develop an understanding of the key parameters of minienvironment design, operation, and control, to investigate energy performance of the minienvironment air system, and to investigate opportunities in improving its energy performance. The focus of the study was on an air system associated with minienvironment space, as distinct from process equipment used for product manufacturing.

This case study focused on energy performance and electric power usage of the air delivery system of a typical minienvironment in a ballroom setting. The report includes analysis of measured energy performance of a minienvironment' s air system, comparisons of the energy performance of the minienvironment studied with that of a cleanroom, and suggestions of energy savings opportunities associated with the use of minienvironment.

4.2 Approaches and Equipment Setup

The case study was designed to measure airflow rates, electric power usage, and air pressures in the minienvironment under various operating conditions. The conditions measured covered the full range of operating points (airflow delivery) that the minienvironment's air system could handle. The key parameters included the following: electric power usage, airflow and air change rate, pressure difference between the space inside the minienvironment and the space surrounding the minienvironment, and energy performance index (EPI). Electric power usage and power factors of the minienvironment air system were measured concurrently with airflow rates and pressures under the range of testing conditions. Analysis was then performed to investigate the correlations among power usage, airflow, pressure control, EPI, power factor, and the size of exhaust opening.

4.2.1 Electric Power Measurement

The power meter used in this study was a true RMS energy analyzer with an uncertainty of ±3% [11]. The meter records the electric current, voltage, power factor, and actual power supplied to air delivery system for the minienvironment. The power meter was used with 0.01-amp to 10-amp current transducers (HA100 Current Probe, uncertainty ±2%) and voltage transducers was used to measure the electric current, voltage, power factor, and actual power supplied to air delivery system for the minienvironment. Parallel current and voltage transducers were also installed to measure concurrent electric power supplied to the motors of fan-filter units. The power meter used in this study was a true RMS energy analyzer (Powersight, uncertainty ±3%).

4.2.2 Airflow and Pressure Measurement

A VelGrid attached to the electronic micro-manometer [12] measured the average speeds of the airflow delivered out of the face of the fan-filter units (FFUs), which were installed at the ceiling

of the minienvironment. The size of individual FFU and HEPA filters was 1 ft by 2 ft. The measurement uncertainty in airflow speeds was ± 3% of reading plus ± 7 fpm from 50 fpm to 2500 fpm. Pressures were measured using a Pitot tube, with a measurement uncertainty of ±2% of reading plus 0.001-inch water column (0.25 Pa) from 0.05-inch water column to 50.00-inch water column (or 0.125 Pa to 12500 Pa). The VelGrid samples 16 points over a 1 ft x1 ft area to determine average airflow speeds. Airflow speed-readings were automatically corrected for the density effect of barometric pressure and temperature. Readings were displayed as local density and true air speeds.

4.2.3 Air Leaks

In normal operation, the accurate control of airflows and pressure difference between minienvironment and the surrounding spaces was realized through accurate system information and controlling exhaust openings. To investigate the relationship among power consumption, airflow, pressure difference, and system efficiency, it was necessary to eliminate unintentional air leaks from the minienvironment and its air system. In this study, the minienvironment enclosure was sealed carefully to avoid leaks and to prevent uncontrollable airflow from the process bay to the chase, or vice versa. During the experiment, we controlled and adjusted the size of opening in the front side of the minienvironment barrier.

4.2.4 Setup Diagrams

Figure 1 Minienvironment

Figure 2 Power measurement setup

4.2.5 Measurements and Data Analysis

The exhaust opening of about 20% of the total FFU surfaces or floor area was initially set while we concurrently measured the maximal airflow rates through FFUs, electric power (Watts) supplied to the FFU speed controller, and the pressure difference between the space inside the minienvironment and the space surrounding the minienvironment. At the 20% exhaust opening, airflow rates through FFUs were then adjusted from zero up to the maximal value so that performance curves can be developed to represent the whole range of actual operation conditions. The inter-correlations among power usage, airflow, pressure control, EPI, power factor can then be analyzed. In addition, a set of parallel measurements was taken for various sizes of exhaust opening. The analysis also includes an examination of the correlation of energy performance indices and other performance metrics of the minienvironment air system with the relative exhaust opening.

4.3 Results

4.3.1 Characteristics of the Minienvironment

The minienvironment in this study was a stand-alone open-loop system, with airflow supplied from the surrounding cleanroom space (Figure 1). The supplied air was filtered through four FFUs, each of which was 1 ft by 2 ft with a depth of two feet. The floor size of the minienvironment was 2 ft by 4 ft with an inner space height of seven feet and seven inches. The supply air was from the top of the minienvironment and there was an opening in the front toward the bottom. The FFUs were integrated with a unidirectional flow shield that was tied to the HEPAs. There was no air recirculation path within the minienvironment; therefore it was an open-loop minienvironment system. The outgoing airflow path allowed uncontrolled mixing with the external environment. In normal operation, a pod was attached to the minienvironment. Such a pod was essentially a box containing a cassette of wafers used in conjunction with a standard mechanical interface (SMIF). The front-open unified pod (FOUP) I/O device was a material handling unit and isolated the cassette of wafers while maintaining the integrity of the wafer environment. The integrated pod transported wafers in and out of the minienvironment without exposure to the surrounding environment.

Four identical 1 ft by 2 ft fan-filter unit (FFU) were used in the minienvironment' s air system. Each of the FFUs was designed with a single-phase AC motor with adjustable airflow rates or airflow speeds controlled by a Silicon Controlled Rectifier (SCR) controller.

4.3.2 Fan-speed Controller

Reducing the operating airflow speeds not only can save FFU fan power, but also may lower noise and be beneficial to the operating life of the fan. Normally, the fan speeds could be controlled manually or by sensing the air pressure in the minienvironment enclosure and reducing the fan supply voltage. In this study, fan speeds were controlled manually to measure the key parameters for a range of operating conditions. The fan integrated in the fan-filter units

was powered by a single-phase AC motor. The fan speed was controlled via a silicon-controlled rectifier (SCR).

Because the purpose was to develop an understanding of the key parameters of minienvironment design, operation, and control, to investigate energy performance of the minienvironment air system, the supply airflow rates into the minienvironment and the pressure difference were controlled to represent various operating conditions commonly observed in the minienvironment air systems. The concurrent power consumption of the minienvironment air delivery system was measured. The data obtained can be used to develop performance curves representing relationships of power, efficiency, airflow rate, and pressures difference. The range of airflow speeds was from zero to 110 fpm, with the range of pressure difference being zero to 0.3-inch water column. In pharmaceutical applications, the minimal airflow speed is typical set at 90 fpm.

4.3.3 Electric Power and Airflow Rates

Reducing the operating airflow speed not only can reduce FFU fan power, but also may improve cleanliness, lower noise, and improve operating life of the fan. Normally one would expect fan power consumption to increase with an increase in airflow rates. Figure 3 shows that when the airflow speed was less than 95 fpm, total electric power supplied to the FFU increases with the increase in airflow rates. In addition, the rate of electric power increase with the airflow rate goes down with airflows when the airflow speed was below 95 fpm (or 760 fpm), at which the total electric power input reaches to a peak. In contrast, when the airflow speed was above 95 fpm, the total electric power decreases with the increase in airflow rate. This indicates that it takes less fan power for the minienvironment air system to run at a higher airflow rate than it does at a lower airflow rate.

The trends observed in the figure also confirm that with this speed controller, once the initial resistance was overcome, the air delivery becomes easier (and therefore, more efficient) for the system to move the same airflow rate through the air system.

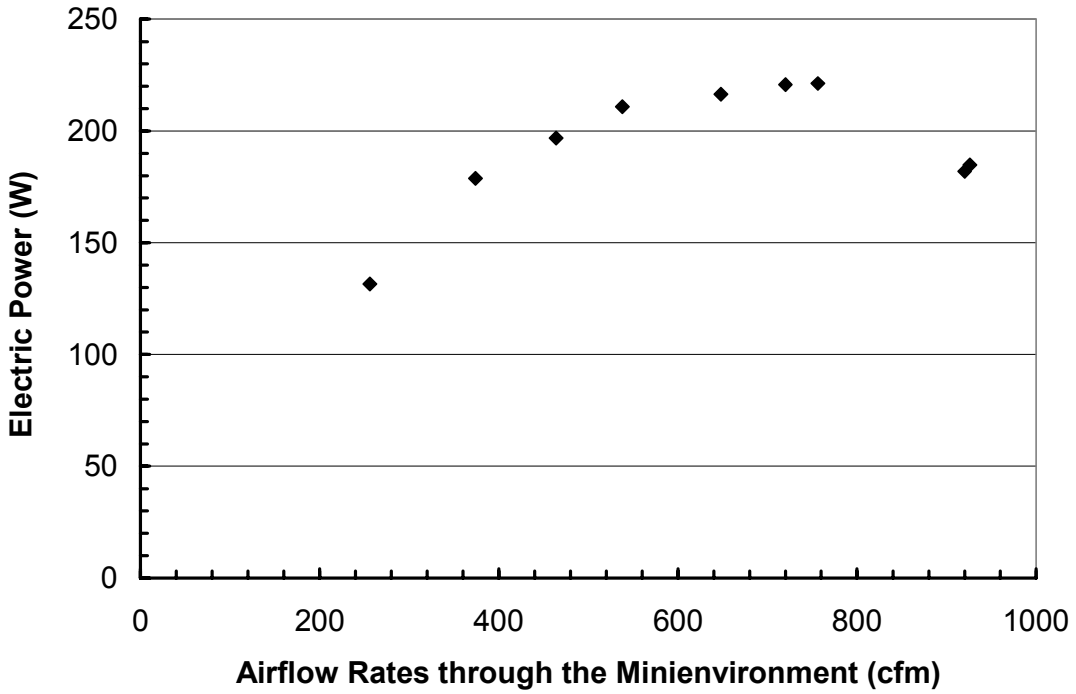

Figure 3 Electric Power and Airflow Rates

4.3.4 Power Factor and Airflow Rate

Associated with the true power consumption, power factor was another metric affecting the power efficiency of the air delivery system. Similar to true power consumption, the figure shows that when the airflow speed was lower than 95 fpm, the power factor increased with CFM in a trend that was close to linear. When the airflow speed was around 95 fpm, the power factor reached a peak at around 0.74. Interestingly, when the airflow speed was higher than 95 fpm, the power factor decreased with the increase in airflow rates. This indicates that when the airflow speeds converged toward 95 fpm, the fan motor power system for air delivery was more efficient in delivering air. The lower end of the power factor was about 0.60, when the airflow speed was about 30 fpm or 110 fpm. At higher airflow speeds (i.e., 110 fpm), the power factor decreased accordingly.

4.3.5 Energy Performance Index

In this study, the energy performance index (EPI) of a minienvironment air system is defined as the total electric power supplied to the fan system divided by the flowrate of the delivered air. A higher EPI means more power is needed for the same airflow rates supplied to and through the minienvironment, corresponding to lower air delivery efficiency in the minienvironment. Figure 4 shows the results in air system's energy performance index, with the EPI ranging from 0.20 W/cfm to 0.42 W/cfm corresponding to the range of airflow speeds from approximately 60 fpm

to 110 fpm. This was within or lower than the overall benchmarked range observed in many large cleanrooms (ISO Class 4 or Class 5) [13]. The re-circulation air system efficiency for ISO Class 4 and 5 cleanrooms ranged from approximately 1,100 cfm/kW to 10,500 cfm/kW, corresponding to the approximate range of EPI values of 0.10 W/cfm to 0.90 W/cfm for all recirculation air systems. Compared to the FFU systems in cleanrooms with ISO Cleanliness Class 5 or lower cleanliness classes, the energy performance index of the minienvironment system appeared to be higher, indicating a less energy-efficient air system in the minienvironment. This may suggest opportunities to improve its air systems' delivery efficiency.

The measured airflow speeds corresponded to airflow rates in the range of approximately 460 to 900 cfm, and a positive air pressure inside the minienvironment in the range of 0.01-inch water column to 0.03-inch water column (or 2.5 Pa to 7.5 Pa). By controlling the airflow, a positive pressure was created to prevent introduction of potential contaminants from the surrounding environment. For common airflow speeds ranging from 50 fpm to 90 fpm, the measured EPI ranged from 0.30 W/cfm to 0.45 W/cfm.

In general, the EPI values decreased with the delivered airflow rates. The rate of the EPI decreasing was almost constant - indicating an almost linear correlation between EPI and airflow rates. The trend indicates that the air system EPI value became lower (more efficient in delivering the air) when the airflow rate through the minienvironment increases.

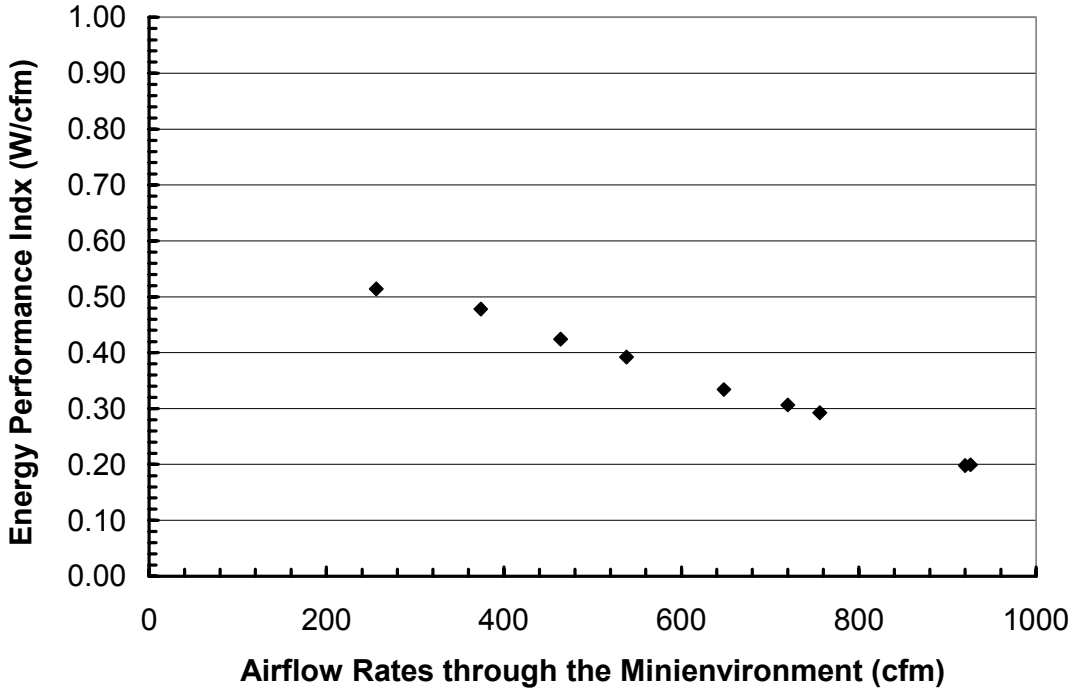

Figure 4 EPI and Airflow Rate

4.3.6 Pressure Difference

The air pressure difference was the difference between air pressure in the minienvironment's internal space and that of its ambient surrounding. The purpose of maintaining a positive air pressure in minienvironment relative to the air in the surrounding spaces was to prevent the less-clean air from being transported to the minienvironment and therefore contaminates the process.

According to IEST CC- RP 028.1 [1], microelectronic minienvironments spanning between process bay and services chase should be designed to maintain a differential pressure, with a typical process-bay pressure exceeding the service-chase pressure by 0.01-inch water column to 0.05-inch water column (or 2.5 Pa to 12.5 Pa). A rule of thumb in the industry is to commonly control the pressure differential with a minimal value of 0.01-inch water column (2.5 Pa) up to 0.03-inch water column (7.5 Pa). However, the ranges seem to be experiential and there is no scientific data to specifically support such ranges.

Figure 5 shows that as expected, pressure-differential increased with delivered airflow rates, and that the increase-rate of pressure-differential was almost constant. This indicates an almost linear correlation between airflow rate and pressure differential between the minienvironment and its surrounding space. A higher airflow produced higher pressure-differential. For example, corresponding to airflow speeds from 50 fpm to 90 fpm, the pressure differential ranged from 0.008-inch water column to 0.02-inch water column (2.0 Pa to 5.0 Pa); corresponding with airflow speeds from 60 fpm to 110 fpm, the pressure differential ranged from 0.01-inch water column to 0.03-inch water column (2.5 Pa to 7.5 Pa).

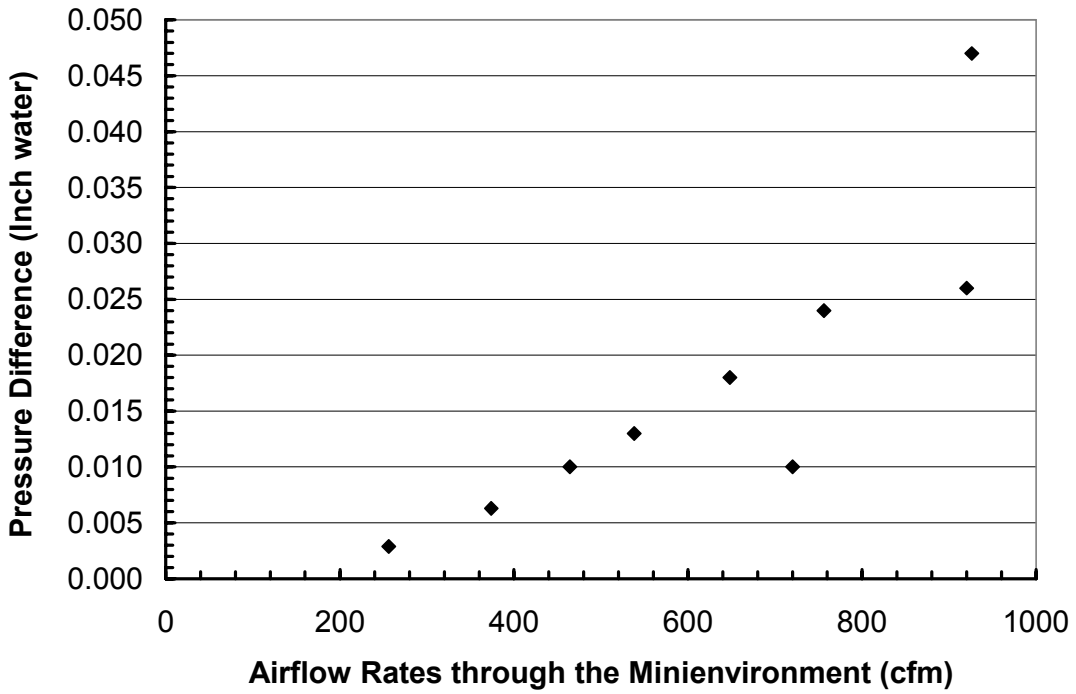

Figure 5 Pressure Difference

4.3.7 Electric Power Density

Figure 6 shows that electric power density changed with airflow speed and pressure differential. Corresponding to the tested operating ranges (30 fpm to 110 fpm) for this minienvironment, power density changed from 16.5 W/ft^2 to 23.0 W/ft^2, with a peak of 27.7 W/ ft^2 when the air speed was 95 fpm.

This range actually fell within the range of fan power density from previously measured ISO Cleanliness Class 4 cleanrooms that was in the range of 16 W/ft^2 to 38 W/ft^2 [13,14]. Given a same airflow speed in general, the power density of the minienvironment tended to be slightly higher than those of cleanrooms of similar cleanliness requirements, especially when the cleanrooms were not fully covered by HEPA filters.

Because of the much smaller minienvironment volume compared to that of full-scale cleanrooms (e.g., ballroom), the amount of airflow rate supplied to a minienvironment was significantly reduced. This result suggests opportunities for a significant overall energy savings potential in through

1) optimization of airflow speeds required in the minienvironment may reduce power consumption of minienvironment while satisfying contamination control; and

2) reduction of the total cleanroom airflow rate by introducing minienvironment due to the vastly smaller volumes of air that must be moved, conditioned, and filtered. The optimization of airflow speeds can be achieved through optimizing fan-filter unit efficiency and the airflow pathway design including exhaust opening.

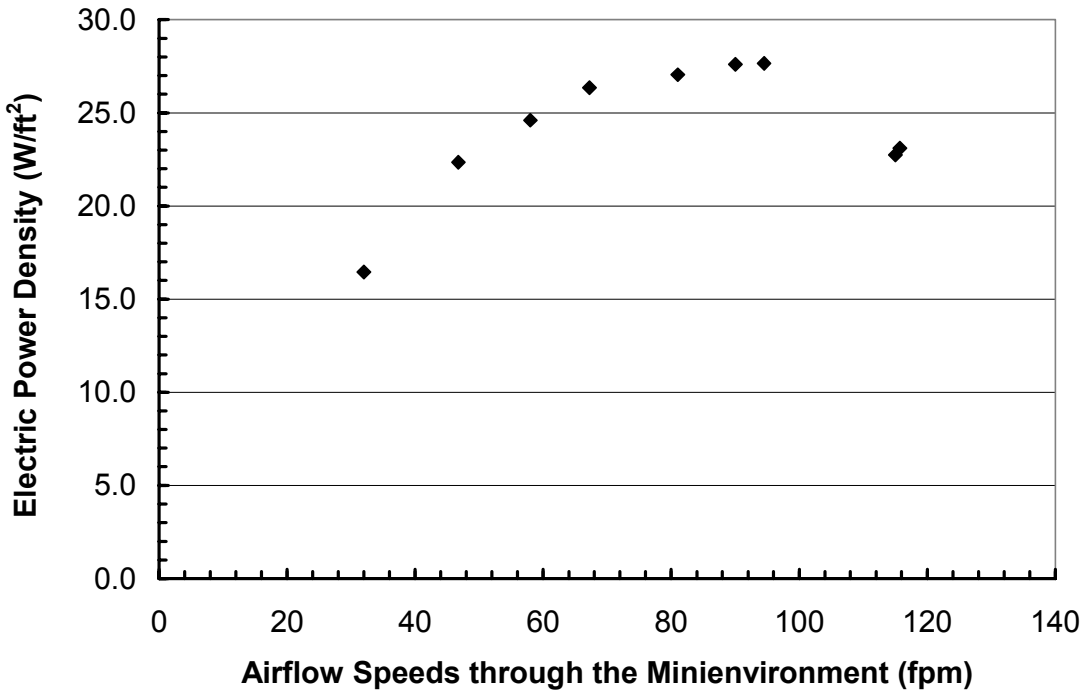

Figure 6 Power Density and Airflow Speeds

4.3.8 Air Change Rates

In semiconductor wafer manufacturing, the air supply for a large "ballroom" with cleanliness of ISO Cleanliness Class 4 or ISO Cleanliness Class 5 was filtered and recirculated at rates as high as 500- or 600- air changes per hour, while the wafer manufacturing only takes place in a relatively smaller area within the whole cleanroom space.

In this case study, the minienvironment typically operated with once-through airflow speeds in the range of 60 fpm and 100 fpm, which was consistent with the airflow speeds commonly

observed in conventional large clean spaces. The HEPA/ULPA filter coverage in the minienvironment was 100% while other cleanrooms can have a coverage ranging from 20% up to 100%. If we convert the airflows into actual air change rates for the minienvironment studied, the actual air change rates ranged from 480 to 800 air changes per hour corresponding to the airflow speeds ranging from 60 fpm to 100 fpm. The air change rate range of the minienvironment was higher than the range observed from those of ISO Cleanliness Class 4 cleanrooms, which were in the range of 385 to 680 air changes per hour corresponding to airflow speeds ranging from approximately 60 fpm to 120 fpm [13].

4.3.9 Exhaust Opening

Given that all other conditions were the same, one would expect that the size of air exhaust from minienvironment could affect the resistance in the airflow pathway, airflow patterns and distribution, and pressure differential between inside and the surrounding environment. It would therefore influence the overall operating performance of the minienvironment air system, including power consumption and energy performance index.

In this case study, the exhaust size of the minienvironment air system was controlled to follow a sequence of changes in terms of its relative size to the FFU coverage, or in this case - simply the minienvironment's floor area (i.e., percentage). The relative exhaust size was calculated as the ratio of actual exhaust size to the floor area of the minienvironment.

With the relative exhaust sizes at approximately 10%, 20%, 30%, and 85% of the minienvironment floor area, the power consumption, power factor, EPI, pressure differential, and airflow rates were measured for the maximal airflows for the minienvironment operating at each exhaust size. Among these parameters, power factor, EPI, pressure differential, and airflow rates were plotted against the relative exhaust opening (Figure 7). From the figure, we can see that the achievable maximal airflow rates went up very slightly when the relative exhaust opening increased from below 10% up to 85%. In the meanwhile, the power factors almost were maintained within a range without significant increasing or decreasing trend. The corresponding EPI values did not decrease significantly as the relative exhaust opening increased from 10% to 85%.

The pressure differential however exhibited a significant drop when the relative exhaust opening increased from below 10% to 20%. The pressure differential surprisingly increased when the relative exhaust opening increased to around 30%. This indicates that differential pressure ranges were largely influenced by the relative exhaust opening; therefore, selecting a certain exhaust opening could help in tuning the differential pressures for control purposes.

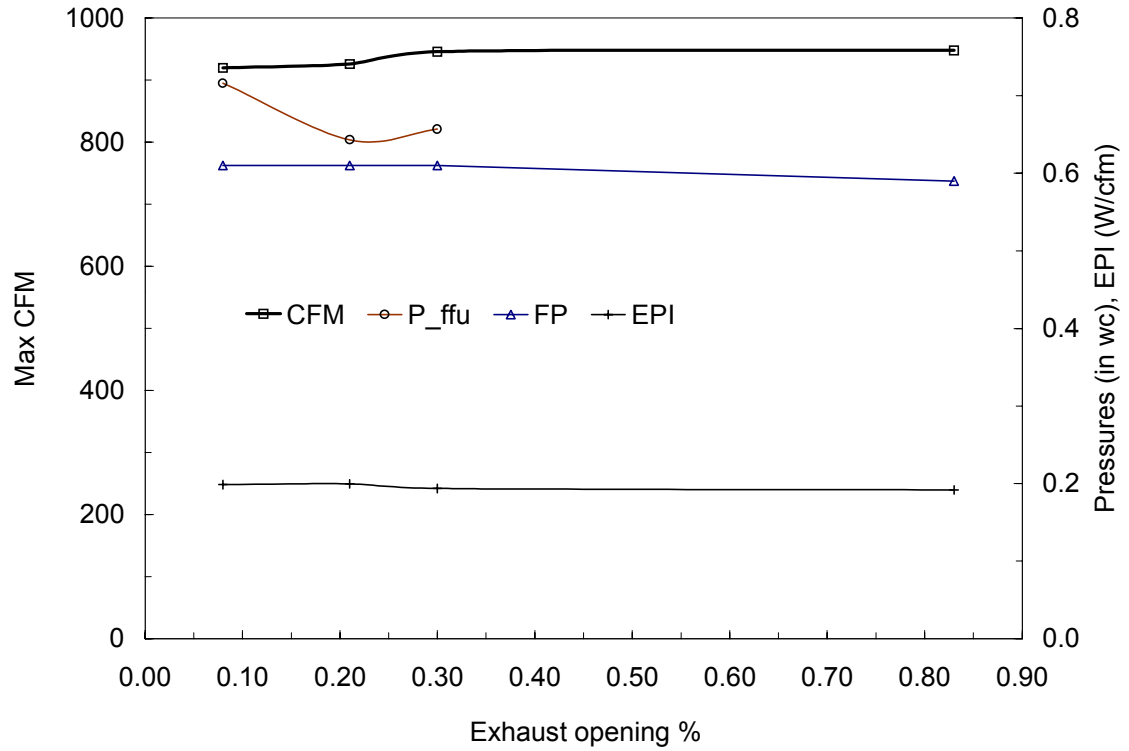

Figure 7 Performance of Minienvironment System as a Function of Relative Exhaust Opening

5. Conclusion and recommendations

Based upon the experimental observation and analysis in this case study, it can be concluded that minienvironment applications largely influence the future design, construction, and operation of cleanroom spaces, depending on their specific contamination control requirements. It is important to determine overall contamination requirements for types of clean spaces through specifying the right cleanliness and control. It is common that designers and users tend to associate higher cleanliness with higher airflows or pressure difference. However, a thorough understanding of contamination control requirements for specific activities is vital because the design, construction, and operation of clean spaces will largely influence the energy consumption as well as effectiveness in contamination control.

Providing measured data to quantify energy performance of the minienvironment, this study shows that the energy performance index of a minienvironment for typical operation tended to be in the vicinity of or higher than that of its counterparts in traditional cleanrooms. In addition, electric power density of the air system in such a minienvironment could be higher than that of normal cleanroom systems. Based upon the analysis, implementing minienvironments as a means of contamination control may produce overall savings in electric power consumption.

A new performance metric was developed in this case study - energy performance index that is based upon electric power usage per airflow rate to characterize the energy efficiency of airflow systems applicable to minienvironments. A lower energy performance index corresponds to a

more energy-efficient airflow delivery system. This case study concludes that the energy efficiency of devices used in air systems such as the FFUs and their control mechanism largely affected the overall air delivery efficiency, and could vary largely. On the other hand, the filtration efficiency could be affected by airflow speeds, the design, geometry, and material of filters used in the minienvironment. Optimal contamination control for minienvironments could be realized by regulating airflow rates and air pressure differentials between minienvironment space and its surrounding space to achieve effective and efficient particulate filtration control.

Recommendations from this study include investigating power density as well as energy performance indices of minienvironment as compared to that of traditional cleanroom systems, integration of minienvironments in cleanrooms, and further analysis of savings potential for integrated design, construction, operation, and management of clean spaces.

Last but not the least, there is a need to develop strategies and best practices that can be used by the industry for energy efficiency improvement. Future activities may include providing technical education, interactions and engagement in developing or improving IEST Recommended Practice and design guidelines, and collaboration with industrial leaders such as leading minienvironment suppliers or users. The improvement in energy savings and minienvironment system performance will be beneficial to sustainable development in this sector.

6. Acknowledgement

The project was funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program (http://www.energy.ca.gov/). This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098. The author would like to thank Chris Humphrey, Pete Hansen, Rich Dow, and Carl Buice of Asyst Technologies, Inc. for their support in this project. Special appreciations are extended to Dr. Sameer Abu-Zaid, Dr. Haifeng Zhang, and Bruce Richardson of Asyst Technologies, Inc., for assistance and valuable input to this report.

7. References

1. The Institute of Environmental Sciences and Technology (IEST). 2002. IEST-RP-CC028.1: Minienvironments. Handbook of Recommended Practices, Contamination Control Division. 5005 Newport Drive, Suite 506, Rolling Meadows, IL 60008-3841, USA

2. International Organization for Standardization (ISO). 2005. ISO 14644-7, Cleanrooms and associate controlled environments - Part 7: Separative Devices (clean air hoods, gloveboxes, isolators and minienvironments). The Institute of Environmental Sciences and Technology (IEST). 5005 Newport Drive, Suite 506, Rolling Meadows, IL 60008-3841, USA

3. International Organization for Standardization (ISO). 1999. ISO 14644-1, Cleanrooms and associate controlled environments - Part 1—Classification of Air Cleanliness**.** 18 pages, The Institute of Environmental Sciences and Technology (IEST). 5005 Newport Drive, Suite 506, Rolling Meadows, IL 60008-3841, USA.

4. International Organization for Standardization (ISO). 2001. ISO 14644-4, Cleanrooms and associate controlled environments - Part 4: Design, Construction and Start-up. The Institute of Environmental Sciences and Technology (IEST). 54 pages. 5005 Newport Drive, Suite 506, Rolling Meadows, IL 60008-3841, USA

5. International Organization for Standardization (ISO). 2002. ISO 14644-3, Cleanrooms and associate controlled environments - Part 3: Metrology and Test Methods. 65 pages. The Institute of Environmental Sciences and Technology (IEST). 5005 Newport Drive, Suite 506, Rolling Meadows, IL 60008-3841, USA

6. Tannous, G. 1996. Air Flow Simulation in a Minienvironment, Solid State Technology, July 1996, pp. 201-207.

7. Olim, M. 1998. "Optimizing Minienvironment Air Flow Patterns Using CFD," FSI International Chaska, Minn., Semiconductor International. Cleanroom Magazine.

8. Shu, S. 1995. The Impact of Minienvironment of Fab Capacity and Yield, Institute of Environmental Science and Technology, 41st Annual Technology Meeting, Anaheim, Calif. pp. 118-123.

9. Rappa, M. 1994. Factors Influencing the Rate of Adoption of Minienvironment Technology in United States, Japan, Korea, and Taiwan. 1994 Proceedings, 12th International Symposium on Contamination Control, Yokohama, Japan, pp. 503-507.

10. International SEMATECH. 1999. Integrated Minienvironment Design Best Practices, Technology Transfer # 99033693A-ENG. International SEMATECH, Inc., Austin, Texas.

11. PowerSight® Energy Analyzer, Summit Technology, 2717 N. Main St., Suite 15, Walnut Creek, CA 94596; http://www.summittechnology.com/pk314m.shtml

12. Shortridge Instruments, Inc., ADM-860C Airdata Multimeter, Backpressure Compensated Air Balance System. 7855 East Redfield Road • Scottsdale, Arizona 85260. Phone (480) 991-6744 • Fax (480) 443-1267 • www.shortridge.com • info@shortridge.com. http://www.shortridge.com

13. Xu, T. 2003. Performance Evaluation of Cleanroom Environmental Systems, Journal of the IEST, 2003 Edition, Vol. 46: 66-73, LBNL-53282.

14. Xu, T., W. Tschudi, G. Bell, E. Mill, and D. Sartor. 2001. "High-tech Buildings – Market Transformation Projects," Final report to California Institute for Energy Efficiency (CIEE) and Pacific Gas and Electric Company (PG&E), California, LBNL-49112.

Lawrence Berkeley National Laboratory Report LBNL-57503

Energy Implications of Minienvironments in Clean Spaces:

A Case Study on Minienvironment Energy End-use and Performance

Final Report

To The

California Energy Commission

2005

Tengfang Xu

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Contents

1. Background 6
2. Introduction 6
3. Objectives 9
4. Approach 9
4.1 Electric Power Demand 10
4.2 Airflow Speed and Pressure Differential 10
4.3 Particle Concentration 11
5. Results 11
5.1 Characteristics of the cleanroom 11
5.2 Characteristics of the minienvironments 13
5.3 Minienvironment Energy Performance 14
5.3.1 Airflows and Air-Change Rates 15
5.3.2 Energy Performance Index (EPI) 18
5.3.3 Electric Power Density 20
5.4 Minienvironment Environmental Performance 22
5.4.1 Pressure Differential 23
5.4.2 Particle Concentration 23
6. Discussion 24
6.1 Case-based Estimation 24
6.2 Design-based Estimation 25
6.3 Summary of the Discussion 26
7. Conclusions and Recommendations 27
7.1 Conclusions 27
7.2 Recommendations 28

8. Acknowledgements 29

9. References 30

10. Appendices 33

10.1 Appendix - Schematic Diagram of Minienvironments in the Cleanroom 33

10.2 Appendix - Case-Based Estimate of Energy-savings 34

10.3 Appendix - Design-Based Estimate of Energy-savings 34

10.4 Appendix - Summary of the Workshop on Minienvironment 35

10.4.1 Outline of the Minienvironment Workshop 35

10.4.2 Topics Covered in the Minienvironment Workshop 35

10.4.3 Participant Forum on Minienvironments 36

10.4.4 Workshop Presentation and Recommendations from the Case Studies (Power Point slides) 36

Table of Figures

Figure 1 Air Change Rates and Airflow Speed 17

Figure 2 Energy Performance Index and Airflow Speed..... 19

Figure 3. Electric Power Density and Airflow Speeds for Five Different Minienvironments and the surrounding Cleanroom..... 21

Table of Tables

Table 1 Cleanroom Airflows and Electric Power Demand 12

Table 2 Characteristics of Sample Minienvironments..... 14

Table 3 Minienvironment Airflow and Electric Power Demand..... 15

Table 4 Minienvironment Environmental Performances..... 23

Energy Implications of Minienvironments in Clean Spaces:

A Case Study on Minienvironment Energy End-use and Performance

1. Background

Cleanroom air-recirculation systems typically account for a significant portion of the HVAC energy use in cleanrooms. Fan's electric power density is defined as the fan's electric power demand divided by the cleanroom floor area. The electric power density for fans to deliver and recirculate air is normally high because of high airflow rates that are required for contamination control in cleanrooms.

With the demand for better contamination control in specific applications, e.g., higher cleanliness within a localized and relatively small space, it is important to optimize the design of clean spaces and their airflows in order to achieve energy savings. The optimization of airflow, layout, and sizing of clean spaces may potentially lead to energy savings.

A minienvironment is a localized environment created by an enclosure to isolate a product or process from the surrounding environment [1][2]. Such a minienvironment is normally used to maintain a stringent level of cleanliness in a tightened clean space. Some minienvironments provide various device and physical configurations to actively or passively direct air from the surrounding cleanroom to and from the minienvironments. Some other minienvironments include independent temperature control, humidity control, and chemical filtration as part of their operation. For these, energy use can become more intensive. In order to understand energy saving implications, it is useful to obtain information on energy and environmental performance of minienvironments. At the same time, it is important to understand their design and field installation, and to identify potential energy-saving opportunities associated with minienvironments and cleanrooms.

This report summarizes a field study on the performance of a group of minienvironments installed in a semiconductor cleanroom facility. The report presents characteristic and performance information about the minienvironments and the cleanroom that encompassed the minienvironments. It also discusses energy-saving implications from applying the minienvironments, and opportunities in integrating minienvironments with the cleanroom as a way of achieving energy-savings. Based upon the findings and conclusions, this study identifies research gaps and recommendations for future work that is necessary to address the gap and to advance the design and operation of minienvironments in clean spaces. LBNL collaborated with PG&E to host a half-day workshop focused on minienvironments. A summary of the minienvironment workshop is included in the Appendix section of this final report.

2. Introduction

The purpose of a minienvironment is to achieve effective contamination control in a localized space, often through maintaining desired pressure differential or supplying unidirectional airflows needed for maintaining cleanliness levels within the space [1]. The dimensions of the minienvironment spaces may vary depending on specific applications. A recent research

provided quantitative information on the performance of a minienvironment system [3][4]. A further study indicated that energy efficiency opportunities exist through optimizing design and operation of minienvironment air systems [5].

The advantages in using minienvironments include:

- Minienvironments may create cleanliness-class upgrade [6], better contamination control, and process integration.

- Minienvironments may maintain better contamination control by controlling pressure differentials or providing unidirectional airflows.

- Minienvironments may potentially reduce energy costs.

Past studies focused on design optimizing of minienvironments and contamination control [7][8][9] [10][11][12][13]. In addition, IEST publishes the methods or protocols on construction and operation of minienvironments [14][15][16]. Other studies or benchmarking activities addressed the impact of production yields by adopting minienvironments[17][18][19]. For example, a benchmarking study on minienvironments provided performance data but excluded information on the energy impact of minienvironments on the enclosing cleanroom facility [19]. Unfortunately, none of the above-mentioned studies [7-19] addressed energy issues associated with minienvironment applications, nor was quantitative data about energy performance provided in any of those studies.

Because minienvironments typically use fan-filter units, the energy intensity may be increased for the space served compared with similar cleanrooms without minienvironments. At the same time, appropriate integration of minienvironments with the surrounding cleanroom may also help to alleviate the overall electric power demand for the facility. A recent study quantified the electric power density of the air system in a minienvironment as a function of airflow speeds and pressure differential [5]. Corresponding to operating ranges for the minienvironment studied, electric power density ranged approximately from 17 W/ft^2 to 28 W/ft^2 (183 W/m^2 to 300 W/m^2) with the air speeds from 30 fpm to 110 fpm (0.15 m/s to 0.55 m/s). This range actually fell within the range of fan power density from previously measured ISO-Cleanliness-Class-4 cleanrooms, i.e., in the range of 16 W/ft^2 to 38 W/ft^2, or 172 W/m^2 to 409 W/m^2 [20][21].

With the goal of achieving the same cleanliness level within a minienvironment as that of the surrounding cleanroom, the airflow rate supplied to the minienvironment can be significantly lower than the airflow rate supplied to a full-scale cleanroom because of the significantly smaller volume of a minienvironment. This presents potential opportunities for energy savings when the required airflow rates for minienvironments could be reduced, i.e., the much smaller volumes of air that must be moved, conditioned, and filtered in a given time [3][4].

Prior to this study, virtually no quantitative data associated with the use of minienvironments in operation was available to quantify the actual energy-savings potential. In order to understand actual energy implications of incorporating minienvironments, it is necessary to quantify the magnitude of electric power demand or energy end-use of various minienvironments and that of

the surrounding cleanroom. This is important to understand the overall energy implications of a cleanroom housing minienvironments.

3. Objectives

This case study was an investigation of the energy and environmental performance of a group of minienvironments in a cleanroom under normal operation. The main objective of this study is to develop field information to understand energy and environmental performance of minienvironments. This report provides discussion of energy-savings potential from adopting minienvironments for contamination control and improving the energy efficiency of minienvironment systems. The information and recommendations developed from this study can be used to identify energy-savings potential, research gaps, and future investigations in achieving efficient and effective (E^2) minienvironments in the industries that use them.

Specifically, the technical objectives of this study include:

1) Understand the energy and environmental performance of the minienvironment systems.

2) Compare the energy performance of the minienvironments and that of cleanrooms.

3) Discuss and estimate energy-savings potential by applying energy-efficient minienvironments within cleanrooms for effective contamination control.

4) Identify research gaps and develop recommendations for future research and investigations.

4. Approach

A minienvironment is used to maintain a certain cleanliness level by controlling particle concentrations in the localized space. In this study, various localized spaces within an ISO-Cleanliness-Class-4 cleanroom were installed and used to achieve a cleanliness level equivalent to that of an ISO-Cleanliness-Class-3 or Class-4 clean space.

In the cleanroom, there were various activities that required different environmental conditions depending on the process or locality within the cleanroom, i.e., ISO-Cleanliness-Class-3 and/or ISO-Cleanliness-Class-4 localized spaces. A number of minienvironments with a cleanliness level equivalent to ISO-Cleanliness-Class-3 spaces in the cleanroom were installed in the facility. Specifically, these stand-alone, self-powered minienvironments were used to provide filtered air through localized HEPA or ULPA filters at certain airflow speeds for various processes or product-testing activities. Another group of minienvironments within the same ISO-Cleanliness-Class-4 cleanroom was designed and installed to provide physical barriers and they contained no additional fan-powered device such as a fan-filter unit. These passive, non-fan-powered minienvironments were used to present physical barriers to isolate the process and activities from contamination, which could be affected by unexpected changes in ambient conditions, local disturbance of airflow patterns, or pollutants from the human occupants working in the ISO-Cleanliness-Class-4 cleanroom that housed the minienvironments.

The measured parameters included electric power demand (representing energy end-use), airflow rates, air pressures, and particle concentrations in and around the minienvironments under normal operating conditions. The information developed in this case study included electric power demand of five different stand-alone, fan-powered minienvironments, the energy performance of these minienvironments, and the effectiveness of contamination control in such device. Key performance metrics were developed and calculated to characterize the overall performance of the minienvironments. This study also compared the performance of the five minienvironments with that of the enclosing cleanroom and other cleanrooms that were previously studied. Based upon the measured data, the study discussed and estimated the potential energy-savings from implementing energy-efficient minienvironments.

4.1 Electric Power Demand

The power meter used in this study was a true RMS energy analyzer with a measurement uncertainty of ±3% [22]. The meter recorded the electric current, voltage, power factor, the actual power supplied to the air delivery systems of the minienvironments in the cleanroom, and the power supplied to the air-handling-unit systems for the cleanroom. The power meter was used with various current transducers (uncertainty ±2%) and voltage transducers to measure the electric current, voltage, power factor, and actual power demand of the air delivery systems. The air delivery systems were the fan-filter units (FFUs) serving the five ISO-Cleanliness-Class-3 minienvironments, and two types of air-handling units serving the ISO-Cleanliness-Class-4 cleanroom. The measured power demand was used to quantify the energy performance of the air systems for the operating minienvironments as well as that of the cleanroom.

4.2 Airflow Speed and Pressure Differential

A backpressure-compensated device attached to an electronic micro-manometer [23] measured the average speeds of the airflow delivered out of the face of the fan-filter units, which were installed at the top of the stand-alone minienvironments. The actual sizes of individual FFUs and HEPA filters varied from minienvironment to minienvironment.

The measurement uncertainty in airflow speeds was ± 3% of reading plus ± 7 fpm (3.5 cm/s) from 50 to 2500 fpm (0.25 m/s to 12.5 m/s). An airflow measurement device was used to sample 16 points over a 1 ft x1 ft (30 cm x 30 cm) area to determine average airflow speeds at a distance of 2.5 inches (6.3 cm) downstream away from the face of the filter frames. Airflow-speed readings were automatically corrected for the density effect of barometric pressure and temperature. Readings were displayed as local density and true air speeds.

Pressures were measured using a Pitot tube with a multi-meter. The multi-meter measures a wide range of pressures from 0.0001-inch-water column (0.025 Pa) to over 60.00-inch-water column (15,000 Pa), with a measurement uncertainty of ±2% of reading plus 0.001-inch-water column (0.25 Pa) from 0.05-inch-water column to 50.00-inch-water column (0.125 Pa to 12,500 Pa). The air pressure differential between the space inside the minienvironment and the space surrounding the minienvironment was recorded for each minienvironment, concurrent to the airflow measurements under the normal operating conditions.

4.3 Particle Concentration

In addition to measuring electric power demand, airflow speeds, air pressure differential between the space inside the minienvironments and the space surrounding the minienvironments, particle concentration levels were measured concurrently to evaluate environmental performance of the minienvironments, i.e., particle concentration inside and outside of the minienvironments.

According to the definition of Airborne Particulate Cleanliness Classes in ISO Standard 14644 [6], the classification of air cleanliness in cleanrooms and associated controlled environments is defined in terms of concentration of airborne particles within the space. For example, a cleanroom with an ISO-Cleanliness-Class-4 level corresponds to no more than 10,000 counts of particles per cubic meter with particle sizes of 0.1-µm or larger, or 352 counts of particles per cubic meter with particle sizes of 0.5-µm or larger, in the space of concern. Using this concept, a minienvironment with an ISO-Cleanliness-Class-3 level corresponds to no more than 1,000 counts of particles sizing 0.1-µm or larger per cubic meter, or 35 counts of particles sizing 0.5-µm or larger per cubic meter, of the minienvironment space.

Laser particle counters [24] were used to measure the particle concentration within the minienvironments. The laser-based particle counter discriminated and counted particles with sizes of 0.1-µm, 0.2-µm, 0.3-µm, 0.5-µm, 1.0-µm, and 3.0-µm. The airflow rate used for particle sampling was 2 cfm (56.6 L/min) supplied by an internal carbon-vane pump in the counters. In general, a higher airflow rate for particle sampling in the chamber of a particle counter indicates higher capacity of sensing particles traveling into the counter and better accuracy in particle counts during transitional (or unsteady-state) sampling.

5. Results

5.1 Characteristics of the cleanroom

The cleanroom housing the minienvironments in this study was located on the second floor of a two-story semiconductor manufacturing facility in Southern California. The ISO-Cleanliness-Class-4 cleanroom had a total floor area of 4,065 ft^2 (378 m^2) with a ceiling height of 10 ft (3.0 m), and operated 24 hours a day and 365 days a year. In addition to one make-up air system, two types of recirculation air systems served the cleanroom: ducted-HEPA-filter and pressurized-plenum.

The fans in the recirculation air-handling units for the cleanroom were originally designed to deal with possible future expansion, which was expected during the original design and installation. For example, in the original design, airflow rates for recirculation consisted of a) 216,000 cfm (2,702 m^3/min) to be supplied through a total of four air-handling units (176 kW) connected to the ducted-HEPA filters, and b) 131,100 cfm (1,811 m^3/min) to be supplied by a total of three additional air-handling units (121 kW) connected to the pressurized plenum.

The air-handling units connected to the ducted-HEPA-filter systems were designed to cover approximately 2,290 ft^2 (213 m^2) of the primary cleanroom space[1], while the other three-air-handling units serving the pressurized plenum covered approximately 1,390 ft^2 (129 m^2) of the primary cleanroom space. The total floor area of the primary cleanroom space was 3,680 ft^2 (or 342 m^2). The cleanroom had a secondary space for return air, which covered a floor area of approximately 385 ft^2 (36 m^2).

Table 1 shows the physical size of the cleanroom, airflow rates, electric power demand, air-system efficiency, air-change rate, and electric power density for the air-recirculation systems, and make-up-air systems in its normal operation.

Table 1 Cleanroom Airflows and Electric Power Demand

Air-handling Systems in ISO Class 4 Cleanroom	Units	Recirculation Air (Ducted HEPA Filters)	Recirculation Air (Pressurized Plenum)	Recirculation Air (Combined)	Make-up Air
Floor Area Served	m^2	213	129	342	342
	ft^2	2,290	1,390	3,680	3,680
Airflow Rate	m^3/min	2,702	1,811	4,513	424
	cfm	95,406	63,963	159,369	14,960
Electric Power	kW	24	13	38	11
Airflow Rate per Power Demand	m^3/min/W	0.11	0.14	0.12	0.04
	cfm/kW	3,915	4,871	4,250	1,324
Average Cleanroom Airflow Speed	m/s	0.21	0.23	0.22	0.02
	feet per minute (FPM)	42	46	43	-
Air-change Rate	m^3air/(hr-m^3room)	250	276	260	24
	ft^3air/(hr-ft^3room)	250	276	260	24
Electric Power Density	W/m^2	115	102	110	33
	W/ft^2	11	9	10	3

In actual operation, the airflow rates from the ducted-HEPA-filter systems and the pressurized-plenum systems were measured to be 95,406 cfm (2,702 m^3/min) and 63,963 cfm (1,811 m^3/min), respectively. The total of the actual recirculation airflow rate was 159,369 cfm (4,513 m^3/min), which was about 46% of the design airflow rate. This corresponded to the fan power of 38 kW, which was approximately 13% of the designed fan power for recirculation air. The average of measured recirculation-fan power density for the cleanroom was approximately 10 W/ft^2 (110 W/m^2), which was lower than the designed 81 W/ft^2 (872 W/m^2). Apparently, the

[1] Note: The floor area served is listed as estimation based upon the assumption that both the ducted-HEPA systems and pressurized-plenum systems provided the same airflow speed at the design condition.

system, including the fan motors were oversized and therefore a much lower pressure drop was entailed in the air recirculation systems than originally expected and designed.

The overall recirculation airflow rate per fan power demand was 4,250 cfm/kW (0.12 m^3/min/W), which was about three times higher than the design (1,165 cfm/kW, or 0.03 m^3/min/W). Overall, the average cleanroom airflow speed was 43 fpm (or 0.22 m/s) compared to the design value of 94 fpm (or 0.47 m/s), with a recirculation-air-change rate of 260 air-volume/hr-room-volume (260 m^3air/hr-m^3room).

5.2 Characteristics of the minienvironments

The minienvironments in this study were partially enclosed spaces in the ISO-Cleanliness-Class-4 cleanroom. Two types of minienvironments were identified in this cleanroom:

1) A stand-alone minienvironment with an open-loop air system, within which airflow was drawn from the surrounding cleanroom space, through fan-filter units that were attached at the top of the minienvironment. The filtered air was then supplied into the minienvironment to maintain a higher cleanliness level within the localized space, i.e., ISO-Cleanliness-Class-3 space.

2) A passive minienvironment to which no additional fan was attached. This was contrary to the case for fan-filter units on top of a stand-alone minienvironment. A passive minienvironment mainly served as a physical barrier to provide a buffer zone from the surrounding space to minimize external disturbance. Normally without any additional filter, it provided separation in order to maintain a cleanliness level within the separate minienvironment, i.e., ISO-Cleanliness-Class-4 space.

A schematic diagram of the minienvironments in the cleanroom is included in the Appendices of this report. In a stand-alone, open-loop minienvironment, the supply air was filtered through FFUs located on top of the minienvironment. Additional flow shields were installed underneath the HEPA filters of the FFUs to create downward unidirectional airflows inside the minienvironment. The outgoing airflows from the minienvironment may then mix with the surrounding air within the cleanroom space.

Because a passive minienvironment did not directly affect overall electric power demand for airflow delivery to the minienvironment, this study focused only on a group of stand-alone, open-loop minienvironment systems.

Table 2 shows the physical size of the inner-space of the stand-alone, open-loop minienvironments that were selected and measured in this study.

Table 2 Characteristics of Sample Minienvironments

Minienvironments	Units	A	B	C	D	E
Floor Area	m^2	6.3	1.2	1.7	0.7	4.1
	ft^2	68	13	18	8	44
Height	cm	178	259	230	216	240
	inch	70	102	91	85	95

Overall, eight minienvironments with the size equivalent to that of "A" listed in Table 2 were located in the ISO-Cleanliness-Class-4 cleanroom. Among these, five minienvironments were stand-alone, open-looped systems that were designed to create ISO-Cleanliness-Class-3 spaces, while three others were passive minienvironments without fans to deliver the airflow from the cleanroom into the minienvironments. Additional minienvironments, including minienvironments B, C, D, and E, were located within the same cleanroom. The total of net floor area of the stand-alone, open-looped minienvironments was estimated as approximately 424 ft^2 (39 m^2), which represented approximately 12% of the cleanroom primary floor area.

5.3 Minienvironment Energy Performance

The operating efficiency of the FFUs in the minienvironments can vary considerably at various operating conditions. Optimizing the airflow speed and air pressure in a minienvironment not only can improve FFU operating efficiency, but also may improve space cleanliness, noise and vibration characteristics, and operating life of the fans.

The minienvironments normally operated continuously - 24 hours a day and seven days a week. Spot measurements were taken at the five minienvironments that were selected to quantify and evaluate their energy and environmental performance.

Table 3 shows the measurement results including airflow rate, airflow speed, electric power, air-system efficiency, energy performance index (EPI), air-change rate, and electric power density for the five minienvironments.

Table 3 Minienvironment Airflow and Electric Power Demand

Minienvironments	Units	A	B	C	D	E	A-E Sum	Average
Airflow Rate	m^3/min	141	21	26	22	106	317	-
	cfm	4,988	745	927	792	3,730	11,182	-
Electric Power	kW	2.1	0.4	0.4	0.3	1.1	4.3	-
Airflow Rate per Power Demand	m^3/min/W	0.07	0.06	0.06	0.09	0.09	-	0.07
	cfm/kW	2,353	1,961	2,250	3,106	3,272	-	2,588
EPI	W/(m^3/min)	15.0	18.0	15.7	11.4	10.8	-	14.2
	W/cfm	0.43	0.51	0.44	0.32	0.31	-	0.40
Average Airflow Speed	m/s	0.37	0.30	0.26	0.50	0.43		0.37
	fpm	73	58	52	99	84	-	73
Air-change Rate	m^3air/(hr-m^3room)	752	412	410	839	642	-	611
Electric Power Density	W/m^2	335	320	246	343	277	-	304
	W/ft^2	31	30	23	32	26	-	28

5.3.1 Airflows and Air-Change Rates

For the selected minienvironments within the ISO-Cleanliness-Class-4 space, the recirculation air was supplied through the FFUs to the minienvironments.

While there were variations in the floor area of the minienvironments ranging from eight ft^2 to 68 ft^2 (0.7 m^2 to 6.3 m^2), the minienvironments exhibited a wider range of airflow rates, namely, ranging significantly from 745 cfm to 4,988 cfm (21 m^3/min to 141 m^3/min) [Table 3]. This wide variation in airflow rates corresponded to different airflow speeds in the various minienvironments, in addition to the various floor areas.

The average airflow speed inside each minienvironment ranged from 52 fpm to 99 fpm (or 0.27 m/s to 0.50 m/s), with an average of 73 fpm (or 0.37 m/s). The airflow speeds were generally higher than the average airflow speed in the surrounding cleanroom, which was 43 fpm (or 0.22 m/s) as shown in Table 1.

The air-change rate of a minienvironments was calculated as the total airflow rate supplied to the minienvironment divided by the actual room volume of the minienvironment. The value of an air-change rate would be the same, regardless of the actual units used (i.e., IP or SI), For example, an air-change rate of 410 m^3air/hr-m^3room (SI unit) represents the same level of air-change rate quantified as 410 ft^3air/hr-ft^3room (IP unit). For simplicity, we only denote one unit convention for air-change rate in the rest of this paper.

The air-change rates of the five minienvironments differed from 410 m^3air/hr-m^3room to 752 m^3air/hr-m^3room, exhibiting a similar range to the operating range of a typical stand-alone, open-looped minienvironment in a previous study [3][4]. In that study, the operating range of air-change rates for the minienvironment was between 480 m^3air/hr-m^3room and 800 m^3air/hr-m^3room, corresponding to airflow speeds ranging from 60 fpm to 100 fpm (or 0.30 m/s to 0.50 m/s) in the minienvironment.

In summary, the air-change rates of the five minienvironments tested in this study were significantly higher than that of the ISO-Cleanliness-Class-4 cleanroom housing the minienvironments, i.e., 260 m^3air/hr-m^3room.

As shown in Figure 1, it is clear that higher average airflow speeds, higher HEPA/ULPA filter coverage in the five minienvironments (i.e., 100%), and lower ceiling heights of the minienvironments collectively contributed to the higher air-change rates within the minienvironments than that of the surrounding cleanroom.

When compared with the average airflow speeds in other ISO-Cleanliness-Class-4 cleanrooms from a previous study [20], the magnitude of airflow speeds from these ISO-Cleanliness-Class-3 minienvironments generally exhibited a similar or lower range. In addition, within similar airflow speed range, the air-change rates of the five minienvironments exhibited a slightly wider range than that of ISO-Cleanliness-Class-4 cleanrooms, which was between 385 and 680 m^3air/hr-m^3room corresponding to airflow speeds ranging from approximately 60 fpm to 120 fpm (or 0.30 m/s to 0.60 m/s) [20]. In general, the HEPA/ULPA filter coverage in the minienvironments was 100% while the ISO-Cleanliness-Class-4 or ISO-Cleanliness-Class-5 cleanrooms may have a lower coverage from 25% up to 100%.

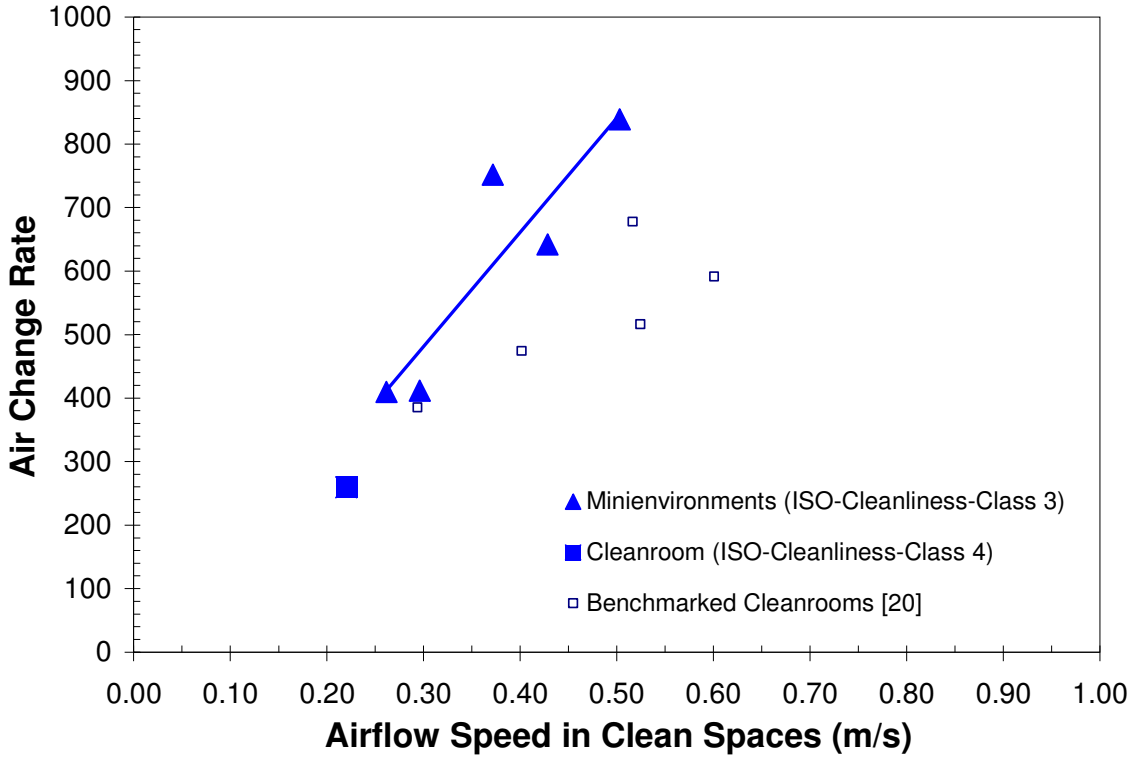

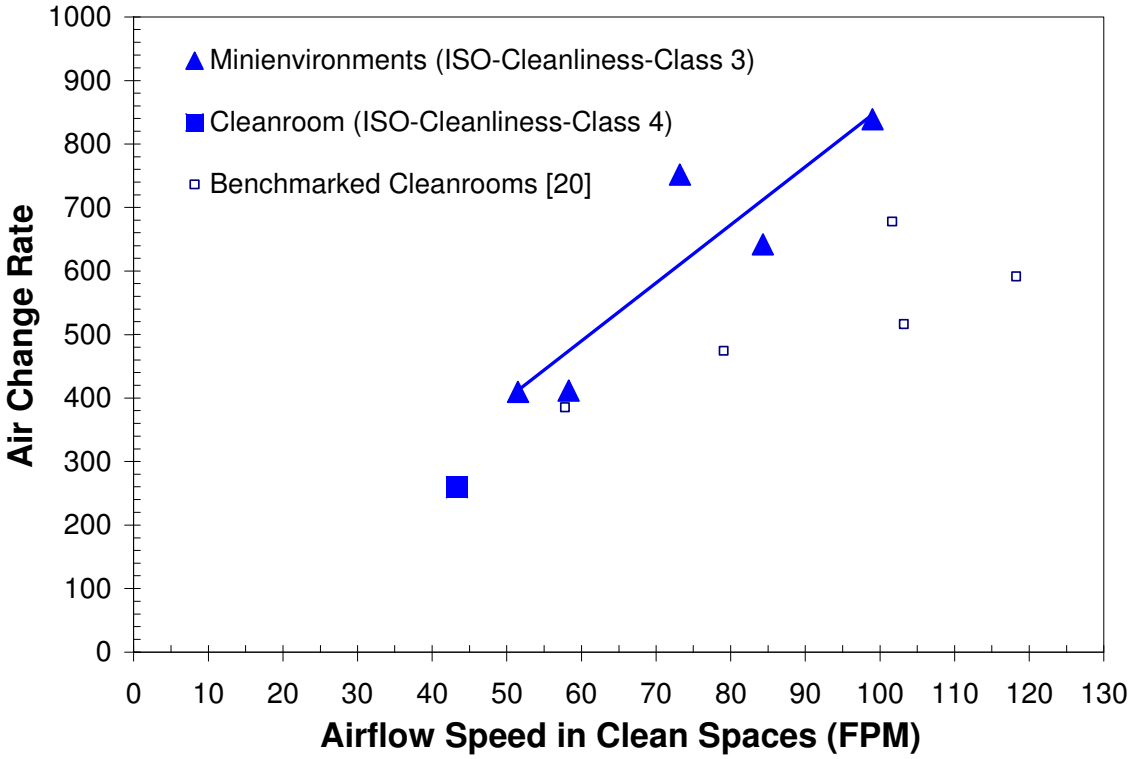

Figure 1 Air Change Rates and Airflow Speed

5.3.2 Energy Performance Index (EPI)

The energy performance index (EPI) of a minienvironment air system is defined as the total electric power supplied to the fan system divided by the airflow rate in the minienvironment [3][4]. A higher EPI value under the same operating condition means that the demand for electric power is higher to supply the same airflow rate to the minienvironment, thus corresponding to lower energy efficiency of the air-delivery systems in the minienvironment.

Figure 2 shows the measured EPI values of the five minienvironment systems compared to that of the surrounding cleanroom (ISO-Cleanliness-Class 4). The air systems' EPI values of the five minienvironment systems (designed as ISO-Cleanliness-Class 3) showed a wide range, i.e., ranging from 0.31 W/cfm to 0.51 W/cfm (10.8 W per m^3/min to 18.0 W per m^3/min), corresponding to airflow speeds ranging from approximately 52 fpm to 99 fpm (or 0.27 m/s to 0.50 m/s). In addition, the EPI values of the ISO-Cleanliness-Class-3 minienvironments were consistently higher than the surrounding ISO-Cleanliness-Class-4 cleanroom, of which the EPI value was 0.24 W/cfm (8.5 W per m^3/min) corresponding with a lower airflow speed.

It is clear that the EPI values among these minienvironments tended to decrease with the increase in the airflow speed (or airflow rates normalized by minienvironment floor area) inside the minienvironments. This trend indicates that within the measured operating range, lower EPI values (more efficient in delivering the air) tended to correlate with higher airflow speeds among the five minienvironments. This trend was similar to the finding from a previous study on a typical stand-alone, open-looped minienvironment system, which exhibited an operating range from 60 fpm to 100 fpm (or 0.30 m/s to 0.50 m/s) in the minienvironment [3][4]. However, EPI values of the minienvironments in this study were slightly higher when compared with that of the other minienvironment operating with the similar airflow speeds [3][4].

Figure 2 also includes the measured EPI values for the various ISO-Cleanliness-Class-4 cleanrooms that were previously studied [20], which ranged between 0.21 W/cfm to 0.53 W/cfm (7.4 to 18.7 W per m^3/min). The EPI values for the minienvironments, which generally operated at a similar or lower airflow speed, were generally higher than those of cleanrooms.

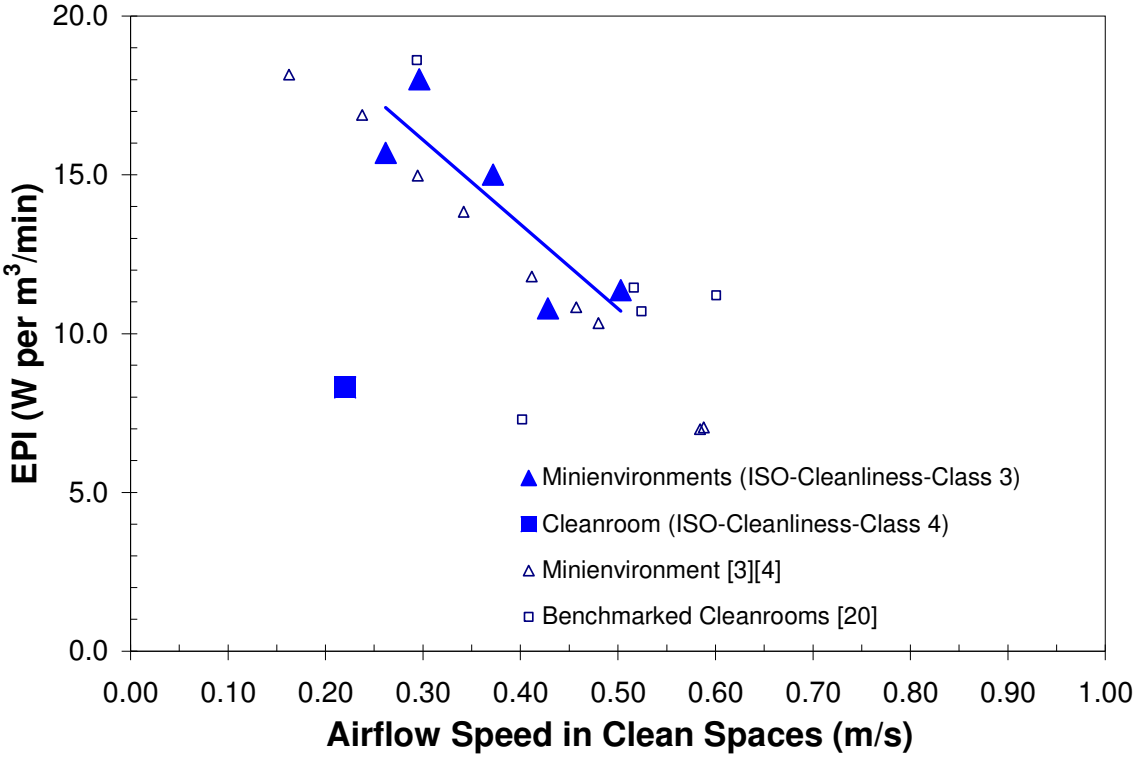

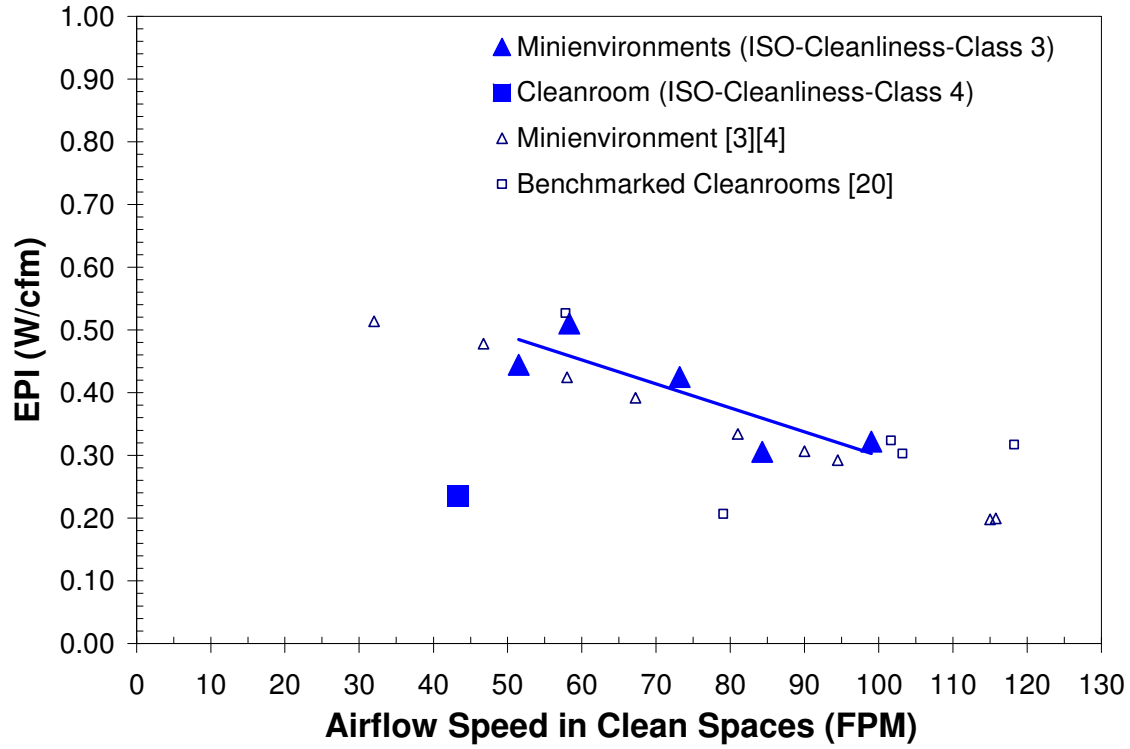

Figure 2 Energy Performance Index and Airflow Speed

5.3.3 Electric Power Density

Electric power density is defined as the electric power demand, which is required for supplying airflow to the clean space such as a minienvironment or a cleanroom, divided by the floor area of the primary clean space intended for contamination control, i.e., floor area of an individual minienvironment or the primary floor area of a cleanroom.

Figure 3 shows the electric power density of the air systems for the five ISO-Cleanliness-Class-3 minienvironments and the air-recirculation fans serving the ISO-Cleanliness-Class-4 cleanroom in this study.

The air-recirculation systems of the ISO-Cleanliness-Class-4 cleanroom in this study included pressurized-plenum and ducted-HEPA recirculation air systems. They collectively exhibited a much lower level of electric fan-power density that those of the minienvironments. Specifically, the electric power density of the air supply systems for five minienvironments ranged from 26 W/ft^2 to 32 W/ft^2 (280 W/m^2 to 344 W/m^2) with an average of 28.3 W/ft^2 (304 W/m^2), while the electric power density of the air-recirculation fans for the ISO-Cleanliness-Class-4 cleanroom was 10.2 W/ft^2 (110 W/m^2). In addition, each value of the electric power density correlated with an average airflow speed inside the clean space. The higher electric power density of the minienvironments corresponded to the airflow speeds ranging from approximately 52 fpm to 99 fpm (or 0.27 m/s to 0.50 m/s), while the lower electric power density of the cleanroom corresponded to an average airflow speed of 43 fpm (0.22 m/s) in the cleanroom. A combination of the following reasons probably contributed to the higher power density in the minienvironments:

- The average airflow speeds in the minienvironments were higher than the average air speed in the surrounding cleanroom.

- The stand-alone minienvironment air systems (FFU systems) with smaller fans were less energy-efficient in delivering air to the intended space, compared to the air-recirculation systems consisting of pressurized-plenum or ducted-HEPA systems typically with larger fans serving the cleanroom.

- The ceiling of all of the minienvironments was fully covered with HEPA filters while the ceiling of the cleanroom was not fully covered with HEPA filters.

It is clear that the electric-power-density values of the minienvironments tended to increase with the increase in the delivered airflow speed (or airflow rate divided by minienvironment floor area) inside the minienvironments. This trend indicates that within the measured operating range, higher values of electric-power-density for the minienvironments (more energy intensive in delivering the air) correlated to higher airflow speeds in the minienvironments. This trend was similar to the finding from a previous study on a minienvironment [3][4] within a certain airflow range (i.e., up to 0.50 m/s).

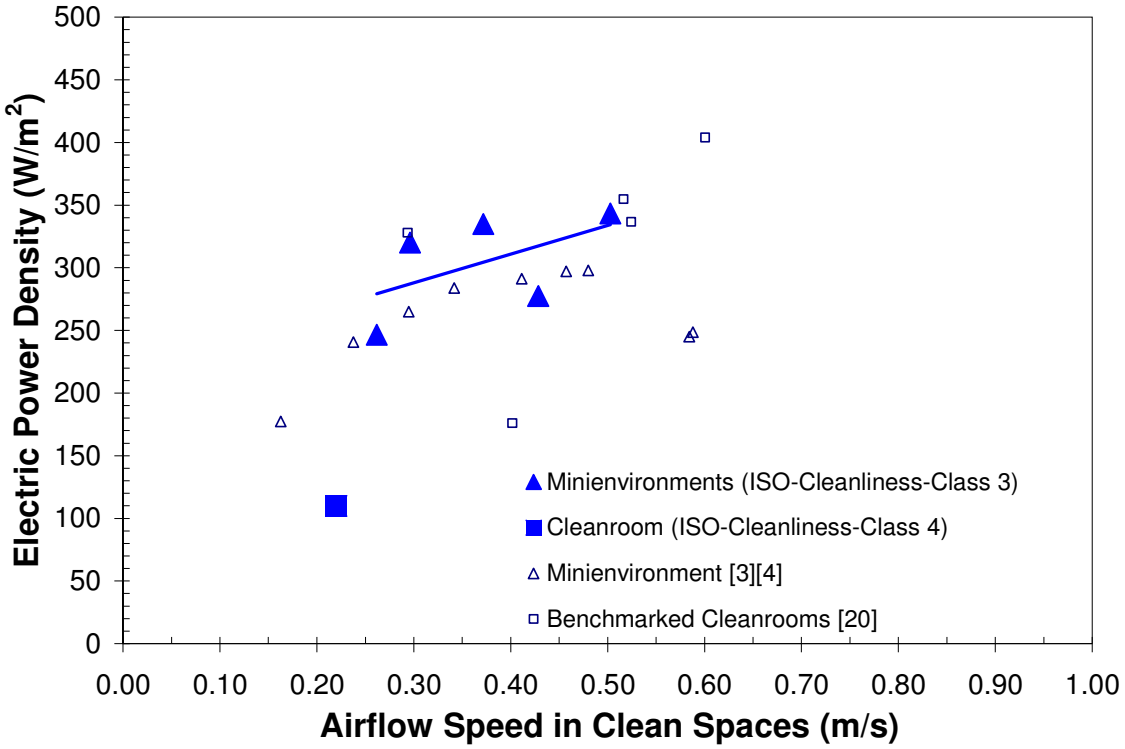

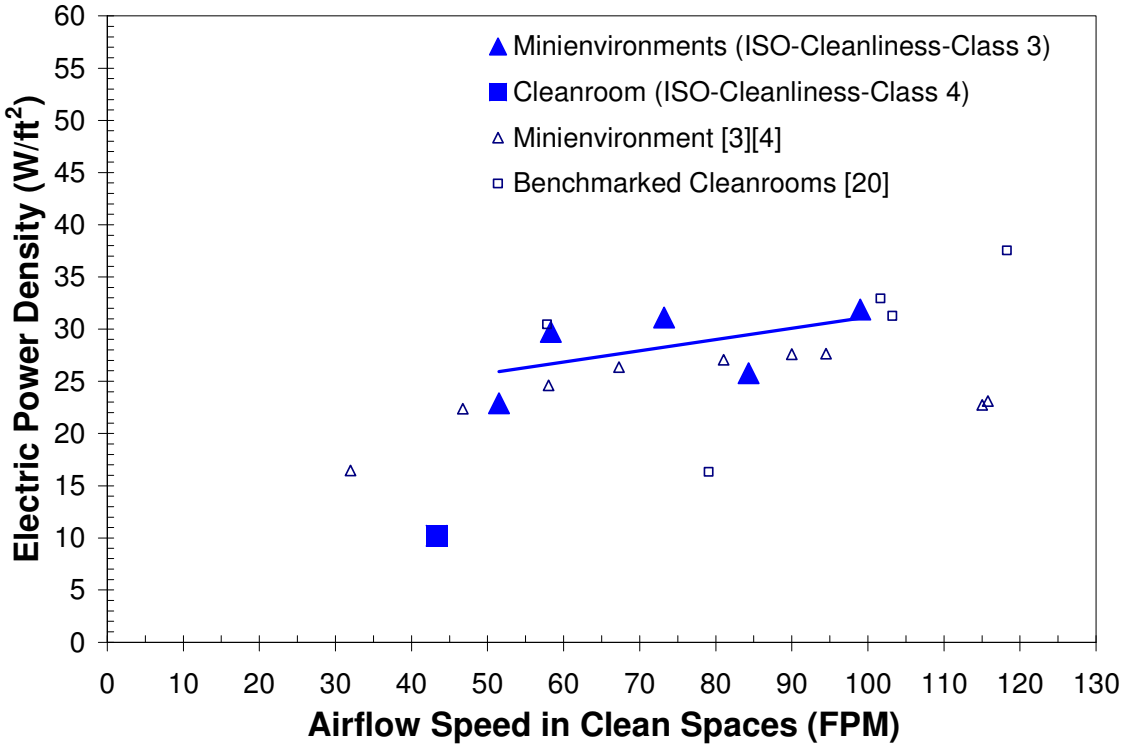

Figure 3. Electric Power Density and Airflow Speeds for Five Different Minienvironments and the surrounding Cleanroom

Furthermore, the figure also includes the electric power density of the air systems reported in previous studies [3][4][20]. While the air-recirculation systems of the ISO-Cleanliness-Class-4 cleanroom in this study collectively exhibited a much lower level of electric fan-power density that those of the minienvironments, they appeared to have lower fan-power density when compared with the group of ISO-Cleanliness-Class-4 cleanrooms with a range of 16 W/ft^2 to 38 W/ft^2 (172 to 409 W/m^2) in a previous study [20]. Those cleanrooms were operating at airflow speeds ranging from 80 fpm to 120 fpm (or 0.40 m/s to 0.60/m/s), higher than the average speed of 43 fpm (0.22 m/s) for the ISO-Cleanliness-Class-4 cleanroom in the current study.

It is also clear that the electric-power-density values of most minienvironments in this study were slightly higher when compared with that of the other minienvironment under the similar range of airflow speeds [3][4]. Given that electric power density of FFU device typically ranged from 20 W/ft2 to 33 W/ft^2 (215 W/m^2 to 355 W/m^2) at the airflow speeds in the vicinity of 50 fpm (0.25 m/s) [25], the minienvironments in this study exhibited similar power density levels of some of the fan-filter units.

In summary, the actual performance data shown in Figure 3 suggests that 1) within the range of airflow speeds measured from the five minienvironments (52 fpm to 99 fpm, or 0.27 m/s to 0.50 m/s), the electric power density of minienvironments typically increased with the increase of average airflow speeds; and 2) the electric power density of the five minienvironments were higher than that of cleanrooms. This indicates that there could be opportunities in optimizing the efficiency of the fan-filter units in the minienvironments, such as optimizing airflow speeds in addition to improve the unit's efficiency. Magnitudes of savings potential are estimated in the discussion section of this report.

5.4 Minienvironment Environmental Performance

The purpose of a minienvironment is to provide contamination control through physical barriers, and it uses filtration to locally control the particle concentration below a certain level within the minienvironment space. It is important to ensure that the enclosed space achieves the required cleanliness class.

The filtration efficiency of HEPA/ULPA filters could be affected by airflow speeds, the design, geometry, and material of filters used in the minienvironment [4]. Optimal contamination control for minienvironments can be realized by regulating airflow rates and air pressure differentials between the minienvironment space and its surrounding space. The benefits of optimal contamination control would include improved effectiveness and efficiency of particulate filtration control.

In common practice, maintaining positive air pressure in a minienvironment relative to the air in the surrounding space may prevent the less-clean air from being transported to the minienvironment and therefore contaminating the process inside the minienvironment.

In the five minienvironments studied, the pressure differential and particle concentration was measured. Table 4 shows the measured results for minienvironments A through E.

Table 4 Minienvironment Environmental Performances

Minienvironments	Units	A	B	C	D	E
Pressure Differential	Pascal	0.15	0.15	0.025	0.025	0.175
	Inch water column	0.0006	0.0006	0.0001	0.0001	0.0007
Space Volume	m^3	11.3	3.1	3.8	1.6	9.9
	ft^3	398	109	136	57	348
Particle Concentration within Minienvironment	Particle count per cubic meter	0	0	0	0	0

5.4.1 Pressure Differential

The pressure differential is the air-pressure difference between minienvironment internal space and its surrounding space. By adjusting the airflow rates, a positive pressure differential for minienvironments may be created to prevent introduction of potential contaminants from the surrounding cleanroom.

Table 4 shows that the measured pressure differential ranged from 0.025 Pa to 0.175 Pa among the five minienvironments. This was lower by several levels of magnitude when compared to the recommended ranges [1], which recommends a typical process-bay pressure exceeding the service-chase pressure by 0.01- to 0.05-inch-water column (or 2.5 Pa to 12.5 Pa) in microelectronic minienvironments. In addition, the measured pressure differential was also much lower than the the range recommended in British Standard [26].

In a recent minienvironment study, the pressure differential ranged from 0.003-inch-water column to 0.024-inch-water column (0.75 Pa to 6 Pa) [4], corresponding to airflow speeds ranging from 32 fpm to 95 fpm (or 0.16 m/s to 0.48 m/s). It is apparent that the actual pressure differential between each minienvironment and the enclosing cleanroom was much lower than the recommended range or the rule of thumb. This was due to large open areas for outgoing airflows through the minienvironments. The observed operation was largely dependent on the function or design of the minienvironment. Less opening area could be achievable by the use of closeable doors at the local area but it was not adopted at the facility site studied.

In summary, while the spot measurements of pressure differential might not be sufficient to represent overall pressure distribution or control for the minienvironments studied, the findings however illustrated that the rule of thumb and the IEST Recommended Practice for the pressure differential in minienvironments may have suggested a higher range than necessary for some of the minienvironment applications. Additional research should look into how to better record and measure spatial pressure distributions over time, and document the acceptable range of pressure differentials in minienvironments.

5.4.2 Particle Concentration

Particle concentration was measured for particles with the sizes ranging from 0.1 micron-meter to three micron-meters within the five minienvironments studied. The particle counter was set to run 30-second samples with a 3-second delay between samples. The sampled particle counts per

space volume were then averaged as reported in Table 4. The measurable concentration was rounded as zero. This was below the particle concentration thresholds for minienvironments with ISO-Cleanliness-Class-3 rating, i.e., no more than 1,000 counts of 0.1-µm particles per cubic meter, or 35 counts of 0.5-µm particles per cubic meter, of the minienvironment space [6]. This indicates that the five minienvironments that were tested all satisfied or even surpassed the minimal environmental requirements for ISO-Cleanliness-Class 3 at the time of particle measurements.

In this case study, supplying and controlling the measured airflow rates through the HEPA filters of the fan-filter units in the minienvironment was sufficient to maintain particle concentration within the required range for the ISO-Cleanliness-Class 3 spaces, even though the actual pressure differential between each minienvironment and the enclosing cleanroom was much lower than the IEST recommended range or British Standard.

6. Discussion

Based upon the measurements in this case study, the average of electric power density of the selected sample ISO-Cleanliness-Class-3 minienvironments was 28.3 W/ft^2 (or 304 W/m^2), while the electric power density for air-recirculation systems in the surrounding ISO-Cleanliness-Class-4 cleanroom was 10.2 W/ft^2 (or 110 W/m^2). As a result, the overall electric power density of the air-recirculation systems for the stand-alone open-looped minienvironments and the ISO-Cleanliness-Class-4 cleanroom as a whole was therefore estimated to be 13.4 W/ft^2 (or 145 W/m^2).

The following includes two approaches for estimating the magnitude of energy savings when implementing energy efficient minienvironments and integrating them with a surrounding cleanroom of various grades of cleanroom cleanliness. The first approach, termed "case-based," is based upon the measurements from this case study, while the second approach, termed "design-based," is based upon the assumptions for various designs and measurements from the relevant studies, including this case study.

6.1 Case-based Estimation

The case-based estimation was built upon actual measurements in this case study. The overall electric power density was13.4 W/ft^2 (or 145 W/m^2) for the minienvironments and the cleanroom as a whole. This level of electric power density was used as the base case for energy-performance comparison.

First, improving the energy efficiency of the minienvironments would create energy-saving opportunities for the overall cleanroom facility. For example, assuming that 40-50% reduction in the minienvironments' power demand would be possible, the electric power savings would be approximately 10-12% compared to the base case, as is illustrated in Appendices. Second, if the electric power density of the fans for cleanroom recirculation air could be reduced by one-third, the overall power savings resulting from implementing the minienvironments and reduced fan power would be 25%.

In another scenario, if the minienvironments were to operate within a non-cleanroom space, meaning that the surrounding cleanliness level (ISO-Cleanliness-Class 4) was not implemented, the overall power savings resulting from implementing the minienvironments and reduced fan power (e.g., office environment) would be approximately 61%. This estimate illustrates that significant energy savings can be achieved by eliminating surrounding cleanliness requirement (i.e., the requirement for ISO-Cleanliness-Class-4 cleanroom cleanliness being relinquished) while assuming that the effective contamination control in ISO-Cleanliness-Class-3 minienvironments could be achieved. The challenge, however, lies in whether or not it is feasible to undergo such change. For example, such a change could be that all processes were to be carried out in minienvironments, achieving equally effective or better contamination control.

6.2 Design-based Estimation

Previous studies indicated that the fan-power density of air recirculation systems in cleanrooms tended to go up with tighter requirements for ISO Cleanliness Class [20]. For example, the fan-power densities of a group of ISO-Cleanliness-Class-4 cleanrooms ranged from 16 W/ft^2 to 38 W/ft^2 (172 W/m^2 to 409 W/m^2), with an average of approximate 30 W/ft^2 (320W/m^2). This range was generally higher than that of the group of ISO-Cleanliness-Class-5 cleanrooms, which equaled to 13.2 W/ft^2 (142 W/m^2) [20]. The fan-power densities of cleanrooms and minienvironments are listed in the Appendices. Because there was no measured data of the fan power density for an ISO-Cleanliness-Class-3 cleanroom, a simplified assumption is taken here, i.e., an ISO-Cleanliness-Class-3 cleanroom was designed to have a fan-power density of 38 W/ft^2 (409 W/m^2), which was the upper range of the ISO-Cleanliness-Class-4 cleanrooms previously studied.

Assuming an ISO-Cleanliness-Class-3 cleanroom is designed with no additional ISO-Cleanliness-Class-3 minienvironment in the facility, we may use this as the base case to estimate energy savings from various designs that would implement minienvironments. For example, we may estimate the energy savings from implementing ISO-Cleanliness-Class-3 minienvironments in such a cleanroom while making the surrounding cleanroom less stringent in terms of its ISO Cleanliness Class (e.g., from Class 3 to Class 4, Class 5, or non-cleanroom, respectively).

If the minienvironments occupy 12% of the total cleanroom floor while they were to operate within an ISO-Cleanliness-Class 4 cleanroom, the overall power savings resulted from implementing the minienvironments and the change in cleanliness requirement (thus reduced fan power) would be approximately 13%. Similarly, if they were to operate within an ISO-Cleanliness-Class 5 cleanroom, e.g., with the airflow rate reduced appropriately, the overall power savings due to the minienvironment implementation would be approximately 57%. Furthermore, if they were to operate within a non-cleanroom space, meaning that the surrounding cleanliness level is not implemented, the overall power savings from implementing the minienvironments and largely reduced fan power (e.g., office environment) would be approximately 86%. This estimate illustrates that eliminating cleanliness requirements for the surrounding space (i.e., the requirement for ISO-Cleanliness-Class-3 cleanliness being relinquished) while maintaining the effective contamination control within the ISO-Cleanliness-Class-3 minienvironments could result in significant energy savings.

In the case study, electric power density of air-recirculation systems in the ISO-Cleanliness-Class-4 cleanroom was measured as 10 W/ft^2 (or 110 W/m^2). A number of minienvironments were located in this ISO-Cleanliness-Class-4 cleanroom. The measured electric power density of the air-recirculation systems was much lower compared to the group of ISO-Cleanliness-Class-4 cleanrooms with a range of 16 W/ft^2 to 38 W/ft^2 (172 W/m^2 to 409 W/m^2) in a previous study [20]. The results illustrate that the electric power savings from adopting the same ISO-Cleanliness-Class-3 minienvironments and reducing the power density of air recirculation systems in the surrounding cleanroom could become even greater.

6.3 Summary of the Discussion

In summary, reducing the electric power density of the cleanroom, implementing energy-efficient minienvironments, and optimizing facility design can collectively contribute to energy savings from operating clean spaces.

Specifically, reducing the fan power density as well as optimizing floor area of the minienvironments and cleanrooms can lead to overall energy savings. Because of the much smaller minienvironment volume compared to that of full-scale cleanrooms (e.g., ballroom), the amount of airflow supplied to the minienvironments for any given time could be significantly reduced. This may present potential opportunities for a significant overall energy savings because of the vastly smaller volumes of airflow that must be moved, conditioned, and filtered within a given time.

In general, in order to create opportunities for significant overall energy savings, measures should be taken to reduce fan power for both minienvironments and cleanrooms. Based upon this study, the following approaches are recommended:

- Reduce electric power demand of the minienvironments.
  - Optimize minienvironment operation, e.g., reduce the airflow and pressurization inside the minienvironments.
  - Optimize the minienvironment design, e.g., geometry and size.
  - Improve the energy efficiency of minienvironment air systems, e.g., FFU efficiency.
- Reduce electric power demand of the primary cleanroom housing the minienvironments.
  - Optimize the control of airflow rate, air-change rates, and pressurization.
  - Optimize the type and size of recirculation air systems.
  - Adopt use of variable-speed-drive motors in air systems.
  - Minimize air system resistance.
  - Optimize the size and cleanliness class of clean space.

7. Conclusions and Recommendations

Through literature reviews and interactions with industry throughout the project, it is clear that there has been some movement in how the industry approaches the design and retrofits of cleanroom facilities used in microelectronics, i.e., toward adopting minienvironments as an emerging technology to achieve effective contamination control. In addition, a similar shift is being observed in facilities for manufacturing drugs, medical device, and other healthcare products, such as using separative devices, bio-safety cabinets, glove boxes, etc. When being appropriately implemented and integrated with large cleanrooms, minienvironments may present many advantages such as better control, facility integration, and lower airflow rates required for a similar production output.

Prior to this case study including a previous study [3][4], there was virtually no quantitative information that was publicly available to provide knowledge or scientific understanding of the energy performance of minienvironments in operation. This case study has further provided quantitative data to characterize the energy performance of minienvironments. The results have been presented to industry stakeholders including an electric utility company, cleanroom facility engineers or managers, and consultants in contamination control (see Appendix on the summary of a workshop).

The following summarizes the conclusions drawn from the study and recommendations for future work.

7.1 Conclusions

This study investigated energy and environmental performance of ISO-Cleanliness-Class-3 minienvironments housed in a traditional, larger ISO-Cleanliness-Class-4 cleanroom used in the microelectronic industry. The measured parameters included electric power demand for minienvironments as compared to cleanrooms of various cleanliness grades. The study also estimated the energy-saving potential of the design, operation, and management of clean spaces when minienvironments were integrated with a traditional, large cleanroom. Based upon the experimental measurements, analysis, and discussion in this case study, the following conclusions are made:

- Energy efficiency levels of the minienvironments in this study were found to be lower when compared with their cleanroom counterparts. Optimal contamination control for minienvironments could be realized by optimizing minienvironment design, regulating airflow rates and/or air-pressure differentials between minienvironment space and its surrounding space to achieve effective and efficient particulate filtration control.

- Minienvironments selected in this study were effective in maintaining particle-concentration levels within what was intended. In addition, the minienvironments exhibited variations in physical sizes, airflow speeds, air-change rates, energy performance index, and electric power density, while air-pressure differentials between minienvironment space and its surrounding space were considerably lower than the suggestions in relevant industrial standard or recommended practice.

- Optimizing the airflow rates for minienvironments and the surrounding cleanroom could result in energy savings. For example, providing the minienvironment with the much smaller volumes of air that must be moved, conditioned, and filtered in a given time. Additional energy savings could be achieved through reducing electric power demand of recirculation systems in the surrounding cleanroom space.

- Adopting the minienvironment concept as a means of contamination control will likely continue to influence the future design, construction, and operation of cleanroom spaces. Successfully applying minienvironments to create large potential energy savings is, however, an emerging concept, if not new, for improved energy efficiency.

7.2 Recommendations

This case study characterized the energy performance of selected minienvironments with the assumption that they were all operating at steady states. In actual applications, it is possible that additional factors, e.g., non-steady-state operation, could make this assumption invalid for certain circumstances. Additional investigation and research is necessary for further understanding the minienvironment technology, and to realize the energy-saving potentials by applying minienvironments to achieve effective cleanroom contamination control.

Recommendations for future investigations and improvement in energy-efficiency practice include the following:

- Develop methods and approaches to determine the cleanliness requirements for contamination control for both minienvironments and the surrounding cleanroom. Optimal cleanliness levels shall meet minimal requirements for controlling particle concentrations, but should not be more stringent than what the process occurring in the cleanroom requires.

- Evaluate acceptable ranges of airflow speeds and air-change rates in minienvironments, and their association with cleanliness levels. Develop optimal airflow rates in the surrounding cleanroom areas, and where possible, reduce supply airflow rates. Using optimal air-change rates will allow designers to lower construction costs as well as to reduce energy costs while maintaining the level of air cleanliness required in cleanroom facilities.

- Investigate and understand acceptable ranges of pressure differential between minienvironments and the surrounding spaces, e.g., optimal airflows through the minienvironments.

- Develop scientific basis of optimal designs and demonstrate energy-saving opportunities for adopting minienvironments as an emerging technology in effective contamination control and develop strategies to improve energy savings.

- Develop additional case studies or benchmarking studies to quantify energy and environmental performance of minienvironment systems in various applications, e.g., separative device, enclosure, glove box, bio-safety cabinet used in different industries.

- Develop scientific understanding of how to improve the effectiveness and energy efficiency of the minienvironment systems including energy efficient fan-filter units and optimal speed control in such applications. Develop approaches and methods to understand and improve filtration effectiveness and energy efficiency.

- Develop information and understanding of dynamic behavior of particle contaminants in minienvironments and the surrounding cleanroom as a function of relevant parameters (e.g. particle size, volatile organic compound, airborne molecular contaminants, etc.) and airflow conditions (e.g., design, pressure differential, human movement, product movements, heat), and its impact on energy management.

- Use computational-fluid-dynamics (CFD) modeling, particle-count monitoring, and experiments as the tools to evaluate the environmental and energy performance of minienvironments, and to assist in the design process, qualification, and validation in actual cleanroom/minienvironment planning and operation.

- Develop and disseminate scientific information on particulate contamination, including temporal and spatial dispersion of contaminants, effectiveness of HEPA/ULPA filters in particle-concentration control, and its relevance to the energy performance of air-delivery systems in clean spaces.

- Develop market information or surveys on the industries and scientific communities using minienvironment concepts and investigate such applications to quantify energy-savings potential.

- Integrate the new knowledge and information in national or international guiding documents in future editions, such as ANSI-accredited IEST's Recommended Practice CC028.1 – Minienvironments, to maximize its usefulness and to benefit sustainable development of the industries using minienvironments.

8. Acknowledgements

The project was funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program (http://www.energy.ca.gov/). This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

9. References

[1] The Institute of Environmental Sciences and Technology (IEST). 2002. IEST-RP-CC028.1: Minienvironments. Handbook of Recommended Practices, Contamination Control Division. 5005 Newport Drive, Suite 506, Rolling Meadows, IL 60008-3841, USA.

[2] International Organization for Standardization (ISO). 2004. ISO 14644-7, Cleanrooms and associate controlled environments - Part 7: Separative Devices (clean air hoods, glove boxes, isolators and minienvironments). The Institute of Environmental Sciences and Technology (IEST). 5005 Newport Drive, Suite 506, Rolling Meadows, IL 60008-3841, USA.

[3] Xu, T. 2005. Opportunities for sustainable design and operation of clean spaces: A Case Study on Minienvironment System Performance. LBNL Final Report, LBNL-56312.

[4] Xu, T. 2005. Investigating the performance of a minienvironment system. Proceedings of 51st ESTECH Conference, Schaumburg, Illinois, The Institute of Environmental Sciences and Technology (IEST). LBNL-56767.

[5] Xu, T. 2006. "A Study on the Operation Performance of a Minienvironment System." *Journal of the IEST*, Lawrence Berkeley National Laboratory Report, LBNL-58743.

[6] International Organization for Standardization (ISO). 1999. ISO 14644-1, Cleanrooms and associate controlled environments - Part 1—Classification of Air Cleanliness**.** 18 Pages, The Institute of Environmental Sciences and Technology (IEST). 5005 Newport Drive, Suite 506, Rolling Meadows, IL 60008-3841, USA.

[7] Tannous, G. 1996. Airflow simulation in a minienvironment, Solid State Technology, July 1996, pp. 201-207.

[8] Tannous, G. 1997. Optimization of a minienvironment design using computational fluid dynamics. Journal of the IEST, Vol. 40 (1): 29-34.

[9] Olim, M. 1998. "Optimizing minienvironment air flow patterns using CFD," FSI International Chaska, Minn., Semiconductor International. Cleanroom Magazine.

[10] Hu, S.C., Y.K. Chuah, M.C. Yen. 2002. Design and evaluation of a minienvironment for semiconductor manufacture processes, Building and Environment, Vol. 37 (2): 201-208. February 2002, Elsevier Science.

[11] Hu, S. C., R.H. Tong. 2005. Particle dynamics in a front-opening unified pod/load port unit minienvironment in the presence of a 300 mm wafer in various positions, Aerosol Science and Technology, Vol. 39 (3): 185-195.

[12] Hu S.C.; Chuah Y.K. Deterministic simulation and assessment of air-recirculation performance of unidirectional-flow cleanrooms that incorporate age of air concept.

Building and Environment, April 2003, vol. 38, no. 4, pp. 563-570(8)

[13] International SEMATECH. 1999. Integrated minienvironment design best practices, Technology Transfer # 99033693A-ENG. International SEMATECH, Inc., Austin, Texas.

[14] International Organization for Standardization (ISO). 1999. ISO 14644-1, Cleanrooms and associate controlled environments - Part 1—Classification of Air Cleanliness**.** 18 Pages, The Institute of Environmental Sciences and Technology (IEST). 5005 Newport Drive, Suite 506, Rolling Meadows, IL 60008-3841, USA.

[15] International Organization for Standardization (ISO). 2001. ISO 14644-4, Cleanrooms and associate controlled environments - Part 4: Design, Construction and Start-up. The Institute of Environmental Sciences and Technology (IEST). 54 Pages. 5005 Newport Drive, Suite 506, Rolling Meadows, IL 60008-3841, USA.

[16] International Organization for Standardization (ISO). 2002. ISO 14644-3, Cleanrooms and associate controlled environments - Part 3: Metrology and Test Methods. 65 Pages. The Institute of Environmental Sciences and Technology (IEST). 5005 Newport Drive, Suite 506, Rolling Meadows, IL 60008-3841, USA.

[17] Shu, S. 1995. The impact of minienvironment of Fab capacity and yield, Proceedings of the 41st Annual Technology Meeting: 118-123, Anaheim, Calif. The Institute of Environmental Sciences and Technology (IEST).

[18] Bonora, Anthony. 1997. There and back again. Solid State Technology. May 1997 edition, Penn Well Publishing Company.

[19] Rothman, L., R. Miller, R. Wang, T. Baechle, S. Silverman; and D. Cooper. 1995. SEMATECH minienvironment benchmarking project. Journal of the IEST, May/Jun 1995, Vol. 38 (3): 36-41.

[20] Xu, T. 2003. Performance evaluation of cleanroom environmental systems. Journal of the IEST, Vol. 46: 66-73. Institute of Environmental Sciences and Technology (IEST), Rolling Meadows, Illinois. Lawrence Berkeley National Laboratory Report, LBNL-53282. Berkeley, Calif., USA

[21] Xu, T. 2004. Considerations for efficient airflow design in cleanrooms, Journal of the IEST, Volume 47: 24-28, October 2004, Institute of Environmental Sciences and Technology (IEST), Rolling Meadows, Illinois. Lawrence Berkeley National Laboratory Report, LBNL-55970. Berkeley, Calif., USA

[22] Power Sight Energy Analyzer, Summit Technology, 2717 N. Main St., Suite 15, Walnut Creek, CA 94596; http://www.summittechnology.com/pk314m.shtml

[23] ADM-860C Airdata Multimeter - Backpressure Compensated Air Balance System. Shortridge Instruments, Inc., 7855 East Redfield Road, Scottsdale, Arizona 85260. http://www.shortridge.com

[24] MetOne laser airborne particle counter, Hach Ultra Analytics, http://www.particle.com/

[25] Xu, T., and M. Jeng. 2004. Laboratory evaluation of fan-filter units' aerodynamic and energy performance. Journal of the IEST, Volume 47: 116-120. The Institute of Environmental Sciences and Technology (IEST), 5005 Newport Drive, Rolling Meadows, Illinois, October 2004.

[26] British Standard 5295 (1989). BSI Standards, 389 Chiswick High Road, London W44 AL, UK.

10. Appendices

10.1 Appendix - Schematic Diagram of Minienvironments in the Cleanroom

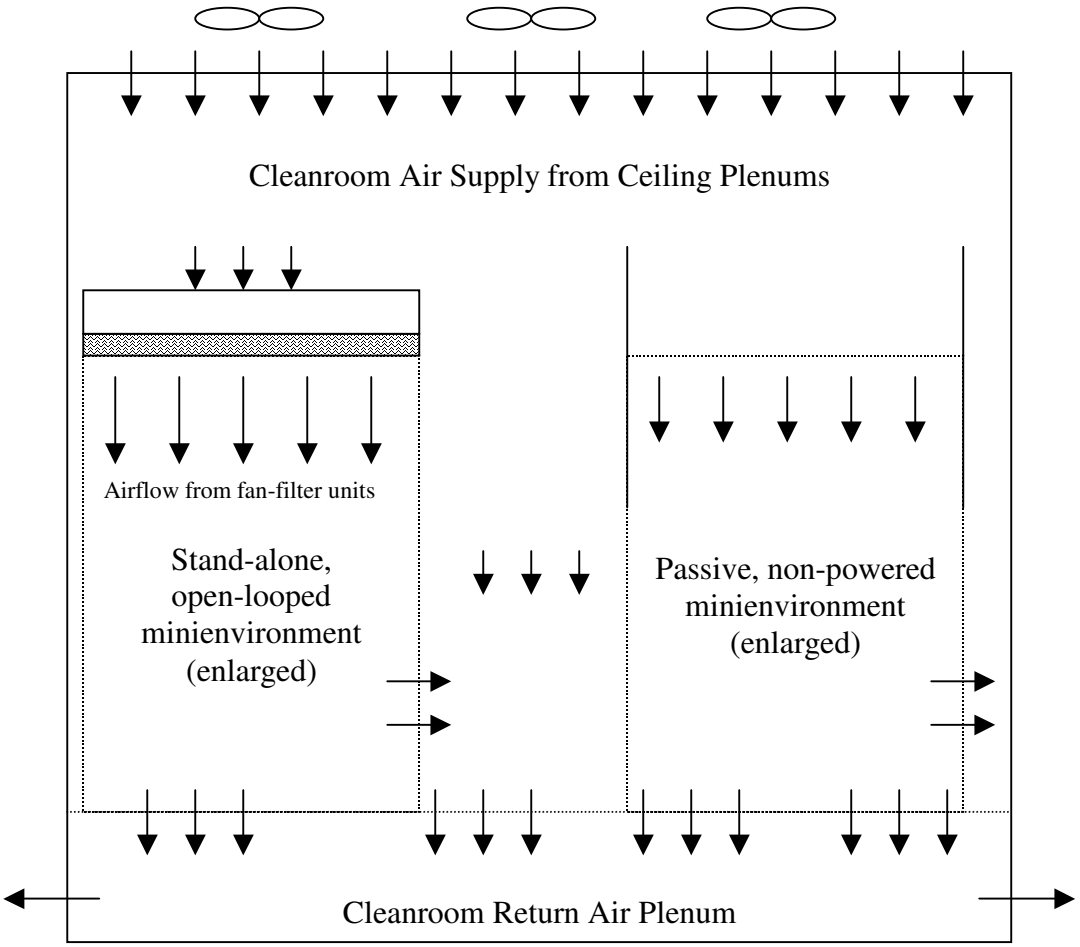

10.2 Appendix - Case-Based Estimate of Energy-savings

Case Study Based	Floor Area (ft^2) / (m^2)	Total Electric Power (kW)	Electric Power Density (W/ft^2)/(W/m^2)	Estimated Energy-savings
ISO-Cleanliness-Class-3 Minienvironment	424/39	12.0	28.3/304	-
ISO-Cleanliness-Class-4 Cleanroom	3680/342	37.5	10.2/110	-
ISO-Cleanliness-Class-4 Cleanroom and ISO-Cleanliness-Class-3 Minienvironment	3680/342	49.5	13.4/145	Base case
ISO-Cleanliness-Class-4 Cleanroom and Improved ISO-Cleanliness-Class-3 Minienvironment (by 50%)	3680/342	43.5	11.8/127	12%
Improved ISO-Cleanliness-Class-4 Cleanroom (by 33%) and ISO-Cleanliness-Class-3 Minienvironment	3680/342	37.0	10.1/108	25%
Non-Cleanroom and ISO-Cleanliness-Class-3 Minienvironment	3680/342	12.0	5.3/57	61%

10.3 Appendix - Design-Based Estimate of Energy-savings

Design based	Floor Area (ft^2) / (m^2)	Total Electric Power (kW)	Electric Power Density (W/ft^2) / (W/m^2)	Estimated Energy-savings
ISO-Cleanliness-Class-3 Minienvironment @12% occupancy	424/39	12.0	28.3/304	-
ISO-Cleanliness-Class-3 Cleanroom	3680/342	139.8	38.0/409	Base case
ISO-Cleanliness-Class-4 Cleanroom	3680/342	109.3	29.7/320	-
ISO-Cleanliness-Class-5 Cleanroom	3680/342	48.6	13.2/142	-
ISO-Cleanliness-Class-4 Cleanroom and ISO-Cleanliness-Class-3 Minienvironment	3680/342	121.3	33.0/355	13%
ISO-Cleanliness-Class-5 Cleanroom and ISO-Cleanliness-Class-3 Minienvironment	3680/342	60.6	16.5/177	57%
Non-Cleanroom and ISO-Cleanliness-Class-3 Minienvironment	3680/342	19.3	5.3/57	86%

10.4 Appendix - Summary of the Workshop on Minienvironment

LBNL developed and conducted a workshop based upon the case studies on minienvironments. The workshop was titled "Trends in Cleanroom Technology and Energy Savings Opportunities," and was organized in collaboration with PG&E and the city of San Jose. It was held in San Jose Martin Luther King Library, on September 13, 2005. The subtitle of the workshop was termed as "Best Practices for Energy Efficient Design, Construction and Operation of Minienvironment/Cleanroom."

10.4.1 Outline of the Minienvironment Workshop

Lawrence Berkeley National Laboratory, the California Energy Commission, Pacific Gas and Electric Company hosted a special workshop on the best practices for energy efficient clean spaces, using minienvironments that are gaining popularity in a wide range of industries.

A minienvironment is a localized (usually minimized) clean environment created by an enclosure to isolate a product or process from the surrounding environment. It carries various names, such as Separative device, safety cabinet, isolator, etc. The goal of the Minienvironment Workshop was to provide a forum for sharing knowledge and to stimulate discussion among participants about emerging technologies and strategies of achieving energy efficiency while maintaining effective contamination control in cleanrooms.

This workshop was designed for suppliers, end-users, designers, facility managers, consultants, and strategic managers. The participants included utility managers, facility engineers, facility managers, consultants, and users.

10.4.2 Topics Covered in the Minienvironment Workshop

- A preview of R&D activities in Minienvironment and Cleanroom Contamination Control sponsored by California Energy Commission PIER program. Presenter: Paul Roggensack, PE, California Energy Commission, Calif.

- Case Studies on Minienvironment Energy Performance: Approaches, Findings, Opportunities, and Recommendations. Presenter: Dr. Tengfang Xu, PE, Lawrence Berkeley National Laboratory, Berkeley, Calif.

- Best Practice and Lessons Learned from Minienvironment Planning and Installation. Presenter: Michael D. Jue, PE, Jazz Semiconductor, Newport Beach, Calif.

- Energy Efficiency Programs in Calif. Presenter: Bill Dunckel, PG&E, San Francisco, Calif.

10.4.3 Participant Forum on Minienvironments

Forum Leader: Dr. Tengfang Xu, PE, Lawrence Berkeley National Laboratory, Berkeley, Calif.

The participants as a group discussed the industrial trends, emerging cleanroom technologies, and strategies for better efficiency in cleanrooms. The concept of Efficient and Effective (E^2) Minienvironment proposed and presented by Dr. Xu was well received and supported by the participants.

The team discussed the significance of and explored the future opportunities in energy efficiency and sustainable development in controlled environments. The team provided suggestions about how to make LBNL's research on E^2 minienvironment known to a wider range of industries. The consensus from the team discussion included 1) Enhancing the visibility of LBNL's research and case studies will be necessary through further marketing to various industries; 2) Collaboration among industries including users, suppliers, utility companies, and government entities such as CEC will be important to facilitate the applications of research and technology, and 3) Further R&D investigations and activities as outlined in the recommended list will be necessary and important for the future success in promoting E^2 minienvironments in cleanroom applications.

10.4.4 Workshop Presentation and Recommendations from the Case Studies (Power Point slides)

Attachment V: Side-by-side Fume Hood Testing

Lawrence Berkeley National Laboratory

California Energy Commission

"Side by Side" Fume Hood Tests Using ASHRAE 110 and "Human As Mannequin" (HAM) Test Methodology to Compare Performance of a Jamestown Conventional Fume Hood and a LBNL High Performance Hood

March 21, 2005

ECT, Inc.

Exposure Control Technologies, Inc.
231 C East Johnson Street
Cary, NC 27513
Phone: (919) 319-4290
Fax: (919) 319-4291

CONTENTS

INTRODUCTION AND SUMMARY 3

SUMMARY OF RESULTS 3

METHODOLOGY 6

HOOD PERFORMANCE TESTS 7

RECOMMENDED PERFORMANCE CRITERIA 12

TEST RESULTS 13

FACE VELOCITY TESTS 13

CROSS DRAFT VELOCITY TESTS 13

SMOKE TESTS 13

STATIC MANNEQUIN – TRACER GAS TESTS 14

HUMAN AS MANNEQUIN TEST 14

ANALYSIS OF RE-ENTRAINMENT AND IMPACT ON HOOD PERFORMANCE 14

List of Tables

Table 1. ECT Qualitative smoke rating chart 12

Table 2 Summary of Static Mannequin Tracer Gas Test Data 16

Table 3 Summary of tracer gas data and sequence time for HAM tests for each hood and test location 24

Table 4 Calculation of exhaust flow, stack velocity and discharge concentration for each hood system 43

INTRODUCTION and SUMMARY

Exposure Control Technologies, Inc. was contracted by the Lawrence Berkeley National Laboratory (LBNL) to conduct a series of modified ASHRAE 110 tests to compare performance of a conventional 6-ft, Jamestown Isolator fume hood (standard hood) and a 6-ft, high performance, fume hood designed by LBNL. With the assistance of Geoffrey Bell and Doug Sullivan of LBNL, the performance tests were conducted on January 11-13, 2005 in Building 63 on the LBNL campus located in Berkeley, CA. The hoods were installed side by side and exhausted by independent exhaust systems of similar design.

Both hoods were tested with the vertical sashes raised to full open that provided nearly equivalent opening areas of approximately 12 ft^2. The standard hood was tested at an exhaust flow of approximately 1200 cfm and the LBNL fume hood had an exhaust flow of approximately 670 cfm resulting in measured and calculated inflow face velocities of 98 fpm and 56 fpm, respectively. The performance tests included:

- Face velocity tests,
- Cross draft velocity tests,
- Airflow visualization (smoke tests),
- Static Mannequin tracer gas tests and
- "Human as Mannequin" tracer gas tests.

The tests were used to define the operating conditions and compare performance of the hoods when operating under similar environmental conditions. Where applicable, the tests were conducted according to the ANSI/ASHRAE 110 "Method of Testing Performance of Laboratory Fume Hoods." However, the "Human as Mannequin" tests were unique and conducted according to procedures developed by LBNL in consultation with the California Occupational Safety and Health Administration (CAL/OSHA). The "Human as Mannequin" tests involved standing in place of the mannequin and conducting a series of choreographed procedures while generating tracer gas in the hood and sampling air near the breathing zone. Concentrations detected in the breathing zone could indicate escape from the hood and potential for inhalation and exposure. Therefore, the results of the static mannequin (SM) tracer gas tests and the "Human as Mannequin" (HAM) tracer gas tests were used to evaluate and compare hood performance. Refer to the section entitled "Methodology" for a more complete description of test procedures and performance criteria.

Summary of Results

The results indicate that when tested side by side and under the prevailing operating conditions, the performance of the LBNL hood operating at 43% less exhaust flow was equivalent to the performance of the standard or conventional fume hood. Average tracer gas concentrations from the static mannequin tests and the HAM tests were significantly less than 0.1 ppm recommended by the American Conference of Governmental Industrial Hygienists (ACGIH) and the AIHA/ANSI Z9.5 American National Standard for Laboratory Ventilation. Average concentrations for both hoods were less than 0.05 ppm and the differences between the hoods were within the limits of instrument accuracy and expected error considering the variable challenge posed by exhaust re-

entrainment and other factors. Refer to Figure 1 for results of the static mannequin tests and Figure 2 for results of the HAM tests.

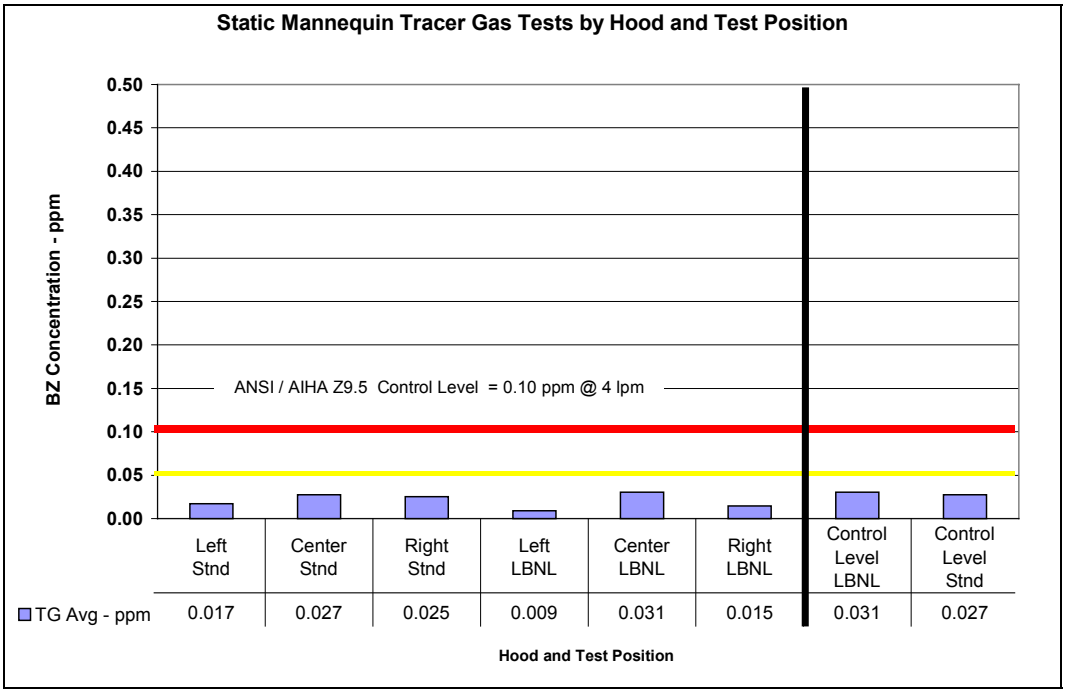

Figure 1 Results of static mannequin tracer gas test results showing average concentrations and control level for each hood and test location.

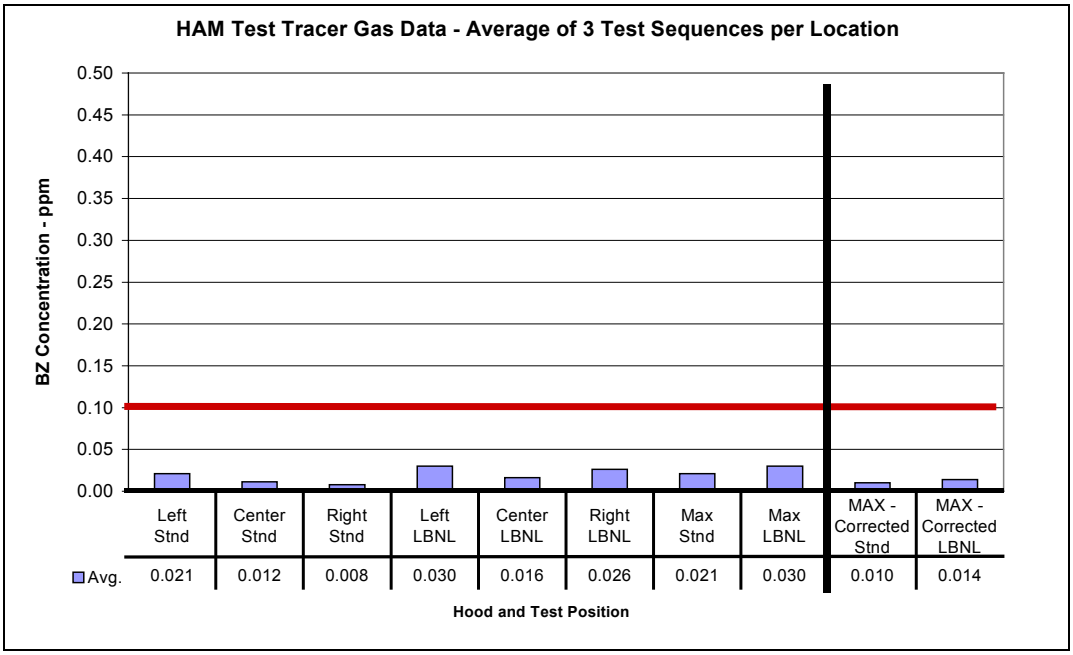

Figure 2 Results of HAM tracer gas tests showing average concentrations and control level for each hood and test location.

Despite the low average concentrations, the test results were very revealing and demonstrated the necessity of tracer gas tests to evaluate performance of a hood system. In addition, the HAM test

with slight modification has particular value as a training tool to better understand the impact of work practices on hood containment.

Refer to the section "Results" for a more complete description of test results, tabulated summaries of data and plots of breathing zone concentrations for each tracer gas test.

Data contained in this report are indicative of performance of the laboratory hoods and ventilation systems at the time and date of the investigation. Any questions concerning information or conclusions contained within this report should be referred to Exposure Control Technologies. We welcome your questions, comments and suggestions concerning the report and we thank the Lawrence Berkeley National Laboratory and the California Energy Commission for the opportunity to participate in this project.

Testing and evaluation were performed by:

Thomas C. Smith
President
Exposure Control Technologies, Inc.

METHODOLOGY

The objective of the study was to compare performance of a commercially available fume hood of conventional design and a high performance, fume hood designed by LBNL. Both hoods were fabricated by Jamestown Metal Products and installed side by side in Building 63 on the LBNL campus located in Berkeley, CA. The hood representing conventional design (Standard) was a 6-ft, Jamestown Isolator Fume Hood having a vertical sliding sash that enabled a maximum opening area of approximately 12.06 square feet (62 inches wide by 28 inches high). The standard hood was equipped with all components of a typical conventional fume hood including an airfoil sill, three slot baffle and angled entries. The face velocity for the standard hood was approximately 98-fpm with a calculated exhaust flow of approximately 1180-cfm.

The LBNL hood was of similar size but of different design. The opening area with the vertical sash full open was approximately 12.01 square feet (61.75 inches wide by 28 inches high). Air is supplied at the top and bottom of the opening using four small fans that draw air from the room to provide an air curtain (air divider) near the plane of the sash. In theory, the air curtain helps to prevent escape of materials generated within the hood and thereby reduces the need for high face velocities to provide capture and containment. The intent of this design is to reduce energy use by reducing the total volume of exhaust required to maintain containment. The exhaust flow for the LBNL hood was approximately 670 cfm or 43% less flow than the standard hood. However, the supply of air at the plane of the sash obviates the typical measurement of face velocity due to the non-perpendicular velocity vectors produced by the air curtain. Dividing the exhaust flow by the opening area provides an equivalent face velocity of approximately 56-fpm.

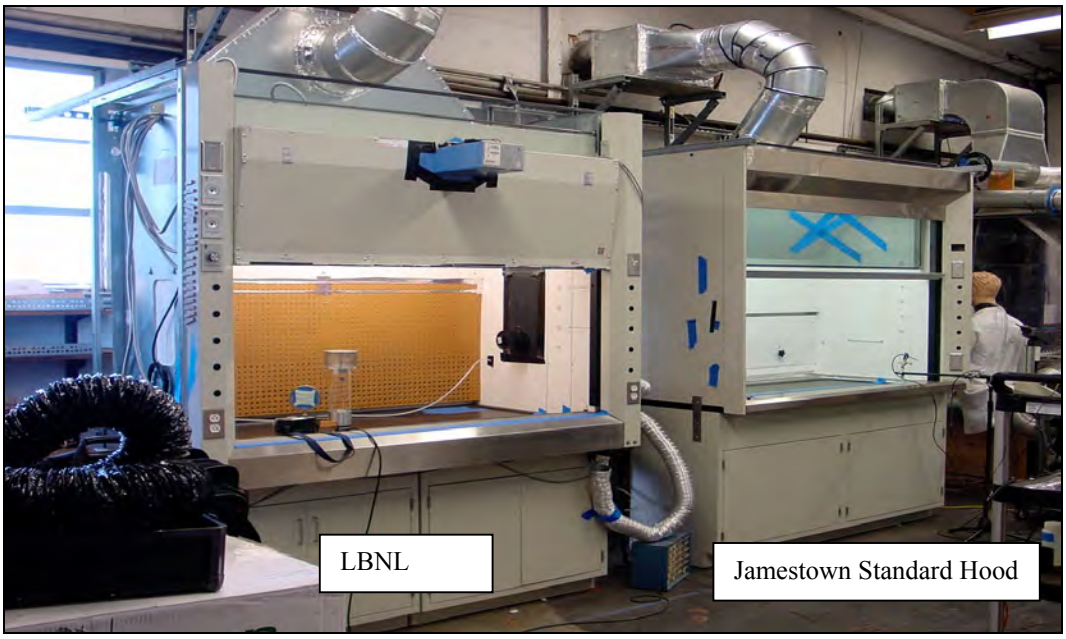

Figure 3 Photo of the LBNL hood and the Jamestown Standard Hood installed for side by side testing in the LBNL test lab.

Hood Performance Tests

The performance tests were conducted on both hoods over several long days from January 11, 2005 through January 13, 2005. The performance tests included:

- Measurement of Face Velocity
- Measurement of Cross Draft Velocities
- Visualization of Airflow Patterns
- Measurement of Tracer Gas Containment
  - Static Mannequin Tests
  - Human as Mannequin Tests

The test lab is located in a metal-sided building that was cold in the morning and warmed as the day progressed. It was suspected that the test results could be affected by the changing climatic conditions. To minimize differences between operating conditions and reduce potential variables associated with changing climatic conditions, attempts were made to conduct each test sequentially on both hoods. With the exception of the Human as Mannequin Test, the tests were conducted according to procedures described in the ANSI/ASHRAE 110 "Method of Testing Performance of Laboratory Fume Hoods." The Human as Mannequin Tests were conducted according to the LBNL procedure developed in consultation with CAL/OSHA. The individual test procedures are briefly described below.

Face Velocity Tests

Face velocity is measured to determine the speed of air entering the hood opening. The face velocity tests were performed by dividing the sash opening into equal area grids of no more than one square foot and measuring the velocity using a fixed probe located at the center of each grid (Figure 4). Air velocities were recorded every second for twenty seconds at each grid location using a calibrated thermoanemometer. The means of the grid velocities were averaged to determine the overall average face velocity. Although face velocities were measured on the LBNL hood, the perimeter air supply prevented accurate measurement and determination of average face velocity. Therefore, the equivalent face velocity for the LBNL hood was determined by dividing the exhaust flow by the opening area. The supply velocities at the top and bottom of the opening were not measured as it was assumed that the supply flow was adjusted to LBNL specifications.

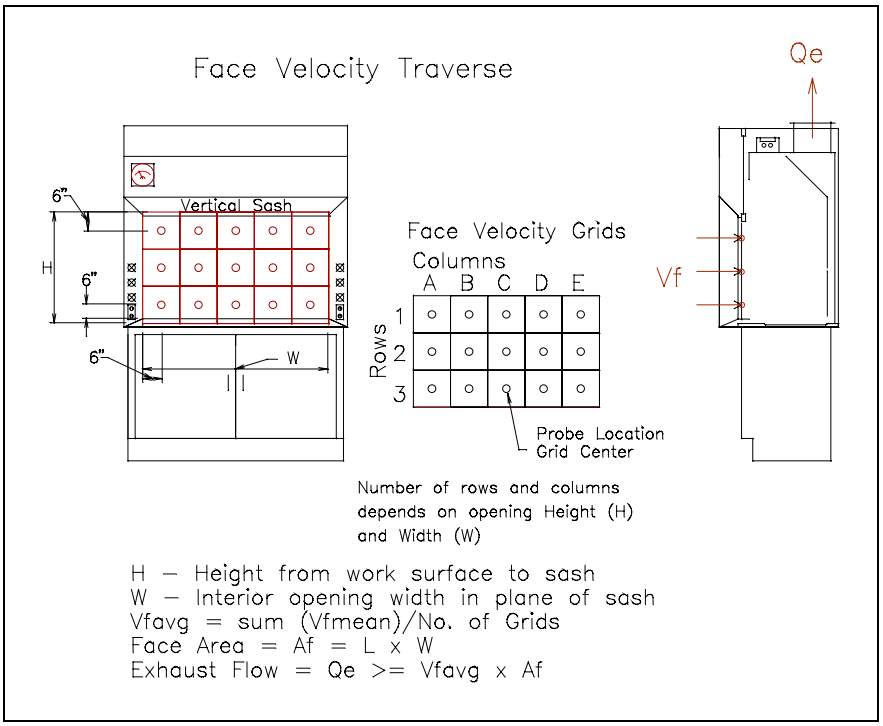

Figure 4: Diagram of hood opening showing face velocity traverse locations

Cross Draft Tests

Room air currents can pose a significant challenge to hood performance. Cross draft air velocities are measured to provide an indication of comparable operating conditions. Vertical and horizontal cross drafts were measured in front of the hood openings using two calibrated thermoanemometers. Measurements were made with the probes positioned parallel to the side edge of the hood approximately five feet above the floor and approximately 18 inches in front of the hood opening. Cross draft velocities were recorded every second for thirty seconds at each test location.

Airflow Visualization Tests

Smoke was generated in the hood and around the periphery of the opening using a Rosco Fog Generator and TiCl4 smoke sticks. Smoke was generated along the periphery of the opening and inside the hood at approximately six inches inside the plane of the sash. Smoke patterns were observed to identify escape and reverse flow zones. Observations were rated subjectively based on the subjective ECT rating table provided in the following section "Recommended Performance Criteria."

Static Mannequin Tracer Gas Tests

The static mannequin tracer gas containment tests were conducted to evaluate and quantify the potential for escape with a mannequin located in front of the hood opening. A mannequin is used

to simulate the presence of a hood user. The mannequin was positioned at the left, center and right side of the hood opening with the nose approximately three inches in front of the plane of the sash. The mannequin height was approximately 67 inches.

A tracer gas, Sulfur Hexafluoride, was released at 4 liters per minute (lpm) through an ASHRAE 110 ejector. The ejector was located approximately 6 inches behind the plane of the sash directly in front of the mannequin. Breathing zone concentrations were measured every second for five minutes in each test location using a calibrated ITI Model 200 Leakmeter. The Leakmeter was calibrated against a calibration gas standard immediately prior to the start of tests. The limit of detection was 0.01 ppm with a maximum scale reading of approximately 1.9 ppm (actual concentrations may be higher than the maximum scale reading of approximately 1.9 ppm).

The static mannequin tracer gas tests were repeated various times at the center test position to evaluate stability of performance, impact of environmental conditions and to characterize exhaust re-entrainment. The results of all tests on the center hood were averaged for the comparison of hood performance. The highest average concentration found at any test location is reported as the control level.

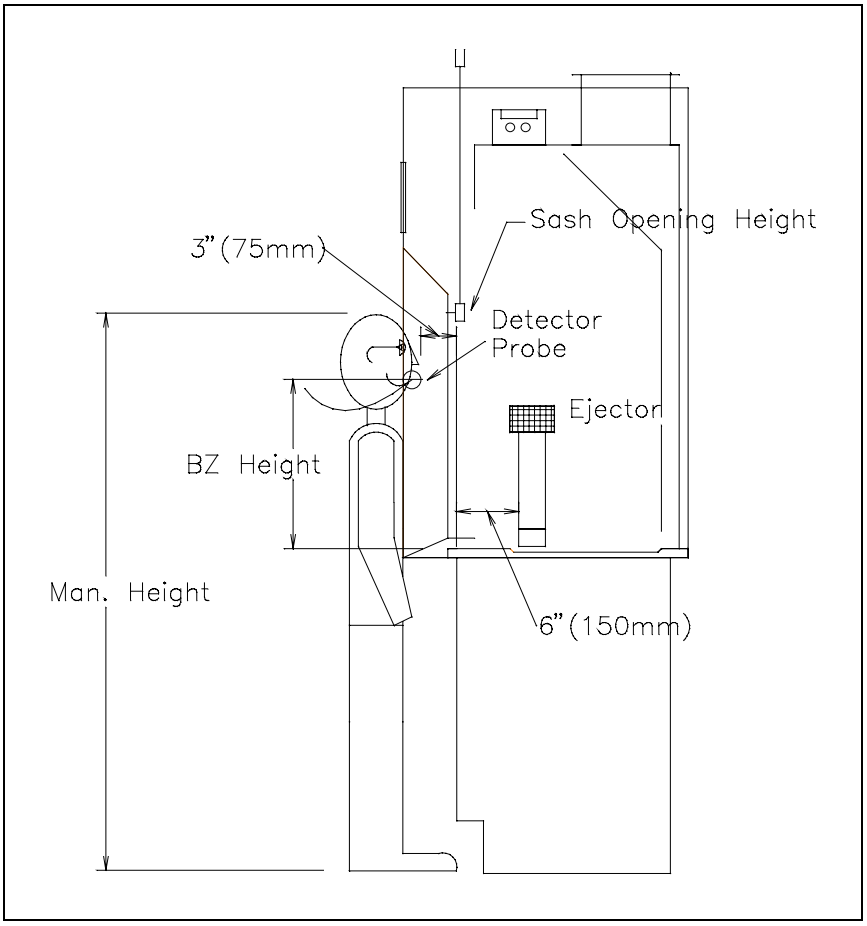

Figure 5 Configuration of mannequin at hood opening during tracer gas tests.

Human as Mannequin Tracer Gas Tests (HAM)

The HAM tests were conducted to provide a dynamic challenge to hood performance by using a person in place of a mannequin and conducting a series of choreographed activities in the hood. Tracer gas is generated using the ASHRAE 110 ejector in the same locations and at the same rate as the static mannequin tests. A short length of tubing is connected to the tracer gas detector and fixed near the mouth of the person using the microphone boom on a hands-free telephone headset. The photos in Figures 6 and 7show Thomas Smith of ECT conducting the HAM test.

The HAM tests involve conducting a series of choreographed activities using objects located within the hood. The objects consist of a beaker, two bottles and a small box that are arranged in specific patterns depending on the test location and manipulated according to a scripted sequence (See photo of objects in Figure 8). The timed sequence of activities includes:

1. Stand at hood opening with arms to side.
2. Insert and remove hands and arms
3. Move objects from six inch line to twelve inch line
4. Exchange position of objects
5. Transfer liquid from 250 ml bottle to 500 ml beaker
6. Grasp bottle, remove from hood and rotate 90 degrees from opening.

Each sequence of activities is conducted over a period of approximately 150 seconds. For more information refer to the LBNL HAM Test procedures described in the September 8, 2004 LBNL report entitled *Side by Side Fume Hood Testing: Human-as-Mannequin Report*.

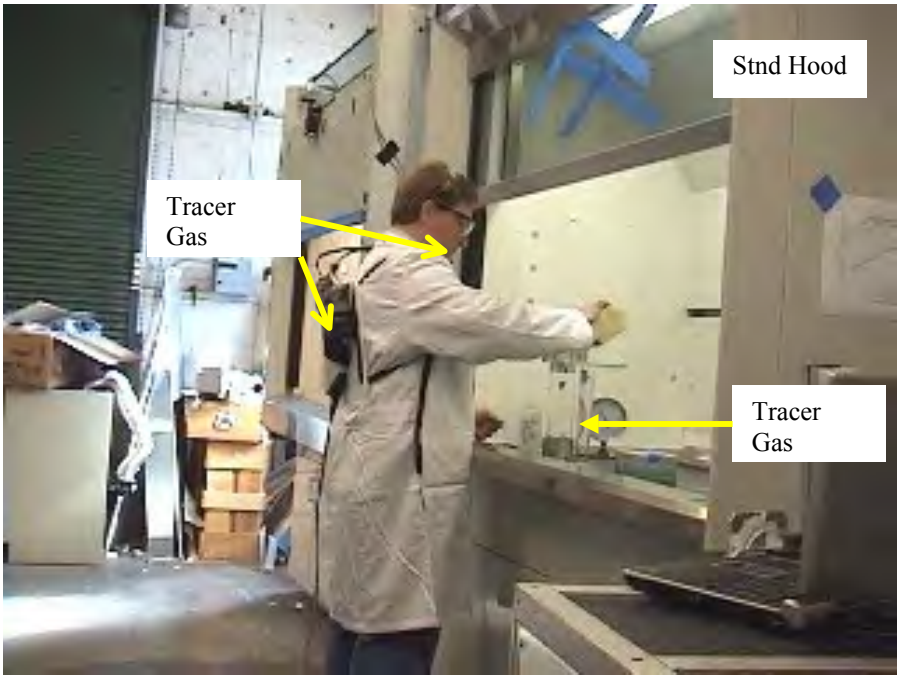

Figure 6 Photo of Tom Smith of ECT pouring liquids during a HAM test.

Each test sequence is timed and concentrations of tracer gas are measured every second for the duration of the test sequence. The sequence of tests is repeated three times at the left, center and right test locations. The data from the three test sequences are combined and analyzed to determine the average breathing zone concentration for each test location. For academic purposes, the procedures were reversed at the completion of each 150 seconds sequence to extend the test over five-minutes so that data could be compared to the static mannequin tests. Future modification of the HAM test procedure should consider a sequence of activities that require five minutes to complete.

Figure 7 Photo of Tom Smith moving objects in hood during HAM Test.

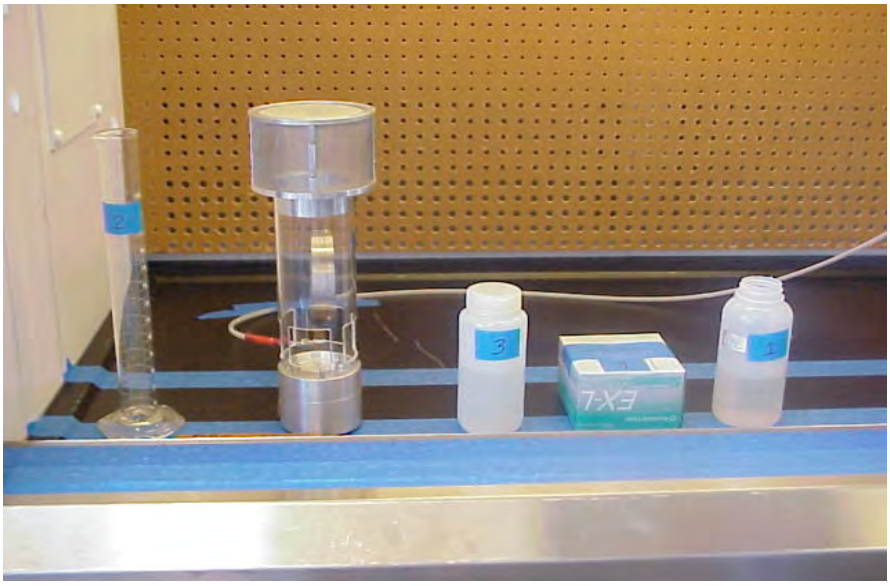

Figure 8 Photo of objects located in hood and manipulated during HAM Test procedures

RECOMMENDED PERFORMANCE CRITERIA

The following criteria were gathered from industry consensus standards (ANSI Z9.5, ACGIH and ASHRAE) or used by ECT in the evaluation of hood performance. The criteria are provided here to assist with evaluating test data and define the criteria used to evaluate and compare performance. The criteria apply to the following tests:

Smoke Visibility Test - Hood must provide complete containment of the smoke generated within the hood. Containment is determined visually and in the absence of the mannequin. ECT rates the observations of airflow patterns as a qualitative judgment of airflow distribution according to the following rating guide:

Table 1. ECT qualitative smoke visualization rating chart

RATING	DESCRIPTION
FAIL	• Smoke was visually observed escaping from the hood
POOR (Low Pass)	• Reverse flow of smoke is evident within six inches of the plane of the sash when generated at least six inches behind the plane of the sash. • Lazy flow into hood along openings • Slow capture and clearance – greater than two minutes for clearance • Observed potential for escape
FAIR (Pass)	• Some reverse flow in hood not within six inches of opening • Smoke is captured and clears readily from interior of hood – less than two minutes • No visible escape
GOOD (High Pass)	• Good capture and relatively quick clearance – approximately 1 minute or less • No Reverse Flow Regions • No Lazy Flow • No visible escape

Face Velocity Test – Ventilation system should be capable of maintaining stable exhaust flow (no more than ± 5% variation) and achieve the target average face velocity within 5 fpm. The face velocity readings should not vary more than 20% of the mean between grid locations.

Cross Draft Test – Room air currents within the lab should not exceed 30 fpm within 24 inches of the hood opening.

Tracer Gas Challenge Test - Hood should provide containment of tracer gas below an average of 0.1 ppm measured for a period of five minutes in the breathing zone of a stationary mannequin positioned directly in front of the ejector location.

Human as Mannequin Test – There are no industry standards or published criteria for HAM tests. During consultation with Cal/OSHA, it was agreed that criteria for a high performance hood not meeting the face velocity requirements was to perform at least as well as a standard hood meeting the face velocity requirement or to provide containment below an average of 0.1 ppm.

TEST RESULTS

The results of the performance tests indicated that both hoods met industry recommended criteria for ASHRAE 110 tracer gas tests (static mannequin tests) and provided comparable performance during the HAM tests. The average concentrations for all tracer gas tests were below 0.1 ppm with average concentrations for many tests equal to or below the minimum limit of detection (approximately 0.01 ppm). The test results are summarized below.

Face Velocity Tests

The standard hood was operating at a face velocity of 98 fpm and a calculated exhaust flow of approximately 1,180-cfm. The LBNL hood was operating at approximately 670 cfm with a calculated face velocity of 56 fpm. The exhaust flow for the LBNL hoods was approximately 43% less than the Standard Hood.

Cross Draft Velocity Tests

The hoods were tested side by side in the test lab. The cross draft velocities were comparable with average velocities less than 30 fpm near the hood openings. The highest velocity was 25 fpm in the horizontal direction in front of the LBNL hood and the highest velocity was 22 fpm in the vertical direction near the Standard hood.

Smoke Tests

Both hoods provided containment of smoke generated at least six inches inside the plane of the sash. According to the subjective ECT rating scheme, the Standard hood rated “low pass” during both the low volume and the high volume smoke tests. The LBNL hood rated “low pass” during the low volume challenge and had a rating of “pass” during the high volume challenge.

Observation of the airflow patterns in the LBNL hood indicated reasonably uniform flow and capture at the baffle perforations and slots. The hood received a low pass rating during the low volume challenge as some minor reverse flow was observed within six inches of the opening when generated at the work surface (primarily due to the recess in the work surface). In addition, a vortex was observed inside the top of the hood with relatively slow but satisfactory clearance. Entrainment of smoke in the vortex region required generation of smoke in the top of the hood.

Observation of smoke patterns in the Standard hood indicated rather poor airflow patterns with considerable reverse flow above the work surface and some eddying above the airfoil sill. In addition, a very turbulent vortex was observed behind the raised sash with significant eddying observed below the sash pull. Furthermore, smoke generated at the work surface was entrained in the vortex region at the top of the hood.

Static Mannequin – Tracer Gas Tests

Table 2 (below) provides a summary of the results from the static mannequin tracer gas tests and Figures 9-15 show the plots of breathing zone concentrations for each test. The results indicate that both hoods met the performance criteria with average concentrations less than 0.1 ppm. The highest concentration for both hoods occurred at the center test position producing equivalent control levels of 0.031 ppm for the LBNL hood and 0.027 ppm for the Standard Hood. The difference between average concentrations is not significant as the results are well within expected error considering the minimum limit of detection for the tracer gas detector was approximately 0.01 ppm and the results were confounded by re-entrainment of exhaust that varied depending on ambient wind conditions. Re-entrainment caused higher background concentrations during tests on the LBNL hood and slightly higher average concentrations. Refer to the section below entitled "Analysis of Re-entrainment" for more information about the impact of re-entrainment on test results.

Human As Mannequin Tests

The results of the HAM tests are provided in Table 3 and the plots of breathing zone concentrations are shown in Figures 16-33. The results indicate that the LBNL hood provided comparable performance to the Jamestown Standard hood during the HAM tests. The maximum series averages for the LBNL hood and the Standard hood were 0.030 ppm and 0.021 ppm, respectively. To account for the bias caused by re-entrainment; a correction factor was applied to both hoods resulting in a maximum average breathing zone concentration of 0.014 ppm for the LBNL hood and 0.010 ppm for the Jamestown Standard hood. Average concentrations found during the HAM tests were less than the static mannequin tests due to the impact of exhaust re-entrainment and the limited duration of the test sequences.

Analysis of Re-entrainment and Impact on Hood Performance

Although all of the tracer gas results indicated that average breathing zone concentrations were less than 0.10 ppm, re-entrainment of exhaust caused the average concentrations to be higher on the LBNL hood than the Standard hood. In addition, the magnitude of re-entrainment and the impact on average concentrations varied depending on the ambient wind conditions. A limited analysis of several tests conducted on both hoods was undertaken to determine the significance of the difference between average concentrations for the two hoods. The results were used to determine a correction factor that was applied to the results of the HAM tests.

The exhaust systems serving the hoods were of similar design and had the same stack height and stack discharge diameter. However, the LBNL hood was operating at approximately 43% less exhaust flow than the standard hood. The reduced flow resulted in a higher discharge concentration and a 43% lower stack velocity (Refer to Table 4). The increasing background concentrations were attributed to re-entrainment of the exhaust and the variability was attributed to changes in wind speed and wind direction. Wind speed and direction were affected by weather conditions and the time of day.

Smoke was used to visualize the exhaust plume from the individual exhaust stacks serving each hood. The photo in Figure 34 shows the upward directed plume discharged from the exhaust stacks under relatively calm wind conditions. The plume from the standard hood system was discharged with higher momentum and had less potential to envelop the building. The photo in Figure 35 shows the plume from both stacks with the wind blowing the plume over the building. The plume from the LBNL hood appears to have a higher potential to envelop the building and increase the potential for re-entrainment. The re-entrained concentrations are variable but the impact is to increase or bias the average tracer gas concentration.

Data from a series of tracer gas tests were analyzed to determine a correction factor that could be applied to the results of the HAM tests. The results of the limited analysis indicated that concentrations from re-entrainment could increase background concentrations on average of 0.016 ppm for the LBNL hood and 0.011 ppm for the standard hood during the 150 second HAM test sequence. The data used to determine the correction factor are plotted and tabulated in Figure 36. The plot in Figure 37 shows the potential difference between re-entrained concentrations for the LBNL hood and the standard hood.

Table 2 Summary of Static Mannequin Tracer Gas Test Data

Hood ID	Sash Opening	Sash Open Area - sq ft.	Avg. Fvel fpm	Calc. Exh. Flow-cfm	Left TG Avg ppm	Left TG Peak ppm	Center TG Avg ppm	Center TG Peak ppm	Right TG Avg ppm	Right TG Peak ppm	Horiz CD Max Avg fpm	Vert CD Max Avg	Smoke Rating	Other Notes
6-ft LBNL HP Hood	Vertical Full Open	12.01	56	670	0.009	0.081	0.045	0.111	0.015	0.150	25	17	Pass (Fair)	1/11/05 - Very Windy Conditions – Test 1
							0.017	0.042						1/13/05 – Cold in lab. – Test 2
Multiple Test Average							0.031	0.076						
6-ft JT STND Hood	Vertical Full Open	12.06	98	1181	0.017	0.150	0.098	1.974	0.025	0.150	11	22	Low Pass (Poor)	1/11/05 - Very Windy Conditions - Test 1
							0.016	0.032						1/12/05 at 10:30 a.m. – Test 2
							0.006	0.186						1/13/05 -9:25 a.m. Overcast, cool and relatively calm winds. – Test 3
							0.012	0.533						1/13/05 -9:25 a.m. Overcast, cool and relatively calm winds. – Test 4
							0.005	0.224						1/13/05 - 11:25 a.m. - Still cold in lab and outside but becoming windier. – Test 5
Multiple Test Average							0.027	0.590						

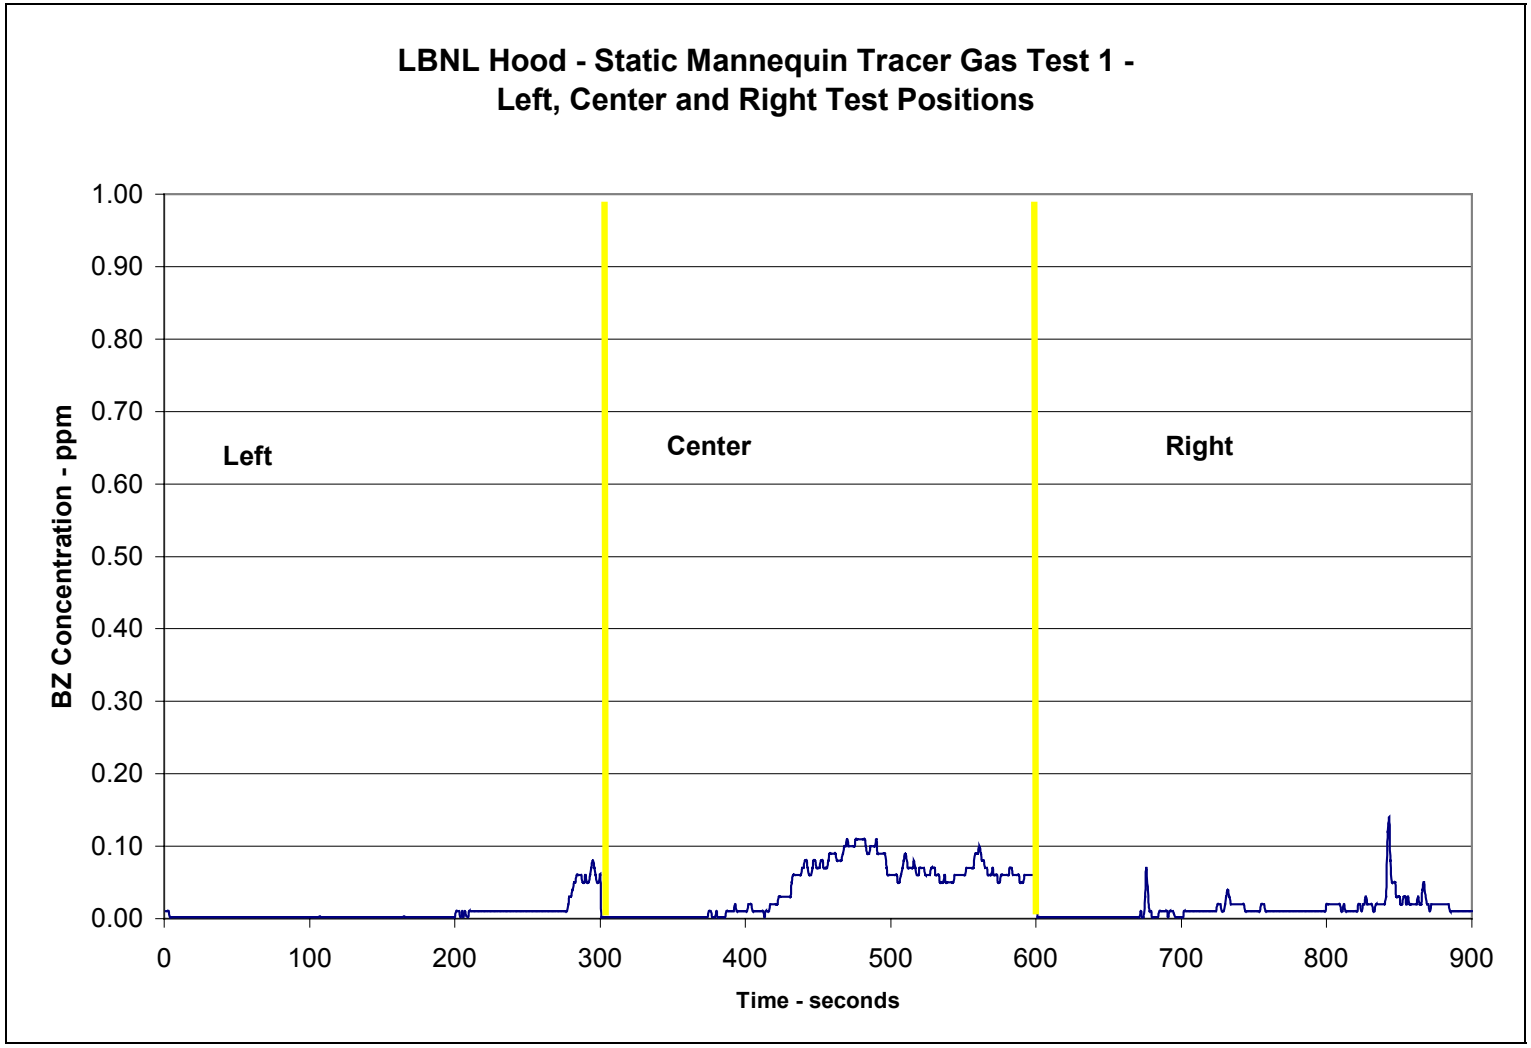

Figure 9 Plot of concentrations measured at each test position during a static mannequin tracer gas test on the LBNL hood.

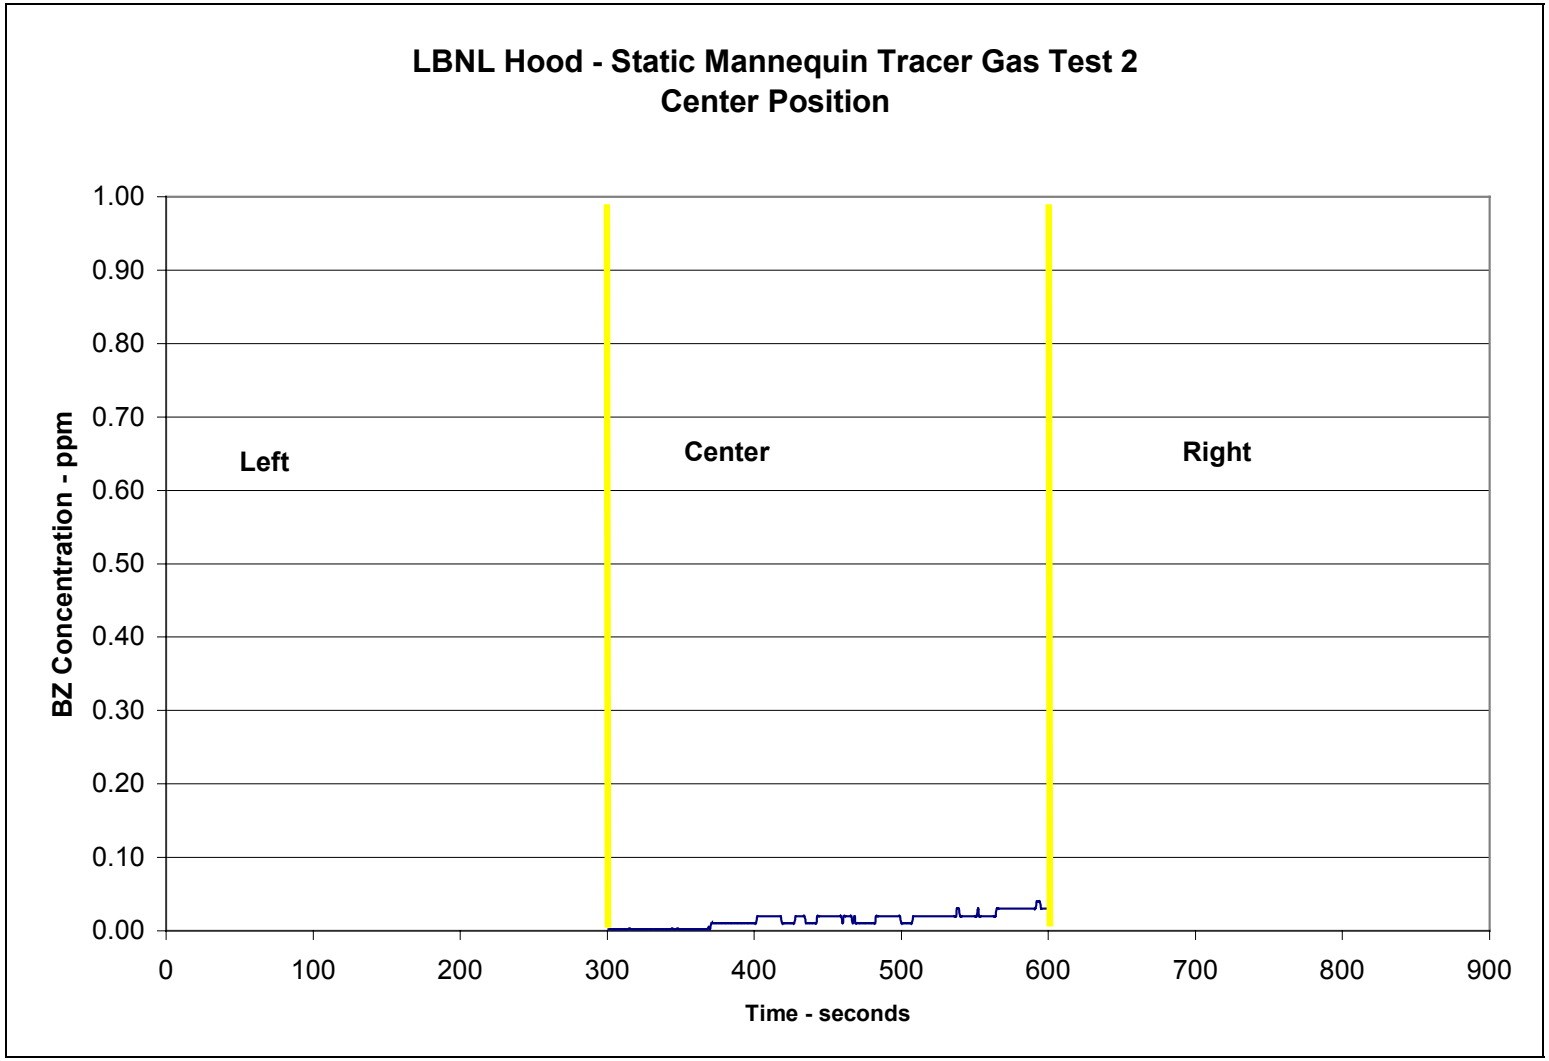

Figure 10 Plot of concentrations measured during a re-test at the center position during a static mannequin tracer gas test on the LBNL hood.

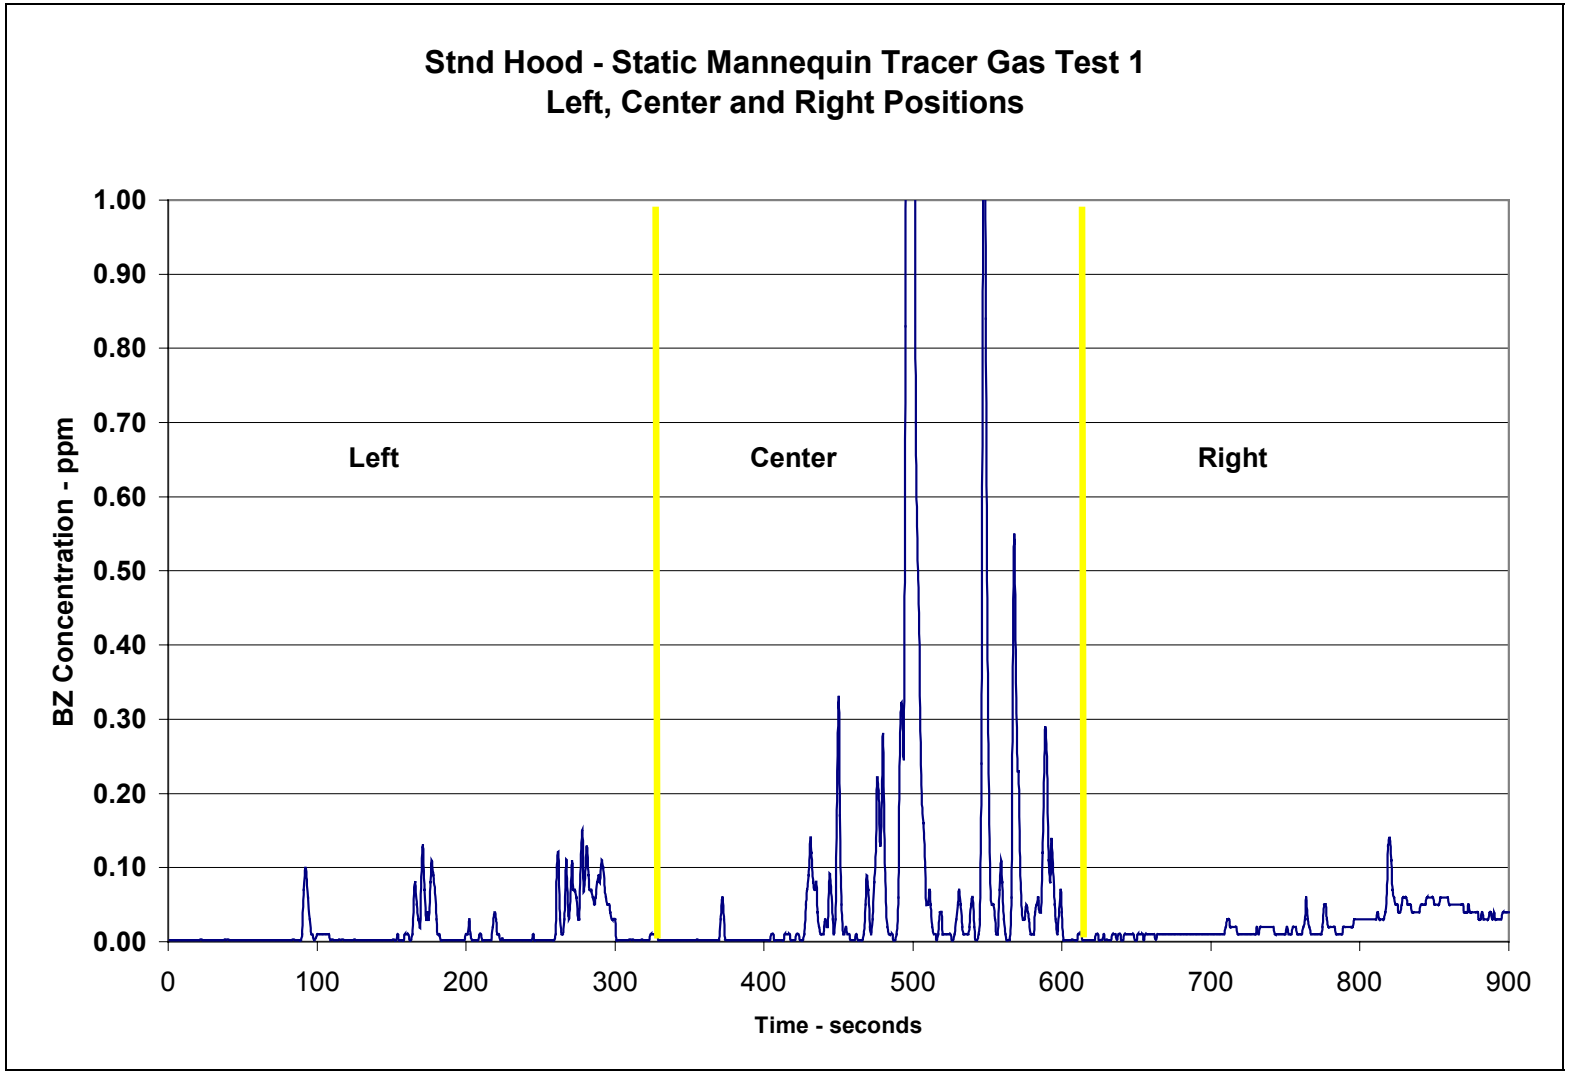

Figure 11 Plot of concentrations measured at each test position during a static mannequin tracer gas test on the standard hood.

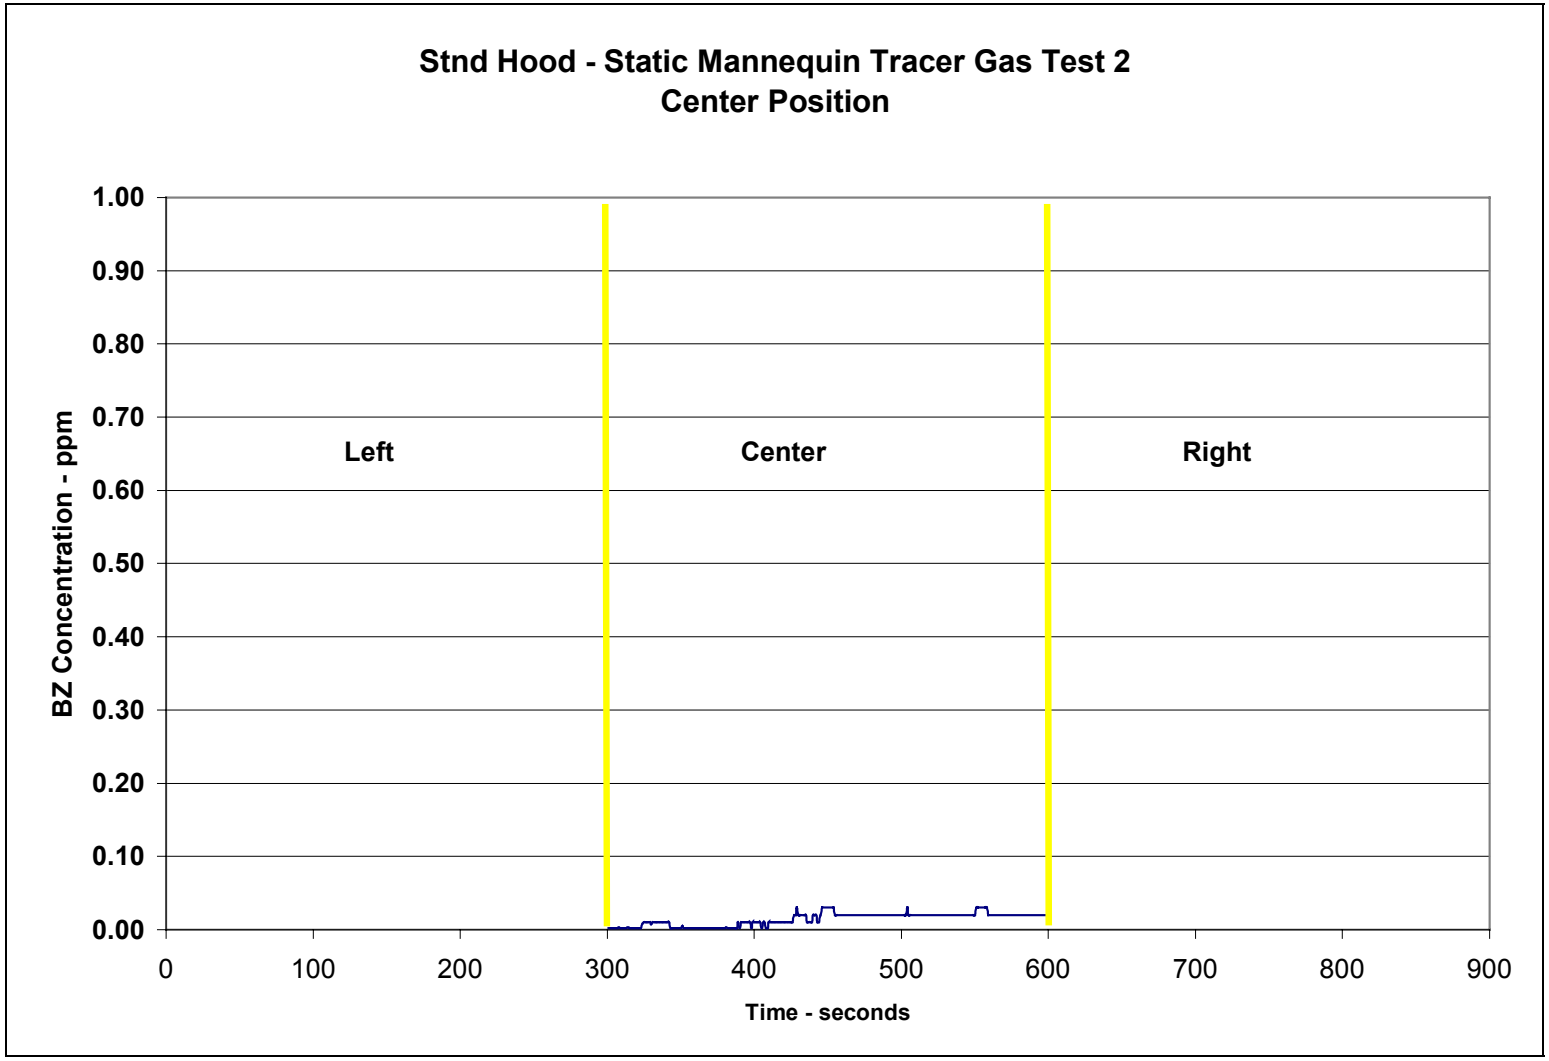

Figure 12 Plot of concentrations measured during a re-test at the center position during a static mannequin tracer gas test on the standard hood.

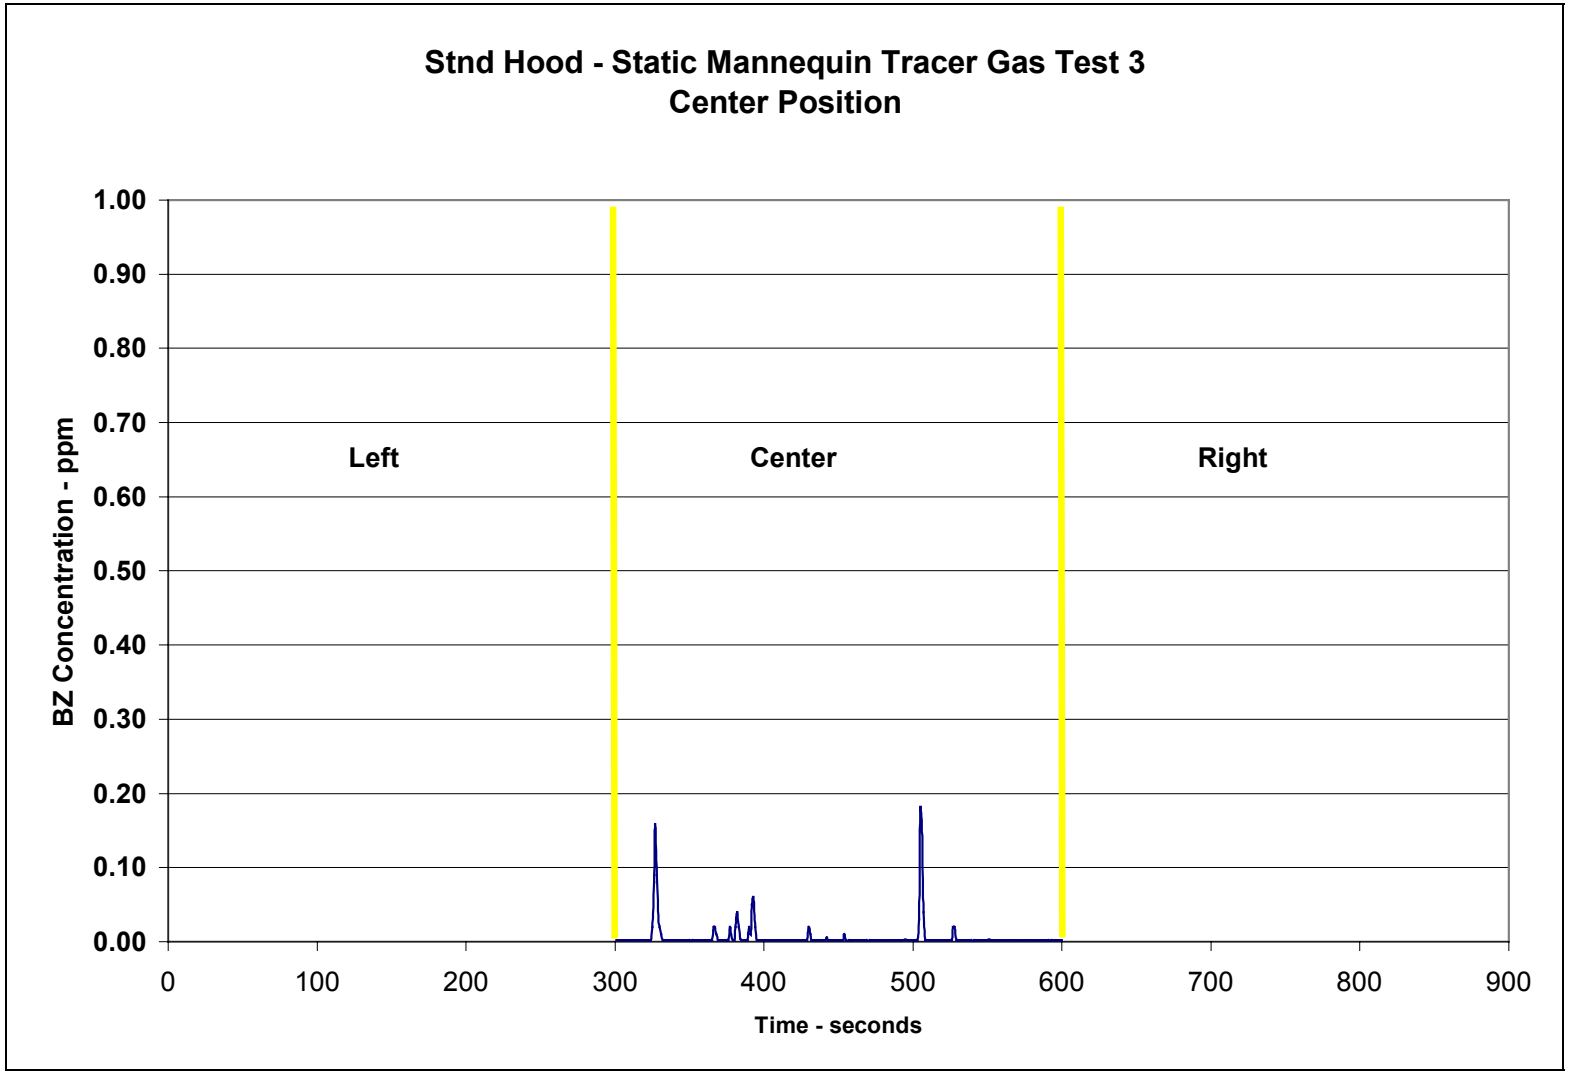

Figure 13 Plot of concentrations measured during a re-test at the center position during a static mannequin tracer gas test on the LBNL hood.

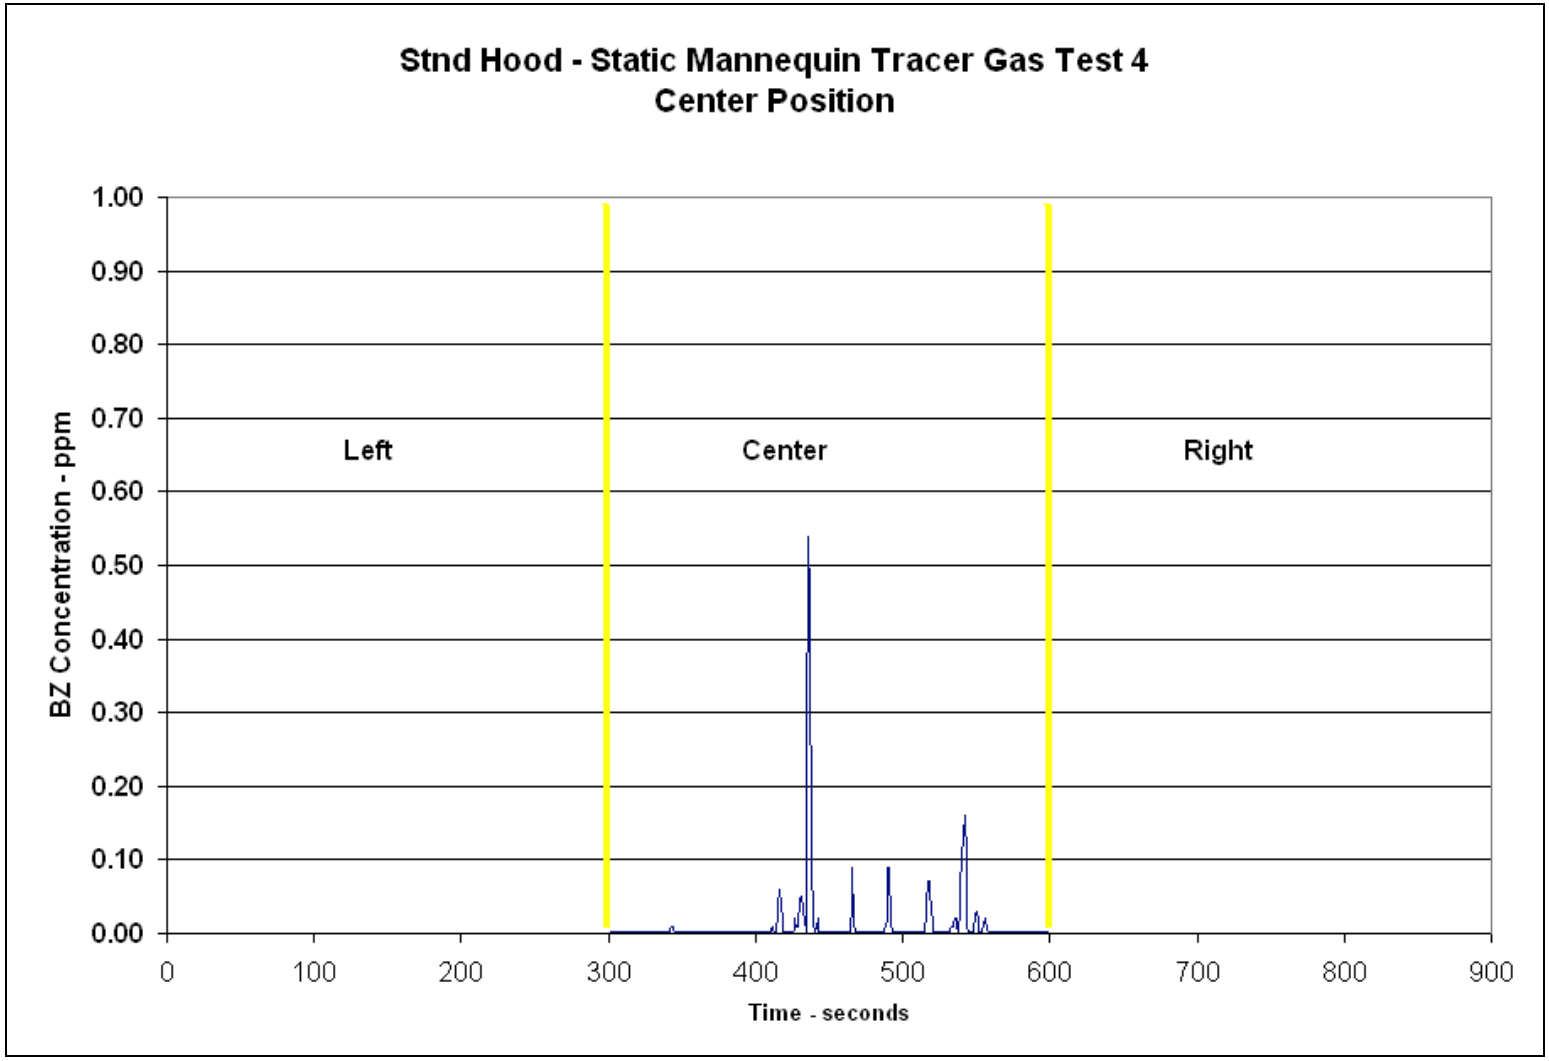

Figure 14 Plot of concentrations measured during a re-test at the center position during a static mannequin tracer gas test on the LBNL hood.

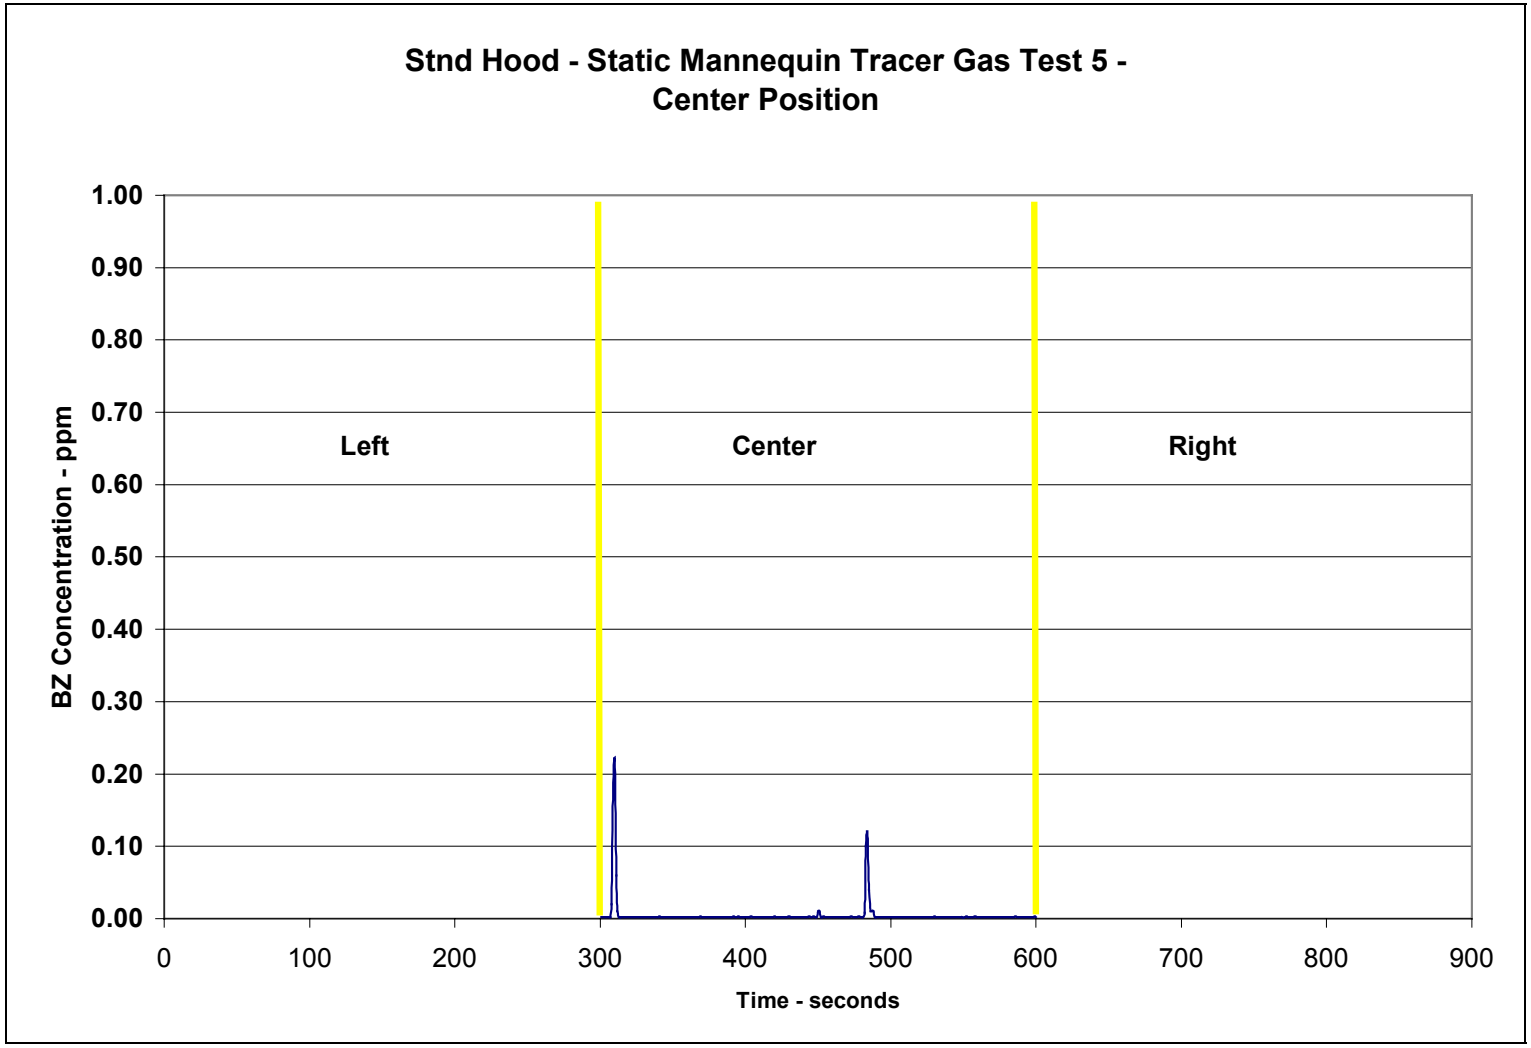

Figure 15 Plot of concentrations measured during a re-test at the center position during a static mannequin tracer gas test on the LBNL hood.

