CALIFORNIA
ENERGY COMMISSION

Energy Research and Development Division

FINAL PROJECT REPORT

Powernet: Cloud-Based Method to Manage Distributed Energy Resources

Gavin Newsom, Governor
April 2021 | CEC-500-2021-030

PREPARED BY:

Primary Authors:
Marie-Louise Arlt, Gustavo Cezar, David Chassin, Claudio Rivetta, and Laura Schelhas

SLAC National Accelerator Laboratory
2575 Sand Hill Rd.
Menlo Park CA 94025
(650) 926-8529
https://www6.slac.stanford.edu

Contract Number: EPC-15-047

PREPARED FOR:

California Energy Commission

Chie-Hong Yee-Yang
Project Manager

Fernando Piña
Office Manager
ENERGY SYSTEMS RESEARCH OFFICE

Laurie ten Hope
Deputy Director
ENERGY RESEARCH AND DEVELOPMENT DIVISION

Drew Bohan
Executive Director

ACKNOWLEDGEMENTS

The authors acknowledge the partners and collaborators who made this project possible.

- California Energy Commission: The California Energy Commission provided primary funding to make this project possible, and Commission Agreement Manager Jamie Patterson and Project Manager Chie-Hong Yee-Yang provided strategic project stewardship.
- SLAC National Accelerator Laboratory and Stanford University: SLAC and Stanford University provided the administrative support, organized the recruitment and selection of participant residents in the project, conducted the purchasing of the equipment and helped with legal issues. Additionally, the use of the Laboratory facilities was important for the development and testing of the hardware and software of the system.
- City of Fremont: The City of Fremont helped and participated in the advertisement of the project to the community.
- Sonnen, Inc.: Sonnen provided the battery energy storage system and helped in part with the selection of possible alternative sites to deploy the system.

PREFACE

The California Energy Commission's (CEC) Energy Research and Development Division supports energy research and development programs to spur innovation in energy efficiency, renewable energy and advanced clean generation, energy-related environmental protection, energy transmission and distribution and transportation.

In 2012, the Electric Program Investment Charge (EPIC) was established by the California Public Utilities Commission to fund public investments in research to create and advance new energy solutions, foster regional innovation and bring ideas from the lab to the marketplace. The CEC and the state's three largest investor-owned utilities — Pacific Gas and Electric Company, San Diego Gas & Electric Company and Southern California Edison Company — were selected to administer the EPIC funds and advance novel technologies, tools, and strategies that provide benefits to their electric ratepayers.

The CEC is committed to ensuring public participation in its research and development programs that promote greater reliability, lower costs, and increase safety for the California electric ratepayer and include:

- Providing societal benefits.
- Reducing greenhouse gas emission in the electricity sector at the lowest possible cost.
- Supporting California's loading order to meet energy needs first with energy efficiency and demand response, next with renewable energy (distributed generation and utility scale), and finally with clean, conventional electricity supply.
- Supporting low-emission vehicles and transportation.
- Providing economic development.
- Using ratepayer funds efficiently.

Powernet: Cloud-Based Method to Manage Distributed Energy Resources is the final report for Contract Number EPC-15-047 with SLAC National Accelerator Laboratory. The information from this project contributes to the Energy Research and Development Division's EPIC Program.

For more information about the Energy Research and Development Division, please visit the CEC's research website (www.energy.ca.gov/research/) or contact the CEC at ERDD@energy.ca.gov.

ABSTRACT

This report describes implementation of a Cloud-based control strategy for managing distributed energy resources on California's electric transmission grid and electric distribution network. It is based on a connection of information networks with power networks to control distributed energy resources. The strategy is based on transactive energy-control techniques, which this project demonstrates to be a feasible and simple way to operate distributed energy resources.

This project assesses distributed energy resource control through a large-scale simulation model of the state's electric system, including detailed modeling of distributed energy resources in residences. These distributed resources are primarily solar photovoltaic panels, home-storage systems, electric-vehicle chargers, and heating, ventilation, and air-conditioning systems. distributed energy resources are controlled through an auction-level market for agents participating in the market by distributing distributed energy resources DER output to the electricity network. A mathematical model of the system was developed to document the essential characteristics and behaviors of the electric distribution network. It was essential to assess and quantify different characteristics of the system including stability, stability robustness, and rejection of perturbations. Results from the large-scale simulation show the impact of multiple distributed energy resources on the system under different penetrations of renewable resources, as well as various operation conditions and capacity limits in the distribution network. Additionally, it allows both residents and utilities to compare renewable resource benefits. A small deployment of the hardware and the software in the City of Fremont has proved the feasibility of this Cloud-based strategy and operation of the communication and information exchange between distributed energy resources and the Cloud.

Keywords: Distributed energy resources, DERs, transactive energy control, cloud-based management of DERs.

Please use the following citation for this report:

Rivetta, Claudio, Marie-Louise Arlt, Gustavo Cezar, David Chassin, and Laura Schelhas. 2021. *Powernet: Cloud-Based Method to Manage Distributed Energy Resources*. California Energy Commission. Publication Number: CEC-500-2021-030.

TABLE OF CONTENTS

LIST OF FIGURES

Page

LIST OF TABLES

Page

EXECUTIVE SUMMARY

Introduction

California's aggressive goals and strategies to address climate change have made the state a world leader in easing this escalating crisis. State initiatives and mandated environmental targets include efforts to slow the pace of climate change (mitigation) and increase resiliency in the face of its inevitable and fast-approaching impacts (adaptation). A key mitigation measure is to increase electricity generation from renewable resources like solar and wind, which in turn reduces the fossil-fueled greenhouse gas emissions that significantly drive climate change. Adaptive functions include providing emergency power for critical services during disasters such as large storms and other meteorological anomalies intensified by climate change.

The increasing number of distributed energy resources connecting to California's electric grid raises operational challenges for the network operators charged with balancing the state's electric transmission and distribution systems. Electric vehicles, heat pumps, flexible loads, storage units, and distributed generation (including renewable generation devices such as wind turbines and photovoltaics), are generally known as distributed energy resources. To mitigate these challenges, these resources must more actively participate in grid coordination and operations.

There are many technological barriers to achieving California's ambitious energy goals. These include streamlining the integration of distributed energy resources, developing their distributed and hierarchical communication and controls, developing the codes and standards that support their capabilities, and building and testing market structures that provide cost-effective ways for grid-edge devices to successfully deliver their products and services.

Challenges with the state's electric grid attract much public attention, and solutions run the gamut from developing new control methods for individual-component operations to radical rethinking of system operations as a whole. Traditionally, the electric grid has operated securely through centralized control. However, with current coordination requirements between transmission and distribution system operators, it is not yet possible to control large numbers of distributed energy resources since the grid control system is centralized by design and cannot integrate them on a meaningful scale.

Project Purpose

This project was built on the guiding principle of managing distributed energy resources by connecting information networks with the greater statewide electric transmission and distribution systems. The primary goal was to demonstrate the feasibility of transactive energy control techniques and simple strategies to control distributed energy resources. Other goals explored the performance and limitations of the control technique, evaluated the impact of residents' system performance preferences, assessed the operation of the decentralized system, and evaluated future implementation and the transition to greater numbers of residences.

The project addressed the control of distributed energy resources by creating a large-scale simulation model of the state's transmission and distribution electric systems, including

detailed modeling of distributed energy resources installed in residences. This novel system was controlled by an auction-level market that maximized agents' benefits in the market and for distributed energy resources. The project used hardware and software to validate the communication and information exchanges between distributed energy resources and the Cloud coordinator.

Project Approach

The Grid Integration, Systems and Mobility group at SLAC National Accelerator Laboratory (formerly the Stanford Linear Accelerator Center, a United States Department of Energy laboratory operated by Stanford University) led the project, with Stanford University as the primary collaborator. The Grid Integration, Systems and Mobility group managed the project, coordinated all subcontractors and assisted with organizing equipment and storage unit purchases and recruiting residents in the City of Fremont.

This project had three primary components, each critical for developing a large-scale project:

1. Define the mathematical (or analytical) model of the electrical system describing the DER dynamics controlled by an auction market.
2. Develop a high-level simulation platform to analyze a large-scale system capable of accommodating different amounts of renewable energy resources and operating conditions.
3. Employ small-scale lab and field use of the hardware and software required to demonstrate operations.

Twelve residents participated in the project. Equipment installed in the residences included a power monitor system, a smart thermostat, and a set of energy storage units. The project demonstrated a cloud-based distributed energy resource management system and considered distributed energy resource integration with legacy systems to monitor and control all functions within communities. This project was essentially the first step in the design and use of the hardware and software (firmware and algorithms) required to manage demand response by controlling renewable energy, storage, and loads at the appliance level by an auction market. The project team then evaluated an electricity system simulation of a residential area that included large number of photovoltaic sources, storage units, and electric vehicles.

Project Results

Project results were evaluated in the simulation using a large-scale mockup tool and simplified models (reduced models) representing residences and devices. The outcome summarized the stability and performance of an aggregated system of between 1,000 and 10,000 houses for different scenarios, markets, and levels of renewable-resource numbers. In this first step to evaluate the hardware and firmware at the residence level, only 12 houses were enrolled in testing the communication and performance of the control scheme. Combining both parts of the project created a good estimate (or extrapolation) of performance results from the 1,000 to 10,000 houses, reducing the risk of validating this control technique for wider populations in the future.

One of the results of this project is the development of a mathematical model of the feedback system, controlled by the market, was developed to define system stability and performance. It also allowed, through bidding into the market, evaluation of the impact of residents' preferences on the system. The advantage of conducting this analysis through analytical models provided the option of quantifying the system's stability and performance limitations. Those limits are otherwise cumbersome and difficult to determine, requiring multiple time-domain simulations to evaluate different scenarios and operation conditions. This model therefore provided valuable insight into market design at the distribution level.

Applying the system-performance results to a larger-scale system was determined from time-domain simulations. A simulation platform that allowed integration of the physical and economic layers of a distribution system was developed for this project's analyses. The goal was to expand market simulation capabilities to more fully understand the market context on a larger scale and study a local energy market's ability to control load flexibility and distributed energy resources at the distribution level. The simulations evaluated cases with 50 percent or higher penetrations of intermittent solar generation while minimizing costs and addressing challenges of security and economic scalability.

The project's proposed technology was applied to a small number of houses in the City of Fremont to test communication between different devices, the Cloud, and algorithms. The tests yielded similar results to a test conducted by Pacific Northwest National Laboratory with communication times between the Cloud and devices negligible (less than 0.3 of a second) to the market's 5-minute operational clearing time. The devices reverted to normal mode, a non-market control mode, when communication was interrupted between devices and the Cloud. The proposed system is scalable because the Cloud's auction-market function is the only element that must be augmented for greater distributed energy resource participation and flexible market loads. Because residences only communicated with the Cloud, defining additional parallel connections did not delay communications. A final point is that the platform was able to integrate devices with public application programming interfaces and open access to the data and exclude units with closed application programming interfaces.

Technology/Knowledge Transfer/Market Adoption (Advancing the Research to Market)

The cloud-based technology developed to create the Powernet platform could be transferred almost directly to a commercial product. Its control scheme would, however, require more interactions to fulfill utility requirements for DER management on a large scale. Examples of these requirements include operations during emergency conditions (islanding), integration with utility control rooms, and distribution system voltage management.0.

Follow-on work building on these Powernet results has already been funded by the United States Department of Energy and the project, the Transactive Energy Service System, is an electric utility planning and operation platform that supports the design, use, and operation of a price-based transactive energy system.

The mathematical model of the auction-based controller developed in this project is novel and can be categorized as a hybrid or variable-control structure that defines methods to assess system stability and performance. Aspects of this research are currently being submitted to

publication in technical and industry journals (Applied Energy, Elsevier Ltd.) and they are under revision.

Benefits to California

This project marks the beginning of solutions to those challenges by developing new control methods for individual operation through radical rethinking of current electric system operations. The ultimate impact of adopting these technologies will benefit the state by lowering energy costs, increasing clean, renewable resource generation, and reducing fossil-fueled greenhouse gas emissions and other pollutants. To allow massive penetration of distributed energy resources and flexible loads, it is necessary to address these challenges in incremental steps to systematically understand the new dynamics of the emerging power system and adopt both new technologies to coordinate distributed energy resources and the training required for operators who control the grid.

Recommendations

A final planning and operation platform based on auction markets must be able to answer the multiple questions that will eventually spur development of the features that will pave the way for the platform to become a robust future tool. One of these steps is integrating the platform with electric utility control rooms, an issue currently being addressed in the Transactive Energy Service System Project. Additional research points needed to extend this platform are to:

1. Analyze control of reactive energy and voltage limits in the electric distribution network for high penetration potential for distributed energy resources.
2. Tailor the system to accommodate emergency conditions, or islanding. To maximize system flexibility, combine the transactive control strategy with other control strategies to reconfigure the distribution network during emergencies.
3. Evaluate other products (for example, storage energy) for transactions beyond the capacity and power used in the present strategy.
4. Conduct a larger sample after previous points have been addressed to effectively demonstrate the entirety of the platform’s features.

CHAPTER 1: Introduction

Background

California has aggressive goals and strategies to slow the advance of climate change, including innovative ways to reduce the greenhouse gas emissions that largely contribute to this environmental crisis. An important component of the state's mandated environmental targets is to transform the electric grid into a hybrid system able to accommodate large penetrations of both decentralized intermittent resources such as solar photovoltaic (PV); batteries; combined heat and power and centralized intermittent generation (solar PV and wind power); storage of two-way power flows; and the overall integration of sensing, communication, computing and controlling load; generation; and storage. California is a national leader in establishing and advancing ambitious energy policies and practices to generate, transmit, distribute, and deliver clean, non-polluting energy and will contribute nationally to the grid transformation, as the U.S. as a whole faces challenges caused by ever decreasing costs of intermittent carbon-free generation and storage; increasing cyber-physical security attacks; aging of electric grid infrastructure; emergence of de-carbonization policies; and increasing extreme weather-related events. There is a growing need to coordinate these resources to minimize costs; increase consumer Quality of Service (QoS); preserve physical, logical, and financial stability; and offer services to the grid in an economically sustainable way.

Powernet, a cloud-based method to manage distributed energy resources (DERs), was built on the principle of connecting information networks to the power grid. Powernet implements a feedback control system based on a retail-level market for managing DERs and integrates those assets with wholesale electricity markets. By leveraging both simulation results and lessons learned from small-scale deployments, this project provides valuable engineering information to extend DERs controlled by retail-markets into large communities. This report provides background and details about implementing the project, and includes:

- Project Overview (Chapter 2): Provides a description of the project, organization, implementation and validation plan to correlate the mathematical model, simulation platform, and residences participating in the project. Chapter 2 also analyzes the project's goals and limitations.
- Control of Distributed Power Resources by Retail Markets (Chapter 3): Analyzes how a retail market can control the energy and DERs in an electricity system. Defines a mathematical model for the complete system, including the market as a control element. More details about the analytical models are presented in Appendix A.
- Retail-Level Markets (Chapter 4): Presents details about the retail market and its implementation as a controller of the system and its operation.
- Simulation Platform (Chapter 5): Provides a complete description of the simulation platform, combining GridLAB-D tools and Python routines. Appendix B gives details of the models and the data used to run the simulations.

- Results (Chapter 6): Presents different cases of the market controlling large assets of renewable resources through simulation. Some results are validated with measurements using the resources installed in the laboratories and residences. Appendix C describes the design and implementation of the units to link resources and the Cloud and development of the complete system deployed at SLAC National Accelerator Laboratory and Stanford University, and residences in the City of Fremont.
- Conclusions, Recommendations, and Extension to Larger Demonstrations (Chapter 7): Summarizes major findings from the project and evaluates next steps for a larger demonstration to validate and replicate simulation results.

The remaining sections of this chapter provide a project description, a description of the purpose and need for the project, and a description of the project team.

Brief Project Description

The project was built on the principle of creating a system to connect information networks with the statewide electricity network. To assemble this system, a physical architecture was used, as depicted in Figure 1 and Figure 2. Local measurements of the voltages and currents per sub-circuits at each house, together with a resident's preferences, were used to effectively control DERs.

Figure 1: New Physical Architecture of the Electrical Grid

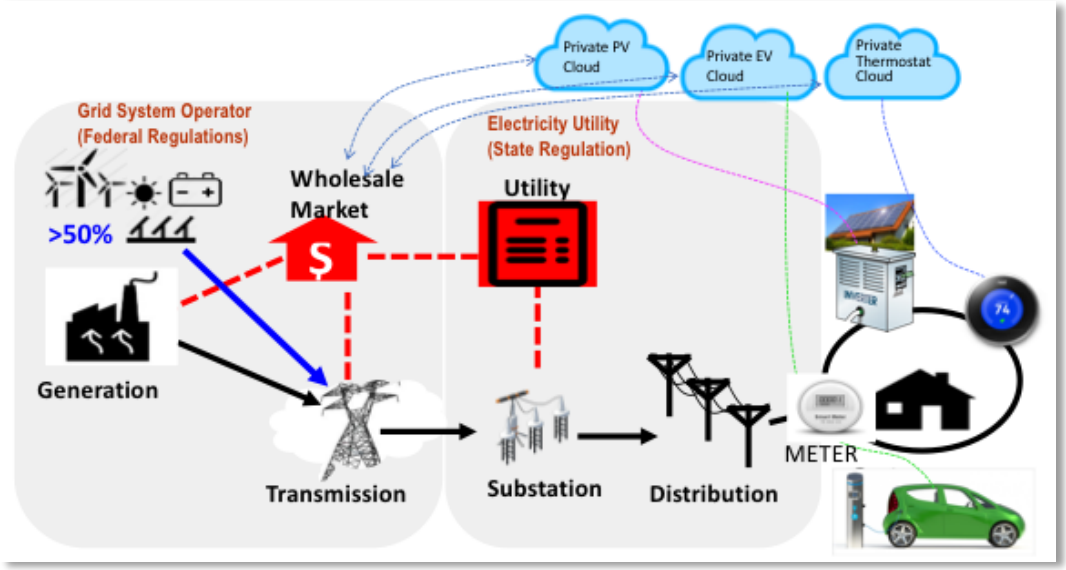

Source: SLAC National Accelerator Laboratory

Figure 2: Cloud-Based Control of the Electrical Grid

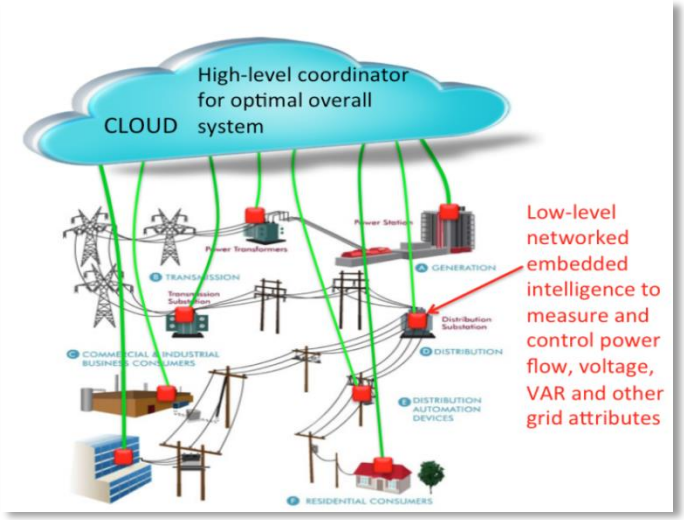

Source: SLAC National Accelerator Laboratory

A decentralized topology was the control topology for this project. In this architecture, a home hub (HH) aggregated behind-the-meter resources and communicated that information to a cloud coordinator (CC), as shown in Figure 3. The cloud coordinator was the core of the Powernet platform and functioned as the market-based controller of the feedback system that extended to the electrical network. The CC controlled the distributed resources participating in the retail energy market.

Figure 3: Powernet Topology

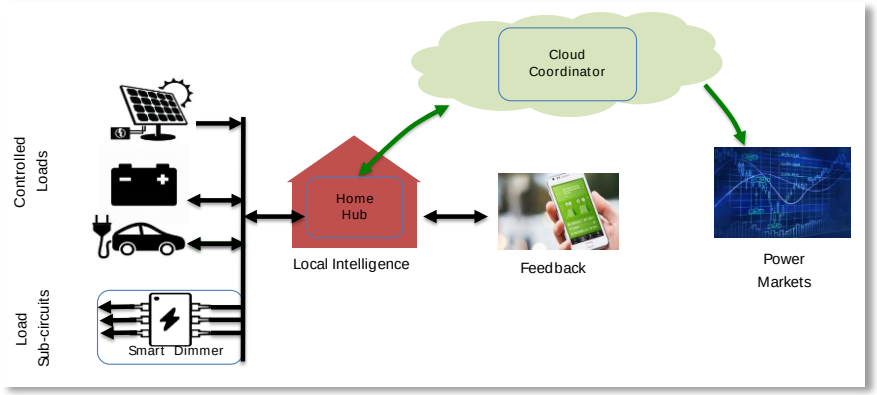

Source: SLAC National Accelerator Laboratory

The main features of this topology were to minimize the information exchange between the cloud coordinator and the controllable resources, maintain operation if communication with the cloud coordinator is lost, and minimize the requirement that residents set preferences.

Purpose and Need

The widespread and general acceptance of this strategy, its deployment, and its use to control DERs in large cities involves multiple steps to both validate the features and understand their limitations. This project is the first step in that process. This project moved the state-of-the-art from the system's basic concepts and mathematical model to an analysis through simulation of a large system incorporating more than a thousand houses. To validate communication and

control of DERs and hardware, multiple devices were also deployed in two laboratories at SLAC National Accelerator Laboratory and Stanford University and in houses in the City of Fremont (Figure 4). This minimal deployment allowed the project team to evaluate the performance of the hardware and algorithms used to link the cloud coordinator with the devices installed in participating houses. It provided an avenue to test the features and characteristics of the market controlling the system, and complemented results of simulations extended to larger-scale areas.

A second stage is being planned based on this experience. The United States Department of Energy has funded initial use of the Transactive Energy Services System (TESS) platform, which is based, in part, on the experience and technical capabilities developed and demonstrated in this project. The initial version of TESS is expected to include a retail market system to determine the real-time price of electricity for customers with controllable loads, batteries, rooftop PV, and electric vehicles.

Figure 4: Aerial View of the City of Fremont

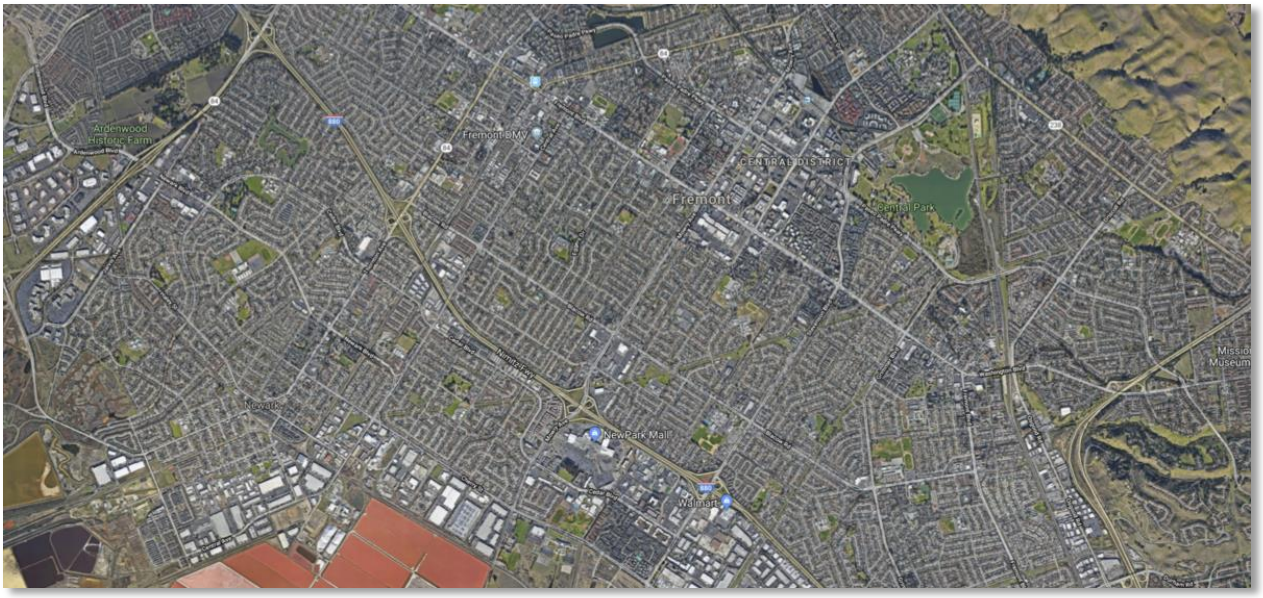

Source: SLAC National Accelerator Laboratory

Project Objectives

The objectives of this project were to:

1. Show penetration of 50 percent or higher of intermittent DERs while minimizing cost and addressing security and economic scalability challenges.
2. Develop mathematical dynamic models of the distribution electrical system controlled by an auction market.
3. Expand market-simulation capabilities to more fully understand the market context of these systems in larger scales.
4. Demonstrate cloud-based DER management systems and consider their integration with legacy systems to monitor and control all functions within energy-smart communities, including smart inverters and smart meter functions that handle DER intermittency issues.

5. Test low-cost plug-and-play integration platforms for DERs installed in buildings.

Project Benefits

The increasing number of DERs connected to power systems raises operational challenges for network operators and general management of those assets. There are many technology barriers to realizing California's mandated energy goals. These include streamlining the integration of DERs, developing distributed and hierarchical communication and controls for DERs, developing codes and standards to support these capabilities, and building and testing the market context that will ultimately allow grid-edge devices to successfully and cost-effectively deliver their services. The control and coordination of large numbers of DERs require innovative approaches.

This project is the beginning of solving these challenges through radical rethinking of system operations and development of new control methods for individual components. The ultimate impacts of applying these innovative technologies will benefit the state and the nation by lowering energy costs, increasing reliance on clean, renewable energy sources, and reducing the greenhouse gas emissions that are major contributors to the existential global-warming crisis. Yet to introduce a massive penetration of renewable energy and flexible loads, it is necessary to first address the challenges in small steps to understand the dynamics of the new power system, the technologies necessary to coordinate and deploy DERs, and the training required for operators who control grid operations.

Project Team

The team was composed of researchers, staff, and students of SLAC National Accelerator Laboratory and Stanford University (SU). Their role was to formulate and solve problems associated with the design of the system and model and develop a simulation platform based on GridLAB-D. Procurement of the devices and supervision of their deployment and testing were conducted by SLAC and SU staff.

The City of Fremont's sustainability office helped recruit the residents who participated in this demonstration project.

In this first stage, 12 residents participated in the project. Equipment installed in the residences included a power monitor system, a smart thermostat, and a set of energy storage units.

- Smart Thermostat: Twelve Nest units were installed in all participating residences.
- Power Monitor System: Twelve eGauge units were installed in all of the residences. Each unit monitored power in the main circuit branches of the house.
- Battery System: Four batteries with a capacity of 10 kilowatt-hours (kWh) were purchased from Sonnen, Inc., and installed in selected residences.

CHAPTER 2:
Project Description

Introduction

The platform architecture is shown in Figure 5. A home hub (HH) managed resources behind the meter and connected them to the cloud coordinator (CC). The HH enabled every residence to operate either connected to or unconnected from the electric grid. It minimized information exchanges and had built-in fail-safe and security mechanisms. It controlled solar, batteries, local generation, and smart loads in the participating houses. The HH acted as an agent and processed all information measured in the houses, including resident preferences, and communicated with the CC. The control algorithm of the system resided with the CC. It was based on an auction market that maximized an individual agent's benefits and the aggregate benefits of both producers and consumers. This market-based controller also selected the participation or connection of resources to the grid. Powernet adapted to different distribution systems while implementing efficient power-sharing in the electric network. Powernet provided the platform that utilized the net-load control that provided grid services in a cost-effective manner and enabled buildings to simultaneously reap the benefits of dynamic pricing and revenues from grid services for the assets they chose to share. Powernet is both open-source and open-design to encourage adoption and standardization.

In the long term, smart homes are expected to improve occupant comfort, reduce energy use, and reduce energy costs. This vision is far from fully realized, however, partly due to hardware costs and performance limitations and partly due to a lack of financial incentives for the key resources required for continuous system operation (Rasch 2013).