Table 3 Summary of tracer gas data and sequence time for HAM tests for each hood and test location

		Left		Center		Right		MAX Breathing Zone Concentration (Uncorrected) - ppm	MAX Breathing Zone Concentration (Corrected) - ppm
		BZ Conc. ppm	SEQ Duration Time-sec.	BZ Conc. ppm	SEQ Duration Time-sec.	BZ Conc. ppm	SEQ Duration Time-sec.		
LBNL - HP	SEQ 1	0.011	153	0.024	154	0.037	160	0.030	0.014
	SEQ 2	0.052	159	0.013	160	0.027	153		
	SEQ 3	0.027	147	0.012	169	0.015	155		
Series Avg. / Total Time		0.030	459	0.016	483	0.026	468		
Standard Hood	SEQ 1	0.036	154	0.007	163	0.010	156	0.021	0.010
	SEQ 2	0.015	135	0.007	157	0.006	150		
	SEQ 3	0.011	145	0.020	158	0.007	148		
Series Avg / Total Time		0.021	434	0.011	478	0.008	454		
Aggregate SEQ Time (MAX) - mm:ss		7:39		8:03		7:48			
Aggregate SEQ Time (AVG) - mm:ss		7:27		8:01		7:41			
Aggregate SEQ Time (MIN) - mm:ss		7:14		7:58		7:34			

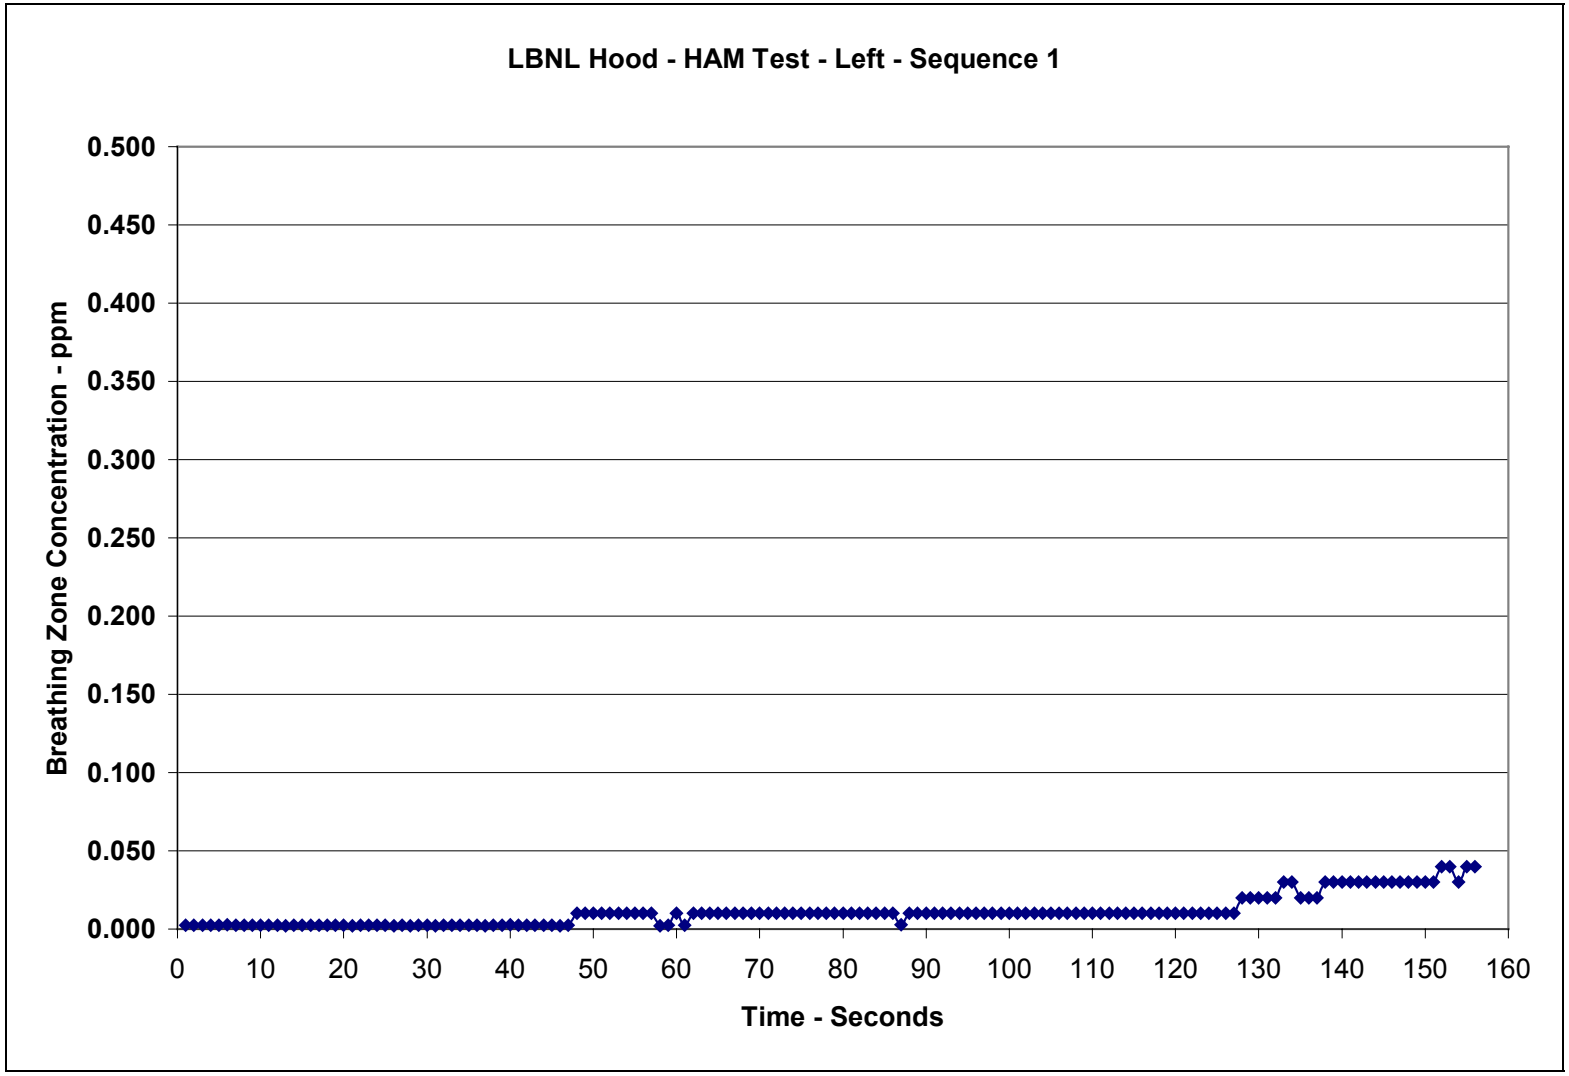

Figure 16 Tracer gas concentrations measured during test sequence 1 on the left side of the LBNL hood

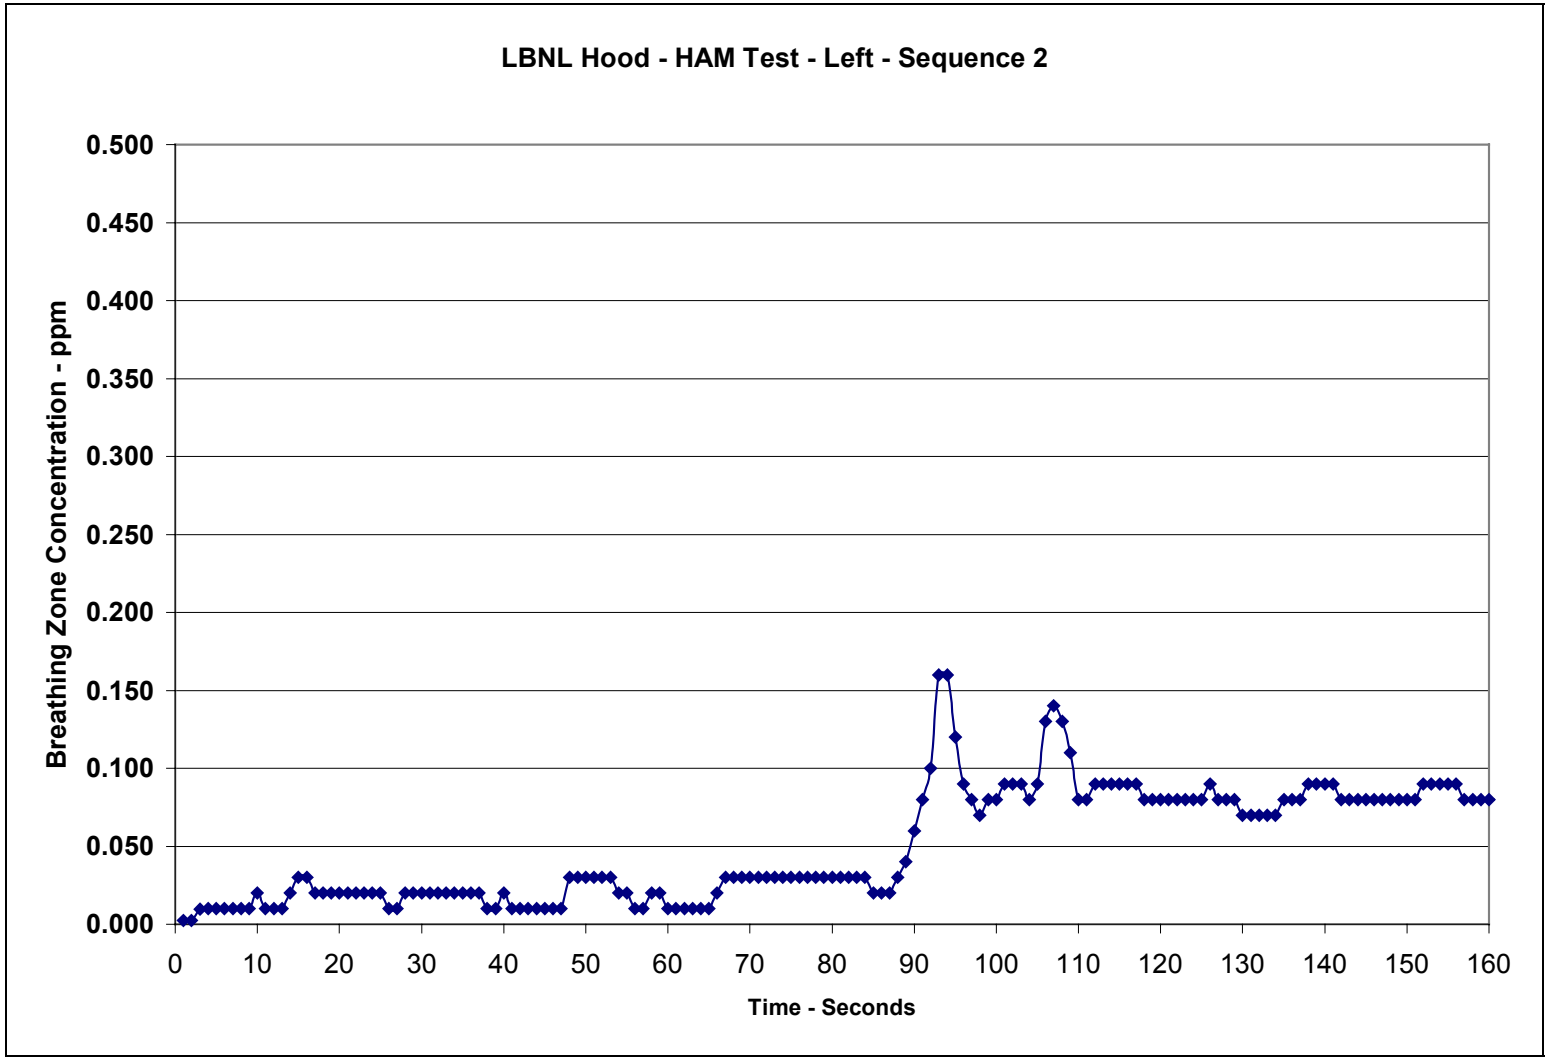

Figure 17 Tracer gas concentrations measured during test sequence 2 on the left side of the LBNL hood

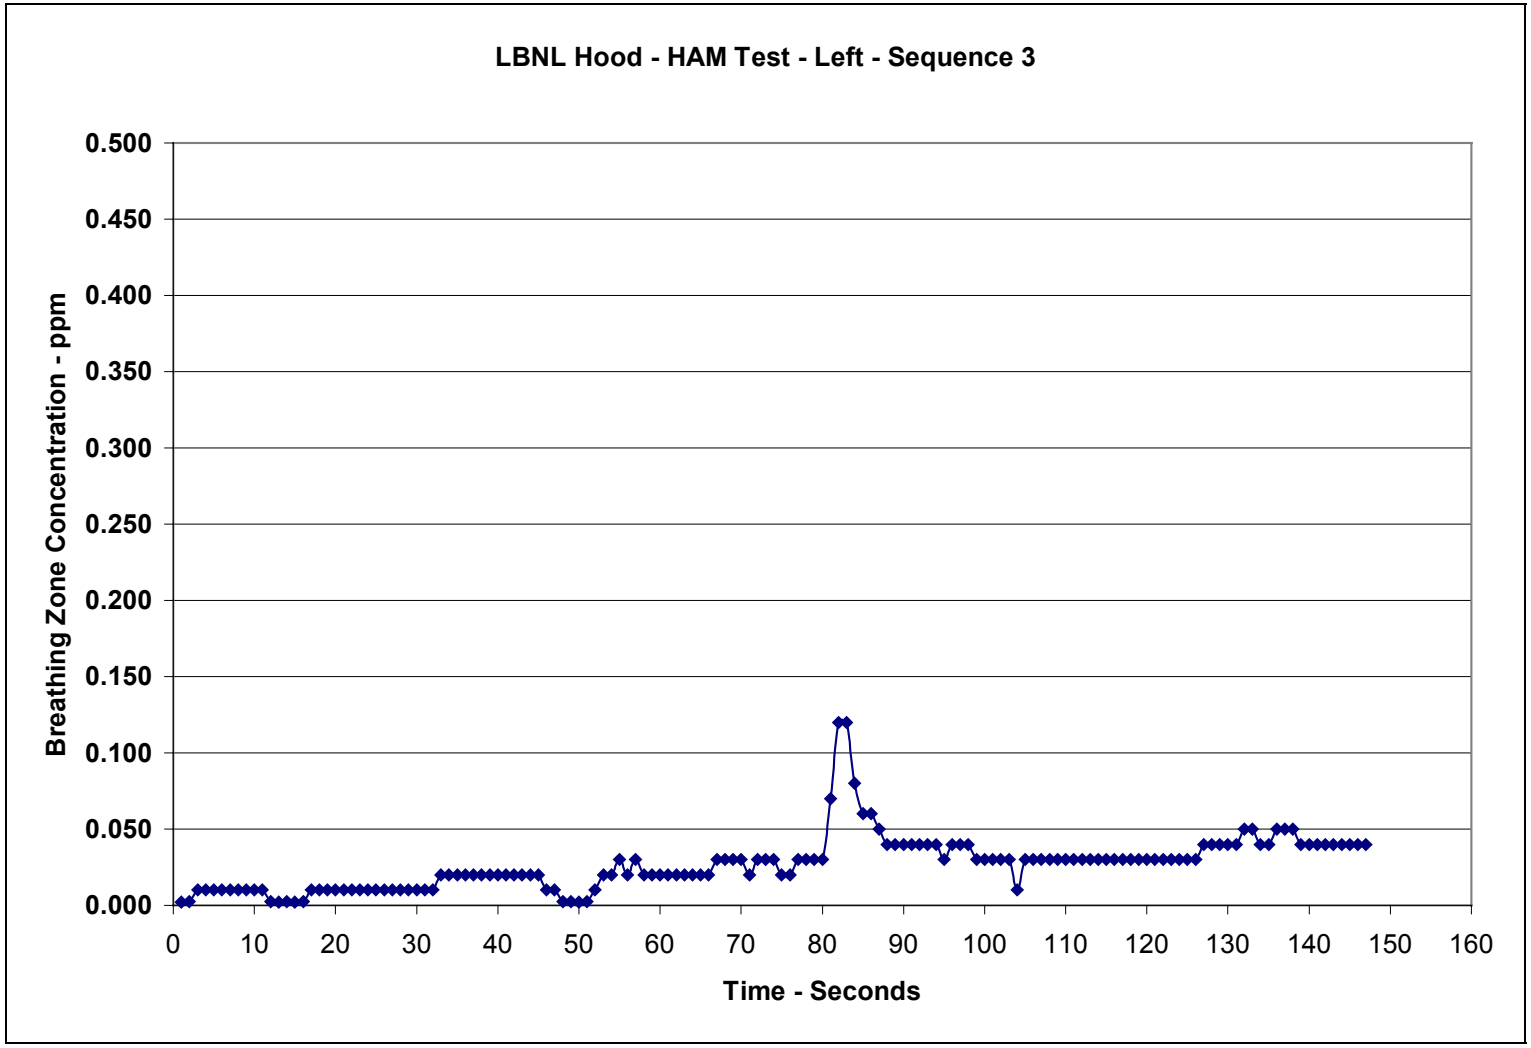

Figure 18 Tracer gas concentrations measured during test sequence 3 on the left side of the LBNL hood

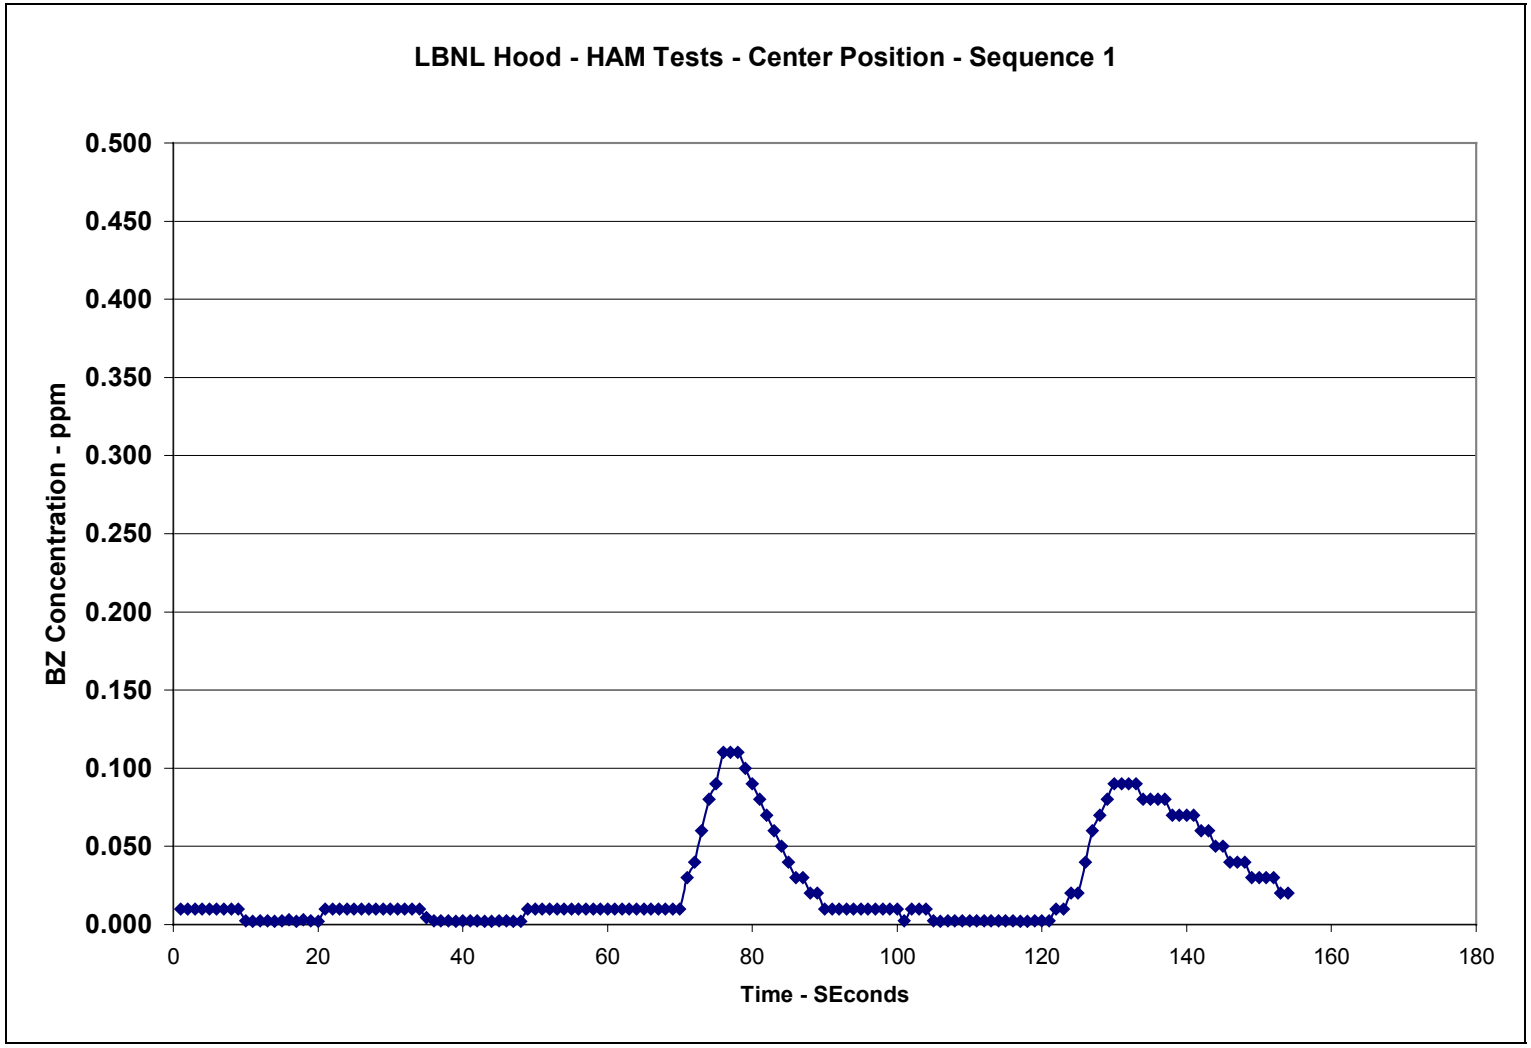

Figure 19 Tracer gas concentrations measured during test sequence 1 in the center position of the LBNL hood

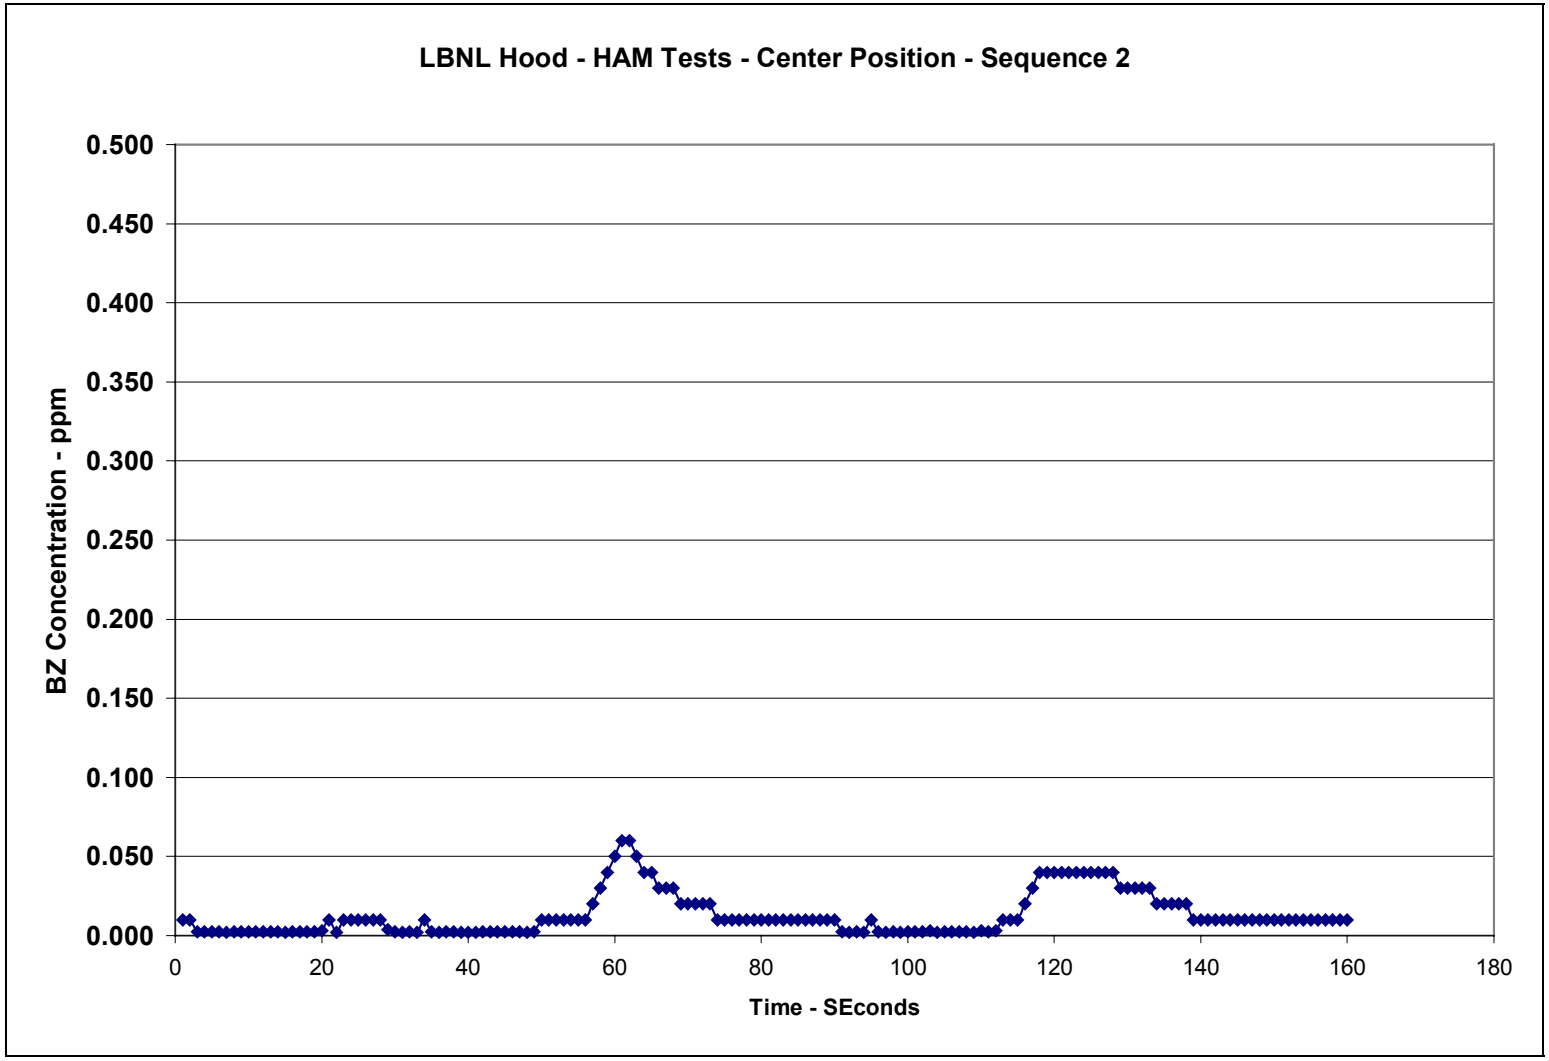

Figure 20 Tracer gas concentrations measured during test sequence 2 in the center position of the LBNL hood

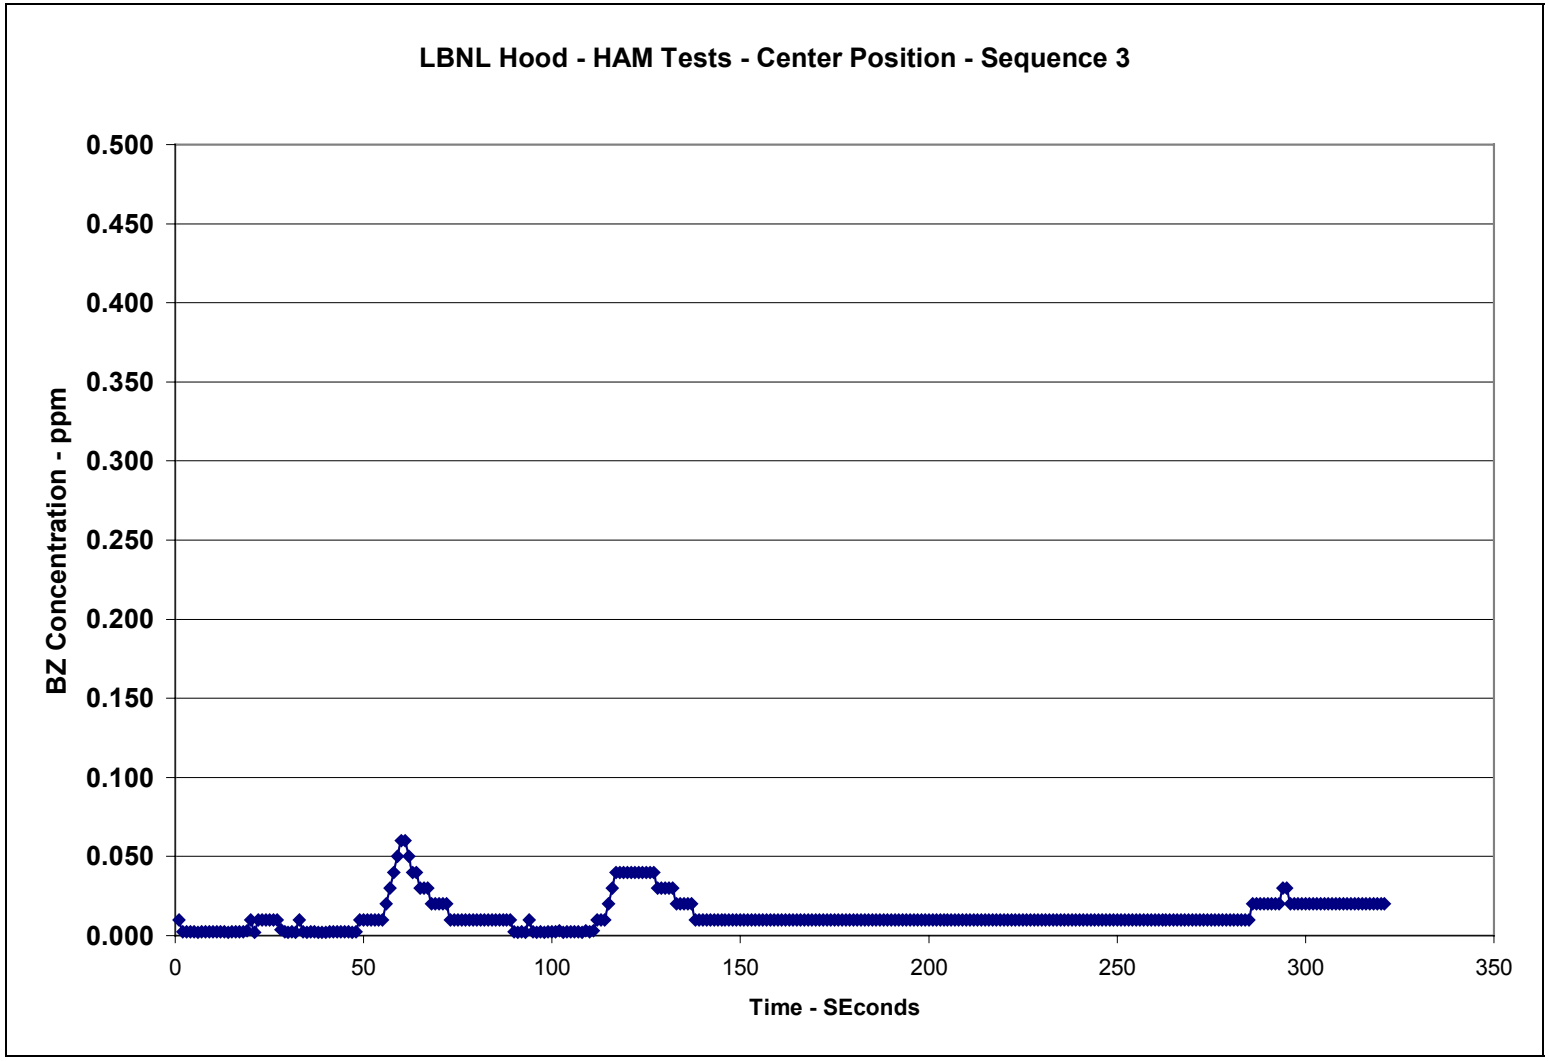

Figure 21 Tracer gas concentrations measured during test sequence 3 in the center position of the LBNL hood

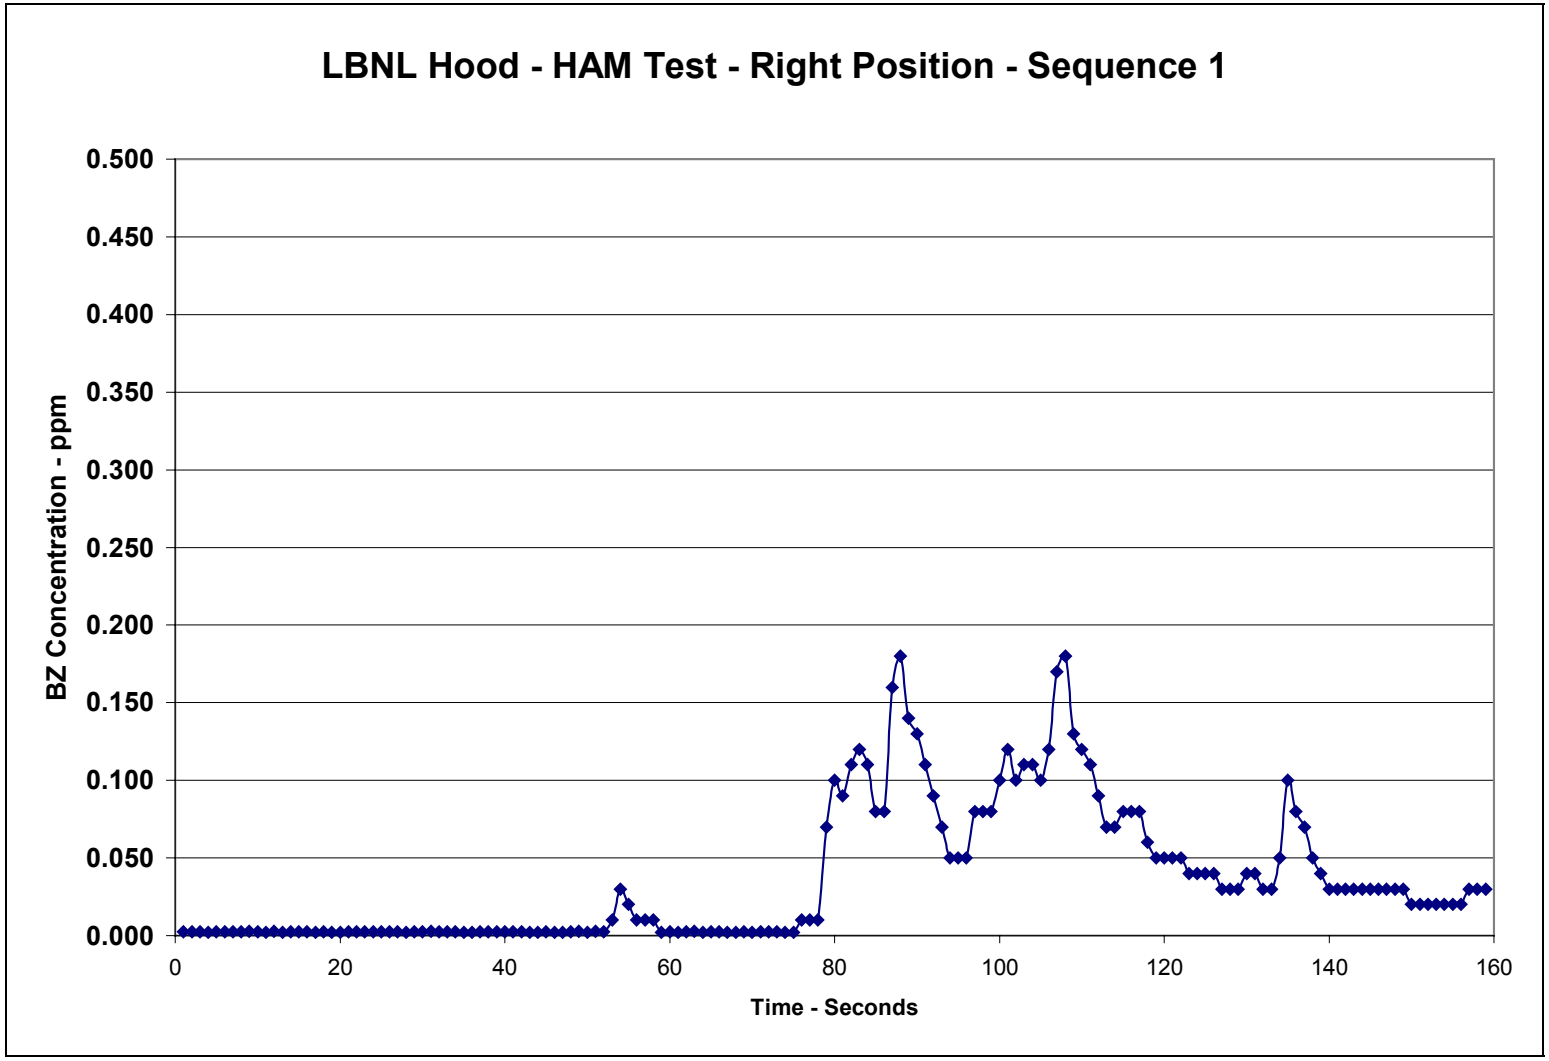

Figure 22 Tracer gas concentrations measured during test sequence 1 on the right side of the LBNL hood

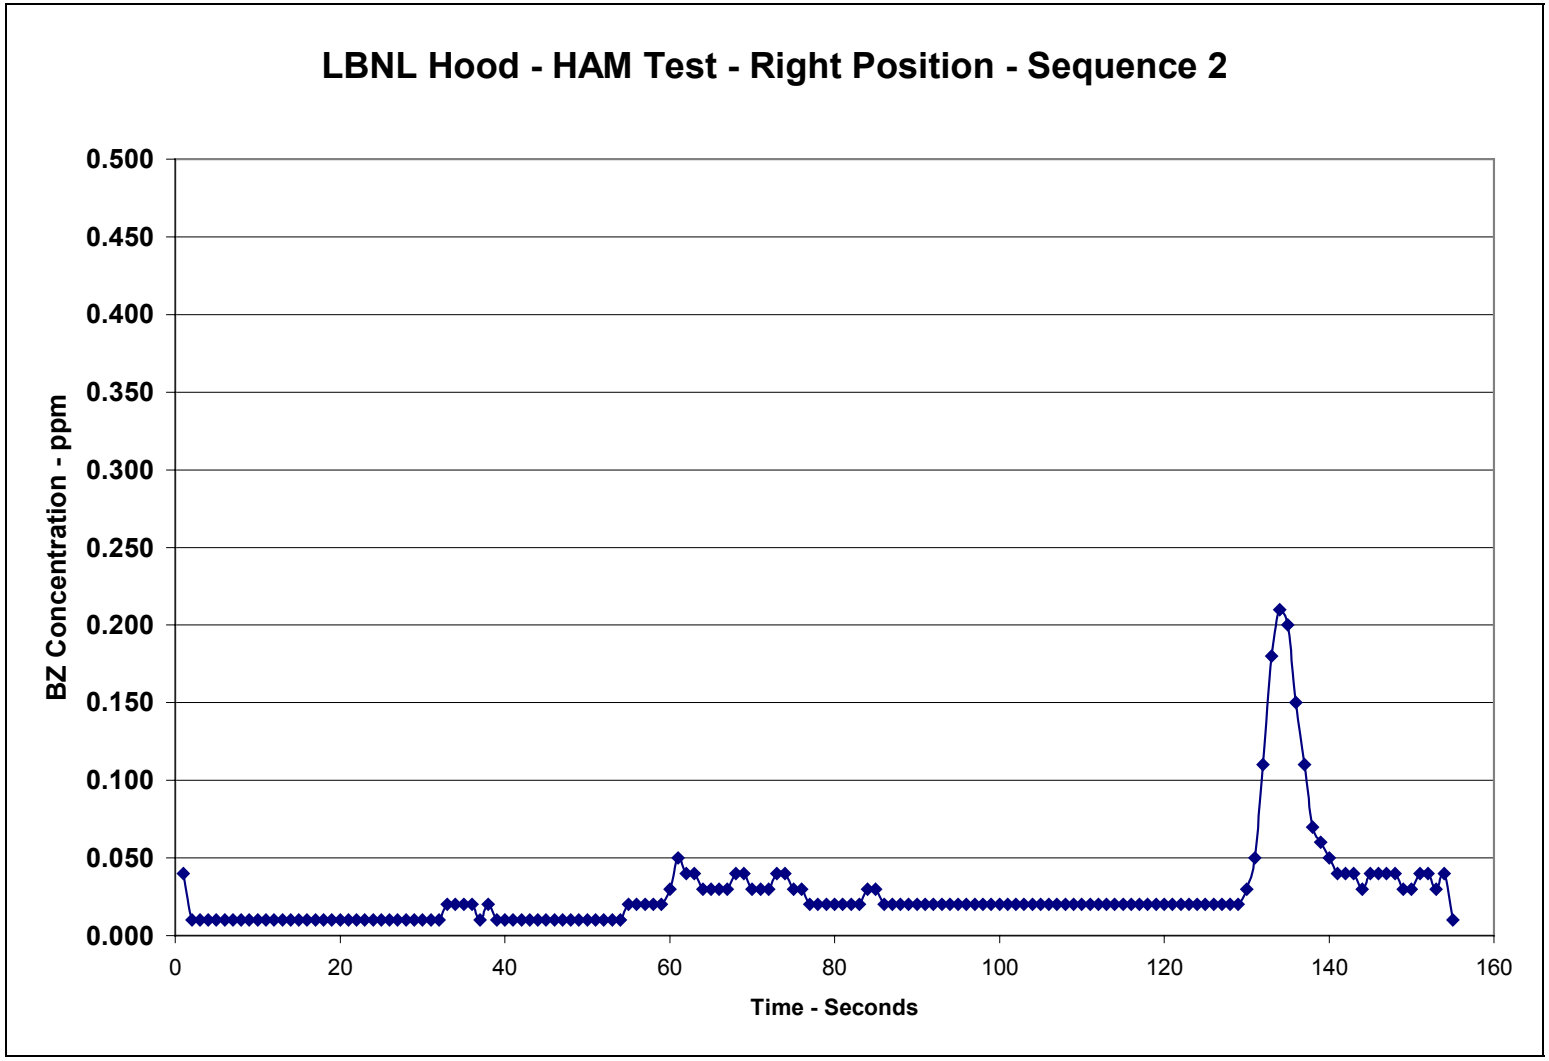

Figure 23 Tracer gas concentrations measured during test sequence 2 on the right side of the LBNL hood

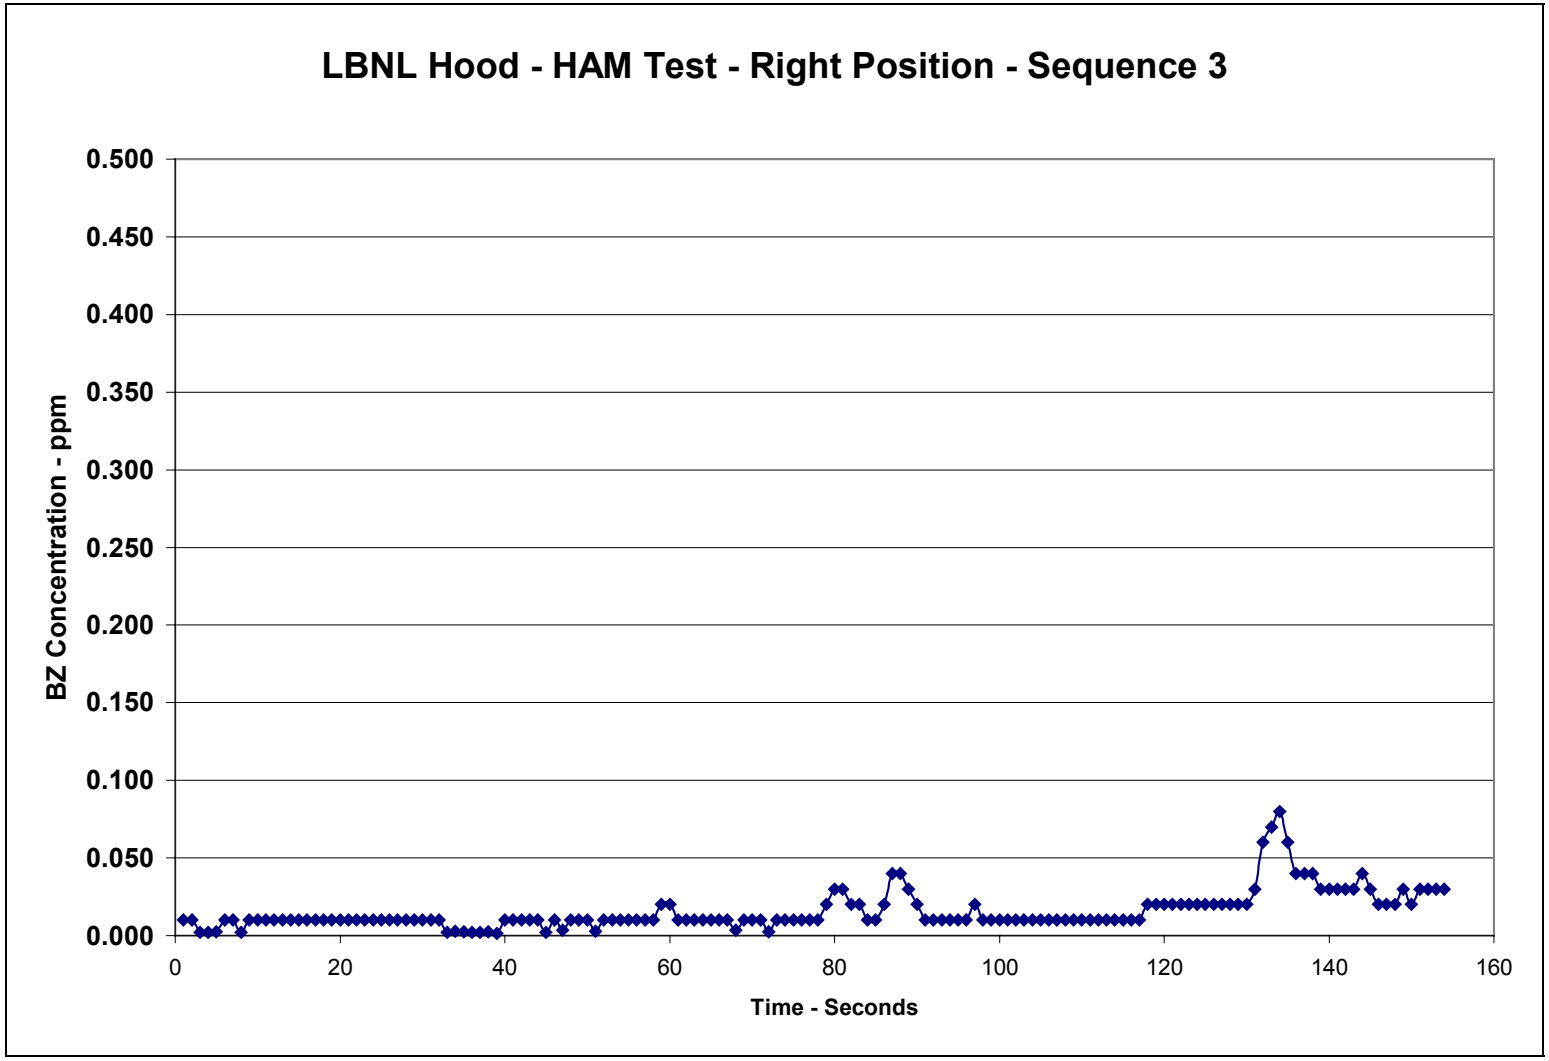

Figure 24 Tracer gas concentrations measured during test sequence 3 on the right side of the LBNL hood

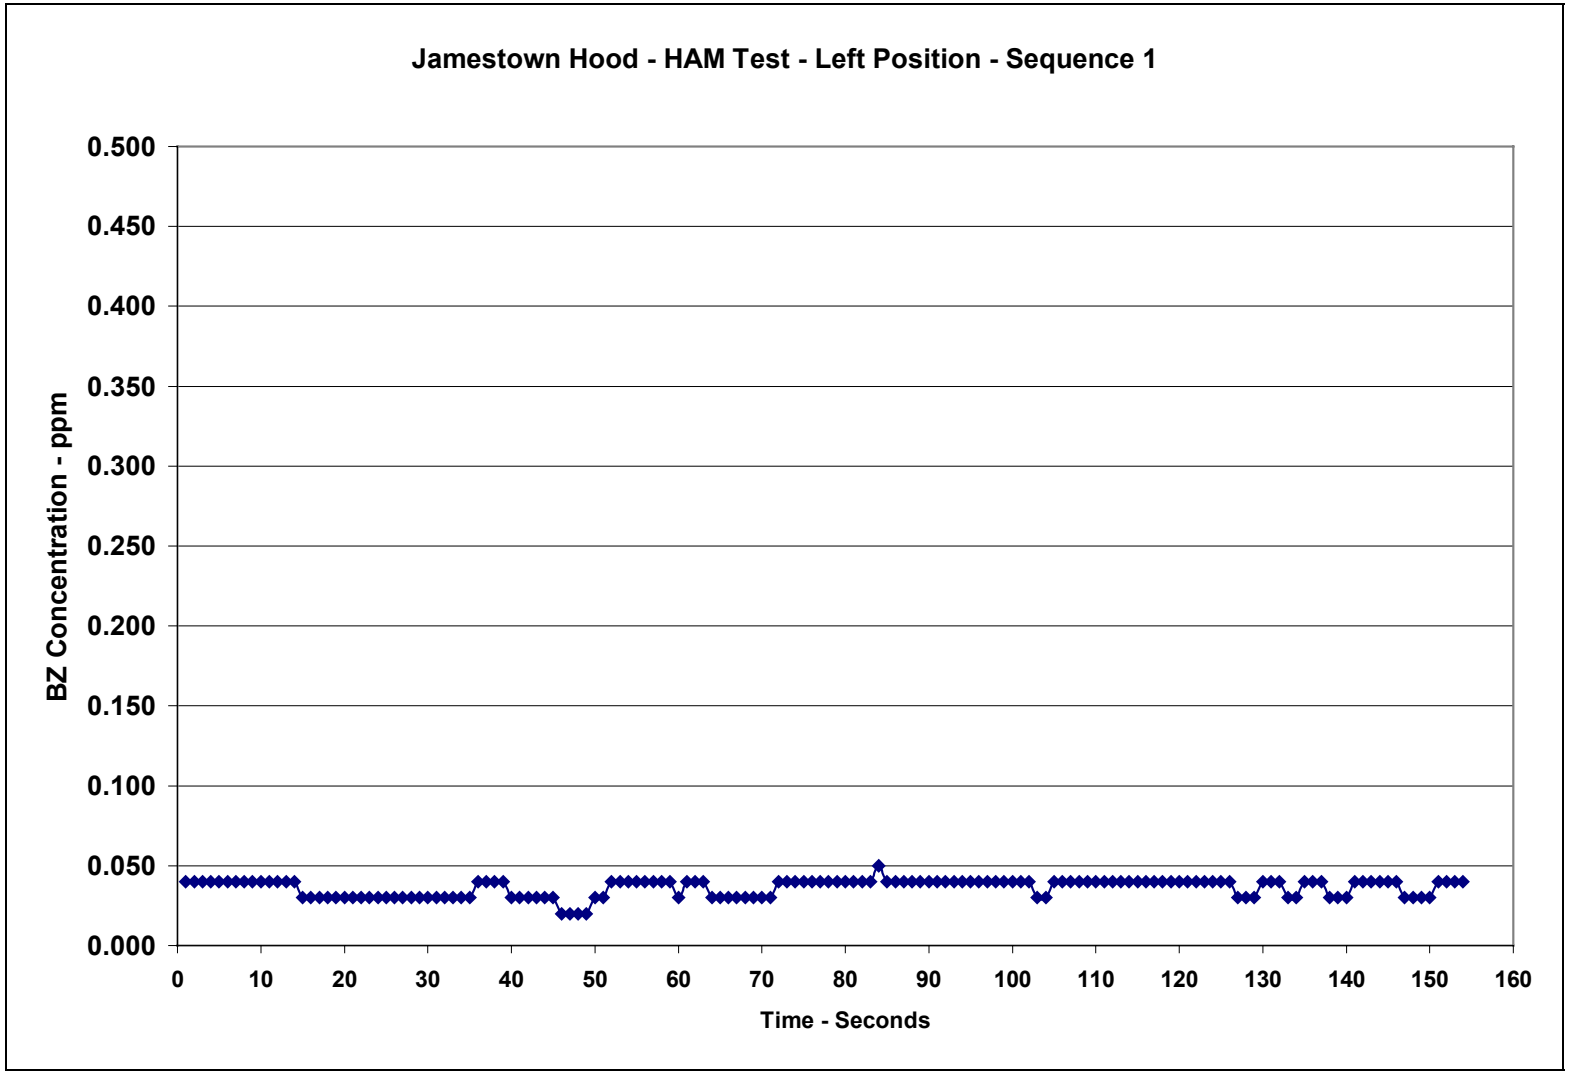

Figure 25 Tracer gas concentrations measured during test sequence 1 on the left side of the Standard hood

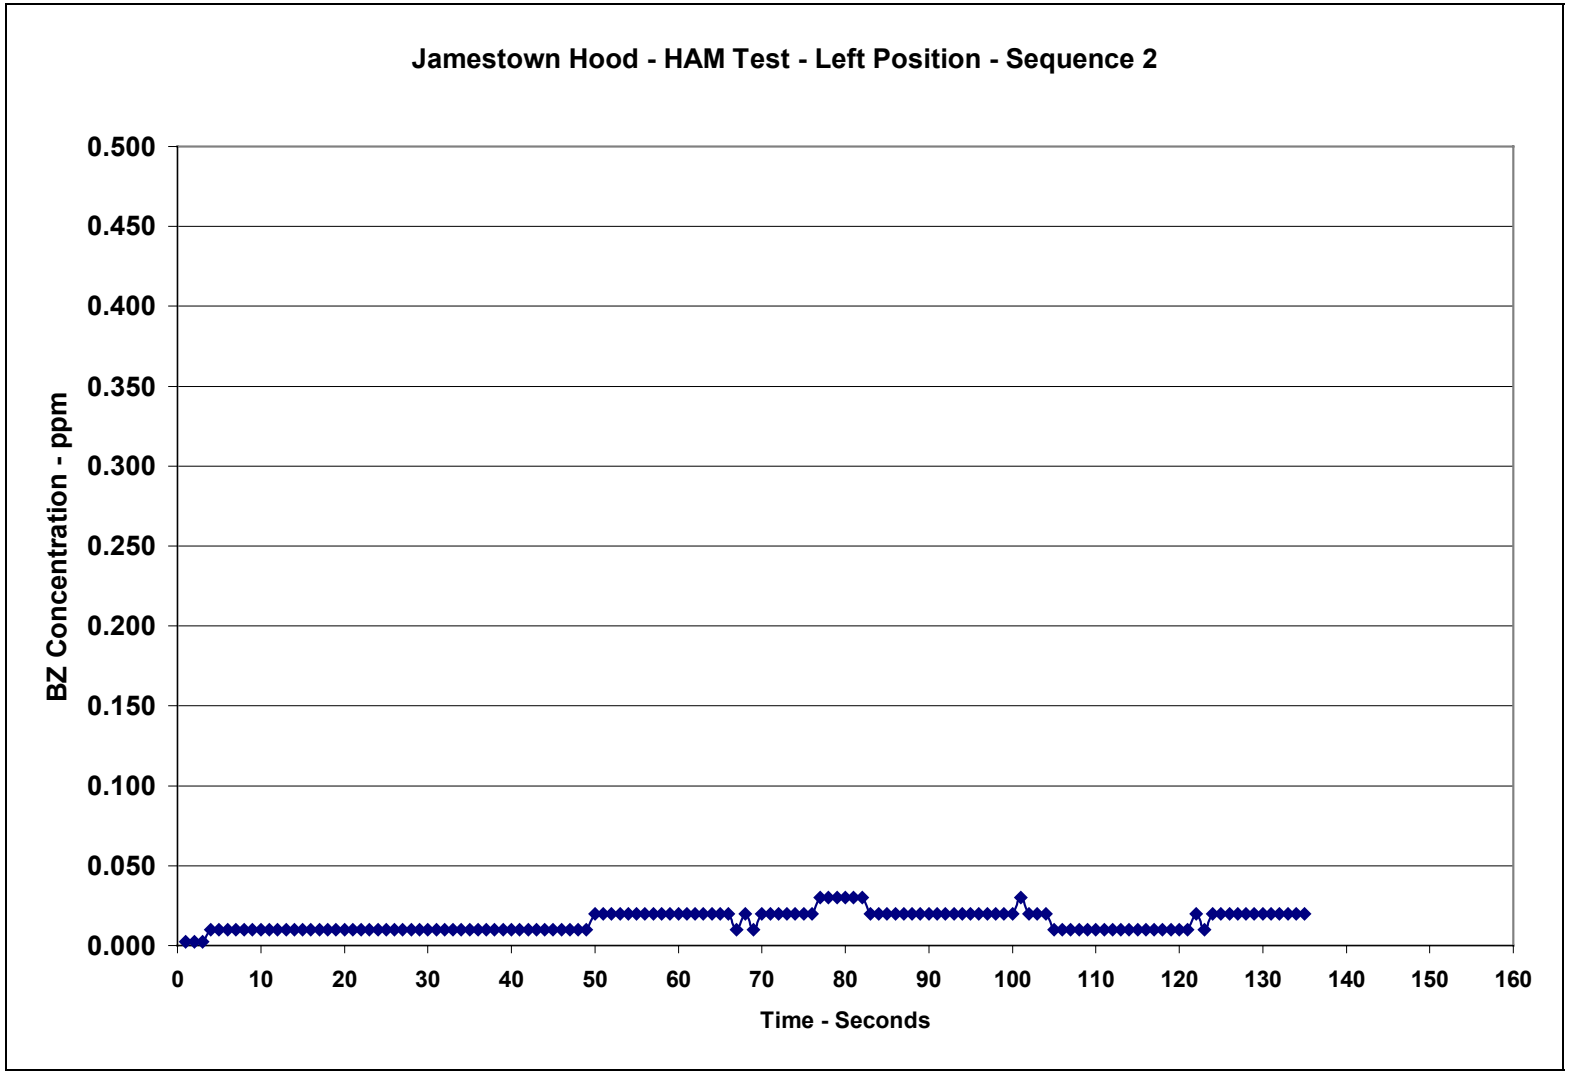

Figure 26 Tracer gas concentrations measured during test sequence 2 on the left side of the standard hood

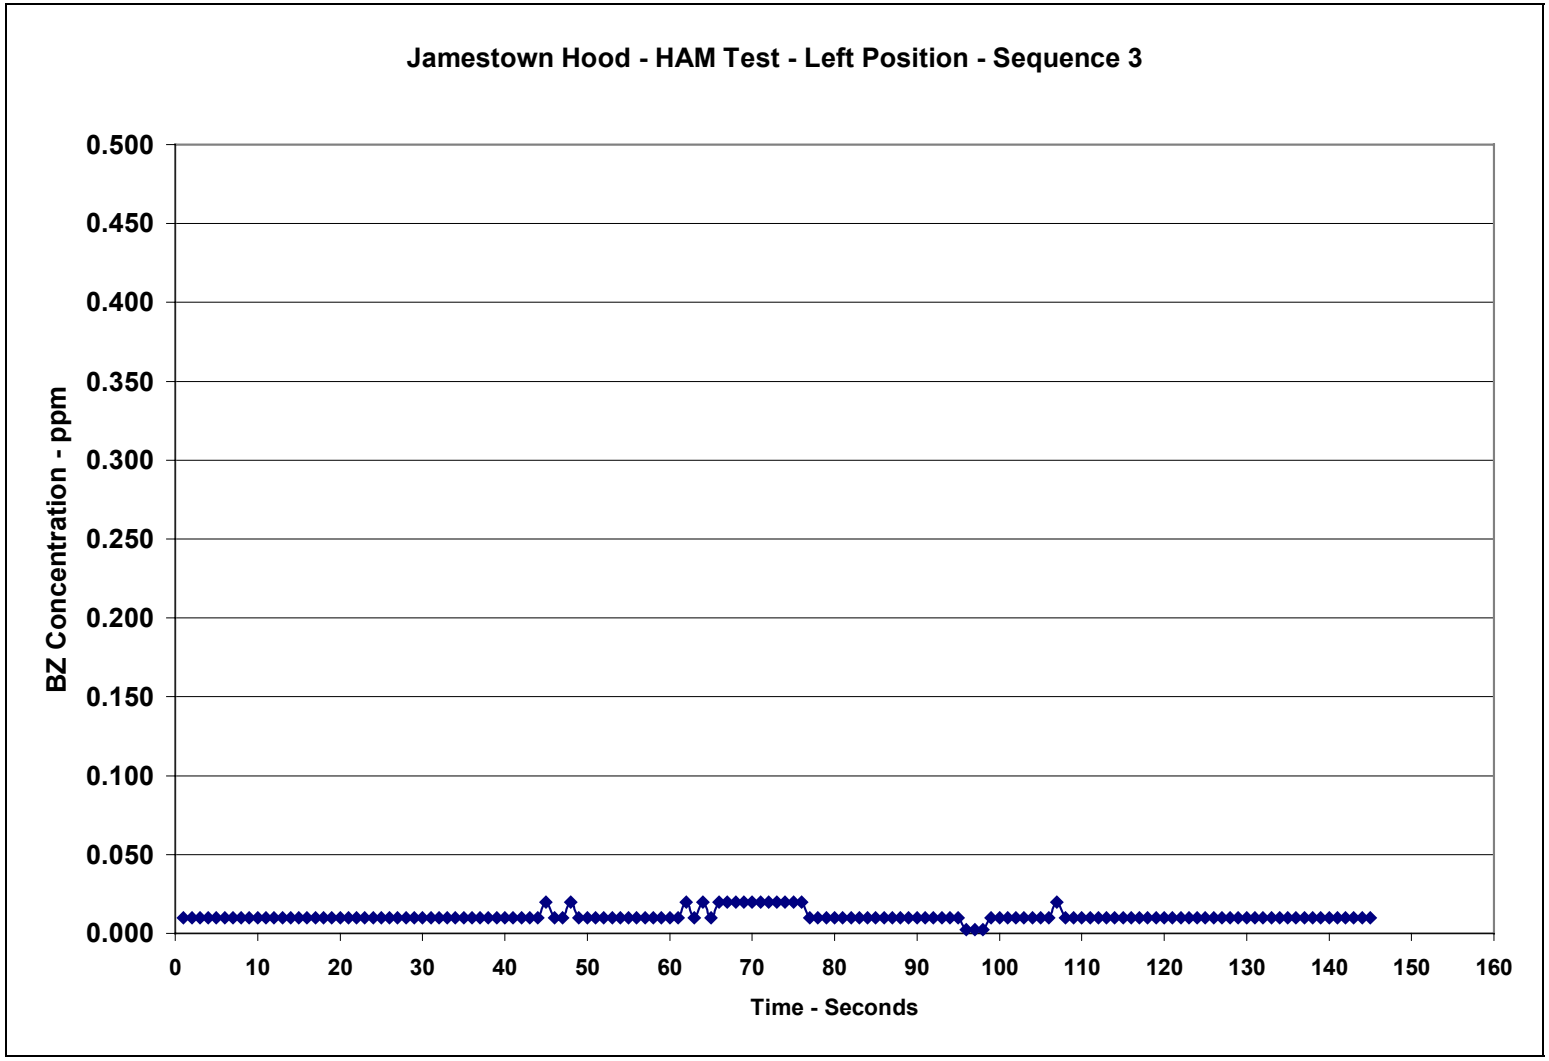

Figure 27 Tracer gas concentrations measured during test sequence 3 on the left side of the standard hood

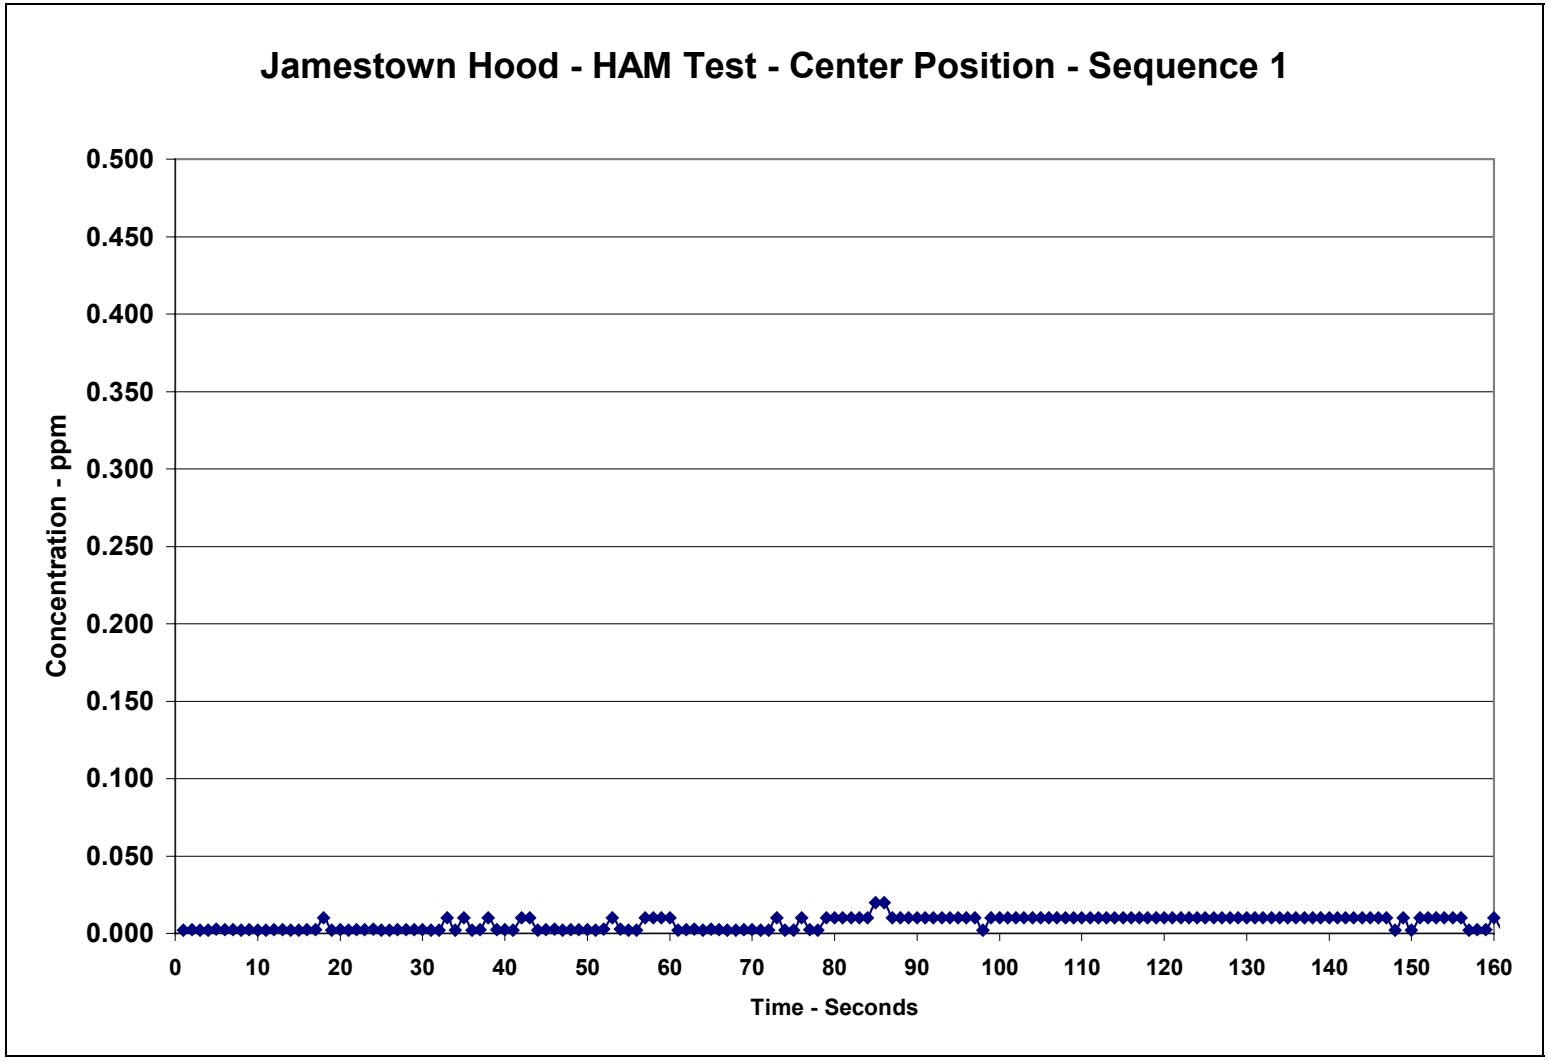

Figure 28 Tracer gas concentrations measured during test sequence 1 in the center test position of the standard hood

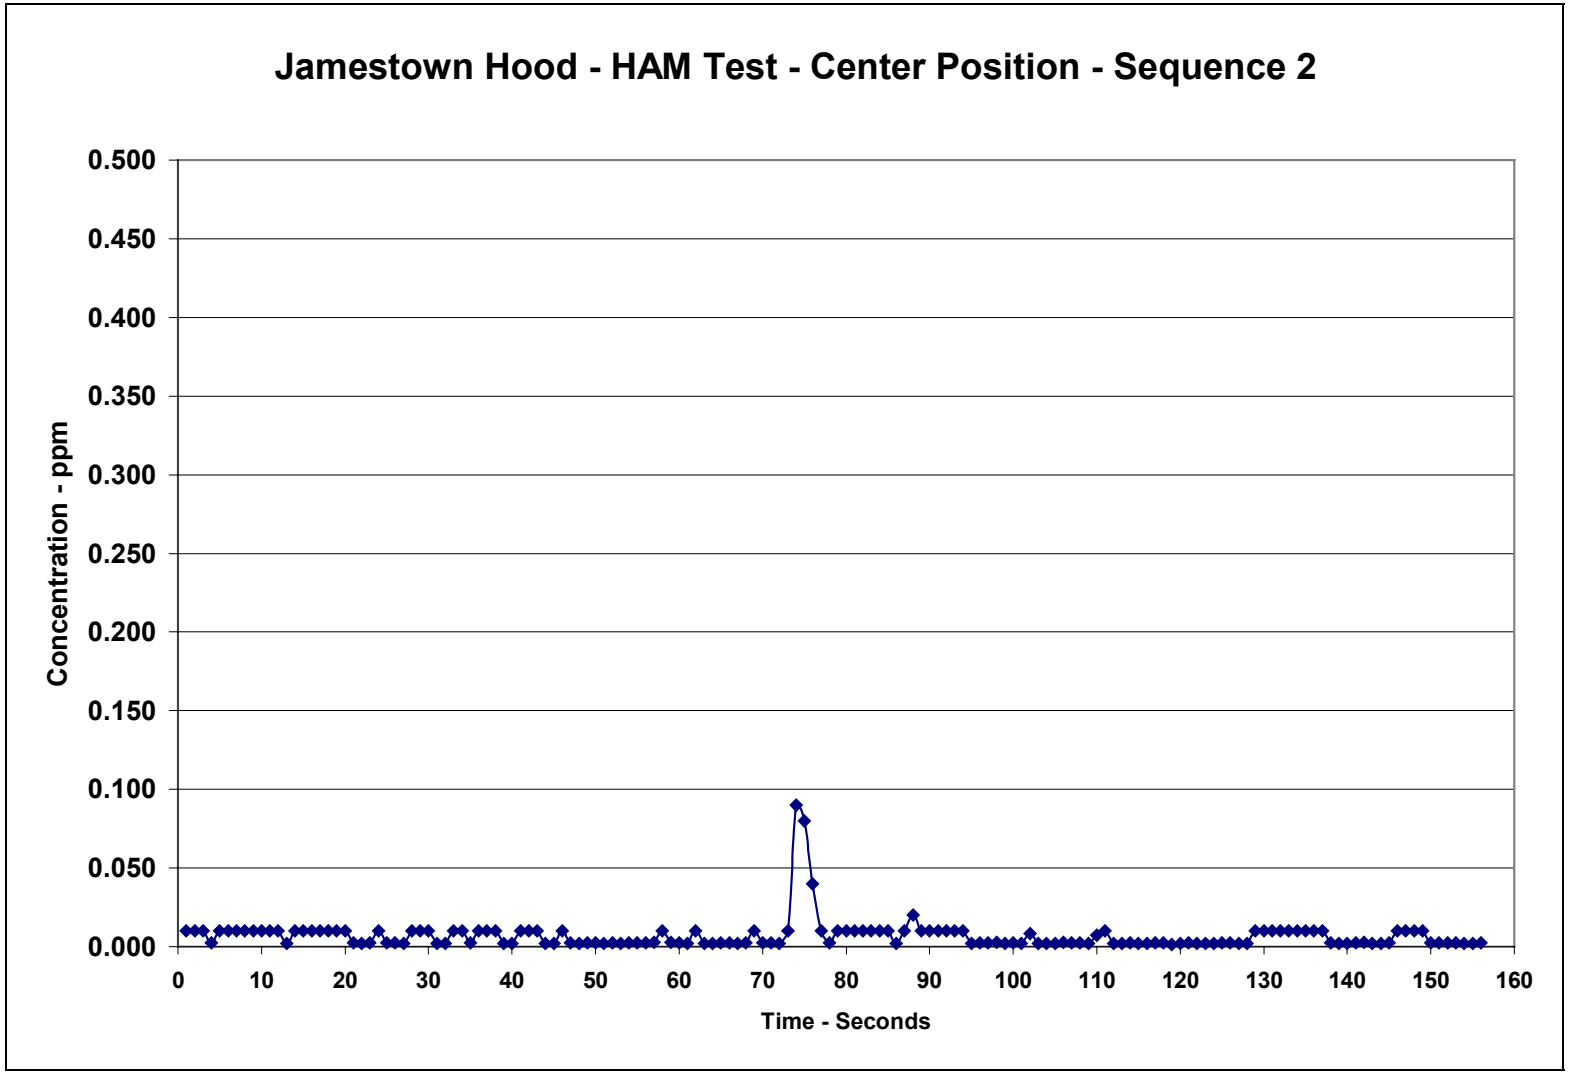

Figure 29 Tracer gas concentrations measured during test sequence 2 in the center test position of the standard hood

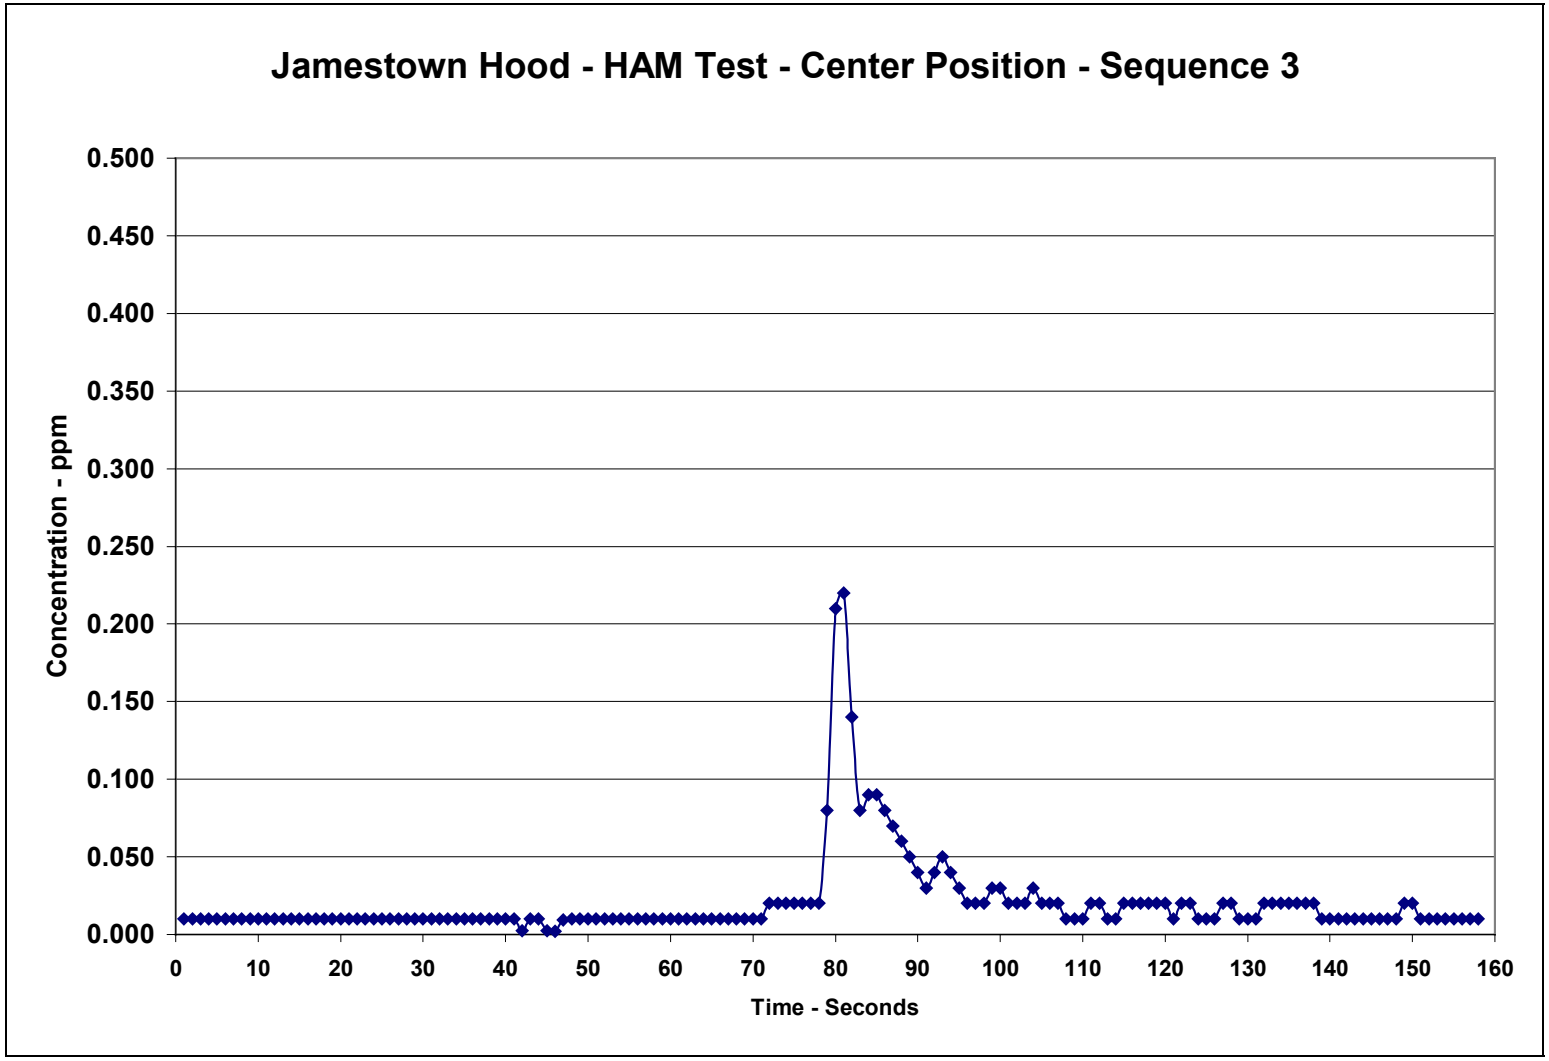

Figure 30 Tracer gas concentrations measured during test sequence 3 in the center test position of the standard hood

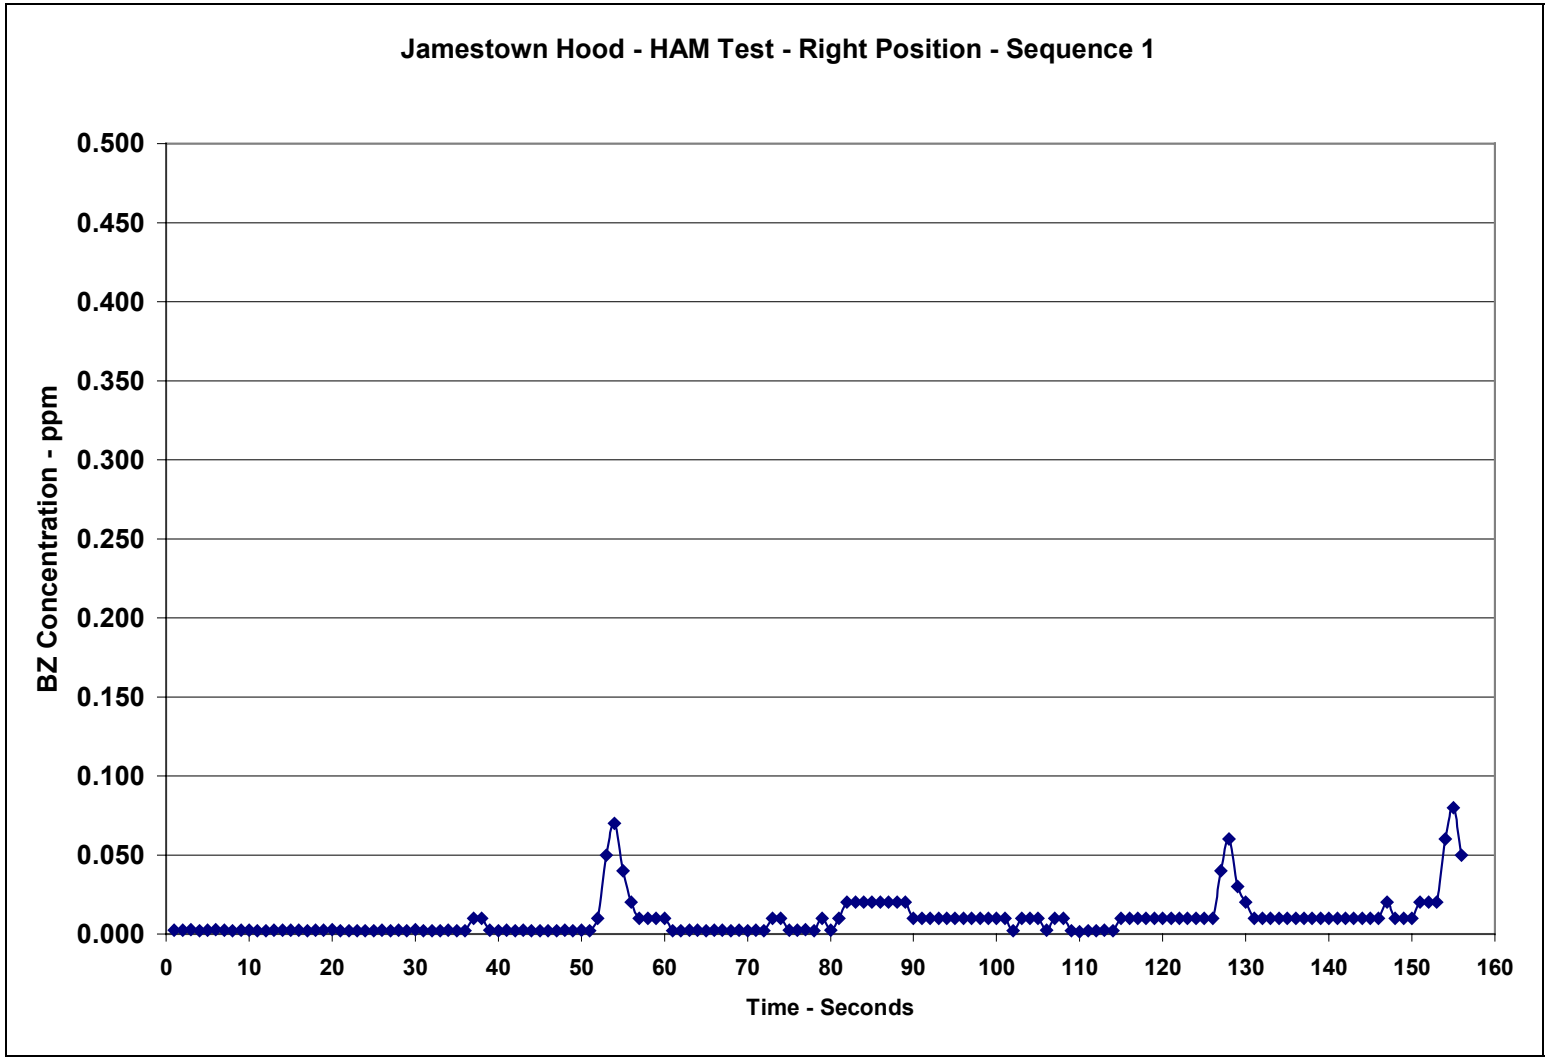

Figure 31 Tracer gas concentrations measured during test sequence 1 on the right side of the standard hood

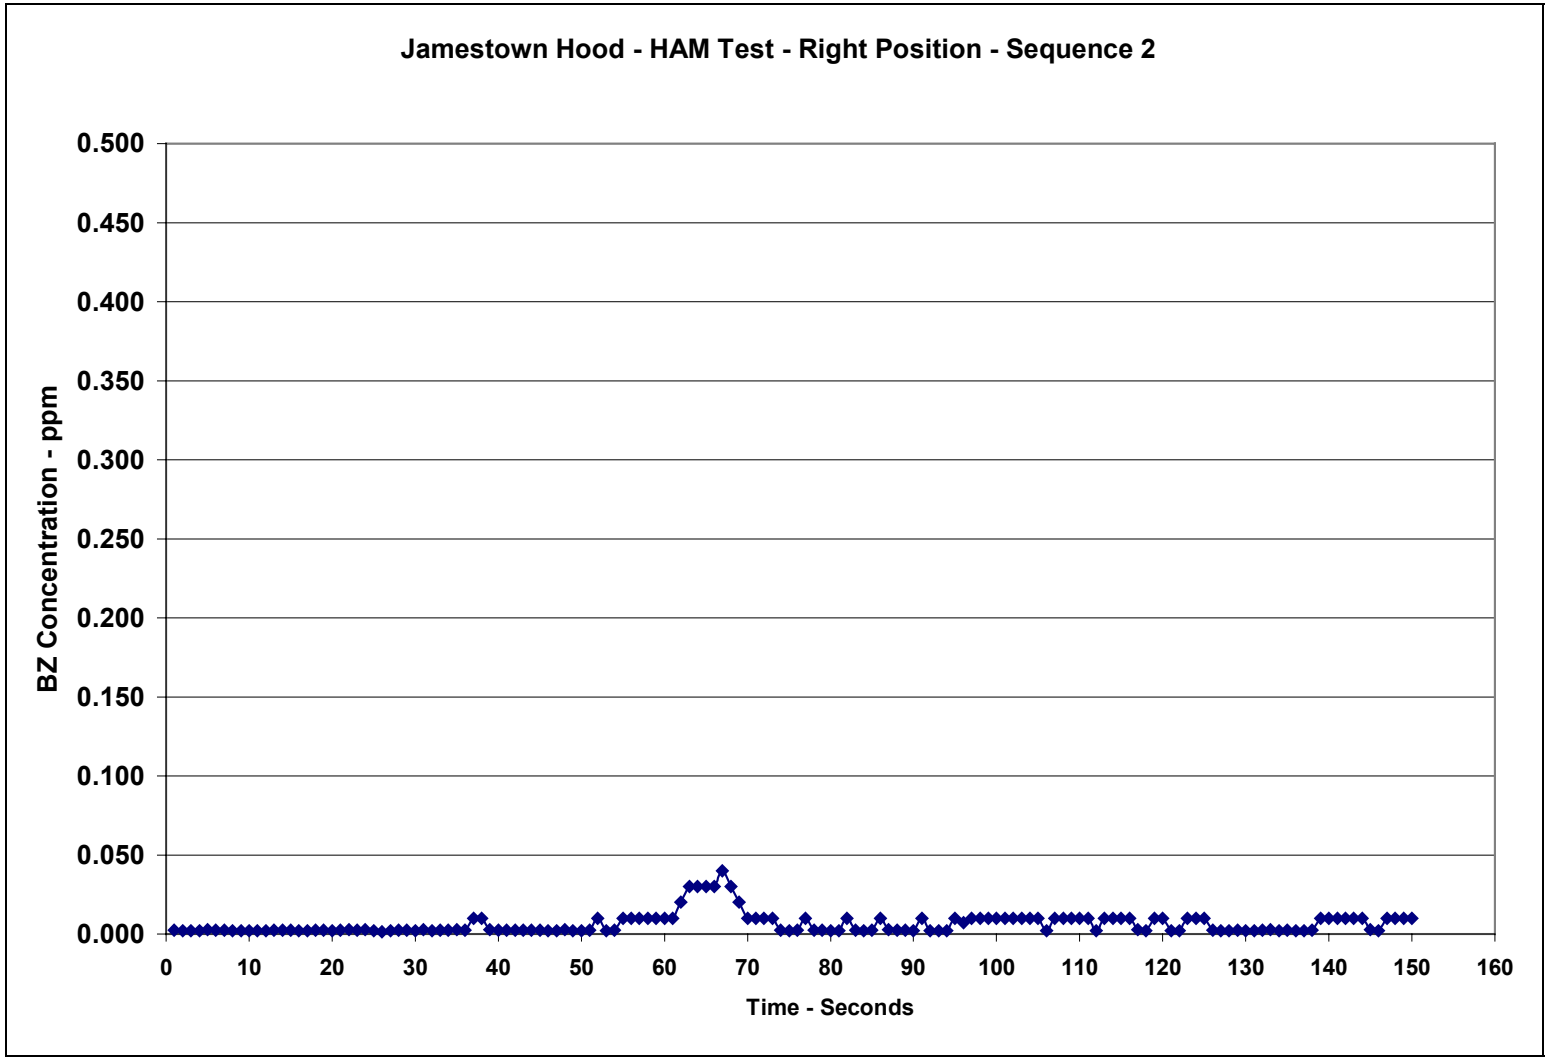

Figure 32 Tracer gas concentrations measured during test sequence 2 on the right side of the standard hood

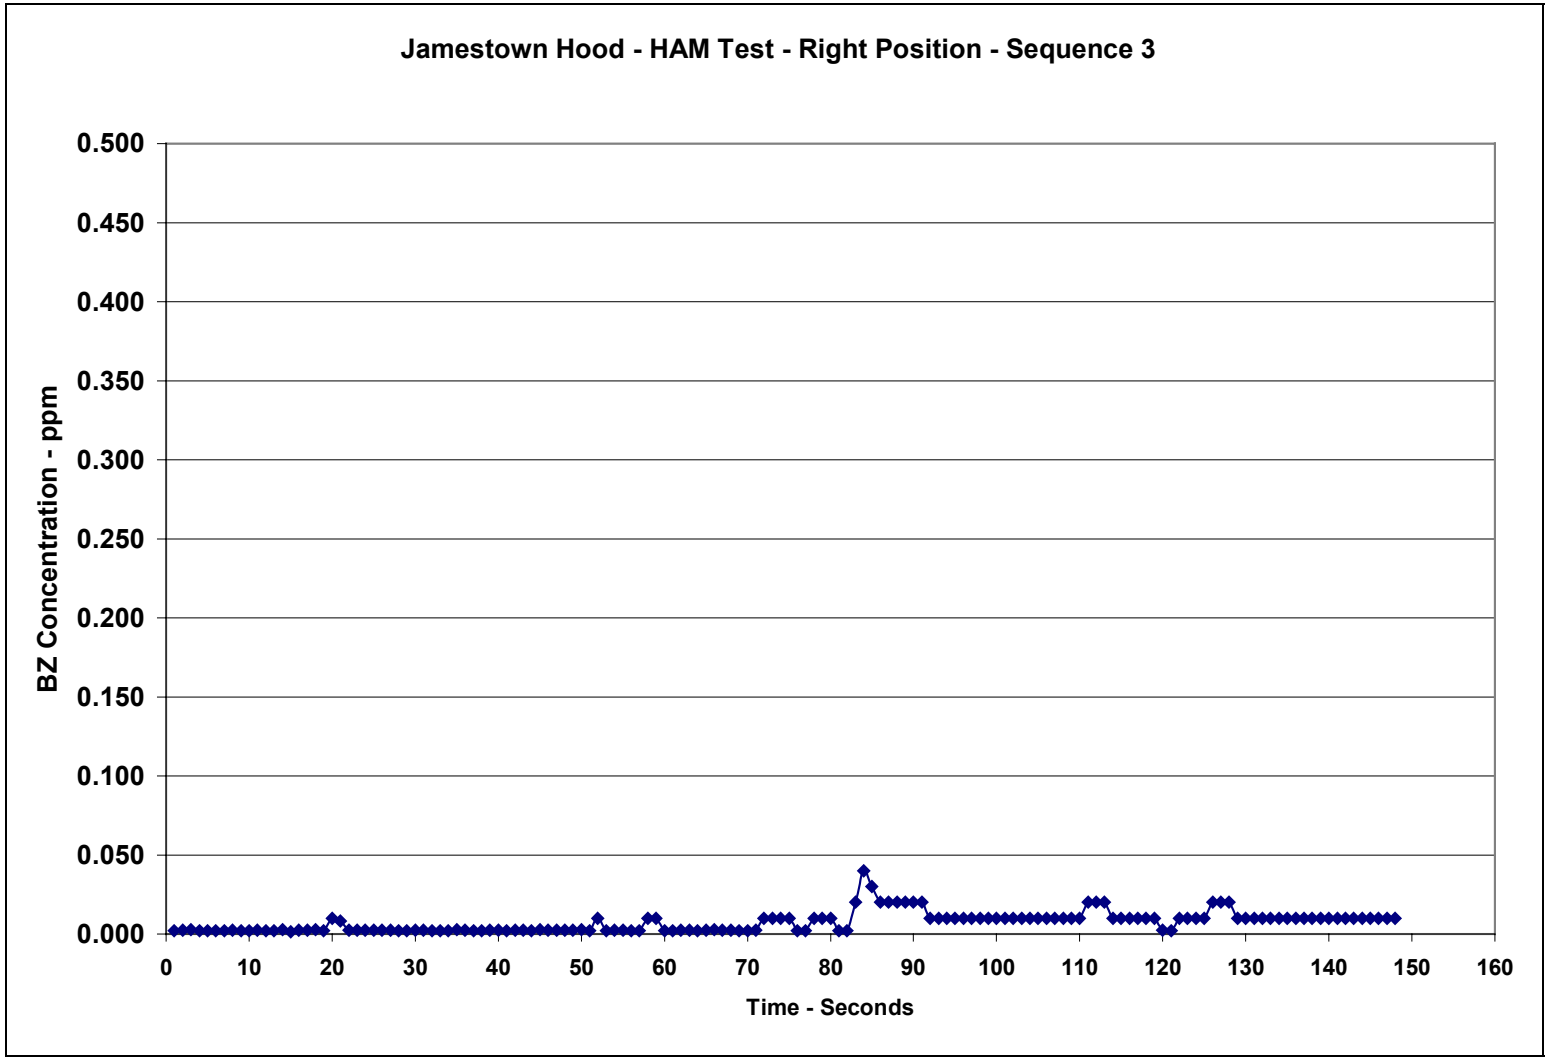

Figure 33 Tracer gas concentrations measured during test sequence 3 on the right side of the standard hood

Table 4 Calculation of exhaust flow, stack velocity and discharge concentration for each hood system

Hood	Flow cfm	Stack Velocity fpm	Exhaust Concentration ppm
LBNL	670	1228	194
Jamestown Standard	1200	2202	108

Figure 34 Photo of exhaust discharge from the exhaust stacks serving the LBNL hood and the standard hood.

Figure 35 Photo of exhaust discharge from the exhaust stacks serving the LBNL hood and the standard hood.

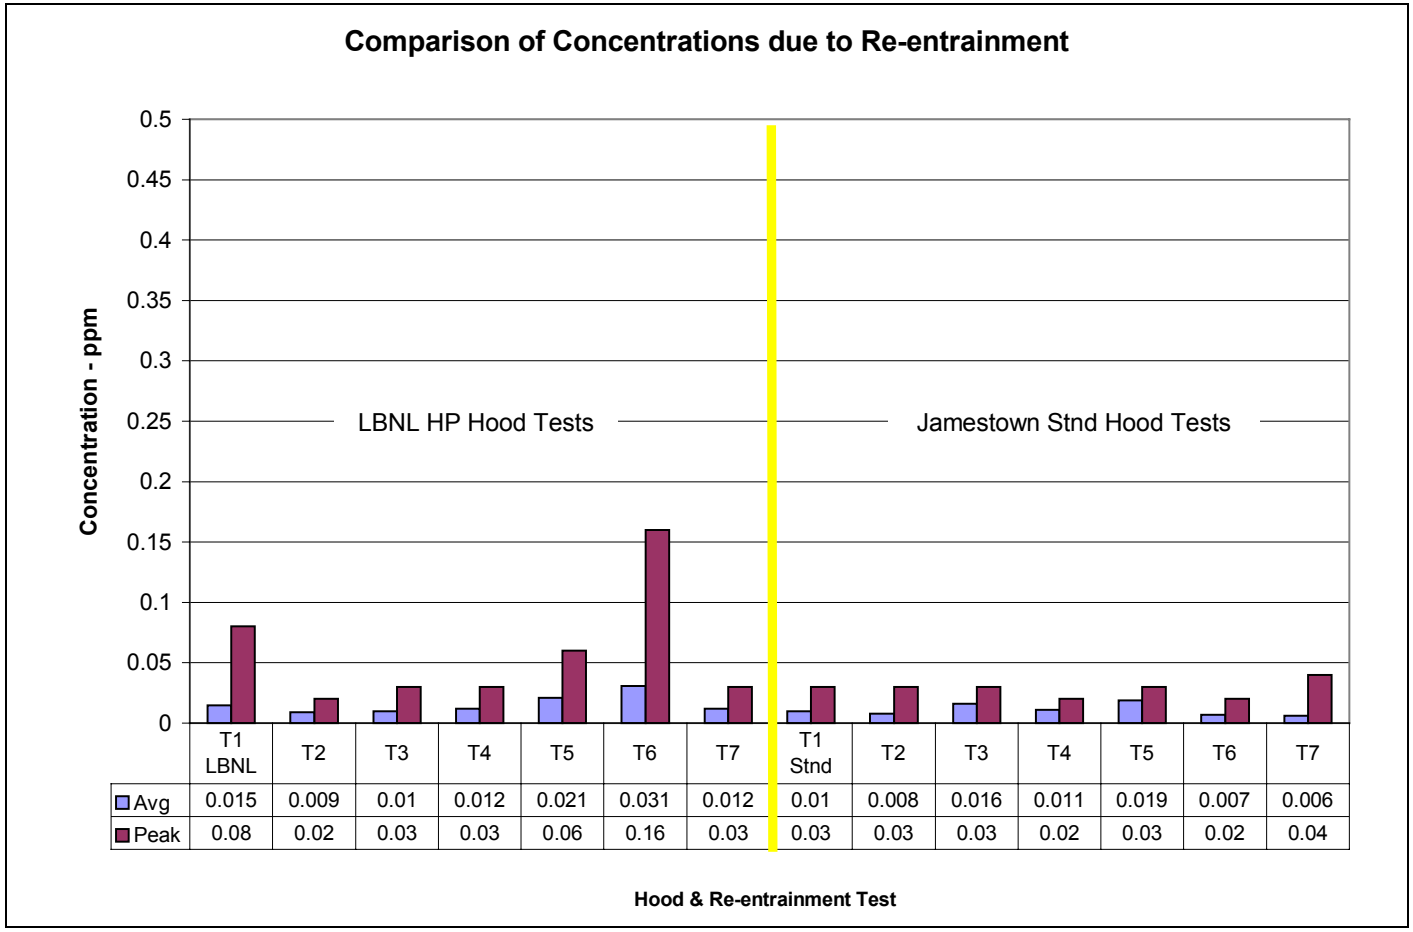

	T1 LBNL	T2	T3	T4	T5	T6	T7	T1 Stnd	T2	T3	T4	T5	T6	T7
Avg	0.015	0.009	0.01	0.012	0.021	0.031	0.012	0.01	0.008	0.016	0.011	0.019	0.007	0.006
Peak	0.08	0.02	0.03	0.03	0.06	0.16	0.03	0.03	0.03	0.03	0.02	0.03	0.02	0.04

Figure 36 Tracer gas concentrations attributed to re-entrainment during a series of tests conducted on the LBNL hood and the Jamestown standard hood.

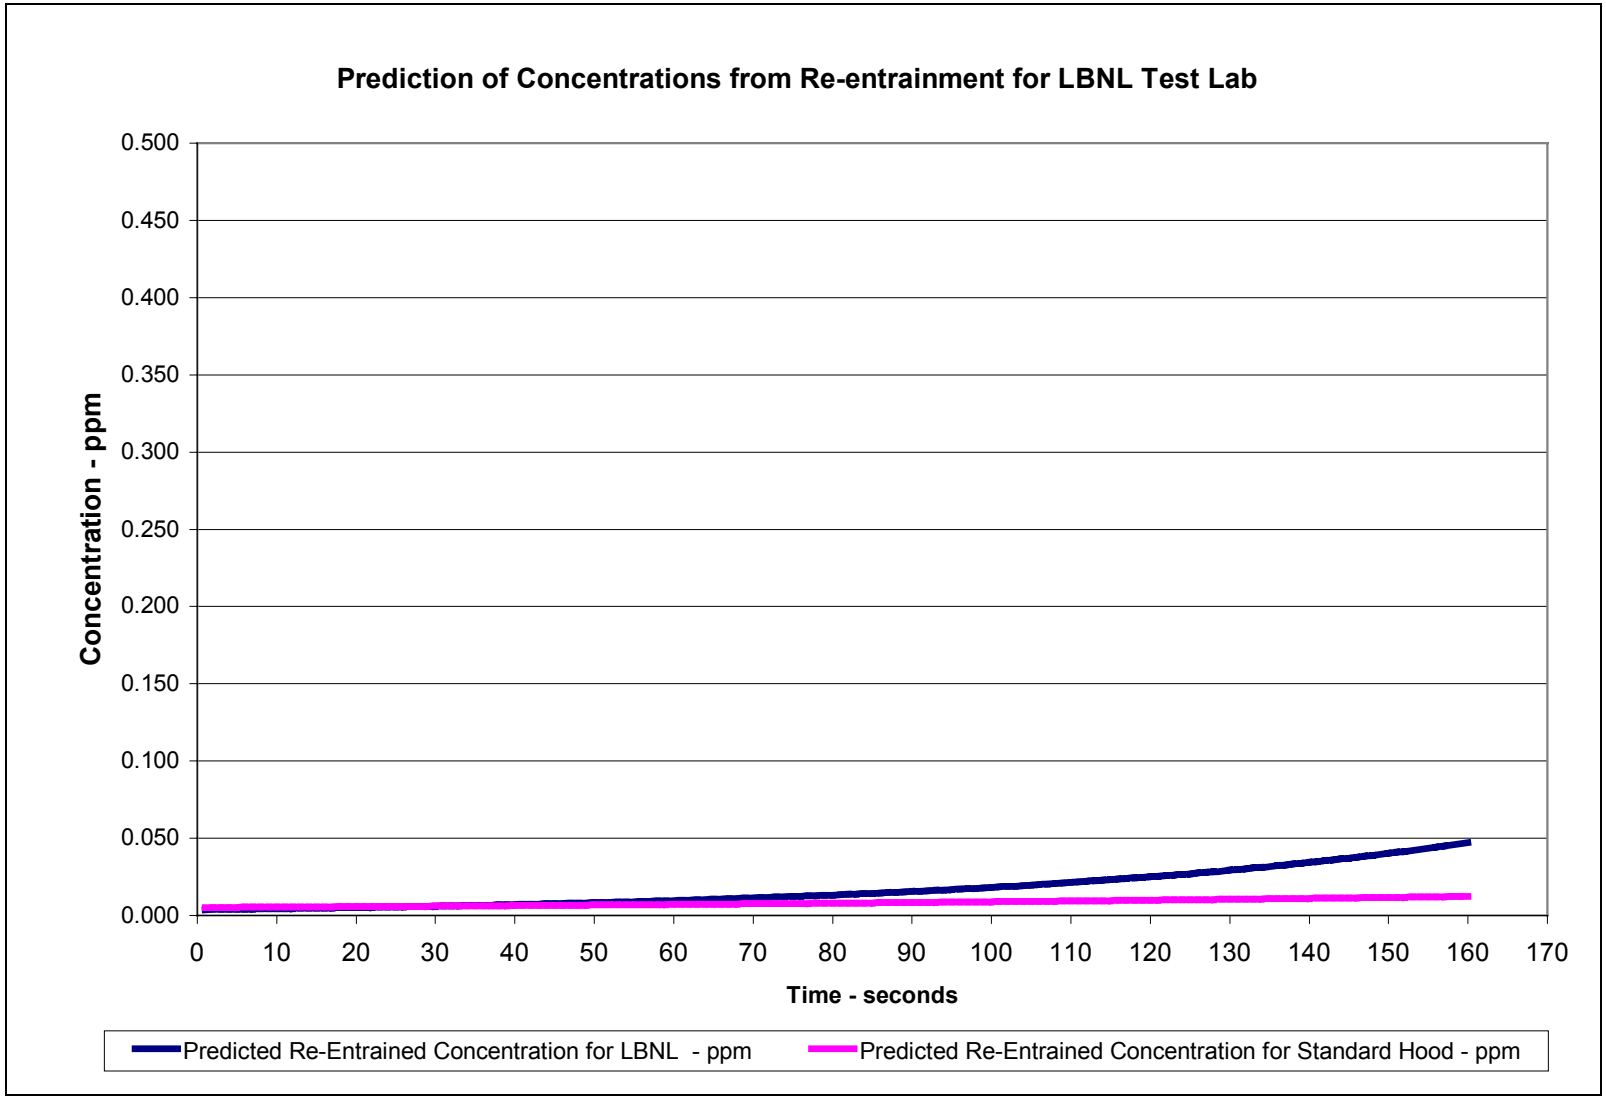

Figure 37 Curve of background concentrations predicted from a series of tests to evaluate re-entrainment.

Attachment VI: Data Center Benchmarking

Data Center Energy Benchmarking:

Part 5 - Case Studies on a Corporate Data Center (No. 22)

Final Report

August 2007

Final Report Prepared by

Tengfang Xu and Steve Greenberg
Lawrence Berkeley National Laboratory (LBNL)
Berkeley CA 94720

Draft provided by

EYP Mission Critical Facilities
Los Angeles, CA 90064
and
Landsberg Engineering, P.C.
Clifton Park, NY 12065

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Table of Contents

1 EXECUTIVE SUMMARY 1

2 REVIEW OF SITE CHARACTERISTICS 4

2.1 ELECTRICAL EQUIPMENT AND BACKUP POWER SYSTEM 4
2.2 MECHANICAL SYSTEM 5

3 ELECTRIC POWER CONSUMPTION CHARACTERISTICS 9

3.1 POWER SYSTEM 10
3.2 CHILLER SYSTEM 11
3.3 PUMPING SYSTEM 12
3.4 COOLING TOWERS AND PUMPS 13
3.5 EMERGENCY GENERATORS 14
3.6 AHU POWER CONSUMPTION 14

4 SYSTEM OPERATION 15

4.1 CHILLED WATER SUPPLY AND RETURN TEMPERATURES 16
4.2 AIR HANDLER UNIT SUPPLY AND RETURN TEMPERATURES AND RELATIVE HUMIDITY 16
4.3 DATA CENTER AIR CONDITIONS 19

5 OBSERVATIONS AND RECOMMENDATIONS 22

6 ACKNOWLEDGEMENTS 23

7 APPENDIX A: DATA FACILITY DEFINITIONS AND METRICS 24

8 APPENDIX B: FACILITY DIAGRAMS 26

1 Executive Summary

The data center in this study had a total floor area of 10,000 square feet (ft^2) with one-foot raised-floors. The data center housed 377 computer racks, and was located in a 110,000-ft^2 office building in Pasadena, California. However, the raised-floor was not utilized for cold air distribution. Communications and power wiring and fire sprinkler were located within the space above the ceiling. There were two standby generators, each rated at 1500 kW/kVA providing backup power supporting all building loads.

The building was served by a single 480-volt utility incoming service from Pasadena Water and Power Department (PWPD) serving the 4000-amp main distribution switchgear (see Figure 30 in Appendix B). The main switchgear fed four distribution switchboards. Two of the switchboards were supporting the critical loads via pairs of parallel-operated 500 kVA UPS modules. The other two switchboards supported the mechanical equipment and other miscellaneous building loads.

The data center operated on a 24 hour per day, year-round cycle, and users had all hour full access to the data center facility. Data center cooling was supplied by the main chilled water system to four air-handling units fitted with VSDs. The four air-handling units provided top flow ducted air supply to data center space. The AHU fan room return air was directly connected to the data center space, i.e., the return-air grille was on the wall separating the computer room from the return-air plenum. The AHUs had 4-inch throwaway air filters with 85% efficiency located at the front of the unit. The chilled water to the building was supplied by two Trane water-cooled chillers.

The study found that more than 80% of the electrical load was consumed by the data center: data center computer load accounted for 54% of the overall building electrical load. 21% of the load was consumed by AHUs, 5% by UPS', and 1% of the power was consumed by the data center lighting. The density of installed computer loads (rack load) in the data center was 57 W/ft^2. In addition, the data center accounted for additional cooling load provided by the chiller plants shared with the rest of the building. Power consumption density for all data center allocated load (including cooling and lighting) was 94 W/ft^2, approximately 10 times the average overall power density for the whole building.

For the data center, 60% of the overall electric power was the rack critical loads, 11% of the power was consumed by chillers, 22% by CRAH units, 6% by UPS', and 1% by lighting system.

General recommendations for improving overall data center energy efficiency include improving the lighting control, airflow optimization, control of mechanical systems serving the data center in actual operation. This includes chilled water system, airflow management and

control in data centers. Additional specific recommendations or considerations to improve energy efficiency are provided in this report.

The data center had a total power density of approximately 94 W/ft^2, while the installed computer power density was 57 W/ft^2. The highest density of energy usage is from the racks. Therefore, it is important to reduce the rack power usage. General recommendations for improving overall data center energy efficiency include improving the lighting control, design, operation, and control of mechanical systems serving the data center in actual operation. This includes chilled water system, airflow management and control in data centers.

The following includes potential measures to improve the energy efficiency of the data center.

- Optimize air balance including reducing fan-wheel speeds as necessary
- Review proper control algorithms to ensure proper controls for the VSDs and AHUs. In particular, the supply air temperatures were generally lower than needed at the IT equipment. Improved air management would allow optimization with large chiller plant energy savings, and perhaps fan energy savings as well.
- Review operation of AHU-3. This unit had a relatively lower return air temperature which might indicate a possible "short-circuiting" of supply air to return air
- Optimize air management by reviewing placement of supply and return registers
- Consider raising chilled water supply water temperature. This can be done automatically using a worst-zone reset strategy to operate at the highest chilled water temperature that would meet the load. If chilled water temperature can be raised, the chiller plant would be more efficient and it would reduce incidental dehumidification, thus reducing the need for humidification as well.
- Optimize the operation of both cooling tower fans – both running together may produce the same heat rejection with much less energy than operating one single fan. In addition, consider resetting the cooling tower water temperature set points according to chiller manufacturer recommendations and/or outside air wet-bulb temperature plus an approach temperature. For any given load and wet-bulb temperature, there is an optimal combination of flow and temperature that can minimize overall chiller plant energy use. Software is available from manufacturers or consultants to optimize this algorithm.
- Install VSDs on the CHW primary pumps and the condenser water pumps and control per a chiller plant optimization routine. Check the operation of the generator jacket heaters. At least one heater appears to operate much more than necessary. The controls should be adjusted or replaced as needed.

- Check to see if the double-conversion UPS units can be operated in a bypass mode without sacrificing protection from power interruption.

2 Review of Site Characteristics

The corporate data center in this study had a total floor area of 10,000 square feet (ft^2) with one-foot raised-floors. The data center housed 377 computer racks, and was located in a 110,000-ft^2 office building in Pasadena, California. However, the raised-floor was not utilized for cold air distribution. Communications and power wiring and fire sprinkler were located within the space above the ceiling. There were two standby generators each rated at 1500 kW/kVA providing backup power supporting all building loads.

The building was served by a single 480-volt utility service from Pasadena Water and Power Department (PWPD) serving the 4000-amp main distribution switchgear (Figure in Appendix B). The main switchgear fed four distribution switchboards. Two of the switchboards supported the critical loads via pairs of parallel-operated 500 kVA UPS modules. The other two switchboards supported the mechanical equipment and other miscellaneous building loads.

The data center operated on a 24 hour per day, year-round cycle, and users had all hour full access to the data center facility. The data center did not have a dedicated chiller system but was served by the main building chiller plant connected to four air-handling units (AHUs) fitted with VSDs. The four air-handling units provided top flow ducted air supply to data center space. The AHU fan room return air was directly connected to the data center space, i.e., the return-air grille was on the wall separating the computer room from the return-air plenum. The AHUs had 4-inch throwaway air filters with 85% efficiency located at the front of the unit. The chilled water to the building was supplied by two Trane water-cooled chillers.

The data for this benchmarking exercise was collected during a one week monitoring period in December 2004. The data was collected using the following measurement instruments: HOBO loggers, Elite loggers, Fluke 41B power meter, K-20 and the building's Automatic Logic Corporation (ALC) system.

2.1 Electrical Equipment and Backup Power System

Electrical power to Data Center #22 was served by the building electrical distribution system, which received PWPD power via a three-phase, four-wire, 480/277-V main service, terminated to a 4,000A main distribution switchboard "MS." The main switchboard then sub-fed four separate distribution boards: "TPSA", "TPSB", "TPSC" and "TPSD." Distribution boards "TPSA" and "TPSB" supported the building data center critical loads, each distribution board serving four 500 kVA, double-conversion UPS modules. The diagram in Appendix B indicates two 500 kVA. The conditioned power from the UPS modules was distributed to the critical loads via two 3,000A, paralleling switchboard: "UPSA" and "UPSB". Distribution boards

"TPSC" and "TPSD" supported the mechanical equipment, lighting and other miscellaneous loads within the building.

The backup generation system consisted of two 1,500kW standby generators connected to a 4,000A paralleling switchgear "PS". There were two main distribution feeds from the paralleling switchgear, each terminated to a standby distribution bus located in the main switchboard "MS". The backup generation power supported the critical loads and the mechanical loads via normally-open circuit breakers that would be automatically closed during the normal power service interruption. There was no data available for PDUs.

A load monitoring device at the time of survey at the service entrance feeder of the main switchboard "MS" had indicated a reading average demand load from the utility service of 1,040 kW.

The backup generation system consists of two 1,500kW standby diesel generators - Model 3512B Engine, Model SR-4B.connected to a 4,000-A paralleling switchgear "PS". The generators had a total of four 15-kW block heaters. There were two 500 kW resistive load banks associated with each generator for testing. There were two main distribution feeds from the paralleling switchgear, each terminated to a standby distribution bus located in the main switchboard "MS". The backup generation power supported the critical loads and the mechanical loads via normally-open circuit breakers that would be automatically closed during the normal power service interruption.

2.2 Mechanical System

2.2.1 Chiller

The data center, along with the remainder of the building, was designed to be cooled by two 320-ton Trane CVHE 0450 multi-stage centrifugal direct drive water-cooled chillers. The chillers were arranged in a parallel configuration. The chillers were manually operated in a lead-lag alternate manner: only one chiller was in operation at the time of the study. The chillers were equipped with VSDs and were rated at 0.55 kW per cooling ton at 100% load, 0.44 kW per cooling ton at 75% load, and 0.33 kW per cooling ton at 30% load. The chiller had an operating chilled water temperature set point of 42°F and was controlled by the chilled water return temperature set point of 50°F through a Trane Tracer Summit system. The chilled water plant used chilled water reset in an effort to reduce chiller energy consumption.

2.2.2 Chilled Water Pump

Primary chilled water was circulated by two parallel Bell and Gossett centrifugal pumps, each with a motor capacity of 30 HP. The primary chilled water pumps were identified as CHWP-6 and CHWP-7. Only CHWP-7 was operating at the time of survey. The pump’s discharge pressure was 48 psig and the suction pressure was 25 psig. The pumps were constant speed and

were controlled by temperature sensors through the facilities building management system. The building operator manually switched the chiller from lead to lag operation. No information on specific sequence of operation regarding chiller and pump operation was available.

Figure 1 shows the chillers and chilled water pumps. Chilled water was circulated through the building secondary pumps by two 60-hp parallel–in line Bell and Gossett centrifugal pumps. The pumps were identified as CHWP-4 and CHWP-5. Both of the pumps were operating at the time of the survey. The pump discharge pressure gage reading was 44 psig and the suction pressure was 23 psig. The pumps were provided with variable speed drives, and were controlled via BMS based on chiller operation. Exact control sequence was not provided.

Figure 1 Chillers and chilled water pumps

2.2.3 Condenser Water Pump

Condenser water was supplied to the cooling towers by two 25-hp Bell and Gossett split case condenser water pumps. The pumps were arranged in parallel and identified as CWP-1 and CWP-2. Both pumps were constant speed and only CWP-2 was operating at the time of the survey. The condenser water pump discharge pressure gage reading was 68 psig and a suction pressure was 34 psig. The pumps were controlled by BMS based on chiller operation. Exact sequence of control was not recorded. At time of measurements, only one chiller was in operation.

2.2.4 Cooling tower

There are two 320-ton blow-through, closed-loop Evapco cooling towers Model UBW 207 with 30-hp fan motors. The cooling towers shown in Figure 2 were arranged in parallel and fitted with 30-hp fans and VSDs. The towers are identified as CT-1 and CT-2. Each tower had a capacity of 4,838 MBH with a design wet bulb temperature of 76°F, with entering water temperature of 97°F and leaving water temperature of 85°F. CT-1 was operating at the time of survey. The cooling tower fan speed was controlled by the tower's leaving water temperature. The exact set point reset schedule of condenser water supply temperature was however not available.

Figure 2 Cooling towers

2.2.5 Air Handling Unit

The data center was served by four Cel-Air air-handling units. Each air-handling unit was powered by a 40-hp fan motor fitted with a VSD. Each air handler had a cooling capacity of 350 MBH. Each air-handling unit had an air supply capacity of 30,000 cfm and the measured W/cfm was 1.7. The exact sequence of operation was not available. The air-handling units were identified as AHU-1 to AHU-4 and located at the fan room provided with a Very Early Smoke Detection Alarm (VESDA) smoke detection system.

The space temperature set point was 70°F. The air handlers are fitted with electric steam humidifiers. Exact control sequence for space relative humidity control was not provided. Each air handling unit supplied cold air to the ceiling space used as a supply plenum and cold air was distributed to the data center by overhead ceiling diffusers. There were no under-floor

perforated tiles to supply cold air to this data center. The air return louvers of the fan room were directly connected to data center space facing south as shown in Appendix B.

AHU-1 had an average supply air temperature of 52°F and an average supply relative humidity of 61%. The return air temperature averaged 74°F with 31% RH. AHU-2 had an average supply air temperature of 53°F and an average supply relative humidity of 64%. The return air temperature averaged 74°F with 26% relative humidity. AHU-3 had an average supply air temperature of 53°F and an average supply relative humidity of 66%. The return air temperature averaged 69°F with 36% relative humidity. AHU-4 had an average supply air temperature of 53°F and an average supply relative humidity of 64%. The return air temperature averaged 73°F with 31% relative humidity. Figure 3 shows the AHU return-air plenum.

Figure 3 AHU return-air plenum

The UPS Room was served by two 15-hp air-handling units fitted with VSDs. Both units - AHU-5 and AHU-6 were operating at the time of the survey. Each AHUs had a design capacity of 22,500 cfm and measured W/cfm of 0.41.

The UPS/Battery Room was served by two 7.5-hp air handling units fitted with VSDs. Both units were operating at the time of the survey. AHU-7 and AHU-8 had capacities of 6,500 cfm each and measured W/cfm of 0.19.

The Electric Room was served by two 3-hp air-handling units fitted with VSDs. Both units were operating at the time of the survey. AHU-9 and AHU-10 had capacities of 3,750 cfm each and measured W/cfm of 0.45.

3 Electric Power Consumption Characteristics

Table 1 shows the end-use electricity demand of the building housing the data center in this study, based on the measurements taken during the monitoring period. The average building electrical load of 1039 kW was recorded from building instruments. The table also includes power density for the square foot area served by each load. Table 2 further shows the power demand and power density of the load within the data center.

From these measurements, it was observed that more than 80% of the electrical load was consumed by the data center: 54% of the load was consumed by the data center computer equipment, 21% of the load was consumed by AHUs, 5% by UPS', and 1% of the power was consumed by the data center lighting. In addition, the data center accounted for additional cooling load provided by the chiller plants shared with the rest of the building.

Table 2 indicates that for the data center, 60% of the overall electric power was the rack critical loads, 11% of the power was consumed by chillers, 22% by all AHUs, 6% by UPS', and 1% by lighting system. The density of installed computer loads (rack load) in the data center was 57 W/ft^2. Power consumption density for all data center load (including cooling and lighting) was 94 W/ft^2, approximately ten times the average power density of the overall building.

Table 1. End-Use of Electricity of the Data Center Building

Description	Electric power demand (kW)	Share of electric energy use (%)	Floor Space (ft2)	Electric power density (W/ft2)
Overall Building Load	**1,039**	**100%**	**110,000**	**9.5**
Data Center Rack Power	**565**	**54%**	**10,000**	**56.5**
Data Center UPS Losses	**44**	**4%**	**10,000**	**4.4**
AHUs Data Center	**205**	**20%**	**10,000**	**20.5**
AHUs UPS/Battery/Elec	**12**	**1%**	**10,000**	**1.2**
Total Chiller Plant	**142**	**14%**	**110,000**	**1.3**
Building Chillers	82	**8%**	110,000	0.8
Building Pumps	49	**5%**	110,000	0.5
Building Cooling Towers	10	**1%**	110,000	0.1
Building Generator Heaters	**59**	**6%**	**110,000**	**0.5**
Data Center Lighting	**13**	**1%**	**10,000**	**1.2**

Table 2. End-Use of Electricity of the Data Center

Description	Electric power demand (kW)	Share of electric energy use (%)	Floor Space (ft2)	Electric power density (W/ft2)
Data Center Rack Power	**565**	**60%**	**10,000**	**56.5**
Data Center UPS Losses	**44**	**5%**	**10,000**	**4.4**
AHUs Data Center	**205**	**22%**	**10,000**	**20.5**
AHUs UPS/Battery/Elec	**12.2**	**1%**	**10,000**	**1.2**
Building Chiller Plant (A)	**106**	**11%**	**10,000**	**10.6**
Building Chillers (A)	**61**	**7%**	**10,000**	**6.1**
Building Pumps (A)	37	**4%**	10,000	3.7
Building Cooling Towers (A)	8	**1%**	10,000	0.8
Data Center Lighting	13	**1%**	10,000	1.3
Total Data Center Only	**944**	**100%**	**10,000**	**94.4**

3.1 Power System

The main electrical incoming service provided support to the critical loads via a pair of redundant 3,000A, 480V main service buses, feeding multiple panels within the building. The

data center was fed through four 500 kVA UPS units. The chillers, pumps, CRAC units lighting and miscellaneous loads were fed from other panels.

3.2 Chiller System

Electric power demand was monitored for the two operating chillers within a one-week period. Chiller CH-1 provided the majority of the cooling load with average power consumption of 82 kW. Chiller CH-2 was in operation occasionally. The total chiller power consumption for the period average 82 kW. For the monitoring period, the ambient air temperature averaged 61°F, while the ambient relative humidity averaged 61%.

Figure 4 shows the scatter plot of total chiller power consumption and outside air temperature. Except for a few occasions when two chillers were simultaneously on, the variations in total chiller power demand showed essentially non-significant correlation with the outdoor air temperature. As the outside air temperature changed, the actual chiller power demand didn't exhibit similar changes. Since most of the load on the chiller is as a result of the data center, the impact of ambient temperature on cooling load is minimal.

The primary chilled water loop was designed for 640 gpm water flow rate. With an average temperature differential over the monitoring period of 8.3°F, the average cooling load was about 220 tons.

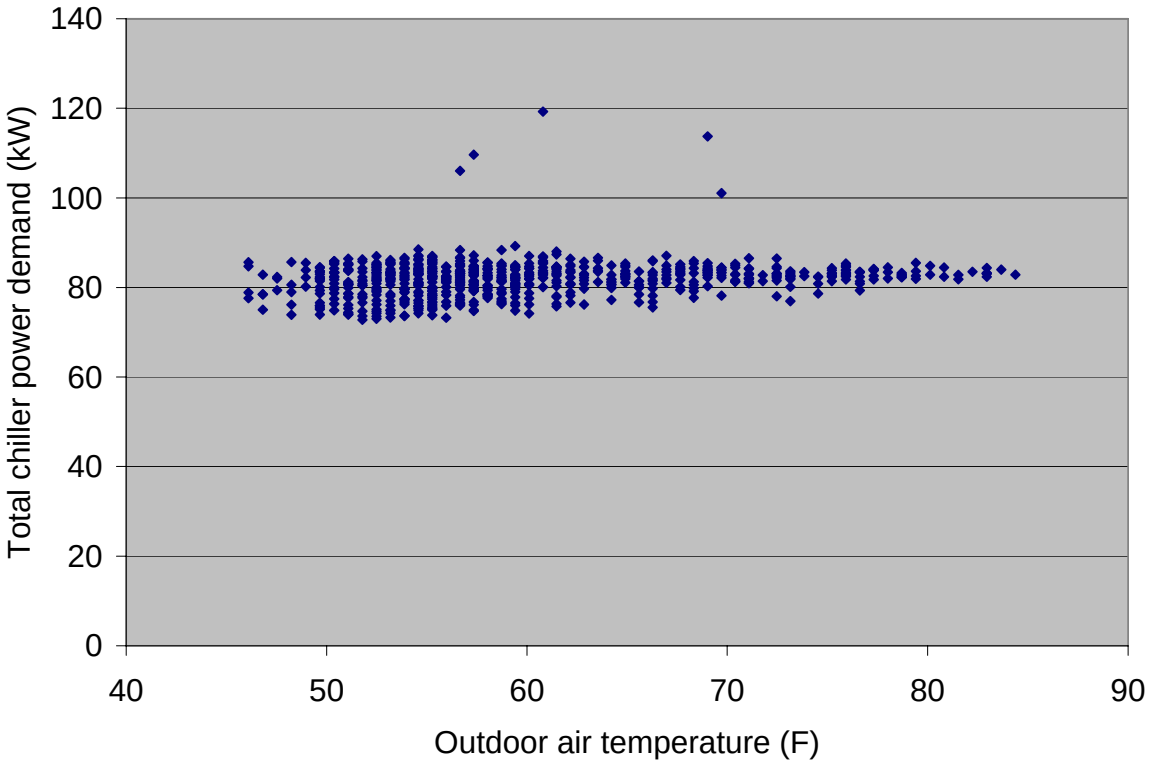

Figure 4 Correlation of chiller cooling power and outdoor air temperature

3.3 Pumping System

The building was served by two 30-hp primary chilled water pumps: CHWP-6 and CHWP-7. Only one pump was operational at any given time. For the pump monitored (CHWP-7 controlled with Chiller 2), the power demand of the primary pump was 19 kW when in operation . The primary chilled water loop was designed at 640 gpm.

The secondary chilled water pumps (CHWP-5 and CHWP-6) were 60-hp units fitted with VSDs. The secondary chilled water loop was designed at 1,280 gpm at full load. The average power demand was 8 kW for CHWP-5 and 10 kW for CHWP-6 (Figure 5). The recorded data shows that the chilled water pumps were in operation in majority of the time during the study but one of the pump appeared to stop operation since Monday afternoon (Jan 24, 2005), while the other pump joined the idleness starting in the mid-week, Wednesday afternoon (January 26, 2005). This indicates that it's advisable to check the operation and data acquisition system.

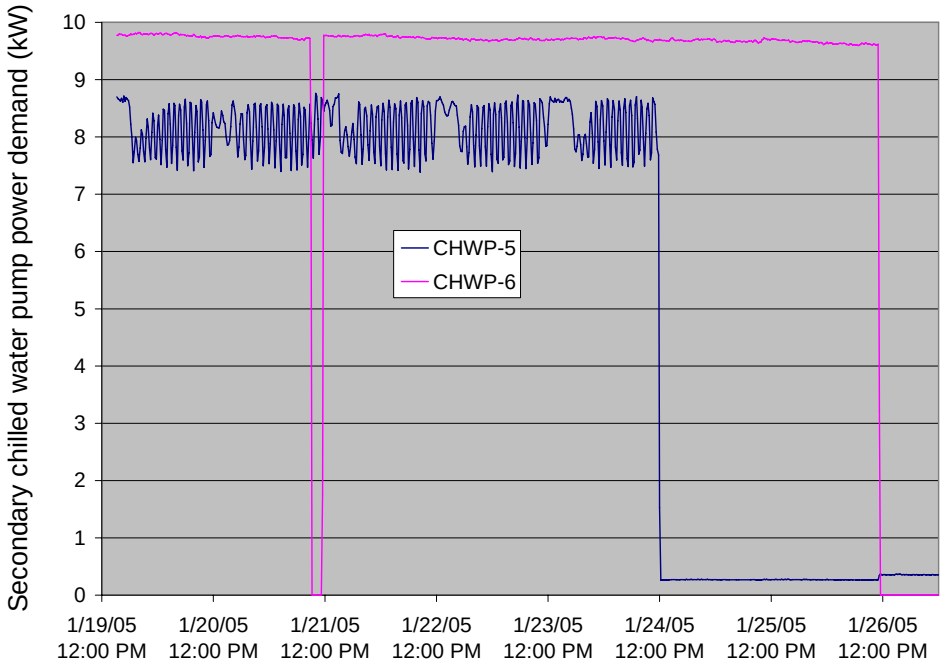

Figure 5 Chilled Water Pump Power Demand

Condenser water pumps (CWP-1 and CWP-2) were designed for 800 gpm at full load. Both pumps were 25-hp units fitted with motor starters. The pumps operated alternately, with an average power demand of approximately 17kW when the pumps were operating (Figure 6).

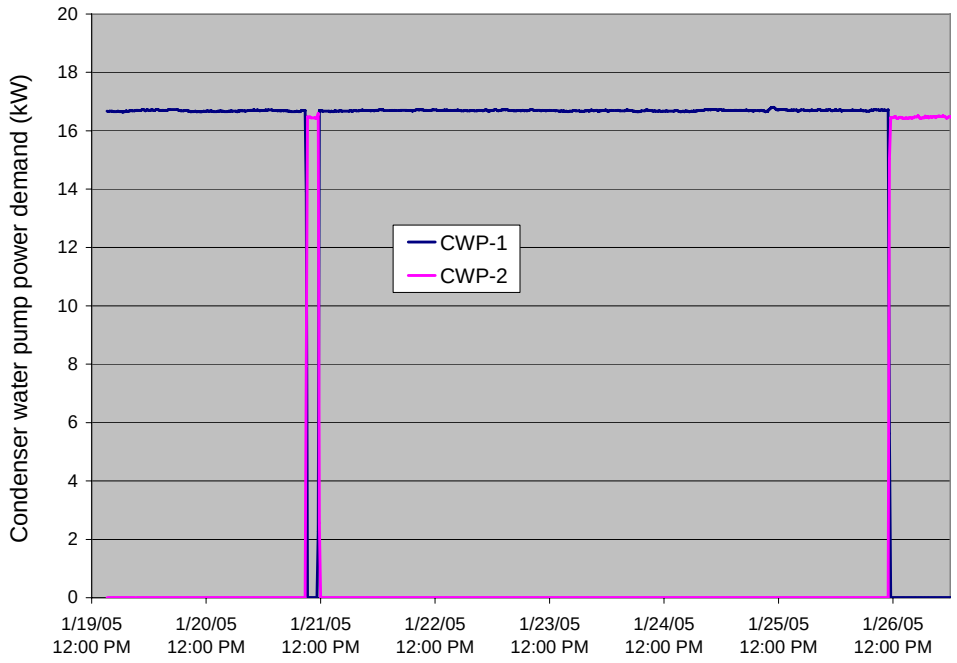

Figure 6 Condenser Water Pump Power Demand

3.4 Cooling Towers and Pumps

The two cooling towers, CT-1 and CT-2 each had one 30-hp fan fitted with a VSD. The units operated alternately at the time of the survey. Figure 7 shows the power demand of the two cooling tower fans that were in alternate operation during the study. Diurnal variation in power consumption was apparent. Power consumption averaged 10.4 kW for the monitoring period, with highs near 25 kW during afternoon peak hours, and lows under 4 kW toward early mornings.

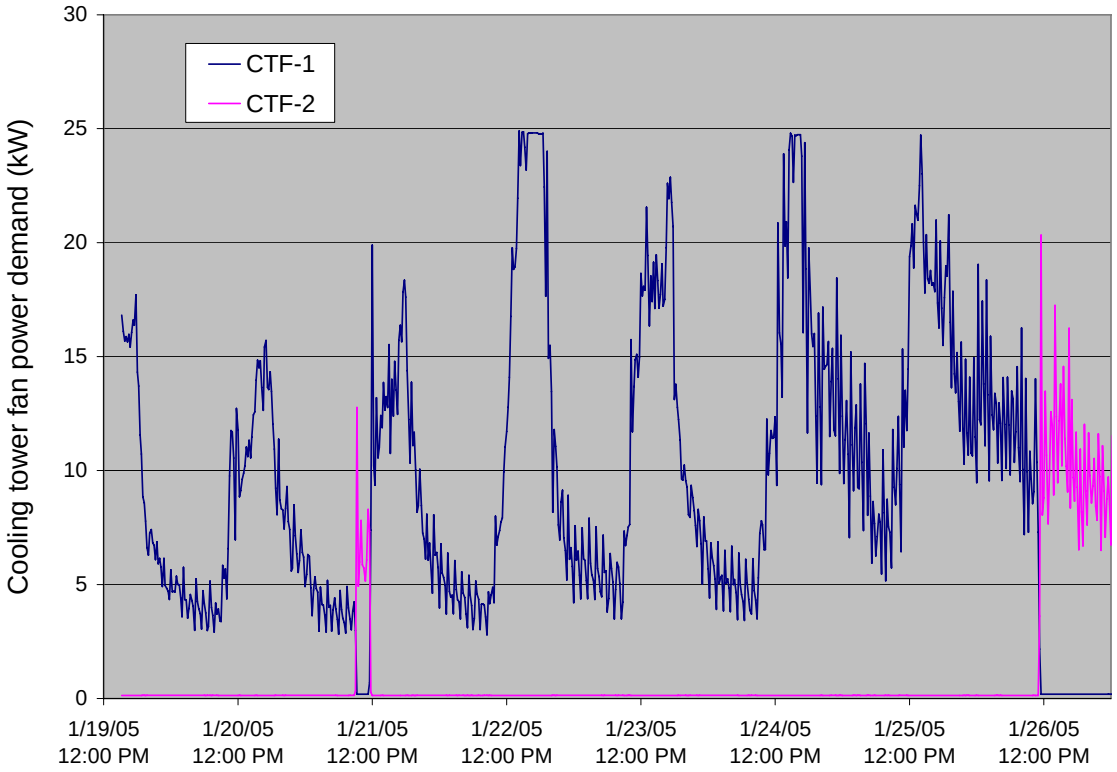

Figure 7 Cooling tower fan power

3.5 Emergency Generators

There were two 1,500 kW Caterpillar engine electric generators. The units were fitted with two 15-kW water jacket heaters each. We monitored the power use of one of the jacket heaters, which were on almost continually and averaged 14.8 kW over the monitoring period. Assuming that four heaters were on in the similar pattern, the total power demand from the heaters was 59 kW.

3.6 AHU Power Consumption

The data center was served by four air handling units, supplying 30,000 cfm each and fitted with 40-hp fans and VSDs. All four units were in operation at the time of the study. Power demand of the four units were recorded and showed an average of 1.7 W per cfm supplied.

The UPS Room was served by two 15-hp air-handling units fitted with VSDs. Both AHUs (AHU-5 and AHU-6) were in operation at the time of the study. AHU-5 and AHU-6 had capacities of 22,500 cfm each and measured W/cfm of 0.41.

The UPS/Battery Room was served by two 7.5-hp air handling units fitted with VSDs. Both units (AHU-7 and AHU-8) were in operation at the time of the survey. AHU-7 and AHU-8 had capacities of 6,500 cfm each and measured W/cfm of 0.17. Figure 8 shows the power demand of AHU 8 over the monitoring period, with a minimum of 1.05 kW to a maximum of 1.17 kW and an average of 1.1 kW.

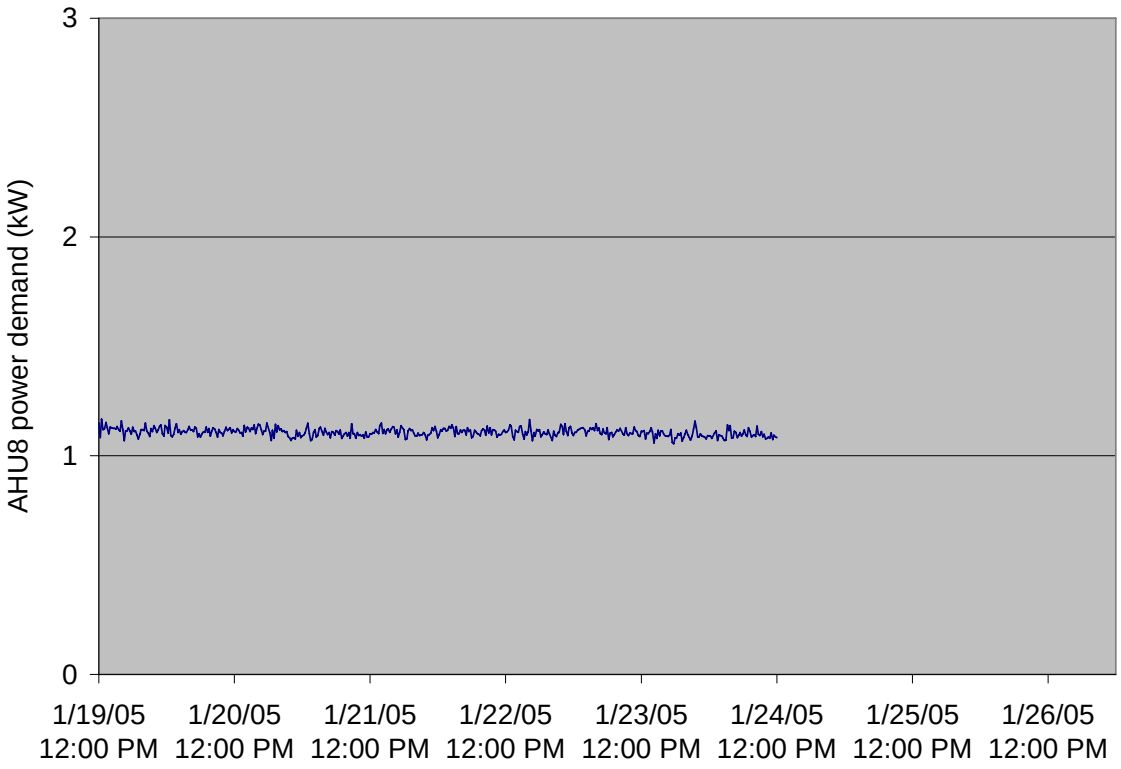

Figure 8 AHU8 power demand

The Electric Room was served by two 3-hp air handling units fitted with VSDs. Both units (AHU-9 and AHU-10) were in operation at the time of the study. AHU-9 and AHU-10 had capacities of 3,750 cfm.

4 System Operation

During the one-week monitoring period, the following HVAC equipment was operating:

- Primary chilled water pumps CHWP-6 and CHWP-7 (alternating)
- Secondary chilled water pumps CHWP-4 and CHWP-5
- Chillers CH-1, CH-2 (alternating)
- Condenser Pumps CWP-1 and CWP-2 (alternating)
- Cooling tower fans CTF-1 and CTF-2 (alternating)

- All four data center air handling units

4.1 Chilled Water Supply and Return Temperatures

The chilled water supply and return temperatures were recorded. For the monitoring period, the average chilled water supply temperature was 42°F while the average chilled water return temperature was 50°F. This produced an average temperature differential of 8°F. The temperatures were stable during the monitoring period.

4.2 Air Handler Unit Supply and Return Temperatures and Relative Humidity

Figure 9 shows supply air temperature and relative humidity for AHU-1. The supply air temperature averaged 52°F while the supply relative humidity averaged 67% RH. The return air temperature averaged 74°F, while the return air relative humidity averaged 31% RH.

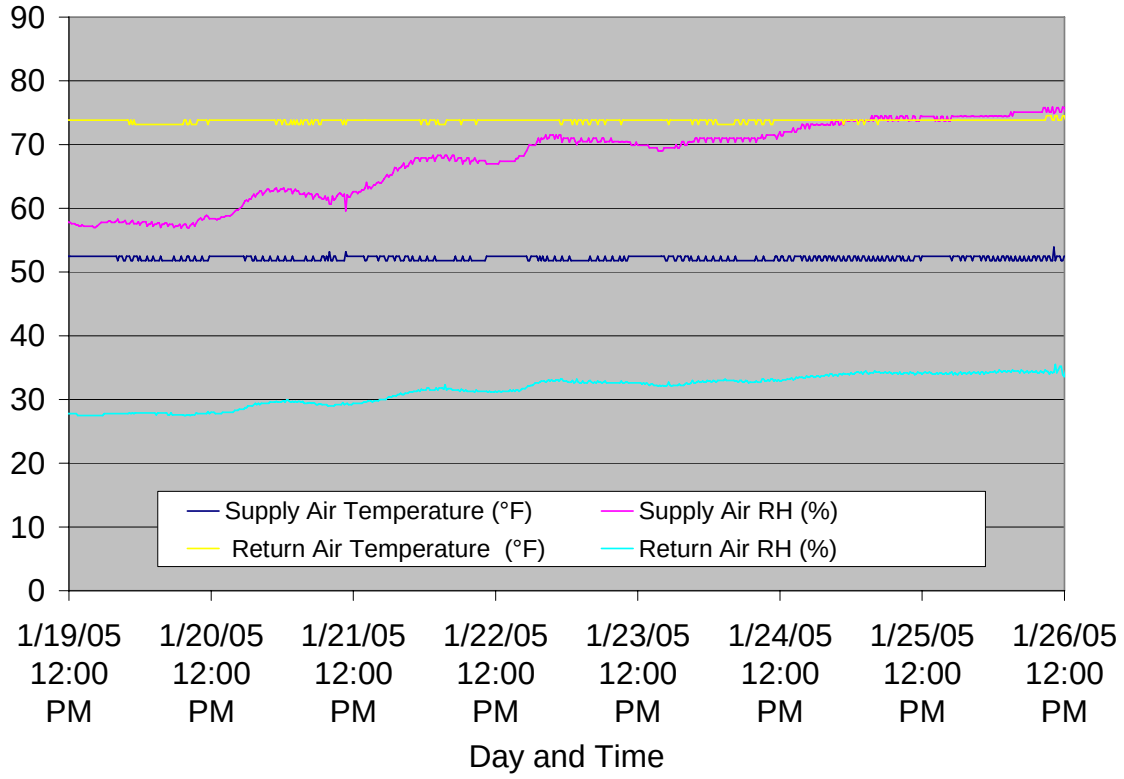

Figure 9 AHU-1 Air Temperature and Humidity

Figure 10 shows supply air temperature and relative humidity for AHU-2. The supply air temperature averaged 53°F while the supply relative humidity averaged 64%. The return air temperature averaged 74°F, while the return air relative humidity averaged 26%.

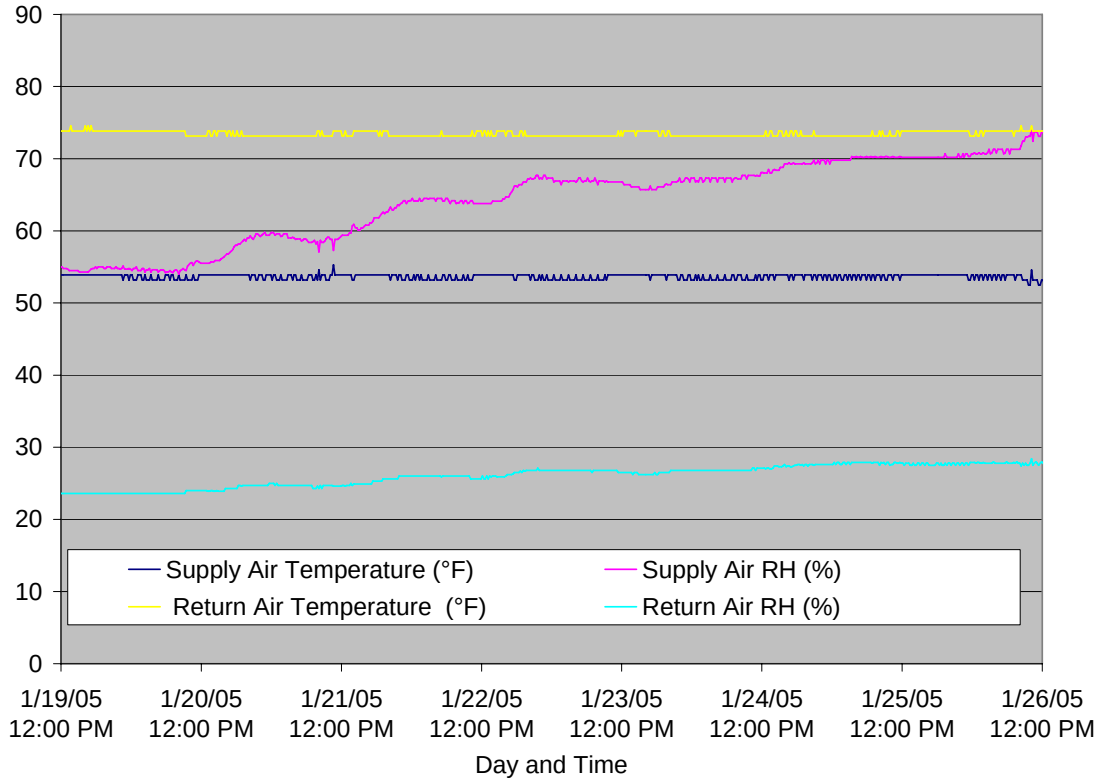

Figure 10 AHU-2 Air Temperature and Humidity

Figure 11 shows supply air temperature and relative humidity for AHU-3. The supply air temperature averaged 53°F while the supply relative humidity averaged 66%. The return air temperature averaged 69°F, while the return air relative humidity averaged 36%.

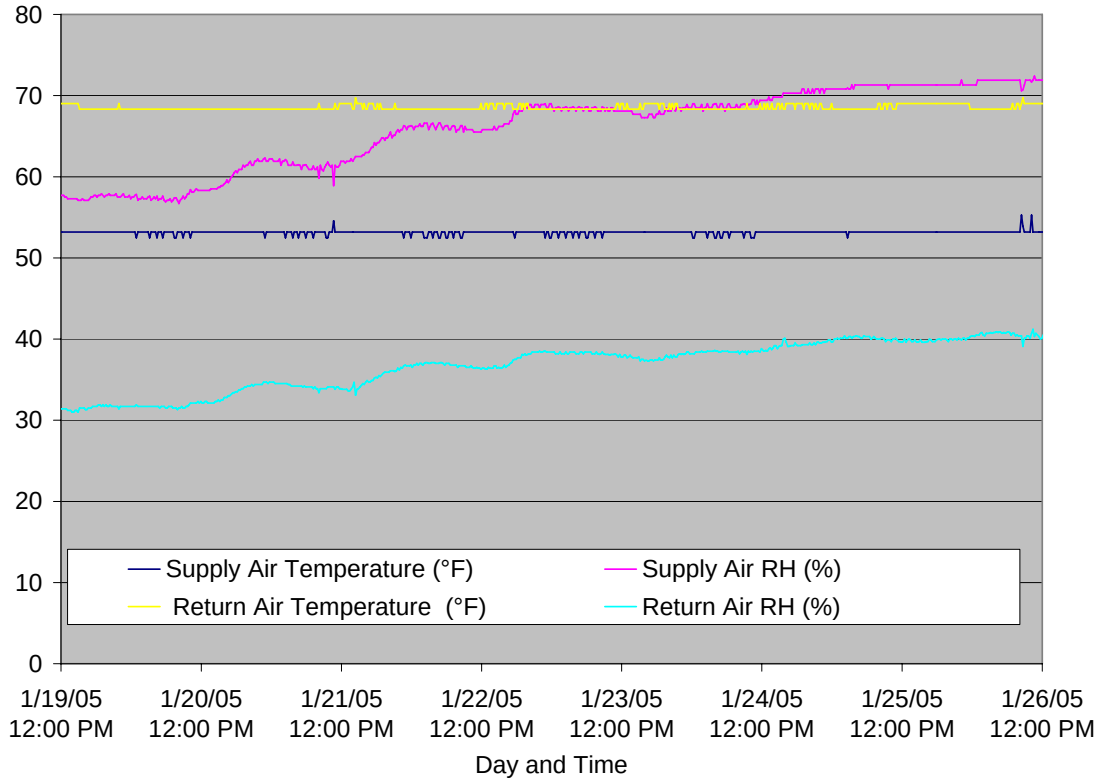

Figure 11 AHU-3 Air Temperature and Humidity

Figure 12 shows the return air temperature and relative humidity for AHU-4. The supply air temperature averaged 53°F while the supply relative humidity averaged 64%. The return air temperature averaged 73°F, while the return air relative humidity averaged 31%.

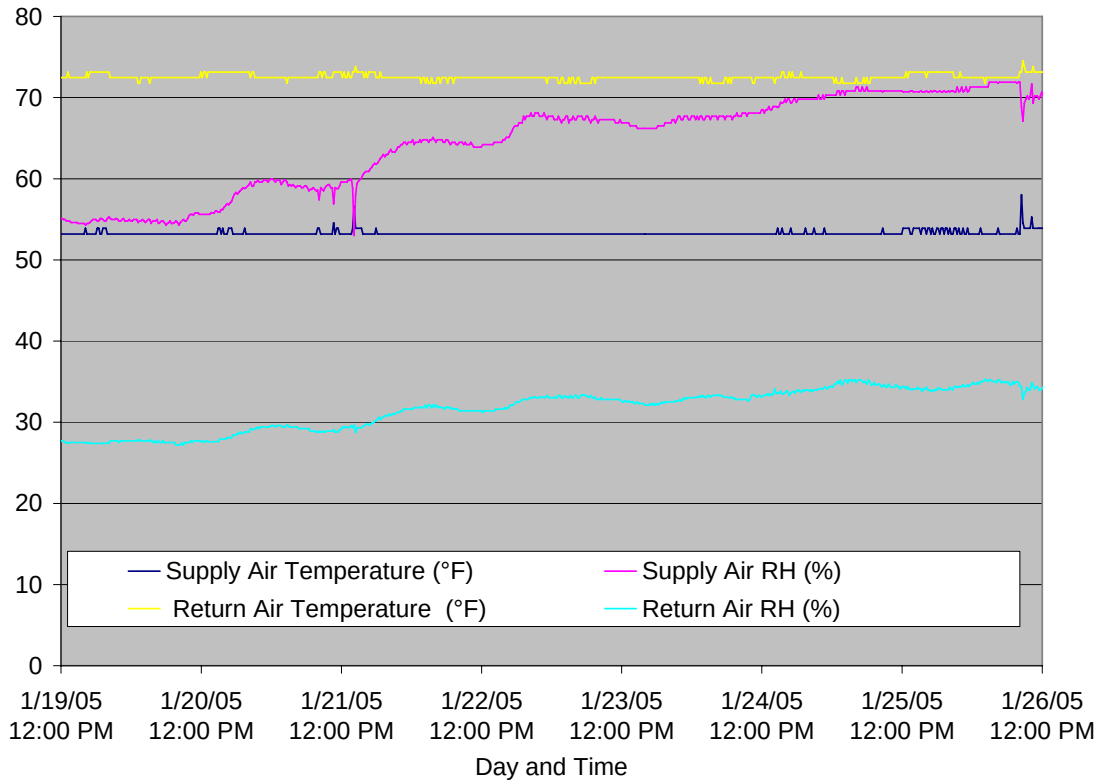

Figure 12 AHU-4 Air Temperature and Humidity

4.3 Data Center Air Conditions

Air temperature and relative humidity was recorded at two locations in the data center, as shown in Figure 13. The average space conditions were 67°F and 35% RH at one location and 69°F and 38% RH at the other location.

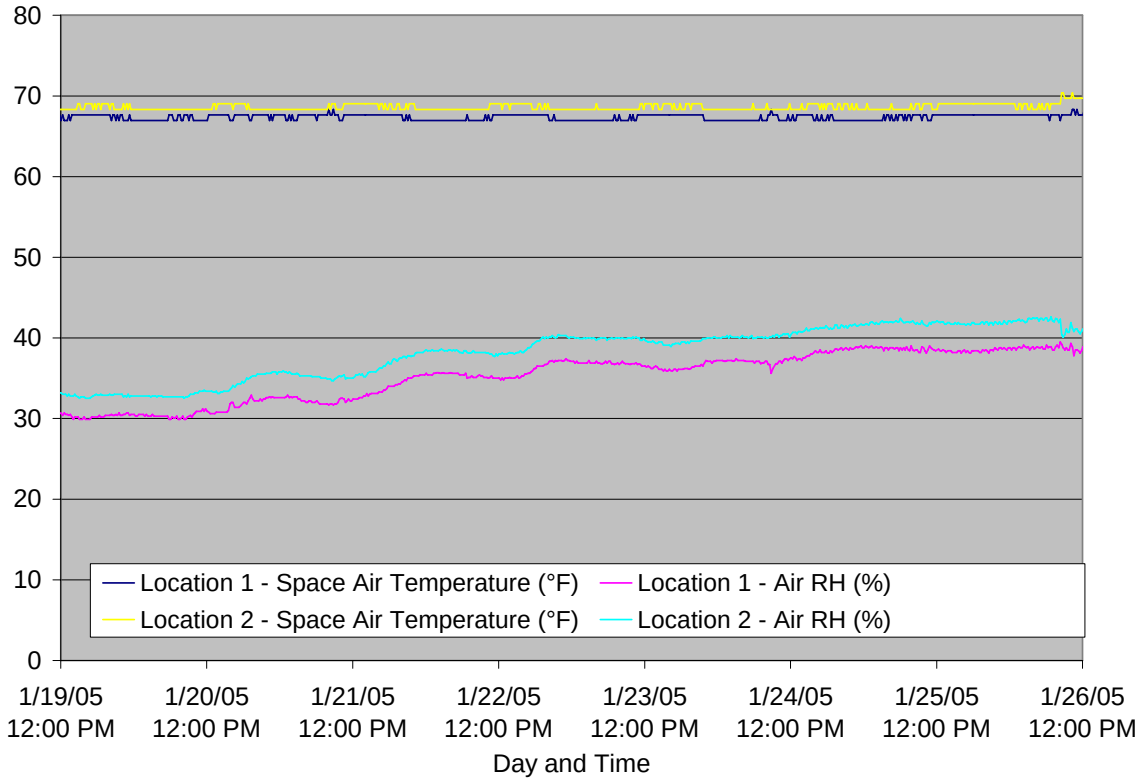

Figure 13 Space Air Temperature and Humidity at Two Locations

Generator coolant inlet and outlet temperatures were monitored in the study, as shown in Figure 14 and Figure 15, respectively. The temperatures averaged 65°F and 64°F, respectively.

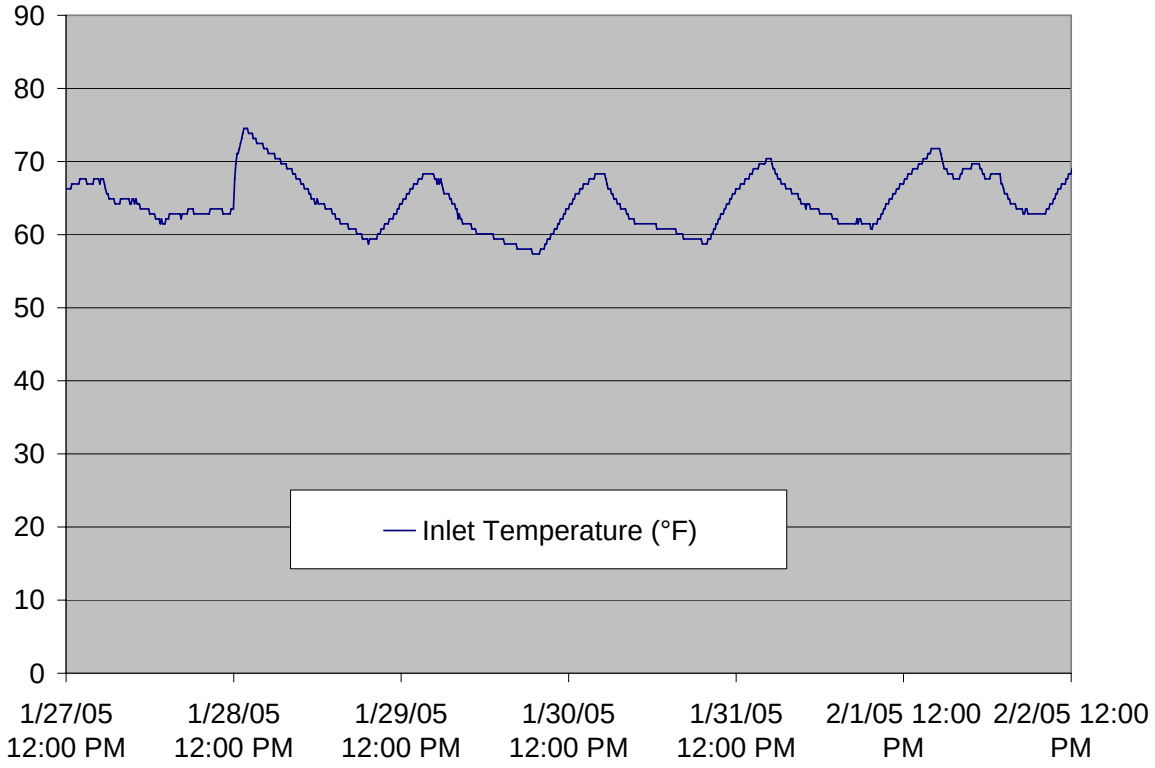

Figure 14 Generator Coolant Inlet Temperature

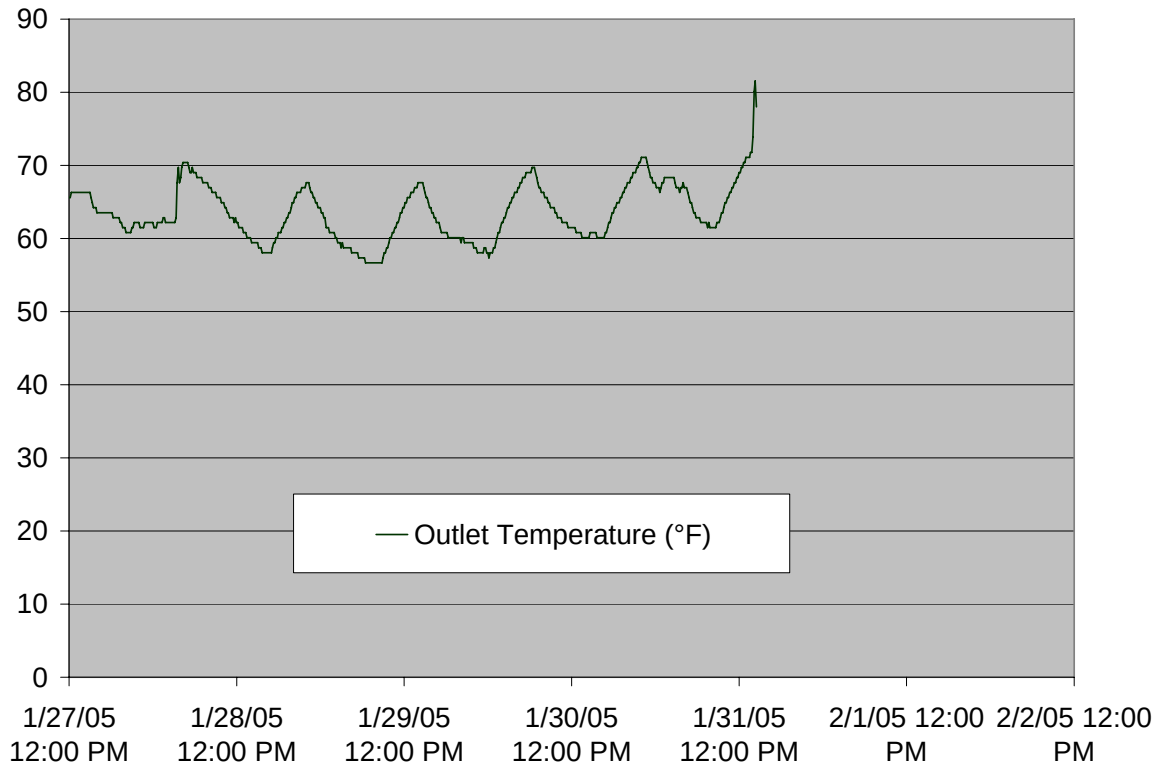

Figure 15 Generator Coolant Outlet Temperature

5 Observations and Recommendations

The data center had a total power density of approximately 94 W/ft^2, while the installed computer power density was 57 W/ft^2. The highest density of energy usage is from the racks. Therefore, it is important to reduce the rack power usage.

General recommendations for improving overall data center energy efficiency include improving the lighting control; design, operation, and control of mechanical systems serving the data center in actual operation. This includes the air-handling equipment, the chilled water system, configuration and control of airflow management in the data center.

The second highest power intensity was the air handling units serving the data center (21 W/ft^2). Several measures may be considered to reduce the power intensity of the AHU systems:

- Optimize air balance including reducing fan-wheel speeds as necessary

- Review proper control algorithms to ensure proper controls for the VSDs and AHUs. In particular, the supply air temperatures were generally lower than needed at the IT

equipment. Improved air management would allow optimization with large chiller plant energy savings, and perhaps fan energy savings as well.

- Review operation of AHU-3. This unit had a relatively lower return air temperature which might indicate a possible “short-circuiting” of supply air to return air

- Optimize air management by reviewing placement of supply and return registers

The following includes potential measures to improve the energy efficiency of the data center.

- Consider raising chilled water supply water temperature. This can be done automatically using a worst-zone reset strategy to operate at the highest chilled water temperature that would meet the load. If chilled water temperature can be raised, the chiller plant would be more efficient and it would reduce incidental dehumidification, thus reducing the need for humidification as well.

- Optimize the operation of both cooling tower fans – both running together may produce the same heat rejection with much less energy than operating one single fan. In addition, consider resetting the cooling tower water temperature set points according to chiller manufacturer recommendations and/or outside air wet-bulb temperature plus an approach temperature. For any given load and wet-bulb temperature, there is an optimal combination of flow and temperature that can minimize overall chiller plant energy use. Software is available from manufacturers or consultants to optimize this algorithm.

- Install VSDs on the CHW primary pumps and the condenser water pumps and control per a chiller plant optimization routine. Check the operation of the generator jacket heaters. At least one heater appears to operate much more than necessary. The controls should be adjusted or replaced as needed.

- Check to see if the double-conversion UPS units can be operated in a bypass mode without sacrificing protection from power interruption.

6 Acknowledgements

This report on data center energy benchmarking was finalized based upon the field data collection performed by EYP Mission Critical Facilities and Landsberg Engineering and the draft report produced in the course of performing work subcontracted for the Lawrence Berkeley National Laboratory (LBNL).

The project was funded by the California Energy Commission’s Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

7 Appendix A: Data Facility Definitions and Metrics

The following definitions and metrics are used to characterize data centers:

Air Flow Density	The air flow (cfm) in a given area (sf).
Air Handler Efficiency 1	The air flow (cfm) per power used (kW) by the CRAC unit fan.
Air Handler Efficiency 2	The power used (kW), per ton of cooling achieved by the air-handling unit.
Chiller Efficiency	The power used (kW), per ton of cooling produced by the chiller.
Computer Load Density – Rack Footprint	Measured Data Center Server Load in watts (W) divided by the total area that the racks occupy, or the "rack footprint".
Computer Load Density per Rack	Ratio of actual measured Data Center Server Load in watts (W) per rack. This is the average density per rack.
Computer/Server Load Measured Energy Density	Ratio of actual measured Data Center Server Load in watts (W) to the square foot area (sf) of Data Center Floor. Includes vacant space in floor area.
Computer/Server Load Projected Energy Density	Ratio of forecasted Data Center Server Load in watts (W) to the square foot area (sf) of the Data Center Floor if the Data Center Floor were fully occupied. The Data Center Server Load is inflated by the percentage of currently occupied space.
Cooling Load – Tons	A unit used to measure the amount of cooling being done. One ton of cooling is equal to 12,000 British Thermal Units (BTUs) per hour.
Data Center Cooling	Electrical power devoted to cooling equipment for the Data Center Floor space.
Data Center Server/Computer Load	Electrical power devoted to equipment on the Data Center Floor. Typically the power measured upstream of power distribution units or panels. Includes servers,

	switches, routers, storage equipment, monitors and other equipment.
Data Center Facility	A facility that contains both central communications and equipment, and data storage and processing equipment (servers) associated with a concentration of data cables. Can be used interchangeably with Server Farm Facility.
Data Center Floor/Space	Total footprint area of controlled access space devoted to company/customer equipment. Includes aisle ways, caged space, cooling units electrical panels, fire suppression equipment and other support equipment. Per the Uptime Institute Definitions, this gross floor space is what is typically used by facility engineers in calculating a computer load density (W/sf).
Data Center Occupancy	This is based on a qualitative estimate of how physically loaded the data centers are.
Server Farm Facility	A facility that contains both central communications and equipment, and data storage and processing equipment (servers) associated with a concentration of data cables. Can be used interchangeably with Data Center Facility. Also defined as a common physical space on the Data Center Floor where server equipment is located (i.e. server farm).

8 Appendix B: Facility Diagrams

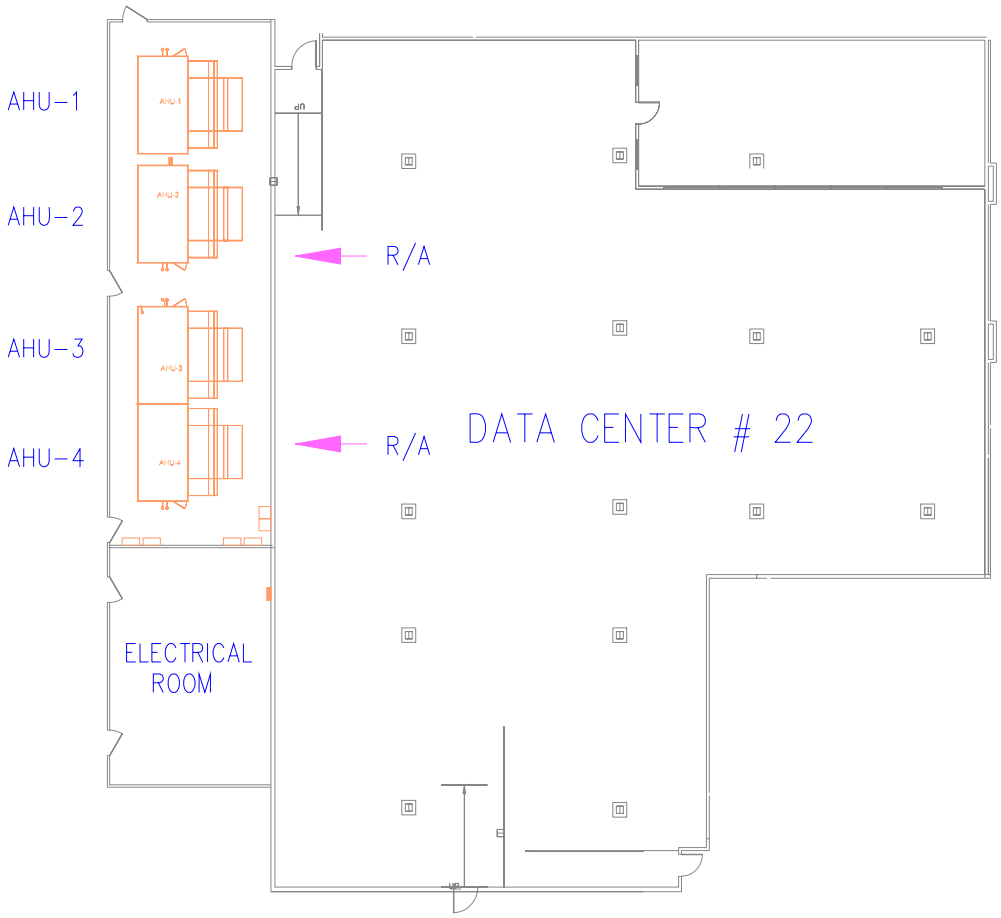

Figure 16. Data Center Plan Layout

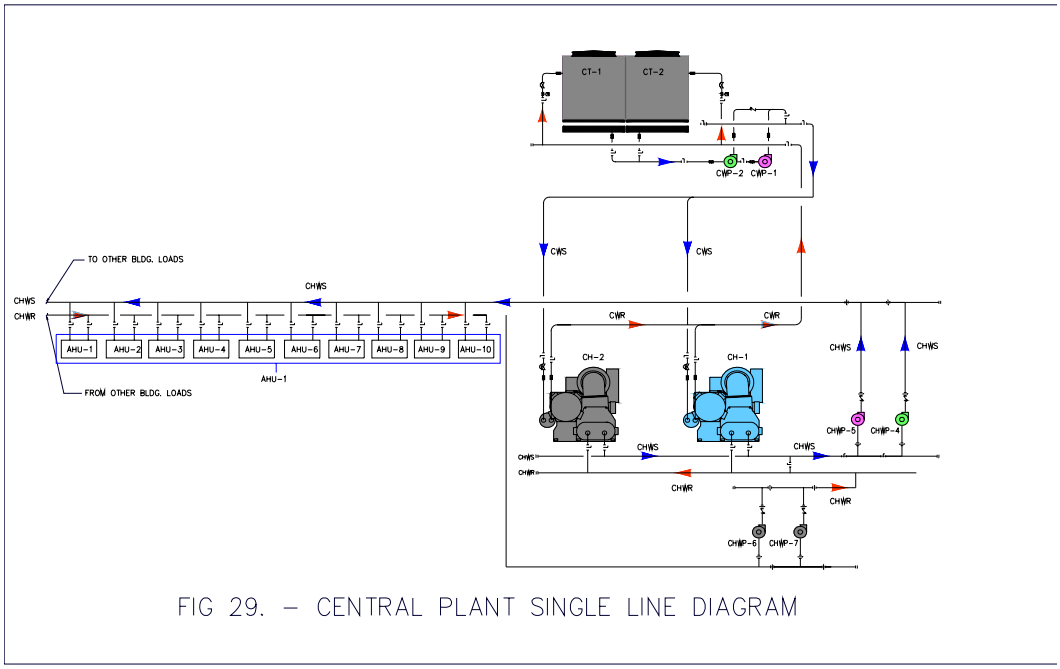

Figure 17. Chilled Water System

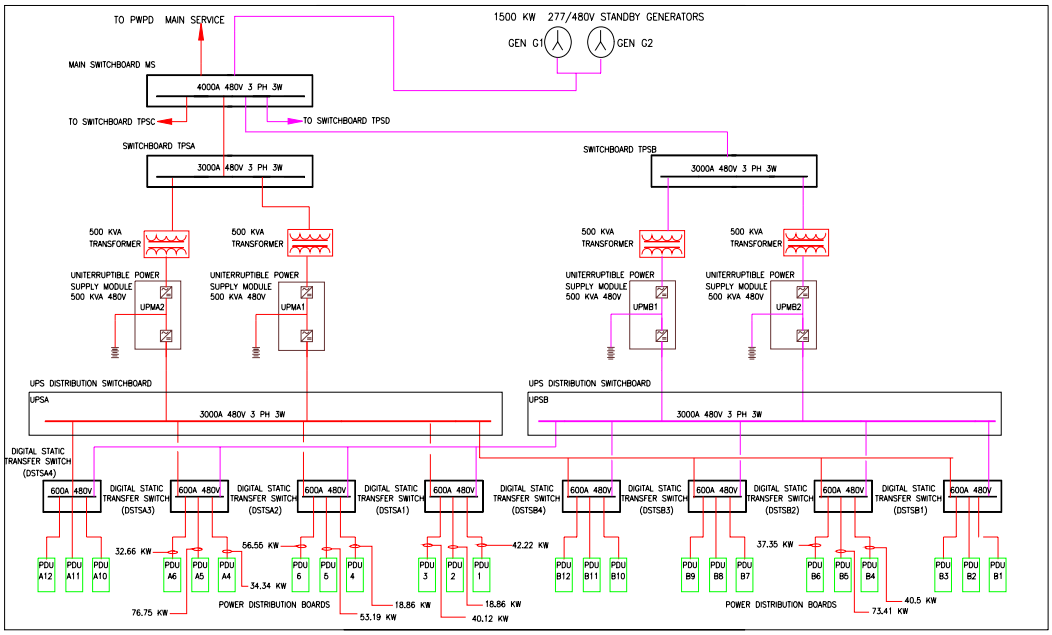

FIGURE 30 – ELECTRICAL SYSTEM SCHEMATIC

Figure 18. Electrical System Schematic

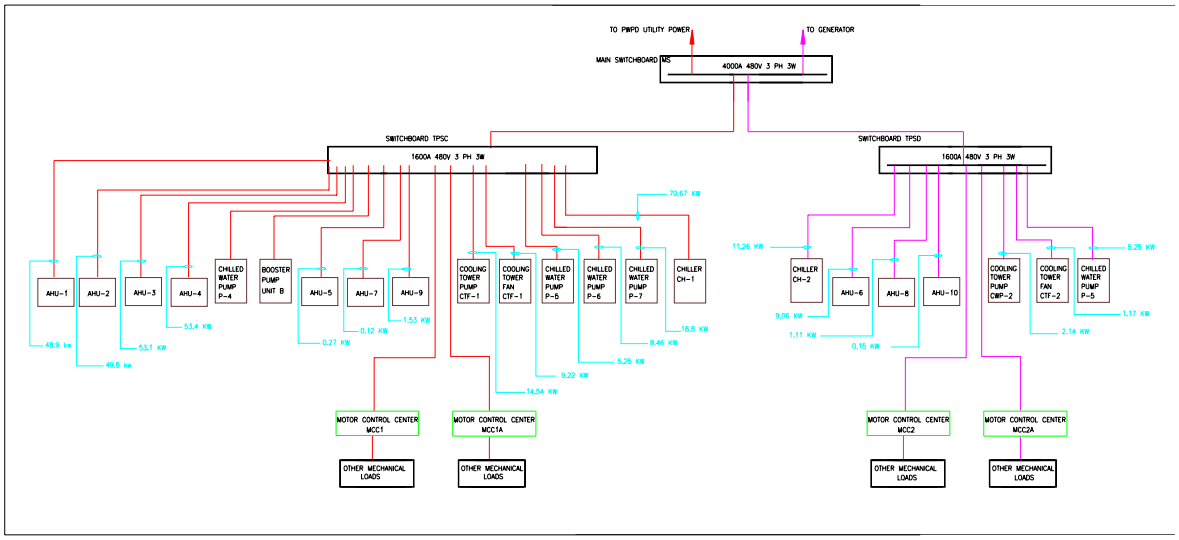

FIGURE 31 – ELECTRICAL SYSTEM SCHEMATIC TO MECHANICAL LOADS

Figure 19. Mechanical Loads and Electrical System

Data Center Energy Benchmarking:

Part 4 - Case Study on a Computer-testing Center (No. 21)

Final Report

August 2007

Final Report Prepared by
Tengfang Xu and Steve Greenberg
Lawrence Berkeley National Laboratory (LBNL)
Berkeley CA 94720

Draft provided by

EYP Mission Critical Facilities
Los Angeles, CA 90064
and
Landsberg Engineering, P.C.
Clifton Park, NY 12065

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Table of Contents

1 EXECUTIVE SUMMARY 1

2 REVIEW OF SITE CHARACTERISTICS 2

2.1 ELECTRICAL EQUIPMENT AND BACKUP POWER SYSTEM 3
2.2 MECHANICAL SYSTEM 4

3 ELECTRIC POWER CONSUMPTION CHARACTERISTICS 6

3.1 POWER SYSTEM 8
3.2 CHILLER SYSTEM 9
3.3 PUMPING SYSTEM 11
3.4 COOLING TOWERS AND PUMPS 11
3.5 STANDBY GENERATORS, UPS AND DATA CENTER PDU'S 12
3.6 COMPUTER ROOM AIR HANDLERS 12
3.7 DATA CENTER POWER SUPPLY 12

4 SYSTEM OPERATION 12

4.1 AMBIENT AIR TEMPERATURE AND HUMIDITY 13
4.2 CHILLED WATER SUPPLY AND RETURN TEMPERATURES 13
4.3 AIR HANDLING UNIT SUPPLY AND RETURN TEMPERATURES AND RELATIVE HUMIDITY 14
4.4 DATA CENTER SPACE AIR TEMPERATURE AND RELATIVE HUMIDITY 16

5 OBSERVATIONS AND RECOMMENDATIONS 17

5.1 LIGHTING 18
5.2 AIRFLOW OPTIMIZATION 18
5.3 AIR TEMPERATURE AND RELATIVE HUMIDITY 20
5.4 CHILLED WATER SYSTEM 20

6 ACKNOWLEDGEMENTS 21

7 APPENDIX A: DATA FACILITY DEFINITIONS AND METRICS 22

8 APPENDIX B: FACILITY DIAGRAMS 24

1 Executive Summary

The data center in this study had a total floor area of 8,580 square feet (ft^2) with one-foot raised-floors. It was a rack lab with 440 racks, and was located in a 208,240 ft^2 multi-story office building in San Jose, California. Since the data center was used only for testing equipment, it was not configured as a critical facility in terms of electrical and cooling supply. It did not have a dedicated chiller system but served by the main building chiller plant and make-up air system. Additionally, it was served by a single electrical supply with no provision for backup power. The data center operated on a 24 hour per day, year-round cycle, and users had all hour full access to the data center facility.

The study found that data center computer load accounted for 23% of the overall building electrical load, while the total power consumption attributable to the data center including allocated cooling load and lighting was 30% of the total facility load. The density of installed computer loads (rack load) in the data center was 63 W/ft^2. Power consumption density for all data center allocated load (including cooling and lighting) was 84 W/ft^2, approximately 12 times the average overall power density in rest of the building (non-data center portion).

For the data center, 75% of the overall electric power was the rack critical loads, 11% of the power was consumed by chillers, 9% by CRAH units, 1% by lighting system, and about 4% of the power was consumed by pumps. The ratio of HVAC to IT power demand in the data center in this study was approximately 0.32.

General recommendations for improving overall data center energy efficiency include improving the lighting control, airflow optimization, and control of mechanical systems serving the data center in actual operation. This includes chilled water system, airflow management and control in data centers. Additional specific recommendations or considerations to improve energy efficiency are provided in this report.

2 Review of Site Characteristics

The data center (DC #21) in this study had a total floor area of 8,580 square feet (ft^2) with one-foot raised-floors. It was a rack lab with 440 racks, and was located in a 208,240 ft^2 multi-story office building in San Jose, California. Since the data center was used only for testing equipment, it was not configured as a critical facility in terms of electrical and cooling supply. The data center did not have a dedicated chiller system but was served by the main building chiller plant and a make-up air system. Additionally it was served by only a single electrical supply with no provision for backup power in the event of a power outage. The center operated on a 24 hour per day, year-round cycle, and users had full access to the data center facility at all hours.

Electric power was supplied to the office building from the utility to a single three-phase, 21 kV primary utility service that was transformed through two 2500 kVA transformers to two 4,000A, 480V/277 V main service switchboards. The switchboards provided building electrical distribution to the data center. The data center is fed from main switchboards MSB-16.1 and MSB-16.2. Each switchboard fed two 480V - 208/120V transformers that each supply distribution panels within the data center. There was no standby generator, uninterruptible power supply (UPS), static switch, or electrical reliability component associated with typical data centers. Communication and power wiring was installed at overhead ceiling. Fire sprinklers were provided under the raised floor and at the ceiling.

Cooling for the data center facility was served by the building's main chiller plant and the make-up air system. The building chilled water system included five air-cooled chillers with chilled water pumps. The cooling system inside the data center included eight Computer Room Air Handling (CRAH) units. Seven CRAH units were in operation during the study and received chilled water from the chiller plant serving the whole building. The CRAH's used a down-flow type supply for air distribution, and return to the top without ducts. Figure 1 shows a view of a typical rack lineup in the data center.

The data for this benchmarking exercise was collected during a one week monitoring period in December 2004. The data was collected using the following measurement instruments: HOBO loggers, Elite loggers, Fluke 41B power meter and the installed Automated Logic Corporation (ALC) direct digital control (DDC) system.

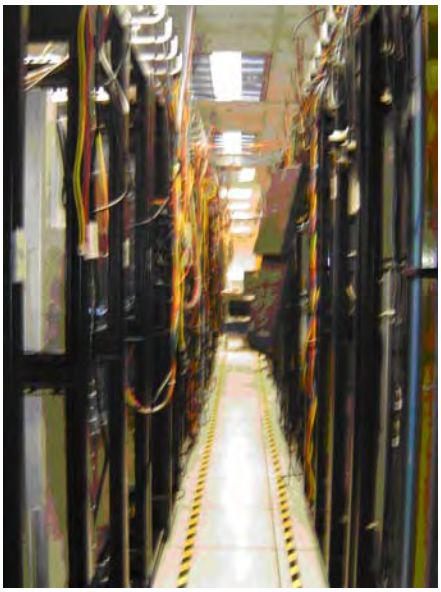

Figure 1 Typical computer rack lineup in the data center

2.1 Electrical Equipment and Backup Power System

Electrical power to Data Center #21 was served by the building electrical distribution system, which received power from Pacific Gas and Electric Company via a three-phase, three-wire, 21-kV main service. The 21 kV utility power was routed through utility metering to a 21 kV switchgear lineup and was transformed to 480/277 V power via two 2500 kVA transformers, which fed two 4,000-A, 480/277 V, 3-phase, 4-wire main service switchboards (MSB 16.1 and 16.2), as shown in the Electrical System One-line Diagram in Appendix B.

As part of the study, the overall building power consumption was monitored over an one-week period in December 2004. During that period, the average amount of power consumed in the building during peak usage (workday) periods was 2,380 kW.

The main switchboards, MSB 16.1 and MSB 16.2, supply power to the data center computer racks through four 225 kVA transformers, two supplied from each switchboard. Each transformer fed an associated 800-Ampere 120/208-V distribution panel, and each of which fed multiple sub-panels within the data center.

The chillers, pumps, CRAH units, lighting and miscellaneous loads are fed from MSB 16.1 via MCC's and other panels.

2.2 Mechanical System

2.2.1 Chiller

The data center, along with the remainder of the building, was designed to be cooled by two 500-ton York Millennium Model Number YTJ1C3E2 water-cooled centrifugal chillers. A third chiller, CH-3, was being installed at the time of the survey. The chillers were arranged in a parallel configuration. CH-1 and CH-2 were operating at the time of the survey. The chilling system flow diagram is shown in Appendix B.

These chillers were installed on the roof of the building along with the other chiller plant components including cooling towers and pumps. The chillers were single-stage centrifugal compressor with modulating capacity of design air conditioning loads (10% to 100%). The chillers were set to maintain a leaving water temperature of 44 °F and were controlled via a stand-alone microprocessor based control to provide optimal chiller efficiency based on a variety of factors including condensing water temperature, evaporator temperature, chilled water set point, motor speed, and pre-rotation vane position.

The chilled water temperature reset was internal through programming of the stand-alone chiller controller. The chiller design parameters were based on 800 gpm of evaporator flow (1.6 gpm/ton) and 1,500 gpm of condenser flow (3 gpm/ton) with an entering (return) chilled water temperature of 59°F and a leaving (supply) water temperature of 44°F. The chilled water average supply and return temperatures measured during the monitoring period were 44.3°F and 52.3 °F respectively, equating to an average temperature differential of 8°F. The delta-T corresponded to 267 cooling tons per chiller, or a total of 533 $tons_R$ for the two chillers. This represented approximately 53% of the design capacity for the two chillers in operation. The chillers were provided with variable speed drives and thus operated efficiently at low loads. The design chiller power demand was 0.42 kW per cooling ton, while the actual power demand during the monitoring period was 0.44 kW per cooling ton.

2.2.2 Chilled Water Pumps

Primary chilled water was circulated by three parallel in-line Bell & Gossett pumps, each with a motor capacity of 15 HP. The pumps were constant speed and identified as CHP16-1, CHP16-2 and CHP16-3. One pump per chiller operated continuously when chillers were in operation, and one pump was provided as a reserve unit. An additional primary pump CHP16-7 was being added at the time of the survey in conjunction with installation of the third chiller. The primary pumps were controlled through the facility Automated Logic Control (ALC) and start/stop with the chillers.

The building chilled water was circulated by the secondary chilled water pumps that consisted of two parallel, in-line, 50-HP, Bell and Gossett centrifugal pumps identified as CHP16-4 and CHP16-5. Both of the secondary chilled water pumps were fitted with VSDs. An additional

50-hp secondary pump, CH 16-6 was being added at the time of survey. The secondary chilled water pumps were controlled in stages based on building cooling demand and controlled based on chilled water return temperature as sensed by the facility Automated Logic Control.

2.2.3 Cooling tower

There were four, Baltimore Air Coil model 3485 series V, cooling towers; CT16-1, CT-16-2, CT16-3 and CT16-4. Each tower was fitted with a 30-hp fan operating on a variable speed drive. The speed of the fans was controlled on cooling tower leaving water temperature (condenser water supply temperature). There was no temperature reset strategy on the cooling towers. The cooling towers were induced draft towers with a design-wet bulb temperature of 68°F, and a leaving and entering water temperature of 73°F and 82°F respectively. Two cooling towers were operational at the time of the study

The cooling towers were served by three, 40-HP, parallel, base-mounted, Bell and Gossett centrifugal pumps. The pumps were identified as CTP16-1, CTP16-2 and CTP16-3. A fourth cooling tower pump, CTP16-4 was being added at the time of the survey. The cooling tower pumps were constant speed and controlled by the condenser water return temperature.

2.2.4 Computer Room Air Handling Units

Seven out of eight 30-ton Computer Room Air Handling (CRAH) units were in operating in the data center, supplying the data center with cold air from the one-foot raised floor. The CRAH's were Pomona Air Model # PW 3000 units. No reheat coils or humidifiers were included in the CRAH units. The CRAH units had 4-inch throwaway air filters located at the top of the unit, rated at 85% efficiency. Each CRAH unit's internal controls were set to maintain temperature and relative humidity set-points of 70°F and 20% RH, respectively, measured at the unit's return air intake.

Based on manufacturer's specification information supplied by the Building Engineer, each air handler has a capacity of approximately 15,000 cfm (cubic feet per minute). The measured power consumption per cfm of air flow of the CRAH's was 0.62 W/cfm. Therefore, the average Air Handler Efficiency-1 was 1620 cfm/kW, based on design airflow. The actual airflow rate was not measured but was calculated as described below.

The space temperature set points were 71°F and 50% relative humidity (RH). Space temperature and RH were recorded at two different locations in the data center, the first location averaged 71°F and 25% RH while the second location averaged 79°F and 25% RH. The average supply temperature from all seven units was 53°F and the average return air temperature was 74°F. Supply relative humidity averaged 64%, while return relative humidity averaged 32%.

Based on average supply and return temperatures, cooling load attributable to the data center racks (175 tons), CRAH's, and lighting power consumption, average airflow rate was estimated to be 13,400 cfm per CRAH.

Average supply and return air temperature and relative humidity for the individual CRAH units are shown in Table 1. These were average readings taken over several days of monitoring.

Table 1 CRAH Unit Supply and Return Air Temperature and Relative Humidity

CRAH Unit	Supply Air (°F)	Supply RH (%)	Return Air (°F)	Return RH (%)
16CAH4	53	61	78	23
16CAH5	51	n/a	75	29
16CAH6	51	n/a	75	29
16CAH7	52	64	71	31
16CAH8	53	n/a	72	35
16CAH13	55	n/a	76	48
16CAH14	53	67	77	34

3 Electric Power Consumption Characteristics

Table 2 shows the end-use electricity demand of the building housing the data center in this study, based on the measurements taken during the monitoring period. The average building electrical load of 2,380 kW was recorded from building instruments. From these measurements, it was observed that 61% of the electrical load was consumed by areas of the building other than the data center and chiller plant, 14% by the chillers and chilled water plant, 23% of the load was consumed by the data center computer equipment, 3% of the load was consumed by the data center CRAH units, and less than 1% of the power was consumed by the data center lighting.

Table 2 Power demand breakdown in the data center building

Power Demand Breakdown	Average Power Consumption (kW)	% Power Consumption
Overall Building Load	2380	100%
Building – non data center & non-chilling Load	1441	60.5%
Data Center Computer Load	540	22.7%
Computer Room Air Handlers	65	2.7%
Building Chiller Plant Pumps and Cooling Towers	92	3.9%
Building Chillers	232	9.7%
Data Center Lighting	10	0.4%

Table 3 shows the power density for the square foot area served by each load. About 70% of the electrical load was consumed by other areas of the building than the data center, including total cooling systems. For the remaining 30% of power consumption, approximately 23% of the building load was consumed by the data center computer equipment and about 7% of the load was by the CRAH units and the data center allocation of chiller and chilled water pump loads. The density of installed computer loads (rack load) in the data center was 63 W/ft^2. Power consumption density for all data center load (including cooling and lighting) was 84 W/ft^2, approximately twelve times the average overall power density in rest of the building (non-data center portion).

Table 3. End-Use of Electricity of the Data Center Building

Description	Electric power demand (kW)	Share of electric energy use (%)	Floor Space (ft2)	Electric power density (W/ft2)
Overall Building Load	2380	100%	208,240	**11**
Chillers – (non Data Center chiller load)	156	6.6%	199,660	0.8
Chiller Plant Pumps & Towers – (non Data Center load)	62	2.6%	199,660	0.3
Building Other - (non DC load &non-HVAC building loads)	1441	60.5%	199,660	7
Total Non-Data Center Load	1659	69.7%	199,660	8
Data Center Computer Load	540	22.7%	8,580	**63**
Data Center CRAH Units	65	2.7%	8,580	7.6
Chillers – (Data Center load)	76	3.2%	8,580	8.9
Chiller Plant Pumps & Towers – (Data Center load)	30	1.3%	8,580	3.5
Data Center Lighting	10	0.4%	8,580	1.2
Total Data Center Load	721	30.3%	8580	**84**

The end-use breakdown for the data center's electric power demand is also shown in Table 4. For the data center, 75% of the overall electric power was the rack critical loads, 11% of the power was consumed by chillers, 9% by CRAH units, 1% by lighting system, and about 4% of the power was consumed by pumps. The ratios of HVAC to IT power demand in the data centers in this study were approximately 0.32.

Table 4. End-Use of Electricity of the Data Center

Description	Electrical Power Consumption (kW)	Percent of Total Data Center Load
Data Center Computer Load	540 kW	75%
Data Center CRAH Units	65 kW	9%
Chillers – (Data Center load)	76 kW	11%
Chiller Plant Pumps & Towers – (Data Center load)	30 kW	4%
Data Center Lighting	10 kW	1%
Total Data Center Load	721 kW	100%

3.1 Power System

Electrical power for the data center was supplied from Main Switchboards MSB-16.1 & 16.2 (Appendix B includes the electrical single-line diagram). Each switchboard fed two 225 kVA, 480V:120/208V Cutler Hammer, dry-type, transformers (model Number N48M28T22A). Each transformer fed an 800A, 120/208V, 3-phase, 4-wire distribution panel. The 800A distribution panels each fed multiple sub-panels for computer racks within the data center.

The chillers, pumps, CRAH units, lighting and miscellaneous loads were fed from MCC's and other electrical panels.

3.2 Chiller System

Electric power demand was monitored for the two parallel operating chillers within a two-week period. CH-1 was operating at the time of survey, however CH-2 was off for the first six days of the monitoring period and then it was placed in operation. Chillers CH-1 and CH-2 had relatively equal loading, when they were both in operation, with average power consumption of 113 kW and 119 kW, respectively. The total chiller power consumption for the period averaged 232 kW.

Total average chiller power demand for the period was 83 kW, which represented the chilling load of the whole building. CH-1 and CH-2 provided the majority of the cooling load with average power demand of 38 and 32 kW, respectively. The chiller CH-5 operated with a power demand of 13 kW.

Figure 2 shows the total chiller power demand and outside air temperatures within the two-week period. The chiller power consumption has been divided by 10 for the plot in order to provide better illustration in scaling between ambient temperature and power. As the outside air temperature changed, the actual chiller power demand also changed accordingly.

In addition, Figure 3 further shows the general correlation of total chiller power consumption and outside air temperature. While showing somewhat correlation between ambient temperature and chiller power consumption, both figures indicate that the chiller load was not only affected by the ambient temperature but influenced by other factors.

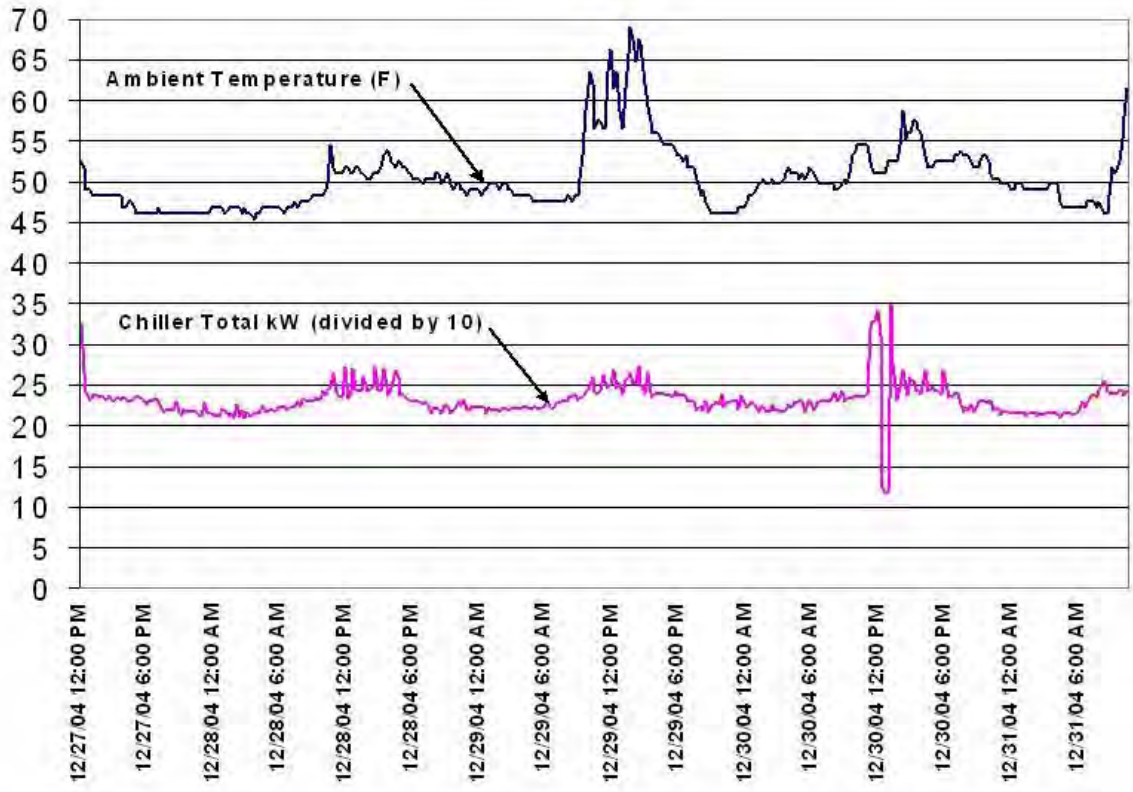

Figure 2 Chiller cooling power and outdoor air temperature

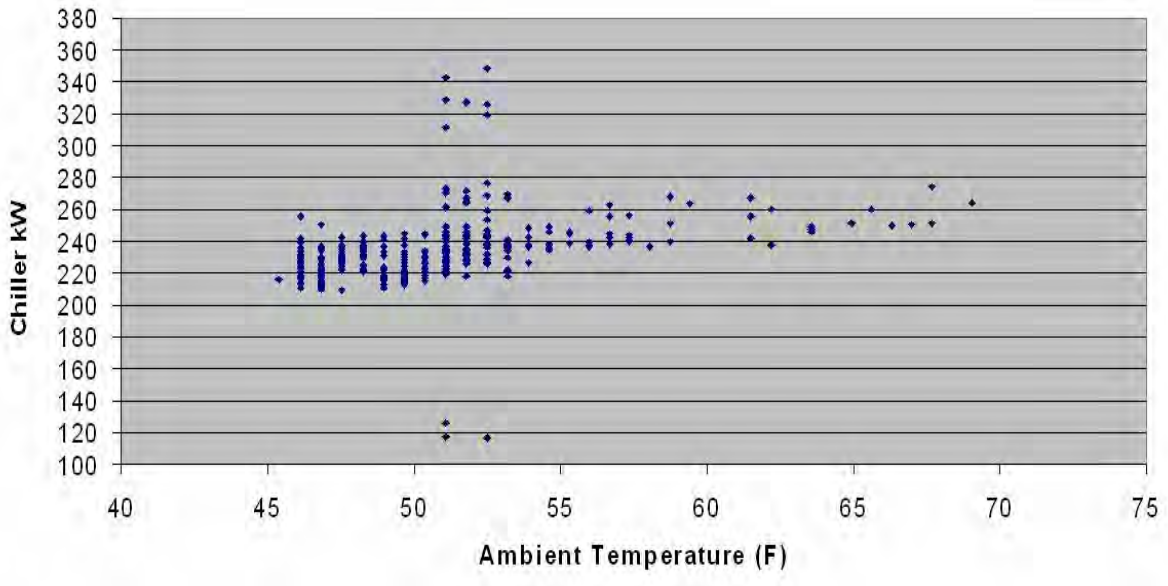

Figure 3 Correlation of chiller cooling power and outdoor air temperature

3.3 Pumping System

The central plant primary chilled water pumping system consisted of three 15-hp primary chilled water pumps: CHP16-1, CHP16-2 and CHP16-3. One pump was designated as a standby unit. An additional primary pump (CHP16-7) was being added at the time of the survey to serve the third chiller that was being installed.

The secondary chilled water pumping system consisted of two 50-hp centrifugal pumps: CHP16-4 and CHP16-5. A third secondary pump (CHP16-6) was being added at the time of survey. The secondary chilled water pumps are fitted with variable speed drives.

3.4 Cooling Towers and Pumps

There were four induced-draft type cooling towers with 30-hp fan motors: CT16-1, CT16-2 CT16-3 and CT16-4. The tower fans were provided with variable speed drives controlled by tower leaving water temperature (condenser water supply temperature). Two cooling towers were in operation at the time of the study.

The cooling towers were served by three 40-hp cooling tower pumps: CTP16-1, CTP16-2, and CTP16-3. A fourth cooling tower pump, CTP16-4 was being added at the time of the survey. The cooling tower pumps were constant speed and stage controlled by the building condenser

water temperature demand through the facility ALC. The average central plant pump and tower fan power demand for the monitoring period was 92 kW.

3.5 Standby Generators, UPS and Data Center PDU's

There was no standby generator or UPS associated with this data center. Likewise, there was no Power Distribution Unit typically observed in a data center.

3.6 Computer Room Air Handlers

The data center was served by three 20-ton CRAH units in operation during the survey period. Each unit had two 7.5-hp constant speed fan motors. Based on data taken at the time of the survey, CRAH power consumption for unit CAH 1.5 was 0.84 watts per unit of airflow rate at one ft^3/min. The variation of minimum to maximum power draw was under 4% during the period.

The data center was served by seven 30-ton CRAH units. Each unit has two 7.5-hp constant speed fans and delivers approximately 15,000 cfm based on manufacturer's specification (13,400 cfm by calculation). Based upon data taken at the time of the survey, AHU power consumption was 0.62 W/cfm using the rated airflow, and 0.69 W/cfm using the calculated airflow rate.

3.7 Data Center Power Supply

The power to the racks was provided by a number of panel boards within the data center supplied from four distribution panels via four distribution transformers. The distribution transformers were rated at 225 kVA each (or 220 kW at an assumed data center power factor of 92%).

The total power consumption of the 440 data center racks, including transformer losses was 520 kW. The average power per rack was 1.2 kW. One rack was monitored over a week period. The average power consumption of the rack being monitored was 0.925 kW. For the selected rack, the minimum value of power consumption was 99% of the maximum value, indicating that the rack power was essentially constant.

4 System Operation

During the two-week monitoring period, the following HVAC equipment was operating:

- Two of three primary chilled water pumps
- One of two secondary chilled water pumps
- Chillers CH-1 and CH-2
- Two of three cooling tower pumps

- Two of three cooling towers
- All seven data center air handling units

4.1 Ambient Air Temperature and humidity

Figure 4 shows the ambient air temperature and relative humidity during the monitoring period. The average ambient temperature was 55°F, with highs of 77°F and lows of 36°F. The average relative humidity was 64% RH.

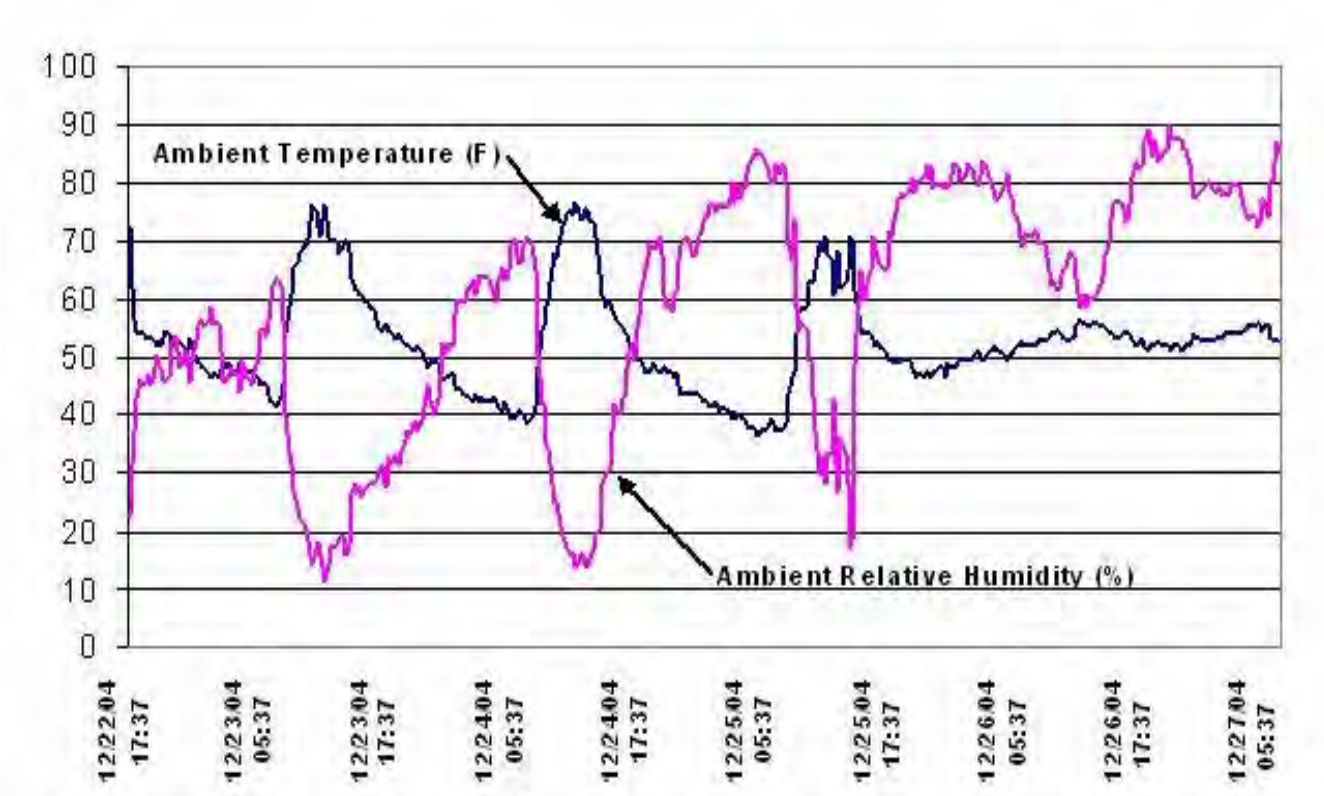

Figure 4 Ambient air temperature and humidity

4.2 Chilled Water Supply and Return Temperatures

The chilled water supply and return temperatures were monitored in the study. The data center chilled water temperatures were shown in Figure 5 along with the corresponding ambient temperature.