Figure 5: Powernet With Market System Architecture

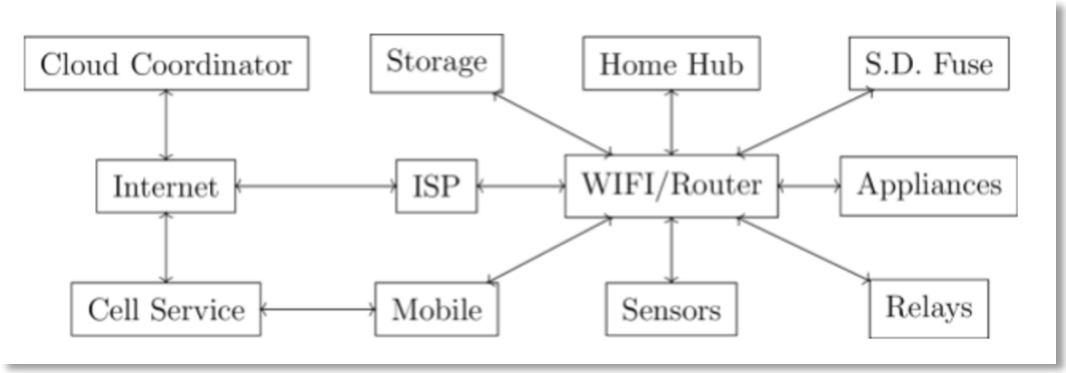

Source: SLAC National Accelerator Laboratory

Background

Increasing the number of DERs connected to power systems raises operational challenges for network operators. DERs need to become more actively involved in grid coordination and operational tasks. The control and coordination of large numbers of DERs require innovative approaches. Powernet systems contributed several significant innovations to this project. First, the integration of market mechanisms, controls, and power electronics enabled the novel

functionality of stable connections and disconnections to and from the grid, in addition to local and global power sharing and grid services. Second, based on a power-simulation tool, a market strategy was developed and integrated, creating an improved platform that enabled both the design and performance of subsystems and the system as a whole. This layered structure enabled operators to apply Powernet to a variety of grid purposes and service offerings, with the assurance that an economically viable operating point would be maintained. In addition, the system was robust and secure by design. Finally, open-source standards were adopted and an open protocol was used for the platform, which enabled the scalable engagement of future devices.

There are many technology barriers to achieving California's ambitious energy goals including streamlining DER integration, developing distributed and hierarchical communication and controls for DERs, developing the codes and standards that support these capabilities, and building and testing the market context that will ultimately allow successful, cost-effective ways for grid-edge devices to deliver their services. This project, with industry engagement, offers solutions to these barriers by developing required hardware, software, and tools.

System Architecture

The physical and platform architecture of the project is shown in Figure 3 and Figure 5. Two main hardware and software components were the HH, which aggregated resources behind the meter and connected them to a CC, and the CC, which operated as the market-based controller of the feedback system. Both elements are described in more detail in this chapter. Implementation of the simulation platform is addressed in Chapter 5, and its application in the laboratories and residences is presented in Appendix B.

Cloud Coordinator

The CC is the core of the Powernet platform. This is where HH's provide information and get set points to operate locally, based on services the platform is required to perform. The primary functions and requirements of the CC are:

- Registering and managing HHs.
- Detecting and mitigating anomalies.
- Controlling resources on the electrical grid.

At its core, the CC is a "Software as a Service" (SaaS) cloud- and subscription-based system. This core is hosted in a cloud-service environment where communication transpires between various instances within this environment. Breaking the platform down into multiple small instances additionally simplifies development and allows easy scalability.

The main algorithm in the cloud coordinator controls responsive resources on the electrical grid, based on an auction market. Each responsive device participating in the network communicates bidding requests or offers through its HH to the Cloud. The real-time market in the CC algorithm defines, at the clearing time (for example five minutes), the price and the allocation of resources (sources and loads) that satisfy the supply-demand power equilibrium. The result of these resource connections to the electric distribution system is then communicated back to the residences via the HH. An additional function of the cloud

coordinator, operating as a resource aggregator, is to communicate with the bulk system operator to receive prices and submit load forecasts or bids, if any.

From an infrastructure point of view, the platform consists of three different environments; Dev, Test, and Production are clearly separate between the operating regimes. This separation also advances the platform's long-term development. Figure 6 shows the CC infrastructure.

Figure 6: Cloud Coordinator Software Architecture Diagram

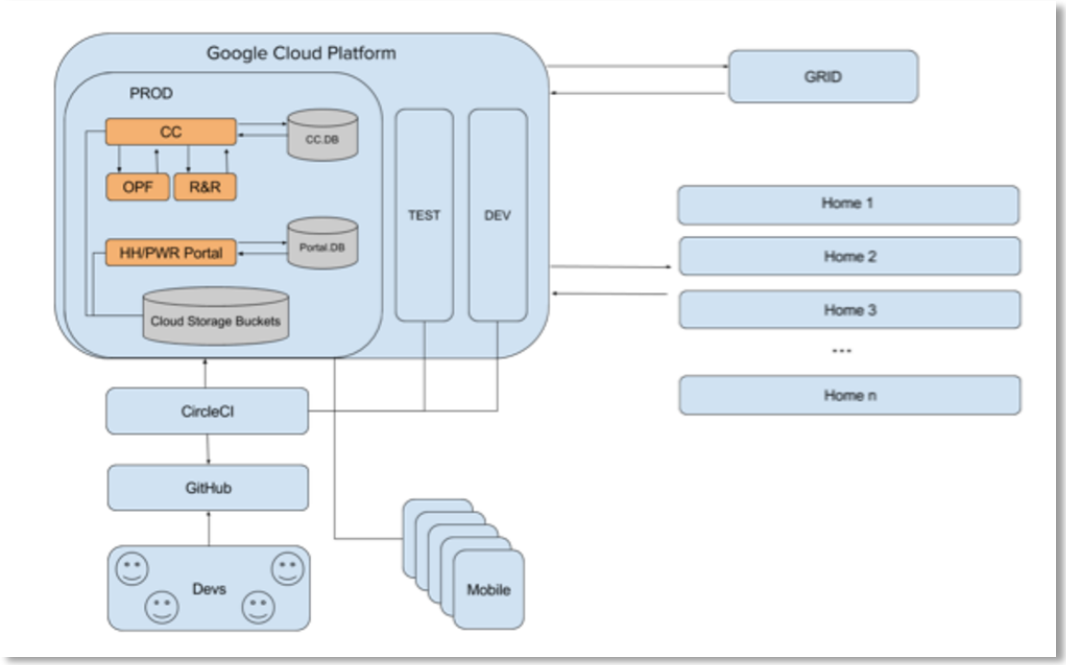

Source: SLAC National Accelerator Laboratory

Home Hub

To leverage capabilities of the grid edge resources for market participation, and to provide reliable and resilient power while preserving asset-owner privacy and comfort, a system with local intelligence capable of integrating this information and meeting required communication needs must be in place. This HH system connects to the Internet and acts as a client to a back-end system in the Cloud.

The HH system provides intelligence to resources at the edge of the grid (e.g., home, commercial, and industrial facilities). (Cezar, et al. 2020). Home Hub key development functions include:

- Integrating all home and facility asset data and states and conducting real-time measurements of main electrical variables and devices in participating houses.
- Measuring voltage, current, and power at device and sub-circuit levels in the houses.
- Sending and receiving information to and from the CC.
- Accounting for user preferences and comfort when setting bids.
- Performing local coordination of controllable assets that follow specific set points (such as price and power) sent by the CC.
- Providing the consumer user interface.

The HH concept within the scope of the Powernet platform was described in an earlier publication (Radavanovic, Rajagopal and Kiliccote 2016), and conceptually its user and grid benefits remain the same. However, there are important changes to its hardware component, its implementation, and its basic functionality. The HH hardware described in the earlier publication (Radavanovic, Rajagopal and Kiliccote 2016) was embedded in an inverter-based design, with both DC and AC power ports. In this project two implementations were evaluated and tested. These two options were possible because the HH functional requirements did not require either a large storage capacity or high computational power to solve complex algorithms. The Home Hub V1.0 – Local Intelligence, represents a case where the HH is fully implemented in a microprocessor unit installed in each participating residence, while The Home Hub V2.0 – Cloud Intelligence, shares the implementation and resources between commercial units installed in the residences and the Cloud. Descriptions of both HHs follow.

The Home Hub V1.0 – Local Intelligence

A simple, low-cost and small-form factor hardware device is highly desirable in allowing the large deployment of HHs into electric distribution systems. HH hardware needs to have the internet capability to send and receive information to and from the Cloud. It is also necessary that it handles all communications with local devices, wired and wireless. Battery and solar inverters, sensors, relays, and HVAC systems are a few examples of devices that can be wire-connected to an HH. The device also needs to sense currents and power in a house's main subcircuits. Finally, HHs need to be flexible enough to interface with different protocols and physical communication layers since different devices communicate differently from one another. Small microprocessors boards are therefore good candidates for this first research project iteration. Figure 7 shows the HH architecture for this option.

Figure 7: Home Hub V1.0 Architecture—Local Intelligence

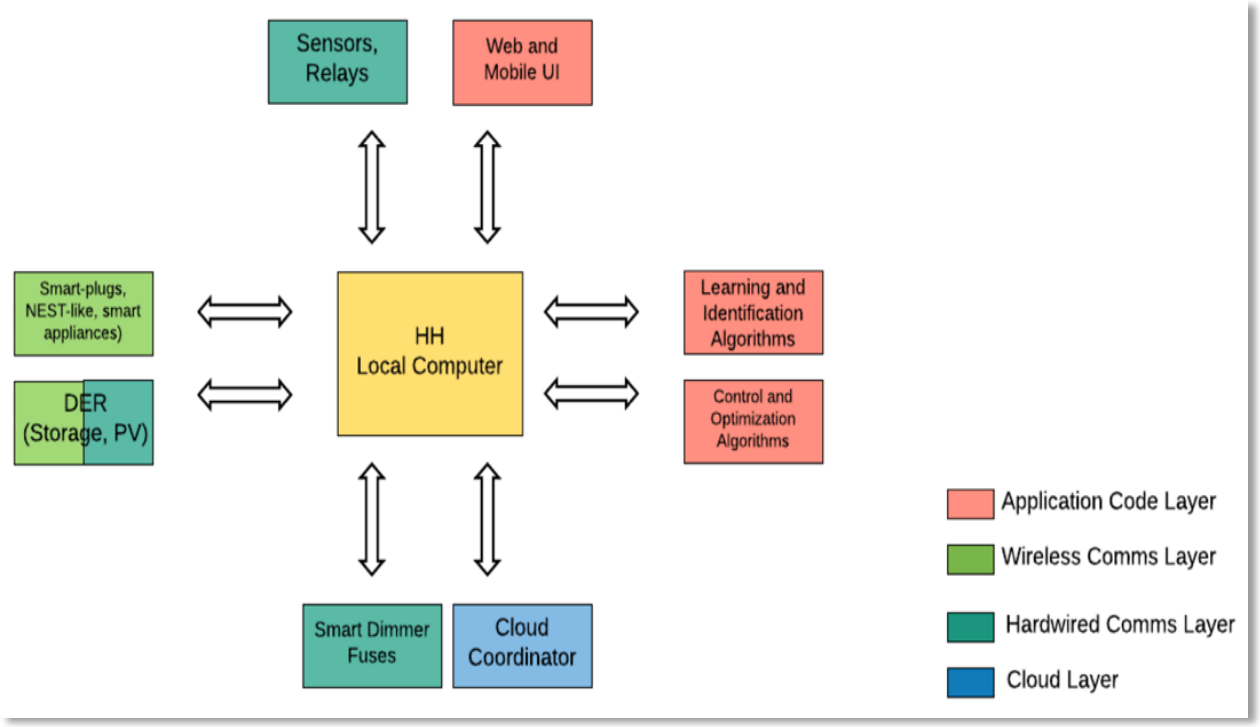

Source: SLAC National Accelerator Laboratory

The Home Hub V2.0 – Cloud Intelligence

This option was based on the ability of participating devices to connect with the Cloud. It was necessary to install a commercial power-monitoring system to measure, in real time, the power consumed by subcircuits in the house. Based on data accessible to each participating house, it was possible to independently implement each unit's algorithms in the Cloud.

The cloud-based HH was highly scalable and dependent only upon commercial devices. This strategy was adopted for participating houses in the City of Fremont. Figure 8 shows the cloud-based HH architecture.

Figure 8: Home Hub V2.0 Architecture–Cloud Intelligence

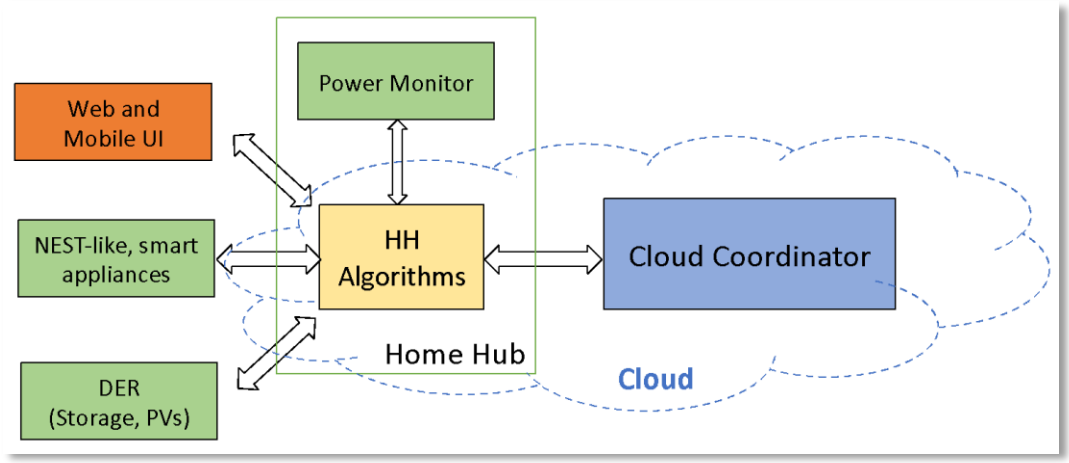

Source: SLAC National Accelerator Laboratory

Feedback Control of the Electrical Network via Markets

The main goal of this project was to demonstrate that transactive energy-control techniques are feasible strategies to control DERs. A set of economic rules, combined with control mechanisms to manage DERs, created a dynamic balance between supply and demand across the entire electrical infrastructure.

The roles of the HH and the CC in local energy markets were multi-faceted. The CC was the market operator. This role can be, but does not have to be, identical with the entity managing local resources. The HHs, based on settings provided by the customers, were additionally designed to both economically evaluate the dispatch of appliances and send resulting bids to the CC. The HH also read the resulting market price after the market clearing and dispatched local appliances accordingly.

Analyzing this scheme from a dynamic point of view, the electric system (combining both controllable DERs and loads) was controlled by a market strategy based upon value as its key parameter. This controller distributed participating DERs and loads in the market.

System Validation

Validation Plan and Tests – Generalities

In general, development of large-scale projects required three main elements for the modeling and validation of power-grid-based systems:

- An analytical or reduced dynamical model
- A large-scale simulation tool (high-level simulation)
- The real system (electric grid)

The goals of the validation plan and testing were two-fold:

- Achieving agreement between model dynamics and the real system.
- Defining generic tests to verify system stability and performance availability.

The first preceding goal was critical to obtaining models that could replicate, in the analysis and design stage, results from when controls were deployed to the grid. The second goal defined different scenarios to analyze and quantify system performance and availability when operating under the proposed controls.

The Powernet project was the first step in designing and implementing the hardware, firmware, and algorithms that controlled the demand-responses that then managed The DERs, storage, and loads at the appliance level. One important outcome of the project was to quantify the stability and performance of an aggregated system for different scenarios, markets, and DER-penetration levels. Because it was the first step in the process of designing and evaluating the system, direct results from deploying equipment in 1,000 to 10,000 houses was neither anticipated nor performed. This result will instead be evaluated through simulations using a large-scale simulation tool and appropriately reduced models. The strategy was rather to estimate the performance and availability of a large-scale system using high-level simulators, based on validated models of the houses, to deploy the technology in fewer houses to prove its operation. These concepts are summarized in Figure 9.

Figure 9: Goals of the Validation Plan

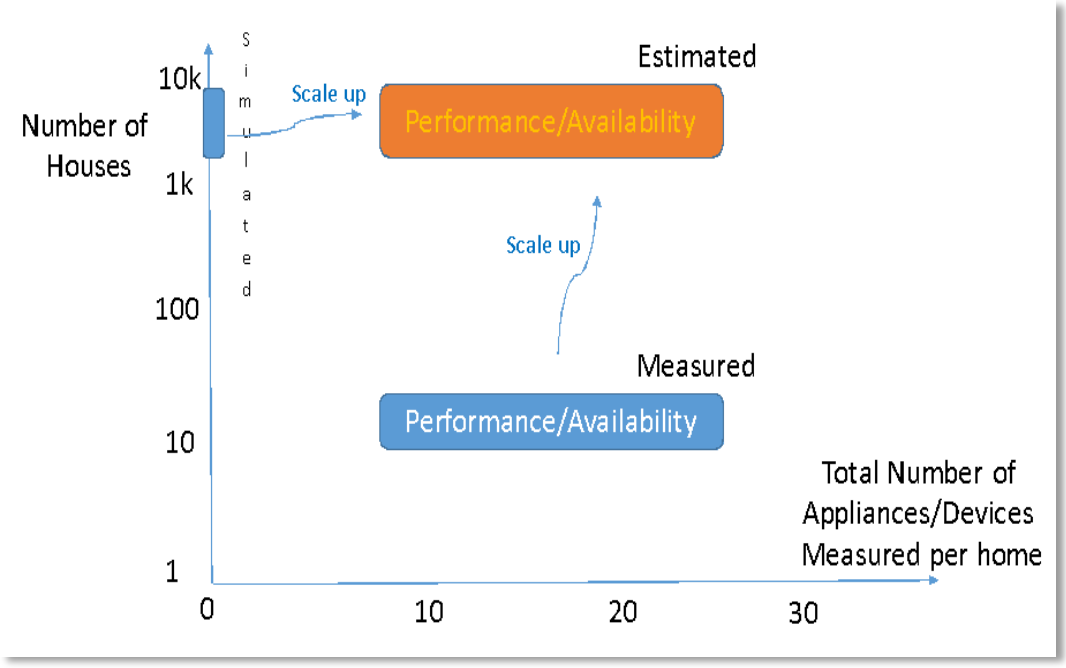

Source: SLAC National Accelerator Laboratory

Because this is the first step of the project, the risk and budget allowed only the installation of hardware and firmware in a small number (10-20) of houses and a simulation for a larger number of houses. To combine both parts of the project and obtain an accurate estimate (or extrapolation) of the results, a model validation was required. The main point of the plot depicted is to answer the question, "Is the estimation close enough to a real-life outcome to confidently apply the full technology to a larger scale?" The validation plan answered that question while attempting to minimize the estimation error in all the steps of the validation process.

The first perspective of this method was that minimizing errors in the modeling was sufficient for extrapolating results for an aggregated system. This concept was valid from a control-design point of view. Though it is critical to have an accurate model, it is also important to keep in mind that uncertainty abounds; robust control strategies must accommodate those uncertainties in both the model estimation and in actual operations. There is another perspective that makes extrapolation to more houses (1,000-10,000) more credible and considers a reliability point of view. Rare disruptive events are always possible in a much larger system of houses, events that can be overlooked in an analysis of system performance under different scenarios. Those unanticipated events are very important for aggregators and utilities when evaluating the merits and limitations of proposed technologies.

In summary, the validation plan proved the proper operation of the hardware, firmware, and software corresponding to the home hub and the monitoring system. Those devices were installed in participating houses and in the laboratories and tested under different conditions. Those devices controlled multiple appliances and devices and communicated with the cloud coordinator. Basic algorithms in the cloud coordinator represented the markets tested only on a small scale within the greater system. The tests on 12 houses provided limited information about their performance. Cloud coordinator algorithms and the markets were studied extensively using the GridLAB-D simulator, with a larger number of houses and scenarios and with different penetrations of renewable resources.

Validation of Device – Appliance Models in Simulators

One of the goals of the validation plan was to reach agreement between the model dynamics and those in the actual system. This goal defined the baseline where individual device and appliance models were compared with the actual ones. This was achieved based on both measurements in the system and its different devices. This methodology allowed a direct comparison of models and system devices, providing a continuous interpretation of results from analytical models. Additionally, it then made it possible to generalize or extend simulation results to an actual system that integrated with the grid.

As depicted in Figure 10, control design was based on system analytical models. After a control algorithm was designed, further tests were conducted through high-level simulations to capture more system detail. An important goal was that the high-level simulation results be close to a future outcome when the system is ultimately deployed to the actual grid. To fulfill this goal, future validation tests should be run to establish close agreement between the analytical representation, the models in the high-level simulator, and the actual electric grid.

Figure 10: Validation Plan

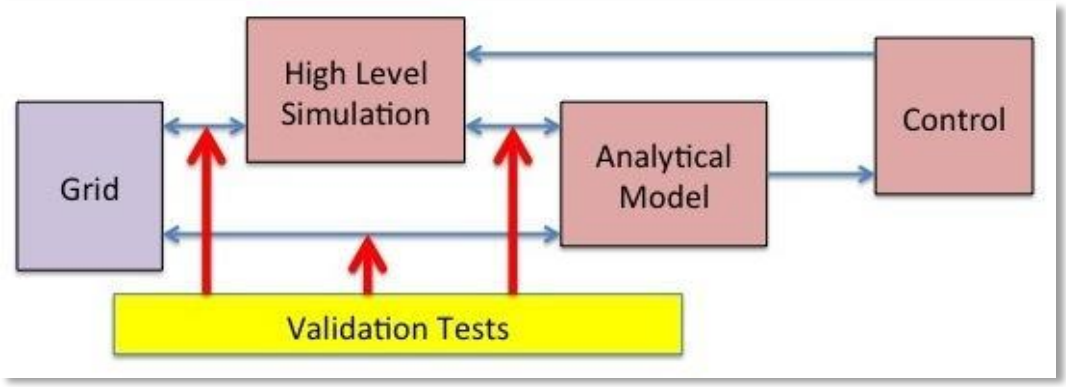

Validation Plan: Coordination among different levels in the project

Source: SLAC National Accelerator Laboratory

Analytical Models

The Volt-Current-Power-Reactive power (V-I-P-Q) characteristics of the multiple individual devices in a home (for example, ZIP models) had to be defined. Those devices included different appliances (including storage devices) and PV cells. It was important to define not only the main dynamic model of the loads, but also their daily time-of-use to estimate their combined effects on aggregation.

The plan had several key steps:

- Evaluated existing analytical models for different devices in a house and quantified them based on direct measurements conducted in the laboratories.
- Evaluated data from multi-house experiments to analyze use patterns for different appliances. Data from Pecan Street, Co. Dataport was used to run the simulations (Pecan Street 2015).
- Installed monitors at the residences to evaluate their appliance use. Used either this data or Pecan Street data to validate models in the simulation.

High-Level Simulation – GridLAB-D

GridLAB-D was used as the agent-based platform to simulate the larger-scale system. The simulation appropriated models for the houses, including storage, appliances, PVs and hardware-in-the-loop (HIL). These models were validated using lab measurements, the houses, or high-level models created in other national laboratories.

Evaluation of Stability, Performance, and Availability

The overall performance of the system was evaluated using predefined conditions for demands and measuring responses to those stimuli. To evaluate the system, different scenarios were defined based on the combination of four factors:

- Different solar penetrations (25-, 50-, and 75-percent).
- Different storage unit levels (15-, 25-, and 50-percent).
- Normal, peak, and outage conditions.
- Introducing grid congestion.

It was pointed out in this report's introduction that quantifying the impact of outages was important for assessing the performance and availability of the system. Some outage situations studied include the temporary loss of houses participating in the program and the main failure of large electric distribution equipment, among other scenarios.

All the house models in the high-level simulation tool (GridLAB-D) were validated, following different schemes. To extrapolate the analysis to between 1k-10k homes, models were created using base data from validated house models. Analytical models representing some set of aggregated houses were defined, allowing the application of future analyses to a larger number of houses. These reduced models were validated with similar results using the high-level simulator.

CHAPTER 3:
Control of Distributed Power Resources by Retail Markets

Introduction

Integrating renewable energy and transportation into California's power grid is driven by policies that drastically reduce the use of fossil fuels and gas emissions. This mandated shift to clean energy is limited, however, by the stability of the operation of the power system when high penetrations of renewable energy generators and electric transportation are added to the grid's energy mix. The interconnected power grid is a complex and large-scale system with rich nonlinear dynamics, and the steadily increasing integration of power electronics interfaced devices (PEIDs), such as renewable energy generation, like wind and solar, energy storage, electric vehicle (EV) chargers, and others into the power network, driven by both economic and climate factors, can lead to disruptive changes in power system dynamics.

The multiple resources distributed in the network make up a redundant system where the amount of total available power (feeder + locally generated) is, under normal conditions, greater than power demand. That condition determines generation choices from among the multiple solutions for power distribution. Additionally, if the electric load can be controlled, as described in the previous chapter, that same concept applies to loads. These options open the possibility of different combinations of sources and load connections to satisfy the instantaneous equilibrium between supply and demand on the state's power grid.

An electric system tasked with combining controllable renewable sources and loads must design a closed-loop feedback control strategy to balance power and consumption. To achieve this goal, the operation point of the system must be stable, so the feedback system must also reject perturbations to satisfy performance criteria in the controller design.

The topology of the "new" electric grid of the future will necessarily include multiple PEIDs. To assess the stability of the operation point, it must consider the full dynamics of the system. In general, this system admits a multi-scale separation of the dynamics. There is a fast dynamic associated with interactions among power electronics interfaced devices on the grid while other, slower dynamics control system variables and the distribution of energy among sources and loads. The control and coordination of a large number of distributed-energy assets require innovative approaches. For this project, slow dynamics were controlled by a decentralized feedback scheme based on an auction market called either a transactive energy control or a local energy market. Additionally, it was assumed that interactions among PEIDs in the electric network were stable so that the study could focus on assessing the stability and performance of the slow dynamics within a frequency range defined by the clearing time of the market. As a first step in the analysis of market-based control for the electric network, a discrete auction order book with a clearing time of five minutes was implemented. Figure 11 shows the impact of feedback systems based on this transactive energy control in time-scale dynamics.

Figure 11: Impact of Transactive Energy Control in Time-Scale Dynamics

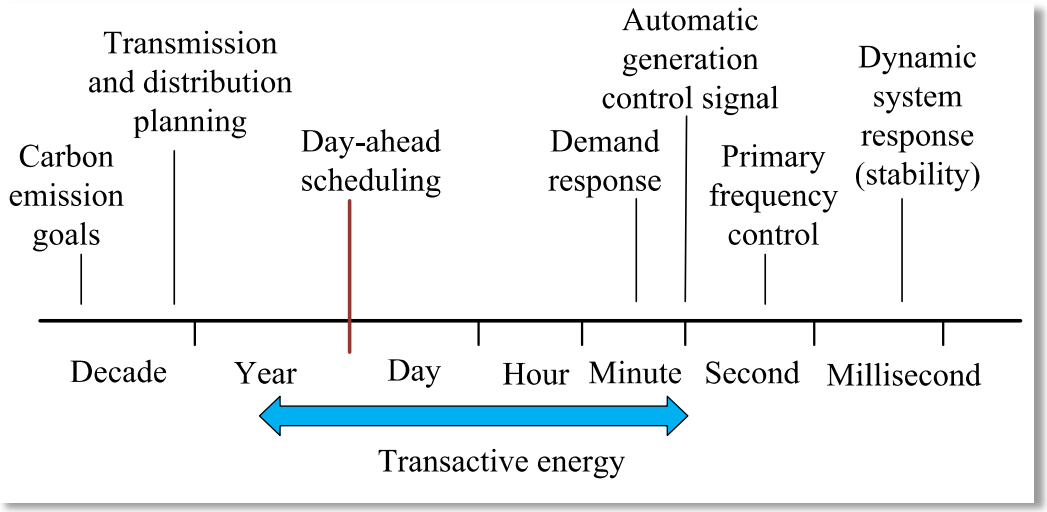

Source: SLAC National Accelerator Laboratory

Transactive energy control has received much positive attention because of its decentralized decision making and transparency. It is defined as "a set of economic rules and control mechanisms that allows the dynamic balance of supply and demand across the entire electrical infrastructure using value as a key operational parameter" (Hu, et al. 2017, Hammerstron, et al. 2007). In this scheme, through efficient and scalable algorithms, mid- to small-sized electricity-consuming and electricity-producing devices automatically negotiate their actions with one another, with devices in the physical network, and with dispatch systems of energy suppliers.