The spikes of high supply and return chilled water temperatures between 12/14/04 and 12/15/04 as shown in Figure 5 were the result of chiller change-over that took place during the

hours. The change-over process included a time delay as the chiller went through its self test start-up sequence.

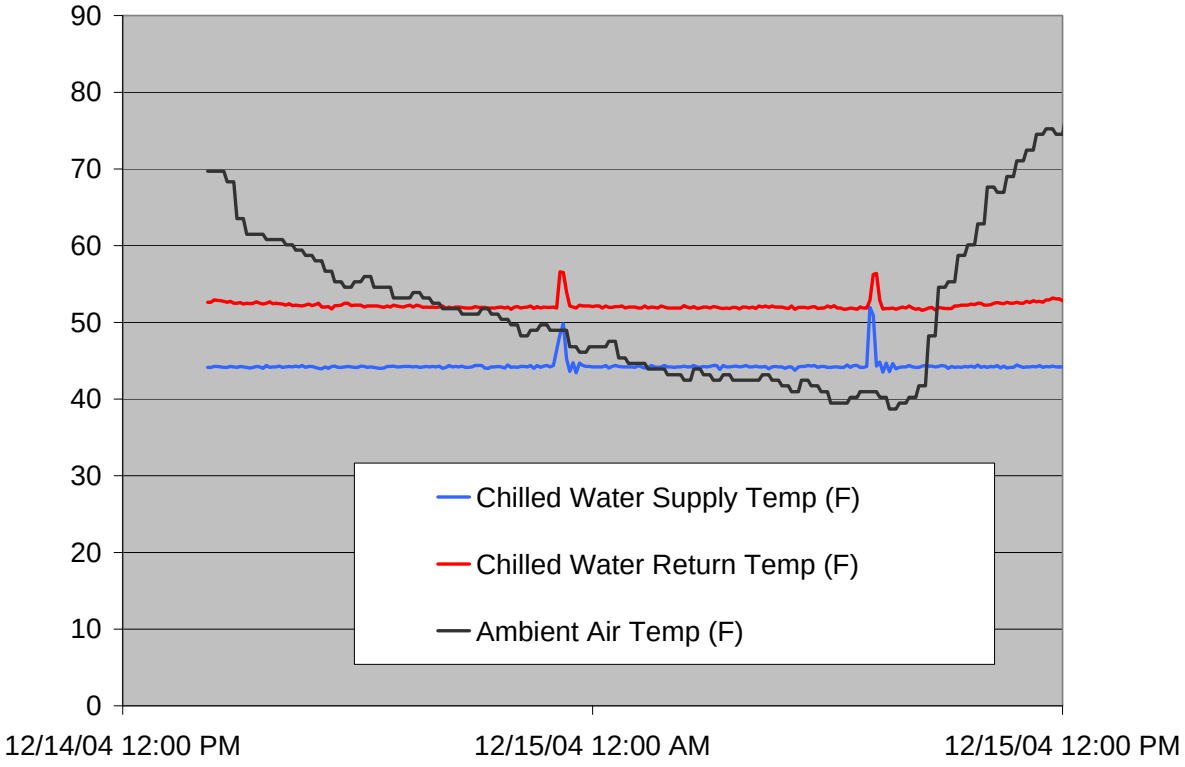

Figure 5 Chilled water and ambient air temperatures

4.3 Air Handling Unit Supply and Return Temperatures and Relative Humidity

Air temperatures and relative humidity for the supply and return airflows were monitored for one week during the study. The average temperatures are shown in Table 1.

Figure 6 shows air temperatures and relative humidity trending for 16CAH-5. The supply air temperature averaged 51°F, the return air temperature averaged 75°F, and the return air relative humidity averaged 28%. Similarly, Figure 7 shows air temperatures and relative humidity trending for 16CAH-6. The supply air temperature averaged 51°F, the return air temperature averaged 75°F, and the return air relative humidity averaged 29%.

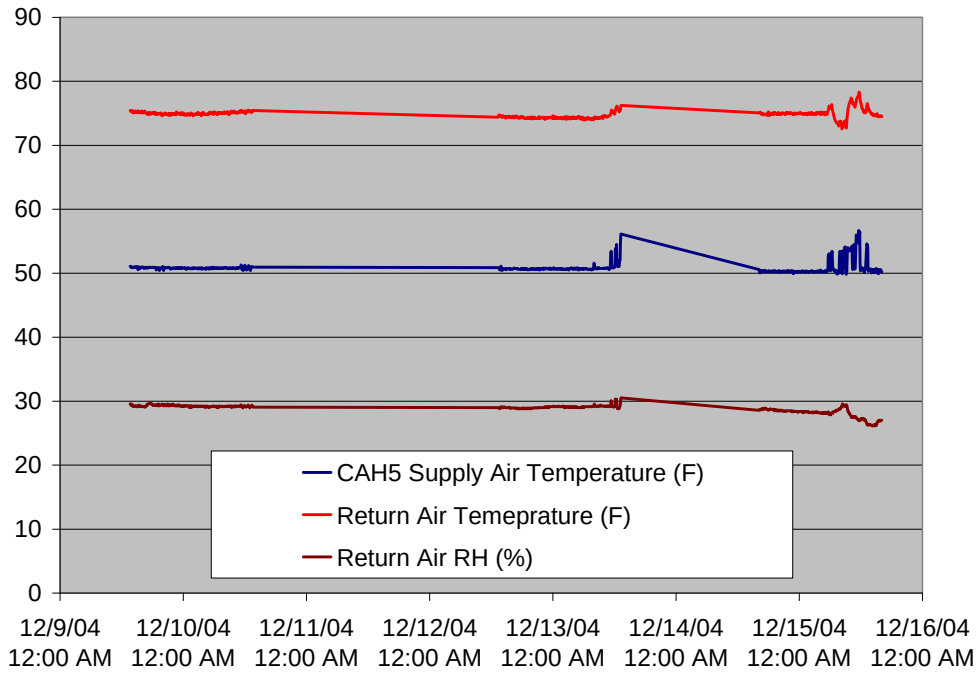

Figure 6 CAHU16-5 Supply, Return Temperatures and Relative Humidity

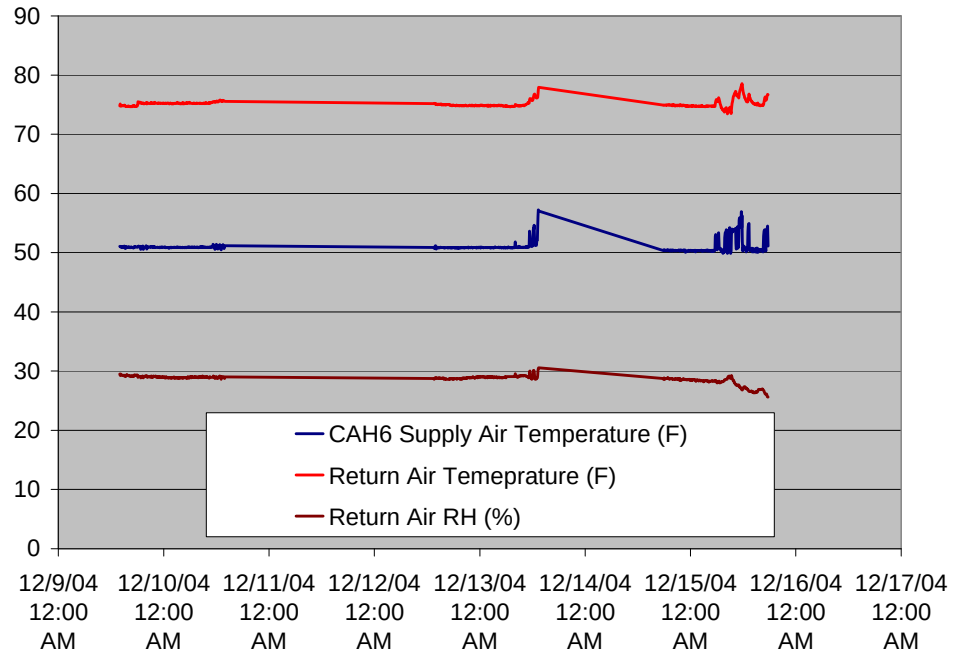

Figure 7 CAHU16-6 Supply, Return Temperatures and Relative Humidity

In addition, Figure 8 shows the temperature and relative humidity trending for 16CAH-8 The supply air temperature averaged 55°F, the return air temperature averaged 73°F, and the return air relative humidity averaged 36%.

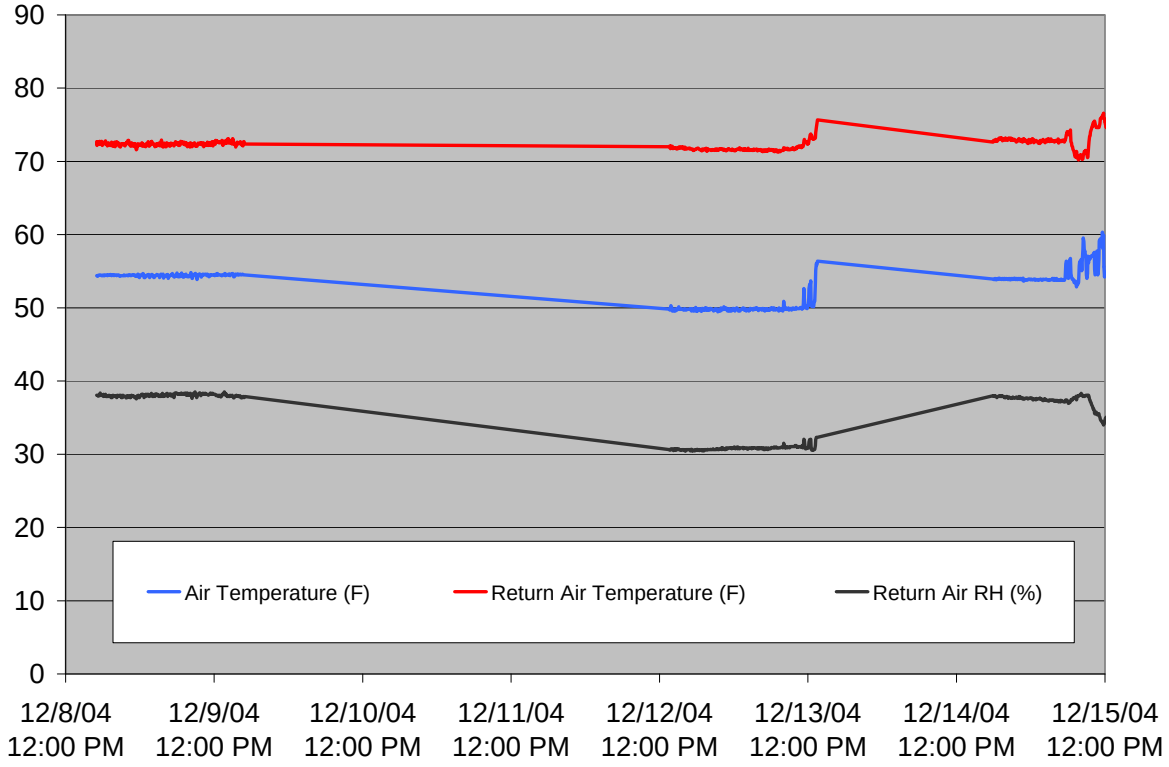

Figure 8 CAHU16-8 Supply, Return Temperatures and Relative Humidity

4.4 Data Center Space Air Temperature and Relative Humidity

The space temperature and relative humidity were monitored and shown in Figure 10. At one location the average space temperature and relative humidity was 71°F and 25% RH, respectively. At another location, the air temperature was 79°F with relative humidity of the 25% RH. The overall average space temperature was 75°F, and average relative humidity was 25%.

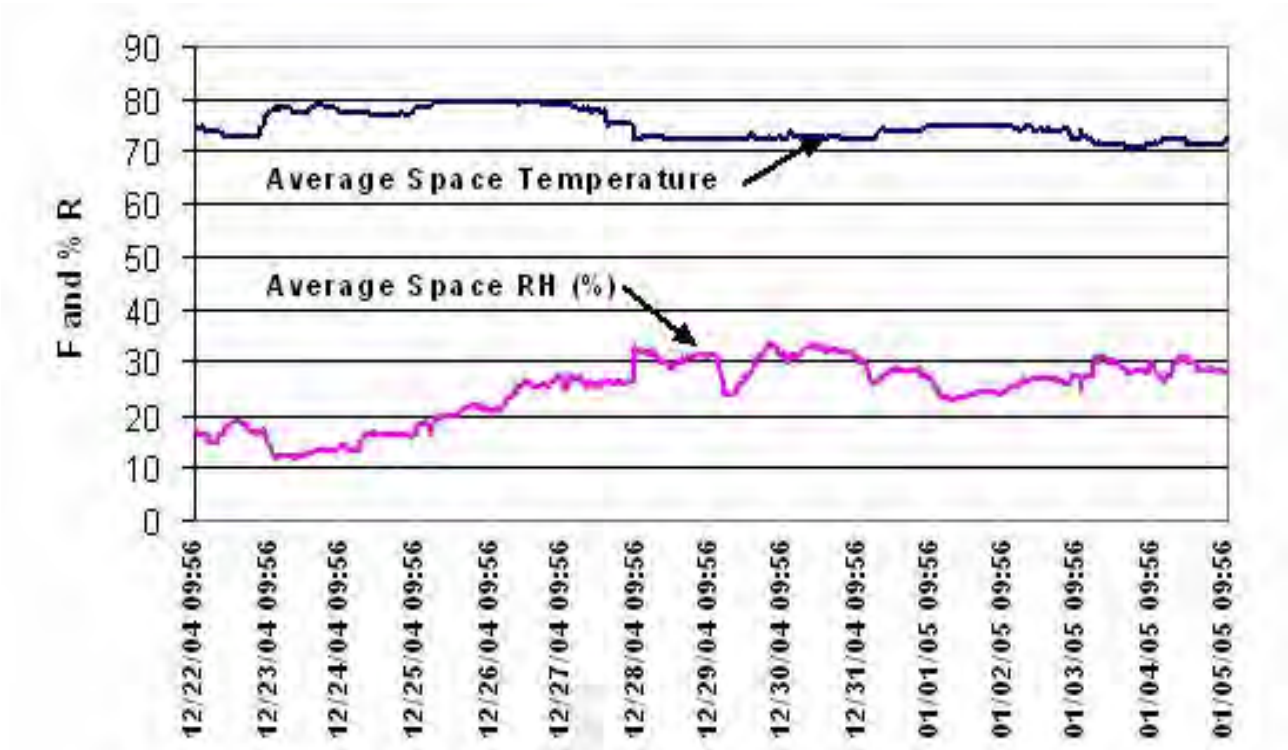

Figure 9 Data center air temperature and humidity

5 Observations and Recommendations

This data center was a rack test lab and was similar to DC 20. Both data centers didn't have electrical losses associated with UPS, standby generators, and redundant on-line equipment. DC 20 had a total power density of approximately 88 W/ft^2, while DC 21 had a total power density of approximately 84 W/ft^2. The IT power density was 61 W/ft^2 and 63 W/ft^2 for DC 20 and DC 21, respectively. watts/sf for DC 21. Corresponding to a slightly higher computer (IT) power density in DC 21, the overall power density for DC 21 was in fact lower than the overall power density in DC 20. This was in part due to 1) higher-efficiency water cooled chiller system that resulted in lower power density (i.e., 9 W/ft^2 in DC 21 compared with 14 W/ft^2 in DC 20); and 2) lower data center lighting power density in DC 21. In addition, The total power consumption per ton of heat load for the chillers and the central plant for DC 20 was approximately 22% greater than for DC 21 (0.73 kW/ton versus 0.6 kW/Ton). Interestingly, the power consumption per ton of heat load for the CRAH units of DC 20 (0.48 kW/ton) was 30% higher than that for the CRAH units of DC 21 (0.37 kW/ton).

The density of installed computer loads (rack load) in the data center studied was 61 W/ft^2. The building and its data center cooling system was provided with various energy optimizing systems that included the following:

- Varying chilled water flow rate through variable speed drives on the primary pumps.
- No energy losses due to nonexistence of UPS or standby generators.
- Minimized under-floor obstruction that affects the delivery efficiency of supply air.
- Elimination of dehumidification/humidification within the CRAH units.

General recommendations for improving overall data center energy efficiency include improving the lighting control, design, operation, and control of mechanical systems serving the data center in actual operation. This includes chilled water system, airflow management and control in data centers. The following additional techniques should result in significant improvements in energy efficiency, effective operation, or both.

5.1 Lighting

The measured lighting load in the Data center was 10 kW with an intensity of 1.2W/ft^2. The lighting power can be reduced by considering the following energy control measures: install lighting zone occupancy sensors; and task lighting in appropriate areas and disable portions of overhead lights where light was not needed.

5.2 Airflow Optimization

5.2.1 Floor Tile Rearrangement

An analysis in airflows through tiles was performed based on the CRAH unit locations, perforated tile locations, and the computer rack locations. At some locations in the data center, airflow rates through the perforated tiles was relatively low, potentially creating areas of higher temperatures - hot spots due to inadequate heat removal. In Figure 10, the small colored squares represent the perforated tiles. The colors of the squares indicate the relative airflow rates through the tiles: darker blue indicates higher airflow rates; while lighter blue indicates lower airflow rates, and yellow-to-amber indicates even lower airflow rates. From the analysis, it is recommended that some of the perforated tiles in the data center should be re-arranged to induce more effective air distribution in the space, reducing occurrence of hot spots. In addition, additional analysis of the airflow should be conducted based on modifications to tile arrangements.

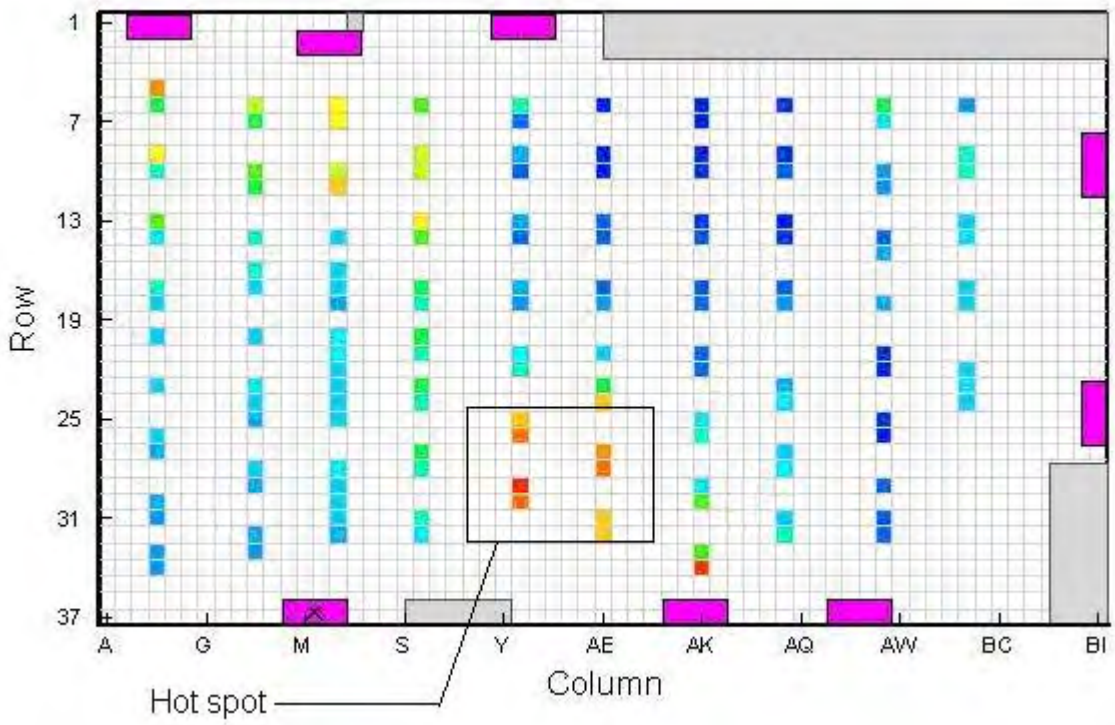

Figure 10 Airflow rates through perforated tiles and hot spots

In summary, cold air was un-evenly distributed throughout the data center. Either blocking unwanted openings on the raised floor or reducing airflow rates using adjustable dampers, or increasing tile perforation in the areas with limited air supply would result in a more even air distribution, thus reducing potential hot spots in the data center.

5.2.2 Wiring Configuration

Cables hanging in front of computer racks caused undesirable airflow deviations in cooling the rack equipment. These communication cables should be properly managed in front of the server or re-routed to the back of the equipment in order to reduce air circulation restrictions, because cold air is drawn through the front of the server to the backside of the server. Adding blank-offs within and between racks could prevent air bypasses and undesired mixing between hot and cold air flows.

5.2.3 Rack Air Management

The lack of delineated hot aisles and cold aisles layout was a significant impediment to proper air distribution in the data center. In hot-cold aisle configurations, racks would be installed in face-to-face and back-to-back configurations, with perforated tiles located below the face-to-face aisles, and no perforated tiles in the back-to-back aisles. Warmer return air would be

drawn from the hot aisles by the CRAH unit supply fans, routed through the cooling coil and supplied back into the raised floor plenum.

A primary recommendation for the rack layout is that cold supply air should flow from the front to the rear of the IT equipment. The study recommends to arrange fronts of equipment on each side of a cold aisle face each other and the backs of equipment in adjacent aisles (hot) facing each other. The cool air entering the front of the computer servers forms a common cold aisle and, warm air discharging at the rear of the servers forming a common hot aisle.

5.3 Air Temperature and Relative Humidity

The recorded data center ambient air's relative humidity ranged between 15% and 35% RH. Changing chilled water supply temperatures may affect relative humidity in the data center: Higher supply water temperatures may correspond to higher air humidity in the space, and would increase the chiller efficiency. The desired air temperatures in the data center were between 71°F and 73°F and a relative humidity between 45 and 50%. To achieve these conditions it is recommended that the supply air temperature from the CRAH units be maintained at 55°F. This would require that the current control on return air temperature be changed to control on supply air temperature. The IT equipment specifications should be reviewed regarding inlet air temperature requirements. With improved air management, it is likely that supply air temperatures from the CRAH could be increased, with resultant chiller plant efficiency improvement and better humidity control.

The CRAH unit manufacturer's specifications should be compared with the present CRAH unit operating conditions to determine what the cooling capacity would be under the proposed operating condition (higher supply air set point and higher chilled water temperature). Additionally, the IT equipment specifications about air humidity requirements should be compared against the actual operating conditions to determine whether further humidification would be required. If no equipment reliability problems had been experienced due to the lower humidity than 45%RH, then it is probably unnecessary to humidify the space within the data center.

5.4 Chilled Water System

Chiller plant optimization should be considered, including control tuning and reset of chilled water and tower water temperatures. For example, setting the chilled water supply temperature to 50°F may provide sufficient sensible cooling in the data center. In the meanwhile, chiller energy consumption would be reduced due to improved operating efficiency. In addition, chiller plant optimization also involves optimizing the interaction between tower fan, condenser water pump, and chiller power to achieve the lowest overall chiller plant power at any given load and ambient condition. As part of this strategy, consideration should be given to adding variable speed drives to the primary chilled water pumps and the tower water pumps.

6 Acknowledgements

This report on data center energy benchmarking was finalized based upon the field data collection performed by EYP Mission Critical Facilities and Landsberg Engineering and the draft report produced in the course of performing work subcontracted for the Lawrence Berkeley National Laboratory (LBNL).

The project was funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

7 Appendix A: Data Facility Definitions and Metrics

The following definitions and metrics are used to characterize data centers:

Air Flow Density	The air flow (cfm) in a given area (sf).
Air Handler Efficiency 1	The air flow (cfm) per power used (kW) by the CRAC unit fan.
Air Handler Efficiency 2	The power used (kW), per ton of cooling achieved by the air-handling unit.
Chiller Efficiency	The power used (kW), per ton of cooling produced by the chiller.
Computer Load Density – Rack Footprint	Measured Data Center Server Load in watts (W) divided by the total area that the racks occupy, or the "rack footprint".
Computer Load Density per Rack	Ratio of actual measured Data Center Server Load in watts (W) per rack. This is the average density per rack.
Computer/Server Load Measured Energy Density	Ratio of actual measured Data Center Server Load in watts (W) to the square foot area (sf) of Data Center Floor. Includes vacant space in floor area.
Computer/Server Load Projected Energy Density	Ratio of forecasted Data Center Server Load in watts (W) to the square foot area (sf) of the Data Center Floor if the Data Center Floor were fully occupied. The Data Center Server Load is inflated by the percentage of currently occupied space.
Cooling Load – Tons	A unit used to measure the amount of cooling being done. One ton of cooling is equal to 12,000 British Thermal Units (BTUs) per hour.
Data Center Cooling	Electrical power devoted to cooling equipment for the Data Center Floor space.
Data Center Server/Computer Load	Electrical power devoted to equipment on the Data Center Floor. Typically the power measured upstream of power distribution units or panels. Includes servers,

	switches, routers, storage equipment, monitors and other equipment.
Data Center Facility	A facility that contains both central communications and equipment, and data storage and processing equipment (servers) associated with a concentration of data cables. Can be used interchangeably with Server Farm Facility.
Data Center Floor/Space	Total footprint area of controlled access space devoted to company/customer equipment. Includes aisle ways, caged space, cooling units electrical panels, fire suppression equipment and other support equipment. Per the Uptime Institute Definitions, this gross floor space is what is typically used by facility engineers in calculating a computer load density (W/sf).
Data Center Occupancy	This is based on a qualitative estimate of how physically loaded the data centers are.
Server Farm Facility	A facility that contains both central communications and equipment, and data storage and processing equipment (servers) associated with a concentration of data cables. Can be used interchangeably with Data Center Facility. Also defined as a common physical space on the Data Center Floor where server equipment is located (i.e. server farm).

8 Appendix B: Facility Diagrams

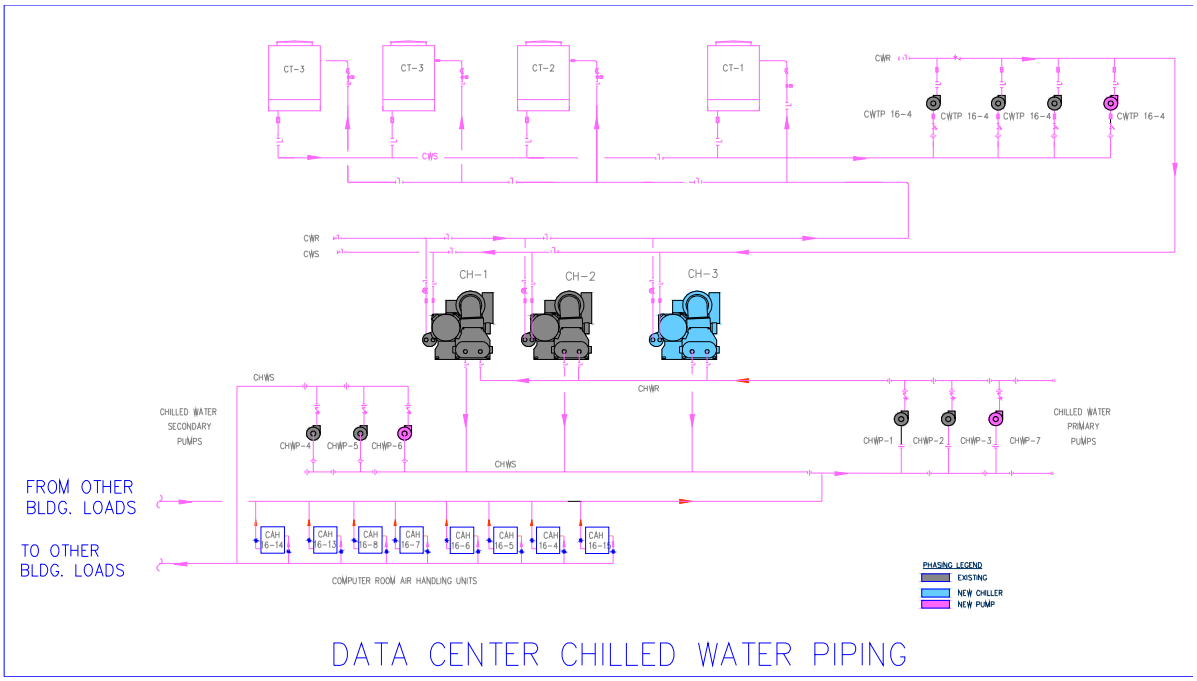

Figure 11 Chilled Water System

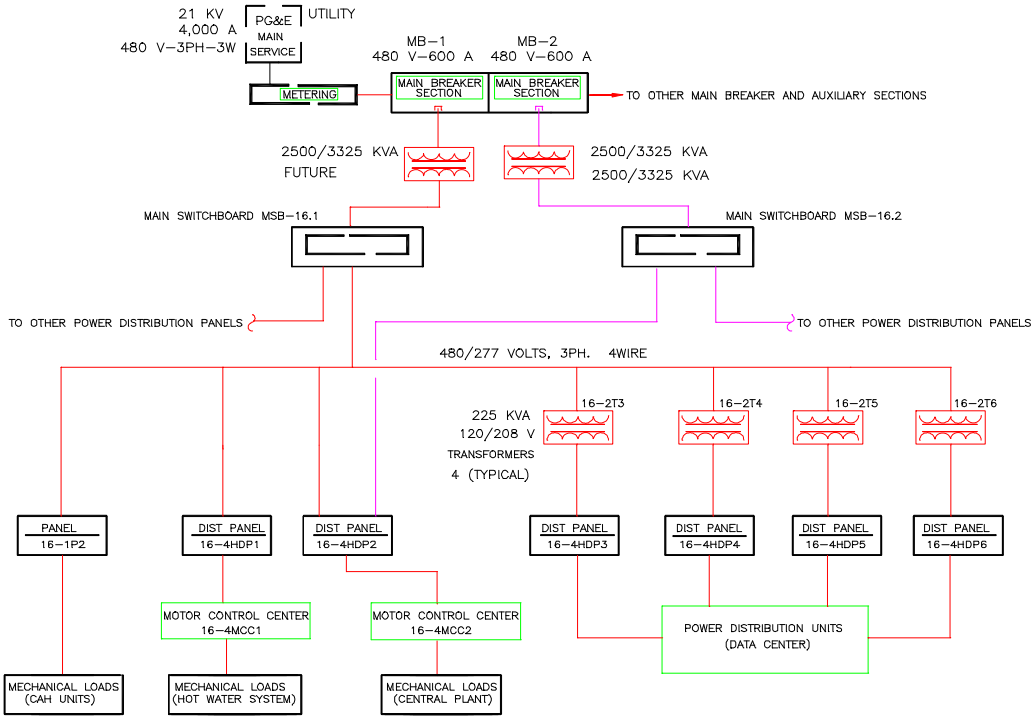

Figure 12 Electrical System Schematic

Data Center Energy Benchmarking:

Part 3 - Case Study on an IT Equipment-testing Center (No. 20)

Final Report

July 2007

Tengfang Xu and Steve Greenberg
Lawrence Berkeley National Laboratory (LBNL)
Berkeley CA 94720

Draft provided by

EYP Mission Critical Facilities
Los Angeles, CA 90064
and
Landsberg Engineering, P.C.
Clifton Park, NY 12065

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Table of Contents

1 EXECUTIVE SUMMARY 1

2 REVIEW OF SITE CHARACTERISTICS 2

2.1 ELECTRICAL EQUIPMENT AND BACKUP POWER SYSTEM 3
2.2 MECHANICAL SYSTEM 3

3 ELECTRIC POWER CONSUMPTION CHARACTERISTICS 4

3.1 CHILLER SYSTEM 7
3.2 PUMPING SYSTEM 8
3.3 COMPUTER ROOM AIR HANDLERS 9

4 SYSTEM OPERATION 10

4.1 CHILLED WATER SUPPLY AND RETURN TEMPERATURES 10
4.2 CRAH SUPPLY AND RETURN AIR 12
4.3 DATA CENTER SPACE AIR TEMPERATURE AND RELATIVE HUMIDITY 13

5 RECOMMENDATIONS 14

5.1 LIGHTING 15
5.2 AIRFLOW OPTIMIZATION 15
5.3 HVAC CONTROLS 15
5.4 CHILLED WATER SYSTEM 16

6 ACKNOWLEDGEMENTS 16

7 APPENDIX A: DATA FACILITY DEFINITIONS AND METRICS 17

8 APPENDIX B: FACILITY DIAGRAMS 19

1 Executive Summary

The data center in this study had a total floor area of 3,024 square feet (ft^2) with one-foot raised-floors. It was a rack lab with 147 racks, and was located in a 96,000 ft^2 multi-story office building in San Jose, California. Since the data center was used only for testing equipment, it was not configured as a critical facility in terms of electrical and cooling supply. It did not have a dedicated chiller system but was served by the main building chiller plant and make-up air system. Additionally it was served by only a single electrical supply with no provision for backup power in the event of a power outage. The Data Center operated on a 24 hour per day, year-round cycle, and users had full-hour access to the data center facility.

The study found that data center computer load accounted for 15% of the overall building electrical load, while the total power consumption attributable to the data center including allocated cooling load and lighting was 22% of the total facility load. The density of installed computer loads (rack load) in the data center was 61 W/ft^2. Power consumption density for all data center allocated load (including cooling and lighting) was 88 W/ft^2, approximately eight times the average overall power density in rest of the building (non-data center portion). The building and its data center cooling system was provided with various energy optimizing systems that included the following

- Varying chilled water flowrate through variable speed drives on the primary pumps.
- No energy losses due to nonexistence of UPS or standby generators.
- Minimized under-floor obstruction that affects the delivery efficiency of supply air.
- Elimination of dehumidification/humidification within the CRAH units.

For the data center, 70% of the overall electric power was the rack critical loads, 14% of the power was consumed by chillers, 12% by CRAH units, 2% by lighting system, and about 2% of the power was consumed by chilled water pumps.

General recommendations for improving overall data center energy efficiency include improving the lighting control, airflow optimization, control of mechanical systems serving the data center in actual operation.. This includes chilled water system, airflow management and control in the data center. Additional specific recommendations or considerations to improve energy efficiency are provided in this report.

2 Review of Site Characteristics

The data center (DC #20) in this study had a total floor area of 3,024 square feet (ft^2) with one-foot raised-floors. It was a rack lab with 147 racks, and was located in a 96,000 ft^2 multi-story office building in San Jose, California. Since the data center was used only for testing equipment, it was not configured as a critical facility in terms of electrical and cooling supply. The data center did not have a dedicated chiller system but was served by the main building's chiller plant and make-up air system. Additionally it was served by only a single electrical supply with no provision for backup power in the event of a power outage. The center operated on a 24 hour per day, year-round cycle, and users had full access to the data center facility at all hours.

Electric power was supplied to the office building from the utility to a single 3,000A, 480V main service switchboard (MSB-H). the MSB fed a 2,000A, 480V building electrical distribution system switchboard (SB-1) via current transducer auto metering. Data center computer power is provided directly from SB-1 (via a 225kVA transformer), without the use of Uninterruptible Power Supplies (UPS's) normally associated with typical data centers. In addition, there was no standby generator serving the facility. Communication and power wiring was installed above ceiling. Fire sprinklers were provided under the raised floor and at the ceiling.

Cooling for the data center facility was served by the building's main chiller plant and the make-up air system. The building chilled water system included five air-cooled chillers with chilled water pumps. The cooling system inside the data center included five Computer Room Air Handling (CRAH) units that received chilled water from the chiller plant serving complete building. Three of the CRAH units were monitored in this study. The CRAH's used down-flow type supply for air distribution, and return to the top with or without ducts. Two (Tag numbers CRAH-1.7 and CRAH-1.5) of the three CRAH's monitored used ducted return to the ceiling plenum. Figure 1shows a view of a typical rack lineup in the data center.

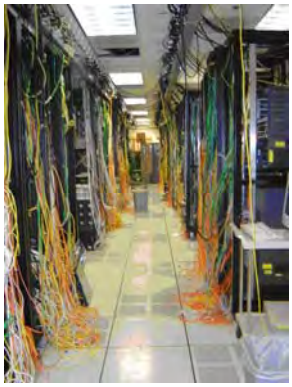

Figure 1 Typical computer rack lineup in the data center

2.1 Electrical Equipment and Backup Power System

Electrical power to Data Center #20 was served by the building electrical distribution system, which received power from Pacific Gas and Electric Company via a three-phase, 480-V, 3000-Ampere main service drop to a main switchboard (MSB-H) as shown on the Electrical System One-line Diagram in Figure 12 of Appendix B.

As part of the study, the overall building power consumption was monitored over a two-week period in December 2004. During that period, the average amount of power consumed in the building during peak usage (workday) periods was 1,238 kW.

The building main switchboard, MSB-H, fed a sub-switchboard SB-1. SB-1 was a three-phase, 480/277-V, four-wire 2,000-Ampere Bus switchboard. The SB-1 switchboard served Data Center #20 through a 225-kVA transformer, identified as HFT7, and three-phase, a 120/208V, four-wire 800-Ampere distribution panel fed by HFT7. All of the server racks received power from the 800-Ampere panel. There was no uninterruptible power supply, static transfer switch, or standby generator serving the data center.

Switchboard SB-1 also fed the building chillers. Power to each chiller was supplied through its own 150-amp circuit breaker located in SB-1. The building chilled water pumps and the CRAH units were fed from a motor control center, H2-MCC1. The H2-MCC1 motor control center received its power directly from main switchboard MSB-H.

The lighting and miscellaneous loads were fed from other electrical panels.

2.2 Mechanical System

2.2.1 Chiller

The data center, along with the remainder of the building, was designed to be cooled by five 60-ton air-cooled chillers. Each chiller had a design power consumption of 0.9 kW per ton of produced refrigeration.

Three chillers were operating during a majority of the monitoring period, with a fourth being placed on-line at no load during a small percentage of the time. Throughout the monitoring period the chillers operated at approximately 49% to 54% of total motor capacity - Chiller #1 carrying 46% of the cooling load, chiller #2 carrying 38%, and chiller #5 carrying about 16%.

2.2.2 Chilled Water Pumps

Primary chilled water was circulated by six parallel in-line Bell & Gossett pumps, each with a motor capacity of 7.5 HP, and a design volume flow rate of 135 GPM. Each pump was sized to provide the flow required for one chiller. The chilled water pumps are identified as CHP-1H through CHP-6H. Each chilled water pump was provided with a variable-speed drive. There was no secondary pump in the building chilled water system.

One of the six pumps, CHP-6H, serves as a redundant pump for backup. CHP-5H and CHP-6H were not in operation at the time of this study. The observed pump discharge pressure (with four pumps in operation) was 50 psig and the suction pressure was 46 psig per the installed pressure gauges. The chilled water pump operation was automatically controlled by the building cooling load demand through the facility's Trane Trace building management system.

2.2.3 Computer Room Air Handling Units

Three out of five package Computer Room Air Handling (CRAH) units were in operation in Data center #20, supplying the data center with cold air from the one-foot raised floor. The CRAH's were Pomona Air Model # PW 3000 units. No reheat coils or humidifiers were included in the CRAH units. The CRAH units had 4-inch throwaway air filters located at the top of the unit, rated at 85% efficiency. Each CRAH unit's internal controls were set to maintain temperature and relative humidity set-points of 70°F and 20% RH, respectively, measured at the unit's return air intake.

Figure 2 below shows a photograph of the CRAH units in data center #20. CAH-1.5 and CAH-1.6 can be seen on the left in the foreground and CAH-1.7 is barely visible at the end of the room. As can be seen in the photo, CAH-1.5 has a ducted return from the ceiling plenum and CAH-1.6 has a non-ducted return and pulls directly from the room.

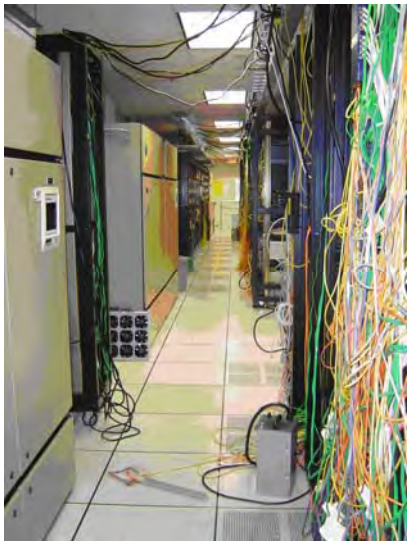

Figure 2 CRAH units in the data center

3 Electric Power Consumption Characteristics

Data center computer power consumption was monitored and shown in Figure 3. The power demand load was fairly steady between 178 and 190 kW, averaging 185 kW for the monitored period, and varying no more than 4% from the average. This data was recorded continuously and shows that there is no discernable cycle in load variation between day and night or between weekdays and weekends (the weekend days were 12/4 & 5 and 12/11 & 12).

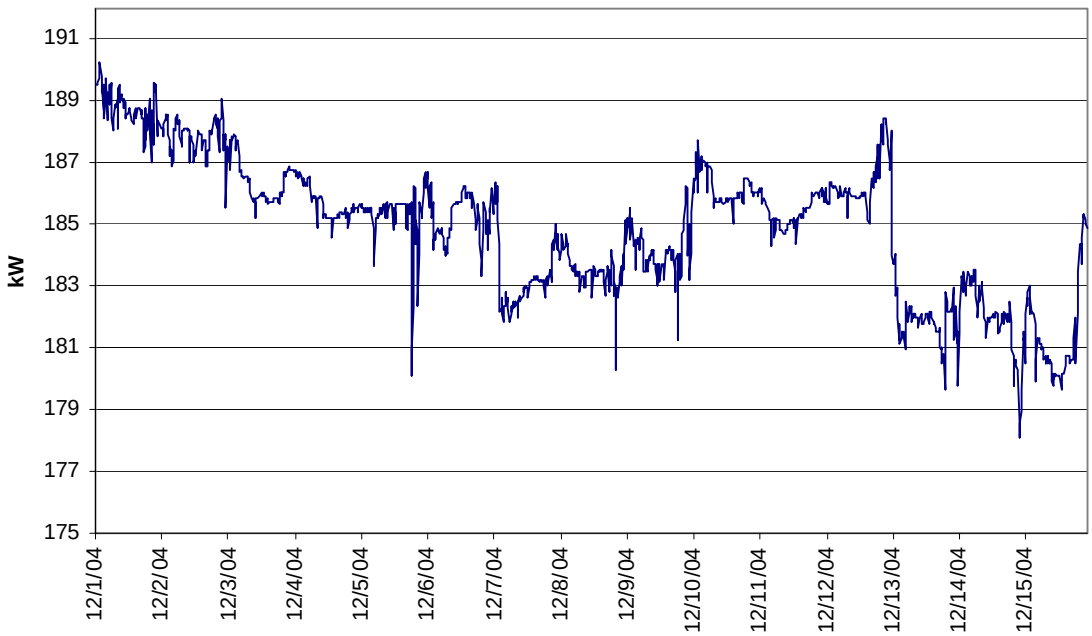

Figure 3 Data center rack power demand over two weeks

Table 1 shows the end-use electricity demand of the building housing the data center in this study. The table shows the power density for the square foot area served by each load. The average building electrical load of 1,238 kW was recorded from building instruments. From the measurements, about 79% of the electrical load was consumed by other areas of the building than the data center, including total cooling systems. For the remaining 22% of power consumption, approximately 15% of the building load was consumed by the data center computer equipment and about 7% of the load was by the CRAH units and the data center allocation of chiller and chilled water pump loads. The density of installed computer loads (rack load) in DC #20 was 61 W/ft^2. Power consumption density for all data center allocated load (including cooling and lighting) was 88 W/ft^2, approximately eight times the average overall power density in rest of the building (non-data center portion).

The ratio of HVAC to IT power demand in the data centers in this study was approximately 0.4. An estimate of "rack-cooling load" was calculated based upon the data center critical power load, assuming 100% of the critical power would become cooling load. For example, using the critical power of 185 kW in the data center, the rack-cooling load of the data center would then be approximately 53 tons. By calculation, the sensible load from the racks, lighting, and CRAH units was about 65 tons. This was used to allocate the portion of the chiller plant serving the data center. The portion of the chilling plant load allocated to the data center was estimated as 42 kW for the chillers and 5 kW for the chilled water pumps.

Table 1. End-Use of Electricity of the Data Center Building

Description	Electric power demand (kW)	Share of electric energy use (%)	Floor Space (ft2)	Electric power density (W/ft2)
Overall Building Load	1238	100%	96,000	13
Chillers – (non Data Center load)	41.5	3.4%	92,976	0.5
CHW Pumps – (non DC load)	4.5	0.4%	92,976	0.1
Building - (other non DC load)	926	74.8%	92,976	10
Total Non-Data Center Load	972	78.5%	92,976	11
Data Center Computer Load	185	14.9%	3,024	61
Data Center CRAH Units	30	2.4%	3,024	10
Chillers – (Data Center load)	41.5	3.4%	3,024	14
CHW Pumps – (Data Center load)	4.5	0.4%	3,024	1
Data Center UPS Losses	0	0.0%	3,024	0
Data Center Lighting	5	0.4%	3,024	2
Total Data Center Load	266	21.5%	3,024	88

The end-use breakdown for the data center's electric power demand is also shown in Table 2. For the data center, 70% of the overall electric power was the rack critical loads, 14% of the power was consumed by chillers, 12% by CRAH units, 2% by lighting system, and about 2% of the power was consumed by pumps.

Table 2. End-Use of Electricity of the Data Center

Description	**Electric power demand (kW)**	**Share of electric energy use (%)**
Data Center Rack Load	185	70%
Data Center CRAH Units	30	12%
Chillers (DC portion)	41.5	14%
CHW Pumps (DC portion)	4.5	2%
Standby Generator	none	0%
Data Center UPS Losses	none	0%
Data Center Lighting	5	2%
Total Data Center Only	**266**	**100%**

3.1 Chiller System

Electric power demand was monitored for the four operating chillers (CH-1, CH-2, CH-3, and CH-5) within a two-week period. Chiller CH-3 was only operational for three days (at very low loads) and chiller CH-4 never ran. The remaining three chillers operated continually for the monitoring period.

Total average chiller kW for the period was 83 kW, which represented the chilling load of the whole building. CH-1 and CH-2 provided the majority of the cooling load with average kW power consumption of 37.8 and 31.8 respectively. The chiller CH-5 operated with a power consumption of 13.4 kW.

Figure 4 shows the total chiller power demand and outside air temperatures within the two-week period. As the outside air temperature changed, the actual chiller power demand also changed accordingly. In addition, Figure 5 further shows the same data points illustrating the correlation of total chiller power consumption and outside air temperatures. The variations in power demand and air temperatures in both figures indicate that the chiller load was not solely affected by the ambient temperature but also was influenced by other factors.

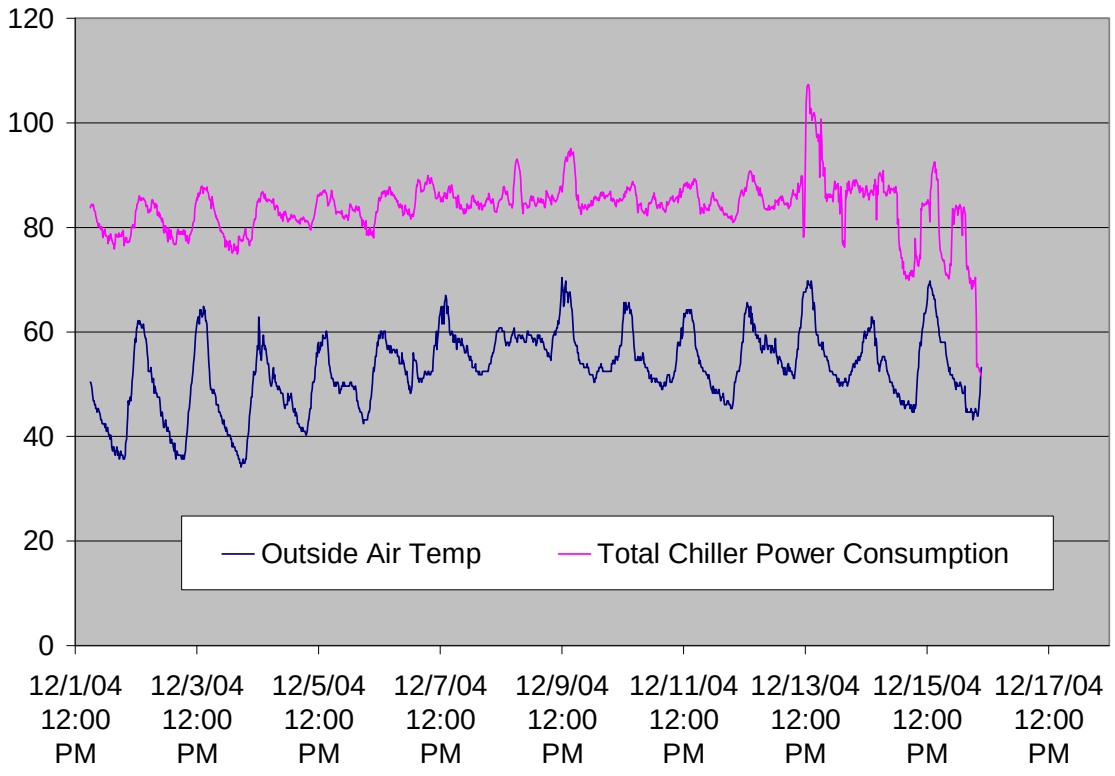

Figure 4 Chiller cooling power and outdoor air temperature

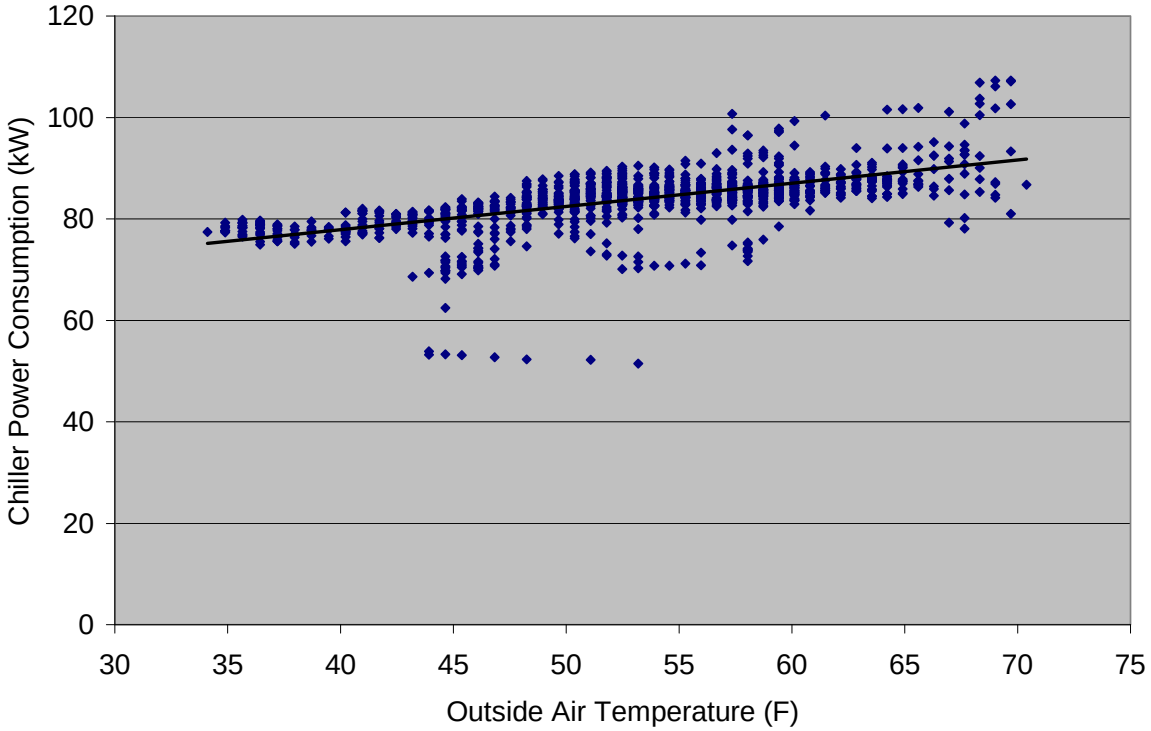

Figure 5 Correlation of chiller cooling power and outdoor air temperature

3.2 Pumping System

The building was served by six 7.5-hp chilled water pumps: CHP-1H through CHP-6H. CHP-1H through 3H were operational at the time of monitoring, while CHP-4H was operational for the first half period of the monitoring. Pumps CHP-5H and CHP-6H were for reserve and were not in operation.

The pumps were fitted with variable speed drives (VSDs) and controlled to modulate from 30HZ to 60HZ. Control sequence was not known. The power consumption of the pumps was monitored and shown in Figure 6. The average total pumps kW for the monitoring period was 9 kW.

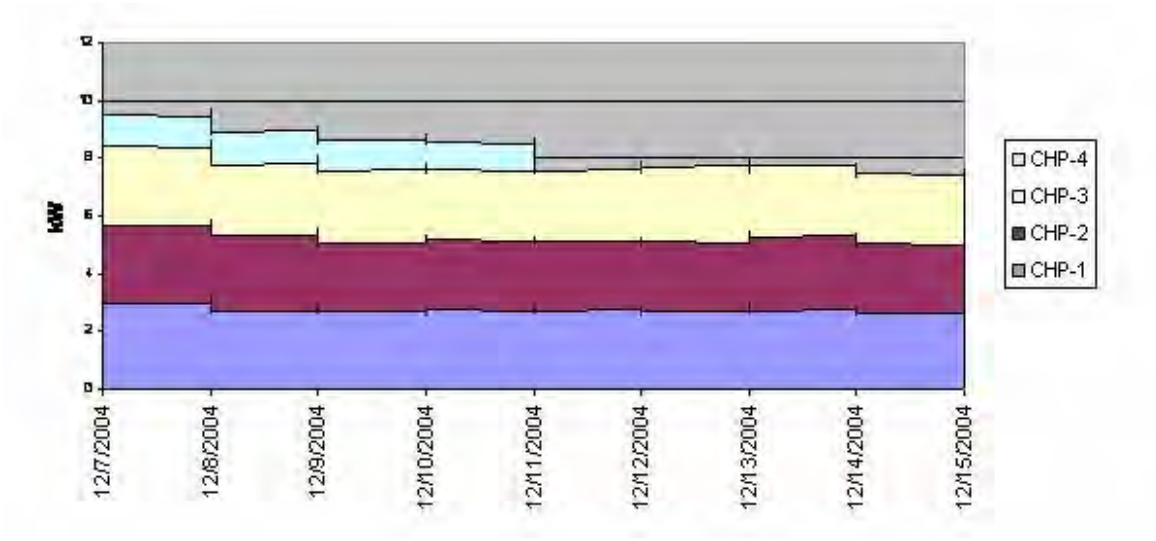

Figure 6 Chilled Water Pumps

3.3 Computer Room Air Handlers

The data center was actively served by three 20-ton CRAH units during the survey period. Each unit had two 7.5-hp constant-speed fan motors.

The power consumption of CRAH unit CAH-1.5 is shown in Figure 7. Based on data taken at the time of the survey, CRAH power consumption was 0.84 watts/cfm. The variation of minimum to maximum power draw was under 4% during the period.

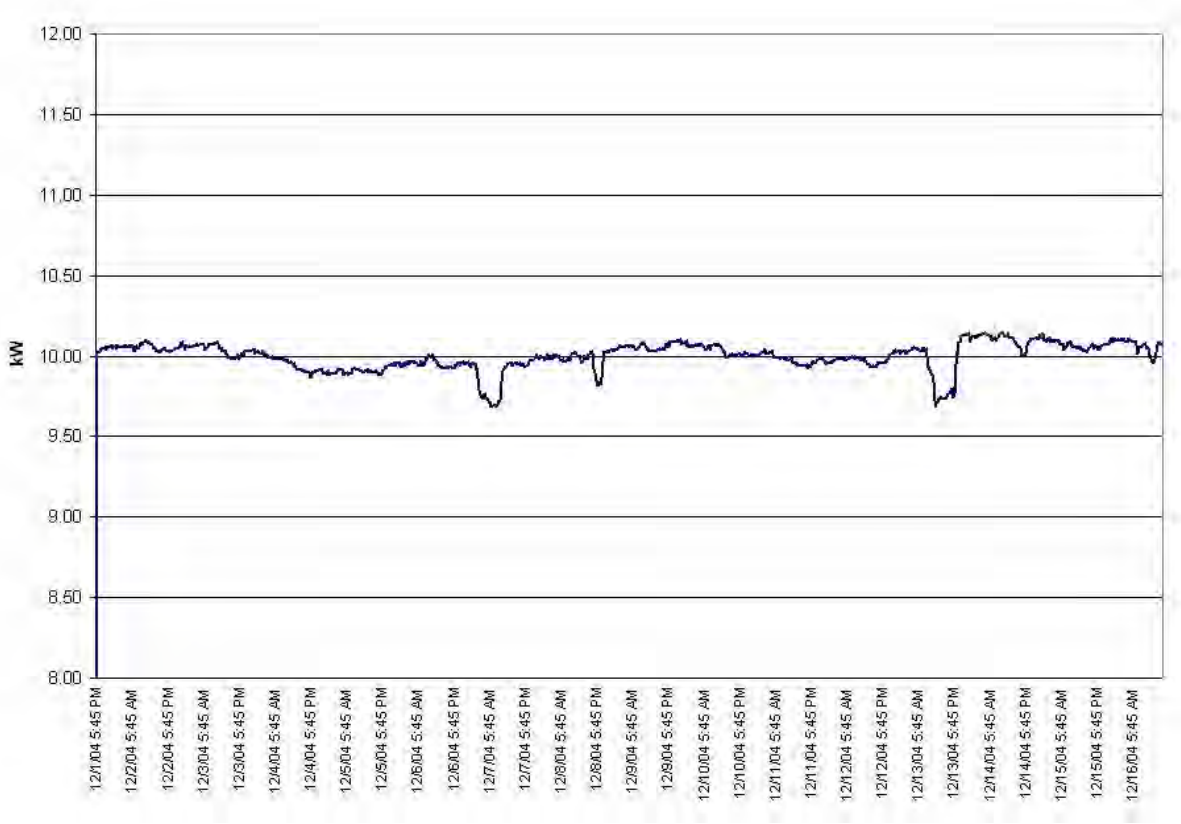

Figure 7 CRAH unit CAH-1.5 Power Consumption

4 System Operation

During the two-week monitoring period, the following HVAC equipment was operating:

- Chilled water pumps CHP-1H, CHP-2H, CHP-3H, and CHP-4H
- Chillers CH-1, CH-2, CH-3 and CH-5.
- All three monitored CRAH units CAH-1.5, CAH-1.6 and CAH-1.7 were operating, while the other two units were off.

4.1 Chilled Water Supply and Return Temperatures

The chilled water supply and return temperatures were monitored in the study.

The data center branch-chilled water return temperature was shown in Figure 8 along with the corresponding ambient temperature and CRAH supply air temperature (for CAH 1.5). The recorded chilled water supply temperature readings were almost identical to the ambient

temperature. In addition, the recorded chilled water supply temperature showed that it rose well above the CRAH supply air temperature every day (which is not possible) in concert with the ambient temperature. It was found that the temperature sensor used for recording chilled water supply temperature was surface-mounted on the piping and was apparently much more influenced by ambient air temperature around the piping than it was by the temperature of the chilled water flowing in the pipe. This indicates that the readings were not right and a verification or calibration should be pursued to correct the error.

The chilled water return temperature correlated with the ambient temperature to some degree. The maximum lows for each day occurred in the morning hours (between 5:00 AM and 8:00 AM) while the peak highs occurred in the early afternoon hours (between 12:00 PM and 3:00 PM). For the monitoring period, the average chilled water return temperature was 71°F, and the average CRAH supply temperature for CAH-1.5 was 55°F.

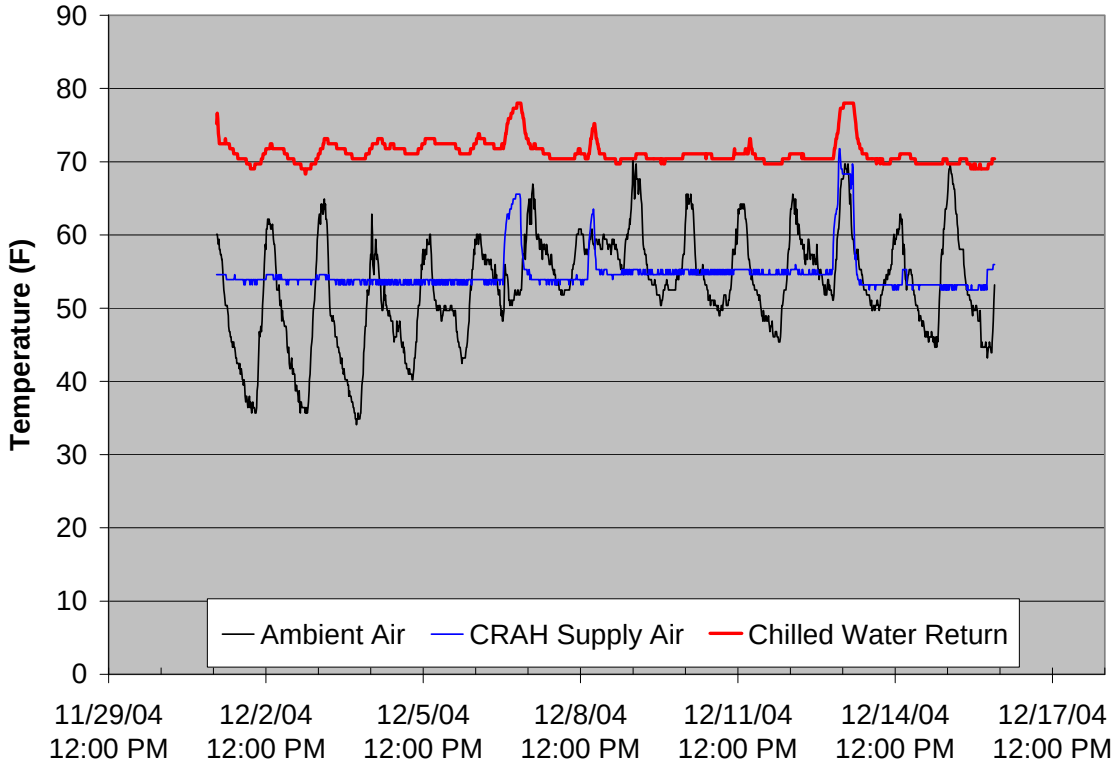

Figure 8. Water and air temperatures

4.2 CRAH Supply and Return Air

Temperatures and relative humidity for the supply and return airflows were monitored for two weeks. Figure 9 shows the supply air temperatures for air handlers CAH-1.5, and CAH-1.6 as well as the return air temperature and relative humidity for CAH-1.7. Supply air temperature excursions were experienced on three different occasions as shown in Figure 9. According to the Building Engineer, the excursions were caused by problems with the modulation of the chilled water control valve. The exact location of the control valve was not known. Therefore, the control valves should be located, checked, and repaired or replaced as needed to ensure stable and reliable control, reducing the power demand for operating the units.

Each CRAH unit was equipped with two 7.5-hp fan motors and was rated to deliver 12,000 ft^3/m (CFM) of cold air to the under floor plenum with a net sensible cooling capacity of 20-tons (after deducting fan heat). During the survey, three operating CRAH's (CAH-1.5, CAH-1.6, and CAH-1.7) were monitored and measured. The average supply temperature for CAH-1.5 was 55°F. CAH-1.6 had an average supply temperature of 58°F and an average supply relative humidity of 51% RH. CAH-1.7 had an average return air temperature of 79°F and an average return relative humidity of 26% RH.

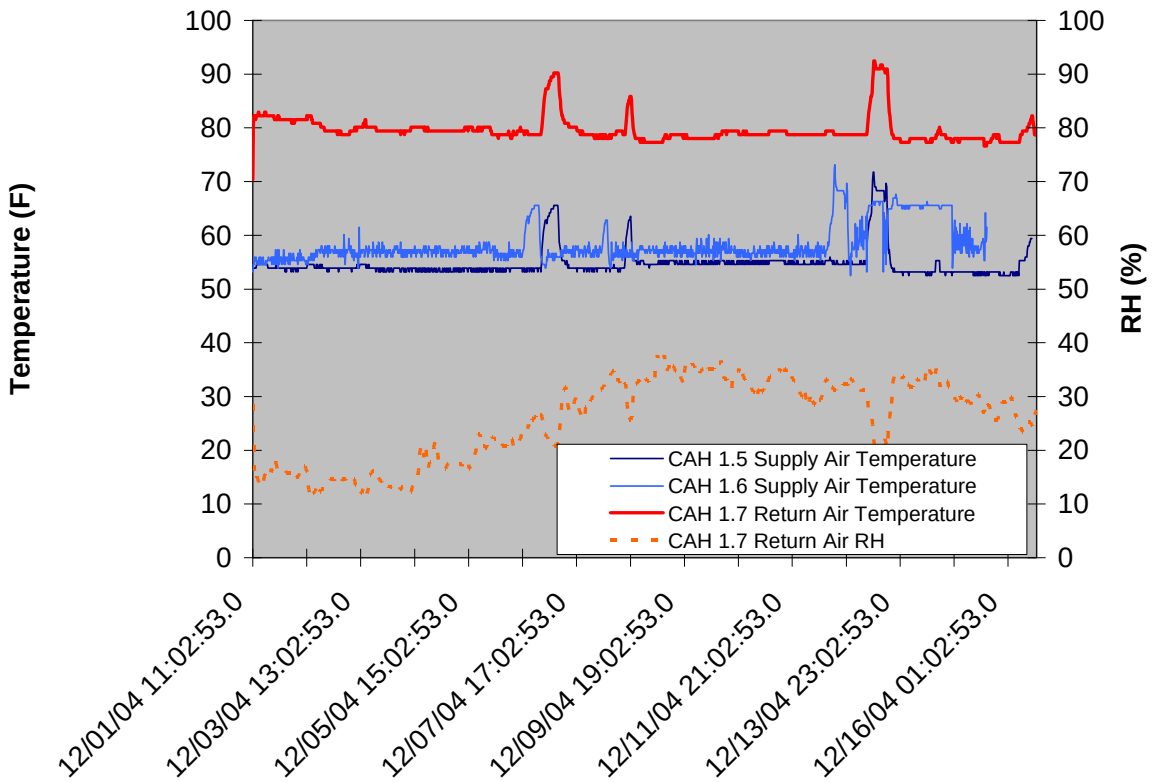

Figure 9 DC Computer Room Air Handling Unit Supply, Return Temperatures and Relative Humidity

4.3 Data Center Space Air Temperature and Relative Humidity

The space temperature and relative humidity were monitored and shown in Figure 10. At one location the average space temperature and relative humidity was 74°F and 32% RH, respectively.

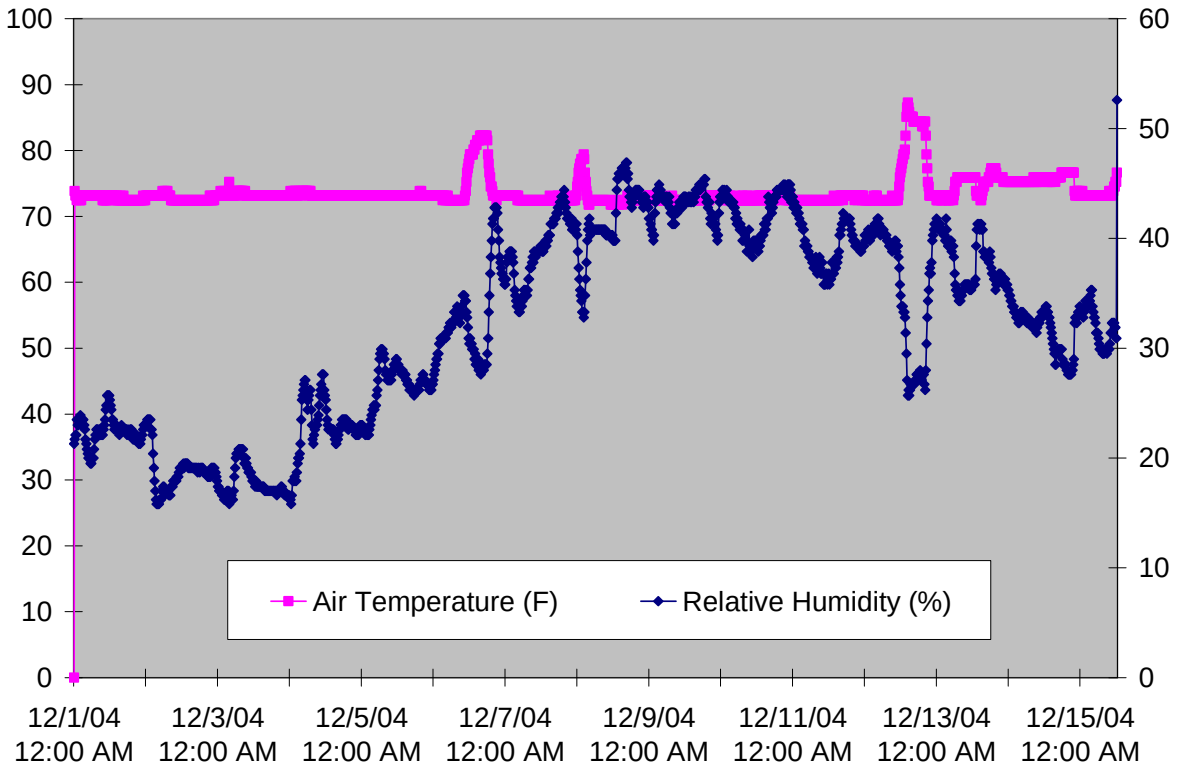

Figure 10 DC air temperature and humidity

5 Recommendations

The density of installed computer loads (rack load) in the data center was 61 W/ft^2. The building and its data center cooling system was provided with various energy optimizing systems that included the following:

- Varying chilled water flow rate through variable speed drives on the primary pumps.
- No energy losses due to nonexistence of UPS or standby generators.
- Minimized under-floor obstruction that affects the delivery efficiency of supply air.
- Elimination of dehumidification/humidification within the CRAH units.

General recommendations for improving overall data center energy efficiency include improving the lighting control, design, operation, and control of mechanical systems serving the data center in actual operation. This includes chilled water system, airflow management and control in data centers. The following additional techniques should result in significant improvements in energy efficiency, effective operation, or both.

5.1 Lighting

The measured lighting load in the Data center was 5 kW with an intensity of $2W/ft^2$. The lighting load can be reduced by considering the following energy control measures: install lighting zone occupancy sensors; and task lighting in appropriate areas and disable portions of overhead lights where light is not needed.

5.2 Airflow Optimization

5.2.1 Floor Tile Rearrangement

Cold air was un-evenly distributed throughout the data center. Either blocking unwanted openings on the raised floor or reducing airflow rates using adjustable dampers (or lower % open perforated tiles) on the areas with excessive air flow would result in a more even air distribution, thus reducing potential hot spots in the data center.

5.2.2 Rack Air Management

A primary recommendation for the rack layout is that cold supply air should flow from the front to the rear of the IT equipment. The study recommends to arrange fronts of equipment on each side of a cold aisle face each other and the backs of equipment in adjacent aisles (hot) facing each other. The cool air entering the front of the computer servers forms a common cold aisle and, warm air discharging at the rear of the servers forming a common hot aisle.

5.2.3 Wiring Configuration

Cables hanging in front of computer racks caused undesirable airflow deviations to cool the rack equipment. These communication cables should be properly managed in front of the server to reduce air circulation restrictions as cold air is drawn through the front of the server to the backside of the server. An additional recommendation would be to utilize a server rack with a bottom opening for cold air and perforated tile in the bottom of the computer server. Adding blank-offs within and between racks could prevent air bypasses and undesired mixing between hot and cold air flows.

5.3 HVAC Controls

The space relative humidity in the data center varies considerably as shown in figures 16 & 17. The CRAH units can be globally controlled by an energy management system through the establishment of temperature set point(s) in the space. This can be created by adding input and output control points at the EMS.

- The current control sequence of chilled water pump was not known. Chilled water pump flow can be improved by providing a control that is based on differential pressure (DP) in the system.

- Check and tune controls and instrumentation for chillers and CRAH's.

- Check the electrical loading on the main and the circuit transducer; in places it appeared to be above the 80% continuous loading allowed by Code.

5.4 Chilled Water System

Consideration should be given towards resetting the chilled water supply temperature to a higher set point. For example, setting the chilled water supply temperature to 50°F may provide sufficient sensible cooling in the data center. In the meanwhile, chiller energy consumption would be reduced due to improved thermal efficiency. This measure can be implemented in steps, raising the temperature set point by 1 or 2°F at a time, while verifying that there is no hot spot in critical locations. In addition, employing evaporative pre-coolers for the air-cooled chiller condensers could increase chiller efficiency, especially at peak conditions.

6 Acknowledgements

This report on data center energy benchmarking was finalized based upon the field data collection performed by EYP Mission Critical Facilities and Landsberg Engineering and the draft report produced in the course of performing work subcontracted for the Lawrence Berkeley National Laboratory (LBNL).

The project was funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

7 Appendix A: Data Facility Definitions and Metrics

The following definitions and metrics are used to characterize data centers:

Air Flow Density	The air flow (cfm) in a given area (sf).
Air Handler Efficiency 1	The air flow (cfm) per power used (kW) by the CRAC unit fan.
Air Handler Efficiency 2	The power used (kW), per ton of cooling achieved by the air-handling unit.
Chiller Efficiency	The power used (kW), per ton of cooling produced by the chiller.
Computer Load Density – Rack Footprint	Measured Data Center Server Load in watts (W) divided by the total area that the racks occupy, or the "rack footprint".
Computer Load Density per Rack	Ratio of actual measured Data Center Server Load in watts (W) per rack. This is the average density per rack.
Computer/Server Load Measured Energy Density	Ratio of actual measured Data Center Server Load in watts (W) to the square foot area (sf) of Data Center Floor. Includes vacant space in floor area.
Computer/Server Load Projected Energy Density	Ratio of forecasted Data Center Server Load in watts (W) to the square foot area (sf) of the Data Center Floor if the Data Center Floor were fully occupied. The Data Center Server Load is inflated by the percentage of currently occupied space.
Cooling Load – Tons	A unit used to measure the amount of cooling being done. One ton of cooling is equal to 12,000 British Thermal Units (BTUs) per hour.
Data Center Cooling	Electrical power devoted to cooling equipment for the Data Center Floor space.
Data Center Server/Computer Load	Electrical power devoted to equipment on the Data Center Floor. Typically the power measured upstream of power distribution units or panels. Includes servers, switches, routers, storage equipment, monitors and other equipment.

Data Center Facility	A facility that contains both central communications and equipment, and data storage and processing equipment (servers) associated with a concentration of data cables. Can be used interchangeably with Server Farm Facility.
Data Center Floor/Space	Total footprint area of controlled access space devoted to company/customer equipment. Includes aisle ways, caged space, cooling units electrical panels, fire suppression equipment and other support equipment. Per the Uptime Institute Definitions, this gross floor space is what is typically used by facility engineers in calculating a computer load density (W/sf).
Data Center Occupancy	This is based on a qualitative estimate of how physically loaded the data centers are.
Server Farm Facility	A facility that contains both central communications and equipment, and data storage and processing equipment (servers) associated with a concentration of data cables. Can be used interchangeably with Data Center Facility. Also defined as a common physical space on the Data Center Floor where server equipment is located (i.e. server farm).

8 Appendix B: Facility Diagrams

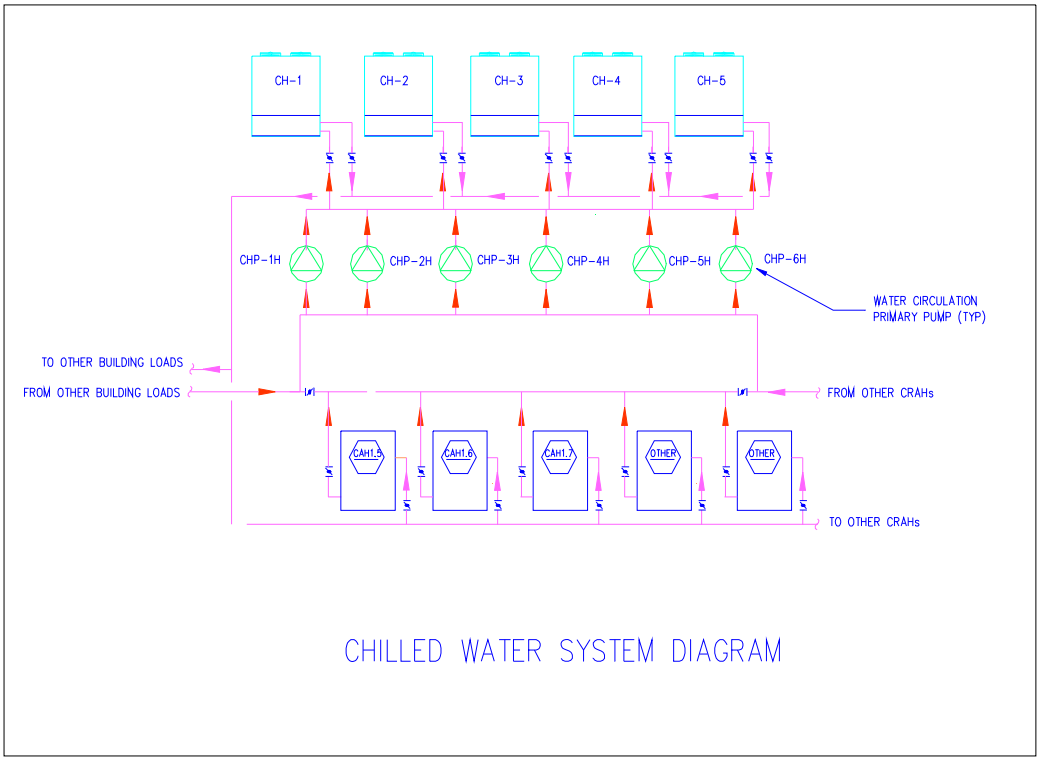

Figure 11 Chilled Water System

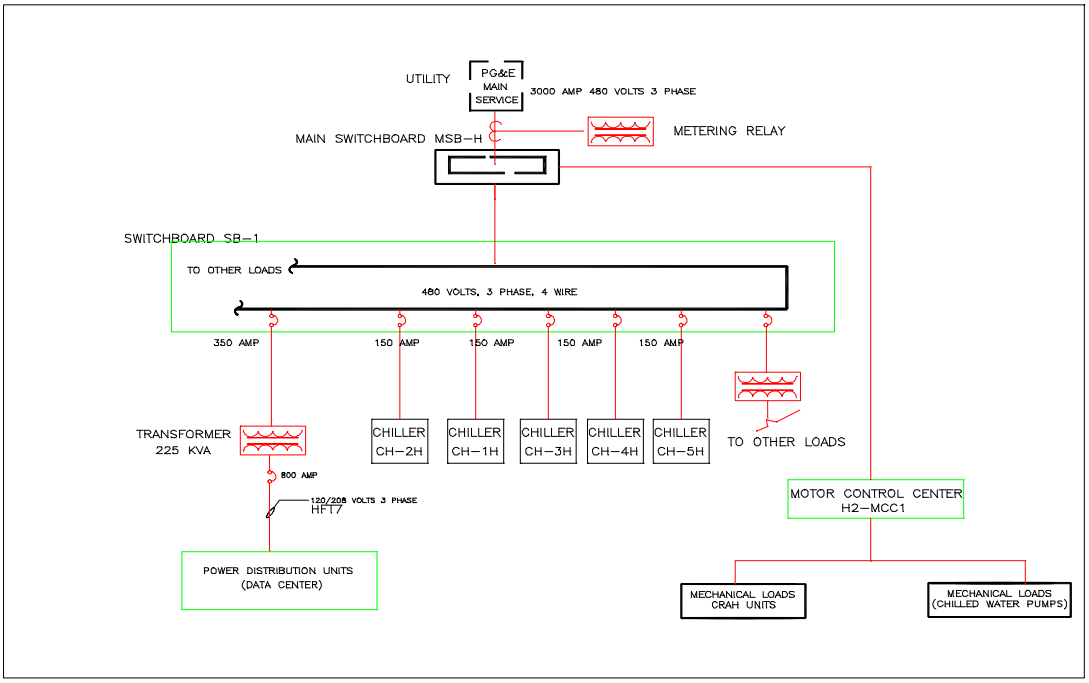

ELECTRICAL SYSTEM SCHEMATIC

Figure 12 Electrical System Schematic

Data Center Energy Benchmarking:

Part 2 - Case Studies on Two Co-location Network Data Centers (No. 18 and 19)

Final Report

August 2007

Final Report prepared by

Tengfang Xu and Steve Greenberg
Lawrence Berkeley National Laboratory (LBNL)
Berkeley CA 94720

Draft provided by

EYP Mission Critical Facilities
Los Angeles, CA 90064
and
Landsberg Engineering, P.C.
Clifton Park, NY 12065

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Table of Contents

1 EXECUTIVE SUMMARY 1

2 REVIEW OF SITE CHARACTERISTICS 2

2.1 ELECTRICAL EQUIPMENT AND BACKUP POWER SYSTEM 2
2.2 MECHANICAL SYSTEM 2

3 ELECTRIC POWER CONSUMPTION CHARACTERISTICS 3

3.1 PDU SYSTEM 7
3.2 EMERGENCY GENERATORS 7

4 MECHANICAL SYSTEM 7

4.1 CHILLER SYSTEM 8
4.2 PUMPING SYSTEM 9
4.3 AHU SYSTEM 10

5 SYSTEM OPERATION 10

5.1 CHILLED WATER FLOW 10
5.2 CHILLED WATER SUPPLY AND RETURN TEMPERATURES 11
5.3 AIR HANDLER UNIT SUPPLY AND RETURN AIR 11
5.4 GENERATOR JACKET AMBIENT TEMPERATURES 12

6 RECOMMENDATIONS 12

6.1 CHILLED WATER SYSTEM 13
6.2 AIR SYSTEM 13
6.3 LIGHTING 14
6.4 METERING AND POWER CONDITIONING EQUIPMENT 14

7 ACKNOWLEDGEMENTS 14

8 APPENDIX A: DATA FACILITY DEFINITIONS AND METRICS 15

9 APPENDIX B: FACILITY DIAGRAMS 17

1 Executive Summary

Two data centers in this study were within a co-location facility located on the sixth floor of a multi-story building in downtown Los Angeles, California. The facility had 37,758 gross square feet floor area with 2-foot raised-floors in the data services area. The two data centers were designated as the west data center (DC #18) and the east data center (DC #19).

The study found that 56% of the overall electric power was consumed by sixth floor critical loads in both data centers, 33% of the power was consumed by HVAC systems, 3% of the power was consumed by UPS units, 3% of the power was for generator losses, and the remaining 5% was used by lighting and miscellaneous loads in the building.

The power density of installed computer loads (rack load) in the two data centers was 20 W/ft^2 and 56 W/ft^2, respectively. The power density was relatively lower in DC #18 compared to other data centers previously studied. In addition, HVAC to IT power demand ratio was 0.6 in DC #18 in this study, and was 0.4 in DC #19.

Two out of three chillers were running at a low partial load, making the operation very energy inefficient. The operation and control of the chillers and air-handling units should be optimized while providing sufficient cooling to the data centers. Although arranging hot aisle/cold aisle design to separate airflow streams would be difficult in such a co-location data center, optimizing air distribution should be pursued.

General recommendations for improving overall data center energy efficiency include improving the design, operation, and control of mechanical systems serving the data centers with various critical loads in place. This includes chiller operation, chilled water system, AHUs, airflow management and control in data centers. Additional specific recommendations or considerations to improve energy efficiency are provided in this report.

2 Review of Site Characteristics

Data Centers # 18 and #19 were located on the sixth floor in a multi-story building in Downtown Los Angeles, California. The data center facility had a total floor area of 37,758 gross square feet (ft^2) with 2-foot raised-floors in the data services area. The data centers were designed to provide co-location data services in areas that are environmentally controlled and monitored. The data center space on the sixth floor was divided into east and west sections, each conditioned by five separate air-handling units. The air-handling units(AHUs) were controlled in unison to cool their respective sections. Chilled water was produced by three 315-ton air-cooled chillers and distributed via primary and secondary pumping water systems.

Energy monitoring was performed during the time of the study conducted between October 27 and November 3, 2004, Data Center #18 (west section) and Data Center #19 (east section) were in operation. Both data centers operated 24 hours per day year-round. The users of the data centers had 24-hour full access to and from their caged spaces. Security requirement was very high. Electric power for both data centers was supplied through three 5,000A, 480V main buses, each with a 750 kVA uninterruptible power supply (UPS).

2.1 Electrical Equipment and Backup Power System

Data Centers # 18 and #19 were served by LA Department of Water and Power and had three main service drops of 5,000A, 480V, 3-phase power to each floor, as shown in Appendix B. Main meter monitoring was not permitted, but the load of the occupied floor was 1,300 kVA.

Each of the service drops fed power to three 750 kVA uninterruptible power supply (UPS) units, which in turn supplied power to the power distribution units (PDUs) feeding the computer racks. There were separate power feeds and generator backup for lighting and HVAC panels. Each power distribution unit was supplied by two UPS feeds. The UPS units were originally designed to operate at 33% capacity but operated at 25% capacity at the time of this study.

2.2 Mechanical System

2.2.1 Chillers

Three Technical System air-cooled chillers (Model # 30A0TSM400) with a cooling capacity of 315-tons each are located on the building roof. The three chillers were designated as CH-1, CH-2 and CH-3. Each chiller had four variable capacity screw compressors. Unloading from 25% to 100% of compressor capacity was via internal slide valves. The chillers had condenser fan staging head pressure control for low ambient conditions down to 20°F. The chillers were sized such that one unit operating at 75% capacity could condition one floor but for normal operation of the plant two chillers were typically operating at partial load. At the time of survey chillers CH-1 and CH-2 were operating, CH-3 was a redundant unit.

2.2.2 Primary Chilled Water Pumps

Primary chilled water was circulated by two centrifugal Bell & Gossett pumps. The primary pumps were identified as (P1P and P2P). Each pump had a motor capacity of 25 HP and a design volume flow rate of 1,040 GPM. Only one primary pump was running at the time of this study. According to the on-site gauges, the primary pump had a discharge pressure of 55 psig and suction pressure of 36 psig with 19 psi of differential pressure.

2.2.3 Secondary Chilled Water Pumps

At the time of the study, secondary chilled water was supplied to the building by two parallel centrifugal pumps rated 40-HP. The pumps were identified as (P1S and P2S) Bell & Gossett fitted with variable speed drives. The pumps variable speed drives were controlled by return water temperature through the building EMCS.

2.2.4 Air Handling Units

The 6th floor space was divided into east (DC #19) and west (DC #18) sections. The east section (DC #19) was more heavily loaded with critical load at the time of the study. Each section was served by five Carrier central station air handling units (Model 39T) that supplied air to ceiling diffusers located above aisles separating the racks. There was no provision of outside air supply nor supply air humidification. There were four reserve air-handling units, as shown in Appendix B.

Each of the air-handling units had a sensible cooling capacity of 600 MBH and was capable of supplying 30,000 cubic feet per minute (ft^3/m). Each air-handling unit included 2" throwaway filters rated at 85% filtration efficiency, a chilled-water cooling coil, and a 30-hp supply-fan motor with a variable speed drive. The five air-handling units were controlled in unison for each section (east and west, respectively). The supply-air fan speed in the air-handling unit was controlled to maintain the static air pressure in the duct. The static pressure was one- to two-inch water columns (or 250-500 Pascal).

The air humidity ranged between 25%RH and 65%RH in both data centers. In the east section (DC #19), the supply air temperature averaged 60°F with 62%RH, while the return air averaged 67°F and 48% RH. In the west section (DC #18), the supply air temperatures averaged 68°F, while the return air averaged 70°F and 44%RH.

3 Electric Power Consumption Characteristics

The following table summarizes the power consumption measured at the facility during this study.

Table 1A. End-Use of Electricity of the Data Center Building

Description	Electric power demand (kW)	Share of electric energy use (%)	Floor Space (ft^2)	Electric power density (W/ft^2)
Overall Building Load	**1000**	**100%**	**37,758**	**25.6**
Data Center 6th Floor Overall Load Data	**564**	**56%**	**14,850**	**37.9**
6th Fl Critical Load PDU 1	197	20%	14,850	13.3
6th Fl Critical Load PDU 2	162	16%	14,850	10.9
6th Fl Critical Load PDU 3	205	21%	14,850	13.8
HVAC Systems	**331**	**33%**	**14,850**	**22.3**
Air Handlers AHU 1-10	65	7%	14,850	4.4
Pumps P2P and P2S	25	3%	14,850	1.7
Chillers	241	24%	14,850	16.2
Generator Losses	**26**	**3%**	**37,758**	**0.7**
Lighting	**45**	**5%**	**37,758**	**1.2**
UPS Losses	**34**	**3%**	**14,850**	**2.3**

A total building power demand of 1,000 kW was recorded from building instruments. The reading resulted in a power factor of approximately 0.76, suggesting that power factor correction was warranted. From these measurements, 56% of the overall electric power was consumed by sixth floor critical loads in both data centers, 33% of the power was consumed by HVAC systems, 3% of the power was consumed by UPS units, 3% of the power was for generator losses, and the remaining 5% was created by lighting and miscellaneous loads in the building.

The end-use breakdown for both data centers' electric power demand is shown in Table 1B. For both data centers combined, 65% of the overall electric power was the rack critical loads, 28% of the power was consumed by HVAC systems, 4% of the power was consumed by UPS units, 1% of the power was for generator losses, and the remaining 2% was for data center lighting.

Table 1B. End-Use of Electricity of the two Data Centers Only

Description	Electric power demand (kW)	Share of electric energy use (%)	Floor Space (ft^2)	Electric power density (W/ft^2)
Data Center Rack Power	**564**	**65%**	**14850**	**38**
6 Floor- PDU - 1	197	23%	14850	13.3
6 Floor- PDU - 2	162	19%	14850	10.9
6 Floor- PDU - 3	205	24%	14850	13.8
HVAC Systems	**246**	**28%**	**14850**	**16.6**
Air Handlers	65	7%	14850	4.4
Pumps P2P, P2S	17	2%	14850	1.1
Chillers	164	19%	14850	11
Generator Losses	**10**	**1%**	**14850**	**0.7**
Data Center Lighting	**17**	**2%**	**14850**	**1.2**
UPS Losses	**34**	**4%**	**14850**	**2.3**
Total Data Center (only)	**871**	**100%**	**14850**	**58.7**

The following explains how the energy use was estimated for each data center, individually. Using the frequency of the VFD for each center (44.3 HZ for the east side and 54.6 HZ for the west side), we estimated the airflow circulation, based on fan laws. Knowing the entering and

leaving air conditions to the cooling coils allowed an estimate of the cooling load for each individual center. From the cooling load in each center, and the respective energy use of the chiller (1.05 kW/ton), we estimated the chiller electrical load for each individual center. The electrical load for the pumps, generator losses, UPS losses, and PDUs were proportioned according to rack load.

Power demand break-down for each data center is shown in Table 2 and Table 3, respectively. The density of installed computer loads (rack load) in DC#18 and DC#19 was 20 W/ft^{2}and 56 W/ft^2, respectively. The ratios of HVAC to IT power demand in each of the data centers in this study were approximately 0.6 in DC #18 and 0.4 in DC #19.

Table 2. End-Use of Electricity of Data Center 18

Description	Electric power demand (kW)	Share of electric energy use (%)	Floor Space (ft^2)	Electric power density (W/ft^2)
Data Center Rack Power	148	59%	7425	20
HVAC Systems	82	33%	7425	11
Generator Losses	3	1%	7425	0.4
Data Center Lighting	9	3%	7425	1.1
UPS Losses	9	4%	7425	1.2
Total Data Center 18 (only)	250	100%	7425	33.7

Table 3. End-Use of Electricity of Data Center 19

Description	Electric power demand (kW)	Share of electric energy use (%)	Floor Space (ft^2)	Electric power density (W/ft^2)
Data Center Rack Power	416	67%	7425	56
HVAC Systems	164	26%	7425	22.1
Generator Losses	7	1%	7425	1
Data Center Lighting	9	1%	7425	1.1
UPS Losses	25	4%	7425	3.4
Total Data Center 19 (only)	621	100%	7425	83.6

An estimate of "rack-cooling load" was calculated based upon the data center critical power load, assuming 100% of the critical power becomes cooling load. For example, using the critical power of 148 kW and 416 kW in each data center, the rack-cooling loads of the data centers #18 and #19 would be approximately 42 tons and 118 tons, respectively.

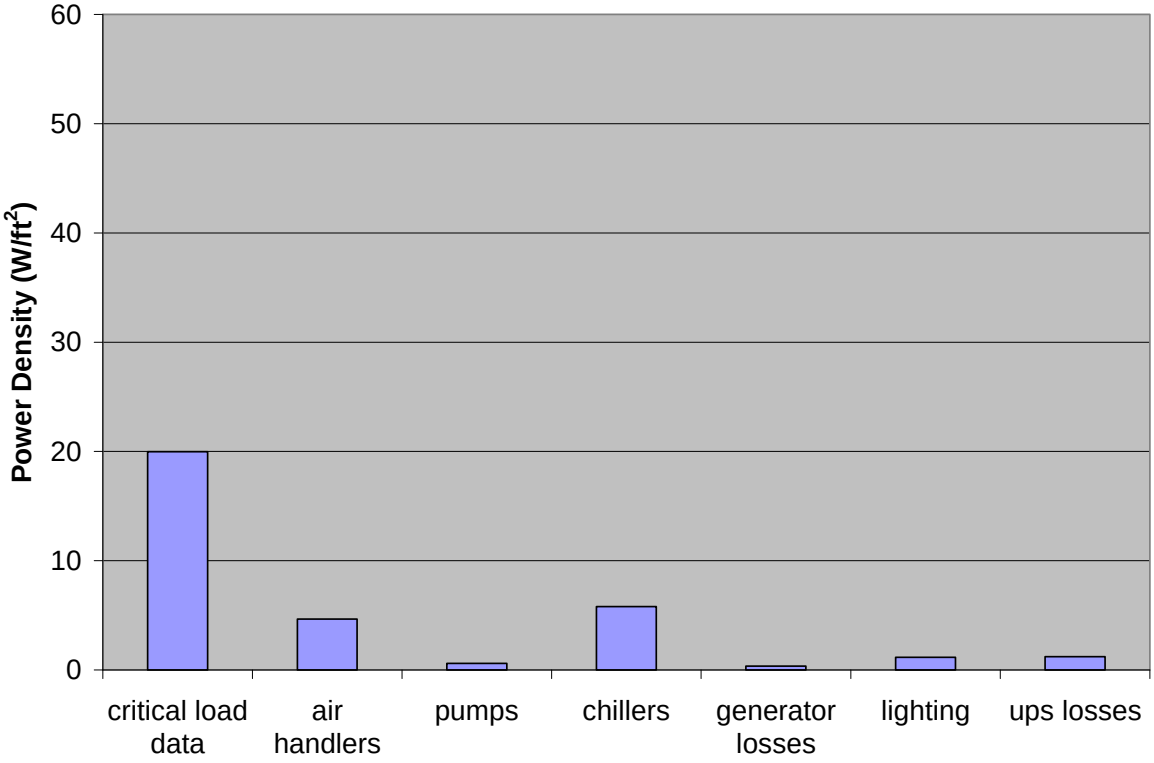

Figure 1. Data Center 18 Power Density

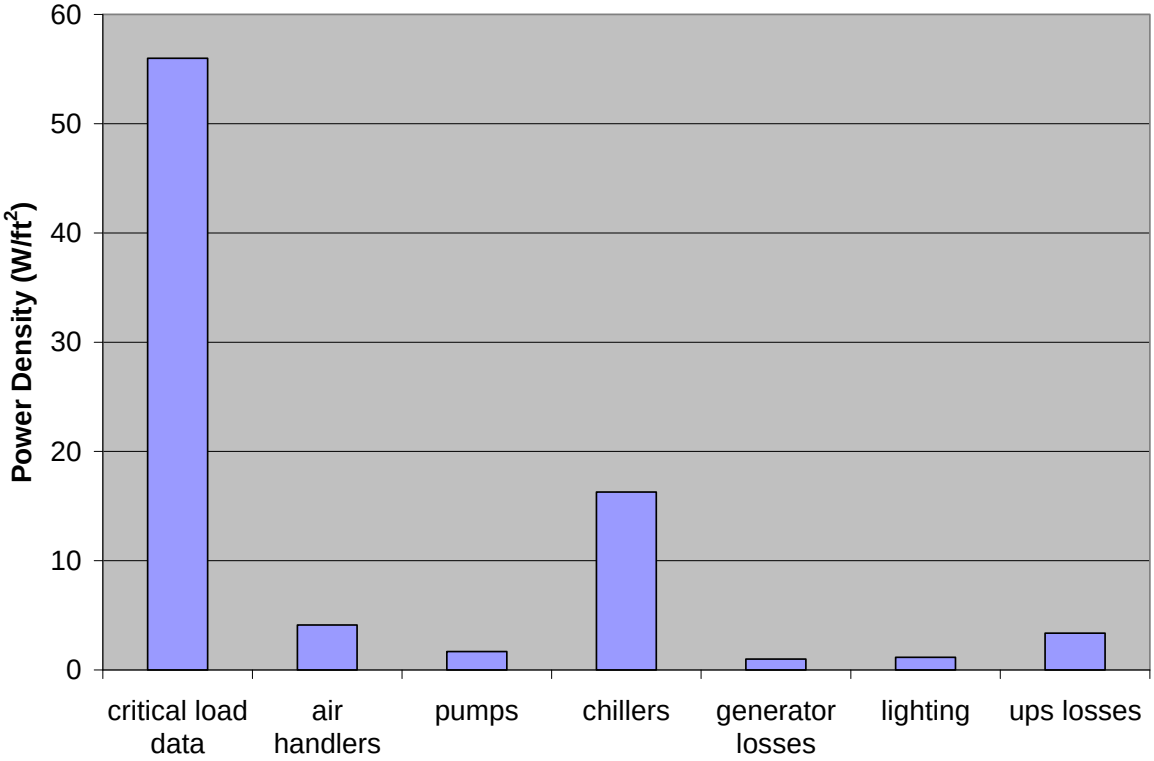

Figure 2. Data Center 19 Power Density

Figures 1&2 show the power density of various components in the facility, including critical power loads, essential mechanical loads, losses from UPS' serving the 6th floor data centers.

3.1 PDU System

Critical electrical power to both data centers on 6th floor was distributed to 12 PDUs fed by UPS 1, UPS 2 and UPS 3.

3.2 Emergency Generators

The three emergency generators had average standby losses of 26 kW during the monitoring period. Emergency generator losses included jacket heat, battery chargers, transformer switches, fuel management system and control.

4 Mechanical System

During the one-week monitoring period, the following HVAC equipment was operating:

- Two Chillers
- Primary chilled water pump P2P, constant speed

- Secondary chilled water pump P2S, on variable speed drive (VSD)
- All ten air handling units (AHUs) with variable speed drives (VSDs): AHU (1-5) west section (DC#18) and AHU (6-10) east section (DC #19).

4.1 Chiller System

Figure 3 shows electric power demand monitored on the two operating chillers for a one week period (October 27 to November 3, 2004). The low chiller power usage around October 28 was the result of alternating operation among the three chillers.

Using the average water temperature rise and the chilled water flow rate, the calculated cooling tonnage was ($Qcooling= \rho GPMC_p \Delta T*60/12000$, in ton). Based on the measured temperatures rise of 7.7°F and averaged water flow rate of 735 GPM, assuming water density ρ of 8.32 lb/gal, the estimated total cooling produced by the chillers was within approximately 235 cooling tons. This was approximately 37 % of the designed cooling capacity of the two chiller at full design load. The portion of actual cooling required for data center rack load was approximately 68% of one chiller.

Chillers' CH-1 and CH-2 power consumption averaged 96 kW and 146 kW respectively during the monitoring period. Therefore, the actual chiller operating efficiency was calculated as approximately 1.0 kW/Ton for the two operating chillers recorded during the monitoring period.

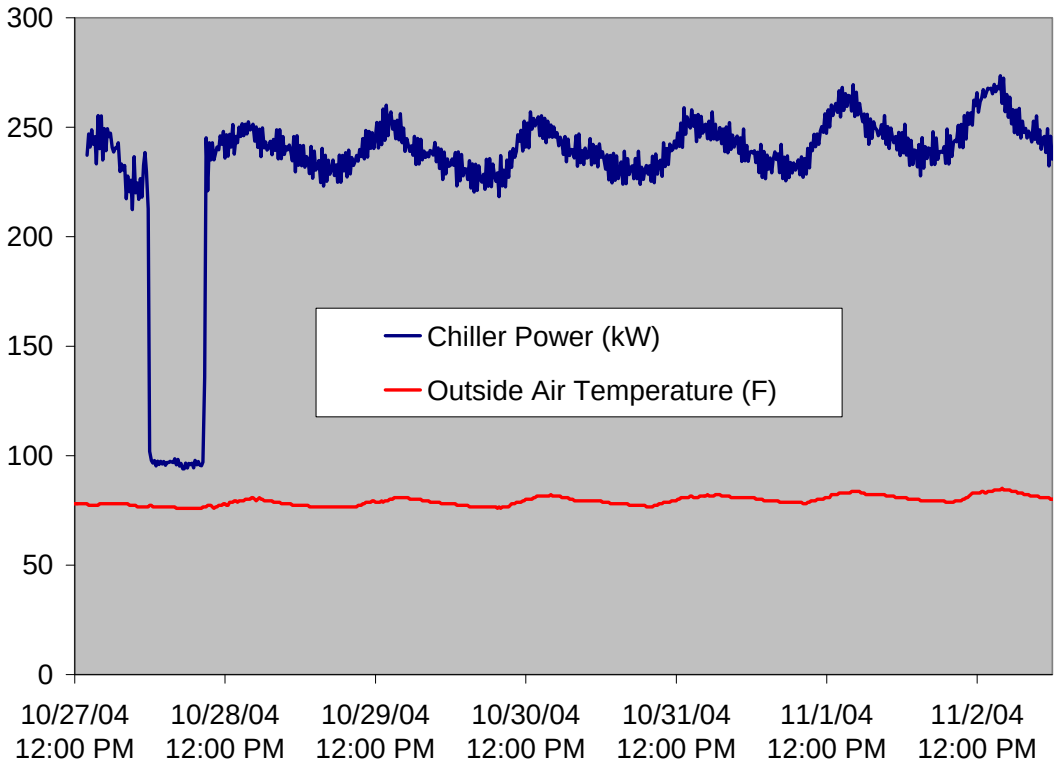

Figure 3. Chiller power demand and outdoor temperature

4.2 Pumping System

Primary chilled water pump P2P and secondary chilled water pump P2S were in operation during the monitoring period. The primary chilled water pump P2P was a constant speed 25-hp pump, and the secondary P2S was a 40-hp pump fitted with a variable speed drive (VSD). The average power consumption for the primary pump (P2P) was 16 kW, while the secondary pump (P2S) was 9 kW. The power consumption for both pumps is shown on Figure 4.

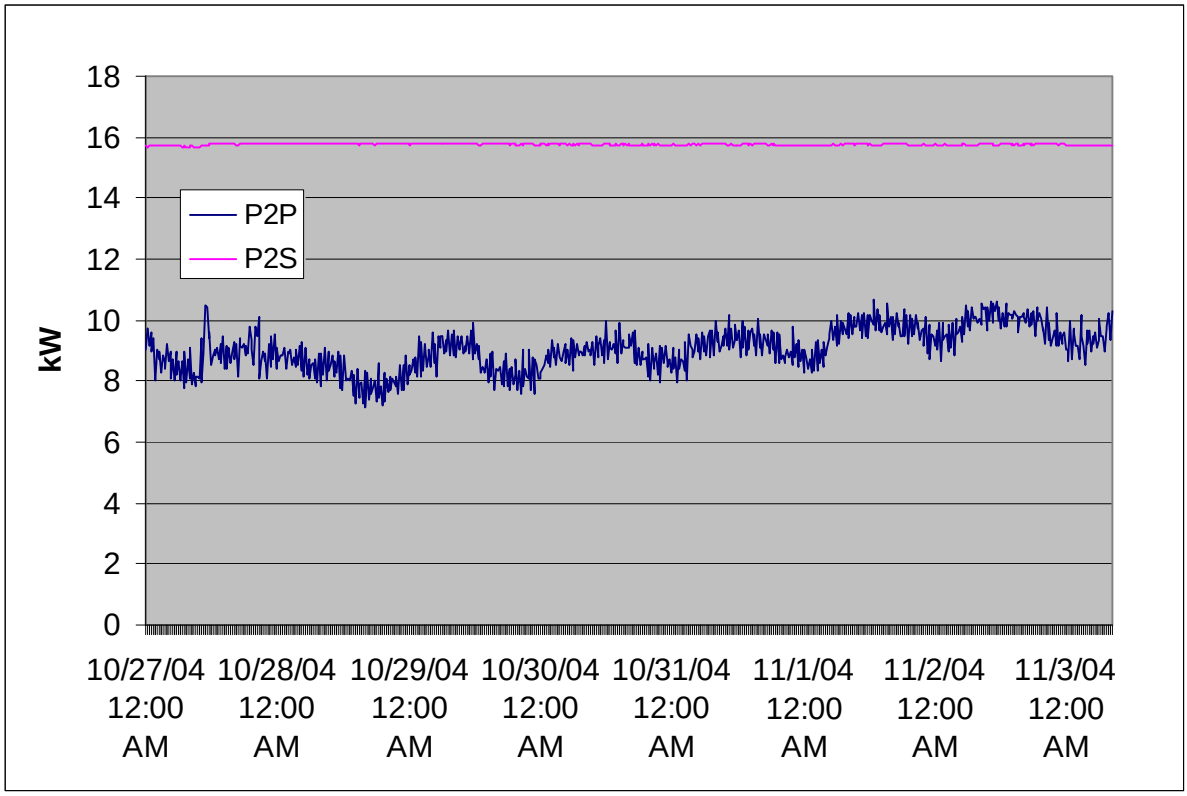

Figure 4. Power Demand for Chilled Water Pumps

4.3 AHU System

The five AHUs (AHU 1-5) serving the west side of the floor (DC #18) were controlled in unison. They were independent from the other five AHUs (AHU 6-10) that were also controlled in unison to serve the east side of the floor (DC #19).

At the time of the survey, EMS printouts of AHU motor operating frequency showed that the west side air handlers were operating at an average of 55 Hz, while the east side air handlers were operating at an average of 44 Hz. AHU 2 served the west side of the 6th floor, and had average power demand of 7 kW with large fluctuations, suggesting a need for tuning the control loop. AHU 10 served the east side of the 6th floor, and had average power demand of 6 kW with little variation.

5 System Operation

5.1 Chilled Water Flow

The total secondary chilled water flow rate was monitored and the results are shown in Figure 5. For the monitoring period, the average flow rate was 735 GPM.

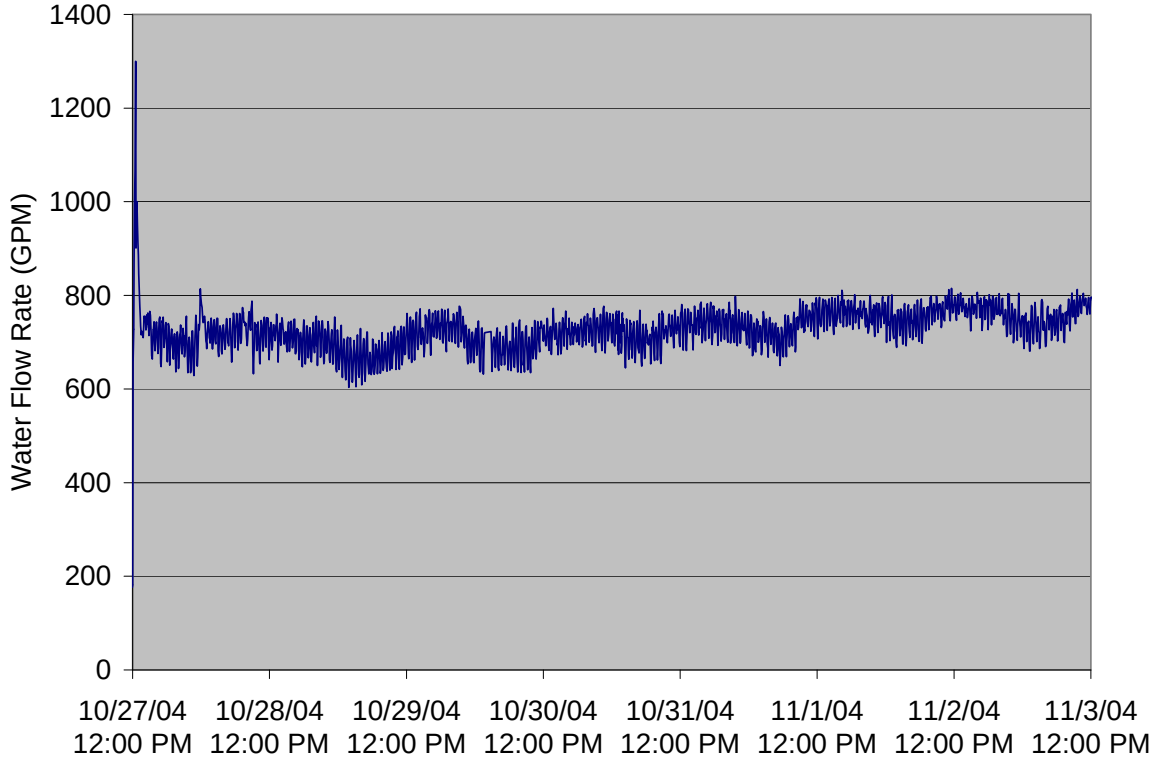

Figure 5 Chilled Water Flow Rate

5.2 Chilled Water Supply and Return Temperatures

The chilled water supply and return temperatures were monitored in the study. For the monitoring period, the average chilled water supply temperature was 44°F, while the average chilled water return temperature was 52°F. The chilled water temperature was controlled by chilled water return temperature set point. This produced an average temperature differential of 7.7°F and an average cooling capacity of 235 tons.

5.3 Air Handler Unit Supply and Return Air

Temperatures and relative humidity for the east and west supply and return air plenums were monitored for a week. In DC #18, average supply and return air temperatures were 68°F and 70°F, respectively. The temperature differential was only 2°F for the period. In DC #19, average supply and return air temperatures were 60°F and 67°F, respectively. The temperature differential was 7°F. The AHU power demand for DC #18 was higher than that for DC #19. This indicates that there was noticeable deficiency in cooling effectiveness induced by operating AHUs for DC #18. Therefore, control of the five AHUs for DC # 18 should be optimized to reduce the power demand for operating the units.

5.4 Generator Jacket Ambient Temperatures

The generator jacket ambient temperatures for Generators 1 and 2 were monitored, as shown in Figure 6. The average ambient temperature was 78°F for generator 1 and 79°F for generator 2.

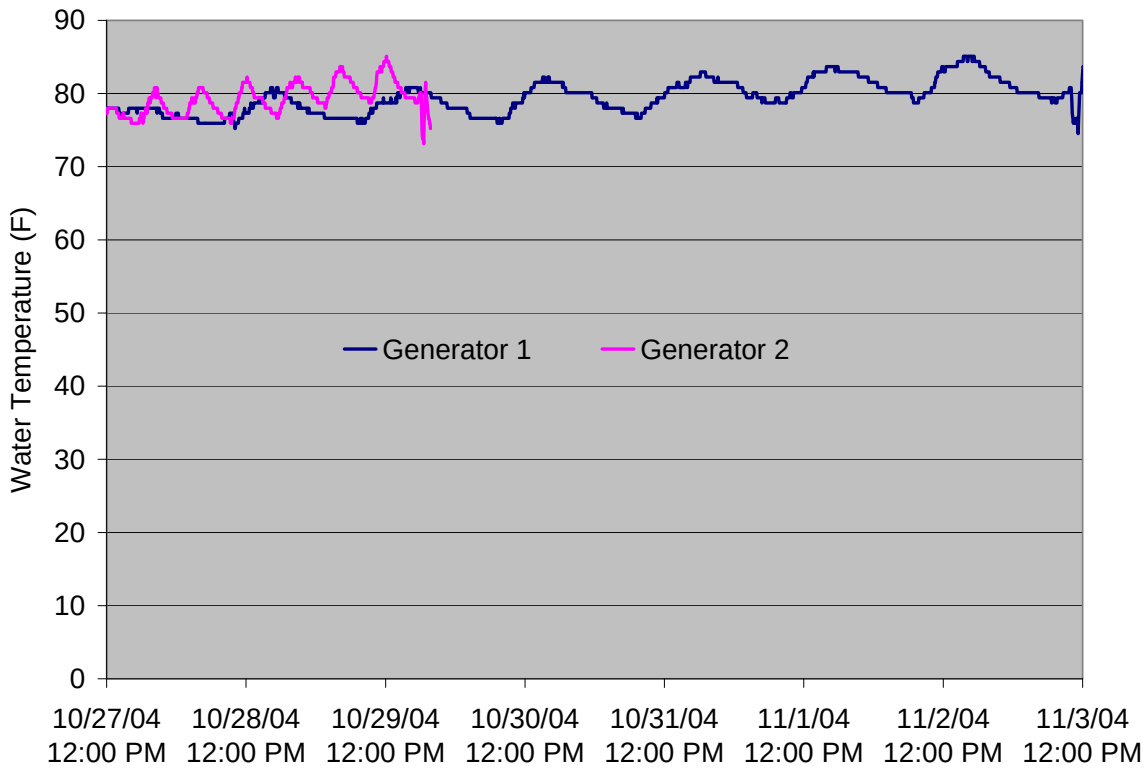

Figure 6 Generator Ambient Temperatures

6 Recommendations

The density of installed computer loads (rack load) in the two data center was 20 and 56 W/ft^2 for DC #18 and DC #19, respectively. The power density of IT equipment in DC #18 was relatively lower compared to other data centers previously studied. In addition, with an HVAC to IT power demand ratio of 0.6 in DC #18 in this study, actual mechanical systems serving the critical load in DC #18 seemed to be oversized and operating less efficiently.

Both chillers were running at a low partial load, making the operation very energy inefficient. Therefore, the operation and control of the chillers and AHUs should be optimized while providing sufficient cooling to the data center. Although arranging hot aisle/cold aisle design to separate airflow streams would be difficult in such a co-location data center, optimizing air distribution should be pursued.

General recommendations for improving overall data center energy efficiency include improving the design, operation, and control of mechanical systems serving the data centers in

actual operation. This includes chiller operation, chilled water system, AHUs, airflow management and control in data centers. Additional specific recommendations or considerations are provided in the following.

6.1 Chilled Water System

Consideration should be given towards resetting the chilled water supply temperature to a higher set point. Lower chilled water supply temperatures may lead to dehumidifying the space air, thus requiring additional re-humidification which would cause energy penalty and yet the existing system does not provide humidification. For example, setting the chilled water supply temperature to 50°F or higher may still provide sufficient sensible cooling in the data center. In the meanwhile, chiller energy consumption would be reduced due to improved thermal efficiency. This measure can be implemented in steps, raising the temperature set point by 2°F at a time, while verifying that no hot spots in critical locations. During these steps, the secondary chilled water pump and air handling unit fan speeds should be monitored, while ensuring that chiller energy savings are not offset by greater energy usage of these mechanical components.

Employing evaporative pre-coolers for the air-cooled chiller condensers may significantly increase chiller efficiency, especially at peak conditions. Resetting secondary chilled water pump speed based on AHU valve positions, keeping one valve 90% open may save energy.

Integrating VSD device and operation in chilled water systems can improve the efficiency. This would be more useful, especially when the future cooling load increases. In addition, optimizing water temperature differential and pump head required would collectively contribute to minimizing total power demand for water systems.

6.2 Air System

Optimize the control of supply and/or return air temperatures and airflow rate from the AHUs, and air distribution.

Optimize air distribution through carefully placing perforated tiles, cable pass-through, and equipment layout. The benefits include achieving greater cooling effectiveness. The temperature difference of supply and return air in DC # 18 was 2°F, indicating a rather low cooling load in this zone.

Recommendation should be given to significantly reduce supply airflow rate of the five air-handling units, while optimizing air distribution within the operating data center space. For example, turn off some AHUs and control airflow rates using VFDs. The possibility of providing more efficient air flow within the tenant spaces, through the use of hot aisle/cold aisle arrangement of computer racks, is limited by the architecture of the tenant cages and the desire to allow these tenants flexibility of use within their own spaces.

6.3 Lighting

The on-site Energy control measures should be considered to reduce lighting load in the data center, this measures should include: Installing lighting zones occupancy sensors; and adding task lighting in appropriate areas and disabling portions of overhead lights.

6.4 Metering and Power Conditioning Equipment

EMCS sensors should be checked and calibrated to provide more accurate readings and monitoring. This includes temperature, pressure, humidity, and power sensors. For example, the reading of power input to PDUs from EMCS was lower than the output power. It's necessary to calibrate the power metering device. In addition, power factor correction device should be provided to improve the accuracy of existing power factor output.

7 Acknowledgements

This report on data center energy benchmarking was finalized based upon the field data collection performed by EYP Mission Critical Facilities and Landsberg Engineering and the draft report produced in the course of performing work subcontracted for the Lawrence Berkeley National Laboratory (LBNL).

The project was funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

8 Appendix A: Data Facility Definitions and Metrics

The following definitions and metrics are used to characterize data centers:

Air Flow Density	The air flow (cfm) in a given area (sf).
Air Handler Efficiency 1	The air flow (cfm) per power used (kW) by the CRAC unit fan.
Air Handler Efficiency 2	The power used (kW), per ton of cooling achieved by the air-handling unit.
Chiller Efficiency	The power used (kW), per ton of cooling produced by the chiller.
Computer Load Density – Rack Footprint	Measured Data Center Server Load in watts (W) divided by the total area that the racks occupy, or the "rack footprint".
Computer Load Density per Rack	Ratio of actual measured Data Center Server Load in watts (W) per rack. This is the average density per rack.
Computer/Server Load Measured Energy Density	Ratio of actual measured Data Center Server Load in watts (W) to the square foot area (sf) of Data Center Floor. Includes vacant space in floor area.
Computer/Server Load Projected Energy Density	Ratio of forecasted Data Center Server Load in watts (W) to the square foot area (sf) of the Data Center Floor if the Data Center Floor were fully occupied. The Data Center Server Load is inflated by the percentage of currently occupied space.
Cooling Load – Tons	A unit used to measure the amount of cooling being done. One ton of cooling is equal to 12,000 British Thermal Units (BTUs) per hour.
Data Center Cooling	Electrical power devoted to cooling equipment for the Data Center Floor space.
Data Center Server/Computer Load	Electrical power devoted to equipment on the Data Center Floor. Typically the power measured upstream of power distribution units or panels. Includes servers, switches, routers, storage equipment, monitors and other equipment.

Data Center Facility	A facility that contains both central communications and equipment, and data storage and processing equipment (servers) associated with a concentration of data cables. Can be used interchangeably with Server Farm Facility.
Data Center Floor/Space	Total footprint area of controlled access space devoted to company/customer equipment. Includes aisle ways, caged space, cooling units electrical panels, fire suppression equipment and other support equipment. Per the Uptime Institute Definitions, this gross floor space is what is typically used by facility engineers in calculating a computer load density (W/sf).
Data Center Occupancy	This is based on a qualitative estimate of how physically loaded the data centers are.
Server Farm Facility	A facility that contains both central communications and equipment, and data storage and processing equipment (servers) associated with a concentration of data cables. Can be used interchangeably with Data Center Facility. Also defined as a common physical space on the Data Center Floor where server equipment is located (i.e. server farm).

9 Appendix B: Facility Diagrams

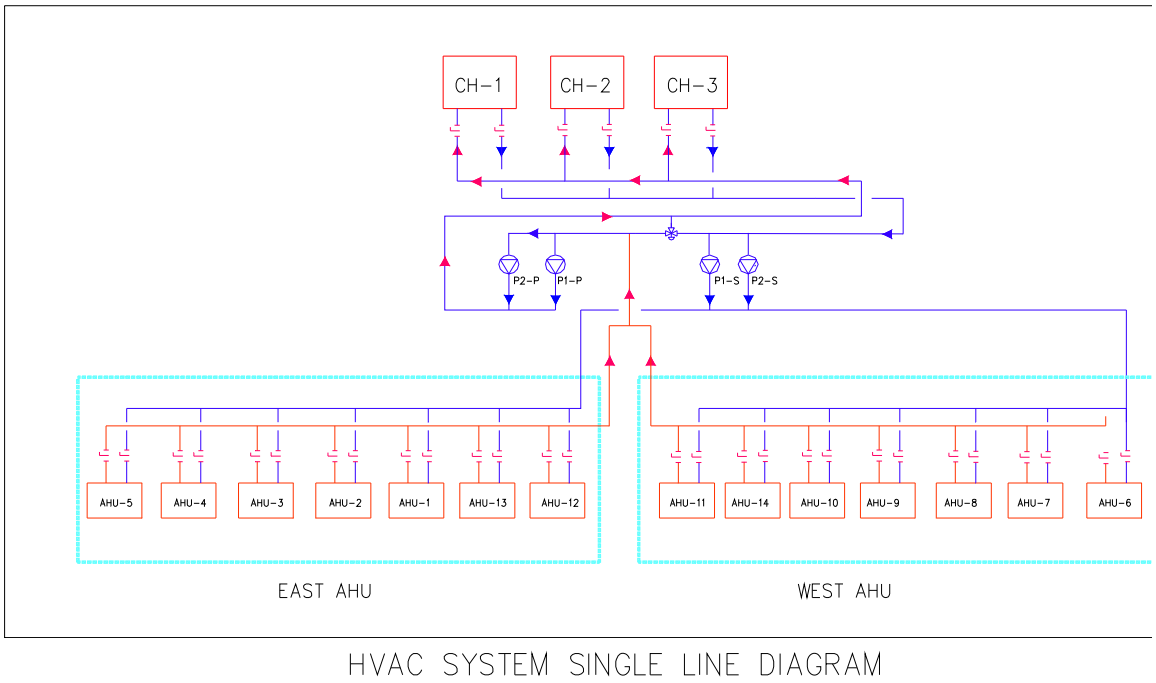

Figure 7 HVAC System

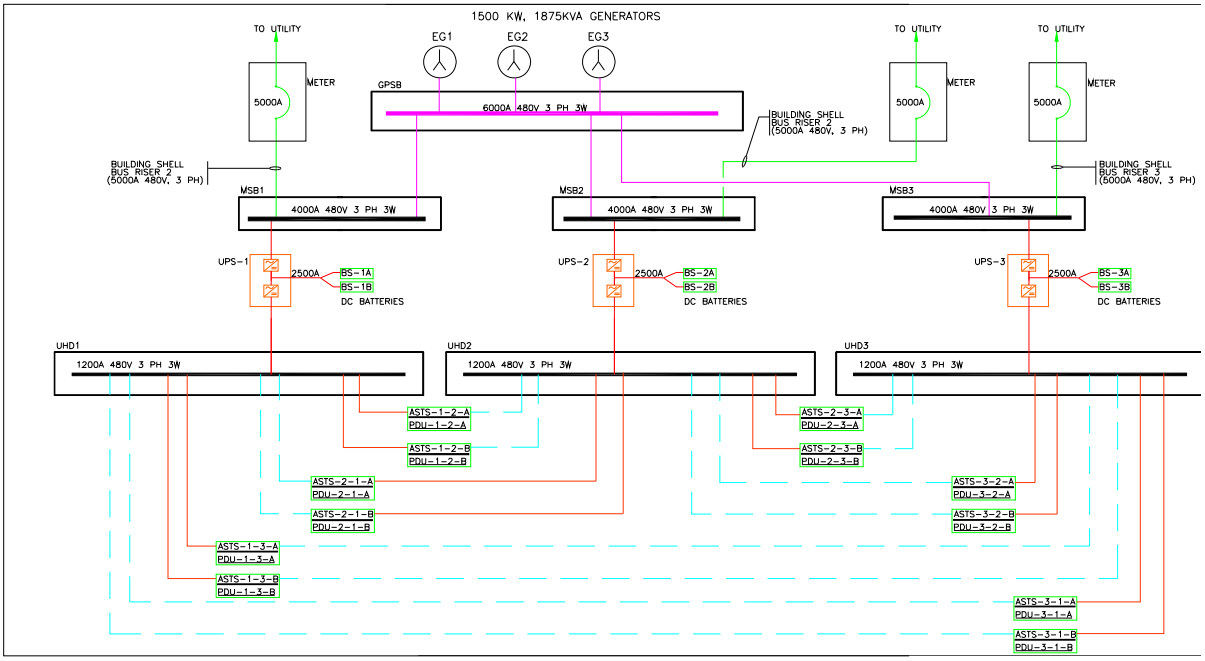

ELECTRICAL SYSTEM SCHEMATIC

Figure 8 Electrical System Schematic

Data Center Energy Benchmarking:

Part 1 - Case Studies on Two Co-location Data Centers (No. 16 and 17)

Final Report

August 2007

Final Report prepared by

Tengfang Xu and Steve Greenberg
Lawrence Berkeley National Laboratory (LBNL)
Berkeley CA 94720

Draft provided by

EYP Mission Critical Facilities
Los Angeles, CA 90064
and
Landsberg Engineering, P.C.
Clifton Park, NY 12065

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Table of Contents

1 EXECUTIVE SUMMARY 1

2 REVIEW OF SITE CHARACTERISTICS 2

2.1 ELECTRICAL EQUIPMENT AND BACKUP POWER SYSTEM 2

2.2 MECHANICAL SYSTEM 3

3 ELECTRIC POWER CONSUMPTION CHARACTERISTICS 8

4 MECHANICAL SYSTEM OPERATION 13

5 RECOMMENDATIONS 19

6 ACKNOWLEDGEMENTS 21

7 APPENDIX A: DATA FACILITY DEFINITIONS AND METRICS 22

8 APPENDIX B: FACILITY DIAGRAMS 24

1 Executive Summary

Data Centers # 16 and #17 were located in a four-story building in San Francisco, California. The data center building had a total floor area of approximately 97,900 ft^2 with 2-foot raised-floors in the data services area. Two out of eight data centers in the building were occupied by computers and equipment, and were in operation at the time of the study conducted between October 15 and October 22, 2004.

Electric power for both data centers was supplied through HITEC power conditioning units without any battery system. Cooling for both data centers was through multiple water-cooled computer room air conditioning (CRAC) units connected to heat exchangers served by cooling towers. Of the total electric demand for both data centers, 40% went to critical computer and equipment load, 32% went to mechanical systems, 9% to UPS losses, and the remaining 19% to miscellaneous systems.

General recommendations for improving overall energy efficiency of the data centers included improving the design, operation, and control of mechanical systems serving the data centers in actual operation. This included primary condenser water system, secondary condenser water system, CRAC units, and airflow management and control in data centers. Specific recommendations for options of improving energy efficiency of the data centers are developed and provided in this report.

- A significant number of CRAC units could be turn off while the rest of the units would be able to provide sufficient cooling to the critical cooling requirements for the data centers.

- Optimize the actual air temperature and humidity set points, e.g., extending the permitted range.

- Optimize the control of supply and/or return air temperature from the CRAC units.

- Optimize air distribution through carefully placing perforated tiles, cable pass-through, equipment layout, and actual operation or non-operation of CRAC units.

- Evaluate and calibrate the monitoring system including the power metering system and secondary water system, e.g., data acquisition and sensing through the EMCS systems.

- Optimize secondary condenser water supply and supply temperatures, e.g., adjusting the control set points, using variable speed drives on the secondary condenser water pumps.

2 Review of Site Characteristics

Data Centers # 16 and #17 were located in a four-story building in San Francisco, California. The data center facility had a total floor area of approximately 97,900 ft^2 with 2-foot raised-floors in the data services area. The data centers were designed to provide co-location data services in areas that are environmentally controlled and monitored. The building includes eight separate data center rooms with raised floors, with two data centers on each of four floors.

The co-location area was designed to house and operate eight data centers simultaneously. At the time of the study, only two out of the eight data centers, both located on the fourth floor - Data Center #16 (Room # 7) and Data Center #17(Room #8)) were occupied and in operation. Each of the data centers housed approximately 10,000-ft^2 floor area for networking equipment.

Both data centers operated 24 hours per day all year-round. The users of the data centers had 24-hour full access to and from their caged spaces. The building has 16,000 ft^2 office space, or about 16% of the total floor area. The equipment in this facility was only two years old at the time of study and the existing building cooling load was primarily limited to the fourth floor data centers.

2.1 Electrical Equipment and Backup Power System

Data Centers #16 and #17 have a main service drop from the electric utility of 1200 A at 38 kV, as shown in Appendix B. This service is stepped down to 480V and distributed to five main switchboards and one reserve switchboard (SWBD).

At the time of the study, only one SWBD-M4 was active due to the limited requirement for load level. The normal real power demand monitored by the site Automatic Logic Corporation (ALC) system was 1,360 kW.

Before the monitoring starting on October 15, 2004, the peak power demand for both centers was observed and recorded as 3,120 kW on October 14, 2004. The reason for observed spike in power demand on October 14 was unknown but may be linked to possible temporary operational testing requirements for the building or simply a reading error.

The electric current to run SWBD-M4 was distributed into two paths: 1) essential power for HVAC and lighting loads, and 2) critical power for computer and equipment loads.

Essential power was fed directly to two SWBDs - D7 and D8, and was then distributed to their loads.

Critical power passed through two HITEC synchronous generators (Model # D85051, 2840 kVA rating). The facility utilized eight HITEC synchronous generators as a uninterruptible power supply (UPS) system without any storage battery. Only two HITECH synchronous generators were in operation at the time of the study. As shown in Figure 1, these UPS units maintained critical power supply for computer loads in the event of electric power outage or disturbance. The UPS power conditioning was a no-break system capable of providing a

regulated output while using the utility supply. In the event of power loss, UPS' diesel-powered generators would start and supply electric power to the essential and critical power paths (D7 and D8). While the generator's engine is starting, the unit's flywheel drives the generator to maintain current for the critical loads.

Figure 1. Typical HITEC Diesel UPS/Backup Generator.

2.2 Mechanical System

2.2.1 Cooling Tower

Data Centers # 16 and #17 were cooled by water-cooled computer room air conditioning (CRAC) units that were supplied by a decoupled condenser water system. Appendix B includes a condenser water flow diagram for the system. The office space was conditioned by water-source heat pumps, which rejected heat into the same condenser water system.

Heat rejection for these centers was designed to be provided by six cooling towers, each with a cooling capacity of 800-ton. The Baltimore Air Coil Series V open circuit towers were designed with 91°F entering water temperature and 76°F leaving water temperature with 72 °F wet bulb temperature. The design water flow rate for each cooling tower was 1,280 gallon per minute (GPM). The cooling towers were labeled as CT5-1 to CT5-6 (Figure 2). They were located at the penthouse section of the building. Each cooling tower has two belt-driven centrifugal fans, arranged in a blow-through configuration and driven by 30-hp motors. The fan motor speeds were controlled by variable frequency drives to maintain condenser water supply at a certain set point temperature. Each cooling tower could be isolated from the

condenser water system by automatic butterfly valves. These valves and their associated towers are controlled through the ALC system by site personnel.

Figure 2. Cooling towers

2.2.2 Primary Condenser Water Pumps

Primary condenser water (in the cooling tower loop) was circulated by six centrifugal vertical in-line Armstrong pumps (P4-1, P4-2, P4-3, P4-4, P4-5 and P4-6) arranged in parallel (Figure 3). Each pump has a motor capacity of 50 HP and a design volume flow rate of 1,280 GPM. These pumps are energized by site personnel, and run at constant speed. Only one primary pump was running at the time of this study.

Figure 3. Primary condenser water pumps.

2.2.3 Heat Exchangers

Condenser water was circulated through up to three Alfa Laval model #M30-FG plate/frame heat exchangers labeled HX4-1 to HX4-3 (Figure 4), depending on load. Each Alfa Laval heat exchanger had a rated capacity of 24,000,000 BTU per hour of total heat rejection of 2,000 tons (actual net cooling capacity of 1,600 tons plus the rejection of compressor heat). One of the three heat exchangers, HX4-3, was in operation at the time of the study and dissipated heat from both data centers on the fourth floor.

The primary, cold-side water temperature supply and return from the cooling tower showed a temperature increase (ΔT) of 3.5-4.8°F for primary condenser water across the heat exchanger, compared to the design temperature increase (ΔT) of 15°F. In addition, for the secondary condenser water across the heat exchanger serving the building cooling load system, the monitored temperatures of water supply and return from the cooling tower showed a very small temperature increase (ΔT) of less than 1ºF.

The actual approach temperature, defined as the temperature difference between the secondary water supply and the primary condenser water supply temperature, was mostly within 5ºF. Each heat exchanger was isolated from both the primary and secondary condenser water systems by automatic valves that were controlled by site personnel through the ALC control system.

Figure 4. Heat Exchanger for condenser water.

2.2.4 Secondary Condenser Water Pumps

The secondary condenser water pumps consisted of seven parallel in-line Armstrong centrifugal pumps with a rated capacity of 50 HP. The pumps were identified as P4-7 to P4-13 (Figure 5).

P4-12 was the only pump in active operation at the time of the study. The pressure differential between the pump's suction and discharge was 28 psi, delivering water at the flow rate of 1,680 GPM between the secondary, hot side of heat exchanger and the building cooling units. The pump's operation was set to be either on or off, and was manually controlled through the facility's control system.

Figure 5. Secondary condenser water pumps.

2.2.5 CRAC Units

The Data Center #16 (Room # 7) and #17 (Room # 8) on the fourth floor have 2-ft raised floors through which cold air is supplied and circulated via packaged CRAC units: Data-Aire Model# DAWD-26-34.

Each of the 21 CRAC units was designed to deliver 10,000 CFM conditioned air, with 11 CRAC units serving Data Center #16 and 10 CRAC units serving Data Center #17. Figure 6 shows typical CRAC units in the data center.

Figure 6. Sample CRAC Unit

The CRAC units' supply fans were all operating at their constant speeds. These units are located along the north and south walls of the data center rooms. Each CRAC unit had a net sensible cooling capacity of 20 tons excluding fan heat. The CRAC units had 4-inch throwaway pleated air filters rated at 30% efficiency, and two water-cooled refrigeration systems, consisting of compressors, water-cooled condensers, and controls. No reheat coils or humidifiers were present in these CRACs. Each CRAC unit's internal controls were used to maintain temperature and relative humidity within a range. The unit's control set point for air temperature was 71°F, at the units' return air intake. No explicit humidity control was performed by these units, and there was no remote monitoring of them. The CRAC units delivered conditioned air to the raised floor plenum and returned warm air from upper spaces in the data centers. At the time of the study, all of the CRAC units operated continuously in both data centers. Among all CRAC units, two CRAC units were monitored in Data Center # 16, and one in Data Center # 17.

Conditioned outside air for the data center was provided by two Mammoth model #VCX-252-GXS water-cooled heat pumps, each rated for 8000 cfm. This air was distributed to each data center room through supply ducts. Humidification for this outside air is provided by two operating Nortec model #NH-150 electric steam generators, each of which was rated at 150 pounds per hour. There were two standby humidifiers. Steam generated by these humidifiers was injected into the make-up air supplied to units' supply ducts.

3 Electric Power Consumption Characteristics

The end-use breakdown for both data centers' electric power demand is shown in Table 1.

Table 1. End-Use of Electricity of the Data Centers (16&17 combined)

Description	Electric power demand (kW)	Share of electric energy use (%)	Floor Space (ft^2)	Electric power density (W/ft^2)
Rack Load (Data Center 16)	377	28%	10,000	37.7
Mech Essential Load* (Data Center 16)	300	22%	10,000	30
Rack Load (Data Center 17)	165	12%	10,000	16.5
Mech Essential Load* (Data Center 17)	132	10%	10,000	13.2
Power Losses to UPS'	123	9%	20,000	6.2
Subtotal Loads (Rack, Essential, and Losses Loads)	1097	81%	20,000	54.9
Other	263	19%	97,878	2.7
Overall Building Load	1360	100%	97,878	13.9

*Mechanical essential loads include all HVAC equipment, including CRACs, cooling towers, and condenser water pumps.

A total facility electrical load of 1,360 kW was recorded from building instruments. The power supply to fourth-floor Switchboard M4, including essential and critical loads, was recorded with building instruments (Square D Power Logic). The electrical losses in the UPS units were calculated by subtracting the essential and critical loads of both data centers from the total power supply to Switchboard M4.

Both data centers on the fourth floor housed a total of 502 computer racks, and with an average power demand of 0.75 kW per rack. The highest rack power demand was reported to be 4 kW. In DC #17, the critical equipment was located in just one half of the space while all of the ten CRAC units were in operation. Consideration should be given to turning off CRACs in unoccupied areas of the floor, and blocking off perforated floor tiles in this area.

From these measurements, it was observed that 40% of the overall electric power was consumed by fourth floor critical loads in both data centers, 32% of the power was consumed by HVAC systems, and 9% of the power was consumed by UPS units, and the remaining 19% was created by lighting, office, and miscellaneous loads in the building.

Power demand breakdown for each data center is shown in Table 2 and Table 3. The ratios of HVAC to IT power demand in each of the data centers in this study were approximately 0.8. The density of installed computer loads (rack load) in DC#16 and DC#17 was 38 W/ft^2 and 16 W/ft^2, respectively. This was relatively lower compared to other data centers previously studied. In addition, the actual mechanical infrastructure in place to serve the critical loads seemed to be relatively high.

Table 2 End-Use of Electricity of Data Center 16

Description	Electric power demand (kW)	Share of electric energy use (%)	Floor Space (ft^2)	Electric power density (W/ft^2)
Total Load-DC 16	763	100%	10000	76.3
Rack Load	377	49%	10,000	37.7
Mech Essential Load	300	40%	10,000	30
UPS	86	11%	10,000	8.6

Table 3 End-Use of Electricity of Data Center 17

Description	Electric power demand (kW)	Share of electric energy use (%)	Floor Space (ft^2)	Electric power density (W/ft^2)
Total Load DC 17	334	100%	10,000	33.4
Rack Load	165	49%	10,000	16.5
Mech Essential Load	132	40%	10,000	13.2
UPS	37	11%	10,000	3.7

An estimate of "rack-cooling load" may be calculated based upon the data center critical power load, assuming 100% of the critical power becomes heat to be rejected by cooling. For example, Q = kW * 3413 / 12000 (ton). Using the critical power of 377 kW and 165 kW in each data center, the rack-cooling loads of the data centers would be approximately 110 ton and 45 ton, respectively. This indicates that for both data centers, a significant number of CRAC units could be turn off while the rest of the units would be able to provide sufficient cooling to the critical cooling requirements.

Figures 7&8 show the power density of critical power loads, essential mechanical loads, losses from UPS' serving the fourth floor data centers. The power density was presented in terms of Watts per square foot of raised-floor.

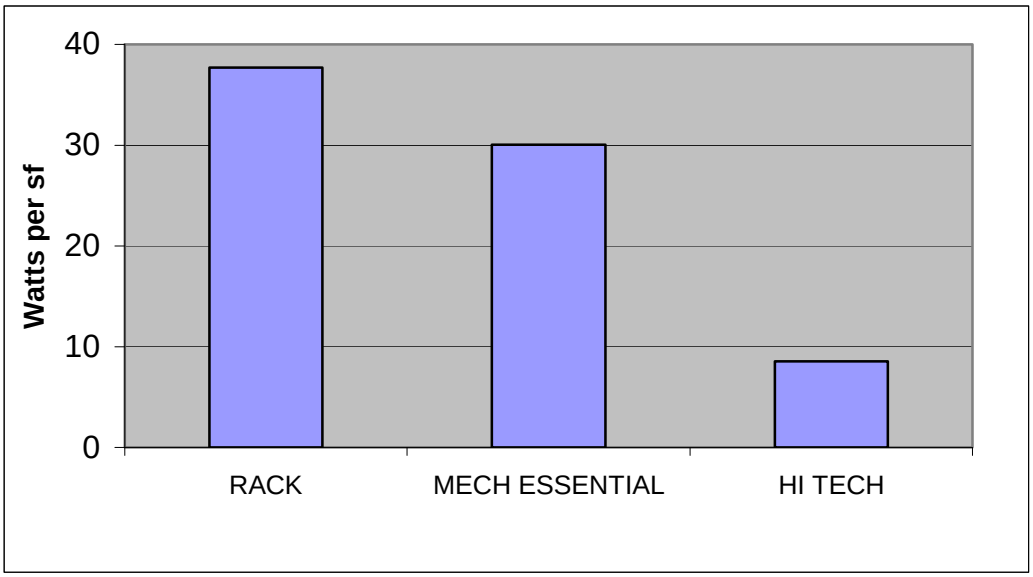

Figure 7. Data Center 16 Power Density

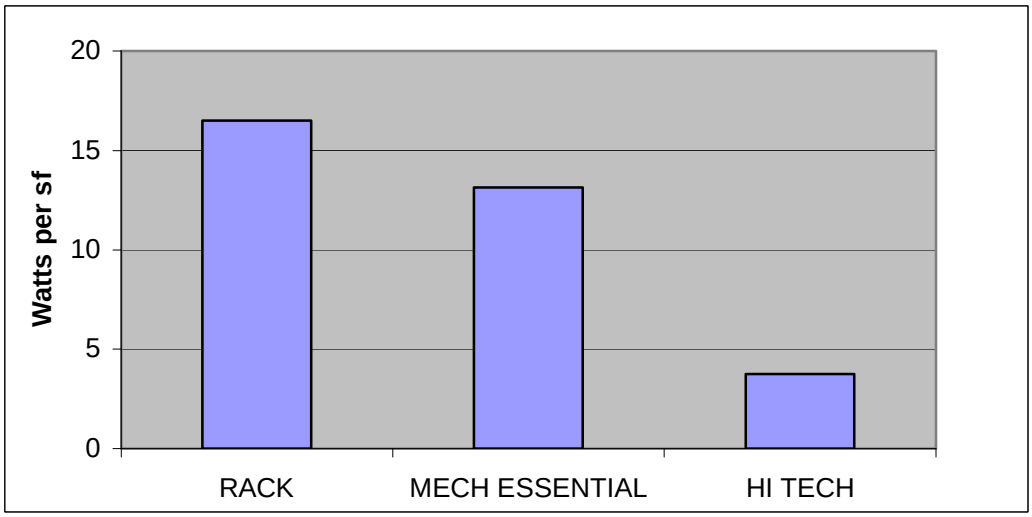

Figure 8. Data Center 17 Power Density

Critical electric power supplied to each data center was through six power distribution units (PDUs) located within each of the data centers: PDU-7A thru PDU-7F in DC16 and PDU-8A thru PDU-8F in DC 17. All of the 200 kVA-rated PDUs had Level 3 Model # RPC-1C-200-BD at 480/208 volts.

Figure 9 shows one of these power distribution units. Typically, the PDUs were recorded to be 96%.

Figure 9. Typical PDU

4 Mechanical System Operation

During the one-week monitoring period, the following HVAC equipment was operating:

- Primary condenser water pump (constant speed)
- Secondary condenser water pump (constant speed)
- Cooling tower (s), with fans on variable speed drives (VFDs). Only one tower, with two fans, was operating at a time to serve the loads
- Plate/Frame heat exchanger
- All CRAC units in both data centers (DC #16 and DC #17)

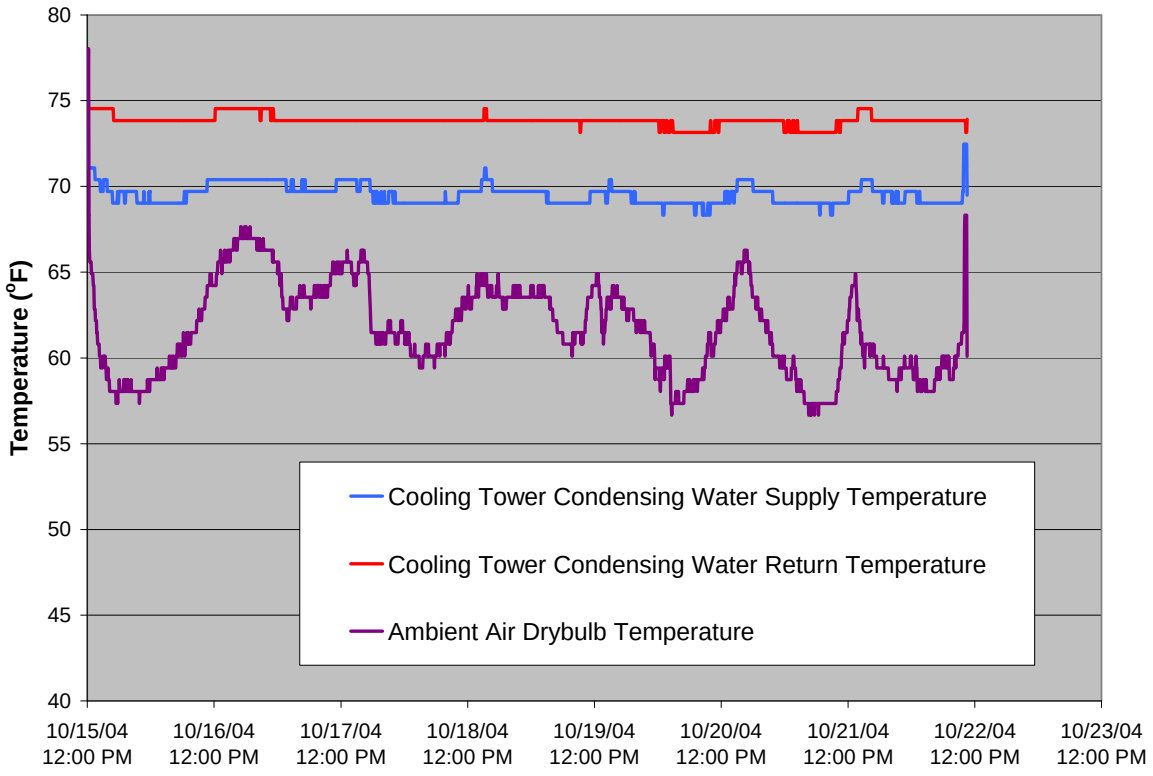

Figure 10. Cooling Tower Condenser Water Temperatures

Figure 10 shows the primary condenser water temperatures and outside air temperature monitored and recorded for a period of one week. The primary (cooling tower) condenser

water supply and return temperature exhibited typical temperature differential ranging from 3.5 to 4.8°F.

The cooling tower fan speeds were controlled by the primary condenser supply water temperature through a variable speed drive. The exact temperature set point was not known. The VFD was operating at within a range around approximately 40 Hz.

Using the average water temperature rise and the primary pump water flow rate, the calculated cooling tonnage can be calculated by the following equation:

$$Tonnage = \frac{60 \rho Q C_p \Delta T}{12000}$$

Where

P: water density in lb/gal, 8.32 lb/gal

Q: water flow rate in gallon per minute, 1280gpm

ΔT: water temperatures rise in °F

Cp: water thermal conductance, 1BTU/lb°F.

With the ΔT ranging from 3.5 to 4.8°F, the estimated total cooling produced by the cooling tower was within approximately 190-260 cooling tons. This was approximately one quarter to one-third of the designed cooling capacity of a cooling tower at the design water flow rate.

Figure 11 shows the recorded water temperatures for secondary (building) condenser water system along with outside air temperature during the monitoring period. Little difference was observed in the recorded supply and return water temperatures, i.e., both near 74°F with the difference mostly within 1°F. Given certain heat transfer efficiency, the heat transfer at both sides of the heat exchangers must be balanced. Apparently, there were errors in temperature sensors or EMCS system monitoring signals concerning the secondary water temperatures. Therefore, we suggest that the monitoring system be examined and calibrated, e.g., data acquisition through the EMCS systems.

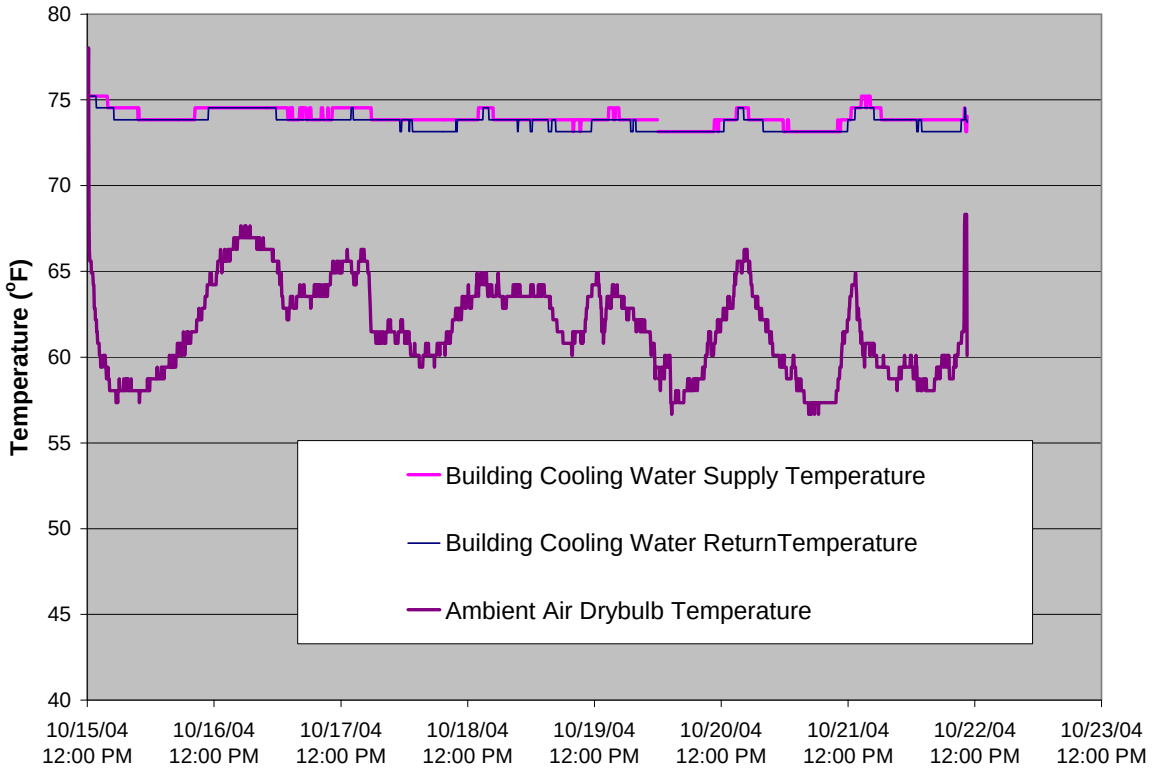

Figure 11. Secondary Condenser Water Temperatures

Air temperature monitoring for three selected sample CRAC units was taken for a period of one week (October 15 to October 22). Room air temperature and relative humidity were measured in the center of the data centers at a height of six feet above the raised floor.

Figure 12 shows supply and return air temperatures for one of the CRAC units, along with space air temperature and relative humidity in Data Center #16 (CRAC 7-3). During the monitoring period (October 15 to October 22), the return air temperatures were constant when the HVAC systems were in normal operation. When the supply air temperature fluctuated, the return air temperature also fluctuated although within a smaller range. In the meanwhile, the data center room RH also changed significantly. When supply air temperature was maintained at a more constant range, the return air temperature to the unit and the room RH became less fluctuated. Most of the time, RH was within 50-60% range. This indicates that temperature control of supply and return air to the individual CRAC unit was significant in maintaining the stability of room air temperature and relative humidity.

In addition, the temperature of return air to the CRAC unit was consistently lower that the space air temperature by approximately 5-6°F. This exhibits large difference between room temperature and return air temperature, perhaps partly due to a “short-circuit” of cold and hot air surrounding this CRAC unit. This indicates that there is noticeable deficiency in cooling

effectiveness induced by operating this CRAC unit. Therefore, the air management of this CRAC unit and perhaps others should be optimized to reduce the waste of cooling provided by the units.

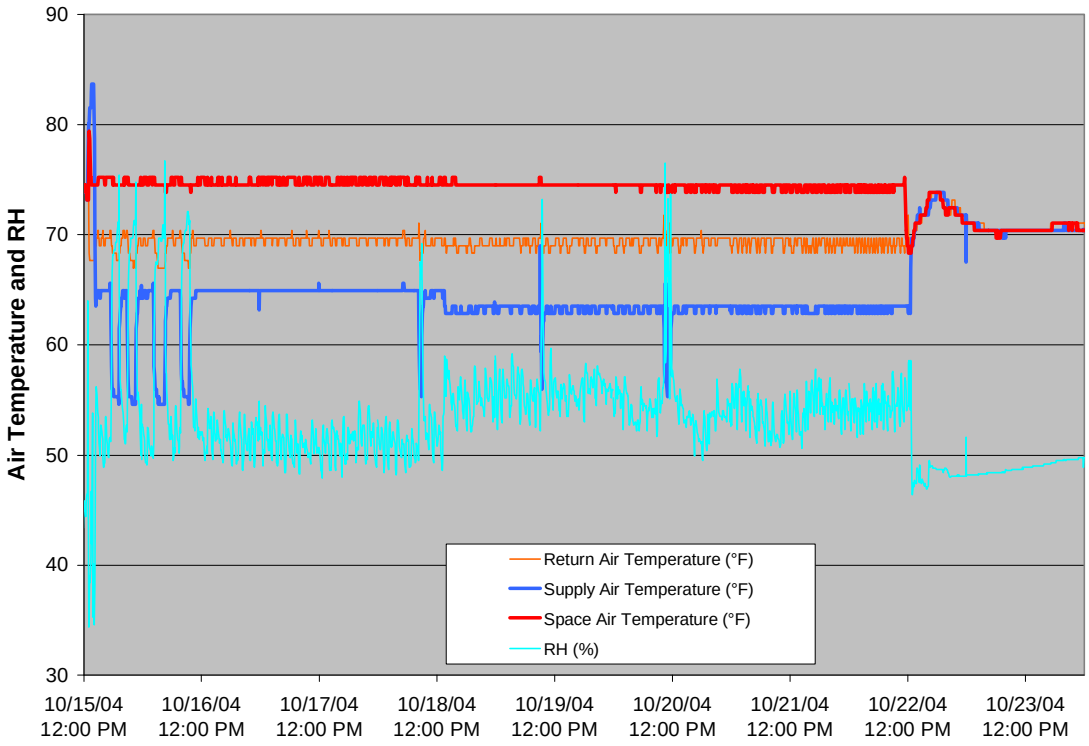

Figure 12. DC #16 CRAC 7-3 Air Temperature and Humidity

Figure 13 shows another CRAC unit's supply air and return air temperatures and the room air temperature in Data Center #16 (CRAC unit 7-12). During the monitoring period (October 15 to October 22 noontime), the return air temperatures were constant when the HVAC systems were in normal operation. Different from CRAC unit 7-3 in the same data center, the temperature of return air to CRAC unit 7-12 was consistently closer to the room air temperature, i.e., mostly within 1-2°F. This suggests that cooling induced by this CRAC unit was more effective in removing heat than was CRAC unit 7-3 in the same data center. However, the large temperature differential between supply and return air temperatures may indicate 1) that the supply air temperature could be elevated to improve heat-exchanging efficiency, and 2) that operation and layout of this CRAC could be further improved to avoid overcool or short-circuiting cold air with return, warmer air.

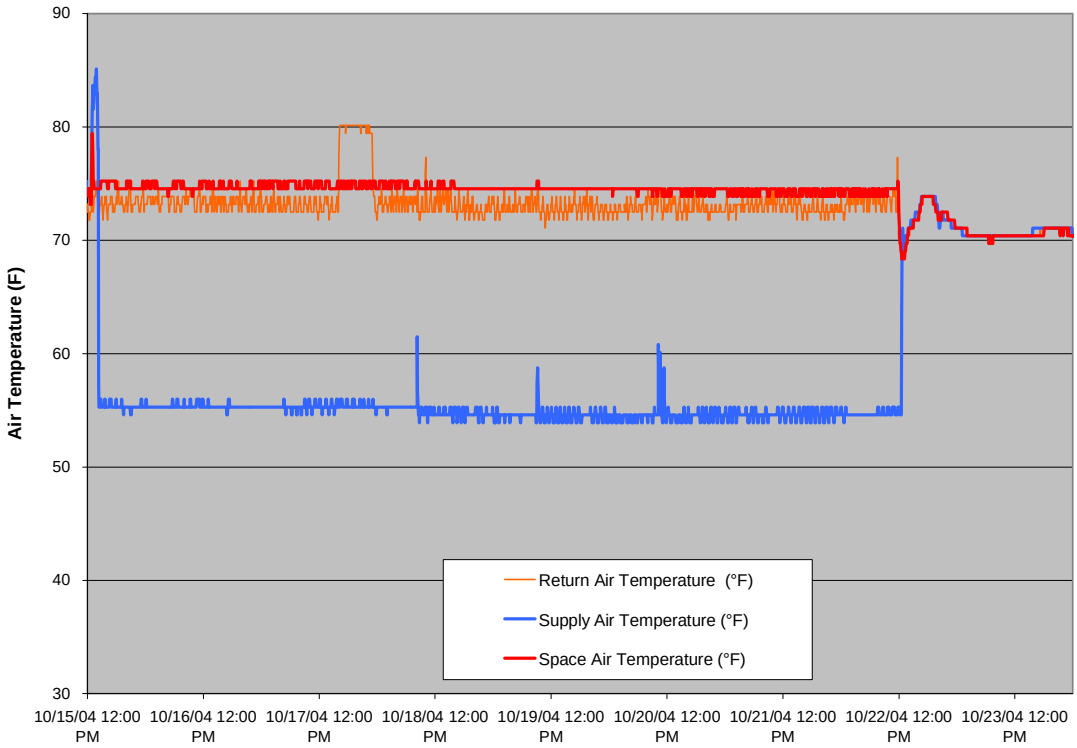

Figure 13. DC #16 CRAC 7-12 Air Temperature

Figure 14 shows the trending of an additional CRAC unit's supply air and return air temperatures and the room ambient temperature in data center #17 (CRAC unit 8-3). Similar to the CRAC unit 7-12 in DC #16, during the monitoring period (October 15 to October 22 noontime), the temperature of return air to CRAC unit 8-3 in this data center was consistently close to the space air temperature, i.e., within 2-3°F, while the supply air temperature was around 55°F. This suggests that while the cooling induced by this CRAC unit was effective, the large temperature differential between supply and return air temperature may indicate that the supply air temperature could be elevated and that operation and layout of this CRAC unit could as well be improved.

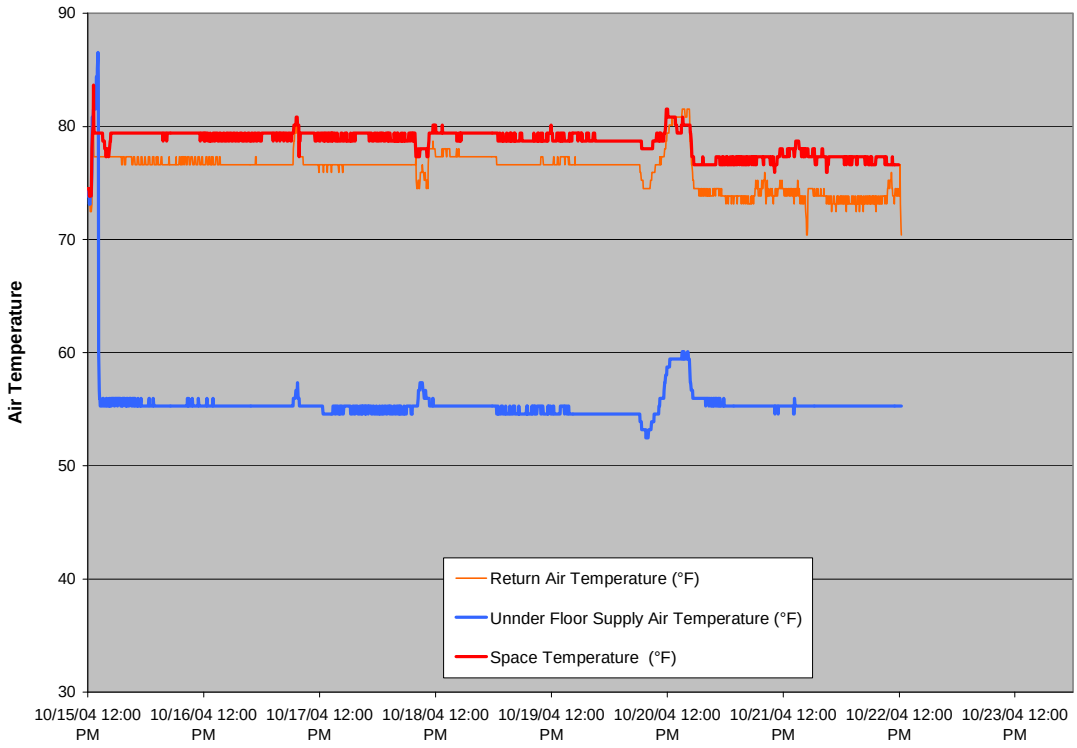

Figure 14. DC# 17 CRAC 8-3 Air Temperatures

In summary, the temperature difference between the data center air and the return air to CRAC unit was found to be significant in one out of three CRAC units selected in the study. This suggests perhaps short-circuiting or mixing of the cold supply air into the return air to the CRAC unit(s). Although arranging hot aisle/cold aisle design to separate airflow streams would be difficult in such a co-location data center, optimizing air distribution should be pursued and would be possible through carefully placing perforated tiles, cable pass-through, and CRAC units. The benefits include achieving greater CRAC effectiveness, in the meanwhile perhaps less humidification and cooling would be required.

Figure 15 shows the cooling tower supply and return temperatures and the fan VFD speed, for a 24-hour period. A slight diurnal variation can be seen in the action of the VFD. The drive was also operating at frequencies varying between 20 HZ and 60 HZ at the beginning of the period, which may indicate a control problem that needed tuning.

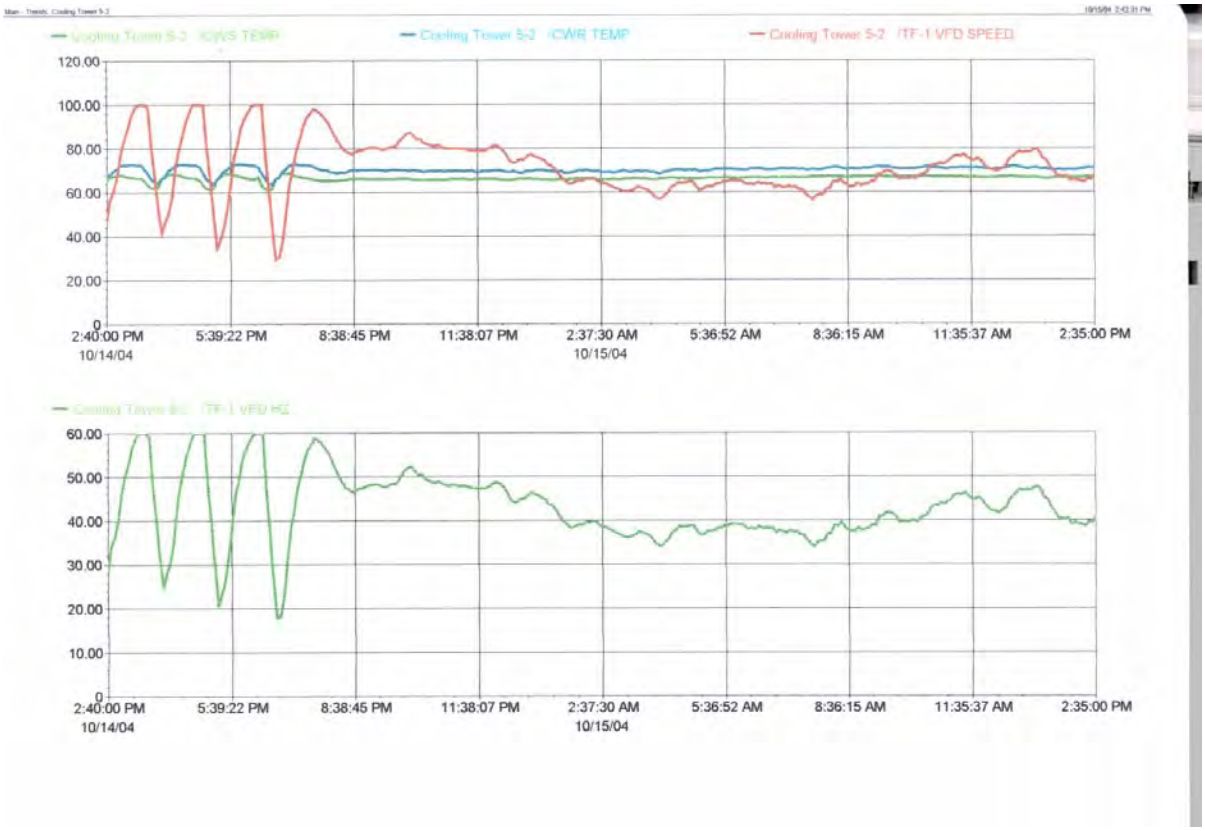

Figure 15. Cooling Tower Supply and Return Temperatures and Fans VFD (Hz)

5 Recommendations

The density of installed computer loads (rack load) was 38 W/ft^2 in data center 16 and 16 W/ft^2 in data center 17, respectively. This was relatively lower compared to other data centers previously studied. In addition, the actual mechanical infrastructure serving the critical loads seemed to be relatively high, with an HVAC to IT power demand ratio of 0.8 in each of the data centers in this study. A significant number of CRAC units could be turn off while the rest of the units would be able to provide sufficient cooling to the critical cooling requirements.

In addition, general recommendations for improving overall data center energy efficiency include improving the design, operation, and control of mechanical systems serving the data centers in actual operation. This includes primary condenser water system, secondary condenser water system, CRAC units, and airflow management and control in data centers.

For the primary condenser water system, cooling plant optimization strategy should be developed. For example, control logic could be improved for cooling tower operation sequences. Operating both fans at lower speeds in a tower may be typically more efficient than operating one tower staged with another. Integrating VFD device and operation in cooling tower water system can improve the efficiency. This would be more useful, especially when the cooling load increases. Similarly, using more than one heat exchanger in parallel may lower

pumps' power demand. In addition, optimizing water temperature differential and pump head required would collectively contribute to minimizing total power demand for water systems.

For the building (secondary) condenser water system, supply water temperature and water flow rate from the heat exchanger may be optimized by providing variable-speed drives to the building condenser water pumps. The variable-speed drives on the secondary pumps can be controlled to provide a differential pressure control across the supply and return runs located at the end of the lines. The installed CRAC units were equipped with two-way modulating valves in the condensers controlled by compressor head pressures. Therefore, at lower cooling loads, these valves could reduce the flow rate of condenser water to the units.

Additional specific recommendations include:

- Optimize the actual air temperature and humidity set points, e.g., extending the permitted range. The make-up air unit's humidification system should be checked to ensure it is operating to maintain a minimum of 35% RH in the space.

- Optimize the control of supply and/or return air temperature from the CRAC units.

- Optimize air distribution through carefully placing perforated tiles, cable pass-through, equipment layout, and actual operation or non-operation of CRAC units. The benefits include achieving greater CRAC effectiveness, in the meanwhile perhaps less humidification and cooling would be required. For example, re-arrange CRAC DC 7 – 3 location and optimize its control. In addition, some CRAC units could be turned off.

- There was an observed difference in power demand readings before monitoring (3120 kW) and during monitoring period (1360 kW). It is worthwhile looking into calibration of the power metering device and/or fine-tuning operation to avoid or minimize electric demand charges for the data centers.

- Observing a discrepancy between the products of flow and temperature difference on the two sides of the heat exchanger, we suggest that the monitoring system be examined and calibrated, e.g., data acquisition for the secondary condenser system through the EMCS systems.

- Re-adjust the secondary condenser water supply temperature set point may be necessary, e.g., based upon outdoor air wet bulb temperature. This strategy allows a lower condenser water temperature to be delivered to the CRACs during most of the year, when the outdoor wet bulb temperature is lower than design conditions. A lower condenser water supply temperature would make it possible to lower the water flow rate to the CRAC units through reducing energy demand for water pumps. Considering the partial occupancy of the facility, the building cooling supply water temperature may be optimized using variable speed drives on the secondary condenser water pumps.

- The existing blow-through cooling towers were inefficient; therefore, considerations should be given to replace lead units with induced-draft towers.

6 Acknowledgements

This report on data center energy benchmarking was finalized based upon the field data collection performed by EYP Mission Critical Facilities and Landsberg Engineering and the draft report produced in the course of performing work subcontracted for the Lawrence Berkeley National Laboratory (LBNL).

The project was funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

7 Appendix A: Data Facility Definitions and Metrics

The following definitions and metrics are used to characterize data centers:

Air Flow Density	The airflow (cfm) in a given area (ft^2).
Air Handler Efficiency 1	The airflow (cfm) per power used (kW) by the CRAC unit fan.
Air Handler Efficiency 2	The power used (kW), per ton of cooling achieved by the air-handling unit.
Chiller Efficiency	The power used (kW), per ton of cooling produced by the chiller.
Computer Load Density – Rack Footprint	Measured Data Center Server Load in watts (W) divided by the total area that the racks occupy, or the "rack footprint".
Computer Load Density per Rack	Ratio of actual measured Data Center Server Load in watts (W) per rack. This is the average density per rack.
Computer/Server Load Measured Energy Density	Ratio of actual measured Data Center Server Load in watts (W) to the square foot area (sf) of Data Center Floor. Includes vacant space in floor area.
Computer/Server Load Projected Energy Density	Ratio of forecasted Data Center Server Load in watts (W) to the square foot area (sf) of the Data Center Floor if the Data Center Floor were fully occupied. The Data Center Server Load is inflated by the percentage of currently occupied space.
Cooling Load – Tons	A unit used to measure the amount of cooling being done. One ton of cooling is equal to 12,000 British Thermal Units (BTUs) per hour.
Data Center Cooling	Electrical power devoted to cooling equipment for the Data Center Floor space.
Data Center Server/Computer Load	Electrical power devoted to equipment on the Data Center Floor. Typically the power measured upstream of power distribution units or panels. Includes servers, switches, routers, storage equipment, monitors and other equipment.

Data Center Facility	A facility that contains both central communications and equipment, and data storage and processing equipment (servers) associated with a concentration of data cables. Can be used interchangeably with Server Farm Facility.
Data Center Floor/Space	Total footprint area of controlled access space devoted to company/customer equipment. Includes aisle ways, caged space, cooling units electrical panels, fire suppression equipment and other support equipment. Per the Uptime Institute Definitions, this gross floor space is what is typically used by facility engineers in calculating a computer load density (W/sf).
Data Center Occupancy	This is based on a qualitative estimate of how physically loaded the data centers are.
Server Farm Facility	A facility that contains both central communications and equipment, and data storage and processing equipment (servers) associated with a concentration of data cables. Can be used interchangeably with Data Center Facility. Also defined as a common physical space on the Data Center Floor where server equipment is located (i.e. server farm).

8 Appendix B: Facility Diagrams

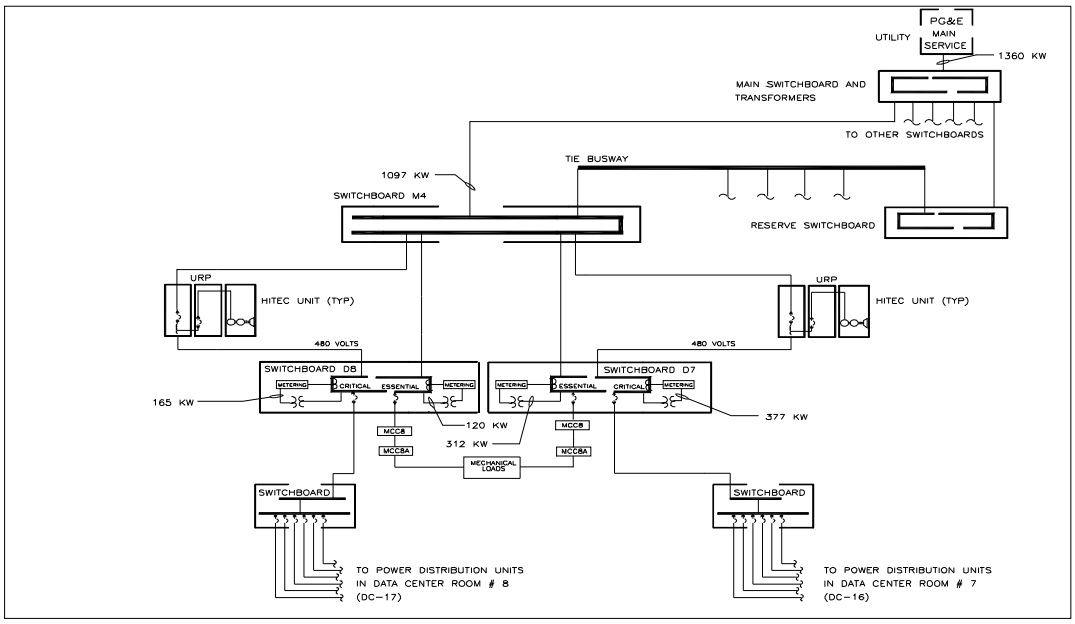

Figure 16. Electrical System Schematic

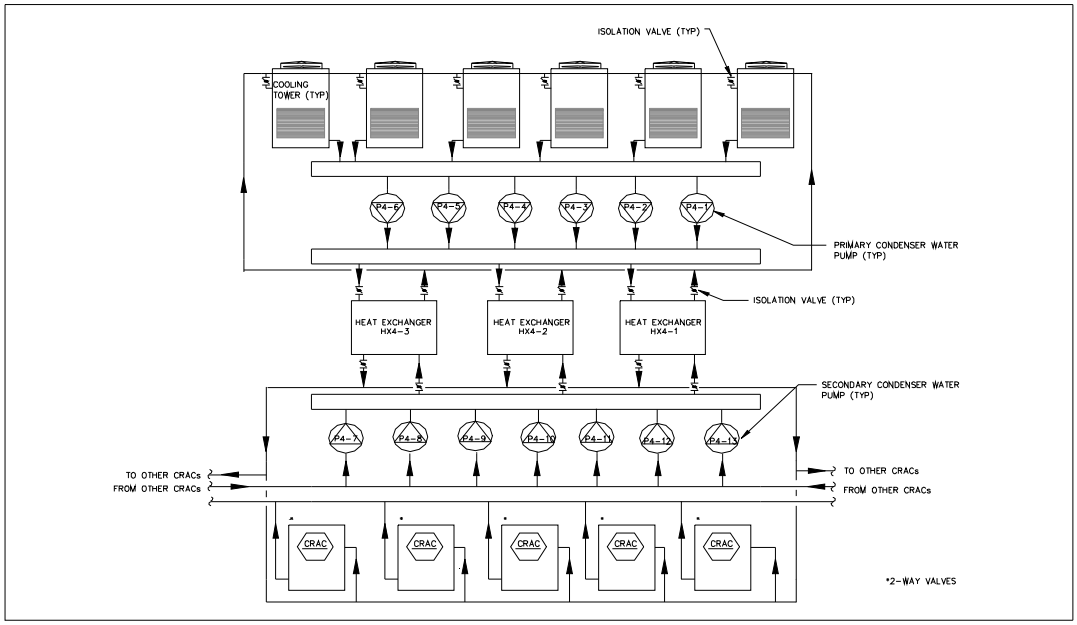

Figure 17. Condenser Water Flow Diagram

HIGH PERFORMANCE DATA CENTERS

A Design Guidelines Sourcebook

January 2006

Pacific Gas and Electric Company®

TABLE OF CONTENTS

Introduction.....2

1. Air Management.....3

2. Air-side Economizer.....11

3. Centralized Air Handling.....17

4. Cooling Plant Optimization.....24

5. Direct Liquid Cooling.....31

6. Free Cooling via Water Side Economizer.....37

7. Humidification Controls Alternatives.....42

8. Power Supplies.....47

9. Self Generation.....52

10. Uninterruptible Power Supply Systems.....55

INTRODUCTION

Data centers can consume 25 to 50 times as much electricity as standard office spaces. With such large power consumption, they are prime targets for energy efficient design measures that can save money and reduce electricity use. But the critical nature of data center loads elevates many design criteria -- chiefly reliability and high power density capacity – far above efficiency. Short design cycles often leave little time to fully assess efficient design opportunities or consider first cost versus life cycle cost issues. This can lead to designs that are simply scaled up versions of standard office space approaches or that re-use strategies and specifications that worked "good enough" in the past without regard for energy performance. The Data Center Design Gudelines have been created to provide viable alternatives to inefficient building practices.

Based upon benchmark measurements of operating data centers and input from practicing designers and operators, the Design Guidelines are intended to provide a set of efficient baseline design approaches for data center systems. In many cases, the Design Guidelines can also be used to identify cost-effective saving opportunities in operating facilities. No design guide can offer 'the one correct way' to design a data center, but the Design Guidelines offer efficient design suggestions that provide efficiency benefits in a wide variety of data center design situations. In some areas, promising technologies are also identified for possible future design consideration.

Data center design is a relatively new field that houses a dynamic and evolving technology. The most efficient and effective data center designs use relatively new design fundamentals to create the required high energy density, high reliability environment. The following Best Practices capture many of the new 'standard' approaches used as a starting point by successful and efficient data centers.

1. AIR MANAGEMENT

Modern data center equipment racks can produce very concentrated heat loads. In facilities of all sizes, from small data center supporting office buildings to dedicated co-location facilities, designing to achieve precise control of the air flow through the room that collects and removes equipment waste heat has a significant impact on energy efficiency and equipment reliability.

Air management for data centers entails all the design and configuration details that go into minimizing or eliminating mixing between the cooling air supplied to equipment and the hot air rejected from the equipment. When designed correctly, an air management system can reduce operating costs, reduce first cost equipment investment, increase the data center's density (W/sf) capacity, and reduce heat related processing interruptions or failures. A few key design issues include the location of supply and returns, the configuration of equipment's air intake and heat exhaust ports and the large scale airflow patterns in the room.

PRINCIPLES

- Use of best-practices air management, such as strict hot aisle/cold aisle configuration, can double the computer server cooling capacity of a data center.
- Combined with an airside economizer, air management can reduce data center cooling costs by over 60%1.
- Removing hot air immediately as it exits the equipment allows for higher capacity and much higher efficiency than mixing the hot exhaust air with the cooling air being drawn into the equipment. Equipment environmental temperature specifications refer primarily to the air being drawn in to cool the system.
- A higher difference between the return air and supply air temperatures increases the maximum load density possible in the space and can help reduce the size of the cooling equipment required, particularly when lower-cost mass produced package air handling units are used.
- Poor airflow management will reduce both the efficiency and capacity of computer room cooling equipment. Examples of common problems that can decrease a Computer Room Air Conditioner (CRAC) unit's usable capacity by 50%[2] or more are: leaking floor tiles/cable openings, poorly placed overhead supplies, underfloor plenum obstructions, and inappropriately oriented rack exhausts.

APPROACH

Improved airflow management requires optimal positioning of the data center equipment, location and sizing of air opening and the design and upkeep of the HVAC system. While the application can vary widely, one overall objective is simple: to remove hot air exhaust from the equipment before the exhaust, and the heat it carries, is mixed with cool supply air and recirculated back into the equipment. Countless design strategies can be used to achieve this

objective. They include: hot aisle/cold aisle rack layout; flexible barriers; ventilated racks; and optimized supply/return grills and/or floor tiles. Energy savings are realized by extending economizer savings into higher outdoor air temperatures (up to 80-85°F) and/or reducing fan airflow and power costs in spaces running at less than design cooling capacity.

Increased economization is realized by utilizing a control algorithm that brings in outside air whenever it is appreciably cooler than the return air and when humidity conditions are acceptable (see Airside Economizer Chapter for further detail on economizer control optimization). In order to save energy, the temperature outside does not need to be below the data center's temperature setpoint; it only has to be cooler than the return air that is exhausted from the room. As the return air temperature is increased through the use of good air management, the temperature at which economization will save energy is correspondingly increased. Designing for a higher return air temperature increases the number of hours that outside air, or a waterside economizer/free cooling, can be used to save energy.

Fan energy savings are realized by reducing fan speeds to only supply as much air as a given space requires. There are a number different design strategies that reduce fan speeds, but the most common is a fan speed control loop controlling the cold aisles' temperature at the most critical locations – the top of racks for underfloor supply systems, the bottom of racks for overhead systems, end of aisles, etc. Note that many Computer Room Air Conditioners use the return air temperature to indicate the space temperature, an approach that does not work in a hot aisle/cold aisle configuration where the return air is at a very different temperature than the cold aisle air being supplied to the equipment. Control of the fan speed based on the space temperature is critical to achieving savings.

Higher return air temperature also makes better use of the capacity of standard package units, which are designed to condition office loads. This means that a portion of their cooling capacity is configured to serve humidity (latent) loads. Data centers typically have very few occupants and small outside air requirements, and therefore have negligible latent loads. While the best course of action is to select a unit designed for sensible-cooling loads only or to increase the airflow, an increased return air temperature can convert some of a standard package unit's latent capacity into usable sensible capacity very economically. This may reduce the size and/or number of units required.

HOT AISLE/COLD AISLE

A basic hot aisle/cold aisle configuration is created when the equipment racks and the cooling system's air supply and return are designed to prevent mixing of the hot rack exhaust air and the cool supply air drawn into the racks. As the name implies, the data center equipment is laid out in rows of racks with alternating cold (rack air intake side) and hot (rack air heat exhaust side) aisles between them. The aisles are typically wide enough to allow for maintenance access to the racks and meet any code requirements. All equipment is installed into the racks to achieve a front-to-back airflow pattern that draws conditioned air in from cold aisles, located in front of the equipment, and rejects heat out through the hot aisles

behind the racks. Equipment with non-standard exhaust directions must be addressed in some way (shrouds, ducts, etc.) to achieve a front-to-back airflow. The rows of racks are placed back-to-back, and holes through the rack (vacant equipment slots) are blocked off on the intake side to create barriers that reduce recirculation, as shown in the graphic below. A raised floor system would be the same, except with the supply coming from tiles in the cold aisle. With proper isolation, the temperature of the hot aisle no longer impacts the temperature of the racks or the reliable operation of the data center; the hot aisle becomes a heat exhaust. The HVAC system is configured to supply cold air exclusively to the cold aisles and pull return air only from the hot aisles.

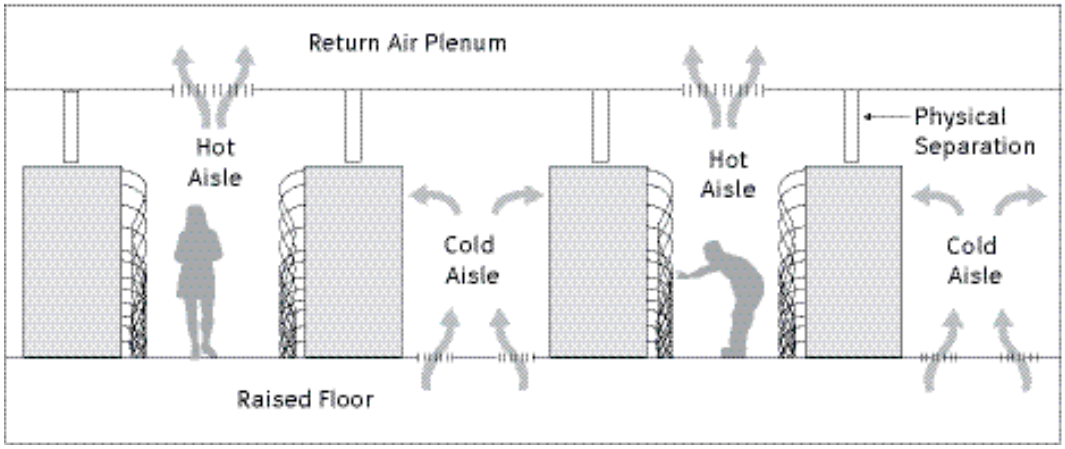

FIGURE 1
Hot Aisle/Cold Aisle Arrangement

The hot rack exhaust air is not mixed with cooling supply air and therefore can be directly returned to the air handler through various collection schemes, returning air at a higher temperature, often 85°F or higher. The higher return temperature extends economization hours significantly and/or allows for a control algorithm that reduces supply air volume, saving fan power. In addition to energy savings, higher equipment power densities are also better supported by this configuration. The significant increase in economization afforded by hot aisle/cold aisle configuration can improve equipment reliability in mild climates by providing emergency compressor-free data center operation during outdoor air temperatures up to the data center equipment's top operating temperature (typically 90°F-95°F). Greater economization also can reduce central plant run-hour related maintenance costs.

Hot aisle/cold aisle configurations can be served by overhead or underfloor air distribution systems. When an overhead system is used, supply outlets that 'dump' the air directly down should be used in place of traditional office diffusers that throw air to the sides, which results in undesirable mixing and recirculation with the hot aisles. In some cases, return grills or simply open ducts, have been used. Underfloor distribution systems should have supply tiles in front of the racks. Open tiles may be provided underneath the racks, serving air directly into the equipment. However, it is unlikely that supply into the bottom of a rack alone will adequately cool equipment at the top of the rack without careful rack design.

The potential for cooling air to be short-circuited to the hot aisle should be evaluated on a rack-by-rack basis, particularly in lightly loaded racks. Floor tile leakage into the hot aisles

represents wasted cooling and lost capacity and should be regularly checked and corrected. Operators should be properly educated on the importance of conserving cooling air in the hot aisles, to prevent misguided attempts to 'fix the hot spots,' and to encourage correction of leaks in raised floor assemblies. Operator education is very important since a hot aisle/cold aisle configuration is non-intuitive to many data center personnel who are often trained to eliminate 'hot spots,' not deliberately create them in the form of hot aisles. The fact that only the air temperature at the inlet of equipment must be controlled is the basis of good air management.

The hot aisle/cold aisle configuration is rapidly gaining wide acceptance due to its ability to serve high density racks better than traditional, more mixed flow configurations. As the power consumption of a single loaded rack continues to climb, exceeding 14 kW in some cases, the physical configuration of the rack's cooling air intake and hot air exhaust becomes crucial. Data center operators have discovered that exhausting large heat loads directly onto a rack of equipment can lead to overheating alarms and equipment failure regardless of the amount of room cooling available. First and foremost, a hot aisle/cold aisle configuration is an equipment layout that improves reliable operation.

A hot aisle/cold aisle design approach requires close coordination between the mechanical engineer designing the cooling system for the data center space and the end users that will be occupying the space. Successful air management goes beyond the design of the room and requires the direct cooperation of the room occupants, who select and install the heat generating equipment.

Underfloor air supply systems have a few unique concerns. The underfloor plenum serves both as a duct and a wiring chase. Coordination throughout design and into construction is necessary since paths for airflow can be blocked by uncoordinated electrical or data trays and conduit. The location of supply tiles needs to be carefully considered to prevent short circuiting of supply air and checked periodically if users are likely to reconfigure them. Removing tiles to 'fix' hot spots can cause problems throughout the system.

Light fixtures and overhead cable trays should be laid out in coordination with the HVAC air supply to ensure no obstructions interrupt the delivery and removal of air to the rows. Hanging fixtures or trays directly under an air supply should be avoided.

FLEXIBLE BARRIERS

Using flexible clear plastic barriers, such as plastic supermarket refrigeration covers or other physical barriers, to seal the space between the tops of the rack and the ceiling or air return location can greatly improve hot aisle/cold aisle isolation while allowing flexibility in accessing, operating, and maintaining the computer equipment below. One recommended design configuration supplies cool air via an underfloor plenum to the racks; the air then passes through the equipment in the rack and enters a separated, semi-sealed area for return to an overhead plenum. This displacement system does not require that air be accurately directed or superchilled. This approach uses a baffle panel or barrier above the top of the rack

and at the ends of the cold aisles to eliminate "short-circuiting" (the mixing of hot and cold air). These changes should reduce fan energy requirements by 20–25 percent, and could result in a 20 percent energy savings on the chiller side.

With an upflow CRAC unit, combining pairs of racks with a permeable barrier creates a system in which hot air can be immediately exhausted to the plenum. Unfortunately, if the hot-cool aisle placement is reversed (with the cold aisles being the ducted aisles), the working (human) spaces would be hot—at temperatures up to or even above 90ºF[3].

VENTILATED RACKS

The ideal air management system would duct cooling air directly to the intake side of the rack and draw hot air from the exhaust side, without diffusing it through the data center room space at all. Specialized rack products that utilize integral rack plenums that closely approximate this ideal operation are beginning to appear on the market. Custom solutions can also be designed using the well defined design principles used for heat and fume exhaust systems.

Such designs should be evaluated on the basis of their effectiveness in capturing hot exhaust air with a minimum of ambient air mixing (typically achieved by placing the capture opening very close to the hot exhaust) and factoring in any fan energy costs associated with the systems. Exhaust systems typically have far higher fan energy costs than standard returns, so the use of small diameter ducting or hoses and multiple small fans should be carefully evaluated to ensure that additional fan power cost does not seriously reduce or eliminate the savings anticipated from improved air management.

OPTIMIZED SUPPLY/RETURN CONFIGURATION

All of the design methods discussed above are approaches to optimizing the airflow through a data center to minimize the mixing of cool supply air and hot waste heat from the equipment. A comprehensive design approach to air management is the single best approach to improving efficiency; however, in retrofit situations or where no resources are available to properly implement airflow control, some simple, low-cost steps can help a data center operate slightly more efficiently.

Diffusers that dump air straight down should be selected and located directly in front of racks, not above or behind. Unlike an office space design, diffusers should be selected and placed in order to dump air directly to where it can be drawn in the equipment, rather than to provide a fully mixed room without human-sensible drafts. The thermostat should be located in an area in front of the computer equipment, not on a wall behind the equipment. Finally, where a rooftop unit is being used, it should be located centrally over the served area – the required reduction in ductwork will lower cost and slightly improve efficiency. While maintenance and roof leak concerns may preclude locating the unit directly over

data center space, often a relatively central location over an adjacent hall or support area is appropriate.

BENCHMARKING FINDINGS/CASE STUDIES

An existing data center room cooled by an underfloor system was having trouble maintaining temperature. Chilled-water cooled Computer Room Air Conditioners (CRACs) with a total capacity of 407 tons were installed and operating in the room. All available floor space for CRACs had been used in an attempt to regain control of the room. The chilled water loop serving the data center, a district cooling system with an installed capacity of 4,250 tons, had the chilled water temperature reset down to 41°F, primarily to assist in cooling this single data center facility. Measurements of the room revealed the air flow condition seen in the figure below.

FIGURE 2
POOR AIRFLOW CONDITION

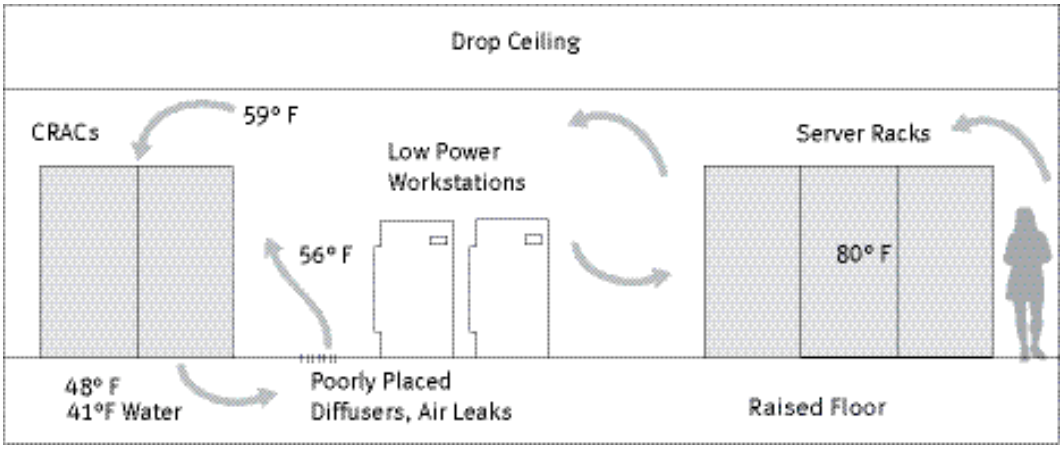

A lack of capacity was the suspect issue. However, air temperature measurements quickly suggested that the actual problem was not the installed capacity, but the airflow management. The majority of the CRACs were located at the end of the room farthest from the highest heat density racks, and they used non-ducted "through-the-space" returns.

Between the racks and the CRAC, there were a number of diffuser floor tiles supplying cooling air to workstations with rather low heat loads. Additionally, there were also loose tiles and

FIGURE 3
CRAC RETURN AIR TEMPERATURE VS CAPACITY

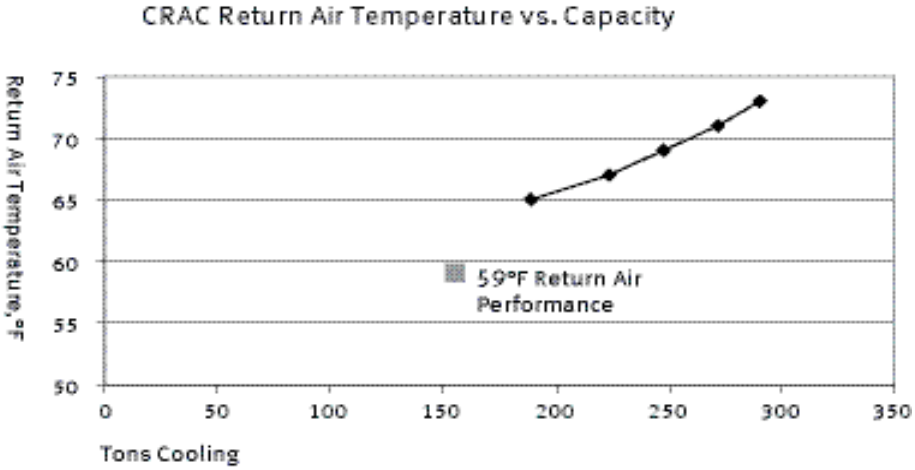

significant air leaks in this area. The result was that a large percentage of cooling air never made it to the racks. Instead, the cooling air bypassed the main load and was returned directly to the CRAC units. The return air temperature at the CRACs was low enough to serve as supply air in many facilities. The result of the low return air temperature is seen in the graph below – the CRAC capacity was derated to almost 50% below its name plate capacity due to the low return air temperature.

In this situation, the ideal solution was not practical. Reconfiguring the data center to implement a hot aisle/cold aisle design and moving the CRAC units closer to the load would have been prohibitively expensive and would have required unacceptable amounts of down time. An alternate solution was proposed to improve airflow as shown below.