The overall slow dynamics of the system were comprised by the aggregation of electric devices in participating residences and the auction market, which acted as the controller. The equilibrium point of this system was balanced between load demand and locally generated power bound for integration into the grid. The equilibrium point must be stable for grid operations. Fast dynamics were assumed to be stable since its presumed application was for relatively small networks connected to a main distribution port. The stability of slow dynamics was assessed in the system's market design and its control implementation. More details about time-scale separation in a dynamic system appear in Appendix A.

Appliances and other devices installed in the residences acted as either non-responsive/inflexible or responsive/flexible loads. The first group formed part of all the loads connected to the electric network, whenever necessary and independent of market behavior. This group represented residents' consumption, which was primarily defined as a priority independent of the tariff associated with the energy's cost. Responsive loads were those where power consumption was determined by both human behavior and the market. It defined a means to control load demand based on the electricity tariff, but without affecting either the preferences or comfort of residents. This last point was extremely important in the design of the control strategies because the participation of appliances and devices in the aggregated flexible load group depended strongly on resident trade-offs between benefits and commitments. Additionally, the modern houses had devices capable of both generating and storing electricity. In a similar sense with loads, generators and storage devices operated as

both non-responsive/non-committed or responsive sources. Based on these concepts for the houses, in aggregation, the goal was to control the total power consumed/generated by the set of houses and the individual power consumed/generated per residence. This control signal was generated by the market and communicated to each house. The market acted as the system controller and generated control signals based on preferences defined by the consumers and the dynamic market price.

In this project, the market was based on dynamic or real-time pricing. It was characterized by determining the price that best reflected wholesale market prices for consumers. The simple electricity market model had three participants: suppliers, consumers, and the California Independent System Operator (California ISO). Suppliers and consumers are price-taking, profit-maximizing agents. The California ISO is an independent, profit-neutral player in charge of clearing the market - that is, matching supply and demand subject to network constraints. Figure 12 shows a block diagram of the feedback system where participating sources and loads were controlled by the market.

Figure 12: Block Diagram of the Feedback System

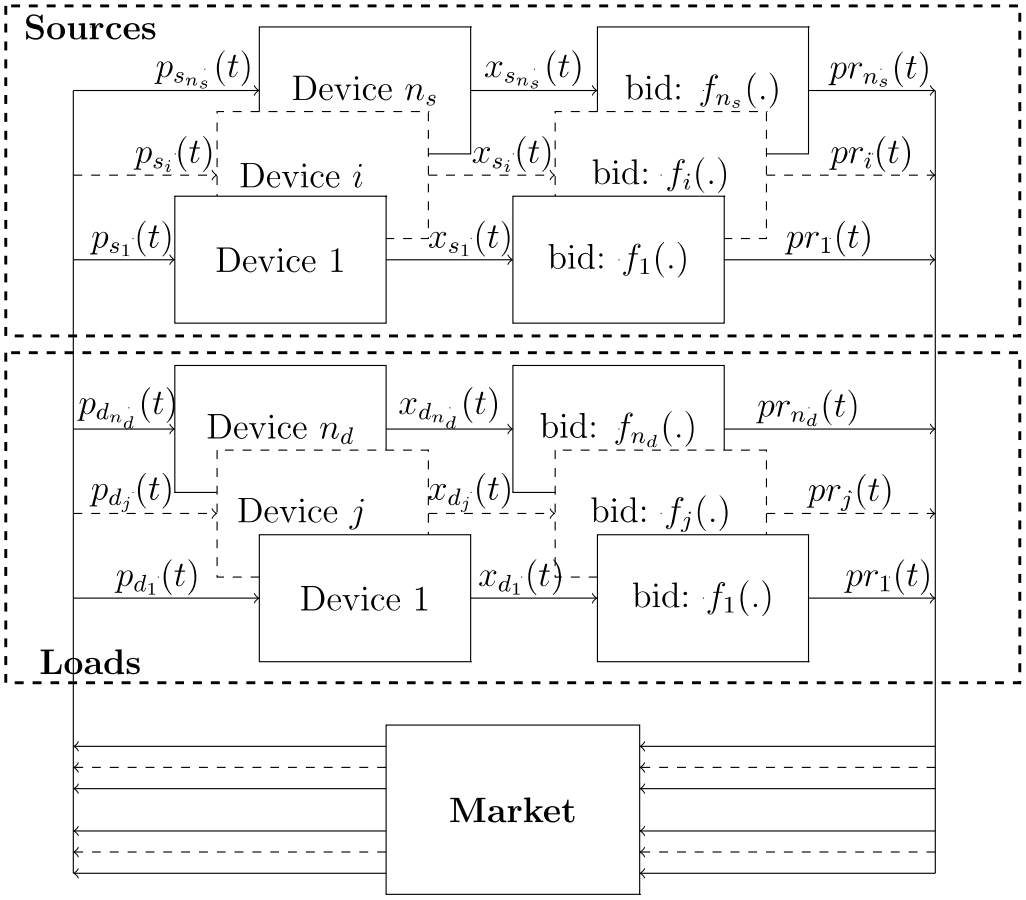

Source: SLAC National Accelerator Laboratory

From Figure 12, it is important to note a forward path in the system that related a control signal (total power, power-per-supplier or consumer, $p_{S_i}(t)$, $p_{d_j}(t)$) to a supplier's or a consumer's preference. ($pr_i(t)$, $pr_j(t)$). In this path, part of the dynamics was defined by the equipment, appliances, and devices installed in the residences, and also by the preferences of

suppliers and consumers. The last feature was part of an indirect control that involved not only the preferences and their relationship with the electric tariff, but also defined the level of participation of devices in the program.

Dynamic Model of Loads and Sources

It is important to understand the dynamics of each device in the residences that participated in the market. The aggregation of these devices defined the overall dynamics of the system, including the auction market as the controller.

The full thermal-dynamic model of the house was described precisely by a partial differential equation. This is a very detailed description of the system involving the aggregated behavior of multiple houses within the electric network. More simplified versions, or reduced models, were used to describe both the temperature in some parts of the house and the power consumed by this thermal equipment that maintained the temperature and comfort inside the house. The complexity of much more detailed thermal models used in most building simulations was reduced to a simple model based on simplified thermal heat flow dynamics. It lumped the heat flow path and thermal mass elements in a few elements, reducing the original partial differential equation to an ordinary differential equation of order one or two. This reduced the number of details in the residence parameter design and therefore the complexity of the overall aggregated system.

The reduced model for the thermal system is shown in Figure 13, where the elements C_A and C_M represent the air and bulk house thermal masses, respectively, U_A represents the conductance of the thermal envelope of the house, and H_M the conductance between the air and the envelope of the house (for example, the walls, windows, ceilings, and floors). T_A represents the air temperature, T_o the external temperature, and T_M the temperature of the envelope of the house (Chassin and PNNL 2008) (Tesfatsion and Bettula 2019). In the model, there were two energy sources; Q_M, the heat from appliances and solar radiation directly delivered to the mass of the house, Q_A the energy delivered to the air from the heating/cooling system, and also a fraction of heat from appliances and solar radiation coupled through the air.

Figure 13: Reduced Model of the Thermal System

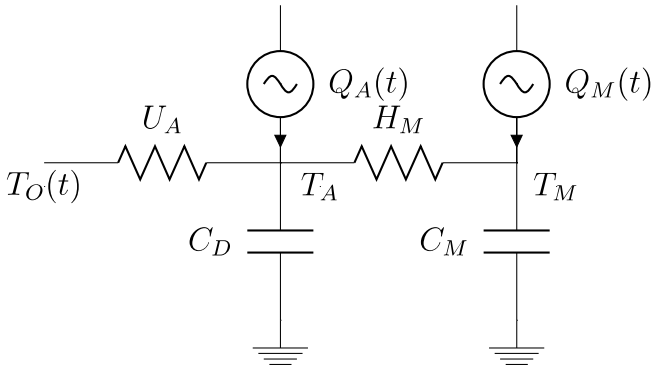

Source: SLAC National Accelerator Laboratory

The second order differential equation describing this system was

$$C_A \frac{dT_A}{dt} = -U_A(T_A - T_O) - H_M(T_A - T_M) + Q_A$$

$$C_M \frac{dT_M}{dt} = -H_M(T_M - T_A) + Q_M$$

This equation can be further reduced to a first-order differential equation if the time constants follow the inequality $U_A/C_A \ll H_M/C_M$. The ambient temperature T_A was controlled by the power Q_A by switching either between the maximum heating power and zero or the maximum cooling power and zero (switching on or off). The average thermal power was controlled by adjusting the duty cycle between these two power levels. The power taken from the grid to generate both the cooling and heating energy was proportional to the level of maximum cooling or heating power required. When the system was switched off, the unit still took a small amount of power from grid to power fans and auxiliary equipment.

The storage system was comprised of the battery units and the power inverters that linked the DC units to the AC grid. Batteries were charged or discharged by setting the power level in the inverter. In this system, an important parameter was the state of charge of the battery, defined by the ratio between the actual charge and battery's maximum capacity, $SOC = {}^{Q}/_{Q_{MAX}}$. The block diagram of the storage unit is depicted in Figure 14. The dynamic model of the system is described by the first order differential equation:

$$\frac{dQ(t)}{dt} = \frac{\eta p(t)}{E} \quad or \quad \frac{d\,SOC(t)}{dt} = \frac{\eta p(t)}{E Q_{MAX}} ,$$

where E is the voltage at the battery terminals, p is the power taken from the grid and η the efficiency of the power converter. In these inverters, the power during battery charge or discharge could be controlled.

Figure 14: Model of Battery Charger

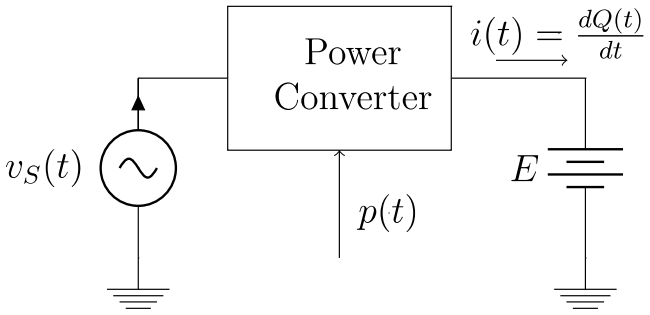

Source: SLAC National Accelerator Laboratory

The electric vehicle battery was similar to those in the storage units installed in the residences. There were slight differences; for some of the residential electric vehicle battery chargers the charging power level was fixed and could not be adjusted (Type I, II chargers). Additionally, in the control strategy of the electric vehicle battery, it was necessary to account for restrictions imposed by the use of the car. The car arrived at home during the evening having a low-charge level in the battery and had to be fully charged by early in the morning. Assuming the

minimum state of charge (SOC) of the battery was about 20 percent, any charge scheme for the vehicle battery had to be within the boundaries shown in Figure 15, where $\bar{p}, \underline{p}$, represent the maximum and minimum power of the charger, respectively, to charge the battery within that time period $t_f - t_{ini}$.

Figure 15: Power Locus to Charge the Electric Vehicle Battery

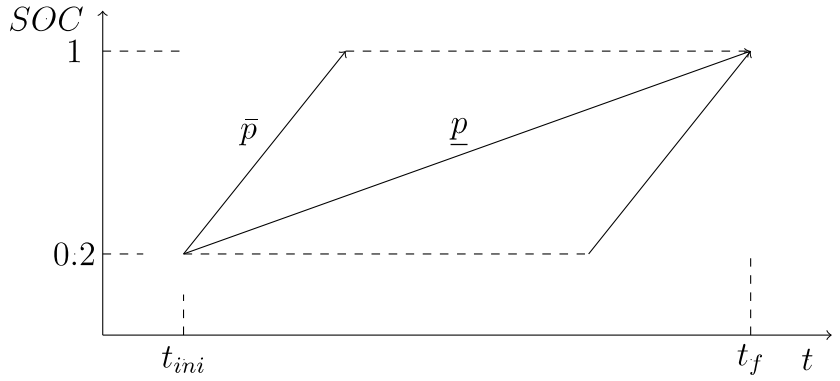

Source: SLAC National Accelerator Laboratory

The group of non-responsive loads and sources was comprised of PV cell panels, general appliances and lights, among other sources. In particular, the power transmitted to the grid by PV panels could be controlled by setting the bias of the cells; but in residential applications that bias is fixed and set to the point where the panel transmits maximum power to the network. In this analysis, it was assumed that the power consumed by general appliances was not controlled by the market. For these inflexible loads and sources, the power consumed and generated could be represented as a time-domain function in the slow-dynamics model.

Auction Market as Multi-Input Multi-Output (MIMO) Controller of Distributed Resources

In the auction market the resident can simultaneously act as consumer, producer, or consumer and producer, depending upon the resources offered for participation in the transaction. The auction mechanism is such that agents can bid into the market for either consuming or supplying energy after agents have already been selected to connect their resources with the electric network. The assets from those agents are called active demands or supplies (Roozbehani, Dahleh and Mitter 2010).

Let $D = \{1, \dots, n_d\}$ and $S = \{1, \dots, n_S\}$ denote the index sets of energy demand (loads) and energy supply (sources), respectively. Let $D_A = \{1, \dots, n_{d_A}\}$ and $S_A = \{1, \dots, n_{s_A}\}$ be subsets of D and S, respectively, where the number of elements represents active demands and supplies. These are the units able to connect to the system after the market's clearing time. As such, $n_{d_A} \leq n_d$ and $n_{s_A} \leq n_s$.

Each demand $j \in D$ is associated with a value function $v_j(.) : R^+ \to R$, where $v_j(x)$ represents the value that consumer j derives from consuming x units of the resource (electricity). Similarly, each supply $i \in S$, is associated with a cost function $c_i(.) : R^+ \to R^+$ representing the cost-per-unit production of the resource. Let $d_j : R^+ \to R^+, j \in D$, and $s_i : R^+ \to R^+, i \in S$ denote

C^1 functions that map price to consumption and production, respectively. In the framework of price taking, each agent maximizes the net benefit derived from the market. Therefore:

$$d_j(\lambda) = \arg\max_{x \in R^+} \; v_j(x) - \lambda x, \quad j \in D,$$

$$s_i(\lambda) = \arg\max_{x \in R^+} \; \lambda x - c_i(x), \quad i \in S.$$

The functions $d_j(\lambda)$ and $s_i(\lambda)$ define the resource units where an agent maximizes the net benefit for a given price λ. These maximization problems have a unique solution in R^+, and the functions $d_j(.)$ and $s_i(.)$ are well defined:

$$d_j(\lambda) = \max\{0, \{x \mid \dot{v}_j(x) = \lambda\}\}.$$

$$s_i(\lambda) = \max\{0, \{x \mid \dot{c}_i(x) = \lambda\}\}.$$

Figure 16 shows these conditions for demand and supply resources. From that figure, it is assumed for this analysis that $d_j(\lambda) = \dot{v}^{-1}(\lambda)$ and $s_i(\lambda) = \dot{c}^{-1}(\lambda)$.

Figure 16: Value Function for Consumers and Suppliers

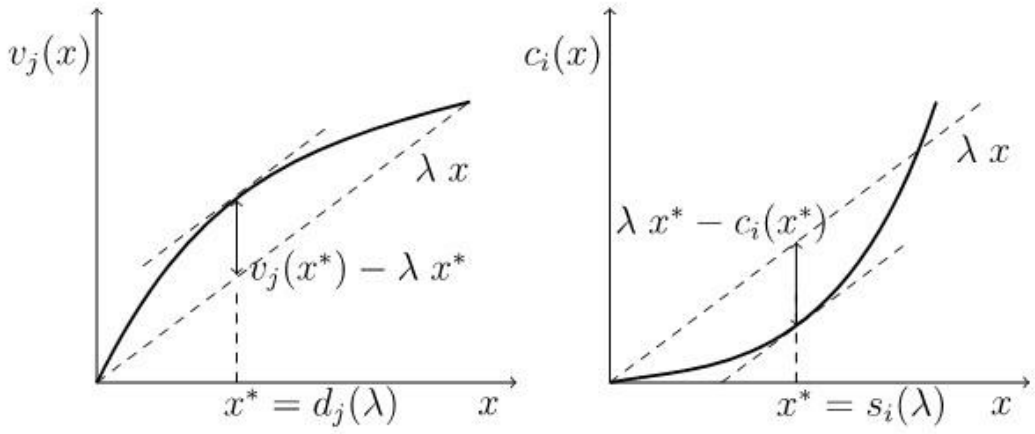

Source: SLAC National Accelerator Laboratory

The social welfare $\mathcal{SW}$ is the aggregate benefit of the producers and consumers:

$$\mathcal{SW} = \sum_{j \in D} (v_j(d_j) - \lambda_j d_j) + \sum_{i \in S} (\lambda_i s_i - c_i(s_i))$$

When the system is at equilibrium, in the sense that the total supply equals the total demand and there is a unique clearing price λ for the entire system, then:

$$\mathcal{SW} = \sum_{j \in D_A} v_j(d_j) - \sum_{i \in S_A} c_i(s_i)$$

Then, the California ISO's optimization problem is:

$$\max \sum_{j \in D} v_j(d_j) - \sum_{i \in S} c_i(s_i), \quad s.t.: \sum_{j \in D_A} d_j = \sum_{i \in S_A} s_i$$

A solution to this problem is characterized by the common price λ^*. The implication is that by setting the market price to λ^*, the system operator creates a competitive environment in which the aggregated surplus is maximized and each agent or participant device maximizes its own net benefit.

Interpretation of the Results Based on Marginal Cost and Willingness to Pay

Assuming $d_j > 0$ and $s_i > 0$, let $pr_{d_j} = {}^{v_j(d_j)}\big/_{d_j}$ and $pr_{s_i} = {}^{c_i(s_i)}\big/_{s_i}$ be the willingness to pay and the marginal cost, respectively. Then the social welfare $\mathcal{SW}$ can be expressed as a function of marginal cost and willingness to pay:

$$\mathcal{SW} = \sum_{j\in D}\left(\frac{v_j\,(d_j\,)}{d_j} - \lambda_j\right)d_j + \sum_{i\in S}\left(\lambda_i - \frac{c_i(s_i)}{s_i}\right)s_i$$

$$= \sum_{j\in D}\left(pr_{d_j} - \lambda_j\right)d_j + \sum_{i\in S}(\lambda_i - pr_{s_i})s_i$$

The willingness to pay can be plotted for each demand d_j and the marginal cost for each supply s_i as depicted in Figure 17. Then maximum social welfare is obtained at a common price $\lambda_j = \lambda_i = \lambda^*$, defined by the intersection of the curves. This intersection also defines the total quantity $p_T = \sum_{j\in D_A} d_j = \sum_{i\in S_A} s_i$. It is important that for $\lambda_i = \lambda_j = \lambda^*$, in Figure 17, the first summand represents the upper shaded area and the second summand the lower area. Both areas are maximum at that particular price, so are also at maximum social welfare. From this plot it is possible to deduce that, after a particular clearing time, from the resources participating in the market, the demands of willingness to pay are larger than the price λ^*, and the sources with marginal cost lower than the price λ^* will be supplying power to the grid.

Figure 17: Demand-Supply Curves

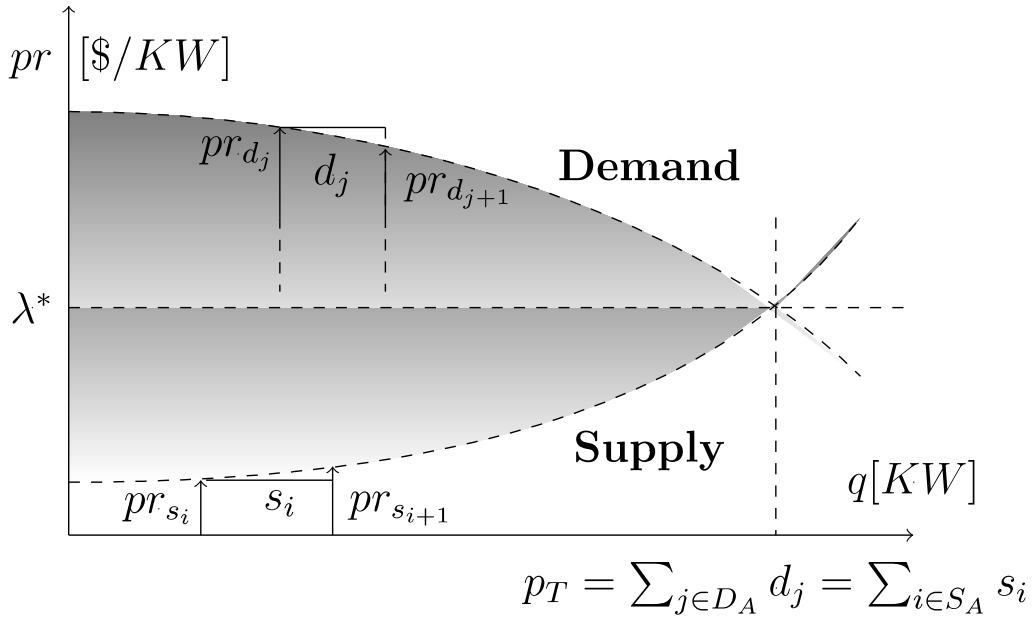

Source: SLAC National Accelerator Laboratory

Implementation of the Control

Based on previous interpretation of the optimization problem, it is possible to implement a multi-input multi-output (MIMO) controller for the system. It is important to know that active loads are bidding greater than the price λ^* defined by the market, and that active sources are bidding less than λ^*. This market mechanism defines a control strategy based on individual biddings pr_{d_j} and pr_{s_i}:

$$if\ pr_{d_j} > \lambda^* \Rightarrow\ load\ ON.$$

$$if\ pr_{d_j} < \lambda^* \Rightarrow\ load\ OFF.$$

$$if\ pr_{s_i} > \lambda^* \Rightarrow source\ OFF.$$

$$if\ pr_{s_i} < \lambda^* \Rightarrow source\ ON.$$

The result defined by load (either on or off) and source (either on or off) means that when the load is on it is connected to the grid and consuming a power $d_j = p_{d_j}$, and when the load is off the consumed power is zero or negligible. This concept is similarly extended to the source. This condition is evaluated by the market at every clearing time and communicated to all the agents or devices participating in the auction.

The sign or step Heaviside function is defined as:

$$y = sign(x) = \begin{cases} 1, & if\ x > 0 \\ 0, & if\ x < 0 \end{cases}$$

Those conditions and results can then be represented at a particular time by the switch functions:

$p_{d_j} = d_j\ sign(pr_{d_j} - \lambda^*)$ for the loads.

$p_{s_i} = s_i\ sign(\lambda^* - pr_{s_i})$ for the sources.

Here's a snapshot of how the controller works. Bids are submitted to the market by the agents. The bids contain information on the maximum willingness to pay (for demand), or the minimum price to accept (for supply), and the desired power level of consumption or generation (kW). Based on individual biddings, (pr_{d_j}, d_j) and (pr_{s_i}, s_i) the market defines two magnitudes: the price λ^* and the total quantity $p_T = \sum_{j \in D_A} d_j = \sum_{i \in S_A} s_i$ (total power). The individual biddings are then compared against the price λ^*, using the switch functions. This defines the status and power usage of the individual devices. The block diagram corresponding to this control is shown in Figure 18. The control defines two global variables, the price λ^* and the total power p_T, and those variables determine the power contribution of the distributed resources.

Figure 18: Block Diagram of the Market-Based Controller

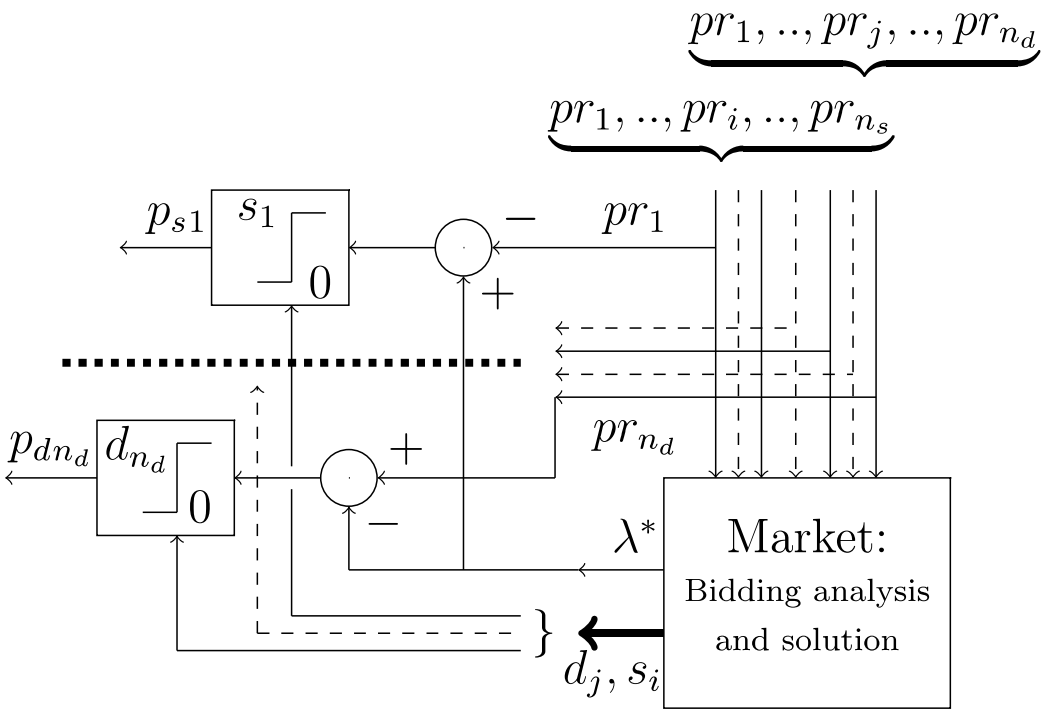

Source: SLAC National Accelerator Laboratory

The control relies on the switch function. This particular type of system is called a variable structure control or hybrid system, which contains two kinds of components: subsystems with continuous dynamics and subsystems with discrete dynamics. When the continuous and discrete dynamics coexist and interact with one other, it is important when developing models that the dynamic behavior of such hybrid system is accurately described. It is only with this dynamic that it is possible to develop designs that fully consider the relations and interactions of the system's continuous and discrete parts.

Example of Modeling and Control of Thermal-Controlled Loads

The thermal system in a residence is a classic example of a hybrid system. It can either be controlled individually by a thermostat, or it can be part of an aggregated system controlled by an auction market. The characteristic of this system is that the energy is supplied to the home in discrete events (switched "on" for heating and cooling power, or "off" for zero or idling power) and at different times, setting an intrinsic combination of both continuous and discrete subsystems. The reduced thermal model for the home thermal system just described for the j^{th} home, it can be described in compact form by:

$$\frac{dT_j}{dt} = A_j T_j + B_{1_j} Q_{EXT_j} + B_{2_j}\, u_j$$

where $T_j = \left[T_{A_j}\, T_{M_j}\right]^T$ represents the temperature vector for the j^{th} house air and mass temperatures $Q_{EXT_j} = \left[T_{o_j}\, Q_{M_j}\right]^T$. This vector includes the perturbations of the external temperature T_{o_j} and the energy Q_{M_j} represents the heat from appliances and solar radiation directly delivered to mass of the house, and $u_j = Q_A$ represents the energy delivered to the air

from the heating and cooling systems and also a fraction of heat from appliances and solar radiation coupled through the air. Q_A can be expressed by $Q_A = \gamma Q_M + p_d$, where the first term models the fraction of the heat energy coupled to the air, and p_d shows the power delivered by the HVAC unit, defining the output $y_j = T_{A_j}$, then $y_j = [1\ 0]\ T_j = C_j T_j$. Assuming the resident relates his or her comfort to the willingness to pay, it is possible to assume that the bidding price in the market for this unit, pr_j, is a function of the air temperature T_{A_j}, $pr_j = f_j\left(T_{A_j}\right) = f_j(y_j)$,where $f_j(.): R \rightarrow R^+$. The full model of the thermal system for one home, including the preference of the resident, was:

$$\frac{dT_j}{dt} = A_j T_j + B_{1_j} Q_{EXT_j} + B_{2_j}\, u_j, \quad y_j = C_j\, T_j, \ \ pr_j = f_j(y_j).$$

Figure 19 depicts the bidding function for an HVAC system, combining both heating and cooling systems. In this case, the function $pr_j = f_j(y_j)$ decreases when the HVAC unit operates as a heater, $pr_j^1 = f_j(T_A)$, for $T_A \in (T_{min}, T_{he})$, and increases when it operates as a cooling unit, $pr_j^3 = f_j(T_A)$, for $T_A \in (T_{co}, T_{max})$, and is equal to zero, $pr_j^2 = 0$, for $T_A \in (T_{he}, T_{co})$ when the air temperature is within the comfort zone.