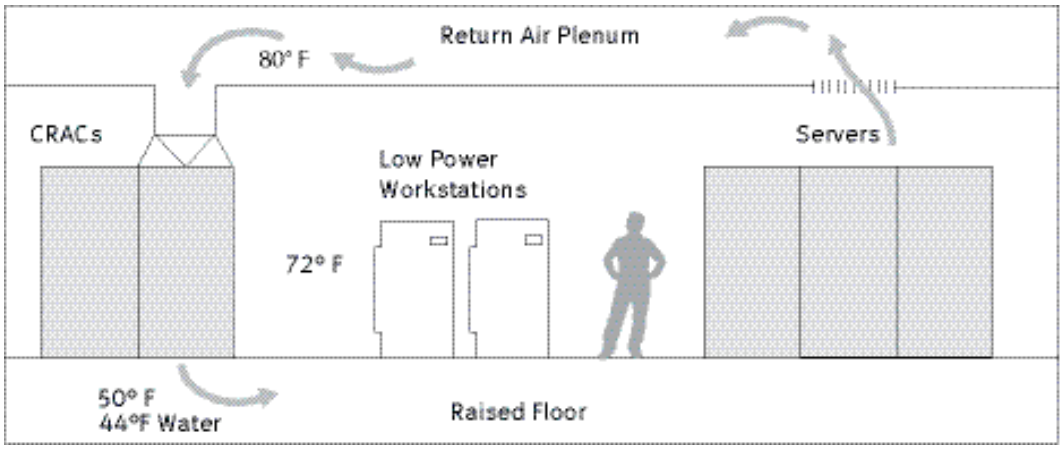

FIGURE 4
ALTERNATE AIRFLOW SOLUTION

The main objective is to collect waste heat as close as possible to the source and then isolate it from the space in the return air stream. Floor supply tiles are replaced by solid tiles in the low load area to increase the air supplied to the farthest, and largest, rack loads. The return air is no longer drawn through the space. Return air travels through a ceiling plenum to avoid mixing with the cool air being served to the room. Note that only stratification is being used to scavenge the hot rack exhaust, resulting in a lower return air temperature than would be expected with a well executed isolation strategy such as hot aisle/cold aisle.

This design literally doubles the cooling capacity of the installed CRACs and allows the district plant chilled water temperature to be increased to 44°F, which would substantially increase the efficiency of the entire 4,250 ton district chilled water system. In this case, correcting the airflow management offered a solution to cooling problems when the traditional approach of adding more cooling equipment did not.

RELATED CHAPTERS

- Air-side Economizer
- Direct Liquid Cooling of Racks
- Free Cooling via Waterside Economizer

REFERENCES

1) Compared to a standard system with no economizer and a return temperature of 78º rather than 85º.

2) CRAC capacity is proportional to the difference between return and supply temperature in standard units.

3) Source: Rocky Mountain Institute (www.rmi.org)

RESOURCES

- Thermal Guidelines for Data Processing Environments, TC9.9 Mission Critical Facilities, ASHRAE, 2004.
- Data Processing and Electronic Areas, Chapter 17, ASHRAE HVAC Applications, 2003.
- Design Recommendations for High Performance Data Centers, Rocky Mountain Institute, 2003.

2. AIRSIDE ECONOMIZER

Data centers operate 24-hours a day with a large, constant cooling load that is independent of the outdoor air temperature. Most nights and during mild winter conditions, the lowest cost option to cool data centers is an air-side economizer. Simply using a the standard office system economizer offered on Title 24 compliant units is not advised until a proper engineering evaluation of the local climate conditions and the space requirements can be completed. Due to the humidity and contamination concerns associated with data centers, careful control and design work may be required to ensure that cooling savings are not lost because of excessive humidification requirements.

Providing adequate access to outdoor air for economization is not an insignificant architectural design challenge. Central air handling units with roof intakes or sidewall louvers are the most commonly used, although some internally-located Computer Room Air Conditioning (CRAC) units offer economizer capabilities when installed with appropriate intake ducting. Rooftop intakes are more secure than ground-level sidewalls, where security bars behind the louver may be appropriate in some unstaffed and/or high risk facilities. Depending on local climate conditions, outdoor air humidity sensors may be required to permit lockout of economization during very dry, high humidification-load, conditions. For control of particulates and contamination, appropriate low-pressure drop filtration should be provided in order to maintain a clean data center environment while not imposing excessive fan power energy costs. Other contamination concerns such as salt or corrosive should be evaluated[1]. In most areas, use of outside air is beneficial, however in critical applications local risk factors should be known and addressed.

PRINCIPLES

- An economizer can cut data center cooling costs by over 60%[2] using standard, commonly available low-cost equipment. Depending on the climate, the steady, 24-hour cooling load of a data center is well suited to take advantage of seasonal and nighttime temperature variations to cool the space.
- A standard data center cooling system can remove heat from the data center it serves only by running compressor(s), a major electrical cost. With an economizer, when the outside air is cooler than the return air, hot return air is exhausted and replaced with cooler, filtered outside air – essentially 'opening the windows' for free cooling.
- Economization must be engineered into the air handling system. Small data centers may be economically served by low cost, mass produced package units. Larger data centers typically justify a more efficient chilled water system with central air handlers.
- In dry climates, controls should include redundant outdoor air humidity sensors to stop economization when the absolute humidity (or dewpoint) is too low to prevent causing an artificially expensive humidification load on very cold days. Dry climates can often realize excellent savings from an evaporative cooling or water-side economizer approach.

- In small data centers located in mixed-use buildings, some energy savings may be realized by maximizing the use of a house, office or support area system that is equipped with an economizer.

APPROACH

An outdoor economizing system is best implemented starting at the schematic design stage, where the required architectural accommodations can be made with little or no additional cost. An air handler equipped with an airside economizer will draw outside air into the space when the outside air (OSA) is cooler than the return air. During economization, the return air system operates as a heat exhaust. Depending on the space load, at about 55°F OSA or below, the compressor(s) should not be required to run at all. During economization, the supply air picks up heat from the space and is exhausted to the outside instead of being recirculated. Energy is saved by simply exhausting the heat rather than removing it mechanically via a compressor. Cooler outside air is then drawn in, cooling the space. While economization is usually implemented with a central air handler configuration, many Computer Room Air Conditioning (CRAC) units (direct expansion or water coil cooled) offer optional economizers and controls.

Data centers in temperate climates with no space humidity control requirements may be able to use standard economizer controls, which do not consider humidity and operated based only on the drybulb temperature. In the more common situation where some level of humidity control is specified, controls may need to be added to lock out the economizer based on low OSA humidity (usually below about 48ºF dewpoint, depending on space humidity and temperature setpoints). The humidity lockout should use an absolute humidity setpoint, either dewpoint or humidity ratio (mass of water vapor/mass of air) for stable operation.

Alternatively, in areas with long periods of cold, dry weather a somewhat more complex humidity control system may be justified to maximize economization. As described below, a sidestream adiabatic humidification system scavenging waste computer heat may be suitable.

Humidity control algorithms should be designed to account for the local conditions, in particular the rate of humidity change. In a data center situation, when a rapid change in outdoor humidity is sensed, economization can simply be stopped and returned to standard mechanical control usually in a matter of minutes. Rapidly changing outdoor air humidity is more of a control problem in critical facilities that do not have the option of closing the outdoor air dampers and returning to a recirculation configuration, such as such as cleanroom fabrication plants with a high exhaust air makeup requirement.

Economization hours can be efficiently extended into dry winter months if computer waste heat is recovered to provide humidification energy. One approach is to use an adiabatic humidifier to humidify a small amount of warm return air, which is then injected into the supply stream. An adiabatic humidifier exposes water to the air stream and uses the heat energy of the air stream itself to evaporate the water. Water is usually exposed through the surface of a wet media, spraying through mist nozzles or by producing an

ultrasonic-generated fog. The return air carrying the waste heat from the computers is cooled evaporatively as it is humidified before it is introduced into the supply air stream; the common direct evaporative cooler (or 'swamp cooler') widely used for greenhouses and in dry climates is a form of adiabatic humidifier. Adiabatic bypass humidifiers can greatly extend the use of, and savings from, economizers for data centers in dry, cold climates.

The figure below shows a simplified plan view of a small rooftop data center system that uses an adiabatic humidifier scavenging waste heat from the return air. As shown, the outside air is slightly too cold to be humidified to the desired dewpoint by direct adiabatic humidification, hence the use of return air, which has more than enough heat energy to evaporate enough moisture. Depending on climate, it can be beneficial to configure the adiabatic humidification to act directly on the outside air to allow for evaporative cooling during hot, dry periods; the example below is appropriate for cold climates.

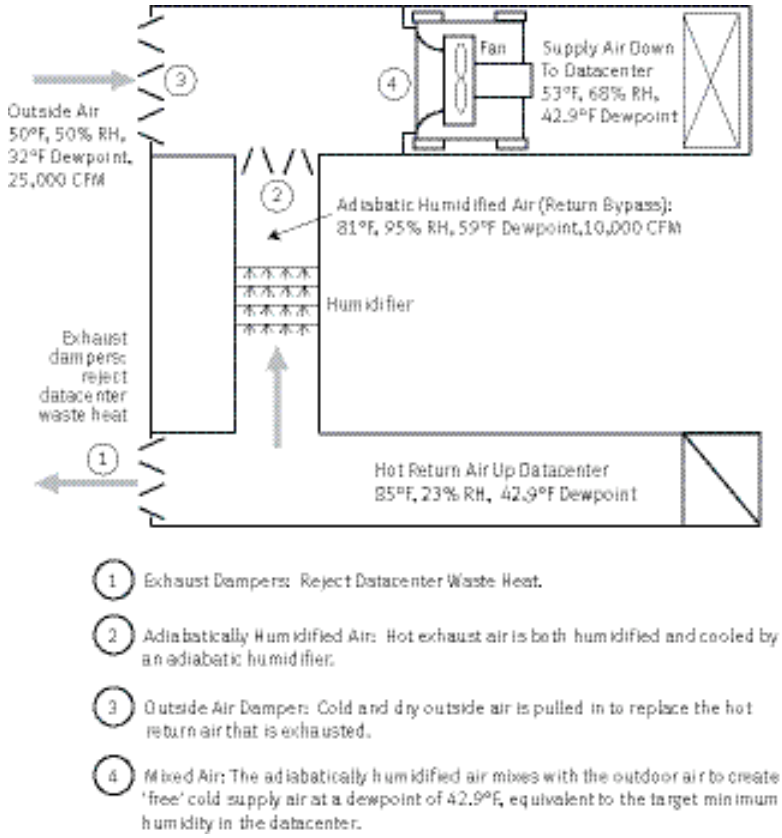

FIGURE 1
ROOFTOP DATA CENTER SYSTEM USING AN ADIABATIC HUMIDIFIER

Integrated economizers provide additional savings compared to non-integrated ones and are the standard option on most modern package units since they are required by California Title 24 code for units over 2,500 CFM and 75,000 btu/hr capacity sold to the large office space market. Rather than waiting until the outside air temperature falls to the supply air temperature setpoint, an integrated economizer opens when the outside air temperature falls below the economization setpoint, which for data centers is ideally the current return temperature (usually plus one or two degrees to account for potential sensor inaccuracy). The compressor(s) run as required and additionally cool the outside air temperature to the supply

air temperature setpoint. The total cooling required from the compressors is reduced by using outside air rather than warmer return air.

An airside economizer offers substantially better savings when paired with a hot aisle/cold aisle configuration that increases the return air temperature, up to 85F or higher in well executed layouts. A higher data center room—or, ideally, cold aisle – temperature setpoint, such as 78F, also increases the potential savings from an economizer. Techniques to increase the return air temperature are discussed in the Air Management chapter. With a properly configured integrated economizer, savings are realized whenever the return air temperature, not the space setpoint temperature, is above the outside air temperature. A well optimized data center layout allows an economizer system to serve as a heat exhaust, collecting waste heat at the source and directly exhausting it. In a well configured data center with a high return air temperature, the economizer can actually provide an additional level of redundancy to the mechanical cooling equipment for much of the year.

Most data centers will require a reasonable level of filtration of outside air. Ideally, the filtration system can be configured to filter outside air prior to combining it with the recirculation air stream, to eliminate unnecessary filtering of the cleaner recirculation air. The fan power required for filtration can be significant if care is not taken to design low face velocity filtration and to use extended media filters.

Other contamination concerns, such as salt or corrosives entrainment, should be evaluated. Corrosive atmospheres tend to be a localized risk and therefore require a case by case evaluation and treatment design. It should be noted that while outdoor contaminants are a site specific concern, a minimal rate of ventilation is recommended for all facilities to control internally generated contaminates[3].

Smaller data centers located in large office buildings may realize significant energy efficiency by using the main office system during normal hours as a primary source of cooling. In particular, data centers that are served by air-cooled split systems or air-cooled Computer Room Air Conditioners (CRACs) present good opportunities for energy savings. In California, due to Title 24 requirements, the main office system probably has an economization capability. In addition, the larger house systems are often more efficient than the standard, air-cooled data center systems, particularly if the house systems use a central chilled water plant. During off-hours, the dedicated computer room system can provide for data center cooling since the main house system would typically be oversized for the small data center alone.

BENCHMARKING FINDINGS/CASE STUDIES

A 1,400 sf data center located in a large office building converted to a system with full economization capability when its dedicated 25-ton air cooled chiller required replacement. Due to load growth and reliability problems, portable air conditioners were also being used to maintain control in the data center space. The existing system used chilled water fed CRAC units located in the data center. During normal operating hours, a large house chilled water plant system served the space, while the dedicated chiller was used for off hour (night and weekend) cooling and backup.

An air-cooled package unit was chosen to replace both the air-cooled chiller and the associated chilled-water CRAC units. The package unit, shown in the picture to the right, included full airside economization capability. With this unit, exhaust air is ejected from one end of the air handler, while outside air is drawn in from the opposite end of the unit through louvers in the screenwall. This system had multiple benefits for the data center: it replaced the failing chiller, allowed for airside economization, and eliminated all chilled water piping from within the data center envelope. The main house air handler system, based on a large water cooled chiller plant, is used during the day and can be started and used at anytime to provide emergency backup.

FIGURE 2
Data Center Converted to Use an Airside Economizer

This system also freed up floorspace in the data center by removing the CRAC units. The removal of the CRAC units effectively enlarged the space available for computing equipment, reducing the need for costly future expansions of the space. The data center can still be served from the central plant, but now through cooled air from the house air handler rather than chilled water from the house loop. The data center uses a hot aisle/cold aisle configuration to increase the return air temperature, which both extends economizer operation into correspondingly higher outdoor air temperatures and increases the amount of heat that is exhausted during economizer operation (see Figure 3).

FIGURE 4
SCHEMATIC DIAGRAM OF AIR DISTRIBUTION SYSTEM

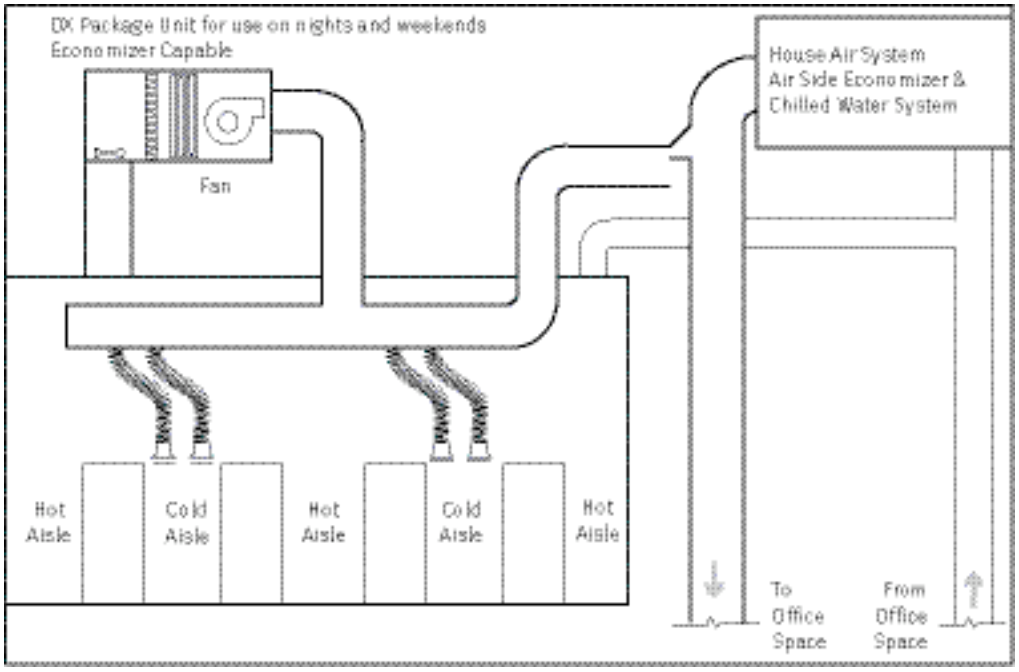

Measurements of the data center showed it to be lightly loaded overall at less than 30 w/sf. Hourly simulation of the data center was performed using a DOE2 model and demonstrated total energy savings of 45% (254,000 kWh/yr). Approximately one quarter of the savings (66,000 kWh/yr) were due to the economizer alone (the balance of savings were from reduced fan power, pumping power, and integration with the central house system). Addition of the economizer to the package unit was a negligible incremental cost item in this project where the primary objective was replacement of the failing chilled water system.

RELATED CHAPTERS

- Air Management
- Freecooling via Waterside Economizer

RESOURCES

- Data Processing and Electronic Areas, Chapter 17, ASHRAE HVAC Applications, 2003.

REFERENCES

1) ASHRAE, 1999 HVAC Applications, pg 16.1

2) Assuming San Jose, California typical weather year, 78°F return temperature and 55°F supply to a 24 hour facility (70% fan efficiency with a standard motor serving a 2" w.g. duct system assumed).

3) ASHRAE, 1999 HVAC Applications, pg 16.1

3. CENTRALIZED AIR HANDLING

FIGURE 1
Interior and Exterior Air Handling Units

Better performance has been observed in data center air systems that utilize specifically-designed central air handler systems. A centralized system offers many advantages over the traditional multiple distributed unit system that evolved as an easy, drop-in computer room cooling appliance. Centralized systems use larger motors and fans, and can be more efficient. They are also well suited for variable volume operation through the use of Variable Speed Drives (VSDs, also referred to as Variable Frequency Drives or VFDs). Most data center loads do not vary appreciably over the course of the day, and the cooling system is typically oversized with significant reserve capacity. A centralized air handling system can improve efficiency by taking advantage of surplus and redundant capacity to actually improve efficiency. The maintenance benefits of a central system are well known, and the reduced footprint and maintenance traffic in the data center are additional benefits.

PRINCIPLES

- A central system allows redundancy to be implemented in a manner that provides maximum reliability (a spinning reserve) and increases normal operating system efficiency. System maintenance is also simplified through centralization.
- In most California climates, air handlers can be located on the roof (ideally in a central location not directly over the data center the space), allowing significant cost savings compared to Computer Room Air Conditioners (CRACs) by reducing the data center floorspace required by the air conditioning.
- Fans and motors tend to be more efficient in larger systems.
- Large air handlers equipped with variable air volume fans tend to have better efficiency when underloaded, as opposed to Computer Room Air Conditioner (CRAC) units where efficiency suffers at part loads. Data center systems are typically operated at part-load to ensure maximum temperature and humidity control stability, reliability, and margin for future increases in load.

- While data center loads tend to be constant 24 hours a day, the loading across the data center floor can vary significantly. A central system can reduce fan power use and save energy by taking advantage of this variance. A low-pressure drop design ('oversized' ductwork or a generous underfloor) is essential to optimizing energy efficiency and long-term buildout flexibility.
- Piping for condensate, humidification, chilled or condenser water and/or refrigerant is reduced or eliminated within the data center envelope.
- Implementation of an airside economizer system is simplified with a central air handler system. Optimized air management, such as that provided by hot aisle/cold aisle configurations, is also easily implemented with a ducted central system.

APPROACH

Early in the evolution of data centers, the typical cooling system involved multiple small air-cooled split systems with small vertical air handlers and independent integrated controls that stood in the data center room and provided cooling. Such a system was easy to install in existing buildings that were initially constructed without consideration for the high density sensible heat loads of modern electronic equipment. Now that the loads and conditioning requirements of data centers are relatively well understood, purpose-built central air handler systems can be designed to meet typical data center requirements with greater efficiency than the traditional multiple distributed units design seen in the figure below.

FIGURE 2
Traditional Multiple Distributed Unit Air Handling System Utilizing Computer Room Air Conditioner (CRAC) Units

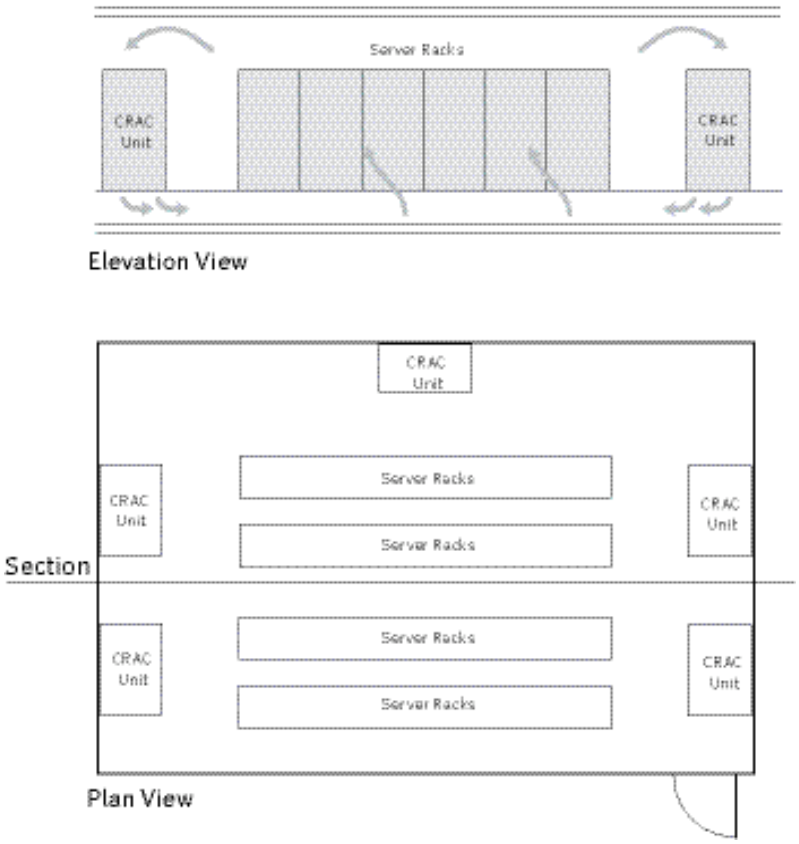

A centralized air handler system can be used to provide efficient data center cooling. The typical approach to a centralized system for a data center is to use a single supply air plenum fed by a few large air handlers. Depending on the size of the space and frequency of data center load changes, variable airflow boxes may be used within the distribution system to actively direct the airflow to the highest load regions. Alternatively, the system may simply be statically balanced after every major change in load. The figure below shows a centralized air handler system where the data center distribution system is fed from above.

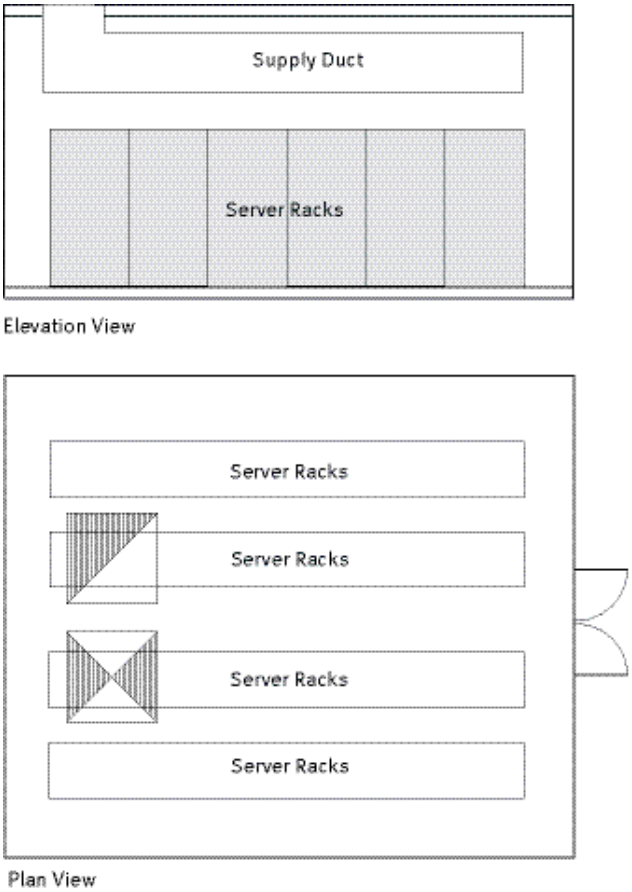

FIGURE 3
TRADITIONAL MULTIPLE DISTRIBUTED UNIT AIR HANDLING SYSTEM UTILIZING COMPUTER ROOM AIR CONDITIONER (CRAC) UNITS

The distribution system must be designed for low pressure drop in order to maximize energy savings. Ducts should be sized significantly larger than typical office systems, since 24 hour operation of the data center increases the value of energy use over time relative to first cost. Since loads often only change when new servers or racks are added or removed, a static balance approach can be quite effective due to the constant, 24 hour nature of most loads. The distribution system may be either underfloor or a traditional overhead ducted configuration – in determining the load carrying capacity, proper air management is usually more important than the supply configuration (see Air Management chapter).

This concept utilizes a smaller number of large premium efficiency motors, fans, and Variable Frequency Drives (VFDs) rather than many smaller motors and fans. Central air handlers should be much larger than the units used in a conventional multiple-distributed system using CRAC units. Larger fans and motors improve in efficiency as they are scaled up; for

example, a NEMA Premium efficiency 10 hp motor is 91% efficient or better, while a 30 hp motor with the same premium designation is 93% efficient. In energy terms, for a 30 hp nameplate motor, a 2% efficiency difference equals about 4,000 kWh per year in energy savings – savings possible simply through the consolidation of three small 10 hp motors into a single 30 hp motor[1]. Fans are more variable in performance, but they too follow a similar trend of higher efficiencies being achieved with larger fans. Since air handlers can be easily located outside of the expensive data center floor area, there is more flexibility to optimize fan inlet conditions to avoid efficiency-robbing system effects. In addition, VFDs can be used more economically with fewer large fans to control fan volume, allowing the system to actively adjust to the exact demand of the data center.

The use of VFDs on data center systems offers significant savings, since oversizing the air handlers, which is standard data center design practice, can actually result in greater system efficiency. While oversized chiller or air-conditioner compressors commonly prove detrimental to energy performance, oversized air handler systems can provide surplus future capacity and redundancy, and can be designed to realize significant part-load energy savings. Until the data center facility is fully built out or loaded with more energy intensive devices, with all future expansion and safety factors exhausted, the system will operate at part-load with surplus capacity available. The surplus capacity offers the ability to save energy in two ways: Part-load operation allows the supply air volume to be reduced; and larger coiling coil areas allow the chilled water system temperature to be increased in order to improve the chiller plant efficiency.

Air handler fan speeds should be controlled in order to reduce the air volume supplied during part-load conditions. Also, all air handlers, including redundant air handling unit(s), should be configured to operate in parallel. Until the unit fans are operating at the minimum speed allowed by the drive and motor, more operating units result in higher air system efficiency. This also provides an always-on spinning reserve that can seamlessly react to a fan failure while resulting in a lower total fan power use[2]. Improvements in fan system efficiency add up to significant savings. The HVAC fan power consumption for a 100 w/sf design data center can range from 11 W/sf to over 22 W/sf[3]. Reducing the supply air volume offers a large, non-linear reduction of this fan energy use – just a 20% reduction in airflow volume reduces the power by 45-50%.

Using the larger coil area available during part-load conditions to increase the chilled water temperature setpoint is discussed in the Cooling Plant Optimization Design Guidelines document. A centralized air handler system in the data center should use temperature sensors, located in at the intakes of the IT equipment, where cooling air is drawn into the racks and a standard temperature must be maintained in order to meet the equipment's operating requirements. The return air stream should not be used as an indication of the space temperature. Due to the highly concentrated loads common in data centers, correspondence between the return air temperature and the cooling air being drawn into data center

equipment is neither reliable nor accurate. Temperature sensors can be mounted on appropriately located pillars, walls, or rack faces, or even hung from overhead cable racks. For accurate and efficient control of the space, they must be located in the spaces where racks draw in cooling air, not located behind equipment in the heat exhaust or in pockets of relatively stagnant air at end walkways.

Centralized air handling systems simplify the implementation and maintenance of outside air economization. In most climates, outside air economization can cut cooling energy use significantly while also offering a degree of emergency cooling; see the Airside Economizer Design Guidelines document for further information. A centralized air handler system can usually be very economically configured to provide the access to outside air required for an economizer system. In many cases, the air handlers can be ordered with a standard, low-cost, economizer module that only requires minimal custom controls to provide economization to a data center. Some CRAC units now offer economization as an option, but even when units with this option are used, providing the outside air intake for each unit can be challenging due to the nature of a multiple distributed system. While both system types can usually benefit from economization, the centralized air handler approach is best configured to do so with minimal first cost.

A well-known benefit of centralized systems is reduced maintenance. Main mechanical components can be located in a single area outside of the data center envelope, where preventive maintenance and regular diagnostics to detect signs of impending failure require less time. Another benefit is that centralized systems simply have fewer parts to maintain. It is not unusual to substitute nine small fans in a multiple distributed unit system with two large fans in a centralized air handling system. The reduction in the number of parts has an almost directly proportional reduction in required maintenance, since the maintenance time requirements for one 60 hp fan are little more than for one 10 hp fan. Locating the cooling equipment outside of the data center room also immediately opens up additional floor area for IT equipment equal to the footprint of the smaller multiple units and their required maintenance clearances. A less quantifiable benefit is that maintenance personnel do not need to traverse the data center with tools to maintain the units. Also, the utilities associated with CRAC units (condensate lines, chilled water lines or refrigerant line sets, etc.) can be more easily located away from the data center equipment.

While a multiple-distributed unit system has a certain level of inherent redundancy, a central air handler system can be designed to offer the same or better levels of redundancy. A distributed unit often is ‘backed up’ by the surplus capacity of one or two adjacent units; every unit must have significant surplus capacity to provide this backup. The redundancy of a distributed system can disappear in areas where IT equipment loads creep above the design values, gradually leaving the space without appropriate redundancy. In contrast, a central air handling system can be designed to provide the same or greater level of redundancy with less total installed capacity. In a large data center, the addition of a single air handler offers N+1

redundancy across the full space, with an actual redundancy of N+2 or greater likely until the space is fully loaded to the design energy density (W/sf).

BENCHMARKING FINDINGS/CASE STUDIES

The pie charts below show the electricity consumption distribution for two data centers. Both are large facilities, with approximately equivalent data center equipment loads, located in adjacent buildings, and operated by the same company. The facility on the left uses a multiple distributed unit system based on air-cooled CRAC units, while the facility on the right uses a central air handler system. An ideal data center would use 100% of its electricity to operate data center equipment – energy used to operate the fans, compressors and power systems that support the data center is strictly overhead cost. The data center supported by a centralized air system (on the right) uses almost two thirds of the input power to operate revenue-generating data center equipment, compared to the multiple small unit system that uses just over one third of its power to operate the actual data center equipment. The trend seen here has been consistently supported by benchmarking data. The two most significant energy saving methods are water cooled equipment and efficient centralized air handler systems.

FIGURE 4
Electricity Consumption Distribution for Two Data Centers

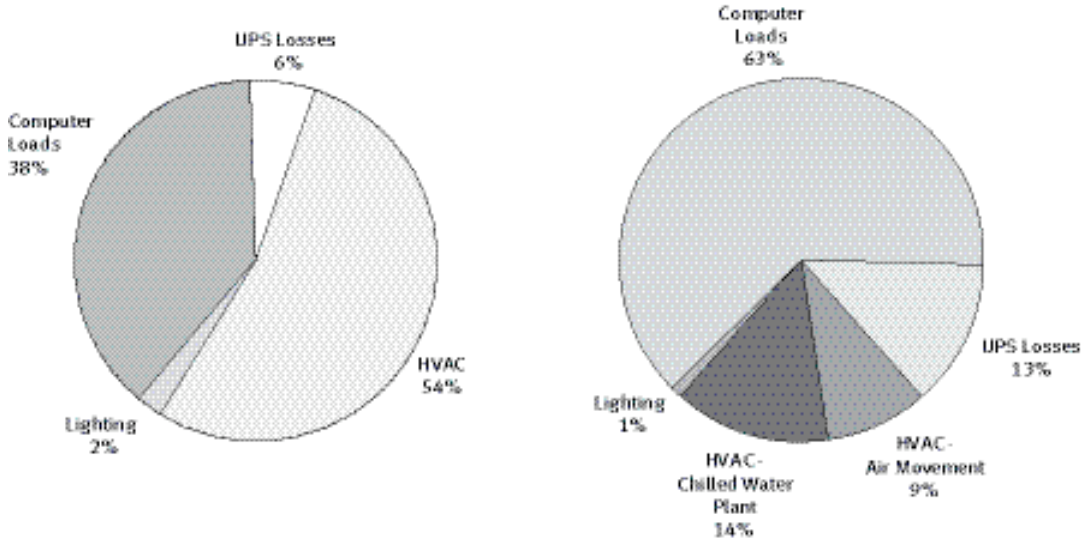

RELATED CHAPTERS

- Air-side Economizer
- Cooling Plant Optimization
- Free Cooling via Waterside Economizer
- Humidification Controls Alternatives

REFERENCES

1) The larger sizes of CRAC units commonly seen used in a multiple small unit system configuration will often use three 10 hp motors driving parallel fans.

2) This can be somewhat counterintuitive, that operating three fans will use less power than operating only two, but it is a natural result of power requirements being related to the cube of the volume. Considering only the pressure drop of the air handler, which typically is the highest pressure drop component in the supply system, operating three 30,000 CFM air handlers to provide 60,000 total CFM will require about half the total power of operating just two air handlers for the same total 60,000 CFM flow. This fan law relationship is not exact, but is a reasonable estimate of actual performance.

3) Assuming a typical fan system efficiency of 60% and a typical supply/return temperature difference of 17°F; low end of the range represents a 3 in. w.g. total pressure system, high end is a 6 in. w.g. total pressure system.

RESOURCES

- Data Processing and Electronic Office Areas, Chapter 17, ASHRAE HVAC Applications Handbook, 2003.

4. COOLING PLANT OPTIMIZATION

FIGURE 1
Cooling Towers from a Chilled Water Cooling Plant

Data centers offer a number of opportunities in central plant optimization, both in design and operation. A medium-temperature chilled water loop design using 55°F chilled water provides improved chiller efficiency and eliminates uncontrolled phantom dehumidification loads (see Humidification Chapter). The condenser loop should also be optimized; a 5-7°F approach cooling tower plant with a condenser water temperature reset pairs nicely with a variable speed (VFD) chiller to offer large energy savings.
A primary-only variable volume pumping system is well matched to modern chiller equipment and offers fewer points of failure, lower first cost, and energy savings. Thermal energy storage can be a good option, and is particularly suited for critical facilities where a ready store of cooling can have reliability benefits as well as peak demand savings. Finally, monitoring the efficiency of the chilled water plant is a requirement for optimization and basic reliable energy and load monitoring sensors can quickly pay for themselves in energy savings. If efficiency (or at least cooling power use) is not independently measured, achieving it is almost as much a matter of luck as design.

PRINCIPLES

- Design for medium temperature chilled water (55°F) in order to eliminate uncontrolled dehumidification and reduce plant operating costs.
- Use aggressive chilled and condenser water temperature resets to maximize plant efficiency. Specify cooling towers for a 5-7°F approach in order to economically improve chiller performance.
- Design hydronic loops to operate chillers near design temperature differential (or dT), typically achieved by using a variable flow evaporator design and staging controls.
- Primary-only variable flow pumping systems have fewer single points of failure, have a lower first cost (half as many pumps are required), are more efficient, and are more suitable for modern chillers than primary-secondary configurations.
- Thermal storage can peak electrical demand savings and improved chilled water system reliability. Thermal storage can be an economical alternative to additional mechanical cooling capacity.
- Use efficient water-cooled chillers in a central chilled water plant. A high efficiency VFD-equipped chiller with an appropriate condenser water reset is typically the most efficient cooling option for large facilities. The VFD optimizes performance as the load on the

compressor varies. While data center space load typically does not change over the course of a day or week, the load on the compressor does change as the condenser water supply temperature varies.

- For peak efficiency and to allow for preventive maintenance, monitor chiller efficiency.

APPROACH

For large data center facilities, a chilled water system served by a central plant is the most efficient approach to providing mechanical cooling. There are many design decisions that impact the efficiency of a central plant; the issues discussed here are selected due to their prevalence in typical data center operation.

USE NON-CONDENSING CHILLED WATER TEMPERATURE

A medium temperature chilled water loop supply temperature setpoint of 55°F or higher should be used to improve chiller efficiency and to prevent unnecessary and wasteful dehumidification. The temperature of the chilled water produced by the chiller plant has a significant impact on the efficiency of the chiller. Chillers are used to force heat from the chilled water loop into the condenser water loop. The temperature difference between the chilled water loop and the condenser water loop, also referred to as the "lift," impacts the chiller's efficiency – the bigger the lift, the lower the chiller's efficiency. Increasing the chilled water loop temperature helps reduce the lift the chiller must overcome. When centrifugal chillers are used, for every degree F increase in chilled water temperature, the chiller's efficiency improves by about 1 - 2%[1]. For example, an increase in chilled water temperature from 44°F to 54°F can be expected to cut chiller power use by 10-20%. In some cases by raising the temperature, the initial chiller selection can be altered since a smaller compressor and motor can be used on the chiller to provide the same capacity.

Medium chilled water temperature also prevents uncontrolled or 'phantom' dehumidification. Data centers are usually specified to operate with a space humidity setpoint that corresponds to a dewpoint of 52°F – 62°F[2] (controlled to around 50% RH). Any coil surface that is at a temperature lower than the space air's dewpoint will dehumidify to some extent, wasting cooling and harming humidity control of the space. To maintain humidity control, sensible cooling coils serving sensible cooling loads, such as data center cooling equipment, should be served with a chilled water temperature setpoint at or higher than the target humidity's dewpoint. Serving a data center load with a 44°F chilled water loop can result in continuous uncontrolled dehumidification of the space down to almost 35% RH (74°F drybulb, 46°F dewpoint) – near the bottom of the acceptable envelope for most facilities.

Dehumidification, when required, is best centralized and handled by the ventilation air system, while sensible cooling, the large majority of the load, is served by medium temperature chilled water at 50-60°F. This is standard procedure for controlling critical facilities that require tight humidity control to ensure profitable process yields, such as semiconductor photolithography cleanrooms or pharmaceutical manufacturing facilities. Assigning sole humidity control duties to the ventilation system offers both high efficiency and control accuracy. Ideally, a

direct expansion (DX) unit with a water-cooled condenser, potentially using chilled water return if tower water is unavailable for the condenser, can be used. An air-cooled DX unit may still be justified, since the load served by the inefficient air cooled unit is outweighed by the energy savings of keeping the central plant at a medium temperature setpoint. In some climates, it is appropriate to reset a central plant's chilled water temperature down during short periods that require dehumidification, but the additional energy cost should be considered. With a chilled water dehumidification reset, the chilled water plant efficiency will drop during peak power periods in dehumidification driven climates, resulting in higher demand charges for some data center operators.

DESIGN AND CONTROL COOLING TOWER SYSTEM FOR LOW CONDENSER TEMPERATURES

Reducing the lift to optimize chiller efficiency involves both a higher chilled water temperature and a lower condenser water temperature. Cooling towers with an approach of 5-7°F should be used and a reset implemented to maintain a condenser water temperature of 5-7°F above the ambient wetbulb temperature . Low air pressure drop cooling towers (typically draw-through) equipped with variable speed fans should be used and the maximum number of towers and/or cells should be operated in parallel at any given time. Minimum water flow requirements for proper tower media wetting typically determine the maximum number of towers and/or cells that can operate. The minimum allowed condenser temperature should be determined by the chiller manufacturer during selection, and usually is around 50 – 60°F. Use of a constant condenser water temperature of 70°F, or even higher, is a major cause of wasted energy in electrical chiller plants. Typically, chillers that allow a lower minimum condenser water temperature offer higher efficiency, a factor that should be considered when selecting a chiller; a 'more efficient' chiller may actually yield poorer energy performance than a slightly less efficient chiller that is better able to capitalize on low condenser water temperatures available at night or during the winter. The condenser water temperature actually available is determined by the climate, design and control of the cooling tower system.

Figure 2 shows the reduction in cooling energy, measured in kilowatts of power required to produce a ton of cooling (12,000 btu/hour), as the Condenser Water Temperature (CWT) is reduced. The single top blue curve is a baseline chiller operating without a condenser water reset. The curves below the baseline are the performance of an equivalent VFD chiller at various CWT. A VFD tends to allow the greatest utilization of low condenser water temperatures and provides better performance at part loads. Unlike most office applications, data centers have a large cooling load 24 hours a day. The large load even during cool outside air temperatures is a huge efficiency opportunity – as seen by the chiller performance data in Figure 2, chillers' power use can be literally halved by taking advantage of the cooler condenser water that can be produced by towers operating in cool conditions. Economizer and/or Free Cooling are other options to be considered that can optimize plant operation during mild outdoor conditions.

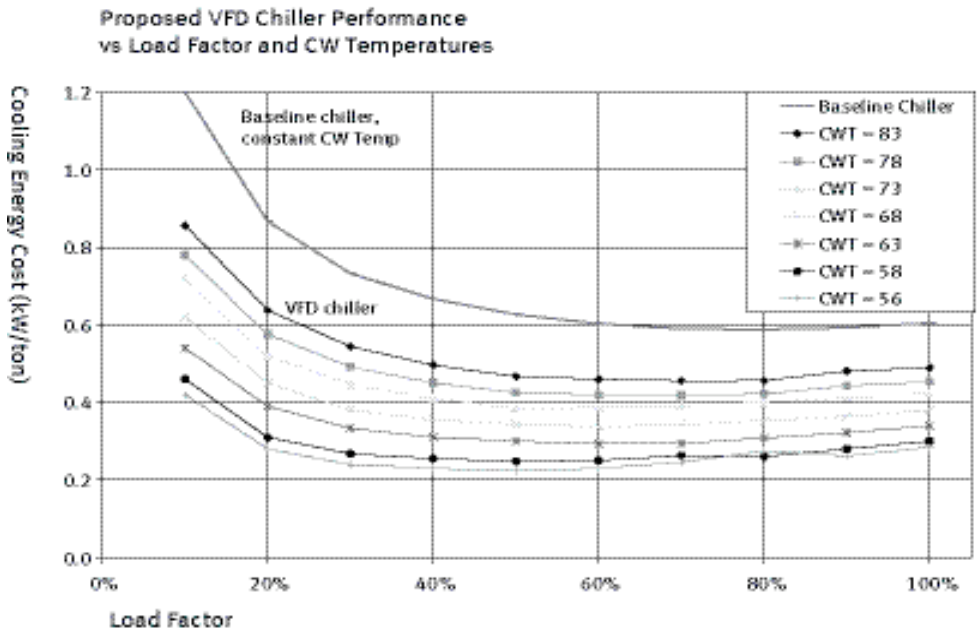

FIGURE 5
PROPOSED VFD CHILLER PERFORMANCE VS LOAD FACTOR AND CW TEMPERATURES

DESIGN COOLING WATER SYSTEM TO OPERATE NEAR DESIGN CHILLED WATER DELTA T AT PART-LOAD

Chillers are optimized to operate with a specific supply and return temperature difference for example, a return water temperature of 65°F and a supply water temperature of 55°F, which would be referred to as a "10°F delta-T." An efficient chilled water system design will operate the chiller at or near its design delta-T over all the expected load conditions, even at part load conditions. In data centers, the high-reliability cooling requirements of the space and the unpredictability of future demands tend to cause designers to specify oversized cooling systems that operate at part load. Full load is typically only reached when the data center is heavily built out and weather conditions result in high condenser water temperatures that add to the chiller's actual compressor load. There are several design steps required to achieve a good part-load delta-T system; most include eliminating unnecessary bypasses (particularly three-way coil valves), and using a pumping system that allows the chiller to operate at or near the design delta-T during the expected part load operation.

Variable flow pumping is required in order to allow the chiller to operate at design delta T during part load conditions. Traditional chiller design maintains a constant flow through the chiller, which inevitably results in the delta T (the difference between the warm water entering the chiller versus the chilled water leaving the chiller) being directly proportional to the load. A 50% load will result in a delta T that is 50% of the design value. This type of operation results in unnecessary pumping power use and often leads to inefficient chiller staging control, with an additional chiller (and condenser water pumps and towers) being staged on before the operating units are fully loaded.

Primary only variable speed pumping is quickly gaining in popularity. It is more efficient, costs less to build and has fewer points of failure. Figure 3 shows a standard primary pumping

configuration. Since data center facilities always maintain a high minimum load, in operation, the single bypass is always closed, indicating that none of the pumping energy is being wasted on powering bypass flow that never leaves the plant room. Much has been written in ASHRAE literature and other sources on this strategy. This strategy is very effective on modern chillers, whose electronic controls allow them to operate well with variable flow.

FIGURE 3
PRIMARY ONLY CHILLED WATER PUMPING SCHEMATIC

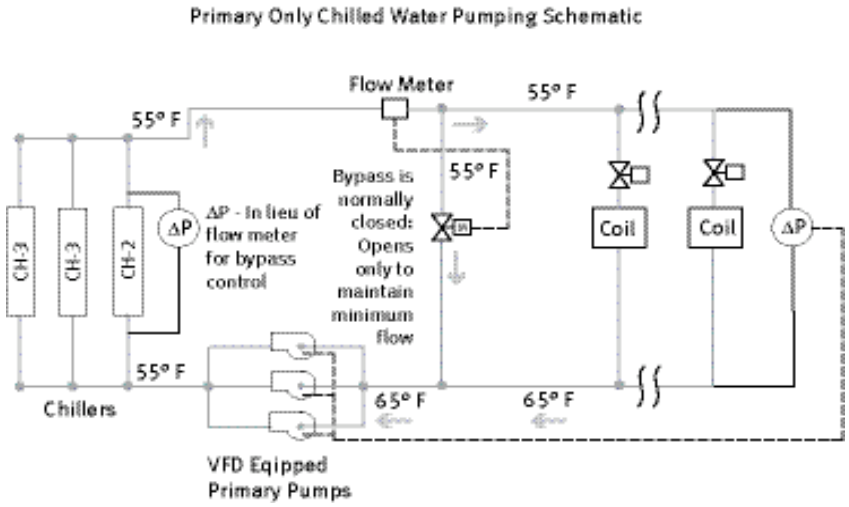

Often chilled water pumping costs can be reduced by 30% to 50%. The staging and control of a primary-only pumping system is no more complex than the traditional primary-secondary approach, but it is different. ASHRAE publications and white papers from chiller manufacturer (such as Trane) are good sources of information on this configuration.

THERMAL STORAGE

Properly designed thermal storage can offer an efficient method of increasing plant capacity without adding chillers and should be considered for critical facilities. Thermal storage has three main benefits. First, it takes advantage of off-peak electric rates which are typically significantly lower. Second, chilled water systems operate more efficiently when outside air temperatures are lower. Third, chilled water storage adds redundancy and can often substitute for one or more back up chillers. Water storage is preferred over ice because water is simpler, cheaper, more efficient, and more reliable – although it requires more space. Using multiple tanks will provide system redundancy and emergency backup cooling potential. Thermal storage can be linked to free cooling systems, such as water side economizers using cooling towers. In some dry climates, a comprehensive system design can combine free cooling and thermal storage to provide an almost full-time lifeline (maintaining space within the over-temperature but permissible range) cooling backup to mechanical cooling.

USE VARIABLE SPEED CHILLERS

Variable speed compressor chillers, often referred to as VSD or Variable Frequency Drive (VFD) chillers, currently offer the best performance for moderate to large data center loads (loads where centrifugal chillers are widely available). In order to capitalize on a VFD chiller's part load performance capabilities, a condenser water reset and low approach cooling tower system,

as discussed above, are required. While data center internal loads do not vary much, the actual load on the chiller compressor does vary as the weather impacts the condenser water temperature. It is the reduction of condenser water temperature during non-design day conditions (which are 99.9% or more of the year for a critical facility plant) that allow the VFD chiller to offload and run more economically at part load.

OTHER CHILLER SELECTION CONSIDERATIONS

Data on chiller efficiency should be requested from manufacturers during chiller selection and the chillers performance, including at part load, be a required part of equipment submittals. It is very important that the efficiency be compared at the condenser water temperatures and part-load ratios at which the plant is expected to operate. The standard ARI performance rating conditions assume a typical office building load that varies proportionally with the condenser water temperature/outside air conditions, which is not the case with data center operations. Internal data center loads are largely unaffected by the outside air conditions that impact the available condenser water temperature. Performance runs should be requested for a number of constant condenser water temperatures in order to allow a bin analysis that compares the chiller's efficiency over a year of operation serving the data center's actual predicted loads. Local weather data should be used for the bin analysis to capture the impact of varying condenser water temperature availability. VFD chillers tend to offer the best part load performance when compared to an equivalent (same condenser and evaporator heat exchanger configuration) constant speed chiller. While the chiller is not the only important element in an efficient plant design, it is the largest energy consumption; whole-plant design optimization and life cycle cost analysis should be utilized when selecting a chiller.

MONITOR PLANT EFFICIENCY

Chilled water plants consume a large amount of energy, yet are rarely monitored in any way in order to verify design and operating efficiency. The small initial expense of installing basic monitoring is usually necessary to achieve and maintain design efficiency. Frequently the efficiency of even a brand new chiller is degraded by minor equipment or installation defect, such as uncalibrated internal sensors, incorrect refrigerant charges, loose wiring harness connections, or non-optimal compressor mapping. It is common to find that chiller energy use is 25-100% higher than specified due to a minor problem that could easily be fixed if it were recognized. Finding and correcting such errors can provide an immediate payback for permanent monitoring equipment. Continuous monitoring also can help to rapidly diagnose operational problems or pending equipment failures.

At a minimum, a monitoring system should be provided that determine and display the chillers kW/ton performance in real-time. Monitoring of the full plant kW/ton offers additional optimization opportunity and can often be achieved for minimal additional cost. A true-power kW sensor, which incorporates voltage, amperage and power factor measurements, should be selected to monitor chiller power. Plant delta-T should be determined using a matched set of stable temperature sensors that provide an accuracy of +/- 0.36°F or better. The delta-T is

often in the range of only 4-9°F, so a closely matched and/or high accuracy pair of temperature sensors is required to achieve reasonable accuracy. For whole plant monitoring, VFD drives often offer an economical way to monitor power consumption of smaller, lower-priority loads such as pumps and towers.

Flow meters are the traditional weak link in plant monitoring equipment since high quality, high stability flow meters tend to be higher cost. Insertion-type flow meters with moving parts have been observed to foul unacceptably rapidly even in well managed closed-loop fluid streams. To provide diagnostic value, the flow meter must be reliable enough that 'odd' readings indicating a plant problem are not dismissed as an inaccurate flow meter. A flow meter based on electromagnetic induction or ultrasonic sensing provides the highest accuracy with exceptional long term stability and is recommended for reliable and accurate plant monitoring. A good practice is to ask the chiller manufacturer the type of flow meter used for their factory tests and use the same type for plant monitoring. This approach can eliminate finger pointing if the chiller is found to not be meeting submittal performance requirements as installed.

RELATED CHAPTERS

- Air-Side Economizer
- Free Cooling via Waterside Economizer
- Humidification Controls Alternatives

REFERENCES

1) This is a rule of thumb. For more accurate data, consult the manufacturer of your chiller for a comprehensive performance selection.

2) Allowing a 62ºF dewpoint dehumidification setpoint (65% RH) or even higher would be an efficient, best practice approach and maintain the space well within common equipment specifications. See Humidification Controls for additional information.

RESOURCES

- Variable-Primary-Flow Systems Revisited, Schwedler P.E., Mick, Trane Engineers Newsletter, Volume 31, No.4, 2002.
- Thermal Guidelines for Data Processing Environments, TC9.9 Mission Critical Facilities, ASHRAE, 2004.
- Data Processing and Electronic Areas, Chapter 17, ASHRAE HVAC Applications, 2003.
- Supervisory Controls Strategies and Optimization, Chapter 41, ASHRAE Applications Handbook, 2003
- ARI Standard 550/590- 2003, Water Chilling Packages Using the Vapor Compression

Cycle, Air-Conditioning and Refrigeration Institute, 2003

5. DIRECT LIQUID COOLING

Direct liquid cooling refers to a number of different cooling approaches that all share the same characteristic of transferring waste heat to a fluid at or very near the point it is generated, rather than transferring it to room air and then conditioning the room air. There are several approaches to implementing liquid cooling for data center equipment. One option, currently available as an option from many rack manufacturers, installs cooling coils directly onto the rack to capture and remove waste heat. The underfloor is often used to run the coolant lines that connect to the rack coil via flexible hoses. Many other approaches are available or being pursued, ranging from water cooling of component heatsinks to bathing components with dielectric fluid cooled via a heat exchanger.

Liquid cooling can service higher heat densities and be much more efficient than traditional air cooling air cooling. A high efficiency approach to cooling data center equipment takes advantage of liquid cooling's ability to efficiently move heat by transferring the waste heat from racks to a liquid loop as close to the heat source as possible, at the highest possible temperature. Currently, the most common approach is to use a chilled water coil integrated in some manner into the rack itself. Liquid cooling is adopted for reasons beyond efficiency; it can also serve higher power densities (W/sf). Energy efficiencies will be realized when such systems allow the use of a medium temperature chilled water loop (55-60°F rather than 44°F) and by reducing the size and power consumption of fans serving the data center.

In the future, products may become available that allow for direct liquid cooling of data center servers and equipment more directly, via methods ranging from fluid passages in chip heatsinks to submersion in a dielectric fluid. While not currently widely available, such approaches hold promise and should be evaluated as they continue to mature and are commercialized for the data center equipment market.

PRINCIPLES

- Water flow is a very efficient method of transporting heat. On a volume basis, it carries approximately 3,500 times as much heat as air.
- Cooling racks of IT equipment reliably and economically is the main purpose of the data center cooling system; conditioning the data center room without the rack load is a minor task in both difficulty and importance.
- Capturing heat at a high temperature directly from the racks allows for much greater use of waterside economizer free cooling, which can reduce cooling plant energy use by 70%[1] or more when operating.
- Transferring heat from a small volume of hot air directly off the equipment to a chilled water loop is more efficient than mixing hot air with a large volume of ambient air and removing heat from the entire mixed volume. A watercooled rack is equivalent to an almost perfect hot aisle/cold aisle configuration, where recirculation of waste heat is eliminated and

the 'hot aisle' heat exhaust can be run at very high temperatures with no impact on workers or equipment in the data center room.

- Direct liquid cooling of components offers the greatest cooling system efficiency by eliminating airflow needs entirely. When direct liquid component systems become available, they should be evaluated on a case-by-case basis.

APPROACH

A cooling system using water as the heat transport medium can conduct about 3,500 more heat on a volume basis than a system using air. For example, one gallon of water per minute (not a particularly large water flow) will remove the thermal equivalent of a kilowatt with a seven-degree temperature rise. By comparison, removing a kilowatt with air requires passing 140 cubic feet (1,050 gallons) of air per minute through the rack at a 23-degree temperature rise. In addition, the low temperature rise of the water flow would allow for relatively warm water to be used for cooling; modern cooling equipment is more efficient when it is producing warmer temperatures, and use of free cooling or a waterside economizer could eliminate high-energy-use compressor cooling from the cooling system altogether. While liquid has a huge advantage in heat removal efficiency, and was a common approach in the dawn of the mainframe era, the use of liquids in data centers is a hotly debated issue. However, as chip power continues to rise and equipment loads continue to increase in density, liquid cooling becomes increasingly inevitable as a strategy for cooling the rapidly increasing equipment load density (measured in W/sf).

FIGURE 1
SERVER RACK PROVIDED WITH AN INTEGRATED CHILLED WATER COIL

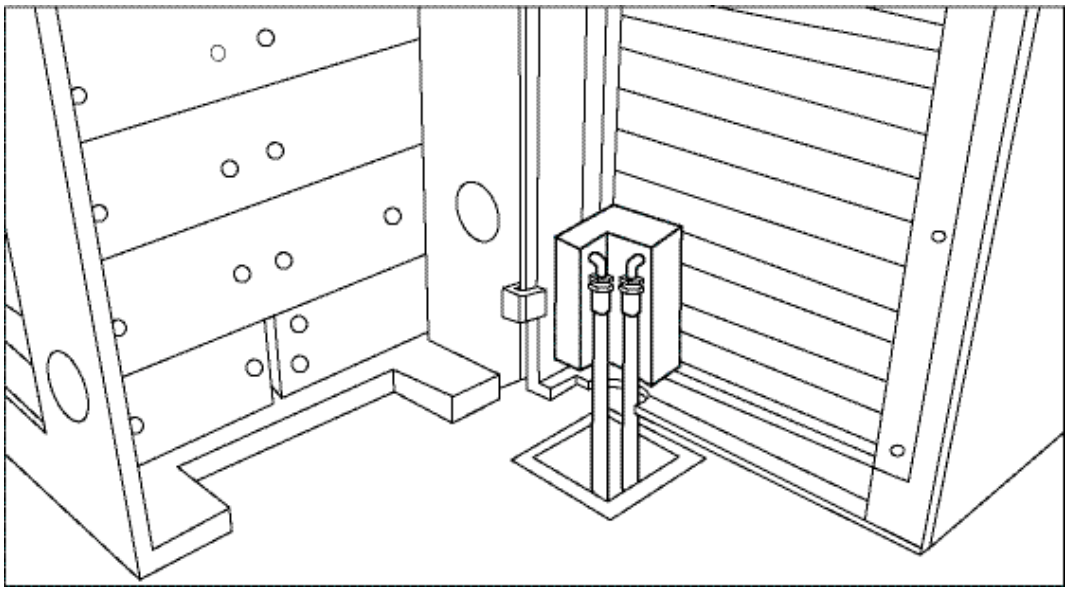

One mature approach to direct liquid cooling is to use a rack provided with an integrated chilled water coil. Figure 1 shows a rack with an open back panel. To the left would be the backplates of the servers installed in the rack. To the right is a hinged door that contains a cooling coil. The flexible hoses supply the coil in the hinged door with cooling water. In use,

hot air from the backs of the server equipment would be blown out the coil (or drawn through by a fan) allowing it to be cooled while it is at its highest temperature. When the rack is opened for service, it would simply revert to using room air for cooling – the room should be designed to service the load of a few racks running open to accommodate expected service accesses.

The first cost of the coil option and additional required piping should be compared to an analysis of expected energy, mechanical equipment, and/or architectural savings. The most significant energy savings are achieved when the liquid cooling system is paired with a waterside economizer or free cooling. Direct liquid cooling directly replaces air handler or Computer Room Air Conditioner (CRAC) capacity. A data center that fully utilizes direct liquid cooling can be served like a standard office space by the airside system. Architectural savings are less common, but can be significant if the use of water cooled racks allows for a reduction in the square footage of the data center by allowing a higher equipment density. In small data centers serving existing buildings, watercooling may allow a standard 'office' space to be used as a data center without the need for costly space retrofits.

Another configuration, shown in the figure below, consists of a water cooling coil placed in the bottom of the rack. Small rack fans circulate hot air from the servers down, through the coil, and then bring the cooled air back to the server intakes. Rack-integrated water cooling places the coil directly in the path of the high temperature air exiting the equipment. This positioning greatly reduces the volume of air that must be moved by removing the waste heat before it is mixed into the room. Rather than cooling 100 CFM of 78°F air down to 60°F, the coil can cool 50 CFM of 110°F air down to 74°F to remove the exact same amount of heat. Since the air does not have to be cooled to as low a temperature to compensate for mixing and heat recirculation, warmer water, possibly from a free cooling system (see Free Cooling via Waterside Economizer), can be used.

Integrating cooling into the racks can allow for the ideal implementation of the hot aisle/cold aisle configuration. Integrated chilled water coil racks are currently available on the commercial market and use common, mature technologies. The designs offered vary and range from simple chilled water systems with integrated fans to systems that combine cabinet cooling with component coolers using manifolds and heat exchangers. The additional components increase the first cost of the system, but can significantly reduce operating costs. They also offer the ability to support higher load densities than is possible with traditional air cooling systems, which can allow the reduction of costly data center square footage.

Many of these systems are designed with sophisticated leak and condensation detection and mitigation systems that address equipment protection concerns. They can be supported with redundant power sources and redundant fans. Some are being designed to accommodate rapid exchange or replacement, in order to further minimize the effects of downtime. A common design goal is to support significantly higher power densities than are possible with standard room air cooling systems, not only within the rack but over the entire data center floor since

FIGURE 2
SCHEMATIC OF INTEGRATED COOLING IN SERVER RACK

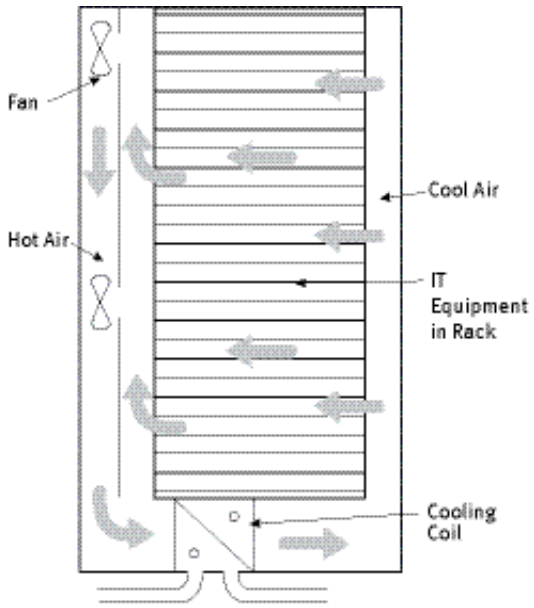

large aisles for cooling air are not needed, nor are large, freestanding CRAC units. Designs should be evaluated to ensure that efficiency has not been sacrificed; in particular, the use of multiple small, constant speed fans or small liquid chillers between the house chilled water plant and the rack coolant loop can reduce system efficiency.

Chilled water systems are common in larger modern and legacy data centers, and the cost of installation of chilled water cooled equipment systems in these types of centers is relatively low. Where existing chilled water systems support CRACs, plant capacity would be unchanged, as the number of air conditioners would naturally be reduced with the introduction of integrated cabinet cooling. However, in centers with either no chilled water, or systems of limited capacity, the initial cost of installation would be significant.

Chilled water piping in the data center envelope can be problematic in facilities that require frequent reconfigurations. The use of flexible hoses to connect to rack coils provides a degree of configuration flexibility and allows racks to be moved more easily. Cooling water lines are typically run underfloor to reduce the danger of leaks or damage.

In many temperate climates, data center loads may not require any dehumidification and should be served by medium temperature chilled water to ensure dry coils and avoid uncontrolled dehumidification. Medium temperature chilled water (55°F or higher) should be used to eliminate any condensation potential and avoid the need for insulation on hookup lines. To maximize efficiency, medium temperature chilled water should be supplied from a loop equipped with free cooling (waterside economizer), either directly or through a close-approach heat exchanger. Depending on the climate at the location, an evaporative free cooling system could supply the majority of a data center's cooling hours without the need for an energy intensive chiller. The ideal opportunity is a site where the mechanical cooling can be eliminated entirely, potentially making the more efficient liquid cooling have the same or even a lower first cost than a standard system.

One advantage of the cabinet mounted solution is that water is kept away from the electrical components, yet it is still close enough to the equipment to provide an efficient method of delivering cooled air. The fans used to move air across the cooling coils and through the equipment can be made both redundant and efficient, as they can be actively controlled to operate at variable speeds based on the rack's actual cooling demand. The air can be evenly distributed or actively controlled across the rack, so that the temperature stratification within

the rack that is so evident in traditional raised floor or overhead distribution systems is virtually eliminated. Poor airflow leading to stratification in the cabinet has been cited as causing three times the rate of failure for equipment mounted at the top of cabinets[2].

While a typical data center boasts the ability to cool loads of up to approximately 5 kW per rack, and careful airflow management can increase that capacity further, integrated water cooled rack systems have demonstrated the ability to efficiently cool loads up to 20 kW per rack[3]. The advantages are obvious; end users are able to cool far more equipment in smaller spaces while reducing their operating costs for an equivalent amount of processing capacity. The more even distribution of cool air through data center equipment increases the useful life of that equipment. Further, maintenance costs of the cabinet integrated systems can be significantly lower than the costs of maintaining traditional air systems capable of supporting equivalent heat loads.

The most efficient possible cooling approach would be to directly cool the individual components with liquid. Research is currently investigating the feasibility of liquid cooling modern servers with methods ranging from water passages to in-chip heatsinks (utilized in the past and with some commercial availability[4]) to spraying the electronics directly with an inert dielectric liquid refrigerant[5]. At this time, products using these methods are not widely available on the general commercial server market; however, they are being actively commercialized and hold significant promise for efficient cooling.

Beyond a traditional coil mounted on the rack to cool the exhaust airstream, other approaches under development tend to keep liquid well away from the chips by using various non-liquid-based methods to move heat from the electronics to the liquid in an off-board or even outside-the-rack location. Heat pipes, carbon fibers, and a few other non-liquid media can transfer heat to a backplane liquid system, eliminating the need to have the water near the chips. Separating the two helps relieve anxieties about mixing water, or any liquid, with electricity. Direct liquid cooling of components is not currently available in the general, non-supercomputer, server market, but is also an approach that has a high potential for realizing significant energy saving opportunities and may be used in future commercial offerings. Direct liquid cooling through the use of chip heatsinks with liquid passages, while not common commercially at the moment, is a mature technology that can offer significant efficiency savings if available.

RELATED CHAPTERS

- Air Management
- Air-side Economizer
- Free Cooling via Waterside Economizer

REFERENCES

1) Baseline study measurement of critical facility plants indicate that chiller power represents approximately 70% of total energy consumption.

2) Data Processing and Electronic Areas, Chapter 17, ASHRAE HVAC Applications, 2003.

3) Christopher Kurkjian, and Doug McLellan, High Density Server Cooling, Engineered Systems, 10/12/2005, http://www.esmagazine.com/CDA/ArticleInformation/coverstory/BNPCoverStoryItem/0,2500,158187,00.html

4) http://www.water-cooled-cpu.com/en/index.html

5) http://www.spraycool.com/flash/WP_Datasheet_Fundamentals.pdf

RESOURCES

- http://www.esmagazine.com/CDA/ArticleInformation/coverstory/BNPCoverStory-Item/0,2500,158187,00.html
- Gain, Bruce, Processor, Throwing Water at the Heat Problem, "http://www.processor.com/editorial/article.asp?article=articles%2Fp2736%2F31p36%2F31p36.asp&guid=&searchtype=&WordList=&bJumpTo=True
- IBM Heat eXchanger water cooled rack, http://www-03.ibm.com/servers/eserver/xseries/storage/pdf/IBM_RDHx_Spec_Sheet.pdf
- Knurr CoolTherm water cooled rack, http://www.knuerr.com/web/en/index_e.html?service/faq/faq_cooltherm.html~mainFrame
- Fujitsu Siemens Primecenter water cooled rack, http://vilpublic.fujitsu-siemens.com/vil/pc/vil/primergy/building_blocks/rack_systems/lc_racks/ds_primecenter_lc.pdf
- Thermal Guidelines for Data Processing Environments, TC9.9 Mission Critical Facilities, ASHRAE, 2004.

6. FREE COOLING VIA WATERSIDE ECONOMIZER

FIGURE 1
Use of a Cooling Tower to Provide Waterside Economization

Data centers present an almost constant, 24-hour, internal cooling load. Free cooling can be provided via a waterside economizer, which uses the evaporative cooling capacity of a cooling tower to indirectly produce chilled water to cool the data center during mild outdoor conditions (particularly at night in hot climates). While a bin analysis using local weather data is required to properly assess the potential, free cooling is usually best suited for climates that have wetbulb temperatures lower than 55°F for 3,000 or more hours per year. It most effectively serves chilled water loops designed for 50°F and above chilled water, or lower temperature chilled water loops with significant surplus air handler capacity in normal operation. Often, existing data centers can capitalize on redundant air handler capacity with chilled water temperature reset controls to retrofit free cooling.

At least 3000 hours per year where wet bulb temperature is below:	Applicability
55°F	Free Cooling Specified Air Handlers or CRACs; Many Low Temperature Systems with Appropriate Chilled Water Reset
45°F	Retrofit Reusing Most Existing Air Handlers Designed for 42°F Chilled Water Without Need for Reset

FIGURE 2
Free Cooling Applicability

PRINCIPLES

- While free cooling is operating, chilled water plant energy consumption costs are cut by up to 70%[1].
- Data centers require cooling 24 hours a day every day of the year – even when it is cold outside. This makes data centers very well suited to waterside economization. For example, in San Jose, free cooling would be expected to operate for over a third of the year, significantly reducing cooling bills and noticeably reducing chiller run hours and maintenance costs.
- Free cooling utilizing a waterside economizer can usually be economically retrofitted to existing chilled water cooled facilities.
- Isolation between the space air and outside air is not impacted by waterside free cooling, making it an alternative to airside economization when this is a concern.
- A flat plate heat exchanger is used to isolate the chilled water loop from the open tower condenser water to prevent fouling of coils.
- A low approach temperature cooling tower plant design is critical for best results. Use of redundant tower capacity can provide low approach temperature operation in a high-reliability 'spinning reserve' operation configuration.
- A traditional chiller is used to provide cooling during hot periods and as an always-available emergency backup. For a portion of the year, free cooling increases reliability by offering a non-compressor based backup to the traditional chiller, particularly at night when plant monitoring and staffing are liable to be lower.

APPROACH

Free cooling operates on the principle that during cool weather conditions, particularly at night, data center cooling loads can be served with chilled water produced by the cooling tower alone, entirely bypassing an energy intensive chiller. In this scenario, the cooling tower produces low temperature water. A heat exchanger is used to cool the building loop while keeping it clean and isolated from the relatively dirty cooling tower water. Free cooling reduces or eliminates chiller power consumption while efficiently maintaining strict temperature and humidity requirements. Other approaches to free cooling include: closed-circuit cooling towers (dry coolers) with glycol for climates with extended freezing conditions; earth coupled heat pumps; and/or independent free cooling loops serving dedicated free cooling coils in air handlers or Computer Room Air Conditioners (CRACs).

Data centers often have redundancy in their cooling tower plants. Through the use of VFDs and common condenser water header systems and/or sumps, the redundant tower capacity can be used to achieve a lower approach temperature. With variable speed fans, it is efficient to operate as many towers as the tower minimum flow requirements allow, maximizing the natural convective cooling while achieving a lower approach capability.

The majority of server computer equipment has a specified allowable operating range of 20% to 80% relative humidity, but ASHRAE recommends an operating range of 40 to 55%[2] and many data centers control the humidity more tightly. The minimum humidity is often controlled to prevent the risk of static discharge damage and the maximum humidity controlled to ensure a non-condensing environment and human comfort. Free cooling is well suited to such facilities, since it allows the space to be fully isolated from the exterior environment's humidity yet still reject heat without use of energy intensive compressor based cooling.

Free cooling can also offer an additional level of redundancy by providing a non-compressor cooling solution for portions of the year. In particular, free cooling can often provide a backup to compressor chillers during cooler nighttime hours when plant staffing may be lower. When the weather allows, free cooling replaces the complex mechanism of a chiller with a simple, non-mechanical heat exchanger.

Use of 'medium temperature' chilled water, in the range of 50°F and higher, maximizes the potential savings from a free cooling system. An efficient data center system is likely to be designed for use of chilled water in this temperature range already, since use of traditional (for office systems) 44°F chilled water is likely to result in uncontrolled dehumidification in high load areas (leading to either artificially low humidity or wasteful simultaneous dehumidification/humidification operation). A typical data center maintained at 72°F and 50% RH is susceptible to uncontrolled dehumidification when the supply water temperature is lower than 52°F.

New data center facilities that are specified for medium temperature chilled water can be designed. Medium temperature chilled water systems also reduce energy use and sometimes compressor sizing of normal chillers at the cost of somewhat larger coils and piping. In existing facilities, free cooling operation can be retrofitted and optimized through the use of an aggressive chilled water temperature reset. Typically, data centers have surplus air handler capacity, in the form of redundant units or overdesigned systems. The surplus capacity can allow the use of higher temperature chilled water than the original design through the use of a reset. A chilled water reset increases the chilled water supply setpoint in order to take advantage of the surplus capacity. The most efficient reset strategy is to have the central control system monitor the control valves on each air handler and/or CRAC and use the highest valve command (i.e., the most open valve) as the input to a control loop. The control loop resets the chilled water temperature, up or down, until the most open valve is at 90-95%. This approach automatically maximizes the use of installed coil capacity and automatically accommodates variations in data center load – increased cooling needs will result in the chilled water temperature being lowered if required to meet the load.

A cooling tower used for free cooling should be selected to provide an approach temperature (Leaving Tower Water Temperature minus Wet Bulb Temperature) between 5 and 7°F (see Cooling Plant Optimization chapter). A lower approach temperature generally results in a physically larger tower since more surface area in the tower is required for the heat rejection

process. A variable speed drive (VSD) for the fan motor should be used in a free cooling tower to minimize on-off cycling and maximize controllability and energy savings.

Figure 3 below shows a typical free cooling setup. Note that the free cooling heat exchanger is placed inline with the chillers. This allows for free cooling to support a portion of the load while the chillers can provide the last few degrees of cooling (referred to as integrated operation). An critical design element is that the chillers automatically and seamlessly provide 100% backup to the free cooling system. Failure of the free cooling system will simply result in the chillers seeing the full data center load; no control loop is required to switch between free cooling and chiller cooling, ensuring no loss in reliability from the addition of a free cooling system.

FIGURE 3
FREE COOLING LOOP SCHEMATIC

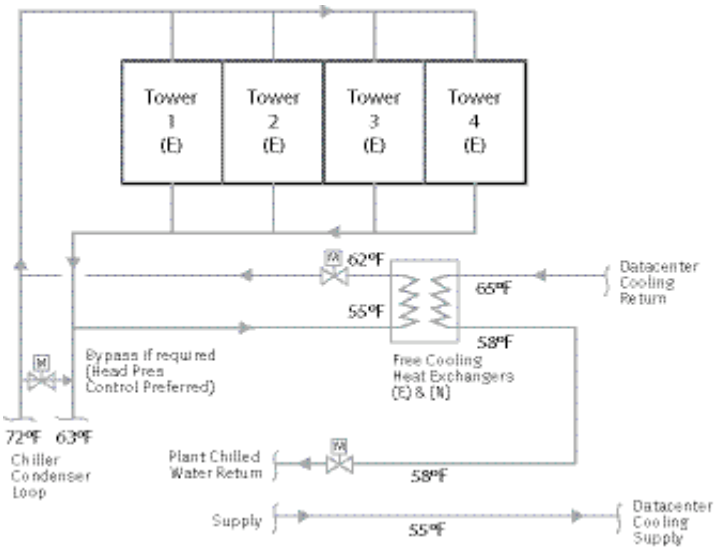

Free cooling requires that the cooling tower produce low temperature water, often lower than a chiller will accept for condenser water. There are two common design approaches to address this concern. One approach is to hydraulically isolate a tower and dedicate it to free cooling only. This is the best approach, but requires careful piping configuration and control valve placement and operation. A redundant backup tower can be provided with automatic isolation valves and used for free cooling. Since free cooling operates during low temperature weather conditions, the chilled water plant load is often low enough that even non-backup towers are available for free cooling use provided the proper automatic valving and a control sequence that gives the chillers priority for tower use (in case of free cooling failure) is implemented.

The other common approach is to share a single condenser water loop and towers between free cooling and the chillers by running the loop at a low, free cooling temperature and providing a bypass at the chillers, as shown in Figure 3. Locating the tower bypass (a standard feature in cooler climates) at the chiller end of the loop instead of at the cooling tower brings low temperature water into the main plant area, in many cases greatly reducing the cost of piping to implement free cooling. The bypass is used to mix warm condenser water leaving the chiller with low temperature condenser water to produce a suitable condenser water supply

temperature – the standard tower bypass control loop. In cooler climates, a tower bypass is usually located directly next to the tower plant with an identical control algorithm to allow starting up the chillers during cold temperatures with a cold tower sump and chiller operation in very low temperatures. This approach is popular in retrofit situations or where the pipe run to the cooling towers is too long to economically allow a second set of pipes for free cooling. Some efficiency is lost by producing lower temperature water for the chillers than is used, but typically this is far outweighed by the reduced chiller compressor energy consumption.

Added costs for a waterside economizer result from controls, heat exchangers, and piping. Some installation will also incur additional costs for additional plant floor space or an additional pump. In typical critical facilities installation, no additional cooling tower capacity is required since the non-critical free cooling system components do not require any redundancy.

BENCHMARKING FINDINGS/CASE STUDIES

The psychrometric chart shown in Figure 4 represents a typical year's weather conditions in San Jose, California. Each hour is plotted as a point on the chart. A data center facility designed to utilize free cooling with 53°F water with a 10°F differentiated temperature could operate without any chillers for 2,800 hours per year. Partial cooling allows a total of 6,200 hours of cooling assistance for a predicted annual reduction in chiller energy use of 52%. A drier climate would yield larger savings.

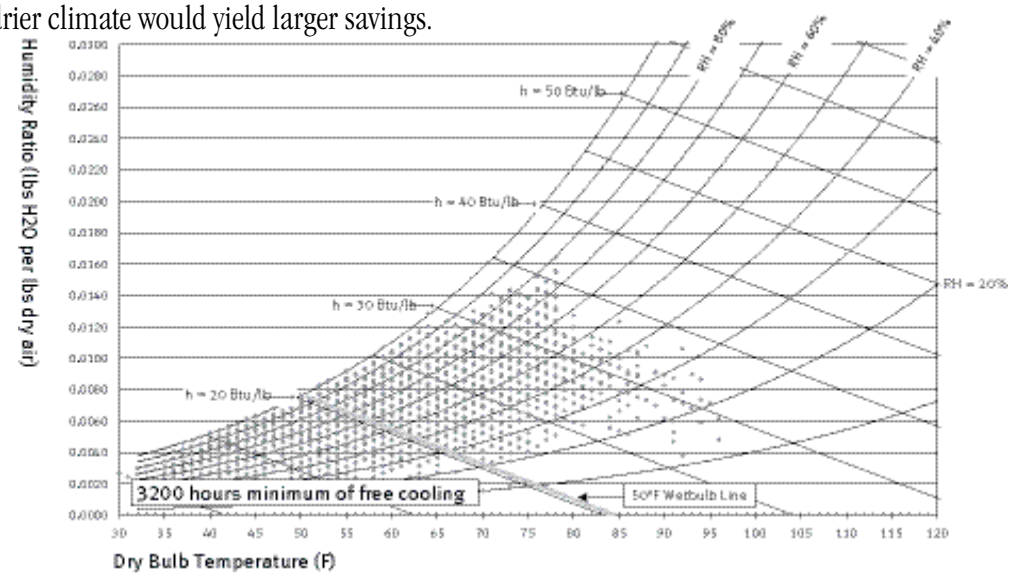

FIGURE 4
PSYCHROMETRIC CHART FOR SAN JOSE, CALIFORNIA

RELATED CHAPTERS

- Air-side Economizer
- Centralized Air Handling
- Cooling Plant Optimization

REFERENCES

1) Baseline study measurement of critical facility plants indicate that chiller power represents approximately 70% of total energy consumption.
2) ASHRAE, TC 9.9, Thermal Guidelines for Data Processing Environments

7. HUMIDIFICATION CONTROLS

Data centers often over-control humidity, which results in no real operational benefits and increases energy use, as humidification consumes large amounts of energy. Humidity controls are frequently not centralized. This can result in adjacent units serving the same space fighting, with one humidifying while the other is dehumidifying. Humidity sensor drift can also contribute to control problems if sensors are not regularly recalibrated. Low-energy humidification techniques can replace electricity-consuming steam generators with an adiabatic approach that uses the heat present in the air or recovered from the computer heat load for humidification. Ultrasonic humidifiers, wetted media ('swamp coolers') and microdroplet spray are some examples of adiabatic humidifiers.

PRINCIPLES

- Humidification is very energy intensive. Dehumidification incurs even higher energy costs since the air is usually cooled to below 45°F to condense out water and then is reheated by an electric heater to ensure the supply air is not too cold.
- Modern servers do not require extremely tight humidity control and typical data centers cannot actually provide tight humidity control due to sensor drift.
- Centralized humidity control can keep all units serving the same space in the same humidification mode, eliminating simultaneous humidification/dehumidification common when using independent Computer Room Air Conditioners (CRACs) controls.
- Utilize adiabatic humidifiers and evaporative cooling for humidification whenever possible. Waste heat in the return air stream can be used to drive adiabatic humidification 'for free' when the outside air is too cold for adiabatic systems.
- Computers do not emit humidity, nor do they require 'fresh' outdoor ventilation air, so a well designed and controlled data center should have minimal humidity loads. Ensure outside air economizer, if present, is properly controlled to prevent unnecessary humidification loads. Provide a tight architectural envelope and optimized pressurization controls to minimize humid or dry outside air (OSA) infiltration.

APPROACH

Humidity control is very energy intensive and should be minimized. Dehumidification requires that the air be cooled to such low temperatures that electric reheat is commonly used to maintain a supply temperature around 50-55°F. Removing moisture requires that a large amount of energy be removed, about 1,000 Btus for every pint of moisture condensed. One thousand btus is equivalent to the energy required to heat all the air in a 100 square foot room from 32°F to 100°F . Beyond the large energy cost of moving that quantity of heat, using energy to simultaneously heat and cool air is an inefficient but common part of the dehumidification process. Standard humidification also adds to the data center cooling load, adding in heat to essentially boiling water and produce steam. Finally, in a chilled water plant

system, dehumidification can reduce the efficiency of the entire plant by requiring a lower chilled water temperature.

Humidity requirements can and should be set to meet the actual equipment specifications. Modern server equipment is designed to operate in a wide humidity range, usually 20% to 80% relative humidity (RH). A somewhat tighter control band is expected in order to account for static control and allow for sensor error; however, specifying a tight humidity band, such as +/-5% is very rarely justified and equally rarely achieved even when specified. In a configuration where multiple Computer Room Air Conditioner (CRAC) units are used with individual humidity control, it is common to find one CRAC unit humidifying right next to another that is dehumidifying. This happens when humidity ranges are tight as the humidity sensors inevitably drift over time. Precise humidity control is not required by data center equipment, so poorly operating humidity control systems, and their associated energy wastage, are often unnoticed or ignored. This and similar humidity control problems can be solved by:

- Using the widest reasonable humidity range, such as 30% - 70% RH
- Recalibrate humidity sensors on a regular basis as appropriate for the control band
- When using multiple CRAC units to control humidity use a centralized control signal so that all units serving a space are in the same mode – humidification or dehumidification.

Modern data centers do not have sensitive punch card processing mechanisms as in the past, and rarely if ever require tight humidity control, with the ASHRAE recommended control range 40-55% RH[2]. To realize energy savings, the humidity controls need to allow the humidity to truly float within an appropriate range. Depending on the climate, good efficient humidity control will maintain humidity at the low end of the allowed range during dry or very cold weather conditions and at the high end of the range during hot and humid weather conditions. A controlled swing of humidity level with seasonal outdoor air conditions is a sign of an effective and efficient humidity control algorithm.

The minimum humidity setpoint is typically set at a level high enough to prevent static discharge risks. A humidity level of 20% RH has been used in facilities that use strict operator grounding procedures (mandatory wrist grounding strap usage) and do not have tape drives. A higher minimum humidity of 30% RH has been used for standard data centers without personnel grounding protocols, and is recommended for the most efficient operation. For the most sensitive data center equipment, particularly in facilities with extensive use of tape drive systems, ASHRAE Thermal Guidelines for Data Processing Environments recommends a minimum humidity of 40% RH. A maximum humidity of 55% - 70% RH is usually appropriate, with 55% as the conservative upper range recommended by ASHRAE. Condensation is only a threat on surfaces that are actively refrigerated, not on the ventilated heatsinks used in servers. Unless there are unusual design features, such as uninsulated chilled water pipes in the data center or an uninsulated metal roof in a harsh winter climate, condensation is rarely a concern.

Barring high infiltration, most cooling systems will automatically control the maximum humidity to about 60% to 70% RH due to the fundamental nature of their operation. Supply air at 60°F saturated with as much moisture as it can possibly carry has a humidity of only 66% RH at 72°F (room temperature). Over-humidity problems are more likely to indicate excessive infiltration (i.e., a broken outside air damper) or malfunctioning humidifiers than a dehumidification control or capacity problem.

Humidity sensors are possibly the least stable sensor technologies in common HVAC use. Without regular calibration, attempts to control to a tight humidity band (1% - 2%) are not meaningful. The problem with expending large amounts of energy to maintain a 50% RH +/-1% when the humidity sensor is off by at least 5% is clear. Even with a data center-appropriate control band such as 30-60% RH, humidity sensors should be regularly calibrated. Humidity control consumes a large amount of energy and requires costly equipment to achieve; if the first cost was justified, then calibrating the humidity sensors should be considered as much an operating cost as the electric bill in order to ensure the first-cost investment is actually providing the desired space condition.

Sensors that are out of calibration, combined with a tight control band and multiple CRAC units with their own independent humidity control, often lead to wasteful simultaneous humidification/dehumidification. There are very few humidity loads in a data center and little need for localized 'spot' dehumidification. Even in large rooms, all the CRAC units serving the room should be in the same mode, either adding or removing moisture. Yet it is not uncommon to find CRAC units serving the same space with returns and supplies less than 10' apart where one is boiling water to add humidity to the air while the other is subcooling and reheating to remove humidity. Without any significant point dehumidification loads, such operation indicates an out-of-control space in addition to significant energy waste. It should be corrected immediately by either widening the humidity band to what the units are actually capable of sensing or calibrating the humidity sensors.

In situations where units have been added over time, there can also be a significant amount of excess humidity control capacity, complicating control. Usually, the initial cooling system is sized to meet the humidity load of the full envelope and ventilation. The addition of equipment load through a gradual buildout may result in the need for additional CRAC units; however, no additional humidity load is needed since data center equipment adds no humidity load. If the added units are all installed with humidity control, perhaps intended to be identical to the original equipment, then there can be a significant amount of surplus humidification capacity. Too high a capacity can result in unstable humidity and inefficient control if systems continually overshoot the setpoint. The control problem can be greatly magnified when many independent CRAC control loops are involved. Many humidifiers also tend to complicate control due to significant lag between the call for humidification and the steam generator warming up.

Low-energy humidification techniques replace electric steam generators with ultrasonic humidifiers, microdroplet spray, and other low-energy technologies. This is evaporative cooling with no reheat; it provides evaporative cooling on the way to reaching the humidity set-point. Most importantly, this eliminates the parasitic heat gain of generating steam to humidify the space while providing "free cooling" in the process of reaching the humidity setpoint.

Most manufacturers of computer room cooling equipment only sell isothermal humidifiers that vaporize water, some with heat lamps others with immersion heaters. Using steam to humidify a data center is a net energy loss because it uses electricity (or gas) and adds heat to the conditioned air—one step forward, two steps back. Adiabatic humidification, designed to eliminate any risk of droplet carryover, is a more efficient alternative. It is the process of injecting water into an incoming air stream using natural convection or forced spray from nozzles. The exiting air has both a lower drybulb temperature (due to evaporative cooling) and a greater moisture content.

When using an air-side economizer in cold conditions, humidification might be required to maintain minimum humidity specifications. Adiabatic humidification can use significantly less energy compared to electric or gas humidifiers. As opposed to electricity or gas driven isothermal humidifiers, adiabatic humidifiers can be designed to use the waste heat from the data center itself to provide the energy required for water evaporation. See the Economizer best practices for a discussion of using a small amount of hot return air bypassed through an adiabatic humidifier into the supply air to provide heat recovery humidification. Assuming humidification is required, adiabatic humidification is the answer to the problem of how to humidify dry winter air and provide free summer evaporative cooling, dependent on the local climate. The figure on the next page shows a sample configuration for a cold climate.

In some cases, the humidification need may be so small that it is more economical to eliminate humidification and simply lockout the economizer during very dry outdoor air conditions. An economizer lockout should be based upon the outdoor air's dewpoint temperature or humidity ratio, not the outside relative humidity. As its name implies, relative humidity is relative to the temperature. Dewpoint or humidity ratio are measurements of the water in the air that do not change when the air is heated or cooled to the data center's operating temperature. A common economizer lockout could be to override the economizer damper to minimum OSA when the outdoor dewpoint is below 40°F, a condition that would correspond to summer peaks of 95°F drybulb temperature and 15% humidity and winter lows below 40°F drybulb.

FIGURE 1
SAMPLE COLD CLIMATE CONFIGURATION

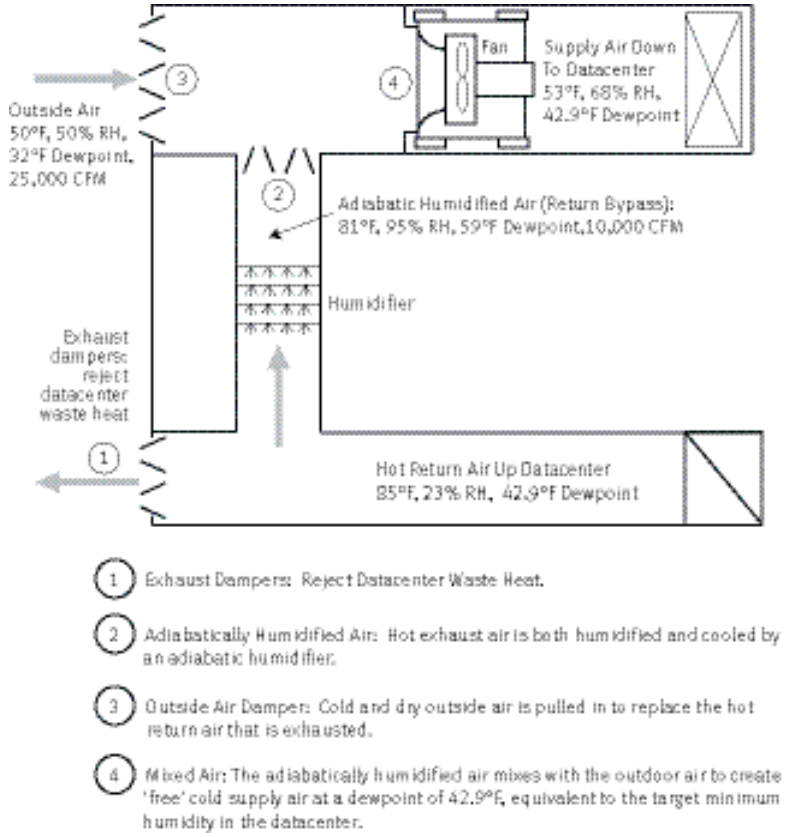

When a data center uses a properly controlled economizer, the latent load (dehumidification or humidification) should be minimal. There are few sources in a data center that add or eliminate humidity – a high humidification load is likely due to economizer control or unintended dehumidification from too low a chilled water temperature (see Cooling Plant Optimization chapter). A dedicated outside air unit should be used to dehumidify or humidify outside air before it is brought into the building. Humidity produced in the space by operators can be easily controlled by supplying outside air at a slightly lower humidity than the setpoint.

RELATED CHAPTERS

- Air-side Economizer
- Centralized Air Handling
- Cooling Plant Optimization

RESOURCES

- Psychometrics, Chapter 6, ASHRAE HVAC Fundamentals Handbook, 2005.
- Thermal Guidelines for Data Processing Environments, TC9.9 Mission Critical Facilities, ASHRAE, 2004.
- Data Processing and Electronic Areas, Chapter 17, ASHRAE HVAC Applications, 2003.

8. POWER SUPPLIES

Most data center equipment uses internal or rack mounted AC-DC power supplies. Using higher efficiency power supplies will directly lower a data center's power bills, and indirectly reduce cooling system cost and rack overheating issues. Savings of $2,700 to $6,500 per year per rack are possible just from the use of more efficient power supplies. Efficient power supplies usually have a minimal incremental cost at the server level, however, management intervention may be required to encourage equipment purchasers to select efficient models. The purchasers need to be given a stake in reducing operating cost and the first cost of the electrical and conditioning infrastructure, or at least be made aware of these costs, in order to make a rational selection. Power supplies that meet the recommended efficiency guidelines of the Server System Infrastructure (SSI) Initiative[1] should be selected. The impact of real operating loads should also be considered to select power supplies that offer the best efficiency at the load level at which they are expected to most frequently operate.

PRINCIPLES

- Specify and utilize high efficiency power supplies in Information Technology (IT) computing equipment. High efficiency supplies are commercially available and will pay for themselves in very short timeframes when the total cost of ownership is evaluated.
- For a modern, heavily loaded installation with 100 racks, use of high efficiency power supplies alone could save $270,000-$570,000[2] per year and decrease the square-footage required for the IT equipment by allowing more servers to be packed into a single rack footprint before encountering heat dissipation limits.
- Cooling load and redundant power requirements related to IT equipment can be reduced by over 10 – 20%, allowing more computing equipment density without additional support equipment (UPSs, cooling, generators, etc.).
- In new construction, downsizing of the mechanical cooling equipment and/or electrical supply can significantly reduce first cost and lower the mechanical and electrical footprint.
- When ordering servers, power supplies that meet at least the minimum efficiency recommendations by the SSI Initiative (SSI members include Dell, Intel, and IBM).
- When appropriate, limit power supply oversizing to ensure higher – and more efficient – load factors.

APPROACH

The individuals specifying data center equipment should consider the efficiency of the power supply. Frequently, there is little connection between the group selecting data center equipment and the group that is aware of (and responsible for paying) the equipment's energy costs. To encourage the use of more efficient power supplies, an organization must illustrate the clear connection between equipment energy usage and operating cost to the people who make the

equipment purchasing decisions. With many types of equipment becoming commodity items and with small difference in price heavily impacting the selection, it is essential that the total cost of ownership of a low efficiency power supply be recognized in the selection process.

One approach to doing this might be to offer an internal budget incentive to be applied to the purchases of IT equipment that meet or exceed the minimum efficiency recommendations set forth by the Server System Infrastructure (SSI) initiative.

Equipment selection should include an evaluation of the power supply efficiency. For servers with integrated power supplies, a high efficiency option should be considered if available. The table below shows the impact of increasing the power supply efficiency from 72% (SSI Minimum) to 83% (SSI Recommended Minimum) on annual operating cost. This level of efficiency improvement is the difference between the SSI Required Minimum efficiency to allow proper cooling and the Recommended Minimum for reasonable energy performance. Only direct electrical savings are shown at an average (including peak charges) energy cost of $0.10 per kWh. Significant additional savings would accrue from the lower cooling requirement of a more efficient supply. When looking at the table, note that the typical server power supply costs about $20 to $40 to manufacture[3] – in most cases, a more efficient power supply would pay for itself in a single year even if its manufacturing costs were doubled.

FIGURE 1
Sample Per Rack Electrical Cost Savings from More Efficient Power Supplies

Power Supplied (Watts Delivered)	Annual Savings From Using a SSI Recommended Minimum Efficiency Supply[1]	Annual Savings Including Typical Cooling Energy[2]
200	$ 37	$ 64
300	$ 55	$ 96
400	$ 73	$ 128
470	$ 86	$ 151

1. Assuming $0.10/kWh, 8760 hr/yr, 83% efficient UPS supply, 72% efficiency baseline PS

2. Cooling electrical demand is estimated at 75% of rack demand, the average ratio of 12 benchmarked datacenter facilities

When a data center is housed and operated by the equipment owners, encouraging more efficient power supplies is usually just a matter of basic management to coordinate the actions (and budgets) of the data center equipment specifiers and those responsible for the infrastructure including those paying the electrical and cooling bills. Often, the department that selects and procures the data center equipment deals with hardware and software deployment issues and has little if any interaction with the department responsible for paying operating costs. With no feedback regarding operating cost, energy efficiency is naturally overlooked by the selectors – unless, perhaps, they have had negative, usually expensive, experience with overheating racks and the associated risk of equipment failure.

A simple management approach to address the disconnect between equipment purchases and the resulting energy bills and extra first-cost for facility infrastructure would be to offer an internal budget incentive equal to one to three years of energy savings for the purchase of equipment meeting the high efficiency performance level. For example, when purchasing servers that consume 200 watts, the purchaser's budget would be increased by $70-$150 per server if the server met the SSI Initiative's recommended efficiency levels. The 'incentive' money would be recovered from the reduced operating cost. Such a system provides much needed feedback between the equipment purchases (and by extension, manufacturers) and the total cost of ownership by allowing equipment purchases to be quickly and rationally evaluated on their true cost. The opportunity for energy savings from using high efficiency power supplies in a data center can be investigated using the simplified calculation tools developed as part of the LBNL project on data centers (see http://hightech.lbl.gov/server-ps-tool.html).

Recognizing the fiscal benefits of high-efficiency power supplies and considering them when selecting data center equipment is vital if high-efficiency power supplies are to make an impact in the market and on data center energy use. Higher efficiency power supplies are transparent to the user – they look the same and operate the same, the only noticeable difference is the lower heat and lower power bills. Without a deliberate effort to evaluate the efficiencies, the market has responded by providing minimum efficiency power supplies, sacrificing even very low cost, high return efficiency features to compete in a commodity market.

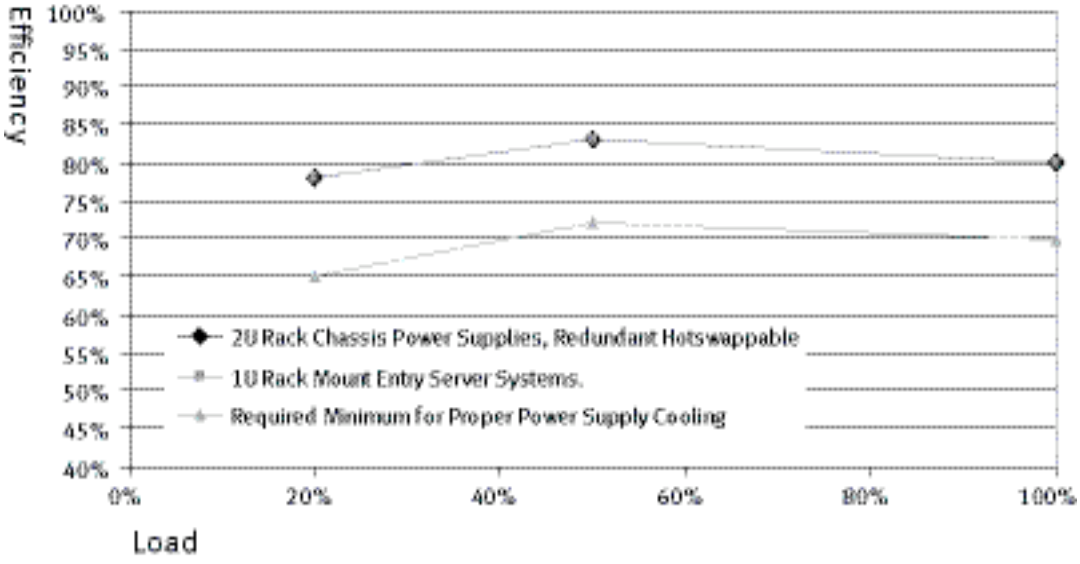

FIGURE 2
SAMPLE PER RACK ELECTRICAL COST SAVINGS FROM MORE EFFICIENT POWER SUPPLIES

Current recommended efficiency levels of server power supplies are outlined in the power supply design guidelines developed by the SSI Initiative. The SSI Initiative's goal is to deliver a set of specifications covering two primary server elements: power supplies and electronics bays. Members of SSI are Intel, Dell, HP, Silicon Graphics, and IBM. The current efficiency levels for power supplies used in servers based on the SSI design guidelines are shown in the following figure. Note that the lower line is the minimum performance required for cooling. The upper line is the efficiency recommended as the best balance between efficiency and first cost for

entry level servers. Considerable work is being done to encourage the production of high efficiency power supplies so the availability of equipment that meets the higher, recommended efficiency performance is expected to increase rapidly.

As much as possible, power supplies should be matched to meet the system load. In recent studies4, measured data of various manufacturers' power supplies showed that most were 70-75% efficient at full load. Typically, however, power supplies operate at between 15 and 40% of the full load; in this range, the efficiency drops off to roughly 50-65%.

It is technically and economically feasible today to design an 85% efficient power supply with an output voltage of 3.3V or less. Higher voltage, higher wattage designs are already achieving peak efficiencies of 90-92%. As seen in the consumer PC market, the technology to make more efficient power supplies exists in a form that could be economically brought to the commercial market. Blade server manufacturers have investigated offering more economical systems in the past, but a lack of market demand stalled the technology. Currently, the efforts of the SSI initiative and various governmental agencies ranging from the California Energy Commission to the EPA are actively encouraging the power supply market to focus design effort on improving efficiency. Taking advantage of the efficiency advances as they become available is both an efficiency and business best practice.

BENCHMARKING FINDINGS/CASE STUDIES

A good example of a server that illustrated many of the principals discussed here was the original RLX "ServerBlade 633" blade server that came on the market in 2000. It was a low-power server designed to meet specific requirements of the web-server market only. Based on a Transmeta chip, it was up to 10 times more power-efficient than servers based on processors from Intel. Because they generated less heat, RLX's servers could be stacked closer with less danger of overheating. Incredibly, RLX demonstrated that they could run a rack containing 336 of their servers safely. The rack consumed only about 3.3 kilowatts at average load, 2.4 kilowatts at idle, and 4.5 kilowatts at absolute peak, even though that was highly unlikely to happen. That's just over 13 watts per server at peak, fully loaded with everything spinning. By comparison, typical server racks today—containing far fewer servers—operate at 7.5 to 14 kilowatts, with talk of reaching 20 kilowatts by the end of 2004. RLX no longer offers the Transmeta based product, and current RLX blade servers consume 180 to 190 watts apiece.

RELATED CHAPTERS

- Uninterrupted Power Supply Systems

REFERENCES

1) Industry group, see http://www.ssiforum.org

2) Depending on cooling system efficiency and rack loading. Top end assumes a rack load of 16.8 kw of server power and a cooling power use equal to rack power use (lower quartile of cooling performance measured in recent benchmarking).

3) Energy Design Resources, RMI Design Brief, Data center Best Practices

4) Ecos Consulting, EPRI Solutions, LBNL, High Performance Buildings: Data Centers, Server Power Supplies, LBNL, 2005.

RESOURCES

- http://www.ssiforum.org
- http://www.efficientpowersupplies.org
- http://www.80plus.org
- http://www.hightech.lbl.gov
- Data Processing and Electronic Areas, Chapter 17, ASHRAE HVAC Applications, 2003.

9. SELF GENERATION

The combination of a nearly constant electrical load and the need for a high degree of reliability make large data centers well suited for self generation. To reduce first costs, self generation equipment should replace the backup generator system. It provides both an alternative to grid power and waste heat that can be used to meet nearby heating needs or harvested to cool the data center through absorption or adsorption chiller technologies. In some situations, the surplus and redundant capacity of the self generation plant can be operated to sell power back to the grid, offsetting the generation plant capital cost.

PRINCIPLES

- Self generation can improve efficiency by allowing the capture and use of waste heat.
- Waste heat can be used to supply cooling required by the data center through the use of absorption or adsorption chillers, reducing chilled water plant energy costs by well over 50%.
- High reliability generation systems can be sized and designed to be the primary power source while utilizing the grid for backup, thereby eliminating the need for emergency generators and, in some cases, even uninterruptible power supply (UPS) systems.

APPROACH

Data centers typically require sufficient emergency generation capacity on-site to support all the data center equipment and its infrastructure. Making this generator capacity the primary source of power for the facility—using efficient technologies—provides numerous benefits. The ideal primary power supply for a data center is an on-site generation system with short and simple distribution paths, and double redundancy at a minimum with waste heat recovered and used to power the cooling system.

Using waste heat for cooling can increase site efficiency and improve reliability for large data centers; in most situations, the use of waste heat is required to make site generation financially attractive. While large data centers have little need for space heating, waste heat from onsite co-generation can drive thermally based cooling systems. This strategy reduces the overall electrical energy requirements of the mechanical system by eliminating electricity use from the thermal component, leaving only the electricity requirements of the auxiliary pumps and cooling tower plant.

Absorbers use low-grade waste heat to thermally compress the chiller vapor in lieu of the mechanical compression used by conventional chillers. Rather than refrigerant and a compressor, a desiccant that absorbs and releases water, in the process absorbing and releasing heat, is used to remove heat from the chilled water loop and reject it to the condenser loop. The electrically driven compressor is replaced by a heat driven dessicant cycle. Single stage, lithium bromide dessicant based chillers are capable of using the low grade waste heat that can be recovered from common onsite power generation options including microturbines, fuel

cells, and natural gas reciprocating engines. Although absorption chillers have low Coefficient Of Performance (COP) ratings compared to mechanical chillers, utilizing 'free' waste heat from a generating plant to drive them increases the overall system efficiency. Absorbers are a very mature technology available from several major manufacturers, although they are less efficient than absorbers in converting low grade (low temperature) waste heat to cooling.

A potentially more efficient thermally driven technology that has begun making inroads in the domestic market is the absorber chiller. An absorber is a desiccant-based cooling system that uses waste heat to regenerate the desiccant and cooling towers to dissipate the removed heat. An absorption chiller minimizes its auxiliary loads by eliminating the absorbent pump and decreasing the run times of the vacuum and refrigerant pumps, thus further limiting the electricity requirements while maintaining a similar thermal COP. The silica gel based system uses water as the refrigerant and is able to use lower temperature waste heat than a lithium bromide based absorption chiller. While absorption chillers have been in production for about 20 years, they have only recently been introduced on the American market.

An appropriate purpose-designed self generation system could eliminate the need for a UPS system, with the attendant first cost and efficiency implications. While some companies have offered such high reliability systems, they have not been widely implemented as of this writing. However, with the proper redundancy and design, data center facilities can eliminate UPS systems and achieve significant efficiency benefits at reasonable cost. The current market offerings for high reliability power should be evaluated for new data centers.

The recommended system would typically be sized to cover the full site load (as allowed by local utility and codes) and connected to the grid to ensure reliability and to improve payback. The grid would serve as the backup to the self-generation plant. The key to successful connection with the utility is two very fast circuit breakers or static switches on each generating bus to quickly disconnect the on-site generator and prevent any possible damage associated with reverse fault flows during times of grid failure, when the on-site generator must operate in an "island" mode.

Any self-generation system would need to be designed carefully to meet all local codes and requirements, including air emission limits. Storage of backup fuel for natural gas systems can also be a code and technical challenge, with propane and dual-fuel capable generators often used to create an onsite emergency fuel storage solution.

Frequently, self-generation systems are sized to only supply a baseline quantity of power, offering the benefits of waste heat reclamation at a reduced first cost. However, depending on the specific generation equipment used, there can be a significant delay between operation in a baseline mode with the utility grid always used to 'top up' to the actual required load, and operation in an island mode, with standard backup generators or load shedding used in place of the grid. It is the transfer time from operating in a grid connected mode to operating on the generator plant alone that often necessitates the continued installation of UPS systems.

Ideally, an important benefit of interconnection with the utility would be that unused generating capacity could be sold back onto the grid. This would offer the flexibility to keep generators running at full load, thus making them optimally efficient (globally more fossil fuel efficient). The export of spare power could be an additional revenue source, shortening the payback period of the total investment. While the economic benefits of selling power back are considerable, the combination of power export and high-reliability operation is challenging. A cooperative distribution utility and careful design is required to achieve the level of reliability required by data center facilities.

RESOURCES

- Data Processing and Electronic Areas, Chapter 17, ASHRAE HVAC Applications, 2003.
- Casestudy: http://www.energyvortex.com/files/MissionPlasticsCaseStudy.pdf

REFERENCES

1) Assuming 40% electrical efficiency natural gas generator with 70% of waste heat recovered and used to drive a chiller with a 0.7 CUP (at $.014/kwh and $1.00/therm).

10. UNINTERRUPTIBLE POWER SUPPLY SYSTEMS

FIGURE 1
TYPICAL UPS UNITS

Uninterruptible Power Supply (UPS) systems provide backup power to data centers, and can be based on battery banks, rotary machines, fuel cells, or other technologies. A portion of all the power supplied by the UPS to operate the data center equipment is lost to inefficiencies in the system. These inefficiencies can total hundreds of thousands of wasted kilowatt hours per year. UPSs differ in their efficiency levels; this should be taken into consideration when selecting a UPS system. Beyond the UPS system efficiency, the electrical design can also affect system efficiency by impacting the typical load factor of UPS operation. Battery-based UPS systems are more efficient at a high load factor – at least 40% of their rated capacity or higher. The UPS system configuration (line reactive versus double conversion) also impacts efficiency. Greater power conditioning capability often entails greater amounts of electrical waste and adds additional heat loads that the mechanical cooling system must remove.

PRINCIPLES

- Select the most efficient UPS system that meets the data center's needs. Among double conversion systems (the most commonly used data center system), UPS efficiency ranges from 86% to 95%. Simply selecting a 5% higher efficiency model of UPS can save over $38,000 per year in a 15,000 square foot data center, with no discernable impact on the data center's operation beyond the energy savings. In addition, mechanical cooling energy use and equipment cost can be reduced.
- For battery-based UPS systems, use a design approach that keeps the UPS load factor as high as possible. This usually requires using multiple smaller units. Redundancy in particular requires design attention; operating a single large UPS in parallel with a 100% capacity identical redundant UPS unit (n+1 design redundancy) results in very low load factor operation, at best no more than 50% at full design buildout.
- Evaluate the need for power conditioning. Line reactive systems often provide enough power conditioning for servers, and some traditional double conversion UPS systems (which offer the highest degree of power conditioning) have the ability to operate in the more efficient line conditioning mode, usually advertised as 'economy' or 'eco' mode.

- Consider new technologies currently being proven on the market, such as flywheel systems. Such systems eliminate replacement and disposal concerns associated with conventional lead-acid based UPS systems and the added costs of special ventilation systems and often conditioning systems required to maintain temperature requirements to assure battery life.

APPROACH

Increasing UPS system efficiency offers direct, 24-hour-a-day, energy savings, both within the UPS itself and indirectly through lower heat loads and even reduced building transformer losses. When a full data center equipment load is served through a UPS system, even a small improvement in the efficiency of the system can yield a large annual cost savings. For example, a 15,000 square foot data center with IT equipment operating at 50 w/sf requires 6,900,000 kWh of energy annually for the IT equipment. If the UPS system supplying that power has its efficiency improved by just 5 percentage points, the annual energy bill will be reduced by 384,000 kWh, or about $38,000 at $0.10 / kWh, plus significant additional savings from the reduced cooling load.

FIGURE 2
Factory Measurements of UPS Efficiency (Tested Using Linear Loads). Note the Wide Range In efficiencies Among the Common Double-Conversion UPS Topography[1]

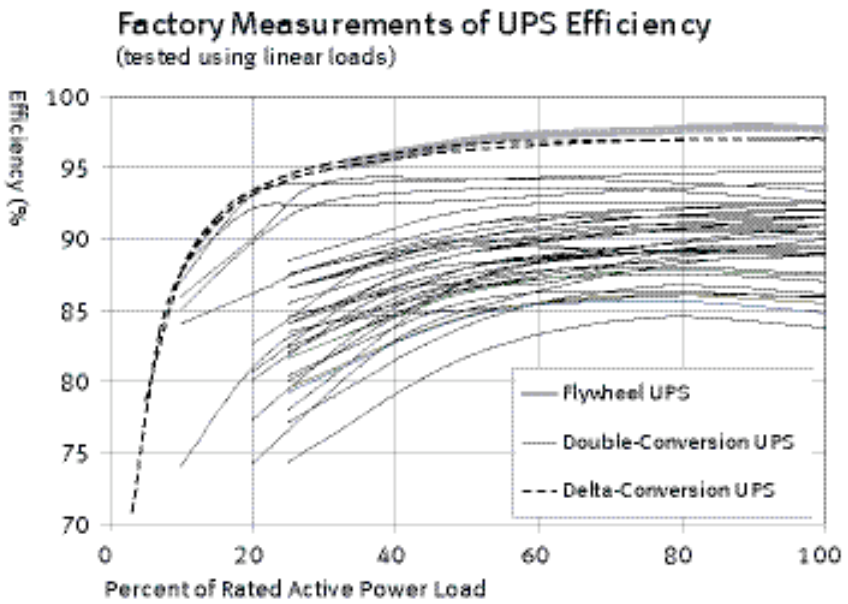

Figure 2 shows measured efficiencies of a large number of different UPS offerings as reported by EPRI Solutions as part of a larger data center energy efficiency project. Notice the large efficiency range, from 86% to almost 95% at full load, and even greater at the partloads where most systems are operated, among the double conversion UPS systems tested. Double conversion UPS systems are the most commonly used UPS in data centers.

UPS system efficiency is improved by operating the system at a load factor of at least 40%. System efficiency drops rapidly when the UPS is lightly loaded, as seen in Figure 4 that shows measured data from a number of UPS systems operating in critical facilities. The load factor can be influenced by both the design approach in sizing the systems and the configuration of redundant units.

For a data center, the UPS system should be configured to efficiently accommodate operation at partial loads. For example, it is common for data centers to operate at less than half of their full design load as the data center is populated. If the UPS configuration shown on the left of the figure below is operating at 50% of its design load (300 kW), each UPS is operating at a load factor of only 25% (150 kW). However, for the same total equipment load of 50% in the configuration on the right, each UPS would need to operate at 33% (100 kW). An efficiency gain of approximately 5% would be realized just from operating a UPS at 33% versus 25% of full load. Both configurations maintain the same level of redundancy.

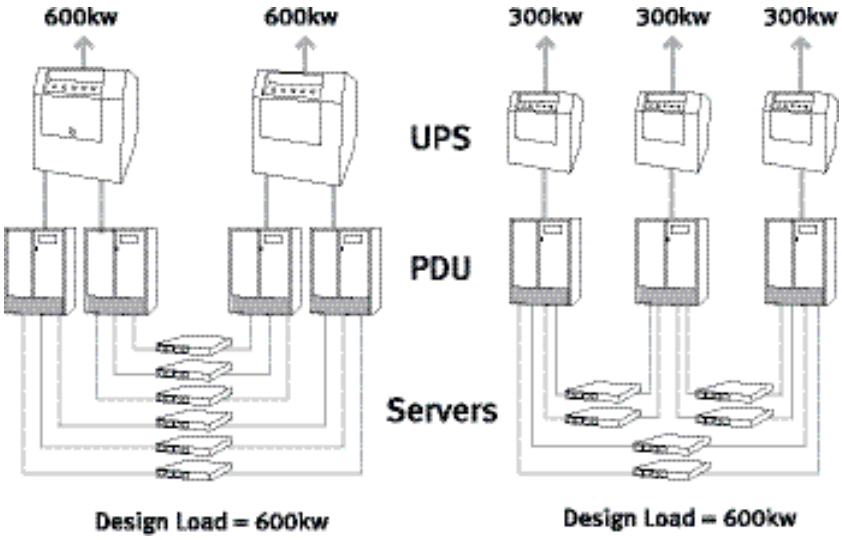

FIGURE 3
An Efficiency Gain is Realized when Operating Smaller UPS Systems at Partial Loads as Opposed to Larger Systems at Partial Loads[2]

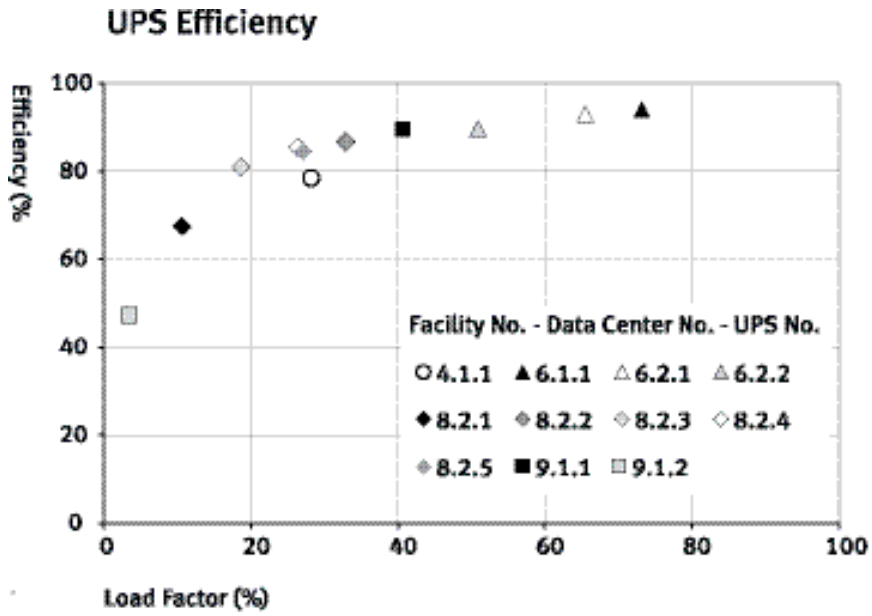

FIGURE 4
Measured Efficiencies of UPS Systems[3]

Accurate IT equipment load estimates can prevent gross oversizing and the resultant underloading of UPS systems. Data centers are often designed for two to four times their actual measured operating equipment load. When designing a data center, using actual measurements of the load of an equivalent operating data center can help develop a rational

design load, and a modular UPS system design approach can easily allow for future growth and delay the first cost of additional capacity until it is actually required. A modular UPS system can allow the total operating capacity to be optimized to match the actual connected load. After any significant equipment installation, the equipment load should be evaluated and all unnecessary UPS units shut off.

Since the UPS system is operating 24 hours a day, even small improvements in unit efficiency can have a significant impact on annual energy use. System efficiency should be evaluated for a number of part load points and for the type of electrical load expected. Payback or Return On Investment (ROI) calculations should be used to estimate the comparative economic impacts of competing equivalent UPS systems. Specifications must include the required efficiencies at set load points, and in many cases specifying witnessed factory testing to confirm the efficiency is justified.

FIGURE 5
Battery UPS Configuration Power Conditioning and Efficiency

UPS Configuration	Power Conditioning	Efficiency
Passive Standby	None	High
Line Reactive	Moderate	Moderate
Double Conversion	High	Low

The internal UPS system design, or topology, and how it treats incoming utility power has a large impact on the efficiency of the system. The most efficient topology is a passive standby configuration, where the equipment loads are connected directly to the utility power and only switched to the UPS system's battery supplied power during a power failure. A topology that provides a degree of continuous power conditioning between the utility and data center loads is a line reactive system. Finally, a double conversion configuration continuously converts all incoming utility power to DC and then uses the inverter section of the UPS system to convert it back to AC and to supply the data center. The table below summarizes how these technologies generally compare in regards to the relative degree of power conditioning they provide and their efficiency potential.

As a general rule, the greater the power conditioning offered, the larger the number of components involved, and the lower the efficiency. Data centers typically require some degree of power conditioning – this precludes the use of passive standby UPS configurations for any critical data center equipment. However, most mass-produced data center servers and equipment can operate through significant power disruptions; most server power supplies have a specified input voltage range of more than +/- 10% relative to the nominal voltage. Computing equipment power supplies are often designed to operate in non-data center

environments, and therefore utilize internal AC to DC power supply converter designs that can accept typical utility power disturbances. Data centers often specify double conversion systems, the gold standard of power conditioning performance (and cost) by default when line reactive systems would be appropriate, more efficient, and reduce both first costs and operating costs.

As a middle point between line reactive and double conversion systems, some UPS manufacturers offer double conversion UPS systems with an 'economy' or 'eco' high efficiency mode. When in 'eco' high efficiency mode, the system typically operates in a line interactive configuration, filtering the incoming utility and supplying it directly to the equipment in normal operation. Bypassing the double conversion circuitry can increase efficiency by about 5%, decreasing the wasted energy by 50 watts for every kilowatt of computer power supplied. The system can be switched to operate in double conversion mode when extra poor conditioning may be beneficial, for example an extended period operation on generator or during periods of known poor power quality, such as announced summer brownouts, rolling blackouts, or severe seasonal storms. Additional conditioning provided by some Power Distribution Units (PDUs) should also be considered when evaluating UPS requirements.

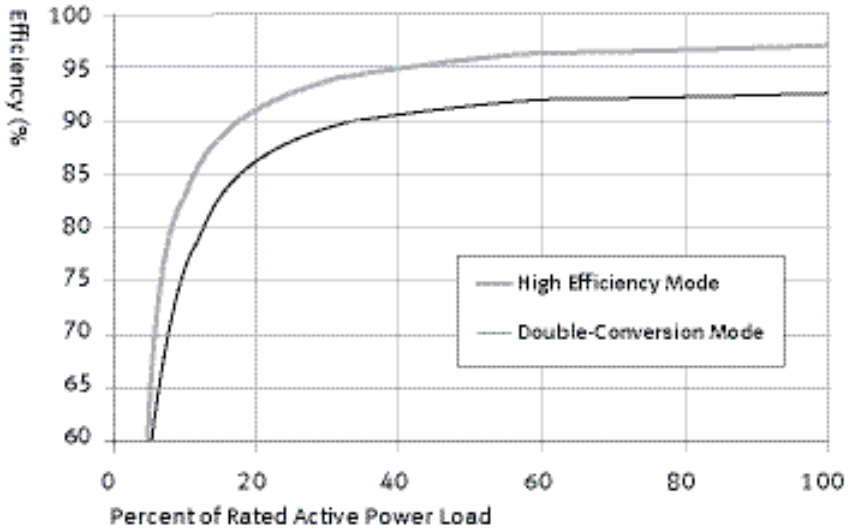

FIGURE 6
EFFECT OF HIGH EFFICIENCY MODE[4]

Beyond the traditional battery powered UPS, there are some new technologies currently being commercialized that eliminate the battery/inverter design approach entirely, offering the opportunity for greater efficiency. One is the rotary UPS, which utilizes a high speed, very low friction rotating flywheel usually coupled with a backup combustion generator that can start up instantaneously to provide emergency power. When power fails, the stored inertial energy of the flywheel is used to drive a generator until the fast start generator can take over the load. Typically, flywheel systems offer a shorter ride-through time than battery based systems, potentially impacting the selection and redundancy of backup generators. Flywheel systems offer the very high efficiency of line-interactive devices, in excess of 95%. The reliability of these systems compared to inverter systems has not yet been fully proven in the market, however the system is commercially available and rapidly gaining operating hours in a wide variety of critical facility applications. This rapidly maturing technology should be considered when selecting an UPS system.

Fuel cell systems that have the potential to blur the line between a UPS system and an onsite power generation plant cells are continuing to mature, and a number of technologies are now being offered on the market. Availability of fuel cell products is limited, but they hold promise as they continue to develop. Another developing technology is the use of very high availability onsite power generation plants to both reap savings from cogeneration and, combined with a utility power interconnect, provide a UPS capability (see Self Generation chapter).

BENCHMARKING FINDINGS/CASE STUDIES

Recent testing conducted by EPRI PEAC for a LBNL data center project confirmed an increase in power efficiency of about 5% when a double conversion UPS system was run in its optional 'high efficiency' line interactive mode and the standard double conversion operation.

In one test, transferring the power supply from the utility line to the batteries was seen to cause only a very brief voltage sag (1/2-1 cycle). A UPS with a line reactive mode (often referred to as a high efficiency or 'eco' mode) is a good energy efficiency option if the data center equipment can tolerate the switchover time of the high efficiency mode, as most standard computer power supplies can. The system should also be checked to verify its compatibility with the installed backup generator system. Specifying a dual mode capable UPS system offers considerable flexibility in accommodating future equipment that may have unusually sensitive power supplies requiring the highest degree of power conditioning provided by a double conversion unit. These systems can also be switched to double conversion mode in known extreme conditions, such as during extended brownouts or periods of generator operation.

FIGURE 7
Interruption of Power when Switching from Passive Standby Mode to High Efficiency Mode

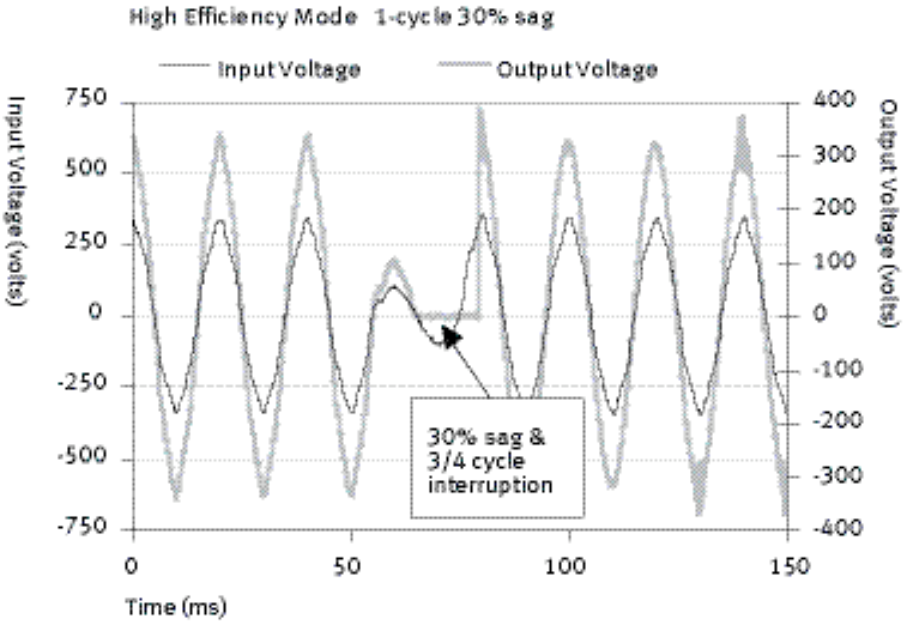

RELATED CHAPTERS

- Power Supplies
- Self Generation

REFERENCES

1) Ecos Consulting, EPRI Solutions, LBNL, High Performance Buildings: Data Centers, Server Power Supplies, LBNL, 2005.

2) ibid

3) LBNL/Rumsey Engineers Cleanroom Benchmarking Study

4) Ecos Consulting, EPRI Solutions, LBNL, High Performance Buildings: Data Centers, Server Power Supplies, LBNL, 2005.

5) ibid

RESOURCES

- Data Processing and Electronic Areas, Chapter 17, ASHRAE HVAC Applications, 2003.
- http://data centers.lbl.gov

This education product was developed utilizing Energy Design Resources funding. Energy Design Resources provides information and design tools to architects, engineers, lighting designers, and building owners and developers. Energy Design Resources is funded by California utility customers and administered by Pacific Gas and Electric Company, San Diego Gas and Electric, Southern California Edison, and Southern California Gas Company, under the auspices of the California Public Utilities Commission. To learn more about Energy Design Resources, visit www.energydesignresources.com.

This Design Guidelines Sourcebook was developed by Rumsey Engineers, in part, based upon work by Lawrence Berkeley National Laboratory supported by the California Energy Commission's Public Interest Energy Research program.

Pacific Gas and Electric Company®

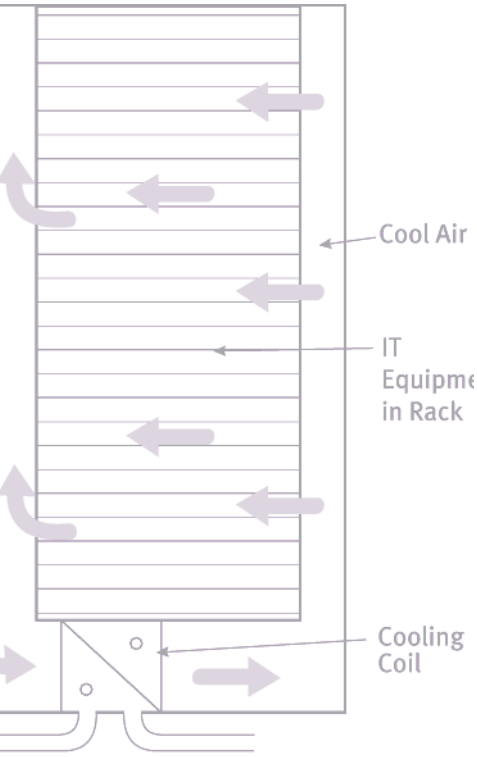

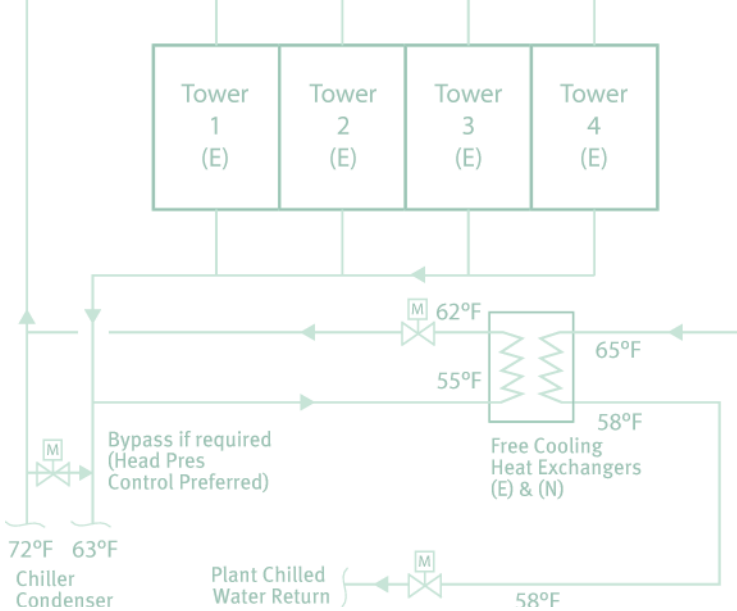

Self Benchmarking Guide for Data Center Energy Performance

Version 1.0

Ernest Orlando Lawrence Berkeley National Laboratory

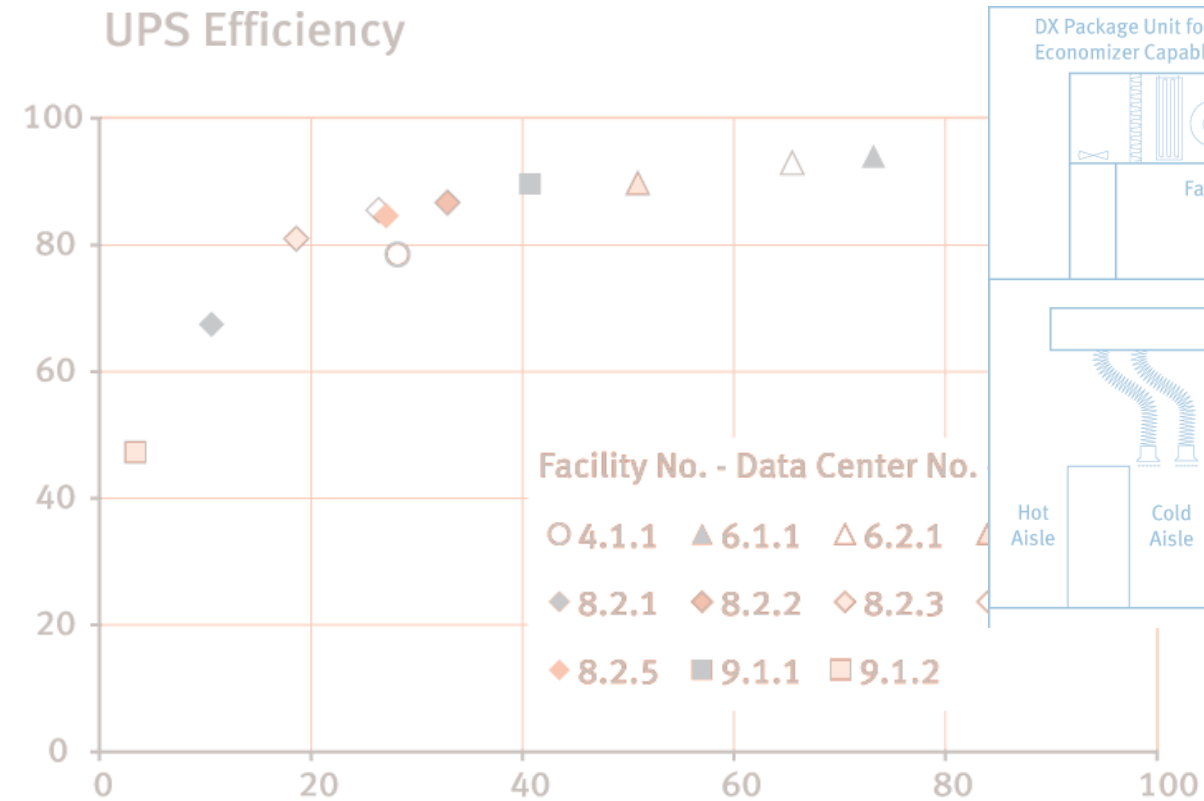

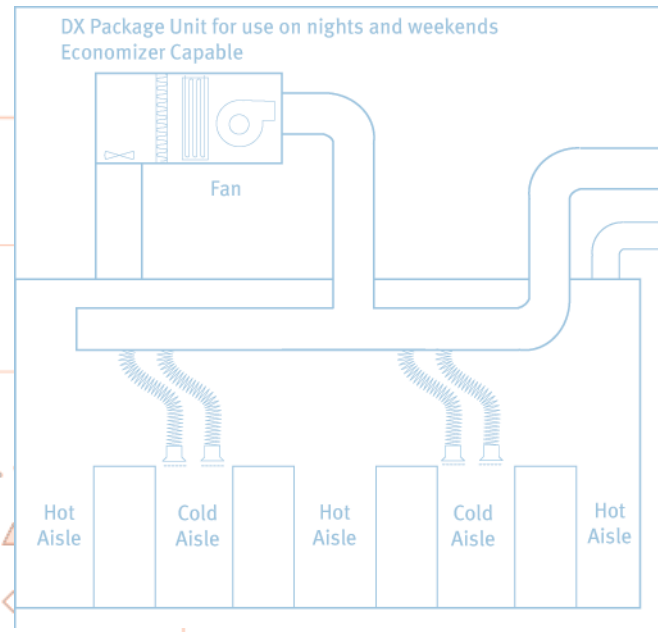

Self Benchmarking Guide for Data Center Energy Performance Version 1.0

Ernest Orlando Lawrence Berkeley National Laboratory

May, 2006

Authors:

Steve Greenberg, William Tschudi,
LBNL

John Weale,
Rumsey Engineers

Disclaimer: This document was prepared as an account of work sponsored by the California Institute for Energy and the Environment through a grant from the California Energy Commission Public Interest Energy Research Industrial Program. While this document is believed to contain correct information, neither the sponsoring organizations nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the Sponsoring Organizations or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the Sponsoring Organizations or any agency thereof or The Regents of the University of California.

Contents

Introduction....1

1. Getting Started- Identify Benchmarking Goals....3

2. Orient IT and Facilities Staff....4

3. Obtain Documentation....4

4. Create Site-Specific Monitoring Plan....6

5. Obtain Monitoring Equipment....7

6. Install Monitoring Equipment....7

7. Collect Data and Remove Monitoring Equipment....12

8. Analyze and Reduce Data....14

9. Compare Findings to Existing Benchmarks....15

10. Compare Data Center to Best Practices....16

11. Share Results....16

Conclusion....17

Appendix A: Data Center Benchmarking Data Tables....19

Appendix B: Checklists and Data Collections Forms....25

Introduction

In order to optimize any system, or combination of systems working together such as those found in a typical data center, it is crucial to accurately assess its current performance. This guide is intended to show data center operators, or their contractors, how to perform a comprehensive measurement, or benchmarking, of their own facilities' energy use.

Energy benchmarking offers a very effective way to evaluate the performance of a data center facility and compare it to similar facilities. Comparisons are a quick and easy way to identify poorly operating areas, which typically have the highest potential for economical modifications that reduce operating cost and/or increase the load capacity of a data center. Benchmarking is particularly valuable as a precursor to an expansion of a facility, because increases in system capacity and identification of best practice design approaches can be used to reduce the cost and increase the energy efficiency of new space.

The following steps guide the operator through the benchmarking process, from the initial definition of the scope of the benchmarking through the initial evaluation of the results. Widely used by most manufacturing industries to improve product quality, benchmarking is also a powerful tool for improving the quality, reliability, and performance of data center facilities. Once the magnitude of power use is identified on a system basis, resources can be prioritized for the areas where the greatest savings can be achieved.

The charts below are part of the results of benchmarking performed on typical datacenter facilities. They show the large differences that can occur between systems supporting the same space requirements. Since the conditioning requirements of data centers are usually similar, with roughly equivalent loads and setpoints, large discrepancies between auxiliary equipment energy costs can clearly identify system designs and/or operations that are inefficient and good targets for improvements in cost and energy use.

Figure 1: Comparison of actual benchmarking data from two data centers showing breakdown of total electricity use per watt of server. Facility 2 uses much less power on supporting infrastructure.

Facility 1

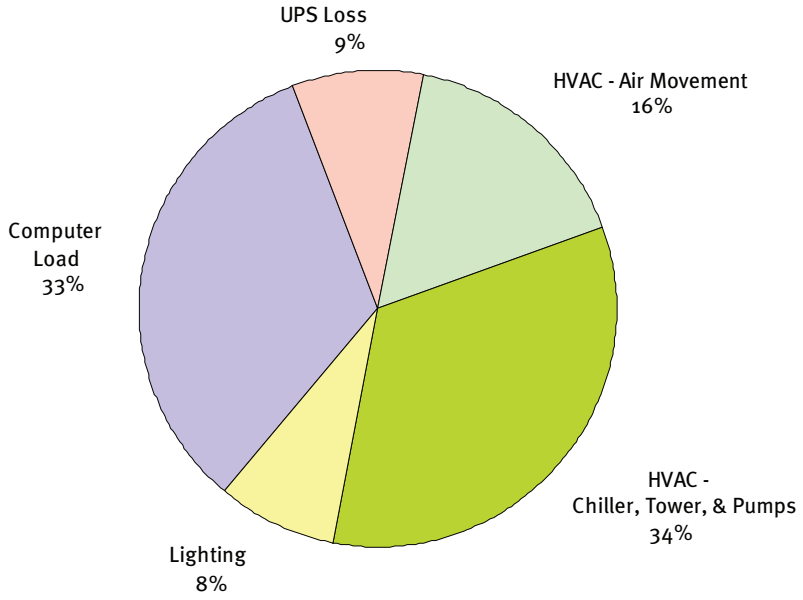

Facility 2

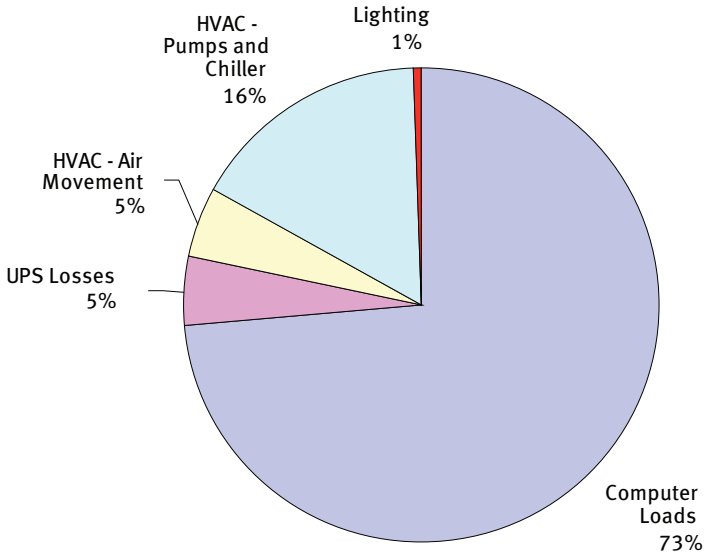

1. Getting started- identify benchmarking goals

The first step in benchmarking your facility is to determine the scope and overall business objectives of the benchmarking project. There are many reasons to benchmark a data center. A few common ones are:

• Establishing a base line performance as a diagnostic tool
Comparing trend data over time to baseline performance can help predict and avoid equipment failure, improving long-term reliability. Efficiency can also benefit by identifying and therefore allowing the correction of typical performance decay that occurs as systems age and calibrations are lost.

• Identifying operational or maintenance issues
In particular, to assist in diagnosing the root cause of hot spots, heat related equipment failure, lack of overall capacity, and other common operational problems. Due to the critical nature of data center environments, such problems are often addressed in a very non-optimal, "Band-Aid" manner due to the need for an immediate fix. Benchmarking can identify those "quick fixes" that should be revisited in the interests of lower operating cost or long-term reliability.

• Helping to plan future improvements
The areas that show the poorest performance relative to other data center facilities usually offer the greatest, most economical opportunity for energy cost savings. Improvements can range from simply changing setpoints in order to realize an immediate payback, to replacing full systems in order to realize energy savings that will show payback over the course of a few years.

• Developing design standards for future facilities
Benchmarking performed over many facilities in recent years has suggested there are some "best practice" design approaches that offer fundamentally lower cost, more efficient facilities. Benchmarking can help identify: 1.) Company-specific best practices that should be duplicated to reduce the cost of future facilities; and 2.) Less efficient design approaches that should be avoided.

• Improving management
The cliché, "You can't manage what you don't measure" applies to data centers. For example, knowing the relationship between server wattage requirements and mechanical cooling costs can help in determining the value of purchasing slightly more expensive but much more efficient server power supplies and IT equipment. Also, operating costs of double-conversion UPSs can be compared to those of line-reactive units to determine if the (possibly unnecessary) additional conditioning is financially justifiable.

An attainable scope for the benchmarking should be defined. For example, if the goal is to establish a benchmarking team within a company with several data centers, it would be useful to start with a relatively smaller and simpler data center known to have good documentation. Centers with high absolute or system energy costs, known performance concerns, or "Band-Aid" fixes might be good candidates as well since they are most likely to produce quick payback (or "low hanging fruit") opportunities. At this point in the process the goals of the benchmarking should be established to help guide the study. A review of the metrics in the appendix of this document can help in selecting the areas that are typically addressed in data center benchmarking (and where significant comparison data are available).

2. Orient the IT and facilities staff

The next step is to engage Information Technology and Facilities personnel at the data center. Benchmarking is often motivated by facilities staff, who must monitor operations, identify and resolve trouble spots and ensure smooth operation; however, it's also important to involve IT staff so that the data center's mechanical infrastructure and IT equipment can be considered as a whole system. This is best achieved through meetings involving IT and Facilities staff, where the benchmarking process and its benefits can be explained; each group can be informed of the documentation and access they need to provide during benchmarking; and the expected results can be identified and communicated to the group.

For example, purchasing more energy efficient IT equipment, such as servers that use "80 plus" certified higher performing power supplies, would at least reduce long-term power costs. Quantifying the savings to the operational budget can justify a minor increase in the equipment infrastructure budget that will result in valuable efficiency gains.

Another issue is airflow management, which requires close control of the layout of equipment in the racks and the exhaust flow configuration of the equipment. While air management has a very large impact on the efficiency of mechanical systems, it can also directly impact IT personnel through reliability improvement, or its opposite, increased heat-related failure rates. The most economical and long-term way to achieve optimal performance is to view and treat the data center infrastructure and the IT equipment in the way it actually operates in real-time: as an integrated, whole system.

There are also significant benefits to the operational facilities staff. Benchmarking can also identify easily corrected problems that can increase the heat density capacity of a space without requiring additional cooling equipment. While useful at nearly all points in a data center's life cycle, benchmarking can be most useful and cost-effective before expansion to identify the most efficient systems and facilitate optimal Return on Investment decisions.

3. Obtain documentation

The process starts with gathering information on utility bills. This could simply be a quick snapshot- looking at the latest bills - but a longer load (and cost) history is useful in evaluating growth trends. Ideally, a year or more of utility bills should be obtained and displayed graphically. The Energy Monitoring and Control System (EMCS) and Supervisory Control and Data Acquisition (SCADA) systems can provide much of the benchmark data for both snapshots of current operation and historical performance, if trend data is available. Where the data trends are available, it can be used to minimize the cost of additional monitoring equipment.

Other resources include drawings of the facility, Test, Adjust, and Balance (TAB) reports, and commissioning or retro-commissioning reports. These documents will be essential to understanding system performance and efficiency opportunities and will provide much of the benchmark design or previously measured performance information. This will enable comparisons of current measured performance (determined by direct measurement or through building management systems) to the original design, or to latest measured performance.

Facilities and IT personnel should review the current data center configuration and operation to better understand its current performance and opportunities for improvement. Documentation

of the electrical and mechanical systems is useful in understanding the data center and creating a monitoring plan. Available documentation varies greatly, from the ideal case of up-to-date as-built drawings, reports and lists, to the worst case of no documentation. Typically, there will be substantial documentation available, but it will often be incomplete and imperfect and require some physical verification. The following documents tend to be available and useful to the process.

1) Floor plans
Floor plans of the center will show the location of the IT equipment racks, computer room air conditioners (CRACs) or computer room air handlers (CRAHs), power distribution units (PDUs), uninterruptible power supplies (UPSs), etc.

2) One-line electrical drawings
These will show the main power distribution configuration, including utility power feed, transformers, switchgear, UPS systems, and distribution to major loads and panels.

3) Mechanical piping and duct work drawings
These show configurations, locations, sizes, and design flows for air and water systems.

4) Mechanical equipment schedules
These are comprised of one or more drawings with tables, typically showing the manufacturer and model numbers of the mechanical equipment, as well as design loads, temperatures, flows, pressure drops, etc.

5) Controls drawings
In particular, a "P&ID" drawing shows a schematic of the mechanical systems being controlled as well as the control points for those systems. The sequence of operations (how the control system is programmed) may be found on the control drawings, in the Specifications, or in a Maintenance Manual submittal.

6) Utilities bills
Ideally, at least one year of monthly electricity, gas, and water bills should be obtained. If the site does not have ready access to them, the Utility can usually provide them upon request by the Utility customer. It is also useful to get several years' data if available in order to trend the center's overall usage over time (as the IT equipment load changes).

7) EMCS and SCADA system points lists
The Energy Monitoring and Control System and Supervisory Control and Data Acquisition system often monitor equipment and systems that are relevant to the benchmarking process. Get a list of the available monitoring points from the system operator(s). Also verify that the points can be logged (or "trended") and that it is reasonably easy to get the data from the system as a spreadsheet or equivalent (comma separated value or csv, e.g.) file.

8) Air and water balance and commissioning reports
These are done by balancing and commissioning contractors and show (for balance reports) air and water flows and temperatures for mechanical equipment and systems. Commissioning reports tend to also show system performance data at real-time operating conditions rather than design conditions. Where systems cannot be easily measured in the operating center, such reports provide the best data; otherwise design data must be used.

4. Create the site-specific monitoring plan

Using all of the information collected in step 3, a center specific monitoring plan can be developed. Once the plan is established and agreed to by the IT and Facilities staff, data collection can be done by taking measurements, recording information from EMCS systems and using monitoring equipment temporarily installed in the building.

1) Select Metrics

Use the LBNL template or create a site-specific list of metrics (several lists of monitoring points and metrics used for prior benchmarking studies are provided in Appendix A and Appendix B). Areas of interest can be selected from the overall list, if needed, to assign priority in the common situation where access and resources are limited. Three key parameters are recommended as fundamental for defining the space: 1.) IT equipment total power and power density, 2.) The ratio of IT energy to HVAC energy (since that is key to determining how efficiently the IT equipment is cooled-see Figure 2), and 3.) The ratio of IT energy to total energy (since it is key to the overall efficiency of the center). Additional metrics should be selected to target specific systems, known problem areas, or to prepare for specific future work. It is common to select at least a few high-level metrics, such as kW/ton for chilled water systems, to serve as a check for gross efficiency or operational problems. Additional measurements are often included during the benchmarking process to identify the cause of unusual measurements as they are taken.

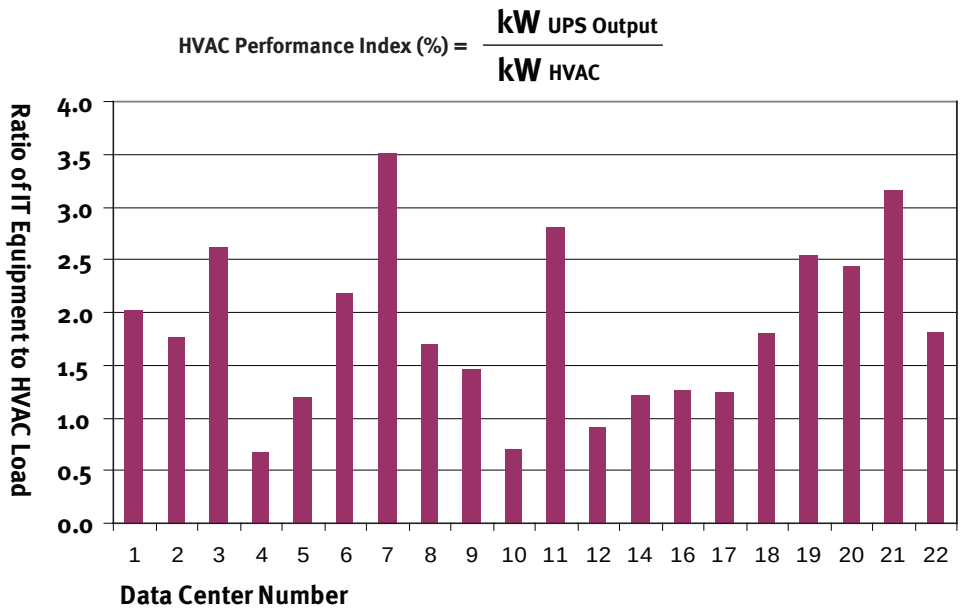

Figure 2: HVAC performance index. By comparing the energy used for cooling the data center (HVAC power in kW) to the UPS output, (which should closely resemble the computer loads) an indicator of HVAC system performance is obtained. The graph shows a comparison of data for 14 system configurations, indicating wide variation in system design and energy efficiency.

2) List data required for metrics

List specific points to spot check for spot measurement trending with EMCS or SCADA system, or use measurements with data loggers. Based on past experience, some of the measurements vary little over time, so they only need to be checked once (see Appendix). Spot checks can be done with either portable instrumentation or by observing instruments that are part of the existing system (thermometers, gauges, control panels, or EMCS/SCADA screens). For more critical spot checks, installed gauges and sensors should have their calibration checked or be compared against a portable calibrated instrument once in order to verify their accuracy. Other items need to be trended over time in order to properly identify operational performance: this can be done with either the EMCS or SCADA, if they have the capability, or with portable data logging equipment.

3) List monitoring equipment needed to collect data

Depending on the scope of the study, the number of systems and pieces of equipment, and how much instrumentation is built into the center, a wide range of monitoring equipment may be needed to complete the benchmarking. For each point, determine how the measurements will be taken and whether the equipment is available in-house. One important point to consider for the power measurements is the size of the current sensors (often referred to as "CTs" even when they are not technically current transformers) needed, both in current rating and physical size (will they be going around multiple large conductors, bus bars, etc.).

5. Obtain monitoring equipment

From the list generated in the previous step, determine how the monitoring equipment will be obtained. Ideally the benchmarking process will become an ongoing part of the management of the data center. In this event it may be advantageous to purchase the monitoring equipment needed for the study. Equipment can also be rented. Other possible sources include consultants or public utilities. For instance, California public utilities provide free monitoring equipment loans through tool lending libraries such as Pacific Gas and Electric Company's at http://www.pge.com/003_save_energy/003c_edu_train/pec/toolbox/tll/tll_home.shtml.

The accuracy of the monitoring equipment should be assessed, including calibration status. Calibration accuracy should support the desired accuracy of results expected from the benchmarking. The measurement range should be carefully considered when determining minimum sensor accuracy. For example, a pair of +/-1.5 F degree temperature sensors are almost useless for determining the waterside ΔT across a typical chiller, which under normal operation may be as low as 5 F degrees.

6. Install monitoring equipment

Access to the data center and coordination must be arranged. Many centers have extensive change control and access procedures: it is critical that these be followed and accounted for in the benchmarking plan. Most power measurements are made on live panels due to the impracticality of any power interruption to the data center, so an electrician needs to be available to place monitoring equipment in or around live electrical panels. Observe all safety procedures, particularly "hot work" requirements, in order to maintain the safety of personnel, protect equipment, and ensure continuous operation of the data center. If electrical panels cannot be fully closed due to the placement of temporary logging equipment, the area must be blocked off with appropriate signage and caution tape, or as otherwise required by the site, for the duration of the required trending period.

1) Set up data logging equipment

Ideally, this step is as simple as configuring trends on the control system. In cases where temporary loggers are used in place of permanent control system sensors, configure the loggers as needed to take the planned measurements. Data loggers vary in their configuration procedures, but typically there is software that runs on a Windows-based PC that allows the recording interval, type of circuit, and other parameters to be set up in the logger.

The trending interval is of particular importance. It should be set as short as is practical, with one minute being the ideal. Longer intervals are often required due to limited logger memory, and will still provide a great deal of valuable information; however, shorter intervals tend to be required in order to identify control problems, such as a control valve cycling from open to closed once every 2 minutes (which may be caused by a PID tuning or sizing issue), or a system switching from humidification to dehumidification every minute (which may be caused by a

bad sensor, inappropriate humidifier type, or unrealistically tight control band setting). Setting all loggers that measure related parameters to the same interval can simplify data analysis and facilitate combining collected data into a single spreadsheet.

Logging equipment setup can be done in the field; however, prior setup can help staff to familiarize themselves with the equipment and afford the opportunity to make test runs to verify that the equipment gathers data and records it back to the computer. Ensure that self-powered loggers have fresh batteries, that sufficient memory is available for the desired durations, and that appropriate size connectors are available. While familiarizing the team with the equipment, create a simple checklist for the logger connections that will be performed at each point (see Appendix B). The checklist should include common quality checks, such as downloading a five minute trend to verify proper operation. Trended data is most valuable when taken simultaneously; the loss of a single point can necessitate repeating many others in order to achieve a complete picture of the data center operation.

For points that will be monitored from the EMCS or SCADA system, verify that the logging and data download capabilities are well understood.

See the Sample Monitoring Setup Checklists in Appendix B for a useful checklist for setup.

2) Logging at CRACs/CRAHs/AHUs
Where consistent with the plan, start logging the computer room conditions of the air conditioners, computer room air handlers, or other air-handling units serving the data center. This may be either through use of the EMCS or with portable instrumentation. The supply and return air conditions (both temperature and relative humidity) are typically monitored over time. Be aware of stratification at the point of measurement; it is possible for the air temperature to vary several degrees between the top and bottom of supply air ducts. Air flow measurement at the units may be taken or flow from an existing test and balance (TAB) report can be used if available. As a last resort, the design airflow can be assumed.

3) Monitor temperature and relative humidity in the data center
One or more loggers can be used to determine the space temperature and humidity conditions. Logger placement depends upon the goals of the study; hot spots at IT equipment inlets; effectiveness of the hot aisle/cold aisle configuration (return air temperature to the cooling unit); bypasses; and mixing. Any EMCS sensors (including those accessible at the CRAC or air handler intakes) should also be trended. In particular, the temperature difference, ΔT, at the CRAC units and the temperature at the outlet of the CRAC units or air handling units should be recorded since this directly indicates the capacity capability and the airflow utilization efficiency of the entire system. Monitoring points may be determined as suggested in the *ASHRAE Thermal Guidelines for Data Center Environments Book*.

4) Inventory data center equipment
To get average rack power, count the number of racks in the data center, and assign a percentage full for each rack. The percentage full may appear to be somewhat arbitrary, but it can be very useful when the same team is used to estimate future data center loads. If desired, a complete inventory of the IT equipment could be obtained, or a partial inventory of 'typical' equipment. Rack loading will vary over time, so a per-rack (rather than typical rack) is most accurate. This inventory can be useful in tracking load growth trends and planning for future expansions.

5) Select racks for individual monitoring
There are several types of racks that may be of interest. Selecting a typical rack allows for projections of total rack load and load density, while selecting a peak load rack offers guidance as to the airflow distribution required for cooling. At a minimum, measure a rack with a typical mixture of equipment to get a sense of the variability of the power requirement. The rack needs to have its own electrical circuit and preferably has just one electrical circuit feeding it. If dif-

ferent types of servers are used, select a rack with the most common type of server.

6) kW monitoring at panel supplying rack
Unless there are built-in meters on the rack-feed circuit, install a portable kW logger at the feed to the rack. If using a built-in meter, start trend log in monitoring software. A significant length of trend time (usually a week) is required to capture data if there are any processing task load impacts on the overall data center power usage.

7) Check kW monitor with one-time reading
Using a portable kW meter, verify the reading on the data logger. This step is generally applicable to any kW logging, whether by portable loggers, EMCS, or SCADA system. If there are significant differences, check the logger or system setup (voltage setting, range, multiplier, etc.). Power meters are a very mature technology, so highly accurate results are to be expected.

8) Determine computer floor area and support areas.
A.) Total. This is the total computer room floor area, whether occupied by equipment or not.
B.) Electrically Active IT Equipment Footprint. This space consists of computers and telecommunications equipment including communication frames, servers, rack-mounted equipment, DASD, and tapes.

9) Obtain raised floor and ceiling height
The distance from the top of the sub floor to the top of the raised floor, the thickness of the raised floor tile (height of plenum blocked by tile and frame) and the ceiling height above the floor containing the IT equipment should be recorded. These measurements are useful for characterizing the data center and getting an idea of the power density it can support. While taking the raised floor measurement, look around under the floor in a few places and note the extent to which the spaces, particularly the plenum spaces used to move air, are constricted by cabling, piping, and conduits. Such observations can be very valuable when identifying and troubleshooting based on benchmarking and should be collected in a single place – useful information formats can range from a three-ring binder section for small projects all the way up to a spreadsheet with fields to sort by system, area, noted digital pictures taken, locations of documentation, and other information, for larger multi-building projects.

10) Check CRACs, CRAHs, or AHUs for status
Status for these types of equipment means whether they are cooling, heating, humidifying, or dehumidifying. Determine the status of all units at approximately the same time and record. This activity should be repeated several times at different times of day during the benchmarking to look for abnormal operation. In particular, note the humidification operation of units serving the same space.

11) Check CRACs, CRAHs, or AHUs for control setpoints and deadbands
Check and record the temperature and humidity setpoints and deadbands (minimum and maximum numbers for both). These may be at the units (typical for CRACs and CRAHs) or the EMCS programming for AHUs. Also note:
1.) Whether the system can humidify (and if so how: steam or electric ultrasonic, infrared, boiling), and 2.) If there is reheat and whether it is electric, steam, or hot water.

12) Measure PDU electrical power loss
To determine electrical power losses in the Power Distribution Units (or equivalent distribution transformers in systems without PDUs), measure the input and output power at the same time. Some PDUs are equipped with internal meters to allow this, though they should be checked with portable meters to assess their accuracy. Note that kVA (or the product of voltage and current) is not a reliable substitute for kW. Design data can be substituted if there's no easy access for power measurements.

Figure 3: UPS efficiency at different benchmarked data center facilities.

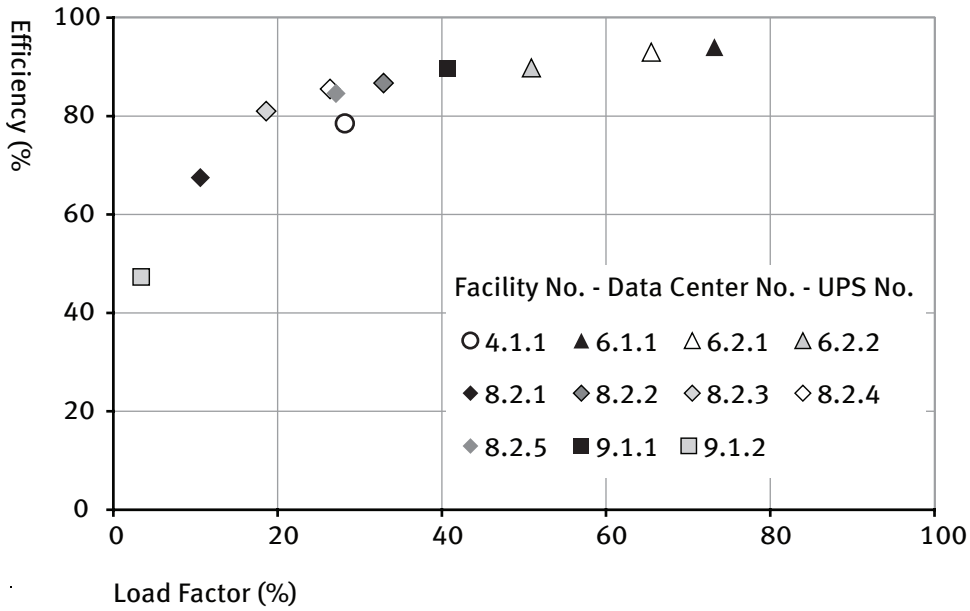

13) Measure UPS electrical power loss

To determine electrical power losses in Uninterruptible Power Supplies, measure input and output power at the same time. Some UPSs are equipped with internal meters that display this, though they should be checked with portable meters to verify their accuracy. Note that kVA (or the product of voltage and current) is not a reliable substitute for true kW meters, which measure the sometimes significant impact of power factor on the total power used. Determine and record the output power and the rated maximum for the UPS and calculate the percent load factor by dividing the output load by the rated maximum. The load factor has a significant impact on the measured UPS efficiency. The graph in Figure 3 summarizes UPS efficiency at different benchmarked sites. The exact shape of the partload curve varies by manufacturer, but a drop at low load factors is a common characteristic of battery-based systems. Similar model-specific performance at part-load information should be available for all major UPS systems.

14) Measure Standby generator electrical power loss

Losses in standby generators include jacket heaters and battery chargers. Measure and log these electrical power loads and the ambient temperature at the generator(s). Also note the jacket water temperature setpoint and measure the temperature (one reading is adequate).

15) Measure Full data center IT equipment electrical power

In order to reduce the number of points required to determine the total IT equipment electrical load, use the most upstream feeders available and feasible to access, serving only IT equipment. If measured upstream of UPSs and PDUs, be sure to subtract off these losses in the data reduction phase (see #9). Typically, a single meter can be used to capture the power usage of many racks or even a small data center.

16) HVAC equipment measurement

A.) Chiller plant. The chiller plant includes the chiller(s), cooling tower(s), and pump(s). Sometimes it is easiest to monitor the entire motor control center feeding the chiller plant as one kW point, and any other loads on it can be monitored for subtraction (or, if constant, can be read with a one-time reading for subtraction). Pumps without VFDs can be spot checked or trended for a short time and assumed to be constant power devices. The chiller and cooling tower systems tend to have high variations in load over time and should be trended. Variable speed pumps should also be trended.

In order to optimize the chiller plant, each component should be logged separately, as well as chilled water and condenser water supply and return temperatures and chilled water flow to allow for the determination of tons of cooling supplied. An accurate tonnage measurement requires highly accurate temperature and flow sensors. The difference between the supply and return temperatures (ΔT) is often 5-10 degrees, or less. Use of non-matched, typical temperature sensors that are only accurate to +/-1 could introduce an uncertainty of almost 50% into the ΔT measurement, rendering it useless. Flow can be determined by clamp-on ultrasonic or by an installed meter, which would allow for continuous monitoring. Care should be taken when selecting a flow meter for installation to select a unit with minimal chance of fouling and calibration drift.
Ideally, plant cooling load will be measured simultaneously with load in the data center.

Plant tonnage can be determined from the following equation (valid for lower temperature water, 40 – 100 F deg):

$$\text{Tons} = \frac{(\text{Return Temperature °F - Supply Temperature °F}) \times \text{Flow GPM}}{24}$$

B.) CRAC/Direct expansion cooler measurement. Each air conditioning unit should be monitored for power use simultaneously with the equipment load in the room served. Monitor power usage of both the inside fan unit (including humidity control) and the external compressor and condenser units. Where water cooled condenser systems that locate the compressor indoors with the supply fan (as opposed to an outdoor compressor/condenser unit) are used, note that it may not be feasible to isolate the fan power energy from the cooler energy.

C.) Air handling equipment. Aside from the temperatures and flows noted above, log the input power to this equipment. It can be done in aggregate if the electrical system configuration is suitable. Note that for most air cooled CRAC units, the interior unit power usage should be assigned to fan power.

D.) Boiler plant (where usage is non-negligible). The boiler plant includes the boiler(s) and pump(s). Log the input electrical power to these as well as the supply and return temperatures. Significant boiler plant usage generally indicates a lower efficiency isothermal humidification system (as opposed to an adiabatic evaporative system, similar in operation principle to the simple “swamp cooler”).

17) Measure lighting power
For the lighting in the data center, measure and log the electrical power consumption. Often this can be done at a single lighting feeder. Make a note of any lighting controls present and their settings (approximate lumens setting, sensor sensitivity, or schedule).

7. Collect data and remove monitoring equipment

Using portable equipment, gather the one-time reads per the monitoring plan. Let the logging and trending run for at least a full week. At the end of the week, download all trends for analysis. Collect data from any central control system trends in an appropriate format (ascii, comma delimited, .xls, etc.) and store it with the stand-alone trending equipment data.

Using the same personnel and safety procedures as where equipment was installed, gather and record the data from the data loggers, EMCS, and SCADA system. Verify that the data are complete. Remove loggers, and let the facilities personnel know that they can stop trending the points that were set up for the study.

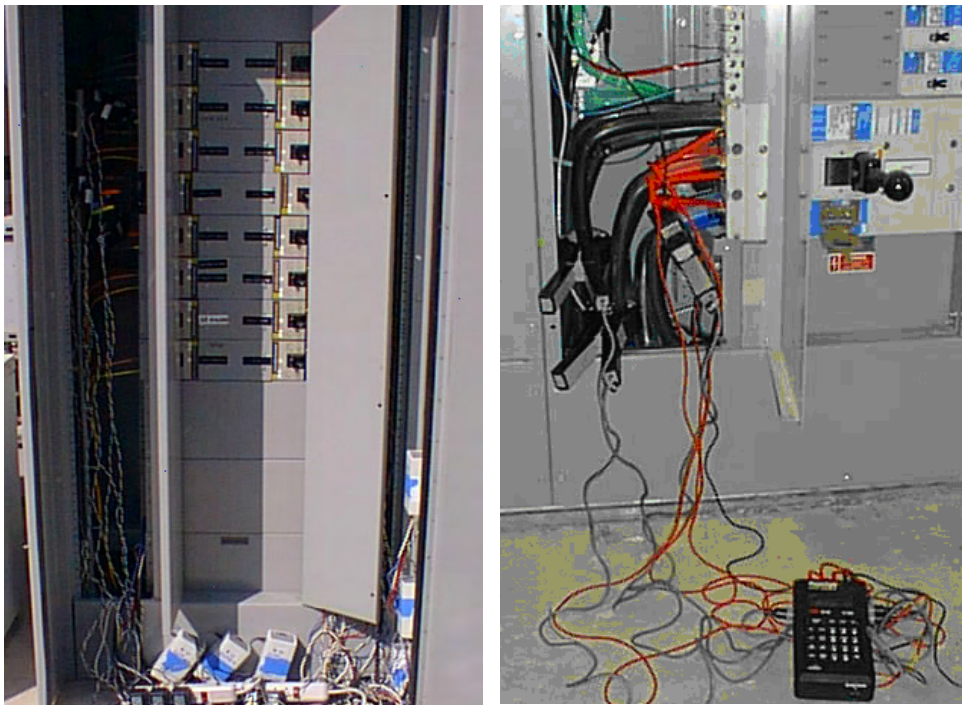

Figure 4: Two kinds of power meters deployed to trend panels (five circuits on left, one on right). Note: safety tape and barriers are not shown.

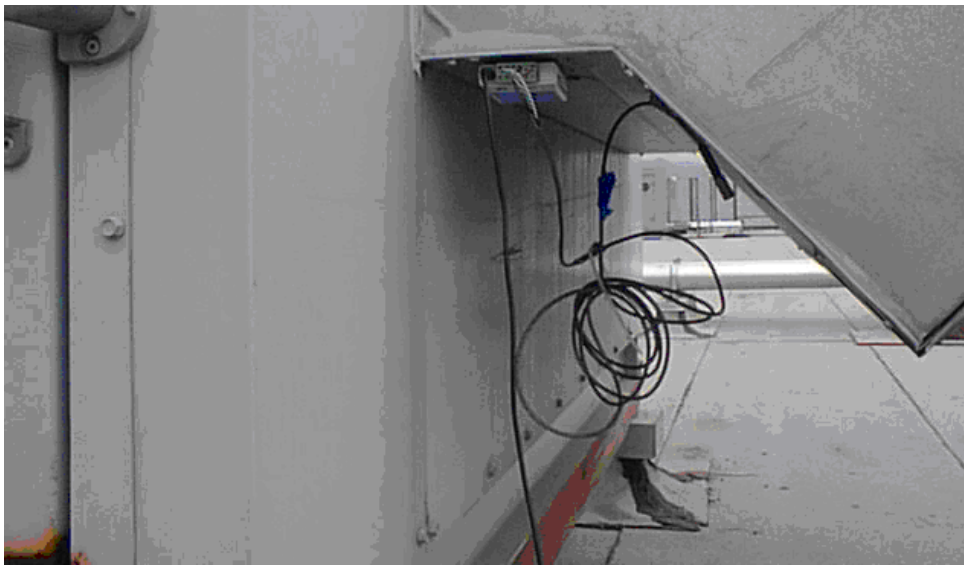

Figure 5: Outside air temperature and humidity logger and sensor deployed in shaded area under outside air intake hood (same sensor and logger can be deployed inside data center).

Figure 6: Ultrasonic 'strap-on' temporary flow meters.

Figure 7: Water temperature sensor installed in test port ("Pete's plug") with logger; note masking tape on logger used to field note time of installation, removal time, sensor used and measurement point name.

8. Analyze and reduce data

After a period of gathering data (defined in the plan), the monitoring equipment is removed, and the data are analyzed to create an end-use breakdown of the data center's energy use.

Using spreadsheets, convert the raw data collected into useful end-use data, consistent with the goals of the survey. Again, three very useful benchmarks are: 1.) The W/sf of IT equipment, 2.) The ratio of IT energy to HVAC energy, and 3.) The ratio of IT energy to total energy. Also commonly included are the cooling plant performance (e.g. overall kW/ton for the chillers, towers, and pumps) and an end-use breakdown for the data center, i.e. lighting, HVAC, IT equipment, and losses in the PDUs, UPSs, and standby generators. Other common benchmarks are listed in the Appendix.

Trends of data should be graphed for examination. In particular, check for rapidly cycling loads, fan operation or humidification/dehumidification operation. These are signs of poor control loop tuning or system sizing that can lead to premature equipment failure, poorer control resolution, and significant energy waste. The temperature trends in the computer room should be relatively flat, with the temperature at the intake of the computer the key control parameter: temperatures 'behind' the racks in the hot aisle do not impact the equipment and should be significantly higher than at the intakes. Essentially all of the trended data is most easily and valuably interpreted graphically: control signals to valves, cycling of compressors, fan speed, and power usage of racks.

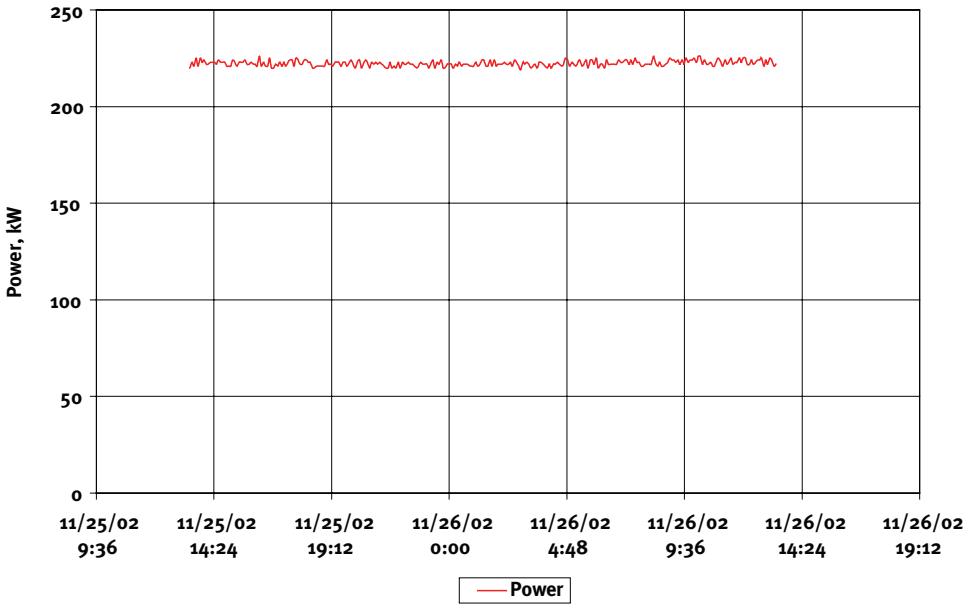

Figure 8: Sample whole-facility data center load variation over 24 hours.

After analyzing the benchmarking results, staff can meet to identify opportunities and determine appropriate next steps to correct problems found, immediately implement low or no-cost system improvements; and document any future design or IT equipment acquisition guidelines and optimization projects.

9. Compare findings to existing benchmarks

Once the data are reduced and the benchmark numbers determined, one can compare the benchmarks to those of other data centers, both within the company and for data centers previously studied. If the LBNL template was followed, comparisons to other benchmark data collected by LBNL is now possible. Any significant deviations from the normal comparison data should be investigated and identified. Significantly greater energy use usually indicates a very good opportunity for savings, while significantly lower energy use indicates a potential best practice approach that should be emulated in associated and future data centers. (A spreadsheet of basic benchmarking data for prior data centers available on the internet at http://hightech.lbl.gov/benchmarking_dc.html).

Beyond simply allowing the comparison of system efficiency, often operational problems are recognized. For example, the graph in Figure 9 suggests a control problem causing unnecessary compressor cycling - a reliability and efficiency concern.

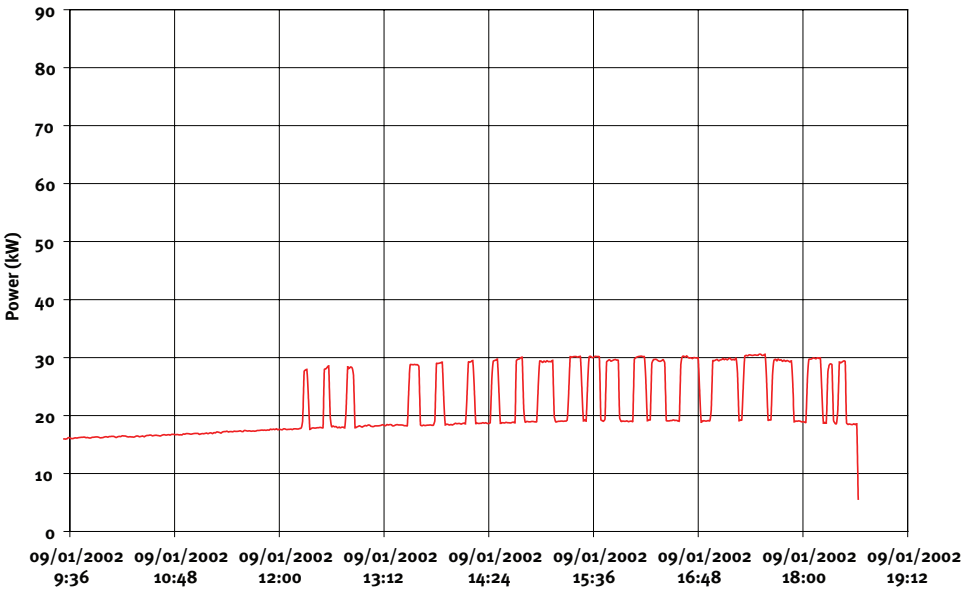

Figure 9: Facility A data center 1 Chiller 1 power.

Air handler coil control loops often show control loops that have not been tuned for the specific conditions of datacenter use. The chart in Figure 10 shows the results of monitoring an air handler where the chilled water and hot water coil are cycling excessively considering the near-constant datacenter load they serve. Unstable control behavior from the systems serving a stable datacenter load are inefficient at best and instigators of equipment failure and/or loss of space control at worst.

Figure 10: Data Center Facility 5 AHU-2 Conditions.

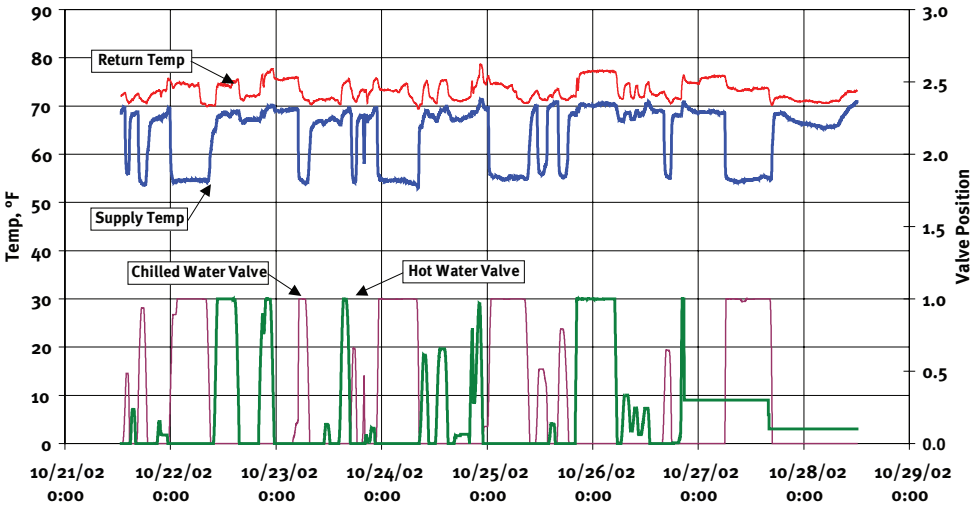

10. Compare data center to best practices

Several Design Guides based upon observed best practices have been prepared for Pacific Gas and Electric Company and are available through their Energy Design Resources website. Technologies include:

1) Air And Water-Side Economizers/Free Cooling
2) Cooling Source And Distribution Systems
3) Optimized Centralized Air Handling
4) Chiller Plant Optimization
5) Air Flow Management
6) Temperature Setpoints For Chilled Water And Supply Air
7) Humidity Controls
8) Evaporative Cooling
9) Reduced Losses In UPSs
10) On-Site Cogeneration
11) Efficient Lighting Sources And Controls
12) Reduced Losses In Standby Generators
13) Efficient Power Supplies And Power Distribution
14) Direct Liquid Cooling Of IT Equipment Opportunities

11. Share results

Consider sharing your results, or an anonymous version of your report with the data center community. At the very least, anonymous contribution of measured benchmarks helps the benchmarking resource to continue to grow. Design best practices are continuously and incrementally improved as the quantity and quality of available data grows. Sharing of best practices helps innovative ideas come from a variety of sources, and helps energy efficient design standards and technologies to evolve.

Conclusion

Optimizing a complex system such as a data center in order to minimize operating costs and maximize capacity and reliability requires that its performance be measured. The benchmarking procedure described in this guide has shown a simple approach to making a detailed measurement of data center performance.

The snapshot of data center performance provided by the benchmarking procedure described here allows a facility's efficiency to be determined and compared against similar facilities on a system-by-system basis. Improved performance and cost savings can be realized by identifying and duplicating the most efficient systems and equipment in the future, and by correcting inefficient systems and equipment. Also, this procedure is intended to convey the sense of the data center as a whole system, rather than a collection of diverse pieces of equipment. Specific metrics and design strategies can contribute to as well as shape this whole system perception.

In the rapidly maturing data center field, benchmarking allows facilities operators to evolve beyond the crude rule-of-thumb design practice of "if its cool, its good enough" to optimize their systems for maximum reliability, minimum first costs and lower operating costs – all while maintaining a "Much-better-than-good-enough" critical environment.

References

Thermal Guidelines for Data Processing Environments; American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc., 2004, Atlanta, GA

Appendix A: Data Center Benchmarking Data Tables

The following data tables will guide the data collection process. Proxy data from design documents or testing results may be used for data points for which direct measurement is not possible.

All time series data for each site can be compiled into a single spreadsheet workbook. Trended data can be reported with in-service average, minimum, and maximum values.

I. Data Center Metrics

All metrics will be calculated using data collected below.

Description	Units
AHU Efficiency	kW/cfm
Annual Fuel Usage	MBtu/sf/yr
Annual Electricity Usage	kWh/sf/yr
Annual Energy Usage	MBtu/sf/yr
Cooling Load Density	sf/ton
IT Equipment Power Density	W/sf
Lighting Power Density	W/sf

Data Center Data Points

Target	Units	Data Source (circle one)	Duration
Total Recirculation Fan (Total CRAC) Usage	kW	From electrical panels	Spot (trended where possible)
Total Make-Up Air Handler Usage	kW	From electrical panels	Spot
Total IT equipment Power Usage	kW	From electrical panels	Spot (trended where possible)
Chilled Water Plant (if chilled water plant is on same meter or in same building as data center that is being measured)	kW	From electrical panel or from chilled water plant measurements	1 week
Rack power usage, 1 typical	kW	From electrical panels	1 week
Number of racks	number	Observation	Spot
Rack power usage, average	kW	Calculated	N/A
Other Power Usage	kW	From electrical panels	Spot
Data Center Temperatures (Approx. 6, located to identify range of temperatures and air distribution issues)	Deg F	Temperature Sensor	1 week
Humidity Conditions	R.H.	Humidity Sensor	1 week
Annual Electricity Use, one year	kWh/yr	Utility bills	N/A
Annual Fuel Use, one year	Therm/yr	Utility bills	N/A
Annual Electricity Use, 3 prior years	kWh/yr	Utility bills	N/A
Annual Fuel Use, 3 prior years	Therm/yr	Utility bills	N/A
Peak Power	kW	Utility bills	N/A
Average Power Factor	%	Utility bills	N/A
Facility (Total Building) Area	sf	Drawings	N/A
Data center Area (“Electrically Active Floor Space”)	sf	Drawings	N/A
Fraction of Data Center in use (Fullness factor; taking into account available floor space and available space within existing racks)	%	Area measurements and rack observations	Spot
Extrapolated IT equipment Power Usage for Fully Utilized Data Center	kW	Calculated	N/A

Note: collect data for each pump as applicable (primary, secondary, tertiary, etc.

CRAC Metrics

All metrics will be calculated using data collected below.

Description	Units
Fan Efficiency	kW/CFM
Cooling Efficiency	kW/Ton
System ΔT (Return T - Supply T)	˚F

CRAC Unit Data Points

Measure a minimum of 3 sampling CRAC units or air handlers per data center.

Target	Units	Data Source	Duration
Air Flow	cfm	(Designed, TAB report)	N/A
Fan Power	kW	3Φ True Power	Spot
VFD Speed	Hz	VFD	Spot
Setpoint Temperature	Deg F	Control System	Spot
Return Air Temperature	Deg F	10k Thermistor	1 week
Supply Air Temperature	Deg F	10k Thermistor	1 week
RH Setpoint	RH	Control System	Spot
Supply RH	RH	RH Sensor	1 week
Return RH	RH	RH Sensor	1 week
Status (cooling, heating, humidifying, dehumidifying)	Misc.	Observation (trend if on-site monitoring in place)	Spot (One week trend if possible)
Cooling Load	Tons	Calculated	N/A

II. Chiller Metrics

All metrics will be calculated using data collected below.

Description	Units
Chiller Efficiency	kW/ton
Tower Efficiency	kW/ton
Condenser Water Pumps Efficiency	kW/ton
Chilled Water Pumps Efficiency	kW/ton
Total Chilled Water Plant Efficiency	kW/ton
Plant Efficiency While Free Cooling	kW/ton

Chiller Data Points

Target	Units	Data Source	Duration
Chiller Power	kW	3Φ True Power	1 week
Primary Chilled Water Pump Power	kW	3Φ True Power	Spot
Secondary Chilled Water Pump Power	kW	3Φ True Power	1 week
Chilled Water Supply Temperature	Deg F	10k Thermistor	1 week
Chilled Water Return Temperature	Deg F	10k Thermistor	1 week
Chilled Water Pump Head (ΔP)	Feet of head	Pressure Transducer or single gauge	Spot
Chilled Water Flow	gpm	Ultrasonic Flow	1 week
Cooling Tower Power	kW	3Φ True Power	1 week
Condenser Water Pump Head	Feet of head	Pressure Transducer or single gauge	Spot
Condenser Water Pump Power	kW	3Φ True Power	Spot
Condenser Water Supply Temperature	Deg F	10k Thermistor	1 week
Chiller Cooling Load	Tons	Calculated	N/A

Boiler Plant Metrics

All metrics will be calculated using data collected below.

Description	Units
Hot Water Pumping Efficiency	kW/MBtu

Boiler Plant Data Points

Target	Units	Data Source	Duration
Hot Water Supply Temperature	Deg F	10k Thermistor	1 week
Hot Water Return Temperature	Deg F	10k Thermistor	1 week
Hot Water Pumping Power	kW	3Φ True Power	1 week
Hot Water Flow	gpm	Ultrasonic Flow	1 week
Boiler Gas Input	CFH	From Gas Meter (if available)	1 week
Boiler Efficiency	%	Design	N/A
Total Gas Use	Therms	Calculated	N/A

III. Other Metrics

Description	Units
Standby Generator Losses	kW loss / kVA rating of generator
UPS Losses	kW loss/ kW Usage
Lighting	W/sf
Height of raised floor	inches

Utilities Data Points

Target	Units	Data Source	Duration
Backup generator(s) size(s)	kVA	Label Observation	N/A
Backup generator standby loss (heaters, chargers, fuel system, controls)	kW	Power Measurement	1 week
Backup generator ambient temp	Deg F	Temp Sensor	1 week
Backup generator heater setpoint	Deg F	Observation	Spot
Backup generator water jacket temperature	Deg F	Temp Sensor	1 week
Backup generator test consumption	Gallons of Diesel?	Site records or measured during a test	
UPS Load	kW	UPS Interface Panel	Spot
UPS Rating	kVA	Label observation	Spot
UPS Loss	kW	UPS interface panel or measurement	Spot
PDU Load	kW	PDU Interface Panel	Spot
PDU Rating	kVA	Label observation	Spot
PDU Loss	kW	PDU interface panel or measurement	Spot

Weather Data Points

Target	Units	Data Source	Duration
Outside Air Dry Bulb Temperature	Deg F	Temp	1 week
Outside Air Wet Bulb Temperature	Deg F	RH Sensor	1 week

Appendix B: Checklists and Data Collection Forms

Sample Monitoring Setup Checklists (Develop similar checklists for actual tools used at site)

PowerSight Power Measurement

__ Verify that the three phases of current are approximately equal. Verify the three phases of voltage are approximately equal. A significant imbalance suggests a poor CT or voltage probe connection. A significant imbalance can also be a (generally undesirable) property of the electrical system (VFD, supply, etc.).

__ Verify monitoring recording is begun and will not end until after the desired time.

__ Verify monitoring mode is recording correct parameters (True Power) and is set to end when full (not 'overwrite when full' or 'continue until stopped').

__ Verify unit is plugged in and receiving outlet power for long trends (top LED).

Temperature and RH with Pace XR440 Pocket Logger

__ Verify channels are set to the correct sensor type

__ Verify the linear scale for the RH logging is set to values on the RH wire's tag.

__ Verify the channels are set on

__ Verify the channel names are correct. It is difficult to change these names and they are likely to appear in any graphs of the data included in the final report.

__ Verify monitoring recording is begun and will not end until after the desired time.

Elite Pro Power Logger

__ Verify white lead is (+) and black lead is (-) on CTs.

__ Verify clock is synchronized with PC clock.

__ Verify "memory type" is set to linear.

__ Verify CT arrows are facing against the flow of power.

__ Verify white alligator clip is connected to a good ground or neutral.

__ Verify unit is plugged into outlet.

__ Verify unit is logging; LED will blink ~ once every 5 seconds.

__ Verify equipment is yielding reasonable data on a real time reading.

__ Verify unit is labeled with equipment ID, time, and panel #.

Generic Data Center Data Collection Form

Date: _____

Team Member(s): _____

Electric Panel No. & Designation	Category (Data Center Cooling, Office Conditioning, etc.)	Measured Watts (kW)	Power Factor	Voltage Va / Vb / Vc	Measured Amps Ia / Ib / Ic / (In)	Circuit Rating (Amp)	Notes

Attachment VII: Improvement of Uninterruptible Power Supplies

High Performance Buildings: Data Centers

Uninterruptible Power Supplies (UPS)

Table of Contents

1. Executive Summary 1
1. Current UPS Efficiencies 1
2. Proposed UPS Efficiency Levels and Labeling 1
3. UPS Energy Use and Potential Savings 2
2. Current UPS Market 5
2.1 Assessment of UPS Topologies 5
2.1.1 Overview of Offline or Standby UPS Topology 5
2.1.2 Overview of Inline or Line-Interactive UPS Topology 6
2.1.3 Overview of Delta-Conversion Topology 8
2.1.4 Overview of Double-Conversion or Online UPS Topology 8
2.1.5 Overview of Flywheel or Kinetic UPS Topology 9
2.1.6 Overview of the U.S. UPS Market (Data Center/IT Sector) 10
2.2 UPS Test Protocol 14
2.2.1 Factory/Laboratory UPS Efficiency Test Method 14
2.2.2 Field UPS Efficiency Test Method 15
3. UPS Efficiency Data 16
3.1 UPS Efficiencies 16
3.1.1 UPS Efficiency and Percent Load 16
3.1.2 UPS Efficiency and Load Characteristics 17
3.1.3 Results of UPS Efficiency Testing: Field Measurements 18
3.1.4 Results of UPS Efficiency Testing: Factory/Laboratory Measurements 19
3.1.5 Investigation of Special High Efficiency Modes in Select UPSs 21
3.2 Redundancy Evaluation 26
3.2.1 Capacity or "N" Configuration 26
3.2.2 Isolated Redundant Configuration 27
3.2.3 "N+1" Parallel, Single-Bus Configuration 27
3.2.4 "2N" and "2(N+1)" Dual-Bus Configurations 28
4. Proposed Efficiency Specifications 30
4.1 UPS Operation 30
4.1.1 Total Cost of Ownership 30
4.2 Proposed Efficiency Label 31
4.2.1 Labeling UPS Efficiency 31
4.2.2 Proposed UPS Efficiency Standard for the U.S. Market 33
5. DC Architecture Scoping Study 40
5.1 Market Trends in IT Load Power Configuration 40
5.2 Efficiency of AC- and DC-based Systems 41
5.3 Demonstration Plans for DC Architecture 42

1. Executive Summary

As part of the California Energy Commission's PIER (Public Interest Energy Research) initiative on efficient data centers, Ecos Consulting and EPRI Solutions were tasked with investigating uninterruptible power supplies (UPSs) used in data centers and other high-tech facilities to guard against interruptions in services, primarily from electrical outages. The goals for this project include a complete analysis of distribution of UPS design efficiencies, which included the development of an accepted UPS test protocol, and the broad circulation of efficiency findings to the industry through the CEC, PIER, Lawrence Berkeley National Laboratories (LBNL), and other industry and efficiency venues, such as www.efficientpowersupplies.org. Using the efficiency findings and market data, we also estimated the overall energy consumption of UPSs in the United States, especially in the State of California, as well as the potential savings from the use of more efficient units. Finally, we also proposed efficiency levels and a labeling scheme for various types of UPS units as a way to encourage the use of more efficient UPSs.

1. Current UPS Efficiencies

Our results indicate that there can be fairly wide variation in efficiency even between various UPS configurations. The table below summarizes the characteristic efficiency of a number of UPS topologies at various load conditions and shows the average efficiency for all of the UPSs measured. The table below also shows that regardless of the configuration or type, UPSs tend to be more efficient at full rather than part load.

Table ES1: Characteristic Efficiency of UPS Topologies

UPS Topology	Efficiency at 25% Load	Efficiency at 50% Load	Efficiency at 75% Load	Efficiency at 100% Load
Delta-Conversion	93% - 94%	96% - 97%	97%	97%
Double-conversion	81% - 93%	85% - 94%	86% - 95%	86% - 95%
Flywheel	unknown[1]	97%	98%	98%
AVERAGE OF ALL UNITS	86%	89%	90%	90%

Due to the scope of this research effort, it was impractical to determine what specific elements of a model's design contributed to its measured efficiency; however, our results indicate that some broader design decisions, such as the general UPS topology, can be indicative of efficiency. In particular, flywheel and delta-conversion UPSs exhibited best-in-class efficiency compared to the double-conversion UPSs that we measured. *We stress that no strong conclusions should be drawn about the inherent efficiency of either the flywheel or delta-conversion topologies because of the relatively small sample size in our study.* Research on the efficiency of both flywheel and delta-conversion UPSs should continue as manufacturers continue to improve their designs and diversify product lines.

2. Proposed UPS Efficiency Levels and Labeling

In the early stages of this project, Ecos Consulting and EPRI Solutions uncovered a proposed energy efficiency and power quality labeling scheme for single-phase, "small" UPSs being developed by the Swiss Federal Office of Energy. The Swiss' proposed label, documents the measured power losses of the UPS in different modes of operation and provides an estimate of the annual energy consumed by energy losses in the UPS design. The label also reports the UPS's ability to handle different types of power disturbances and the power quality (power factor and total harmonic distortion) that the UPS

[1] Data on flywheel efficiency was only presented down to about 30% load.

presents to the grid. It was designed to match the style of existing EU labels for other electric appliances.

In the interests of harmonizing with international efforts to report and promote UPS energy efficiency in a standardized format, our team created a modified version of the proposed Swiss label and a draft report summarizing the criteria that would be used to evaluate UPS efficiency in the U.S. market under the label. Our proposed label, shown below in Figure ES1 includes the following changes:

- The label reports differing levels of energy conversion efficiency rather than energy losses. It was our opinion that reporting efficiency was more appealing to a potential UPS purchaser.
- We have eliminated a field in the label that reports energy losses incurred by operating the UPS for 2000 hours with no load.
- To help direct the focus of the label solely on energy efficiency and to simplify its design, we decided to eliminate the various tables that report on filtering of power disturbances and power quality.

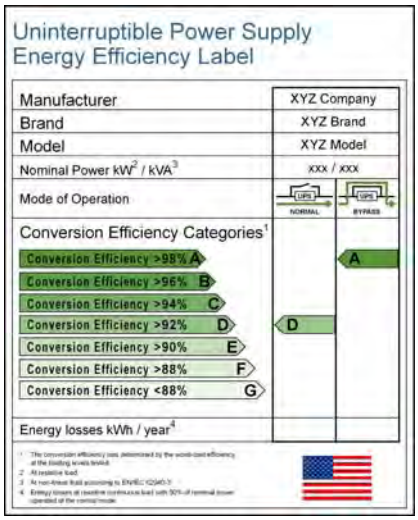

Figure ES1

3. UPS Energy Use and Potential Savings

Using efficiency data and our estimates of UPS stock, we constructed a first-order estimate of UPS energy use for the US. This energy use and savings estimate is limited to the data center/IT sector due to the scope of our research, in which we examined the efficiency of data centers only.