Figure 19: Willingness-to-Pay Functions for Heating, Ventilation, and Air Conditioning Systems

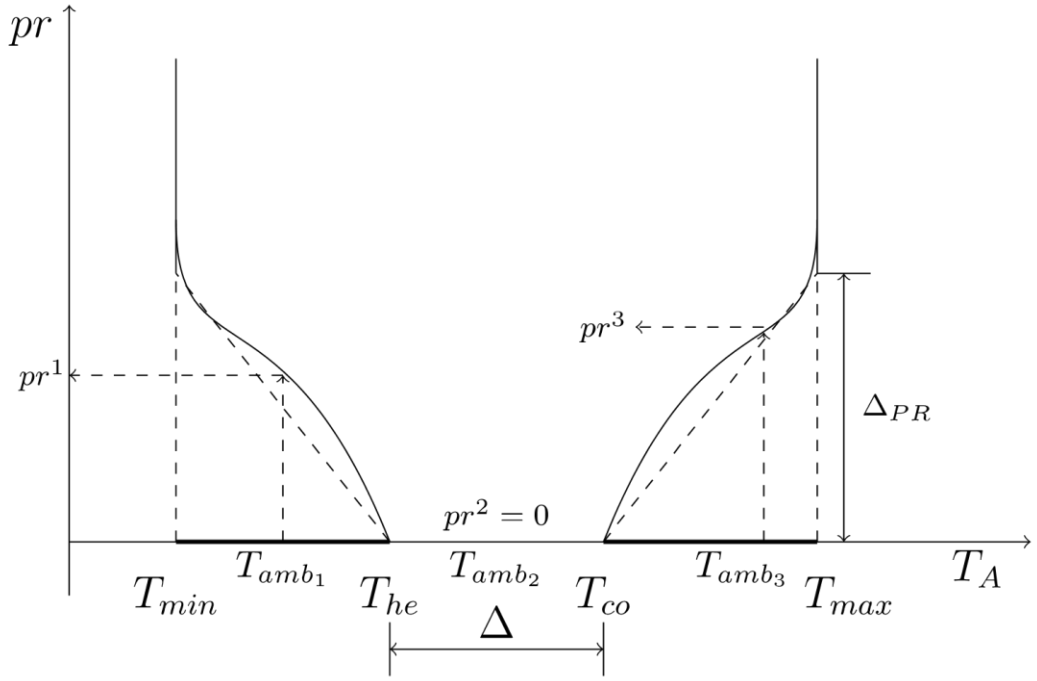

Source: SLAC National Accelerator Laboratory

Residents' Comfort Preferences

This model can be extended to represent the thermal system for multiple homes $\in D = \{1, \ldots., j, \ldots. n_d\}$. It is represented by:

$$\frac{dT}{dt} = A_{th} T + B_1 Q_{EXT} + B_2\, u, \quad y = C\ T, \ \ pr_d = f(y),$$

where $T = [T_1 \quad \cdots \quad T_{n_d}]^T, u = [u_1 \quad \cdots \quad u_{n_d}]^T,$ and $pr_d = [pr_1 \quad \cdots pr_j \cdots \quad pr_{n_d}]^T$. More details about the full model are presented in Appendix A.

The forward system, including thermal-controlled loads and biddings, based on residence temperatures, was shown in Figure 12. This system was MIMO, where n_d input signals (powers: $p_{d1}, .., p_{d_j}, \ldots, p_{d_{n_d}}$) control n_d bids (willingness to pay: $pr_1, .., pr_j, \ldots, pr_{n_d}$) sent to the market. Closing the forward system with the market-based controller, the complete feedback system is shown in Figure 20.

Figure 20: Block Diagram of a Full Closed-Loop System of Thermal Loads

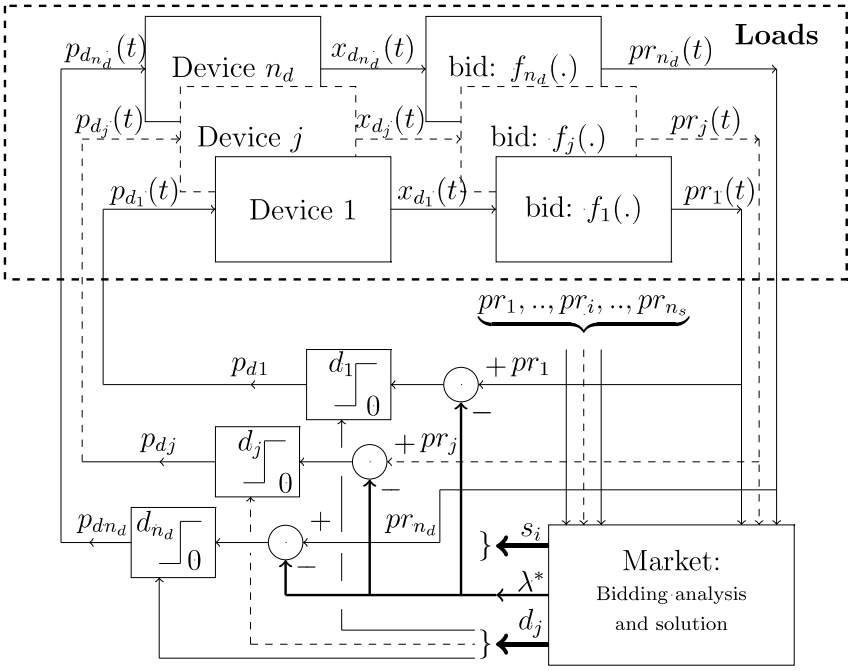

Source: SLAC National Accelerator Laboratory

Both the thermal system of the houses and the comfort preferences of residents were independent of one another. So given the market price λ^*, each house was controlled independently. The control signal defined by the market, for each unit, was:

$$u_j = p_{d_j} = d_j \, sign\big(pr_j - \lambda^*\big), \quad \text{with } j = 1, \ldots, n_d,$$

and represented the action of the controller on the forward system. Combining the equations for the forward system and the controller, the closed loop dynamics for the j^{th} residence were represented by the equations:

$$\frac{dT_j}{dt} = A_j T_j + B_{1_j} Q_{EXT_j} + B_{2_j} u_j, \quad y_j = C_j \, T_j$$

$$pr_j = f_j\big(y_j\big), \qquad u_j = p_{d_j} = d_j \, sign\big(pr_j - \lambda^*\big).$$

This model can be extended to represent the thermal system of multiple homes $\in D = \{1, \dots., j, \dots. n_d\}$. In compact form, this is represented by:

$$\frac{dT}{dt} = A_{th}T + B_1 Q_{EXT} + B_2\, u, \quad y = C\, T, \quad pr_d = f(y),$$

$$u = \left[d_1\, sign(pr_1 - \lambda^*), \dots, d_j\, sign\left(pr_j - \lambda^*\right), \dots, d_{n_d}\, sign\left(pr_{n_d} - \lambda^*\right)\right]^T.$$

If the control signal u is substituted in the differential equation $^{dT}/_{dt}$, the right-half side is discontinuous due to the $sign(.)$ function. The differential equation was not a unique solution because of its discontinuous right-hand side. It therefore failed to satisfy the conventional existence and uniqueness results of differential equation theory.

This particular type of system is called a variable structure control or hybrid system. These systems have been extensively studied and applied as strategies to regulate or control servomechanism systems. They are composed of two types of components: subsystems with continuous dynamics, and subsystems with discrete dynamics that interact with each other. In Appendix A, a description of the variable structure control (VSC) provides details on the sliding mode control, which is one control strategy (Hung, Gao and Hung 1993) (DeCarlo, Zak and Matthews 1988) (Utkin, Sliding Modes and Their Application in Variable Structure Control 1978). Sliding mode controls are characterized by a discontinuous feedback control law that switches the structure of the system during the evolution of the state vector to maintain the state trajectories in a predefined subspace. The state-feedback control law is not a continuous function of time; it can switch from one continuous structure to another based on current states in the state-space. The purpose of the switching control law is to steer the system's state trajectory onto a prespecified surface in the state-space, and to maintain the plant's state trajectory on this surface for subsequent times. The surface is called a switching surface, or sliding surface.

The VSC control design breaks down into two phases. The first one chooses a switching surface so that the plant state, restricted to the surface, has desired dynamics. The second phase is to design a switched control law that will drive the plant state to the switching surface and maintain it on the surface upon interception (DeCarlo, Zak and Matthews 1988). From Figure 9, associating a n_d number of houses in the aggregated distribution, the order of the systems was $2n_d$, which corresponded to temperatures in the air and within the envelope of each house. From that figure, it is possible to define a $(2n_d + 1)$-dimensional switching surface or equilibrium manifold $\sigma(T_A, t) = 0$ for the case corresponding to houses participating with the thermal system in the market. It is given by:

$$\mathcal{H} = \{(T_A, t) \in R^{2n_d+1} \mid \sigma(T_A, t) = 0\}$$

where

$$\sigma(T_A, t) = [\sigma_1(T_A, t), \quad \dots, \sigma_j(T_A, t), \dots, \quad \sigma_{n_d}(T_A, t)]^T$$
$$= \left[pr_1(T_{A_1}, t) - \lambda^*(T_A, t), \quad \dots, \quad pr_{n_d}(T_{A_{n_d}}, t) - \lambda^*(T_A, t)\right]^T.$$

$\lambda^*(T_A, t)$, is the price defined by the local energy market at the clearing time and $pr_j(T_{A_j}, t)$ is the willingness to pay by the j^{th} agent at time *t.* These surfaces are designed so that the system

state trajectories, restricted to $\sigma(T_A, t) = 0$, have a desired behavior such as stability or tracking.

The switching law $u(T,t) = \left[u_1(T,t), \ldots, u_{n_d}(T,t)\right]^T$, is defined such:

$$u_j = \begin{cases} u_j^+(T,t) & \text{when } \sigma_j(T_A,t) > 0 \\ u_j^-(T,t) & \text{when } \sigma_j(T_A,t) < 0 \end{cases}$$

then the control for each house is $u_j = d_j \, sign\left(\sigma_j(T_A,t)\right)$, assuming that the reaching conditions for each house are satisfied. See more details in Appendix A.

For the aggregated system, when all the trajectories reach the sliding surface $\mathcal{H}: \sigma(T_A,t) = 0$, all the houses are setting the comfort level or internal temperature based on the common price $\lambda^*(T_A,t)$ defined by the market. This result is important because it means that in steady-state the overall system converges to the state where the individual willingness to pay per house is equal to the market price. It is:

$$\sigma(T_A,t) = 0 \;\rightarrow \sigma_j(T_A,t) = 0, \;\; \forall\, j = 1, \ldots., n_d$$

$$pr_j\left(T_{A_j}, t\right) - \lambda^*(T_A,t) = 0 \;\rightarrow pr_j\left(T_{A_j}, t\right) = \lambda^*(T_A,t),$$

and

$$T_{A_j} = f_j^{-1}(\lambda^*(T_A,t)), \quad \forall\, j = 1, \ldots., n_d$$

assuming that $f_j^{-1}(.)$ exists.

This ideal solution of the discontinuous differential equation, which is the convergence to the sliding surface $\mathcal{H}: \sigma(T_A,t) = 0$, is achieved if the system switches between the different modes at infinite frequency. That solution is not possible in practical cases because it implies that the HVAC system has to turn either on or off infinitely fast. In reality, some approximations of the sliding surface $\mathcal{H}$ are used so that the system trajectory converges and remains in the new surface and the equipment is switching on or off at finite frequency. In the implementation of the system, using as controller the auction-market with a clearance time T_c, the switching time of the thermal power units is defined in units of T_c, $t = k\,T_s$ with $k = 0, 1, \ldots, \infty$. The minimum switching period is T_c. The sliding surface is redefined as $\mathcal{H}: |\sigma(T_A,t)| \leq \epsilon$ and the control strategy is:

$$u_j(T_A,k) = \begin{cases} d_j \rightarrow \sigma_j(T_A,k) = pr_j\left(T_{A_j}, k\right) - \lambda^*(T_A,k) > 0 \\ 0 \rightarrow \sigma_j(T_A,k) = pr_j\left(T_{A_j}, k\right) - \lambda^*(T_A,k) < 0 \end{cases} \quad \text{for } k = 0, 1, \ldots, \infty.$$

The control law is similar to the ideal one described here, but the inequality is checked at defined-time instants $t = k\,T_s$, with $k = 0, 1, \ldots, \infty$. This particular implementation is called discrete-time variable structure control (Gao, Wang and Homaifa 1995), and Figure 20 depicts the simulation results of the air temperature for multiple houses in a system controlled by an auction market.

Figure 21 shows, as an example, the response of air and mass temperatures, T_A, T_M, respectively, that corresponded to five houses. These houses covered the extrema in the

parameter space used in the simulations presented in this report. They mainly include 1- and 2-story residences with floor areas between 1,000 and 4,000 square feet. The comfort level in each residence was set to maintain air temperatures between 65˚F and 70˚F.

Figure 21: Air and Mass Temperatures for Some Residences

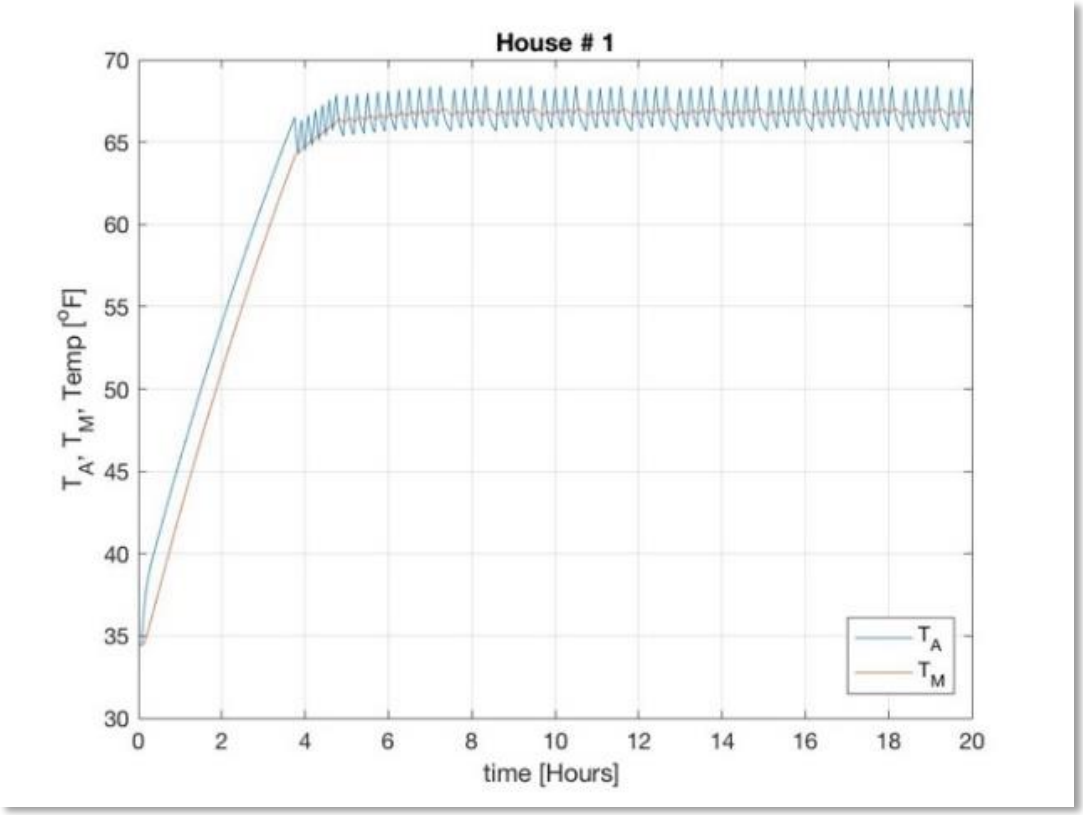

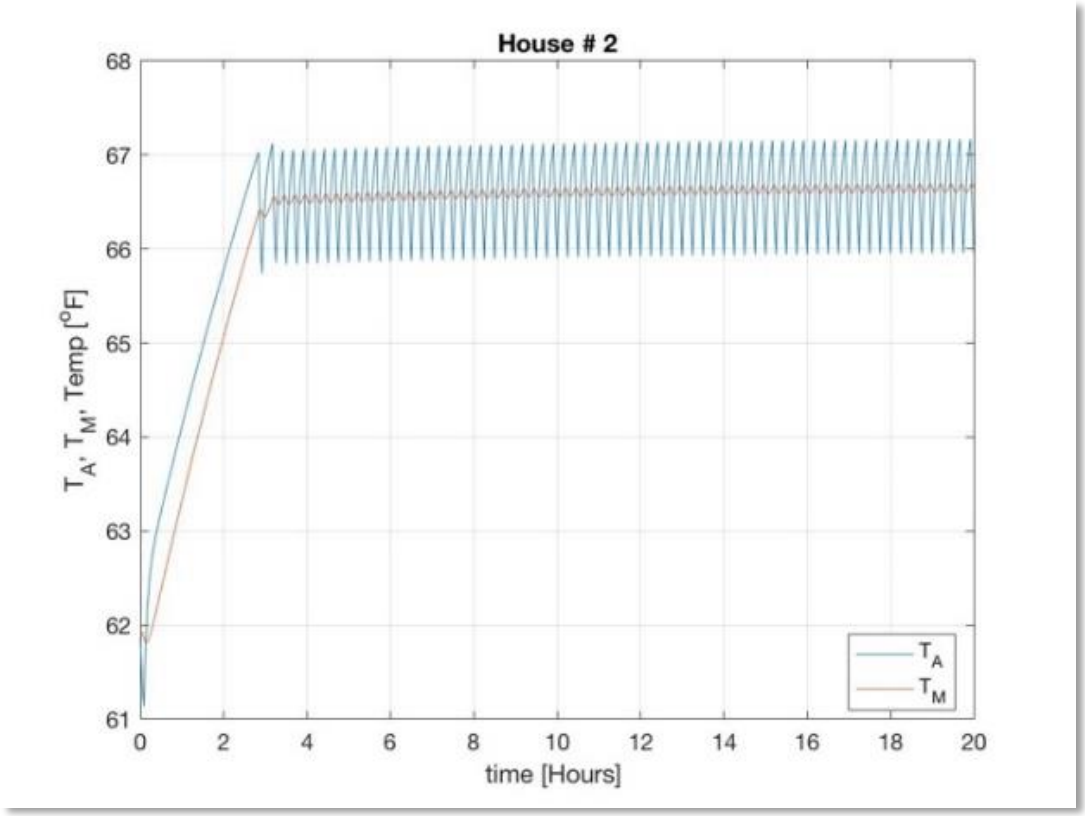
House # 2
68
67
66
65
64
63
62
61
0
2
4
6
8
10
12
14
16
18
20
time [Hours]
T_A, T_M, Temp [°F]
T_A
T_M

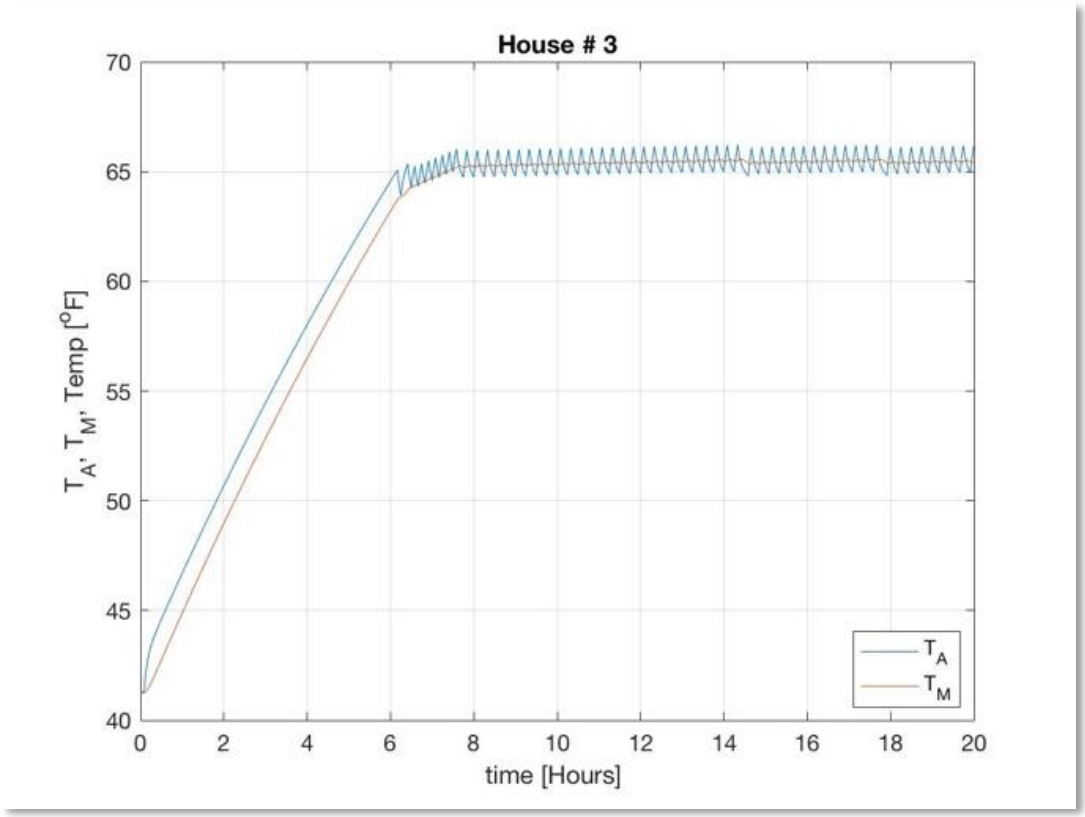
House # 3
70
65
60
55
50
45
40
0
2
4
6
8
10
12
14
16
18
20
time [Hours]
T_A, T_M, Temp [°F]
T_A
T_M

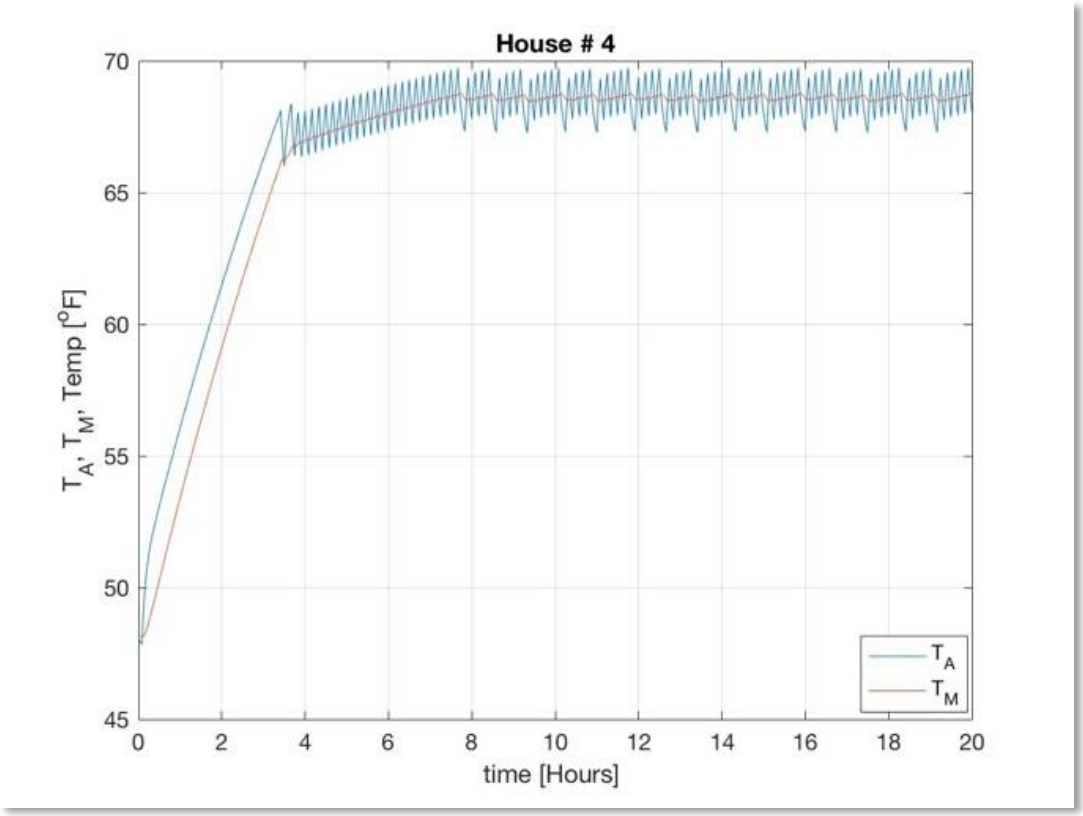

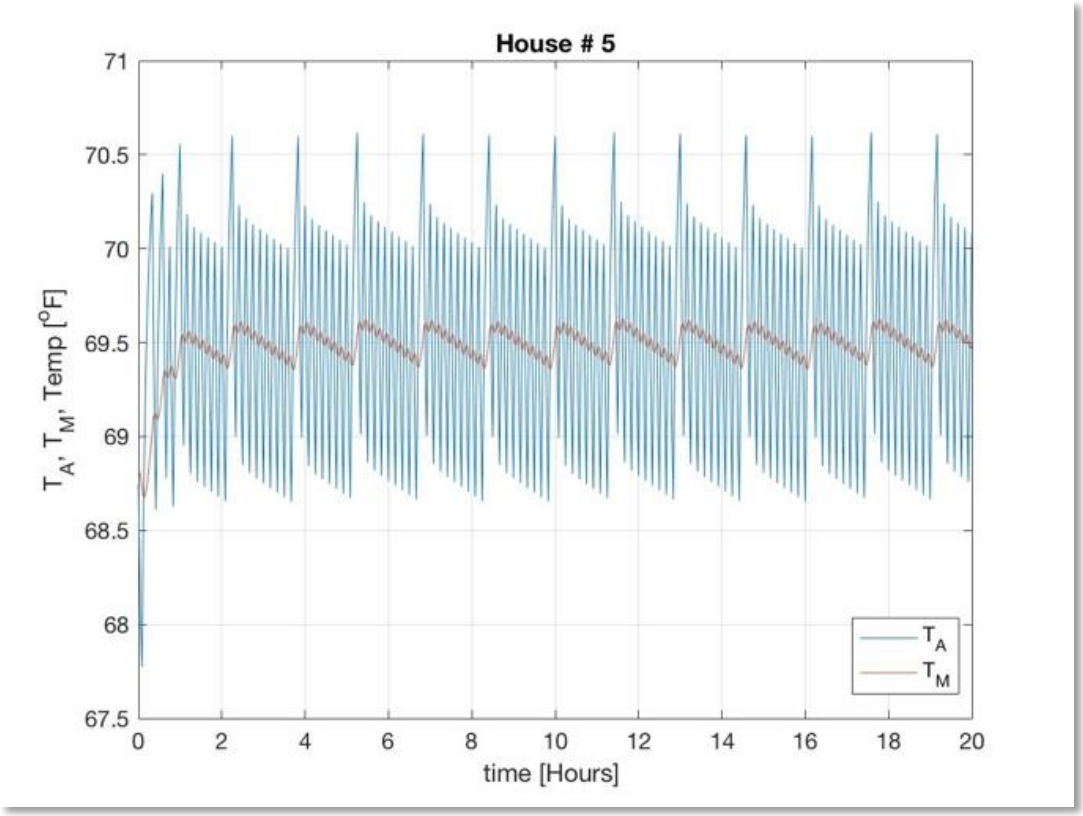

Parameters:

House#1: 1 story residence, 1000 sq. ft., window-wall ratio = 0.18, airchange_per_hour = 0.1, Ca = 423.71 $\mathbf{Btu/^oF}$, Cm = 1717.52 $\mathbf{Btu/^oF}$, Ua = 236.35 $\mathbf{Btu/^oF\ hr}$.., Hm = 4844 $\mathbf{Btu/^oF\ hr}$.

House#2: 2 story residence, 4000 sq. ft., window-wall ratio = 0.1, airchange_per_hour = 0.95, Ca = 1694.85 $\mathbf{Btu/^oF}$, Cm = 6870 $\mathbf{Btu/^oF}$, Ua = 981.44 $\mathbf{Btu/^oF\ hr}$, Hm = 15962 $\mathbf{Btu/^oF\ hr}$.