Based on our estimates, the data center/IT sector currently consumes roughly 7.1 TWh (7.1 billion kWh) of electricity per year. We estimate that the State of California alone, which contains roughly 15% of the nation's data center floor space,[2] consumes about 1 billion kWh of electricity and spends about $100 million per year in electric bills due (mostly) to power conversion losses in data center UPSs.[3] Our national energy use figures are noticeably higher than past estimates made by Arthur D. Little for the U.S. Department of Energy, which estimated UPS energy consumption in the U.S. "IT/telecom" sector at 5.8 TWh per year.[4] Both Ecos' and Arthur D. Little's energy use estimates are shown in Table ES2..

---

[2] R. Juarez, M. Alic, K. Chetan and B. Johnson. *SpaceDex III – Hosting Space: Not All Hosting Space is Created Equal – Smart, Complex Space Takes Data Center Stage.* Boston: Robertson Stephens Inc. 2001.

[3] Because California contains 15% of the nation's data center floor space, its UPSs consume 15% of the energy consumed by all UPSs in the US.

[4] K. Roth, F. Goldstein and J. Kleinman. *Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings, Volume I: Energy Consumption Baseline.* Cambridge, MA: Arthur D. Little, 2002. p. 84

Table ES 2: Comparison of UPS Energy Use Estimates

	DOE/ADL 2002 Estimate of UPS Annual Energy Consumption in U.S. IT/telecom Sector	**Ecos/EPRI 2004 Estimate of UPS Annual Energy Consumption in U.S. IT/data center Sector**
Estimated National UPS Stock	151,761	250,343
Nationwide Annual Energy Consumption (TWh)	5.8	7.1
Estimated California UPS Stock	n/a	37,551
California Annual Energy Consumption (TWh)	n/a	1.1

Using a "high efficiency" and "typical" cases for UPS energy consumption, we have estimated the energy savings potential for the U.S. The introduction of a voluntary UPS efficiency specification that would be administered under a program such as ENERGY STAR or other institutions may result in a percentage of the UPS market meeting efficiency levels, resulting in partial realization of the savings potential.

If, for example, 25% of the UPSs installed in the U.S. were able to comply with the proposed UPS efficiency levels, the nation would save about 700 million kWh of electricity per year, amounting to about $70 million in cost savings. We have summarized our energy savings scenarios for mandatory and voluntary specifications in Table ES3, along with similar energy savings estimates that apply only to the State of California, which we assume contains about 15% of the nation's UPSs.[5]

The energy savings figures shown here could be significantly larger when examining the *entire* UPS market and not just the data center/IT sector. For UPSs above 5 kVA in power output, the data center/IT sector only accounts for roughly two thirds of the installed units. An additional one third of UPSs installed in small offices, industrial facilities, hospitals and the like have not been included in our energy use and savings estimates. Our analysis also excludes the below 5 kVA UPS market segment, more typical of small office applications, residential installations, etc. The combined sales of UPSs below 5 kVA amounted to over 8 million units in 2004, close to 200 times the unit sales volume of UPSs greater than 5 kVA in power output.[6] Although these are lower power devices than the larger UPSs examined in this report and would thus have smaller per-unit power losses, they might represent an attractive energy savings opportunity simply due to the large number of units in operation that should be studied in the future.

[5] Based on estimates that California contains 15% of the nation's data center floor space. R. Juarez et al. 2001.
[6] Frost and Sullivan. 2004.

Table ES3: Energy Savings Scenarios for UPS Efficiency

	Mandatory Specification – 100% of Installed Units Comply	**Voluntary Specification – 25% of Installed Units Comply**
Nationwide Annual Energy Savings (TWh)	2.8	0.71
Nationwide Annual Utility Cost Savings (million USD)	$280	$71
California Annual Energy Savings (TWh)	0.42	0.11
California Annual Utility Cost Savings (million USD)	$42	$11

It should be noted that additional energy savings can be realized from reduced facility cooling needs because less waste heat is generated from the conversion and storage equipment. It has been estimated that the secondary savings from reduced cooling needs as a result of more efficient facilities is on the same order of magnitude as the direct, primary savings from efficient conversions. Finally, there is also additional savings from reduced capital investment in equipment if a more efficient system is selected, as well as reduced real estate investment from smaller equipment footprints. These additional savings can be realized up front rather than through reduced operational costs.

2. Current UPS Market

Assessment of UPS Topologies

Our first task was to survey the various types of UPS design topologies available on the market today – particularly those that are commonly used in data centers – and assess the efficiency of various designs through a data collection effort. A general description of the topologies follows in this section, whereas our analysis of the efficiency data gathered from the field and through manufacturer participation appears in Section 1.3. Some of the major questions that we hoped to answer in our overview of UPS topologies included:

- What specific factors separate one design topology from another?
- Which topologies are likely to be relevant to data center applications?
- How many UPSs of the various topologies are installed in data center/IT facilities across the U.S.?

Overview of Offline or Standby UPS Topology

The offline or standby UPS is the simplest and oldest UPS topology. The operating principle of a standby UPS is, as the name suggests, to stand by in case of a disruption in power to a critical load. The UPS only delivers power to its connected load when power from the utility does not meet specifications. The vast majority of the time when utility power is deemed acceptable, the UPS allows the connected load to draw power directly from AC mains, using a small amount of power to trickle charge the UPS batteries.

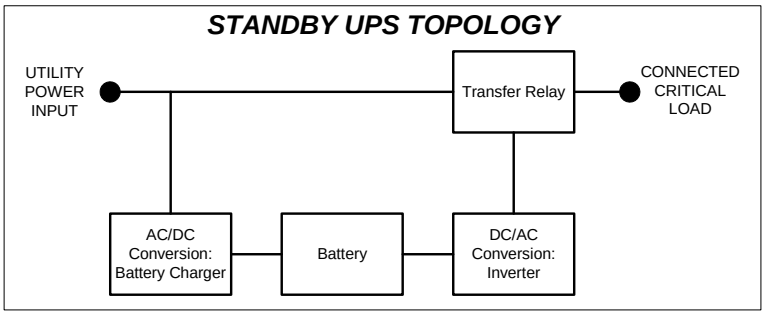

Figure 1

Figure 1 illustrates the basic subsystems of a standby UPS. During "normal mode," when utility power is available and within acceptable quality tolerances, the UPS supplies the connected load directly with utility power. A small amount of AC power is used to trickle charge the UPS's battery. For example, small standby UPSs designed to power individual workstations (rated less than 1 kVA) typically have fixed losses associated with this battery charging between 5 and 10 watts. An illustration of normal mode for standby UPSs is shown in Figure 2.

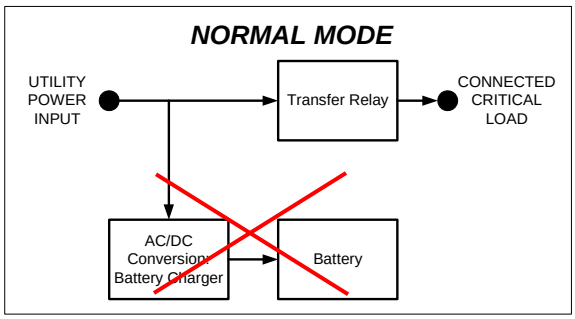

Figure 2

When there is an outage or the utility power goes beyond predefined quality specifications, the transfer relay isolates the load from utility power, and the UPS provides AC electricity through the battery and inverter in “stored energy mode.” This is illustrated in Figure 3.

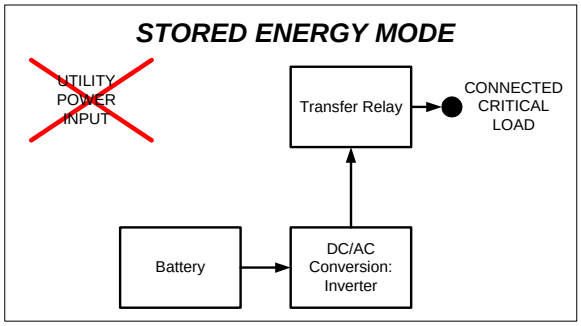

Figure 3

For most office and individual workstation backup applications, the UPS delivers unconditioned power to the connected load during normal mode. In some higher quality designs, utility power might be slightly conditioned upstream of the load so that it meets tighter voltage specifications.

Standby UPSs are most common in low-power applications (0 – 2 kVA) such as individual workstation backup power, small office server backup power, etc. Even though some higher quality designs might be capable of filtering utility power or protecting loads from power surges, the lack of complete power conditioning and load isolation generally make standby UPSs unsuitable to critical loads such as industrial facilities or data centers.

Overview of Inline or Line-Interactive UPS Topology

So-called line-interactive UPSs differ from standby UPSs because they are able to provide improved conditioning to utility power by interacting with the incoming electricity. They achieve this by placing inverter/battery charging circuitry or transformers in parallel with the AC utility signal. This design allows a line-interactive UPS to compensate for over- or under- voltages in the incoming utility power, but often allows commode mode noise and perturbations in AC frequency to pass through to the load.

The line-interactive topology is illustrated in Figure 4 below.

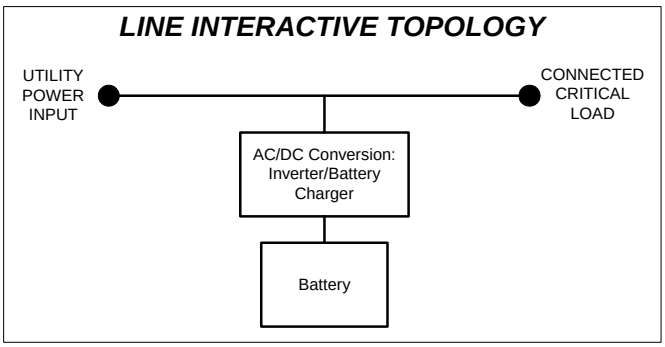

Figure 4

As with a standby UPS, a line-interactive UPS allows utility power to flow to the connected load in normal mode, but because of the parallel configuration of the inverter/battery charger in this topology, the UPS can condition incoming utility power for over- or under-voltages by interacting with the AC main. Small perturbations in frequency and common mode noise usually cannot be filtered out. The battery is continuously trickle charged during normal operation through the inverter/battery charger similar to the standby UPS topology. The normal mode of operation for a line-interactive UPS is illustrated in Figure 5 below.

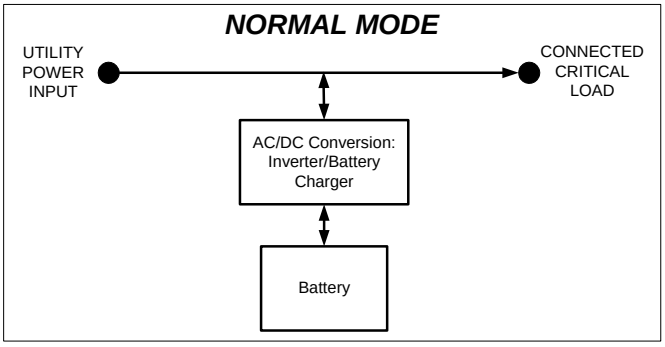

Figure 5

When utility power is unavailable or reaches unacceptable limits, a line-interactive UPS, like a standby UPS, will enter stored energy mode. The UPS disconnects the load from utility power and reroutes this load with a static switch to backup AC power, provided by the battery through the inverter. This is illustrated in Figure 6 below.

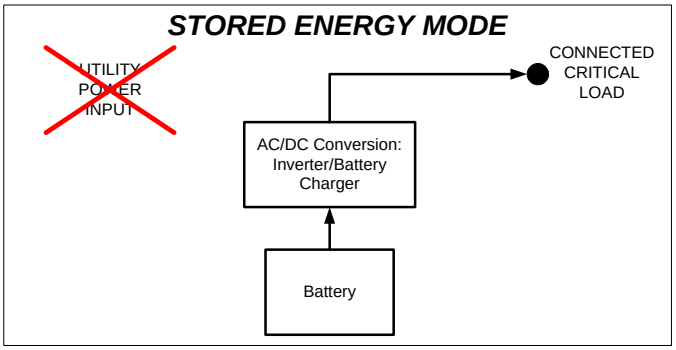

Figure 6

Most line-interactive UPSs are found in low-power applications ranging from 0 – 5 kVA.

Overview of Delta-Conversion Topology

The delta-conversion topology is a proprietary topology developed by American Power Conversion (APC), sold in sizes large enough to accommodate multiple racks or even an entire data center room. It is a unique design compared to line-interactive and double-conversion topologies. Delta-conversion UPSs use a special transformer configuration to interface between the load utility power, with a "delta" inverter in the transformer secondary to regulate input current and power. With this configuration, the UPS can regulate the magnitude, wave shape, and power factor of the current supplied at the UPS input, while still controlling the voltage very accurately at the load. This results in effective load isolation and very high power factor at the UPS input.

The basic modes of operation for delta-conversion UPSs are similar to those shown below for the double-conversion topology.

Overview of Double-Conversion or Online UPS Topology

Online or double-conversion UPSs are the most commonly used battery-based UPSs in data center environments because they are capable of completely isolating sensitive IT loads from unconditioned utility power. They receive their name, predictably, because they convert unconditioned utility power two times under normal operating conditions: first from AC to DC electricity and then back again from DC electricity into a highly conditioned AC signal. Double-conversion UPSs always provide the load with a high quality, conditioned AC signal, even during normal operation when utility power is available (Figure 7). For this reason, double-conversion UPSs are more common in high-availability, high-power mission critical applications such as industrial facilities and data centers. Typical output power ratings used in a data center setting for these types of UPSs range from 10 to over 1,000 kVA.

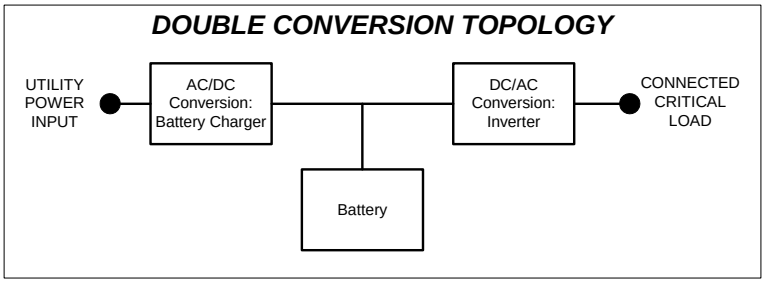

Figure 7

During normal mode, utility AC power is converted into DC power and is then converted back to a conditioned AC output for critical loads. During this phase of operation, the battery receives a continuous trickle charge to keep it at maximum charge. A flow chart of this mode of operation is illustrated in Figure 8 below.

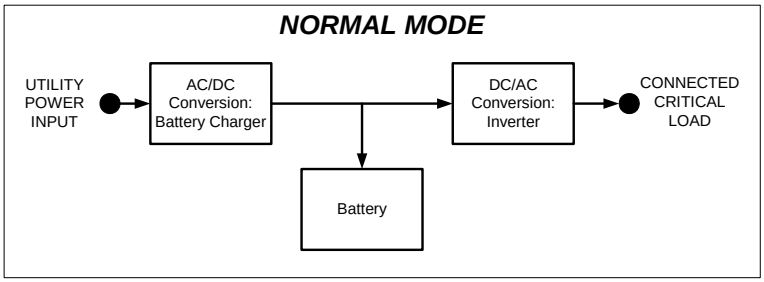

Figure 8

When utility power goes beyond acceptable limits, the double-conversion UPS enters stored energy mode, which is identical to stored energy mode in the line-interactive topology. The inverter draws DC power from the battery and continues to put out a conditioned AC signal to critical loads. In other words, the output of a double-conversion UPS is *always* conditioned. This often makes double-

conversion UPSs the topology of choice for highly sensitive loads in mission critical applications, although delta conversion line-interactive UPSs are making inroads.

Overview of Flywheel or Kinetic UPS Topology

Flywheel UPSs represent quite a departure from the three topologies described above due to the absence of a chemical battery for energy storage.[7] Flywheel UPSs, often referred to as "kinetic" UPSs, instead store backup energy in the motion of a rapidly spinning disc or flywheel. The flywheel itself must be made of a high-strength, dense material to resist disintegration at the high speeds at which it spins–typically in the thousands of rotations per minute. The wheel usually spins upon high-grade bearings in a vacuum chamber to eliminate friction as much as possible and preserve the amount of stored energy. In backup operation, the UPS converts the energy stored in the motion of the flywheel into electricity in a manner similar to a gas-powered generator, providing 10 to 30 seconds of conditioned AC electricity to critical loads. This is just enough "ride-through" time for backup power generators to come up to working condition and provide power to the critical load. A typical configuration of the flywheel or kinetic UPS topology is illustrated in Figure 9 below. Note how this topology is very similar to the double-conversion topology, except that the battery has been replaced with a flywheel and an AC/DC flywheel converter.

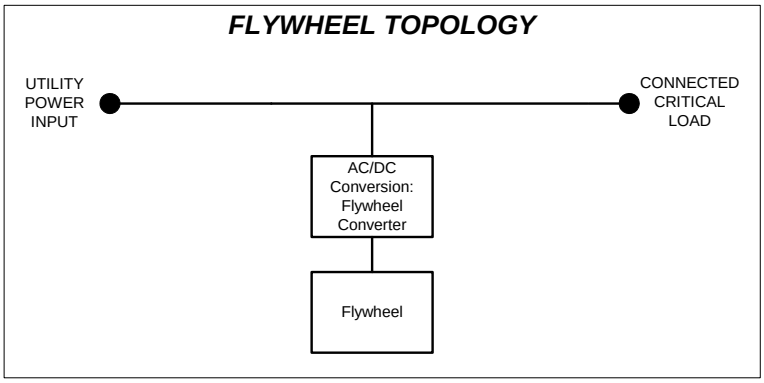

Figure 9

During normal mode, utility power is conditioned and delivered to the critical load through the rectifier and the inverter. The details of this conversion may vary from model to model, but the end goal is the same as with the double-conversion topology: keep the critical load isolated from unconditioned utility power at all times. In this mode when utility power is available, energy is initially consumed to bring the flywheel up to operating speeds. Thereafter, a properly designed flywheel system only requires a small amount of energy to keep the flywheel spinning at its target speed. This equivalent to charging the batteries in the other UPS topologies; however, well-designed flywheel systems tend to be more efficient than battery-based systems because, once brought up to running speed, they lose very little of their speed or "charge." Figure 10 illustrates the normal mode of operation in a flywheel UPS.

[7] For official EN/IEC definitions of the three common, battery-based UPS topologies previously mentioned, refer to: *IEC 62040-3, Testing Procedures for UPS Systems*. International Electrotechnical Committee. April 30, 2004.

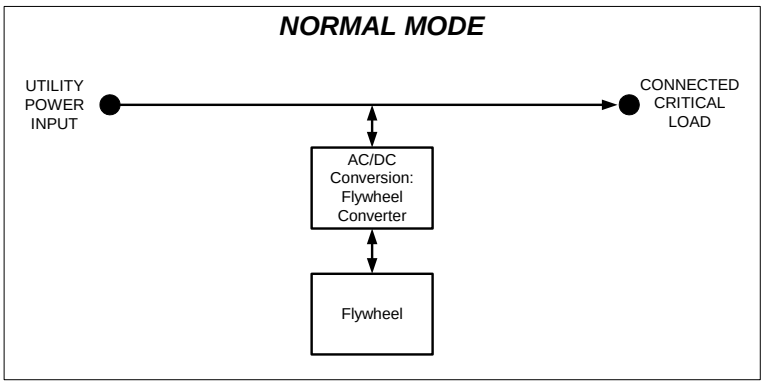

Figure 10

Stored energy mode in a flywheel UPS is similar to other UPS topologies except that it typically acts as a bridge between utility power and backup diesel-generated power. Power is provided through the kinetic energy stored in the UPS's flywheel only long enough to allow for generators to spin up, similar to some configurations of battery-based UPSs. Figure 11 illustrates backup mode in a flywheel UPS.

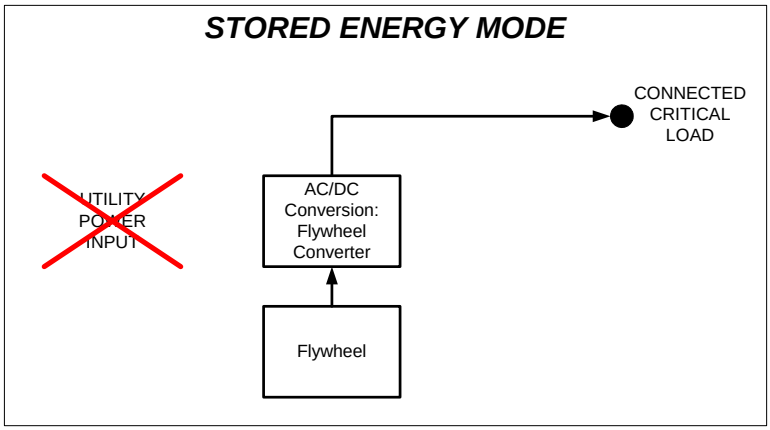

Figure 11

Overview of the U.S. UPS Market (Data Center/IT Sector)

We recognized early in the process of determining the size of the UPS market and the installed stock of UPSs in the data center/IT sector that new market research would be required to provide an accurate estimate. The tech boom of the late 1990s translated into rapid development of data center facilities around the country and along with it, increased purchases of UPSs to help protect mission-critical IT loads; however, this growth did not sustain itself. The "bursting" of the "tech bubble" in 2001 combined with market declines associated with the September 2001 terrorist attacks on the United States had a dramatic effect on the data center/IT sector for UPSs, which experienced two consecutive years of negative growth.[8] Many earlier estimates of UPS stocks based their calculations on the rapid tech boom growth curve of the late 90's and were unable to take into account the effects of the tech downturn in 2001. Figure 12 illustrates just how serious the downturn was for the UPS market in the U.S. Sales are not expected to recover to pre-2001 levels until after 2006.[9]

[8] Frost and Sullivan. *World UPS Markets*. 2004. p. 1-6

[9] Data source: Arthur D. Little for sales data from 1990 – 2000, Frost and Sullivan for sales data from 2001 – 2010.

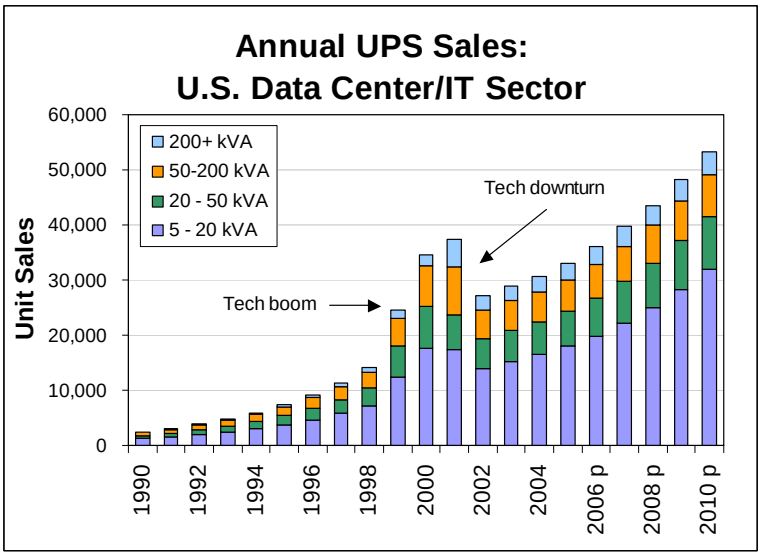

Figure 12

To provide a more accurate estimate of UPS stocks, it was necessary to take the tech downturn of 2001 into account by obtaining up-to-date UPS sales figures. We have chosen to utilize a Frost and Sullivan report published in 2004 entitled *World UPS Market* to gain access to more current market information. In general, the report describes a two-year period of negative growth in the global UPS market followed by a rebound in sales and renewed growth in 2003. The report cites a few trends that contributed to this growth, including an upswing in IT spending/development and lessons learned from the costly 2003 northeast blackout about the reliability of the aging U.S. power grid.

We were able to dissect the report to determine what percent of UPS sales occurred in the U.S. data center/IT sector and, from these sales figures, built a model to estimate the stock of installed data center/IT UPSs in the nation. The report provided detailed sales information for several different ranges of UPS sizes, measured in kVA, and broke down percent of sales based on geographic region (allowed us to characterize North American market) and end-user market (allowed us to characterize the percent of sales dedicated to data center/IT applications).

Several basic assumptions were required in order to estimate unit sales in the U.S. data center/IT sector. First, we assumed that the U.S. market comprises 90% of all North American UPS sales, in keeping with past estimates made for the U.S. Department of Energy by Arthur D. Little on the size of the U.S. UPS market.[10] Secondly, we excluded UPSs below 5 kVA in our analysis because these units are small enough that they would not likely be used to support data center equipment. Thirdly, we assume that the percent of sales by end-user market presented in the Frost and Sullivan report apply to all geographic regions. In other words, we assume that data center/IT spending is relatively similar in all regions of the world and particularly that the global data center/IT market share presented by Frost and Sullivan is representative of purchasing behavior in the United States. Finally, in order to arrive at an estimate of the stock of installed UPSs in the data center/IT sector, we assumed that the average lifetime for a UPS is 10 years. This number is based on estimates provided in both the Frost and Sullivan 2004 market report as well in the aforementioned Arthur D. Little report on office equipment energy use.

The sales data presented in the Frost and Sullivan report only begins in 2001, and because we have assumed an average UPS lifetime of 10 years, it was necessary to obtain earlier sales data provided in Arthur D. Little's 2002 analysis to be able to sum the cumulative sales over 10 years. Our stock estimate, thus, begins in the year 1997 and is projected out to the year 2010, even though our supplemental sales data begins in 1988.

[10] K. Roth, F. Goldstein and J. Kleinman. *Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings, Volume I: Energy Consumption Baseline.* Cambridge, MA: Arthur D. Little, 2002. p. 84

We estimate that there are about 225,000 UPSs in operation in the U.S. data center/IT sector today, about 15% of which reside in the State of California. Half of the national UPS stock are units in the 5.1 to 20 kVA range of apparent power output. We assume that many of these units would not be used to power entire data center facilities but rather could be used to support individual server racks and the like. UPSs in the last two categories – 50.1 to 200 kVA and 200+ kVA – would likely be used to power entire data center facilities. Table 1 provides an approximate outline of the current distribution of UPSs grouped by their rated output power along with typical price ranges for each category.[11] Although we expect the stock numbers of UPSs to continue to rise, we do not expect the percent distribution of these categories of UPSs to change very much in the near future. Figure 13 plots the expected growth in the stock from the years 1997 through 2010. The dashed line indicates the projected growth in stock that would have occurred had the tech boom of the late 90's continued past the turn of the century.

Table 1: Summary of UPS Stock, U.S. Data Center/IT Sector, 2004

	Stock of Installed Units	% of Stock	Typical Price Range (USD)
5.1 to 20 kVA	114,431	51%	$4,000 - $16,500
20.1 to 50 kVA	46,159	20%	$8,000 - $37,000
50.1 to 200 kVA	46,639	20%	$19,000 - $98,000
200 + kVA	19,080	9%	$30,000 - $207,000

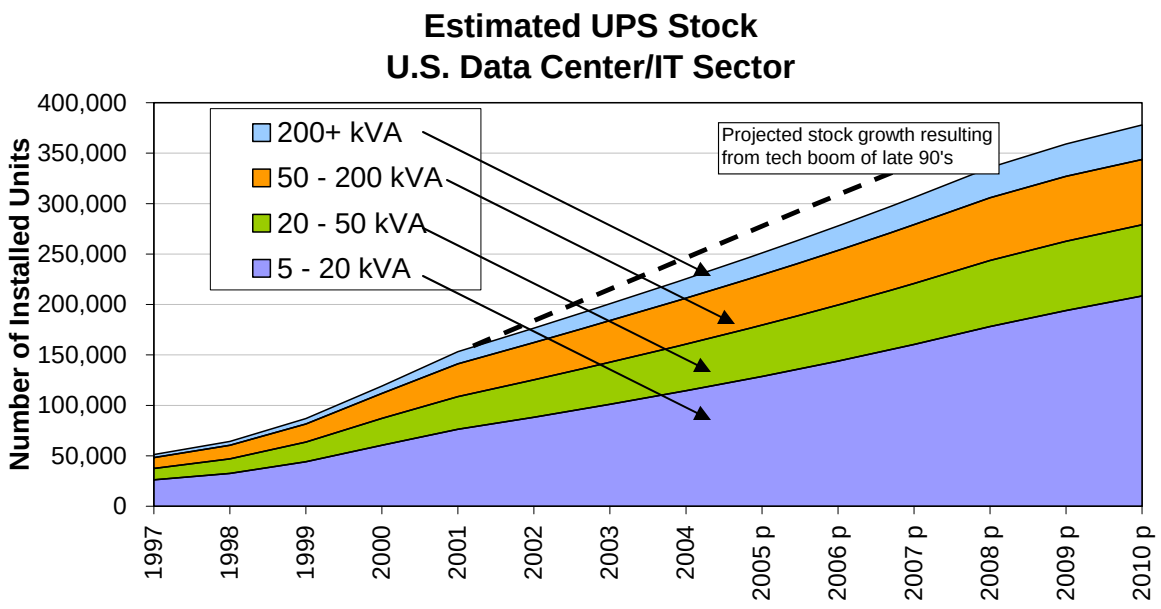

Figure 13

Using the estimated stock data presented in Figure 13 and a few assumptions about the composition of the UPS market, we have further estimated the basic composition of the UPS stock by design topology. Based on information from the 2004 Frost and Sullivan market report, the 2002 Arthur D. Little report on office equipment energy use, and common practices in data centers, we assume that

[11] Source of prices: Frost and Sullivan.

there are essentially no standby UPSs in use in the data center/IT sector. This topology simply does not provide the type of availability that the market requires for protecting mission-critical loads.

Because flywheel UPSs are still in their infancy in data center/IT applications, they constitute less than 1% of of UPSs in operation for units greater than 50 kVA. It should be noted that we are aware of flywheel UPS systems from several manufacturers that are in use in data centers and industrial facilities around the country and that this topology is likely to continue making inroads in mission-critical facility applications as the technology matures and prices come down. The topology has special appeal in sizes above 200 kVA because it is battery-free and, thus, can avoid the cost and downtime associated with servicing large banks of lead-acid batteries.[12]

Finally, it was necessary to make some assumptions about the split in sales between the double-conversion/online and offline/line-interactive topologies. For UPSs in the 5.1 to 20 kVA category, we assume that 81.3% of the stock is double-conversion and 18.7% is line-interactive, based on past sales data.[13] For UPS sizes above 20 kVA, some past literature has assumed that all sales are in double-conversion topologies; however, APC has been making inroads with its delta-conversion in this power range, and our analysis acknowledges this fact by breaking out delta-conversion UPSs as part of the UPS stock above 20 kVA. Since APC is the only manufacturer producing delta-conversion UPSs in sizes above 20 kVA (the company manufactures delta-conversion units up to 1000 kVA in size) and since the company sells exclusively delta-conversion units at these sizes, we assume that the percent of delta-conversion UPSs in the stock matches APC's market share for the 20.1 – 50 kVA, 50.1 – 200 kVA, and 200+ kVA market segments.[14]

Table 2 outlines our assumptions about the relative distributions of various UPS topologies in the U.S. data center/IT sector. The estimates are presented in terms of percent of stock and number of installed units for several size categories. Figure 14 graphically illustrates how the stock breaks down by topology. For UPSs larger than 5 kVA used in data centers, the market is essentially dominated by one UPS topology, the double-conversion/online topology. The data labels in the two right-hand columns indicate the estimated number of delta-conversion units installed in the 50 to 200 kVA and 200+ kVA ranges.

Table 2: Composition of U.S. Data Center/IT Sector UPS Stock by Topology

	Line-interactive (% stock) (# units)	Delta-conversion (% stock) (# units)	Double-conversion (% stock) (# units)	Flywheel (% stock) (# units)	TOTAL UPSs in Stock (# units)
5.1 - 20 kVA	18.7%	0%	81.3%	0%	114,431
	21,399	0	93,033	0	
20.1 - 50 kVA	0%	11.1%	88.9%	0%	46,159
	0	5,124	41,035	0	
50.1 – 200 kVA	0%	2.0%	97.5%	< 0.5%	45,639
	0	913	44,498	< 228	
200+ kVA	0%	0.6%	99.0%	< 0.5%	19,080
	0	114	18,889	< 95	

[12] Frost and Sullivan. p. 8-5

[13] Taylor J. and Hutchinson J. "Uninterruptible Power Supplies, Parts 1, 2, and 3." *Electronic Buyers News Power Supplement.* September 2000.

[14] Frost and Sullivan. pp. 6-19, 7-18, and 8-20.

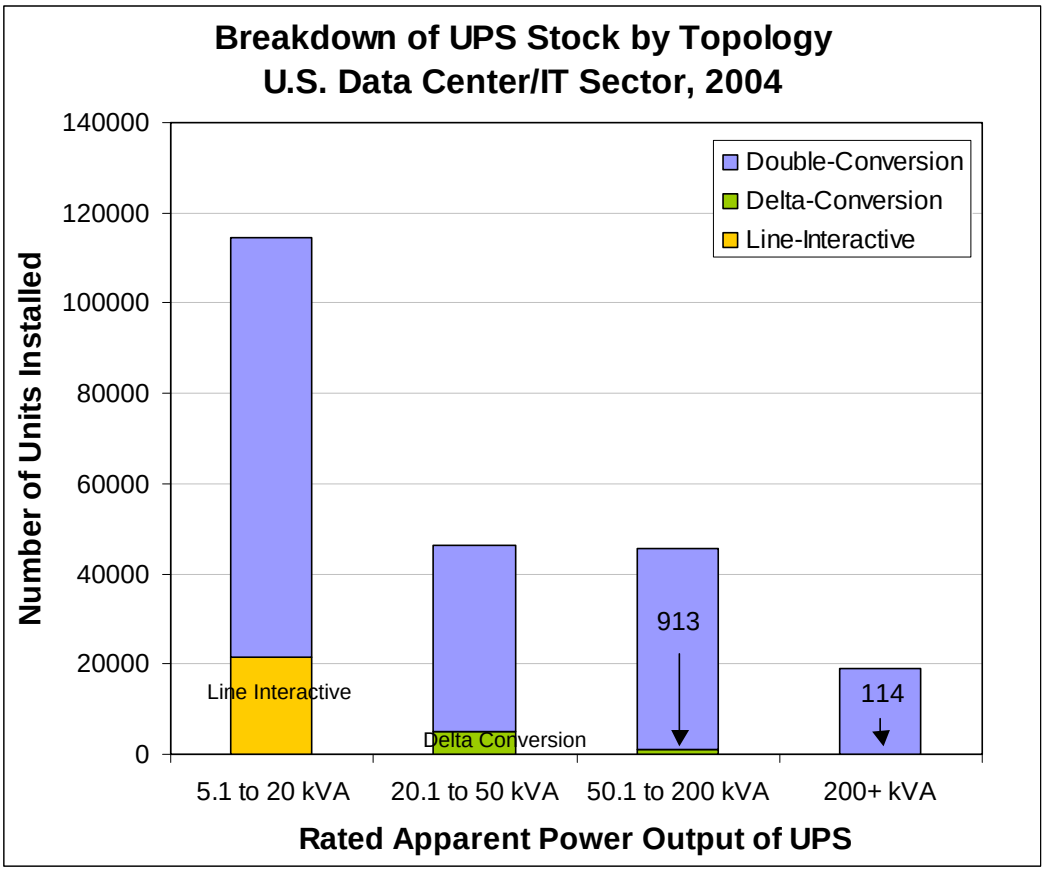

Figure 14

UPS Test Protocol

One of our most important tasks was to develop clearly written UPS efficiency test protocols that were to be used either in a factory/laboratory setting or a field setting to gather credible, comparable data on UPS efficiency, tested at resistive and non-linear loads. The International Electrotechnical Committee (IEC) has created a detailed and comprehensive international standard (IEC 62040-3) to guide manufacturers in testing their UPS systems,[15] which devotes 2 pages to UPS efficiency testing. The standard discusses the definition of efficiency, discusses approved linear and non-linear loads, and provides a suggested format for recording the data, but unfortunately only recommends testing for efficiency at 50%, 75% and 100% of the UPS's rated active and apparent power outputs. Past research indicates that UPSs in a data center environment typically operate at or below 50% of their rated active power output, where the efficiency of the system can decline significantly, and therefore it was of particular interest to our team to develop a UPS test method that was more specific about measuring efficiency at part loads.

Factory/Laboratory UPS Efficiency Test Method

In order to gather data to benchmark data center UPS efficiency, encourage manufacturers to measure UPS efficiency at part loads, and to raise awareness of the importance of properly characterizing UPS efficiency, Ecos Consulting and EPRI Solutions developed a test procedure that exclusively guides manufacturer testing of UPSs that are designed for data center and mission-critical applications, entitled *Guidelines for Manufacturers' Testing to Benchmark Data Center UPS*

[15] *IEC 62040-3, Testing Procedures for UPS Systems*. International Electrotechnical Committee. April 30, 2004.

Efficiency.[16] The test method was circulated widely to UPS manufacturers in mid-2004 to spur on an effort to obtain more accurate UPS efficiency data. Five of these manufacturers have agreed to participate in our project by sharing the results of UPS efficiency tested guided by our test procedure. These results are described in Section 1.3 below.

The test method is largely based on the IEC 62040-3 standard and references the document regarding the appropriate load characteristics for resistive and non-linear loads, the definition of UPS efficiency, and the types of UPS testing typically performed by manufacturers. The EPRI Solutions/Ecos test method expands upon IEC 62040-3 by recommending efficiency testing at no load (stand-by), 10%, and 20% load conditions, which has allowed us to better characterize the shape of UPS efficiency curves at part loads. The manufacturer test method also provides several pages of useful tables and forms to aid manufacturers in reporting detailed efficiency and power quality information about their models.

Field UPS Efficiency Test Method

In the absence of any formal test procedure for measuring the power consumption and efficiency of field-installed UPSs, EPRI Solutions and Ecos Consulting developed a field test method as a compliment to our manufacturer guidelines, entitled *Guidelines for Field Data Collection to Benchmark Data Center UPS Efficiency.*[17] Because the design implementation of UPSs in the field can vary and because there are a number of different ways that power/efficiency data could be obtained from an installed unit, the method needed to provide explicit guidance for many different types of configurations. The document provides guidelines for measuring field UPS efficiency through a number of means, including:

- Visual measurements from a front-panel display
- Data logging using UPS monitoring software
- Data logging using permanently connected power meters
- Data logging using a building energy management system
- Direct measurement with power meters

This field test method has been used as a guideline for testing UPS loading and efficiency at several data centers and other mission-critical facilities in the U.S. The results of this testing are presented in Section 1.3 below.

[16] *Guidelines for Manufacturers' Testing to Benchmark Data Center UPS Efficiency*. June 7, 2004. <http://hightech.lbl.gov/Documents/UPS/UPS_Efficiency_Measurement.pdf>

[17] *Guidelines for Field Data Collection to Benchmark Data Center UPS Efficiency.* June 7, 2004. <http://hightech.lbl.gov/Documents/UPS/UPS_Field_Data.pdf>

3. UPS Efficiency Data

UPS Efficiencies

EPRI Solutions, Ecos Consulting, and Lawrence-Berkeley National Laboratories have gathered measurements on the efficiency of a variety of commercially available UPSs under controlled, factory conditions as well as field conditions using the test methods described above in Section 1.2.

The efficiency of a UPS, as defined by the International Electrotechnical Committee, is "the ratio of (active) output power to (active) input power under defined operating conditions,"[18] where defined operating conditions refer to a specific percent load and load type (linear/resistive versus non-linear). We are primarily concerned with UPS efficiency while the unit is in its normal mode of operation, when utility power is within accepted tolerances and the UPS can use this power to feed the critical load. UPSs spend the vast majority of their time in this mode of operation, so this is where efficiency can have the most impact on energy use and total cost of ownership.

The efficiency of a UPS can be impacted by a number of factors that can be controlled during testing, including the percent of load being supplied by the unit and the power quality characteristics of the load (resistive vs. non-linear loads). We describe how these factors can affect UPS efficiency below, followed by an analysis of the data that manufacturers have shared with our team.

UPS Efficiency and Percent Load

UPS efficiency varies with the amount of active power being supplied to a load and tends to decrease steadily at part loads. UPS efficiency curves are usually relatively flat between loads of 50% and 100%, but decrease steadily below 50% and precipitously below 20%. The efficiency of a UPS can vary by 7 to 12 percentage points depending on whether it is partly or fully loaded. The lower efficiency at part loads can result significant conversion losses compared to the same UPS operating at close to full load. Figure 15 illustrates a typical load curve of a UPS in our current data set. Note that there is no "typical" operating range for any UPS.

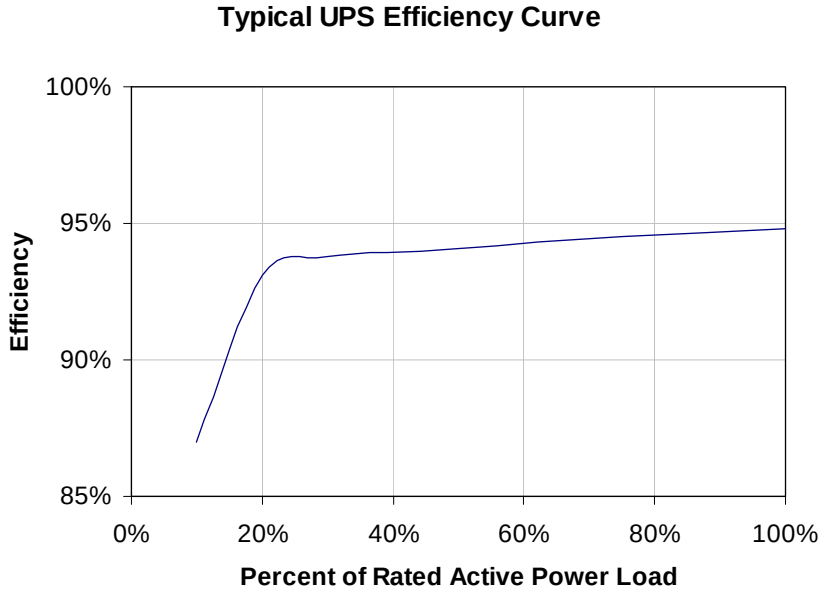

Figure 15

[18] *IEC 62040-3…* p. 52

UPSs in mission critical environments typically operate at points between 30% and 50% of nominal and in situations where the data center is not fully populated this can be even lower resulting in very large inefficiencies. Therefore knowing the efficiency of a UPS at loads below 50% is extremely important in estimating real world energy use. As the figure above shows, efficiency tends to decline steadily below about 50% load, so understanding the behavior of the efficiency curve at these low loads is crucial to making accurate estimates of energy use and operational costs. Unfortunately, manufacturers usually only report UPS efficiency at 100% load, with occasional reporting at 50% and 75% load. This is one of the main reasons why EPRI Solutions' and Ecos Consulting's attempt to characterize UPS efficiency across a wide range of loading was so important.

There are a number of reasons why data center UPSs operate at part loads, one of which is that UPSs are commonly sized to meet the maximum utilization of space in a data center, even though maximum utilization rarely occurs.[19] This results in oversizing of the UPS for the actual amount of installed mission-critical infrastructure.

UPSs in mission-critical facilities operate in redundant configurations, which also act to reduce the theoretical percent loading of the unit to below 50%. In a redundant UPS configuration, the load of the facility is shared between two or more UPSs so that, if one of the units fails, the other(s) will still be available to supply the facility's critical load. Any one UPS in a redundant configuration must be capable of supplying the *entire load* required by the facility. Because the load is shared equally between two or more of these full-sized UPSs, the largest load that a data center UPS could theoretically experience under most conditions is 50% of its rated active power load. A more thorough discussion of redundant UPS configurations and their effects on overall efficiency is provided in Section 1.4.

UPS Efficiency and Load Characteristics

The characteristics of a UPS load can have a noticeable effect on the measured efficiency of the unit. Manufacturers often report UPS efficiency tested with resistive or linear loads; however, UPSs in many mission critical installations usually have to power highly non-linear loads, like the switch mode power supplies (SMPS) used in desktop and low-end servers. These power supplies often have poor power quality characteristics, including lower power factor (0.7) and high total harmonic distortion (in excess of 100%). Low power factor loads like a SMPS cause the UPS to deliver more peak current, changing the way in which the system is loaded. A non-linear load like a SMPS can also create harmonic currents in distribution wiring, which can dramatically increase the amount of current in distribution wiring if the UPS does not include special filters.

These different characteristics of non-linear devices change the way the UPS is loaded and can have an impact on efficiency. Figure 15 shows how the measured efficiency of a given UPS in our data set differed by 1% to 2% between a linear and non-linear load. The additional current demands of the non-linear load drives the efficiency down by a small but measurable amount. This is important because, as mentioned earlier, UPSs in mission critical environments very often support non-linear loads. Testing with non-linear loads will be most indicative of operational efficiency.

[19] Madsen J. "Continuous UPS Availability: How Important is it to Your Company?" *Energy User News.* August 11, 2000. < http://www.energyusernews.com/CDA/ArticleInformation/features/BNP__Features__Item/0,2584,14489,00.html>

Figure 15

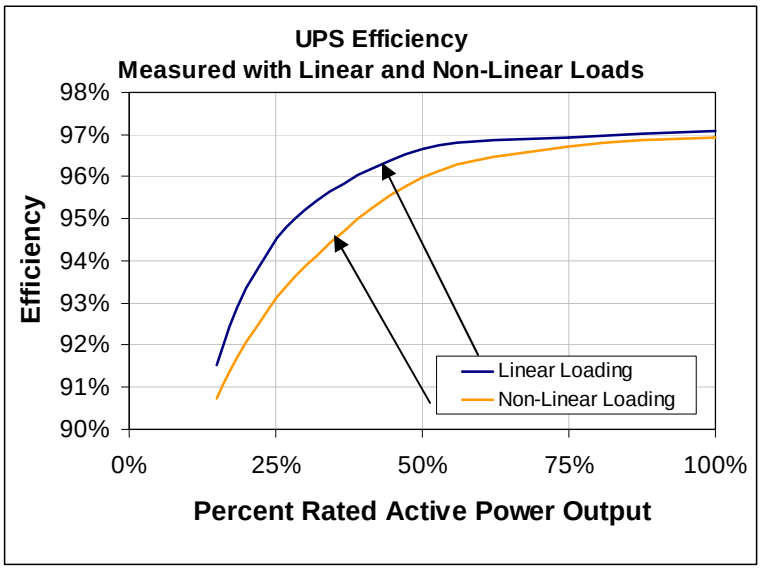

Results of UPS Efficiency Testing: Field Measurements

Results of recent data center UPS efficiency testing support the claim that part load efficiency – efficiency measured at less than 50% of a UPS's rated active power load – is of greatest importance in estimating the operational costs and total cost of ownership of a UPS. EPRI Solutions, Ecos Consulting, and LBNL benchmarked the efficiency of UPSs operating under real world conditions in data centers. Ecos and EPRI Solutions built upon research conducted in 2001 in which LBNL benchmarked the performance of 14 data center facilities along with their UPS equipment. We added the UPSs from two new UPSs, both supporting IT loads in the EPRI Solutions data center.

Over 80% of the UPS systems measured were operating at below 50% of their rated load. Many of these units were partially loaded because they were sharing their load in a redundancy configuration, which we denote in the chart; however, four of the partially loaded units were not operating redundantly. They were simply oversized and underutilized. Figure 16 displays the results of our testing, plotting both the percent of the UPS's nominal load utilized in the field as well as efficiency of the UPS under real world conditions. Table 3 summarizes the average load factor and operational efficiency for UPSs tested in the field.

Table 3: Average Loading and Efficiency of UPSs in the Field

Average Load Factor of UPSs Tested in the Field	Average Efficiency of UPSs Tested in the Field
37.8%	85.2%

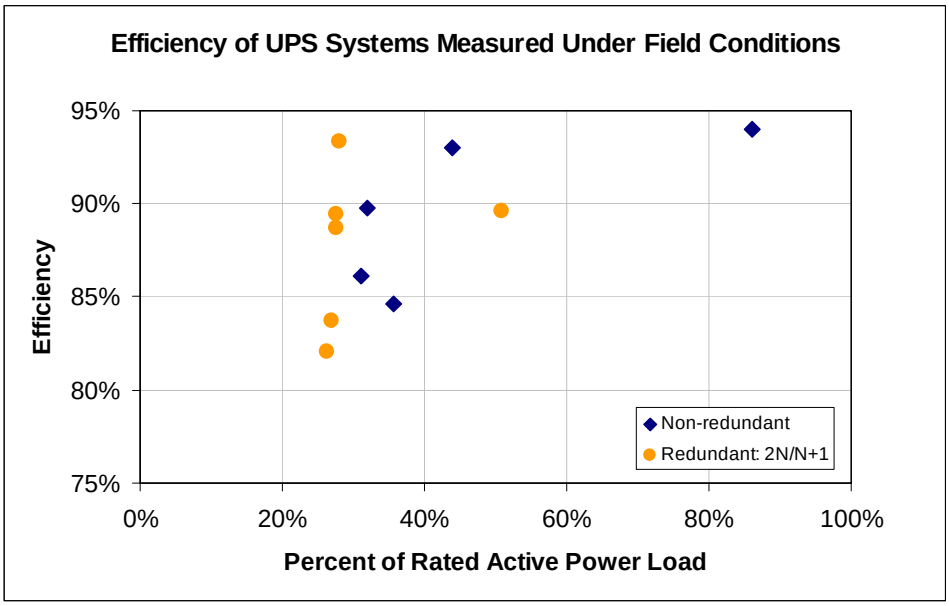

Figure 16

Results of UPS Efficiency Testing: Factory/Laboratory Measurements

Our results show that UPS efficiency varies significantly between different models. At low load conditions (20% of rated active power load), our findings indicate that UPSs are anywhere from 81% to 93% efficient. At full load conditions, our measurements show that UPSs are anywhere from 86% to almost 98% efficient. Although these figures span a range of less than 15 percentage points, upgrading from one of the lower efficiency UPSs to a model at the high end of the range could cut UPS operating costs by over a half. The range in UPS efficiency also suggests that there may be an opportunity to encourage higher efficiency designs through product labeling and standards, a technique which has been applied in the past to other high-power devices but which has not yet been adopted for UPSs.

Figure 17 summarizes UPS efficiency test results that have been shared with EPRI Solutions and Ecos Consulting by manufacturers. The results presented in this chart were conducted under factory conditions with linear loads and were guided by the manufacturer test guidelines mentioned earlier in the report. A second chart, Figure 18, presents the results of testing conducted using non-linear loads.[20] The efficiency curves are identified by their general topology, including flywheel, double-conversion, and delta-conversion.[21]

[20] Note: some manufacturers reported efficiency measurements conducted under both linear and non-linear loads; however, a large number of the UPSs were only tested with resistive loads, hence the larger number of lines in Figure 17.

[21] Note: due to the scope of this project, only UPSs common in data center applications were tested. Standby or offline UPSs are uncommon in high-power, mission critical applications and, thus, were not tested.

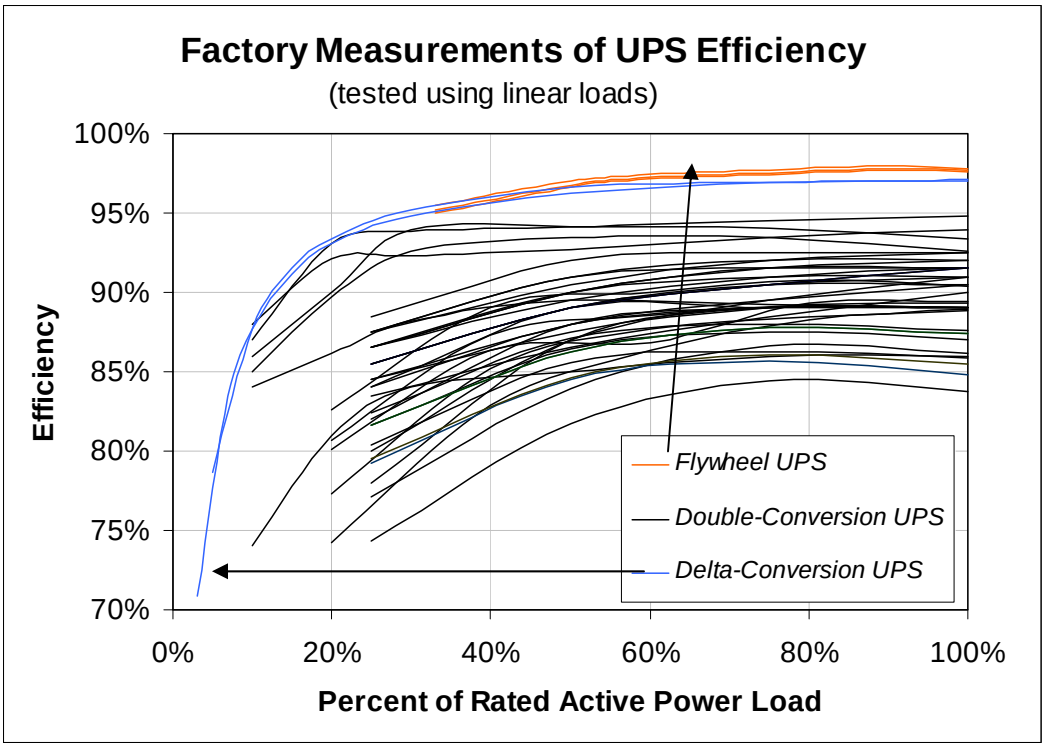

Figure 17

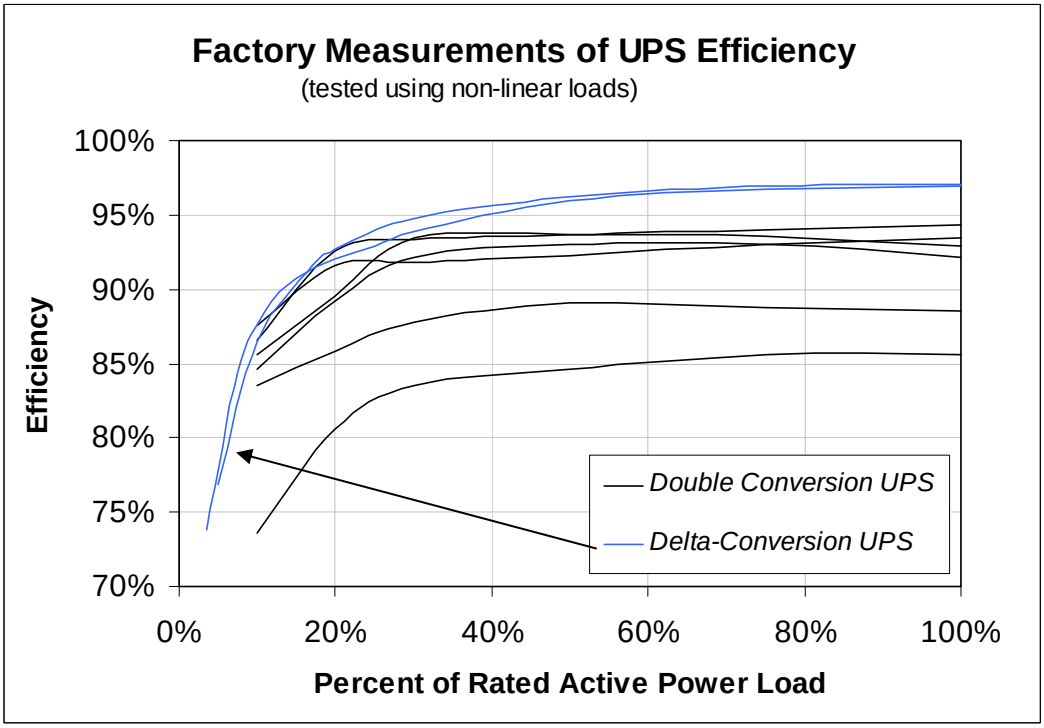

Figure 18

Our results indicate that there can be fairly wide variation in efficiency even between UPSs of the same general topology design, such as the double-conversion UPSs plotted in Figure 18 above, which vary from 86% to almost 95% efficiency at full load. The table below summarizes the characteristic efficiency of a number of UPS topologies at various percent load conditions and shows the average efficiency for all of the UPSs measured.

Table 4: Characteristic Efficiency of UPS Topologies

UPS Topology	Efficiency at 25% Load	Efficiency at 50% Load	Efficiency at 75% Load	Efficiency at 100% Load
Delta-Conversion	93% - 94%	96% - 97%	97%	97%
Double-conversion	81% - 93%	85% - 94%	86% - 95%	86% - 95%
Flywheel	unknown[22]	97%	98%	98%
AVERAGE OF ALL UNITS	86%	89%	90%	90%

Due to the scope of this research effort, it was impractical to determine what specific elements of a model's design contributed to its measured efficiency; however, our results indicate that some broader design decisions, such as the general UPS topology, can be indicative of efficiency. In particular, flywheel and delta-conversion UPSs exhibited best-in-class efficiency compared to the double-conversion UPSs that we measured. The three flywheel UPS designs and two delta-conversion designs had consistently higher efficiency than the double-conversion models for loads above 30% of nominal and similar performance in the 10% to 20% range compared to the highest efficiency double-conversion models. *We stress that no strong conclusions should be drawn about the inherent efficiency of either the flywheel or delta-conversion topologies because of the relatively small sample size in our study.* Research on the efficiency of both flywheel and delta-conversion UPSs should continue as manufacturers continue to improve their designs and diversify product lines. Nevertheless, the high efficiency of these UPS toplogies in our current data set suggest that they deserve serious consideration by facility managers looking to lower operating costs in mission critical facilities that utilize large UPSs (greater than 50 kVA in capacity) and policy makers wishing to identify the most efficient technologies.

Desire for high efficiency should naturally be balanced with concerns over load isolation and reliability. Although the delta-conversion UPSs that we tested performed better in regards to overall efficiency compared to double-conversion units, some manufacturers have argued that delta-conversion UPSs do not provide the same load isolation as a "true" double-conversion UPS. This report does not attempt to answer the question as to which topology is the overall best choice for mission critical applications. We simply report the observed efficiencies of the various topologies, and facility managers will ultimately need to judge whether a particular UPS can handle load disturbances and provide an acceptable level of load isolation.

Investigation of Special High Efficiency Modes in Select UPSs

A handful of UPS manufacturers now advertise double-conversion units with programmable high efficiency modes, sometimes referred to as "power-saver" or "eco" modes. Ecos Consulting identified the following units in Table 5 as having some means of enabling higher efficiency through a user control. Although most of the units occupy the sub-50 kVA range of output and would not be capable of powering entire data centers, the Chloride 90-NET models are available in sizes ranging up through 800 kVA, proving that such technology could already be applied to entire facilities.

[22] Data on flywheel efficiency was only presented down to about 30% load.

Table 5: Summary of UPS Models with High Efficiency Modes

UPS Manufacturer	Model	Power Range (kVA)	Manf.-Reported Double-Conversion Efficiency	Manf.-Reported Energy-Saver Mode Efficiency
MFR#1	9170+ Online	3 – 18	88%	97%
MFR#2	90-NET	60 – 800	92% - 93%	97%
MFR#2	70-NET	10 - 60	91% - 92%	97%
MFR#3	Galaxy 3000	10 – 30	85% - 86%	unknown

EPRI Solutions was able to test a 3 kVA model of the Powerware 9170+ Online UPS. The unit is, by design, a double-conversion UPS but offers a "power-saver" high efficiency mode that the manufacturer claims can boost efficiency from the 88% range to 97%. EPRI Solutions conducted controlled laboratory efficiency measurements on the 9170+ to validate the manufacturer's claims. Although the high efficiency mode only allows for the reported 97% efficiency at full load, our tests found that the feature did enable 5% efficiency improvements across the range of percent loads, as illustrated in Figures 19 (resistive load testing) and 20 (non-linear load testing).

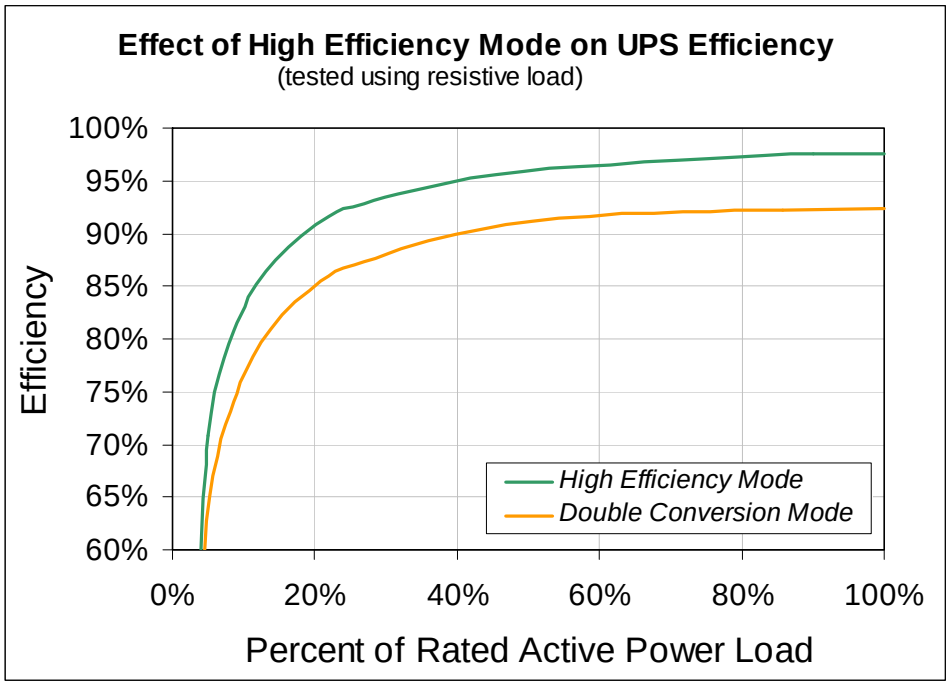

Figure 19

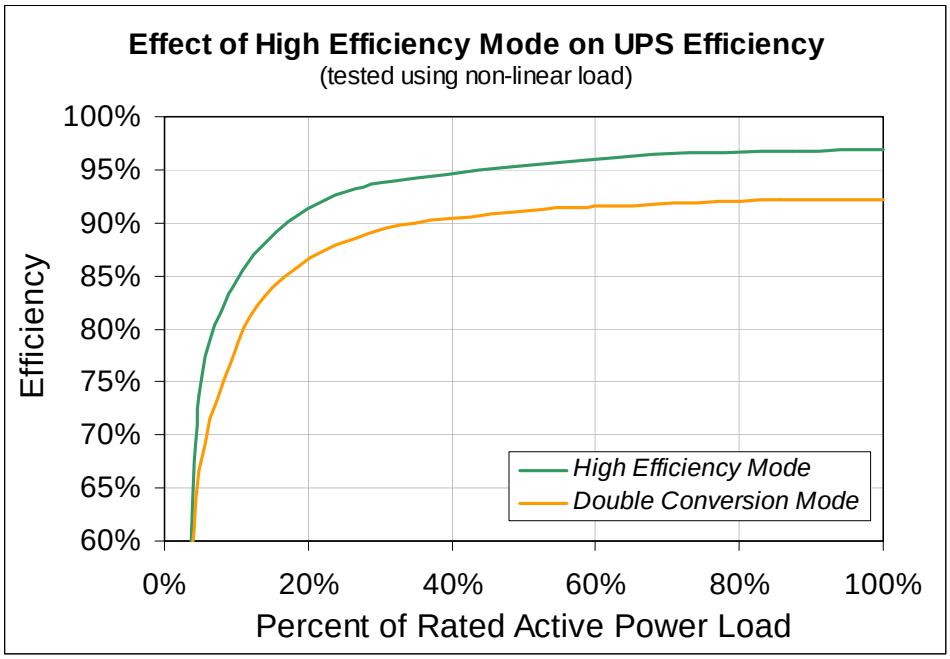

Figure 20

One of the potential drawbacks of enabling high efficiency modes in double-conversion UPSs is that the UPS no longer operates in double-conversion mode, potentially leaving mission critical loads more exposed to power abnormalities like voltage sags and line noise. It would seem that there is an inherent trade-off between high efficiency and complete load isolation, but is this trade-off significant enough that data center managers should choose not to enable high efficiency modes for fear of exposing their IT equipment to more power irregularities?

EPRI Solutions investigated these concerns by conducting additional tests on the Powerware 9170+ Online UPS, exposing the unit to power abnormalities in both double-conversion and high efficiency modes. The tests sought to determine the response of the UPS output when the UPS input (i.e. utility power) experienced an 80% and 30% sag in voltage for various durations of time. In particular, these experiments were designed to answer questions such as:

- Is the output of the UPS seamless regardless of voltage sags at the input?
- Does the UPS transfer over to battery when it sees power abnormalities at the input?
- If the output of the UPS is interrupted, how long does the interruption last, and would this interruption pose a threat to mission-critical infrastructure like servers?

In general, the performance of the 9170+ in its normal double-conversion mode – what Powerware calls “Auto” mode – was flawless for both 80% and 30% voltage sags. The AC output maintained a constant frequency and voltage regardless of disturbances to the input whether the UPS transferred over to battery backup power. Figure 21 illustrates one case in which the UPS responds to a 1-cycle (16.7 ms) 30% voltage sag. The unit transferred to battery backup power in this case, but maintained a seamless AC output. The behavior shown in Figure 22 is generally indicative of the unit’s performance in 80% voltage sag tests as well, with the main difference being that, in the 80% sag tests, the UPS maintained a seamless AC output without transferring to battery backup power.

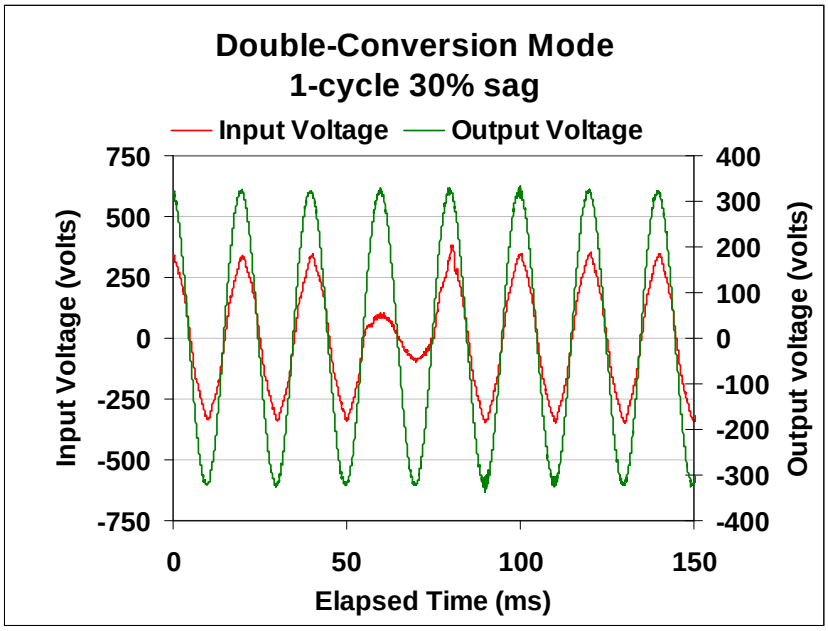

Figure 21

Unfortunately, the 9170+ does not appear to have the same "flawless" response to power abnormalities when its high efficiency mode is enabled. Take the case of the 1-cycle 30% voltage sag test. In double-conversion or "Auto" mode, the 9170+ was able to ride out this disturbance with no noticeable change to its output. As illustrated in Figure 22, the unit handles the disturbance much differently when high efficiency is enabled by the user. The output voltage first undergoes a 30% drop for one half of a cycle, followed by a 3/4-cycle disturbance while the UPS transfers over to its batteries for power. In total, this adds up to 20.9 ms of unconditioned power at the output of the UPS for a voltage disturbance that only lasts one AC cycle or 16.7 ms. EPRI Solutions experienced the same type of behavior when conducting 80% voltage sag tests on the 9170+, except that the duration of the voltage sag on the UPS output was longer (2.5 AC cycles) before the unit made the transition.

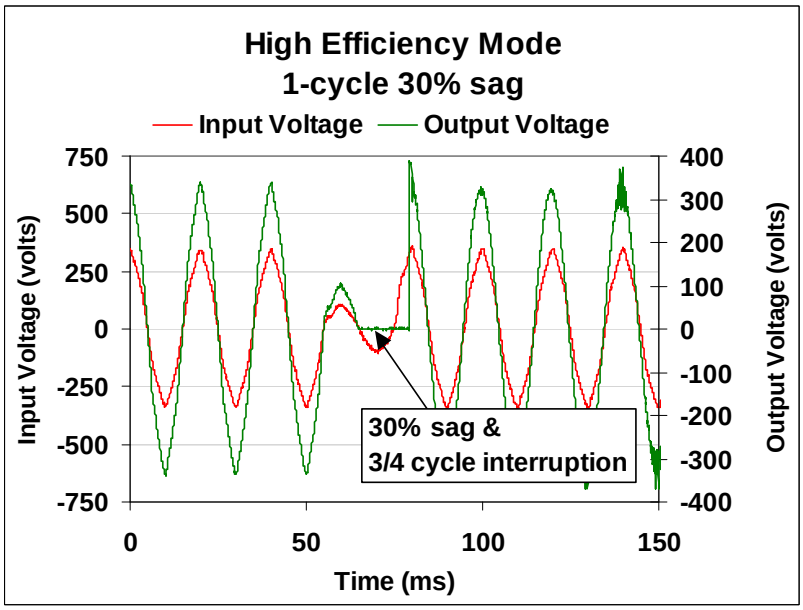

Figure 22

The tests not only indicated that the high efficiency mode would introduce delays in transferring to backup power, exposing connected loads to unconditioned power for brief periods of time, but they also suggested that the unit would have to transfer to backup power more frequently because it was unable to "ride out" small voltage disturbances. Recall that the 9170+, operating in double-conversion mode, was able to maintain a seamless output during the 80% voltage sag tests for disturbances up to 30 AC cycles or half a second in duration without transferring to battery power. When the same tests were performed with the unit set to its high efficiency mode, the unit switched to backup power for disturbances of 10 AC cycles or longer.

Tables 6 and 7 summarize all of the voltage sag testing performed on the 9170+ Online UPS operated in double-conversion and high efficiency modes, respectively. We have included a column that indicates whether a server power supply designed to Server System Infrastructure (SSI) specifications would be able to ride out the disturbance to the UPS output.[23]

Table 6: Summary of Voltage Sag Testing, Double-Conversion Mode

% Input Voltage Sag	Sag Duration (AC cycles & ms)	UPS Action	Length of Disturbance to Output (AC cycles & ms)	Disturbance Can be Handled by SSI PS
80%	1 cycle, 16.7 ms	No battery transfer	0 cycles, 0 ms	Yes
	10 cycles, 166.7 ms	No battery transfer	0 cycles, 0 ms	Yes
	30 cycles, 500 ms	No battery transfer	0 cycles, 0 ms	Yes
30%	1 cycle, 16.7 ms	Battery transfer	0 cycles, 0 ms	Yes
	10 cycles, 166.7 ms	Battery transfer	0 cycles, 0 ms	Yes
	30 cycles, 500 ms	Battery transfer	0 cycles, 0 ms	Yes

Table 7: Summary of Voltage Sag Testing, High Efficiency Mode

% Input Voltage Sag	Sag Duration (AC cycles & ms)	UPS Action	Length of Disturbance to Output (AC cycles & ms)	Disturbance Can be Handled by SSI PS
80%	1 cycle, 16.7 ms	No battery transfer	1 cycle, 16.7 ms	Yes
	10 cycles, 166.7 ms	Battery transfer	3.25 cycles, 54.2 ms	No
	30 cycles, 500 ms	Battery transfer	3.25 cycles, 54.2 ms	No
30%	1 cycle, 16.7 ms	Battery transfer	1.25 cycles, 20.9 ms	Yes
	10 cycles, 166.7 ms	Battery transfer	1.25 cycles, 20.9 ms	Yes
	30 cycles, 500 ms	Battery transfer	1.25 cycles, 20.9 ms	Yes

In general, the test results have lead EPRI Solutions and Ecos Consulting to the following conclusions about the 9170+ Online UPS. We stress that these conclusions should not necessarily apply to all of the UPS models that we have identified above that have a user-enabled high efficiency mode:

- The unit will tend to transfer to battery backup power more frequently when the high efficiency mode is enabled.
- Because the output is not fully conditioned, small or momentary abnormalities in power may pass through to the critical load when the unit is operating in high efficiency mode.

[23] According to SSI, power supplies for servers must be able to ride out voltage sags less than 1 cycle in duration.

- The unit currently shows a fairly consistent 3/4 –cycle delay in transferring to battery backup power for large input power disturbances, combined with 1/2- to 2.5-cycle disturbances to the unit's output before it recognizes the need to transfer to backup power.
- Although many of the disturbances to the output power that we measured while the UPS was in high efficiency mode lasted for only about one cycle, we have to conclude that the power disturbances are significant enough that even a properly designed, SSI-compliant server power supply would not be able to ride them through resulting in power failure to connected servers.
- It is unlikely that data center managers would be willing to operate a double-conversion UPS in high efficiency mode. This mode effectively turns a fully isolated double-conversion UPS into a less isolated line-interactive UPS. If customers are paying for the full performance of a double-conversion unit, they are likely to want all of the benefits that go along with it, including complete load isolation even though there is an energy penalty.[24]

Redundancy Evaluation

Today's facility managers expect previously unheard-of levels of availability from their UPS systems. Due to the expense of down time, world class facilities strive for UPS systems with six nines of availability (99.9999% available systems), meaning that mission-critical loads should have an available supply of power 99.9999% of the time. To put this in perspective, a UPS system with six nines of availability can only have 32 seconds of down time per year. Unfortunately, a single UPS system is typically capable of 99.92% availability, leaving a data center at risk of about 7 hours of down time every year.[25]

Redundant UPS configurations help to secure higher levels of availability by preventing a facility's power system from being interrupted by a single UPS failure. This is typically done by sharing the facility's total load between two or three UPSs so that, even if one or two units simultaneously fail, there will still be another backup UPS to condition utility power and supply backup power in the case of power outages. Each individual UPS in a redundant configuration usually must be capable of supplying the entire critical load by itself to achieve the highest levels of availability. Because of redundancy configurations, a data center's load is commonly shared equally between two or more UPSs that, individually, could power the entire facility by themselves. This load sharing reduces the percent load in individual UPSs and has the effect of decreasing their operational efficiency.

Below we provide an overview of some common types of UPS configurations along with an indication as to the level of availability that they provide and how the overall configuration affects efficiency.

Capacity or "N" Configuration

In the simplest UPS configuration, a single UPS supplies power to the entire data center facility. "N" refers to the predicted size of the load protected by the UPS. Since the configuration only requires one UPS capable of supplying this load, the name for the configuration is simply "N". This configuration is highly uncommon in mission critical facilities like data centers because there are no backup UPSs available to protect the load in case the main UPS fails or requires maintenance. As mentioned above, the estimated availability of a single UPS is 99.92%, meaning that the critical load will be unprotected from power abnormalities and failures for about 7 hours every year.[26]

The main advantage of the "N" configuration is that it has no effect on the percent load of the UPS and, thus, no effect on the operational efficiency.

---

[24] Personal communication with Tom Geist of EPRI Solutions. May 10, 2005.

[25] K. McCarthy. "Comparing UPS System Design Configurations." American Power Conversion. 2004. < ftp://www.apcmedia.com/salestools/SADE-5TPL8X_R0_EN.pdf>

[26] K. McCarthy. "Comparing UPS System…"

Isolated Redundant Configuration

The next level of complexity in the design of UPS systems is the isolated redundant configuration. This design is fairly similar to the "N" configuration except that it adds a second UPS to protect the critical load in the case that the main UPS fails in some way or needs to be bypassed for routine maintenance. Rather than sharing power between the two UPSs, the main UPS handles all of the critical load unless it fails or requires maintenance, in which case a static transfer switch (STS) quickly and automatically transfers the load over to the second UPS. The STS must be sufficiently fast in its switching (< 1AC cycle) so that the critical load receives fairly continuous power. Figure 23 shows a basic block diagram view of an isolated redundant configuration.

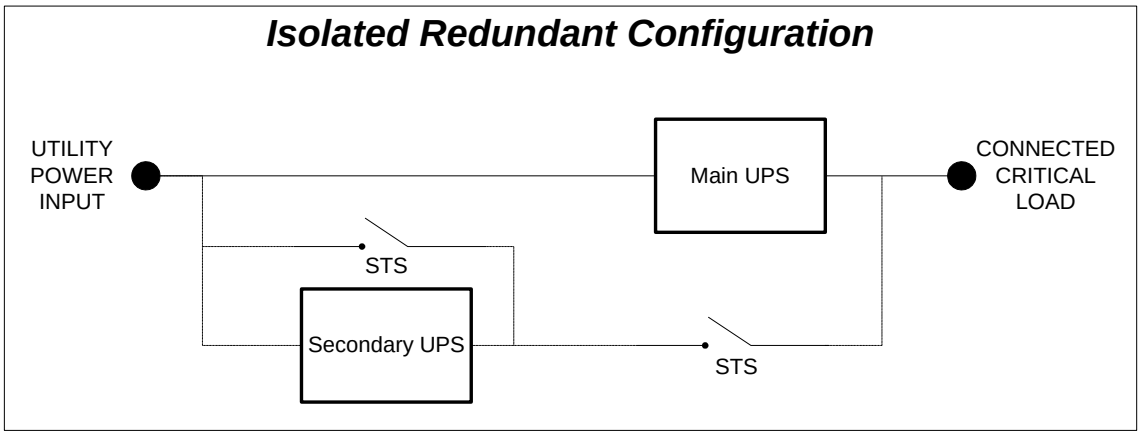

Figure 23

For the additional capital cost of the secondary UPS, the data center manager can make some small gains in overall availability. The availability of an isolated redundant configuration is estimated at 99.93%, which may not seem like much of an improvement over the "N" configuration, but which reduces the amount of unprotected load time to about 6 hours.[27]

As with the "N" configuration, the percent loading of the system remains relatively unchanged with this configuration because the load is bourn solely by the main UPS. Although the load and operational efficiency of the main UPS remain unchanged, the addition of the secondary UPS means that there will be some fixed energy losses required to keep batteries charged and keep the unit operating.

"N+1" Parallel, Single-Bus Configuration

In the "N+1" configuration, higher redundancy is achieved by sharing the data center load between two UPSs arranged in parallel. Both UPSs have an output rating of "N," meaning that either one is capable of supplying the entire data center load. A system of STS's are not required for this system design because both UPSs are tied to the same inputs and outputs; however, the UPSs usually need to be the same make and model in order to maintain a synchronous output to the load. Figure 24 illustrates the "N+1" configuration in a block diagram.

[27] K. McCarthy. "Comparing UPS System…"

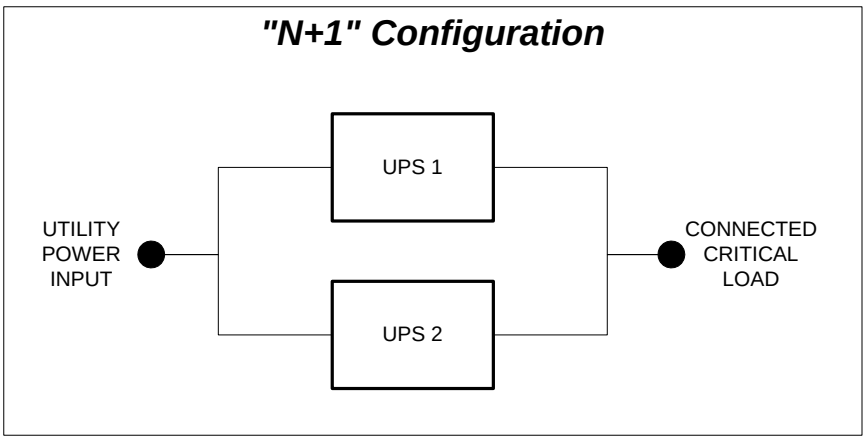

Figure 24

Industry experts estimate that "N+1" UPS configurations have an availability of 99.93%, similar to the isolated redundant configuration described above, without the need for some of the more complicated switch gear.[28]

One of the major disadvantages is that, by sharing the load between two UPSs, the "N+1" configuration has lower operational efficiencies than the isolated redundant and "N" configurations. When a load is shared equally between two UPSs, each of which is rated to handle the full load by itself, then the largest percent load that could theoretically be achieved by either UPS is 50%. Because UPSs are typically oversized for the loads they support, the percent load and operational efficiency of units arranged in the "N+1" configuration drop further. Based on actual data center measurements (see Section 1.3.3), we estimate that typical operational load factors range from 30% to 50% for individual UPSs arranged in an "N+1" configuration, bringing operational efficiency down into the 80% to 90% range.

"2N" and "2(N+1)" Dual-Bus Configurations

Many servers today employ redundant power supplies, which essentially contain two separate, equally rated power supplies that share the server's DC load. Should one "half" of the power supply fail, the other will assume the full load. This is analogous to the "N+1" UPS configuration described above in which two UPSs share a common load. Redundant power supplies do not, however, share their AC input through one cord. Rather, they contain two cords – one for each "half" of the unit – so that the unit can be powered by two separate AC buses.

This is where the "2N" UPS configuration comes in. This configuration allows servers with redundant power supplies to be powered by two completely separate AC sources by protecting the power on two independent utility lines. Rather than sharing the load placed on one utility line, the "2N" configuration places one UPS on each of two incoming utility lines, thus preventing single points of failure. All devices downstream of the UPSs have redundant power supplies, receiving half of their power from the first utility line and half from the second. Should both the utility power *and* the UPS on one of the lines fail, the redundant power supply will still have the other protected utility line as a backup. Similarly, if one of the UPSs requires servicing, critical loads can still be protected by the second UPS. Figure 25 shows a block diagram configuration of the "2N" configuration, which is the most common.

[28] K. McCarthy. "Comparing UPS System…"

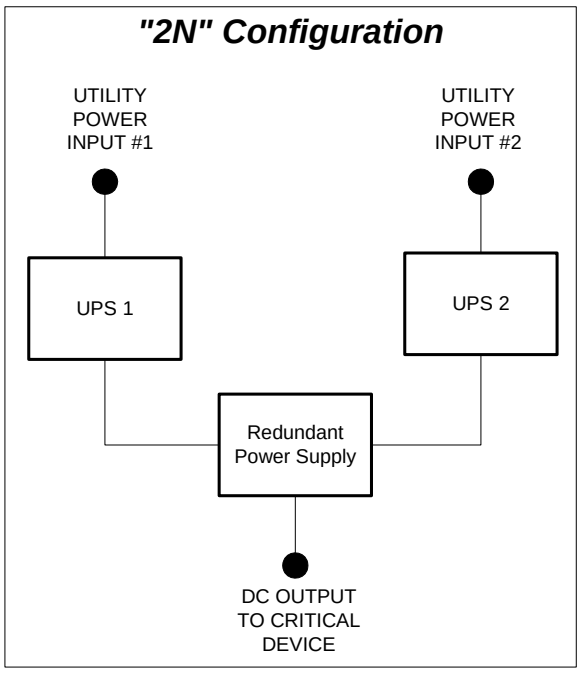

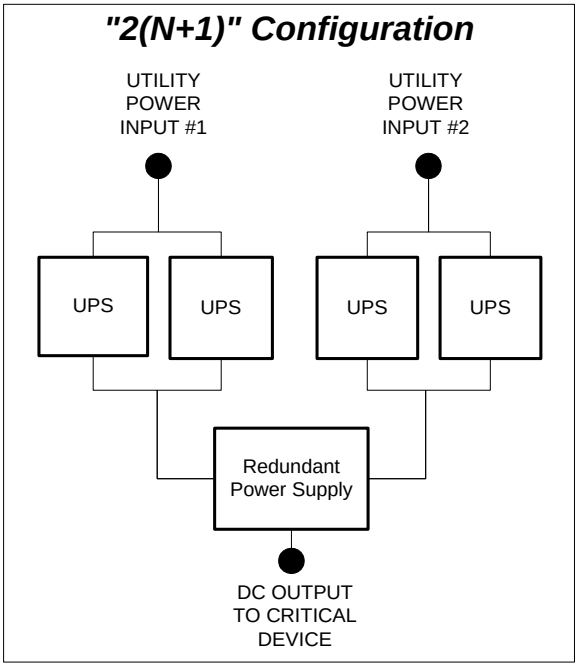

Figure 25 **Figure 26**

An easy way in which to improve the availability of this system is through the "2(N+1)" configuration. This approach takes the same, dual-bus approach used in "2N" configurations, but adds a full, "N+1" configuration on each utility line for a total of four UPSs. Figure 26 illustrates the approach in block diagram form.

Because the "2N" configuration uses redundant utility sources of power, it further improves on the availability offered by "N," "N+1," and isolated redundant configurations, making over four nines of availability (> 99.99% available) possible. With the added redundancy provided through a "2(N+1)" configuration, theoretical redundancy can exceed 99.9999%.[29]

As with the "N+1" configuration, operational efficiency suffers with improved redundancy, because of load sharing and a reduction in percent load for each individual UPS. From an efficiency standpoint, the "2N" configuration is identical to an "N+1" setup, because the data center load is shared between two UPSs. This reduces theoretical maximum loading for each UPS in the "2N" setup to 50%. Field measurements of UPSs operating in "2N" configurations indicate that operational load factors range from 30% to 50% of the UPS's rating. The maximum theoretical percent load is even further lowered in the "2(N+1)" configuration, where load is shared between four UPSs. The theoretical maximum load factor for each individual UPS is 25%, and operational load factors are likely in the range of 15% to 25%.[30] At these low percent loads, UPS efficiency drops off precipitously from the 90% range to below 80% (see figures 19 and 20, for example).

[29] K. McCarthy. "Comparing UPS System…"

[30] We have not measured "2(N+1)" UPS configurations under operating conditions in the past, so our estimate of load factor is extrapolated from measurements on "2N" and "N+1" configurations.

4. Proposed Efficiency Specifications

UPS Operation

Total Cost of Ownership

EPRI Solutions developed a calculator for evaluating rack-level energy savings in data centers derived from higher efficiency power conversion. The calculator takes efficiency-related energy losses into account in several key power conversion devices: the data center UPS, individual server power supplies, and motherboard-level voltage regulator modules for processors. The calculator compares the energy use of a high efficiency case to a base case and reports energy savings through improved efficiency in each of the devices mentioned above. The user can modify various efficiency and power use assumptions to customize the calculations to their needs and can view annual dollar savings as well as net present value of savings over a user-defined period of time.

Ecos and EPRI Solutions further revised this calculator and developed it into a more complete total cost of ownership (TCO) tool that focuses on energy savings for an entire data center facility rather than just one rack. We have used the results of recent UPS and power supply efficiency testing to inform our assumptions for the base case and high efficiency case. We have also built in options which allow the user to specify different UPS redundancy configurations, which as we have shown above, can have a dramatic effect on efficiency and, likewise, operational costs. The user now can select to evaluate the effects of efficiency in any combination of the above devices (UPSs, power supplies, and VRM's). The calculator now allows the user to define and adjust the following parameters through an interface like the one shown in Figure 27. This calculator is currently available on LBNL's website devoted to efficient high tech centers (www.hightech.lbl.gov):

- Number of racks in data center
- Number of servers per rack
- Cost of electricity
- Duration to calculate net present value of energy savings
- Discount rate to calculate net present value of energy savings
- Cooling system efficiency
- Power supply nominal power rating
- UPS nominal power rating
- Redundancy configuration of both the base and high efficiency cases

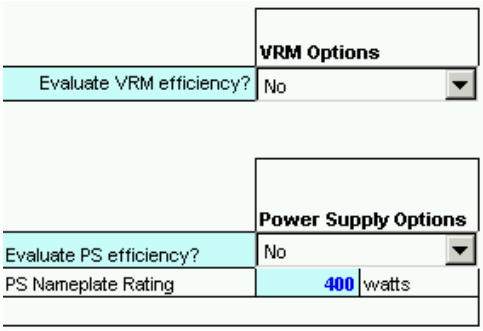

Figure 27: example interface for TCO calculator

Proposed Efficiency Label

Labeling UPS Efficiency

In the early stages of this project, Ecos Consulting and EPRI Solutions uncovered a proposed energy efficiency and power quality labeling scheme for single-phase, “small” UPSs being developed by the Swiss Federal Office of Energy. The Swiss’ proposed label, shown below, documents the measured power losses of the UPS in different modes of operation and provides an estimate of the annual energy consumed by energy losses in the UPS design. The label also reports the UPS’s ability to handle different types of power disturbances and the power quality (power factor and total harmonic distortion) that the UPS presents to the grid. It was designed to match the style of existing EU labels for other electric appliances and was intended to be used by manufacturers to voluntarily report efficiency and power quality information to potential UPS purchasers.[31] The Swiss Federal Office of Energy is currently investigating the possible mandatory use of the label and is also working with members of the European Union on developing a voluntary Code of Conduct (CoC) for UPS systems.

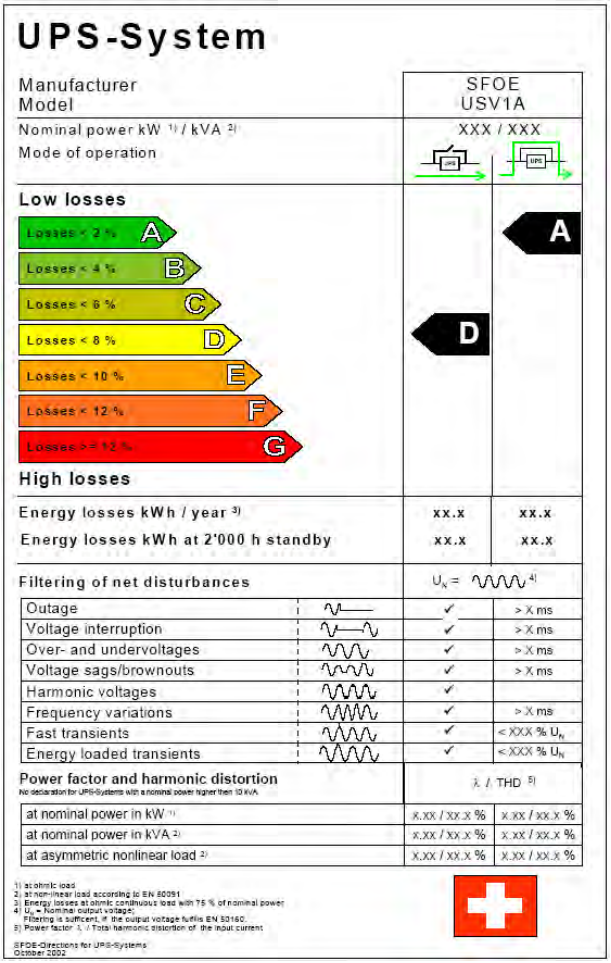

Figure 28

In the interests of harmonizing with international efforts to report and promote UPS energy efficiency in a standardized format, our team created a modified version of the proposed Swiss label and a draft report summarizing the criteria that would be used to evaluate UPS efficiency in the U.S. market

[31] Schnyder Engineers Ltd. *Label for UPS Systems*. The Swiss Federal Office of Energy. October 2002. pp. 1 - 3

under the label. Our proposed label, shown below in Figure 29, differs from the Swiss version in several distinct ways:

- We have modified the label so that it reports differing levels of energy conversion efficiency rather than energy losses. It was our opinion that reporting efficiency was more appealing to a potential UPS purchaser.
- We have eliminated a field in the label that reports energy losses incurred by operating the UPS for 2000 hours with no load. This information did not seem of much use to UPS purchasers, especially data center users who will likely keep their UPS at least partially loaded with servers year round.
- To help direct the focus of the label solely on energy efficiency and to simplify its design, we decided to eliminate the various tables that report on filtering of power disturbances and power quality. Manufacturers readily report this type of information in product literature should potential purchasers wish to compare different UPSs based on these criteria.

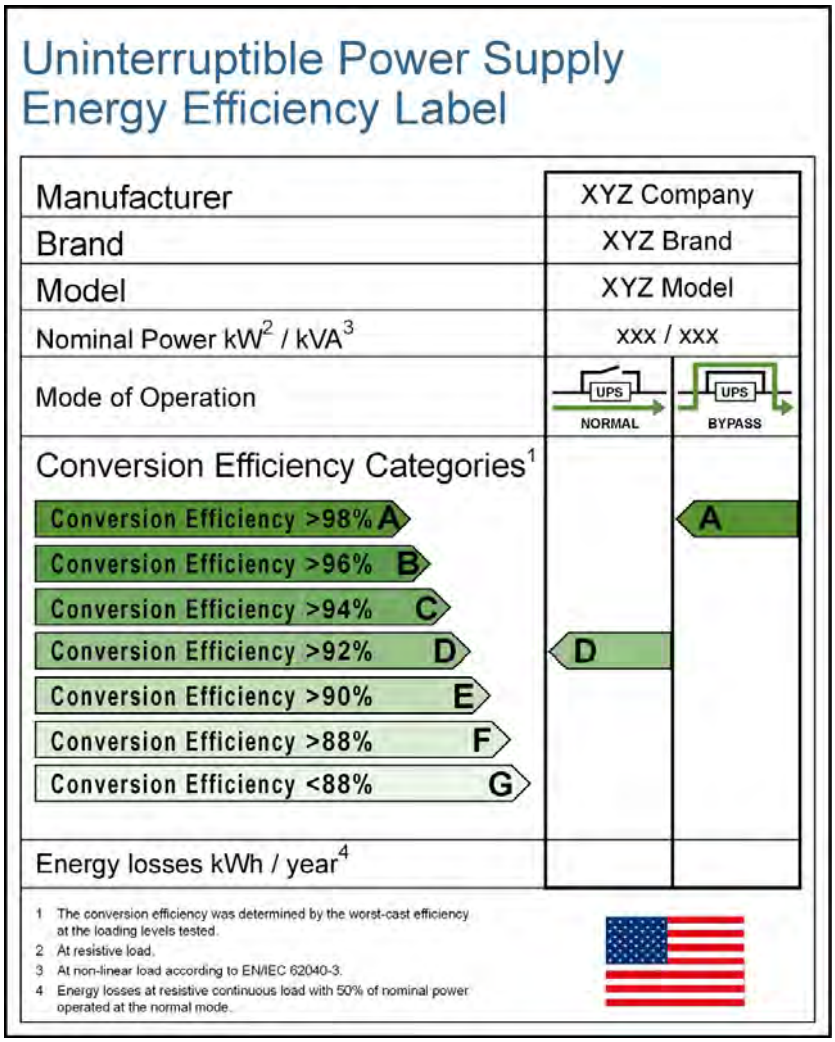

Figure 29

Manufacturer Comments on Proposed UPS Label

The draft label and report were distributed to numerous leading UPS manufacturers to seek comment on the label's design and to ensure that the documentation for the label would provide ample guidance for manufacturers to measure and report UPS efficiency in the proposed format. Several manufacturers responded with comments within a three-week period. We have distilled several common threads of discussion from the various comments received and summarize them below. The main concerns from manufacturers regarding the proposed label were:

Display of performance-related information: Most UPS manufacturers that submitted comments shared a concern that the proposed energy efficiency label does not provide enough information about the UPS's general design and performance. Manufacturers not only want potential purchasers to judge their models based on efficiency, but also performance issues relating to the quality of the UPS output, the ability of the unit to filter out power disturbances, the UPS input power factor, etc. In short, manufacturers want consumers to be able to judge for themselves the potential tradeoffs between efficiency and performance. They feel that consumers should be able to use the label to distinguish the highest efficiency UPS models of a given level of performance.

Bypass mode of operation: Several manufacturers recommended that the column in the UPS label used to show UPS efficiency in the bypass mode of operation be removed altogether. They argue (rightly) that the bypass mode is used to periodically take the UPS out of operation for servicing, which happens only several hours out of the year. During this period of time, the UPS cannot perform its design function: to provide backup power and power conditioning to critical loads. As a result, they argue that bypass mode is not indicative of typical UPS operation and should not be included in the label.

Loading conditions: Manufacturers had two major concerns about testing and loading conditions, namely that lower load levels below 50% of nominal load should be required in testing and that the displacement power factor of linear loads is specified in the labeling guidelines. Manufacturers felt that lower load levels should be required in testing mainly because, as our research has shown, the vast majority of UPSs operate at below 50% of their nominal load. Since the efficiency of a UPS at these lower percent loads can be significantly lower than at 50%, 75% and 100% load, manufacturers thought that an additional loading point in the 20% to 30% range should be required. Due to the effects that load power factor can have on UPS efficiency measurements (one manufacturer cites 2% lower efficiency with increased load power factor), one UPS manufacturer thought that the power factor of the UPS test load should be explicitly stated on the label.

The categorical letter scale: Several manufacturers were quick to point out the past failings of the letter scale used in the proposed UPS label. Recall that the original Swiss UPS label upon which our proposed label was modeled was developed to match European-style efficiency labels used in the EU to report refrigerator and air conditioner efficiency to consumers. These labels all use a categorical letter scale with "A" corresponding to the top efficiency level. Manufacturers noted that, as refrigerators and other appliances have become more efficient in Europe, the original letter scale was no longer adequate to separate out the most efficient models. Over time, additional categories such as "A+" and "A++" have been added to distinguish the top performers. Manufacturers are worried that, as UPS efficiency improves, the categorical scale used in our proposed label could become similarly obsolete. Some also argued that the categorical scale was not the correct approach for the U.S. market because it would be unfamiliar to U.S. consumers and provided too little information about UPS efficiency over a wide range of loads. One manufacturer suggested listing efficiency in a table at 50%, 75% and 100% load as an alternative, although this is not as visual of an approach.

Proposed UPS Efficiency Standard for the U.S. Market

In addition to examining a potential mandatory label for UPS systems sold in the U.S. market, Ecos Consulting and EPRI Solutions also developed and investigated the energy savings potential for a voluntary/mandatory energy efficiency specification. Voluntary labeling programs, such as the U.S. Environmental Protection Agency's ENERGY STAR® program, and mandatory efficiency programs,

such as the California Energy Commission's title standards, have been highly successful in encouraging the efficiency of other electric devices such as lighting, ceiling fans, and refrigerators; however, these programs have never focused attention on uninterruptible power supplies. Programs may have avoided larger UPSs used in mission-critical facilities in the past because these devices are not typically purchased by most consumers.

Below we examine the hurdles to creating a UPS efficiency specification, propose a sample specification based on our measured UPS efficiency data, and estimate the energy savings achievable through the specification.

Challenges to a UPS Efficiency Specification

Creating market transformation mechanisms for UPSs is not as straightforward a task as for other power conversion devices. In many simpler commodity-type power conversion devices, such as external power supplies, there might be a large variation in efficiency between the most and least efficient models, but very little variation in the basic performance and function of the devices. For example, server power supplies are now sold in common form factors designated by the Server System Infrastructure group (SSI) that helps to ensure similar performance characteristics, output voltages, adapter pin-outs, etc. Typical power ratings on server power supplies are also fairly similar and fall within a range from 200 to 600 watts DC.

However in UPSs, the size and performance characteristics of the devices can vary drastically depending on model, manufacturer, etc. UPSs come in a wide variety of sizes ranging from models designed to power individual workstations (< 1kVA units) to those designed to power entire facilities (> 100 kVA), meaning that the largest models can handle two orders of magnitude more power than the smallest models. The way that UPSs deal with power abnormalities and the extent to which UPSs can isolate a load from interruptions in utility power also vary drastically. Standby UPSs and many line-interactive UPSs can pass along voltage sags, line noise, or changes in AC frequency to connected loads, whereas most double-conversion UPSs can filter out all of these disturbances.

Because of the wide variation in performance among UPSs and the equally wide range of available sizes, it becomes increasingly difficult to create a one-size-fits-all UPS specification. UPSs can pay a penalty in overall efficiency by providing better load isolation, filtering, and other types of improved performance, and therefore it would be unfair to compare UPSs with a high-quality output to those with a lower-quality output. Furthermore, smaller power conversion devices typically cannot achieve the same levels of efficiency as larger devices, and so it would be equally unfair to compare the efficiency of a 100 kVA, facility-level UPS to that of a 1 kVA, workstation-level UPS. As a result, any proposed UPS efficiency specification should take the size and performance of the UPS into account.

Sample Efficiency Specification

Ecos Consulting used our data set of factory and laboratory UPS efficiency measurements to create a UPS efficiency specification that could be used by mandatory and voluntary programs. We have mimicked the approach taken by SSI with server power supplies by recommending efficiency levels at 20%, 50% and 100% of the UPS's nominal power rating. Based on the results of field testing of UPS efficiency, we believe that the 20%- and 50%-load efficiency levels are the most important for obtaining real world energy savings, because these points are closest to the range of percent loads at which most UPSs operate.

As described above, it would be difficult to create a "one-size-fits-all" UPS efficiency specification due to large variation in the size and features, both of which have a measurable effect on efficiency. Although we do not currently have enough information to devise an unbiased mechanism that accounts for differences in UPS performance, we were able to account for the UPS size issue by creating three slightly different specifications for UPSs in different size ranges: those with nominal power ratings less than or equal to 20 kVA, those greater than 20 kVA and less than or equal to100 kVA, and those greater than 100 kVA. These three ranges of power ratings correspond to the regions

in Figure 30 outlined with dashed lines. The average efficiency measurements[32] of UPSs in a given region are fairly similar to each other (usually to within a few percent) and do not exhibit any large variations in efficiency as a function of the UPS's rated power output, at least according to our current data set. Our "bracketing" of the UPS specification into three different size groups is merely indicative of the approach that we recommend be taken in formal specification development. In the development of a more formal UPS efficiency specification, a more careful statistical study might need to be performed, coupled with industry stakeholder consultation to determine exactly where the boundaries between different size brackets should be.

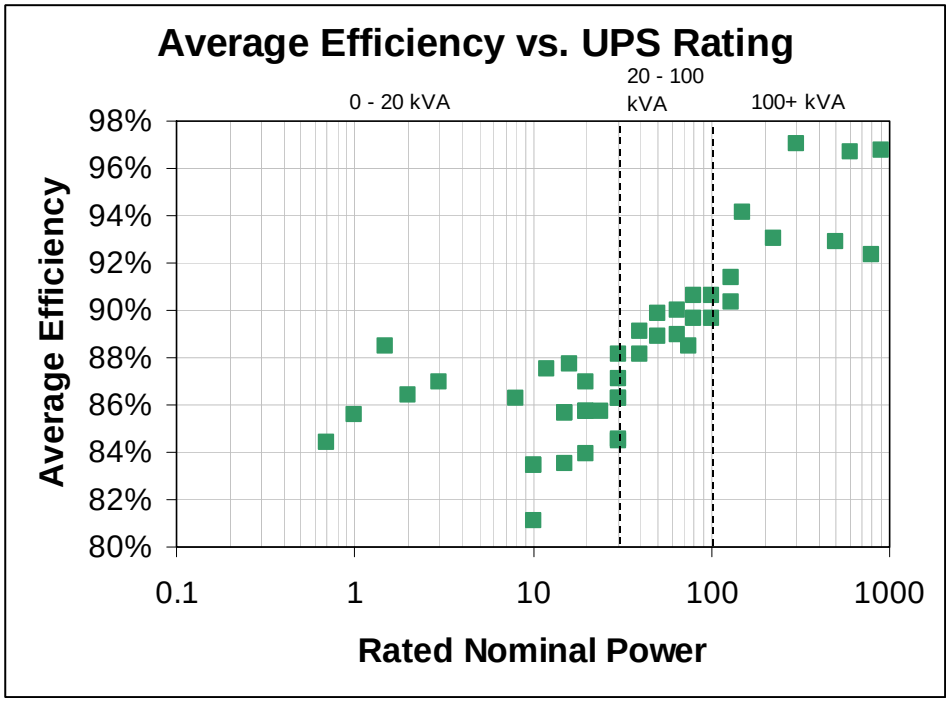

Figure 30

By examining the collected efficiency measurements in our data set within each of the UPS size groups described above, we developed recommended efficiency levels for our proposed specification, shown in Table 8.[33] In order to pass the specification, a UPS must exceed the proposed efficiency levels at 20%, 50%, and 100% load. We designed the specification such that roughly the top 20% to 30% most efficient UPSs in our data set would pass the criteria, which ensures that our proposed efficiency specification is technologically achievable even today. Charts provided in Appendix A plot the various efficiency levels alongside our UPS efficiency measurements.

Table 8: Proposed UPS Efficiency Specification Levels

UPS Nominal Power Output	Efficiency at 20% Load	Efficiency at 50% Load	Efficiency at 100% Load
Below 20 kVA	80%	87%	88%
20 kVA to 100 kVA	83%	89%	90%

[32] Note: average efficiency is calculated by taking the mean efficiency at 25%, 50%, 75%, and 100% load. In several cases where manufacturers did not measure efficiency at these exact load points, it was necessary to use another load point within +/- 10% of 25%, 50%, 75% or 100%.

[33] We used UPS efficiency measured with linear loads to determine specification levels, because the majority of bulk of our data set is comprised of these types of measurements. In a formal specification, non-linear loads should be used to measure efficiency (described in IEC-62040-3) because these loads more closely resemble the types of loads seen by most UPSs in operating environments.

Above 100 kVA	90%	95%	95%

Estimated Energy Savings of Proposed Specification

Ecos estimates that our recommended improvements to the efficiency of uninterruptible power supplies could save millions of dollars in utility costs every year by trimming power conversion energy losses. We have performed a first-order energy use and savings analysis to determine the effectiveness and potential energy savings of the proposed UPS efficiency standard and have built this analysis upon several key pieces of research that have emerged from work on UPS efficiency:

- Estimates of UPS stock presented in Section 1.1.5 for the US data center/IT sector
- Field measurements of average UPS load factors and operational efficiency, presented in Section 1.3.3.
- Factory and laboratory measurements of UPS efficiency across a wide range of loads, presented in Section 1.3.4.

Our general approach was to compare the energy losses or energy consumption of "typical" or "average" UPSs to those operating at our proposed efficiency levels within each of the three size/power categories described above. We first required several simplifying assumptions about UPS loading and operational efficiency. Based on our field measurements of data center UPSs, we assumed that most large UPSs servicing IT equipment and data centers operate at about 38% load and with an average efficiency of about 85%. Since UPSs smaller than 20 kVA were not tested in the field, we had to refer to our factory and laboratory measurements submitted by manufacturers to determine an average efficiency for this category of UPS – 83.6% at our typical percent load of about 38%.[34] Due to the wide range of UPS sizes, we also needed to assume a typical size for UPSs in each of the three categories. We used the midpoint of the range to approximate average UPS size in the < 20 kVA and 20 – 100 kVA ranges, and we assumed a typical value of 200 kVA for UPSs larger than 100 kVA.[35] Finally, we assume that the amount of active output power (in kW) that a typical UPS can provide to loads is roughly 80% of the UPS rating in kVA (this is the apparent power rating and is always larger than the apparent power rating). This assumption is based on typical UPS nameplate ratings from manufacturer data sheets. Table 9 presents some of the key assumptions about "typical" or "average" UPSs.

Table 9: Assumptions for Typical UPS

Output Power Range (kVA)	Assumed Output Power (kVA)	Assumed Output Power (kW)	Load Factor During Operation	Efficiency at Operational Load Factor
< 20 kVA	10	8	38%	83.6%
20 – 100 kVA	70	56	38%	85.2%
> 100 kVA	200	160	38%	85.2%

In the case of high efficiency UPSs, we kept assumptions regarding output power and load factor the same but raised the UPS efficiency based on the levels in the proposed efficiency specification in the section above. Since the efficiency specification does not specifically require an efficiency level at average operational loads (38%), we linearly interpolated between the 20% and 50% efficiency levels to estimate the kind of operational efficiency that one could expect from a UPS that complies with the standard. Table 10 shows our complete assumptions for high efficiency UPSs.

[34] This average efficiency had to be linearly interpolated based on efficiency measurements taken at other load points.

[35] This assumption is based on past sales data from Frost & Sullivan, Roth, and others that indicate that most UPS sales above 100 kVA cluster heavily in the 100 to 200 kVA range. 200 kVA seems to be an acceptable median value.

Table 10: Assumptions for High Efficiency UPS

Output Power Range (kVA)	Assumed Output Power (kVA)	Assumed Output Power (kW)	Load Factor During Operation	Efficiency at Operational Load Factor
< 20 kVA	10	8	38%	84.2%
20 – 100 kVA	70	56	38%	86.6%
> 100 kVA	200	160	38%	93.0%

With the given assumptions about the output power of the UPS, the typical percent load placed on the UPS, and the efficiency of the UPS at operational loads, we were able to estimate average power losses in different sizes of UPS for average and high efficiency UPSs. Table 11 shows the estimated power losses and estimated annual energy consumption (AEC) based on year-round operation.

Table 11: UPS Annual Energy Consumption Estimates

Average UPS		
Output Power Range (kVA)	Power Losses (kW)	Annual Energy Consumption (kWh)
< 20 kVA	0.59	5,197
20 – 100 kVA	3.68	32,211
> 100 kVA	10.51	92,032
High Efficiency UPS		
Output Power Range (kVA)	Power Losses (kW)	Annual Energy Consumption (kWh)
< 20 kVA	0.57	4,971
20 – 100 kVA	3.28	28,693
> 100 kVA	4.55	39,878

Using this information and our estimates of UPS stock, we constructed a first-order estimate of UPS energy use. This energy use and savings estimate is limited to the data center/IT sector due to the scope of our research, in which we measured the efficiency of data center UPSs both in the lab and in the field. Based on our estimates, the data center/IT sector currently consumes roughly 7.1 TWh (7.1 billion kWh) of electricity per year as a result of power conversion losses in UPSs – an amount equal to about 0.3% of combined commercial and industrial electricity consumption in the US.[36] Assuming an average electric utility charge of $0.10/kWh, this amounts to over $700 million spent every year by the data center/IT sector in the US to pay for electricity that is simply converted to heat inside the UPS.[37] We estimate that the State of California alone, which contains roughly 15% of the nation's data center floor space,[38] consumes about 1 billion kWh of electricity and spends about $100 million per year in electric bills due to power conversion losses in data center UPSs.[39] Our national

[36] Percentages based on "No. 916. Electric Energy Sales by Class of Service and State: 2001." *Statistical Abstract of the United States: 2003.* U.S. Census Bureau. 2003. p. 593

[37] Note: this figure does not include electricity expenses for the devices connected to the UPS

[38] R. Juarez, M. Alic, K. Chetan and B. Johnson. *SpaceDex III – Hosting Space: Not All Hosting Space is Created Equal – Smart, Complex Space Takes Data Center Stage.* Boston: Robertson Stephens Inc. 2001.

[39] Because California contains 15% of the nation's data center floor space, its UPSs consume 15% of the energy consumed by all UPSs in the US.

energy use figures are noticeably higher than past estimates made by Arthur D. Little for the U.S. Department of Energy, which estimated UPS energy consumption in the U.S. "IT/telecom" sector at 5.8 TWh per year.[40] Both Ecos' and Arthur D. Little's energy use estimates are shown in Table 12. Also included are Ecos Consulting's estimate of UPS stock and energy use in California, even though the 2002 DOE/ADL report did not make such an estimate. There are several reasons for the estimated increase in UPS energy consumption:

- The UPS stock has grown significantly since the publication of the Arthur D. Little report, due to heavy sales at the tail end of the tech boom and continued, although weaker, sales of UPSs in the three years since. Overall, our stock estimate is 65% higher in terms of numbers of units.
- Research conducted by Ecos Consulting, EPRI Solutions, and LBNL indicate that UPSs typically operate at lower load factors and with lower efficiency than was assumed in the Arthur D. Little estimate, which assumed roughly 50% loading. We now know that UPSs more often operate in the 30% to 50% range of loads, where power conversion efficiency begins to decline.

Table 12: Comparison of UPS Energy Use Estimates

	DOE/ADL 2002 Estimate of UPS Annual Energy Consumption in U.S. IT/telecom Sector	**Ecos/EPRI 2004 Estimate of UPS Annual Energy Consumption in U.S. IT/data center Sector**
Estimated National UPS Stock	151,761	250,343
Nationwide Annual Energy Consumption (TWh)	5.8	7.1
Estimated California UPS Stock	n/a	37,551
California Annual Energy Consumption (TWh)	n/a	1.1

Based on a comparison of Ecos' "high efficiency" and "typical" cases for UPS energy consumption, we have estimated the energy savings potential for enacting our proposed UPS efficiency standard in the U.S. Again, due to the scope of our work, this estimate applies to the data center/IT sector and does not indicate savings that might be achievable in residential or industrial UPS installations. Although it might not be likely that U.S. federal authorities would implement a national, mandatory standard governing the efficiency of UPSs, we estimate that our proposed efficiency measures would have saved 2.8 TWh of electricity in 2004 if all UPSs in use in the U.S. were forced to comply. This would result in about $280 million in savings from lower electric utility bills. Applying the proposed efficiency specification on a mandatory basis could cut nationwide energy use by UPSs by about 40%; however, the proposed savings would amount to only about 0.1% of combined industrial and commercial annual electricity consumption.

A potentially more feasible way to encourage energy savings in UPSs would be the introduction of a voluntary UPS efficiency specification that would be administered under a program such as ENERGY STAR. In this energy savings scenario, it is likely that only a fraction of the UPS market will meet the proposed efficiency levels. If, for example, 25% of the UPSs installed in the U.S. were able to comply with the proposed UPS efficiency levels, the nation would save about 700 million kWh of electricity per year, amounting to about $70 million utility cost savings. We have summarized our energy savings scenarios for mandatory and voluntary specifications in Table 13, along with similar energy

[40] K. Roth, F. Goldstein and J. Kleinman. *Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings, Volume I: Energy Consumption Baseline.* Cambridge, MA: Arthur D. Little, 2002. p. 84

savings estimates that apply only to the State of California, which we assume contains about 15% of the nation's UPSs.[41]

Table 13: Energy Savings Scenarios for Proposed UPS Efficiency Specification

	Mandatory Specification – 100% of Installed Units Comply	**Voluntary Specification – 25% of Installed Units Comply**
Nationwide Annual Energy Savings (TWh)	2.8	0.71
Nationwide Annual Utility Cost Savings (million USD)	$280	$71
California Annual Energy Savings (TWh)	0.42	0.11
California Annual Utility Cost Savings (million USD)	$42	$11

The energy savings figures shown here could be significantly larger when examining the *entire* UPS market and not just the data center/IT sector. For UPSs above 5 kVA in power output, the data center/IT sector only accounts for roughly two thirds of the installed units. An additional one third of UPSs installed in small offices, industrial facilities, hospitals and the like have not been included in our energy use and savings estimates. Our analysis also excludes the below 5 kVA UPS market segment, more typical of small office applications, residential installations, etc. The combined sales of UPSs below 5 kVA amounted to over 8 million units in 2004, close to 200 times the unit sales volume of UPSs greater than 5 kVA in power output.[42] Although these are lower power devices than the larger UPSs examined in this report and would thus have smaller per-unit power losses, they might represent an attractive energy savings opportunity simply due to the large number of units in operation that should be studied in the future.

It should be noted that additional energy savings can be realized from reduced facility cooling needs because less waste heat is generated from the conversion and storage equipment. It has been estimated that the secondary savings from reduced cooling needs as a result of more efficient facilities is on the same order of magnitude as the direct, primary savings from efficient conversions. Finally, there is also additional savings from reduced capital investment in equipment if a more efficient system is selected, as well as reduced real estate investment from smaller equipment footprints. These additional savings can be realized up front rather than through reduced operational costs.

[41] Based on estimates that California contains 15% of the nation's data center floor space. R. Juarez et al. 2001.
[42] Frost and Sullivan. 2004.

5. DC Architecture Scoping Study

Market Trends in IT Load Power Configuration

New generation data processing equipment used in data centers and general business services provides ever-increasing processing speed and bandwidth utilization benefits, at the same time, increasing growth in Internet-based communication and commerce activities continually stimulate the need for more performance. As a result of this self-reinforcing spiral, the amount and concentration of electrical energy being dedicated toward data processing and storage infrastructure is also on an exponential growth path. This growth creates new challenges regarding energy and resources to power these systems.

Data processing equipment is generally built around microprocessor cores, or individual server units in dense racks and clusters. Data centers use AC in a conventional distribution to deliver power to the data processing equipment (servers, storage disk arrays, etc.). This typical delivery scheme involves power conversion devices (power supplies) to convert AC to DC at the equipment component and processor levels. In the past these microprocessors-based server unit would typically utilize less than 100W of energy per unit. However, modern equipment quite often exceed 200W per unit, depending on configurations, with future equipment forecast to go well in excess of this value. As these channels are paralleled together into clusters, the total energy being utilized presents unique challenges.

Inside the server boxes, power supplies can provide power factor correction as well as load isolation from the incoming power line. This conversion generally involves at least two stages. In addition, most modern microprocessors require very low voltages at fairly high currents, such as 1.1V at 100A. The precision of the voltage required is such that voltage regulation circuitry must be located directly next to the microprocessor. In order to effectively realize this circuitry, most processors require that an intermediate DC voltage, such as 12VDC, be delivered to the processor/local regulator combination. Thus, from the power supply, there can be up to six or more power conversion stages between facility power entry and the microprocessor. (480VAC to 208VAC, 208VAC to 400VDC, 400VDC to 208VAC, 208VAC to 400VDC, 400VDC to 12VDC, 12VDC to 1.1VDC)

Depending on processor and server loading, which can dynamically range from 0% to 100%, the efficiencies of power supplies can be lower at lower load levels and can significantly impact overall system efficiency. In many cases, redundant power supplies are used to deliver this power, either on standby or in load-sharing configurations inside servers, both of which reduce individual power supply loading.

The need for reliability also necessitates the use of UPS units to condition the AC power provided to the servers. The central AC UPS generally involves conversion from incoming raw AC power to DC, and then reconverting from DC back to AC. Moreover, since the UPS generally requires a maintenance bypass switch, facility-level entry power usually must be transformed from 480VAC down to 208VAC prior to the UPS. The central battery is then connected as an alternate input to the DC-AC converter so that if incoming raw AC power is interrupted, the system automatically switches over to battery power. Uninterruptible AC power is then passed through an AC power distribution grid and fanned out to individual data processing apparatus. The servers receive the AC output of the UPS (converted from DC) and then convert again to DC with its power supply.

Typically, uninterruptible power is viewed at the facility level, as opposed to the equipment level. This perspective provides an easy division between the facility power equipment and data processing equipment, with each focusing on a different part of the power delivery. However, this view also makes it difficult to ascertain and optimize overall operating efficiency and total cost of ownership, since losses from each of these conversion processes directly translate into heat, adding to both the server cooling load and the overall data center's cooling load.

Efficiency of AC- and DC-based Systems

Inside the server units, the limiting factors on power conversion efficiency are voltage rating of semiconductors, along with their corresponding conduction losses. In addition, power converter package size limitations, and the economics of power supply design and manufacturing can also play a part. All of these factors combine to ultimately limit available efficiency improvements for the power conversion processes. Power conversion efficiency for best-in-class computer grade AC-DC power supplies currently do not provide efficiency higher than about 80% (See related report on “best-in-class” testing results for server power supplies). For servers, on-board non-isolated converters that provide final processor power conversion can range as high as 98% efficiency.

The other area of complexity within large data processing installations is uninterrupted AC power distribution. Facility-level AC UPS systems generally have a centralized circuit breaker panel with power monitoring and bypass provisions for UPS and battery servicing. Often, these units are located at some distance away from the equipment being powered, leading to possible confusion regarding which equipment is fed by what breaker and introducing the possibility of inadvertent information technology equipment shut down due to operator confusion. In some cases the nature of the UPS being utilized is very sensitive to load current harmonics, resulting in a great level of care being required to assure data processing equipment and UPS compatibility.

The centralized battery plant utilized in facility-level UPS systems can be a compromise between what is required for successful system realization and component limitations. The actual battery run time required for acceptable system operation is often just a few minutes. The time represents the delay needed to switch to alternate utility power feeds, or to bring an auxiliary source of power on line (such as a motor generator). However, when batteries for centralized AC UPS systems are sized, the nature of the voltages required or battery type chosen can often result in hold up times well in excess of what is required, resulting in wasted energy storage and delivery capabilities. Additionally, the lower efficiency presented by the off-line power supplies utilized by data center equipment produces an extra load on batteries that only goes into producing heat instead of power conversion.

Using the typical efficiencies of the different conversion steps listed below – from facility entry AC to the processor:

- Step Down Transformer (99.5%)
- AC UPS (85%)
- Computer Power Supply (70%)
- Point of Load Converter (98%)

An estimate of the overall net power efficiency of about 58% can be estimated as follows:

Step Down Transformer (99.5%) x AC UPS (85%) x Computer Power Supply (70%) x Point of Load Converter (98%) = 58%

The above result indicates that for every watt of power utilized to process data, another 0.4W to 0.5W is required to support power conversion. In addition, for air-conditioned facility cooling, a conservative estimate suggests that another watt of power will be required for each watt utilized to cool the power conversion equipment. While additional cooling power can seem insignificant at the individual microprocessor level, when overall data processing activities reach power usage levels on the order of 200kW, or more (for example, a large Internet 4 hub installation), almost 100 kW of power would be wasted in the overall power conversion process, not including additional cooling loads.

Demonstration Plans for DC Architecture

5.3.1 Demonstration Objectives

The demonstration has the following goals:

1. Show that DC-powered server(s) and/or server rack exists in the same form factor or can be built and operated from existing components with minimal effort.
2. Show that DC-powered server(s) and/or server rack can provide the same level of functionality and computing performance when compared to similarly configured and operating server (and/or server rack) containing AC power supplies, as measured with industry standard measurement devices and software benchmarking tools.
3. Measure and document any efficiency gains from the elimination of multiple conversion steps in the delivery of DC power to the server hardware.
4. Identify areas for follow up investigations.

5.3.2 Approach

The proposed demonstration will compare two power distribution schemes to servers and/or server racks that perform the same data processing functions. The two systems are shown in Figures 1 and 2 below.

- System #1 is comprised of a double conversion UPS (local to the rack) powering a server rack with AC.
- System #2 utilizes a single rectifier (local to the rack) to provide DC to the server rack containing DC powered server(s).

Although many data centers use a centralized UPS, and a data center DC distribution design might be based on the plant telecom system which will provide similar functionality, for the purposes of this demonstration, a local system will be used in order to better control all the variables to be monitored.

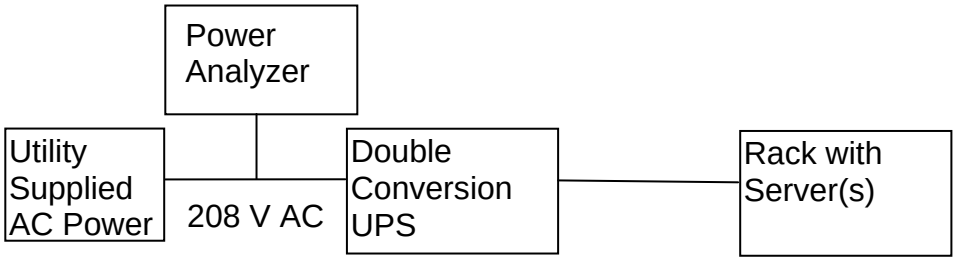

Figure 31: System #1, Utilizing A UPS And AC Powered Servers.

Power analyzers on the input of each system will provide power consumption data as the servers go through a standard benchmarking that uses a controllable and repeatable data processing routine.

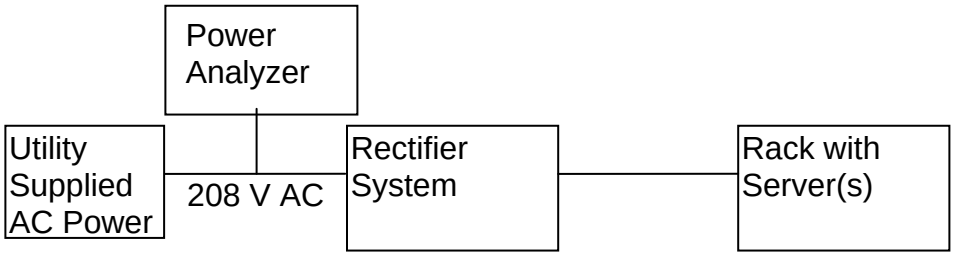

Figure 32: System #2, Utilizing Single Rectifier and DC Powered Servers.

The power consumption by System #1 will be compared to the power used by System #2, to determine if there is any difference in the power consumption by the two systems while performing identical tasks. The results of this demonstration will allow comparison of power consumption by the two approaches, and to determine whether increases in efficiency are possible.

Attachment VIII: Investigation of Computer Performance Metrics

IT Performance Metrics — Project Summary

August 28, 2007

Bruce Nordman, LBNL

While the majority of the data center portion of this project was focused on the power and cooling infrastructure, the actual electronic (IT) loads are also a source of potential energy savings. Early on in our project it was recognized that the lack of good metrics of IT equipment energy efficiency was a barrier to energy savings and were needed by IT manufacturers, purchasers, users, energy analysis, and utilities. We undertook three main activities to advance this topic: reviewing the relevant literature and opportunities in this area, working with industry and others to develop standard metrics and test procedures, and gathering additional information about typical utilization patterns and levels.

The focus of work to date has been on servers, as they are estimated to be the majority (about 70%[1]) of data center IT load. The other components — external storage and network equipment are left for future work.

Literature

Earlier in this project we reviewed academic literature and data from manufacturers about server energy consumption, with a particular focus on performance benchmarks, energy measurement, and utilization. This was informed by early indications that commercial servers often operated at low levels of utilization of their computational capacity. This resulted in a discussion paper showing the importance and potential for reducing server power consumption at low levels of utilization, and metrics which accurately show this behavior. It also reviewed the available data on utilization which showed very low average levels (e.g. XX-XX%), and that there were less data of this form than desirable.

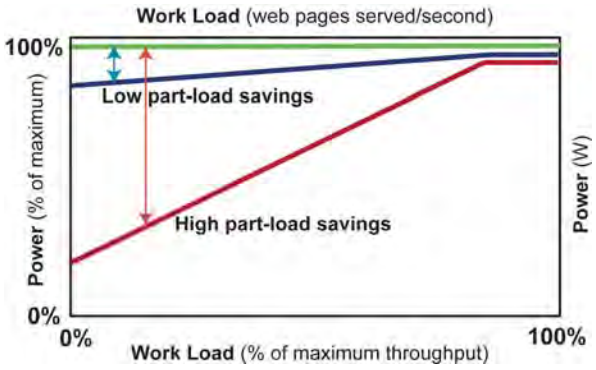

The review is presented in Appendix A.

Test Procedures

At the January 31, 2006 and March 27, 2006 EPA-sponsored workshops with industry, there was widespread agreement that the largest barrier to improving the energy efficiency of servers was the lack of good common metrics of server energy efficiency. The Standard Performance Evaluation Corporation (SPEC) had already begun considering creating a power-performance benchmark project, and the consensus driven by the EPA meetings helped make that initial inquiry into a full-fledged project. Benchmark development is a

[1] "EPA Report to Congress on Server and Data Center Energy Efficiency", PL 109-431, August 2, 2007

slow process and it was desired to have a test procedure in place as soon as possible rather than wait until the SPEC process was done. Under the lead of LBNL via Jon Koomey, and the auspices of EPA, an initial procedure was developed[2] (see Appendix B). This procedure is a structure for testing a server using some existing benchmark; which benchmark(s) to use is not specified

In the 18 months since that time, the industry committee (SPECpower, www.spec.org/specpower) has devoted considerable resources to making that benchmark a reality; it is scheduled for release in late 2007. LBNL has been a member of the SPECpower committee since the beginning, providing advice and expertise in energy/power metering, test procedures, and in what data and approaches are most useful to the energy efficiency community. Deliberations and data within SPEC committees are confidential unless voted otherwise, but once the benchmark and associated reporting are complete, they will be released for public view and be fully transparent.

SPECpower has made clear from the beginning that it will assess systems at varying levels of utilization, to provide energy consumption levels for each utilization level. This allows a purchaser to pay attention to those levels most appropriate to their operation. It is also necessary to determine the maximum performance level (and associated power consumption) to determine what the lower performance levels are. The graphic at left (from the SPEC web page) shows the stair-step pattern typical of running a server through defined points of performance. The initial SPEC product will be a single workload, with others expected to follow. The "workload" (or particular software application and way it is exercised) is critically important as different servers or configurations of servers can be optimized for one type of workload or another, depending on the balance of raw processing, floating-point arithmetic, memory usage, disk access, etc. that is needed. Comparing two systems, one may seem to be more energy efficient with one workload and the other for a different workload.

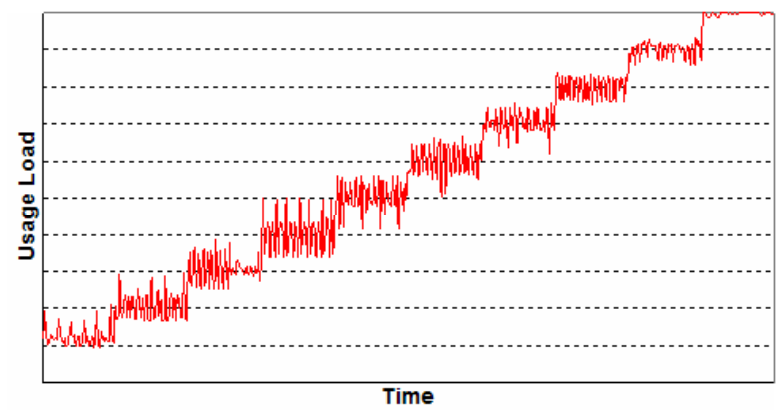

Utilization

The most important thing a data center operator wants to know about their server's computational capacity relative to their actual use is that it is "enough". Not having sufficient capacity is tantamount to failure for a commercial facility. The downsides to having excess capacity are small by comparison — one of these is extra energy

[2] Koomey, Jon et al., "Server Energy Measurement Protocol", November 3, 2006, version 1.0, available at: http://energystar.gov/index.cfm?c=prod_development.server_efficiency_benchmark

consumption. There are many reasons for average utilization levels to be relatively low: systems need to be designed for periodic or hoped-for peak traffic, extra systems need to be available to accommodate some amount of hardware failure, development systems need to be separate from production systems, maintenance activities may take some portion of systems offline periodically, keeping one application per server (to avoid a crashing application from taking others down with it), and people may not have a good idea of the levels at which some system resource will become scarce and lead to performance problems, naturally leading to staying away from any level which could cause problems.

Another issue is the question of what measure of utilization to use. Most people use processor utilization as reported by operating system utilities, but these are known to be increasingly problematic, particularly as systems have multiple processors and cores, and as they vary their clock rate. Nonetheless, a replacement statistic has not yet emerged.

In asking anecdotal assessments of typical utilization from many people in the course of this project, the most common figures reported were in the 15%-20% range, with the highest being 30% and lowest 7%. Other sources of data (e.g. standard reference web traces) showed peak throughput about 4 times the average, making 25% the most that average utilization could be (less if the peak didn't require 100% of system resources). Other traces showed peak values over 20 times the typical, so that average utilization would be less than 5%.

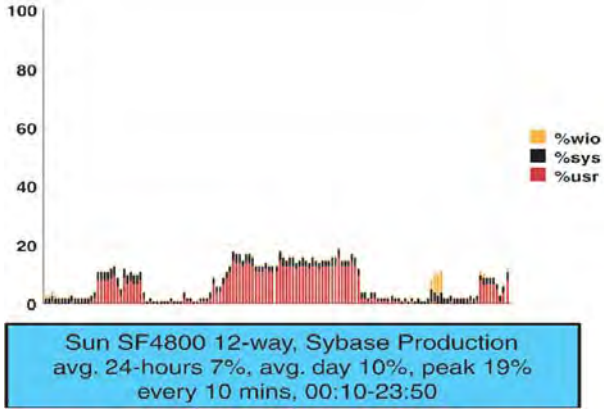

Three sources of data were assessed that included definite direct measures of utilization. The first was from a consultant that assessed hundreds of thousands of servers, and found an average over these of 4.5%. A large financial institution shared detailed utilization measures of their systems (a figure they track closely) and for several hundred servers reported 7% as their long-run average. A set of papers from IBM reported 15-20% utilization for unix systems, and less than 5% for Windows systems, with "less than 10%" an overall assessment (see figure for example)

One driver of low utilization is the practice of assigning only one application to each volume (small) server, to avoid having an application that fails taking down other applications with it if the operating system needs to be rebooted. Virtualization can alleviate this problem by having separate operating system instances for each application on a server to isolate them. This has great theoretical benefit to reduce the number of servers and hence energy use, but does not necessarily reduce the value of reducing power at low levels of utilization. For example, even if we assume that virtualization is wildly successful and reduces the server count by 75%, it would only increase average utilization from 5% to 20%.

Conclusions and Recommendations

We believe that the emergence of metrics of server power and performance has already affected design decisions of server manufacturers as they know that decisions made today will reflect well or badly on the power/performance of systems they sell for years to come. Participation by the energy efficiency community in the development of such metrics has several advantages: helps industry understand that efficiency community is an ally in the process, it increases the quality and usefulness of the product, and ensures that there is knowledge about benchmark design and operation in the efficiency community. Once the metrics are widely available, we can expect to see them used in advertising and utility rebates so that an ability to ensure they are being used reasonably will be important.

While there are commercial software products for tracking utilization across servers in a data center, doing so is more cumbersome and expensive than it should be. There is a public role for working towards standard mechanisms for IT (and perhaps non-IT) equipment in data centers to routinely report key performance and power characteristics to the data center operator.

There is much more to be mined in the area of system utilization as a theme to gain energy savings, both in increasing average utilization, and driving the market to systems that are more responsive in energy use to compute load. In particular, the practice of powering down servers when capacity needs are lower is the least-well explored mechanism, and one which needs metrics tailored to it.

In the course of preparing the EPA Report to Congress on Servers and Data Centers, it became clear that we lacked a good overall measure of compute capacity for the country as a whole, to assess the role of increasing demand for computation as a factor in driving up data center energy use. We also lack a good model for understanding what drives that inexorable increase in computation, though some important factors are certainly well known.

Metrics for external storage and network equipment need to be developed. In addition, high-performance (scientific) computing has specific needs for metrics quite distinct from those of general-purpose business computing.

Attachment IX: Energy Efficient Power Supplies

High Performance Buildings: Data Centers

Server Power Supplies

December, 2005

My Ton
Ecos Consulting
Portland, Oregon

Brian Fortenbury
EPRI Solutions
Knoxville, TN

The work described in this report was performed as part of Lawrence Berkeley National Laboratory's (LBNL) "High-Performance High-Tech Buildings" project focused on improving energy efficiency in Laboratory, Cleanroom, and Data Center facilities. This project was sponsored by the California Energy Commission's Public Interest Energy Research (PIER) Industrial program with technical leadership by LBNL. Ecos Consulting and EPRI Solutions performed the research reported herein collaboratively under subcontract to LBNL and with input from industry advisors and UPS manufacturers.

Legal Notice

This report was prepared as a result of work sponsored by the California Energy Commission (Energy Commission). It does not necessarily represent the views of the Energy Commission, its employees, or the State of California. The Energy Commission, the State of California, its employees, contractors, and subcontractors make no warranty, express or implied, and assume no legal liability for the information in this report; nor does any party represent that the use of this information will not infringe upon privately owned rights. This report has not been approved or disapproved by the Energy Commission, nor has the Energy Commission passed upon the accuracy or adequacy of this information in this report.

Table of Contents

1 Executive Summary ... 4
1.1 Current Power Supply Efficiencies ... 4
1.1.1 Laboratory Testing ... 4
1.1.2 Field Testing ... 4
1.1.3 CPU Utilization ... 5
1.1.4 Proposed Power Supply Efficiency Standards ... 5
1.2 Energy Saving Potentials ... 5
1.2.2 Annual Energy Consumption – Other Data Center Devices ... 7
1.2.3 Energy Savings Potential – Other Data Center Devices ... 7
2 Existing Power Supply Efficiency ... 10
2.1 Server Power Supply Topologies ... 10
2.2 Test Protocol for Power Supplies ... 11
2.3 Laboratory Testing and Results ... 12
2.3.1 Power Supplies Tested ... 12
2.3.2 Efficiency Test Results ... 13
2.3.3 Manufacturer-Reported Efficiency Data ... 15
3 Field Testing ... 16
3.1 Field Testing Protocol ... 16
3.2 EPRI, Ecos, and LBNL Field Tests ... 17
3.2.1 Servers Tested ... 17
3.2.2 Server and Power Supply Loading Curves ... 19
3.2.3 CPU Utilization ... 21
4 Proposed Power Supply Efficiency Specifications ... 23
4.1 Coordination with SSI ... 23
4.2 Market Penetration of Servers and Power Supplies ... 24
4.2.1 Servers ... 24
4.2.2 Power Supplies ... 29
4.3 Energy Savings Potential of Proposed Efficiency Levels ... 30
4.3.1 Annual Energy Consumption of Servers ... 31
4.3.2 Existing vs. Proposed Power Supply Efficiency Levels ... 32
4.3.3 Energy Savings Estimates ... 33
5 Power Supply Energy Savings in Other Devices ... 35
5.1 Power Supply Topologies in Other Devices ... 35
5.1.1 Routers ... 36
5.1.2 Switches ... 39
5.1.3 Hubs ... 41
5.1.4 External Data Storage ... 43
5.2 Existing Efficiency of Power Supplies in Other Devices ... 45
5.2.1 Manufacturer Data ... 45
5.3 Market Penetration of Other Devices and Energy Savings Potential ... 47
5.3.1 AEC ... 48
5.3.2 Adjusted AEC ... 49
5.3.3 Efficient Power Supply Energy Savings Potential ... 49

1 Executive Summary

As part of the California Energy Commission's PIER (Public Interest Energy Research) initiative on efficient data centers, Ecos Consulting and EPRI Solutions were tasked with investigating and characterizing server power supplies as well as the power supplies used by other devices in data centers.[1] The goals for this project include an analysis of power supply efficiencies, which included the development of an accepted test protocol for server power supplies, lab and field testing of a broad range of server power supplies and documenting the results. Other objectives of the project include the wide circulation of efficiency findings to the industry through the CEC, PIER, Lawrence Berkeley National Laboratories (LBNL), and other industry and efficiency venues, such as SSI – the Server System Infrastructure group, the PSMA – the Power Sources Manufacturers Association, the American Society of Heating, Refrigeration and Air Conditioning Engineers – ASHRAE, and www.efficientpowersupplies.org. Using the efficiency findings and market data, we also estimated the overall energy consumption of servers in the United States, especially in the State of California, as well as the potential savings from the use of more efficient units. Finally, we also worked with industry groups to press the case for more efficient power supplies.

1.1 Current Power Supply Efficiencies

1.1.1 Laboratory Testing

Our test results showed that most server power supplies' efficiency at converting AC to DC typically peaks at loads between 50-60% and drops off dramatically at loads under 30%. The tested power supplies have efficiencies in the 70-75% range (at 50% load). The most efficient power supply tested demonstrate that significantly more efficient server power supply designs do exist on the market compared to the average performers, and this may represent significant energy savings potential for data centers or enterprise computing where servers are operated on a "24/7" basis. Our findings also show that server power supply designs with poor efficiency are still available and can result in unnecessary power consumption and excess heat generation while in use.

1.1.2 Field Testing

To investigate the relationship between server operation and power supply loading, we developed testing protocols and collected field data on servers from a number of facilities, including LBNL (scientific computing), EPRI Solutions (commercial/high tech facility), and Ecos Consulting (general business applications). Our finding from these field tests shows that even highly utilized machines like the LBNL servers do not fully utilize the capacity of their power supplies. It also suggests that server power supplies are oversized for the actual requirements of the machines in which they are being used.

Most servers power supplies measured were operating somewhere between about 20% and 50% of their rated load all of the time, which is the exact portion of power supply curves at which efficiency begins to dramatically decrease. None of the servers tested ever exceeded 50% of its rated output. The combination of low efficiencies and oversizing of server power supplies can have dramatic effects on the net energy consumption of those power supplies. An oversized, inefficient power supply would often waste two to three times as much net AC power[2] to meet that load as a properly sized, efficient one.

[1] At the most basic level, power supplies are used to convert alternating current to direct current needed by the processors in servers as well as other electronic components.
[2] Here net AC power is defined as the difference between the AC power input of the power supply and its DC power output.

1.1.3 CPU Utilization

In addition to measuring power consumption, we also extracted CPU utilization data from each server's log for a 24-hour period during field testing.[3] A CPU utilization curve for each server was constructed representing the percent of time the server's CPU utilization exceeded a given percent. It is difficult to draw any conclusions from such a small sample, but it appears that for the most part power supply loading has very low correlation with "server activity."

The servers we measured do not exhibit any activity-based power management, whereby power consumption is reduced when the server is operating below full capacity. Activity-based power management attempts to scale processor power consumption to the demand for processing resources placed on the CPU. The absence of this technology in the servers that we measured might help to explain the lack of any correlation between the utilization curve of the CPU and the load duration curves of the power supplies.

1.1.4 Proposed Power Supply Efficiency Standards

Based on the efficiency levels documented in existing equipment and the field studies conducted, we closely coordinated with SSI and Intel[4] to propose changes to the SSI specifications for higher power supply efficiencies.

At the Intel Developer Forum in March 2005, in a joint presentation, Intel announced new proposed SSI Industry specifications to encourage more efficient power supplies. The new specifications include testing conditions which align with ATX12V, as well as required and recommended efficiency levels at 20%, 50%, and 100% loading. SSI is also now supporting a power supply technology known as Power Supply Management Interface (PSMI). This new industry standard provides a basic internal mechanism for server power supplies to report power consumption and efficiency data directly to the server's motherboard with 5% to 10% accuracy. The technology could conceivably be used for a number of purposes. Power and efficiency data could be logged on the server's hard drive or reported to HVAC equipment so that fans and cooling equipment could scale their output to the heat output of IT equipment.

1.2 Energy Saving Potentials

We estimated the energy savings potential of more efficient server power supplies in three steps:

1. We estimated the current AEC (Annual Energy Consumption) of servers in the United States using the basic methodology developed in the ADL/US DOE report.[5]
2. We determined the percentage of current server energy use that could be saved due to more efficient power supplies, based on the new SSI recommended specifications.
3. By applying the percentages established in #2 to the AEC estimates established in #1, we were able to estimate the overall energy savings potential of the new recommended efficiency levels.

For California-specific estimates, we used a range of 10% to 15% of estimated US results.

1.2.1 Annual Energy Consumption of Servers

Starting from the basic methodology established in the ADL/US DOE study, we constructed a revised estimate of the annual energy consumption (AEC) of servers in the United States.

---

[3] CPU utilization describes how busy a server's processor or processors are; it tells us the percent of the CPU's processing capacity, or throughput, that is being used to perform tasks.
[4] Brian Griffith, Intel Power Server Architect, EPG (SSI Coordinator).
[5] Roth, Kurt W., Fred Goldstein, and Jonathan Kleinman (Arthur D. Little). "Office Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings. Volume I: Energy Consumption Baseline." Prepared for US DOE, Building Technologies Program, January 2002. NTIS Number: PB2002-101438.

Figure ES 1 shows our revised estimate of server AEC for 2004. The revised AEC is 14.6 TWh, which represents a 45% increase over the ADL/US DOE estimate of 10.1 TWh for 2000. This increase is attributable to growth in both the number and average power draw of low-end servers, which comprise the majority of units. The estimate for California is in the range of 1.5 TWh to 2.2 TWh.

Figure ES1. Server Annual Energy Consumption (AEC)

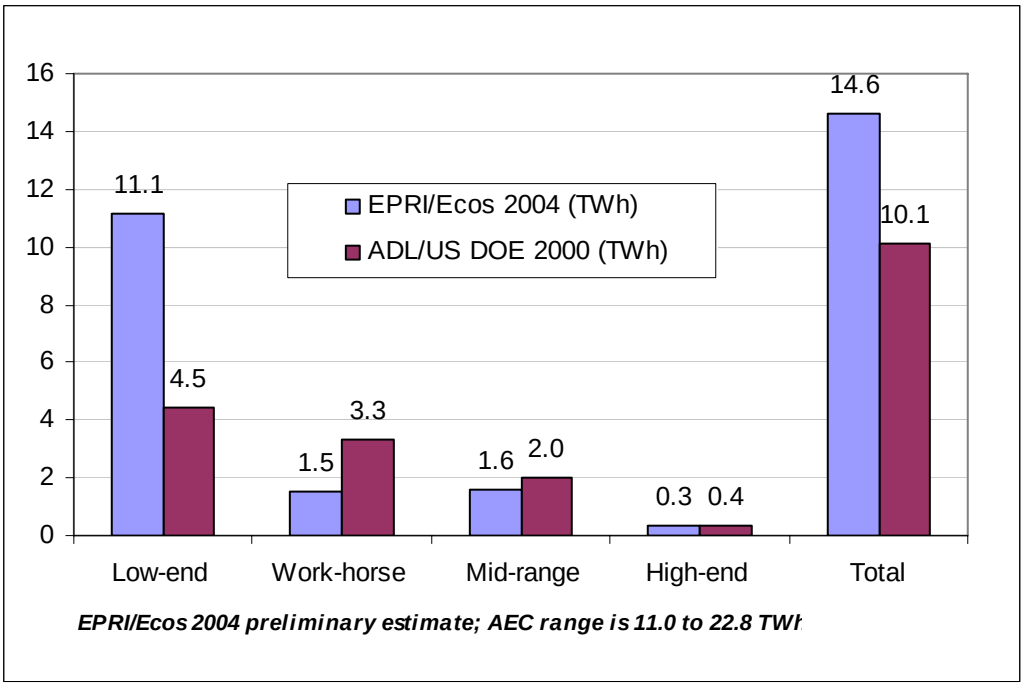

1.2.1 Energy Saving Potentials – Servers

Based on our overall estimate of server AEC, we estimate the amount of electricity that can be saved due to more efficient power supplies would be approximately 1.5TWh. We also project a high-efficiency case, where power supply efficiency is 83% and electricity savings of 2.3 TWh are realized. An estimate for the potential savings for California is between 10% and 15% of the overall US savings potential, or between 230 GWh and 345 GWh.

Table ES1. Energy Savings Potentials

Server Category	Number of Servers (Millions)	Annual Energy Consumption in TWh @ 70% PS Efficiency (Current)	Annual Energy Consumption in TWh @ 78% PS Efficiency (New SSI)	Annual Energy Consumption in TWh @ 83% PS Efficiency (High Adoption)
Low-end	6,587,061	11.1	10.0	9.4
Work-horse	506,470	1.5	1.4	1.8
Mid-range	151,678	1.6	1.5	1.4
High-end	14,730	0.3	0.3	0.4
Total	**7,259,939**	**14.6**	**13.1**	**12.3**

Improved power supply efficiency will also yield compounded savings from reduced air conditioning loads and UPS losses not included here. We also note that the potential for the infrastructure load reduction (i.e. improved overall efficiency and reduced cooling load) is on the same order of magnitude as the power supply efficiency improvement.

1.2.2 Annual Energy Consumption – Other Data Center Devices

We also investigated the range of other devices used in data centers, and arrived at a first order estimate of their energy consumption. Our estimates used a similar methodology to the one employed by ADL/US DOE (2002). This methodology, however, does not account for existing power supply efficiency, nor does it account for the fact that devices normally operate at a fraction of their stated output rating. Therefore, revised our initial AEC estimates to reflect an average of 30% loading and 70% power supply efficiency. Table ES2 presents the "adjusted" AEC.

Table ES2. AEC, adjusted for 70% PS Efficiency and 30% Loading

Segment	AEC (TWh)	Adjusted AEC (TWh)	% of Adjusted AEC
Routers	1.3	0.6	8%
Switches - LAN	9.0	3.9	59%
Switches - WAN	0.4	0.2	3%
Hubs	1.3	0.6	8%
Storage Devices	3.3	1.4	22%
Total	15.3	6.6	100%

Note: AEC is adjusted by a factor of .43 =.3/.7 to account for loading and PS efficiency.

Together, these "other devices" consume a little less than half of the energy that servers do. LAN switches contribute the most to AEC simply due to their large numbers. Remember that, as with servers, not all of the energy consumed goes to power the device itself. Some energy is "lost" or consumed by the power supply itself; the amount of energy consumed depends on the efficiency of the power supply. In the typical existing case, where power supply efficiency is 70%, 30% or about 2.0 TWh is consumed by the power supply (i.e., energy losses)

1.2.3 Energy Savings Potential – Other Data Center Devices

As with servers, energy savings can be achieved by increasing the efficiency (or reducing the energy losses) of the power supply of other devices. These high-level estimates show that there is significant energy savings potential from improving efficiency of power supplies in devices other than servers. While the savings potential is not as large as in servers in absolute terms, the extension of efforts to these "other' devices would be relatively straightforward. Much of the groundwork for improving power supply efficiency and developing specifications for servers, desktops, etc. has already been done or is in process. This work could serve as a starting point for efforts targeting other devices.

The main challenge to encouraging efficiency in the power supplies of other IT equipment like routers and switches is that there is currently no industry body like SSI coordinating standards and efficiency improvements in these products. Power supply designs for this type of equipment are usually customized for a particular product and may vary from manufacturer to manufacturer, whereas in servers, only a few distinct form factors for power supplies exist. The diversity of power supply designs in equipment like routers and switches may complicate efforts to uniformly encourage efficiency improvements.

1.3 Conclusions and Recommendations

This project served well as a first foray into this particular area of high-tech buildings and data centers, and has provided a good overview of server supplies utilization and current efficiency levels. Because of the broad scope, however, we were only able to identify areas for further investigations, but not able to carry out in-depth investigations of any particular server category or application. Below are a number of observations and recommendations based on the findings of this study:

Server Loading: While we got fairly consistent results from our small sample indicating most server power supplies are not fully taxed in everyday use, it would be difficult to project what the loading

would be for servers across all data centers. We recommend a more rigorous sampling of servers and their various tasks as a logical next step.

Server Replacement Cycles: We have not investigated this issue, or the associated issue of what type of market penetration can be expected as these new systems go into the field if new SSI standards are adopted. We currently do not have such information, other than anecdotal – this is another area that we would recommend for additional research

Retrofit Market for Server Power Supplies: We have also not investigated whether or not there exists a retrofit market for server power supplies. Similar to the replacement issue, this is another area that we would recommend for additional research.

Implications of Blade Servers: This server class is currently experiencing high growth, and does not require a power supply to be associated with it. While this category has the potential to increase efficiency due to the fact that many blades can share one power supply, the efficiency of blade power supplies has not been adequately considered.

Continue Work with the SSI Group: Given that the SSI had already published some new standards, this one area that we can best support the industry by helping to promote continuous energy efficiency improvement industry wide. Our recommendation is that the project team continue the relationships that we developed.

More, Better Industry Data: We were not able to obtain regional data on product shipments, thus we were only able to approximate the California market based on population estimates. The server locations that we tested are representative of locations found in CA, and two locations out of the three that were tested employ servers that belong in the “low-end” category. This is another area where additional research is recommended.

Documenting Efficient Server Power Supply Designs: We have not investigated any power supply design changes under the assumption that the industry will pursue their own configurations that meet new proposed standards. It is worth noting that here are a variety of approaches that can be taken to design power supplies for higher efficiency. For example, from the PIER sponsored power supply primer, we find that the most lossy subsystems include:

- The switching element
- Control IC
- Transformer
- Output rectifier

Losses can further be addressed with:

- More efficient Power Factor Correction chip
- Active Clamp on the main transformer to replace lossy RCD clamp

To illustrate some of these approaches, the PIER funded Efficiency Challenge design competition in 2004 yielded around a dozen designs all using different techniques to achieve greater than 80% efficiency. Some of those techniques included:

- Optimized selection of control IC using a variable off-time technique to lower losses.
- Flyback transformer optimized to reduce leakage inductance and winding resistance
- Operation mode optimized to strike a balance between switching loss and conduction loss
- The MOS switch is carefully selected to reduce switching and conduction losses
- Output rectifier is also carefully selected to reduce switching and conduction losses
- Output rectifier can use synchronous rectification with low loss MOS switches
- Output rectifier can use Silicon Carbide diodes to lower losses
- Synchronous rectifier on the input

- Litz wire used in flyback transformer design
- One method called burst-mode control halts all switching for light loads

A manufacturer can use a few of these methods to achieve improved levels of efficiency; the methods chosen may depend on cost, availability, or designer preference. But in conclusion, the industry has a wide variety of choices available and will likely develop even more. There will not likely be any stifling of creativity as a result of these standards.

2 Existing Power Supply Efficiency

This task's objective was to document the existing efficiencies of server power supplies. To do this we:

- Assessed the power supply topologies in servers.
- Developed the test protocol for testing internal power supplies.
- Conducted laboratory testing of different power supply configurations and loadings.

2.1 Server Power Supply Topologies

Servers come in all shapes and sizes, as do the power supplies that they contain. The Server System Infrastructure[6] (SSI) group led by Intel has identified several distinct categories of server power supplies (Table 1).

Table 1. SSI Power Supply Product Categories

Category	Description	Form Factor
EPS1U	Entry-Level Power Supply	1U
EPS2U	Entry-Level Power Supply	2U
ERP2U	Entry Redundant Power Supply	2U
EPS12V	Entry Non-Redundant Power Supply	PS/2
ERP12V	Entry Redundant Power Supply	Pedestal servers
EPS	Entry Power Supply	PS/2
TPS	Thin (low profile) Power Supply	Low-profile servers
MPS	Midrange Power Supply	Midrange chassis, fits 3 or 4 across within a 19" rack mount system
DPS 2.0	Distributed Power Supply	For system using Distributed Power System (DPS) architecture – delivering 48VDC bulk power. 2.74W x 4.86H x 12.8L (inches)

Based on industry market research (See Section 3.2) and consultation with Intel, SSI members and other industry partners, we determined that the dominant technology in the marketplace is the multi-output front-end ac-dc power supply. These power supplies are commonly found in 1U, 2U, and pedestal servers[7], which account for over 90% of the server market and are found in both offices and data centers alike.

Consequently, we selected three corresponding SSI power supply categories—**EPS1U, EPS2U, and EPS12V**—as the focus of this study. These topologies were chosen mainly because they represent the most common designs used in rack and pedestal servers today. They are also non-redundant power supply designs, meaning that they contain only one path for AC electricity to be converted and

[6] SSI, www.sssiforum.org, is an industry initiative intended to provide ready to use design specifications for common server hardware elements (chassis, power supplies, and racks) to promote and support server industry growth. The initiative is comprised of "Promoters" such as Dell, Intel (SSI's founder), IBM, and Silicon Graphics, and "Adopters" -- companies that have agreed to adopt and utilize final SSI specifications in product design and manufacturing. A complete list of Adopters can be found on the SSI website: http://www.ssiforum.org/membership.aspx.

[7] The "U" in 1U and 2U servers denotes that the server is 1 or 2 units (1.75 or 3.5 inches) in height.

delivered to the server's various components (although there may be more than one unit installed in each server box).

A straightforward test method has been developed to test the efficiency of non-redundant internal power supplies, but this method does not yet address redundant designs, in which there may be two or more paths for AC electricity to flow to the server.[8] For this reason it was most sensible to focus on the above categories for the purposes of this study. A standard test procedure for redundant supplies is needed, but that it would likely have different loading guidelines than non-redundant supplies, as redundant units may spend more time at lower loading levels.

2.2 Test Protocol for Power Supplies

SSI has developed design guides, which include recommended efficiency specifications, for many of its designated power supply categories.[9] Current versions of these specifications are available on SSI's website at http://www.ssiforum.org/specifications.aspx.

Early versions of the SSI specifications did include power supply efficiency recommendations, but only focused on *full load* conditions. Yet, as field tests demonstrate (See Section 2.2), many server power supplies operate at less than full load a majority of the time and at a significantly lower efficiency than they would at full load. Reasons for this might include low "traffic" at a particular points in time (e.g., overnight), oversized power supplies, or (although not the primary focus of this study) the use of multiple power supply units to insure up-time in the event of a failure.

To better document the efficiency of server power supplies at these lower loads, EPRI Solutions and Ecos Consulting worked with the SSI group and other industry stakeholders, including the PSMA – the Power Sources Manufacturers Association to develop a standardized test method for measuring server power supply efficiency across a *range of loads*. Such procedures allow the development of standardized efficiency curves, so purchasers can size power supplies properly and reasonably estimate the efficiencies they will achieve in operation.

EPRI Solutions previously developed a standardized test procedure for internal desktop computer power supplies. (See **Proposed Test Protocol For Calculating The Energy Efficiency of Internal Ac-Dc Power Supplies, Revision 4.0** this document is currently available at:

http://www.efficientpowersupplies.org/pages/Generalized_Internal_Power_Supply_Efficiency_Test_Protocol_R4.pdf.

Building upon this work, and in consultation with Intel and SSI, EPRI Solutions developed a new test protocol for calculating the energy efficiency of internal ac-dc power supplies typically used in computer servers. This document is available at:

http://eetd.lbl.gov/ea/mills/ht/Documents/PS/Server_Draft_PSTestProtocol.pdf.

Detailed loading guidelines were added for different types of server power supplies. A draft test protocol was circulated to interested parties (e.g., all SSI and project advisors) and was also distributed at various industry conferences and meetings (e.g., Intel's Fall IT Symposium and the ASHRAE's Technical Committee 9.9 Annual Meeting in 2004). Based on comments received, the test protocol was revised a second time.

Following is a key excerpt from the protocol:

> *4.3 Power Supply Loading*
>
> *The efficiency of the power supply shall be measured at 20%, 50%, and 100% of rated current. In addition to these three load conditions, other loading conditions may be identified*

[8] The redundant units can be configured in a number of ways, thus creating different load levels on the PS depending on the configurations.

[9] Efficiency in this case is specifically defined as measured input power to the power supply minus measured output by the power supply, and divided by the measured input.

that are relevant to the manufacturer and user of the power supply. Procedures for loading power supplies are described in detail in Section 6.1.1 below[10].

Testing at a load condition below 25% load condition should be guided by IEC 60180 Ed 1.0, Measurement of Standby Power, which establishes the measurement methods for low-power-mode operation of an appliance.

2.3 Laboratory Testing and Results

2.3.1 Power Supplies Tested

Having developed the test protocol, 28 power supplies were obtained from manufacturers and other industry partners, such as the Power Sources Manufacturers Association (PSMA) (Table 2).Some of these power supplies were purchased; others were borrowed and later returned to the manufacturer.

S. No	SSI Type	Model No.	Rated Power
			Watts
1	**TPS1U**	DPS-200PB-118	**203**
2	**EPS1U**	DPS-350-PB	**350**
3	**EPS1U**	DPS-350-PB	**350**
4	**MPS**	DPS-450-CB 1	**450**
5	**MPS**	DPS-450-CB A	**375**
6	**TPS1U**	DPS-125FB	**125**
7	**EPS1U**	ENH-0620	**200**
8	**EPS1U**	P1G-6300 P	**300**
9	**EPS12V**	FSP550-60PLN	**550**
10	**EPS12V**	ENS-0246B	**460**
11	**EPS1U**	FSP460-631U	**460**
12	**EPS12V**	FSP460-60PFN	**460**
13	**EPS1U**	ENH-0635A	**350**
14	**EPS2U**	P2G6460P	**460**
15	**EPS1U**	TC1U35	**350**
16	**EPS1U**	TC1U40	**400**
17	**EPS1U**	FSP350-601U	**350**
18	**EPS2U**	TC2U35	**350**
19	**EPS2U**	TC2U40	**400**
20	**EPS12V**	API4FS06	**550**
21	**EPS12V**	API4FS06	**550**
22	**EPS1U**	AP13FS43	**500**
23	**EPS12V**	PSM6600P	**600**
24	**EPS1U**	M1G6500P	**500**
25	**EPS1U**	M1G6500P	**500**
26	**EPS12V**	HP2-6500P	**400**
27	**EPS12V**	HG2-6400P	**500**
28	**DPS**	DPS-400-GB-1	**400**

Table 2. List of Server Power Supplies Tested

[10] Loading guidelines are needed to ensure consistency when measuring the efficiency of server power supplies with multiple outputs. Guidelines currently do not exist, but may be added in the future. Reference: http://www.ssiforum.org/html/adoptedspecs.asp, for a specification that defines a non-redundant power supply that supports an entry-level server.

Each of the power supplies was tested according to the procedures outlined in the *Proposed Test Protocol for Calculating the Energy Efficiency of Internal AC-DC Server Power Supplies, Review Draft Revision 2.0.* Figure 1 provides a photo of the laboratory testing facility.

Figure 1. Laboratory Test Setup

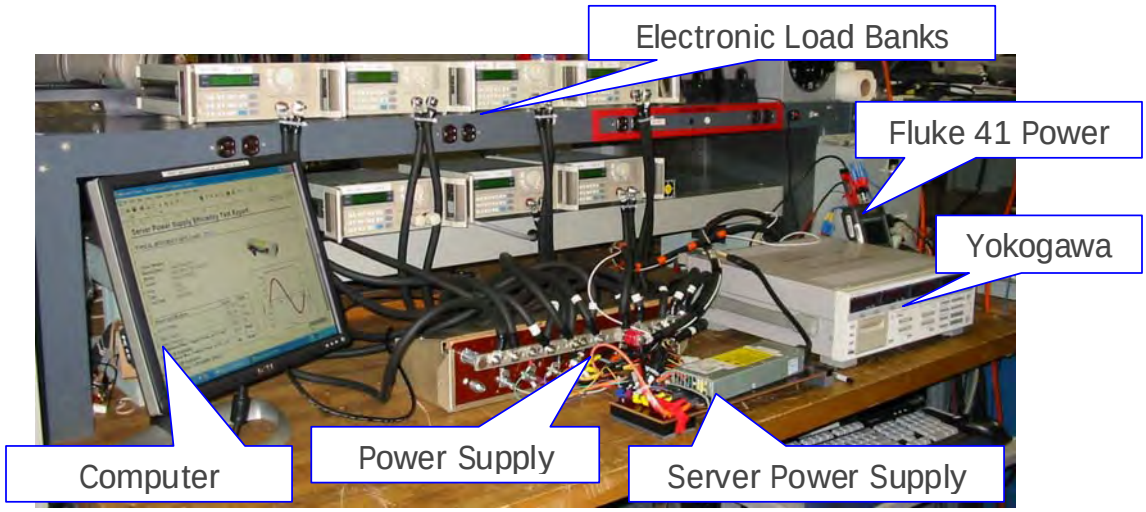

2.3.2 Efficiency Test Results

Figure 2 summarizes the results of the power supply tests.[11] Each black line on this chart represents the measured "efficiency curve" of an individual server power supply: the red line represents the average of all power supplies tested. The largely consistent shape of these curves shows that the measured efficiency typically peaks at loads between 50-60% and drops off dramatically at loads under 30%. Most power supplies are grouped closely together, with efficiencies in the 70-75% range (at 50% load). A few outlying power supplies show relatively high or low efficiency (i.e., the efficiency curves at the top and bottom of the chart). The most efficient power supply tested demonstrate that significantly more efficient server power supply designs do exist on the market compared to the average performers, and this may represent significant energy savings potential for data centers or enterprise computing where servers are operated on a "24/7" basis.

Unfortunately, our findings also show that server power supply designs with poor efficiency are still available and are causing unnecessary power consumption and excess heat generation in data centers., The least efficient power supply measured was a 300W EPS1U unit, which was 52.7% efficient at 20% load, and only 37.0% efficient at 10% load. This means that at 20% load, where many server power supplies are typically loaded, this unit would consume roughly half of the electricity flowing through it simply through inefficiencies in the design.

[11] Individual test reports for each power supply are available at: http://eetd.lbl.gov/ea/mills/ht/Documents/PS/Server_Draft_PSTestProtocol.pdf.

Figure 2. Measured Power Supply Efficiencies

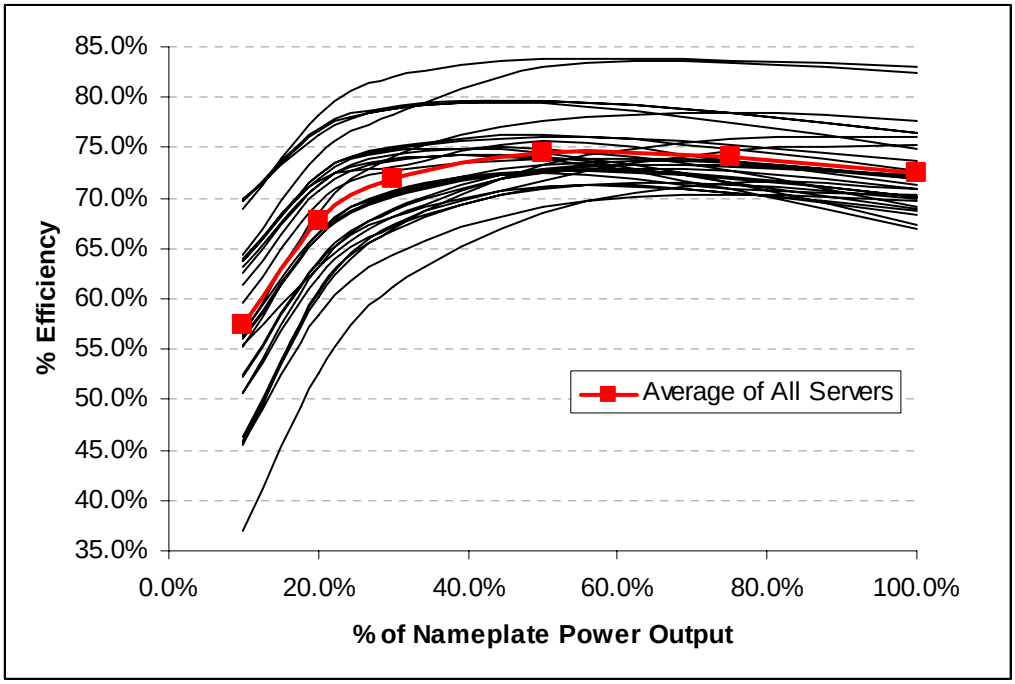

Table 3 shows that the average measured efficiency of all 28 power supplies was only 57.4% at 10%. Peak average measured efficiency for all power supplies was 74.6% at 50% loading.

Table 3. Average Measured Power Supply Efficiency (by SSI Type)

SSI Type	# of PS Tested	Loading %					
		10%	20%	30%	50%	75%	100%
EPS12V	8	60.4%	69.8%	73.3%	75.5%	74.5%	72.6%
EPS1U	12	55.7%	66.3%	70.9%	73.8%	73.8%	72.3%
EPS2U	3	62.9%	71.7%	74.5%	75.0%	72.6%	69.6%
TPS1U	2	53.6%	65.9%	71.6%	75.3%	75.8%	75.0%
MPS	2	48.2%	60.3%	65.8%	70.1%	70.9%	69.9%
DPS	1	64.5%	74.4%	78.2%	82.9%	83.5%	82.4%
Best-in-class	1	69.0%	78.3%	82.1%	83.9%	83.6%	83.0%
Worst-in-class	1	37.0%	52.7%	61.1%	68.6%	71.5%	70.9%
All Power Supplies	28	57.4%	67.7%	71.9%	74.6%	74.2%	72.5%

Note: Efficiency values shown for “All Server PS” correspond to red line in Figure 2.

While the data set is not large enough to state definitively that one power supply form factor is automatically more efficient than another, we can draw the following broad conclusions:

- Power factor correction is much more common in server power supplies than desktop power supplies, though it is becoming increasingly prevalent in high end desktop units
- Surprisingly, average server power supply efficiencies tend to be lower than average desktop power supply efficiencies in our data sets. Part of this is likely due to the recent response by

desktop manufacturers to the market opportunity presented by ENERGY STAR and efforts to push for efficiencies above ENERGY STAR levels (such as 80 Plus)[12] though both also extend to servers.

- The consequences of power supply oversizing remain significant, since power supply efficiency curves are rarely flat. Power supplies designed to operate routinely at 10 to 30% load are generally running about 12 to 15 percentage points less efficiently than if they were sized to operate routinely at 50 to 75% load.

2.3.3 Manufacturer-Reported Efficiency Data

In addition to conducting laboratory tests, we surveyed the manufacturer specification or datasheets for information about power supply efficiency. We discovered that many manufacturers do not report power supply efficiency, and those that do typically only report it at full load (i.e., 100%). Of more than 100 datasheets surveyed, only 71 reported power supply efficiency data. Table 4 provides a summary of the data collected.

Table 4. Manufacturer-Reported Efficiencies (by SSI Type)

SSI Type	# of Power Supplies Surveyed	Average Efficiency (at full load)
EPS 12V	17	67.7%
EPS 1U	26	70.7%
EPS 2U	13	70.1%
ERP12V	0	n/a
ERP2U	0	n/a
EPS	0	n/a
TPS	4	69.8%
MPS	4	65.0%
DPS	7	81.1%
All Power Supplies	71	69.6%

Figure 3 shows the average manufacturer-reported efficiencies by SSI category and how they compare to our laboratory test results at 100% loading. The manufacturer-reported data match quite closely with the power supply test results, likely indicating similarity in test method (at least at full load, which is the simplest way to test a power supply). The similarity also shows that manufacturers are not over-reporting the efficiency of their products. In fact, for all SSI Types except EPS 2U, the test results actually yielded slightly higher average efficiencies than the manufacturer-reported data. In the case of EPS 2U, the difference between the average manufacturer reported efficiency and tested efficiency was negligible. For all power supply categories, average manufacturer-reported efficiency at 100% load was 69.6% versus 72.5% from the tests results.

[12] 80 Plus is a program to encourage manufacturers to voluntary adopt more energy-efficient power supplies using utility rebates. See www.80plus.org for more details.

Figure 3. Manufacturer-Reported Efficiencies and Test Results at 100% Loading (by SSI Category)

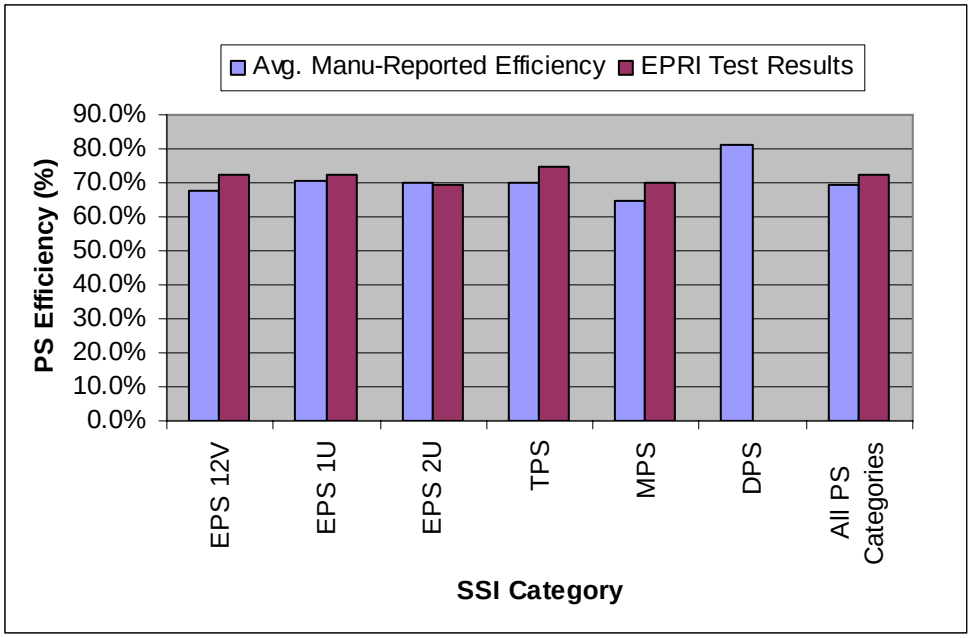

Note: EPRI Solutions did not test any DPS power supplies.

Nevertheless, as Figure 2 (Measured Power Supply Efficiency) makes clear, 100% loading is only one part of the story. Data centers virtually never operate server power supplies at 100% load, so their efficiency at that level is academically interesting, but has no relevance to total cost of ownership. It is critical to expand the consideration of server power supply efficiency to loads below 100% when studying overall energy use and opportunities for energy savings. The next section of this report (Section 2, Field Testing) reinforces this point by examining the typical loading conditions of operating servers.

3 Field Testing

This task's objective was to conduct field testing to determine achievable energy savings from more efficient power supplies in operating server installations. To do this we:

- Developed field-testing protocols.
- Conducted field tests at EPRI Solutions, Ecos Consulting, and LBNL facilities.

3.1 Field Testing Protocol

We developed a simple three-step test protocol/procedure for conducting *in situ* field tests to measure server power consumption.[13]

1. Before taking any measurements, some basic information (Table 5) was recorded about each server's function, technical specifications, power supply size and configuration:

[13] In some high-tech environments (e.g. datacenters) this protocol was not applicable and a customized approach was used to gather data directly from the equipment's log.

Table 5. Server and Power Supply Information

Server Information	Power Supply Information
Make	Number of power supplies
Model	Number of AC input cords
Application (e.g. print, file, e-mail)	Rating of each power supply (watts)
Processor	Power supply configuration (n, n+1)
Memory	Cross reference to SSI specification
Hard Drive	
Operating System	
Power Management enabled/disabled	
Other/Miscellaneous	

2. Each server was measured using a "WattsUp? Pro" power meter that recorded power consumption (in watts) for a continuous period of 15 to 144 hours at data intervals ranging from four seconds to 34 minutes.

3. The data collected from each server were then rank ordered and divided into nine percentiles (1%, 5%, 10%, 25%, 50%, 75%, 90%, 95%, and 99%). In this way, we could determine what level of power was being drawn by the server for what percent of the time.

3.2 EPRI, Ecos, and LBNL Field Tests

Initial field tests were conducted in EPRI Solutions and Ecos Consulting facilities located in Knoxville, Tennessee, and Portland, Oregon, respectively. Both EPRI Solutions and Ecos can be described as "small office" environments—approximately 50 people and 10,000 sq. feet each. The servers located at each of these facilities handle various routine business functions (e.g., print and file server, web server, Microsoft Exchange server, etc.), all of them operate on a 24/7 basis.

Field tests were also conducted at an LBNL data center located in Berkeley, California. The LBNL servers, although similar to some of the others tested in terms of power supply topology and configuration, are custom servers used primarily for distributed scientific computing. As such, these machines operate at close to 100% processor utilization all of the time and have a significantly different usage profile than the office servers tested at Ecos and EPRI Solutions.

3.2.1 Servers Tested

A total of 18 servers—six at EPRI Solutions, nine at Ecos, and three at LBNL—were measured using the field-testing protocol described above. Table 6 lists the servers that were tested and some information about their function and power supply configuration. Power management was not enabled on any of the servers tested.

Table 6. Servers Tested

Server Function	Processor Information	# of Power Supplies	# of AC Cords	Rating of Each Power Supply (W)	Redundancy Config	Cross Ref to SSI
Print& File Server	Single Intel Pentium 3 733MHz	3	3	320	N+2	ERP12V
Accounting Server	Dual XEON Processors @ 2GHz	2	1	500	N+1	ERP12V
Web & SQL Server	P3 ~ 700 MHz	1	1	330	N	EPS
Terminal Server	P3 933MHz	3	1	300	N+2	ERP12V
Exchange Server	Single Intel Pentium 3 - 933 MHz	2	2	330	N+1	ERP12V
PQ Remote Monitoring Server	Dual XEON (P3) 550MHz	3	1	275	N+1	ERP12V
Telephony Server	Intel PIII 700Mhz	2	2	300	N+1	ERP2U
Antivirus Server	Intel Dual PIII 745Mhz	1	1	150	N	EPS1U
Accounting/ADP Payroll Server	Intel Dual PIII 745Mhz	1	1	150	N	EPS1U
GoldMine Sync Server	Intel PIII 1.26Ghz	1	1	200	N	EPS1U
Print & File Server, Domain Controller	Intel Xeon 2Ghz	2	2	350	N+1	ERP2U
MS SQL Server	Intel Xeon 700Mhz	2	2	270	N+1	ERP2U
Web Application Server	Intel PIII 1.26Ghz	1	1	200	N	EPS1U
Mail Server	Intel P450Mhz	1	1	60	N	EPS12V
Mail Server	Intel P450Mhz	1	1	60	N	EPS12V
Scientific Computing	AMD Opteron 1.4 GHz	1	1	350	N	EPS
Scientific Computing	AMD Opteron 1.4 GHz	1	1	350	N	EPS
Scientific Computing	AMD Opteron 1.4 GHz	1	1	350	N	EPS

3.2.2 Server and Power Supply Loading Curves

A load duration curve for each server was constructed from the percentile data representing the percent of time the server power exceeded a given value in watts. We derived a single server load duration curve by averaging these results for the 18 different servers (Figure 4, blue line). Because of the relatively large difference in function, we also constructed separate single server load duration curves for the group of three LBNL servers and for the 15 EPRI Solutions and Ecos server. (Number of servers is shown in parentheses.) As Figure 4 shows, the LBNL servers consumed about 50 watts more, on average, than the EPRI and Ecos servers. As mentioned above, these machines are closer to a scientific mainframe than typical office servers, and this difference in intended use may explain the additional power consumption; however, it may be too early to draw any hard conclusions about this difference in load duration curves. A number of factors beyond the computers' application, including power supply rating, redundancy configuration, processor type, and server utilization could be responsible for these results.

Figure 4. Measured Load Duration of Servers

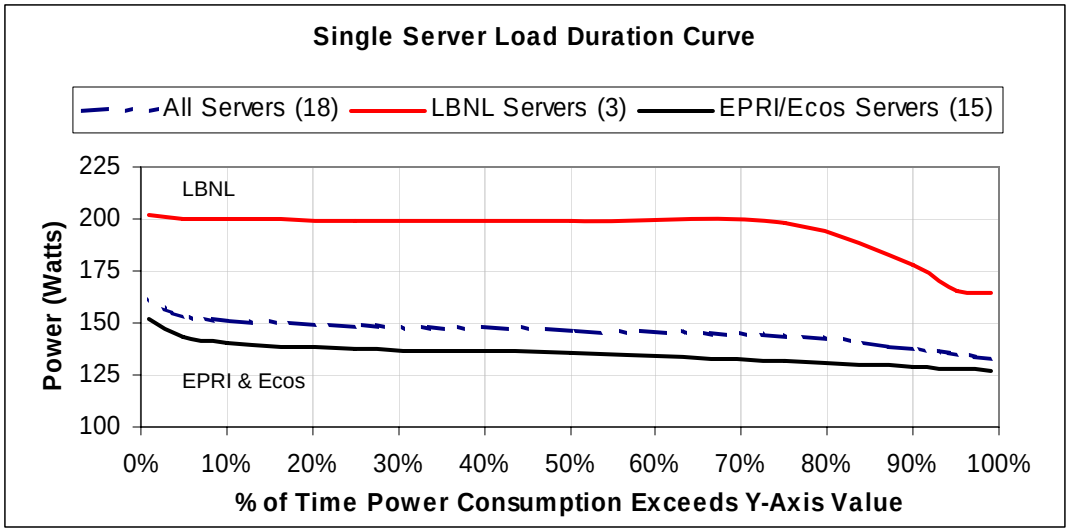

Source: EPRI/Ecos field measurements.

Figure 4 shows that the median (50th percentile) power consumption of a single server is 145.7 watts.[14] By annualizing the server power data, we can estimate total median server power consumption for a small office environment, like EPRI Solutions or Ecos, depending on the number of servers found. For example, Table 7 shows that, a five-server office would consume, on average, 6,384 kWh per year.[15]

Table 7. Estimated Annual Energy Consumption of Servers

# of Servers	Power Draw (W)	kWh per Year
1	145.7	1,277
2	291.5	2,553
3	437.2	3,830
4	583.0	5,107
5	728.7	6,384

[14] This is close to the 125W used to estimate server energy consumption in Roth, Kurt W., Fred Goldstein, and Jonathan Kleinman (Arthur D. Little). "Office Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings. Volume I: Energy Consumption Baseline." Prepared for US DOE, Building Technologies Program, January 2002. NTIS Number: PB2002-101438.

[15] While not included in this calculation, EPRI-Solutions found that the median power consumption for all the other network equipment (e.g., router, network interface card, etc.) needed to support their six-server IT infrastructure was about 150W, or 1,314 kWh per year—slightly more than adding an additional server.

Using the load duration data, the known power supply rating and configuration of each server, and an average of power supply efficiency, we were also able to generate a power supply loading curve showing the percent of time the server exceeded a given percentage of the power supply rating in watts, using the following formula:[16]

$$\%\ of\ Rated\ PS\ Watts = \frac{Watts * 70\%}{PSRating * \#ofPS}.$$

Figure 5 shows the average power supply loading for servers with different redundancy configurations. From our small sample it appears that server power supplies for small office applications are normally loaded in the 15%-35% range. This highlights the importance of light load power supply efficiency, or the need to better match power supplies with their load. Light load efficiency is especially important for redundant systems. The median power supply loading for non-redundant servers (N configuration) is just over 30%; whereas, for redundant power supply configurations—N+1 and N+2—the median loading was only 18.7% and 15.3%, respectively.

Figure 5. Power Supply Loading and Redundancy

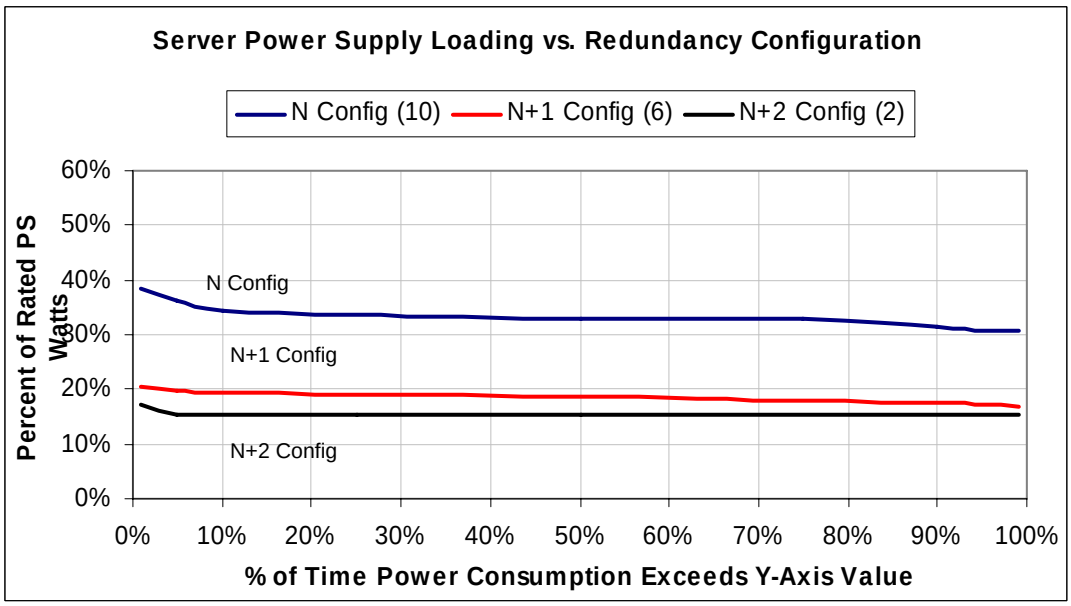

Source: EPRI/Ecos field measurements.

Figure 5 also shows that the loading curves are relatively flat, indicating not only that the server is drawing power or operating 100% of the time, but that there is little variation in the percent loading of its power supply during that time.

To investigate whether server function played any role in power supply loading, we compared the three LBNL servers to seven EPRI/Ecos servers with the same power supply redundancy configuration (N). Figure 6 shows that the LBNL servers were at the upper end, but were not substantially different from the other servers in terms of power supply loading. This suggests that even highly utilized machines like the LBNL servers/mainframes do not fully utilize the capacity of their power supplies. It also suggests that server power supplies are oversized for the actual requirements of the machines in which they are being used. Most servers measured were operating somewhere between about 20% and 50% of their rated load all of the time, which is the exact portion

[16] This formula multiplies the watts recorded by the power meter by the power supply efficiency, 70% or 0.7, to account for 30% energy loss, which is "consumed" by the power supply. It is only an approximation, as power supply efficiency varies with load.

of power supply curves at which efficiency begins to dramatically decrease (see the previous section for more on this topic). None of the servers tested ever exceeded 50% of its rated output.

Figure 6. Loading of Server Power Supplies

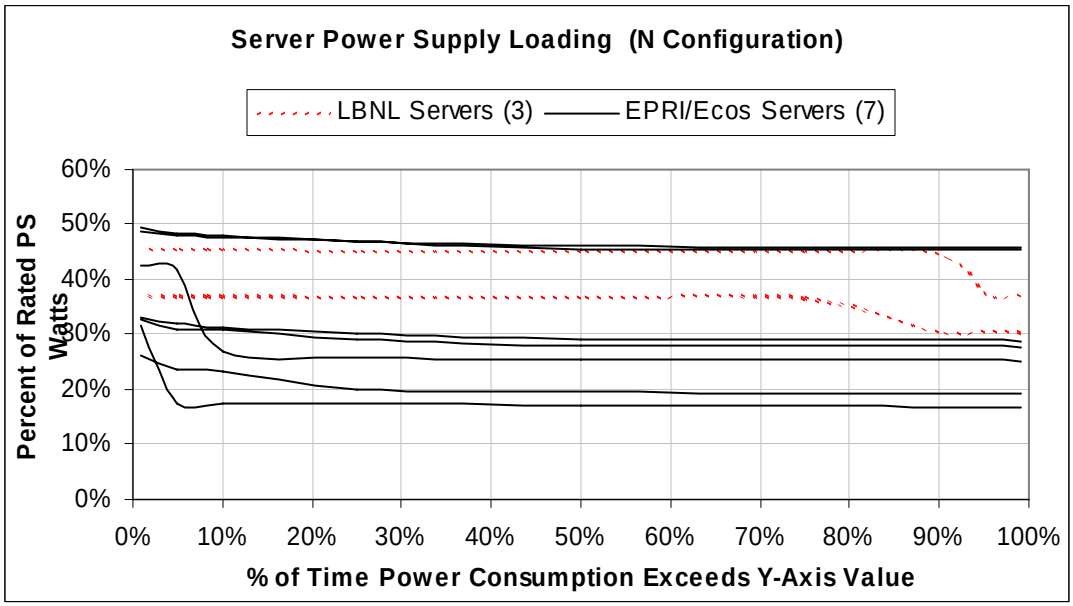

Source: EPRI/Ecos field measurements.

The combination of low efficiencies in, and oversizing of server power supplies can have dramatic effects on the net energy consumption of those power supplies. As Figure 6 illustrates, a dual Xeon processor-based server might typically require 100 to 300 watts DC from its power supply. An oversized, inefficient power supply would often waste two to three times as much net AC power[17] to meet that load as a properly sized, efficient one.

3.2.3 CPU Utilization

In addition to measuring power consumption, we also extracted CPU utilization data from each server's log for a 24-hour period.[18] CPU utilization describes how busy a server's processor or processors are; it tells us the percent of the CPU's processing capacity, or throughput, that is being used to perform tasks.

The utilization data were divided into percentiles, using the same methodology as described for the power data. A CPU utilization curve for each server[19] was constructed representing the percent of time the server's CPU utilization exceeded a given percent.

For comparison, Figure 7 shows the CPU utilization of the LBNL servers (data center/scientific computing) was very different from the Ecos servers (office environment). Average CPU utilization is a remarkably low 5 to 6%.[20] In fact, more than 95%% of the time, these servers operate at less than 30% utilization. In contrast, the LBNL servers spend a majority of their time operating at, or near, full capacity. It is difficult to draw any conclusions from such a small sample, but it appears that for the

[17] Here net AC power is defined as the difference between the AC power input of the power supply and its DC power output.
[18] Note: the time period for which CPU utilization data was collected was not necessarily the same as the time period for which power consumption was measured.
[19] CPU utilization data was collected for 7 Ecos servers, 5 EPRI Solutions servers, and three LBNL servers; three of the Ecos servers were not included.
[20] This median number is somewhat misleading. It could be (and is likely) the case that an organization with multiple servers relies heavily on one or two servers which have very high utilization, but the rest are performing secondary or backup functions and are idle much of the time. By averaging all servers together, we have masked this phenomenon.

most part power supply loading (which was relatively flat, see Figures 5 and 6) has very low correlation with "server activity."

Figure 7. Server Utilization

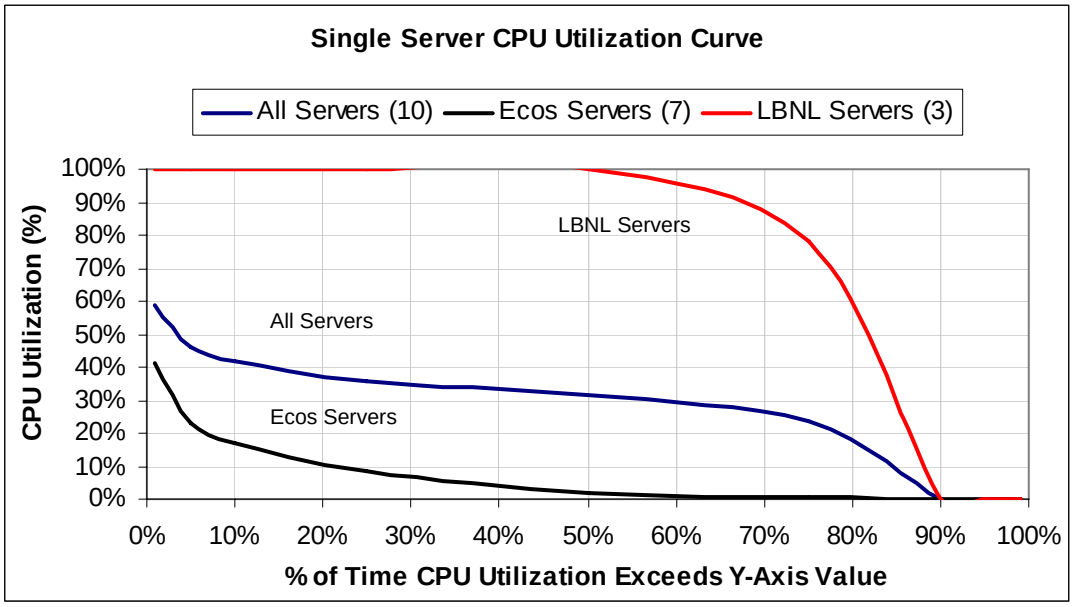

Source: EPRI/Ecos field measurements.

The servers we measured do not exhibit any activity-based power management or demand based switching (DBS), as Intel calls it, whereby power consumption is reduced when the server is operating below full capacity. Activity-based power management attempts to scale processor power consumption to the demand for processing resources placed on the CPU. The absence of this technology in the servers that we measured might help to explain the lack of any correlation between the utilization curve of the CPU and the load duration curves of the power supplies. Figure 8 provides a side-by-side comparison of processor utilization and power consumption in a server, demonstrating the poor correlation between power use and processor use.

Figure 8. AC Power Input vs. % CPU Time

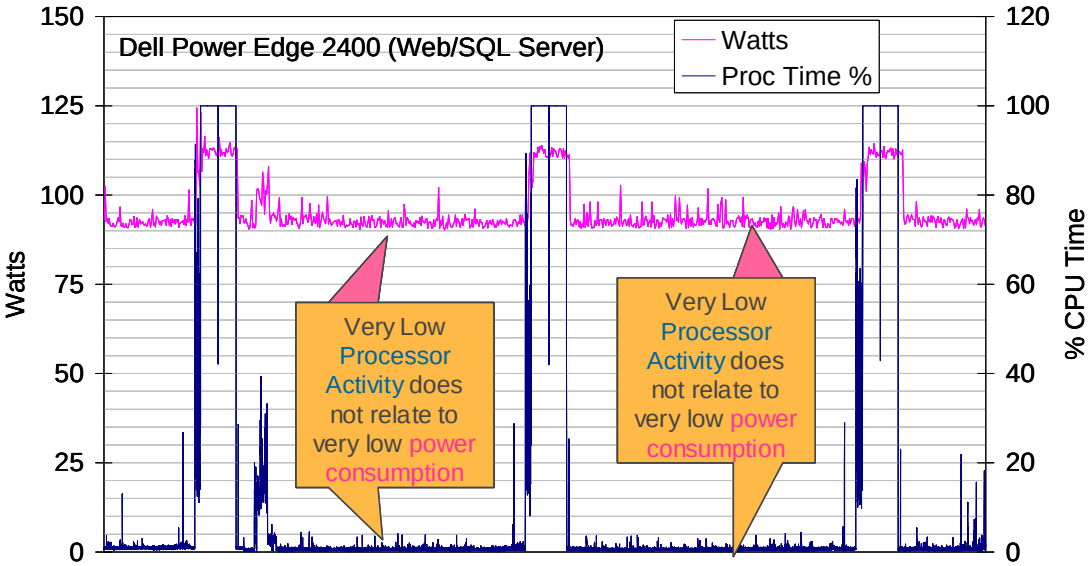

Source: EPRI field measurements.

4 Proposed Power Supply Efficiency Specifications

This task's objective was to recommend new efficiency levels to SSI for consideration in its future specifications. To do this we:

- Coordinated with SSI on recommended new efficiency levels in power supply guidelines.
- Performed market research to assess the market penetration of different server configuration and power supply topologies.
- Estimated the energy savings potential by combining market research with the proposed efficiency levels.

4.1 Coordination with SSI

Based on the efficiency levels documented in existing equipment and the field studies conducted, we closely coordinated with SSI and Intel[21] to propose changes and to the SSI specifications based on the *Proposed Test Protocol.*

At the Intel Developer Forum in March 2005, in a joint presentation, Intel announced new proposed SSI Industry specifications to encourage more efficient power supplies. The new specifications include testing conditions which align with ATX12V, as well as required and recommended efficiency levels at 20%, 50%, and 100% loading. SSI is also now supporting a power supply technology known as Power Supply Management Interface (PSMI). This new industry standard provides a basic internal mechanism for server power supplies to report power consumption and efficiency data directly to the server's motherboard with 5% to 10% accuracy. The technology could conceivably be used for a number of purposes. Power and efficiency data could be logged on the server's hard drive or reported to HVAC equipment so that fans and cooling equipment could scale their output to the heat output of IT equipment.

Table 8 provides a comparison of the old and new SSI specifications with regards to efficiency requirements. Establishing high minimum recommended levels at 20% and 50% loading is particularly critical for improving energy efficiency, since field-testing (Section 2) showed that servers operate in this range most of the time.

[21] Brian Griffith, Intel Power Server Architect, EPG (SSI Coordinator).

Table 8. Comparison of Old and New SSI Specifications

SSI Category		Old Specification	New Specification		
EPS12V		**Rev 2.1**	**Rev 2.8**		
	Loading	**100%**	**100%**	**50%**	**20%**
	Required Efficiency		70%	72%	65%
	Recommended Efficiency	68-72%	77%	80%	75%
ERP12V		**Rev 1.0**	**Draft Rev 1.4**		
	Loading	**100%**	**100%**	**50%**	**20%**
	Required Efficiency		70%	72%	65%
	Recommended Efficiency	68-70%	77%	80%	75%
EPS1U		**Rev 2.1**	**Rev 2.9**		
	Loading	**100%**	**100%**	**50%**	**20%**
	Required Efficiency		70%	72%	65%
	Recommended Efficiency	65-75%	80%	83%	78%
EPS2U		**Rev 2.1**			
	Loading	**100%**	**100%**	**50%**	**20%**
	Required Efficiency				
	Recommended Efficiency	68-72%			
ERP2U		**Rev 2.0**	**Rev 2.2**		
	Loading	**100%**	**100%**	**50%**	**20%**
	Required Efficiency		70%	72%	65%
	Recommended Efficiency	70-82%)	80%	83%	78%

Note: Recommended minimum efficiency in old SSI specification depended upon PS rating (watts).

4.2 Market Penetration of Servers and Power Supplies

We performed market research, through literature review and other industry sources, to assess the market penetration of different server configurations and their power supply types.

This research was performed in two stages:

1. First, we gathered market data on servers to determine which types of configurations and manufacturers had the largest market penetration.
2. Next, we gathered information on the power supplies contained in those units.
3. With respect to California-specific data, this information was not available on any consistent basis, as industry data are not reported for any one state. Thus, we have not attempted to breakout any CA-specific estimates of product shipment or server population. We have incorporated approximations for California where specific estimates may make sense, such as energy savings projection.

4.2.1 Servers

The server market can be sliced many different ways (e.g., by shipments, revenue, geographic region, class, manufacturer, price range, operating system, processor, U-rating,[22] etc.) IDC's Worldwide Quarterly Server Tracker and Forecast tools are probably the most comprehensive data source for statistics on server shipments. The Tracker provides actual quarterly data for more than 15

[22] The "U" rating is based on the height of the server: A 1U server is 1.75-in. tall.

different data categories (including those mentioned above). The Forecast provides fewer data categories, but projects future unit shipments and revenues.[23]

Such data can be purchased from IDC at a substantial cost. However, we felt we could obtain a fairly accurate picture of the U.S. server market relying on limited data from IDC and various other quantitative and anecdotal sources all of which seemed to corroborate.

Worldwide Server Shipments

IDC reported actual 2003 unit shipments of servers of 5.3 million (Table 9), almost a 19% increase over 2002 shipments of 4.4 million.[24]

Table 9. Worldwide Server Market, 1999-2003

Worldwide Server Factory Revenue ($M) and Units Shipped					
	2003	2002	2001	2000	1999
	$M	$M	$M	$M	$M
WW Total	$46,131	$44,649	$50,496	$61,675	$57,708
Units	5,281,231	4,442,690	4,276,119	4,369,840	3,761,141

Source: IDC, 2004 Release as shown in *The Business for Storage Networks. Chapter 1, Industry Landscape: Storage Costs and Consumption*. Cisco Systems, Sept 3, 2004. Available at: http://www.cisco.com/application/pdf/en/us/guest/netsol/ns516/c1272/cdccont_0900aecd80257124.pdf

IDC expects worldwide server shipments to grow steadily through 2008 (Figure 9).

Figure 9. Historical and Forecast Server Shipments, 2000-2008

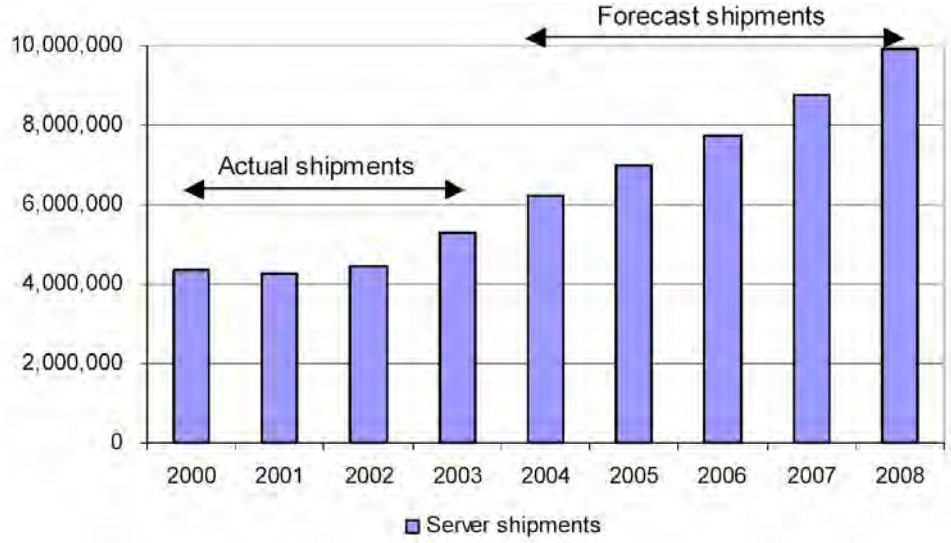

Source: IDC, 2004 as shown in Humpreys, John and Jessica Yang, IDC White Paper sponsored by Rackable Systems. Server Innovations: Examining DC Power as an Alternative for Increasing Data Center Efficiency and Reliability, August 2004

[23] http://www.idc.com./getdoc.jsp?containerId=IDC_P348 and http://www.idc.com./getdoc.jsp?containerId=IDC_P5514.
[24] Gartner Group, another well-know data source, reports similar 2003 shipments of 5.4 million.

U.S. Server Shipments

The United States represents about 41% of the worldwide market in terms of shipments (Table 10). On a population basis, the State of California would account for about 10% to 15% of the overall US shipment of products.

Table 10. US Percentage of Worldwide Server Shipments

Year	Worldwide Shipments (Units)	US Shipments (Units)	US %
2000	4,327,511	1,882,184	43%
2001	4,425,977	1,712,614	39%
2002	4,610,328	1,949,361	42%
2003	5,469,016	2,256,918	41%
2004	6,700,000	2,747,000	41% est.
Average			**41%**

Source: Gartner Group/Dataqwest, with some interpolation by Ecos.
2004 sales from http://uk.news.yahoo.com/050225/221/fd8j9.html

The dominant manufacturers in the U.S. market are Hewlett-Packard, Dell, and IBM, and to a lesser extent Sun Microsystems and Gateway. These five companies represent over 70% of the U.S. market in terms of units shipped (Table 11).

Table 11. U.S. Server Shipments (Units), 2001-2002

Company	2002 Shipments	2002 Market Share (%)	2001 Shipments	2001 Market Share (%)	Growth (%)
Hewlett-Packard	506,589	26	483,938	28.3	4.7
Dell	487,984	25	399,236	23.3	22.2
IBM	225,315	11.6	217,171	12.7	3.8
Sun Microsystems	143,753	7.4	124,056	7.2	15.9
Gateway	18,000	0.9	25,803	1.5	-30.2
Others	567,720	29.1	462,410	27	22.8
Total	1,949,361	100	1,712,614	100	13.8

Note: The data for Hewlett-Packard and Compaq have been combined.
Source: Gartner Dataquest (January 2003)

Volume Servers

Many top manufacturers, such as Dell, emphasize volume or low-end servers. This emphasis parallels recent server sales trends, which reveal that severs priced under $25,000 are the market's bright spot.[25] Lower-end units are becoming ever more capable. "An entry-level box now can do what it took a four-way to do," according to Meta Group analyst Carl Greiner.[26] Manufacturers like these servers not only because their low price point helps boost volumes of sales with small to medium sized business, but also because these servers are increasingly turning up in more high-tech environments, like data centers.

A number of manufacturers are relying on the entry-level server market to drive their overall server business and boost volume sales. For example, in September 2004, Sun Microsystems indicated that the company is renewing its commitment to the low-end server market and increasing its “portfolio of

[25] http://www.newsfactor.com/story.xhtml?story_id=21934
[26] http://www.newsfactor.com/story.xhtml?story_id=21934

offerings in the x86 segment."[27] Volume or low-end servers are thought to be the primary driver of market growth, as indicated by the following market reports:[28]

- In 2002, Gartner Group reported that the U.S. server market experienced a strong "resurgence," with server shipments increasing by 13.8 percent to a total of 1.9 million units. Much of this growth was attributed to low-end server sales.[29]
- In the third quarter of 2003, Vernon Turner, group vice president of IDC's Worldwide Server Group indicated, "Volume servers [those priced less than $25,000] are generating most of the positive momentum in the worldwide server market. This shows that the IT community has embraced volume server deployments as a mainstream technology to meet a wide range of data-processing requirements and to support a wide variety of computing workloads."[30]
- IDC data for the second quarter of 2004 indicated 22.7% year-over-year unit shipment growth for servers, reflecting "strong unit growth in the volume server segment."[31]

Volume or low-end servers also represent over 90% of the market.

- Roth et al. (2002) showed that from 1998 to 2000 low-end servers represented between 88% and 91% of the U.S. market in terms of units shipped (Table 12).

Table 12. U.S. Server Shipments by Class, 1998-2000

Year	Low-end	Work-horse	Mid-range	High-end	Total	Low-end%
1998	1,082,180	104,776	37,813	2,852	1,227,621	88%
1999	1,367,839	119,641	40,340	2,663	1,530,483	89%
2000	1,615,126	121,097	41,314	2,510	1,780,047	91%

Source: ADL/US DOE 2002.

- Gartner Group reported that in the last quarter of 2003, inexpensive dual processor servers containing 32-bit processors from Intel or Advanced Micro Devices accounted for 1.45 million or 91% or of the 1.59 million servers shipped worldwide.[32] IDC data for the same quarter matched closely, reporting that unit shipments of basic x86 servers grew at 23 percent to nearly 1.4 million servers worldwide, with factory revenues growing at 15 percent to $5.5 billion.[33]
- IDC also indicated that single and dual processor capacity servers represented 91% of all server unit shipments in 2003. IDC expects two-processor capacity server unit shipments to account for more than 6.2 million units in 2008, nearly double the 3.5 millions units shipped in 2003.[34]
- In 2004, x86 servers dominated shipments, with 91%, and at current growth rates, the X86 server business will account for over 50% of all server (dollar) sales in 2005.[35]

Rack-Optimized vs. Pedestal Servers

Presently, there is a fairly even split between rack-optimized and non-rack optimized (pedestal) servers, with blade servers accounting for a small fraction of the market (Figure 10). In the future, the number of pedestal servers is expected to remain stable, while the number of rack-optimized, and especially blade servers, grows.[36]

---

[27] http://www.expresscomputeronline.com/20040906/coverstory01.shtml
[28] http://nwc.serverpipeline.com/midrange/46200036
[29] http://www.serverwatch.com/news/article.php/1575051
[30] http://www.prnewswire.com/cgi-bin/stories.pl?ACCT=105&STORY=/www/story/11-26-2003/0002065610
[31] http://www.tekrati.com/T2/Analyst_Research/ResearchAnnouncementsDetails.asp?Newsid=3460
[32] http://news.com.com/2100-1010_3-5149716.html?tag=nefd_top
[33] Bekker, Scott. "IDC: Volume Servers Shaking Up Server Market." ENTNews, March 3, 2004. http://www.entmag.com/news/article.asp?EditorialsID=6151
[34] http://www.internetnews.com/ent-news/print.php/3389211
[35] http://uk.news.yahoo.com/050225/221/fd8j9.html (Yahoo News, UK, Feb 25, 2005)
[36] Note: Blade servers were excluded from our study, as these servers are still relatively new and **currently** represent only a small portion of the market. Nevertheless, blade server technology should be closely watched as it represents an emerging trend in IT and datacenter applications. (See Blade Servers.)

Figure 10. Historical and Forecast Adoption of Pedestal, Rack-Optimized and Blade Servers, 2000-2008

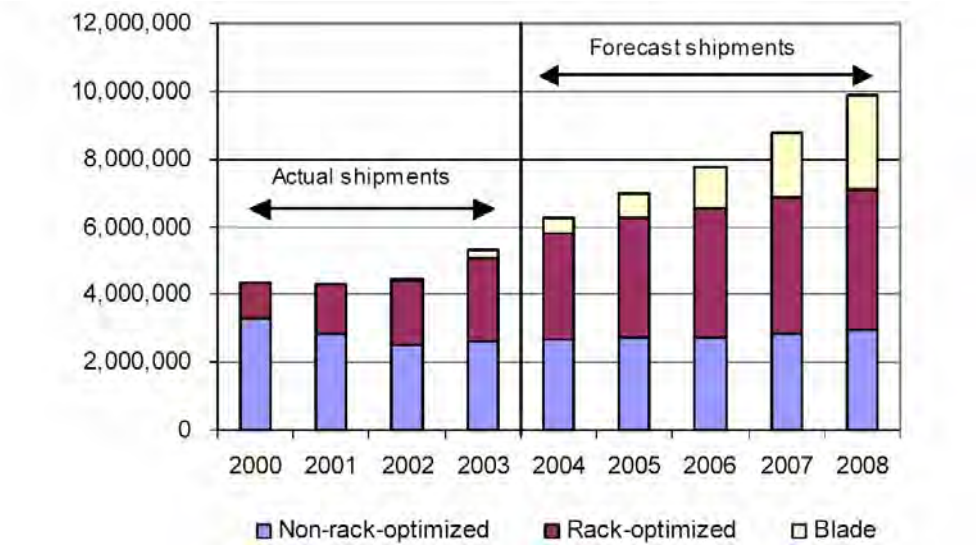

Source: IDC, 2004 as shown in Humpreys, John and Jessica Yang, IDC White Paper sponsored by Rackable Systems. Server Innovations: Examining DC Power as an Alternative for Increasing Data Center Efficiency and Reliability, August 2004

IDC tracks the U-rating[37] of servers. Table 13 shows a break down of the servers shipped in the U.S. in the forth quarter of 2001 by U-rating. 2U, 1U, and 0U (tower/pedestal) servers each account for a little over 20% of the market.

Table 13. U.S. Server Shipments by U-Rating, Q4 2001

Server U-rating	Units (Q401)	%
2	213,558	23.05%
1	209,882	22.65%
0	199,938	21.58%
5	138,040	14.90%
4	79,066	8.53%
7	27,760	3.00%
6	9,610	1.04%
All Others	48,598	5.25%
Total	**926,452**	**100.00%**

Source: IDC, Worldwide Quarterly Server Tracker, demo – additional data can be purchased from IDC.

Historically, large “mainframe” computers dominated the datacenter. Now, even though some high-end models like IBM's zSeries and the HP Superdome are selling quite well, there is a trend toward using smaller servers, including blades.[38] At a recent APC event in Portland, Oregon, one presenter indicated that 5 years ago, the average server in a data center was a 5U, 1 processor machine, whereas, today, the average server is 1.5U, dual processor with at least 2 gigabytes of memory.[39]

[37] The "U" rating is based on the height of the server: A 1U server is 1.75-in. tall.
[38] http://www.cioupdate.com/trends/article.php/3289721
[39] APC Presentation, Multnomah Athletic Club, Portland, OR, July 13, 2004.

Blade Servers

Many analysts believe that blade servers are poised to transform the server industry. According to an article entitled “The Data Center of the Future” by Drew Robb:

> Blade servers offer companies low cost scalability, since it is easy to assign a batch of these servers to a particular application, rather than having to buy a more expensive server which then remains underutilized. Since blade servers stuff a dozen or more servers into a single box, they drastically cut the infrastructure costs for racks, cabling and cooling. Then there is the ease of support. When one goes down it is a simple act to swap out a server card and let the system automatically rebuild.[40]

Roughly 185,000 blade servers were sold in 2003.[41] However, IDC predicts blade servers will account for one out of every four servers sold by 2007.[42]

In fact, IDC forecasts that blade server sales will reach over $1 billion in 2004 on the strength of Dell's new blade offerings and interest from small and medium-sized businesses.[43] The market is fluid, and many manufacturers are revamping their blade product lines.

Of the large vendors, HP was the first to offer (Proliant) blade servers. However, second-quarter 2004 numbers from IDC, reveal that HP is now in second place, shipping 32% of units, while IBM leads the market with 44% of unit shipments.[44] IBM plans to expand its blade offerings and began shipping a new 7U, 14 blade BladeCenter Express Chassis at the end of November 2004. Dell also recently moved forward with its most significant modular server—the PowerEdge 1855, which supports up to 10 servers in a 7U chassis and is geared for companies that rely on large Web farms and high-performance computing clusters.[45]

While Bruce Kornfeld, Director of Enterprise Marketing at Dell, believes that customers are looking at blades for “great server technology, better density and savings in power and cooling to ease management,” many are still “not willing to pay a premium for blade servers over traditional 1U and 2U rack servers."[46] The implications of blade server growth for the power supply market are significant. It will become increasingly common in the future to see one or two high wattage (1000 watts or more) blade server power supplies powering up to six blades, rather than a larger number of smaller wattage power supplies powering individual servers.[47]

4.2.2 Power Supplies

While market data on server shipments are readily available, data on the power supplies contained within those servers are more difficult to find. We sampled some popular server models to get an idea of the size, shape, and configuration of the power supplies they contain. Table 14 provides the power supply information listed for 17 different popular new product offerings.

---

[40] http://www.cioupdate.com/trends/article.php/3289721
[41] http://www.itfacts.biz/index.php?id=P1517
[42] http://www.serverwatch.com/news/article.php/3411471
[43] http://www.serverwatch.com/news/article.php/3440321
[44] http://www.serverwatch.com/news/article.php/3411471
[45] http://www.serverwatch.com/news/article.php/3436331
[46] http://www.serverwatch.com/news/article.php/3436331
[47] Jeremiah P. Bryant, “AC-DC Power Supply Growth Variation in China and North America,” presented at the Applied Power Electronics Conference (APEC) 2005, Darnell Group, March 6-10, 2005.

Table 14. Power Supply Specifications in Selected Popular Server Models

Make	Sever Model	Form Factor	Power Supply Description
Dell	Poweredge 750	5U Tower or Rack	650W non-redundant or 675W hot-plug redundant power 110/220 Volts
Dell	Poweredge 1750	1U	Optional, hot plug redundant 325 W power supplies, 110/220 Volts, Optional -48V DC power supplies
Dell	Poweredge 1800	5U Tower or Rack	650W non-redundant or 675W hot-plug redundant power 110/220 Volts
Dell	Poweredge 1850	1U	550W, optional hot-plug redundant power, 110/220 Volts
Dell	Poweredge 2800	5U Tower or Rack	930W, optional hot-plug redundant power, 110/220 Volts
Dell	Poweredge 2850	2U	700W, optional hot-plug redundant power, 110/220 Volts
HP	Proliant DL360G4	1U	Optional hot plug redundant power supply (460 W)
HP	Proliant DL380G4	2U	575-Watt CE Mark Compliant Optional Hot Plug AC Redundant Power Supply
HP	Proliant DL585	2U	Optional hot plug redundant power supply (800 W)
HP	Proliant ML110	5U Tower or Rack	350W power supply
HP	Proliant ML350G4	5U Tower or Rack	Most models: 725W, optional hot-plug redundant, NHP SCSI models: 460W, non hot plug supply
IBM	eServer 326	1U	411W 1 std/1 max
IBM	Xseries 336	1U	Hot-swappable 585W power supplies, 1 std/2 max
IBM	Xseries 346	2U	Hot-swappable 625W power supplies, 1 std/2 max
IBM	Xseries 306	1U	300W 1 std/1 max, 110 or 220 volt universal auto sensing
Sun	SunFire V210	1U	320W power supply
Sun	SunFire V240	2U	One required, two for redundancy (hot-swappable) with separate power cords (400W)

Source: PC Magazine and technical specifications available at www.dell.com, www.ibm.com, www.hp.com.

This information tracked closely with the power supply types and configurations that we observed in our field testing (See Sections 2.2 and 2.3), and reinforced the decision—made in consultation with Intel, SSI members and other industry partners—to focus this study on multi-output front-end AC-DC power supplies commonly found in 1U, 2U, and pedestal servers (non-redundant PS).

4.3 Energy Savings Potential of Proposed Efficiency Levels

We estimated the energy savings potential of more efficient server power supplies in three steps:

1. We estimated the current AEC (Annual Energy Consumption) of servers in the United States using the basic methodology developed in the ADL/US DOE report “Office Energy

Consumption by Office and Telecommunications Equipment in Commercial Buildings. Volume I: Energy Consumption Baseline."[48]

2. We determined the percentage of current server energy use that could be saved due to more efficient power supplies, based on the new SSI recommended specifications.
3. By applying the percentages established in #2 to the AEC estimates established in #1, we were able to estimate the overall energy savings potential of the new recommended efficiency levels.

For California-specific estimates, we used a range of 10% to 15% of estimated US results.

4.3.1 Annual Energy Consumption of Servers

Starting from the basic methodology established in the ADL/US DOE study, we constructed a revised estimate of the annual energy consumption (AEC) of servers in the United States.

To construct the revised 2004 EPRI/Ecos estimate, we:

- Updated the server stock numbers by adding four additional years of market data (2001-2004) on numbers of units shipped.

Table 15. Estimated Server Units

Server Categories	2000 ADL/US DOE Stock (Units Millions)	2004 EPRI/Ecos Stock (Units Millions)	% Change
Low-end	4,065,145	6,587,061	62%
Work-horse	577,960	506,470	-12%
Mid-range	185,195	151,678	-18%
High-end	16,549	14,730	-11%
Total	**4,844,849**	**7,259,939**	**50%**

- Revised the power draw for low-end and workhorse servers based on a sample of current data gathered from manufacturers, as well as information gathered during laboratory and field tests about typical power supply efficiency and server loading.

Table 16. Estimated PS Characteristics

Server Categories	2000 ADL/US DOE		2004 EPRI/Ecos	
	Max Power Draw (Watts)	Avg Power Draw (Watts)	Max Power Draw (Watts)	Avg Power Draw (Watts)
Low-end	250	125	450	193
Work-horse	1,300	650	800	343
Mid-range	2,450	1,225	2,450	1,225
High-end	5,040	2,520	5,040	2,520

Notes: For 2000 ADL/US DOE average power draw was calculated as 50% of maximum. For 2004 EPRI/Ecos, maximum power draw for Low-end and Workhorse was revised based on sample of manufacturer data. Average power draw was also revised to reflect 30% server loading, and 70% power supply efficiency.

Figure 11 shows our revised estimate of server AEC for 2004. The revised AEC is 14.6 TWh, which represents a 45% increase over the ADL/US DOE estimate of 10.1 TWh. This increase is attributable to growth in both the number and average power draw of low-end servers, which comprise the majority of units. The estimate for California is in the range of 1.5 TWh to 2.2 TWh.

[48] Roth, Kurt W., Fred Goldstein, and Jonathan Kleinman (Arthur D. Little). "Office Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings. Volume I: Energy Consumption Baseline." Prepared for US DOE, Building Technologies Program, January 2002. NTIS Number: PB2002-101438.

Figure 11. Server Annual Energy Consumption (AEC)

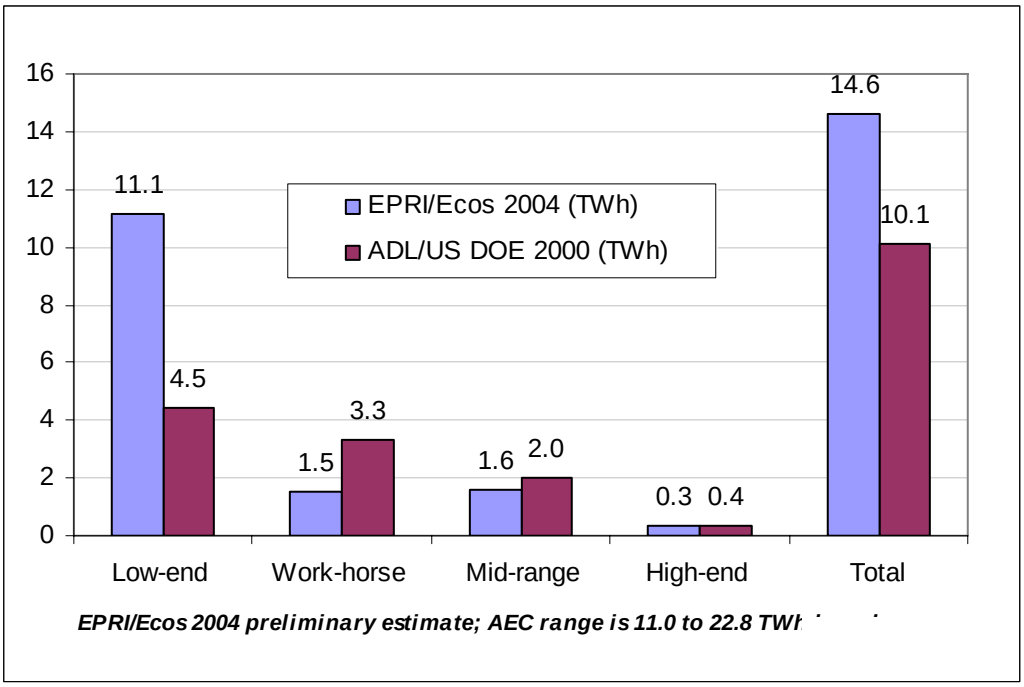

4.3.2 Existing vs. Proposed Power Supply Efficiency Levels

Table 17 compares the existing and recommended efficiencies at different loading points for the EPS 12V and EPS 1U power supplies. The SSI Recommended (2005) efficiency levels, recently endorsed by ENERGY STAR in its first draft specification revision for servers, yield significant efficiency increases, depending upon loading.[49] Likewise, a utility-funded efficiency program currently being operated by Ecos Consulting for desktop and server power supplies, 80 Plus, improves existing low end server power supply efficiencies by at least 14 to 31% at the 20% load condition where server power supplies often operate.[50]

Table 17. Comparison of Existing and Proposed Efficiency

Specification		20% Load	50% Load	100% Load
EPS12V (pedestal)	Existing Server	70%	76%	73%
	SSI Required (2005 proposed)	65%	72%	70%
	SSI Recommended (2005 proposed)	75%	80%	77%
	80 Plus	80%	80%	80%
EPS1U (rack)	Existing Server	61%	71%	72%
	SSI Required (2005 proposed)	65%	72%	70%
	SSI Recommended (2005 proposed)	78%	83%	80%
	80 Plus	80%	80%	80%

[49] See http://www.energystar.gov/index.cfm?c=revisions.computer_spec.
[50] See www.80plus.org for additional information.

4.3.3 Energy Savings Estimates

Table 18 provides an example to demonstrate the energy use of single EPS1U server power supply under various specification scenarios. As the table shows, implementing the SSI Recommended (2005) specifications would result in more than a 50% decrease in power supply electricity use. (Note that this estimate in energy use reduction is based on both utilization data as well as efficiency data showing that power supplies typically spend most of their time at lower load levels, where they are less efficient).

Table 18. Power Supply Energy Use for Single 425W EPS 1U Server

				Power Supply Electricity Use per year	
Specification	**20% Load**	**50% Load**	**100% Load**	**kWh**	**$**
Typical Server	61%	71%	72%	556	$56
SSI Required (2005 proposed)	65%	72%	70%	521	$52
SSI Recommended (2005 proposed)	78%	83%	80%	274	$27
80 Plus	80%	80%	80%	216	$22

Note: Assumes loading in the 20-50% range for a 425W EPS1U power supply and $0.10 per kWh.

Power supply efficiency measures the ratio of output power to input power. Thus, a 70% efficient power supply has energy losses of 30% and must draw more power than actually needed to run a server. Improving power supply efficiency reduces energy losses or the energy "consumed" by the power supply itself. *For example, replacing a 70% efficient power supply with an 85% efficient power supply cuts losses in half (from 30% to 15%)*.

Table 19 shows the percent energy savings for (a) the power supply and (b) the whole server for different efficiency gains.

Table 19.

(a) Reduction in Net Power Supply Energy Use

		New Power Supply Efficiency							
		60%	**65%**	**70%**	**75%**	**78%**	**80%**	**83%**	**85%**
Existing Power Supply Efficiency	**60%**	0%	19%	36%	50%	58%	63%	69%	74%
	65%	-24%	0%	20%	38%	48%	54%	62%	67%
	70%	-56%	-26%	0%	22%	34%	42%	52%	59%
	75%	-100%	-62%	-29%	0%	15%	25%	39%	47%
	78%	-136%	-91%	-52%	-18%	0%	11%	27%	37%
	80%	-167%	-115%	-71%	-33%	-13%	0%	18%	29%
	83%	-225%	-163%	-109%	-63%	-38%	-22%	0%	14%
	85%	-278%	-205%	-143%	-89%	-60%	-42%	-16%	0%

(b) Reduction in Server Energy Use

		New Power Supply Efficiency							
		60%	**65%**	**70%**	**75%**	**78%**	**80%**	**83%**	**85%**
Existing Power Supply Efficiency	**60%**	0%	8%	14%	20%	23%	25%	28%	29%
	65%	-8%	0%	7%	13%	17%	19%	22%	24%
	70%	-17%	-8%	0%	7%	10%	13%	16%	18%
	75%	-25%	-15%	-7%	0%	4%	6%	10%	12%
	78%	-30%	-20%	-11%	-4%	0%	2%	6%	8%
	80%	-33%	-23%	-14%	-7%	-3%	0%	4%	6%
	83%	-38%	-28%	-19%	-11%	-6%	-4%	0%	2%
	85%	-42%	-31%	-21%	-13%	-9%	-6%	-2%	0%

Assuming a baseline efficiency of 70%, Tables 19 (a) and (b) indicate a 34% reduction in power supply electricity use, and a 10% reduction in overall server energy use. Our comparison of existing and proposed efficiency (Table 20) yields, on average, a power supply efficiency increase of 8%.[51]

Based on our overall estimate of server AEC (Section 3.3.1), we estimate the amount of electricity saved due to more efficient power supplies would be approximately 1.5TWh. Figure 12 also shows a high-efficiency case, where power supply efficiency is 83% and electricity savings of 2.3 TWh are realized. An estimate for the potential savings for California is between 10% and 15% of the overall US savings potential, or between 230 GWh and 345 GWh.

Improved power supply efficiency will also yield compounded savings from reduced air conditioning loads and UPS losses not included here. We also note that the potential for the infrastructure load reduction (i.e. improved overall efficiency and reduced cooling load) is on the same order of magnitude as the power supply efficiency improvement.

Table 20. Energy Savings Potentials

Server Category	Number of Servers (Millions)	Annual Energy Consumption in TWh @ 70% PS Efficiency	Annual Energy Consumption in TWh @ 78% PS Efficiency	Annual Energy Consumption in TWh @ 83% PS Efficiency
Low-end	6,587,061	11.1	10.0	9.4
Work-horse	506,470	1.5	1.4	1.8
Mid-range	151,678	1.6	1.5	1.4
High-end	14,730	0.3	0.3	0.4
Total	**7,259,939**	**14.6**	**13.1**	**12.3**

[51] The 8% efficiency gain was derived by assuming that servers operate, on average, at 30% loading (i.e., 20% two-thirds of time and 50% one/third of time) and then calculating a weighted average of efficiency increases at 20% and 50% loading for EPS12V and EPS1U. This also implicitly assumes an even split between pedestal and rack servers, which appears consistent for server sales in 2004 and earlier.

Figure 12. Energy Savings Potential of PS in Servers

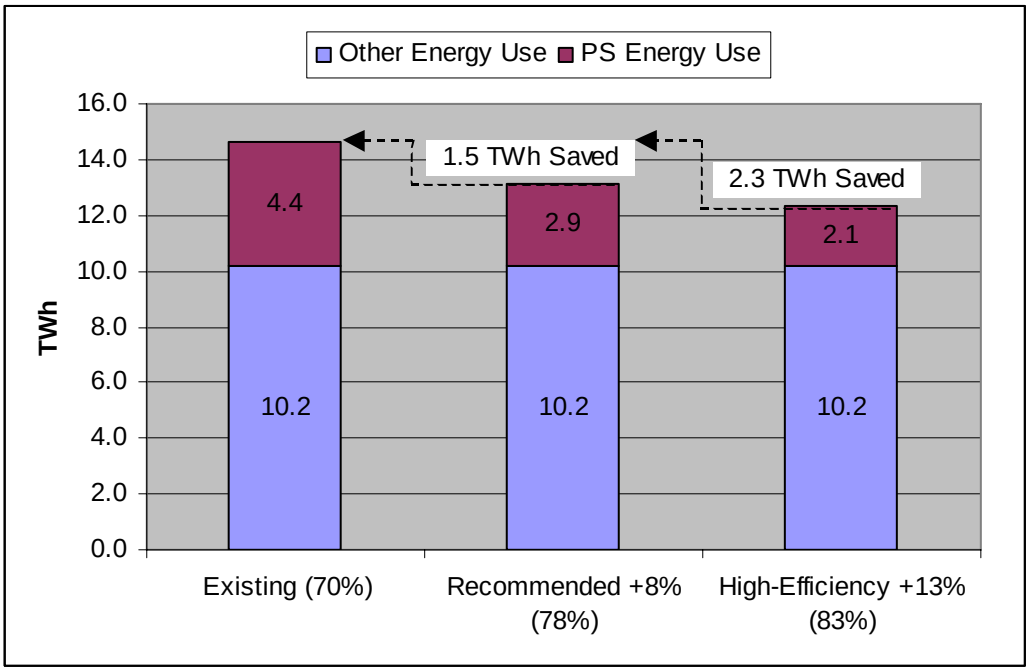

5 Power Supply Energy Savings in Other Devices

This task's objective was to assess other power supply savings opportunities in data centers. To do this we:

- Identified relevant topologies and configurations of "other" power supplies (e.g., routers, switches, hubs, etc.)
- Compiled information about the existing efficiency of power supplies in these devices through testing and manufacturer data.
- Estimated the market penetration of these devices and energy savings potential of more efficient power supplies.

5.1 Power Supply Topologies in Other Devices

Servers are certainly not the only devices in data centers that have potential to yield energy savings through improved power supply efficiency. Other devices such as routers, switches, hubs, and data storage units also contain power supplies. While the saving potential per unit is lower in these components, the sheer number of these devices warrants investigation.

As cited in Aebischer (2003), Mitchell-Johnson (2001) found that the composition of equipment at a typical data center in the United States was 60% servers, 18% switches, 9% disks and 8% routers.[52]

Using this information, a rough estimate of the number of devices in use can be made from our server stock estimates (Section 3.3.1). We estimated that there are approximately 7.2 million servers in use

[52] Aebischer, Bernard, Energy- and Eco-Efficiency of Data Centres, 2003, p. 14. Available at: http://www.cepe.ch/

in the United States. Table 21 shows how this would translate into numbers of other devices if all servers and other data handling equipment were found only in data centers.

Table 21. Other of Devices found in Data Centers

	Stock (units)	% of Equipment found in Data Centers
Servers	7,259,939	60%
Switches	2,177,982	18%
Disks	1,088,991	9%
Routers	967,992	8%
Other	604,995	5%
Total Devices	**12,099,898**	**100%**

Note: Switches represents the number of devices, not number of ports.

We know, however, that servers and other data handling equipment are routinely used in office environments as well, so we conducted more in-depth market research on four broad categories of "other" devices—routers, hubs, switches, and storage devices.

5.1.1 Routers

A router is a device that forwards data packets along networks. Routers are located at gateways, the places where two or more networks connect. Commonly, routers are used to connect two LANs or WANs or a LAN and an ISP's network.[53]

Cisco dominates the router market, claiming a 90 percent market share. Cisco's 2600 router is one of the most popular models on the market and now has added features such as Voice over Internet Protocol (VoIP).[54] Rivals Juniper Networks and 3Com have launched recent efforts to attack Cisco's dominant share. Juniper has started to take away market share with new systems in the terabit range.[55]

Routers in use today are typically stand-alone rack-mounted boxes, ranging in size from 1U to full-rack systems.[56] Some of the more common router deployments (and their power supply types) used by service providers, data centers, and enterprise networks are described in Table 22.

[53] http://www.webopedia.com/TERM/r/router.html
[54] http://news.com.com/Juniper+to+invite+Cisco+to+%27Pepsi%27+challenge/2100-1037_3-5171594.html
[55] http://techrepublic.com.com/5100-22_11-5363222.html
[56] ADL/US DOE (2002) p. 68.

Table 22. Description of Common Router Deployments

Router Class	Site	# of Routers	Examples	Power Supply
Service Provider	Telecommunications or Data Centers	Several dozen routers/switches or more in one facility.	Cisco 6513/7613, Cisco 12000, 12400, Juniper M160, M320, T320. Foundry or Extreme switches.	Mainly dual power supplies, sometimes 3 per system (two active, one in backup). Choice of AC or DC. Large telco would be DC.
Service Provider (Core)	Telecommunications	2 or 4 multiple chassis working together.	Juniper T640, Avici TSR, Cisco CRS-1 *Note: These are rate – a pair of T320s (above) would be more common.*	
Mid-range	Data Centers		Juniper M20, M40e, Cisco 7500, 7600/6500 systems.	Usually redundant dual power supplies AC or DC.
Small	Customer premise or enterprise network	Switches also used	Cisco 2500/2600	Some accommodate dual power supplies, but many do not. Some provide for a DC option but AC is most common.
Very Small	Residential DSL or cable, or small business/LAN			External AC power supplies.

Source: Dave O'Leary, Juniper Networks.

ADL/US DOE (2002) estimated the current stock of routers at 3.2 million by summing the shipments of routers, as reported in ITIC (2000) for the four years 1997 to 2000. We obtained more recent estimates of router shipments (segmented by router class) from Synergy Research Group (Table 23). This data shows that low-end, small and branch-office, routers account for over 80% of the current router stock assuming a four-year life as in ADL/US DOE (2002).

Table 23. Revised Router Stock

Router Class	2001	2002	2003	2004	Total Stock	% of Stock
Small	36,176	15,833	206,836	105,890	364,735	19%
Branch Office	382,928	325,912	273,732	237,688	1,220,260	62%
Mid-Range	81,178	78,131	70,103	58,774	288,186	15%
High-End Enterprise	10,453	11,689	5,874	1,503	29,519	2%
Service-Provider Edge	12,140	8,005	10,172	13,061	43,378	2%
Service-Provider Core	4,238	1,476	998	1,193	7,905	0%
Total	**527,113**	**441,046**	**567,715**	**418,109**	**1,953,983**	**100%**

Note: 2004 shipments are based on 3 quarters of actual data and one quarter of forecasted data.

Source: Synergy Research Group.

Next, we surveyed the manufacturer datasheets of over 65 different routers to determine the average power supply rating (in watts) and configuration for each router class (Table 24).

Table 24. Router Power Consumption

Router Class	Examples	Average PS Rating (Watts)	# of PS
Small	Allied Telesyn AR3XX Family Cisco 800, 1000, 1600 Tasman 1000 Family	15	1
Branch Office	3Com 3XXX Family Allied Telesyn AR4XX Family Cisco 1600, 1700, 2500, 2600 Tasman 4000	57	1 (or 2)
Mid-Range	3Com 5XXX Family Cisco 36XX, 4000	182	2 (or 1)
High-End Enterprise	3Com 6XXX Family Cisco 7500, 7206	513	2
Service-Provider Edge	Cisco 64XX, 7500, 7206, 7600, 10000, 7301, 7400 Juniper M5, M10, M20, ERX, M320 Laurel Networks ST200, ST Redback 400 Smart Edge, SMS 1800SL, SMS 1000SL Nortel MPE 9500	1,215	2 (or more)
Service-Provider Core	Alacatel 7770 OBX Avici TSR, SSR, QSR Caspian Aperio Cisco GSR, 12000 series, CRS-1 Juniper M40x, M160, T320, T640	3,660	2 (or more)

Based on this information, we determined the power consumption for routers of different classes (Table 25).

Table 25. AEC per Unit, Router

Router Class	Average PS Rating (Watts)	Hours / Yr	AEC, Per Unit (kWh)
Small	15	8,760	51
Branch Office	57	8,760	201
Mid-Range	182	8,760	1,275
High-End Enterprise	513	8,760	3,597
Service-Provider Edge	1,215	8,760	8,516
Service-Provider Core	3,660	8,760	25,653
Total			

5.1.2 Switches

LAN Switches

Roth (2002) estimated the stock of LAN switch ports in the United States to be 95 million.[57] This estimate is probably low, as recent shipments of switches are more than double those of just a few years ago.

In 2003, the worldwide LAN switch market grew 16.0% with total port shipments of 193.0 million versus 166.3 million in 2002. Despite this, market revenues decreased from $13.0 billion to $11.4 billion, due to falling average sales prices. In-Stat/MDR is predicting total port shipments of nearly 502.8 million by 2008.[58] Gigabit switches continue be one of the main drivers of this market, having surpassed Fast Ethernet in revenue during the first quarter of 2004. Cisco is by far the market leader in LAN switches, with over two-thirds of the market: HP, Nortel, 3Com, Extreme, and Foundry are also players. [59]

According to Dahlquist and Borovick (2000) [60] and other sources,[61] the United States accounts for about half of worldwide sales. Based on this information, and a two-year product life, we estimate the 2004 stock of U.S. LAN switches at approximately 213.4 million ports (Table 26).

Table 26. LAN Switch Stock – Worldwide and U.S. Market (Ports, millions)

Year	WW Switch Sales	U.S. Switch Sales	U.S. Stock
2001	137.2	68.6	
2002	166.3	83.2	151.7
2003	193.0	96.5	179.7
2004	233.7	116.9	**213.4**
2005	283.1	141.5	258.4
2006	342.8	171.4	312.9
2007	415.2	207.6	379.0
2008	502.8	251.4	459.0

Note: Worldwide sales data from In-Stat/MDR, assumes straight-line growth from 2003 to 2008. Assumes U.S. equals 50% of worldwide market. Stock calculated as two-year rolling average.

To determine the AEC of these switches, we surveyed over manufacturer datasheets for over 100 different products. Table 27 summarizes these results.[62]

Table 27. Power Consumption of a Switch

LAN Switch Type	Power Consumption per Port (Watts)	Hrs/Yr	AEC per Port (kWh/Yr)
Ethernet/Fast Ethernet	4.3	8,760	38
Gigabit	5.4	8,760	47

Note: Manufacturer datasheets did not typically differentiate between Ethernet and Fast Ethernet; therefore these categories were combined.

[57] ADL/US DOE, p. 71.
[58] http://www.instat.com/press.asp?Sku=IN0401449LN&ID=1037
[59] http://www.idc.com/getdoc.jsp?containerId=pr2004_08_19_165431
[60] ADL/US DOE, p. 71. Note: this estimate is probably low, as it is based only two years of sales data. Switches likely have a product lifetime similar to routers, 4 years.
[61] http://www.extremenetworks.com/aboutus/pressroom/releases/pr05_17_00.asp
[62] These numbers are similar to the findings of ADL/US DOE (2002), which estimated 35 kWh/yr per port.

WAN Switches

ADL/US DOE (2002) estimated a total of 50,000 WAN shelves installed in the United States.[63] However, this estimate is likely to be low as it appears to be based on only 1 year of sales data.

More recent market research from Dittberner Associates forecasts that the number of carrier WAN switch *and routers* shipped will increase from 24,841 units in 2001 to 82,157 in 2008, alongside a revenue increase from $3.76 billion to $18.2 billion.[64] However, other data sources describe a very different picture for this market indicating that "multiservice WAN switch sales rebounded nicely in 2004, after three consecutive years of declining sales."[65] The decline in sales was largely the result of service providers cutting capital expenditure budgets: about 90% of total multiservice switch sales are to service providers.[66]

2004 saw a turnaround in worldwide multiservice WAN switch sales. In-Stat/MDR reported sales of $2.3 billion, an increase of 9% over 2003 levels. Growth rates are expected to gradually decline over time as "switches face stiff competition from routers and service providers evolve to MPLS networks."[67] Assuming a 50% U.S. market share (as for LAN switches), Table 28 estimates the U.S. sales of WAN switches through 2008. Over the past three years, ATM/IP/MPLS-related equipment accounted for about 75-80% of WAN multiservice switch sales, while the remaining 20-25% was Frame Relay.[68]

Table 28. WAN Switch Sales –Worldwide and U.S. Market ($, billions)

Year	WW WAN Switch Sales	Projected Growth (%)	U.S. WAN Switch Sales
2001	$3.8		$1.9
2002	$2.3	-41%	$1.1
2003	$2.1	-7%	$1.1
2004	$2.3	9%	$1.2
2005	$2.5	7%	$1.2
2006	$2.6	6%	$1.3
2007	$2.8	6%	$1.4

Worldwide projected growth rates and sales from In-Stat/MDR.

Assuming an average sales price of $155,000[69] and a product life of four years, we estimate the current stock of WAN switches installed in the Unites States to be 34,000 (Table 29.)

Table 29. U.S WAN Switch Sales and Stock, Units

Year	U.S. WAN Switch Sales	U.S. WAN Switch Stock
2001	12,417	
2002	7,326	
2003	6,813	
2004	7,426	**33,983**
2005	7,946	29,512
2006	8,423	30,609
2007	8,928	32,724

Note: U.S. WAN switch estimated based on 50% market share.

[63] ADL/US DOE, p. 72.
[64] http://www.dittberner.com/news/newsrelease2.php
[65] http://www.instat.com/press.asp?ID=1286&sku=IN0501949WN
[66] http://www.instat.com/press.asp?ID=570&sku=IN030636WN
[67] http://www.instat.com/press.asp?ID=1286&sku=IN0501949WN
[68] In-Stat/MDR, Networking Quarterly - Multiservice WAN Switches, 2002-2004. Available at: http://www.instat.com/Catalog/ncatalogue.asp?ID=67&year=2004
[69] Based on Dittberner Associates 2001 data: $3.76 Billion and 24,841 units shipped.

Leading vendors in the multiservice WAN category are Alcatel, Ciena, Cisco, Ericsson, Lucent, Marconi, and Nortel.[70] Table 30 provides some example switching products, their power supply configurations, and approximate power consumption. As the table shows these products typically offer redundant power, with both DC and AC options available as required.

Table 30. Multi-Service WAN Switches

Manufacturer	Model	Watts	Power
Cisco	Cisco MGX 8950 Multiservice Switch	5000	Optional Ac
Cisco	Cisco MGX 8850 Series	1000	Optional Ac
Cisco	MGX8830	1200	Ac
Cisco	MGX8830	1050	Dc
Nortel	Multiservice Switch 6480		Ac or Dc, redundant
Nortel	Multiservice Switch 6440		Ac or Dc, redundant
Nortel	Multiservice Switch 6420		Ac
Nortel	Multiservice Switch 7420		Dc (Ac option), redundant
Nortel	Multiservice Switch 7440		Dc (Ac option), redundant
Nortel	Multiservice Switch 7460		Dc (Ac option), redundant
Nortel	Multiservice Switch 7480		Dc (Ac option), redundant
Nortel	Multiservice Switch 15000		Dc
Lucent	GX 550® Multiservice WAN Switch	1154	Dc
Lucent	GX 550® Multiservice WAN Switch	450	Ac
Lucent	CBX 3500	2880	Ac or DC Redundant chassis power distribution system
Lucent	CBX 500	1400	Redundant AC or Dc
Alcatel	OS/R	375	Ac, optional redundant
Alcatel	OS/R	650	Ac, optional redundant
Alcatel	Omni Switch 512	50	Internal Dc
Ciena	DN 7000	220	Dc
Ciena	DN 7050	350	Dc
Ciena	DN 7100	680	Dc
Ciena	DN 7200	2800	Dc

Source: Manufacturer datasheets.
Note: Nortel did not provide access to hardware specifications free of charge online.

Based on the manufacturer data collected, we determined the average power draw of a WAN switch to be 1,284 watts. Hence, assuming continuous operation the AEC of a WAN switch is approximately 11,250 kWh per year (Table 31).

Table 31. Power Consumption of a WAN Switch

	Power Consumption per Unit (Watts)	Hrs/Yr	AEC per Units (kWh/Yr)
WAN Switch	1,284	8,760	11,247

Power consumption per unit was determined by the average of 15 products which listed wattage ratings.

5.1.3 Hubs

A hub is common connection point for devices in a network. Hubs are commonly used to connect segments of a LAN.

[70] http://www.instat.com/abstract.asp?id=4&SKU=IN0501949WN

Silva (1998) estimated the number of hub ports installed in commercial office buildings in the United States to be 93.5 million.[71] However, over the last 5 years there has been a rapid migration from hubs to switches. In 1999, PC Magazine proclaimed, "The hub is dead….long live the switch!…No one should even think about installing a hub."[72]

Figure 13 shows a decline in worldwide shipments of Ethernet hubs from over 50 Million ports in 1998 to only 30 Million ports in 2003. Likewise, while shipments of Fast Ethernet hubs rose through 2001, these too have started to decline.

Figure 13. Worldwide LAN Hub Market, System Ports (000s)

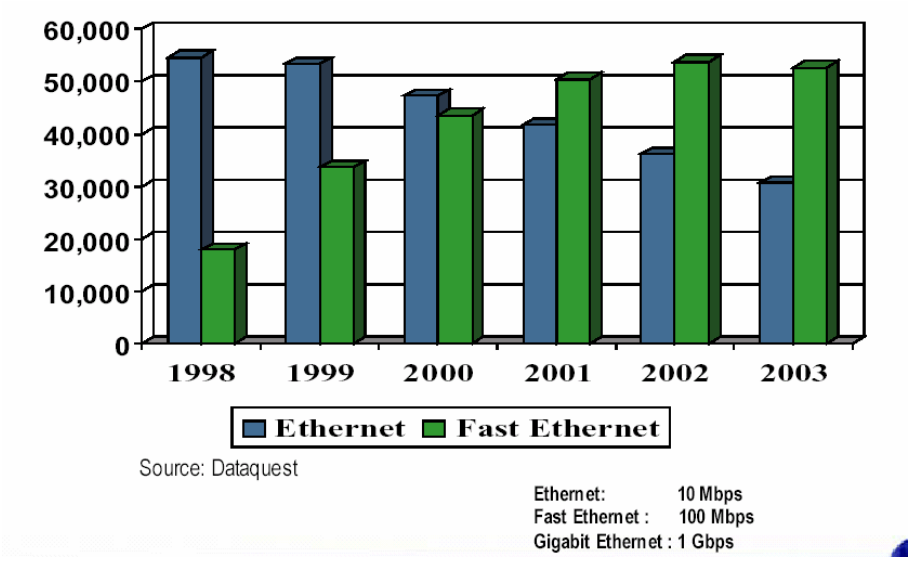

Source: http://www.xilinx.com/esp/consumer/home_networking/pdf_files/ethernet/complete.pdf

Using the above sales data and the same assumptions for hubs as for LAN switches (i.e., the United States represents 50% of the market and a product life of two years), we estimate the current stock of U.S. LAN hubs is approximately 77 million ports (Table 32).

Table 32. LAN Hub Stock – U.S. Market (Ports, millions)

LAN Hub Type	2001	2002	2003	2004e	2004 Stock	% of Stock
Ethernet	20	18	15	13	28	36%
Fast Ethernet	25	27	26	24	50	64%
Total	**45**	**45**	**41**	**36**	**77**	**100%**

Notes: Assume US market is 50%. Source: Ecos estimate based on 2003-2004 market data; 2004 data estimated based on trends in previous year sales.

In 2000, IDC projected that the worldwide *installed* base of 10 Mbps (Ethernet) hubs would drop from 184 million to 153 million ports (17%) and another 21% to 121 Million in 2001.[73] Assuming a 20% drop in subsequent years, by 2004th the worldwide installed base of Ethernet hubs would be projected to approximately 62 Million ports in 2004, and the U.S. portion about 30 million. This is very close to our estimate of 28 million ports in Table 32.

To determine the AEC of hubs, we surveyed over manufacturer datasheets for over 12 different products. Table 33 summarizes these results.[74]

[71] Roth, p. 73.
[72] http://www.onshore.com/downloads/partners/cisco/CIS_Mig_hub.pdf
[73] http://www.onshore.com/downloads/partners/cisco/CIS_Mig_hub.pdf
[74] These numbers are similar to the findings of Roth (2002), which estimated 35 kWh/yr per port.

Table 33. Power Consumption of a Hub

LAN Hub Type	Power Consumption per Port (Watts)	Hrs/Yr	AEC per Port (kWh/Yr)
Ethernet/Fast Ethernet	1.9	8760	17

Note: Manufacturer datasheets did not typically differentiate between Ethernet and Fast Ethernet; therefore these categories were combined.

5.1.4 External Data Storage

The demand for data storage is large.[75] IDC data (Table 34) indicates that in 2003 nearly 400,000 external storage units were shipped worldwide, of which about two-thirds was direct attached storage (DAS) and the remaining one-third was networked, including both network attached storage (NAS) and Storage Area Networks (SAN).[76]

In 2004, IDC reported that it was encouraging to see the "continued acceleration in the annual growth rate for external disk storage systems petabytes, which grew 63% in 2004. With revenue accelerating, but at a slower pace, average pricing remains competitive during the ongoing influx of higher-capacity drives and applications."[77]

EMC was the largest vendor of external storage in 2004, with a 21.1% revenue share. HP and IBM followed with 18.7% and 12.6% share, respectively. Hitachi and Dell are also dominant players.[78]

Table 34. Worldwide External Data Storage Shipments and Sales

	1999	2000	2001	2002	2003	2004est
Units						
DAS	503,608	509,667	425,255	298,264	270,379	301,492
Networked	74,215	138,455	140,902	141,148	128,599	143,397
Total External Disk Storage	577,823	648,122	566,157	439,412	398,978	444,889
Non-OEM Factory Revenue ($M)						
DAS	$13,773	$14,452	$9,357	$5,932	$5,504	$7,112
Networked	$4,368	$7,299	$7,838	$7,165	$8,087	$8,055
Total External Disk Storage	$18,141	$21,751	$17,195	$13,097	$13,591	$14,165

Note: Networked includes NAS and SAN.
Source: 2004 data estimated based 2004 IDC sales data and average prices. 1999-2003 data: IDC, 2004 Release as shown in *The Business for Storage Networks. Chapter 1, Industry Landscape: Storage Costs and Consumption*. Cisco Systems, Sept 3, 2004. Available at:
http://www.cisco.com/application/pdf/en/us/guest/netsol/ns516/c1272/cdccont_0900aecd80257124.pdf

The United States disk storage systems market represents 39% of the worldwide market.[79] Assuming a product life of four years, we estimate the 2004 stock of External Disk Storage to be 721,280 units.

[75] The ADL/US DOE 2002 report did not look at data storage separately; it was combined with servers.
[76] Williams, Bill. The Business Case for Storage Networks. Cisco Press, 2004. Available on: http://www.ciscopress.com/content/images/1587201186/samplechapter/1587201186content.pdf
[77] http://www.idc.com/getdoc.jsp?containerId=pr2005_03_03_154203
[78] http://www.idc.com/getdoc.jsp?containerId=pr2005_03_03_154203
[79] Genereux, Scott. The United State Storage Market: A Perspective from Hitachi Data Systems Corporation.

Table 35. U.S. External Data Storage Stock

Year	Units Shipped			Estimated Stock
	DAS	Networked	Total	
1999	196,407	28,944	225,351	
2000	198,770	53,997	252,768	
2001	165,849	54,952	220,801	
2002	116,323	55,048	171,371	870,290
2003	105,448	50,154	155,601	800,541
2004	117,582	55,925	173,507	**721,280**

Notes: Assumes U.S. is 39% of worldwide market. Stock calculated as 4-yr rolling average.

There is a wide range of storage products. We briefly surveyed some of the storage products offered by Dell, EMC, HP, IBM, and once again we found that many manufacturers do not provide details about power supplies in their product datasheets. Table 36 lists some of the products that did report power supply information.

Table 36. Selected Storage Products

Manu	Model	Form Factor	Watts	Power Supply	Efficiency
Dell	PowerVault 22Xs	3U	600		
Dell	Powervault 745N	1U	280		
EMC	NetWin	1U	250	Auto-sensing 120/240V	
EMC	CLARiiON AX100 - single processor	2U	250	Single power supply	
EMC	CLARiiON AX100 - dual processor	2U	326	Redundant	
EMC	CLARiiON CX700	2U	510	Redundant	
EMC	NS700	8U or 4U	510		
EMC	NS500	6U or 3U	578		
EMC	CLARiiON CX500	2U	618	Redundant	
EMC	Clariion CX300	2U	618	Redundant	
HP	HP ProLiant DL100 Storage Server	1U	250	Single power supply	
HP	HP ProLiant ML110 Storage Server	Tower 5U	350	1 350W non-redundant power supply	
HP	HP StorageWorks DL380-SL Clustered Gateway	2U	575	CE Mark Compliant, Hot Plug Redundant power supply included	
HP	HP ProLiant DL380 G4 Storage Server	2U	575	Redundant Power Supplies; 2 x 575-Watt Hot Plug Power Supplies	73.1-77%
HP	HP ProLiant ML350 G4 Storage Server	Tower 5U	725	1 hot plug 725 Watt Power Supply (Base Model); 2 hot plug 725 Watt Power Supply (Int SCSI Storage)	
HP	HP ProLiant ML370 G4 Storage Server	Tower 5U	775	1 hot plug 775 Watt Power Supply (Base Model); 2 hot plug 775 Watt Power Supply (High Performance Model)	
HP	HP ProLiant DL580 G2 Storage Server	4U rack mount	800	2 x 800-Watt (low line or high line) Hot Plug Power Supplies	70-76%

Manu	Model	Form Factor	Watts	Power Supply	Efficiency
HP	HP ProLiant DL585 Storage Server	4U rack mount	800	2 x 800-Watt (low line or high line) Hot Plug Power Supplies	70-76%
IBM	IBM TotalStorage NAS Gateway 500	Rack-mount 4U	670	Two 670 Watt auto-ranging hot swappable power supplies	

Source: Manufacturer datasheets.

Based on the products listed above, we determine the average AEC of external disk storage units to be 4,591 kWh per year. Table 37 summarizes these results.

Table 37. Power Consumption of External Data Storage

LAN Hub Type	Power Consumption per Unit (Watts)	Hrs/Yr	AEC per Units (kWh/Yr)
External Disk Storage	529	8760	4,591

5.2 Existing Efficiency of Power Supplies in Other Devices

5.2.1 Manufacturer Data

We found that manufacturers of routers, hubs, switches, and storage devices provide very little, if any, information about the device's power supply in their product specifications. Manufacturer datasheets sometimes give the rated output and/or configuration of the power supply, but do not list any information about efficiency. Dave O'Leary of Juniper Networks confirms that this is standard industry practice.[80]

Aebischer (2003) notes that the power supplies contained in these other devices are similar to the ones found in servers, therefore one would expect to see similar results.[81] Aebischer also measured "the workload of several servers, routers and switches…" and found it "…lies typically between 50% and 30%."[82]

Juniper Networks provided us with some limited information about several of their products (Table 38).

[80] Personal communication with Dave O'Leary, Nov 29, 2004.
[81] Aebischer, Bernard, Energy- and Eco-Efficiency of Data Centres, 2003, p. 28. Available at: http://www.cepe.ch/
[82] Aebischer, Bernard, Energy- and Eco-Efficiency of Data Centres, 2003, p. 31. Available at: http://www.cepe.ch/

Table 38. Power Supply Efficiency in Juniper Network Devices

Model	Device Category	PS Rating (Manufacturer? Configuration?)	# of PS	Reported Efficiency (at Full Load)	Tested Efficiency ?
M5, M10	Internet Router	800W (AC); 700W (DC)	two load-sharing, AC or DC		
M7i	Internet Router	293W	one or two load sharing, AC or DC		
M20	Internet Router	750W isolated	two load-sharing, isolated, AC or DC		
M40e	Internet Router	2900W isolated (AC); 3000W nonisolated (DC)	two load-sharing, pass-through, AC or DC		
M160	Internet Router	2400W nonisolated; 3000W nonisolated (enhanced)	two load-sharing, pass-through power supplies, DC		
M320	Internet Router	1750W (AC) 2000W (DC)	four load-sharing AC power supplies, AC or DC		
T320	Internet Router	3200W	two load-sharing DC power supplies		
T640	Internet Routing Node?	3200W	two load-sharing DC power supplies		
TX	Matrix Platform	4560W	two load-sharing DC power supplies		

Source: http://www.juniper.net/techpubs/hardware/

In addition, we found one power supply manufacturer (Delta Electronics) that classifies its product offerings by application (Desktop, Server, Networking, and Other). Figure 14 provides a summary of the power supply ratings and efficiencies of all Delta Electronics *Networking* products. These products are further divided into five subcategories: Network, PC Peripheral, Storage, 1U Application, and Information Application. While there is a very wide distribution of power supply ratings (from about 5W to over 4500W, the majority of these products have efficiencies in the 60-75% range. Only a handful of the Network and Storage power supplies are 80% or more efficient; these appear to be the higher-rated models (700W and above). One of the units claims a remarkable 95% efficiency.

Figure 14. Efficiency of Delta Electronics Networking Power Supplies

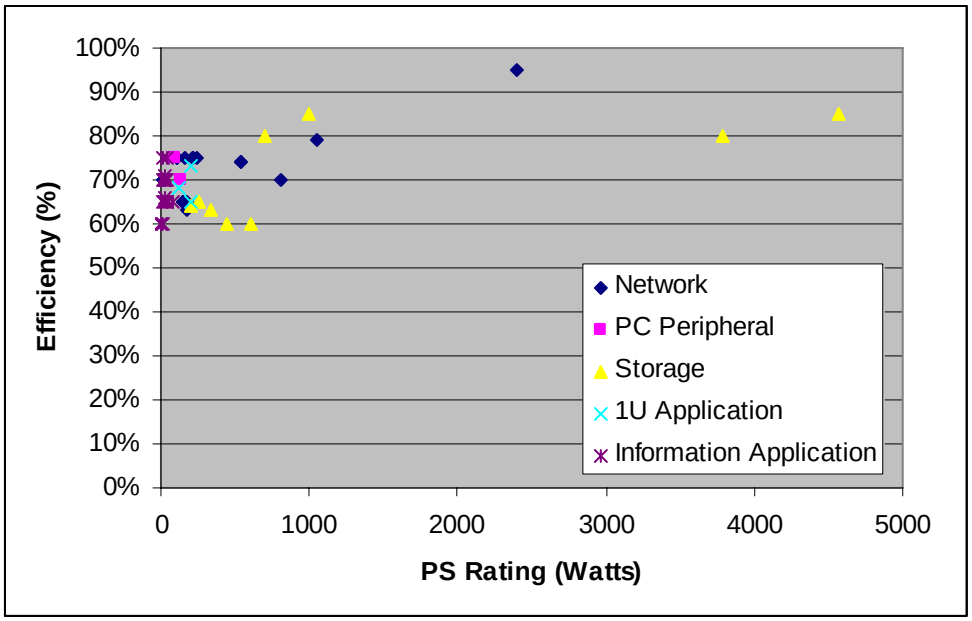

Source: Delta Electronics, product datasheets. Available at: http://www.delta.com.tw/product/ps/sps/sps_main.asp

The average efficiencies for different subcategories of Delta Electronics Networking power supplies are provided in Table 39.

Table 39. Average Efficiency of Delta Electronics Networking Power Supplies

Networking Power Supply Subcategory	# of Power Supplies	Average Efficiency (%)
Network	17	73.2%
PC Peripheral	3	73.3%
Storage	10	70.7%
1U Application	3	68.7%
Information Application	18	67.3%
All	**51**	**70.4%**

Source: Delta Electronics, product datasheets. Available at: http://www.delta.com.tw/product/ps/sps/sps_main.asp

5.3 Market Penetration of Other Devices and Energy Savings Potential

We estimated the energy savings potential of more efficient power supplies for "other" devices in three steps:

1. We developed a preliminary AEC estimate for each category of device from market research gathered on the number of devices and the power supplies, contained there in (Section 4.1).
2. We revised our AEC estimate to more accurately reflect the typical operation and efficiency of these devices (i.e., 30% loading and 70% efficient power supply) (Section 4.2).

3. We then estimated savings potential of more efficient power supplies for two cases (as for servers): a “recommended” 78% efficiency case and a high-efficiency 83% case.

5.3.1 AEC

The AEC for each category of device—router, switch, hub, and data storage—was calculated by multiplying the annual AEC per unit times the total stock. This methodology is similar to the one was used in ADL/US DOE 2002. Comparisons the ADL/US DOE 2002 estimate are provided here for reference.

Routers

The annual energy consumption (AEC) for routers is 1.31 TWh (Table 40). This estimate is similar to the ADL/US DOE 2002 estimate of 1.1 TWh. However, their number was derived very differently, assuming over 1 million more servers and flat average 40 watts per server (about 350 kWh per year).

Table 40. Annual Energy Consumption, Routers

Router Class	AEC per Unit (kWh)	2004 Stock	AEC (TWh)
Small	51	364,735	0.02
Branch Office	201	1,220,260	0.24
Mid-Range	1,275	288,186	0.37
High-End Enterprise	3,597	29,519	0.11
Service-Provider Edge	8,516	43,378	0.37
Service-Provider Core	25,653	7,905	0.20
Total		**1,953,983**	**1.31**

Switches

For LAN switches, we estimate AEC of 9.0 TWh. (Table 41). This nearly three times higher than the AEC estimated by ADL/US DOE (2002) for LAN switches (3.3 TWh). Our estimate of AEC per port is similar; so, the increase in overall AEC is due to the increase in the number of ports installed.

Table 41. AEC of LAN Switches.

LAN Switch Type	AEC per Port (kWh)	2004 Stock (Ports, Millions)	AEC (TWh)
Ethernet/Fast Ethernet	38	106.7	4.0
Gigabit	47	106.7	5.0
Total		**213.4**	**9.0**

Notes: Assumes 50% of market is Gigabit.

For multiservice WAN switches, we estimate 0.4 TWh AEC. (Table 42). Like LAN switches, this nearly three times the AEC estimated by ADL/US DOE (2002) for WAN switches (0.15 TWh). However, unlike LAN switches, for WAN switches the increase in AEC is largely due to a much higher estimate of energy consumption per device.

Table 42. AEC of WAN Switches.

	AEC per Units (kWh/Yr)	2004 Stock	AEC (TWh)
WAN Switch	11,247	33,983	0.4

Hubs

Our AEC estimate for LAN hubs is 1.3 TWh. (Table 43). This is slightly less than the AEC estimated by ADL/USDOE (2002) for hubs (1.6 TWh), which makes sense given the transition from hubs to switches.

Table 43. AEC of LAN Hubs

LAN Hub Type	AEC per Port (kWh/Yr)	2004 Stock (Ports, Millions)	AEC (TWh)
Ethernet	17	28	0.5
Fast Ethernet	17	50	0.8
Total		**77**	**1.3**

External Data Storage

Combining our stock and AEC per unit estimates yields an annual AEC of 3.3 TWh for U.S. external disk storage (Table 44).

Table 44. AEC of External Data Storage

	AEC per Unit (kWh)	2004 Stock	AEC (TWh)
External Data Storage	4,591	721,280	3.3

5.3.2 Adjusted AEC

The AEC estimates above use a similar methodology to the one employed by ADL/US DOE (2002). This methodology, however, does not account for existing power supply efficiency, nor does it account for the fact that devices normally operate at a fraction of their stated output rating.

We revised our initial AEC estimates to reflect an average of 30% loading and 70% power supply efficiency. Table 45 presents the "adjusted" AEC.

Table 45. AEC, adjusted for 70% PS Efficiency and 30% Loading

Segment	AEC (TWh)	Adjusted AEC (TWh)	% of Adjusted AEC
Routers	1.3	0.6	8%
Switches - LAN	9.0	3.9	59%
Switches - WAN	0.4	0.2	3%
Hubs	1.3	0.6	8%
Storage Devices	3.3	1.4	22%
Total	15.3	6.6	100%

Note: AEC is adjusted by a factor of .43 =.3/.7 to account for loading and PS efficiency.

5.3.3 Efficient Power Supply Energy Savings Potential

Together, these "other devices" consume a little less than half of the energy that servers do. LAN switches contribute the most to AEC simply due to their large numbers. Remember that, as with servers, not all of the energy consumed goes to power the device itself. Some energy is "lost" or consumed by the power supply itself; the amount of energy consumed depends on the efficiency of the power supply. In the typical existing case, where power supply efficiency is 70%, 30% or about 2.0 TWh is consumed by the power supply (i.e., energy losses)

As with servers, energy savings can be achieved by increasing the efficiency (or reducing the energy losses) of the power supply. We present two alternative cases to demonstrate the potential energy savings due to more efficient power supplies. The first is a "recommended case" where power supply efficiency is increased from 70% to 78%; the second is a "high-efficiency case" where power supply efficiency reaches 83%. Figure 15 shows that by increasing power supply efficiency by 8%, to 78% we are able to save 0.7 TWh by cutting power supply energy consumption by 34% to 1.3 TWh.

Similarly, under the high-efficiency scenario, savings of 1.1 TWh are realized by cutting power supply energy consumption 52% to 0.9 TWh.

Figure 15. Energy Savings Potential of PS in Other Devices

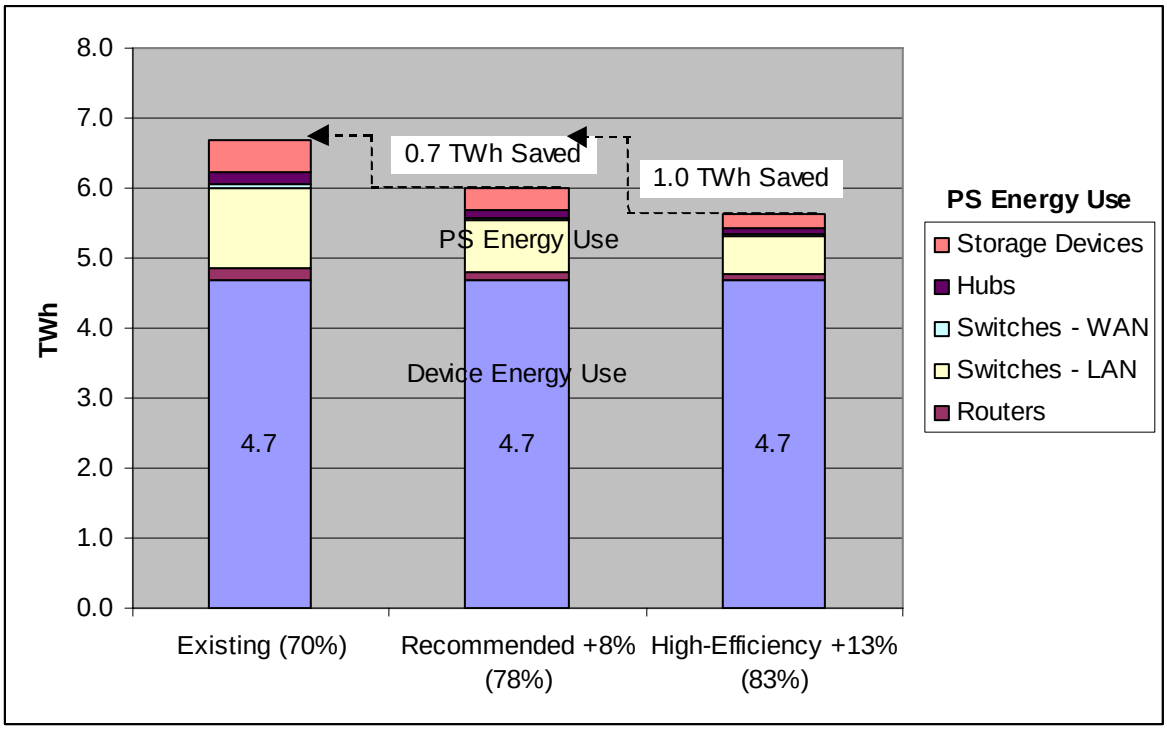

These high-level estimates show that there is significant energy savings potential from improving efficiency of power supplies in devices other than servers. While the savings potential is not as large as in servers in absolute terms, the extension of efforts to these “other’ devices would be relatively straightforward. Much of the groundwork for improving power supply efficiency and developing specifications for servers, desktops, etc. has already been done or is in process. This work could serve as a starting point for efforts targeting other devices.

The main challenge to encouraging efficiency in the power supplies of other IT equipment like routers and switches is that there is currently no industry body like SSI coordinating standards and efficiency improvements in these products. Power supply designs for this type of equipment are usually customized for a particular product and may vary from manufacturer to manufacturer, whereas in servers, only a few distinct form factors for power supplies exist. The diversity of power supply designs in equipment like routers and switches may complicate efforts to uniformly encourage efficiency improvements.

Fortunately, one does not necessarily have to address all of these product categories at once and can instead focus on the “low hanging fruits”. Our estimates help to provide some guidance as to which devices should be prioritized first. For example, LAN switches clearly represent the best energy savings opportunity because of the sheer numbers of devices in use.

Attachment X: Demonstration of Demand Controlled Ventilation

Outreach & Technology Transfer

Prepared for the California Energy Commission
High-Performance Buildings for High-Tech Industries

by

Lawrence Berkeley National Laboratory
Applications Team

UC Research Agreement, Contract #500-02-004
Task Work Plan: Task 6

Status as of June 26, 2007
Task Lead: Evan Mills

For more information on this project, see:
http://hightech.lbl.gov

OBJECTIVES

The Opportunity

The California Energy Commission's PIER Industrial Program has recognized that energy consumed by buildings used in California's high-tech industries and institutions represents an attractive opportunity for energy savings and peak electricity demand reductions, often coupled with significant non-energy benefits. High-tech buildings are characterized by large base-loads operating 24 hours a day with energy intensities much larger than typical office buildings. High-tech buildings include laboratories, cleanrooms, and data centers that are essential to various industries (and to public-sector energy users such as schools and government facilities) important to the state's economy.

The overarching goal of this project is to improve the energy efficiency of high-tech buildings—cleanrooms, laboratories, and data centers—thereby reducing energy use in California, providing financial benefits to California companies, retaining high-tech companies in California, and improving the health and safety of California workers. A number of major technical tasks will be conducted with applications to the following facility types and topics:

Cleanrooms:
- Energy benchmarking
- Opportunities in mini environments
- Development of a fan filter unit test procedure
- Demand-controlled filtration

Laboratories:

Energy benchmarking
High-efficiency laboratory fume hoods

Data Centers:
Energy benchmarking
Efficient power supplies
Efficient uninterruptible power supplies (UPS)
Direct-current powering
Improved air management
Building Equipment – IT Interface

In parallel with the Tech Transfer activities, demonstration projects have been executed, the results reported on the project website, and communicated through appropriate industry venues such as workshops with local industry organizations.

Communications with key constituencies have been central to the success of the project. This report outlines a comprehensive Draft Outreach & Technology Transfer efforts to support the dissemination of results from the aforementioned technical tasks to the public and key stakeholders and decision-makers in the high-tech facilities arena.

Measures of Success

- Short term (this task): Specific goals included communicating results (text, data, graphics) of individual tasks to key audiences and decision-makers likely to utilize the information. As a result, the "high-tech facilities" activities have increased visibility in the energy and facilities management communities.

- Long term: widespread (and actionable) awareness of the products of this project among key decision-makers, and actions taken on the basis of our results leading to significant energy savings in new and existing high-tech facilities.

Applicable Contract Language

The goal of this task is to develop a plan to make the knowledge gained, experimental results and lessons learned available to key decision-makers.

The Performing Institution shall:

- *Prepare a Technology Transfer Plan*
  - *The plan shall explain how the knowledge gained in this project will be made available to the public. The level of detail expected is least for research-related projects and highest for demonstration projects. Key elements from this report shall be included in the Final Report for this project.*

- *Submit the draft Technology Transfer Plan to the Commission Project Manager for review and comment in accordance with the approved Schedule of Deliverables.*
- *Conduct technology transfer activities in accordance with the Technology Transfer Plan. These activities shall be reported in the Monthly Progress Reports.*
- *Submit all deliverables identified in the approved Technology Transfer Plan to the Commission Project Manager for review and comment.*

Deliverables:

- *Draft Technology Transfer Plan (submitted: March 30, 2004)*
- *Final Technology Transfer Plan*
- *All Applicable Technology Transfer Deliverables*

APPROACH

Strategy

We began by developing a detailed portrait of our diverse audiences, and mapped them against the relevant communications channels and techniques. We then utilized the following strategies to achieve technology and information transfer to these audiences:

- Establish "Basic Infrastructure"
- Interact with Project Advisory Committees (PACs)
- Produce Self-published Products
- Disseminate Information Through the Trade and Technical Media
- Seek Visibility Through The Popular Media
- Distribute Results Through The Internet: Web & E-mail
- Distribute Results Through Conferences, Meetings, and Workshops
- Transfer Technology Through Standards & Guidelines; Service on Technical Committees
- Distribute Results Through Curriculum & Training

Target Audiences and Communications Channels

Communications planning appropriately began with an assessment of the intended audiences, and the channels available for reaching them. The matrix shown below exemplifies productive intersections between target audiences and communications channels, and is being used as a roadmap for creating and targeting communications products and activities within this Project.

Communications and Outreach Matrix: High-Performance Facilities for High-Tech Industry

		Communication Channel											
Priority	**Audiences**	**Reports; Presentations**	**Best-Practices Briefs**	**Standards & Guidelines**	**Face-to-Face (conf's, mtgs)**	**Utility Wkshps**	**Trade Journals**	**Tech. Journals**	**In-House Media**	**Popular Media**	**Curriculum & Training**	**Internet**	**Intranet**
*****	End users (Facility owners, operators, designers, etc – public/private)		x	x	x	x	x				x	x	
*****	PACs (one for Labs & Cleanrooms; one for Data Centers)	x			x							x	x
*****	Client (CEC/CIEE)	x			x							x	x
****	Benchmarking Sites	x	x		x								
****	Industry associations (Silicon Valley Mfr's. Group, etc.)		x	x	x	x	x				x	x	
***	Product manufacturers, designers, contractors, etc.		x	x			x	x				x	
***	Utilities		x	x	x	x		x	x		x	x	
***	Educators		x				x		x		x	x	
**	Other Researchers	x			x			x	x			x	
**	General Public								x	x		x	
*	Other spokespeople for efficiency in Hi-Tech facilities (e.g., Labs 21 & USGBC)	x	x	x	x		x	x	x			x	

Note: For first column, one star = lease emphasis; five stars = greatest emphasis

OUTCOMES

Description of Work

Many task-specific technology transfer efforts were covered in their respective Task Description documents. The following discussion pertains to the overall plan and complementary crosscutting activities involving particular audiences and/or communication channels. Special emphasis has been given to the results of the demonstration projects and associated "Best Practices" findings. A comprehensive list of downloadable publications is located here: http://hightech.lbl.gov/library.html

- Establish "Basic Infrastructure"

  Included preparation of business collateral: PowerPoint templates; project and public web sites, etc., as well as identification of key audiences and communications channels (see previous table). This task also involved establishing and maintaining a "Stakeholder" database (currently ~340 names) that has been used for regularly broadcasting outreach products to key sub-groups such as trade press editors, facility managers, etc.

- Interact with Project Advisory Committees (PACs)

  As leaders in their industry, the PAC members have first-hand insight into the sectors in question and ways of effectively targeting and communicating our results. The PACs are also comprised of important constituents for the project results, and thus the PAC meetings—to be held periodically throughout the course of the project—will effectively serve to relay information to this group.

  PAC meetings:

  - (1) Cleanrooms and Laboratories 13 Nov 2003 and 13 May 2004
    (2) Data Centers 16 Dec 2003 (second meeting TBD)

  Establishment of Intranet site for PACs and CEC
(see http://eetd.lbl.gov/hightechintranet)

  - Sponsor, Contractor, and Consultants information; Contact List; Minutes of PAC meetings; Deliverables; Task Descriptions; Presentation Archive.

- Produce Self-published Products

  CEC core deliverables & administrative reporting
  Best-Practices Briefs (large audience)
  Posters (for permanent or portable display)
  LBNL technical reports (small audiences)
  PowerPoint Presentations (mid-sized audience)
  Web content (see below; large audience)

  Brief articles (<500 words) were periodically be published in LBNL's "EETD Newsletter", which reaches approximately 10,000 readers. These also served a dual function of forming the basis for longer trade press articles. EETD articles are also automatically posted on the web.

- Disseminate Information Through the Trade and Technical Media

We periodically issued media advisories to trade journals (as distinct from full articles, which will also be done). Many trade journals do not publish articles from outside authors, but ran "news" stories on topics of interest to their readership.

In addition to publishing many LBNL Technical Reports, we continually sought publication opportunities and other forms of collaborations with trade journals and trade media outlets, i.e.:

- *A2C2*
- *Ascribe Newsletter*
- *Association of Facility Engineers Newsletter*
- *Automated Buildings*
- *Bladeletter*
- *Byte.com*
- *CIO Magazine*
- *Cleanrooms*
- *Consulting-Specifying Engineer*
- *Control Engineering*
- *The Data Center Journal*
- *EETD News*
- *Energy & Power Management*
- *Energy & Power*
- *Energy Vortex*
- *Environmental Valuation and Cost-Benefit News*
- *E-Week*
- *Engineered Systems*
- *FEMP Focus*
- *Forbes*
- *GCN*
- *Green Supply Line*
- *Heating, Piping, and Air Conditioning*
- *The Hosting News*
- *HPAC Engineering*
- *Laboratory Equipment*
- *Power Management Designline*
- *Power Sources Manufacturers Association*
- *SFL.org*
- *Sustainable Industries Journal*
- *TechTarget*
- *TMC.net*
- *Zerodowntime*

We published in various appropriate professional and technical journals or conferences. Examples:

- *ASHRAE Journal*
- *ASHRAE Transactions*
- *Energy*
- *IEEE Proceedings of the IEEE. International Symposium on Electronics and the Environment*
- *IEST Journal*
- *Proceedings 2004 Semiconductor Equipment and Materials International, SEMICON West*
- *Proceedings of the 2004 ACEEE Summer Study on Energy Efficiency in Buildings*
- *Proceedings of the 2006 ACEEE Summer Study on Energy Efficiency in Buildings*
- *Proceedings of The 51st ESTECH Conference*
- *The ASCE Journal of Infrastructure Systems*

Based on a suggestion by Genentech, affirmed by the PAC, we prepared a White Paper for corporate decision-makers on the energy and non-energy benefits of increasing the energy performance of high-tech facilities. Target audiences include CFOs, Supply Managers, Purchasing Managers, and those issuing corporate energy policies. Authors: Gary Shamshoian (Genentech), Michelle Blazek (AT&T), Phil Naughton (SEMETECH), Bob Seese (Critical Facilities Associates), Evan Mills and Bill Tschudi (LBNL). Some of the underpinnings for this are provided in the major report "Business Case for Sustainable Design in Federal Facilities" (Placett et al. 2003). A summary of this work was subsequently published by *Forbes*.

- <u>Seek Visibility Through The Popular Media</u>

  LBNL's news release service was used at selected points where we had particularly notable results. This served as the means of reaching general audiences. We were very successful in having our work reported in the popular or wide-circulation media, including:

  - *CNET News*
  - *CNN Money*
  - *Computerworld*
  - *East Bay Business TImes*
  - *Financial Times*
  - *Information Week*
  - *Market Wire*
  - *Nature*
  - *Red Herring*
  - *Scientific American*
  - *SFGate.com*
  - *Technology Review*
  - *Wall Street Journal*

- Distribute Results Through The Internet: Web & E-mail

We designed and built out a one-stop "umbrella web site" integrating our currently disparate communication products and results (e.g. cleanrooms, FFU draft procedure, fume hoods, benchmarking, Labs 21, etc.). All written products (reports, presentations, etc.) have been posted/linked there. http://hightech.lbl.gov

Publish an electronic newsletter on the Berkeley Hood, entitled "News from the Hood", a "one-screen", periodic publication describing new developments in the high-performance fume hood project (existing mailing list; several "issues" already published, Current circulation approximately 200. This evolved into a more comprehensive newsletter, entitled *High-tech News*, which is distributed to about 340 people plus about 74 media reporters see http://ateam.lbl.gov/hightech/fumehood/Newsletter.html).

We published detailed Benchmarking information for each of our three major facility types (Labs, Cleanrooms, and Datacenters).

We created simplified calculators for Fume Hoods; UPS/PS). Phoenix Controls is utilizing the project's fume hood energy calculator (http://xx) as a tool to help their sales representatives explain energy efficiency options to their customers. PG&E is also encouraging its customer representatives to use the tool.

- Distribute Results Through Conferences, Meetings, and Workshops

Project staff presented results from this project at 96 meetings (as of June 2007). We regularly partnered with industry groups to find experts to contribute to the discussions.

A comprehensive list of presentation titles and venues (including downloadable versions of the talks) can be found here: http://hightech.lbl.gov/events.html

Venues included:

- 1st Annual Roche Energy Conservation Summit
- 24 x 7 Exchange
- ACEEE Summer Study on Energy Efficiency in Buildings
- ACEEE Summer Study on Energy Efficiency in Industry
- Advanced Chemical Safety Meeting
- ASHRAE Annual Meeting
- Austin Energy Rebate Program Rollout for Data Centers
- Building a Sustainable Campus Community (UCSC)
- Building Owners' and Tenants' Summit on Biotech, Medical & Pharmaceutical Lab Facilities
- CBI Forum on Laboratory Construction
- Consortium of Information Systems Executives
- Critical Facilities Roundtable
- Data Center Decisions
- Datacenter Dynamics
- Energy Efficiency Working Group (EEWG)

- Energy Star Pharmaceutical Focus Meeting
- Enterprise Servers and Data Centers: Opportunities for Energy Efficiency
- EPA Workshop on Data Center Performance Metrics
- ESTECH 2007
- Gene Acres Conference
- High-Performance Computing Conference
- IEST Meeting
- International Society for Pharmaceutical Engineering
- Labs 21 Annual Conference
- Labwise Conference
- National Conference on Data Center Infrastructure - Outlook and Technology
- NextGen Data Center
- Northwest Energy Efficiency Alliance Microelectronics Workshop (NEEA)
- Silicon Valley Leaders Group
- Silicon Valley Manufacturers Group
- Trends in Cleanroom Technology and Energy Saving Opportunities: Best Practices for Energy Efficient Design, Construction and Operation of Minienvironment/Cleanroom
- Turnkey Conference
- US Green Buildings Council
- Company-based Meetings
    - AMD (Critical Facilities Roundtable)
    - Abjenix (Biotech Facilities Association Meeting)
    - Cisco Systems (data centers workshop)
    - Sybase (cleanroom and datacenter workshop)
    - Emerson Network Power
    - Intel Developers Forum
    - Intel Technology Symposium
    - Pentadyne (DC Powering of IT Equipment Demonstration meeting)
    - Quantum Consulting (project overview)
    - Sun Microsystems (direct current for data centers)

  Utility Venues:
    - Southern California Edison
    - Sacramento Municipal Utility District
    - Pacific Gas and Electric Company
    - Southern California Gas Company

A variety of activities involved working directly with California Utilities:

- LBNL established and maintained contact with the Emerging Technologies Coordinating Council (ETCC). This included attendance (by invitation) at several of their meetings.
- LBNL's work informed the setting of PG&E's application to the CPUC for their 2006-2008 programs, with explicit mention of new programs to promote Fan Filter Units

Some task leaders organized topical workshops as part of their work, e.g.:

- Mini environments
- Demand Controlled Filtration
- Power Supplies
- Data Center Training
- Cleanroom Training

- Transfer Technology Through Standards & Guidelines; Service on Technical Committees

Various tasks in the project utilized this "channel" for technology transfer. Activities involved in these issues include the Berkeley Hood, the fan filter unit, and mini-environment tasks. Measurement/rating protocols were developed within the UPS, IT Interface, and Efficient Power Supplies tasks. In addition to influencing end-users, standards and guidelines also provide a tool for intermediaries (e.g. utilities) who need reference points for defining activities such as incentive programs, labeling, etc.

Product Rating/Labeling: ENERGY STAR served as another natural form of guideline and several of the tasks that produced analyses and data (e.g. for power supplies) that could support eventual inclusion of high-tech products and/or facilities in the ENERGY STAR programs.

Voluntary Standards and Guidelines (examples for fume hoods and data centers):

- ASHRAE: Work with SPC-110 sub-committee is ongoing. SPC-110 is revising and updating the ASHRAE 110-1995 method.
- ASHRAE: Work with Technical Committee TC 9.9 which issued thermal guidelines for data center facilities and has an ambitious program to address many areas of data center HVAC.
- ANSI/AIHA: Results from side-by-side testing and dynamic testing protocol were helpful to refine their performance-based standard. LBNL work may influence future updates of their standard.
- ACGIH: Results from hood testing were helpful in refining their standard. LBNL work could influence future updates of their standard.

- Industry Technical Committee Work: Project participants serve on a number of industry technical committees, e.g., which provide a key channel for outreach:

  - ACGIH: Results from hood testing will be helpful to refine their standard. LBNL work could influence future updates of their standard [Bell]
  - ANSI/AIHA: Results from side-by-side testing and dynamic testing protocol will be helpful to refine their performance-based standard. LBNL work could influence future updates of their standard [Bell]

- ASHRAE: Work with SPC-110 sub-committee is ongoing. SPC-110 is revising and updating the ASHRAE 110-1995 method. Sub-committee involvement provides an opportunity to have input during the updating process so that the revised method will embrace new fume hood technologies.
- ASHRAE: Work with Technical Committee TC 9.9 which issued thermal guidelines for data center facilities and has an ambitious program to address many areas of data center HVAC.
- ASHRAE Clean-spaces technical committee TC 9.11 and also a corresponding member of the "Mission Critical facilities committee" TC 9.9 (data centers) [Tschudi]
- ASHRAE TC 9.10 (laboratories) [Bell]
- Critical Facilities Round Table (CFRT) Committee on data centers. [Tschudi, Greenberg, Xu]
- IEC technical committee 22: Power electronics. This technical committee is responsible for International Standard IEC 62040-3 on UPS performance and also on power supplies [Mansoor] - Pending
- IEEE 1621 [Chair]. Power Control User Interface Standard [Nordman]
- IEEE Taskforce on Single Phase Harmonics (Will be relevant for the power supply part of the work) [Mansoor]
- IEEE Industrial Application Society (Will be relevant for any power electronics related work both for UPS and Power supplies for data center industry) [Mansoor]
- IEST Working Group 12 on cleanroom design considerations. [Xu]
- IEST Working Group 36 on fan-filter unit test procedure. [Xu]
- IEST WG on Mini environment, which has close relevance with WG12. [Xu] - Pending
- Lead author of the proposed global test procedure for measuring the energy efficiency of single voltage external AC/DC power supplies. This procedure will be used by EPA, the CEC, CECP, CNIS, and the European Code of Conduct process for measuring power supply efficiency. [Calwell]
- USGBC LEED for Labs committee [Mathew]
- We Contributed to the ASHRAE. "Design Considerations for Data and Communications Equipment Centers." American Society of Heating, Refrigerating, and Air-Conditioning Engineers, Technical Committee 9.9. (Chapter 13 deals directly with energy efficiency).

- Distribute Results Through Curriculum & Training

  - We organized our datacenter research materials into an extensive special self-paced "training" website (co-funded by FEMP). See http://hightech.lbl.gov/DCTraining/top.html
  - Several California utilities invited us to lead or participate in training sessions for their constituents. These are cataloged at: http://hightech.lbl.gov/events.html
    Within a project separately funded by the National Science Foundation, we established involvement in vocational training via an LBNL/Peralta-Laney Memorandum of Understanding, which includes a new building operator curriculum for program at Oakland's Laney College (Peralta Community College District), in collaboration with three other California Community colleges (San Jose City College, Fresno City College and Los Medanos College). The National Science Foundation funded this activity at ~$900. A Focus Group was conducted to ascertain industry needs. Participants cited the need for training applicable to high-tech facilities. Laney will create a short-course (2-4 weekends) on high-tech facilities for regular students and continuing education, if LBNL provides some instructors and/or content. This provides an opportunity to disseminate our findings (e.g. best practices) to present and future practitioners. Several utilities have expressed interest in collaborating, and using the resulting curriculum in courses at their Energy Centers.

Implementation Team and Responsibilities

1. Task Leader: Evan Mills
2. Other Staff: Bill Tschudi, Dale Sartor, plus task-specific leaders
3. Other partners: Media Allies (e.g. trade magazine editors), utilities collaborating in trainings, etc.

Attachment XI: Demonstration of Standardized Testing of Fan-filter Units

LBNL – 63420

Demand Controlled Filtration in an Industrial Cleanroom

David Faulkner, Dennis DiBartolomeo, Duo Wang

Environmental Energy Technologies Division
Indoor Environment Department
Lawrence Berkeley National Laboratories
Berkeley, CA 94720

September 2007

This work is governed under U.S. Department of Energy Contract No. DE-AC03-76SF00098 with the Regents of the University of California and funded by the California Energy Commission's Public Interest Energy Research (PIER) Program through agreement with the Regents of the University of California administered by the California Institute for Energy and Environment (CIEE) (under Memorandum Agreement NO. MR-0312).

Demand Controlled Filtration in an Industrial Cleanroom

David Faulkner, Dennis DiBartolomeo, Duo Wang

Abstract

In an industrial cleanroom, significant energy savings were realized by implementing two types of demand controlled filtration (DCF) strategies, one based on particle counts and one on occupancy. With each strategy the speed of the recirculation fan filter units was reduced to save energy. When the control was based on particle counts, the energy use was 60% of the baseline configuration of continuous fan operation. With simple occupancy sensors, the energy usage was 63% of the baseline configuration. During the testing of DCF, no complaints were registered by the operator of the cleanroom concerning processes and products being affected by the DCF implementation.

Introduction

Demand controlled filtration (DCF) is a method of saving energy in cleanrooms by decreasing the recirculation fan speed when demand for filtration is low. A particle counter or an occupancy sensor can measure the demand for filtration. Another method of saving energy in cleanrooms is the use of programmable timers that decrease recirculation fan speeds when occupants are not expected, such as nights and weekends.

DCF should not be confused with demand controlled ventilation, which is the control of outside air entering a building based on the demand due to occupancy and/or processes occurring in the building.

In a past study, we found that industrial cleanroom users were hesitant to implement DCF, often because of concern that large changes in the recirculation fan speeds would produce a burst of particles discharged into the cleanroom. We have not seen this phenomenon in our previous or current study. Another concern is that for large manufacturing facilities, the potential loss of product outweighs any potential energy savings. There are many small research or industrial facilities that are not as product-oriented that can benefit from the potential energy savings from implementing DCF.

In the past few years there has been gathering interest in saving energy in cleanrooms and some businesses are experimenting with DCF. There is a large Southern California manufacturer with Class 100 (ISO Class 5) and 10,000 (ISO Class 7) cleanrooms covering about 15,000 ft^2. During non-working hours, the make-up and recirculation air handler fans at this facility are turned down to recirculate less air. A reduction in fan speed of about 30% at night, when the space is unoccupied, reduces the measured recirculation air handler power consumption from 32 kW to less than 9kW – a reduction of 75%.

Another large manufacturing facility in Southern California is in the process of implementing active, real-time DCF using particle counters. The facility manager has received rebates from the local utility for implementing DCF. The particle counters have been purchased and preliminary testing has been done. The next steps are installing the particle counters and implementing the control routines.

Finally, a large East Coast university recently modernized a building with cleanrooms. In the process they installed particle counters in every cleanroom area and plan to implement DCF.

Energy savings potential

The potential for energy savings is tremendous as fan energy has an approximately cubic relationship with fan speed. Thus, even a small reduction in fan speed can translate into large energy savings. Cleanrooms employing DCF can achieve lasting savings by implementing active control of the recirculation fans based on inputs from either a particle counter or an occupancy sensor.

Also, cleanroom operators may be able to participate in demand response programs offered by utilities. During high electricity demand days, such as hot summer days, many utilities offer incentives to their customers, such as rebates or price reductions, to voluntarily reduce electricity demand at their facilities. Typically these programs ask customers to reduce electricity consumption for periods of 1-6 hours. As long as the cleanroom fans are equipped with variable frequency drives (VFD), then the recirculation fan speeds can be adjusted lower during critical demand days. A 10-20% reduction in recirculation fan speed can translate to a reduction in energy use of over 40%.

Methods

Facility

Experiments were conducted to evaluate air quality in an industrial cleanroom in the San Francisco Bay Area while DCF was implemented. The cleanroom that was studied was one of a suite of four adjacent cleanrooms, each with different classifications from Class 10 (ISO Class 4) to Class 10,000 (ISO Class 7). The suite of cleanrooms was designed so that the inner Class 10 (ISO Class 4) cleanroom (Group-5) is at a higher pressure than the adjacent Class 100 (ISO Class 5) cleanroom (Group-6), which in turn is at higher pressure than the adjacent Class 1,000 (ISO Class 6) cleanroom (Group-3), etc., see Figure 1. There are open grilles for air to flow from the Group-5 cleanroom to the Group-6 cleanroom to the Group-3 cleanroom and finally to the Group-2 cleanroom. The cleanroom in which the experiments were conducted was the Group-2 area and was a Class 10,000 (ISO Class 7).

The floor area of the entire suite is about 1600 ft^2 and the floor area of the Group-2, the experimental cleanroom, was 600 ft^2. The layout of Group-2 is shown in Figure 2. Each cleanroom had air recirculated by ceiling-mounted 2 ft x 4 ft fan filter units with HEPA filters. The Group-2 cleanroom had 16 fan filter units covering about 20% of the ceiling area. There were 4 hoods in the middle of the cleanroom as well as 2-3 flow hoods at each end of the cleanroom. Airflow through these hoods were not controlled as a part of our study.

Outside make-up air was controlled by a separate ventilation system and was not controlled in our study. Environmental conditions, temperature and humidity were controlled by a separate control and also were not part of our study.

Before our study, a laptop computer using a proprietary software program controlled the cleanroom fan filter units. The software allowed control of the fan filter units individually or in groups. This software was used to lower the fan speeds to 50% at night (22:00 to 6:00) and on weekends. This software and corresponding laptop would occasionally fail and the fan filter units would run at the speed of the last valid command; the fan filter units are designed to always run at the last valid command. Thus, occasionally the fan filter units for all of the cleanrooms would run for several days at 50% speed before the problem was discovered.

After implementing the night and weekend energy reduction, particle concentrations at 4 locations in the cleanroom suite were recorded continuously for background data. The particles concentrations were measured with MetOne 237B particle counters (Hach Ultra, Grants Pass, OR) capable of counting particles in 6 size bins from 0.3 to 5 microns at a sample flow rate of 0.1 cfm. Some particle counters were moved to different locations within the cleanroom suite during the background measurements. Based on this data and the desires of the facility manager, the experiments were performed in the Group-2 cleanroom which was the Class 10,000 (ISO Class 7) cleanroom.

After background particle measurements were recorded, a particle counter was placed in the study cleanroom (Group-2) near the center of the room on a countertop about 3 ft above the floor. To implement DCF, the particle counter was connected to a laptop computer running an off-the-shelf software program, LabVIEW Version 5.1 (National Instruments). Control routines in LabVIEW had been developed to record particle counts, and use the particle count information to decide whether to increase or decrease fan speeds. The particle counter sent counts to the laptop every 30 seconds and the particle counts in the 0.3 and 0.5 micron size bins were compared to established upper limits. The fan speed was either increased or decreased until either the maximum or minimum fan speed was reached. A record of the fan speed, particle counts and other parameters were recorded every 60 seconds in LabVIEW.

After implementing DCF with the particle counter, we implemented DCF with control based on occupancy. The strategy was simple; if there were occupants detected, then the fan filter units were set to the programmed maximum speed, and after 30 minutes of no occupancy detected, the fans were set to 50% of programmed maximum speed.

There were 6 wireless infrared battery-operated occupancy sensors placed in the cleanroom. The 6 sensors were mounted on the walls so that each work area and traffic lane was covered by at least one occupancy sensor. A wireless signal was sent from the occupancy sensors to the laptop and the above strategy was implemented.

The data from the particle counter was converted to an equivalent ISO Class using the following equation (ISO 14644-1: 1999(E)):

$$C_n = 10^N \left(\frac{0.1}{D}\right)^{2.08} \qquad (1)$$

where,

C_n is the maximum permitted concentration (in particles per cubic meter of air) of airborne particles that are equal to or larger than the considered particle size. Value is rounded to the nearest whole number, using no more than three significant figures.

N is the ISO classification number, which shall not exceed a value of 9. Intermediate ISO classification numbers may be specified; with 0.1 the smallest permitted increment of N.

D is the considered particle size, in micrometers, which was 0.3 µm in this study.

0.1 is a constant, with a dimension of micrometers.

With algebraic manipulation equation (1) can be rearranged to solve for N, the ISO classification number:

$$N = \log_{10} C_n + 2.08(1 + \log_{10} D) \qquad (2)$$

Results

The data presented in Figures 3-6, show the ISO Class for the particle size of D = 0.3 µm when the Group-2 cleanroom was using DCF controlled by particle counts and by occupancy sensors. The concentration, C_n, is calculated from the counts measured by the particle counter.

During these experiments, the fan filter units were running at about 77% of maximum speed. Thus, in the following discussion, full speed means 77% of maximum speed and half speed means about 38% of maximum speed. Also, in all plots, the greatest fan speed will be 77%.

Figure 3 shows fan speed and particle counts in the Group-2 cleanroom (converted to ISO Class) versus time for a weekday while the fan speed was controlled by particle counts. From the graph, it is clear that the fans are actively being controlled by the particle counts most of the work day from about 6:00 to 17:00. The ISO Class stays around 3.5 to 4 during the working hours. At night the particle counts increase, despite there being no occupants in the cleanroom. The reason for this increase in particle counts while there were no occupants has not been explained by the building manager/cleanroom operator. Pressure measurements made between cleanrooms showed that at night, some of the pressure differences between cleanrooms reduced in magnitude and some actually reversed direction. It was thought that this was possibly due to changes in the whole building HVAC system. The building manager did not have details concerning nighttime setback conditions in the whole building HVAC system. Each morning, as the whole building HVAC system reset and the cleanroom fan filter units increased, the cleanroom, returned to “clean” conditions very quickly.

Figure 4 shows the fan speed and particle count data for a weekday while the cleanroom had DCF control based on occupancy sensors. The fan filter units were either at full speed (77% of maximum) while occupants were detected or at “half” speed (~38% of

maximum) with no occupants detected in the last 30 minutes. After noon, it can be seen that no occupants were detected on 2 or 3 occasions. In response, the fan filter units reduced to half speed and the ISO Class (based on particle counts) increased rapidly from about 3 to 5. But each time occupants were detected, the particle counts decreased quickly.

The rapid response in particle counts to the change in fan filter speed is seen in Figure 5, which shows similar data as in Figure 4, but over a shorter time span. The ISO Class changes from about 5 to about 3 in approximately 3-4 minutes as the fan speed increased from half speed to full speed.

Finally, Figure 6 shows the fan speed and particle count for a weekend day with DCF control based on occupancy sensors. As soon as an occupant was detected around 9:00, the fan speed increased to full speed and the particle counts decreased by about a factor of 300 (ISO Class 6 to Class 3.5). Around 10:10, an absence of occupants was detected and fan speeds were set to "half" speed and the particle counts increased. About 10 minutes later, an occupant was detected and the fan speeds increased to "full" speed and particle counts again decrease, as indicated by the ISO Class.

Energy values are shown in Table 1, for a typical week. There are 5 different configurations that are compared. Two of the configurations are used as baselines for comparison with various energy saving configurations. The baselines are 24/7, which means the fans are running all the time, and the other baseline is the night and weekend setback, in which the fans were reduced at night and on the weekends. Comparing the two baselines, the night and weekend setback saves about 28% in energy as compared to the fans running constantly.

The use of the particle counter to implement DCF saved more than the occupancy sensors, but not by much, 40% vs 36% as compared to 24/7 as a baseline and 17% vs. 12% as compared to the night and weekend setback as a baseline. Thus, using a few occupancy sensors at $30 each is less expensive than a particle counter that may cost $3,000 or more.

Table 1. Energy use (kW-h per week) and percent savings due to demand controlled filtration.

Configuration	**kW-h per week**	**Compared to 24/7**	**Compared to Night & Weekend Setback**
24/7	151		
Night & Weekend Setback	109	28%	
Occupancy Sensor	96	36%	12%
Particle Counter	91	40%	17%
"Half" Speed	70	54%	36%

Conclusions

The data collected indicate two major conclusions. First, DCF was implemented without disruption to the processes in the cleanroom. This implementation was relatively easy and inexpensive, less than $10,000 in equipment. Second, recovery of particle concentration from a "dirty" condition to a "clean" condition occurred in a matter of minutes. Thus, those cleanroom operators in industry need not be concerned about the removal of particles taking days or hours after a reduction in the fan speeds.

The energy savings realized were about 40% with either method as compared to running the cleanroom fans constantly, 24/7. Savings were about 15% with either method as compared to night-time and weekend setback.

Acknowledgements
Bill Tschudi and Mark Sippola of Lawrence Berkeley National Lab, Larry Chu and Peter Rumsey of Rumsey Engineers, and Reza Kazemi of Environmental Filtration, Inc.

References

Faulkner, D., Fisk, W.J., and Walton, J.T. (1996) Energy Savings in Cleanrooms from Demand-Controlled Filtration", Journal of the Institute of Environmental Sciences 39(2): 21-27. LBNL-38869

International Standard, Cleanrooms and associated controlled environments – Part 1: Classification of air cleanliness, ISO 14644-1: 1999(E)

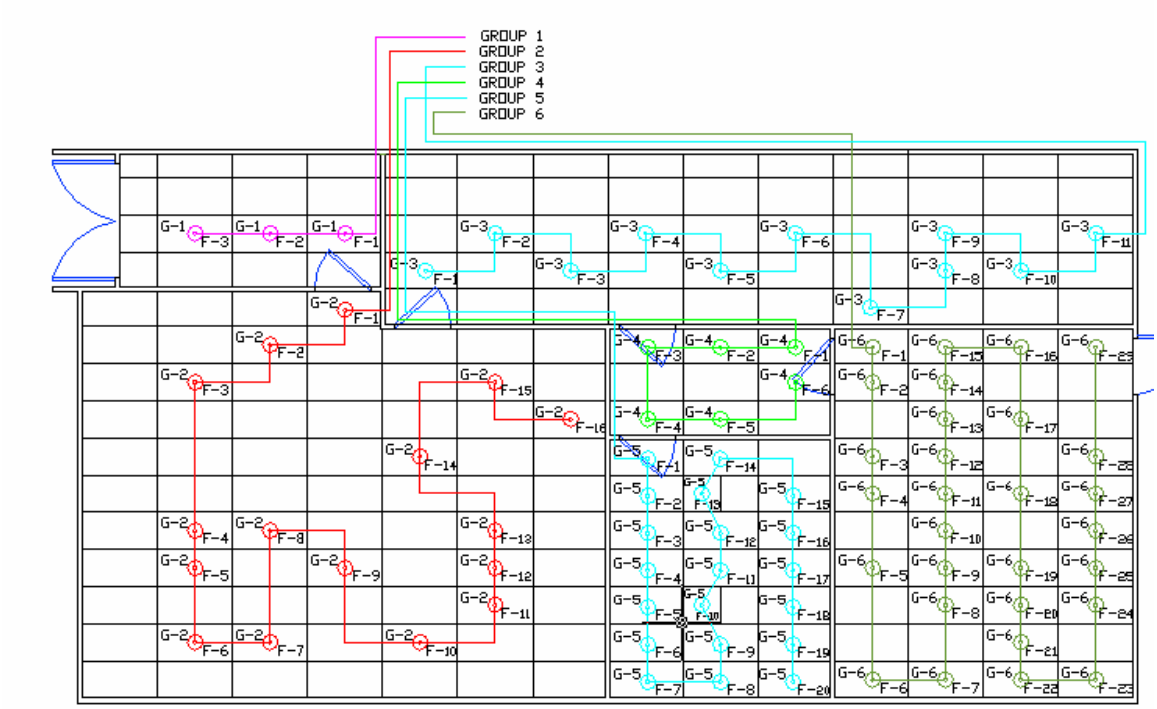

Figure 1. Layout of the suite of cleanrooms in the studied industrial facility. G-5 (Group-5) is an ISO Class 4 cleanroom. G-2 (Group-2) is an ISO Class 7 cleanroom. G-1 (Group-1) and G-4 (Group-4) are gowning rooms. G-3 (Group-3) is an ISO Class 6 cleanroom and G-6 (Group-6) is an ISO Class 5 cleanroom. Fan filter units are labeled as F-1, F-2, etc. Each rectangle is 2 x 4 ft.

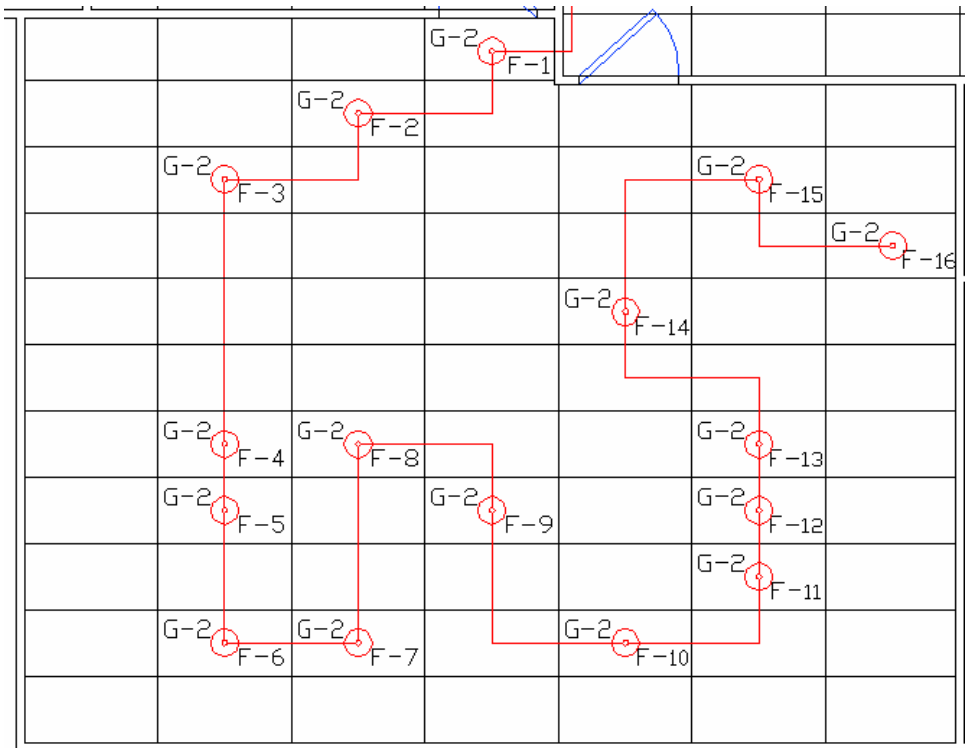

Figure 2. Layout of fan filter units in study cleanroom, Group-2. There are 16 fan filter units, shown as F-1 to F-16. Each rectangle is 2 x 4 ft.

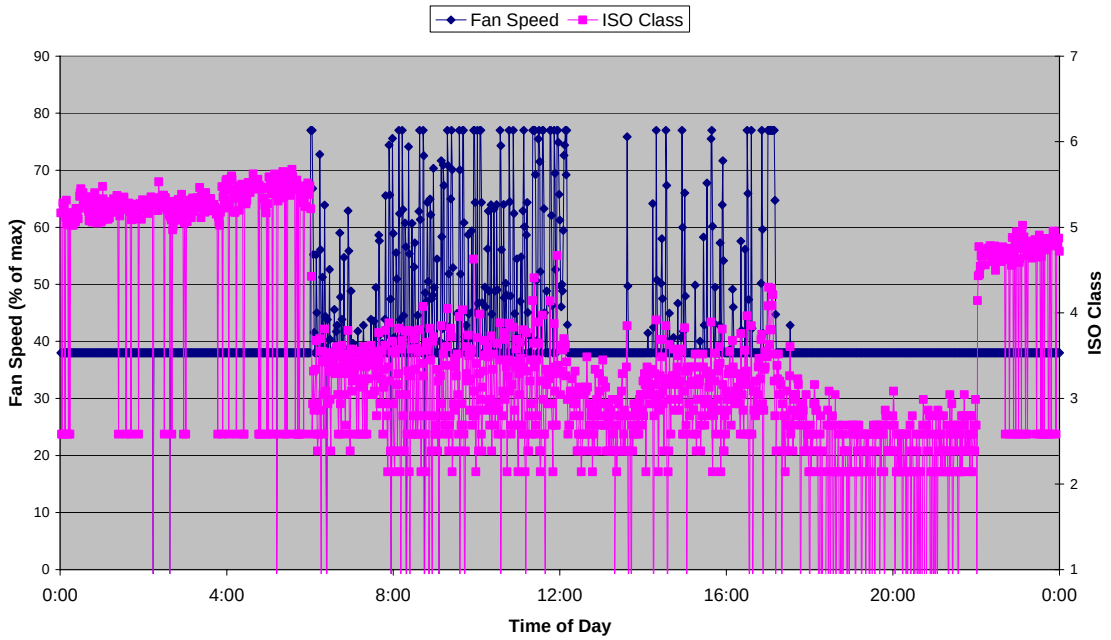

Figure 3. Fan speed and particle counts (converted to ISO Class) versus time for a weekday with DCF control based on particle counts. The cleanroom classification is based on the particle concentration at 0.3 micrometers.

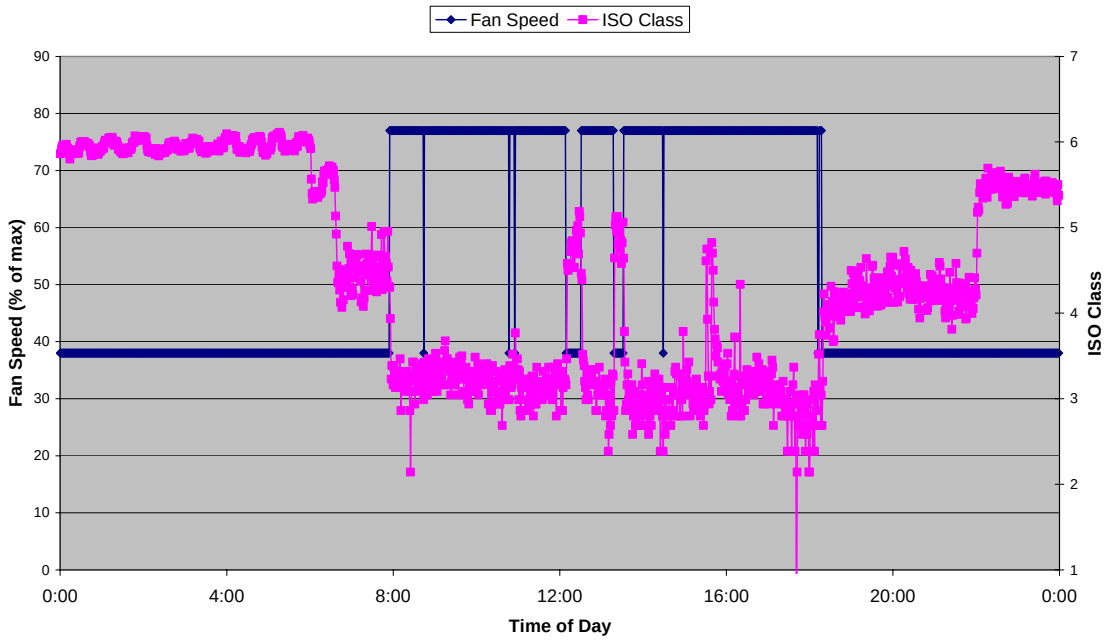

Figure 4. Fan speed and particle counts (converted to ISO Class) versus time for a weekday with DCF control based on occupancy sensors. The cleanroom classification is based on the particle concentration at 0.3 micrometers.

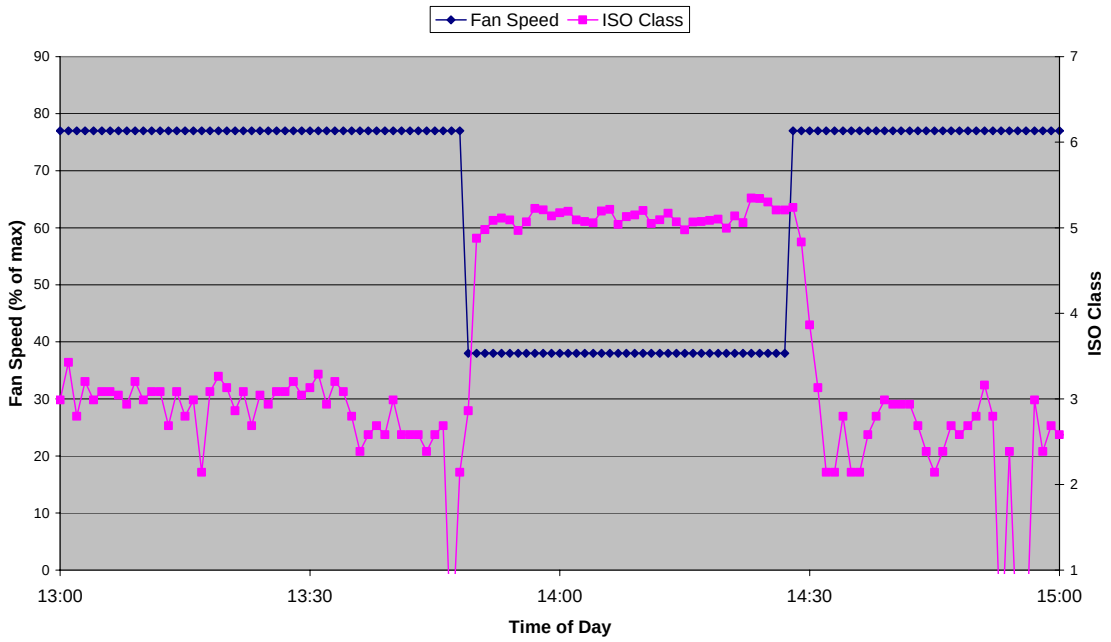

Figure 5. Fan speed and particle counts (converted to ISO Class) versus time for a weekday with DCF control. There is a rapid return to “clean” conditions within 3-4 minutes (data recorded every 1 minute) after the fans increase from a low speed to a high speed. The cleanroom classification is based on the particle concentration at 0.3 micrometers.

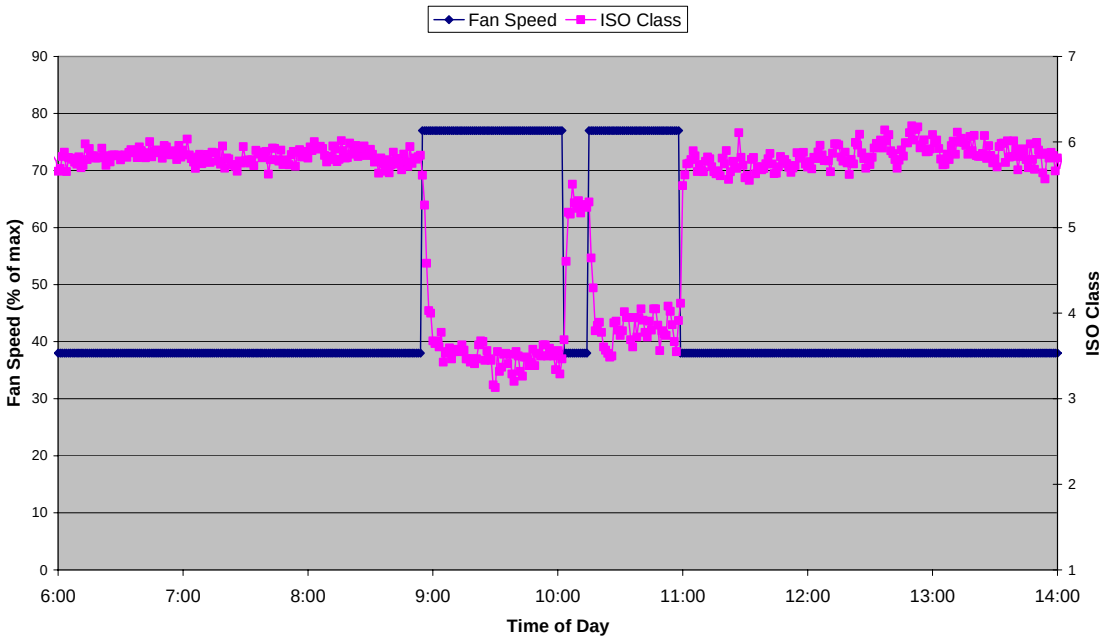

Figure 6. Fan speed and particle counts (converted to ISO Class) for a weekend day with DCF control based on occupancy sensors. Changes in particle concentrations, as expressed by the ISO Class, are very quick. The cleanroom classification is based on the particle concentration at 0.3 micrometers.

Attachment XII: Demonstration of DC Powering of IT Equipment

LBNL- 61684

Demonstration of Advanced Filtration Technologies:

Developing Energy-rebate Criteria through Performing Standard Laboratory Tests and Statistical Analyses

Final Report (LBNL- 61684)

To

California Energy Commission

2007

Tengfang Xu and Duo Wang

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Contents

1. Abstract 5

2. Background 6

3. Goals and Tasks 6

4. Project activities and deliverables 7

5. Report on Findings 12

5.1 Objective 12

5.2 Designing Test Rig 12

5.2.1 Sample fan-filter unit 13

5.2.2 Designing test rig 13

5.3 Developing and Refining Laboratory Method and Procedure 14

5.4 Developing Polynomial Regression Models to Characterize Energy Performance 16

5.5 Findings 18

5.5.1 Airflow rate and pressure differential 18

5.5.2 Total electric power demand 20

5.5.3 Total pressure efficiency 23

5.5.4 Summary of findings 25

5.6 Developing Energy Incentive Criteria 26

5.6.1 Prescription-based approach 26

5.6.2 Performance-based approach 36

5.6.3 Recommending criteria options for energy incentive programs 37

6. Summary 39

7. Acknowledgements 42

8. References 42

9. Appendices 45

9.1 Initial Test Setup 45

9.2 Test Equipment and Device 48

9.3 Leakage Detection and Quantification 50

9.4 Test setup for evaluating repeatability 52

9.4.1 Repeatability of measured performance at different flow restrictions created by placement of an FFU in opposite directions 52

9.4.2 Repeatability of measured performance at different ambient conditions 54

9.4.3 Repeatability of measured performance measured with different lengths of the ductwork leading to an FFU 55

9.5 FFU Characterization Reports 56

Demonstration of Advanced Filtration Technologies:

Developing Energy-rebate Criteria through Performing Standard Laboratory Tests and Statistical Analyses

Tengfang Xu and Duo Wang

1. Abstract

Fan-filter unit systems used for re-circulating clean air in cleanrooms are gaining popularity in California as well as in rest of the world. Under normal operation, fan-filter units require high power demand, typically ranging from 100 to 300 W per square meter of cleanroom floor area (or approximately 10-30 W/ft^2). Operating 7 by 24, they normally consume significant electric energy, while providing required contamination control for cleanrooms in various industries. Previous studies focused on development of a standard test procedure for fan-filter units. This project is to improve the methods, and develop new information to demonstrate the methods can be used to assist the industries to apply more energy-efficient fan-filter units in cleanrooms.

Specifically, this project expands previous developmental activities of a test protocol to characterize a pool of 17-sample fan-filter units (FFUs) recruited from Asia, Europe and North America. Through laboratory experiments and modeling characterization, the project develop and demonstrate means of identifying and applying existing or new filtration technologies with higher energy efficiency in the market. All the FFUs had a nominal size of 61-cm-by-122-cm (2-ft by 4-ft). We established a new testing facility, performed new laboratory test using a refined test method, conducted data analyses, developed models to characterize energy performance of the 17 FFUs. Based upon the laboratory test results, we developed a relative ranking system to compare energy efficiency of the FFUs, and recommended options of formulating initial energy-incentive criteria for consideration and use in utility companies' future rebate programs.

The refined standard methods including the test procedure and regression models were used to characterize dynamic operation of individual fan filter units. Based upon the laboratory results, options for initial energy-incentive criteria are developed for utility companies. The new rebate program(s) can then become a strategic means to identify and promote applications of energy efficient FFUs. In addition, the results can also be used for future product development toward higher-efficiency. FFUs with higher energy performance are identified in the project, which also demonstrates that significant reduction in energy use in cleanroom systems is possible through

adopting energy efficient FFUs. Lastly, this report provides recommendations for future research, development, and implementation of high-performance FFUs.

2. Background

Cleanrooms are essential to a variety of industries important to the state's economy. Cleanroom facilities are characterized by large base-loads operating 24 hours a day with energy intensities much higher than typical commercial buildings. Recognizing that energy consumed by cleanroom buildings in California represents an attractive opportunity for energy savings, Lawrence Berkeley National Laboratory (LBNL) performed research and outreach activities in this sector, and developed technologies, tools, and strategies addressing various aspects of the overall efficiency opportunity in cleanrooms.[1]

In the course to reach their energy savings potential, we have identified areas such as fan-filter unit systems used in re-circulating clean air in cleanrooms. Prior project focused on development of a draft standard test procedure for fan-filter units. FFUs are gaining popularity for use in California cleanrooms as well as in rest of the world. Demonstration of energy efficient filtration technologies is vital to improve application of existing or new energy-efficient technologies and to develop strategies to transform the market toward higher energy efficiency.

The overarching goal of the research was to improve the energy efficiency of cleanrooms and to increase energy savings and demand savings in cleanroom systems. This would provide financial and energy benefits to California companies. Specifically, this project expands the development and outreach activities to promote energy efficiency for cleanrooms that use 61-cm-by-122-cm (2'-by-4') fan-filter units, through demonstrating the development of new standard methods to characterize FFUs, and application in developing energy-rebate program criteria.

3. Goals and Tasks

This project's goal is to demonstrate energy-efficient fan-filter technologies through utilizing the new standard methods developed by LBNL.[2][3] The main tasks were

- to refine the methods to characterize fan-filter unit operation under various control setting,
- to carry out fan-filter tests and to document test procedure
- to develop mathematical models for characterizing FFU energy performance
- to design and conduct additional experiments to improve understanding and robustness of the standard methods developed in this project.

Based upon the data gathered from laboratory tests and modeling, technical information was analyzed and developed. The developed information is expected to quantify energy efficiency of FFUs available in the cleanroom market, and to engage the industries including utility companies, end-users, and suppliers in promoting use of energy efficient FFUs. The results can suggest options of to recommending criteria for possible utility energy incentive programs to promote applications of energy efficient fan-filter technologies.

4. Project activities and deliverables

In order to achieve the goals, we designed and assembled a new test rig with a data acquisition system as the test facility at LBNL. We used the initial test rig to measure and evaluate functionality and performance of 17 fan-filter units. The draft test method was refined and used in this round of experiments. In the new version of standard methods, we developed and included the documentation on equipment requirements (e.g., accuracy and calibration), test rig setup and integrity diagnostics, and details of the testing procedure. The contents of new document have been continuously refined and improved through numerous experimental trials and errors.

In addition, we have been working with industry organizations (e.g., IEST, AMCA, CleanRooms Magazine, and Controlled Environments Magazine) to attempt to gain industry acceptance of the standard methods. In particular, these two leading magazines - CleanRooms Magazine, Controlled Environments Magazines supported the R&D effort by providing free of charge support, i.e., periodic calls for industry participation and updates of the FFU project. Because of collaboration, we have received a number of in-kind support from suppliers and users of FFUs in Asia, Europe, and North America.

In the meanwhile, we continued interactions with IEST contamination control division and working group committee – "WG036 Testing Fan Filter Unit." IEST is an ANSI-accredited organization actively in developing and publishing industrial standards such as RPs that guide practices in cleanroom and controlled environments. The IEST WG 036 committee consists of approximately twenty professionals associated with the industries, standard-setting entities including AMCA, ASHRAE, and government. Over the last two years, the committees were informed of the project outcomes and were engaged in providing comments. So far, the committee has been open to integrate the refined methods and procedures in a new Recommended Practice expected to be accepted and published by ANSI-Accredited IEST. The draft standard [1][2] to characterize energy performance of fan-filter units was approved and adopted by the Working Group 036 (WG036) of IEST Contamination Control Division in its Recommended Practice (RP) Draft. We also communicate with AMCA, which is awaiting official publication of Recommended Practice by IEST. The final publication timeline of RP036.1 is unclear but the committee is planning to vote the RP in November 2007. If approved, the document will undergo editorial work followed by the approval of the IEST Executive Board.

In this phase of research and demonstration, we have developed new characterization methods using new regression models to characterize energy performance of individual fan filter units. The new models allow users to quantify unit's efficiency and power demand levels and functionality. Through additional data analyses, we have developed two options of feasible energy-rebate criteria for the use of California utilities in their possible energy-incentive programs.

At the onset of the project, we evaluated advantages and disadvantages of possible test facilities (e.g., PG&E's Thermal Flux Energy Test Facility in San Ramon, UC Berkeley campus, ITRI – AMCA certified facility in Taiwan, AMCA headquarters facility in Illinois, and other testing labs). As a result, we designed and constructed of an FFU test rig and established a data acquisition system at LBNL Building 63 laboratory space. We then performed calibration and troubling-shooting of the test device, and conducted leakage tests of the FFU test rig at LBNL. Pilot tests and experiments were designed and conducted to identify potential deficiencies or pitfalls of the new test rig or equipment. We improved and validated the integrity of the test rig and identified measurement accuracies of devices.

A standard method in place will enable consistently characterizing and reporting the energy performance of fan-filter units (FFUs). LBNL carried out protocol development, performed experiments and evaluations on various fan-filter units, developed modeling methods to quantify FFU performance, and generated comparative performance information of fan-filter units. [5][6][7][8] [9] during the course of this project, we refined the methods and procedures to include more details and specific guidance as part of the standard methods. The new standard methods include mathematical regression modeling to characterize individual fan filter unit's energy performance and functionality based upon the laboratory work.

Using the new standard methods allows characterizing various FFU models or products, which in turn will help manufacturers to better understand their products, and assist cleanroom owners and designers to make informed decisions concerning selection of fan-filter units for re-circulation systems. In addition, performance data generated by the standardized testing may be used to establish l performance rating criteria should they be desired or needed by interested parties. For example, energy efficiency practice may be encouraged through utility incentive programs, which may be established based upon quantitative information from the standard laboratory tests, modeling, and statistical analyses.

We performed a series of laboratory experiments and mathematical modeling to characterize efficiency of seventeen fan-filter units, which were recruited from and donated by the industries in Asia, Europe, and North America (Canada, USA). We also conducted statistical analyses and modeling to develop a relative ranking system. Using the relative ranking system and measured performance, we have developed options for establishing initial energy-rebate criteria. The criteria can be easily used to identify higher-performers of fan-filter units for rebate qualifications, or to estimate cost effectiveness of efficient and functional units under specific

operating conditions as desired by designer, owner, and the utility company to be involved in determining the magnitude of incentive.

Project activities and deliverables are described in the following:

- Completed the design, construction, and commission of the test rig, we completed the assembling of data acquisition systems and performed trouble-shooting of the test device and facility at LBNL.

- Interacted with two leading magazines – Cleanroom Magazine and Controlled Environments: 1) announcing the demonstration project and recruiting participants and units, 2) providing updates and outcomes of the project. For example, Controlled Environment Magazine (November 2005, April 2006, November 2006) and Cleanrooms Magazine (December 2005 – March 2006, October 2006 - November 2006) reported news and/or outcomes on the PIER/LBNL FFU project. The FFU work under PIER was among the top stories of online news or Contamination Control Newsletter in February, March, October, and November 2006 (Cleanrooms Magazine). We obtained 17 fan-filter units free of charge from the anonymous industrial participants including suppliers and users of fan-filter units. Confidentiality agreements were reached between participants and LBNL. Unless permitted by the participants, the final report and future publications of the test information will not identify any of fan-filter manufacturers or products by name.

- Through interactions with the industry, including suppliers, users, and IEST, LBNL received eighteen 61-cm x 122-cm (2-foot x 4-foot) FFUs, selected, and completed the laboratory tests of seventeen FFUs. We have completed testing two extra units than originally budgeted and allocated within this phase of demonstration.

- Held meetings with PG&E to report project results and discuss ways to be used by PG&E. We presented and disseminated up-to-date test results to PG&E, arranged meetings and discussed options for possible rebate programs to be offered by PG&E, and enlisted recommendations for further R&D beyond this demonstration project.

- Interacted with users and suppliers to promote the use of the standard test method developed at LBNL. Interacted with a number of FFU manufacturers and users concerning their participation in the FFU demonstration project to improve the understanding of FFUs' energy performance and means to save energy. Seagate and Texas Instrument actually required FFU suppliers to provide test results by adopting the LBNL standard method to evaluate energy efficiency of fan-filter units. Several manufacturers are interested in participating in test bigger units, e.g., 122-cm x 122-cm (4-foot x 4-foot) FFUs. The participant in this project includes manufacturers from Asia, Canada, Europe, and US and an end-user in California.

- Performed data analysis, developed, refined, and finalized a module for reporting performance of individual FFUs to participants of the units. As part of agreement with fan-filter suppliers, LBNL produced 17 individual confidential reports and delivered them to the participants who donated their FFUs, respectively. The manufacturers and users participating in the demonstration project are now able to obtain the electric power demand under any operating conditions within the tested range and will be able to compare with other units under a selected operating condition.

- Disseminated publications (research papers and the draft standard test method) to the industry and the project participants.

    - T. Xu participated in IEST's regular WG meetings on "Testing Fan Filter Units,' and has been reinforcing the adoption of measurement methods in the LBNL standard test method.

    - T. Xu offered a tutorial class on applications of energy efficient fan-filter units at ESTECH 2006.

    - LBNL hosted laboratory visits by PG&E, CEC, ITRI, PNNL, EPA, and other project participants.

    - T. Xu present project findings at a PG&E workshop in November 2006.

- LBNL refined the existing draft standard test method (version 1.3, 2005), based upon the additional questions or comments provided by end-users and manufacturers, and the new knowledge and experience gained throughout the experiments in this phase of demonstration tests. LBNL updated the draft standard test procedure for testing fan-filter units, and made it (Version 3.0, 2007) available in the LBNL web page. The documents have stimulated significant interests from the stakeholders including users, suppliers, and IEST, the latter of which has voted to accept test methods of airflow rates, pressure, and electric power.

- LBNL developed technical information to compare FFU energy performance based upon the test data. LBNL performed a series of statistical analyses to quantify performance variation across units and operating conditions, and developed a relative ranking system. The relative ranking system is used to establish energy-rebate criteria for possible use by utility companies in establishing a rebate program. Along with the draft final report on findings, an invited article on the current study was published in Cleanrooms Magazine (November 2006), and project updates appeared in Cleanrooms Magazine, Contamination Control Newsletter, and Controlled Environments Magazine.

- Prepare a draft report of findings following completion of the testing for inclusion as an attachment to the final report.

Much of the energy in cleanrooms (and minienvironments) is consumed by 61-cm x 122-cm (2-foot x 4-foot) fan-filter units-typically located in the cleanroom ceiling (25-100% coverage)-which deliver re-circulated air to the clean spaces. LBNL has developed the first-ever standard energy test method for FFUs, which is being adopted by IEST Working Group 36 as a critical portion of industry standard- IEST RP CC036.1.

5. Report on Findings

5.1 Objective

The main objective of the project is to develop a testing rig and data acquisition, and use the new rig to characterize energy performance of 17 fan-filter units using the draft standard method developed at Lawrence Berkeley National Laboratory [1] in LBNL's Building 63 facility. In addition, we designed and conducted additional experiments to evaluate and improve the robustness of the standard methods including the laboratory testing method, equipment requirements, procedures, and modeling methods. Furthermore, additional data analysis was conducted to provide recommendations for options that utility companies can use for possible rebate programs to encourage the use of energy efficient FFUs.

One of the important outcomes from this project was to develop, refine, and publish an updated, standalone standard method for characterizing energy performance of fan-filter unit, including testing methods, procedures, and performance modeling based upon laboratory data. This report will 1) describe test setup of the demonstration at LBNL, 2) summarize results from experimenting 17 sample fan-filter units using the standard methods being updated, 3) evaluate repeatability of laboratory testing under a variety of conditions, and 4) develop options of energy-rebate criteria based upon the test results and statistical analyses. The energy-rebate criteria can be for the use of utilities in designing rebate programs to promote implementation of FFUs that are more energy efficient.

5.2 Designing Test Rig

Prior to the study, we reviewed open literature including ISO Standards, ASHRAE Standards including AMCA/ASHRAE Standard 51.2. These standards address the requirements for measuring airflow rate and air pressures associated with different HVAC equipment in ducted or un-ducted testing layouts. Some of the standards are applicable to characterize fan performance, or air conditioners in various laboratory setting. For example, fan performance characterization requires large chambers and flow baffles to create uniform airflows entering the fans. However, none of these standards specifically addresses any method for characterizing fan filter units or provides any direct guideline on acceptable laboratory methods or testing rigs to characterize electric power demand, airflow rate, and pressure differential across a fan filter unit in different operation.

An initial test rig was designed, constructed, commissioned in a laboratory of LBNL. The test rig was used to conduct laboratory testing of 17 FFUs recruited from the market. The test rig was also used and adapted to perform additional experiments on selected FFUs in this study. The Appendix Section includes details of the test rig material, sizes of components, sensor locations, and equipment. Additionally, we assessed and refined the existing test methods through

numerous trials and errors. The analyses and evaluations of date produced from the laboratory testing and experiments have created new information on FFU performance, new knowledge of the robustness of various test rig design and device setting, and validated laboratory methods, test procedures, and mathematical characterization of dynamic operation of FFUs. The new standard methods including the testing procedures were developed, and were published separately in the updated, refined standard document [3].

5.2.1 Sample fan-filter unit

Most of the individual FFU samples recruited and used in the demonstration project consisted of a small fan, and a HEPA or ULPA filter, and a sealed box sized to fit into standard cleanroom ceiling grids 61-cm by 122-cm (2 ft by 4 ft) as shown in Figure 1. One of the units has dual identical motors and fan wheels that are placed symmetrically inside the 61-cm x 122-cm (2-foot x 4-foot) FFU.

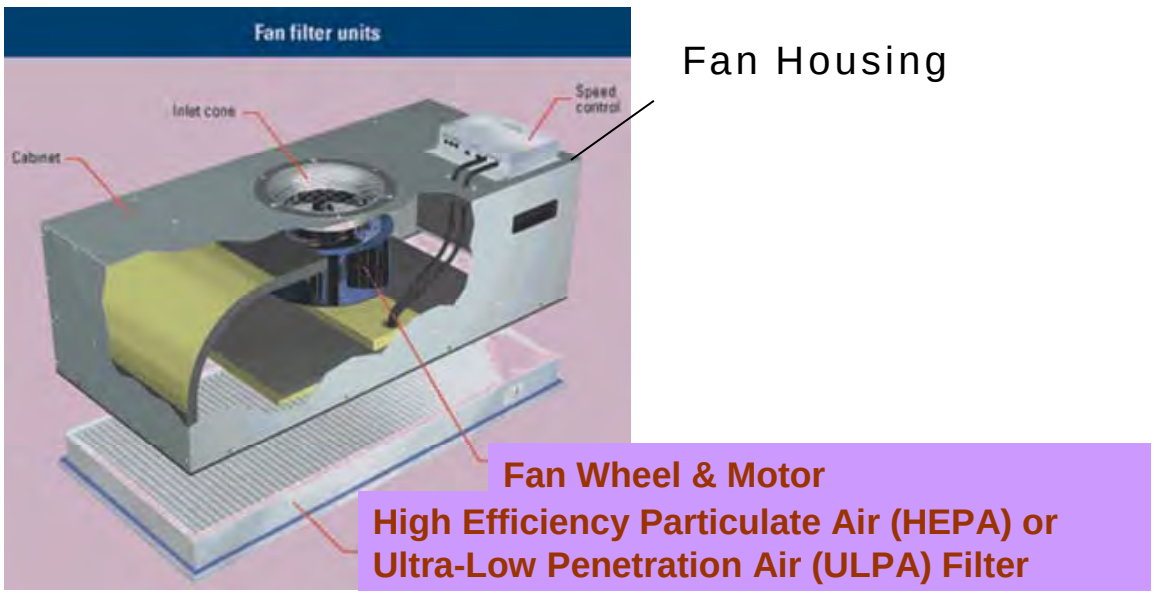

Figure 1 Generic Structure of 61-cm-by-122-cm Fan-filter Unit

5.2.2 Designing test rig

An initial test rig was designed, constructed, and installed according to the test method developed at LBNL. The initial test rig setup contained a combination of requirements by existing ASHRAE/AMCA and ISO Standards. In this test setup, a 61-cm by 122-cm (2 ft by 4 ft) FFU tested was mounted vertically. The FFU inlet was connected at the discharge end of a

straight 61-cm by 122-cm (2 ft by 4 ft) ductwork made of galvanized sheet metal. The ductwork was reinforced using metal rods as well as additional wood poles to withstand the force produced by the pressure difference between inside and outside of the duct.

A booster fan and a damper were installed at the ductwork inlet to modulate air pressures across the FFU while the airflow rates were concurrently controlled. Upstream of the 61-cm by 122-cm (2 ft by 4 ft) ductwork, the rig contains a round duct with a diameter of 25 cm (10 inches). The round duct connect to a standard flow straightener leading to a single-nozzle, 25-cm (10-inches) flow meter that is used to record airflow rates through the tested unit. The initial test setup was used to test and evaluate 17 sample FFUs in this study.

Additionally, we investigated the robustness of the test rig in terms of its physical integrity (e.g., leakage), and repeatability of measured performance under various test rig designs. Specifically, we changed various test rig setting and refined the methods to conduct experiments to evaluate repeatability and explore means to improve consistency. For example, we developed new information through designing and conducting additional experiments while changing test rig sizes and device locations. As a result, analysis of experimental data and information provided better understanding of measurement techniques for airflow rates, pressure change, power demand, and efficiency levels. The results from experimenting and analysis quantified the sensitivity of measured performance to variation in laboratory conditions, such as test rig sizes, locations of pressure taps, air temperatures.

5.3 Developing and Refining Laboratory Method and Procedure

The unit airflow rate, fan-wheel rotational speed, pressure differential across the FFU, and total electric power demand were recorded concurrently at any given operating condition. The operating conditions selected for testing covered typical operating conditions of each unit. Each of the operating conditions was achieved by adjusting the assisting fan wheel speed, the position of the damper installed in the test rig, and speed-control setting (if any) of fan filter unit. Normally, continuous adjustments of the device setting are essential to achieve a range of operating conditions defined by the pressure differential across the unit and the actual airflow rates of the fan filter unit.

In order to generate various testing conditions, the pressure differentials across FFU were normally controlled at various levels for each RPM setting. For example, within each incremental RPM setting, pressure differential was set to be as low as zero to 0.2-inch water column (0-50 Pa) up to 1.5-inch water (375 Pa), when applicable. The total power demand and performance metrics were then obtained for a specific operating condition and/or a specific range of operating condition, e.g., pressure differential of 0.5-inch water (125 Pa) coupled with and actual unit airflow rate of 520 cfm (245 Ls^{-1})) that the unit was capable of supplying. The FFU modulation device when applicable was integrated with the FFU for setting fan-wheel rotation speeds (Rotation per Minute, RPM). In these cases, the fan-wheel motor in the FFU was set at different RPMs during the testing to obtain performance data. At each RPM setting, the FFU was

tested at various pressure-rise across the unit, which was modulated by adjusting Iris damper positions. The ambient conditions and the airflow conditions were recorded and were used for the air density conversion to the equivalent standard condition (i.e., 1 atm, 20°C). We assume that the airflow was isothermal, although a small fraction of heat was generated from fan motors, which was transferred to the airflow. The reported performance data were based upon the standard air condition, i.e., with the air density of 1.20 kg/m^3, in order to directly compare the energy and aerodynamic performance. More details of testing methods, procedures, and device were described in appendices and the updated standard methods.

5.4 Developing Polynomial Regression Models to Characterize Energy Performance

Based upon the test results, we analyzed individual unit's functionality. i.e., operable conditions, and developed mathematical models to characterize dynamic operation and energy performance of individual fan-filter units.

Specifically, total electric power demand and total pressure efficiency were used as the yardsticks for quantifying unit's energy performance. The main approach was to develop polynomial regression models to quantify total power demand and total pressure efficiency of the individual units corresponding to their applicable operable conditions. The models characterize the relationship among energy performance metrics, airflow rates, and pressure differential across the units.

An FFU's total electric power demand includes power demand for fan motor, transformer, speed controller and display, and any accessories attached to the unit. Total electric power demand was measured using a power meter for selected operating conditions. To characterize total electric power demand, a polynomial regression model was established based upon laboratory test data, as it relates to actual operating conditions that are defined by unit airflow rates and pressure differential across the unit. For FFUs with a single-speed drive, the total electric power demand can be calculated using either of the following equations.

$$W_t = C_0 + C_1 \cdot D_p + C_2 \cdot D_p^2, \text{ or } W_t = C_0 + C_1 \cdot Q + C_2 \cdot Q^2$$

Where

D_p is the pressure differential across the fan filter unit.

Q is the airflow rate across the unit under standard atmospheric condition.

C_i (i = 0, 1, 2) is a coefficient developed from experimental data through polynomial regressions.

For FFUs with a multi-speed-drive, total electric power demand can be calculated using the following equation.

$$Wt = C0 + C1 \cdot Dp + C2 \cdot Q + C11 \cdot Dp^2 + C12 \cdot Dp \cdot Q + C22 \cdot Q^2$$

Where

Dp is the pressure differential across the fan filter unit

Q is the airflow rate across the unit under standard atmospheric condition.

$C_{i,j}$ (i, j = 0, 1, 2) is a coefficient developed from experimental data through polynomial regression.

An FFU's total FFU pressure efficiency (η_t) is defined as the actual airflow pressure power divided by the total electric power input to the FFU unit. Total pressure efficiency was calculated based upon laboratory test data, as it relates to operating conditions that are defined by unit airflow rates and pressure differential. For FFUs with a single-speed drive, total pressure efficiency can be calculated using either of the following equations.

$$\eta_t = \frac{Q \cdot D_p}{C_0 + C_1 \cdot D_p + C_2 \cdot D_p^2}, \text{ or } \eta_t = \frac{Q \cdot D_p}{C_0 + C_1 \cdot Q + C_2 \cdot Q^2}$$

Where

D_p is the pressure differential across the fan filter unit.

Q is the airflow rate across the unit under standard atmospheric condition.

C_i (i = 0, 1, 2) is a coefficient developed from experimental data through polynomial regressions.

For FFUs with a multi-speed-drive, the total pressure efficiency may be calculated using the following equation.

$$\eta_t = \frac{Q \cdot D_p}{C_0 + C_1 \cdot D_p + C_2 \cdot Q + C_{11} D_p^2 + C_{12} \cdot D_p \cdot Q + C_{22} \cdot Q^2}$$

Where

D_p is the pressure differential across the fan filter unit

Q is the airflow rate across the unit under standard atmospheric condition.

$C_{i,j}$ (i, j = 0, 1, 2) is a coefficient developed from experimental data through polynomial regressions.

5.5 Findings

In this report, we present the following parameters obtained from laboratory characterization, and include information for performance reporting: Unit's airflow rates (or airflow speeds), pressures rise across the unit, total electric power demand, and total pressure efficiency.

5.5.1 Airflow rate and pressure differential

Figure 2 shows the measured airflow rate as it related to the air pressure differential across a sample FFU with AC power supply and a multi-speed-drive controller. The measurements were taken while fan-wheel rotation speeds were set at three preset RPM levels sequentially.

It is clear that the FFU pressure differential decreased with the increase in airflow rates within the operating range of the FFU for each RPM setting. In addition, the FFU pressure differential across the unit was higher with the higher RPM at any given airflow rate. Similarly, for a given airflow rate, pressure differential across the FFU was maximized by operating the unit at its highest setting for fan-wheel's speed control.

Different from the sample unit shown in Figure 2, Figure 3 shows measured airflow rate as it related to the air pressure differential across another sample FFU with an electric commuted DC motor and a variable-speed-drive speed controller. The measurements were taken while fan-wheel rotation speeds were being set at five different levels respectively.

It is clear that the FFU pressure differential decreased with the increase in airflow rates within the operating range of the FFU for each speed-control RPM setting. In addition, FFU pressure differential across the unit was higher with the higher RPM at any given airflow rate. Similarly, for a given airflow rate, the pressure differential across the FFU was maximized by operating the unit at its highest setting for fan-wheel's speed control.

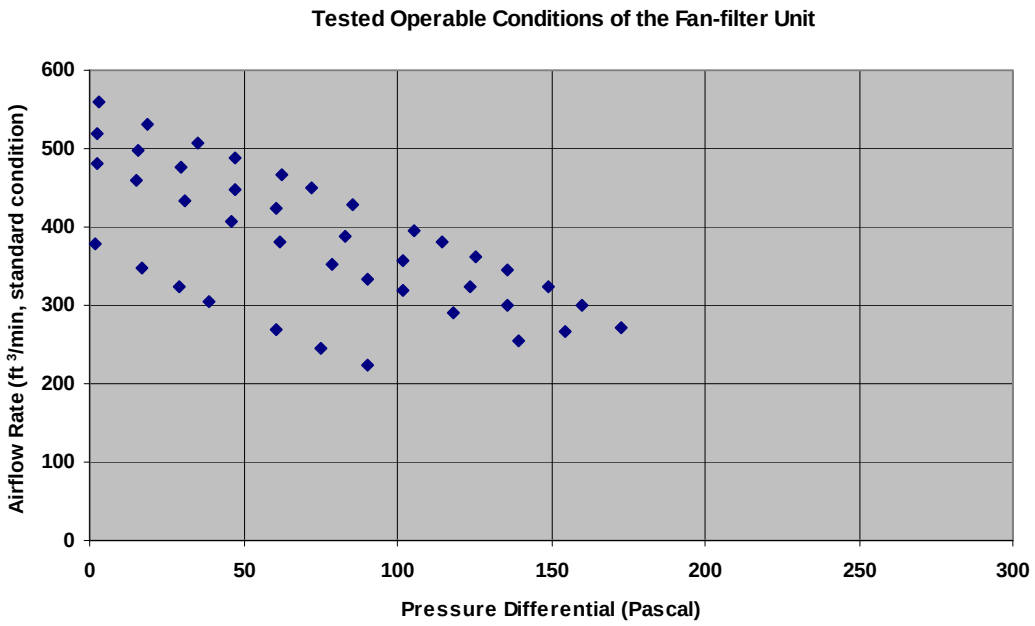

Figure 2 Airflow Rates vs. Pressure Differential (FFU Sample A)

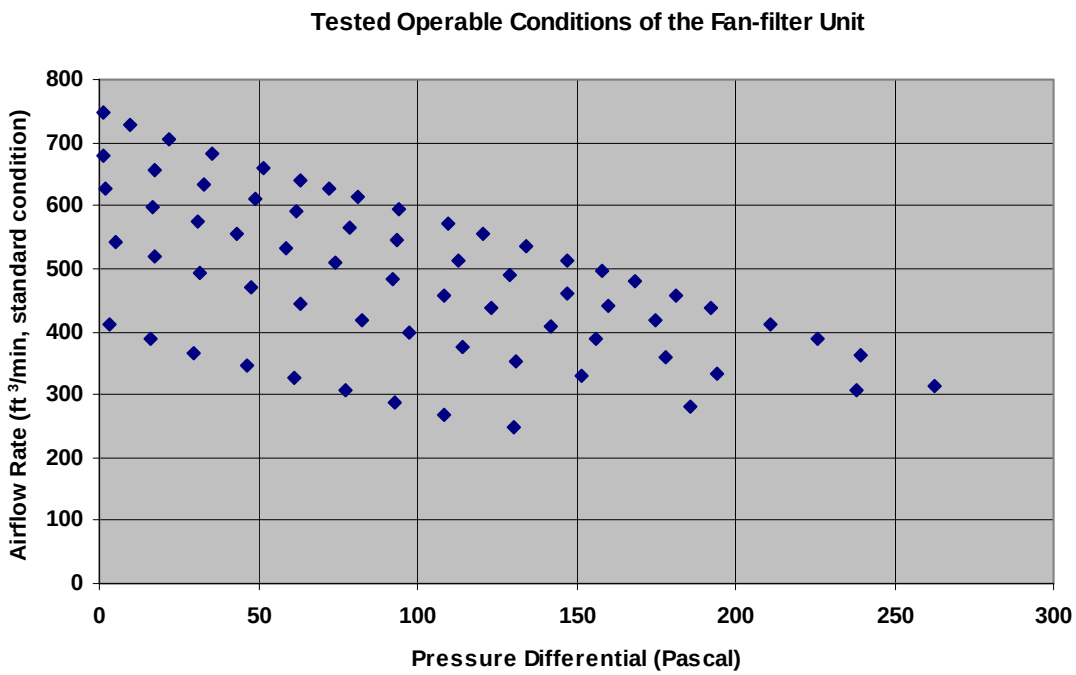

Figure 3 Airflow Rate and Pressure Differential (FFU Sample B)

5.5.2 Total electric power demand

An FFU's total electric power demand includes power demand for fan motor, transformer, speed controller and display, and any accessories attached to the unit.

Figure 4 shows the measured total electric power demand as it corresponded to the airflow rates and pressure differential across a sample FFU (Sample A). Sample A has an AC power supply and a multi-speed-drive controller. Each line in the figure represents the contour of the equivalent electric power demand within the selected operable conditions: 20 Pa ≤ Dp ≤ 150 Pa, Q ≥ 9.9 m3/min (or 0.08-iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm).

Figure 4 shows that Sample A's total electric power demand generally increased with the increase in airflow rates as well as decrease in pressure differential across the unit. Another finding is that total electric power demand was higher with the higher RPM at any given airflow rates and pressure differential within the FFU's operating range.

Figure 5 shows measured airflow rate as it related to the air pressure differential across another sample FFU (Sample B). Sample B has an AC power supply and a variable-speed-drive controller. Each line in the graph represents the contour of the equivalent electric power demand under the selected operable conditions: 20 Pa ≤ Dp ≤ 150 Pa, Q ≥ 9.9 m3/min (or 0.08-iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm).

Figure 5 shows that Sample B's total electric power demand generally increased with the increase in airflow rates as well as decrease in pressure differential across the unit. Compared to Sample A shown in Figure 4, Sample B exhibited higher capacity in producing higher airflow rate range as well as pressure differential across the unit. In addition, the increase of total power demand of Sample B was less sensitive to changes in pressure differential and more sensitive to changes in airflow rate. This indicates that the FFU with VSD motor was more receptive to changes in pressure differential - such a unit possesses higher capability to maintain energy demand relatively more constant over variations in external resistance, when everything else remains the same.

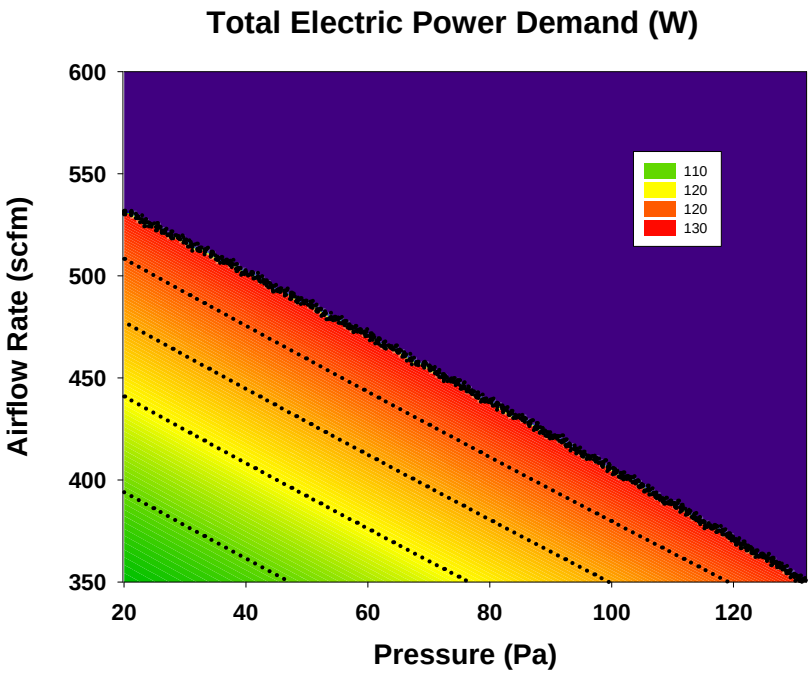

Figure 4 FFU Sample A with multi-speed-drive motor: Total electric power demand of the fan filter unit under selected operable conditions: 20 Pa ≤ Dp ≤ 150 Pa, Q ≥ 9.9 m3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

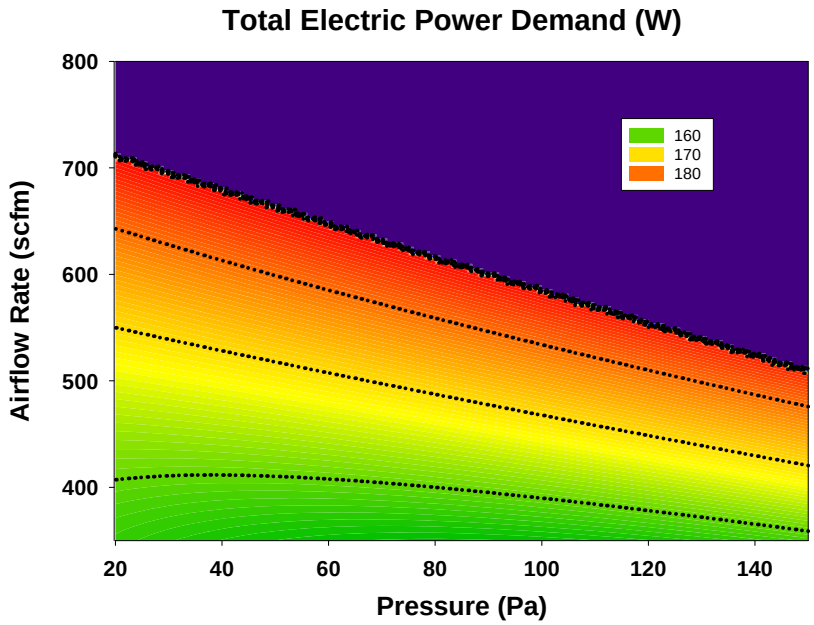

Figure 5 FFU Sample B with various-speed-drive motor: Total electric power demand of the fan filter unit under selected operable conditions: 20 Pa ≤ Dp ≤ 150 Pa, Q ≥ 9.9 m3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

5.5.3 Total pressure efficiency

Total pressure efficiency was calculated based upon laboratory test data, as it relates to operating conditions that are defined by unit airflow rates and pressure differential.

Figure 7 shows the performance curves of the FFU sample A: Total pressure efficiency as it changed with the airflow rates and pressure differential across the unit within a selected range of operating conditions. This sample unit (Sample A) has an AC power supply and a multi-speed-drive controller. Each dotted curve-line in Figure 7 represents the contour of the equivalent total pressure efficiency under selected operable conditions: 20 Pa ≤ Dp ≤ 150 Pa, Q ≥ 9.9 m3/min (or 0.08-iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm).

Figure 7 shows that sample A's total pressure efficiency exhibited a different trend than that of its total electric power demand shown in Figure 4. Total pressure efficiency generally increased with the increase in pressure differential across the unit; however, it remained relatively insensitive to the changes in airflow rates.

We found that the sample FFU's total pressure efficiency exhibited a different trend than that of its total electric power demand: the total pressure efficiency generally increased with the increase

in pressure differential across the unit; in the meanwhile, it also increased with increase of airflow rates. However, total pressure efficiency remained less sensitive to the changes in airflow rates at lower pressure differential setting that exhibited at higher-pressure differential.

Figure 8 shows FFU sample B's total pressure efficiency, as it related to the airflow rates and pressure differential across the unit within a selected range of operating conditions. This sample unit has a various-speed-drive controller. Each line in Figure 8 represents the contour of the equivalent total pressure efficiency under selected operable conditions: 20 Pa ≤ Dp ≤ 150 Pa, Q ≥ 9.9 m3/min (or 0.08-iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm).

Figure 8 shows that the sample B's total pressure efficiency also exhibited a different trend than that of its total electric power demand shown in Figure 5. Total pressure efficiency generally increased with the increase in pressure differential across the unit; in the meanwhile, it also increased with increase of airflow rates. However, total pressure efficiency remained less sensitive to the changes in airflow rates at lower pressure differential setting that exhibited at higher-pressure differential.

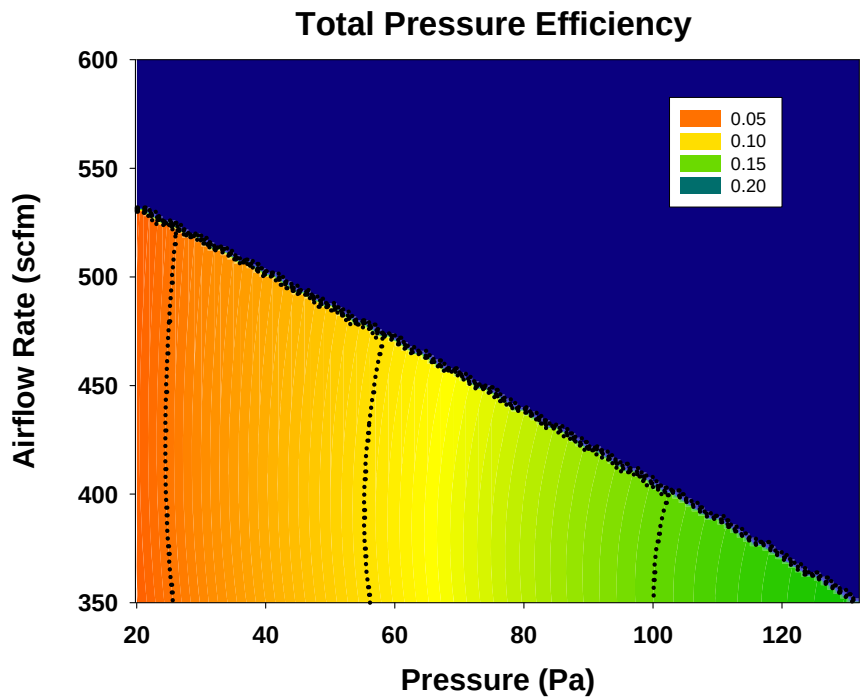

Figure 6

Figure 7 FFU Sample A: Total Pressure Efficiency of the fan filter unit under selected operable conditions: 20 Pa ≤ Dp ≤ 150 Pa, Q ≥ 9.9 m3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

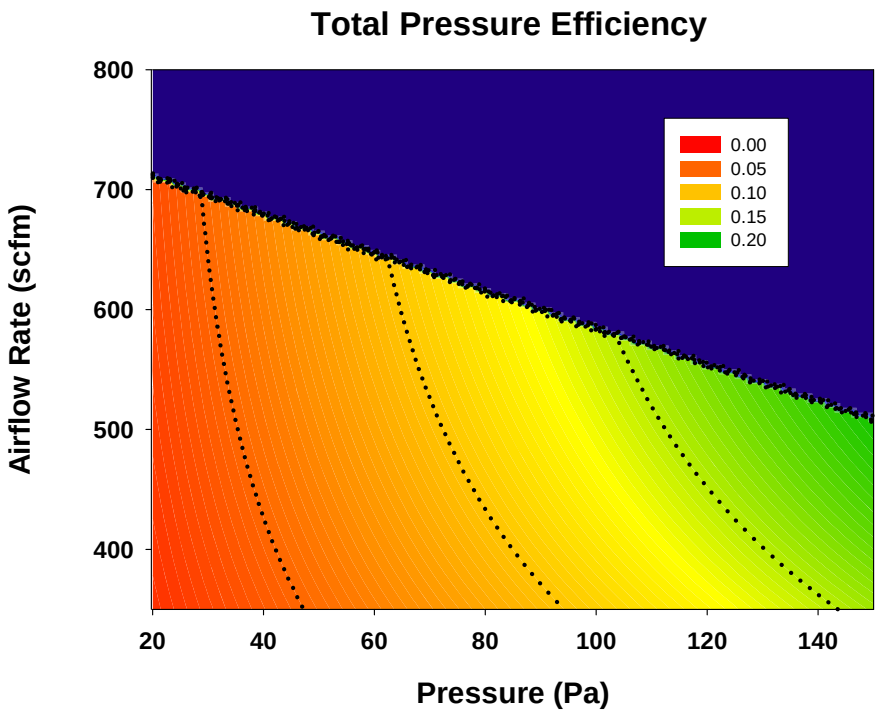

Figure 8 FFU Sample B: Total Pressure Efficiency of the fan filter unit under selected operable conditions: 20 Pa ≤ Dp ≤ 150 Pa, Q ≥ 9.9 m3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

5.5.4 Summary of findings

Using the standard test method, we have performed laboratory experiments on 17 FFUs donated by fan-filter-unit manufacturers in the US, Canada, Europe, and Asia. Operable conditions, total electric power demand, and total pressure efficiency for each individual unit were quantified using the updated standard laboratory test method. The following summarizes the key findings:

The energy performance of each FFU, after being tested within its own operable conditions, may be characterized and mathematically quantified as a function of operating conditions. The operating condition is defined by airflow rate and air pressure differential across the unit.

Among these we have seen electric power demand ranging from 100 to 400 watts per FFU, and significant variations in efficiencies (in excess of ten-times) depending on actual operating conditions, defined in terms of air flow rates (or velocity) and pressure differential across the unit that represents unit's capability to overcome pressure loss elsewhere in recirculation system.

The total pressure efficiency at any given operable condition can be obtained through the laboratory experiment and multi-variable regressions. When evaluating unit's efficiency under the same operating condition or the same range of operating conditions, we have observed a factor of three or more variation in energy efficiency across the 17 sample units.

The significant variations across operating conditions and units have posed challenges in defining or using a single performance metric for comparing performance, and applications in energy-incentive programs. The following presents the development of energy-rebate criteria through further statistical analyses, including the methods for percentile-ranking modeling and a relative ranking system, and recommendations for energy-rebate criteria.

5.6 Developing Energy Incentive Criteria

Cleanroom owners and users are a historically unresponsive market to energy efficiency efforts. As one of California's important and leading industries, this unique customer segment warrants additional attention and effort to motivate owners/users to participate in energy efficiency programs. Energy rebate programs are designed specifically to motivate the cleanroom owner, designer, or manager toward selecting and installing energy efficient FFU products. With product offerings suitable for the cleanrooms, energy-incentive program(s) may provide opportunities to serve two distinct beneficiaries of high-performance FFUs: the owner (user) and suppliers of such units.

The following analysis provides options of basis upon which energy rebate program can be established: prescription- and performance-based rebates. Using prescription-based approach, incentives will be paid based on identified products. The program identifies the more energy efficient units in the market and awards certain incentive to buyers of such units. Using performance-based approach, rebates are awarded for deemed savings measures. This will involve evaluations of cost-effectiveness, product price difference. The users and cleanroom designers will need to work with utility companies to estimate savings considering difference in energy savings as related to product premiums and cost effectiveness.

5.6.1 Prescription-based approach

5.6.1.1 Statistical method

In order to understand patterns of performance variations exhibited by sample units using the same test method, a simplified method would be needed to quantify and categorize energy-efficiency baseline information. Such baseline information needs to be sufficiently robust for utility to adopt in their energy-rebate programs.

First, we have performed a series of statistical analyses to evaluate the significance of difference in total pressure efficiency of all units under various combinations of operable conditions or ranges of operable conditions. Then we examined the significance in efficiency differences. In

particular, we conduct statistical ANOVA-Tests and t-Tests for groups or pairs of total pressure efficiency developed for the pool of sample units. A relative performance ranking system is then developed to quantify the observed difference and to identify rebate-criteria. The relative ranking scores are used to examine the robustness of the suggested initial rebate criteria. For instance, we can find out which units may get recommendations for rebate vs. which ones should not. Once establishing an initial rebate criteria, users may identify an initial set of units that surpass such criteria.

The following presents the results that are relevant to devising, identification, and recommendations for rebate criteria for utilities' rebate program.

5.6.1.2 Percentile-ranking model

The percentile ranking is a function of an array of data, which is determined by a range of total pressure efficiency data that defines relative standing, and the percentile value in the range [0, 1] - inclusive. The percentile-ranking models return the kth (k = 0, 100) percentile of values in the range as exhibited by the total pressure efficiency obtained for all of the sample units tested for a specific range of operating conditions. The 12 sets of operating conditions selected are shown in Table 1 in this report.

Total pressure efficiency of each unit was calculated for 12 sets of operating conditions (selected and specified in Table 1). Statistical percentile-ranking models were established to categorize the magnitudes of total pressure efficiency of the units. These models were used to characterize FFU performance at 12 sets of operating conditions that were defined by airflow rate and pressure differential across the FFUs. In particular, percentile-ranking models were employed to quantify levels of total pressure efficiency corresponding to specific efficiency ranking category that is determined by incremental percentiles derived from the sample FFUs.

Table 1. 12 selected sets of operating conditions for all FFU Samples

Selected Operable Conditions	Pressure rise (Pa)	Pressure rise (in. water)	Airflow range (m^3/min)	Airflow range (ft^3/min)	Nominal Airflow Speed (m/s)	Nominal Airflow Speed (ft/min)
1	50-150	0.2-0.6	9.1-22.7	320-800	0.2-0.5	40-100
2	75-125	0.3-0.5	9.1-22.7	320-800	0.2-0.5	40-100
3	50-75	0.2-0.3	11.3-18.1	400-640	0.3-0.4	50-80
4	75-100	0.3-0.4	11.3-18.1	400-640	0.3-0.4	50-80
5	100-125	0.4-0.5	11.3-18.1	400-640	0.3-0.4	50-80
6	125-150	0.5-0.6	11.3-18.1	400-640	0.3-0.4	50-80
7	75	0.3	11.3-18.1	400-640	0.3-0.4	50-80
8	100	0.4	11.3-18.1	400-640	0.3-0.4	50-80
9	125	0.5	11.3-18.1	400-640	0.3-0.4	50-80
10	75	0.3	14.7	520	0.3	65
11	100	0.4	14.7	520	0.3	65
12	125	0.5	14.7	520	0.3	65

Table 2 shows the outcomes from percentile-ranking models developed from the sample units' total pressure efficiency under each of the selected operating conditions.

Supplemental to the data shown in Table 2, Figure 9 shows levels of total pressure efficiency as a function of percentile ranking developed from the 17 sample FFUs tested in this study. In general, units with lower ranking, e.g., under 50th percentile (or median) tend to have similar and lower total pressure efficiency. On the other hand, units with higher ranking, e.g., above 50th percentile (or median) tend to exhibit higher total pressure efficiency. The higher efficiency levels also exhibits bigger spread across the incremental percentile ranking. This observation holds true for all operating conditions or ranges of conditions, as illustrated by the polynomial-regressed trend-lines with R-square values higher than 0.99.

Table 2. Percentile ranking of total pressure efficiency

Total Pressure Efficiency	Percentile										
Selected Operable Conditions	**0%**	**10%**	**20%**	**30%**	**40%**	**50%**	**60%**	**70%**	**80%**	**90%**	**100%**
1	6.6%	7.8%	9.2%	9.4%	10.9%	11.2%	11.9%	12.7%	13.8%	17.8%	22.2%
2	7.5%	8.1%	9.7%	10.1%	11.4%	12.0%	12.5%	13.4%	14.4%	18.3%	22.6%
3	4.9%	5.4%	6.7%	7.1%	8.1%	8.5%	8.8%	9.5%	10.6%	14.3%	18.0%
4	6.3%	7.2%	9.4%	9.5%	10.2%	11.2%	11.8%	11.9%	13.3%	17.4%	22.3%
5	9.2%	10.1%	11.6%	11.8%	12.4%	13.7%	14.0%	14.5%	15.4%	20.2%	25.6%
6	11.0%	11.7%	13.0%	13.3%	13.3%	15.3%	15.8%	16.5%	18.4%	23.0%	28.2%
7	5.9%	6.3%	8.1%	8.4%	9.2%	10.0%	10.4%	10.8%	12.3%	16.0%	20.3%
8	8.1%	9.3%	10.7%	10.8%	11.8%	12.8%	13.0%	13.3%	15.2%	19.1%	24.1%
9	10.2%	10.7%	12.3%	12.6%	13.0%	14.5%	15.0%	15.5%	17.5%	22.1%	27.0%
10	5.9%	6.9%	8.3%	8.6%	8.9%	9.8%	10.6%	10.8%	13.0%	17.2%	20.6%
11	8.1%	8.7%	10.6%	10.8%	11.3%	12.4%	13.4%	14.2%	16.3%	20.5%	24.4%
12	10.3%	10.3%	11.9%	13.3%	13.9%	14.5%	16.3%	17.8%	20.4%	24.0%	27.3%

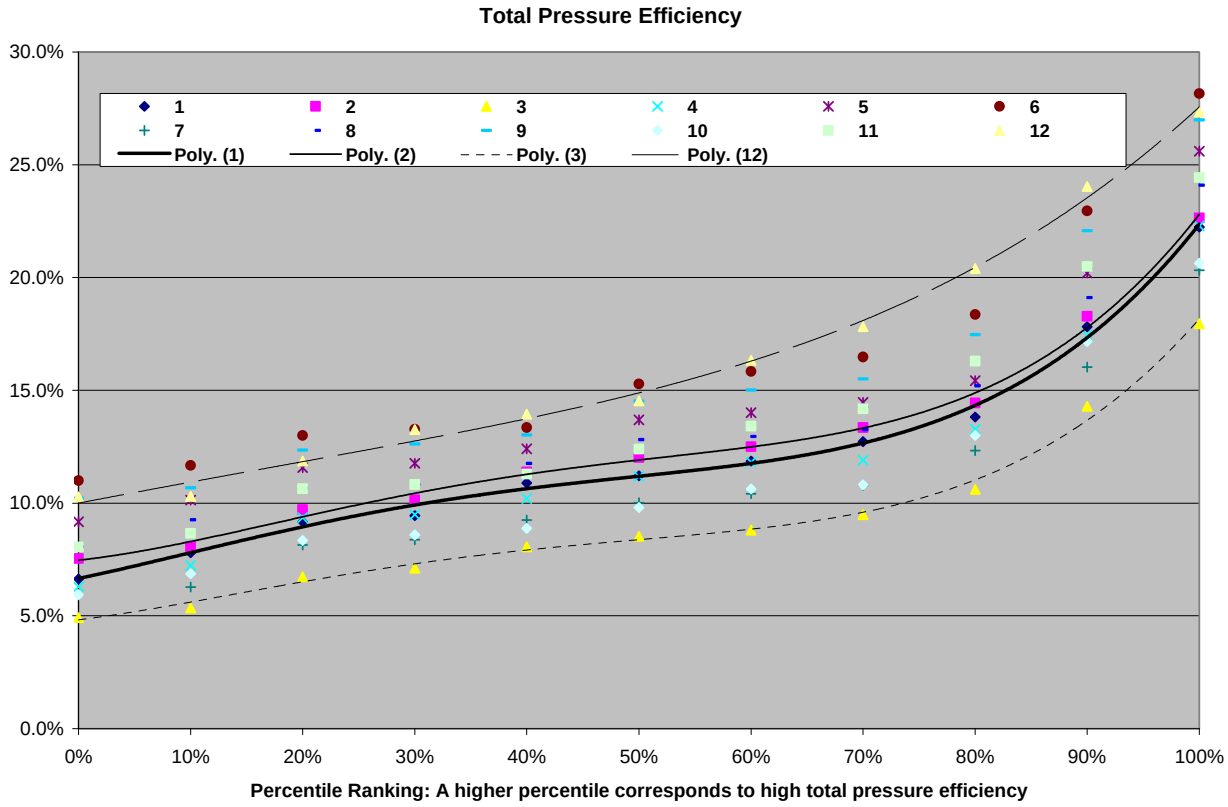

Figure 9 Percentile ranking of total pressure efficiency at selected operating conditions

Table 3 shows additional details about the percentile ranking curves in Figure 9. The table illustrates the equations and fitted significance for polynomial regression of total pressure efficiency based upon percentile ranking under four selected operating conditions in this study, i.e., conditions 1, 2, 3, and 12. Within each operating conditions range, average total pressure efficiency increased as a function of the ranking percentile.

Table 3 Polynomial regression of total pressure efficiency

Selected Operable Conditions	Polynomial Equation Y: Total pressure efficiency; X: Percentile	R-Square
1 (50-150Pa, 9.1-22.7m^3/min)	$y = 0.5166x^4 - 0.6121x^3 + 0.1467x^2 + 0.1058x + 0.0665$	0.9958
2 (75-125Pa, 9.1-22.7m^3/min)	$y = 0.6482x^4 - 0.9104x^3 + 0.3602x^2 + 0.0553x + 0.0747$	0.9962
3 (50-75Pa, 11.3-18.1m^3/min)	$y = 0.5081x^4 - 0.6526x^3 + 0.2145x^2 + 0.0635x + 0.0482$	0.9945
12 (125Pa, 14.7m^3/min)	$y = 0.1002x^4 + 0.0002x^3 - 0.0199x^2 + 0.0951x + 0.1$	0.9959

Based upon this percentile-ranking function, we may initially examine the difference in efficiency levels and the significance. In particular, we examined the significance of difference of average total pressure efficiency among all selected ranges of operating conditions given the same percentile. We then select a likely threshold of acceptance or qualification. Using this threshold of acceptance or qualification as a criterion, we then can identify FFU candidates of which total pressure efficiency is higher than the acceptance criteria.

In order to evaluate the significance of difference in efficiency levels, a series of ANOVA -and t-Tests were performed to test the statistical significance of difference in efficiency levels among the prescribed performance ranking and between operating conditions. Through performing ANOVA analyses for the ranking output across different operating conditions and ranges of the conditions as selected and shown in Table 1, we have found that there are generally significant differences in the average total pressure efficiency across and within the 12 selected conditions. For example, the difference in average total pressure efficiency at various operating conditions was statistically significant, as exhibited in Table 4.

Table 4. ANOVA significance of difference in total pressure efficiency at various conditions

Anova: Single Factor

SUMMARY

Groups	*Count*	*Sum*	*Average*	*Variance*
3	11	1.018576	0.092598	0.001501
4	11	1.304881	0.118626	0.002092
5	11	1.58433	0.14403	0.002252
6	11	1.793494	0.163045	0.002686
7	11	1.178764	0.10716	0.001804
8	11	1.479902	0.134537	0.002131
9	11	1.704436	0.154949	0.002558

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Between Groups	0.043761	6	0.007294	3.398084	0.00533	2.231192
Within Groups	0.150246	70	0.002146			
Total	0.194007	76				

Furthermore, a series of paired-sample t-Tests were performed to examine the significance of difference between average total pressure efficiency at different operating ranges. For example, a t-Test was performed for two different ranges of operating conditions: 1) pressure differential of 50-150 Pa with 9.1-22.7 m^3/min airflow rate and 2) pressure differential of 75-125 Pa with the same 9.1-22.7 m^3/min airflow rate. As shown in Table 5, we found that based upon the selected percentiles (0% up to 100%), the actual *p*-value of the t-Test indicates that there was a significant difference in average efficiency due to the difference between two operating-condition ranges. This significant difference in total pressure efficiency holds true even though the operation-condition ranges, i.e., conditions 1 and 2, had overlapped significantly.

Table 5. Significance of difference in total pressure efficiency

t-Test: Paired Two Sample for Means

	50-150	*75-125*
Mean	0.121463	0.127337
Variance	0.002049	0.001995
Observations	11	11
Pearson Correlation	0.999285	
Hypothesized Mean Difference	0	
df	10	
t Stat	-10.79656	
P(T<=t) one-tail	3.92E-07	
t Critical one-tail	1.812461	
P(T<=t) two-tail	7.84E-07	
t Critical two-tail	2.228139	

Through performing ANOVA- and t-Tests and examining the results, we discovered that there is significant difference in efficiency levels of various units and across operating conditions. In addition, no single set of operating condition or range of such operating conditions should be expected to represent any "typical" operation or efficiency levels., .

In fact, there are a variety of designs and possible ranges of operating conditions for FFU applications in cleanrooms and controlled environments. The levels of energy efficiency of an individual FFU may be quantified through performing laboratory tests and mathematical derivations. However, it appears that no single efficiency level can be used for baseline information, unless the designer or user has selected or required a specific operating condition for its cleanroom application.

5.6.1.3 Relative ranking system

In order to develop energy-rebate criteria, an alternative solution is needed. The following analysis presents an alternative to obtain energy-rebate criteria.

Based upon the percentile-ranking models developed from the test data, we established a relative ranking system among the sample FFUs. Relative ranking scores for all sample units, defined as the actual total pressure efficiency divided by the median value of the array of samples' efficiency values under the same operating condition, were calculated.

Assigning a relative ranking score to each fan-filter unit, this relative ranking system allows identifying better performers vs. worse performers in terms of their total pressure efficiency levels. After identifying individual units that are relatively better performer under various

operating conditions, the next step is to check and validate whether the relative ranking holds the same pattern even though actual efficiency levels have varied from unit to unit and from one operating condition to another condition. If the observed relative ranking pattern exhibits consistency, then using the raking system to filter out better performers vs. worse performers would become a valid approach that can be used to define acceptance criteria for rebate. For example, it can be used to identify units of which total pressure efficiency surpasses the acceptance criteria to be selected and defined as the minimum ranking requirement.

Table 6 Performance ranking scores for each unit under two different operating conditions

Operating conditions	1	2		Operating conditions	3	4	5	6	7	8	9	
Pressure (Pa)	50-150	75-125		Pressure (Pa)	50-75	75-100	100-125	125-150	75	100	125	
Airflow rate (m^3/min)	9.1-22.7	9.1-22.7		Airflow rate (m^3/min)	11.3-18.1	11.3-18.1	11.3-18.1	11.3-18.1	11.3-18.1	11.3-18.1	11.3-18.1	
Fan-filter unit	Relative ranking	Relative ranking	Unit - Accepted or Rejected	Fan-filter unit	Relative ranking	Relative ranking	Relative ranking	Relative ranking	Relative ranking	Relative ranking	Relative ranking	Unit - Accepted or Rejected
1	1.98	1.89	A	1	2.10	1.99	1.87	1.84	2.03	1.88	1.86	A
2	1.73	1.65	A	2	1.84	1.72	1.60	1.56	1.76	1.61	1.58	A
3	1.45	1.40	A	3	1.51	1.40	1.29	1.26	1.44	1.31	1.27	A
4	1.23	1.17	A	4	1.11	1.07	1.03	1.03	1.08	1.02	1.03	A
5	1.16	1.20	A	5	1.24	1.19	1.09		1.23	1.16		A
6	1.11	1.05	A	6	1.12	1.06	1.00	1.00	1.07	1.00	1.00	A
7	1.06	1.03	A	7	1.03	1.06	1.06	1.11	1.04	1.04	1.10	A
8	1.02	1.04	R	8	0.99	0.97	0.93		0.98	0.95		R
9	0.98	0.97	R	9	1.01	1.03	1.02	1.05	1.02	1.00	1.05	R
10	0.97	0.95	R	10	0.94	0.91	0.87	0.87	0.92	0.87	0.87	R
11	0.85	0.83	R	11	0.83	0.85	0.85	0.87	0.84	0.84	0.86	R
12	0.83	0.85	R	12	0.79	0.84	0.86	0.87	0.81	0.84	0.90	R
13	0.82	0.81	R	13	0.83	0.84	0.83	0.84	0.84	0.82	0.84	R
14	0.70	0.68	R	14	0.65	0.67	0.68	0.72	0.65	0.66	0.70	R
15	0.69	0.67	R	15	0.58	0.63	0.67	0.74	0.59	0.63	0.71	R
16	0.59	0.63	R	16	0.61	0.56			0.60			R
17	-	-	R	17		-	-	-	-	-	-	R

Table 6 shows the calculated relative ranking scores for each sample unit under nine different operating conditions, respectively. These nine operating conditions are conditions 1 through 9 as illustrated in Table 1. Relative ranking score was calculated as the ratio of actual total pressure efficiency of the unit to the median value of total pressure efficiency derived from the 17 sample FFUs corresponding to an operating condition.

For a selected operating condition or a selected range of operating conditions as illustrated in Table 6, units with relative ranking scores greater than 1.00 are identified as better performers.

For example, units 1 through 8 were identified as better performers under conditions 1 and 2; and units 1 through 7 and unit 9 were identified as better performer under conditions 3 through 9. In this pool of 17 samples, there are additional units such as units 8, 9, 10 having relative ranking scores near 1.00. Therefore, their acceptance for rebate was more susceptible to the exact cut-off point of the relative ranking scores. For instance, unit 8 would be accepted in conditions 1 and 2, while it will have to be replaced by unit 9 in conditions 3 through 9. Nonetheless, this could be due to relatively small sample size of the FFUs tested or just the fact that quite a few FFUs had relative ranking scores clustered within the vicinity of 1.00.

Using relative ranking scores as acceptance criteria for rebate would become very straightforward because this allows identifying energy-efficient products or models. The criteria can be developed based upon standard laboratory test results and the relative raking system. The qualified products (accepted samples) will be expected to perform at an efficiency level better than or close to median efficiency at a given operable condition and the same range of such conditions. For example, we can safely identify seven units (number 1 through 7 in Table 6) with relative ranking scores over 1.00, and recommend them be initially accepted for energy rebate program. In this case, fewer than half of the units (seven out of 17) tested would be initially accepted as candidates in energy-rebate program. As an alternative, if more units are desired for inclusion for rebate qualification, e.g., units 1 through 10. The expanded set of units (including 8, 9, and 10) were added because their performance was close to average and exhibited sufficient consistency under all the operating conditions listed in Table 6.

In either case, the caveat is that not all units accepted for rebate will be able to operate at all operating ranges. For example, units 5 and 8 will never operate at conditions 6 or 9. Users of this information or criteria recommendation should bear in mind that although unit 5 is made legible for meeting energy-rebate criteria in most of the cases, it might not be an appropriate candidate for the user to purchase, let alone to receive rebate. The bottom line is that unit 5 or 8 should not be selected by the end-user or the designer if the designed cleanrooms are expected to operate at conditions 6 and 9, thereby neither should it be qualified for rebate in this case.

5.6.2 Performance-based approach

As an alternative to statistical approach, a rebate may be awarded on the deemed-saving basis. The performance-based program will be based on certain quantifiable performance, e.g., how much better than average FFU, considering the actual price difference between more efficient and less efficient units. Cleanroom users can earn incentives for using energy efficient units. For example, designers, owners of the facility, and/or utility companies can select certain operating condition and calculate the actual power demand and annual energy consumption of FFUs tested. By examining the potential utility savings from efficient FFUs and the price differential between efficient units and less-efficient unit, the utility companies may wish to customize the energy rebate on a case-by-case basis. For example, they will then decide whether to award any financial incentive rebate to certain products, and the magnitude of such incentive rebate. The reported performance data and the equations will assist them to identify functional and efficient

units, and will provide the base for quantify the magnitude of rebates considering their own budget, deemed saving potential, and the price differential of FFUs.

5.6.3 Recommending criteria options for energy incentive programs

In order to define the criteria used for utility's rebate program, we recommend two options to formulate energy incentives: prescription- and performance-based incentive programs.

5.6.3.1 Criteria for prescription-based incentive program

The criteria of prescription-based program can be established by specifying a set of products that meet or surpass certain threshold of energy performance. Such threshold can be identified and selected based upon relative ranking scores of the FFU energy performance. Analyses were conducted to determine the set of FFUs that would outperform others, e.g., approximately half (or more) of the rest of sample units. Our tests, analyses, and validations show that most of the units with a relative ranking score greater than 1.00 under one of the operating conditions largely exhibited a relative ranking score greater than 1.00 under the other operating conditions listed in the table. Except for some operating conditions that a unit may not achieve under any circumstances, the relative ranking scores of the sample units in this project have exhibited consistent patterns in relative performance compared to a median performer. Therefore, using the relative ranking score is suitable for identifying FFU out-performers in terms of energy efficiency. Based on the statistical analyses, we recommend the following guidelines for utility companies to consider in formulating a prescription-based energy-rebate program:

1) The rebate-criteria for prescription-based incentive program will be the relative ranking scores of the FFU products gathered from the sample units tested.

2) Identify an initial set of higher performers of energy-efficient fan-filter units from the 17 sample units in this study, and reward dollar rebates for procuring these products that surpass the minimum energy-rebate criteria, e.g., relative ranking scores greater than 1.00, or seven out of 17 products tested. The utility companies, at their own discretion and preference, may wish to consider more or less stringent criteria based upon percentile higher or lower than 50%, e.g., 10 out of 17 products tested. In either case, higher-performers (i.e., higher energy-efficiency) in the FFU sample pools may be identified as initial set of candidates qualified for a specific rebate.

3) Identify additional high-performers of energy-efficient units to be included in the rebate program, through requesting and obtaining comparable and standard test results of additional units in the market. Additional units may be tested and evaluated. Such new units may include additional 61-cm by 122-cm (2-ft by 4-ft) units and 122-cm-by-122-cm (4-ft-by-4-ft) units. A quick way to evaluate whether or not a new unit is qualified for rebate is to compare itself with the existing relative ranking system developed from the 17 sample units. As an alternative, especially if there are substantially more units and products tested, a new relative

ranking system would need to be developed using the method presented in this study. Then newer rebate criteria may be established and used for future rebate programs.

4) As more units are tested using the standard laboratory tests, compatible performance data are available for further analyses. Newer rebate criteria may then become necessary over time and can be developed using the statistical method in the study. A periodic update of the list of product qualified for rebates may then be made available for users and utility companies in the future. This new information and dynamic criteria will not only reward the use and purchase of energy-efficient units on the market, but also will encourage suppliers to design, develop, and market more efficient units that would ultimately transform the FFU market toward higher energy efficiency.

5.6.3.2 Criteria for performance-based incentive program

The criteria of performance-based program can be established by calculating energy savings from FFU products for a specific cleanroom application, while taking into utility costs (e.g., electricity charge rate) and actual unit cost differential among the FFU products. The estimates of energy savings will be made based upon the test results applicable to a specific set of design or operating conditions of the FFU for cleanroom applications.

This option of energy rebates will be awarded on the case-by-case basis, considering units' functionality, deemed energy-savings under designed conditions, and the unit cost. For example, designers, owners of the facility, and/or utility companies can agree upon and select a certain operating condition, and calculate the actual power demand and annual energy consumption of available FFUs as compared to the sample FFUs.

By examining the potential cost reduction in energy and power demand from more efficient FFUs and the price differential between efficient units and less-efficient units, the utility companies may customize the energy rebate on a case-by-case basis. Based on the calculation(s), the utility companies can first decide whether to award any financial incentive (rebate) to any products based upon the magnitude of the potential energy and power demand benefits and payback. If it is determined to be appropriate to award an incentive, they will need to specify the magnitude(s) of such incentive rebate and/or a formula to recommend and quantify the magnitudes. The decisions should take into consideration of actual functionality of FFUs under designed operation conditions, potential energy savings, and product price.

The reported performance data, operable conditions, and the equations developed for individual products from laboratory testing in this report will be used to assist in identifying functional and efficient units for specific applications. This will provide the base for quantifying the magnitudes of rebates, considering their own budget, deemed saving potential, and the price differential of individual FFUs.

The LBNL standard has been and is being adopted by users and owners to understand FFU performance. The results are used in their process of selecting and purchasing fan-filter units with better and improved performance. In addition, the outcomes of the standard tests are now being considered by utilities seeking to promote applications of energy efficient FFUs. A successful energy-rebate program would allow a utility to provide financial incentive for end-users to specify and purchase energy-efficient fan-filter units. Furthermore, some end-users have been proactively pursuing ways to reduce cleanroom operating costs and life-cycle costs by selecting energy-efficient fan-filter units. For example, in designing and constructing large cleanrooms, some large companies in the US and Asia including Texas Instruments have required FFU suppliers or bidders to perform and report tests according to the LBNL standard, which allows provision and comparison of performance data in a consistent way.[2][3]

6. Summary

All individual standard laboratory-testing reports have been completed and reported back to anonymous individual participants in this project. Each report on an individual fan filter unit provides rigorous and useful data produced from standard characterization of FFU's operation. In the course of the project, the laboratory method previously developed at LBNL has been under continuous evaluation and update.[2][3] Based upon the updated standard, it becomes feasible and easier for users and suppliers to characterize and evaluate energy performance of FFUs in a consistent way.

Laboratory testing of FFU energy performance has provided useful data for suppliers and end users to understand the functionality and performance of FFU products under a variety of operating conditions. More importantly, suppliers and users can now quantify and compare energy performance of FFU products under any pre-determined and operable conditions. For the design and construction of large cleanrooms, some companies have required FFU suppliers to report laboratory tests based upon the characterization methods developed at LBNL[2] [3]. For example, many manufacturers of fan filter units expect to produce or obtain such data to quantify energy efficiency levels of their products. Ultimately, the standard test method is expected to be used or adopted by more and more suppliers, specifiers, and end-users to understand and improve FFU functionality and performance, as well as utilities seeking to promote applications of energy efficient FFUs.

In this demonstration project, we have refined the characterization methods, performed laboratory experiments to characterize fan-filter units, developed polynomial regression models to quantify the efficiency and power demand as a function of airflow and pressures. We have found significant variations in energy performance from unit to unit, as well as significant variation of units' performance across various operating conditions. The observed significant variations posed challenges in identifying a single efficiency level as baseline information for utility's rebate programs.

In addition, for the benefit of possible energy incentive programs operated by California utility companies, we have performed statistical analyses of the test data, developed a relative ranking system, and examined statistical consistency of observed variation in units' energy performance. Based upon the statistical evaluations of the test data gathered from testing the sample FFUs, we have developed an energy-rebate criteria, which allows identification of an initial set of units and products that are more energy efficient at a given operation condition or within a given range of operating conditions. Utility may use the information in formulating their energy-rebate programs to promote the use and purchase of FFUs that are more energy efficient.

In summary, we have accomplished the following:

- Designed, constructed, and commissioned a new test rig used to perform FFU testing and experiments.

- Developed and performed laboratory tests of 17 different FFUs.

- Produced and delivered 17 individual reports on characterizing dynamic operation of individual FFUs with various controls and operation.

- Developed laboratory experiments to examine, validate, and improve the laboratory test method including procedure that were used to characterize various FFUs.

- Refined the draft standard method previous developed at LBNL; promulgate the application of the new standard methods through interactions and communications with the industries (users, manufacturers, professional society, and magazines).

- Developed mathematical modeling of characterizing energy performance of individual fan-filter units within their operable conditions.

- Performed statistical analyses and developed a relative ranking system that resulted in initial criteria to compare units' performance.

- Developed options of criteria for energy-incentive program, identified higher-performers of energy efficient FFUs, and provided recommendations for utility companies in designing and implementing energy –rebate programs for promoting use of energy-efficient FFUs.

Through an extensive review of open literature, development of the standard methods, and additional experimental validations of the refined methods in this project, the following issues are identified to be important and may need future development.

- Continue to improve rigor and robustness of the new, largely improved standard methods, e.g., characterize design and speed-control impacts on overall performance, e.g., type of motors, internal housing, air path, and size of units.

- Characterize internal pressure distribution, e.g., pressure variations across HEPA/ULPA filters, and within the air-path of the FFU.

- Characterize airflow uniformity.

While the contamination control industries are moving toward tighter contamination control and increasing desire for higher energy efficiency, it is important to strategize the development and implementation of higher-efficiency FFUs in actual cleanrooms. The following includes important actions that need to be taken:

- Disseminate the results among technical and professional societies and across industries.

- Interact with relevant professional societies and standard development bodies to further absorb and adopt the refined standard test method, e.g., assist and facilitate the development and publication of industrial standard - IEST RP036.1: Testing Fan Filter Unit.

- Provide technical assistance to users and manufacturers to use efficient FFUs and improve FFU performance.