House#3: 2 story residence, 4000 sq. ft., window-wall ratio = 0.1, airchange_per_hour = 0.1, Ca = 1694.85 $\mathbf{Btu/^oF}$, Cm = 6870 $\mathbf{Btu/^oF}$, Ua = 501.23 $\mathbf{Btu/^oF\ hr}$, Hm = 15962 $\mathbf{Btu/^oF\ hr}$.

House#4: 2 story residence, 4000 sq. ft., window-wall ratio = 0.18, airchange_per_hour = 0.95, Ca = 1694.85 $\mathbf{Btu/^oF}$, Cm = 6870 $\mathbf{Btu/^oF}$, Ua = 1078.98 $\mathbf{Btu/^oF\ hr}$, Hm = 15620 $\mathbf{Btu/^oF\ hr}$.

House#5: 1 story residence, 2561 sq. ft., window-wall ratio = 0.14, airchange_per_hour = 0.6, Ca = 1085.13 $\mathbf{Btu/^oF}$, Cm = 4398.58 $\mathbf{Btu/^oF}$, Ua = 613.86 $\mathbf{Btu/^oF\ hr}$, Hm = 9320 $\mathbf{Btu/^oF\ hr}$.

Source: SLAC National Accelerator Laboratory

CHAPTER 4:
Retail-Level Markets

Problem Setup

A medium-voltage distribution system with a single connection to the overlying grid, in this case a residential area with hundreds of thousands of houses, was a use case for this project. Figure 22 shows the Institute of Electrical and Electronic Engineers (IEEE) 123 feeder test model, a residential distribution feeder, which served as the showcase for this project's analyses.

The system exhibited a constraint in its connection to the overlying transmission system. At node 150 (the only common point of connection with the electric grid), a limit in power was studied. This power limit's (forced by the market) effects were measured at node 150. While conventional power-system planning has typically sized electrical equipment to accommodate aggregate load with a safety margin, information technologies allow for a more sophisticated real-time control of resources (including shaving peaks), allowing more efficient infrastructure investment. The ability of the control to limit power at critical points leads to savings for the utility and, eventually, for the customer. Such a control is even more important for today's and future electricity systems where decentralized resources like PV systems, batteries, and electric vehicles drive the variability in aggregate loads where conventional system planning would be more expensive.

Figure 22: IEEE 123 Feeder

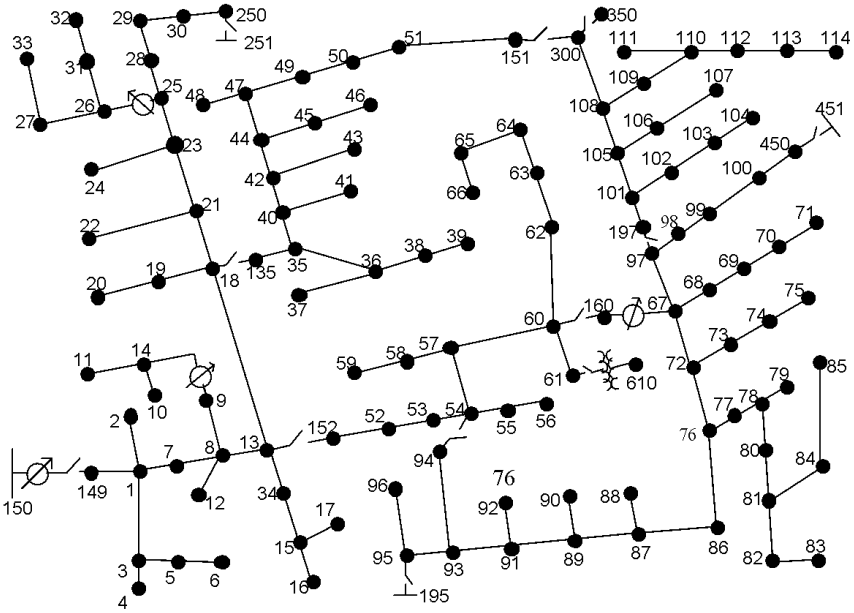

Source: SLAC National Accelerator Laboratory

The figure shows the IEEE 123 feeder, representative of a U.S. residential distribution system, illustrated by the IEEE Distribution System Analysis Subcommittee (IEEE 2010).

The coordination of flexible load and generation resources by a system operator to manage grid operations can follow a variety of principles. Those might include the dispatch according to technical effectiveness, size, or type of operator. Another approach, however, is the dispatch or curtailment of resources according to *economic principles*. This means that customers with a high value from consumption will get dispatched with higher priority than a customer with low value.

Such an approach has several advantages and:

- Ensures that those customers that attach the most value to being dispatched do not get curtailed during system operations.
- Provides information about the value of increasing grid connections.
- Stipulates efficient long-term investment into flexible load resources such as batteries or electric water heating.

The Retail Market as a Place of Economic-Value Discovery

One way to discover the value customers assign to flexible load dispatch is through a central retail marketplace. Market participants submit bids to market operators that have, on the demand side, information about how much energy they are willing to consume over the next market interval, and at what maximum price. For supply, agents submit how much they are able to supply and at what minimum price. These bids show the value that agents assign to the dispatch of their resources, provided there is sufficient competition in the market.

The market operator eventually matches demand and supply to determine the market clearing price and awards the winning bids. Following this procedure, constraints (like the maximum load of a connecting transformer) can be incorporated to restrict supply from the wholesale market by the technical limit. Figure 23 illustrates the interaction between the market, the grid operator, and agents (or consumers).

Figure 23: Setup of Local Energy Market and Interaction With Physical Grid

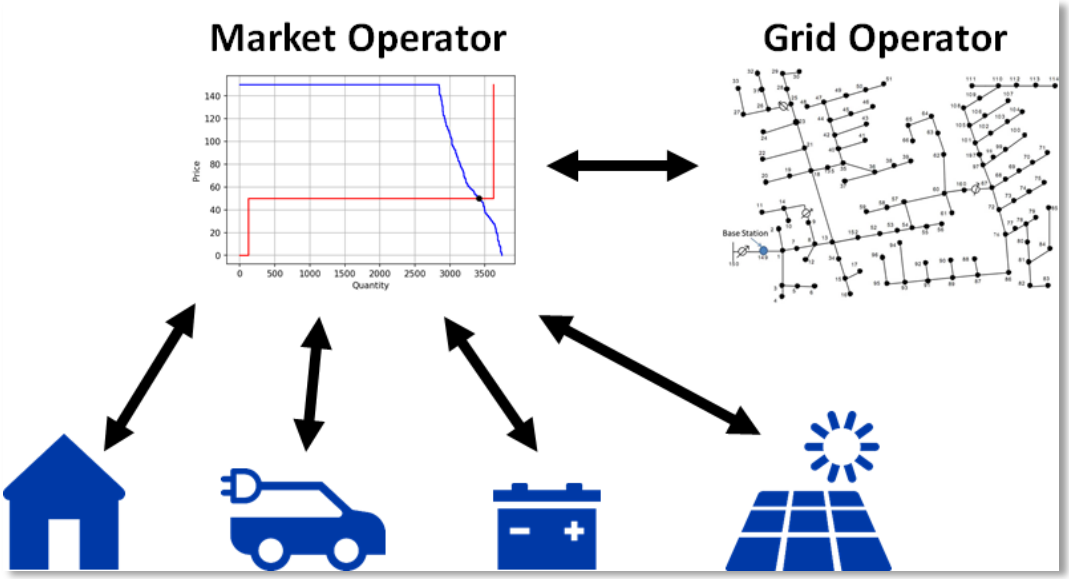

Source: SLAC National Accelerator Laboratory

Local energy markets have been lumped together under the term "transactive systems," described as loads that not only react to signals by the system operator but also actively contribute to price signals and commit to resulting dispatches. Before 2013, there were two largely independent but closely related transactive research programs underway, one in North America led by a team at PNNL, and one in Europe led by a team at ECN in the Netherlands. From about 2001 to 2013, PNNL led a series of research projects to develop the basic concepts and the first system design, called Transactive Control.

One of PNNL's research projects was the Olympic Peninsula Demonstration Project, which operated in 2006 and 2007 and produced groundbreaking results. Those results included 5-minute real-time retail price signals (using a close-loop bid/response mechanism), the discovery that peak loads could be reduced by as much as 60 percent in the short term (for example, less than an hour) and about 15 percent over the long term (for example, 1 year), and that the system was extremely resilient in cases of significant resource loss.

In this project, that work was extended by taking an economic perspective; it presented an institutional framework, analyzed and simulated customer economic objectives in local retail markets, and determined the long- and short-term effects of such a system under different penetrations of renewable sources and flexible loads.

Participating Agents

The general setup of a local energy market and their interactions between agents is illustrated in Figure 3 and includes several key agents.

Consumers and Appliances

Consumers form the demand side of a local energy market. An agent is defined by the combination of an appliance and its owner, a person, or a household.

In general, agents can be grouped as either flexible or inflexible. Flexible agents can shift their operations while inflexible agents cannot (or choose not to). Theoretically, every device can be flexible if, by either automation or manually, the consumer makes consumption responsive to price. It is theorized that local energy markets can only be sustainable if participating loads can be automated to bid on behalf of customers and dispatch according to market signals. Potential flexible consumers include thermostatically controlled devices, batteries, and electric vehicle chargers. All other loads (for example, lights, entertainment, or cooking) as well as appliances which are part of the former list but not automated, are assumed to be inflexible. Other potentially flexible appliances would include scheduled equipment like pool pumps and defrost systems.

Mapping customer preferences to an actual bid price is not trivial. In general, the consumer does not have a preference for the dispatch of the appliance itself, but rather for the service it provides. This needs to be transformed into a bid that can be related to the products traded on the local energy market (for example, KW in the next market interval). A customer j's willingness to pay for the electricity for the appliance is denoted by v_j^s.

Some general principles underlie bidding choices in the subsequent simulation. Chapter 5 will further describe the implementation chosen for this analysis which is, however, only one example of these general principles:

- HVAC Systems. The consumer's willingness to pay should increase with the difference of the indoor air temperature to the desired set point, which should be zero if at the setpoint. Additionally, since a house also provides thermal storage, cost optimization in time can be pursued and the opportunity costs of charging considered. If the agent expects prices to come down soon, it might delay cooling or heating. The quantity bid corresponds to the rated power of the HVAC system; for example, q_{heat} or q_{co} those devices cannot usually modulate their power input continuously.
- Batteries. Batteries can generate profits from arbitrage. In other words, they should charge when prices are low and discharge when prices are high. The quantity part of the bid is restricted to the interval between zero and the maximum rate of charge and discharge u_{max} and depends on the current state of charge.
- Electric Vehicles. If the state of charge is low, then the owner will be, *ceteris paribus*, willing to pay more than when the state of charge is already high. Also, willingness to pay should be higher closer to the time of departure. As for normal batteries, the quantity part of the bid is restricted to the interval between zero and the maximum rate of charge u_{max} and depends on the current state of charge.

Suppliers

Suppliers include the wholesale market as well as owners of decentralized generation and batteries on the distribution level. Those are generally independent from the local utility but, depending on the institutional framework and the regulation of grid and generation levels, are not necessarily so. The marginal cost is defined by mc_i^s as the quantity defined by an additional unit of electricity (in kW) by a supplier i.

- Wholesale Market Supply. The wholesale market supply is provided at the competitive wholesale market price pr^{WS} at the respective node. The supply is constrained by the transformer capacity connecting the distribution system to the upper voltage levels.
- PV Panels. As long as revenue from generation is positive, the PV module is incentivized to supply energy at the current generation. The actual behavior will, in practice, be very dependent on the regulatory framework, such as net metering, taxes, or a feed-in tariff.

Grid Operator

The grid is owned by the utilities. The grid operator is the California Independent System Operator (California ISO) that manages the transmission grid and ensures power-supply quality. For that purpose, the voltage and the load of sensitive appliances in the grid such as transformers is measured by the operator. The operator provides input to the market operations by supplying the control variable, e.g. the load of the transformer connecting the distribution system to the transmission system at higher voltage levels.

Market Operator

The market operator collects the bids of flexible appliances in the distribution system and clears the market. In principle, the grid and the market operator can be identical. The decision of what agent represents either of them can significantly affect the distribution of savings achieved by the local market and the available incentives. For instance, if the grid and the market operator are identical, the operator has an incentive to under-invest in the infrastructure and collect high congestion rents instead. Other options are separated grid and

market operators where the market operator could be a community choice aggregator (CCA) or another independent party.

Market Design

Market design offers a variety of aspects that market operators must choose among. The market design therefore defines the "rules of the game" and determines how consumers will both participate in the market and be rewarded for it.

The following market design aspects (among others) are of importance:

- Product definition (for example, energy or power)
- Type of order book
- Market intervals
- Eligible participants
- Bid structure
- Market-clearing rules
- Outside options

The market design for this project, is described here:

Product Definition

The project considered a power-only market. On the demand side, market participants submitted the power quantities they wished to consume during the upcoming market interval. On the supply side, generators and batteries submitted the power that they were willing to supply. Likewise, the market operator placed a bid representing power potentially supplied by the wholesale market, which is restricted to the available grid capacity. The available grid capacity was the reason that an energy-only market was not chosen.

Type of Order Book

A discrete auction order book is used in this project. At each trading interval, the market operator collects the bids submitted, determines the market results, and empties the order book. Non-served bids need to be re-submitted. An alternative is a continuous order book where orders get cleared as supply and demand match, and unserved bids are kept until they either expire or clear.

Market Interval

An interval of five minutes was used to clear the market. In general, the shorter the interval, the easier it is for market operators to control grid constraints and the better it is for all agents to adjust to new information like changes in unresponsive load, prices, or forecasts, all of which contribute to greater efficiency and faster dynamic system response. On the other hand, many devices are subject to must-run times, which complicates bids when must-run times exceed the market interval. Also, a higher market clearing frequency significantly increases the amount of information to be exchanged, potentially causing problems from communication lags in the information system infrastructure. In this project, the market interval of five minutes was selected because the most significant resource, HVAC systems, had a must-run duration of not greater than five minutes.

Eligible Market Participants

Every consumer who is electrically connected to a geographic area is eligible to participate in the market. Consumers may participate in the market with specific appliances (for example, their electric vehicles) and choose outside options for others.

Bid Structure

Bids are submitted separately for each participating customer and each appliance at each market interval for which the customer desires to consume electricity. The bid contains information on the desired power level of consumption (in kW), the maximum willingness to pay (for demand), or the minimum price to accept (for supply) in dollars per Megawatt-hour ($/MWh).

Market Clearing Rule

The market clearing rule is the rule where market operators match demand and supply to determine the clearing price. This approach for today's energy wholesale markets was to clear the market by surplus maximization, under the assumption that the bids will coincide with the value that agents assign to the good, specifically $pr^d = v^d$ for demand and $pr^s = mc^s$ for supply:

$$SW_t = \sum_j \left(pr^d_{j,t} - \lambda_t\right) q^d_{j,t} + \sum_i \left(\lambda_t - pr^s_{i,t}\right) q^s_{i,t}$$

$$s.t. \sum_i q^s_{i,t} = \sum_j q^d_{j,t}$$

where $pr^d_{j,t}$ is the bid for the *j-th* demand at time t, $pr^s_{i,t}$ is the bid for the *i-th* supply at time t, λ_t is the overall common price at time t and $q^s_{i,t}, q^d_{j,t}$ are the quantities for the *j-th* demand and *i-th* supply, respectively, at time t. The surplus or social welfare (SW) maximization is subject to the equality of cleared demand and supply bids by quantity. Given the discontinuous nature of the demand and supply curves, if they intersect the price is generally determined by the highest allocated supply bid. The price determined is a uniform price. Bidders with supply bids higher than the clearing price and demand bids lower than the clearing price exit the process unserved.

Figure 24: Order Book and Market Clearing

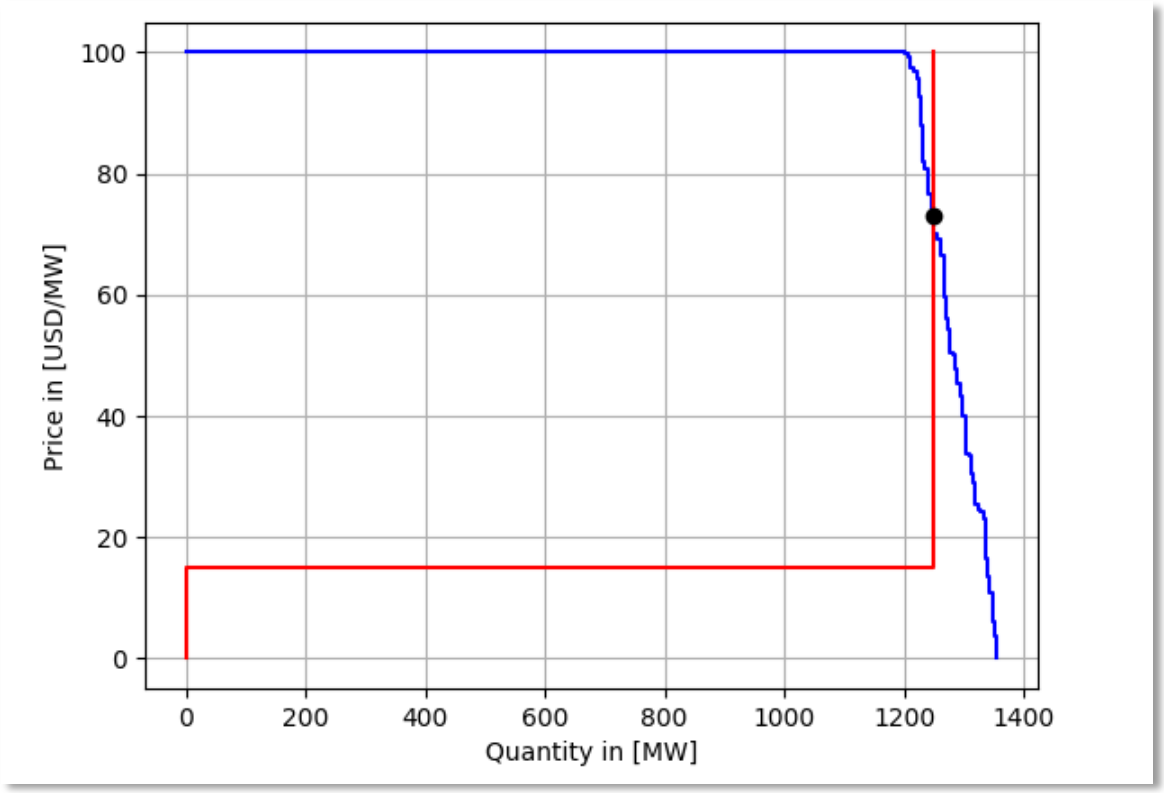

The equilibrium price results at the intersection of supply and demand function

Source: SLAC National Accelerator Laboratory

Outside Option

Markets do not generally have outside options. Considering current regulation and policy objectives, it is reasonable to assume that consumers may choose (for distinct appliances) to stay on a fixed retail tariff. However, flexible appliances and responsive customers have a unique incentive to subscribe to flexible prices if they can take advantage of low price-periods. These options are compared in more detail in Chapter 6.

Welfare Analysis

A local energy market can increase welfare by achieving efficiency in the short-term, through long-term use of scarce grid infrastructure, or the efficient use of supply.

- Short-Term Use of Scarce Grid Infrastructure. The first aspect refers to the fact that, without a retail market, the system operator must decide on load curtailments not based on economic values but rather on other factors, such as safe operations. The retail market allows participants to transparently determine the value they attach to their dispatches so system operators can perform necessary curtailments based on economic principles, reducing overall system costs.
- Long-Term Use of Scarce Grid Infrastructure. The second aspect covers efficient investment in the grid infrastructure. If the retail market is able to consistently decrease the peak consumption in the area of interest, significant infrastructure investments can

be deferred or foregone altogether. This reduces grid investment and, potentially, fixed cost components and grid tariffs for consumers.

- Efficient Use of Supply. Third, a retail market also makes the actual cost of supply visible to customers because the wholesale market supply enters the local market as a supply bid at actual wholesale market costs. If customers are only facing fixed retail tariffs, retail companies are forced to purchase electricity regardless of current wholesale prices. Those costs are then redistributed to customers. If customers reacted to high wholesale market prices, the overall costs of electricity procurement would decrease and customer prices would be reduced. However, the distribution of achievable benefits could vary depending on the number of customers and devices participating in the local energy market and the levels of cross-subsidization between different customer segments. In Chapter 6, welfare estimates are provided. For a full welfare analysis cost, components must also be considered, though those components are not part of this report.
- Upfront Investment. The introduction of a local energy market requires a level of investment, including the costs of setting up information systems and a marketplace as well as maintenance costs by the market operator. The introduction of a local energy market makes sense if the net welfare is expected to be positive.
- Transaction Cost. Transaction costs cover various costs of the operation of a local energy market. This includes information costs for participating consumers (for example, regarding the necessary hardware, the operations, and the settings), the operation and maintenance of the information system, and the market platform.

Distributional Effects

Even if the overall welfare effect is positive, distributional effects are important for the success of local energy markets. If customers are able to switch back to an (unchanged) fixed retail tariff, local energy markets cannot make them worse off. If, however, customers are forced to participate, or if investment costs for a local energy market are significant, customers can be worse off if they do not have the load flexibility to respond to increased local market prices. Additionally, if the grid and the market operator are identical, there are adverse incentives for efficient grid investment. As described in the previous section, the market operator is then able to earn congestion rents (the difference between the supply and the demand price) at a profit without reinvesting the money earned into new grid infrastructure.

CHAPTER 5:
Simulation Platform

Introduction

A modularized simulation platform was created to represent a large-scale system. As described in previous chapters, the decentralized structure proposed to control the electrical network has three major components: the electrical system, composed by the grid and the residences and commercial buildings, the home hub, and the cloud coordinator. The architecture described is summarized in Figure 25.

Figure 25: Setup of the Cloud Coordinator — Home Hub Architecture

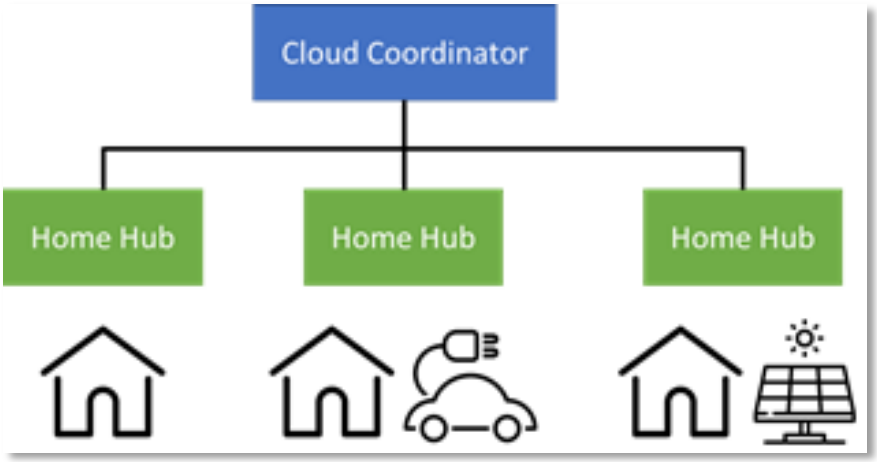

Source: SLAC National Accelerator Laboratory

In the simulation platform, the same system is divided into two main parts: the physical and the economic. The physical layer models the full electrical system, while the economic layer models the retail market, including both home-hub functionality (which calculates the bids) and the cloud coordinator functionality (which implements the market, clears the market, updates prices, and computes costs).

The physical layer is implemented in GridLAB-D (Chassin, Schneider and Gerkensmeyer 2008) (Chassin, Fuller and Djilali, GridLAB-D: An agent-based simulation framework for smart grids 2014) to make use of the multi-purpose models developed to simulate the electrical network, residences and their appliances, renewable sources, and storage units. GridLAB-D is an open-source, well-established tool developed by Pacific Northwest National Laboratory and SLAC National Accelerator Laboratory, with funding from DOE and the Energy Commission. In the simulation tool, the physical states of the system are tracked and updated at each time interval. The economic layer is coded in Python and divided into the agents' and the central market fields. The agents' side includes the HHs of all participating houses. At the beginning of each time interval, the HH reads the current state of the respective house and calculates bids for each appliance participating in the market. All bids, including supply-demand bids and unresponsive loads, are added to a Python market object. Instead of leveraging existing GridLAB-D functionalities, a new Python module called "market" was created, with greater

flexibility and extensive possibilities for the type of bid, clearing rules and tie breaking, handling of edge cases, and plotting. After all bids have been gathered, the market is cleared, and the information published. Based on this information, the HHs determine the dispatch of appliances and pass the new state back to the physical layer of the simulation. Figure 26 depicts the architecture of the simulation platform.

Figure 26: Architecture of the Simulation Platform

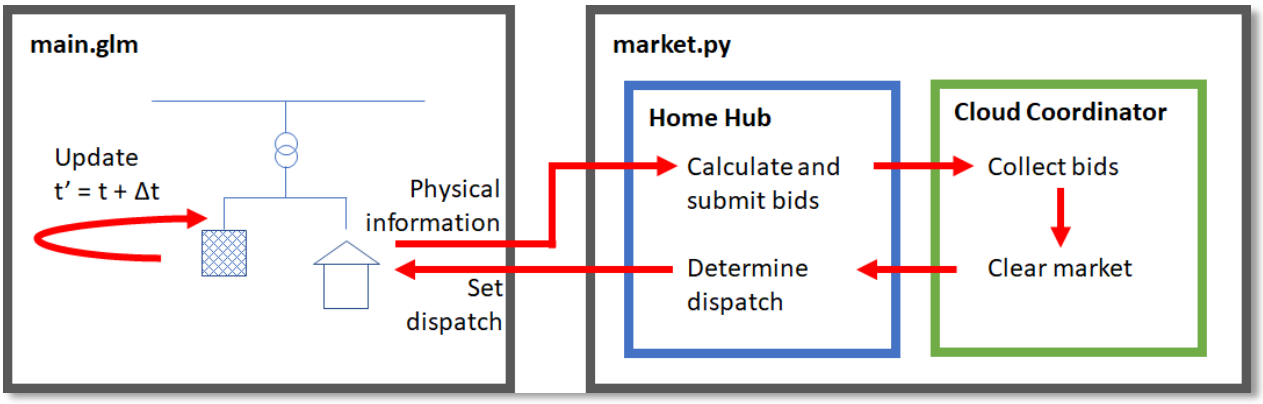

Source: SLAC National Accelerator Laboratory

Implementation

GridLAB-D capabilities have been extensively developed to facilitate communication between the physical and market layers in the simulation platform. As result, GridLAB-D is enabled from a Python shell. This has significant advantages. First, it is possible to use Python to dynamically read out information from, and writing information to, the GridLAB-D model. Second, the Python environment was able to store global variables throughout the simulation. This strategy accelerated the simulation process. Before this improvement, GridLAB-D exchanged information with Python routines through databases. Between simulation steps, each variable had to be stored either in the GridLAB-D model itself or locally in a database or CSV file. The information included the activity in the previous market interval or historical prices, among others. This made the coding easier, more transparent, and more accurate.

To simulate a local market, a new market module in Python was created. The Python framework was used because it is a language used by many people in both the industry and academia, and it easily handles larger data sets, allowing for the efficient calculation of bids and market results. Additionally, the routines used in the HHs and the cloud coordinator deployed in both the laboratories and the residences use Python code. Using Python, in both instances, makes translation of the algorithms (tested in the simulation tool to the final system) simpler.

Furthermore, a docker image was deployed for GridLAB-D, which is a tool that provides software in virtual containers. GridLAB-D has many different components and is sensitive to machine specifications. A docker image is a virtual environment that ensures that different versions work identically across various machines. The simulation was running within the docker container, and the results were dumped to the local machine. Figure 27 illustrates this new setup.