- Assist utilities to establish and implement energy-rebate program to promote FFUs that outperform others.

Identifying and selecting energy efficient units in cleanroom applications can bring about energy and demand savings while maintaining and improving the effectiveness of contamination control. Through this research and demonstration project, it becomes feasible for end users or cleanroom owners to become better informed of the energy performance to aid in their planning and selection for use in new facility construction or renovation. For example, they may now require suppliers to provide the units' performance obtained through the LBNL standard. In addition, more FFU manufacturers are becoming motivated to understand performance of their units, to improve design, operation and controls of their FFUs, and to serve their customers better. Furthermore, utility companies or other public interest programs may use the results and recommendations to establish energy-rebate criteria, and implement programs to encourage the use of efficient units. Last but not the least, the outcomes from this work have and will continue to facilitate and add to the on-going development of an industrial standard, i.e., IEST- RP-CC036.1: Testing Fan Filter Unit.

7. Acknowledgements

The authors wish to acknowledge the participation of Dennis Dibartolomeo, William Fisk, Dale Sartor, Douglas Sullivan, and William Tschudi of LBNL in various phases of this project. LBNL wishes to thank John Haystead of CleanRooms Magazine, and Paul Nesdore of Controlled Environments Magazine for providing assistance in our outreach activities and in unit recruitment; numerous industrial participants (anonymous) for providing fan-filter units in this project free of charge; members of IEST WG36 for providing feedback; and Robert Davis, Chris Condon and Mark Bramfitt of PG&E for their assistance to LBNL activities. Last but not the least, LBNL wishes to acknowledge Paul Roggensack, Anthony Wong, and Pramod Kulkarni of California Energy Commission for their support for this project.

The project is funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

8. References

[1] Xu, T. 2005. Advanced Efficient Filtration: Fan-filter Units. Final Report, Lawrence Berkeley National Laboratory Report, LBNL-57727, Berkeley, California.

[2] Xu, T. 2005. Laboratory Method of Testing Energy Performance of Fan-Filter Units, Version 1.3 (2005). Lawrence Berkeley National Laboratory Report, LBNL-54626, Berkeley, California.

[3] Xu, T. 2007. Standard Methods of Characterizing Performance of Fan Filter Units, Version 3.0. Lawrence Berkeley National Laboratory Report, LBNL-62118, Berkeley, California.

[4] The Institute of Environmental Sciences and Technology (IEST). In progress. Draft IEST-RP-CC036.1: Testing Fan Filter Unit. Rolling Meadows, IL.

[5] Xu, T. and M. Jeng. 2004. Laboratory Evaluation of Fan filter Units' Aerodynamic and Energy Performance. Journal of the IEST, Volume 47: 116 120. The Institute of Environmental Sciences and Technology (IEST), 5005 Newport Drive, Rolling Meadows, Illinois, October 2004. Lawrence Berkeley National Laboratory Report, LBNL 54250, Berkeley, Calif.

[6] Gale, S. 2004. FFUs: Setting a course for energy efficiency, Special Report, Cleanroom Magazine, September 2004, http://cr.pennnet.com/Articles/Article_Display.cfm?Section=Articles&Subsection=Display&Article_ID=211810

[7] Xu, T. 2004. Introducing a Standard Testing Method for FFUs, Controlled Environments (Formerly A2C2), http://www.cemag.us/articles.asp?pid=469, September 2004

[8] M. Jeng, T. Xu, C. Lan. 2004. Toward Green Systems for Cleanrooms: Energy Efficient Fan-Filter Units, Proceedings of SEMI Technical Symposium: Innovations in Semiconductor Manufacturing (STS: ISM), SEMICONÆ West 2004, San Francisco, CA, Semiconductor Equipment and Materials International, July 2004. LBNL 55039.

[9] Xu, T. 2004. Developing a Standard Method For Evaluating Fan-Filter Unit Performance – Applications in Utility Incentive Programs. Poster distributed to three major California utilities, California Energy Commission, and other stakeholders.

[10] Xu, T., C. Lan., and M. Jeng. 2006. "Performance of Large Fan Filter Units for Cleanroom Applications." *Building and Environment.* Lawrence Berkeley National Laboratory Report, LBNL-58220.

[11] Chen, J., C. Lan, M. Jeng, and T. Xu. In press. "The Development of Fan Filter Unit with Flow Rate Feedback Control in a Cleanroom." *Building and Environment*. Lawrence Berkeley National Laboratory Report, LBNL-60530.

[12] ASHRAE. 1999. *ANSI/ASHRAE Standard 51-1999*. Also AMCA 210-99. Laboratory Methods of Testing Fans for Rating.

[13] ASHRAE Standard 41.2: Standard Methods for Laboratory Airflow Measurements

[14] ASHRAE Standard 37: Methods of Testing for Rating Unitary Air Conditioning and Heat Pump Equipment

[15] International Organization for Standardization (ISO). 2003. ISO 5167. “Measurement of fluid flow by means of pressure differential devices inserted in circular cross-section conduits running full.” Part 1: General principles and requirements; Part 2: Orifice plates; Part 3: Nozzles and Venturi nozzles; Part 4: Venturi tubes.”

9. Appendices

9.1 Initial Test Setup

It is essential to obtain accurate measurements of total electric power demand, airflow rates (or actual airflow speeds), and the static (and total) pressure differential across the FFU. Among these, accuracy and complexity in airflow rate, power demand, and pressures, and locations of pressure measurement applicable to FFU are important considerations in test rig setup and equipment selection.

To determine the initial dimensions of the test duct section, we referred to previous work and a number of ASHRAE Standards on airflow and pressure measurements. For example, ASHRAE Standard 51 (i.e., ANSI/AMCA 210): Laboratory Methods for Testing Fans for Aerodynamic Performance Rating provides guidance that specifies a minimal duct length of 10 times of hydraulic diameter [12]. According to this, the ductwork leading to a 61-cm by 122-cm (2-ft by 4-ft) FFU section would require a minimal length of 8-m (27-ft). As an alternative, it is possible if the duct section were to be reduced to a smaller size in order to reduce its minimal length. In the latter case, however, the required length of transition piece would still be equally significant if not longer because the tapering angle is required to be no greater than 7½° (degree).

In addition, some other standards provide different requirements for measurement locations of the external pressures[13][14][15]. For example, one ASHRAE standard specifies test duct section with the same sectional dimensions of the unit being tested, with a duct length no less than 2½ times the square root of the product of the duct dimensions [13]. For a 61-cm by 122-cm (or 2-ft by 4-ft) duct section, the minimal duct length would then be about 2.2-m (7-ft). The pressure taps are located at 4/5th of the duct length away from the tested unit, which is 1.7-m (5.7-ft). It also requires a downstream duct straightener for the flow meter, with a length of five to five and a quarter times the diameter of the cross-section circular duct.

Figure 11 illustrates the initial experimental layout for measuring airflow rate, static (and total) pressure, and total electric power demand. The unit's airflow rate, airflow speed, total electric power demand was concurrently recorded for all operating points adjusted by varying the pressure differential and airflow rate across the FFU. The airflow path was designed as a flow-through to FFU in the test rig. This conceptual setup shows that flow-nozzle(s) are used to measure airflow rate in the upstream of airflow path directed toward the FFU inlet. In this project, the FFU face was installed to be vertical. Airflow from immediate downstream of the FFU was discharged to the atmosphere. Airflow rate measurement contains a single-nozzle for recording airflow rates through the tested unit. Figure 10 shows the actual test rig setup at LBNL.

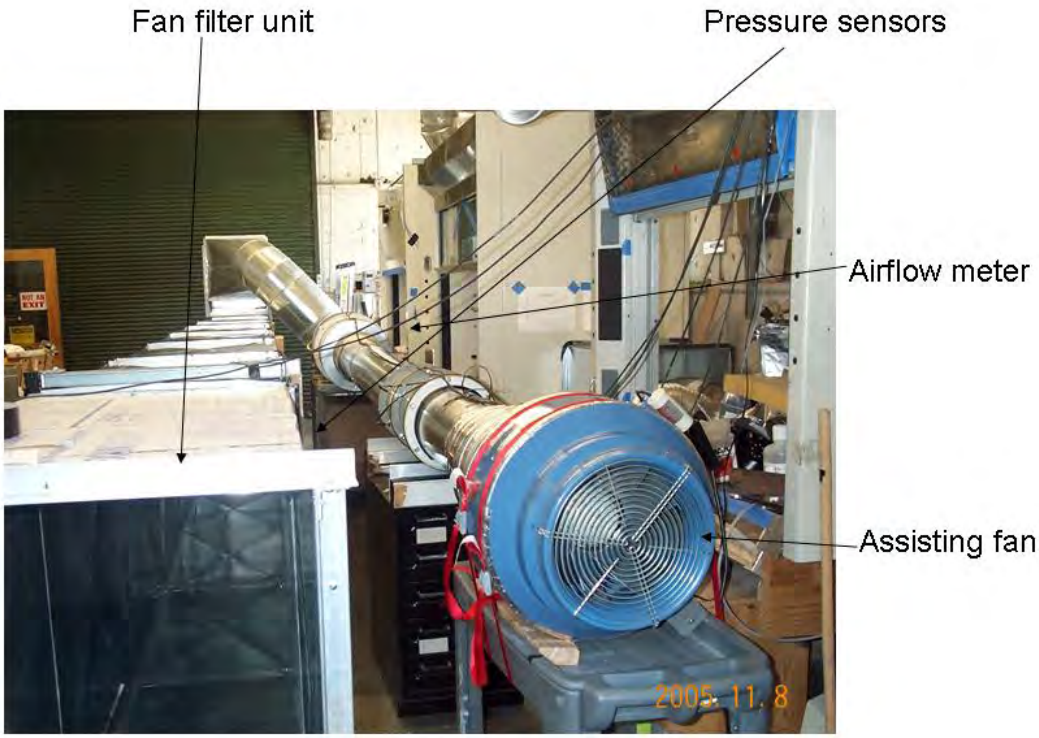

Figure 10. Laboratory test rig at LBNL

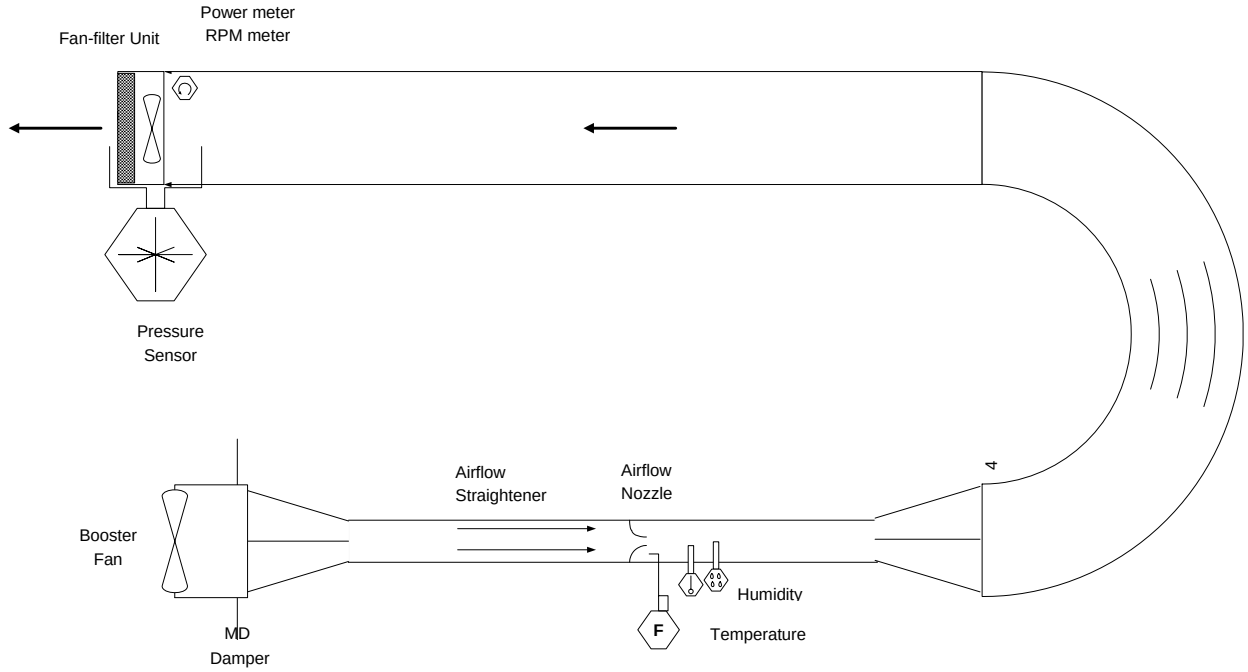

Figure 11 FFU test setup

Rigid metal ductwork chamber was built and installed for the test rig. The FFU tested was mounted vertically on the exit end of a straight duct. The size of ductwork connected to the FFU was 61-cm by 122-cm (or 2-ft by 4-ft), approximately equal to the FFU section. The length of straight duct upstream of FFU was 10-m (33-ft), which was approximately 12 times of the hydraulic diameter of the ductwork connecting to the unit.

Along the duct there are three pressure taps installed using Pitot Tube to monitor the pressure difference between the selected locations inside the ductwork and where the air is discharged from the unit (outside of the duct). The measurement points of the Pitot Tubes were located at the center of four duct sections, with a distance of 0.3-m, 2.6-m, 4.9-m, 8.7-m (1-ft, 8.5-ft, 16-ft, and 29-ft) away from the FFU airflow inlet. Exit airflow from the downstream of the FFU was directed to the atmosphere from the filter face. The pressure taps were installed at the following distance away from the inlet of test FFU: 0.3-m(1-ft), 2.6-m (9-ft), 4.9-m (16-ft), and 8.7-m (29-ft).

Recommended by the flow-nozzle manufacturer, there was a recommendation of minimal duct length for airflow uniformity in flow nozzle. We however initially installed a 5-ft duct straight round duct with a diameter of 10” upstream of the flow meter, and an 8-ft duct straight round duct with a diameter of 10” downstream of the flow meter.

The initial test rig setup contained a combination of conservative requirements by existing ASHRAE/AMCA standards, and was used to evaluate the sample FFUs in this study.

9.2 Test Equipment and Device

In the tests for 61-cm by 122-cm (2-foot by 4-foot) fan filter units, we used a nozzle-based airflow meter to measure airflow rates ranging from 250 to 1300 cfm (7 to 37 m^3min^{-1}), corresponding to pressure signals ranging from 15 to 413 Pa with rated accuracies of 0.5 % of the reading. The airflow rate accuracy would be within 5% even when the pressure measurement error approaches one Pascal or slightly higher. Specifically, a nozzle Pitot airflow sensor was used.[1] It combines integral flow straightener, a flow nozzle and a Pitot averaging array to minimize the effects of turbulence, vortexing, swirl and profile shift. The flow straightener help align the flow and the nozzle conforms the airflow to a known profile shape. The Pitot averaging array installed at the throat of nozzle measured the conditioned velocity profile, the nozzle eliminates distortions in the velocity profile caused by upstream obstructions and doubles the velocity before it is measured by the Pitot averaging array, requiring no upstream straight run with ±0.5% reading accuracy.

An eight-channel electronic data logger was used to record measured parameters including airflow rate, air pressure, temperature, humidity, electric power demand, and fan-wheel speed for each of the testing conditions. Specifically, we used Energy Conservatory Model APT 8.[2] The output signals of data logger including pressure transducers were recorded with a computer-based data acquisition system. The eight-channel electronic differential pressure transducer has a measuring range of ± 400 Pa, with rated accuracies of the larger of ± 0.2 Pa or ±1% of reading. The calibration of the eight-channel pressure transducer system was checked using a micro-manometer that has a micrometer and electrical circuit for precisely measuring the height of the fluid column.

Additional pressure transducers were used to measure and monitor air pressures in various locations in the test rig. Pitot tubes were used to measure and record pressure differential before and after of the fan-filter unit. Pitot tubes were installed at various locations in the ductwork or

---

[1] http://www.thermo.com/com/cda/product/detail/0,1055,14801,00.html; Model NZP1031-10”-1-CF (circular carbon steel, standard flange and duct section, neoprene gaskets; 10-inch aluminum nozzle with 316 stainless steel probe array, aluminum flow straightener; Brandt standard flanges with industry standard bolt hole pattern with a companion flange.

[2] http://www.energyconservatory.com/products/products3.htm

chamber to measure or monitor of pressure changes along the ductwork or chamber leading to the fan-filter unit. Four such static pressure taps were installed at locations of 0.3-m (1-foot), 2.6-m (8.5-foot), 4.9-m (16-foot), 8.7-m (28.5-foot) away from the inlet of an FFU connected at the end the 10-m (33-ft) ductwork.

The fan-wheel speeds of the FFU were recorded using a device for recording number of rotations per minute (RPM) concurrently with other measurements for each of the test conditions. An RPM sensor[3] was installed at a fixed location upstream of the FFU inlet. The front-panel programmable RPM sensor (ACT-3) has an RS232 bi-directional interface. It includes NIST Traceable Certificate of Calibration, measuring 5 RPM - 999,999 RPM, with one or multiple pulses per revolution and scaling from 0.0001 to 99,999, totalizing and counting from 1 to 99,999. It has two alarm set points, latching or non-latching, with programmable hysteresis, low limit lockout, and analog voltage output of 0-5 Vdc with current output of 4-20 mA. The location was selected to allow direct monitoring of the speed of rotating fan-wheel blades. For various fan-filter units, additional painting on their fan-wheel blades, along with necessary accessories such as reflective tape, was necessary to enable proper sensing and recording of fan-wheel speeds.

A portable barometer[4] with digital display was used to record the atmospheric pressure around the test rig. The barometer has accurate pressure sensing with microprocessor-based computations to provide instantaneous pressure readings. The barometer has an accuracy within ± 0.02% F.S. (F.S. = 900 mm Hg abs, or 120 kPa). The device is temperature compensated over its operating temperature range. Sensors used to record the atmospheric pressure, air temperature and humidity were placed at locations within the laboratory that represent the psychometric conditions of the air flowing in and out of the test rig.

The total true mean electric power demand for the whole unit included fan motor, speed control and display device, transformer, and additional accessories. Total electric power demand was measured using an accurate electric power meter for one-phase AC power. It measured true power ranging from 0 to 500 W with an accuracy of the larger of ±0.2% of reading and ±0.04% of the full scale.[5] For example, the accuracy of the power meter measuring power demand of 100

---

[3] http://www.monarchinstrument.com/act.htm#Model%20ACT-3

[4] http://www.novalynx.com/230-355.html

[5] http://www.ohiosemitronics.com/products/gw5.html; http://www.ohiosemitronics.com/pdf/gw5.pdf;

W and 300 W would be ±0.2 W and ±0.6 W, respectively. We calibrated the actual electric power demand readings against a highly accurate power meter with ±0.05% basic accuracy.[6] It was accurate with distorted waveforms and poor power factors and had a built-in integrator for measuring energy maximum demand or averaged values. The maximal reading deviation of the power meter used in FFU testing from that of the highly accurate meter was within 2%.

In order to control the airflow rate and pressure differential across the FFU tested, we used an ancillary booster fan[7] and a damper[8] to modulate static (and total) pressures and airflows across the FFU. The airflow rates under a certain fan-wheel speed setting were adjusted by two means: 1) varying fan-wheel speeds of the booster fan thereby changing airflow rates, and 2) controlling the pressure across the airflow damper of which the position affects the pressure loss as a way to emulate the changing pressure resistance in the external system. The booster fan is used to emulate various external resistances against which the unit may be able to overcome. The booster fan included a variable-speed controller for use of adjusting the airflow rates. Details of selecting operating conditions and the testing procedures are described in the standard test method for FFUs with different speed control techniques: 1) FFU with a single-speed-drive motor, 2) FFU with a multi-speed-drive motor, and 3) FFU with a variable-speed-drive (VSD) motor.

9.3 Leakage Detection and Quantification

The equipment setup and test rig configuration was expected to ensure minimal air leaks between the enclosed ductwork/chamber and its surrounding environment - the ambient environment external to test rig. The leakage should be maintained at a minimal level or accurately quantified so that airflow rate through the FFU can be accurately quantified.

The air leakage was first quantified and evaluated by performing the following 1) sealing the test rig (including FFU), 2) connecting a low-flow flow meter coupled with Pitot tube, 3) measuring leaking airflow rates corresponding to various air pressure differential across the ductwork.

---

[6] http://www.atecorp.com/Equipment/Voltech/PM3000A.htm; Voltech Universal Power Analyzer (PM3000A).

[7] http://www.kbelectronics.com/catalog_fan_oem_wall.htm; Kanalflakt Fan. KBWC-115, 115 VAC 50/60 Hz Rated 15.0 amps.

[8] http://www.continentalfan.com;

The measured airflow rate corresponding to certain pressure differential was used to quantify airflow leakage ratio, defined as leaking airflow rate divided by the total airflow rate at the pressure inside the ductwork of the test rig. Measures to reduce leakage of test rig were then carried out to minimize leakage ratio for tested operating conditions.

Specifically, the leaking airflow rate was measured using a differential-pressure Venturi-type flow meter. The Venturi-type flow meter is a Style VS brass-screwed meter with a diameter of 3.8-cm (1.5-inch) and a length of 10.8-cm (4.25-inch) with an accuracy of ±1%[9]. With a fan-filter unit attached and tested, we detected air leakage in the ductwork of the initial test rig that was constructed. The ductwork including all traverse joints, longitudinal seams, and duct wall penetrations created for wire or cable connections was then manually sealed using sealant.

Additional leakage tests and sealing were performed accordingly to quantify the leakage levels. After some iteration, we reduced the maximal air leakage ratio well within 3% when the airflow rate was 8.5 m^3/min (300 cfm) while the pressure across the ductwork of the test rig was 140 Pa (0.56-inch water). In principle, leaking airflow rates of the ductwork were quantified by the following equation.

Where $Q_{ductleak}$ is the duct leakage rate (cfm or m^3/s), C is the constant reflecting area characteristics of leakage path; ΔP_s is the static pressure differential between interior and exterior of the ductwork; and N is the exponent relating to turbulent or laminar flow in leakage path . Actual leakage ratio became much lower than the measured 3% when the airflow rates through the unit were higher than 8.5 m^3/min (300 cfm), while the pressure across the ductwork was maintained the same or even lowered. For example, the leakage ratio of the test rig became 1.5% down to 1% when the actual airflow rate was set at 17 m^3/min (600 cfm) while the pressure across the ductwork of the test rig was 140 Pa (0.56-inch water) or down to 70 Pa (0.28-inch water), respectively. On the other hand, the leakage ratio of the test rig became 2.2% when the actual airflow rate was set at 17 m^3/min (600 cfm) while the pressure across the ductwork of the test rig was set higher at 250 Pa (1.0-inch water). In the worse case scenario, the leakage ratio became 4.4% if the pressure differential increased to 250 Pa (1.0-inch water) while airflow rate was maintained at the level of 8.5 m^3/min (300 cfm).

Based upon the leakage quantification described above, we would be able to calculate the leakage airflow rates for each of the operating conditions and modify the actual airflows through the FFU tested. However, given that common airflow rates for a 61-cm by 122-cm (2-foot by 4-

[9] http://www.gerandengineeringco.com/

foot) unit would be higher than 8.5 m^3/min (300 cfm), actual leakage airflow rates associated with the test rig would normally appear to be within 3% of the measured airflow rates. The air leakage was therefore considered negligible in this test rig because of the small deviation from actual airflow rate measured by the flow nozzle.

9.4 Test setup for evaluating repeatability

Prior to the study, we reviewed open literature including ISO Standards, AMCA and ASHRAE Standards that address the requirements for measuring airflow rate and air pressures associated with different equipment in ducted or un-ducted testing layouts.[12] [13][14][15].

These standards are either applicable to characterize fan performance, or specific device such as air conditioners or heat pumps in various laboratory setting. For example, the AMCA/ASHRAE standard requires large chambers and flow baffles to create uniform airflows entering the fans to be tested. However, none of these standards addresses any specific method applicable to characterize fan filter units or provides any direct guideline on an acceptable laboratory method to measure power demand, airflow rate, and pressure differential across the fan filter units.

In order to understand the potential influence of various setup, ambient conditions, and layout of the test rig on laboratory measurements, we designed and installed various test rigs to address potential implications of such variations on the repeatability of measured performance. We then designed a set of experiments and evaluated the significance of difference in measured results. In particular, the experiments were to evaluate consistency and repeatability of test results through characterizing the same units at different ambient conditions, at opposite airflow directions, and in test rigs with different sizes such as duct lengths. The following summarize the approach and results.

9.4.1 Repeatability of measured performance at different flow restrictions created by placement of an FFU in opposite directions

In order to evaluate the impacts of airflow restrictions at the inlet and outlet of an FFU on measured performance, we placed selected FFUs with HEPA/ULPA facing into the ductwork with a section size of 61-cm by 122-cm (2-ft by 4-ft). The FFU inlet was un-ducted and took air from the open ambient space through a prefilter. With this FFU setting, we performed experiments on airflow, pressure, and power demand, and compared the results with those from original test setup whereby ducted air was directed into the FFU inlets and flew through the HEPA/ULPA filter that faced open air.

The new test rig setup is illustrated in Figure 12. In this setup, the direction of the FFU tested, the flow meter, and the booster fan were all reversed compared to its original setup illustrated in Figure 11. In this test setup, the supplied airflows were free-flow before entering the FFU inlet while the airflows out of the HEAP/ULPA filter are restrictive by the sized ductwork, i.e., 61-cm

by 122-cm (2-ft by 4-ft) ducted chamber. In this case, the airflow rates were measured downstream of the FFU outlet.

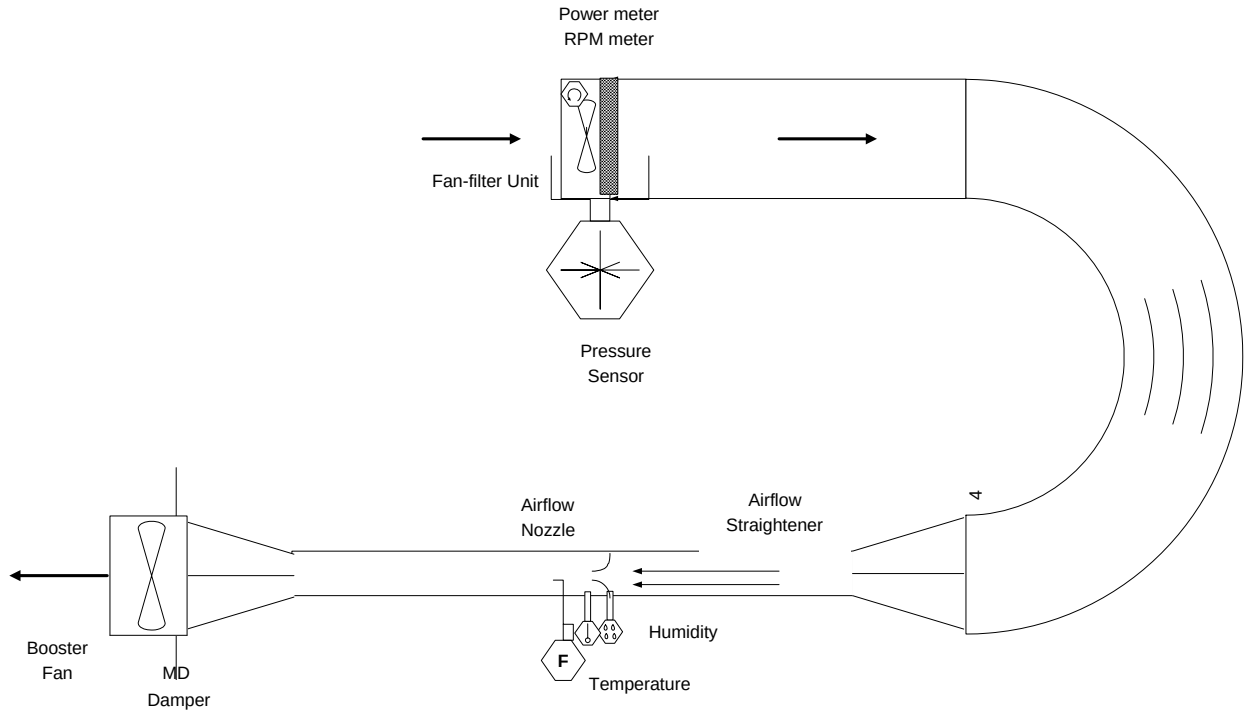

Figure 12 Test rig with varied airflow restrictions at FFU inlet and outlet

A comparison of the measured electric power demand of the two selected FFUs from both test rig setup (Figure 11 and Figure 12) shows that the relative difference of measured power demand was within 2% or less for each FFU sample between the two test-rigs. The relative difference in measured total pressure efficiency is even smaller, mostly within 1% at the operable conditions listed in Table 1 for each unit.

Based on the measured data and evaluations, we have found that the measured performance corresponding to two different flow restrictions at FFU inlet and outlet was repeatable between the two rigs, without inducing variance in uncertainties. The testing results show that while there were variations in airflow restrictions at the inlet and outlet of an FFU, placing an FFU with HEPA/ULPA facing into a restrictive, ducted chamber or to a non-restrictive, open space produced repeatable measurements of FFU airflow rates, pressures, and power demand for a same operating condition. The results in total electric power demand as well as total pressure efficiency were consistent, even though the regression equations differed from each other.

In summary, the influence of the flow directions and restrictions at FFU inlet and outlet on the measured electric power demand or total pressure efficiency appeared to be minimal and was negligible.

9.4.2 Repeatability of measured performance at different ambient conditions

One would expect that variations in motor temperatures might affect the actual efficiency or electric power demand of fan motors. If other parameters remain the same for the same FFU, ambient air temperatures might influence the steady-state motor temperature. Therefore, large difference in ambient temperatures might have some affect on operating motor efficiency.

In order to examine the hypothesis, we selected two FFUs and repeated measurements of energy performance of each unit at different ambient temperatures: the first unit was tested in a relatively warm day in summer where the inlet air temperature was 26°C and was tested again in the fall when it was cooler (21°C); the second unit was tested in a mild summer day when the inlet air temperature was 21°C while it was tested again in a cooler day where the inlet air temperature was 17°C.

For the first unit, we have identified slightly higher power demand for the unit when operating at warmer temperatures (26°C compared to 21°C). This was contrary to what would have been expected in that the fan motor would be more energy-efficient if at all under a warmer condition. However, the measured relative difference was within 3% for power demand (Watt) and was within 2% for the total pressure efficiency under two ambient air temperatures for the operable conditions listed in Table 1. The difference was within measurement uncertainties.

For the second unit, we have identified slightly higher power demand for the unit when operating at a cooler temperature (17°C compared to 21°C). This was in line with what would have been expected in that the fan motor would be less efficient when the motor was cooler. However, the measured relative difference was within 2% for power demand (Watt) and was within 2% for the total pressure efficiency under two ambient air temperatures for the operable conditions listed in Table 1. The difference was within measurement uncertainties.

While large difference in ambient air temperatures might somewhat cast initial influence on the actual power demand of the FFU motor, the experimental results showed insignificant influence of the recorded ambient air temperatures on measured power demand and actual total pressure efficiency measured from the experiments.

From the investigation and performance comparisons, we recommend that measures should be taken to ensure FFU testing be conducted under a relatively constant ambient temperature, which normally would call for warm-up time for the test rig and FFU to operate at a stable condition. In addition, we should avoid testing FFU at extreme air temperatures, and even better, maintain the ambient air temperature for the FFU test rig within a certain range, e.g., normal room temperatures.

9.4.3 Repeatability of measured performance measured with different lengths of the ductwork leading to an FFU

Various ASHRAE and AMCA Standards specify different requirements for duct lengths leading to tested units (air conditioner, coil, fan, etc.) for airflow rate and pressure measurements, but none of them addresses fan-filter units. The initial test rig setup in this study adopted perhaps a more conservative approach in terms of minimum duct length to be recommended. The initial straight ductwork was composed of eight 61-cm by 122-cm (2-ft by 4-ft) sections that were connected in sequence, each having a length of 60-cm (46-inch). This equaled to approximately 12 times of the duct's hydraulic diameter (32-inch).

In order to investigate the effect of lengths of the ductwork leading to FFU inlet or outlet on the measured performance of fan filter units, we reduced the length of the 61-cm by 122-cm (2-ft by 4-ft) ductwork by removing some of the sections, and conducted experiments respectively. The shortened duct lengths ranged from twelve- down to 1.5-times of duct's hydraulic diameter ($12D_h$ down to $1.5D_h$), respectively.

In all these experiments, we placed the selected FFUs with HEPA/ULPA facing into ducted chamber and recorded airflow rates, pressures, power demand. Figure 13 illustrates the test rig setup. In this test setup, airflows were free-flow before entering the FFU inlet while the airflows out of the HEAP/ULPA are restrictive, i.e., 61-cm by 122-cm (2-ft by 4-ft) ducted chamber. The airflow rates were measured downstream of the FFU outlet.

Based on the measured data and comparisons, we have found that the measured performance corresponding to different duct lengths was repeatable without inducing variance of accuracy concerns. The testing results show that while there are variations required duct lengths for different equipment suggested by some ASHRAE/AMCA standards, we found no significant difference in the measured energy performance of fan filter units from our study. The required ductwork length is found to be acceptable at 1.5 D_h level. The results in total electric power demand as well as total pressure efficiency are consistent, even the regression equations differed.

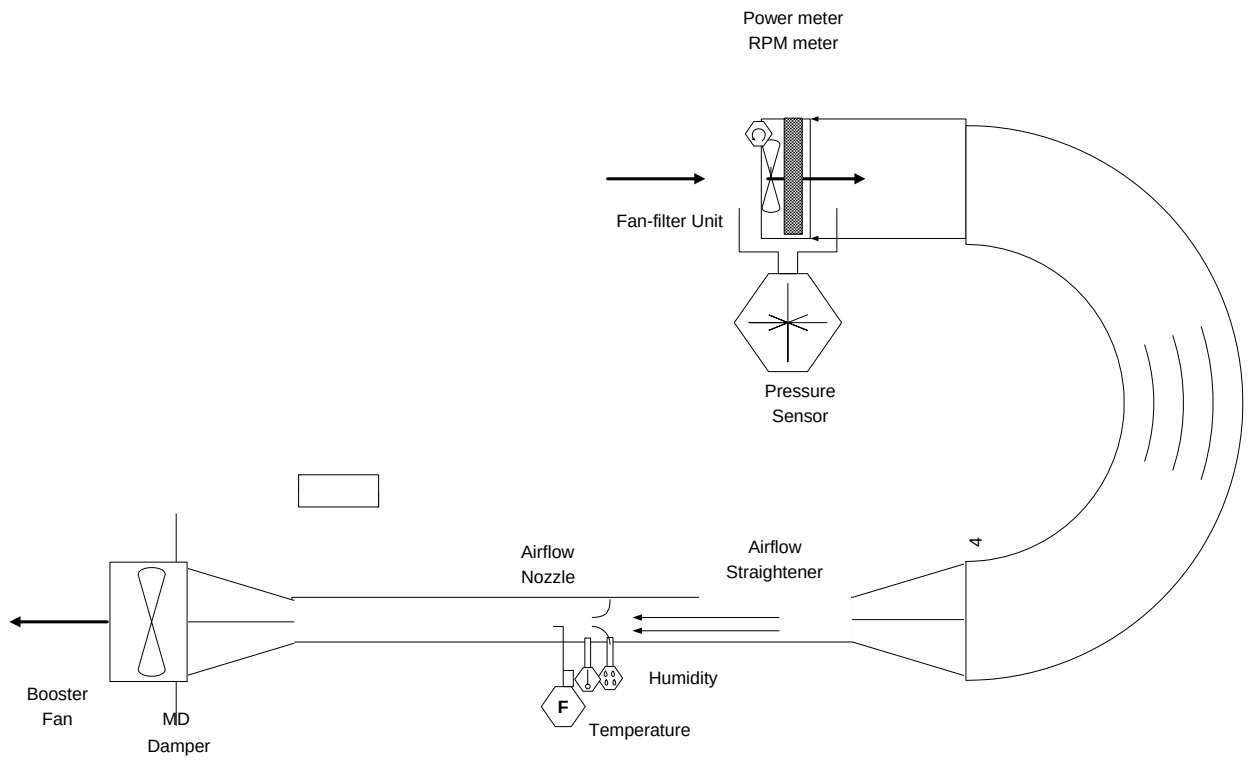

Figure 13 Test rig with a shortened duct length

The findings and evaluations are significant in that 1) the size of test rig can be practically reduced significant, requiring minimal laboratory space to set up the test rig while saving costs (material, labor, space), and 2) having the minimal duct length requirement would allow and emulate FFU application in a cleanroom where the space/duct leading to the FFU is normally minimal.

9.5 FFU Characterization Reports

For each of the FFU tested in LBNL, we generate individual reports on the characterization for each FFUS studied in this project. As a benefit for participants in this research project, each report was submitted to its respective supplier in confidence. This appendix includes the results from laboratory and modeling characterization, and contains no specific information on FFU models or manufacturer names. Each report was numbered or dated randomly, i.e., LBID 2588-1 through 17, without any indication about the order in FFU receipt or testing dates whatsoever.

Report on

Laboratory Evaluation of Fan-filter Unit

LBNL Report, LBID-2588-1

Prepared by

Tengfang Xu and Duo Wang

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

July 25, 2006

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Acknowledgements

The project is funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Report on Laboratory Evaluation of Fan Filter Unit
Lawrence Berkeley National Laboratory, LBID-2588-1

Reporting Characteristics of the Unit Based on Manufacturer's Shipment

Manufacturer
Brand/Model
Serial Number
Unit Size
Fan Motor
Fan Wheel

Filter Efficiency

LBNL Tracking #

Reporting Energy Performance Based on Laboratory Testing

This report is based upon the laboaratory evaluation of energy and airflow characteristics of the fan filter unit that was supplied to Lawrence Berkeley National Laboratory. All measured parameters, when applicable, are converted to their equivalents at the standard atmospheric condition (1 ATM, 20°C, sea level).

Total electric power demand, airflow rates, and pressure differential across the fan filter unit was recorded for a range of operable conditions of the unit. Normally, test conditions were selected with the airflow rates no less than 300 ft^3/min (8.5 m^3/min) under standard condition. The speed controller was set at its highest setting followed by lower setting. Figure 1 shows operable conditions that were tested with various speed control setting.

Figure 1. Operable conditions with various speed control setting

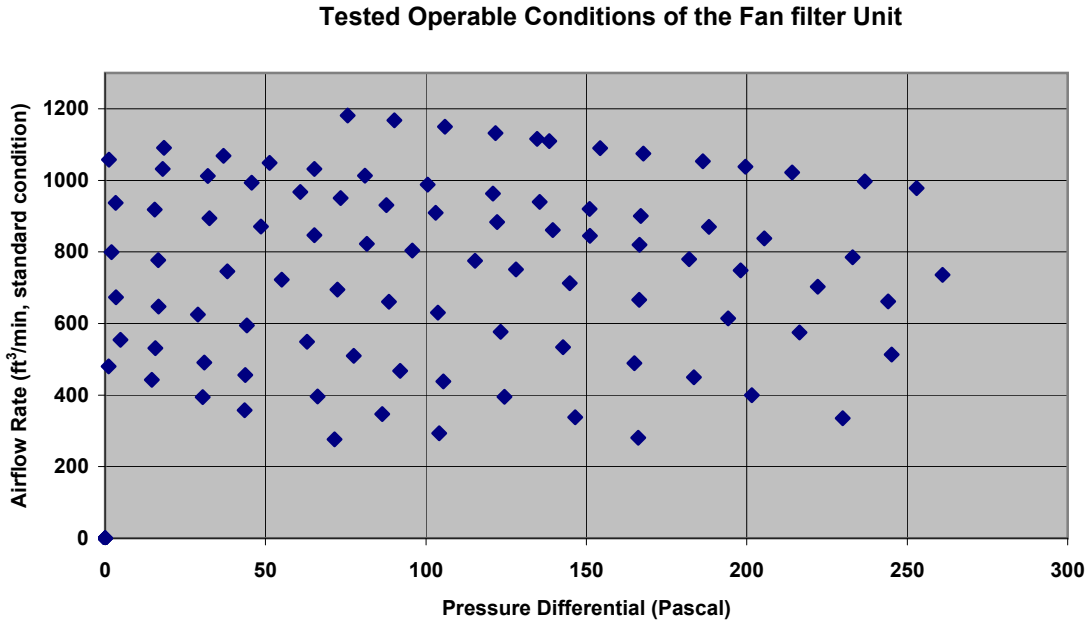

Total Electric Power Demand

Total electirc power demand (EPD) of the fan filter unit includes all the electric power necessary to operate the fan filter unit. It was measured concurrently with airflow rates and pressures under all testing conditions. As a result, the total electirc power demand can be calculated using the following equation.
The equation is derived from laboratory testing results from the operating conditions as shown in Figure 1. The R-square of the regression is included, which explains the statistical significance of the power demand predicted by the equation. A higher R-square number (with the possible maximum of 1) indicates higher degree of confidence in the power demand value derived from the laboratory testing.
Figures 2 illustrates the magnitude of total electric power demand within a selected range of operating conditions.

Electric Power Demand (Watt) = $C_0 + C_1 \times D_p + C_2 \times Q + C_{11} \times D_p^2 + C_{12} \times D_p \times Q + C_{22} \times Q^2$

R^2 = 0.999

where D_p is the pressure differential across the fan filter unit, in Pascal

Q is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)

C_0 =	8.45E+01
C_1 =	-4.21E-02
C_2 =	-2.67E-01
C_{11} =	3.99E-04
C_{12} =	8.12E-04
C_{22} =	4.16E-04

Figure 2 Total electric power demand of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

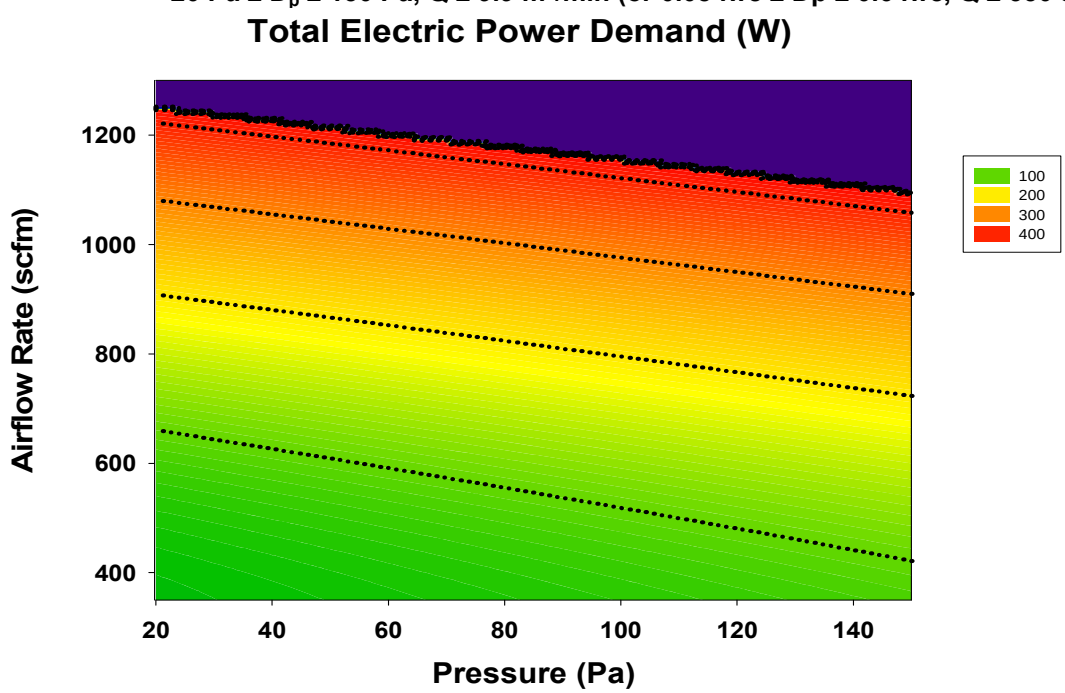

Total Pressure Efficiency

Total Pressure Efficiiency (TPE) is calculated by multiplying airflow rate and pressure differential across the unit then divided by total electric power demand as shown in the following equation.
Figures 3 illustrates the magnitude of total pressure efficiency within a selected range of operating conditions.

Total Pressure Efficiency (%) = 0.000471947443 x Q x D_p / EPD

where	Q	is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)
	D_p	is the pressure differential across the fan filter unit, in Pascal
	EPD	is the total electirc power demand (EPD) of the fan filter unit include all the electric power necessary to operate the fan filter unit, in Watt

Figure 3 Total pressure efficiency of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

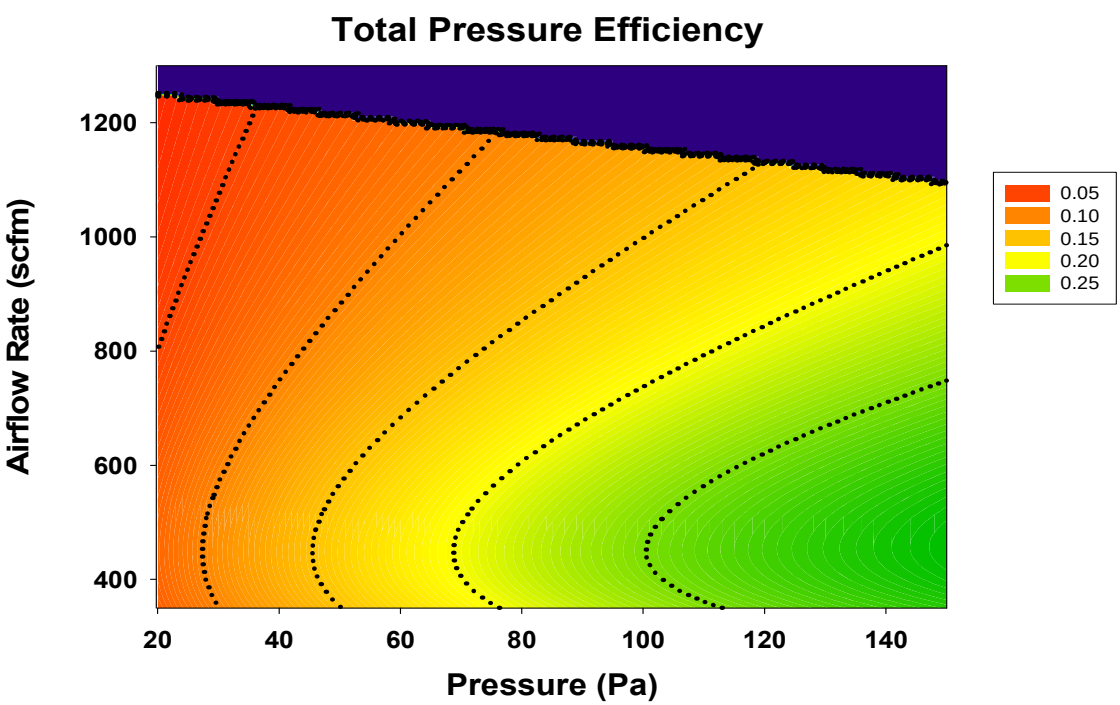

Note of unit conversion:

1 kPa = 4.015 inch water column
1 m/s = 196.9 feet per minute (fpm)
1 m^3/min = 35.3 ft^3/minute (cfm)

Report on

Laboratory Evaluation of Fan-filter Unit

LBNL Report, LBID-2588-2

Prepared by

Tengfang Xu and Duo Wang

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

August 8, 2006

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Acknowledgements

The project is funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Report on Laboratory Evaluation of Fan Filter Unit
Lawrence Berkeley National Laboratory, LBID-2588-2

Reporting Characteristics of the Unit Based on Manufacturer's Shipment

Manufacturer
Brand/Model
Serial Number
Unit Size
Fan Motor
Fan Wheel

Filter Efficiency

LBNL Tracking

Reporting Energy Performance Based on Laboratory Testing

This report is based upon the laboaratory evaluation of energy and airflow characteristics of the fan filter unit that was supplied to Lawrence Berkeley National Laboratory. All measured parameters, when applicable, are converted to their equivalents at the standard atmospheric condition (1 ATM, 20°C, sea level).

Total electric power demand, airflow rates, and pressure differential across the fan filter unit was recorded for a range of operable conditions of the unit. Normally, test conditions were selected with the airflow rates no less than 300 ft^3/min (8.5 m^3/min) under standard condition. The speed controller was set at its highest setting followed by lower setting. Figure 1 shows operable conditions that were tested with various speed control setting.

Figure 1. Operable conditions with various speed control setting
Tested Operable Conditions of the Fan-filter Unit

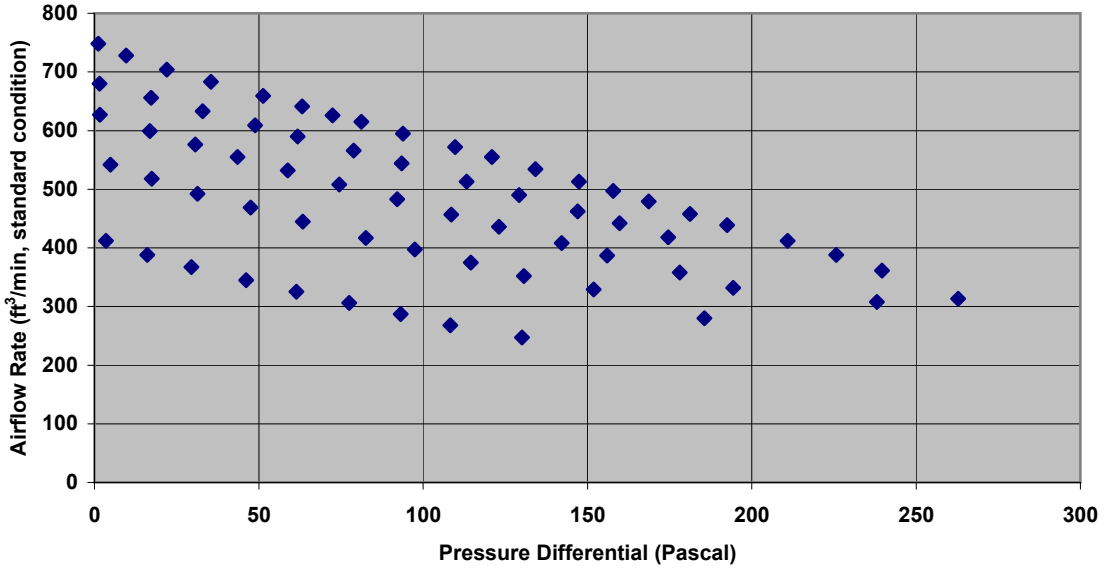

Total Electric Power Demand

Total electirc power demand (EPD) of the fan filter unit includes all the electric power necessary to operate the fan filter unit. It was measured concurrently with airflow rates and pressures under all testing conditions. As a result, the total electirc power demand can be calculated using the following equation.

The equation is derived from laboratory testing results from the operating conditions as shown in Figure 1. The R-square of the regression is included, which explains the statistical significance of the power demand predicted by the equation. A higher R-square number (with the possible maximum of 1) indicates higher degree of confidence in the power demand value derived from the laboratory testing.

Figures 2 illustrates the magnitude of total electric power demand within a selected range of operating conditions.

Electric Power Demand (Watt) = $C_0 + C_1 \times D_p + C_2 \times Q + C_{11} \times D_p^2 + C_{12} \times D_p \times Q + C_{22} \times Q^2$

R^2 = **0.995**

where D_p is the pressure differential across the fan filter unit, in Pascal

Q is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)

C_0 =	1.75E+02
C_1 =	-4.33E-01
C_2 =	-1.00E-01
C_{11} =	6.78E-04
C_{12} =	9.26E-04
C_{22} =	1.58E-04

Figure 2 Total electric power demand of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

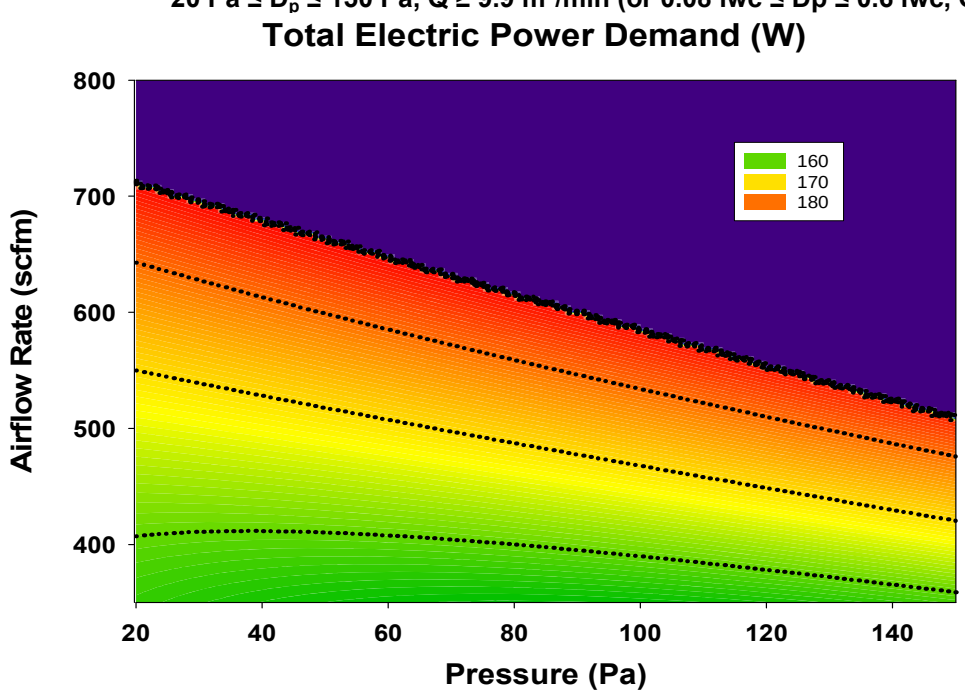

Total Pressure Efficiency

Total Pressure Efficiiency (TPE) is calculated by multiplying airflow rate and pressure differential across the unit then divided by total electric power demand as shown in the following equation.
Figures 3 illustrates the magnitude of total pressure efficiency within a selected range of operating conditions.

Total Pressure Efficiency (%) = 0.000471947443 x Q x D_p / EPD

where	Q	is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)
	D_p	is the pressure differential across the fan filter unit, in Pascal
	EPD	is the total electirc power demand (EPD) of the fan filter unit include all the electric power necessary to operate the fan filter unit, in Watt

Figure 3 Total pressure efficiency of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

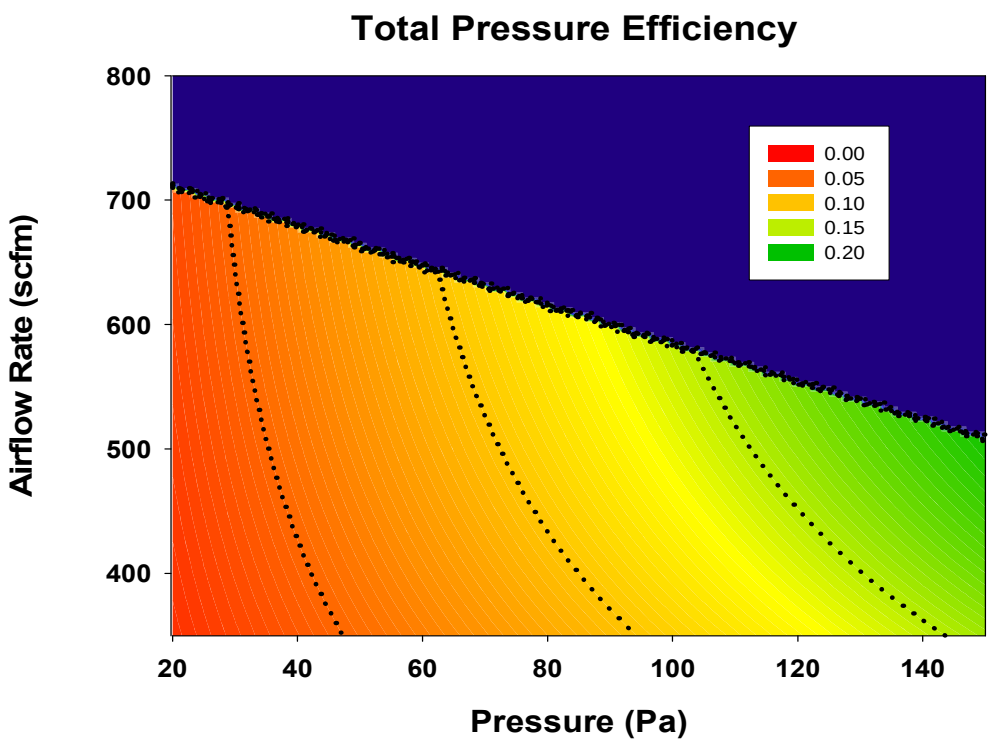

Note of unit conversion:

1 kPa = 4.015 inch water column
1 m/s = 196.9 feet per minute (fpm)
1 m^3/min = 35.3 ft^3/minute (cfm)

Report on

Laboratory Evaluation of Fan-filter Unit

LBNL Report, LBID-2588-3

Prepared by

Tengfang Xu and Duo Wang

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

August 8, 2006

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Acknowledgements

The project is funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Report on Laboratory Evaluation of Fan Filter Unit
Lawrence Berkeley National Laboratory, LBID-2588-3

Reporting Characteristics of the Unit Based on Manufacturer's Shipment

Manufacturer
Brand/Model
Serial Number
Unit Size
Fan Motor
Fan Wheel

Filter Efficiency

LBNL Tracking

Reporting Energy Performance Based on Laboratory Testing

This report is based upon the laboaratory evaluation of energy and airflow characteristics of the fan filter unit that was supplied to Lawrence Berkeley National Laboratory. All measured parameters, when applicable, are converted to their equivalents at the standard atmospheric condition (1 ATM, 20°C, sea level).

Total electric power demand, airflow rates, and pressure differential across the fan filter unit was recorded for a range of operable conditions of the unit. Normally, test conditions were selected with the airflow rates no less than 300 ft^3/min (8.5 m^3/min) under standard condition. The speed controller was set at its highest setting followed by lower setting. Figure 1 shows operable conditions that were tested with various speed control setting.

Figure 1. Operable conditions with various speed control setting

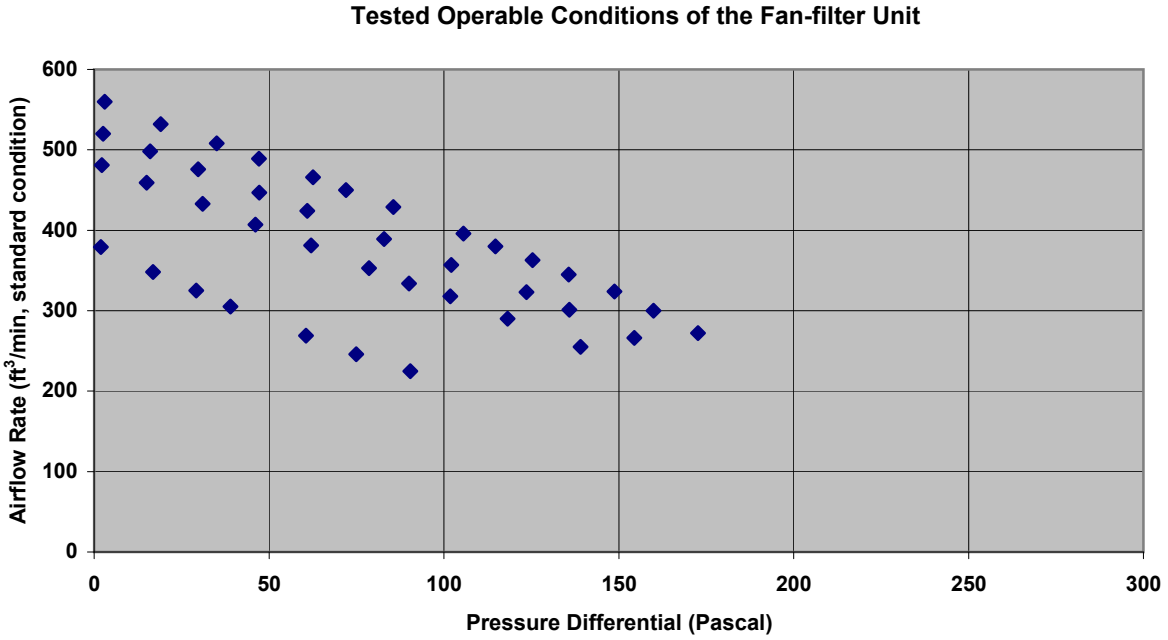

Total Electric Power Demand

Total electirc power demand (EPD) of the fan filter unit includes all the electric power necessary to operate the fan filter unit. It was measured concurrently with airflow rates and pressures under all testing conditions. As a result, the total electirc power demand can be calculated using the following equation.

The equation is derived from laboratory testing results from the operating conditions as shown in Figure 1. The R-square of the regression is included, which explains the statistical significance of the power demand predicted by the equation. A higher R-square number (with the possible maximum of 1) indicates higher degree of confidence in the power demand value derived from the laboratory testing.

Figures 2 illustrates the magnitude of total electric power demand within a selected range of operating conditions.

Electric Power Demand (Watt) = $C_0 + C_1 \times D_p + C_2 \times Q + C_{11} \times D_p^2 + C_{12} \times D_p \times Q + C_{22} \times Q^2$

R^2 = 1.000

where D_p is the pressure differential across the fan filter unit, in Pascal

Q is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)

C_0 =	1.45E+02
C_1 =	-7.18E-01
C_2 =	-4.38E-01
C_{11} =	1.78E-03
C_{12} =	2.40E-03
C_{22} =	7.23E-04

Figure 2 Total electric power demand of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

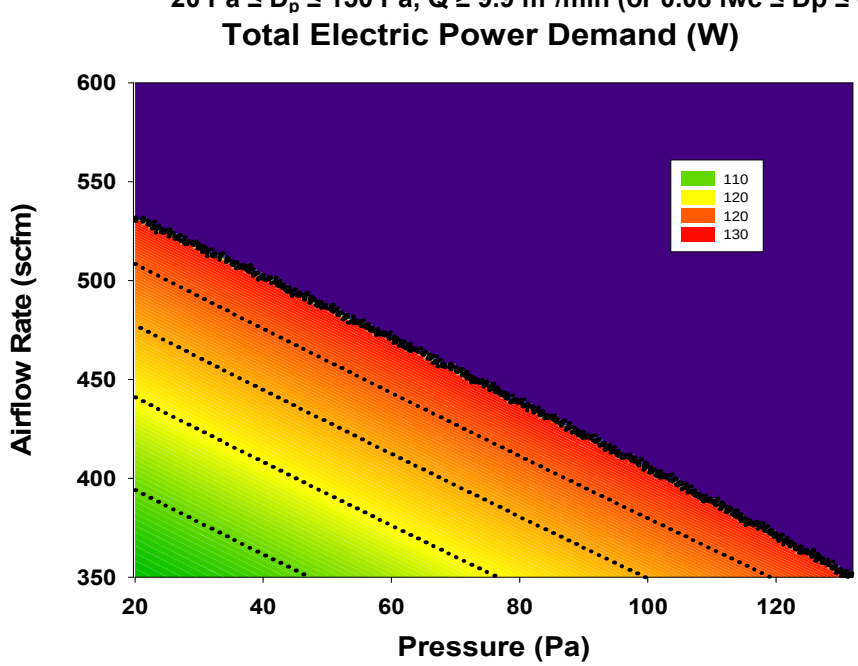

Total Pressure Efficiency

Total Pressure Efficiiency (TPE) is calculated by multiplying airflow rate and pressure differential across the unit then divided by total electric power demand as shown in the following equation.
Figures 3 illustrates the magnitude of total pressure efficiency within a selected range of operating conditions.

Total Pressure Efficiency (%) = 0.000471947443 x Q x D_p / EPD

where	Q	is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)
	D_p	is the pressure differential across the fan filter unit, in Pascal
	EPD	is the total electirc power demand (EPD) of the fan filter unit include all the electric power necessary to operate the fan filter unit, in Watt

Figure 3 Total pressure efficiency of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

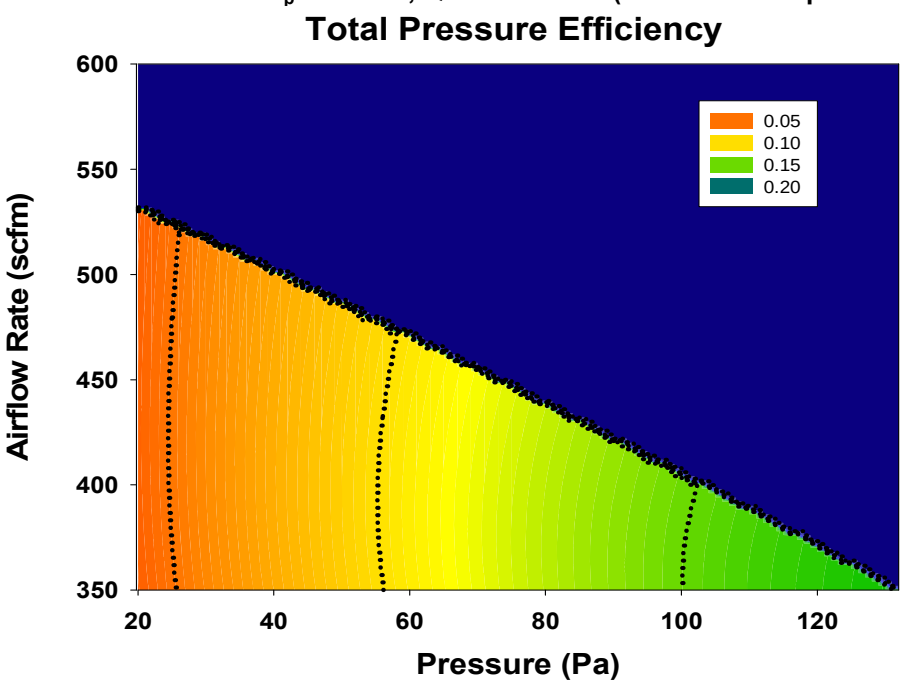

Note of unit conversion:

1 kPa = 4.015 inch water column
1 m/s = 196.9 feet per minute (fpm)
1 m^3/min = 35.3 ft^3/minute (cfm)

Report on

Laboratory Evaluation of Fan-filter Unit

LBNL Report, LBID-2588-4

Prepared by

Tengfang Xu and Duo Wang

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

August 10, 2006

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Acknowledgements

The project is funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Report on Laboratory Evaluation of Fan Filter Unit
Lawrence Berkeley National Laboratory, LBID-2588-4

Reporting Characteristics of the Unit Based on Manufacturer's Shipment
Manufacturer
Brand/Model
Serial Number
Unit Size
Fan Motor
Fan Wheel

Filter Efficiency

LBNL Tracking #

Reporting Energy Performance Based on Laboratory Testing

This report is based upon the laboaratory evaluation of energy and airflow characteristics of the fan filter unit that was supplied to Lawrence Berkeley National Laboratory. All measured parameters, when applicable, are converted to their equivalents at the standard atmospheric condition (1 ATM, 20°C, sea level).

Total electric power demand, airflow rates, and pressure differential across the fan filter unit was recorded for a range of operable conditions of the unit. Normally, test conditions were selected with the airflow rates no less than 300 ft^3/min (8.5 m3/min) under standard condition. The speed controller was set at its highest setting followed by lower setting. Figure 1 shows operable conditions that were tested with various speed control setting.

Figure 1. Operable conditions with various speed control setting

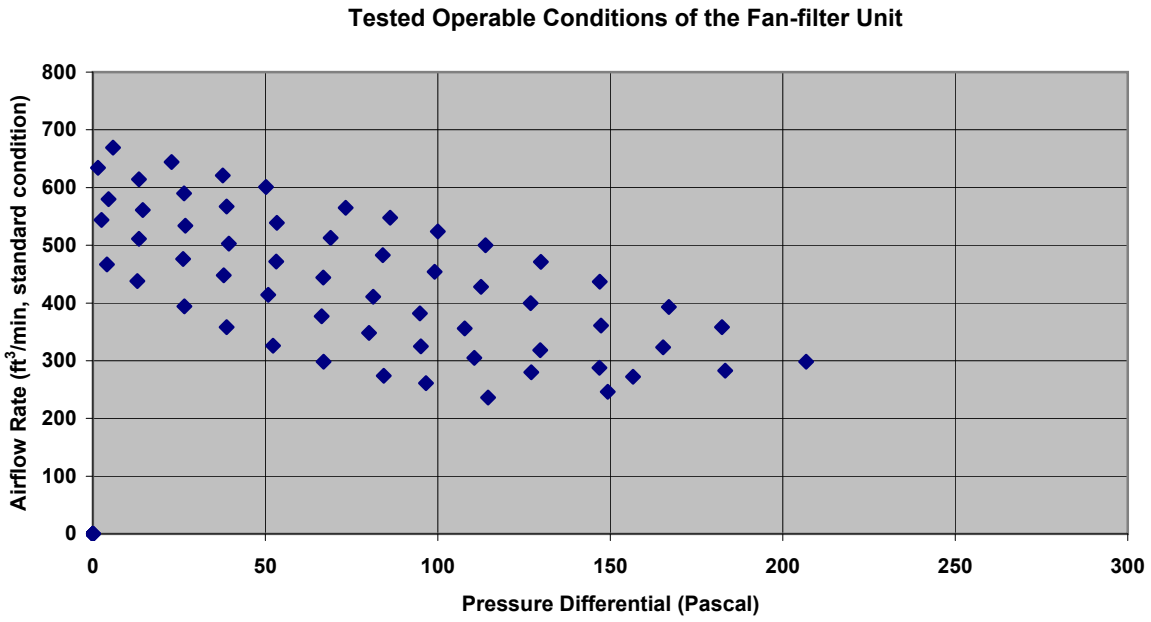

Total Electric Power Demand

Total electirc power demand (EPD) of the fan filter unit includes all the electric power necessary to operate the fan filter unit. It was measured concurrently with airflow rates and pressures under all testing conditions. As a result, the total electirc power demand can be calculated using the following equation.

The equation is derived from laboratory testing results from the operating conditions as shown in Figure 1. The R-square of the regression is included, which explains the statistical significance of the power demand predicted by the equation. A higher R-square number (with the possible maximum of 1) indicates higher degree of confidence in the power demand value derived from the laboratory testing.

Figures 2 illustrates the magnitude of total electric power demand within a selected range of operating conditions.

Electric Power Demand (Watt) = $C_0 + C_1 \times D_p + C_2 \times Q + C_{11} \times D_p^2 + C_{12} \times D_p \times Q + C_{22} \times Q^2$

R^2 = 1.000

where D_p is the pressure differential across the fan filter unit, in Pascal

Q is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)

C_0 =	1.22E+02
C_1 =	7.67E-02
C_2 =	1.37E-01
C_{11} =	-3.32E-04
C_{12} =	1.29E-04
C_{22} =	-1.22E-04

Figure 2 Total electric power demand of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

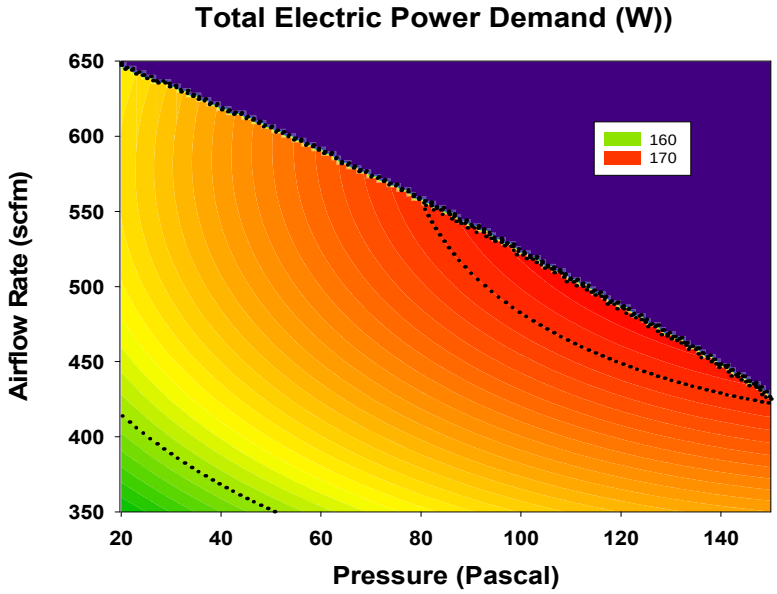

Total Pressure Efficiency

Total Pressure Efficiiency (TPE) is calculated by multiplying airflow rate and pressure differential across the unit then divided by total electric power demand as shown in the following equation.
Figures 3 illustrates the magnitude of total pressure efficiency within a selected range of operating conditions.

Total Pressure Efficiency (%) = 0.000471947443 x Q x D_p / EPD

where	Q	is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)
	D_p	is the pressure differential across the fan filter unit, in Pascal
	EPD	is the total electirc power demand (EPD) of the fan filter unit include all the electric power necessary to operate the fan filter unit, in Watt

Figure 3 Total pressure efficiency of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

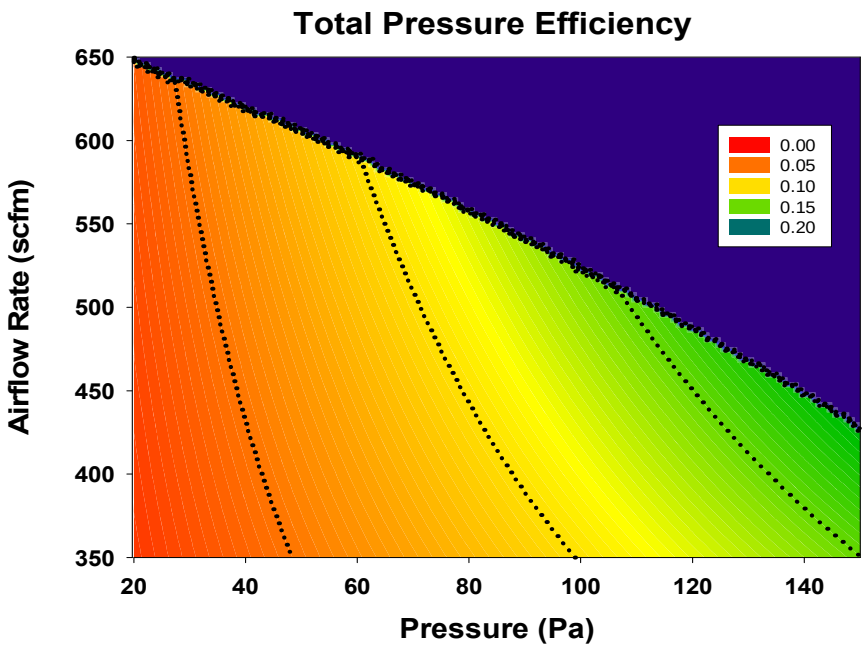

Note of unit conversion:

1 kPa = 4.015 inch water column
1 m/s = 196.9 feet per minute (fpm)
1 m^3/min = 35.3 ft^3/minute (cfm)

Report on

Laboratory Evaluation of Fan-filter Unit

LBNL Report, LBID-2588-5

Prepared by

Tengfang Xu and Duo Wang

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

August 11, 2006

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Acknowledgements

The project is funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Report on Laboratory Evaluation of Fan Filter Unit
Lawrence Berkeley National Laboratory, LBID-2588-5

Reporting Characteristics of the Unit Based on Manufacturer's Shipment
Manufacturer
Brand/Model
Serial Number
Unit Size
Fan Motor
Fan Wheel

Filter Efficiency

LBNL Tracking #

Reporting Energy Performance Based on Laboratory Testing

This report is based upon the laboaratory evaluation of energy and airflow characteristics of the fan filter unit that was supplied to Lawrence Berkeley National Laboratory. All measured parameters, when applicable, are converted to their equivalents at the standard atmospheric condition (1 ATM, 20°C, sea level).

Total electric power demand, airflow rates, and pressure differential across the fan filter unit was recorded for a range of operable conditions of the unit. Normally, test conditions were selected with the airflow rates no less than 300 ft^3/min (8.5 m3/min) under standard condition. The speed controller was set at its highest setting followed by lower setting. Figure 1 shows operable conditions that were tested with various speed control setting.

Figure 1. Operable conditions with various speed control setting

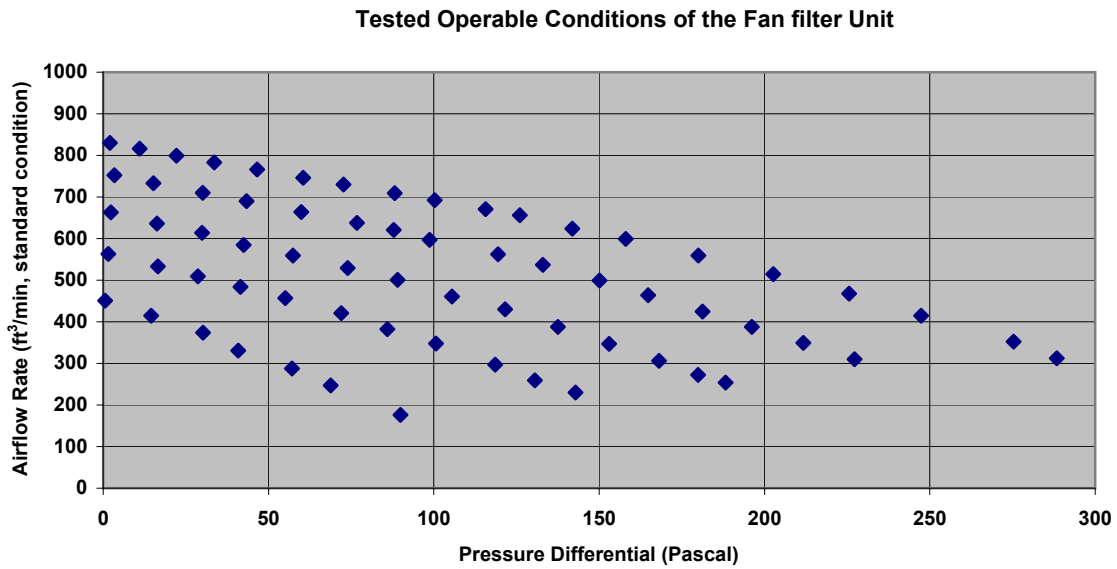

Total Electric Power Demand

Total electirc power demand (EPD) of the fan filter unit includes all the electric power necessary to operate the fan filter unit. It was measured concurrently with airflow rates and pressures under all testing conditions. As a result, the total electirc power demand can be calculated using the following equation.

The equation is derived from laboratory testing results from the operating conditions as shown in Figure 1. The R-square of the regression is included, which explains the statistical significance of the power demand predicted by the equation. A higher R-square number (with the possible maximum of 1) indicates higher degree of confidence in the power demand value derived from the laboratory testing.

Figures 2 illustrates the magnitude of total electric power demand within a selected range of operating conditions.

Electric Power Demand (Watt) = $C_0 + C_1 \times D_p + C_2 \times Q + C_{11} \times D_p^2 + C_{12} \times D_p \times Q + C_{22} \times Q^2$

R^2 = 0.999

where D_p is the pressure differential across the fan filter unit, in Pascal

Q is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)

C_0 =	7.68E+01
C_1 =	-1.34E-01
C_2 =	-3.32E-01
C_{11} =	7.09E-04
C_{12} =	1.44E-03
C_{22} =	6.42E-04

Figure 2 Total electric power demand of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

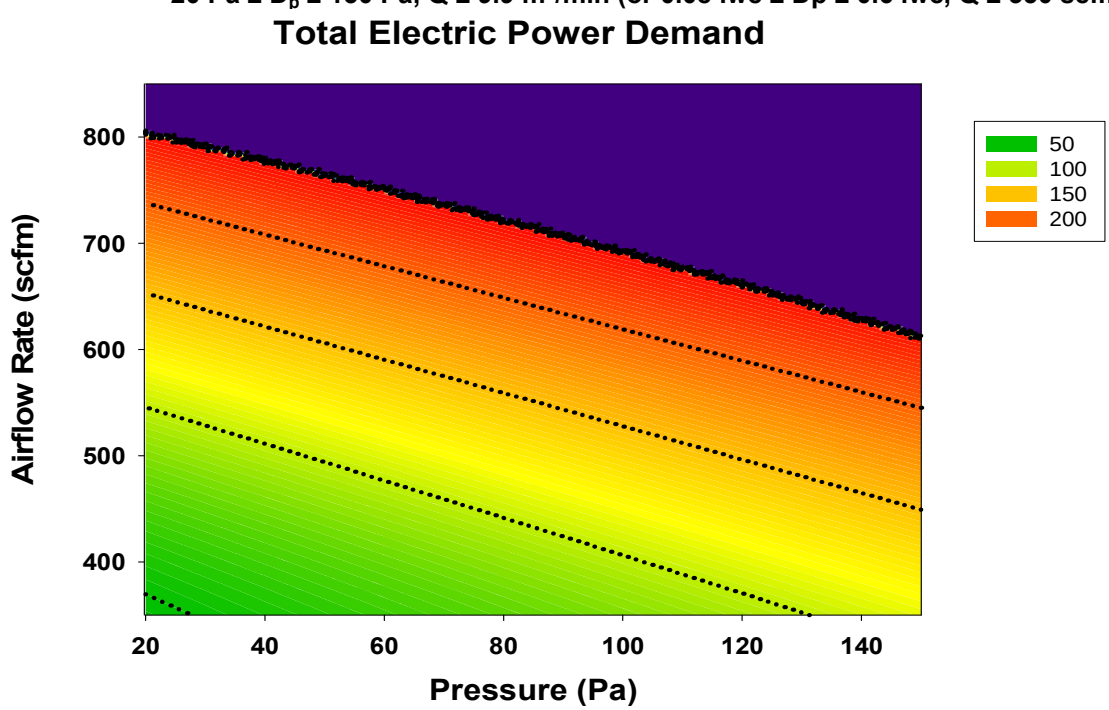

Total Pressure Efficiency

Total Pressure Efficiiency (TPE) is calculated by multiplying airflow rate and pressure differential across the unit then divided by total electric power demand as shown in the following equation.
Figures 3 illustrates the magnitude of total pressure efficiency within a selected range of operating conditions.

Total Pressure Efficiency (%) = 0.000471947443 x Q x D_p / EPD

where	Q	is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)
	D_p	is the pressure differential across the fan filter unit, in Pascal
	EPD	is the total electirc power demand (EPD) of the fan filter unit include all the electric power necessary to operate the fan filter unit, in Watt

Figure 3 Total pressure efficiency of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

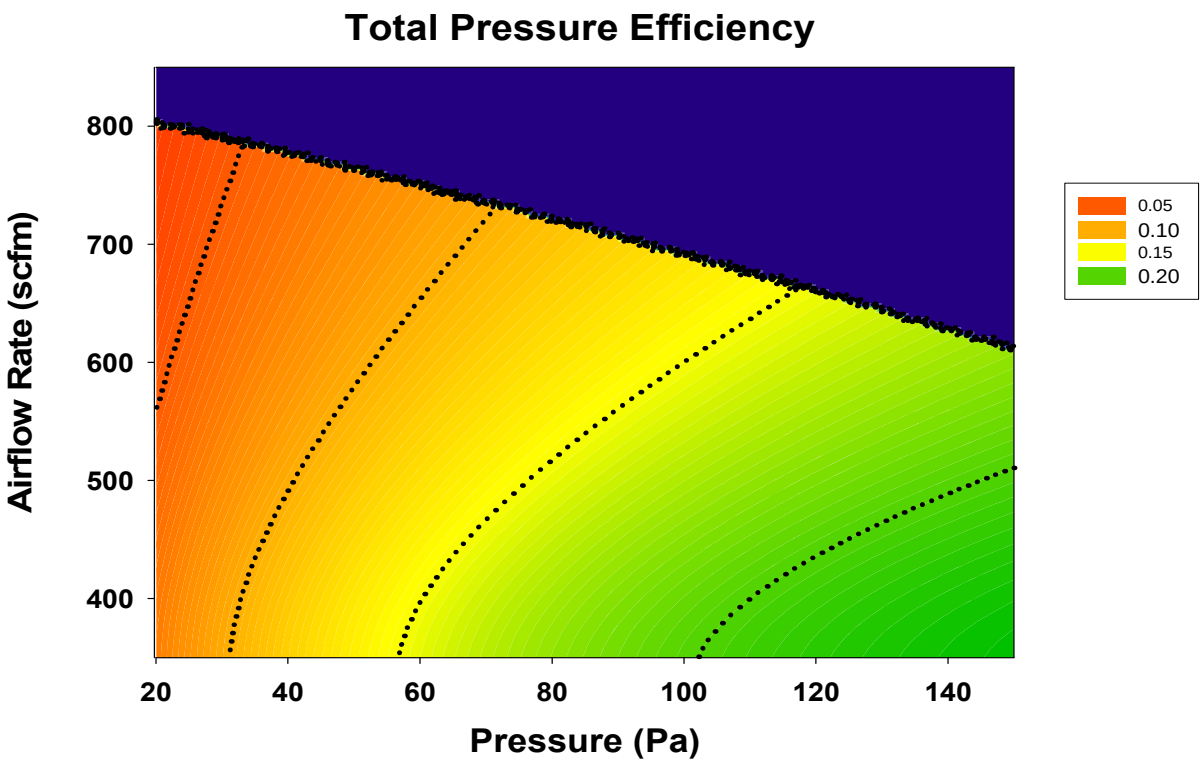

Note of unit conversion:

1 kPa = 4.015 inch water column
1 m/s = 196.9 feet per minute (fpm)
1 m^3/min = 35.3 ft^3/minute (cfm)

Report on

Laboratory Evaluation of Fan-filter Unit

LBNL Report, LBID-2588-6

Prepared by

Tengfang Xu and Duo Wang

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

August 11, 2006

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Acknowledgements

The project is funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Report on Laboratory Evaluation of Fan Filter Unit
Lawrence Berkeley National Laboratory, LBID-2588-6

Reporting Characteristics of the Unit Based on Manufacturer's Shipment

Manufacturer
Brand/Model
Serial Number
Unit Size
Fan Motor
Fan Wheel

Filter Efficiency

LBNL Tracking

Reporting Energy Performance Based on Laboratory Testing

This report is based upon the laboaratory evaluation of energy and airflow characteristics of the fan filter unit that was supplied to Lawrence Berkeley National Laboratory. All measured parameters, when applicable, are converted to their equivalents at the standard atmospheric condition (1 ATM, 20°C, sea level).

Total electric power demand, airflow rates, and pressure differential across the fan filter unit was recorded for a range of operable conditions of the unit. Normally, test conditions were selected with the airflow rates no less than 300 ft^3/min (8.5 m^3/min) under standard condition. The speed controller was set at its highest setting followed by lower setting. Figure 1 shows operable conditions that were tested with various speed control setting.

Figure 1. Operable conditions with various speed control setting

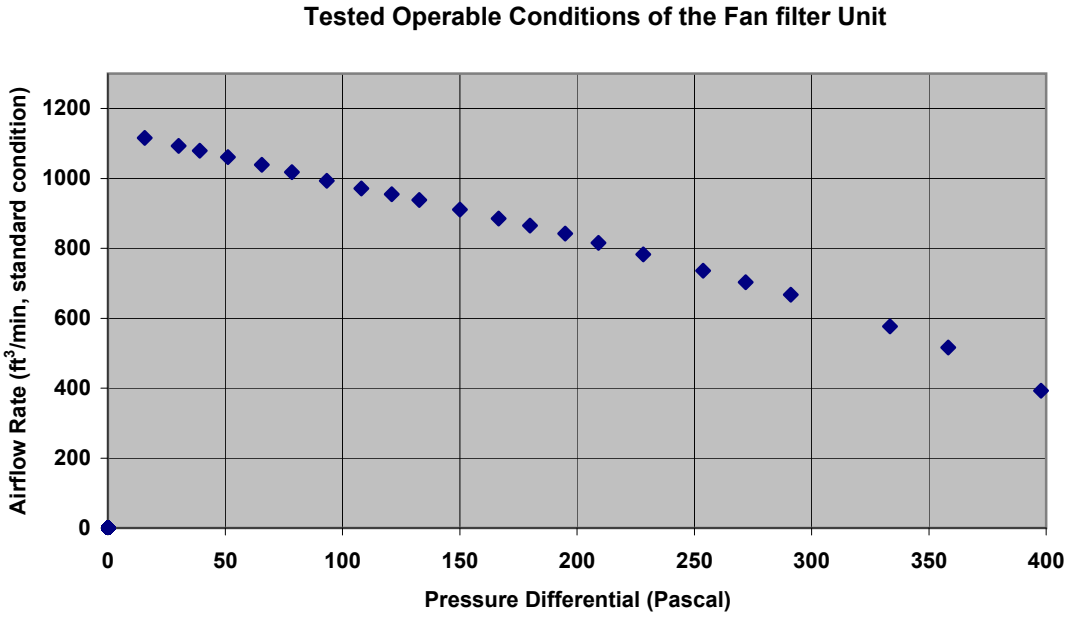

Total Electric Power Demand

Total electirc power demand (EPD) of the fan filter unit includes all the electric power necessary to operate the fan filter unit. It was measured concurrently with airflow rates and pressures under all testing conditions. As a result, the total electirc power demand can be calculated using either of the following equations.

Each of the equations is derived from laboratory testing results from the operating conditions as shown in Figure 1. The R-squares of the regression are included, which explain the statistical significance of the power demand predicted by the equation. A higher R-square number (with the possible maximum of 1) indicates higher degree of confidence in the power demand value derived from the laboratory testing.

Figures 2 illustrates the magnitude of total electric power demand within a selected range of operating conditions.

Electric Power Demand (Watt) = $C_{01} + C_1 \times D_p + C_{11} \times D_p^2$

$R_1^2 =$ **0.983**

Electric Power Demand (Watt) = $C_{02} + C_2 \times Q + C_{22} \times Q^2$

$R_2^2 =$ **0.993**

where D_p is the pressure differential across the fan filter unit, in Pascal

Q is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)

C_{01} =	405.94	C_{02} =	228.58
C_1 =	2.13E-01	C_2 =	0.3972
C_{11} =	-8.00E-04	C_{22} =	-0.0002

Figure 2 Total electric power demand of the fan filter unit under selected operable conditions: Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

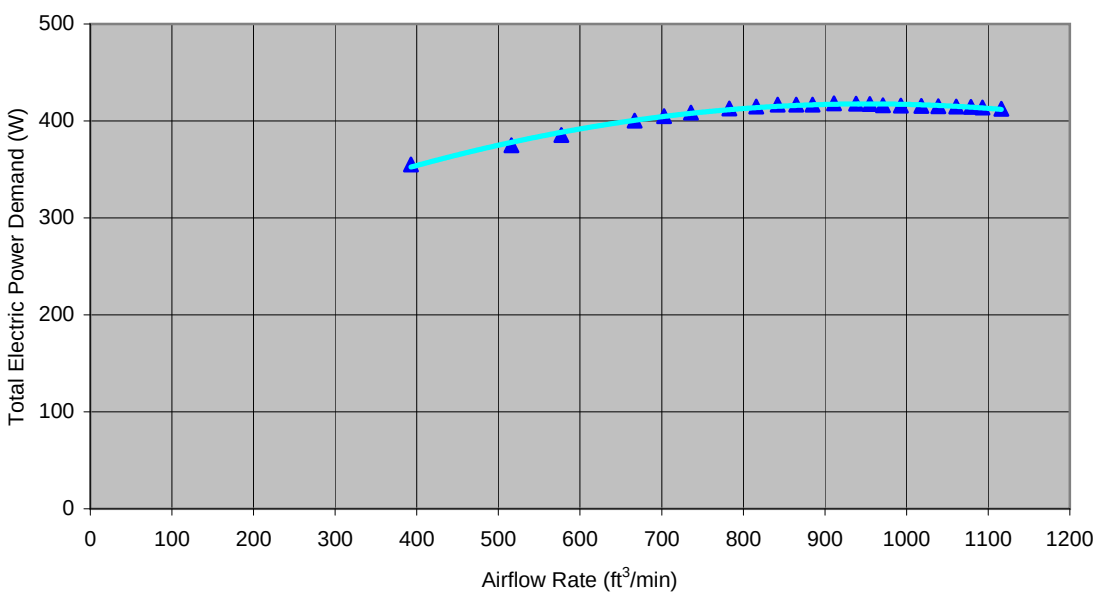

Airflow Rate (ft^3/min)

Total Pressure Efficiency

Total Pressure Efficiiency (TPE) is calculated by multiplying airflow rate and pressure differential across the unit then divided by total electric power demand as shown in the following equation.
Figures 3 illustrates the magnitude of total pressure efficiency within a selected range of operating conditions.

Total Pressure Efficiency (%) = 0.000471947443 x Q x D_p / EPD

where	Q	is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)
	D_p	is the pressure differential across the fan filter unit, in Pascal
	EPD	is the total electirc power demand (EPD) of the fan filter unit include all the electric power necessary to operate the fan filter unit, in Watt

Figure 3 Total pressure efficiency of the fan filter unit under selected operable conditions: Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

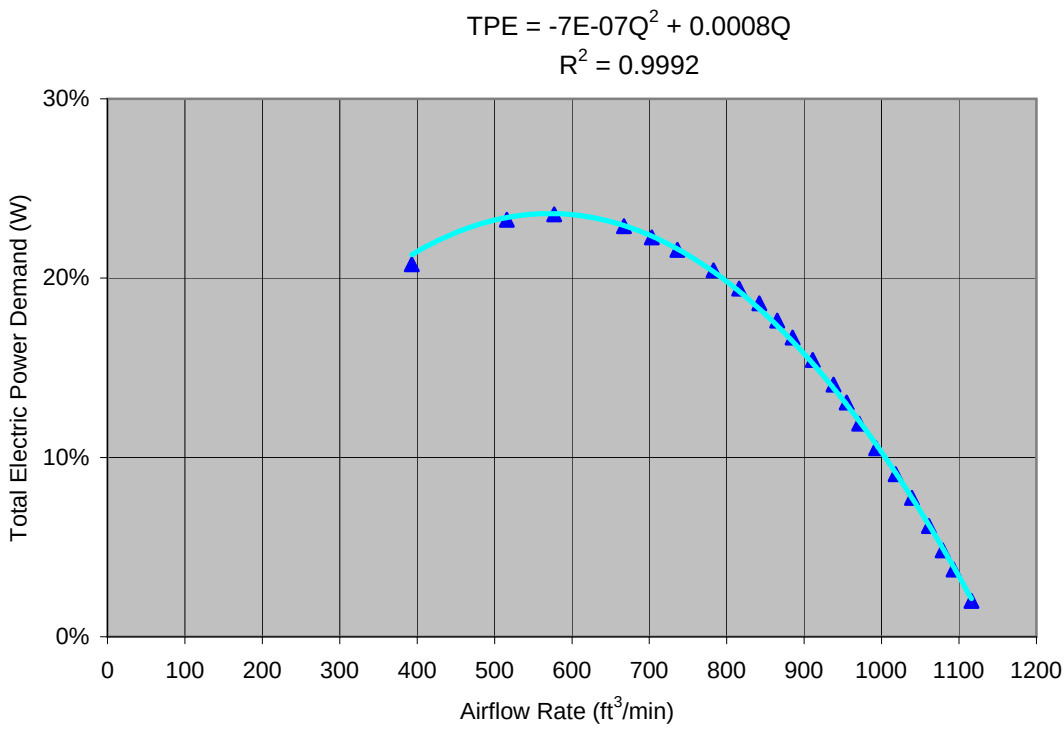

Note of unit conversion:

1 kPa = 4.015 inch water column
1 m/s = 196.9 feet per minute (fpm)
1 m^3/min = 35.3 ft^3/minute (cfm)

Report on

Laboratory Evaluation of Fan-filter Unit

LBNL Report, LBID-2588-7

Prepared by

Tengfang Xu and Duo Wang

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

August 17, 2006

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Acknowledgements

The project is funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Report on Laboratory Evaluation of Fan Filter Unit
Lawrence Berkeley National Laboratory, LBID-2588-7

Reporting Characteristics of the Unit Based on Manufacturer's Shipment

Manufacturer
Brand/Model
Serial Number
Unit Size
Fan Motor
Fan Wheel

Filter Efficiency

LBNL Tracking #

Reporting Energy Performance Based on Laboratory Testing

This report is based upon the laboaratory evaluation of energy and airflow characteristics of the fan filter unit that was supplied to Lawrence Berkeley National Laboratory. All measured parameters, when applicable, are converted to their equivalents at the standard atmospheric condition (1 ATM, 20°C, sea level).

Total electric power demand, airflow rates, and pressure differential across the fan filter unit was recorded for a range of operable conditions of the unit. Normally, test conditions were selected with the airflow rates no less than 300 ft^3/min (8.5 m^3/min) under standard condition. The speed controller was set at its highest setting followed by lower setting. Figure 1 shows operable conditions that were tested with various speed control setting.

Figure 1. Operable conditions with various speed control setting

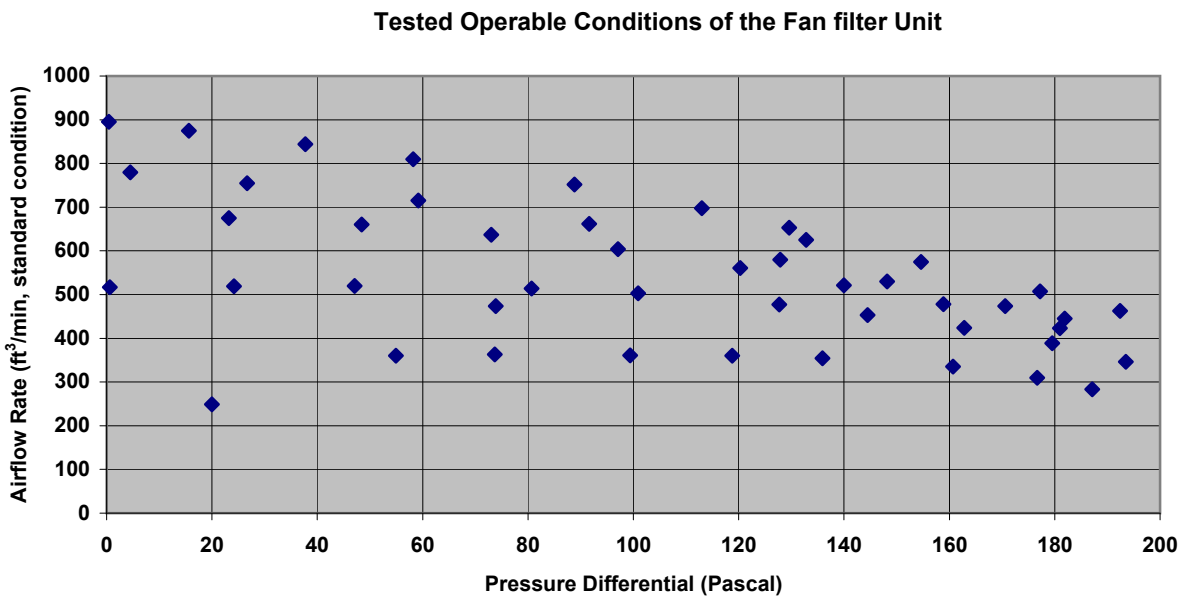

Total Electric Power Demand

Total electirc power demand (EPD) of the fan filter unit includes all the electric power necessary to operate the fan filter unit. It was measured concurrently with airflow rates and pressures under all testing conditions. As a result, the total electirc power demand can be calculated using the following equation.

The equation is derived from laboratory testing results from the operating conditions as shown in Figure 1. The R-square of the regression is included, which explains the statistical significance of the power demand predicted by the equation. A higher R-square number (with the possible maximum of 1) indicates higher degree of confidence in the power demand value derived from the laboratory testing.

Figures 2 illustrates the magnitude of total electric power demand within a selected range of operating conditions.

Electric Power Demand (Watt) = $C_0 + C_1 \times D_p + C_2 \times Q + C_{11} \times D_p^2 + C_{12} \times D_p \times Q + C_{22} \times Q^2$

R^2 = 0.994

where D_p is the pressure differential across the fan filter unit, in Pascal

Q is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)

C_0 =	3.22E+01
C_1 =	-1.01E-01
C_2 =	7.33E-01
C_{11} =	-9.96E-04
C_{12} =	1.28E-04
C_{22} =	-3.51E-04

Figure 2 Total electric power demand of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

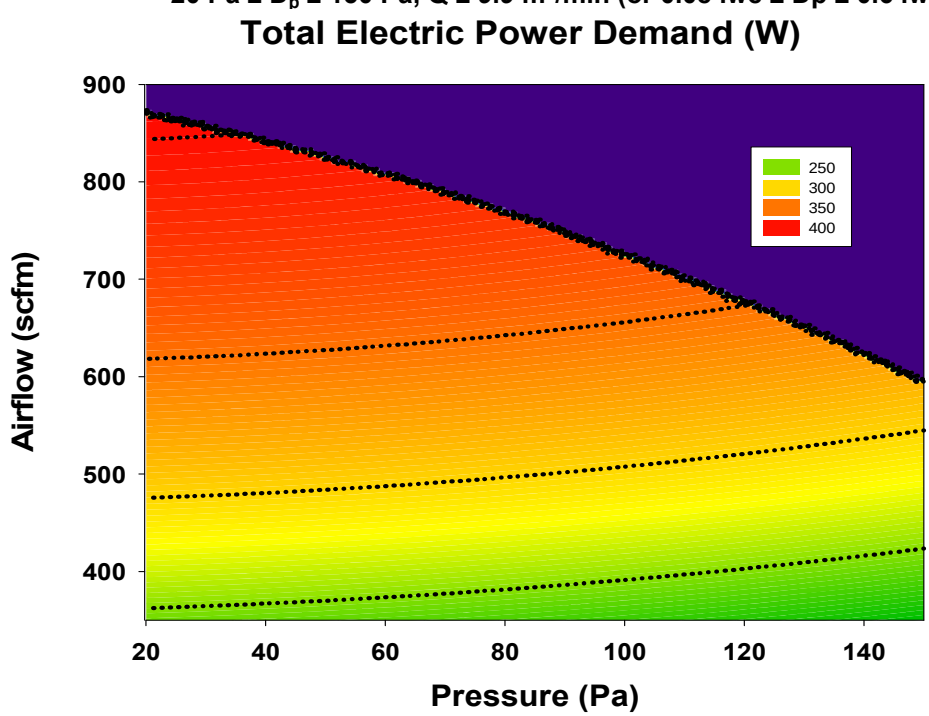

Total Pressure Efficiency

Total Pressure Efficiiency (TPE) is calculated by multiplying airflow rate and pressure differential across the unit then divided by total electric power demand as shown in the following equation.
Figures 3 illustrates the magnitude of total pressure efficiency within a selected range of operating conditions.

Total Pressure Efficiency (%) = 0.000471947443 x Q x D_p / EPD

where	Q	is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)
	D_p	is the pressure differential across the fan filter unit, in Pascal
	EPD	is the total electirc power demand (EPD) of the fan filter unit include all the electric power necessary to operate the fan filter unit, in Watt

Figure 3 Total pressure efficiency of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

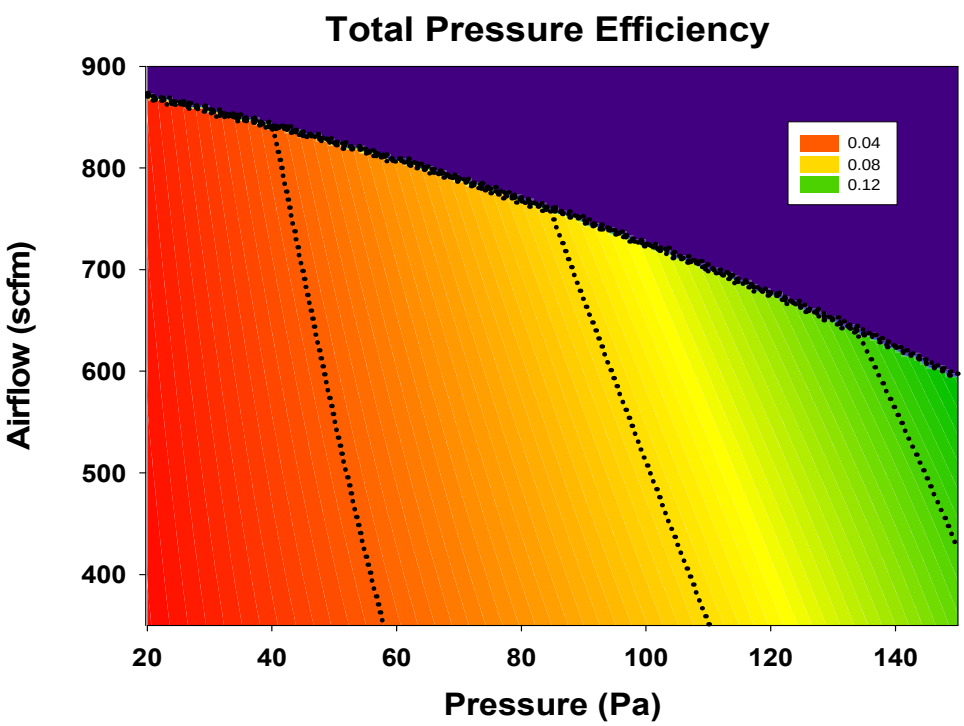

Note of unit conversion:

1 kPa = 4.015 inch water column
1 m/s = 196.9 feet per minute (fpm)
1 m^3/min = 35.3 ft^3/minute (cfm)

Report on

Laboratory Evaluation of Fan-filter Unit

LBNL Report, LBID-2588-8

Prepared by

Tengfang Xu and Duo Wang

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

August 17, 2006

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Acknowledgements

The project is funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Report on Laboratory Evaluation of Fan Filter Unit
Lawrence Berkeley National Laboratory, LBID-2588-8

Reporting Characteristics of the Unit Based on Manufacturer's Shipment

Manufacturer
Brand/Model
Serial Number
Unit Size
Fan Motor
Fan Wheel

Filter Efficiency

LBNL Tracking #

Reporting Energy Performance Based on Laboratory Testing

This report is based upon the laboaratory evaluation of energy and airflow characteristics of the fan filter unit that was supplied to Lawrence Berkeley National Laboratory. All measured parameters, when applicable, are converted to their equivalents at the standard atmospheric condition (1 ATM, 20°C, sea level).

Total electric power demand, airflow rates, and pressure differential across the fan filter unit was recorded for a range of operable conditions of the unit. Normally, test conditions were selected with the airflow rates no less than 300 ft^3/min (8.5 m^3/min) under standard condition. The speed controller was set at its highest setting followed by lower setting. Figure 1 shows operable conditions that were tested with various speed control setting.

Figure 1. Operable conditions with various speed control setting

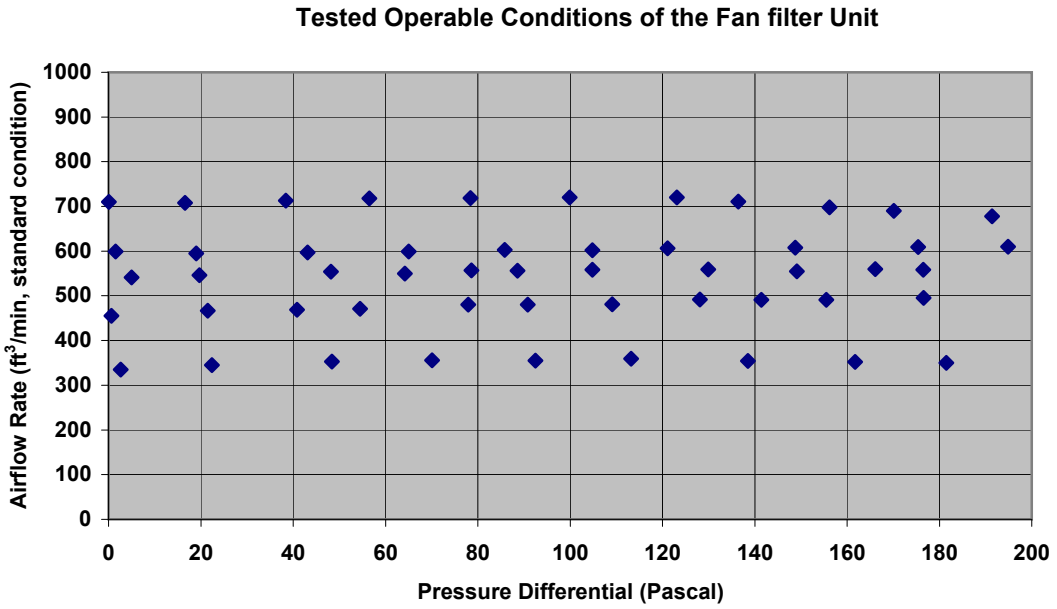

Total Electric Power Demand

Total electirc power demand (EPD) of the fan filter unit includes all the electric power necessary to operate the fan filter unit. It was measured concurrently with airflow rates and pressures under all testing conditions. As a result, the total electirc power demand can be calculated using the following equation.

The equation is derived from laboratory testing results from the operating conditions as shown in Figure 1. The R-square of the regression is included, which explains the statistical significance of the power demand predicted by the equation. A higher R-square number (with the possible maximum of 1) indicates higher degree of confidence in the power demand value derived from the laboratory testing.

Figures 2 illustrates the magnitude of total electric power demand within a selected range of operating conditions.

Electric Power Demand (Watt) = $C_0 + C_1 \times D_p + C_2 \times Q + C_{11} \times D_p^2 + C_{12} \times D_p \times Q + C_{22} \times Q^2$

R^2 = 0.999

where D_p is the pressure differential across the fan filter unit, in Pascal

Q is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)

C_0 =	1.72E+02
C_1 =	-6.45E-02
C_2 =	-7.81E-01
C_{11} =	-5.77E-04
C_{12} =	1.88E-03
C_{22} =	1.26E-03

Figure 2 Total electric power demand of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

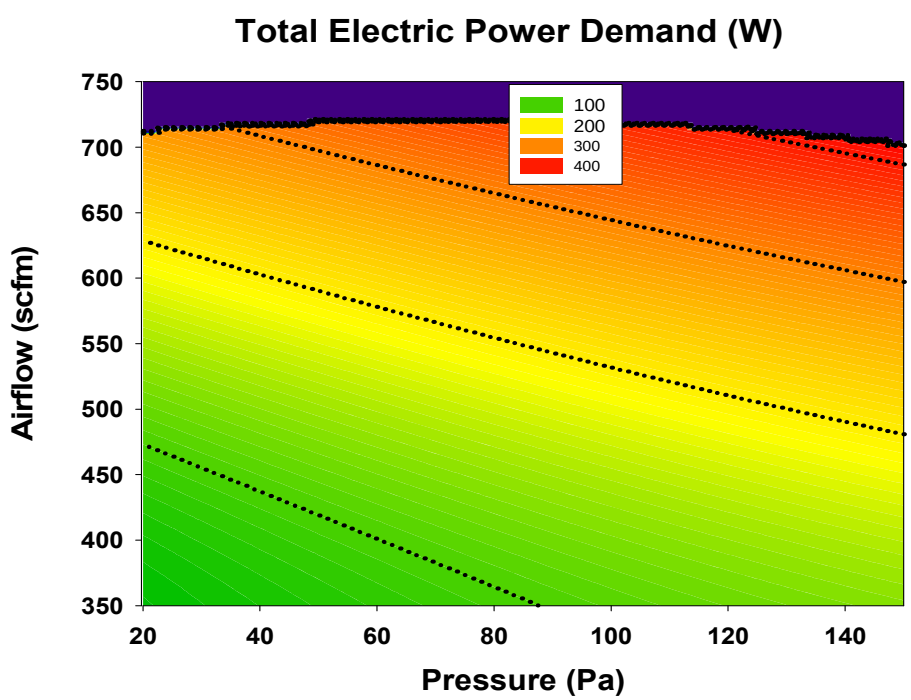

Total Pressure Efficiency

Total Pressure Efficiiency (TPE) is calculated by multiplying airflow rate and pressure differential across the unit then divided by total electric power demand as shown in the following equation.
Figures 3 illustrates the magnitude of total pressure efficiency within a selected range of operating conditions.

Total Pressure Efficiency (%) = 0.000471947443 x Q x D_p / EPD

where	Q	is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)
	D_p	is the pressure differential across the fan filter unit, in Pascal
	EPD	is the total electirc power demand (EPD) of the fan filter unit include all the electric power necessary to operate the fan filter unit, in Watt

Figure 3 Total pressure efficiency of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

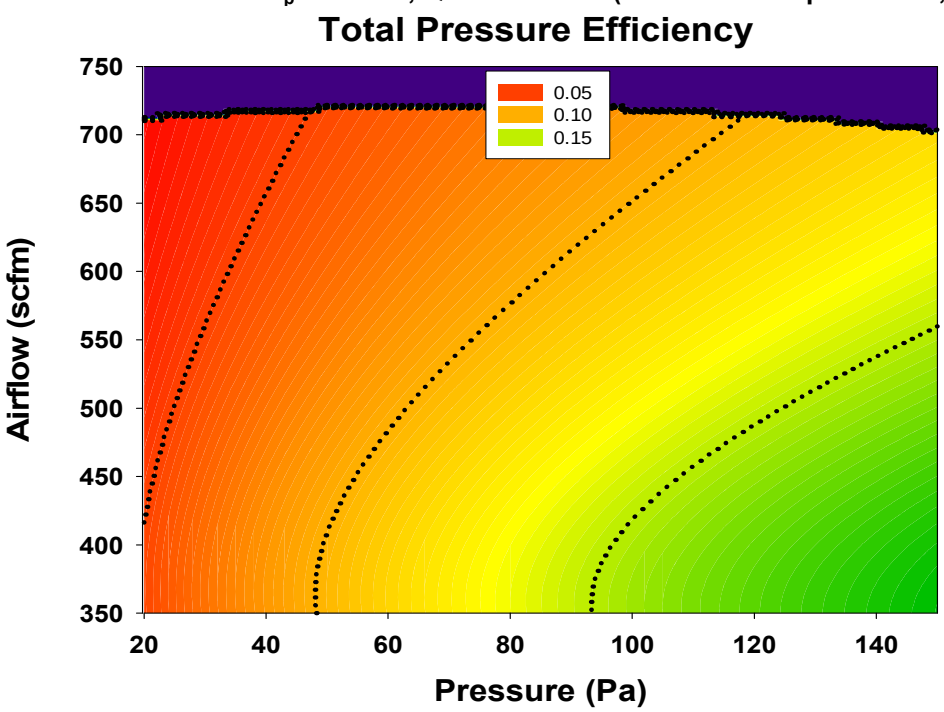

Note of unit conversion:

1 kPa = 4.015 inch water column
1 m/s = 196.9 feet per minute (fpm)
1 m^3/min = 35.3 ft^3/minute (cfm)

Report on

Laboratory Evaluation of Fan-filter Unit

LBNL Report, LBID-2588-9

Prepared by

Tengfang Xu and Duo Wang

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

August 18, 2006

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Acknowledgements

The project is funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Report on Laboratory Evaluation of Fan Filter Unit
Lawrence Berkeley National Laboratory, LBID-2588-9

Reporting Characteristics of the Unit Based on Manufacturer's Shipment

Manufacturer
Brand/Model
Serial Number
Unit Size
Fan Motor
Fan Wheel

Filter Efficiency

LBNL Tracking #

Reporting Energy Performance Based on Laboratory Testing

This report is based upon the laboaratory evaluation of energy and airflow characteristics of the fan filter unit that was supplied to Lawrence Berkeley National Laboratory. All measured parameters, when applicable, are converted to their equivalents at the standard atmospheric condition (1 ATM, 20°C, sea level).

Total electric power demand, airflow rates, and pressure differential across the fan filter unit was recorded for a range of operable conditions of the unit. Normally, test conditions were selected with the airflow rates no less than 300 ft^3/min (8.5 m^3/min) under standard condition. The speed controller was set at its highest setting followed by lower setting. Figure 1 shows operable conditions that were tested with various speed control setting.

Figure 1. Operable conditions with various speed control setting

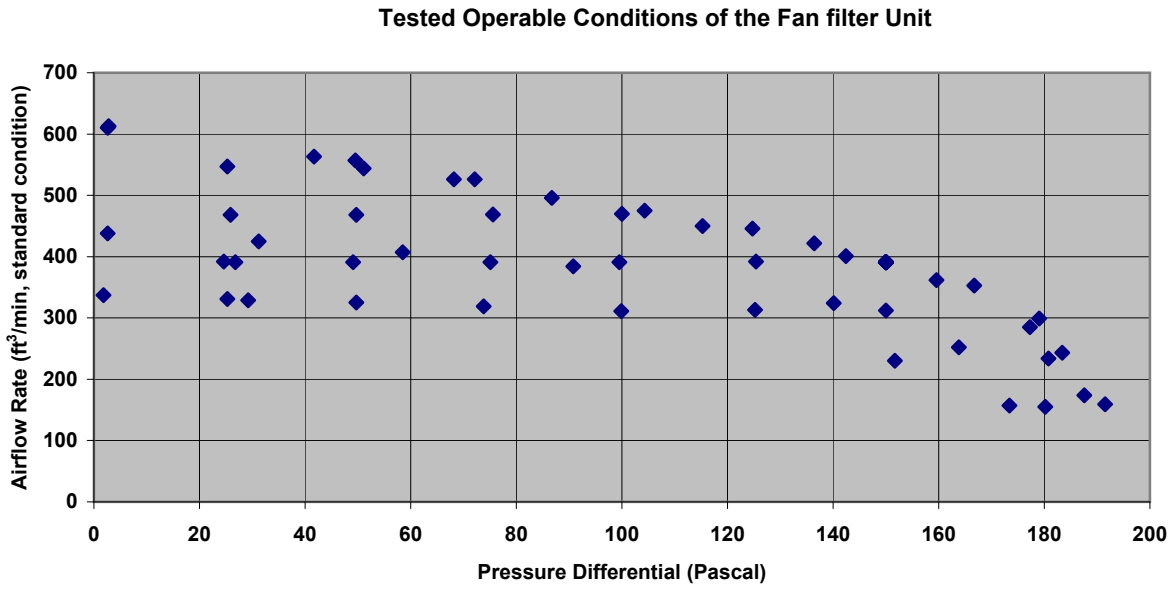

Total Electric Power Demand

Total electirc power demand (EPD) of the fan filter unit includes all the electric power necessary to operate the fan filter unit. It was measured concurrently with airflow rates and pressures under all testing conditions. As a result, the total electirc power demand can be calculated using the following equation.

The equation is derived from laboratory testing results from the operating conditions as shown in Figure 1. The R-square of the regression is included, which explains the statistical significance of the power demand predicted by the equation. A higher R-square number (with the possible maximum of 1) indicates higher degree of confidence in the power demand value derived from the laboratory testing.

Figures 2 illustrates the magnitude of total electric power demand within a selected range of operating conditions.

Electric Power Demand (Watt) = $C_0 + C_1 \times D_p + C_2 \times Q + C_{11} \times D_p^2 + C_{12} \times D_p \times Q + C_{22} \times Q^2$

R^2 = 0.997

where D_p is the pressure differential across the fan filter unit, in Pascal

Q is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)

C_0 =	3.32E+01
C_1 =	-2.01E-01
C_2 =	3.84E-01
C_{11} =	1.43E-03
C_{12} =	-4.37E-05
C_{22} =	-6.15E-06

Figure 2 Total electric power demand of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

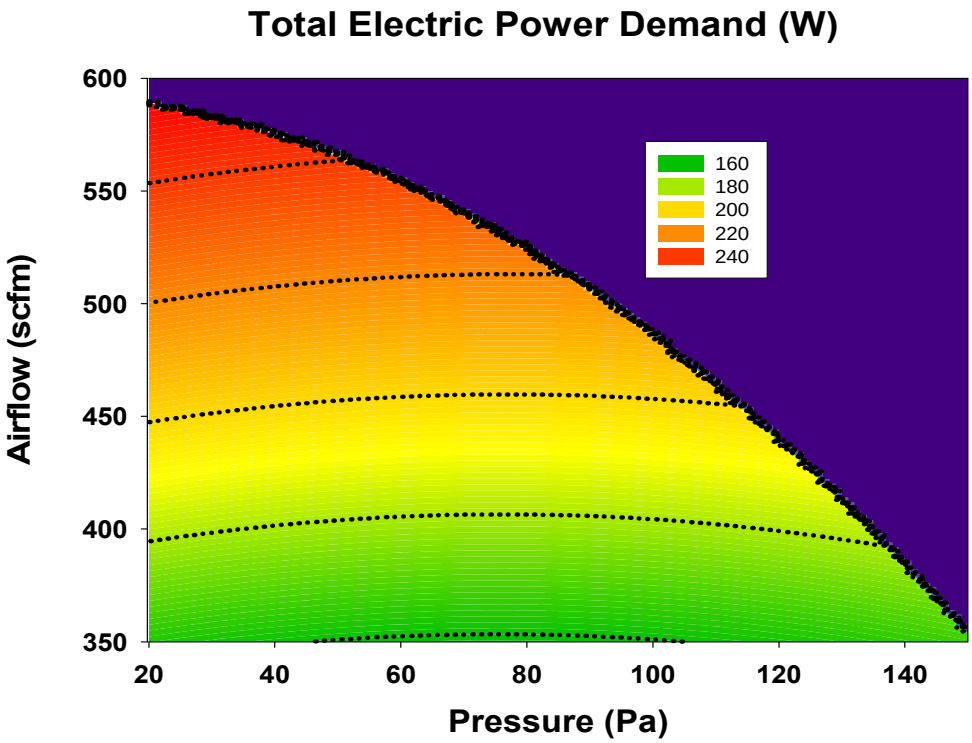

Pressure (Pa)

Total Pressure Efficiency

Total Pressure Efficiiency (TPE) is calculated by multiplying airflow rate and pressure differential across the unit then divided by total electric power demand as shown in the following equation.
Figures 3 illustrates the magnitude of total pressure efficiency within a selected range of operating conditions.

Total Pressure Efficiency (%) = 0.000471947443 x Q x D_p / EPD

where	Q	is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)
	D_p	is the pressure differential across the fan filter unit, in Pascal
	EPD	is the total electirc power demand (EPD) of the fan filter unit include all the electric power necessary to operate the fan filter unit, in Watt

Figure 3 Total pressure efficiency of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

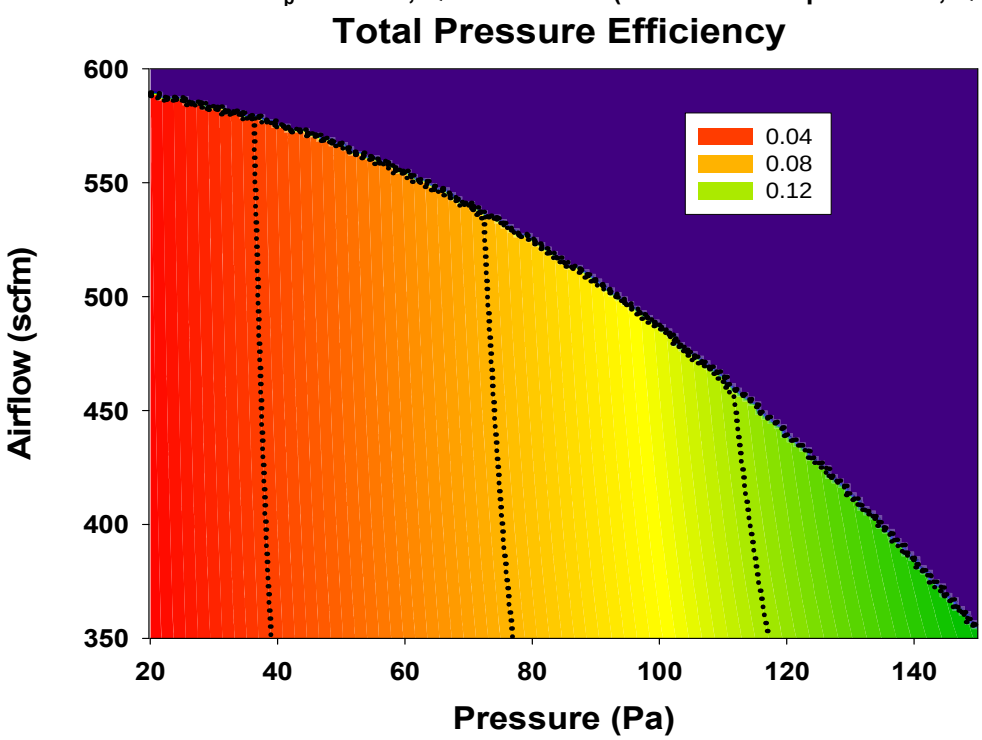

Note of unit conversion:

1 kPa = 4.015 inch water column
1 m/s = 196.9 feet per minute (fpm)
1 m^3/min = 35.3 ft^3/minute (cfm)

Report on

Laboratory Evaluation of Fan-filter Unit

LBNL Report, LBID-2588-10

Prepared by

Tengfang Xu and Duo Wang

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

August 18, 2006

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Acknowledgements

The project is funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Report on Laboratory Evaluation of Fan Filter Unit
Lawrence Berkeley National Laboratory, LBID-2588-10

Reporting Characteristics of the Unit Based on Manufacturer's Shipment

Manufacturer
Brand/Model
Serial Number
Unit Size
Fan Motor
Fan Wheel

Filter Efficiency

LBNL Tracking #

Reporting Energy Performance Based on Laboratory Testing

This report is based upon the laboaratory evaluation of energy and airflow characteristics of the fan filter unit that was supplied to Lawrence Berkeley National Laboratory. All measured parameters, when applicable, are converted to their equivalents at the standard atmospheric condition (1 ATM, 20°C, sea level).

Total electric power demand, airflow rates, and pressure differential across the fan filter unit was recorded for a range of operable conditions of the unit. Normally, test conditions were selected with the airflow rates no less than 300 ft^3/min (8.5 m^3/min) under standard condition. The speed controller was set at its highest setting followed by lower setting. Figure 1 shows operable conditions that were tested with various speed control setting.

Figure 1. Operable conditions with various speed control setting

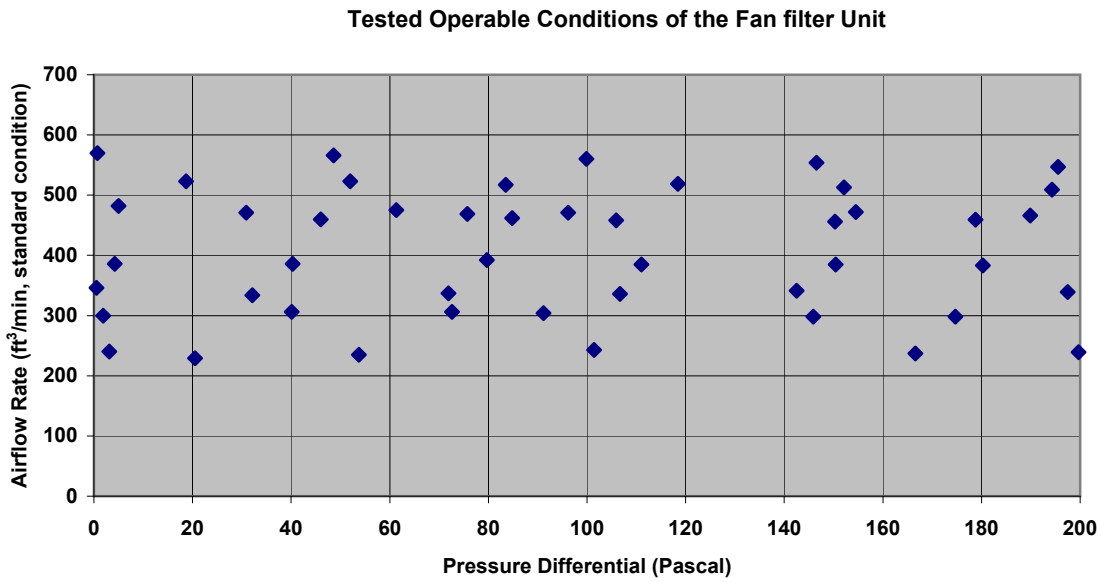

Total Electric Power Demand

Total electirc power demand (EPD) of the fan filter unit includes all the electric power necessary to operate the fan filter unit. It was measured concurrently with airflow rates and pressures under all testing conditions. As a result, the total electirc power demand can be calculated using the following equation.

The equation is derived from laboratory testing results from the operating conditions as shown in Figure 1. The R-square of the regression is included, which explains the statistical significance of the power demand predicted by the equation. A higher R-square number (with the possible maximum of 1) indicates higher degree of confidence in the power demand value derived from the laboratory testing.

Figures 2 illustrates the magnitude of total electric power demand within a selected range of operating conditions.

Electric Power Demand (Watt) = $C_0 + C_1 \times D_p + C_2 \times Q + C_{11} \times D_p^2 + C_{12} \times D_p \times Q + C_{22} \times Q^2$

R^2 = 0.999

where D_p is the pressure differential across the fan filter unit, in Pascal

Q is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)

C_0 =	8.59E+01
C_1 =	1.78E-01
C_2 =	-5.35E-01
C_{11} =	4.97E-05
C_{12} =	9.66E-04
C_{22} =	1.22E-03

Figure 2 Total electric power demand of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

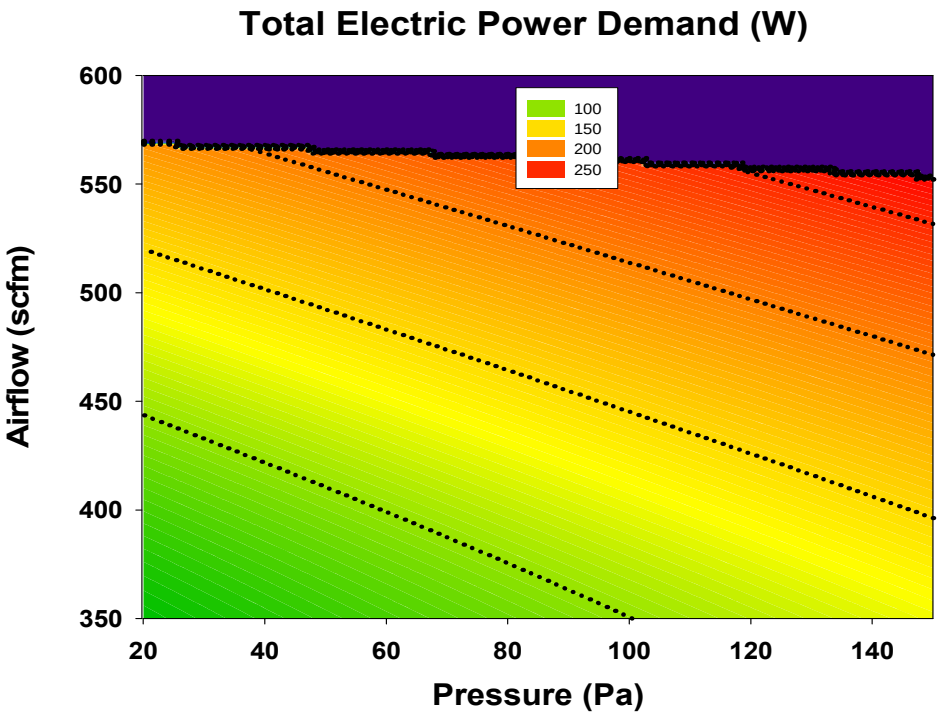

Total Pressure Efficiency

Total Pressure Efficiiency (TPE) is calculated by multiplying airflow rate and pressure differential across the unit then divided by total electric power demand as shown in the following equation.
Figures 3 illustrates the magnitude of total pressure efficiency within a selected range of operating conditions.

Total Pressure Efficiency (%) = 0.000471947443 x Q x D_p / EPD

where	Q	is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)
	D_p	is the pressure differential across the fan filter unit, in Pascal
	EPD	is the total electirc power demand (EPD) of the fan filter unit include all the electric power necessary to operate the fan filter unit, in Watt

Figure 3 Total pressure efficiency of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

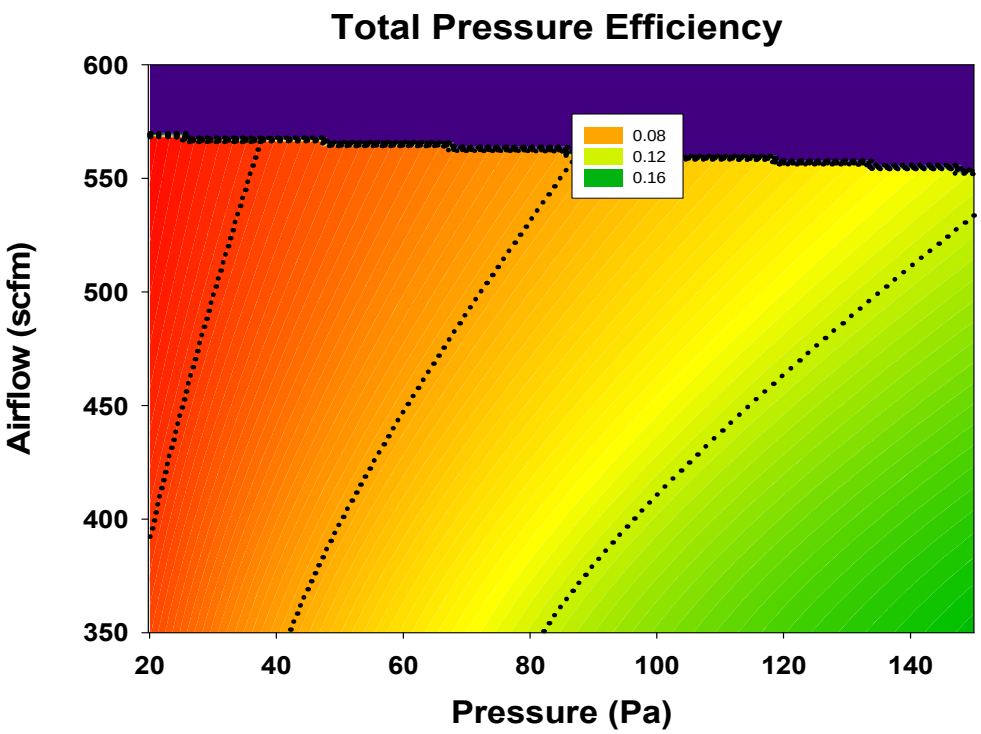

Note of unit conversion:

1 kPa = 4.015 inch water column
1 m/s = 196.9 feet per minute (fpm)
1 m^3/min = 35.3 ft^3/minute (cfm)

Report on

Laboratory Evaluation of Fan-filter Unit

LBNL Report, LBID-2588-11

Prepared by

Tengfang Xu and Duo Wang

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

September 15, 2006

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Acknowledgements

The project is funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Report on Laboratory Evaluation of Fan Filter Unit
Lawrence Berkeley National Laboratory, LBID-2588-11

Reporting Characteristics of the Unit Based on Manufacturer's Shipment

Manufacturer
Brand/Model
Serial Number
Unit Size
Fan Motor
Fan Wheel

Filter Efficiency

LBNL Tracking #

Reporting Energy Performance Based on Laboratory Testing

This report is based upon the laboratory evaluation of energy and airflow characteristics of the fan filter unit that was supplied to Lawrence Berkeley National Laboratory. All measured parameters, when applicable, are converted to their equivalents at the standard atmospheric condition (1 ATM, 20°C, sea level).

Total electric power demand, airflow rates, and pressure differential across the fan filter unit was recorded for a range of operable conditions of the unit. Normally, test conditions were selected with the airflow rates no less than 300 ft^3/min (8.5 m^3/min) under standard condition. The speed controller was set at its highest setting followed by lower setting. Figure 1 shows operable conditions that were tested with various speed control setting.

Figure 1. Operable conditions with various speed control setting

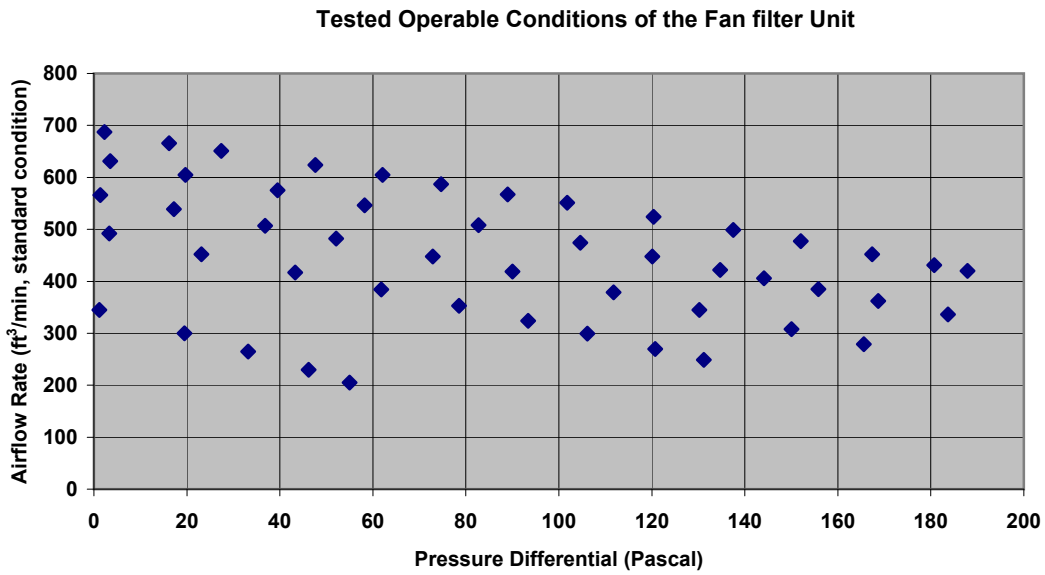

Total Electric Power Demand

Total electric power demand (EPD) of the fan filter unit includes all the electric power necessary to operate the fan filter unit. It was measured concurrently with airflow rates and pressures under all testing conditions. As a result, the total electric power demand can be calculated using the following equation.

The equation is derived from laboratory testing results from the operating conditions as shown in Figure 1. The R-square of the regression is included, which explains the statistical significance of the power demand predicted by the equation. A higher R-square number (with the possible maximum of 1) indicates higher degree of confidence in the power demand value derived from the laboratory testing.

Figures 2 illustrates the magnitude of total electric power demand within a selected range of operating conditions.

Electric Power Demand (Watt) = $C_0 + C_1 \times D_p + C_2 \times Q + C_{11} \times D_p^2 + C_{12} \times D_p \times Q + C_{22} \times Q^2$

R^2 = 0.999

where D_p is the pressure differential across the fan filter unit, in Pascal

Q is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)

C_0 =	1.08E+02
C_1 =	-1.05E-01
C_2 =	-4.38E-01
C_{11} =	6.95E-04
C_{12} =	1.92E-03
C_{22} =	9.44E-04

Figure 2 Total electric power demand of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

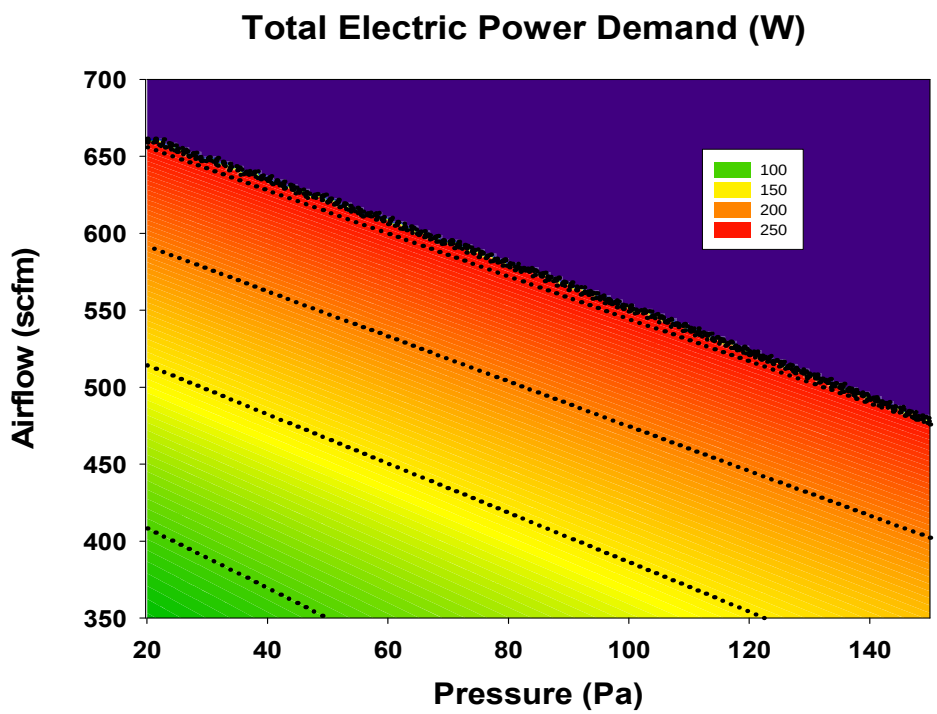

Total Pressure Efficiency

Total Pressure Efficiiency (TPE) is calculated by multiplying airflow rate and pressure differential across the unit then divided by total electric power demand as shown in the following equation.
Figures 3 illustrates the magnitude of total pressure efficiency within a selected range of operating conditions.

Total Pressure Efficiency (%) = 0.000471947443 x Q x D_p / EPD

where	Q	is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)
	D_p	is the pressure differential across the fan filter unit, in Pascal
	EPD	is the total electirc power demand (EPD) of the fan filter unit include all the electric power necessary to operate the fan filter unit, in Watt

Figure 3 Total pressure efficiency of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

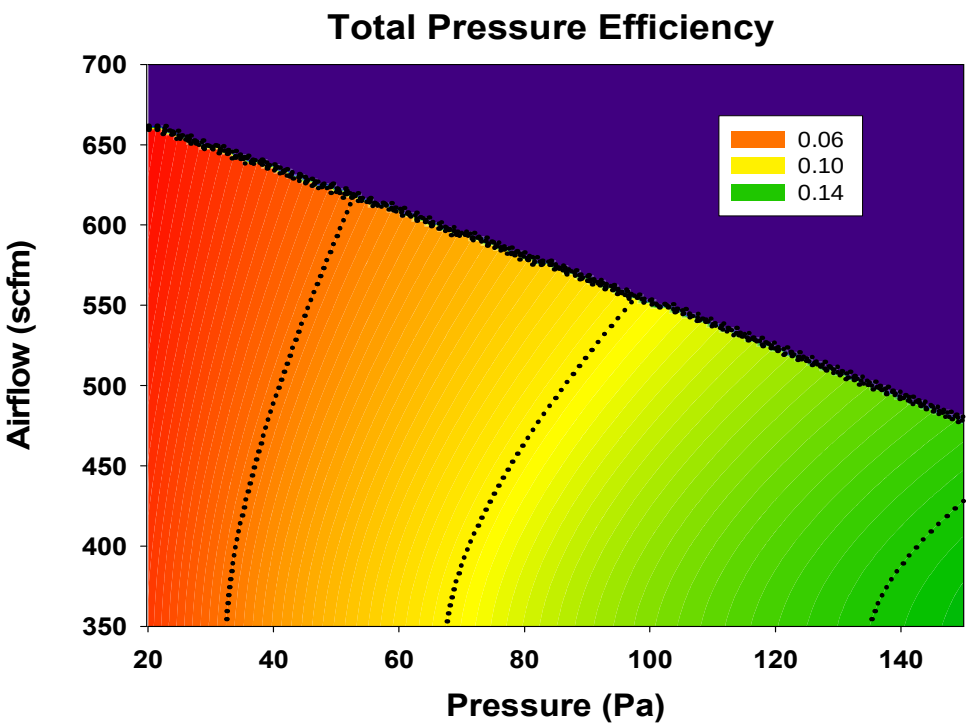

Note of unit conversion:

1 kPa = 4.015 inch water column
1 m/s = 196.9 feet per minute (fpm)
1 m^3/min = 35.3 ft^3/minute (cfm)

Report on

Laboratory Evaluation of Fan-filter Unit

LBNL Report, LBID-2588-12

Prepared by

Tengfang Xu and Duo Wang

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

September 15, 2006

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Acknowledgements

The project is funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Report on Laboratory Evaluation of Fan Filter Unit
Lawrence Berkeley National Laboratory, LBID-2588-12

Reporting Characteristics of the Unit Based on Manufacturer's Shipment

Manufacturer
Brand/Model
Serial Number
Unit Size
Fan Motor
Fan Wheel

Filter Efficiency

LBNL Tracking #

Reporting Energy Performance Based on Laboratory Testing

This report is based upon the laboratory evaluation of energy and airflow characteristics of the fan filter unit that was supplied to Lawrence Berkeley National Laboratory. All measured parameters, when applicable, are converted to their equivalents at the standard atmospheric condition (1 ATM, 20°C, sea level).

Total electric power demand, airflow rates, and pressure differential across the fan filter unit was recorded for a range of operable conditions of the unit. Normally, test conditions were selected with the airflow rates no less than 300 ft^3/min (8.5 m^3/min) under standard condition. The speed controller was set at its highest setting followed by lower setting. Figure 1 shows operable conditions that were tested with various speed control setting.

Figure 1. Operable conditions with various speed control setting

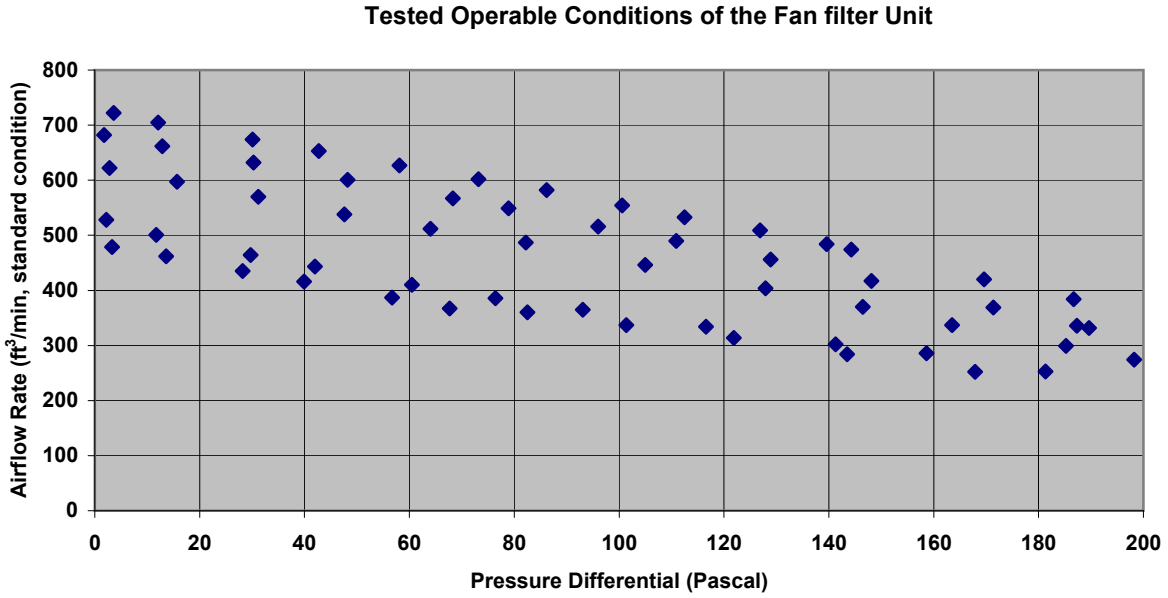

Total Electric Power Demand

Total electric power demand (EPD) of the fan filter unit includes all the electric power necessary to operate the fan filter unit. It was measured concurrently with airflow rates and pressures under all testing conditions. As a result, the total electric power demand can be calculated using the following equation.

The equation is derived from laboratory testing results from the operating conditions as shown in Figure 1. The R-square of the regression is included, which explains the statistical significance of the power demand predicted by the equation. A higher R-square number (with the possible maximum of 1) indicates higher degree of confidence in the power demand value derived from the laboratory testing.

Figures 2 illustrates the magnitude of total electric power demand within a selected range of operating conditions.

Electric Power Demand (Watt) = $C_0 + C_1 \times D_p + C_2 \times Q + C_{11} \times D_p^2 + C_{12} \times D_p \times Q + C_{22} \times Q^2$

R^2 = 0.999

where D_p is the pressure differential across the fan filter unit, in Pascal

Q is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)

C_0 =	5.04E+02
C_1 =	-2.25E+00
C_2 =	-1.14E+00
C_{11} =	3.72E-03
C_{12} =	4.17E-03
C_{22} =	1.06E-03

Figure 2 Total electric power demand of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

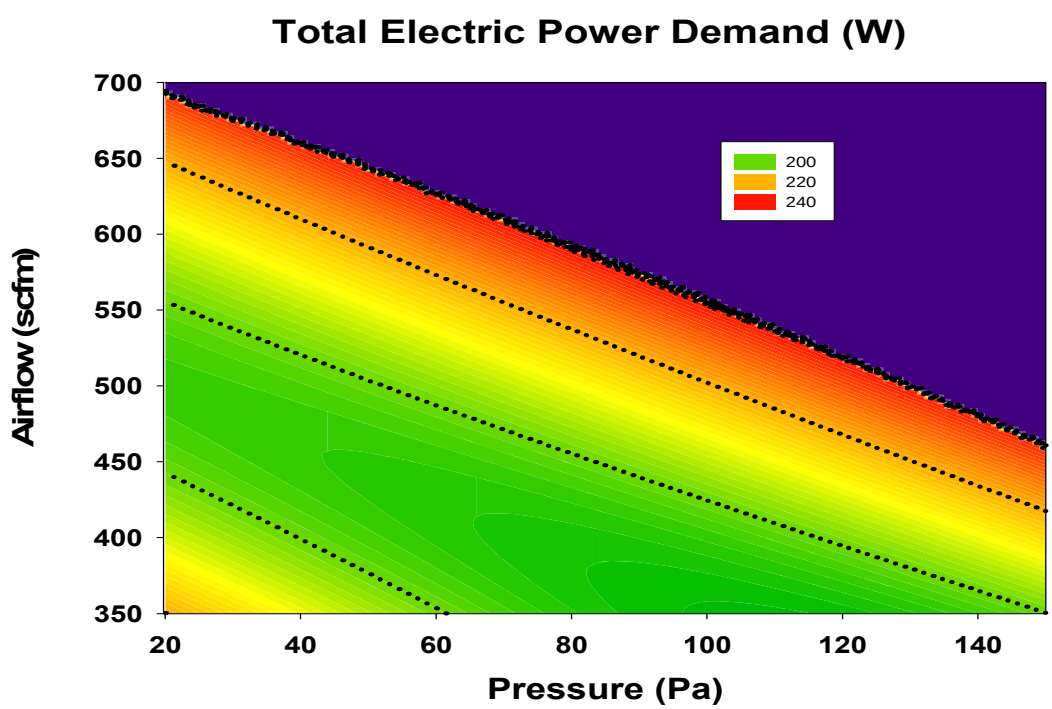

Total Pressure Efficiency

Total Pressure Efficiiency (TPE) is calculated by multiplying airflow rate and pressure differential across the unit then divided by total electric power demand as shown in the following equation.
Figures 3 illustrates the magnitude of total pressure efficiency within a selected range of operating conditions.

Total Pressure Efficiency (%) = 0.000471947443 x Q x D_p / EPD

where	Q	is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)
	D_p	is the pressure differential across the fan filter unit, in Pascal
	EPD	is the total electirc power demand (EPD) of the fan filter unit include all the electric power necessary to operate the fan filter unit, in Watt

Figure 3 Total pressure efficiency of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

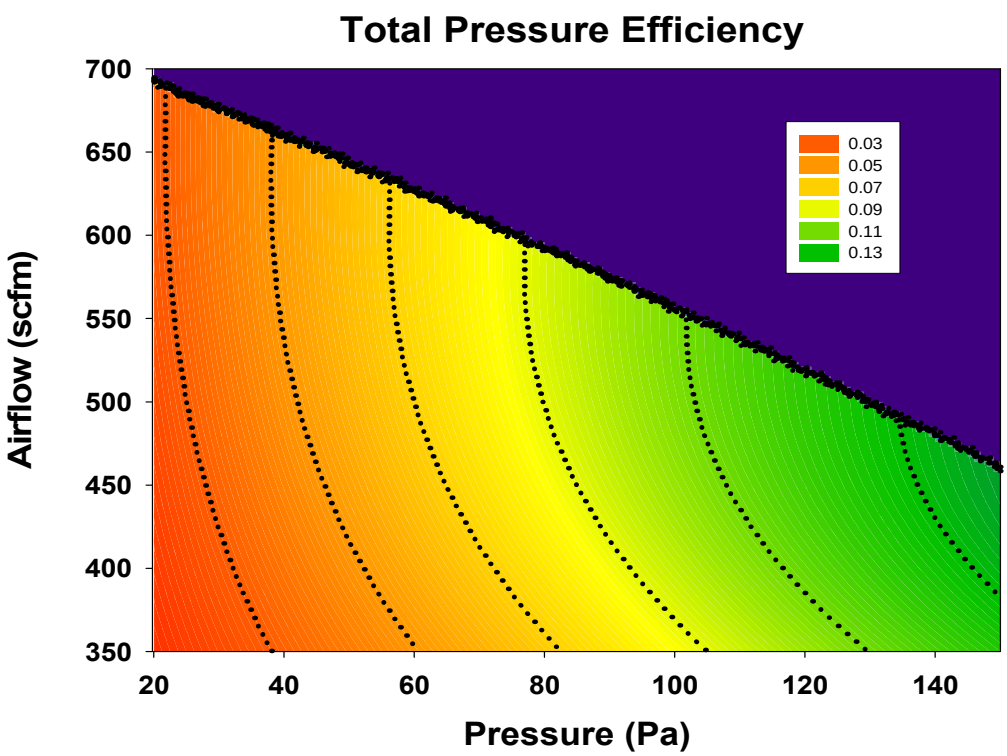

Note of unit conversion:

1 kPa = 4.015 inch water column
1 m/s = 196.9 feet per minute (fpm)
1 m^3/min = 35.3 ft^3/minute (cfm)

Report on

Laboratory Evaluation of Fan-filter Unit

LBNL Report, LBID-2588-13

Prepared by

Tengfang Xu and Duo Wang

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

September 15, 2006

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Acknowledgements

The project is funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Report on Laboratory Evaluation of Fan Filter Unit
Lawrence Berkeley National Laboratory, LBID-2588-13

Reporting Characteristics of the Unit Based on Manufacturer's Shipment

Manufacturer
Brand/Model
Serial Number
Unit Size
Fan Motor
Fan Wheel

Filter Efficiency

LBNL Tracking #

Reporting Energy Performance Based on Laboratory Testing

This report is based upon the laboratory evaluation of energy and airflow characteristics of the fan filter unit that was supplied to Lawrence Berkeley National Laboratory. All measured parameters, when applicable, are converted to their equivalents at the standard atmospheric condition (1 ATM, 20°C, sea level).

Total electric power demand, airflow rates, and pressure differential across the fan filter unit was recorded for a range of operable conditions of the unit. Normally, test conditions were selected with the airflow rates no less than 300 ft^3/min (8.5 m^3/min) under standard condition. The speed controller was set at its highest setting followed by lower setting. Figure 1 shows operable conditions that were tested with various speed control setting.

Figure 1. Operable conditions with various speed control setting

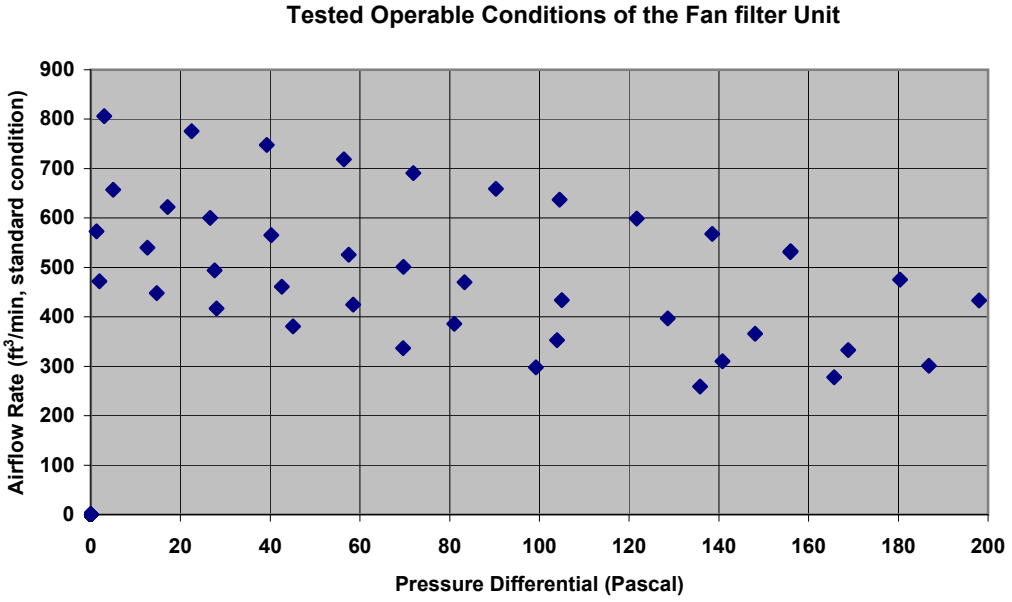

Total Electric Power Demand

Total electric power demand (EPD) of the fan filter unit includes all the electric power necessary to operate the fan filter unit. It was measured concurrently with airflow rates and pressures under all testing conditions. As a result, the total electric power demand can be calculated using the following equation.

The equation is derived from laboratory testing results from the operating conditions as shown in Figure 1. The R-square of the regression is included, which explains the statistical significance of the power demand predicted by the equation. A higher R-square number (with the possible maximum of 1) indicates higher degree of confidence in the power demand value derived from the laboratory testing.

Figures 2 illustrates the magnitude of total electric power demand within a selected range of operating conditions.

Electric Power Demand (Watt) = $C_0 + C_1 \times D_p + C_2 \times Q + C_{11} \times D_p^2 + C_{12} \times D_p \times Q + C_{22} \times Q^2$

R^2 = 0.999

where D_p is the pressure differential across the fan filter unit, in Pascal

Q is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)

C_0 =	2.27E+02
C_1 =	-5.61E-01
C_2 =	-1.84E-01
C_{11} =	1.21E-03
C_{12} =	1.39E-03
C_{22} =	2.51E-04

Figure 2 Total electric power demand of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

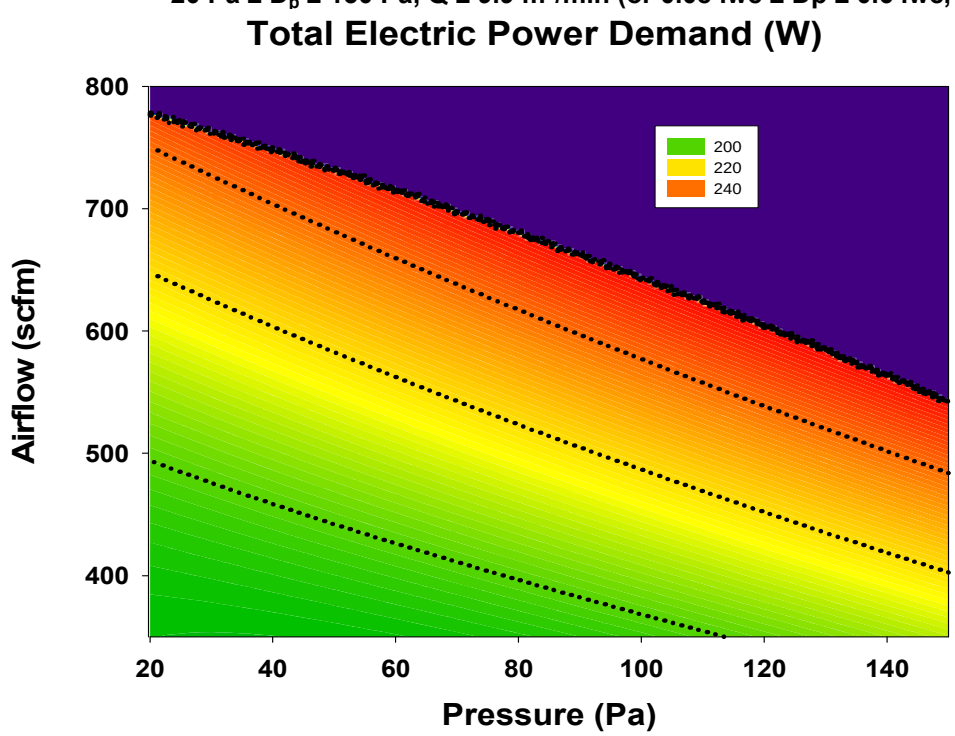

Total Pressure Efficiency

Total Pressure Efficiiency (TPE) is calculated by multiplying airflow rate and pressure differential across the unit then divided by total electric power demand as shown in the following equation.
Figures 3 illustrates the magnitude of total pressure efficiency within a selected range of operating conditions.

Total Pressure Efficiency (%) = 0.000471947443 x Q x D_p / EPD

where	Q	is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)
	D_p	is the pressure differential across the fan filter unit, in Pascal
	EPD	is the total electirc power demand (EPD) of the fan filter unit include all the electric power necessary to operate the fan filter unit, in Watt

Figure 3 Total pressure efficiency of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

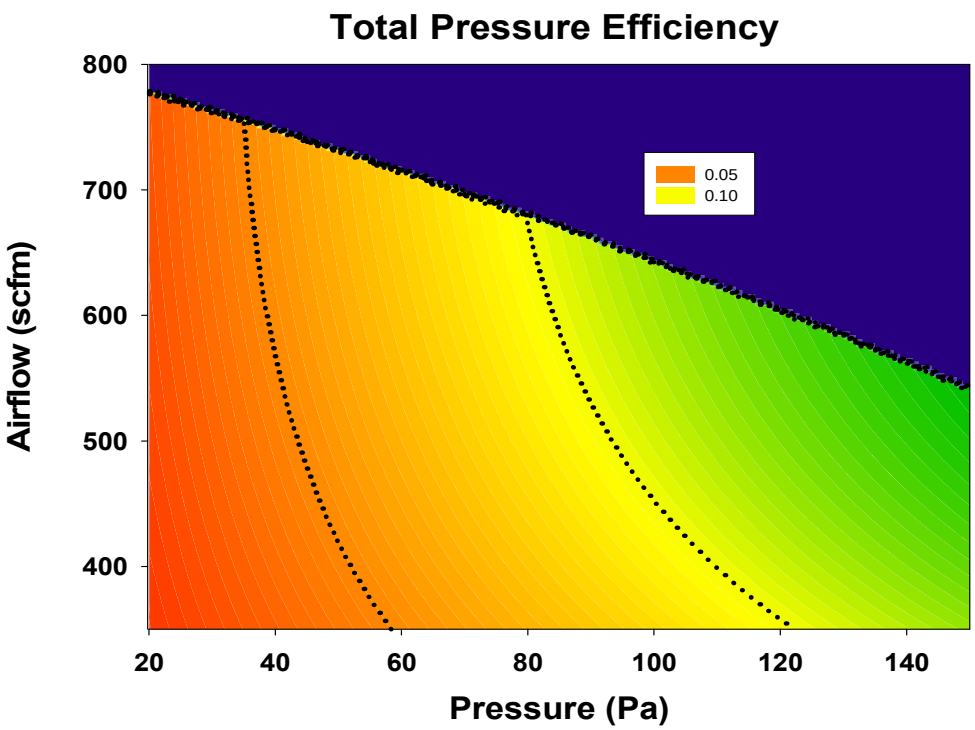

Note of unit conversion:

1 kPa = 4.015 inch water column
1 m/s = 196.9 feet per minute (fpm)
1 m^3/min = 35.3 ft^3/minute (cfm)

Report on

Laboratory Evaluation of Fan-filter Unit

LBNL Report, LBID-2588-14

Prepared by

Tengfang Xu and Duo Wang

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

September 15, 2006

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Acknowledgements

The project is funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Report on Laboratory Evaluation of Fan Filter Unit
Lawrence Berkeley National Laboratory, LBID-2588-14

Reporting Characteristics of the Unit Based on Manufacturer's Shipment

Manufacturer
Brand/Model
Serial Number
Unit Size
Fan Motor
Fan Wheel

Filter Efficiency

LBNL Tracking #

Reporting Energy Performance Based on Laboratory Testing

This report is based upon the laboratory evaluation of energy and airflow characteristics of the fan filter unit that was supplied to Lawrence Berkeley National Laboratory. All measured parameters, when applicable, are converted to their equivalents at the standard atmospheric condition (1 ATM, 20°C, sea level).

Total electric power demand, airflow rates, and pressure differential across the fan filter unit was recorded for a range of operable conditions of the unit. Normally, test conditions were selected with the airflow rates no less than 300 ft^3/min (8.5 m^3/min) under standard condition. The speed controller was set at its highest setting followed by lower setting. Figure 1 shows operable conditions that were tested with various speed control setting.

Figure 1. Operable conditions with various speed control setting

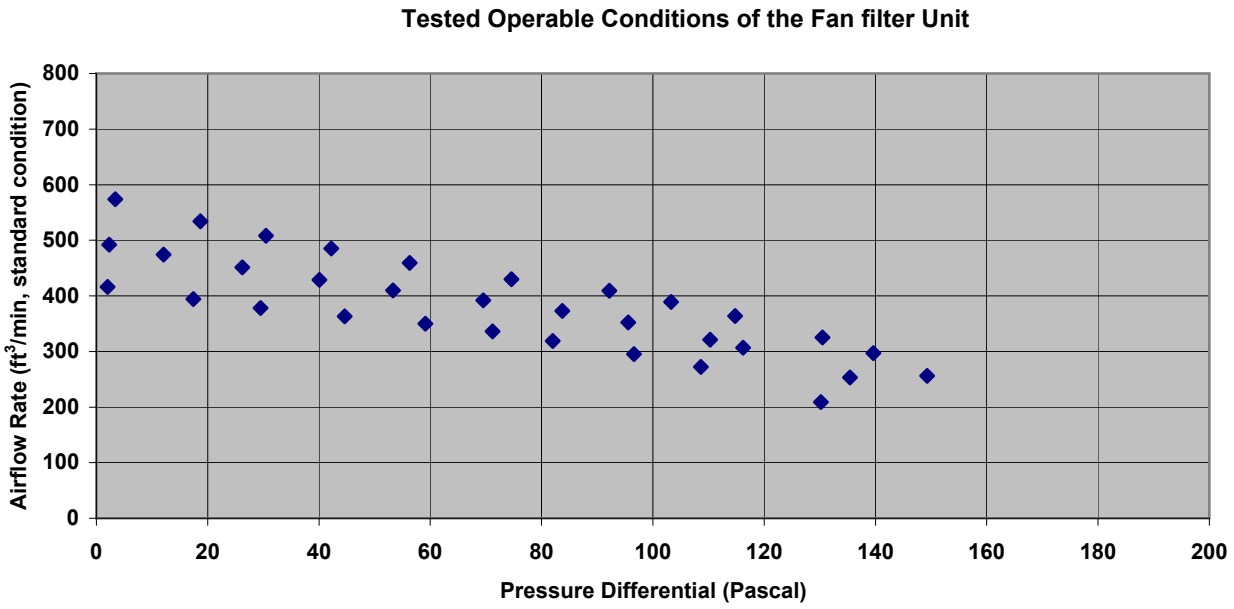

Total Electric Power Demand

Total electric power demand (EPD) of the fan filter unit includes all the electric power necessary to operate the fan filter unit. It was measured concurrently with airflow rates and pressures under all testing conditions. As a result, the total electric power demand can be calculated using the following equation.

The equation is derived from laboratory testing results from the operating conditions as shown in Figure 1. The R-square of the regression is included, which explains the statistical significance of the power demand predicted by the equation. A higher R-square number (with the possible maximum of 1) indicates higher degree of confidence in the power demand value derived from the laboratory testing.

Figures 2 illustrates the magnitude of total electric power demand within a selected range of operating conditions.

Electric Power Demand (Watt) = $C_0 + C_1 \times D_p + C_2 \times Q + C_{11} \times D_p^2 + C_{12} \times D_p \times Q + C_{22} \times Q^2$

R^2 = 0.996

where D_p is the pressure differential across the fan filter unit, in Pascal

Q is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)

C_0 =	4.58E+02
C_1 =	-4.97E+00
C_2 =	-1.74E+00
C_{11} =	1.52E-02
C_{12} =	1.09E-02
C_{22} =	2.63E-03

Figure 2 Total electric power demand of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

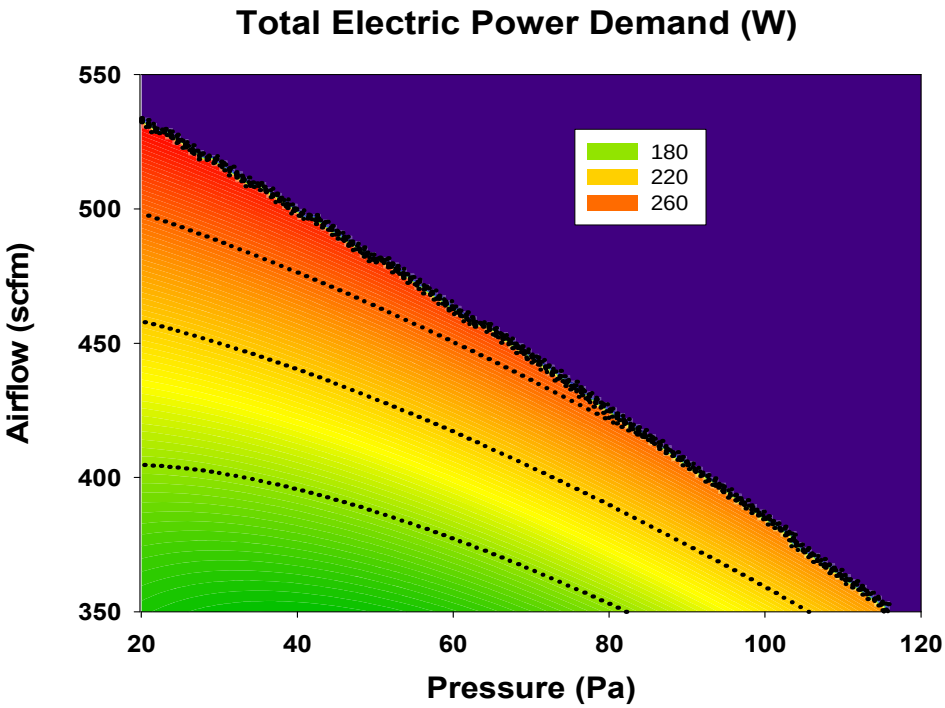

Total Pressure Efficiency

Total Pressure Efficiiency (TPE) is calculated by multiplying airflow rate and pressure differential across the unit then divided by total electric power demand as shown in the following equation.
Figures 3 illustrates the magnitude of total pressure efficiency within a selected range of operating conditions.

Total Pressure Efficiency (%) = 0.000471947443 x Q x D_p / EPD

where	Q	is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)
	D_p	is the pressure differential across the fan filter unit, in Pascal
	EPD	is the total electirc power demand (EPD) of the fan filter unit include all the electric power necessary to operate the fan filter unit, in Watt

Figure 3 Total pressure efficiency of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

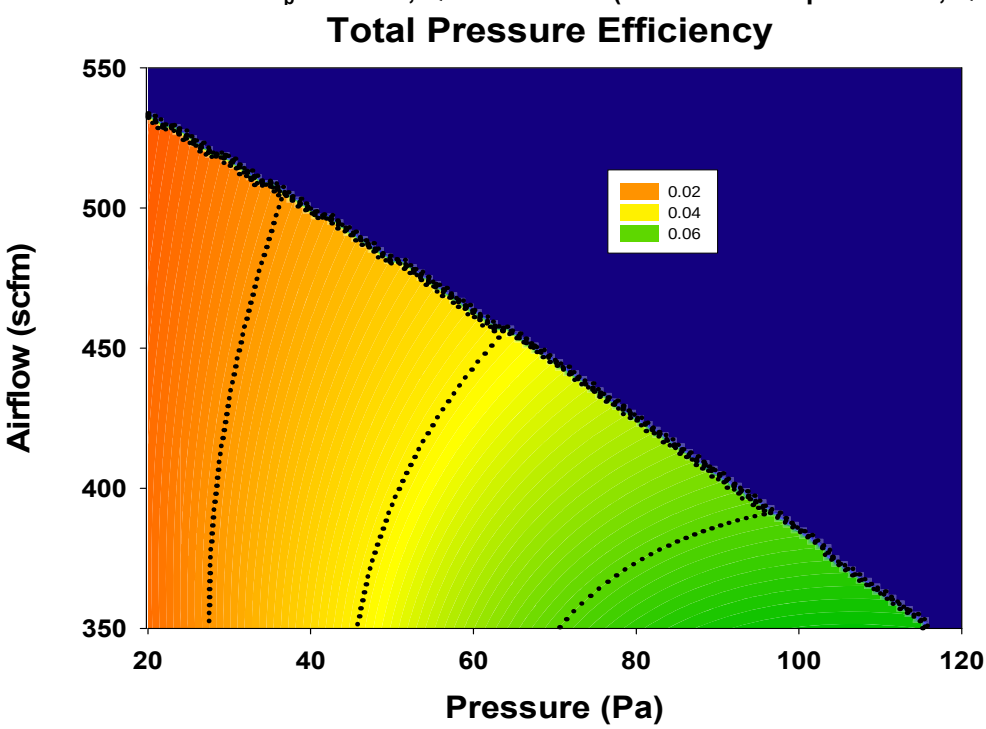

Note of unit conversion:

1 kPa = 4.015 inch water column
1 m/s = 196.9 feet per minute (fpm)
1 m^3/min = 35.3 ft^3/minute (cfm)

Report on

Laboratory Evaluation of Fan-filter Unit

LBNL Report, LBID-2588-15

Prepared by

Tengfang Xu and Duo Wang

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

September 15, 2006

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Acknowledgements

The project is funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Report on Laboratory Evaluation of Fan Filter Unit
Lawrence Berkeley National Laboratory, LBID-2588-15

Reporting Characteristics of the Unit Based on Manufacturer's Shipment

Manufacturer
Brand/Model
Serial Number
Unit Size
Fan Motor
Fan Wheel

Filter Efficiency

LBNL Tracking #

Reporting Energy Performance Based on Laboratory Testing

This report is based upon the laboratory evaluation of energy and airflow characteristics of the fan filter unit that was supplied to Lawrence Berkeley National Laboratory. All measured parameters, when applicable, are converted to their equivalents at the standard atmospheric condition (1 ATM, 20°C, sea level).

Total electric power demand, airflow rates, and pressure differential across the fan filter unit was recorded for a range of operable conditions of the unit. Normally, test conditions were selected with the airflow rates no less than 300 ft^3/min (8.5 m^3/min) under standard condition. The speed controller was set at its highest setting followed by lower setting. Figure 1 shows operable conditions that were tested with various speed control setting.

Figure 1. Operable conditions with various speed control setting

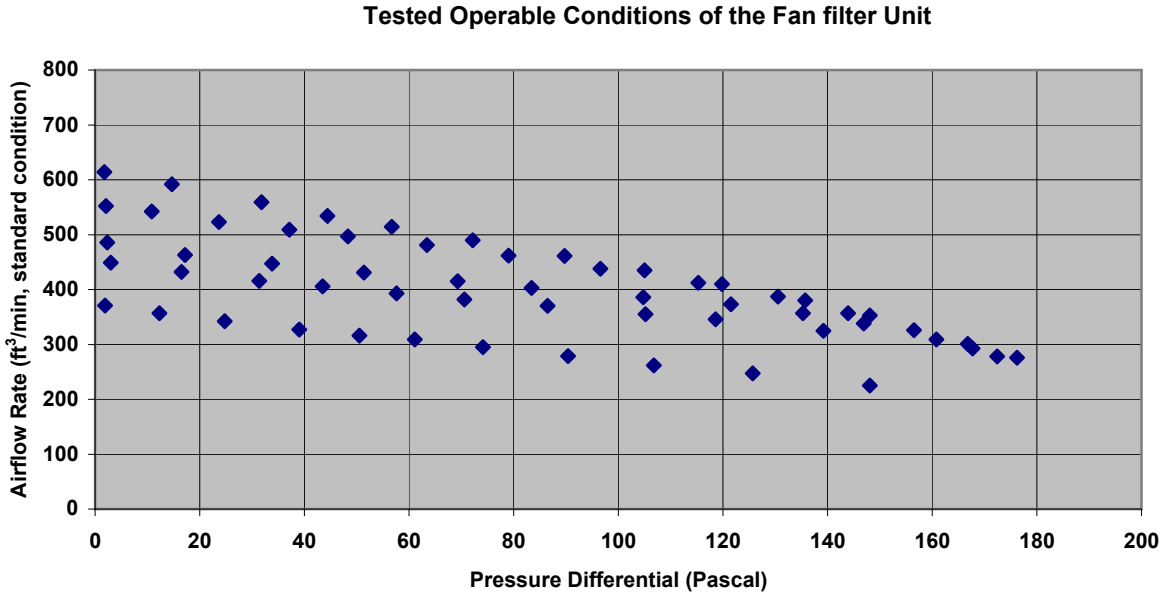

Total Electric Power Demand

Total electric power demand (EPD) of the fan filter unit includes all the electric power necessary to operate the fan filter unit. It was measured concurrently with airflow rates and pressures under all testing conditions. As a result, the total electric power demand can be calculated using the following equation.

The equation is derived from laboratory testing results from the operating conditions as shown in Figure 1. The R-square of the regression is included, which explains the statistical significance of the power demand predicted by the equation. A higher R-square number (with the possible maximum of 1) indicates higher degree of confidence in the power demand value derived from the laboratory testing.

Figures 2 illustrates the magnitude of total electric power demand within a selected range of operating conditions.

Electric Power Demand (Watt) = $C_0 + C_1 \times D_p + C_2 \times Q + C_{11} \times D_p^2 + C_{12} \times D_p \times Q + C_{22} \times Q^2$

R^2 = 0.999

where D_p is the pressure differential across the fan filter unit, in Pascal

Q is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)

C_0 =	1.16E+01
C_1 =	2.92E-01
C_2 =	-1.34E-01
C_{11} =	-1.35E-04
C_{12} =	1.04E-03
C_{22} =	7.70E-04

Figure 2 Total electric power demand of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

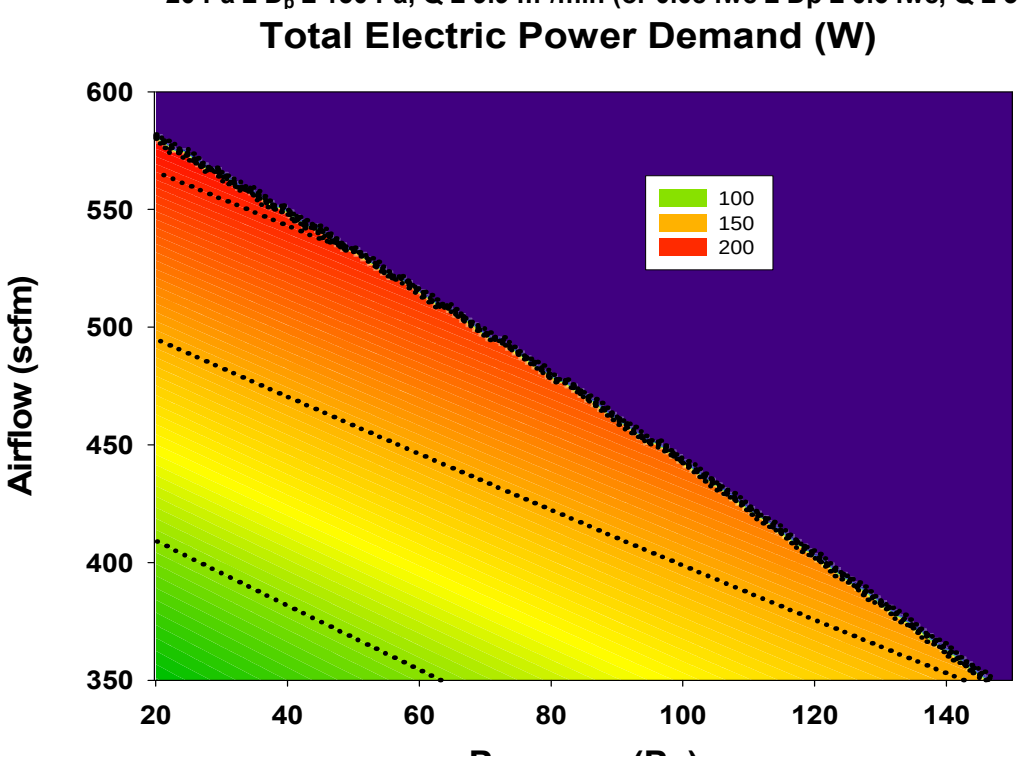

Pressure (Pa)

Total Pressure Efficiency

Total Pressure Efficiiency (TPE) is calculated by multiplying airflow rate and pressure differential across the unit then divided by total electric power demand as shown in the following equation.
Figures 3 illustrates the magnitude of total pressure efficiency within a selected range of operating conditions.

Total Pressure Efficiency (%) = 0.000471947443 x Q x D_p / EPD

where	Q	is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)
	D_p	is the pressure differential across the fan filter unit, in Pascal
	EPD	is the total electirc power demand (EPD) of the fan filter unit include all the electric power necessary to operate the fan filter unit, in Watt

Figure 3 Total pressure efficiency of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

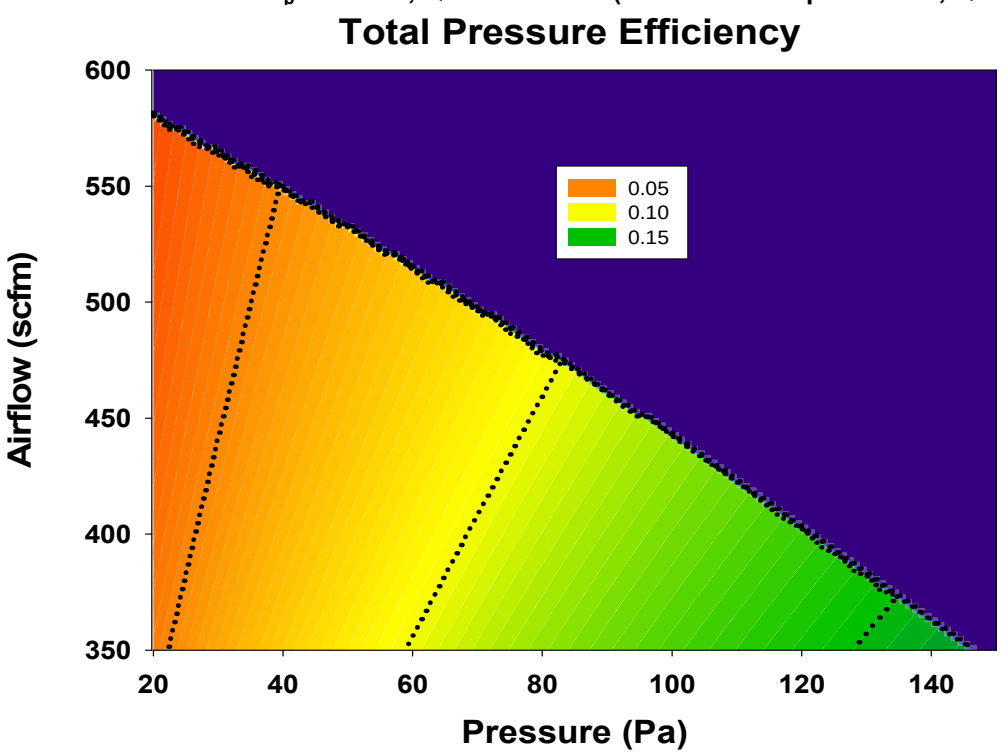

Note of unit conversion:

1 kPa = 4.015 inch water column
1 m/s = 196.9 feet per minute (fpm)
1 m^3/min = 35.3 ft^3/minute (cfm)

Report on

Laboratory Evaluation of Fan-filter Unit

LBNL Report, LBID-2588-16

Prepared by

Tengfang Xu and Duo Wang

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

September 18, 2006

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Acknowledgements

The project is funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Report on Laboratory Evaluation of Fan Filter Unit
Lawrence Berkeley National Laboratory, LBID-2588-16

Reporting Characteristics of the Unit Based on Manufacturer's Shipment

Manufacturer
Brand/Model
Serial Number
Unit Size
Fan Motor
Fan Wheel

Filter Efficiency

LBNL Tracking #

Reporting Energy Performance Based on Laboratory Testing

This report is based upon the laboratory evaluation of energy and airflow characteristics of the fan filter unit that was supplied to Lawrence Berkeley National Laboratory. All measured parameters, when applicable, are converted to their equivalents at the standard atmospheric condition (1 ATM, 20°C, sea level).

Total electric power demand, airflow rates, and pressure differential across the fan filter unit was recorded for a range of operable conditions of the unit. Normally, test conditions were selected with the airflow rates no less than 300 ft^3/min (8.5 m^3/min) under standard condition. The speed controller was set at its highest setting followed by lower setting. Figure 1 shows operable conditions that were tested with various speed control setting.

Figure 1. Operable conditions with various speed control setting

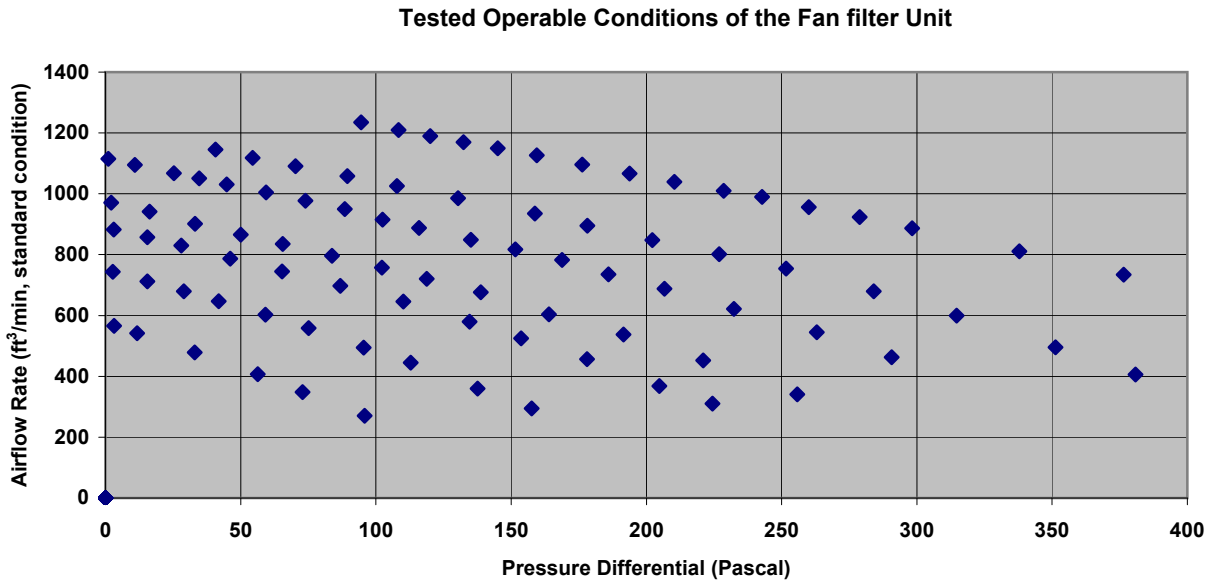

Total Electric Power Demand

Total electric power demand (EPD) of the fan filter unit includes all the electric power necessary to operate the fan filter unit. It was measured concurrently with airflow rates and pressures under all testing conditions. As a result, the total electric power demand can be calculated using the following equation.

The equation is derived from laboratory testing results from the operating conditions as shown in Figure 1. The R-square of the regression is included, which explains the statistical significance of the power demand predicted by the equation. A higher R-square number (with the possible maximum of 1) indicates higher degree of confidence in the power demand value derived from the laboratory testing.

Figures 2 illustrates the magnitude of total electric power demand within a selected range of operating conditions.

Electric Power Demand (Watt) = $C_0 + C_1 \times D_p + C_2 \times Q + C_{11} \times D_p^2 + C_{12} \times D_p \times Q + C_{22} \times Q^2$

R^2 = **0.999**

where D_p is the pressure differential across the fan filter unit, in Pascal

Q is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)

C_0 =	6.68E+01
C_1 =	2.22E-02
C_2 =	-1.76E-01
C_{11} =	3.75E-04
C_{12} =	9.34E-04
C_{22} =	3.24E-04

Figure 2 Total electric power demand of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

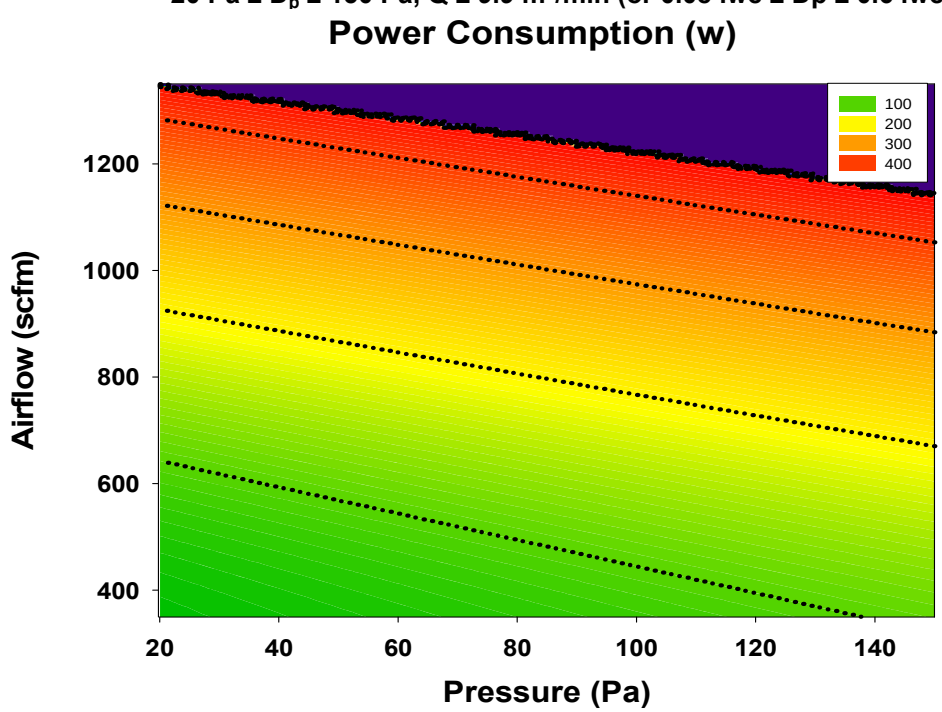

Total Pressure Efficiency

Total Pressure Efficiiency (TPE) is calculated by multiplying airflow rate and pressure differential across the unit then divided by total electric power demand as shown in the following equation.
Figures 3 illustrates the magnitude of total pressure efficiency within a selected range of operating conditions.

Total Pressure Efficiency (%) = 0.000471947443 x Q x D_p / EPD

where	Q	is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)
	D_p	is the pressure differential across the fan filter unit, in Pascal
	EPD	is the total electirc power demand (EPD) of the fan filter unit include all the electric power necessary to operate the fan filter unit, in Watt

Figure 3 Total pressure efficiency of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

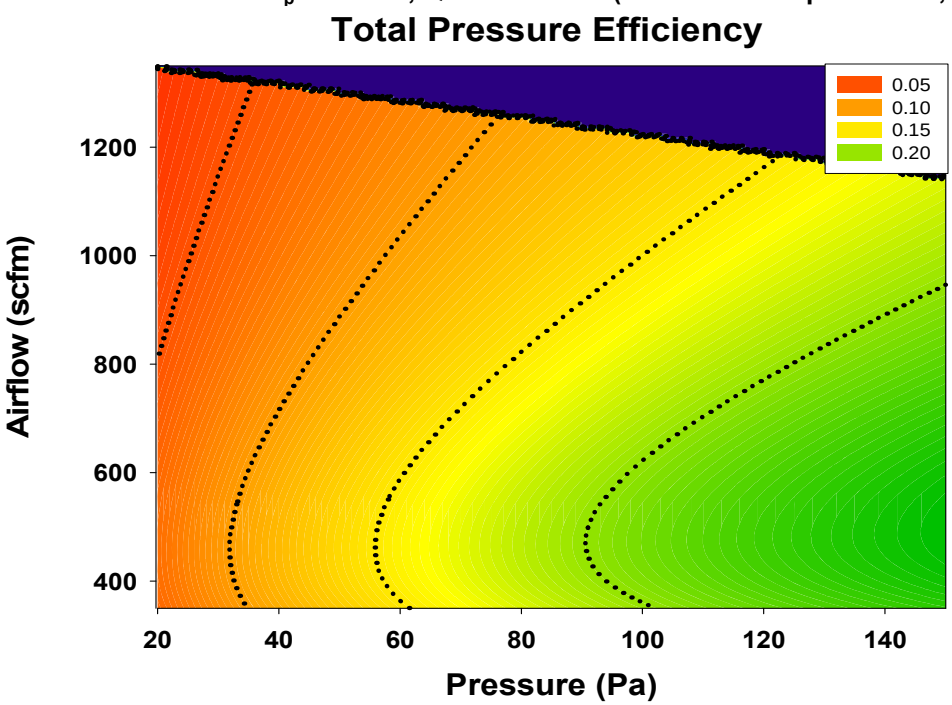

Note of unit conversion:

1 kPa = 4.015 inch water column
1 m/s = 196.9 feet per minute (fpm)
1 m^3/min = 35.3 ft^3/minute (cfm)

Report on

Laboratory Evaluation of Fan-filter Unit

LBNL Report, LBID-2588-17

Prepared by

Tengfang Xu and Duo Wang

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

September 18, 2006

Disclaimer

This document was prepared as an account of work sponsored by the United States Government and California Energy Commission. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor California Energy Commission, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Acknowledgements

The project is funded by the California Energy Commission's Industrial section of the Public Interest Energy Research (PIER) program. This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Building Technology, State, and Community Programs, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Report on Laboratory Evaluation of Fan Filter Unit
Lawrence Berkeley National Laboratory, LBID-2588-17

Reporting Characteristics of the Unit Based on Manufacturer's Shipment

Manufacturer
Brand/Model
Serial Number
Unit Size
Fan Motor
Fan Wheel

Filter Efficiency

LBNL Tracking #

Reporting Energy Performance Based on Laboratory Testing

This report is based upon the laboratory evaluation of energy and airflow characteristics of the fan filter unit that was supplied to Lawrence Berkeley National Laboratory. All measured parameters, when applicable, are converted to their equivalents at the standard atmospheric condition (1 ATM, 20°C, sea level).

Total electric power demand, airflow rates, and pressure differential across the fan filter unit was recorded for a range of operable conditions of the unit. Normally, test conditions were selected with the airflow rates no less than 300 ft^3/min (8.5 m^3/min) under standard condition. The speed controller was set at its highest setting followed by lower setting. Figure 1 shows operable conditions that were tested with various speed control setting.

Figure 1. Operable conditions with various speed control setting

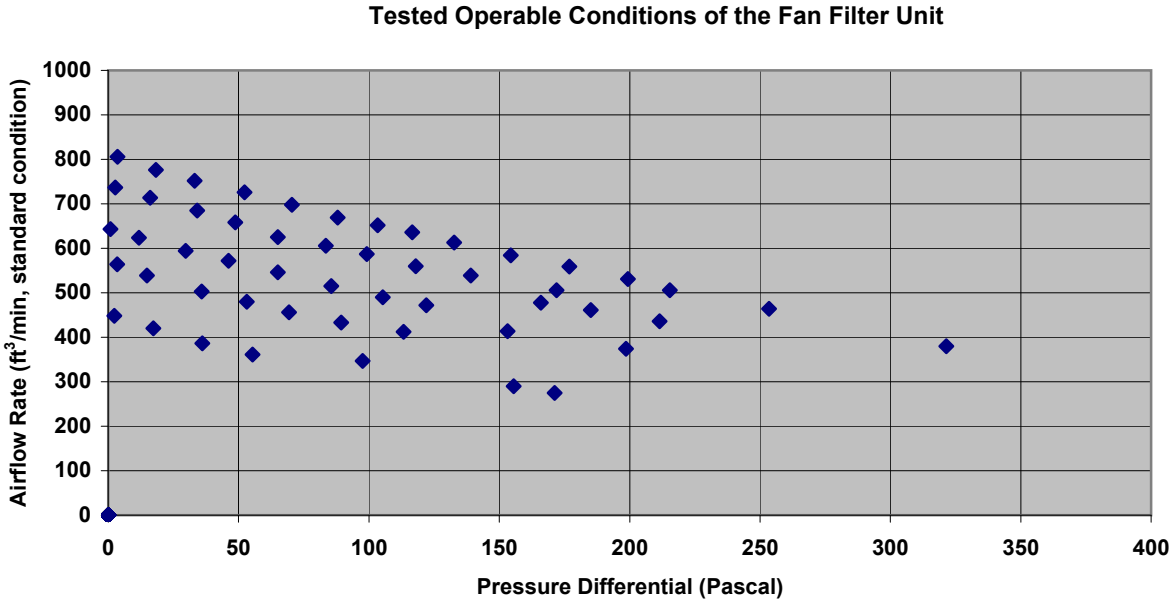

Total Electric Power Demand

Total electric power demand (EPD) of the fan filter unit includes all the electric power necessary to operate the fan filter unit. It was measured concurrently with airflow rates and pressures under all testing conditions. As a result, the total electric power demand can be calculated using the following equation.

The equation is derived from laboratory testing results from the operating conditions as shown in Figure 1. The R-square of the regression is included, which explains the statistical significance of the power demand predicted by the equation. A higher R-square number (with the possible maximum of 1) indicates higher degree of confidence in the power demand value derived from the laboratory testing.

Figures 2 illustrates the magnitude of total electric power demand within a selected range of operating conditions.

Electric Power Demand (Watt) = $C_0 + C_1 \times D_p + C_2 \times Q + C_{11} \times D_p^2 + C_{12} \times D_p \times Q + C_{22} \times Q^2$

R^2 = 0.999

where D_p is the pressure differential across the fan filter unit, in Pascal

Q is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)

C_0 =	-1.04E+00
C_1 =	6.61E-01
C_2 =	6.23E-01
C_{11} =	-8.46E-04
C_{12} =	-3.56E-04
C_{22} =	-2.65E-04

Figure 2 Total electric power demand of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

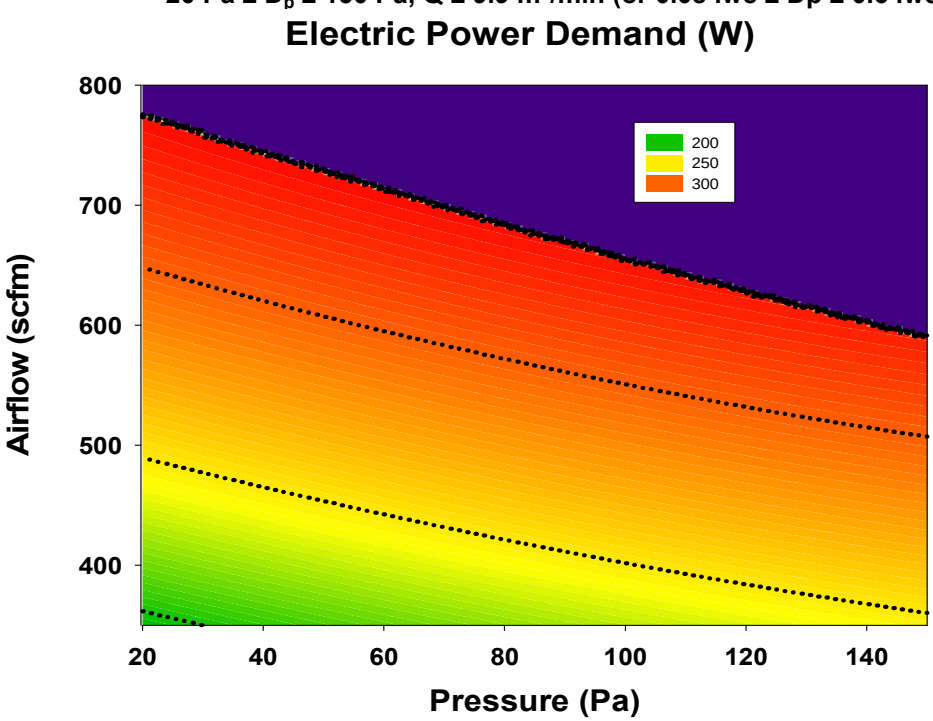

Total Pressure Efficiency

Total Pressure Efficiiency (TPE) is calculated by multiplying airflow rate and pressure differential across the unit then divided by total electric power demand as shown in the following equation.
Figures 3 illustrates the magnitude of total pressure efficiency within a selected range of operating conditions.

Total Pressure Efficiency (%) = 0.000471947443 x Q x D_p / EPD

where	Q	is the airflow rate across the unit under standard atmospheric condition, in scfm (standard ft^3/minute)
	D_p	is the pressure differential across the fan filter unit, in Pascal
	EPD	is the total electirc power demand (EPD) of the fan filter unit include all the electric power necessary to operate the fan filter unit, in Watt

Figure 3 Total pressure efficiency of the fan filter unit under selected operable conditions: 20 Pa ≤ D_p ≤ 150 Pa, Q ≥ 9.9 m^3/min (or 0.08 iwc ≤ Dp ≤ 0.6 iwc, Q ≥ 350 scfm)

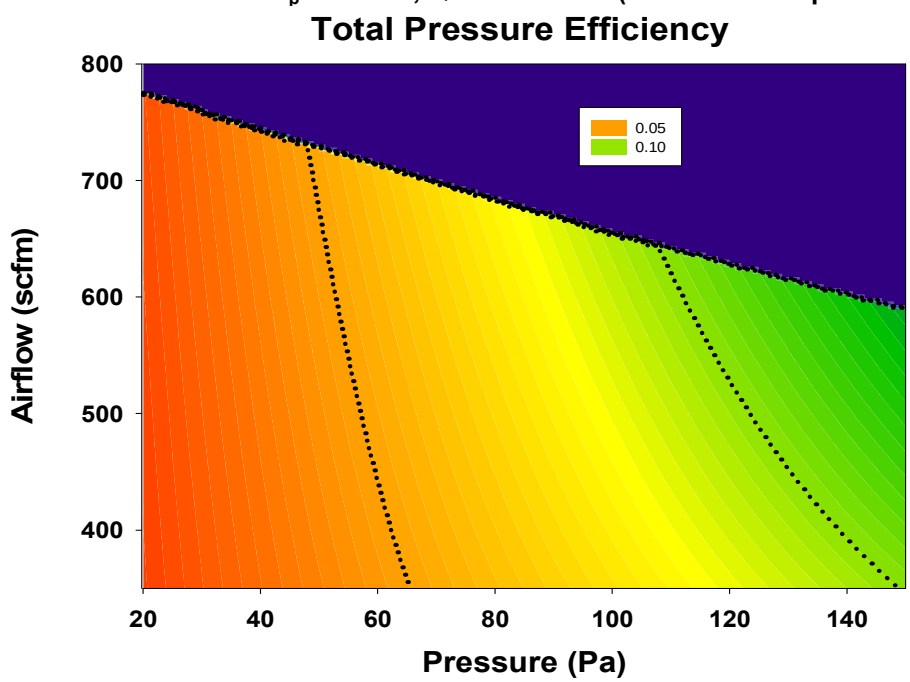

Note of unit conversion:

1 kPa = 4.015 inch water column
1 m/s = 196.9 feet per minute (fpm)
1 m^3/min = 35.3 ft^3/minute (cfm)

Attachment XIII: Energy Efficient Scientific Computing

DC Power for Improved Data Center Efficiency

January 2007

My Ton

Ecos Consulting

Brian Fortenbery

EPRI

William Tschudi

Lawrence Berkeley National Laboratory

The work described in this report was performed as part of Lawrence Berkeley National Laboratory's (LBNL) "High-Performance High-Tech Buildings" project focused on improving energy efficiency in Laboratory, Cleanroom, and Data Center facilities. This project was sponsored by the California Energy Commission's Public Interest Energy Research (PIER) Industrial program with technical leadership by LBNL. Ecos Consulting and EPRI Solutions performed the research reported herein collaboratively under subcontract to LBNL and with input from industry advisors and manufacturers.

Legal Notice

This report was prepared as a result of work sponsored by the California Energy Commission (Energy Commission). It does not necessarily represent the views of the Energy Commission, its employees, or the State of California. The Energy Commission, the State of California, its employees, contractors, and subcontractors make no warranty, express or implied, and assume no legal liability for the information in this report; nor does any party represent that the use of this information will not infringe upon privately owned rights. This report has not been approved or disapproved by the Energy Commission, nor has the Energy Commission passed upon the accuracy or adequacy of this information in this report.

Table of Contents

Table of Contents 3
1. Executive Summary 1
1.1 Introduction 1
1.2 Project Objectives 1
1.3 Project Implementation 2
1.4 Project Results 3
1.5 Conclusions and Recommendations 4
2. Background 7
2.1 Introduction 7
2.2 Project Background 7
2.3 Overview of Data Center Power Use 8
2.4 Overview of Current Data Center Power Delivery System 14
2.5 Demonstration Project Objectives 20
3. Project Chronology and Definition 22
3.1 Timeline 22
3.2 Participants 27
3.3 Power Delivery: Definition 28
3.4 Equipment Definition: AC Reference System 29
3.5 Equipment Definition: Facility-Level 30
3.6 Equipment Definition: Rack-Level 31
3.7 Equipment Definition: Monitoring 32
3.8 Other DC Systems Demonstrated (48V DC) 34
4. Implementation Issues 36
4.1 Commercially available equipment/missing components 36
4.2 Grounding, Bonding and Protection 37
4.2.1 Telecom DC Systems Grounding Example (up to 160V DC) 38
4.2.2 Additional considerations 42
4.2.2 Supplemental grounding and bonding 43
4.3 Reliability 46
4.4 Costs 47
4.5 Measurement Verification 48
4.6 Other Issues 49
5. Safety/Codes & Standards/Industry Associations 51
5.1 Selecting and Locating Suitable Circuit Protection 52
5.2 Grounding Methods to insure suitable operation of circuit protection 52
5.3 Leakage Currents and Shock Hazards 52
5.4 Connection Devices 52
5.5 Power Distribution Equipment within the Rack 53
5.6 Training and Certification of Personnel Working Within Racks 53

5.7 Training and Certification of Personnel Working on Distribution Systems. 53
5.8 Installation and Maintenance of Under Floor Power Distribution..... 54
5.9 Warnings and Labeling 54
5.10 Commissioning Tests..... 54
6. Results..... 55
6.1 AC Reference System Measurements..... 55
6.2 DC System Measurements..... 56
6.3 Efficiencies of AC- and DC-Based Systems 57
6.4 Non-Technical Project Results: Open House Sessions..... 59
7. Conclusions and Recommendations 62
Appendix A: Project Participants..... 65
Participants List 66
Project Organizers..... 66
Industry Partners 66
Appendix B: Location for One-Line Drawings of Demonstration Configurations 68
Appendix C: Participant "Code of Conduct"..... 69
Appendix D: Measurement Details..... 71
1 AC Reference System Measurements..... 71
2 DC System Measurements..... 71
3 Efficiencies of AC- and DC-Based Systems 71

1. Executive Summary

1.1 Introduction

The mission of data centers in use today calls for them to be much more energy intensive than other office buildings. Recent reports prepared by this project team for the California Energy Commission estimated that the US annual energy consumption by the equipment in these centers can use up to 14.6 TWh, with supporting equipment such as uninterruptible power supplies (UPS) accounting for up to 7.1 TWh. The data center market represents an important component of the California economy as well as a considerable and growing source of electrical demand. Preliminary investigations confirmed that research with the objective of improving the efficiency of data center was warranted.

Typical data center power delivery designs use AC (alternating current) power, distributed to the facility at 600V AC or 480V AC. This AC power is then stepped down to 208V AC or 120V AC for distribution to racks for use by servers and other information technology (IT) equipment. An UPS and energy storage system, such as batteries or flywheels is used to isolate equipment from power interruptions or other disturbances. This set up generally involves converting incoming AC power to DC (direct current) for energy storage. The DC power is then converted back to AC for the facility distribution grid and routed to power distribution units (PDUs) for distribution to equipment in racks.

Inside the servers and other IT equipment such as storage or networking units, power supplies convert AC (at 208/120V AC) to DC voltage needed for the digital electronics. Power supplies usually provide power factor correction as well as load isolation from the incoming power line for these sensitive electronic components. Thus, there can be up to six or more power conversion stages between facility power entry and the microprocessor or other data processing circuits.

The power losses due to the use of inefficient power conversion devices from both outside and within equipment result in a large loss of useful electrical power, as well as directly increasing the energy required to remove the heat produced. Thus, for every watt of power utilized to process data, about 0.9W is required to support power conversion. In addition, about 0.6 to 1 watt will be required to cool the power conversion equipment.

This report details a demonstration of alternative approaches for delivering power to computational and network equipment in a data center using DC, and comparing its efficiency to the more traditional approach using AC.

1.2 Project Objectives

The objectives of this demonstration project are to develop and demonstrate a power delivery system that does not contain as many power conversion stages using existing equipment and vendors where possible.

This project implemented a power delivery system that distributes DC to the server racks. The system used a single rectification stage, thereby removing the conventional UPS, transformer, and the rectifier in the server's first stage power supply. A standard AC distribution system is installed next to this DC system, server loads were connected and programmed to run identical routines. For this identical amount of computing work, the input power for whole system was measured and compared.

The following goals were identified for this demonstration project:

1. Show that DC-powered server(s) exists in the same form factor or can be built.
2. Show that DC-powered server(s) provides the same level of functionality.
3. Measure and document any efficiency gains from the elimination of multiple conversion steps in the delivery of DC power to the server hardware.
4. Identify areas requiring further development or follow up investigations.

1.3 Project Implementation

As part of the previous server power supplies and UPS work, the project team identified a number of industry stakeholders and held discussions with them regarding participation in a possible demonstration of an integrated DC power delivery system.

In early 2006, a suitable location for the demonstration project was found in Newark, CA. The project participants also helped to define the three configurations for the demonstration project. They are:

- **AC Reference Configuration**: This configuration is needed to simulate current data center typical set-up, delivering 208/120V AC input to AC-powered servers, and to be used as a reference to compare conversion efficiency.
- **Facility-Level DC Configuration**: This configuration is needed as the proof of concept – the ability to deliver high-voltage DC throughout the facility. This configuration handles the DC conversion/distribution at the building/data center level, converting 480V AC to 380V DC and delivering this directly to the DC-powered server units in the rack.
- **Rack-Level DC Configuration**: This configuration is needed to provide a possible migration option for AC data centers operators wishing to use DC equipment without facility-wide DC power distribution. This configuration accomplishes DC conversion/distribution at the rack level, using a rectifier unit to convert 208/120V AC at the rack, and delivering 380V DC to DC-powered servers.

In addition, a number of conditions were agreed upon by the participants and the project team on implementation, including:

- **Testing and Measurements**: The group agreed on test points and metrics, with emphasis on measuring the efficiency of the configurations – there would be no direct comparison of server equipment performance.
- **DC Input Voltage**: Due to compatibility with existing equipment and devices, the group settled on 380 V DC for the high voltage DC input.

Project participants contributed time, equipment, in kind services or input towards project implementation. The demonstration set up as defined was completed by mid June 2006,

and demonstrated to the press and other industry representatives on June 21. A series of "open houses" were held through out the months of July and August. Over 200 visitors attended these demonstrations.

The interest generated by the project helped to bring about other manufacturers of DC equipment, notably manufacturers of communication and support equipment using 48 V DC. These manufacturers provided additional computing and telecommunication equipment for the use of the demonstration project. The presence of this equipment provided another reminder that DC power distribution is not new, and has been safely and effectively used in telecommunication and data networks. In addition, it also showed that DC and AC power delivery system can co-locate in a data center facility.

1.4 Project Results

Our results indicate that the DC approach does provide an increase in conversion efficiency. We were fortunate enough to have access to two AC distribution systems as well as two DC conversion/distribution systems, and the efficiency ratios were determined for both sets.

Table ES1

System Efficiency	UPS Efficiency	Transformer Efficiency	PS Efficiency	System Efficiency
AC System A: Measured Efficiency	90%	98%	90%	79%
AC System B: Measured Efficiency	90%	98%	90%	79%
DC System A: Measured Efficiency	94%	100%	92%	87%

Energy Consumption	Compute Load (kWh)	Input Load (kWh)	Efficiency Gain
AC System A: Measured Consumption	23.3	26.0	
AC System B: Measured Consumption	23.3	25.9	
DC System A: Measured Consumption	22.7	24.1	
% Energy Consumption Improvement vs. AC System A			**7.3%**
% Energy Consumption Improvement vs. AC System B			**7.0%**

It can be seen (Table ES1) that there is about a 7% decrease in input energy using the first DC system compared to the "best in class" AC systems. With the second DC system, the values are slightly lower, but still about 5% improved over the AC systems (Table ES2).

Table ES2

System Efficiency	UPS Efficiency	Transformer Efficiency	PS Efficiency	System Efficiency
AC System A: Measured Efficiency	90%	98%	90%	79%
AC System B: Measured Efficiency	90%	98%	90%	79%
DC System B: Measured Efficiency	92%	100%	92%	85%

Energy Consumption	Compute Load (kWh)	Input Load (kWh)	Efficiency Gain
AC System A: Measured Consumption	23.3	26.0	
AC System B: Measured Consumption	23.3	25.9	
DC System B: Measured Consumption	22.7	24.6	
% Energy Consumption Improvement vs. AC System A			**5.0%**
% Energy Consumption Improvement vs. AC System B			**4.7%**

It should be noted that both of the AC distribution system used represent the best on the market with regard to efficiency. Both of the AC UPSs are high efficiency units, and the efficiencies of the power supplies in the AC servers – at 90%, are much higher than units currently found in today's data centers. By comparison, a typical AC system in today's data center would have a UPS that was about 85% efficient, and power supplies around 73% efficient.[1] The estimated improvement of the DC system over these "typical" systems is shown in Table ES3 below.

Table ES3

System Efficiency	UPS Efficiency	Transformer Efficiency	PS Efficiency	System Efficiency
AC Typical Distribution Efficiency	85%	98%	73%	61%
DC Distribution Efficiency	92%	100%	92%	85%

Energy Consumption	Compute Load (W)	Input Load (W)	Efficiency Gain
Typical AC Distribution Efficiency	10,000	16,445	
DC Distribution Option (Optimized)	10,000	11,815	
% Energy Consumption Improvement vs. Typical AC Distribution			**28.2%**

In this case, an improvement of over 28% is possible in an average data center. This means the DC distribution system, as demonstrated, will have the potential of using 28% less energy than the typical AC system found in today's data centers. Since data center HVAC loads are typically about the same as the IT load, this means that a 28% improvement in distribution and conversion also means a 28% overall facility level efficiency improvement.

It should be noted that the magnitude of the DC efficiency gain is highly dependent on the AC reference system and AC/DC power supply that it is being compared to. However, exposing the industry to "best in class" systems may also be useful in bringing attention to the need for improving the average efficiency of data centers.

1.5 Conclusions and Recommendations

This demonstration project was able to coordinate the participation of 21 organizations, their equipment, and/or in kind contribution, worked with other organizations' input throughout the implementation process, and assembled equipment and services worth over a million dollars in value. We were also able to conclusively demonstrate to the data center industry (via the 200+ open house attendees and the media coverage) that DC delivery systems are viable, can be 20% or more efficient than current AC delivery systems, be more reliable, and potentially cost less in the long run.

Overall, the project succeeded in meeting the objectives that were set out at the beginning. In particular:

Availability of DC Equipment: The demonstration project showed that DC-powered servers exist in the same form factor as AC servers or can be built and operated from existing components with minimal effort. DC servers currently exist (in the 48V DC form

[1] These are typical numbers that were found in our evaluation of the servers and UPSs markets.

factors), but 380V DC servers and storage equipment could be built and operated from existing components. Further, the demonstration project gave visibility to the DC power conversion and distribution equipment, highlighting two commercially available rectification systems, as well as UL-listed buss bars for DC applications.

DC Functionality: The project also showed that DC-powered servers can provide the same level of functionality and computing performance when compared to similarly configured and operating server containing AC power supplies. The demonstration equipment included storage units as well as DC network equipment that can use a variety of DC voltages.

Demonstrated Gains in Efficiency: The project demonstrated clear efficiency gains from the elimination of multiple conversion steps in the delivery of DC power to the server hardware. Results were measured and documented from two sets of DC delivery systems, and compared to two sets of AC delivery systems. In both cases, the DC delivery system showed a minimum of 5% to 7% efficiency gains without significant optimization over two AC distribution systems that are "best in class" and much more efficient than most systems found in today's data centers. These measured efficiency gains did not include additional gains from a reduction in cooling loads, which can have the potential for additional savings. Raising awareness of the AC - UPS system efficiency will have a benefit even if the DC solution is not embraced.

Follow-Up Investigations: A number of areas for follow up investigations were identified that will help generate industry discussions, and provide useful leverage points to move the industry forward in the direction of DC distribution. These include:

- **Grounding, Protection and Overloading Prevention**: A number of grounding, protection and overload prevention practices for DC data centers are discussed in this report for industry considerations.
- **Reliability**: Anecdotal data shows that DC-powered data centers have the potential to be more reliable than AC-powered data centers. However, data does not exist for DC voltages higher than 48V DC. Leadership will need to come from the industry in adoption or additional testing for the industry to move forward on this area.
- **Costs**: While cost data exists, it has not yet been compiled in a way that direct comparisons can be made for the two distribution systems (or their TCO). The demonstration has generated significant interests from data center designers and system integrators, and further discussion on this area can lead to at least a first-order estimate of DC distribution costs. A related area is the costs of DC components vs. AC components. Currently, AC components may enjoy better economies of scale, but wide-spread DC power adoption may change this equation.
- **Integration with Other Sources**: The use of DC in data center can also simplify the integration of alternative energy sources, such as solar and other forms of renewable energy, as well as fuel cells and distributed generation, which are all DC-based.
- **Other Issues**: Of significant concerns is the lack of industry knowledge of the advantages of DC distribution, as well as misconceptions about DC power. Additional education and outreach efforts will be required if the energy savings potentials of DC powered data centers are to be realized.

Going forward, there remain many barriers to the adoption of DC power distribution, and need additional follow on work. In particular, a number of barriers have been identified during the course of this projects, they include:

- **Increasing Awareness of DC Distribution**: The industry's current knowledge of available options for efficiency and DC distribution is quite low. Further, there is no single, trusted source of information, or an entity dedicated to the promotion of data center energy-efficiency and DC distribution (other than the CEC/PIER efforts). With the current industry support and interest, a "DC Power" association of some sort will help to focus interest and help to elevate awareness among the data center market.
- **Creating a Market for DC**: A number of market barriers still need to be addressed in a consistent, unified manner. Coordinating utilities and other efforts, at least in California, will go far towards getting DC approaches to take hold in the data center market. There is utility interest in establishing a baseline of performance and cost, which can then help to address at least the early adoption barrier of cost. Other efforts are still needed, and strategies to address market transformation used by the conservation movement can be directly applied here. In addition, the US Congress has recognized the potential for energy savings with HR-5646, so that coordination with DOE and EPA is needed to ensure no duplication of efforts.
- **Develop standards to accelerate adoption**: Agreement on distribution voltages, electrical connectors, grounding, DC power quality, and other issues will be important to enable the market to adopt DC distribution on a large scale. The PIER program should facilitate these efforts by bringing together the appropriate industry representatives.
- **Developing Pilot Projects**: Once the Demonstration Project was completed, there are no other places where such a set up can be found. Efforts are needed to continue the Demonstration Project's role in informing the industry of the DC distribution alternative. Discussions are underway with a number of "early adopters," and the project team proposes to establish several pilot projects to:
    - Create demand for DC servers to enable certification efforts to proceed
    - Determine cost factors for DC systems - capital and operating cost
    - Evaluate and resolve any remaining barriers
    - Publicize successful systems in real data centers.

2. Background

2.1 Introduction

This report details a demonstration of alternative approaches for delivering power to computational and network equipment in a data center using DC (direct current), and comparing its efficiency to the more traditional approach using AC (alternating current). The report is divided into five main sections: The first section discusses data centers' energy requirements, power delivery systems and efficiency implications. The second section deals with the project chronology and definition of project objectives and equipment. The third section covers issues encountered during the implementation. The fourth section discusses safety and other associated issues. The fifth section provides the project results, both technical and non-technical. The last section summarizes the conclusions, recommendations, and next steps.

2.2 Project Background

There are many types of "data centers" in use today across the U.S. and around the world.[2] They range from corporate data centers in a wide range of industries: banks, telecommunication facilities, to Internet services facilities, as well as institutions such as research organizations, universities, national laboratories, and government facilities. They generally consist of a collection of servers – high powered computers typically housed in equipment racks, in tandem with networking and storage equipment to store and move large amounts of data. Along with the data processing equipment are the supporting equipment: power distribution, back-up systems, cooling, and other services to maintain data center operations.

A number of other definitions have also been used to describe "data centers" and their functions.[3] At the most basic level, data centers house computers and information technology (IT) equipment to provide functions such as data information storage, data processing, and information dissemination. With the boom associated with the Internet in the late '90's came new names for data centers including "server farms", "collocation facilities" and "telecommunication hotels."

Regardless of their make-up or affiliation, most data centers' mission calls for them to be much more energy intensive than other office buildings. This is due to the high power requirements of not only the computing equipment, but also for the infrastructure needed to support the computing equipment, and their constant availability, or "up time." In fact,

[2] The terms "data centers" and "datacenters" are used interchangeably by the industry. This report uses "data centers." For the purpose of this report, these are dedicated facilities that comprise more than 1,000 square feet of computational space and support equipment.

[3] See, for example, http://en.wikipedia.org/wiki/Data_center, or ACEEE and CECS 2001; Aebischer et al. 2002b; Blount et al. 2001; Brown et al. 2001; Callsen 2000; Elliot 2001; Gruener 2000; Mitchell-Jackson 2001; Planet-TECH 2002; Robertson and Romm 2002.

large data centers more closely resemble industrial facilities in their energy density than typical commercial buildings.

Figure 1. Data Center

Source: UC Berkeley

The data center market has a significant presence in California, and represents an important component of the California economy as well as a considerable and growing source of electrical demand on the state's electric utilities. Preliminary investigations and utility case studies confirmed that research with the objective of reducing the large, continuous electrical loads in data centers was warranted. Because of this, California utilities and the California Energy Commission (CEC) became interested in learning more about the high-tech sector and the data center market.

As part of the California Energy Commission's PIER (Public Interest Energy Research) initiative on improving energy efficiency for high-tech buildings, Lawrence Berkeley National Laboratory (LBNL) was tasked with preparing R&D roadmaps for data centers as well as for cleanrooms and laboratories, focusing on reducing their energy use.[4] Using these roadmaps, sponsors of public sector R&D programs are able to prioritize their investments in the research, development, and demonstration of innovative energy efficient technologies and best practices for the next decade.

2.3 Overview of Data Center Power Use

From the beginning of commercial computing, data centers have been important to industries, institutions, and governmental agencies. However, it was the Internet and the rise of "mission-critical" computing facilities that brought energy consumption in data centers to the forefront. In 1999, a widely cited article by Mark Mills and Peter Huber posited that the "Internet" consumed about 8% of US electricity in 1998, and that this

[4] Tschudi, W.F., Xu, D. Sartor, and D. Stein. 2003. "High Performance Data Centers: A Research Roadmap". LBNL-53483. Available from http://hightech.lbl.gov/library.html#Publications.

would likely grow significantly to account for 30 to 50% of U.S. electricity consumption by 2010 or 2020. They cited the dramatic growth in Internet-related equipment installed base since 1995, particularly servers and computer network equipment needed for the Internet.[5]

Mills and Huber's work generated significant interests by the US Congress, the US Department of Energy (US DOE) and prompted additional investigations by other researchers. Lawrence Berkeley National Laboratory and Arthur D. Little, Inc., who conducted more detailed studies (Koomey et al., 1999; Koomey, 2000; and Roth, Goldstein, and Kleinman, 2002) conclude that this initial estimate was highly overstated, and that computer, office, and network ICT equipment potentially accounted for only about 3 percent of U.S. power consumption in 1999/2000.[6]

Common among these follow on studies is the observation that there is a large variation in energy intensity and energy efficiency of key systems in the various facilities under examination. It is safe to say that current understanding of data center energy use in the industries and institutions that rely on them is very limited. Typically, data center professionals have a thorough understanding of issues related to power quality, reliability, and availability, as these have been high priority items, but not energy consumption or efficiency.

To accurately estimate the power consumption and future needs of data centers in the US, one place to begin is to characterize the stock of existing data centers and their load intensity. These characteristics have turned out to be difficult to estimate. The market is a competitive one, and can be constantly changing. Further, there is no reliable source of market data covering all of the various types of data centers, and companies operating them tend not to publicly disclose their capacities and loading. Load intensity for data centers supporting the Internet fluctuates greatly with the rise and decline of internet-based companies. However, data center load intensity is also affected by the trends in computing capability and energy intensity within IT equipment.

There are three main energy trends that have been identified by studies as drivers for this new awareness regarding data center energy requirements:

- **Advances in computer technologies**: faster chip technology was creating higher heat density in smaller and smaller geometries. The simultaneous compaction and increase in electrical power caused concern over the ability to cool future generations of IT equipment.

[5] Mills, M and P. Huber, "Dig More Coal—the PCs are Coming," Forbes, May 1999.

[6] Roth, K; F. Goldstein; J. Kleinman "Energy Consumption by Commercial Office and Telecommunications Equipment in the U.S." Presentation at "Impact of ICT on the Energy System" 2002.

- **Significant increases in electrical power requests from utilities by facilities**: meeting growing requests by computing facilities could require major changes in current electrical utility generation and distribution infrastructure. [7]

- **Predictions of large increases in electrical demands for future facilities**: IT professionals, data center operators, and facility designers have been predicting even larger requirements for future computing facilities.[8]

Ecos Consulting and EPRI Solutions were tasked with investigating and refining existing estimates of the overall energy requirements by this sector as a follow up to the above studies in this area. In particular, Roth, Goldstein, and Kleinman in 2002 arrived at an estimated Annual Energy Consumption (AEC) for IT equipment using a bottom-up analysis of the stock, usage patterns and power draw (by mode) of IT equipment. They concluded in their analysis that for 2000, the estimated AEC was about 10.1 TWh for US data centers, as shown in Table 1, below.[9] Their methodology is as follows:

- Generate a list of equipment types and collect existing data from literature.
- Develop a preliminary estimate of national energy consumption for each equipment type.
- Select 5 to 10 equipment types for further evaluation, based upon preliminary calculations and perceived growth in future energy consumption. Ideally, the selected equipment types should represent 66% to 75% of all energy consumed by office and telecommunications equipment in the commercial buildings sector.
- Briefly describe the 5 to 10 equipment types selected, with the intent to provide insight into how each equipment type uses energy and function in a commercial office environment, including: physical description, functions performed, and commercialization history.
- Develop refined bottom-up estimates of national energy consumption of each selected equipment type, for Y2000 and projections for Y2005 and Y2010.
- Compare the results with the results of other studies.
- Qualitatively discuss possible indirect impact of commercial office and telecommunications equipment upon energy consumption, e.g. e-commerce, building heating and cooling loads, etc.
- Publish the findings in a report, including feedback from government and industry experts.

It is important to note that the Roth's methodology established here for estimating the AEC, and also the follow up work by Ecos and EPRI only account for the direct energy consumption of the computing equipment (servers/computational equipment and UPSs)

---

[7] Tschudi, W.F., P. Sreedharan, T. Xu, D. Coup, and P. Roggensack. 2003. "Data Centers and Energy Use – Let's Look at the Data." ACEEE Conference Paper.

[8] Markoff, J. and S. Hansel, "Hiding in Plain Sight, Google Seeks More Power" *The New York Times*, June 14, 2006.

[9] 1 TWh = 1,000,000,000,000 Wh, or 1 billion kWh.

and not the additional energy required to support or cool the facilities.[10] Benchmarking performed by LBNL, while illustrating a wide range of performance, shows that it is not uncommon that infrastructure load requirements are of similar magnitude to the IT equipment load.

Table 1. Estimated 2000 Annual Energy Consumption for the IT Sector

Equipment Type	US AEC in TWh	% of Total AEC
Computers (between $25k and $349k)	24	22%
Computers (<$25k, i.e. PCs)	18	16%
Monitors	18	16%
Computers (>$349k)	10	9%
Laser Printers	9.3	8%
Copiers	8.3	7%
Server Computers	7.7	7%
Telecomunications Network Equipment	7.1	6%
Computer Network Equipment	6	5%
Uninterruptible Power Supplies (UPSs)	2.5	2%
Total Equipment	**110.9**	**100%**

Source: Roth et al, 2002.

Ecos' and EPRI Solution's follow up investigations of the AEC primarily focused on the various equipment and components in use in data centers and their energy consumption, with the main focus on server and equipment power supplies as well as uninterruptible power supplies.[11] For server power supplies, the goals for the project included an analysis of power supply efficiencies, which required the development of an accepted test protocol for server power supplies, lab and field testing of a broad range of server power supplies and documenting the results. Using the efficiency findings and market data, we arrived at estimates of overall energy consumption of servers in the United States, especially in the State of California (as well as the potential savings from the use of more efficient units).[12]

Similarly, for the UPS market, the investigation included an analysis of UPS topologies and a complete analysis of distribution of UPS design efficiencies, which included the development of a UPS test protocol. Using the efficiency findings and market data, we estimated the overall energy consumption of UPSs in the United States, especially for the

[10] Note: In this report, several terms are used generically, including "servers" as a generic term to describe the computing equipment used in data centers, including racks of servers and contents; and "UPSs" to describe any number of back-up power sources and approaches to maintaining "up time".

[11] Ecos Consulting and EPRI Solutions. 2005. "High Performance Buildings: Data Centers -- Uninterruptible Power Supplies (UPS), and "Ecos Consulting and EPRI Solutions. 2005. "High-Performance Buildings: Data Centers -- Server Power Supplies."

[12] Other objectives of the project included working with industry groups to press the case for more efficient power supplies. the wide circulation of efficiency findings to the industry through the CEC, PIER, Lawrence Berkeley National Laboratories (LBNL), and other industry and efficiency venues, such as SSI – the Server System Infrastructure group, the PSMA – the Power Sources Manufacturers Association, the American Society of Heating, Refrigeration and Air Conditioning Engineers – ASHRAE, and www.efficientpowersupplies.org.

State of California, as well as the potential savings from the use of more efficient units. Finally, the investigation also proposed efficiency levels and a labeling scheme for various types of UPS units as a way to encourage the use of more efficient UPS units. The results from the server power supplies and UPS studies are presented below.

Annual Energy Consumption of Servers (Data Processing Equipment)

Starting from the basic methodology established in the ADL/US DOE study, Ecos and EPRI Solutions constructed a revised estimate of the annual energy consumption of servers in the United States. Test results for this work showed that most server power supplies' efficiency at converting AC to DC typically peaks at loads between 50-60% and drops off dramatically at loads under 30%. The tested power supplies have efficiencies in the 70-75% range (at 50% load). Figure 2 below shows the efficiency range of the server power supplies tested.

Figure 2. Power Supply Efficiency Curves

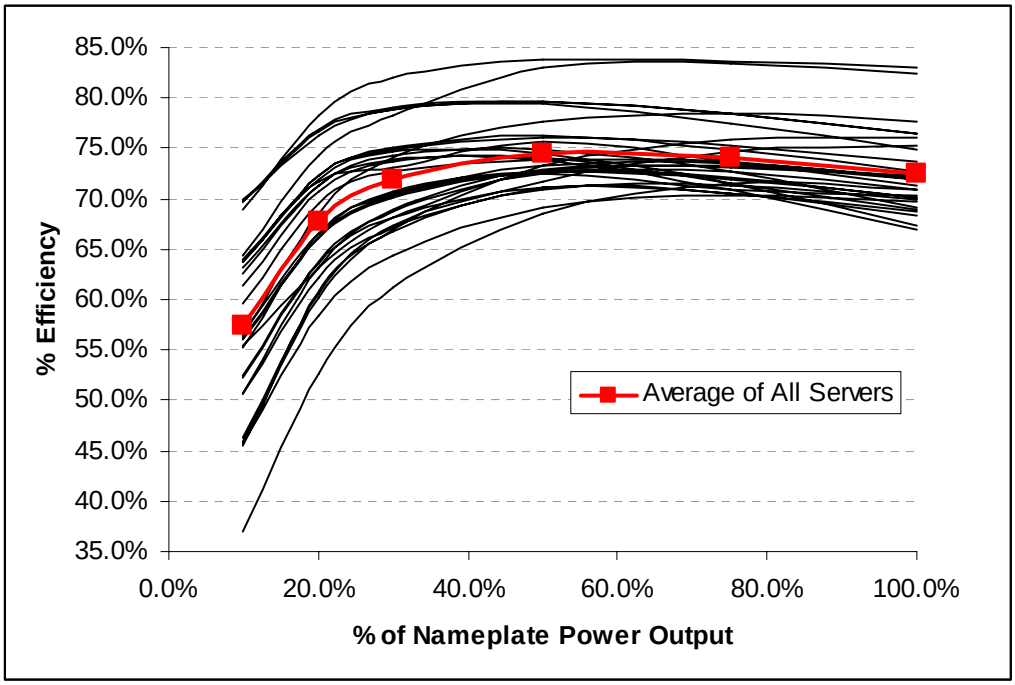

The research also showed that server power supply designs with poor efficiency are still available and can result in unnecessary power consumption and excess heat generation while in use. Further, the servers that were surveyed do not exhibit any activity-based power management. The absence of this process in the servers that were measured might help to explain the lack of any correlation between the utilization curve of the CPU and the load duration curves of the power supplies.

Figure 3 below shows our revised estimate of server AEC for 2004. The revised AEC is 14.6 TWh, which represents a 45% increase over the ADL/US DOE estimate of 10.1 TWh for 2000. This increase is attributable to the growth in both the number and average

power draw of low-end servers, which comprise the majority of units. The estimate for California is in the range of 1.5 TWh to 2.2 TWh.

Figure 3. Estimated Annual Energy Consumption of Servers

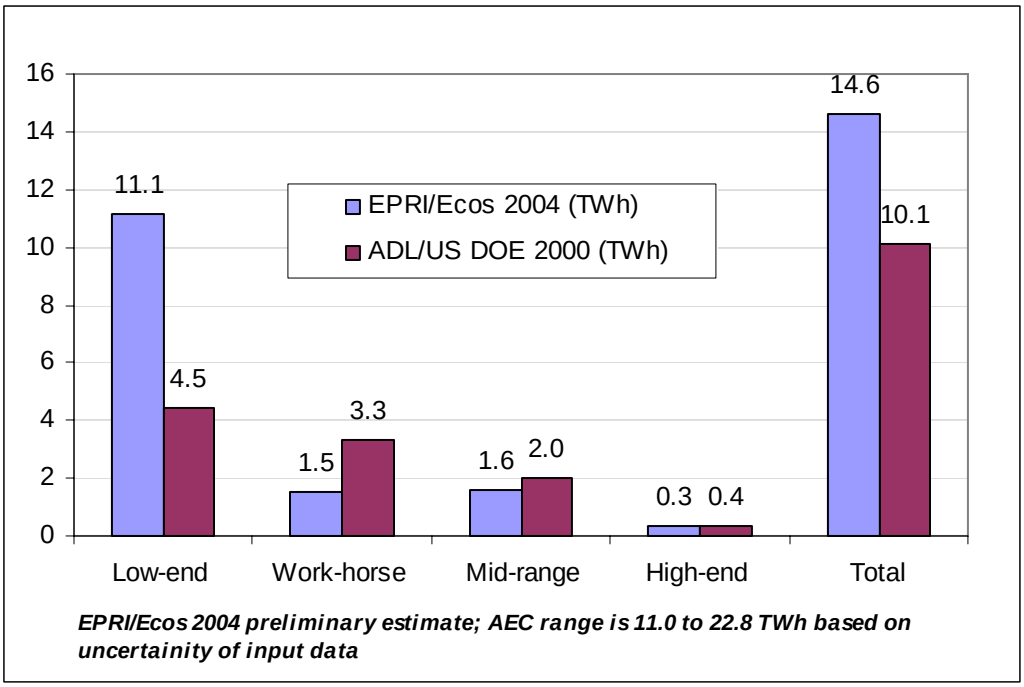

Annual Energy Consumption of UPSs

To estimate UPS power consumption, Ecos and EPRI started by sizing the UPS market and the installed stock of UPSs in the data center/IT sector. Our results indicate that there can be a wide variation in efficiency between various UPS configurations. The table below summarizes the characteristic efficiency of a number of UPS topologies at various load conditions and shows the average efficiency for all of the UPSs measured.

The table below also shows that regardless of the configuration or type, UPSs tend to be more efficient at full rather than part load which is where they typically operate due to over sizing and redundancy strategies.

Table 2. UPS Topologies and Efficiencies

UPS Topology	Efficiency at 25% Load	Efficiency at 50% Load	Efficiency at 75% Load	Efficiency at 100% Load
Delta-Conversion	93% - 94%	96% - 97%	97%	97%
Double-conversion	81% - 93%	85% - 94%	86% - 95%	86% - 95%
Line-Interactive	NA	97% - 98%	98%	98%
Standby	NA	NA	NA	NA
AVERAGE OF ALL UNITS	86%	89%	90%	90%

Note: A straight average was used in the table above.

Using efficiency data and our estimates of UPS stock, we constructed a first-order estimate of UPS energy use for the U.S. This energy use and savings estimate is limited

to the data center/IT sector due to the scope of our research, in which we examined the efficiency of data centers only. Based on our estimates, UPSs in the data center/IT sector currently handle roughly 7.1 TWh (7.1 billion kWh) of electricity per year. We estimate that the State of California alone, which contains roughly 15% of the nation's data center floor space, consumes about 1 billion kWh of electricity and spends about $100 million per year in electric bills and a significant amount is due to power conversion losses. Our national energy use figures are noticeably higher than past estimates made by Arthur D. Little for the U.S. Department of Energy, which estimated UPS energy consumption in the U.S. "IT/telecom" sector at 5.8 TWh per year.

The quantification of the annual energy consumption for data centers – other than providing a plausible range and a first-order estimate – remains an elusive goal for a number of reasons. However, with the additional information gathered from the server power supplies and UPS studies, a number of observations can be made:

1. Data centers, both in the US and California, can account for a significant demand on electrical power, based on the number of servers in operation.

2. When supporting equipment are taken into account, the range of AEC for data centers become very large, even without taking their cooling needs into consideration.

3. Based on findings from our investigations into this market, there are opportunities for more efficient equipment to reduce this very large amount of energy annually needed to operate data centers.

4. Because more efficient equipment tend to reduce losses, especially conversion losses in the power delivery chain, increases in equipment efficiency can translate directly into reduction in cooling loads, leading to more energy savings. Another potential benefit is a reduction in heat-related failure.

2.4 Overview of Current Data Center Power Delivery System

In order to understand the opportunities for improving data center efficiency, it is necessary to examine the current data center power delivery system from end to end. That is, an examination of the delivery chain starting from the facility level and ending at the processor level. There are often several levels of power conversion occurring in data center facilities and within the IT equipment contained therein, as well as redundancies and other accepted practices by data center designers and builders to increase availability, or "up time". This examination will not only help to identify areas where there can be areas for improvements, but may also pinpoint issues and other opportunities for simplification. While we have identified power supplies and UPSs as two significant areas based on our research, they by no means are the only areas where possible efficiency improvements can be made.

Data Center Electrical Distribution

Current data center power delivery designs use AC power, distributed from the utility (or "at the curb") to the facility at 600V AC or 480V AC depending on the size of the facility. This AC power is then stepped down to 208V AC or 120V AC via transformers for distribution to server racks for use by the servers and other data center equipment. AC is also used in powering the ancillary support equipment, such as HVAC and lighting.[13] In the U.S., servers typically take 208V AC (or 120V AC) input. Figure 4, below illustrate this power delivery and use system.

Figure 4. Power Distribution for Data Center Equipment

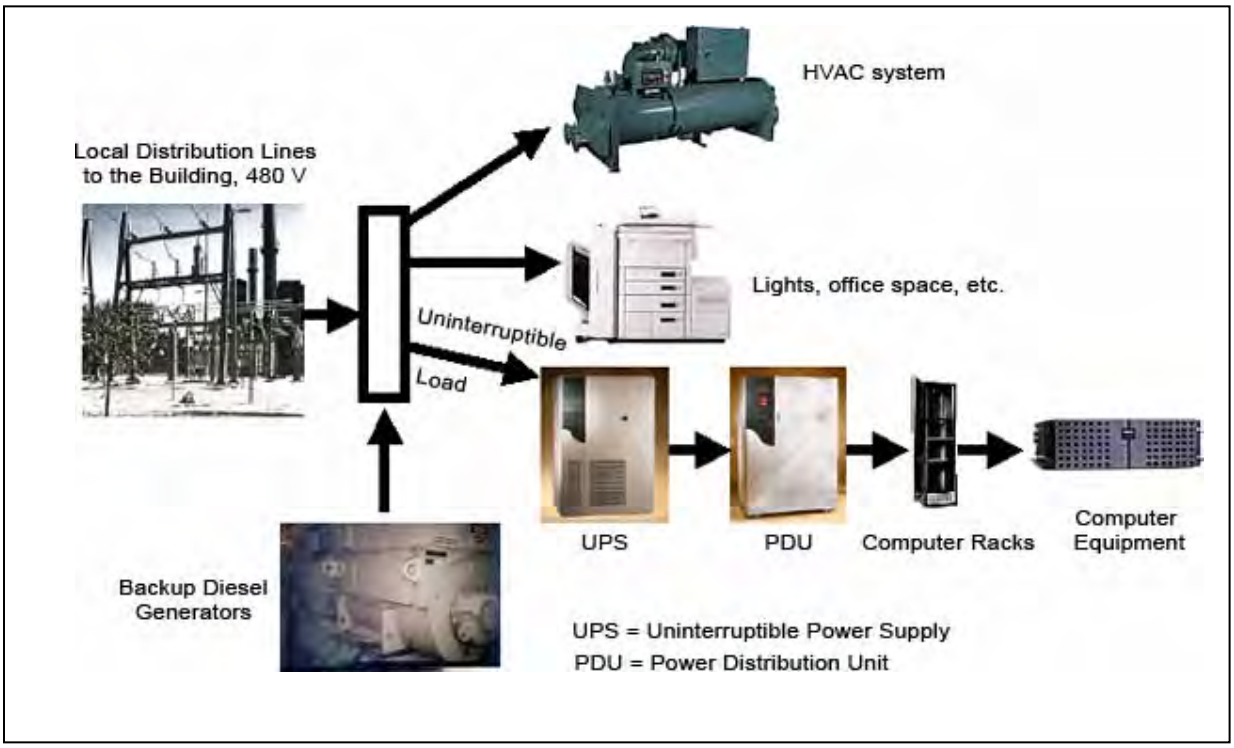

In this power chain, there is usually an uninterruptible power supply (UPS) which is coupled with and energy storage system, such as batteries or flywheels to ensure the servers and associated processing equipment are never exposed to power interruptions or other power line disturbances. Depending on the UPS configuration, the central AC UPS generally involves conversion from incoming AC power to DC for energy storage, and then reconverting from DC back to AC for distribution throughout the facility via power distribution units (PDUs).

Moreover, if the UPS requires a maintenance bypass switch, facility-level entry power must be transformed from 480V AC down to 208V AC prior to input to the UPS. The central battery is then connected as an alternate input to the DC-AC converter so that if incoming raw AC power is interrupted, the system automatically switches over to battery

[13] To insure availability, data center HVAC systems can also be on a UPS protected circuit.

power. Uninterruptible AC power is then passed through an AC power distribution grid and routed to PDUs for distribution to individual data processing apparatus such as equipment racks. The servers and other equipment in the racks receive the AC output of the UPS (converted from DC) and then convert again to DC within its power supply. Figure 5, below presents a simplified view of this power delivery chain.

Figure 5 . Simplified Illustration of Power Conversion Steps

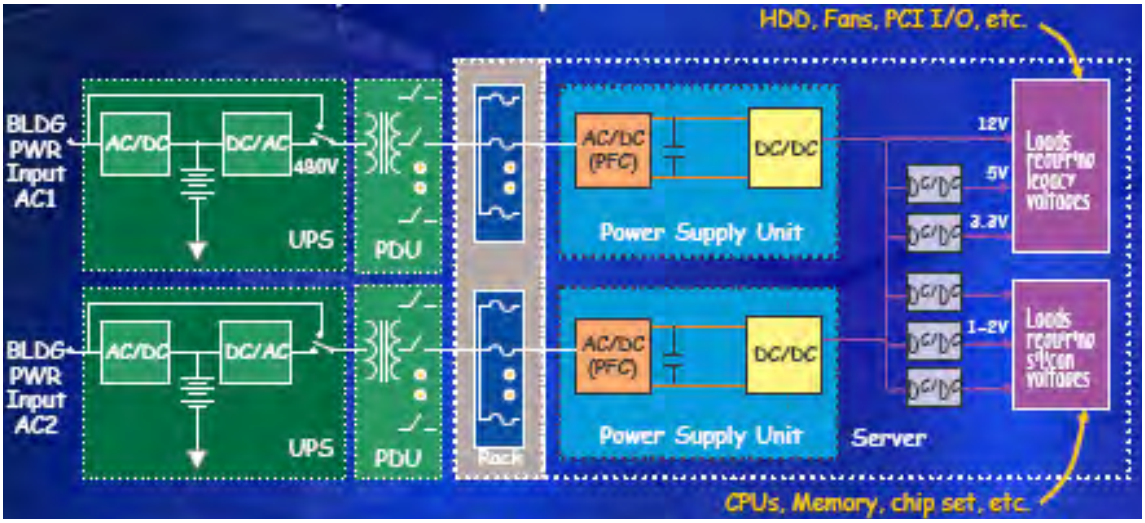

Source: Intel Corp.

The power quality and reliability requirements are expensive, but in data centers, like no other industry, these are paramount, and the extra cost is borne without question. In fact, the PDU is often connected to two parallel UPSs in some data center designs, to provide redundancy for higher availability in case one UPS fails or taken down for maintenance service or repair. In addition to the UPS redundancy, the power supplies in the servers themselves are often redundant, with two or even three power supplies in each server box ready and capable of powering the server completely in case one or more of them fails.

The typical large data center AC power distribution system is very complex and costly. Because of the reliability requirements of these facilities there are many redundancies, power conversions, paralleling controls, static transfer switches and bypass connections built into the system. It is because of these complexities that a data center power system is so much more reliable then a typical commercial building power system. It is also because of these complexities that a current day data center power system is very inefficient in both space utilization and energy efficiency.

A detailed example of the power distribution system approach (one of many) in a data center is shown below.

Figure 6. Power Delivery System Details

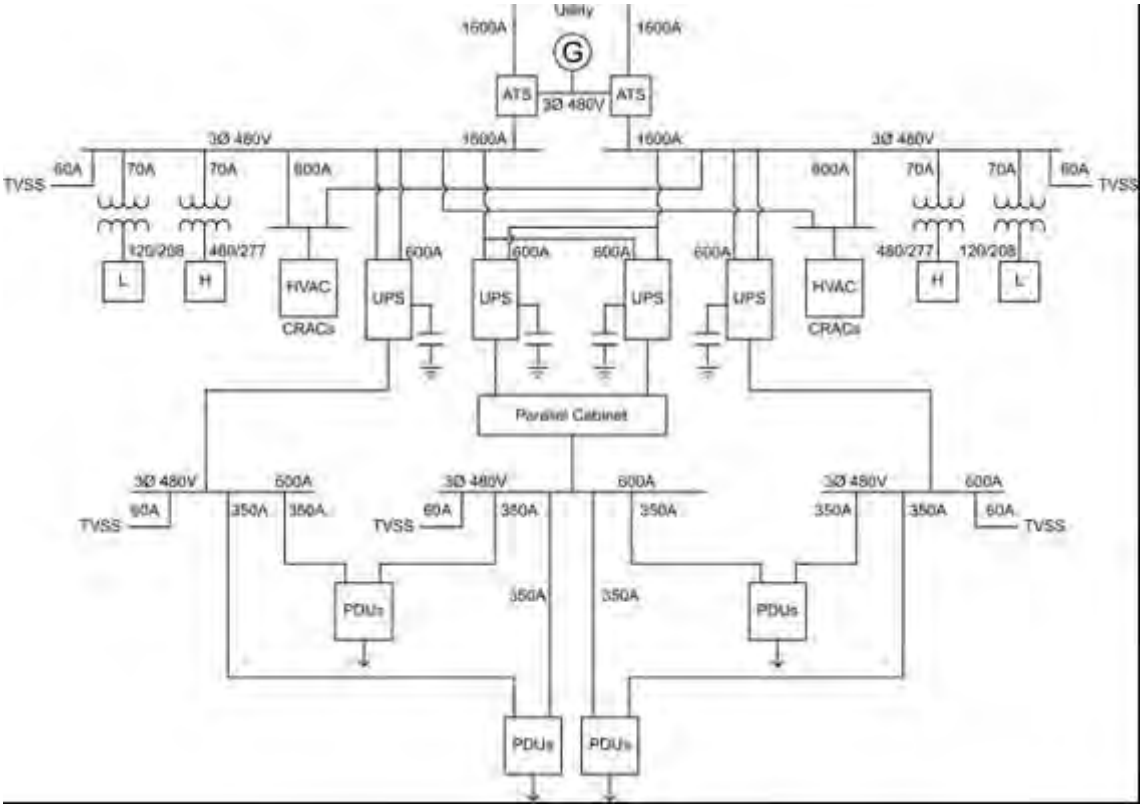

Typically, uninterruptible power is viewed at the facility level, as opposed to the equipment level. This perspective provides an easy division between the facility power equipment and data processing equipment, with each focusing on a different part of the power delivery. However, this view also makes it difficult to ascertain and optimize overall operating efficiency and total cost of ownership, since losses from each of these conversion processes directly translate into heat, adding to both the server cooling load and ultimately, the overall data center's cooling load.

Server Internal Electrical Distribution

Inside the servers and other data processing equipment such as storage units, power supplies convert AC voltage (at 208/120V AC) to DC voltage for use by the digital processing electronics. Power supplies usually provide power factor correction as well as load isolation from the incoming power line for these sensitive electronic components. This conversion generally involves at least two stages. In addition, most modern microprocessors require very low voltages at fairly high currents, such as 1.1V at 100A. The precision of the voltage required, coupled with losses incurred by distributing high power at ~1V, and the high number of different rail voltages needed dictate that voltage regulation circuitry must be located directly next to the microprocessor. In order to

effectively realize this circuitry, most processors require that an intermediate DC voltage, such as 12V DC, be delivered to the processor/local regulator combination.

Thus, from the power supply, there can be up to six or more power conversion stages between facility power entry and the microprocessor.[14] Depending on processor and server loading, which can dynamically range from 30% to 100%, the efficiencies of power supplies will likely be lower at lower load levels and can significantly reduce overall system efficiency. In many cases, redundant power supplies are used to deliver this power, either on standby or in load-sharing configurations inside servers, both of which reduce individual power supply loading and efficiency.

Issues and Limitations of Current Data Center Designs

Due to the many levels of power conversion occurring in current data center facilities and within the IT equipment contained therein, most data centers experience significant electrical power losses in their facility systems' supply and distribution. This also includes losses in transformers, power line conditioners, UPS, line losses, etc.[15] Intel Corporation surveyed these losses at an actual data center, and the magnitudes of these losses for this facility are detailed in Figure 7 below.

Figure 7. Source and Magnitude of Data Center Losses

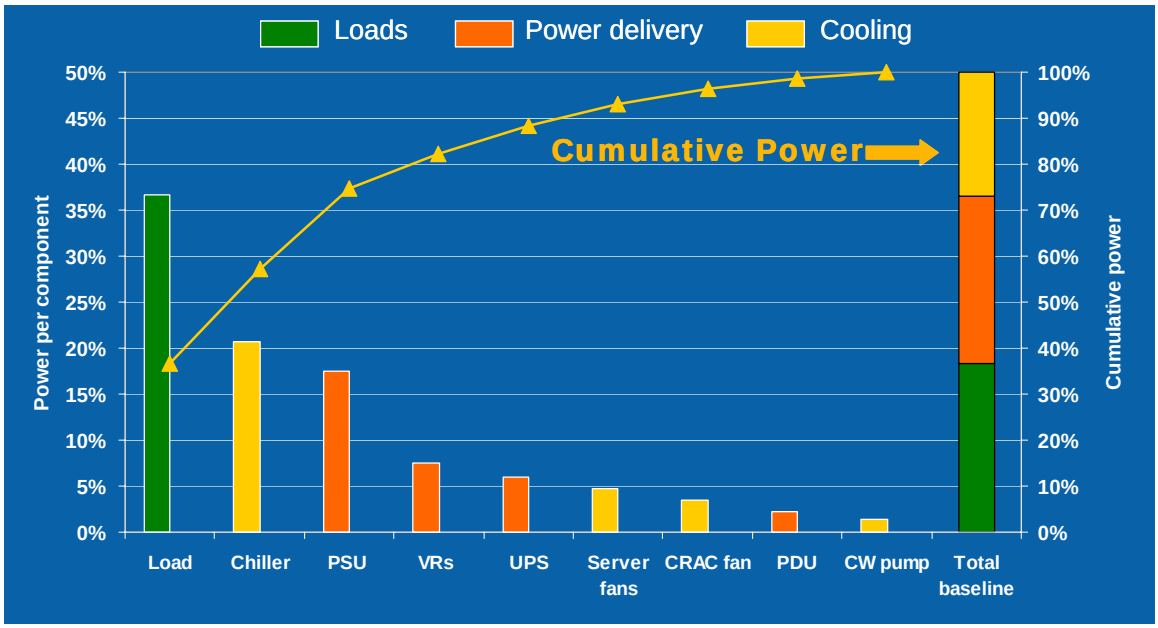

Source: Intel Corp.

The centralized battery plant utilized in facility-level UPS systems can be a compromise between what is required for successful system realization and component limitations. The actual battery run time required for acceptable system operation is often anywhere

[14] Example conversion steps: 480VAC to 208VAC, 208VAC to 400VDC, 400VDC to 208VAC, 208VAC to 400VDC, 400VDC to 12VDC, 12VDC to 1.1VDC.

[15] Note that power line losses are not covered in the scope of this report.

from 10 to 20 seconds to just a few minutes. The time represents the delay needed to switch to alternate utility power feeds, or to bring an auxiliary source of power on line (such as an engine generator). However, when batteries for centralized AC UPS systems are sized, the nature of the voltages required or battery type chosen can often result in hold up times well in excess of what is required, resulting in wasted energy storage and delivery capabilities. [16]

Inside the server units, the limiting factors on power conversion efficiency are voltage rating of semiconductors, along with their corresponding conduction losses.[17] In addition, power converter package size limitations, and the economics of power supply design and manufacturing can also play a part. All of these factors combine to ultimately limit available efficiency improvements for the power conversion processes. Power conversion efficiency for best-in-class computer grade AC-DC power supplies currently do not provide efficiency higher than about 90%. For servers, on-board non-isolated converters that provide final processor power conversion can range over 90% efficiency.

Figure 8. Conversion Losses and Their Effects on Data Center Cooling

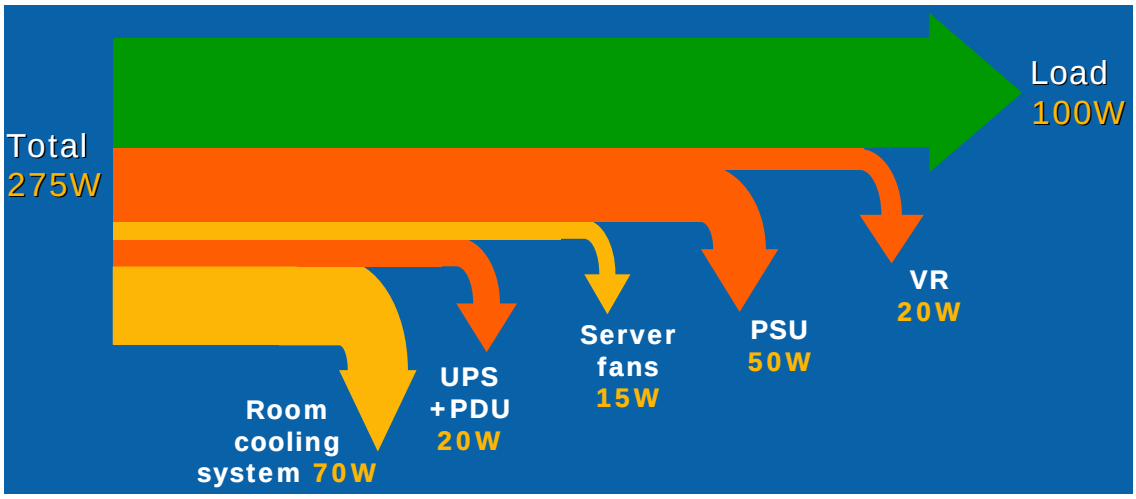

Source: Intel Corp.

The power losses due to inefficient power conversion from both outside and within IT equipment result in a large loss in useful electrical power, as well as directly increasing the energy required for HVAC systems to remove the heat produced in conditioned spaces. Figure 8 above details the conversion losses and the effects on cooling for a given representative computational load.

[16] Additionally, the lower efficiency presented by the off-line power supplies utilized by data center equipment produces an extra load on batteries that only goes into producing heat instead of power conversion.

[17] Note that power conversion efficiency is limited by device technologies used, but include switching and conduction losses of MOSFETs, ESR of capacitors and inductors, etc.

The above result indicates that for every watt of power utilized to process data, another 0.9W is required to support power conversion. In addition, for air-conditioned facility cooling, bench marking by LBNL suggests that another 0.6 to 1 watt (or more) of power will be required for each watt utilized to cool the power conversion equipment. While additional cooling power can seem insignificant at the individual microprocessor level, when overall data processing activities reach power usage levels on the order of 200kW, or more (for example, a large Internet 4 hub installation) this can be a considerable load.

These trends provide an impetus to the industry to support research on better cooling methods and more efficient equipment/components, and this is happening on many fronts. But to look at the compute loads and determine ways to increase their efficiency or otherwise decrease their heat density, will also help alleviate the problem. In fact, from a cost savings point of view, it is even better, since every watt saved in a server means one less watt of cooling required.[18]

2.5 Demonstration Project Objectives

There have been many responses from the different sectors of the datacenter market to cope with this trend in rising power demand.[19] The responses vary in their focus, ranging from increased processor efficiency with demand management and/or multiple processor cores at the processor level, to liquid cooling at the server or rack level, and various HVAC strategies at the facility level. These efforts also include the CEC PIER Program-supported research to inform the IT industry, including our prior research on the efficiencies of server power supplies and UPSs, and EPA's ENERGY STAR® Programs. However, the efforts to date tend to focus on minimizing losses (or improving the efficiency of) specific components within the AC distribution system rather than optimization of the whole data center power delivery system.

An optimal system might integrate the IT equipment with the facility in such a way as to minimize power conversions. For example, the individual power supplies in servers could be eliminated if the correct voltages of DC power could be supplied efficiently from a central system, or in the case of fuel cells, directly from the power source. One industry expert envisions the data center of the future similar to a computer in its case. Taking this idea a step further, the electrical system could be thought of as an integrated system from where it enters the data center to the ultimate end use. When viewed in this manner, optimized systems could be designed so as to optimize energy (distribution and conversion losses), reliability, power quality, and potentially provide additional benefits such as elimination of harmful harmonics.

---

[18] However, given a more efficient server, datacenter operators would be more likely to increase the number of servers for a fixed cooling load rather than reduce the overall electric bill, because the demand for computing power is also rising.

[19] According to the Uptime Institute, the trend in power density for servers, storage devices and communications equipment will continue to rise, reaching new highs every year. From "How to Plan, Justify, and Manage a Major Data Center Project."

The concept of a DC data center is not new – most telephone companies' central offices use 48 VDC, and there are a number of DC datacenter proponents.[20] There are also companies offering DC servers today.[21] However, the current solutions on the market do not offer an integrated approach in the same way that AC datacenters are being designed and built currently. Although the DC solution is not new, unique or exceptional, it is still very difficult for most people to envision such a drastic change from the current norm, and it will take some time for the market to adopt the changes. Therefore, a demonstration of this technology using existing equipment (to the extent possible) will go far towards elevating the advantages of such an approach. In addition, demonstration projects tend to provide a critical step to verify engineering expectations for performance, and build awareness and confidence in the new technologies and practices.

The objectives of this Demonstration Project are to develop a power delivery system that does not contain as many power conversion stages using existing equipment and vendors where possible. One way to accomplish this system is to use DC power distribution. This project implemented a power delivery system that distributes DC to the server racks. Server manufacturers can supply products with power supplies that are ready to operate on DC.[22] The system used a single rectification stage, thereby removing the conventional UPS, transformer, and the rectifier in the server's first stage power supply. The removal of these stages, and the energy losses associated with these them, will achieve some energy savings automatically.

With a standard AC system installed next to this DC system, server loads were connected and programmed to run identical routines. For this identical amount of computing work, the input power for whole system was measured and compared. Taking all of the considerations, and considering the available resources, the project team identified the following goals for the demonstration project:

1. Show that DC-powered server(s) and/or server rack exists in the same form factor or can be built and operated from existing components with minimal effort.
2. Show that DC-powered server(s) and/or server rack provides the same level of functionality and computing performance when compared to similarly configured and operating server(s) (and/or server rack) containing AC power supplies.
3. Measure and document any efficiency gains from the elimination of multiple conversion steps in the delivery of DC power to the server hardware.
4. Identify areas requiring further development or follow up investigations.

[20] Gross, P., and K. L. Godrich, "Total Integrated DC Datacenters." Paper presented at Intelec 2005.

[21] Humphreys, J. and J. Yang, "Server Innovations: Examining DC Power as an Alternative for Increasing Data Center Efficiency and Reliability." IDC White Paper sponsored by Rackable Systems Inc. August 2004.

[22] Servers used in the demonstration project were provided as "beta" equipment, which can be commercially available within six months with minimal certification effort

3. Project Chronology and Definition

As part of the server power supplies and UPS work, we identified a number of stakeholders in the latter part of 2005 and began dialogues with these industry representatives regarding a possible demonstration of an integrated DC power delivery system.

3.1 Timeline

Stakeholders Meeting – Fall 2005

The concept gained momentum through additional discussions and gained the CEC's support in the summer of 2005, an initial meeting for stakeholders was held in Chatsworth, CA in August of 2005, which was attended by over 20 industry representatives. The meeting was hosted by Pentadyne Power Corporation and Baldwin Technology, Inc.[23]

The meeting involved experts in data center design and operation, UPS vendor, controls manufacturers, technology integrators, instrumentation manufacturers, and component suppliers. No data center end users were present, but several were represented by the meeting participants. Several hours of lively discussion helped to identify potential industry partners, and brought out many issues and opportunities to address in the demonstration and eventual implementation.

During the discussion which followed, the results from prior investigations into UPS and power supply efficiency were presented and a number of considerations to be accounted for in a demonstration project were discussed. Barriers to implementing this technology were also introduced. Issues discussed include:

- All major server manufacturers currently have products that operate on 48V DC.[24]
- Commercially available equipment can be used for DC demonstration (no R&D required) including UL listed buss bar to supply the server equipment.
- The server manufacturers will always want to have control over the final conversion part of their equipment because of power quality/surge, safety, etc.

[23] The meeting also featured a demonstration of DC power distribution at the Pentadyne manufacturing facility. In this demonstration DC power was supplied at approximately 500 volts through commercially available components. The flywheel system maintains constant output at preset levels. The 500 VDC power is then converted to 48 VDC which can then be fed to the IT equipment. Depending upon the connected load, the flywheel will maintain this level for varying lengths of time. Baldwin Technologies supplied the DC delivery system, and Dranetz BMI supplied the monitoring instrumentation.

[24] Sun indicated at the meeting that it has swappable AC or DC power supplies.

- 6-7 kW per rack is the approximate limit for air cooling within a data center environment (although ASHRAE has shown an 18 kW rack in air (average today is 2-1/2 to 3 kW per rack).

- Data center user group sponsored by Emerson Network Power reports average of 60 W/square foot (this matches well with LBNL findings).[25]

- What should be the distribution voltage? 500V DC; 380V DC; 48V DC? Existing Sun power supplies can be changed out to run on 380V DC. IBM uses 350V DC. 380V DC was the consensus choice.

- Server manufacturers want their power source to be an integral part of the box – for safety, as well as to prevent others from plugging in.

- Barriers to DC power delivery included:
    - different voltages today in servers and other IT equipment
    - safety concern above 150V DC due to unfamiliarity
    - Above 600 volts - SELV classification
    - a new standard for voltage distribution may be needed

- Possible host sites were discussed

Kick-off meeting – April 2006

Discussions continued through the Fall of 2005 and into 2006, with additional manufacturers and industry representatives expressing interest and/or coming on board with commitment for participation. One outstanding issue crucial to the project implementation was finding suitable location(s) for the demonstration in California.

By Spring 2006, a suitable location for the demonstration project was found in Newark, CA, at Sun Microsystems' manufacturing facility in the San Francisco Area. With the location identified, this allowed the participant group to move forward with discussing the details for the demonstration. A meeting was held at this facility on April 3 and 4, 2006 for all interested parties in the demonstration. This meeting was intended to finalize the equipment list and scheduling for the tentatively named "DC Powering Architecture for Data Centers Demonstration."

A number of agreements were reached at this meeting, including:

> **Demonstration Configurations**: There was general agreement that both a rack-level and a facility-level demonstration could be accomplished, and needed for completeness and credibility.[26] Both could take place at the current Sun location

[25] The group thought this includes HVAC requirements.

[26] "Rack-level" is the term used to describe DC conversion at the server rack. This approach converts the facility's supplied AC into high-voltage DC via a rack-mounted rectifier unit. This approach concentrates the DC conversion into one unit/location in the rack, removing the AC to DC conversion function from the

if needed equipment could be found. For an "apples to apples" comparison, three different configurations were needed (one is needed as reference):

1. Configuration 1: Current data center typical set-up delivering 208/120V AC input to AC-powered servers (Reference Configuration).
2. Configuration 2: DC conversion/distribution at the rack level (Rack-Level Demo), using a rectifier unit to convert 208/120V AC at the rack, and delivering high-voltage DC to DC-powered servers.
3. Configuration 3: DC conversion/distribution at the building/data center level (Facility-Level Demo), converting 480V AC to 380V DC and delivering this directly to the DC-powered server units in the rack.

Equipment Set: Configurations 1 and 2 discussed above required additional equipment (UPS back up, PDU for power delivery) to accurately reflect the conditions at a typical data center, all three may also require additional loads to simulate operating conditions. The AC servers used would be matched to equivalent DC models and the same applications would run on both sets. Configurations 1 and 2 would share an AC buss way, while Configuration 3 would have its own DC buss way.

Testing and Measurements: The group agreed on test points and metrics, with emphasis on measuring the efficiency of the configurations – there would be no direct comparison of server equipment performance.

DC Input Voltage: Due to compatibility with existing equipment and devices, the group settled on 380 V DC for the high voltage DC input.[27] The test DC voltage would be set at 380 ± 5%, with 3% peak to peak for the ripple, which included high frequency and line frequency ripple components. However, the group agreed that this demonstration is about measurements rather than standard-setting. Other related points on DC input voltage were brought up by participants:

- Input DC voltage should be maintained between 300 volts and 400 volts for use with batteries and +/- 5% for use with flywheels.
- When using batteries, additional DC to DC conversion may be required.

DC Power Supply Efficiencies: Manufacturers noted that gains in the efficiencies of DC power supplies will probably be in the order of 1% to 3%. The server manufacturers already offer efficient DC power supplies at around 90%,

---

servers themselves. This approach can be attractive to operators who want to use DC powered servers, but not able to convert their facility exclusively to DC. "Facility-level" is the term used to describe DC conversion at the facility or data center level.

[27] HP and IBM were not available to present alternatives at this meeting (Intel noted that IBM currently uses 380 VDC in their high-end servers).

but end users can buy the less efficient (less than 80%) and less expensive power supplies (~10% less in cost). Thus for DC servers, the increase in on-board power supply efficiencies will come from the elimination of the computer on board power factor correction circuits and rectifiers. (Most of the efficiency gains were expected to come from the elimination of the AC to DC stages of the rack- or facility-level conversion).

Safety: The issue of high-voltage DC safety was brought up. The group agreed that, this was adequately addressed by the use of UL-rated equipment for high voltage DC, and other concerns were evaluated through a safety committee. For the future, we will research DC safety standards in use for other applications, including standards in use for the EU, military, and transit industries.

Other Related Technical Issues: The group also brought up various other issues, some of which were deferred:

- Construction Costs
- Reliability
- Equipment costs
- Grounding
- Environment
- Final Project Report

Equipment Demonstration Sessions: Sun Microsystems suggested that an "Open House" be held on 24 May 2006 to demonstrate the assembled equipment, with regular occurrences after that.

Code of Conduct: As the project was focused only on the demonstration aspect, participants also agreed to a "code of conduct" for cooperation and data sharing to minimize conflicts among competitors. A copy of this agreement is included in Appendix C.

Equipment assembly – May 2006

All participants at the April 06 meeting agreed that the months of April and May should be used for equipment configuration, definition, acquisition testing, and shipping, with the target date to finalize all equipment and transportation to the demonstration location was selected as 1 May 2006. On-site equipment assembly and verification took place in the first 2 weeks of May (1 May to 12 May), and testing/measurements took place on the 3rd week (15 May to 19 May).

Figure 9. Assembly at the demonstration site

In actuality, the gathering of equipment and shipping it to the demonstration facility took quite a bit longer, due to the complexity of the equipment, shipping times, as well as the need for approval of the proposed additional efforts. While a number of large pieces arrived in April and early May, it was not until the last weeks of the month that the assembly process actually started. Equipment assembly, integration, and testing took place in the last two weeks of May, and participants gathered for the first power-up on the first week of June.

Initial "Open House" – June 7, 2006

The first "Open House" was planned for May 24, and open to the press. However, a number of delays caused the date to slip, and with equipment "shakedown" still under way in the first week of June, the initial Open House was a more low-key event, with the press event rescheduled for two weeks later (June 21). Invitations were made only to interested industry representatives and colleagues.

Keeping the event low-key allowed the coordinating team of LBNL, Ecos, and EPRI Solutions to test presentation materials, go over equipment procedures, and trouble shoot any remaining issues prior to opening up the demonstration to a wider audience. The initial Open House also allowed the coordinating team to finalize the logistics and format of these events, which were used for the remainder of the summer.

Press Event: June 21, 2006

In order to achieve maximum media impact for the demonstration, one specific Open House was designated as a Press Event, to ensure that the project receives widespread industry attention. Participating industry representatives worked with their organization's PR staff to identify possible press contacts for invitation to the event. Both trade and popular publications, as well as electronic, television and radio were targeted. LBNL and the CEC and various industry partners' public relations representatives also participated in the process. LBNL's PR office prepared both the press invitation to the event and the press release for the day of the event, with input from other participants' PR representatives.[28]

Open House Events: July 12, 26; Aug 9, 16, 18

Participants agreed upon an every-other week Open House schedule, occurring on Wednesdays for the remaining of the summer, with adjustments for the July 4 holiday

[28] Electronic copies of the invitation and press release can be found at: http://hightech.lbl.gov/dc-powering/media.html.

and other competing events. The end date of the Open House was to coincide with the ending of the Phase 1 work for the LBNL team. Each of the Open House's two sessions (10 AM and 2 PM) averaged between 10 and 20 attendees from the data center industry and other sectors.

Due to unforeseen circumstances, the last Open House was changed to August 16, instead of August 23 as originally planned.[29] A last minute request from the US Armed Forces and a number of contractors resulted in a final, abbreviated Open House on the morning of August 18. After this the equipment set up was completely dismantled and returned to manufacturers or suppliers.

3.2 Participants

The demonstration project started with about 20 participants representing experts in data center design and operation, UPS vendors, controls manufacturers, technology integrators, instrumentation manufacturers, and component suppliers. As the project progressed and became known through a number of venues, the list of participants continued to grow throughout the duration of the project, new participants included data center end users, and a number of cellular carriers. In fact, industry representatives attending the "Open Houses" heeded our call for participation and a number quickly responded with additional equipment and/or in kind services from their respective industry sectors. Appendix A contains a full listing of the participants.

Participants of the Demonstration Project generally belonged in two groups:

- Contributors: Project "contributor" either contributed equipment, in-kind services, or both towards the project. A list of contributors is included below.

[29] The Newark demonstration facility was sold over the summer to a biotech firm, which was assumed control of the facility in September.

Alindeska Electrical Contractors, LLC
American Power Conversion, Inc. (APC)
Baldwin Technologies, Inc.
Cisco Systems
Cupertino Electric, Inc.
Dranetz-BMI
Emerson Network Power
Industrial Electric Manufacturing (IEM)
Intel Corporation
Nextek Power Systems
Panduit Corp.
Pentadyne Power Corporation
Rosendin Electric Inc.
SatCon Power Systems
Solara
Square D/Schneider Electric
Sun Microsystems
TDI Power
UNIVERSAL Electric Corp

- Stakeholders: Project "stakeholders" included industry experts and interested parties who regularly participated in discussions, provided feedback on approaches and other issues. Stakeholders did not contribute equipment to the project, but provided valuable insight into the process. A list of stakeholders and their industry sector is included below.

380voltsdc.com
CCG Facility Integration
Cingular Wireless
Dupont Fabros
EDG2, Inc.
EYP Mission Critical
Gannett
Hewlett Packard
Morrison Hershfield Corporation
NTT Facilities, Inc.
RTKL
SBC Global
Transistor Devices, Inc.
Verizon

3.3 Power Delivery: Definition

As discussed, the project participants helped to define the three configurations for the demonstration project. They are:

- AC Reference Configuration: This configuration is needed to simulate current data center typical set-up, delivering 208/120V AC input to AC-powered servers, and to be used as a reference to compare conversion efficiency.

- Facility-Level DC Configuration: This configuration is needed as the proof of concept – the ability to deliver high-voltage DC throughout the facility. This configuration handles the DC conversion/distribution at the building/data center level, converting 408V AC to 380V DC and delivering this directly to the DC-powered server units in the rack.

- Rack-Level DC Configuration: This configuration is needed to provide a possible migration option for AC data centers operators wishing to use DC equipment without facility-wide DC power distribution. This configuration accomplishes DC conversion/distribution at the rack level, using a rectifier unit to convert 208/120V AC at the rack, and delivering 380V DC to DC-powered servers.

3.4 Equipment Definition: AC Reference System

As detailed above, the AC reference configuration was intended to represent the typical set-up found in today's data center, taking 480 VAC from the utility feed, conditioning and stepping it down to 208/120V AC for input to AC-powered servers. In addition, due to uncertainty regarding manufacturers' abilities to provide sufficient servers to constitute a significant electrical load on the power delivery system, an AC load bank was added.[30] As configured, the set up has the ability to deliver from 15 kW up to 75 kW of electrical power to the server racks, while the projected actual load available for the demonstration was closer to 2 – 4 kW, or about 4% to 8% of the available system power. The load bank would allow the simulation of 30 kW to 40 kW power demand on the system, more accurately reflecting the loading of current and future server racks in data centers. The simplified AC power delivery system is shown below.

Figure 10. AC Reference Configuration

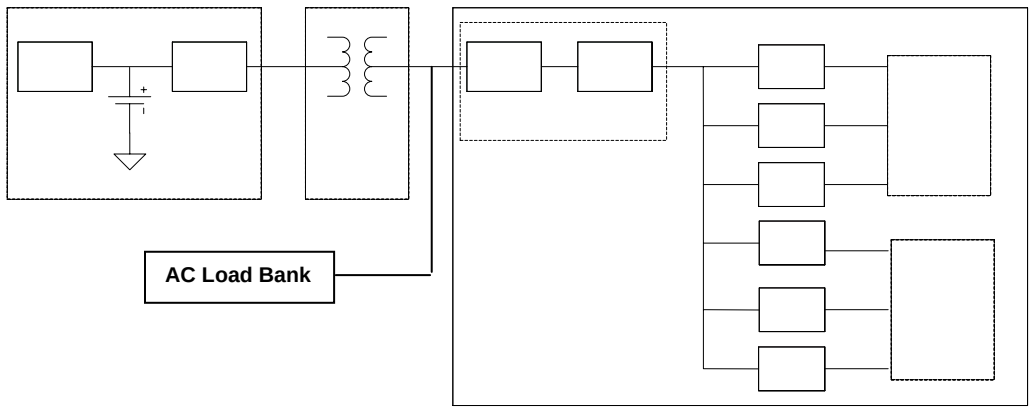

In this diagram the UPS is a double-conversion UPS (typical), preceded by a transformer in a power distribution unit (PDU). Inside the server is an AC/DC conversion at the input to the server power supply. There are still many DC/DC converters at lower voltage levels, known as voltage regulator modules, or VRMs, but those are not the focus of this project. So to implement this system, the complete equipment is listed below:

- Double conversion UPS
- Transformer (for PDU)
- Servers
- Supplemental load banks (to load the UPS to a typical level)
- 600 V rated busway

[30] This was a commercially available product, leased by the Demonstration Project expressly for this purpose.

- Various circuit breakers
- Wire and other hardware

3.5 Equipment Definition: Facility-Level

The facility-level DC configuration was intended to replace the typical AC set-up found in today's data center, taking 480 VAC from the utility feed, convert it to DC voltage, and deliver the DC voltage to the racks for input to DC-powered servers. This configuration is needed as the proof of concept – demonstrating the ability to deliver high-voltage DC throughout a data center facility. As in the AC set up, the uncertainty regarding manufacturers' abilities to provide sufficient DC servers to constitute a significant electrical load on the power delivery system, a DC load bank was added.[31] The DC load bank would also allow the simulation of 15 kW to 40 kW power demand on the system, more accurately reflecting the loading of current and future server racks in data centers.

The next diagram shows the arrangement for DC power delivery to the data center loads.

Figure 11. Facility-Level DC Configuration

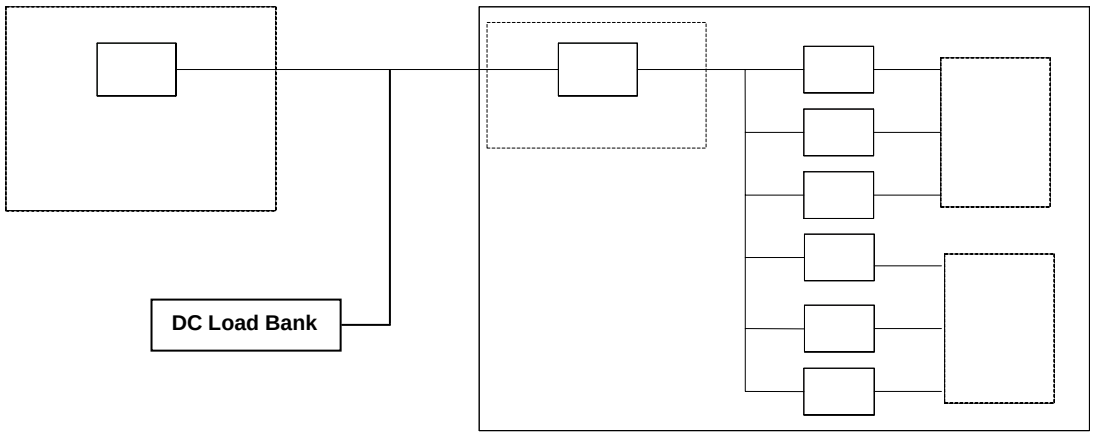

In this layout the only power conversion before the server takes place through one rectifier. Inside the server power supply, only one DC/DC conversion occurs. In addition to the reduction in losses, this arrangement is inherently simpler, which will should lead to lower equipment cost, lower installation cost, and higher system reliability due to lower parts count. The necessary equipment for implementation is listed below:

- System level rectifier
- Servers, modified to take 380 VDC

[31] This was a custom built DC load bank, designed and constructed specifically for the use of this Demonstration Project.

- Supplemental load banks (to load the UPS to a typical level)
- 600 V rated busway
- Various DC circuit breakers
- Wire and other hardware

3.6 Equipment Definition: Rack-Level

The rack level demonstration was designed to show that 380 VDC could be supplied at the rack level and distributed throughout the rack. Instead of using a rectifier at each power supply, only one rectifier would be necessary to convert AC to DC for the entire rack. The advantage for this approach would be to change the heat removal requirements – it may be more economical to remove the heat at the rack level rather than from the servers themselves. Measurement of cooling effectiveness was not in the scope of this project; rather, this arrangement was intended to show feasibility of design. In addition, this configuration is needed to demonstrate that a possible migration option is available for AC data centers operators wishing to use DC equipment without facility-wide DC power distribution.

Figure 12. Rack-Level DC Configuration

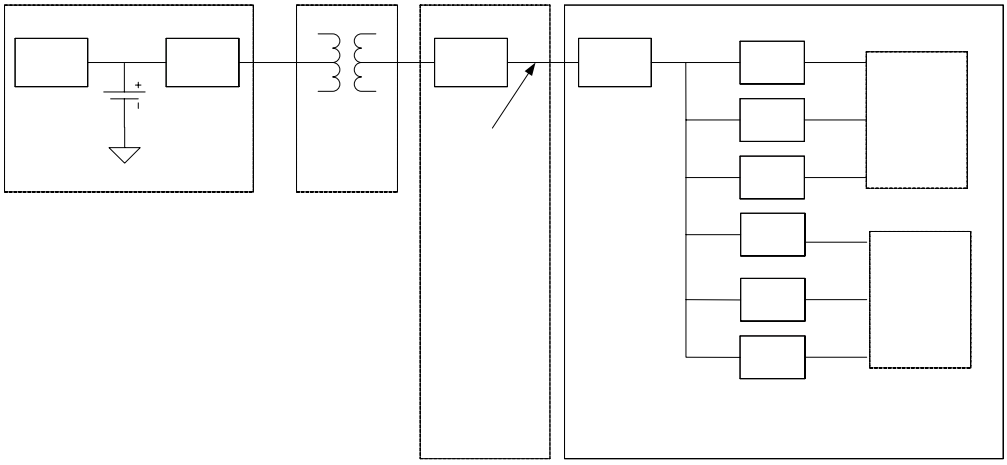

So to implement this system, the complete equipment is listed below:

- Double conversion UPS
- Transformer (for PDU)
- Rack level rectifier
- Servers, modified to take 380 VDC
- Supplemental load banks (to load the UPS to a typical level)
- 600 V rated busway

- Various circuit breakers
- Wire and other hardware

3.7 Equipment Definition: Monitoring

The monitoring system used has high bandwidth data capture capability using voltage probes and current transformers (CTs). The system uses a central node called an info node and multiple individual modules called data nodes that reside close to the point of measurement. The system has the capability to capture waveforms during power line disturbances, and also to capture steady state data. In this case, the system was used to capture steady state voltage, current and power data. Figure 13 below illustrates the layout used by this system. In addition to providing real-time data measurement to the system operators, this system has the ability to provide "live" data display via the internet. In fact, measurements were made available via the LBNL website for the duration of the demonstration project.

Figure 13. Diagram of the Monitoring System

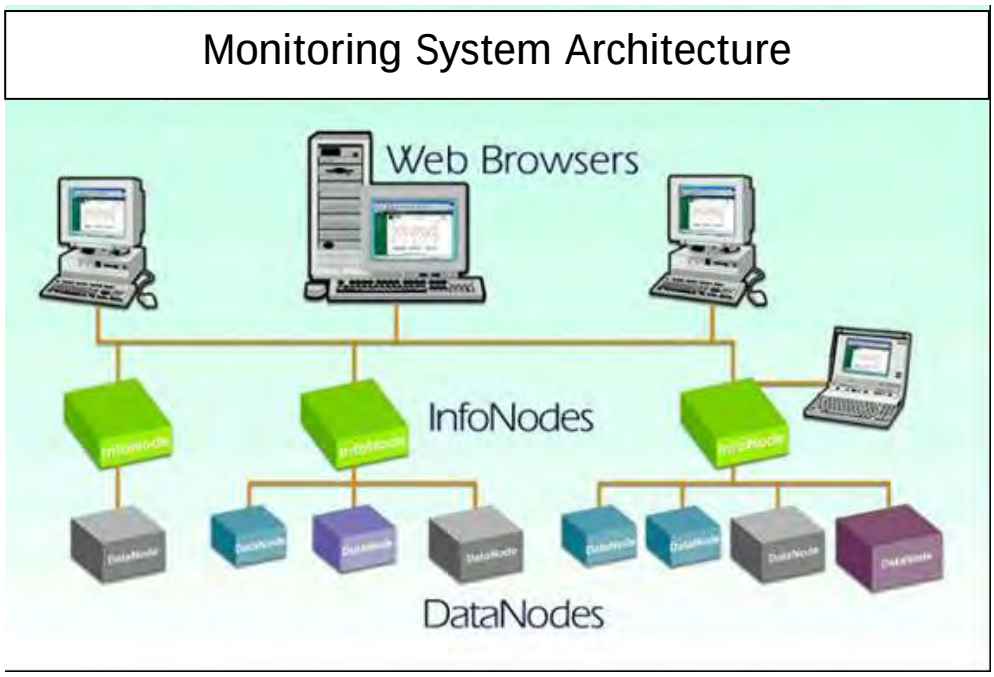

The placement of the monitors can be seen in the one-line diagrams below (a full-page layout is referenced in Appendix B).

Figure 14. AC Distribution Layout and Monitoring Points

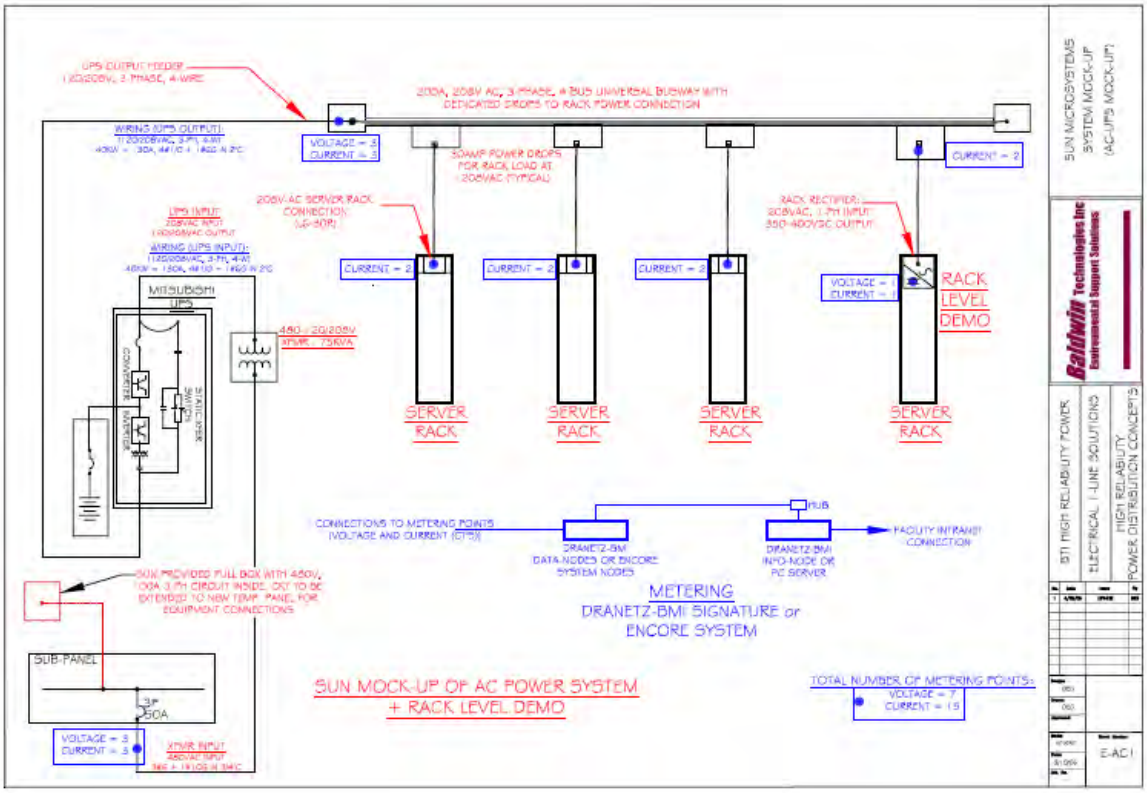

Source: Baldwin Technologies, Inc.

The above diagram shows the AC Reference System and its monitoring points, along with the USP system's batteries and transformer. The system distributes 208 VAC to the server racks through 30 A breakers. Also shown is the rack level demo. The Monitoring system is shown as blue dots with callouts.

There are 7 voltages and 15 currents measured.

Figure 15. DC Distribution Layout and Monitoring Points

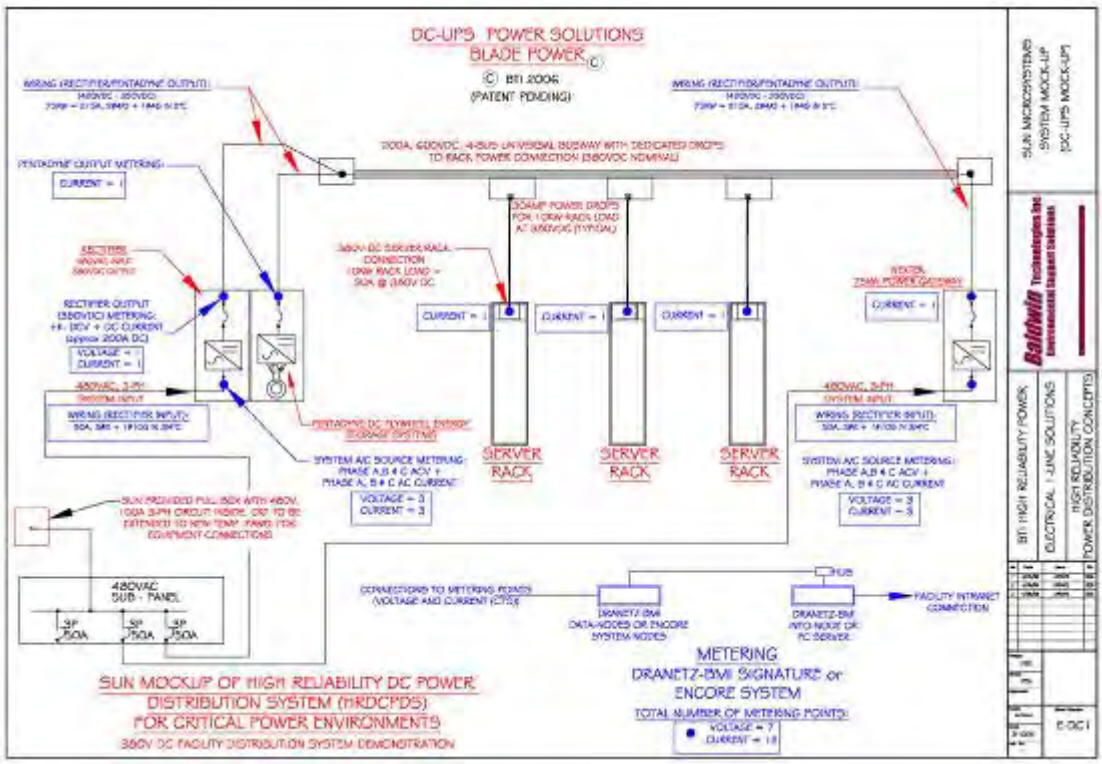

Source: Baldwin Technologies, Inc.

The diagram above shows the monitoring points for the DC Demonstration set up. Note that two main rectifiers are shown in this diagram. Only one is necessary, but a second one provides redundancy and a second vendor was able to participate. A flywheel is used for energy storage (ride through for interruptions), but batteries can also be used with DC distribution. The DC bus is energized at 380 VDC, with individual racks powered through 30 A breakers on the buss way.

The monitoring points are shown on this diagram as blue dots with callouts. There are 7 voltages and 12 currents shown.

3.8 Other DC Systems Demonstrated (48V DC)

The interest generated by the project helped to bring about other manufacturers of DC equipment, notably manufacturers of communication and support equipment using 48 V DC. These manufacturers provided additional computing, telecommunication equipment and racks for the use of the demonstration project. The presence of these manufacturers and their equipment provided another effective reminder that DC power distribution is not a new concept, and has been safely and effectively used in telecommunication and data networks. In addition, it also demonstrated that DC and AC power delivery system can co-locate in a data center facility.

The project coordinating team had initially also investigated the issue of 48V DC power delivery systems and related efficiencies, using available data from manufacturers and other industry experts.[32] This supplemental work was intended to inform the team on selecting an appropriate DC voltage for the demonstration. However, input from the participants, especially the server manufacturers, helped narrow down the demonstration voltage to 380 V DC. This obviated the need for an in-depth comparison of different DC voltages.

These supplemental DC systems were included in the overall demonstration for completeness only. Due to limited resources, monitoring equipment constraints, timeline, and priority for implementing the rack- and facility-level demonstration configurations as well as obtaining measurement, these 48 V DC systems were not monitored or their efficiencies measured. Without comparable AC equipment, it was not possible to carry out the same types of comparison that was being performed on the other configurations.

[32] Gross, P and Godrich, K. L. "Total DC Integrated Data Centers." Paper presented at 2005 Intelec.

4. Implementation Issues

As with any prototyping projects, where discrete equipment and products from different manufacturers have to work together for the first time, there were a number of challenges. In discussion regarding implementation within the group, we had agreed to divide the configurations into two distinct areas: outside of the server rack, and inside the server rack. This allowed the server manufacturers to focus on the DC conversion and adaptation for their products to be used in the server racks, while the facility level team focused on integrating the power delivery infrastructure.

Actual implementation and integration of the demonstration configurations were accomplished in an astonishingly short time, thanks in part to the commitment and enthusiasm of the participants. In fact, a number of industry "veterans" remarked that this Demonstration Project was one of the quickest ones that they had experienced. Although everything came together well, quickly, and every effort was made to use available, off the shelf components, there were a number of issues that we encountered that were not entirely solvable with existing equipment, or needed further investigation. They are outlined below.

4.1 Commercially available equipment/missing components

One of the stated objectives of the project was to accomplish the demonstration with commercially available products to the extent possible, especially with the power delivery infrastructure. The project was mostly successful in this regard with the following exceptions:

Within the server racks, we identified a number of components that are not yet commercially available:

1. Servers: There currently are no servers commercially available for 380V DC operation.[33]

2. Power cord/Power plug or connector: Although there are commercially available power cord and plug/connectors for 380V DC. To enable widespread adoption of this technology, the industry will need to standardize on only one type of connector (similar to an AC cord).[34]

3. Power "strip": There are there no 380 V DC powers "strips" for power distribution within a server rack. Similar to the power connector, it is desirable for the industry to standardize.

[33] Equipment used is not yet safety and EMI certified.

[34] We were able to locate a DC connector after the implementation had taken place, but it is not a universally accepted product. Note that the demonstration project used standard AC connectors and power plugs, but this may have serious implications in a mixed voltage environment, such as hot swapping servers.

Outside of the server racks, we also encountered a number of issues related to lack of commercially available equipment, including

Figure 16. DC Power Drops

1. Power drops/connectors: There are not yet specifically designated connectors for powering server racks. (The team used standard Hubbell connectors and standard APC equipment racks).

2. Circuit breakers: The availability on the market of DC circuit breakers specifically for this application was scarce. In fact, UL listed breakers were one of the more difficult items to obtain for integration.

4.2 Grounding, Bonding and Protection

Grounding, bonding and protection are issues with important implications for the adoption of DC power systems. These received significant discussion at the Open House sessions, and warranted additional research by the team. We identified a number of documents suitable for DC application (although not specifically for 380V DC), which are summarized in Table 3 below.

These industry documents are for the most part accounted for and harmonized to the fullest extent practicable in the IEEE Emerald Book. Formally titled: *IEEE Standard. 1100-2005 – Recommended practice for Powering and Grounding Electronic Equipment*, the Emerald Book provides a useful general reference in terms of the grounding and bonding recommended practices for facilities that simultaneously contain both AC power systems and DC power systems.

The Emerald Book integrates many of the traditional telecommunications recommendations and discusses how to integrate the AC and DC power systems to accomplish the important safety and performance objectives of each. The key concepts related to grounding and bonding have to do with the connection to the buildings grounding electrode system which the Emerald Book refers to as the "common bonding network and when to use a single point of connection to the grounding electrode system to form an "isolated bonding network." These objectives and the concepts employed are detailed in the following sections along with some historical background.

Much of the guidance on grounding and bonding DC power systems for telecommunications and information technology equipment is rooted in the traditional telephone (telecommunications) utility (regulated) industry. The basis of this guidance is supported in IEEE Std. 1100-2005 for the commercial (deregulated) industry - with some modifications made to meet requirements of the commercial market segment.

Table 3: Industry Documents Associated with DC Applications

ANSI	T1.311-1998	DC Power Systems—Telecommunications Environment Protection
ANSI	T1.315-1994	Voltage Levels for DC-Powered Equipment Used in the Telecommunications Environment
ANSI	T1.330-1997 (R2002)	Valve-Regulated Lead-Acid Batteries Used in the Telecommunications Environment
ANSI	T1.333-2001	Grounding and Bonding of Telecommunications Equipment
ANSI	T1/TIA J-STD-607-A-2002	Commercial Building Grounding (Earthing) and Bonding Requirements for Telecommunications
IEEE	Std. 1100-2005	IEEE Recommended Practice for Powering and Grounding Electronic Equipment
NFPA	70-2005	National Electrical Code (NEC®)
NFPA	75-2003	Standard for the Protection of Information Technology Equipment
Telcordia	GR-1089-CORE 2006	Electromagnetic Compatibility and Electrical Safety - Generic Criteria for Network Telecommunications Equipment
Telcordia	GR-295-CORE 2004	Isolated and Mesh Bonding Networks: Definition and Application to Telephone Central Offices
T1A	TIA/T1 J-STD-607-A-2002	Commercial Building Grounding (Earthing) and Bonding Requirements for Telecommunications
T1A	TIA 942	Telecommunications Infrastructure Standard for Data Centers
UL	UL 60950-1 - Edition 1 04/03	Information Technology Equipment - Safety - Part 1: General Requirements

Source: Panduit Corp.

Historically, telecom is DC powered and information technology equipment has been AC powered. Therefore, the grounding and bonding standards for these different types of power systems evolved differently due to the DC being predominantly utilized in a regulated environment considered "under the exclusive control of the utility" (NFPA 70). As such, DC power system operating voltages are affected by factors such as:

1. Battery technology (vented or valve-regulated)
2. Rectifier output regulation in maintaining constant voltage under dynamic loads
3. Voltage drops in DC power conductors
4. Operating voltage limits of various connected loads
5. Derating for environmental conditions such as altitude

4.2.1 Telecom DC Systems Grounding Example (up to 160V DC)

Generally, telecom DC power systems date back to the first telephone systems where DC was used for talk battery, signaling circuits and for operating switching and control relays. Large centralized (bulk) DC power plants (systems) had primary components such as rectifiers, power-boards, primary and secondary distribution feeders, and fuse bays. These were built with reliability in mind, and space and efficiency were secondary

considerations. The DC power system was grounded to earth (often more than once). The grounded conductor was termed the "Return."

Modern DC power systems are more compact, use much smaller footprint components, and are more efficient. Today, a small centralized DC power system can be contained in a single rack. For standards compliance purposes, this equipment is generally classified as information technology equipment (ITE). For the purposes of evaluating DC powered ITE, robustness of the system is a key consideration. Accordingly, the design of grounding, bonding and protection of the installation is based on specific system considerations.[35]

Generally, the modern centralized DC power system is designed to operate as a single-point grounded system. An equipment-grounding (bonding) conductor is installed as part of the distribution circuit to ground any metal parts of the ITE and to clear any ground faults by facilitating the timely operation of the upstream overcurrent protection device. The system grounding (electrode) conductor is termed the DCG and is connected to the Return near the DC power source. The Return is only grounded once.

This arrangement is extremely similar to grounding and bonding an AC power system per NFPA 70-2005 – with some possible variations allowed by UL60950.

Figure 17: Similarity of Recommended Grounding for AC and DC Power Systems and Load Equipment

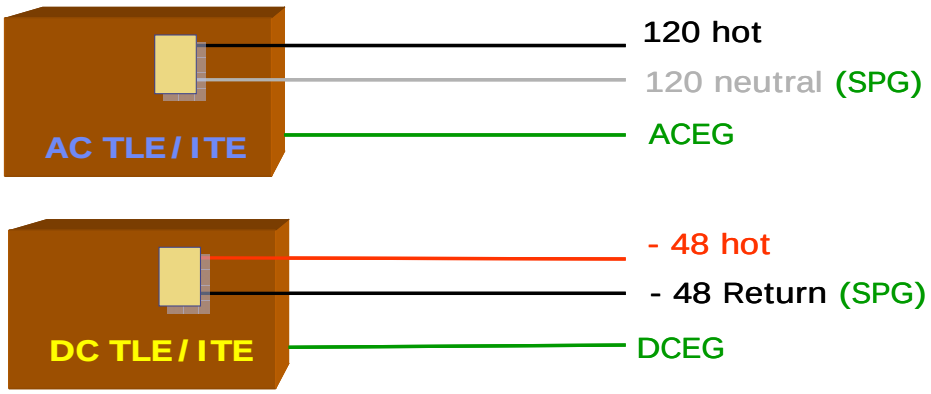

Source: Panduit Corp.

[35] For example, these considerations include:
1. Is the ITE suitably robust (per Telcordia GR-1089-2006) to operate in a Common Bonding Network (CBN)? Or must it be operated in an Isolated Bonding Network (IBN)?
Note: The significant role of the IBN topology is to isolate the ITE from currents expected flowing through the CBN – especially lightning. Availability requirements for ITE in a data center may not need to be specified to the "high-nines" required for a telecommunications service provider (TSP) – such as at a telecommunications central office. The availability specification will determine if an IBN may be appropriate for the data center – or a portion of the data center.
3.Is the DC power supply dedicated to a single type of equipment bonding network (CBN or IBN) or is it to be a shared resource?
4.Where is the planned location for the DC power supply relative to the location of the ITE?
Is any of the ITE required to integrate the Return and DC equipment grounding conductor (DCEG) at the DC power input?

Figure 18. DC Power System Showing Single-Point Grounded Return[36]

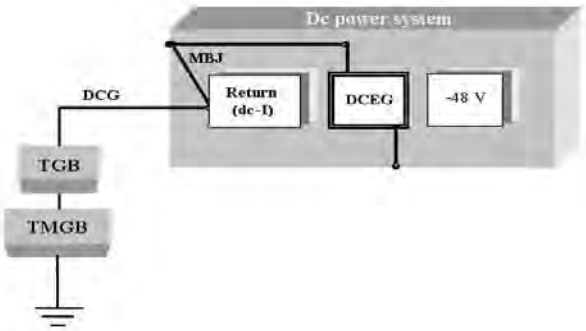

Source: William Bush

Figure 19. Grounding of DC Power Input for IT Equipment[37]

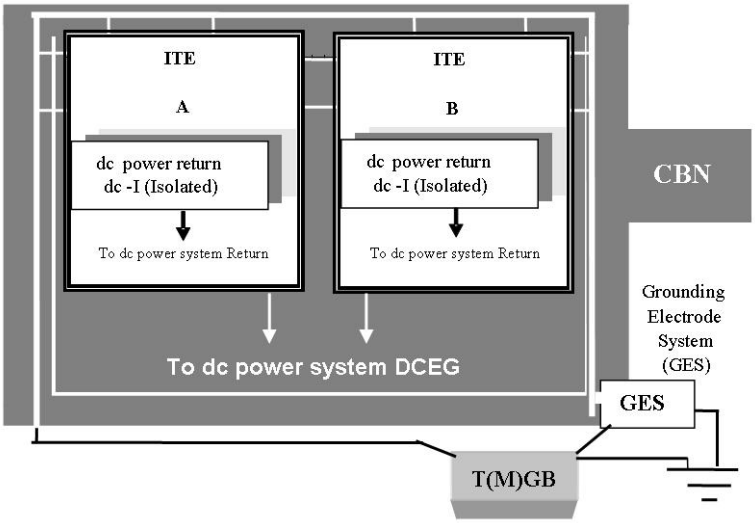

Source: William Bush

[36] SPG = Single point ground; ACEG = AC equipment grounding conductor; DCEG = DC equipment grounding conductor; TLE/ITE = Telecommunications Legacy Equipment / Information Technology Equipment

[37] DCEG = DC Equipment Grounding Conductor; TGB = Telecommunications Grounding Busbar; TMGB = Telecommunications Main Grounding Busbar; MBJ = Main Bonding Jumper; DCG = DC System Grounding Conductor.

Figure 20. Illustration of the Common Bonding Network (CBN)

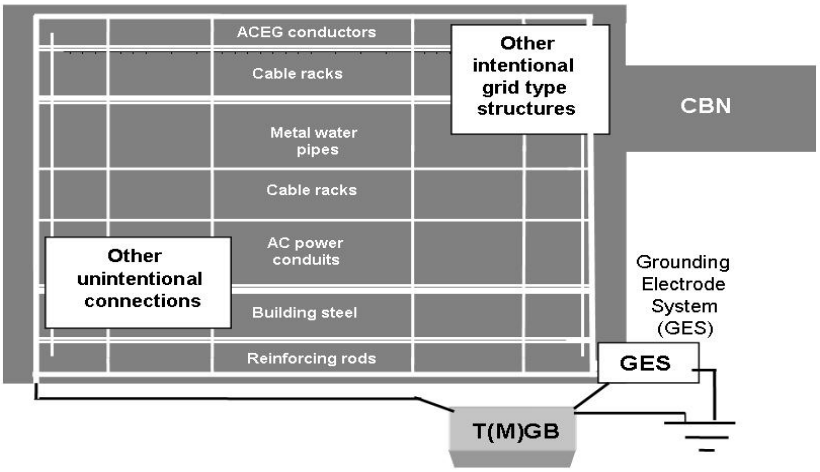

Source: William Bush

Figure 21. Illustration of the Isolated (insulated) Bonding Network

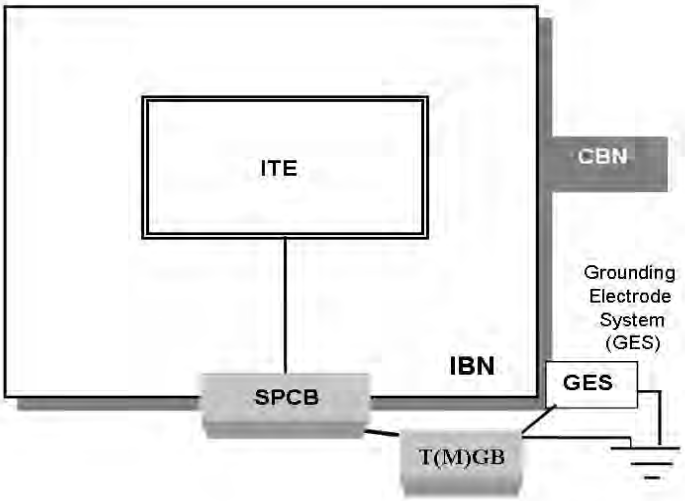

Source: William Bush

4.2.2 Additional considerations

Porting the DC power system telecom utility practices to the non-regulated environment requires additional considerations, they include:

1. The highest DC voltage covered by the telephone/telecom/ITE industry is 160 (ANSI T1.311). The utilized 380V DC (or higher) utilized by the demonstration project is essentially new territory and will require some additional safety and performance investigation. However, certain established principles for DC below are not expected to change.

    a. Over-current protection devices (OPD) and disconnect devices for DC power systems will need further investigation for the greater than 160V DC systems. AC rated devices do not automatically port over to the DC power system at the same rated voltage and current. The established 2:1 protection coordination scheme for DC fuses is not readily applicable since data centers typically utilize circuit breaker technology.

    b. Transients for other than 48V DC systems are not well described – if at all.

    c. Battery disconnects (where batteries are used as standby resource) considerations: 1 pole or 2 poles, and their interrupt ratings

    d. Conductor sizing is not a 100% straightforward utilization of the several Tables usually utilized from NFPA 70-2005 for AC conductors.

2. The rectifier technology is assumed to be switched mode power supply (SMPS) which may involve noise control via power line filters and additional audible noise from cooling fans.

3. Rectifier operation modes at the higher voltages such as 380V DC may need verification for ac input power factor correction, load sharing, paralleling, voltage sensing, current limitation, etc.

4. The DC power distribution topology is historically buss or cabling. The buss version described in this project will likely not be the only topology utilized. Further, there may also be a need for an under-floor topology. Distribution (such as rigid busbar) to withstand under extreme fault current conditions will likely to vary considerably from that for a telecom 48V DC system.

5. Design parameter for voltage drop from standby resource (such as a battery plant) to the load equipment. For telecom, the parameter is typically 1 V per 24 Volts of the supply. 48 V supply = 2 V drop.

6. Battery sizing: For a 380V DC battery plant, the size of the battery rack (if metallic) grounding conductor needs to be considered.

7. Centralized DC power systems are restrained to use in a Restricted Access Area (RAA) by UL 60950. Note: The data center is considered a RAA.

8. Grounding: Should the 380V DC system be operated as a positive grounded system (48V DC telecom) or negative grounded system (24V DC telecom).

9. The location of attachment of the DC system grounding electrode conductor (DCG) to the Centralized DC power system is allowed to occur forward from the source toward the telecom load equipment per UL60950.
10. The DC power system equipment grounding conductor (DCEG) is allowed to be routed separate from the DC supply and Return circuit conductors per UL 60950.
11. The DCEG is permitted to be bonded to the Return at the load equipment per UL 60950.
12. Considerations 9-11 reflect telephone company practices accommodated by first UL 1459 and then UL 60950.
13. The CDCPS historically can serve two recognized equipment bonding topologies.
    a. Common Bonding Network (CBN)
    b. Isolated (Insulated) Bonding Network (IBN)
14. The CBN and IBN equipment bonding topologies were described in IEEE Std. 1100-1999. TIA 942 also notes the CBN topology directly and the IBN through reference to IEEE Std. 1100-2005.
15. Based upon these considerations, the prudent approach is to utilize IEEE Std. 1100-2005 as the base document for grounding and bonding the DC power system in a data center.
    a. Essentially mirrors topology for an AC power system
    b. Single-point grounding of the DC power system at the source location
    c. A co-routed DC equipment grounding conductor with the circuit wiring (supply such as -48 V, Return)
    d. Bonding of DC equipment grounding conductor to the "Return" at the load equipment is prohibited
    e. Fully controlled paths for direct current

4.2.3 Supplemental grounding and bonding

For the data center, a telecommunications grounding and bonding infrastructure in accordance with J-STD-607-A, TIA 942 and IEEE Std. 1100-2005 is expected.[38] This infrastructure is bonded to the electrical power grounding electrode system, to building steel (where accessible) and to the serving AC power panel equipment ground at each floor. Grounding needed for the data center equipment(s) is obtained from the expected ground bar on that floor (such as a Telecommunications Ground Bar – TGB). Note that this grounding and bonding infrastructure is not the same physical structure as the grounding infrastructure that might be placed for the electrical power system.

[38] In no situation should a totally separate grounding system be deployed as this will bring to bear both safety and performance problems.

Figure 22. Illustration of Telecommunications Grounding and Bonding Infrastructure

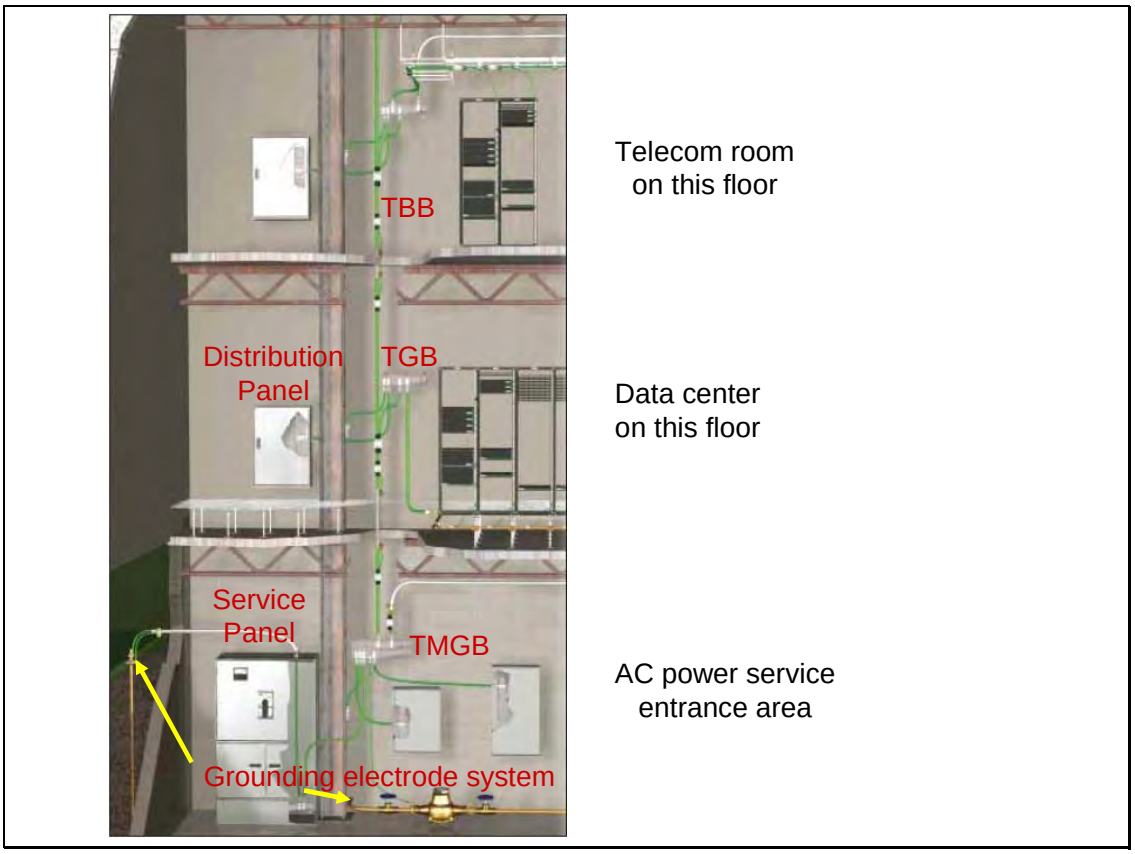

Source: Panduit Corp.

Noise concerns for the data center equipment do involve common mode noise generated and distributed by the power system to the electronic load equipment. Generally, the equipment AC input power supplies are quite tolerant of common mode noise. It should be expected that the server and other manufacturers will design and test equipment (DC input power supplies) to be similarly robust. [Already accomplished for 48V DC telecom equipment meeting the requirements of Telcordia GR-1089-CORE-2006 and placed into a Common Bonding Network (CBN).]

A supplemental signal reference structure (such as an under floor signal reference grid), where properly designed, installed and utilized) can be effective for noise control across low-frequency single-ended signaling links or poorly designed communication links. However, in many cases the design is limited to a few MHz (at best) and severely compromised by large loop areas between the cabling and the reference structure. Modern high bandwidth cabling (especially fiber optics) is significantly robust and the added value of a signal reference grid is very questionable. IEEE Std. 1100-2005 recommends “decoupling” of links for distributed systems since distributed “common”

grounding systems are suspect after low-to-mid frequencies and often impractical after several meters.

As noted in IEEE Std. 1100-2005, the default equipment bonding topology is the Common Bonding Network (CBN). The CBN can be readily utilized for efficient direct bonding of equipment and other apparatus to the grounding system. Such an arrangement provides efficient grounding and inter/intra-unit bonding of metal cabinets, racks and miscellaneous metal objects (especially when they are not powered). Electrostatic charge buildup and dissipation is also greatly aided by the CBN multiple grounding paths.

Figure 23. Data Center Grounding Infrastructure at the Room Level

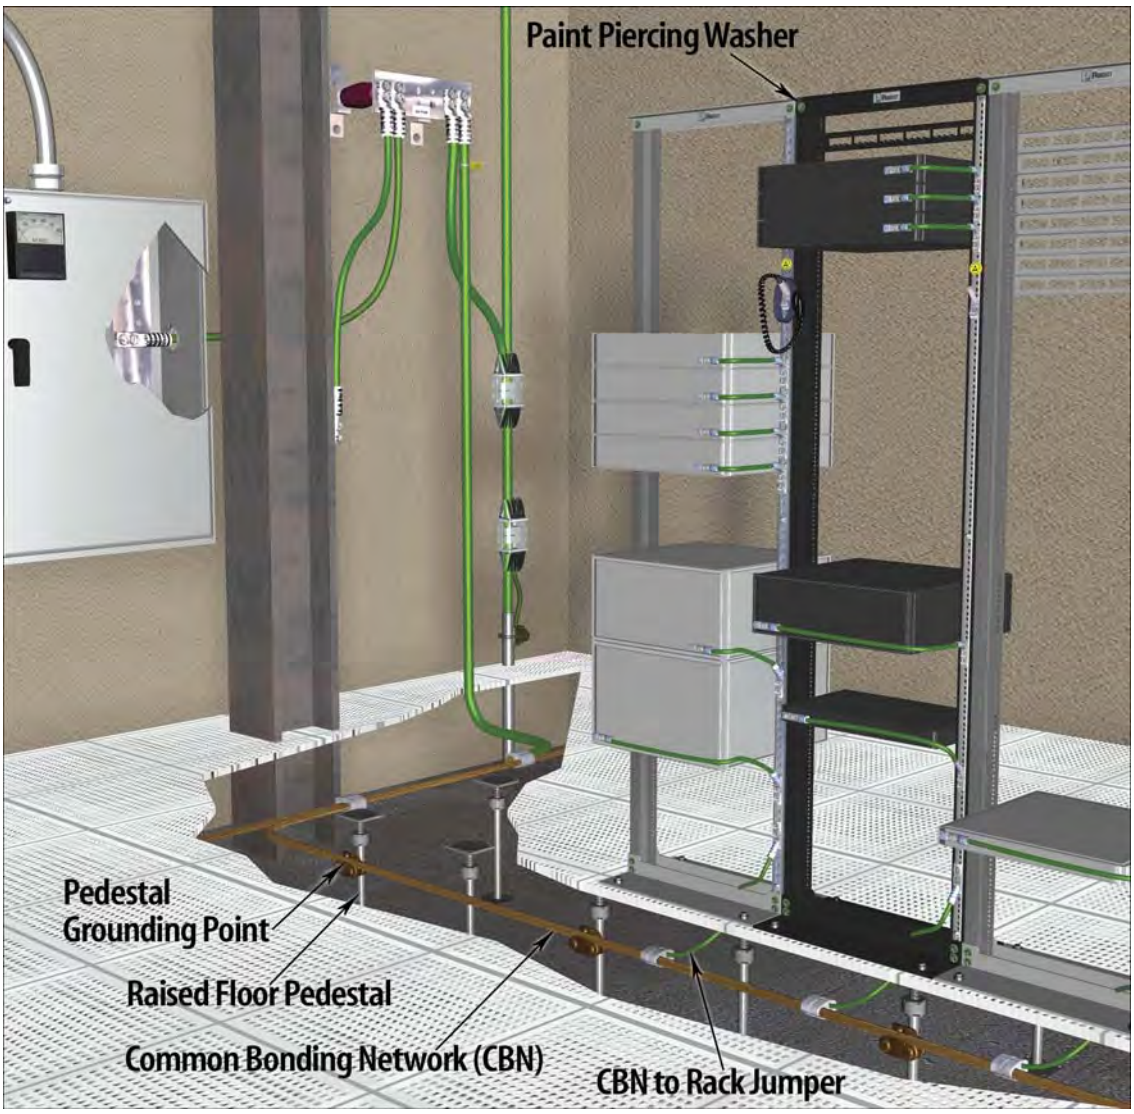

Source: Panduit Corp.

Based upon these considerations, the prudent approach is to utilize IEEE Std. 1100-2005 as the base document for supplemental grounding and bonding of the data center. This approach has the following advantages:

1. Harmonized with J-Std-607-A and ANSI T1 .333
2. Compatible with NFPA 70 and NFPA 75
3. Promotes the use of suitably immune links or decoupled links (fiber optic link is ultimate decoupling method)
4. Promotes the use of equipment with recognized immunity (less susceptibility)
5. Promotes the recognition and proper application of certain equipment bonding topologies (variations of CBN and IBN)
6. Default supplemental grounding is an extended CBN such as an above rack or under floor conductor grid.
    a. Not intended to function as a signal reference grid
7. Where otherwise deemed a requirement [historically unlikely unless for older generation Electronic data process (information technology) equipment], a signal reference grid suitable for up to a few MHz is also described.

4.3 Reliability

We received a number of inquiries from open house visitors about the DC delivery systems' reliability, and this topic certainly warranted additional investigation. There are a number of existing studies covering this topic; however, they all deal with the reliability of the telecom industry's 48V DC systems. We found no available studies dealing with the reliability of high-voltage DC (as compared to 48V DC) systems.

However, there are a number of similarities between the 48V DC systems and our Demonstration set up, and the following inferences can be drawn regarding DC power delivery in general:

- Ability to shift large heat loads out of the servers.
- Reduction of overall power consumption for each server box.
- Reduction in overall component count (from utility to silicon).
- Higher DC UPS reliability record.[39]
- Reduction of power conversion steps.

All of the available studies indicated that DC power distribution has a higher reliability record than AC distribution, often by several orders of magnitude, due to the above factors. NTT Facilities, the largest operator of data centers in Japan with close to ~80%

[39] Åkerlund , J. "DC Powering of Internet Certifies for Telephony." Uninterruptible Power Networks, UPN AB, Sweden, 2004.

of the market, also has similar data supporting increased DC power delivery reliability that they have collected in support of DC power distribution. While a definitive conclusion cannot yet be made regarding DC system reliability without additional data, it seems that a DC approach would not reduce data centers' reliability in any way.

4.4 Costs

Another issue that received a considerable amount of attention and generated considerable discussion is the cost of implementing DC power delivery systems. Unlike the available studies on DC reliability, there is less information on the costs of DC data centers. A detailed analysis of cost was outside the scope of this project, and we were not able to directly compare costs due to the fact that some of the components were not directly available, or are only available on the market in small quantities, which significantly affect costs.

Nevertheless, the project team was able to collect a number of anecdotal evidence on costs, presented below:

According to a recent white paper, the cost of a DC power plant is typically lower than an AC UPS system by 10% to 30%.[40] However, this source also claims that the additional engineering and distribution costs, along with the need to power some AC-only equipment can increase costs of a DC system for a data center.[41] In addition, manufacturers have confirmed that DC servers and storage array costs will likely be higher, at least in the near term due to the lower demand volume.

The sources that we consulted on costs all agree on one area where DC data center costs will likely be lower: TCO, or total costs of operation. This is due to a number of factors, including lower energy costs (from a reduction in both conversion - heat generation and HVAC – heat removal), lower replacement costs (from higher equipment reliability). As about half of total IT costs are related to hardware and software, with the rest in administration and overhead costs, the lower power and maintenance costs also help to reduce the administration costs.[42]

Our preliminary investigation shows that data is available for AC power delivery system costs, so that a related future investigations should consider a direct comparison of costs, from the design, build, and equip to TCO for both AC and DC data centers. Of particular interest would be a comparison of the costs to build and operate AC and DC data centers

[40] "AC vs. DC for Data Centers and Network Rooms" APC White Paper #63, American Power Conversion, 2003.

[41] Note that for our Demonstration project, the DC conversion and distribution was specifically designed to deliver DC voltage to the equipment racks as well as lighting.

[42] Burt, J. "Can DC Power Cut Data Center Costs?" *eWeek*, March 27, 2006. "Thermal Management & Server Density: Critical Issues for Today's Data Center," White Paper, Rackable Systems, 2004.

of a typical size, such as a 5 MW or 10 MW center. Coupled with reliability data, this information will go far to help inform decision makers on future data center designs.

4.5 Measurement Verification

For the Demonstration Project, two sets of measurement were made, one with the current transducers or transformers (CTs) supplied with the Signature System, and one with a commercially available Power Analyzer. This was done to ensure the reliability of the measurements, and to address the fact that one of the potential issues with the demonstration configuration could be the accuracy of the measuring instruments.[43] The approach used for the Demonstration Project is briefly outlined below.[44]

The measurement setups for AC and DC measurements respectively are shown in Figure 24. At each measurement point, voltages and currents are measured and converted to digital outputs. The setups for AC and DC measurements are shown in Figure 24. For both AC and DC, the voltage is applied directly to the voltage measurement pod, which outputs a proportional voltage signal to the data node.

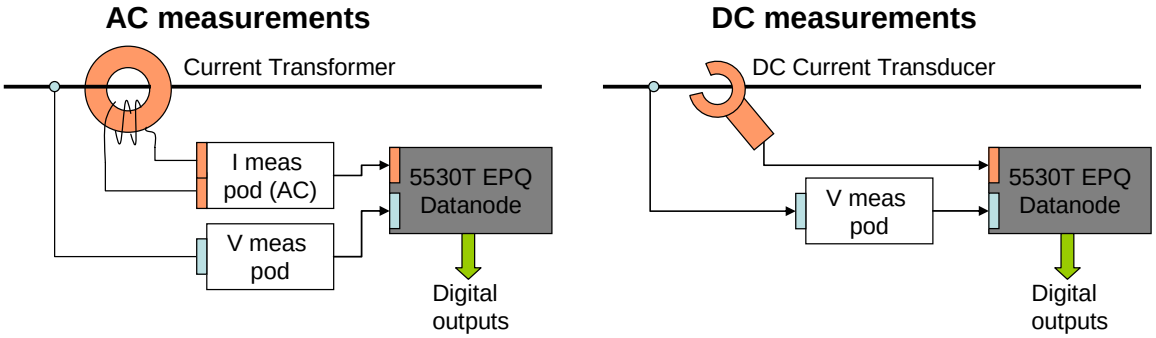

Figure 24. Measurement Points

For AC measurements, the current transformer generates a current proportional to the primary current. This current is an input to the current measurement pod, which outputs a voltage to the data node where the analog voltage is converted to a digital output. For DC current measurement, the current transducer outputs a voltage proportional to the current, and this voltage is input directly to the data node. Table 4 lists the measurement accuracies for each piece of measurement equipment used.

[43] DC current transducers with extremely high accuracy (< 1%) can be harder to locate, especially under the time constraint of the implementation schedule and the quantities required, so standard DC CTs, (~1% accuracy) were used.

[44] Analysis prepared by A. Pratt (Intel), and R. Ignall (Dranetz BMI), August 4, 2006.

Table 4. Measurement Equipment Accuracies

Description	Part Number	Range	Accuracy	Comments
Data node and measurement pod	5530T EPQ Series		±0.1%	Accuracy spec includes both data node and pods
AC Current transformer	AL-151		±1%	150A full range
DC Current transducer	PR150/SP1	15 to 180 A DC	±1%	150A full range
Server input power			±0.3%	

An error is introduced at each step in the measurement process, and the overall measurement accuracy is a function of all these errors. The calculated power measurement accuracy at each measurement point is a function of the overall current measurement accuracy and the voltage measurement accuracy, as shown in Table 5.

Table 5. Calculated Power Measurement Accuracy

Measurement point	Nominal voltag	Power range	Current measurement accuracy	Voltage measurement accuracy	Calculated power accuracy
Server input					±0.3%
AC system input	480V ac		±1.1%	±0.1%	±1.2%
AC UPS output	208V ac		±1.1%	±0.1%	±1.2%
DC system input	480V ac		±1.1%	±0.1%	±1.2%
DC UPS output	380V dc	5.7 to 68.4kV	±1.1%	±0.1%	±1.2%

These measurement errors need to be considered in determining the difference between the ratio of the measured input powers for a DC and an AC system.

Thus, this analysis of the error propagation in the system revealed that the worst case possible error would be on the order of the energy saved. Therefore, a second set of measurements were carried out (at the same locations) to verify original measurements, using the Power Analyzer and CTs with 0.25% accuracy. Both sets of measurements are included in the results section of this report.

4.6 Other Issues

A number of other issues were identified during our Open House sessions, both technical and non-technical in nature. We list a number of them here:

- Awareness: While the demonstration was able to increase industry awareness of more efficient alternatives to today's practices, the industry's current knowledge of available options is quite low. Additional research and outreach by industry respected sources are needed.

- Creating a Market for DC: For DC approaches to take hold in the data center market, a number of market barriers still need to be addressed in a consistent, unified manner.

- On-Going Efforts: The Demonstration Project accomplished its objective of informing the industry of the DC distribution alternative, but once the project is completed, there are few other places where such a set up can be found.

5. Safety/Codes & Standards/Industry Associations

Due to the experimental nature of the demonstration project, as well as the agreed upon DC voltage level, which is higher than typically found elsewhere, the project team took steps to ensure both equipment and personnel safety.[45] This was accomplished though the establishment of a "Safety Committee" consisting of project participants with extensive experience in data center design and construction, participants from the telecommunication industry, power electronics, and other.

The "Safety Committee" reviewed the implementation plans, one-line drawings and discussed equipment specifications and compatibility prior to assembly. Members of the committee also participated directly in the equipment assembly process. The project team also consulted members of the committee regarding issues such as connectors. Participants also agreed that the Committee was primarily responsible for safety issues "outside of the rack," while equipment manufacturers were responsible for the safety "inside of the rack." This clear division of responsibilities helped to reduce confusion and minimized conflicts during the assembly.

In addition to assisting the project team with equipment finalization and assembly, the Committee also identified issues and relevant standards needed for data center DC power practices. The section below summarizes the recommended practices and associated codes and standards that warrant consideration for a 380VDC power system. At a minimum, these considerations should include:

- Selecting and locating suitable circuit protection
- Grounding and bonding methods to insure proper operation of the suitable circuit protection (as detailed in section 4)
- Leakage currents and shock hazards
- Environmental conditions and maintenance
- Connection Devices
- Power distribution equipment within the rack
- Training and certification of personnel working inside racks
- Training and certification of personnel installing building power distribution
- Installation and maintenance of under floor power distribution
- Warnings and labeling
- Commission Testing

[45] The International Electrotechnical Commission and its national counterparts (IEE, IEEE, VDE, etc.) define high voltage circuits as those with more than 1000 V for alternating current and at least 1500 V for direct current, and distinguish it from low voltage (50–1000 V AC or 120–1500 V DC) and extra low voltage (<50 V AC or <120 V DC) circuits. Note that various safety and insurance organizations consider anything outside of the ELV range (i.e. greater than 50 V) to be dangerous and in need of regulation.

5.1 Selecting and Locating Suitable Circuit Protection

As discussed in section 4, suitable circuit protection using breakers and fuses does not directly translate from AC to DC and suitable DC circuit breakers will likely need to be selected. In terms of the breaker ratings, the manufacturers of these devices must be consulted for coordination. Further, experience in the from the FAA critical DC bus power systems work, along with information from photovoltaic inverter industry, the DC drive industry and electromotive (DC train) industry can serve as a reference resource for further information on this topic. Available fault currents may vary depending upon the type and short circuit current availability from any energy storage devices used. Fault clearing capabilities of DC rated protection devices may not be sufficient under all fault conditions.

5.2 Grounding Methods to insure suitable operation of circuit protection

Section 4 detailed the basic considerations for grounding and bonding, but focuses on grounding topologies and techniques that enhance or supplement basic grounding and bonding requirements. In terms of promoting operation of the circuit protective devices during fault conditions, the recommendations set forth in the National Electric Code (NFPA 780) and the IEEE *Emerald Book* contain the fundamental recommendations that can be incorporated with the previous discussion on suitable circuit protection. Also it should be considered that lower voltages than the subject 380V DC demonstration may be considered in the future, therefore, the telecommunications related standards may provide some of the useful materials related to the lower voltage DC systems.

5.3 Leakage Currents and Shock Hazards

Because AC and DC powered information technology equipment can and does contain noise filters and other components that create leakage currents across the ground paths the requirements outlined in the Underwriters Laboratories document *UL 60950-1* must be taken into consideration as a basic listing for ITE. The document entitled: *Information Technology Equipment – Safety – Part 1: General Requirements* are basically related to leakage current, but the issue of electric shocks is well represented. This relates to considerations where the grounding conductor inadvertently becomes loose or disconnected and the rack or the equipment case becomes energized.

5.4 Connection Devices

Market adoption of higher voltage DC systems will require the standardization of DC connectors that can be safely used by untrained personnel. In existing data centers personnel can connect/disconnect servers powered at 120 V without any safety concerns. Because DC current does not have a zero crossing like AC current, it is more difficult to interrupt, and simple disconnection will cause arcing. In order to address this concern,

some connectors have been developed that can contain the arc until it is extinguished. End users will want standard connectors to choose from for these applications.[46]

Figure 25. Tesla coil showing repeated electric discharges.

5.5 Power Distribution Equipment within the Rack

For the same reasons outlined in section 5.4 above, cabling and plug strips used in server racks must also have standard safe designs, so that untrained personnel can reasonably be expected to work in the rack without danger.

5.6 Training and Certification of Personnel Working Within Racks

One issue of concern to datacenter users is the availability of trained personnel to work on DC power systems. Licensed electricians are common for AC power systems, but the requirements for handling DC systems will be different and therefore require special training. To address this concern, training programs must be developed and provided to personnel working on DC facilities. This training should include both work safety requirements and recommissioning after any modifications

5.7 Training and Certification of Personnel Working on Distribution Systems

The same requirements are needed for personnel working on facility distribution system as

[46] Interrupting an existing current flow often produces a low voltage spark or arc. As the contacts are separated, a few small points of contact become the last to separate. The current becomes constricted to these small hot spots, causing them to become incandescent, so that they emit electrons (through thermionic emission). Even a small 9V battery can spark noticeably by this mechanism in a darkened room. The ionized air and metal vapor (from the contacts) form plasma, which temporarily bridges the widening gap. If the power supply and load allow sufficient current to flow, a self-sustaining arc may form. Once formed, an arc may be extended to a significant length before breaking the circuit. Attempting to open an inductive load facilitates the formation of an arc since the inductance provides a high voltage pulse whenever the current is interrupted. AC systems make sustained arcing somewhat less likely since the current returns to zero twice per cycle. The arc is extinguished every time the current goes through a zero crossing, and must reignite during the next half cycle in order to maintain the arc.

personnel working within equipment racks, discussed above in section 5.6.

5.8 Installation and Maintenance of Under Floor Power Distribution

Power cable under the raised floor does not provide the same level of comfort as overhead cabling, as clearances and touch hazards may be more difficult to control. It is very likely that under floor cabling will be a preferred arrangement, so standards for installing these systems will be necessary for end users to be comfortable with this approach.

5.9 Warnings and Labeling

Because these are not yet common systems, there may be a need to develop supplemental warning labels to place on and inside cabinets where personnel might be exposed to high DC voltages and or capacitive voltages when the system is de-energized

5.10 Commissioning Tests

Once the suitable circuit protection has been designed and integrated into the systems it is important that consistent and repeatable procedures and test methods be developed and made available to personnel who will be performing commissioning of the power systems and the protection settings.

6. Results

The magnitude of the DC efficiency gain is highly dependent on the AC reference system and AC/DC power supply that it is being compared to. In the case of this Demonstration Project, we were fortunate enough to have access to two AC distribution systems as well as two DC conversion/distribution systems, and the efficiency ratios were determined for both sets. Our results indicate that the DC approach does provide an increase in conversion efficiency, however exposing the industry to "best in class" systems may also be useful in improving the average efficiency of data centers.

6.1 AC Reference System Measurements

The AC system efficiency was measured with two different UPS systems. The units were both high-efficiency transformer units loaned by industry partners. The initial measurements and set up primarily were primarily carried out on the original loan unit (AC System A). However, another industry partner was invited to submit their system towards the end of the demonstration period, and complied with their highly efficient AC UPS setup (AC System B).

Measurements here are designated as follows:

- Input energy to UPS= *ACEin*
- Output energy to AC loads= *ACEout*
- Input energy to Rectifier= *DCEin*
- Output energy to DC loads= *DCEout*

The figure below shows how the measurements were made for the AC system:

Figure 26. AC Systems Measurement Set Up

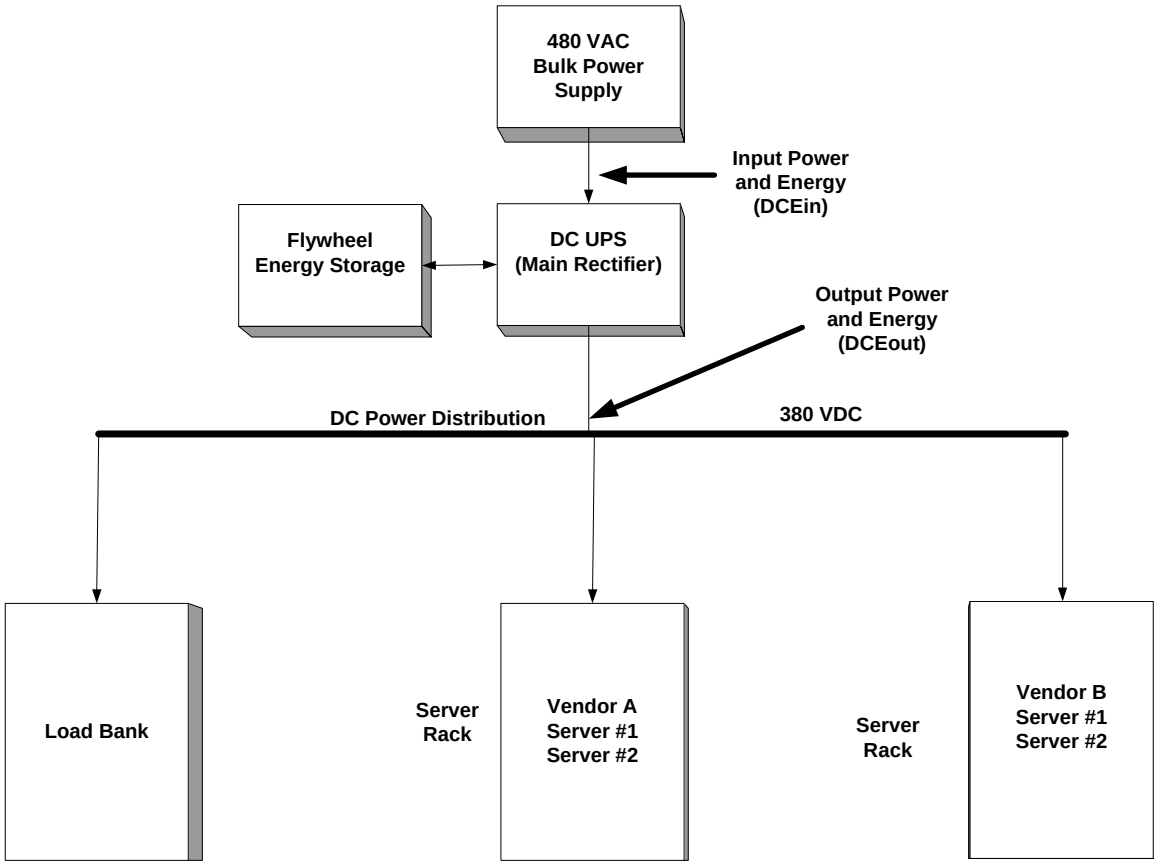

6.2 DC System Measurements

The DC system efficiency was also measured for two different rectifier units, also on loan from industry partners. The efficiency of both rectifier units was measured.

The diagram below shows how the measurements were made for the DC system:

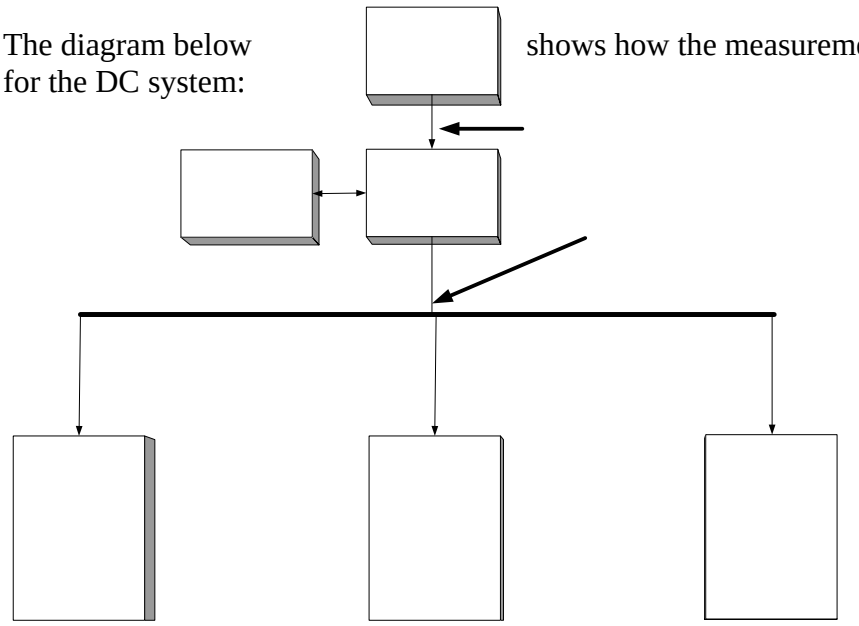

Figure 27. DC Systems Measurement Set Up

6.3 Efficiencies of AC- and DC-Based Systems

For the measurements, the AC and DC distribution used identical storage and server units except for their power supplies. Both DC and AC processing equipment used the same exercising routines to simulate processor and system loading.[47] Although these "simulated" loads were fairly constant, there were still some small fluctuations. In order to ensure these did not skew the results, power measurements were integrated over typical utility demand periods of 15 minutes, using the following formula:

$$E = \int_0^{15} p(t)dt$$

With the server activities constant for both systems, this means the power on the load side of the AC UPS or DC rectifier is nearly equal, a comparison of the input power can then show the difference in energy use. In actuality, there is about a 2% difference in the power loading, due to the higher efficiency of the DC servers.

Therefore,

$$DCEout / ACEout \approx 0.98$$

The objective is to find the percentage improvement, *I*, of the DC system over the AC system.

[47] Server manufacturers provided their own set up and exercise routines for their equipment sets.

$$I = \frac{ACEin - DCEin}{ACEin}$$

The AC and DC power measurements are compared using data gathered with DC System A set up first, then using the DC System B set up. Both sets of results are presented below.

Table 6. DC System A Efficiency Measurements

System Efficiency	UPS Efficiency	Transformer Efficiency	PS Efficiency	System Efficiency
AC System A: Measured Efficiency	90%	98%	90%	79%
AC System B: Measured Efficiency	90%	98%	90%	79%
DC System A: Measured Efficiency	94%	100%	92%	87%

Energy Consumption	Compute Load (kWh)	Input Load (kWh)	Efficiency Gain
AC System A: Measured Consumption	23.3	26.0	
AC System B: Measured Consumption	23.3	25.9	
DC System A: Measured Consumption	22.7	24.1	
% Energy Consumption Improvement vs. AC System A			**7.3%**
% Energy Consumption Improvement vs. AC System B			**7.0%**

Table 7. DC System B Efficiency Measurements

System Efficiency	UPS Efficiency	Transformer Efficiency	PS Efficiency	System Efficiency
AC System A: Measured Efficiency	90%	98%	90%	79%
AC System B: Measured Efficiency	90%	98%	90%	79%
DC System B: Measured Efficiency	92%	100%	92%	85%

Energy Consumption	Compute Load (kWh)	Input Load (kWh)	Efficiency Gain
AC System A: Measured Consumption	23.3	26.0	
AC System B: Measured Consumption	23.3	25.9	
DC System B: Measured Consumption	22.7	24.6	
% Energy Consumption Improvement vs. AC System A			**5.0%**
% Energy Consumption Improvement vs. AC System B			**4.7%**

It can be seen in the first table that there is about a 7% decrease in input energy using the DC system compared to the "best in class" AC system. Using the second DC system, the values are slightly lower, but still about 5% improved over the AC system.

Finally, it should be noted that both of the AC distribution system used represent the best on the market with regard to efficiency. Both of the AC UPSs are high efficiency units, and the efficiencies of the power supplies in the AC servers – at 90%, are much higher than most in the data center industry. By comparison, a typical AC system in today's data center would have a UPS that was about 85% efficient, and power supplies around 73%

efficient.[48] The estimated improvement of the DC system over these "typical" systems is shown in Table 8. System Efficiency Comparisons below.

Table 8. System Efficiency Comparisons

System Efficiency	UPS Efficiency	Transformer Efficiency	PS Efficiency	System Efficiency
AC Typical Distribution Efficiency	85%	98%	73%	61%
DC Distribution Efficiency	92%	100%	92%	85%

Energy Consumption	Compute Load (W)	Input Load (W)	Efficiency Gain
Typical AC Distribution Efficiency	10,000	16,445	
DC Distribution Option (Optimized)	10,000	11,815	
% Energy Consumption Improvement vs. Typical AC Distribution			**28.2%**

In this case, an improvement of over 28% is possible in an average data center. This means the DC distribution system, as demonstrated, will have the potential of using 28% less energy than the typical AC system found in today's data centers. Since LBNL benchmarking has shown that HVAC loads are typically about the same as the IT load (or worse), this means that a 28% improvement in the efficiency of the electrical distribution to the IT equipment also means a 28% overall facility level improvement.

6.4 Non-Technical Project Results: Open House Sessions

A total of seven Open Houses (including one session on August 18) were held during the duration of the project, with over 200 total attendees representing a diverse cross-section of the data center industry. Attendees ranged from manufacturers to users/operators, government, as well as the media.

A number of Open Houses garnered much larger audiences due to the fact that they were planned to coincide with other events taking place in or near the facility. One was the Data Centre Dynamics conference that took place in San Francisco in July; the other was the Critical Facilities Roundtable quarterly meeting, which was hosted at the facility. The Demonstration also received interest from utilities, and a number of utility representatives and government officials attended a "utility" Open House session.

[48] These are typical numbers that were found in our evaluation of the servers and UPSs markets.

Figure 28. Demonstration Site with Installed Equipment

The response to the demonstration's press release was quite spectacular, including coverage by a Bay Area television station, as well as articles about the demonstration project by trade and popular press following the press event. The interest generated by the project was quite broad, with coverage by the press ranged from CNN Money to Byte.com, to *CIO Magazine* and *eWeek.* While there was no direct press tracking by LBNL or participants, as of the end of September, the demonstration project generated at over 20 print and on-line articles, and at least one TV interview.[49]

[49] References to the press coverage can be found at http://hightech.lbl.gov/dc-powering/media-coverage.html.

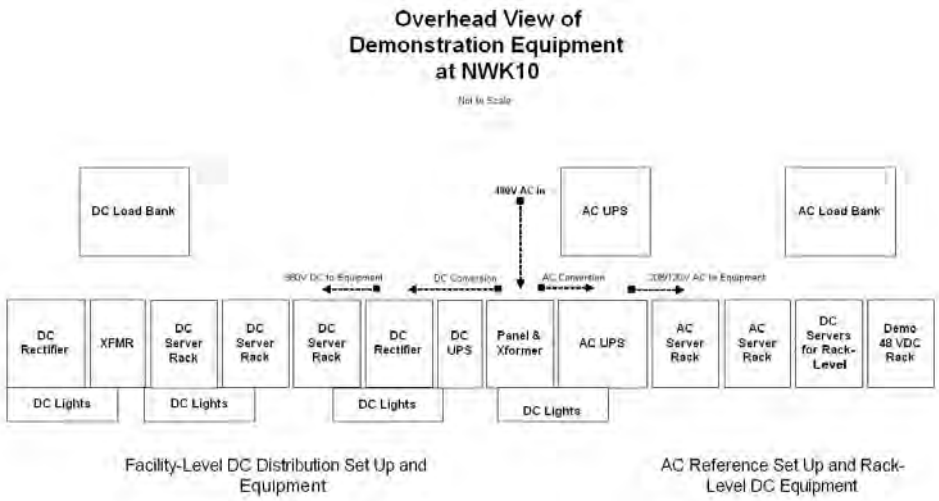

Figure 26. Installed Equipment Map of the Demonstration Site

7. Conclusions and Recommendations

In the final analysis, this Demonstration Project was able to coordinate the participation of over 20 companies and their equipment and/or in kind contribution, worked with other companies in the implementation process, and assembled equipment and services worth over a million dollars in value. We were also able to conclusively demonstrate to the data center industry (via the 200+ attendees and the media) that DC delivery systems are viable, can be 20% or more efficient than current AC delivery systems, be more reliable, and potentially cost less in the long run.

Overall, the project succeeded in meeting the objectives that were set out at the beginning. In particular:

Availability of DC Equipment: The Demonstration Project showed that DC-powered servers exist in the same form factor as AC servers or can be built and operated from existing components with minimal effort. The Project reminded the industry that DC servers currently exist (in the 48V DC form factors), and also showed that 380V DC servers and storage equipment could be built and operated from existing components. Further, the Demonstration Project gave visibility to the DC power conversion and distribution equipment, highlighting two commercially available rectification systems, as well as UL-listed buss bars for DC applications.

DC Functionality: The Project also showed that DC-powered servers can provide the same level of functionality and computing performance when compared to similarly configured and operating server containing AC power supplies. While the DC-powered servers were not compared with AC powered servers using industry standard software benchmarking tools, this was not a significant issue due to the fact that identical server units (except for the power supplies) were provided by manufacturers. Measurements of efficiency were made using industry standard measurement devices. The Demonstration also included storage units, as well as DC network equipment that can use a variety of DC voltages.

Demonstrated Gains in Efficiency: The Project demonstrated efficiency gains from the elimination of multiple conversion steps in the delivery of DC power to the server hardware. Results were measured and documented from two sets of DC delivery systems, and compared to two sets of AC delivery systems. In both cases, the DC delivery system showed a minimum of 5% to 7% efficiency gains without significant optimization over two AC distribution systems that are “best in class” and much more efficient than most systems found in today’s data centers. These measured efficiency gains did not include additional gains from a reduction in cooling loads, which can have the potential for additional savings. Raising awareness of the AC - UPS system efficiency will have a benefit even if the DC solution is not embraced.

Follow-Up Investigations: The Project also identified a number of areas for follow up investigations that will help generate industry discussions, and provide

useful leverage points to move the industry forward in the direction of DC distribution. These identified areas are topics that require industry agreement and/or adoption, which included:

- Grounding, Protection and Overloading Prevention: A number of grounding, protection and overload prevention practices for DC data centers are discussed in this report for industry considerations.
- Reliability: Anecdotal data shows that DC-powered data centers have the potential to be more reliable than AC-powered data centers. However, data does not exist for DC voltages higher than 48V DC. Leadership will need to come from the industry in adoption or additional testing for the industry to move forward on this area.
- Costs: While cost data exists, it has not yet been compiled in a way that direct comparisons can be made for the two distribution systems (or their TCO). The Demonstration has generated significant interests from data center designers and system integrators, and further discussion on this area can lead to at least a first-order estimate of DC distribution costs. A related area is the costs of DC components vs. AC components. Currently, AC components may enjoy better economies of scale, but wide-spread DC power adoption may change this equation.
- Integration with Other Sources: A DC data center can also simplify the integration of alternative energy sources, such as solar and other forms of renewable energy, as well as fuel cells and distributed generation, which are all DC-based.
- Measurement Verification: The accuracy of measurements was confirmed with an additional measurement approach. This helped to address any concerns over introduced error factors from the instrumentation used.
- Other Issues: Of significant concerns is the lack of industry knowledge of the advantages of DC distribution, as well as misconceptions about DC power. Additional education and outreach efforts will be required if the energy savings potentials of DC powered data centers are to be realized.

Going forward, there remain many barriers to the adoption of DC power distribution, and need additional follow on work. In particular, a number of barriers have been identified during the course of this projects, they include:

- Increasing Awareness of DC Distribution: The industry's current knowledge of available options for efficiency and DC distribution is quite low. Further, there is no single, trusted source of information, or an entity dedicated to the promotion of data center energy-efficiency and DC distribution (other than the CEC/PIER efforts). With the current industry support and interest, a "DC Power" association of some sort will help to focus interest and help to elevate awareness among the data center market.

- Creating a Market for DC: A number of market barriers still need to be addressed in a consistent, unified manner. Coordinating utilities and other

efforts, at least in California, will go far towards getting DC approaches to take hold in the data center market. There is utility interest in establishing a baseline of performance and cost, which can then help to address at least the early adoption barrier of cost. Other efforts are still needed, and strategies to address market transformation used by the conservation movement can be directly applied here. In addition, the US Congress has recognized the potential for energy savings with HR-5646, so that coordination with DOE and EPA is needed to ensure no duplication of efforts.

- Developing Pilot Projects: Once the Demonstration Project was completed, there are no other places where such a set up can be found. Efforts are needed to continue the Demonstration Project's role in informing the industry of the DC distribution alternative. Discussions are underway with a number of "early adopters," and the project team proposes to establish several pilot projects to:
    - "prime" the market by creating demand for DC servers to enable certification efforts to proceed
    - Determine cost factors for DC systems - capital and operating cost
    - Evaluate and resolve any remaining barriers
    - Publicize successful systems in real data centers

- Develop standards to accelerate adoption: Agreement on distribution voltages, electrical connectors, grounding, DC power quality, and other issues will be important to enable the market to adopt DC distribution on a large scale. The PIER program should facilitate these efforts by bringing together the appropriate industry representatives.

Appendix A: Project Participants

Participants List

Project Organizers

- Lawrence Berkeley National Laboratory: William Tschudi, Evan Mills, Steve Greenberg
- Ecos Consulting: My Ton
- EPRI Solutions: Brian Fortenbery

Industry Partners

- Alindeska Electrical Contractors
- AMD
- APC
- ANCIS Inc.
- Baldwin Technologies
- CCG Facilities
- California Energy Commission
- Cingular Wireless
- Cisco Systems
- Cupertino Electric
- Data Power Design
- Dell
- Dranetz-BMI
- Dupont Fabros
- EDG2
- Emerson Network Power
- EYP Mission Critical Facilities
- Fairchild Semiconductors
- Gannett
- HP
- IBM
- IEM (Industrial Electric Manufacturing)
- Intel
- Morrison Hershfield
- Nextek Power Systems
- NTT Facilities
- Panduit Corp.
- Pentadyne
- Rosendin Electric
- RTKL
- Rackable Systems
- SatCon Power Systems
- SBC Global
- Solara
- Square D
- Sun Microsystems
- 380voltsDC.com
- TDI Power
- Universal Electric
- Verizon Wireless
- Visa

Appendix B: Location for One-Line Drawings of Demonstration Configurations

Demonstration configurations can be found on line at:

http://hightech.lbl.gov/dc-powering/walkthrough.html

Appendix C: Participant "Code of Conduct"

GENERAL UNDERSTANDING AMONG PARTICIPANTS OF THE DC ARCHITECTURE DEMONSTRATION PROJECT

1. This Demonstration will evaluate the energy efficiency benefits of DC power distribution systems in data centers. It will be carried out as a public-private partnership coordinated by Lawrence Berkeley National Laboratory, Ecos Consulting and EPRI Solutions. The work is sponsored by the California Energy Commission and the Participants.

2. There is no cost to become a project Participant. However, it is understood that Participants will contribute equipment, time, and/or expertise.

3. Participants will be identified in project publications and share in the credit and results of the Demonstration.

4. The Demonstration will evaluate data center AC and DC distribution systems from a power delivery efficiency perspective, comparing traditional AC distribution, DC distribution at the rack level, and DC distribution at the facility level. The methodology will be to exercise any given piece of IT equipment identically for both the AC and DC configurations in order to create similar loading on the power delivery systems.
   a. Performance comparison of AC and DC distribution systems will be obtained.
   b. To the extent possible, the vendors of specific equipment will not be publicly identified. Reporting will focus on the energy implications of the distribution systems rather than on the types of equipment or a particular vendor. For example, Rack 1 equipped with servers A, B, and C the configuration may be identified, while the manufacturers of the servers or power conversion equipment will remain anonymous.
   c. Individual equipment (such as a server) may run any application selected by the manufacturer, as long as operating conditions are identical in the AC and DC trials.

5. A final project report will be made available to all participants for review before external submission/dissemination.

6. Participants agree to communicate a consistent message to outside entities regarding the purpose and goals (per this memorandum) as well as the results of this Demonstration (per the final project report).
   a. LBNL will coordinate publicity with Participants and outside entities. Initial news releases should be published on a single day to be mutually determined.
   b. Any “Press coverage” with respect to the Demonstration should clearly identify the goals and results of the project and clearly indicate the “multi-party” involvement in the project.

7. The host facility will make reasonable arrangements for site access.
   a. During setup and operation, access to equipment used in the demonstration should be provided only on an as-needed basis.
   b. No personnel should be able to conduct internal examination of specific equipment other than their own without explicit consent of the relevant equipment owner.
   c. Host facility to provide access (with escort if required) to all participants, their guests, or other interested parties for showcasing the Demonstration.
   d. Organized open house sessions with one of the LBNL team present will be held rather than ad hoc visits.

8. The host facility will provide direct or remote access to monitor power and change of application programs during defined time periods.

Appendix D: Measurement Details

1 AC Reference System Measurements

The AC system efficiency was measured with two different UPS systems. The units were a Mitsubishi-UPS with a standard 75kVA input transformer unit loaned by Baldwin Technologies, Inc. (it is a unit owned by BTI who on loan to Pentadyne for testing with their flywheel system), and an APC-UPS unit with a high efficiency input transformer loaned by American Power Conversion, Inc. The initial measurements and set up primarily used the Mitsubishi UPS system. However, APC was invited to submit their system towards the end of the demonstration period, and complied with a highly efficient AC UPS setup. For thoroughness of the demonstration, the performance of APC UPS was also measured. [50]

2 DC System Measurements

The DC system efficiency was also measured for two different rectifiers. Initial measurements used a SatCon Inc. rectifier unit, but since the Demonstration Project was provided with two DC rectification systems, a Nextek rectifier unit was also measured.

3 Efficiencies of AC- and DC-Based Systems

For the measurements, the AC and DC distribution used identical storage and server units except for their power supplies. Both DC and AC processing equipment used the same exercising routines to simulate processor and system loading.[51] Although these "simulated" loads were fairly constant, there were still some small fluctuations. In order to ensure these did not skew the results, power measurements were integrated over typical utility demand periods of 15 minutes, using the following formula:

$$E = \int_0^{15} p(t)dt$$

With the server activities constant for both systems, this means the power on the load side of the AC UPS or DC rectifier is nearly equal, a comparison of the input power can then show the difference in energy use. In actuality, there is about a 2% difference in the power loading, due to the higher efficiency of the DC servers.

Therefore,

$$DCEout / ACEout \approx 0.98$$

[50] The measurement and monitoring system used was the Signature System on loan from Dranetz-BMI.
[51] Server manufacturers provided their own set up and exercise routines for their equipment sets.

The objective is to find the percentage improvement, *I*, of the DC system over the AC system.

$$I = \frac{ACEin - DCEin}{ACEin}$$

The AC and DC power measurements are compared using data gathered using the SatCon as the DC rectifier first, then using the Nextek unit as the DC rectifier. Both sets of results are presented below.

Table 9

System Efficiency	UPS Efficiency	Transformer Efficiency	PS Efficiency	System Efficiency
AC System A: Measured Efficiency	90%	98%	90%	79%
AC System B: Measured Efficiency	90%	98%	90%	79%
DC System A: Measured Efficiency	94%	100%	92%	87%

Energy Consumption	Compute Load (kWh)	Input Load (kWh)	Efficiency Gain
AC System A: Measured Consumption	23.3	26.0	
AC System B: Measured Consumption	23.3	25.9	
DC System A: Measured Consumption	22.7	24.1	
% Energy Consumption Improvement vs. AC System A			**7.3%**
% Energy Consumption Improvement vs. AC System B			**7.0%**

Table 10

System Efficiency	UPS Efficiency	Transformer Efficiency	PS Efficiency	System Efficiency
AC System A: Measured Efficiency	90%	98%	90%	79%
AC System B: Measured Efficiency	90%	98%	90%	79%
DC System B: Measured Efficiency	92%	100%	92%	85%

Energy Consumption	Compute Load (kWh)	Input Load (kWh)	Efficiency Gain
AC System A: Measured Consumption	23.3	26.0	
AC System B: Measured Consumption	23.3	25.9	
DC System B: Measured Consumption	22.7	24.6	
% Energy Consumption Improvement vs. AC System A			**5.0%**
% Energy Consumption Improvement vs. AC System B			**4.7%**

It can be seen in the first table that there is about a 7% decrease in input energy using the DC system compared to the "best in class" AC system. Using the Nextek system, the values are slightly lower, but still about 5% improved over the AC system.

Finally, it should be noted that both of the AC distribution system used represent two of the best on the market with regard to efficiency. Both of the AC UPSs are high efficiency units, and the efficiencies of the power supplies in the AC servers – at 90%, are much higher than most in the

data center industry. By comparison, a typical AC system in today's data center would have a UPS that was about 85% efficient, and power supplies around 73% efficient.[52] The estimated improvement of the DC system over these "typical" systems is shown in Table 8. System Efficiency Comparisons are shown below.

Table 11

System Efficiency	UPS Efficiency	Transformer Efficiency	PS Efficiency	System Efficiency
AC Typical Distribution Efficiency	85%	98%	73%	61%
DC Distribution Efficiency	92%	100%	92%	85%

Energy Consumption	Compute Load (W)	Input Load (W)	Efficiency Gain
Typical AC Distribution Efficiency	10,000	16,445	
DC Distribution Option (Optimized)	10,000	11,815	
% Energy Consumption Improvement vs. Typical AC Distribution			**28.2%**

In this case, an improvement of over 28% is possible in an average data center. This means the DC distribution system, as demonstrated, will have the potential of using 28% less energy than the typical AC system found in today's data centers. Since LBNL benchmarking has shown that HVAC loads are typically about the same as the IT load (or worse), this means that a 28% improvement in the efficiency of the electrical distribution to the IT equipment also means a 28% overall facility level improvement.

[52] These are typical numbers that were found in our evaluation of the servers and UPSs markets.

Attachment XIV: Technology Transfer Activities

July 14, 2006

Summary
Computational Research and Theory Building
Energy Design Charrette
June 15 – 16, 2006

Introduction

On June 15th and 16th, 2006 an Energy Design Charrette was held for the planned Lawrence Berkeley National Laboratory (LBNL) Computational Research and Theory Building. This summary reflects the ideas presented during the charrette. The charrette was somewhat atypical in that it occurred following the completion of the Programming Phase of the project before design solutions were developed by the design team. In a traditional charrette, a design concept is typically reviewed and critiqued in order to offer suggestions for improvements. In this case, with the program requirements as the starting point, the charrette explored design concepts, industry trends, emerging technologies, uncertainties, barriers, and creative solutions. This timing was chosen to allow the ideas and concepts from the charrette to be applied initially to the design and hopefully minimize false starts or backtracking to incorporate energy efficient strategies.

The charrette offered a unique opportunity to exchange information and ideas between LBNL's scientific computing, facilities, and Environmental Energy Technologies Divisions, data center design professionals, and Industry representatives. Several representatives from the American Society of Heating Refrigerating and Air-conditioning Engineers' (ASHRAE) technical committee on data centers participated in the Charrette. Attached (Appendix A) is list of attendees and the agenda (Appendix B) for the Charrette.

The objective of the Charrette was to bring together the stakeholders and other technical resources to provide a roadmap for creation of the most energy efficient high performance computer center in the world - a world class building for world class computing. The charrette focused on the energy aspects of the most intensive areas of the building (as opposed to office areas), and explored the expected energy loads and how they could be reduced or efficiently accommodated. Creative solutions, as well as constraints and metrics to assess performance were identified. The information discussed should lead to design strategies to be further investigated

Below is a summary of key points from the two days of discussion:

July 14, 2006

June 15

Overview of Computational Research and Technology goals

Michael Banda of the Computing Sciences Division provided an overview of NERSC and project goals.

- Energy Efficiency is a requirement
- Project must satisfy program requirements
- Office of Science wants to see NERSC relocate to the Lab

Programming Summary

Bart McClelland of RMW, the lead architect for the programming phase, provided an overview of the project as determined through the program interviews. Key information includes:

- 40,000 sf computer floor
- 35,000 sf support space
- 75,000 sf office space
- 350 staff
- M&E service yards: 33,000 sf
- Site is on steep hill at entrance to lab
- Envision air handlers along side of equipment room
- Facility needs to provide reliable, high levels of power and cooling to support multiple generations of advanced computational equipment - yet to be designed
- Goal is 100% uptime
- Cooling towers may be on grade

Dan Smith and Charles Lee of Practicon-Ackerman, the electrical and mechanical engineers for the programming phase discussed preliminary concepts. Are we designing for people comfort or machine comfort?

July 14, 2006

NERSC requirements

Howard Walter provided a summary of NERSC requirements. See attached presentation. Key points are:

- Computer equipment loads:

Year	Computer Load	Total load
2009	19 MW	25 MW
2013	31 MW	45 MW
2015	37 MW	62 MW

- 37 MW ~ 900 W/sf

- Currently Oakland Scientific Facility total load is approximately 2 MW

- LBNL's total electrical load today is about 12 MW. Can utility (WAPA) provide additional 62 MW in this timeframe? Will the NERSC program support a $35 million yearly power cost in the future?

- Infrastructure for 10 years needs to be provided (i.e. 62 MW) - DOE will expect infrastructure to be in place for computer procurements during this time.

- Uninterruptible Power Supply (UPS): 1.2 MW initially going to 2.0 MW. UPS will only support storage devices - 1000 terabytes of storage - not entire IT or facility load.

- Every three years a new computer system is procured – each system runs 5 to 6 years (two computers on floor). Must have ability to operate one side while changing out the other.

- 8' tall equipment racks

- minimum 12' floor to ceiling - no need for dropped ceiling

- Minimum 36"- 42" raised floor - primarily for cabling

- Use of air and water side economizers are possible

- Redundancy is planned to be N + 1. (should N+2 be planned to allow for maintenance?)

- ♦ Include provision for ease of powering 110 V AC, 208 V AC, and 480 V single and three phase AC

- ♦ 3 zones are envisioned – two areas for computer racks, + center storage. Each computer area should be isolatable for ease of changeout. Zones could have different environmental conditions.

What defines a successful project?

The participants were asked to define success for this project and the following ideas were provided. The participants identified the following:

- Maximum computational power and research output
- Usable
- Flexible
- Flexible cooling options - ability to use air or liquid based systems
- Energy Efficiency - a model for such facilities
- Success in competing for future procurements
- Showcase/Replicability
- Scalable
- Supporting Computing capability
- Data center as the core of the building
- Efficiency – e.g. Watts/flop
- World class data center
- Getting NERSC back on the hill
- Ability to explain science and energy efficiency to the public
- Maintainability - high % of time operating and % processor utilization

This list of attributes was to be revisited near the end of the Charrette to ensure that concepts/actions identified supported these goals.

What metrics should be used to evaluate options?

Next, a discussion of metrics applicable to scientific computing was held. The group brainstormed the following as possible metrics of interest:

- Uptime - % of time in operation
- Papers/ MWh or Calculations/MW
- Sustained throughput for scientific computations - sustained system performance (SSP): Watt-hrs/teraflop-yrs
- Sustained flops vs. peak
- Flops/watt
- Computing power/square feet
- Papers/publications
- Awards

- Ability to attract future funding
- Science/$
- Computing power vs. Total power: IT watts/Total watts
- Computing power vs. HVAC system power: IT watts/HVAC watts
- IT Balance (e.g. need memory as well as computation)
- Bandwidth to and from storage
- Storage efficiency (e.g. watts per unit storage)
- Degree of innovation – percent improvement in cooling, bleeding edge vs. cutting edge
- % of time within environmental specifications
- % energy and performance improvement over baseline
- Speed to install and activate new systems

Bruce Nordman brought up that industry benchmark committees (e.g. SPEC) are developing server metrics (peak performance per watt and utilization). The idea is to tie energy use to existing computing benchmark programs. Additional metrics applicable to scientific computing could be developed (e.g. sustained flops/average Watts)
Top 500 supercomputer list could be a starting point for collecting actual energy consumption data relative to number of processors for the largest supercomputers.

Target values for metrics

Tour of Oakland Scientific Center/Trends and predictions for scientific computing

Howard Walter provided a tour of the Oakland Scientific Center prior to a discussion of requirements, trends, and predictions for scientific computing power requirements. Key points are:

- System projections in 1999-2000 were overstated. Last two procurements energy consumption were close to projections.
- Systems are currently air cooled and this could continue for some time. Need to provide flexibility to operate on air and/or liquid for the full cooling load as the industry is moving to liquid solutions. No limit to air cooling is known.
- DARPA is leading government efforts and has focused industry on more accurate estimates.
- Expected rack power loads are:

Name	MW (IT load)	Year	Cooling
N3E	0.8	2003	air

N5A	2.1	2006	air
N5B	4	2007	
N6	8	2009	
N7	18	2012	
N8	18	2015	

- DARPA projections for petascale facility will include 40-50 kW/rack; not air-cooled; 15,000 sf; 15-25 MW

- One new computer spec requires 55 degree air supply - lower than ASHRAE guideline allowable.

- One scenario is 35-40 MW racks in 10-20 increments

LBNL Data Center Research update

Bill Tschudi provided an update on LBNL research activities. See attached presentation. Key points:

- DC Powering demonstration - over 20 firms contributed equipment and expertise to demonstrate how DC power could be distributed at a facility level and at the rack level. The team selected 380 Volts as the distribution voltage. Servers can operate today with this voltage, however Intel and Sun Microsystems modified servers to bypass their AC-DC conversion for further efficiency gain. Overall a 10-20% energy savings over traditional AC systems results. Additional gains possible - e.g. variable speed drives on pumps/fans

- Air Management demonstration - An experiment in the Oakland facility was in progress. By completely isolating the "cold aisle" airflow could be reduced resulting in a six fold fan energy saving. By isolating the hot air return, a similar three fold fan energy saving was seen.

- Air economizer demonstration - LBNL is collecting data on contamination levels and humidity control in centers that use outside air for cooling compared to those that do not. Preliminary results show acceptable conditions in all centers.

- Benchmark results show a wide variation in performance for all systems/components in data centers. CRT design can utilize the results to target higher performing systems and components.

July 14, 2006

Industry trends - and projections of IT equipment load growth

- Don Beaty described the ASHRAE publications related to data centers. The latest ASHRAE trend chart shows up to 35-40kW per rack, 600 W/SF by 2015 as the highest possible load (fully loaded rack), actual will apply diversity of space and equipment
- Intel: customers were complaining about increasing heat levels so decided to limit processor power to 130 W. for foreseeable future and use multi-core technology to improve performance. Memory is becoming increasing heat source.
- Higher efficiency power supplies are being developed. DC powering would eliminate some conversion losses (heat).
- RAM is getting less efficient

IT energy efficiency and procurement strategies

The participants were asked to identify areas where energy efficiency of the IT equipment could be influenced and possible procurement strategies to move the market towards more energy efficient computing since well over half of the total power consumption is for powering IT equipment. NERSC pointed out that supercomputer procurements are typically a "best value" procurement that consider energy use as one of many criteria. As a result there was reluctance to consider many of the charrette suggestions summarized below:

- Procurement could include power and cooling equipment as a complete package
- Operations - need rack level measurement to monitor power uses
- Specify that IT equipment should meet ASHRAE thermal guidelines and power reporting and operate with those ranges
- Specify or select equipment with higher allowable inlet air temperatures (e.g. able to use outside air and cooling tower water (no chillers))
- Tie IT temperature monitoring (inside servers) to building controls and control building systems to maintain conditions inside the servers
- As part of the “Best Value” procurement consider sustained system performance per Watt
- DARPA requires energy performance metric (NRG*)
- Use life-cycle cost in value assessment (elect. cost, power, cooling equipment costs)
- Specify energy consumption budget in the RFP
- Require high-efficiency power supply architecture - including option for high voltage DC power

- Encourage development of liquid cooled systems
- Bottom up cooling (e.g. Blue Gene) may be better than hot and cold isle
- Optimize utilization of equipment (turn off what is not being used)
- Increase weight of life cycle cost/total cost of ownership evaluation in best value procurement
- Purchase IT and cooling equipment as a package
- Outsource and Service Level Agreement the computing requirements
- Consider SSP/Watt as part of best value evaluation
- Get energy input to DARPA to help formulate energy efficient strategy for federal super computers
- Lease to own

Barriers to requirements in procurement specifications aimed at improving IT equipment perform were discussed. Energy is considered as one of many factors when making a "best value" procurement. Adding requirements such as liquid cooling or more efficient power supplies may limit the number of systems that are bid. Super computers are built from standard industry components only at a higher density. It would be difficult and costly to move the market since the supercomputer market represents a small fraction of the overall market. It would be acceptable to signal the need for more efficient equipment but not to require it. Currently, purchase specifications do not include many requirements in order to maximize the number of bids. The minimum requirements include footprint sustained performance, effectiveness, and reliability. The evaluation criteria includes the cost of ownership (life cycle cost) and is normalized for 3 and 6 years. There is also a qualitative assessment of the systems. The primary goal is sustained computing (flops). Currently the flops/$ is improving more than the flops/watt or per square ft.

June 16:

Recap of Day 1

Baseline megawatts - Howard Walter revised IT loads and clarified that these were loads for NERSC computers only. Other computing and infrastructure loads are additive. Building and total loads below need to be developed by the design team (other program requirements not changed)

Year	IT Equipment	Building	Total
2009	19	10	29
2013	31	19	50
2015	37	26	63

Target reduction (%)			

The initial installation of 6 MW may require 55 degree (F) air.

Design Team Initial Concepts

The design team presented preliminary concepts to meet the program requirements (base case). See attached presentations (Bart McClelland - RMW Architects; Dan Smith and Charles Lee of Ackerman-Practicon)

Key points from the conceptual description are:

- Total building: 150k sf 50% for people; 50% computers
- 36 MW load
- 11,000 ton AC
- 8 x 1250 ton chillers (mechanical concept); 10 - 1250 ton chillers each with 1000 kVA load (electrical concept) Comment: this would represent an inefficient chilled water design.
- 5,690,000 cfm at 20 ºF delta T or 3,800,000 cfm for 30°F delta T
- 8' above ceiling
- 8-14' return chase
- 1st floor height 16-20 feet
- Like a class 10 cleanroom
- Water available in service space for local cooling
- 36 MW is 40 mVA
- Utility substation is not required; existing load is 20 MW on 90 MW station
- Rotary UPS systems do not need air conditioning; static UPS systems do but could probably be cooled outside the conditioned space with evaporative cooling
- Power factor of existing systems is 0.98; next system coming in is 0.92

Conceptual design discussion

The participants next critiqued the initial concepts and discussed issues relating to the major challenges outlined in the program. Key points from this discussion follow:

- ♦ Use of vane axial fans for central air systems is a high cost solution. Equivalent cleanroom solutions are $3,000-$4000/sf
- ♦ Added floors over and under the computer floor add cost
- ♦ Hybrid systems may be a solution - liquid cooled racks commercially available or custom designs. This would limit building cost but add cost later.

- Look at air path and building volume. Adding space that doesn't serve a purpose is costly.
- Need to provide flexibility for "old style" systems that do not fit in racks
- Complete isolation of hot and cold airstreams should be the goal. Isolating the hot aisle allows options for heat removal. Isolating the cold aisle can minimize airflow by matching server's needs and get away from overcooling the air.
- Higher delta T is best from an efficiency standpoint
- Many HVAC arrangements including "hybrid" systems are described in ASHRAE design guidelines
- It is impossible to accommodate any load and any cooling method. Consider local rack cooling and/or have IT vendor supply cooling.
- For continued air cooling, the width of the cold and hot aisles will need to increase to accommodate higher airflows and cooling loads, thereby decreasing the useable floor area
- Several design concepts with varying degrees of liquid and air cooling should be developed and cost estimated.
- Likely scenario is starting with air cooled systems, transition to hybrid systems, and then mostly liquid cooled systems.
- ASHRAE thermal guidelines should be the basis for design
- A case study from 1999: 60 MW facility with 18" chilled water piping distributed under the floor w/taps for every other rack
- Ability to accommodate both liquid and air systems is likely to be too costly but providing space for both and considering future tie-in should be possible.
- Distributing power at higher voltages is more efficient. Even if DC systems are not pursued, options to distribute 480 V AC to the rack should be provided. Transformers and other heat producing equipment could be located outside or in areas with minimal conditioning.
- Flexible partitions to zone area of the center
- On-site generation where waste heat could be used for adsorption chillers
- Use of outside air economizers coupled with monitoring and appropriate filtration

Brainstorming radical ideas

The participants, in five groups were asked to come up with three energy efficient approaches or one idea that had not previously been discussed. The results include:

July 14, 2006

- Thermal chimney - Take advantage of natural ventilation and prevailing winds.

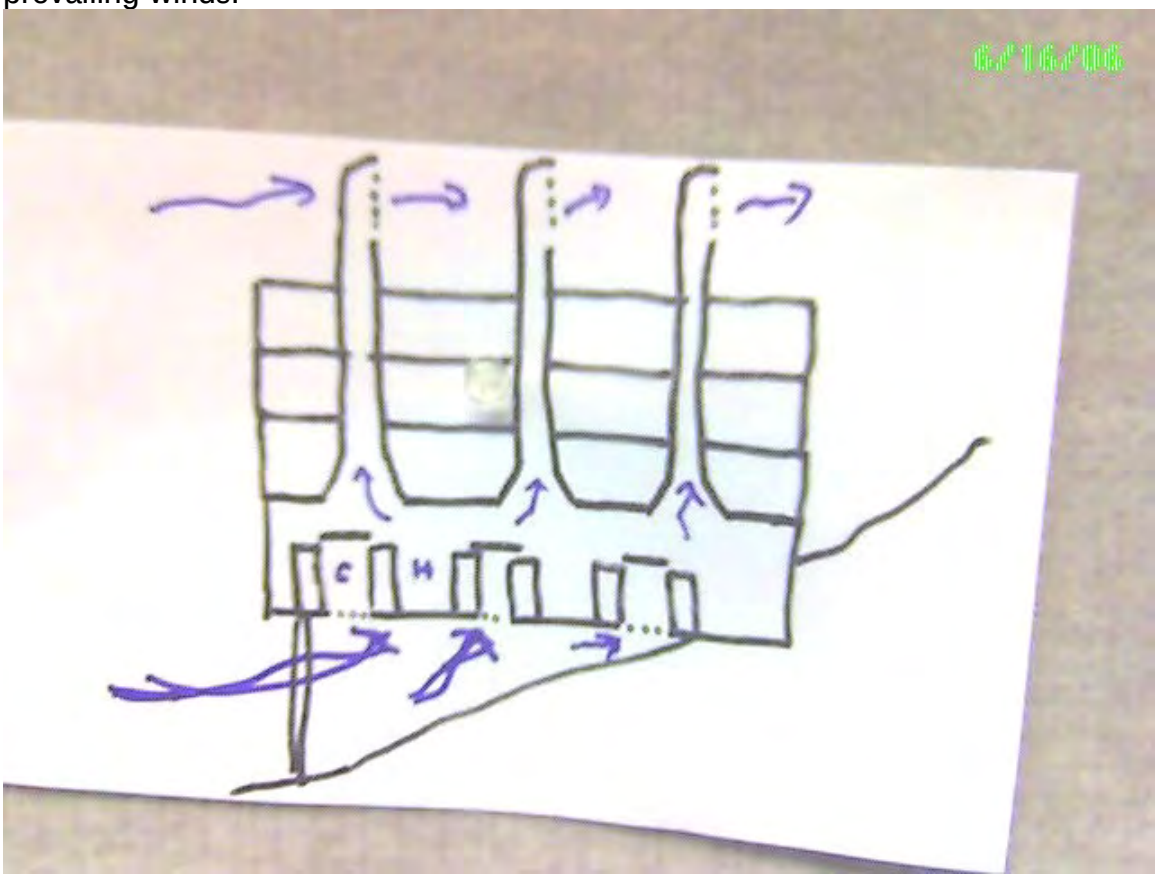

Considerations: Adverse weather lasting days; construction nearby; Diesel particulates from nearby access road; filtration requirements; possibly add humidification under the building; reliance on air cooling

July 14, 2006

- Heat recovery

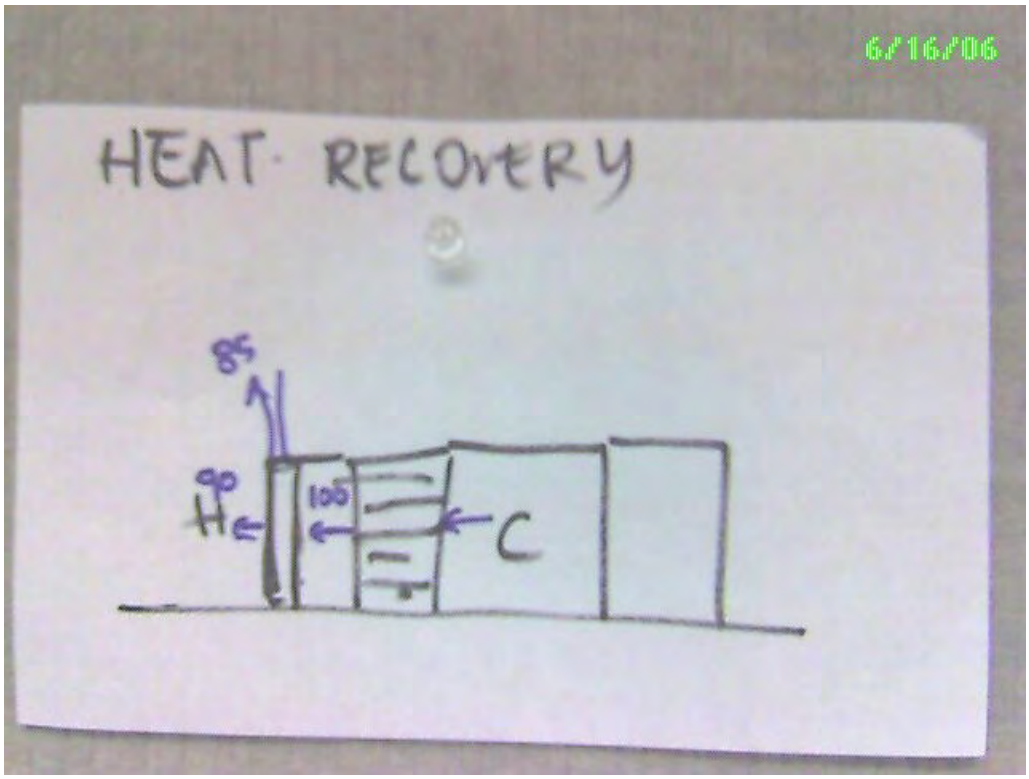

Use the waste heat to preheat outside air and for a central hot water loop to heat adjacent buildings or part of the UC campus. Alternatively the hot air could cascade to the office areas for heating.

July 14, 2006

- Floating tub

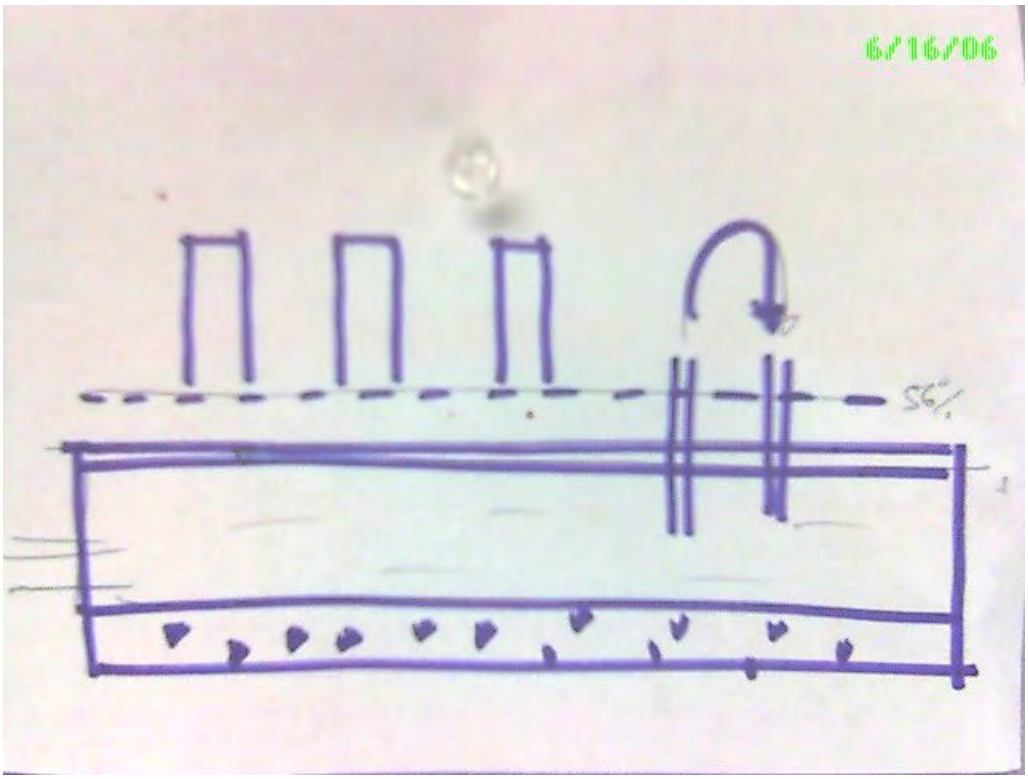

A heat sink consisting of a tank of water with a membrane surface that has pipes penetrating it that allows off-peak cooling of water.

July 14, 2006

- Outside air augmented with fans

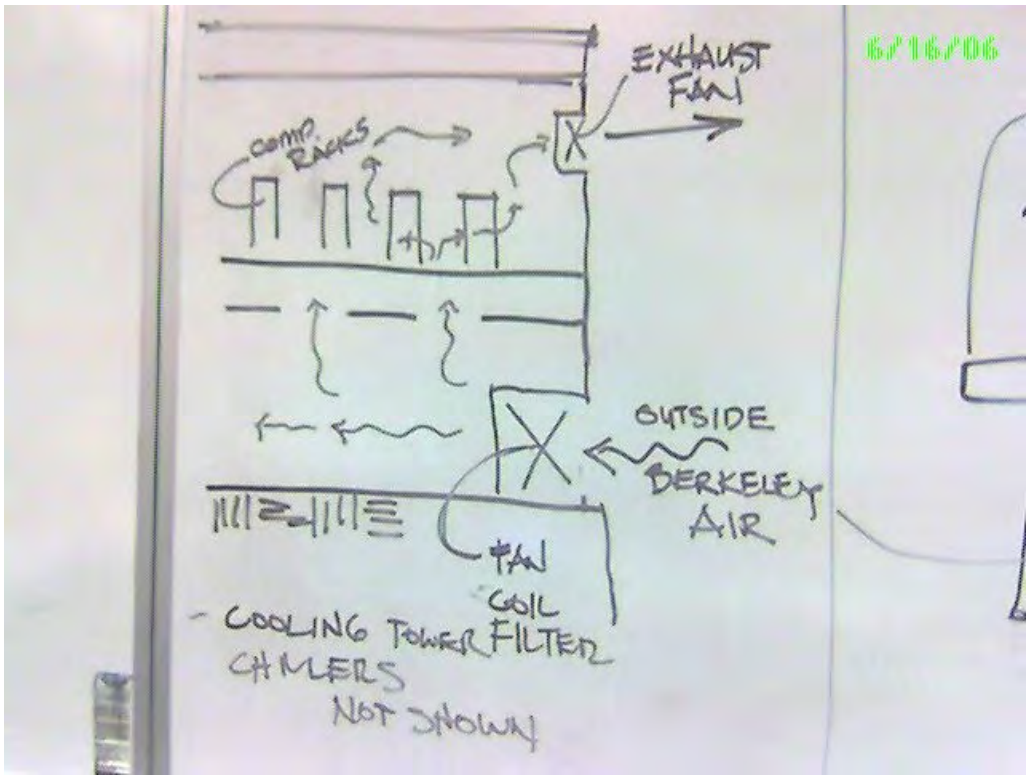

Similar to the thermal chimney only augmented by a fan-coil and filtration. Once through so no return path.

- Use of modular commercial air handlers in service floor

Conversation has moved from how do we do this as a showcase of efficiency to how can we accommodate the given program needs at all (within budget)

<u>Criteria for evaluating options</u>

The following criteria were selected for evaluation of design options:

1. First cost
2. Life Cycle Cost
3. Energy
4. Flexibility
5. Reliability

July 14, 2006

Discussion, evaluation, and prioritization of alternatives

A brief discussion of alternatives resulted in a recommendation to move forward in an evaluation of at least three options. Don Beaty's suggestion was to sufficiently fund design development for different options to obtain cost estimates and evaluate life cycle costs. An evaluation matrix was suggested by Mark Hydeman as shown below:

Base: Central air (LLNL TSF). Economizer. Limited zoning. Service Floor.
Option 1: Central air (LLNL TSF). Economizer
Option 2: No Central Air (localized cooling). Water side economizer
Option 3: Hybrid. 150 w/sf 15-20 KW/rack air cooled, remaining localized cooling. Economizer

Criteria	Weight	Base case (LLNL TSF)	Option 1	Option 2	Option 3
1. First cost		0	-?	+	+ (?)
2. LCC		0	+	+	+
3. Energy		0	+	+	+
4. Flexibility		0	0	-	+
5. Reliability		0	0,-	0,-	-

Building performance metrics and performance targets

The participants identified the following as key metrics for evaluating the supercomputer building performance:

Building Performance Metric	**Group Estimated "Standard"**	**First Draft Target**
KW/ton plant including fans	1 – 1.5	.7 - 1.0
KWh/annual ton hours	1.2 – 1.7	0.3
IT to total	.6 (.74 NERSC)	0.9
IT to HVAC	1.7	9
Reconfigure time/cost per system	6 months	1 month
Spatial efficiency SF IT/total of computer floor	15 – 20%	25%
W/CFM	0.5	0.2
W/GPM	15	10
Rack cooling index (measure of compliance with temp spec)	40-100	100

Power conversion efficiency Substation to rack w/o UPS	0.85	0.95
Heat recovery utilization	0	TBD

Summary

To be completed….