Figure 27: New Setup of Simulation Platform Architecture

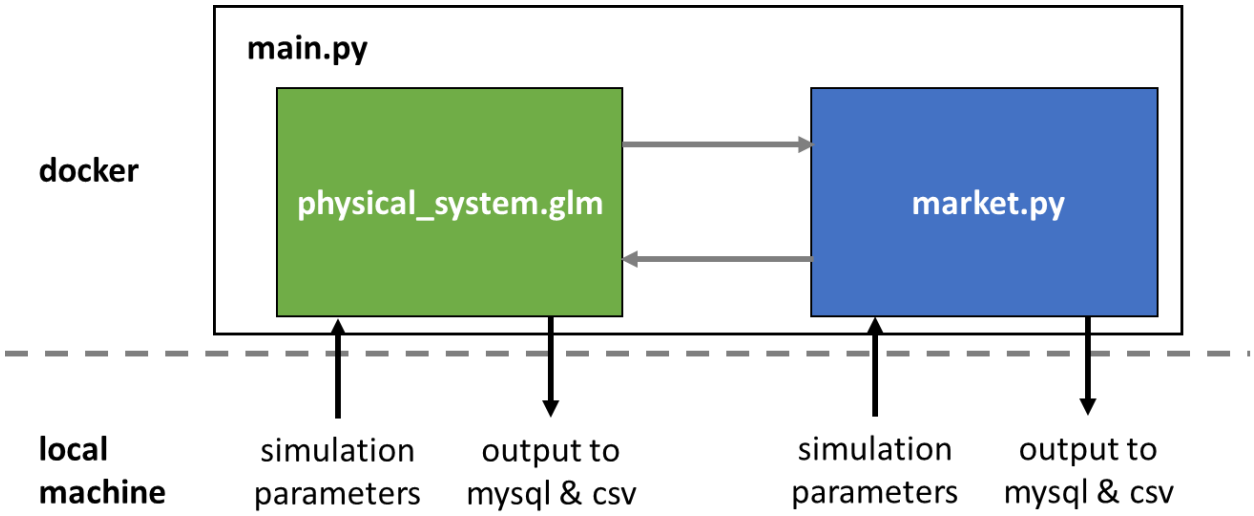

Source: SLAC National Accelerator Laboratory

Physical Layer

To analyze and test the control algorithms in a large-scale system of 1,000 to 10,000 houses, GridLAB-D was used to implement the physical layer of the simulation platform. This setup allowed researchers to include the house model with all the devices and appliances and communication systems in the simulated electric distribution network. It was important to validate the device and appliance models in the house representations in GridLAB-D to obtain simulation results close to the future outcome when the equipment and algorithms are deployed to the actual electric grid. These houses were allocated in an area that can be electrically connected following the network topology of a distribution feeder.

The main characteristics of the physical layer are summarized here and a more detailed explanation is presented in Appendix B.

- Distribution Feeder. The distribution network used in the simulation is the standard IEEE 123 feeder (IEEE 2010). It represented a typical residential radial distribution grid and operated at a nominal voltage of 4.16 kV, with total power of about 3.6 Megawatts (MW). The feeder was connected to the overlaying voltage level by a single transformer, site of most potential congestion. The feeder itself branches out and represented multiple streets with electrical loads connected at regular intervals.
- Residences. To represent the houses, six different models were created based on one- or two-story units equipped with three possible cooling and heating systems. The floor areas and temperature set points were randomly generated. The statistics are based on mean and standard deviation values as reported by the Residential Energy Consumption Survey by the U.S. Energy Information Administration (EIA) (EIA 2009). Furthermore, to specify the thermal characteristics of the building stock, parameters (as reported by the Energy Star ® Home Sealing Specification of 2011) (ENERGY STAR 2011) were used, following specifications for Southern California (IECC climate zones) (PNNL 2012). These include the rate at which air is exchanged with the exterior.

- Residential Load and Appliances. To evaluate residential base load and appliance use, time-series data were created based on smart-meter data published by the Pecan Street Data Project (Pecan Street 2015). While the majority of the project focuses on consumers in Texas, the dataset from San Diego as of 2015 was chosen because, during that time, the greatest number of year-long profiles of townhomes and single-family homes was available.
- Photovoltaics. Houses equipped with PV were randomly drawn. The panels were sized so that they covered the mean daily energy use of a house in July, given the maximum available rooftop area.
- Batteries. Batteries were scaled to store mean daily-energy usage or daily PV production. Batteries were placed only in houses equipped with PV panels since the economic incentives were greatest.
- Electric Vehicles. Stationary charging stations were modeled by batteries connected to the physical system when an electric vehicle connected. Real data for residential customers in California, provided by a large charging operation company, was used to identify connection and charging times and estimate the average charging rate and the SOC. Data available throughout the year were not available, but event series were bundled for the days of interest.

Economic Layer

Home Hub

The HH measured all relevant physical state variables in a house, determined the willingness to pay for each flexible appliance to dispatch, controlled committed loads, and communicated with the cloud coordinator. In Appendix C, the implementation of the HH, the performance evaluation of the communication and the sensing of variables using the eGauge monitoring system are presented.

This subsection elaborates upon which physical state variables were measured and how they entered into the process of willingness to pay, which is the basis of an appliance-specific bid. The focus is on how the HH builds bids for each appliance where the bid is a tuple (a finite list of ordered elements) of a maximum price (willingness to pay), and the desired power at which the device is willing to consume over the following market interval. These bids are sent to the cloud coordinator by the local HH and signal the priority of the appliance to be dispatched. Table 1 lists the relevant electrical appliances that are flexible in dispatch and can therefore be actively controlled through a local market. The signal refers to the state variable in the system used to generate the willingness to pay in the bidding process, while the setting describes the criteria leading to a willingness to pay. The sampling time defines the time interval that the signal is sampled and post-processed during the simulation and display.

Table 1: Relevant Flexible Electrical Appliances

Appliance	Signal	Settings (specified by user)	Measurement Device	Sampling Time / Post-Processing	Use
HVAC system	Ambient temperature	Comfort factor	Nest thermostat	1 min / none	Display / bidding
Battery	SOC Active power (branch 1)	Price thresholds	eGauge	1 min / MA filtering	Display / bidding
Electric vehicle	SOC Active power	Estimated time of departure	Charging station	1 min	Display / bidding
PV system	Active power	—	PV system	1 min	Display

Source: SLAC National Accelerator Laboratory

The dynamic model of each appliance or device participating in the market is described in Chapter 3. In this section, the criteria followed to generate the bid for each appliance are described.

Thermal, Heating, Ventilation, and Air-Conditioning Systems

Thermal systems are able to participate in transactive systems if they are electric. In cases where only one side, either cooling or heating, is electric, only the electric one will be active (for example, electric cooling and gas heating). For the following explanations, the research team assumed that both sides of the unit were electric.

The primary signal used to generate the bid was the ambient temperature T_{amb}. Residents set maximum temperature T_{max}, the minimum temperature T_{min}, and a temperature range ΔT between heating and the cooling set points where the system did not operate. This represents a general case, so that any other bidding situation involving only heating or cooling could be deduced. Additionally, customers specified a curve linking temperature sensitivity to their willingness to pay. For instance, a steeper curve indicates a sensitive customer and means that a small deviation from the set point translated to a high willingness to pay for electricity.

Accordingly, the function used to generate the bid is a function of the error between the ambient temperature T_{amb} and the limiting temperatures of the comfort zone, $T_{amb} - T_{co}$ or $T_{amb} - T_{he}$. The function increases in the case of the bid for the cooling system and decreases for the heating case. The reasoning is that the customer is willing to pay more if the ambient temperature diverges from the comfort zone. Figure 28 depicts this example and Table 2 summarizes the bidding function.

Figure 28: Temperature-Dependent HVAC Bidding Function

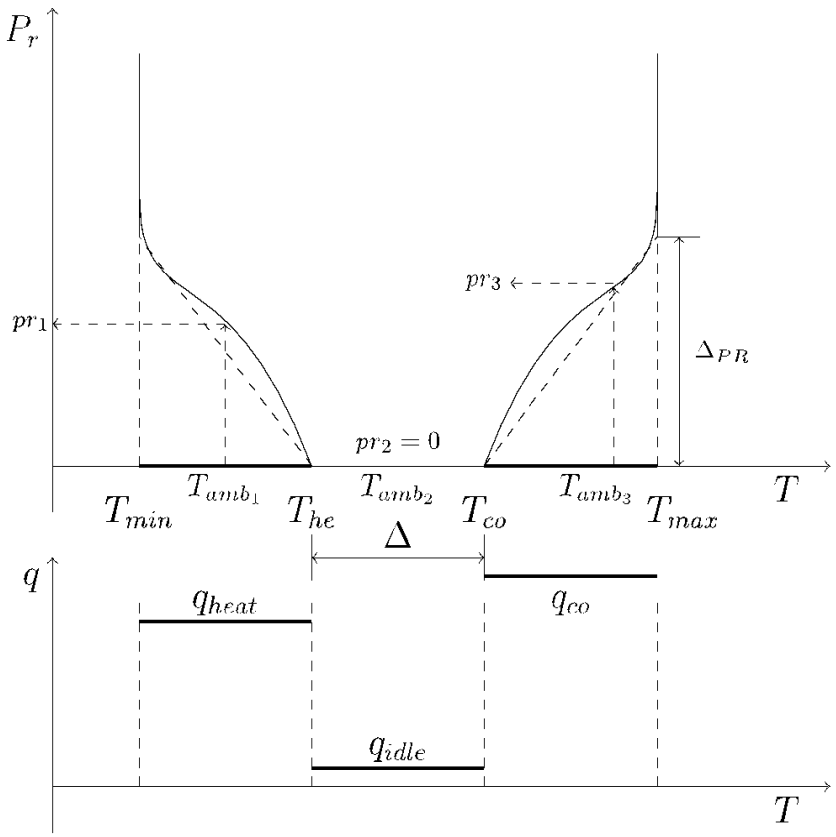

Source: SLAC National Accelerator Laboratory

Table 2: Summary of Bidding Functions for the HVAC System

Operation	T_{amb}	Willingness to pay	
Heating	$T_{min} < T_{amb1} < T_{he}$	$pr_1 = f_H(T_{amb1} - T_{he})$	$f_H(.)$ decreasing function
Comfort Zone	$T_{he} < T_{amb2} < T_{co}$	$pr_2 = 0$	
Cooling	$T_{co} < T_{amb3} < T_{max}$	$pr_3 = f_C(T_{amb3} - T_{co})$	$f_C(.)$ increasing function

Source: SLAC National Accelerator Laboratory

It is possible to deduce from the bidding curves that if the ambient temperature is either $T_{amb} > T_{max}$ or $T_{amb} < T_{min}$, the load becomes nonresponsive. In the implementation of the HH algorithm used in the simulation, the bidding functions $f_C(.)$, $f_H(.)$, are linear with respect to temperature, as shown by the dotted lines in Figure 28. For heating operations, the bid is calculated by:

$$pr = \frac{\Delta_{PR}}{T_{max} - T_{co}}(T_{amb} - T_{co}).$$

For cooling, the bid is defined by

$$pr = \frac{\Delta_{PR}}{T_{he} - T_{min}}(T_{amb} - T_{he}).$$

The thermal system controls the power, q, by switching the power between two states: power on or power off. In Figure 28, power off is associated with the power level, q_{idle}, while power on is defined by the power levels, $q_{heat}$and q_{co} for heating and cooling modes, respectively. In the bidding process, the quantity component equals the maximum power during the heating/cooling period, q_{heat}, q_{co}, respectively. For cooling, bidding is (pr_3, q_{co}).

Battery System

Batteries do not serve a comfort purpose but, from the customer's perspective, they are employed for price arbitrage or cost minimization for demand peak rates or time-of-use rates. Profit-maximizing bids are built upon expected opportunity costs of charging or discharging. While various bidding approaches are possible, in this project the opportunity costs were estimated by ranking day-ahead prices for the upcoming 24 hours and, based on the maximum state of charge of the battery, the hours during which charging and discharging were most profitable. The lowest price during the hours of discharging were determined by a lower threshold of selling bids. Likewise, the highest price of the hours of charging was used as the highest threshold for buying. Then, in the bidding process, the battery submitted two bids (given that the current SOC was between the minimum and maximum SOC), a buy and a sell bid. In this bid, the quantity component equaled the maximum power delivered by the power inverter during charging and discharging the battery and the price bid corresponded to determined price thresholds.

Whenever the price fell below the lower buy bid, the battery was allocated for buying and charging. Whenever the price was above the higher threshold for selling, the battery was allocated for selling and discharging. This algorithm is described in **Error! Reference source not found.**, where it is assuming the clearing time of the market (T_{cl}) is equal to the time interval of the day-ahead price. The approach is illustrated in Figure 29, where the threshold prices for selling and buying are shown. Other bidding strategies could include a price proportional to the SOC or, as a result of optimal scheduling, a price based on the price forecast.

Table 3: Algorithm to Determine Buying/Selling Thresholds

Algorithm Steps	Actions
Step 1	Get Day-Ahead price vector
Step 2	Sort price vector ascending by value (Day-ahead price sample at $T = T_{cl}$)
Step 3	Divide ($SOC_{max} - SOC_{min}$) by the ramping constraint u_{max} and the market interval T_{cl} to determine the number of effective charging/discharging periods N ($\frac{\Delta SOC}{\Delta t} = \frac{SOC_{max} - SOC_{min}}{N T_{cl}} = \frac{\eta p_{max}}{E Q_{max}} = u_{max}$ then $\frac{SOC_{max} - SOC_{min}}{u_{max} T_{cl}} = N$)
Step 4	Use *N-th* position of the price vector as threshold for buying
Step 5	Use *(24 hours / T_{cl} - N)-th* position of the price vector as threshold for selling.

Source: SLAC National Accelerator Laboratory

Figure 29: Price Thresholds for Battery Dispatch

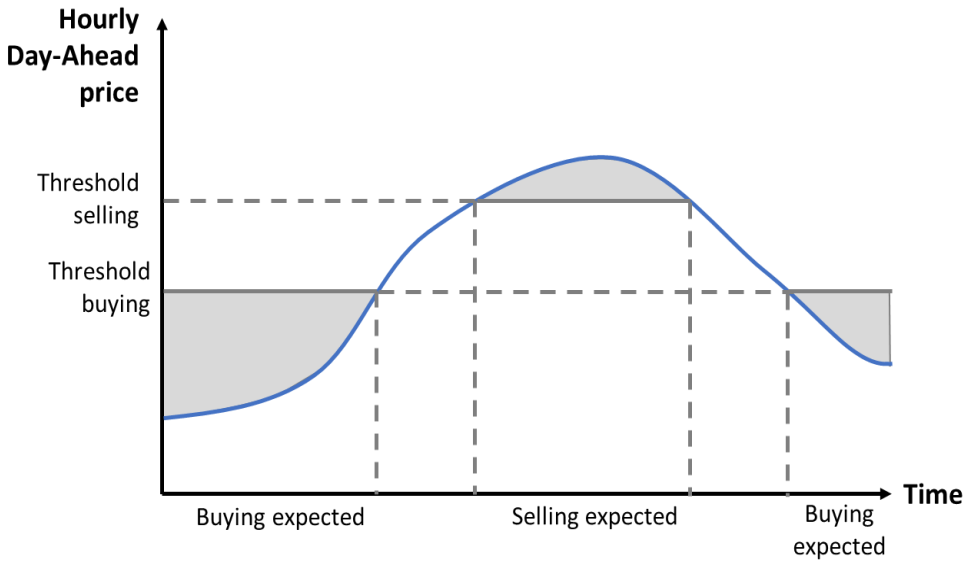

Source: SLAC National Accelerator Laboratory

Electric Vehicle Charging Stations

Electric vehicle charging does not serve the purpose of profit maximization through price arbitrage. Instead, the vehicles should be charged at minimum costs. Two types of vehicle charging stations are described: commercial charging stations and workplace or home charging stations.

For commercial charging, it is assumed that drivers want to charge as quickly as possible. Therefore, the charging station, on behalf of the driver or the respective retailer, bids the maximum price to the market. The quantity component corresponds to the rated power of the charging station.

For workplace or home charging, however, drivers do not generally demand charging as soon as possible so are willing to delay charging to reduce electricity costs. Accordingly, the price bid depends on the SOC of charge of the battery and the expected time of departure. To determine priorities, the ratio between the residual empty storage capacity and the expected time until departure was calculated. The result gave the average power at which the car needed to be charged to be fully charged at the time of departure. If that magnitude is at or above the rated power of charging $\frac{\Delta SOC}{\Delta t} \geq u_{max}$, the car owner will be bidding the maximum price p_{max}. If below, it is willing to pay a lower price, depending on the car owner's price sensitivity k, and the bid is $pr_t = pr_{max} - k \cdot (u_{max} - \frac{\Delta SOC}{\Delta t})$.

Figure 30: Electric Vehicle Bidding

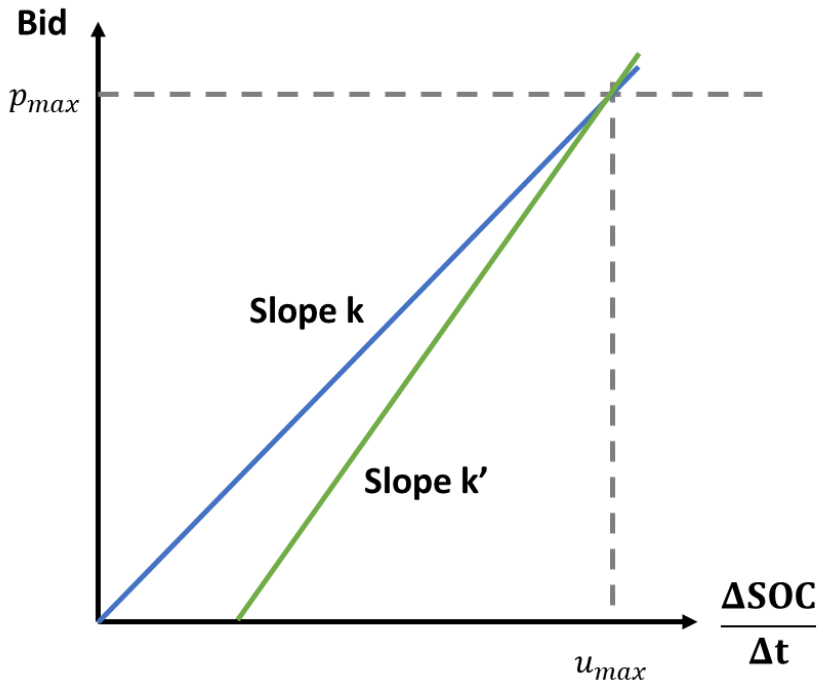

Electric vehicle bidding, depending upon departure time and SOC

Source: SLAC National Accelerator Laboratory

Generation of Bids for Non-Committed or Non-Responsive Generation: PV Systems

In the residential PV system, the maximum power generated by the system was not controllable by the resident. This electricity was variable because of solar irradiation and weather conditions. Additionally, the resident could not connect or disconnect this device from the grid. Because this device is not controllable, it could not participate directly in the bidding process but could be included in the market as additional generation to improve market transparency. The HH communicates to the cloud coordinator the power generated by the PV system every five minutes to be included in the market's generation supply function. The income of the PV owner depends on regulation and might be tied to both the resulting local market price and a regulated fixed price.

Cloud Coordinator

The cloud coordinator is a centralized entity and creates a market that allocates controllable renewable resources. The cloud coordinator collects all bids, forecasts on inelastic demand, and builds the supply and demand functions used to clear local markets that allocate scarce local grid capacity. The result is communicated back to customers, who dispatch accordingly.

This implementation corresponds to a real-time power dispatch market with a discrete order book. First, the cloud coordinator forecasts the inflexible load for the next market interval and includes them into the demand function at maximum price pr_{max}. However, these loads are not facing the resulting market price but rather a fixed retail tariff. Additionally, flexible appliances submit their bids, consisting of tuples of maximum price and quantity (pr_t, q_t). The cloud coordinator sorts the incoming bids in descending order. The blue line in Figure 30

represents this downward-sloping demand function. The flat part of the curve spanning approximately 3000 kW represents the inelastic demand and the bids at maximum price. The downward-sloping part of the demand curve represents flexible and price-sensitive devices.

Similarly, the supply function is built from bids from both local energy supply (from solar panels and batteries) and supplies from the wholesale market at the real-time wholesale price, again consisting of tuples of minimum price bids and quantity (pr_t, q_t). In the example given in Figure 30, the flat part of the red supply curve represents the wholesale market supply at real-time costs. Supply from the wholesale market level is limited by the available grid capacity connecting the local and the overlaying grid level, at 3500 kW in Figure 31. The absence of any zero-price supply indicates that no rooftop PV is available. The absence of any matching quantities for supply and demand indicates that no battery storage is available, either.

Figure 31: Example of Supply and Demand Function in a Local Market

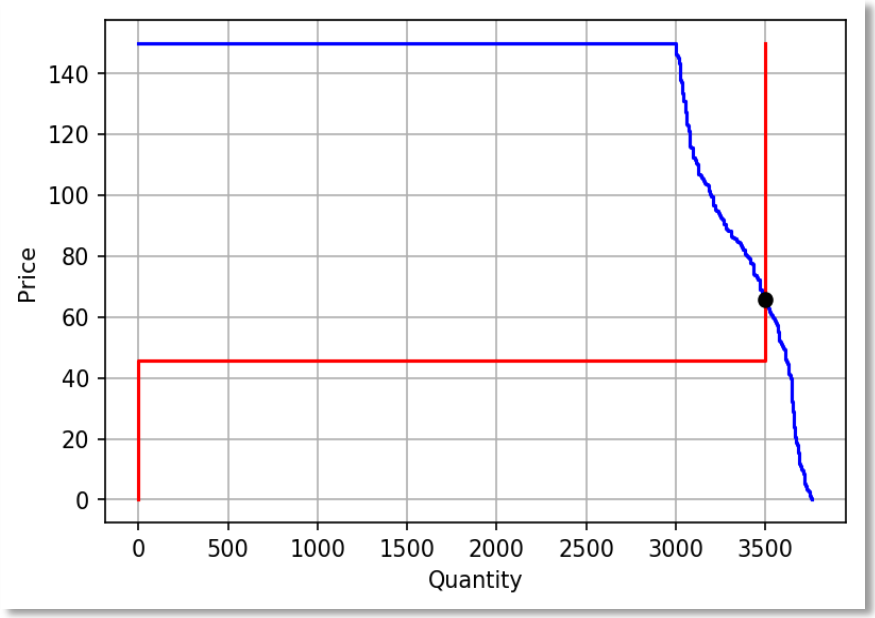

Source: SLAC National Accelerator Laboratory

The market is cleared at the intersection between supply and demand. This is done by sorting the bids according to their price and allocating supply and demand in equal quantities, as long as the demand bids are higher than supply bids. Once the demand price is smaller than the next supply bid (which would have to equal the demand and supply quantities allocated), the algorithm stops.

If the connecting grid capacity is sufficient, the resulting local market price will equal the wholesale market price. However, if there is insufficient local supply capacity, as shown in Figure 32, the local energy market price will exceed the wholesale market price, with the difference being the so called "congestion rent" used to measure the value of additional grid capacity.

After market clearing, the cloud coordinator publishes the result. That can either be through the publication of the clearing price alone or through specific successful bids. Appliances react accordingly and realize the market result, effectively restricting demand on the local level to the degree possible.

In some cases, especially in high-demand times, the market clearing price will saturate at the maximum price, as depicted in Figure 32. Then, the mere publication of an equilibrium price would result in all loads dispatching at the maximum price, violating the technical constraint of 3000 kW. Instead, a clearing rule that calls out cleared appliances ensures that tie-breaking is efficiently done, so that the technical constraint is respected.

Figure 32: Local Market Clearing at Maximum Price

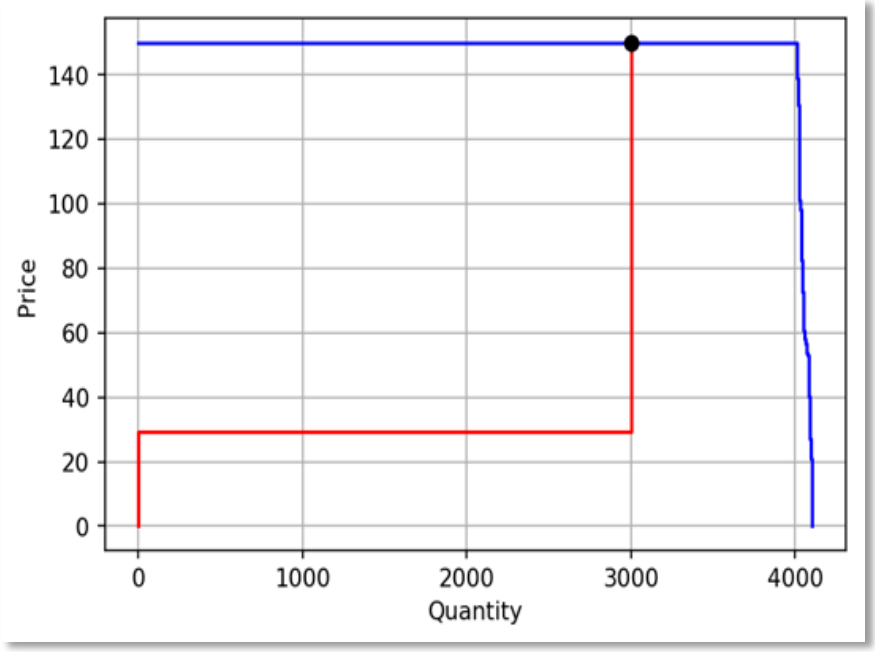

Source: SLAC National Accelerator Laboratory

Given sufficient forecasting quality of inflexible loads, the market will be capable of effectively managing grid constraints. Figure 33 features a time simulation to show how the system load can be effectively controlled around maximum capacity.

Figure 33: System Load on IEEE 123 Feeder With 3-MW Constraint

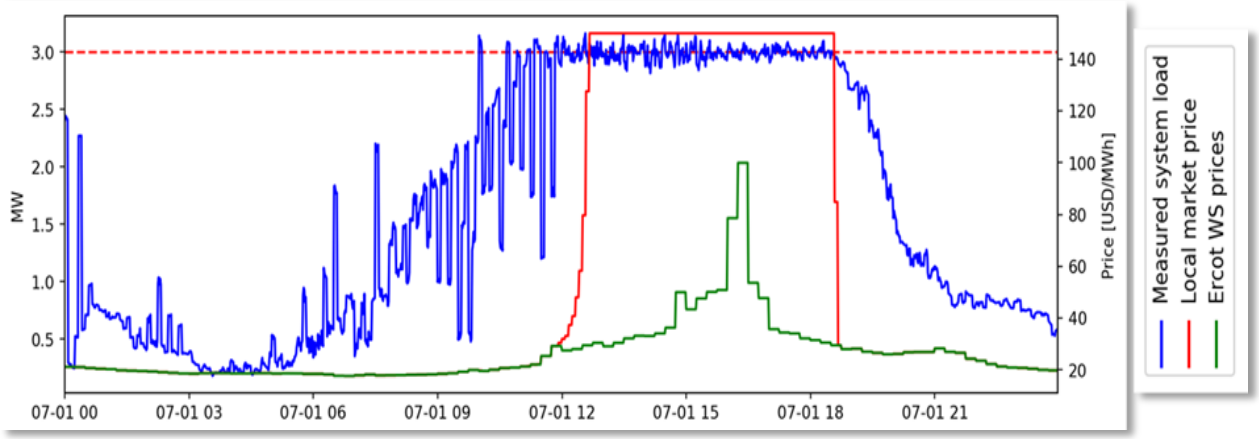

Load and prices for the IEEE 123 systems reaching the 3MW constraint

Source: SLAC National Accelerator Laboratory

The performance of the market depends on a number of other factors:

- A good forecast of the inelastic demand that does not submit direct bids is crucial. If the inelastic load is forecasted too low, too many elastic loads are allocated and the technical constraints could be violated, resulting in inefficiently low local prices. On the other hand, if the inflexible load estimate is too high, the price will be too high and available grid capacity will be under-utilized, decreasing welfare for all customers.
- Violations could also occur if physical load shapes significantly differ from submitted bids.
- Inelastic load alone could surpass the technical capabilities of the system, with no degree of freedom for market allocations. For this reason, in this simulation the system operator was able to randomly select loads to curtail when necessary, equivalent to traditional emergency load shedding.

CHAPTER 6:
Results

Market Results

This section analyzes the state of the system without a local energy market. The system was simulated using electricity consumption, weather, and price data from San Diego, California.

Analysis of Status Quo

As a first step, in the second weeks of January, July, and October and without a market or flexible appliances, an analysis was developed to better understand the electric load profile of a community of 1120 houses in a time of high stress. In this case, HVAC operation was solely regulated by internal temperature controls inside the house and was not price responsive. Figure 34, Figure 35, and Figure 36 show the measured aggregated loads at the connection point to the overlaying grid for a week (Monday through Sunday). Additionally, the figures depict the wholesale market price.

Figure 34: Measured Real Power and Wholesale Market Prices in January

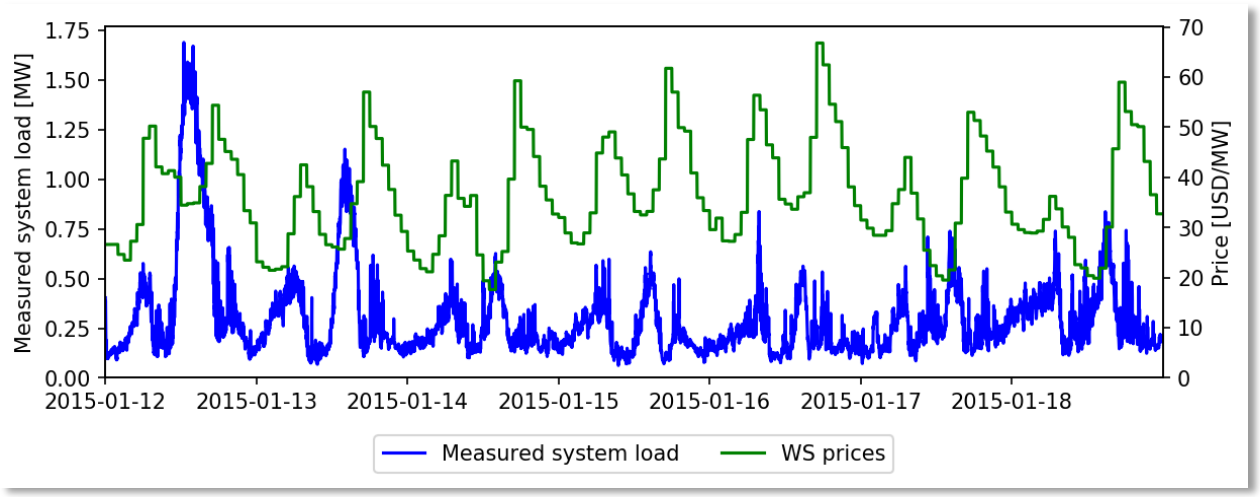

Load and prices for the IEEE 123 systems reaching the 3 MW constraint

Source: SLAC National Accelerator Laboratory

Figure 35: Measured Real Power and Wholesale Market Prices in July

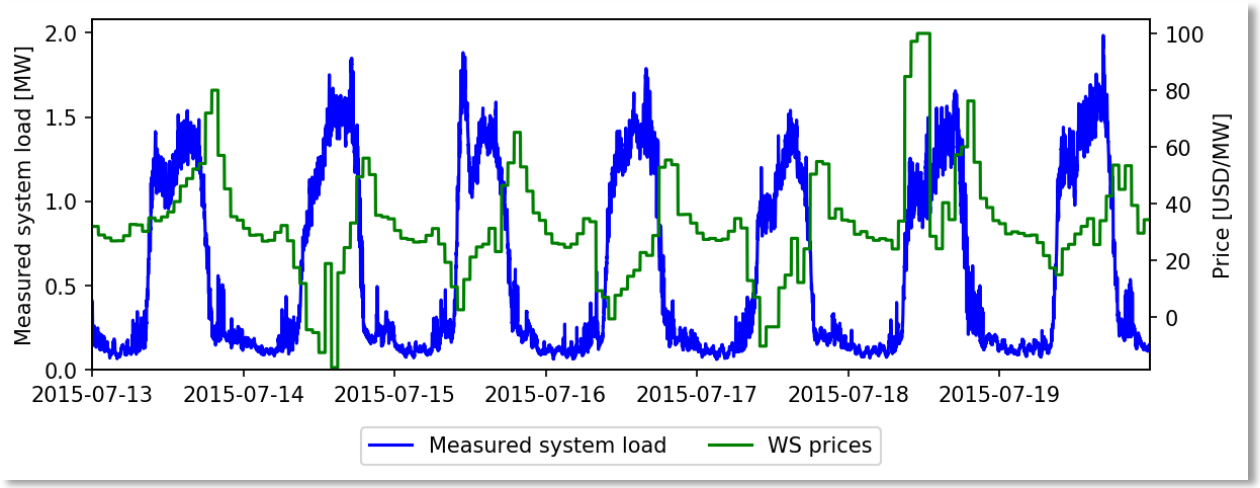

Source: SLAC National Accelerator Laboratory

Figure 36: Measured Real Power and Wholesale Market Prices in October

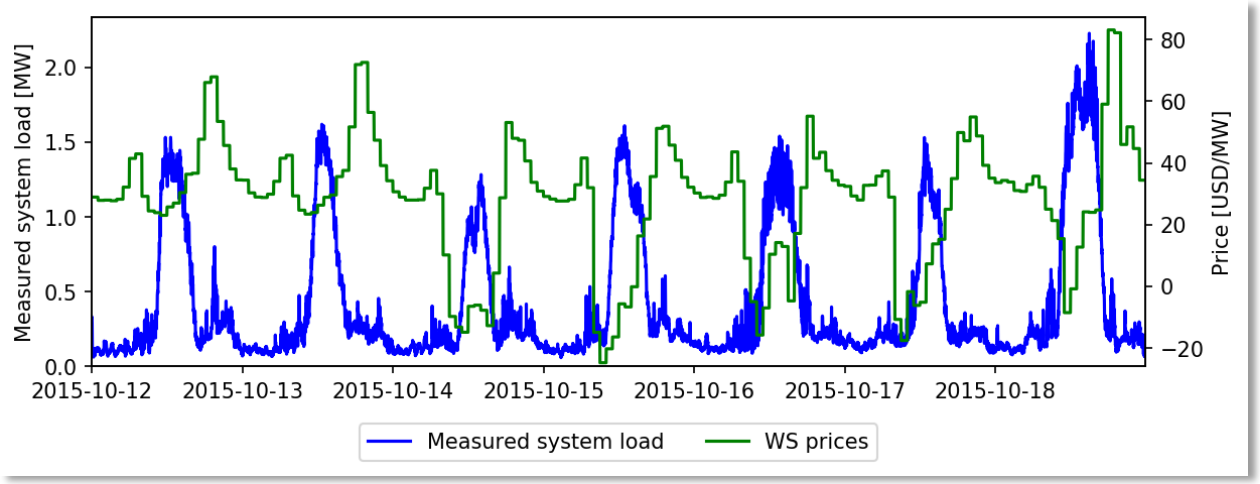

Source: SLAC National Accelerator Laboratory

This analysis provided the following insights. January had an average load of 0.299 MW, with a maximum of 1.688 MW. HVAC operations used a share of 43.0 percent of total energy consumption. In July, the average load was 0.582 MW and the maximum load was 1.984 MW. The energy of HVAC operations was 68.3 percent of total energy consumption. Finally, in October, the average load was 0.418 MW and the maximum reached 2.227 MW. The share of HVAC load was 56.2 percent.

For further analysis, the month of July was chosen due to its highest average load. In particular, July 16, 2015, was selected to investigate sensitivity to other parameters such as share of PV, batteries, or HVAC. In this analysis, two days were simulated, including July 15, and only data from Day 2 were used to ensure that physical appliances were in equilibrium or in a steady state, and that results were not affected by the simulation initialization.

To evaluate the system load under different PV penetrations, the percentage number gives the ratio between the number of houses equipped with PV and the total number of houses. As discussed earlier, PV systems are deterministically distributed to ensure comparability between

scenarios and sized according to the July mean daily energy use of a respective house. Figure 37 shows the impact on the net power delivered to the distribution network when PV penetration in changed from 0 percent to 50 percent on July 16, 2015.

Figure 37: Measured Real Power Under Different PV Penetrations, by Percent

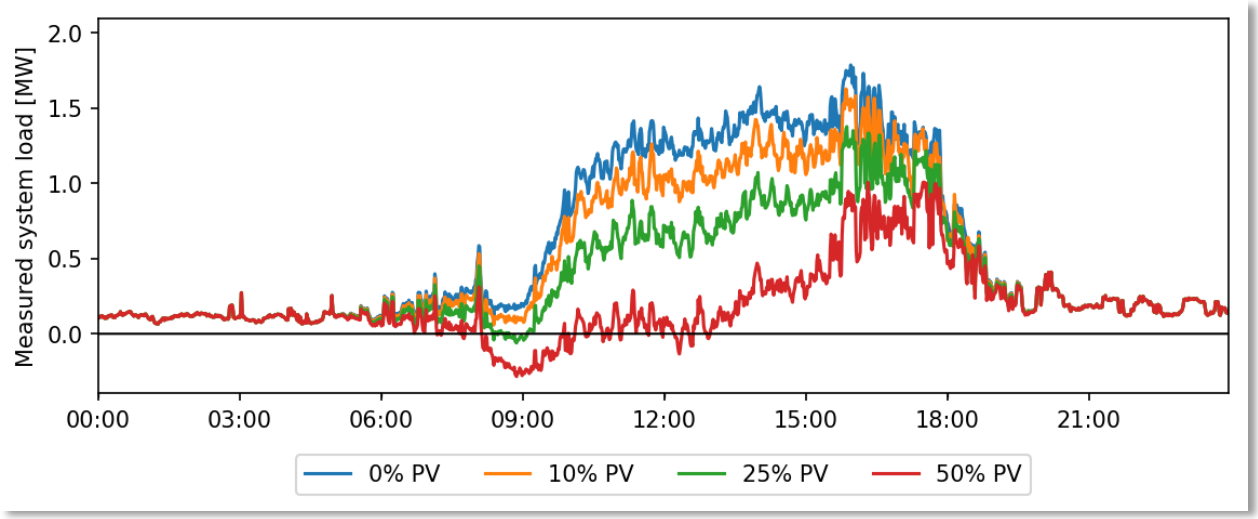

Source: SLAC National Accelerator Laboratory

The maximum net load experienced in the system under the various scenarios is summarized in Table 3. This table reports the maximum PV infeed, the energy produced per day by the PV system and the PV infeed as a share of total energy consumption. The last magnitude is defined by the ratio between the average power produced by the PV system per day and the average power consumed by the loads for the same day (0.58MW). The maximum measured real power will serve as a starting point when the connecting capacity to the overlaying voltage level is restricted.

Table 3: Maximum System Load Under Different PV Penetrations on 07/15/2015

PV Penetration [%]	Maximum Measured Real Power [MW]	Maximum PV Generation [MW]	Average PV Power [MW]	PV Average Power as Share of Total Consumption [%]
0	1.787	0.0	0.0	0.0
10	1.628	0.234	0.073	12.0
25	1.378	0.597	0.188	30.8
50	1.008	1.215	0.382	62.9

Source: SLAC National Accelerator Laboratory

To evaluate the flexibility of the HVAC load, Figure 38 shows both the total power consumed by the HVAC systems regulating the temperature during July 15-16, 2015, and the total power demanded by the HVAC systems. This last curve corresponds to the sum of the power demanded by all HVAC units. It corresponds to the power demanded by all units when they operated at the same time and represents the market's maximum available flexible power. This defines the control spectrum of the market: When all the HVAC units participate in the

market, the flexible load from HVAC systems can ideally change from 0 MW to that limit to regulate the temperature in the houses. If the participation of the HVAC loads in the market decreases, the magnitude of this limit curve would be reduced accordingly.

Figure 38: Power Limits for Flexible HVAC Loads

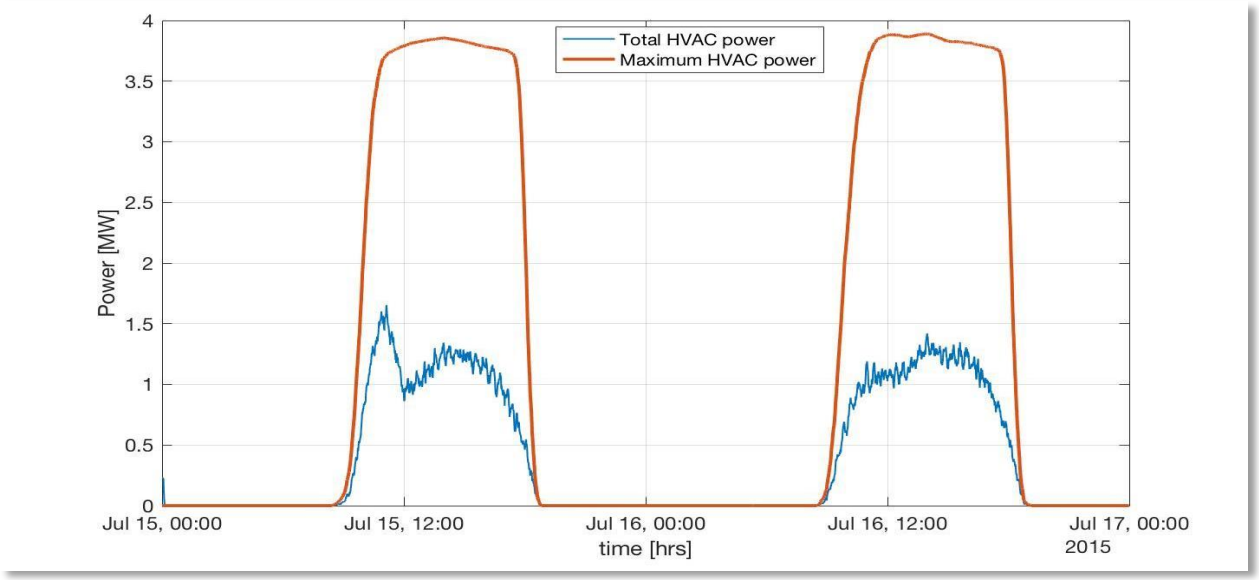

Source: SLAC National Accelerator Laboratory

Local Energy Markets Without Capacity Constraints

In today's electricity markets, customers typically subscribe to a fixed retail tariff and have no incentives to reduce their consumption to reflect actual scarcities in the system, particularly supply scarcity and binding grid constraints. This drives a wedge between consumption at the fixed retail level on one side and the time-dependent state of the system on the other side.

The value of transactive systems and local energy markets covers multiple issues. First, they are able to pass on supply scarcity signals by communicating the wholesale market price and the real cost of energy procurement to consumer devices. As economists have pointed out (Caramanis, Schweppe and Tabors 1983) (Borenstein 2005), letting consumers face real-time prices leads to more efficient dispatch and lowers procurement prices for the retailer, which in turn can then be passed on to consumers. Second, a transactive system can reflect grid constraints and manage additional constraints that occur between the wholesale market and the local distribution system.

In this chapter, how a transactive system reacts in the absence of capacity constraints is analyzed. In this case, real-time wholesale market prices enter the local energy market as unconstrained supply bids, and the price of the local energy market always equals the wholesale market price. It is assumed that not all electric load is subject to the resulting price. Those that do not actively participate face the fixed retail rate. In that case, customers operate their HVAC units using internal temperature regulation controls. The market operator places a demand bid on behalf of those customers by estimating the amount of unresponsive load. Those that do actively participate, however, face the real-time local energy market price.

The results show that if there are no capacity constraints on the system, market-based control of the distribution network does not introduce significant change in aggregate load. Figure 39

depicts this case for different participation of HVAC systems in the market. In this case, no renewable resources or storage units are considered. 0-percent HVAC represents the base case without markets. In this analysis, the temperature at each house is set at approximately the same temperature in all cases.

Figure 39: Measured Real Power Under Different HVAC Participation Shares

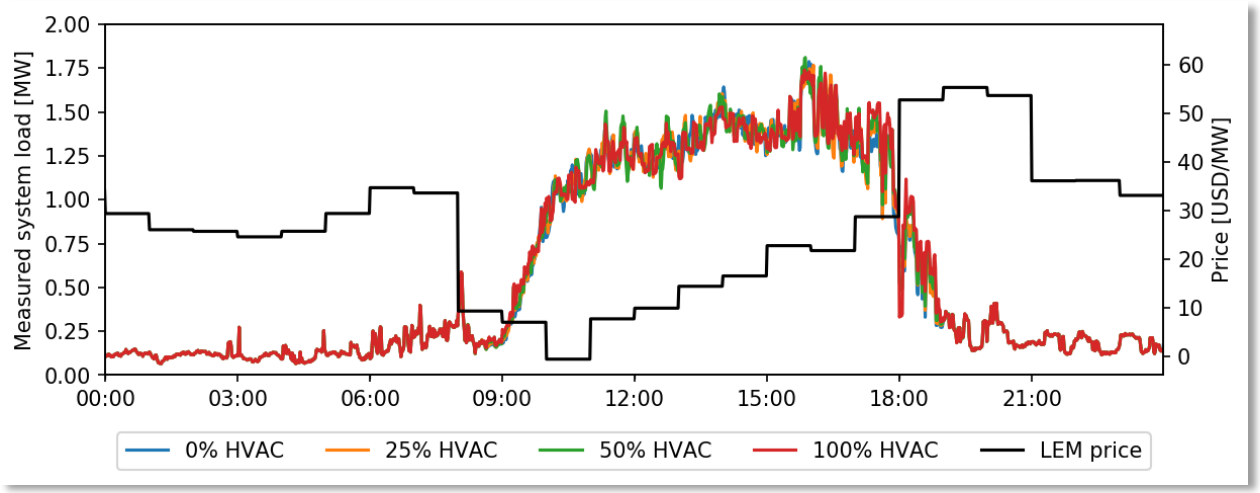

Measured power for different participation of HVAC system, no PVs or batteries. The reference temperature is set approximately the same at each house for all cases. 0% HVAC represents the NO MARKET case.

Source: SLAC National Accelerator Laboratory

The wholesale market price is particularly low during the daytime, between 8 a.m. and 6 p.m., and particularly high between 6 p.m. and 9 p.m. Still, no significant load shifting was observed toward the lower price periods. There are two reasons. First, the wholesale market price goes down when demand for cooling is highest. Second, the myopic control of the HVAC systems, as implemented and described in Chapter 5, limits the ability of the HH to foresee price changes and participate in active load shifting (which one would typically expect around the times of significant price changes around 8 a.m., 6 p.m., and 9 p.m.). However, even with a forward-looking control, the ability to pre-cool or pre-heat is restricted by the thermal characteristics of the houses and most probably restricted to one hour.

It is important to observe, in Figure 39, that the aggregated load responds with a power spike when the local energy market (LEM) price, or equilibrium price of the local transactive system, suddenly jumps at 8 a.m. and 6 p.m. This response is typical in a ”network of oscillators,” like the HVAC systems, where external perturbations (LEM prices) force synchronization of the oscillators and turn on a large group at the same time (8 a.m.), when the price is decreased, and off at 6 p.m., when the price increases. These transients decay over time because the synchronization among units cannot be preserved given the diversity of HVAC systems and thermal responses of the houses. This results in different oscillation frequencies for the HVAC units and synchronous off-and-on of equipment, diversified over time. Figure 40 provides insight into how aggregate load evolves when the price jumps at 8 a.m. (positive) and 6 p.m. (negative). This transient can be very harmful for the power system if the base power of the aggregated system is large at the time the price suddenly changes.

Figure 40: Measured Real Power at LEM Price Jumps

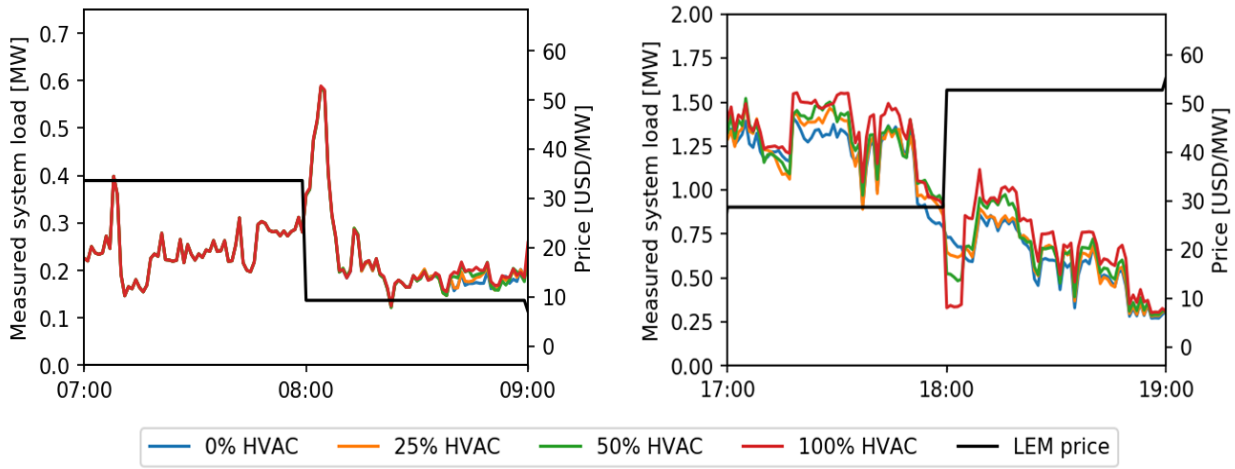

Source: SLAC National Accelerator Laboratory

Another important characteristic is that the behavior of loads does not anticipate or predict the behavior of the LEM price. The consequences of myopic bidding behavior and price controls have been observed in other projects, such as in the Olympic Peninsula, which followed a similar implementation. At 9 a.m., the total load behavior does not change because the HVAC controller builds willingness to pay solely on past prices, nor does it postpone cooling to a later time when prices are lower. Instead, the jump in real power shortly after 9 a.m. seems to be entirely driven by outside temperatures and does not depend on the share of HVAC participation in the market or by responding to real-time prices. Similarly, HVAC agents do not anticipate higher prices at 6 p.m. and build their willingness to pay on past (low) prices. When the price jumps at 6 p.m., their price bids tend to be lower than the costs of wholesale market supply so they will not be cleared and be part of the market allocation, leading to a sudden load decrease of nearly 0.5 MW. From a control point of view, this price increase temporarily synchronizes the HVAC units, but after this transient ends consumers keep demanding power from the grid at high prices and are getting dispatched again. The slow reduction in the power demanded after 6 p.m. is related to decreasing external temperatures.

The existence of the transient as well as the sudden drop in load show the importance of both forward markets for consumers and more sophisticated controls. Also, anticipating price behavior and active load scheduling instead of a myopic control, the costs of power supply will decrease. While the example of San Diego shown in Figure 39 shows a relatively good correlation between prices and power consumption, using the information of forward markets could decrease costs even more. Specifically, the HVAC systems could start operating one hour earlier at 8 a.m. when the price is decreased, and the HVACs could turn off at 6 p.m. when the price is increased. Forward markets have not been implemented in this project.

Table 4 summarizes the maximum load for different participation rates of HVAC systems, as well as for penetrations of PV systems and batteries. In this analysis, customers not participating with an HVAC system in the market control the temperature of the house using the internal regulator of the HVAC unit. These values will be used as the basis for the determination of reasonable capacity constraints on the system. It is possible that flexible loads that respond to real-time prices can increase maximum load in a feeder if lower-than-retail-tariff wholesale market prices are passed on.

Table 4: Maximum System Load Under Different Shares of House HVACs

Flexible Houses [%]	Maximum Measured Real Power, in MW:				
	No PV or Batteries	10% PV	25% PV	10% PV and Batteries	25% PV and Batteries
0	1.787	1.628	1.378	1.378	1.397
25	1.767	1.651	1.383	1.609	1.440
50	1.810	1.647	1.385	1.615	1.440
100	1.740	1.614	1.415	1.609	1.485

Measured total power for different participation of HVAC system for no PVs or batteries, 10-25% PV and 10-25% PV and batteries

Source: SLAC National Accelerator Laboratory

Summarizing, aggregate load behavior is not changed systematically by a local energy market if the supply has no constraints, appliance controls are myopic, and forward markets do not exist. Other aspects, the distribution between the utility, the market operator, and customers, are discussed in the section "Welfare Considerations."

Local Energy Markets with Capacity Constraints

In this section, the second aspect of the value proposition of a local market is analyzed: whether the market is capable of responding to local grid constraints. It is introduced by a binding capacity constraint at the bus connecting the distribution grid to the transmission system. Such a problem structure has already been implemented in real-world pilots such as the Olympic and Columbus demonstrations, and is highly relevant to constrained distribution systems.

The results show that a local energy market is capable of incorporating a local grid constraint. A market with 50 percent of houses with HVAC systems is analyzed first. It is known from the previous analysis that the system experiences a peak of 1.787 MW without a market and 1.81 MW where 50 percent of participating houses have HVAC in the absence of a capacity constraint. Therefore, the market for capacity restriction of 1.6 MW was tested.

The total power of the system measured at the connection point between the distribution network and the transmission grid is depicted in **Error! Reference source not found.** for July 16, 2015. The market successfully controls the aggregate load below the level of 1.6 MW, lowering peak demand by approximately 0.2 MW. In this plot, it was possible to observe the price in the local market increase when the total power reached the constraint value.

Total power is controlled by regulating operation of the flexible load—by participating HVAC systems in the market, as one example. Three constellations on the local energy market are possible. First, if the demand of HVAC systems willing to pay the real-time wholesale market price is less than the constraint of 1.6 MW, the downward-sloping demand curve intersects the supply curve, as illustrated in Figure 42a. In that case, the constraint is not binding. Second, if demand cost at the wholesale market is higher, the price rises to the intersection of the demand and supply curves, as illustrated by Figure 42b. Based on that price, dispatch of the HVAC load can be allocated among all the participants by a single price signal. However, if, third, the demand is so high that the curves intersect at the price cap or saturation limit, as

shown in Figure 42c, a price signal is insufficient to control aggregate load, though appliances must be specifically addressed regarding their dispatch. If the remaining load at the saturation level is unresponsive load only, the market can no longer provide effective guidance and other means like lateral switching, disconnecting the feeder, or curtailment must be exercised. It is important to note that the market defines a lower price if the power constraint is not reached by the distribution system (Figure 42a) than when a power constraint is binding (Figure 42b), or when the demand curve is saturated (Figure 42c). That price corresponds to the intersection point of the horizontal part of the supply curve and the demand, shown in Figure 42.

Figure 41: Measured Real Power and Prices

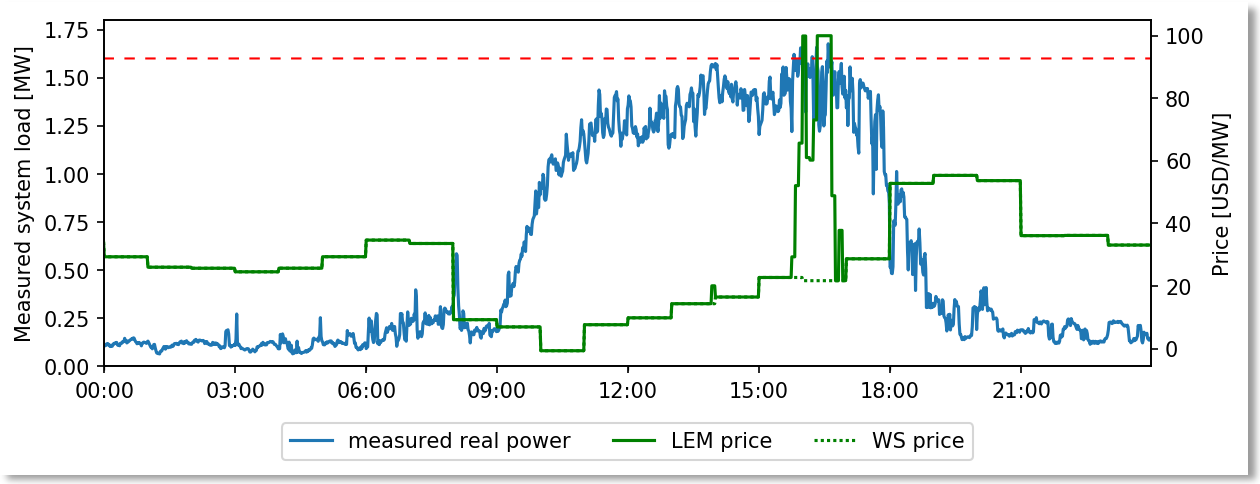

Source: SLAC National Accelerator Laboratory

Figure 42: Market Clearing for Different Operating Conditions

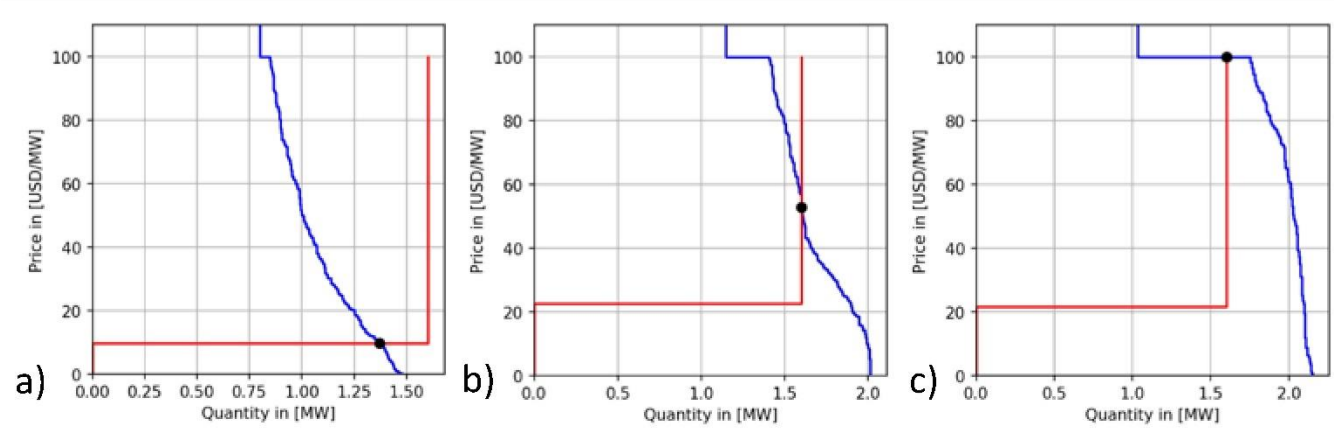

Market Clearing at 12pm (a: unconstrained), at 3.50pm (b: constrained), and 4.30pm (c: saturated)

Source: SLAC National Accelerator Laboratory

Comparing the price for all cases, it is possible to conclude that in the case where the system operates with a power lower than the limiting constraint, the market price is defined by the wholesale price. When the power reaches the constraint, the price is totally defined by the market, which reflects the scarcity of the grid constraint. In this calculation, conducted by the market, the estimate of major magnitudes to define demand and cost curves is important. From figures 41 a-c, it is possible to observe that the estimation of non-flexible load power

directly implicates the market price. Any error in this estimate will affect the price at the clearing time of the market.

Another important conclusion from Figure 42b is that the increased price in the market will dispatch the power among the HVAC units participating in the market and will raise the temperature inside participating houses. It is important to observe that from the thermo-dynamic point of view, if there is a limit in the total power to cool down the houses, it is impossible to keep the level of comfort inside the houses when there is no shortage of power. The market cannot violate that condition. Under such conditions the market dispatches the power among the HVAC units to bids that participants send to the market. If the constrained power is within the power limits of the flexible HVAC units, operators of the distributed system do not need to curtail the HVAC units at particular houses, and the market is in charge of distributing the power among those units, based on the biddings.

Further analysis and quantification of these issues are described here. The impact of the estimation of the un-responsive loads on the market price is analyzed based on the data shown in Figure 41 between 3 p.m. and 6 p.m. In this period, the total power consumed is close to the capacity constraint limit. Figure 43 depicts the unresponsive load as estimated for the market interval, the actual behavior of the inflexible load, the measured real power, and the market clearing volume for a capacity constraint of 1.6 MW. The difference in power between the actual behavior of the inflexible load and the measured real power is equal to the power demanded by the HVAC units. The difference between the actual inflexible load and the estimated unresponsive load is due to the signal processing and filtering of the measured load. The inflexible signal is sampled at a higher rate (30-60 seconds), and after processing is presented to the market at a rate of five minutes. Similarly, the clearing quantity is calculated by the market every five minutes, while the measured real power is acquired at a higher rate. Due to this constraint, the market never clears above the transformer capacity.

Figure 43: Measured Total and Inflexible Real Power Versus Market-Clearing and Unresponsive-Load Forecasts

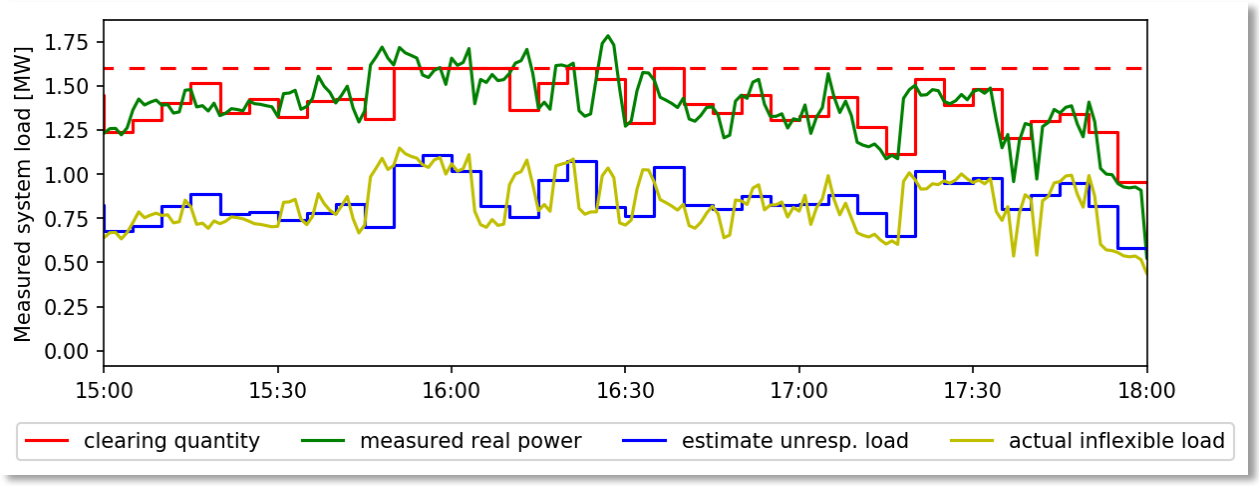

Source: SLAC National Accelerator Laboratory

It is possible to observe from Figure 43 that the market fails to control the aggregate load efficiently if the unresponsive load is incorrectly forecasted. The error between the estimated unresponsive load and the actual inflexible load translates directly into an error between the

clearing quantity and the measured real power. As an example, Figure 44 shows an error in an unresponsive load forecast.

Figure 44: Forecasting Error for Unresponsive Loads

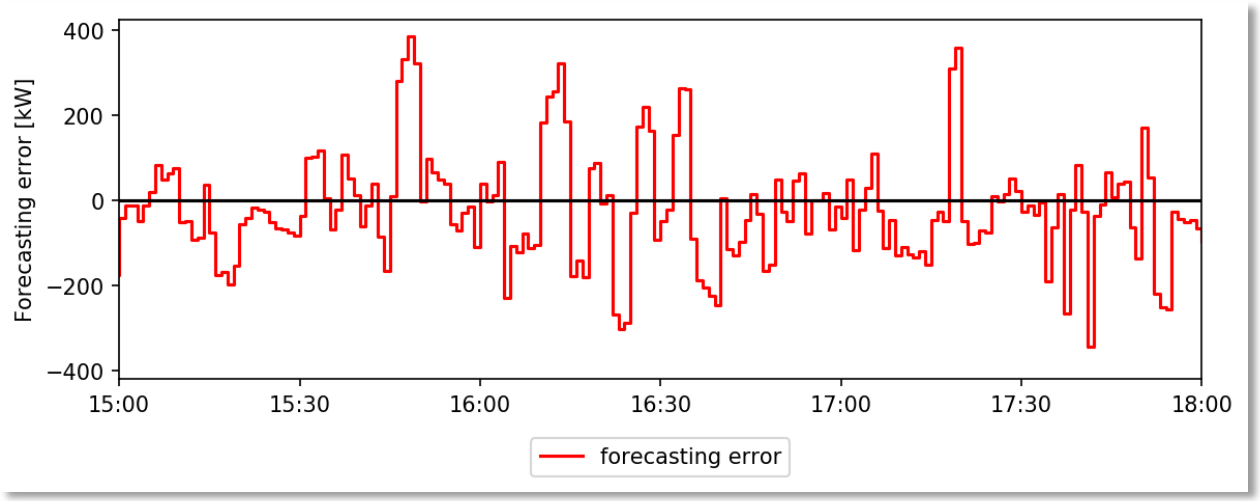

Source: SLAC National Accelerator Laboratory

In local energy markets, estimating unresponsive load is more challenging than in larger electric systems for two reasons; the number of participating appliances is smaller so the variance of the forecast is higher. In the simulation based on the real data supplied, it is possible to see base-load jumps (non-HVAC load) of up to 243.1 percent, as depicted in Figure 45 and indicated by the black dashed line.

Figure 45: Baseload (Non-Heating, Ventilation, and Air Conditioning Load)

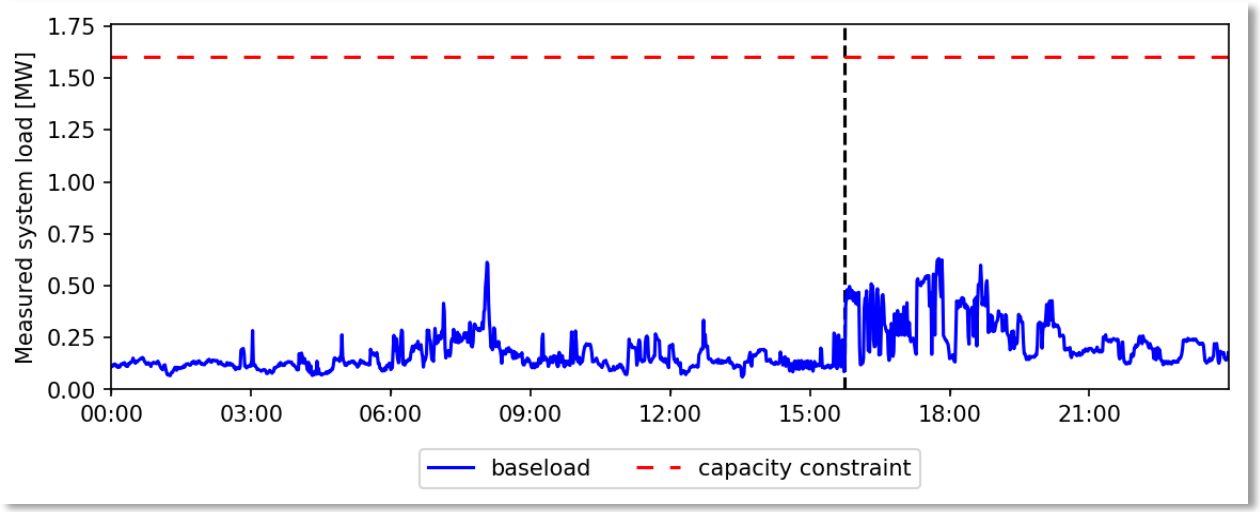

Source: SLAC National Accelerator Laboratory

To study the saturation of the controller and the action of the constraint limit in the system, the statistical analysis of tie-line overloading was conducted. It was important to evaluate the impact of estimates in the bidding quantities and errors in the forecasts of unresponsive loads. This operation area corresponds to the system when it functions between the regions depicted in Figure 42b and Figure 42c. In this analysis, the capacity limit was changed in multiple steps from 1.6 MW to 1.8 MW in the market controller, where the real power on the critical point of

the distribution system connection is measured. The load duration curve was generated by recording the maximum power and the time duration at that level. In these plots, the power level at the critical point in the system (y-axis) is represented as a function of the time elapsed (x-axis) at that given power level, where the power level is plotted in descending order.

Figure 46 shows the resulting load duration curve for a system with 50 percent HVAC participation. The power measurement is sampled at one minute and the time duration is expressed as a percentage of a 24-hour period. The capacity constraints are marked by the red dash lines. For tighter constraints, the respective load curve becomes flatter. However, exceeding the capacity constraint cannot be fully avoided. A capacity constraint of 1.8 MW lasts for 1 minute, for 1.7 MW 8 minutes, and for 1.6 MW 24 minutes. The peak load decreases slightly for a capacity restriction of 1.8 to 1.802 MW and for 1.7 MW to 1.728 MW, but increases between 1.6 MW to 1.664 MW. For the latter, in times when the capacity constraint is exceeded, it does so by a mean of 63.8 kW.

Figure 46: Load Duration Curve of 50 Percent Heating, Ventilation, and Air Conditioning Participation Shares Under Capacity Restrictions of 1.8, 1.7, and 1.6 Megawatts

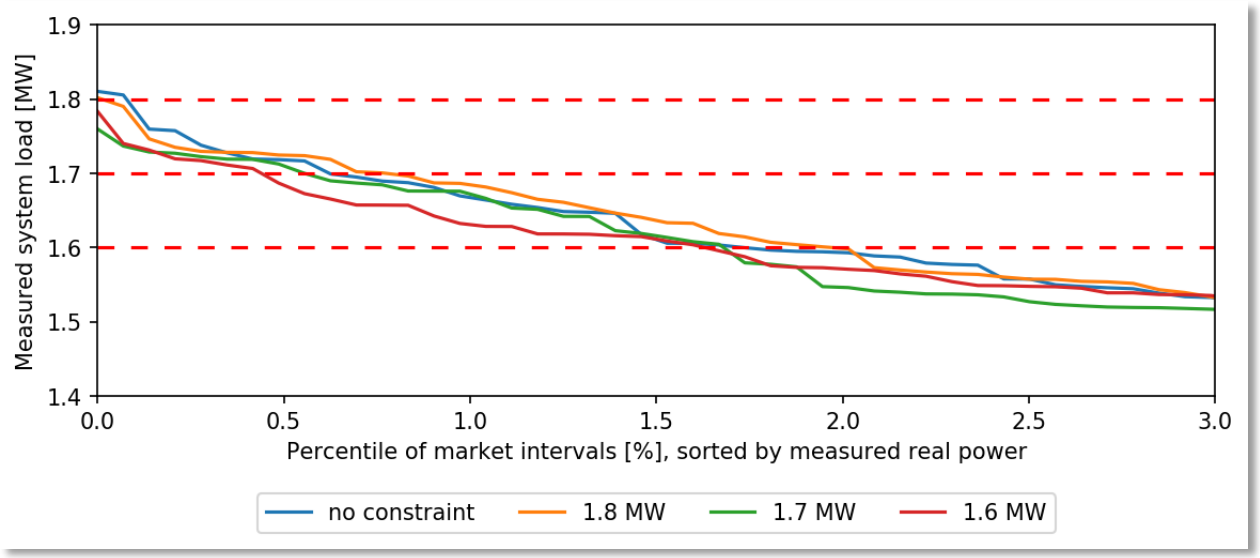

Source: SLAC National Accelerator Laboratory

Different errors and perturbations in the market affect overloading in the system when the controller forces the capacity limit. Errors and perturbations in the market are associated with the power level (quantity) used by the agents in the bidding, and in the forecasting of unresponsive loads.

The market can be imprecise because the actual consumption of loads does not necessarily equal the bidding quantity placed. For instance, it can be found that the actual consumption of HVAC systems deviates from the rated power of the devices between -36.5 percent and 91.5 percent. The reason is that the actual power consumption of the devices does not necessarily comply with the rated power, but rather depends on the state of the appliance and its environment (for example, ambient temperature or downtime). In this work, the bidding quantity (power to be consumed) is set to be equal to the last measurement of the power consumed by a device. This decreases the maximum deviation between the bidding quantity

and the actual consumption to below 4 percent. Further minor deviations from the target capacity level can be attributed to small deviations of appliances from their bid as well as to grid loss. Figure 47 depicts the impact of this error in the load duration curve of the system, assuming a perfect unresponsive load forecast.

Figure 47: Reproduction of Figure 46, With Perfect Unresponsive Load Forecast

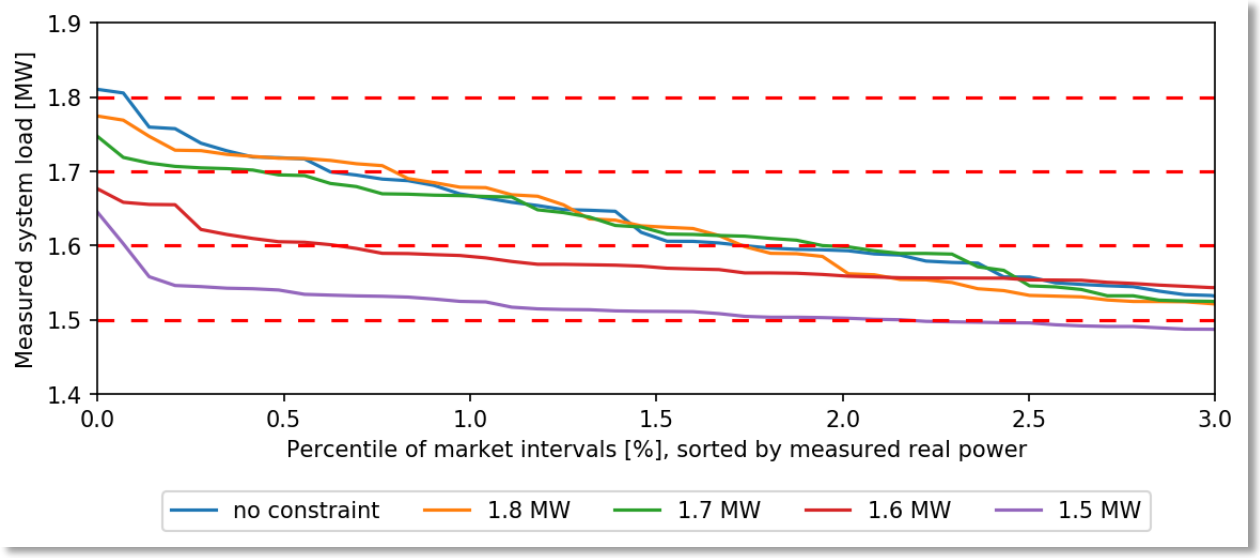

Source: SLAC National Accelerator Laboratory

Results for 50 percent HVAC load participation are summarized and compared in Table 5 for different constraint levels. While the market is able to control the load and sustainably lower it, this comes at a price. The mean price increases from 26.39 USD/MWh to 33.29 USD/MWh for the chosen scenarios. The mean maximum temperature per house only slightly increases from 71.62 °F to 71.92 °F. Houses' participation in the transactive system, rather than subscription to the retail tariff, depends upon the ability to shift load into low-price periods.

Table 5: Results for a Market with 50-Percent Flexible Heating, Ventilation, and Air Conditioning and Different Constraints

Constraint	Minutes When Constraint Is Exceeded	Average Positive Deviation from Constraint [kW]	Mean Price [USD/MWh]	Mean Maximum Temperature in Houses [°F]
None	0	0.0	26.39	71.62
1.8 MW	0	0.0	26.39	71.62
1.7 MW	7	13.3	26.54	71.62
1.6 MW	10	30.0	28.64	71.67
1.5 MW	31	28.4	33.29	71.92

Source: SLAC National Accelerator Laboratory

Results depicted in Figure 48 analyze the case of the system with 100 percent flexibility and different options to define the constraint limit. In this analysis, a constraint is defined at 1.6MW and the effect of a plus-5 percent safety margin. The boundary performance for a case

without errors in the power estimation (perfect baseload forecast) is included in the plot. Safety margins decrease the risk of overshooting but might lead to inefficiencies when the market is frequently cleared below the actual physical constraint, both decreasing customer comfort and inflating prices. On the other hand, system operation measures such as curtailing provide the means to correct dispatch. For that purpose, the transactive systems provide guidance since the bids of the last market clearing can be used to determine customers with the least willingness to pay and a price based on which they can be compensated for in the case of curtailment.

As illustrated by the green line, the safety factor of 5 percent decreases the number of time periods exceeding the capacity constraint from about 1.6 percent (24 minutes) to 1.2 percent (20 minutes). However, above that threshold mean excess power increases from 63.8 kW to 105.8 kW. For the perfect unresponsive load forecast scenario (red line), the constraint is only violated in 7 minutes, by a mean of 4.7 kW**.** This shows that a high forecasting quality for unresponsive load is crucial for effectively controlling aggregate load around or below the capacity constraint.

Figure 48: Load Curves for Different Safety Factors

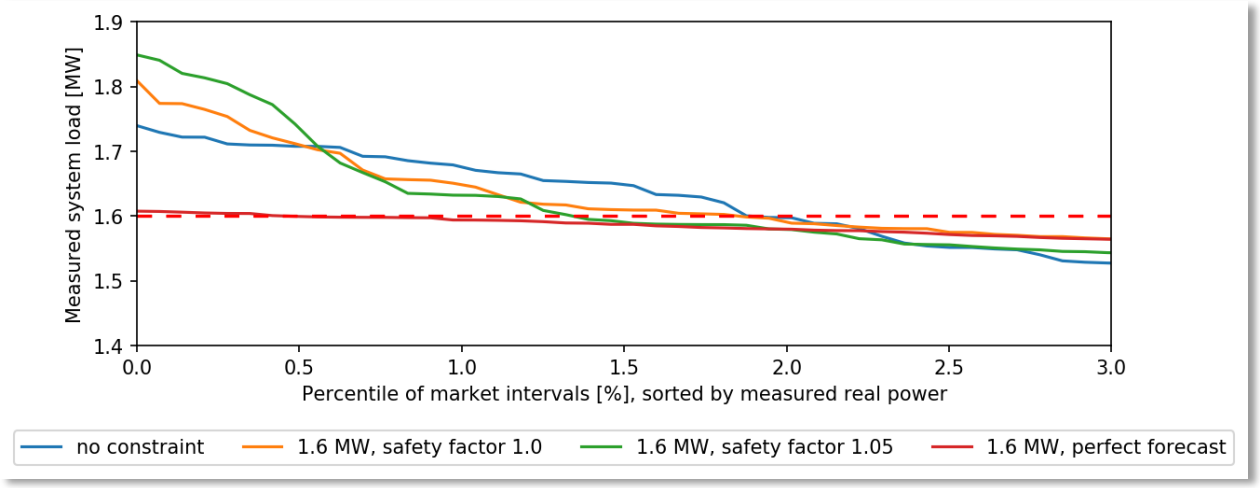

Source: SLAC National Accelerator Laboratory

Local Energy Markets with Batteries

The penetration of PV units in the system decreases the net load at the connection point of the distribution grid to the transmission network. The effect of local generation was quantified and is depicted in Table 1, Table 2, and Figure 37 In the following analysis, different penetration levels of batteries are also included in the electric system. As it was presented in previous chapters, battery units were installed in all houses with PV equipment.

As explained in chapters 4 and 5, the control strategy for the batteries follows a simple scheduling policy where, at the beginning of each day, the HH determines the highest price at which a battery is willing to purchase electricity, and the lowest price at which to sell it. Figure 49 illustrates the resulting charging behavior for a sample battery, with the red horizontal lines representing the price thresholds for selling and buying. Accordingly, the battery does not exhibit any activity in between those thresholds, in this example for prices between approximately 10 and 37 USD/MW. The battery loads during times of low prices, in this case

between 8 a.m. and 10 a.m., and discharges after 6 p.m., when wholesale market prices are highest.

Figure 49: Wholesale and Local Real-Time Prices and Buy/Sell Price Thresholds of a Sample Battery

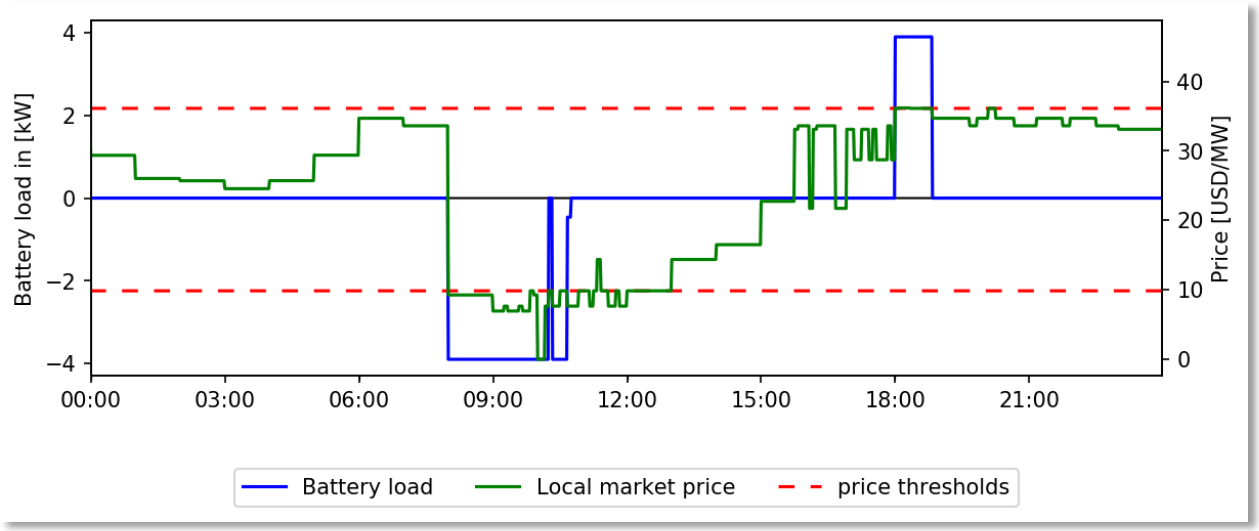

Source: SLAC National Accelerator Laboratory

To test the ability of the market to drive down peak demand, a system configuration comprised of 25 percent PV and 25 percent batteries (with a maximum capacity limit of 1.2 MW) was chosen. As depicted in Figure 49, the load valley produced by the PV generation is filled by the battery load, increasing the measured real power but keeping it mostly below the defined capacity limit. Figure 49 illustrates the power consumed and generated by the batteries.

If the feeder capacity limit is at risk for violation, local energy market prices rise above the wholesale market. For battery penetration of 25 percent, this is the case between approximately 10 a.m. and 11 a.m., and between 4 p.m. and 6 p.m., as shown in Figure 50. Around those times, the measured load also slightly violates the capacity constraint, which can be attributed to another uncertainty—PV generation. On the other hand, the local energy price sinks below the wholesale market level between 6 p.m. and 9 p.m., decreasing procurement costs for local customers.

Figure 50: Measured Real Power for 25 Percent PV Penetration

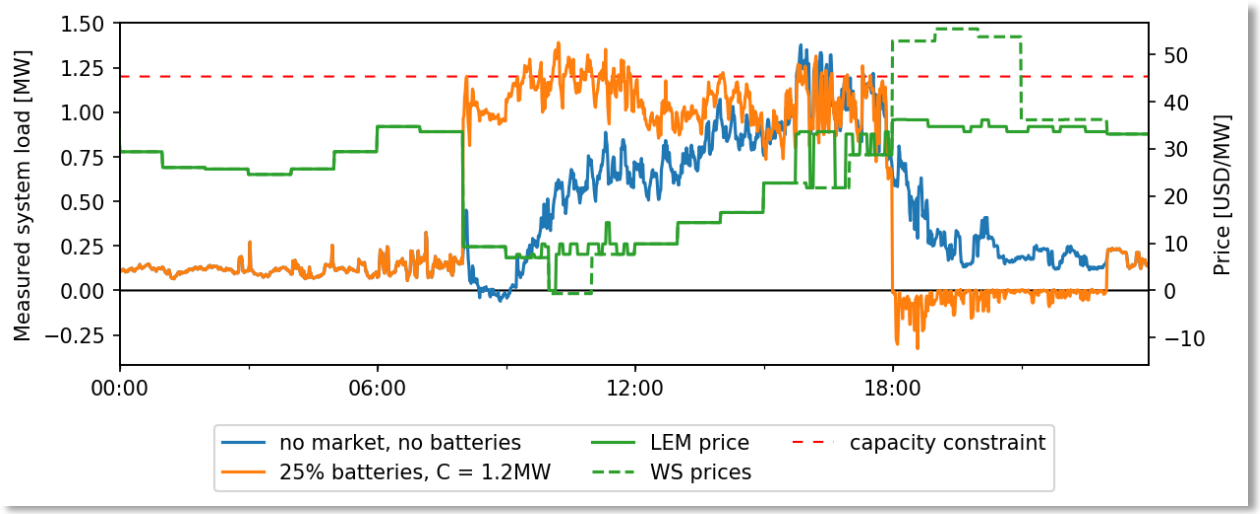

System with 25% PV, 25% batteries, and a capacity constraint, compared to a system with 25% PV only and no constraint. The prices refer to the system with batteries and the capacity constraint.

Source: SLAC National Accelerator Laboratory

Figure 51 depicts the dispatch of the batteries. It is shown that the batteries are charging in the morning hours and discharging in the early evening, supporting the local system.

Figure 51: Measured Battery Power for 25 Percent PV Penetration

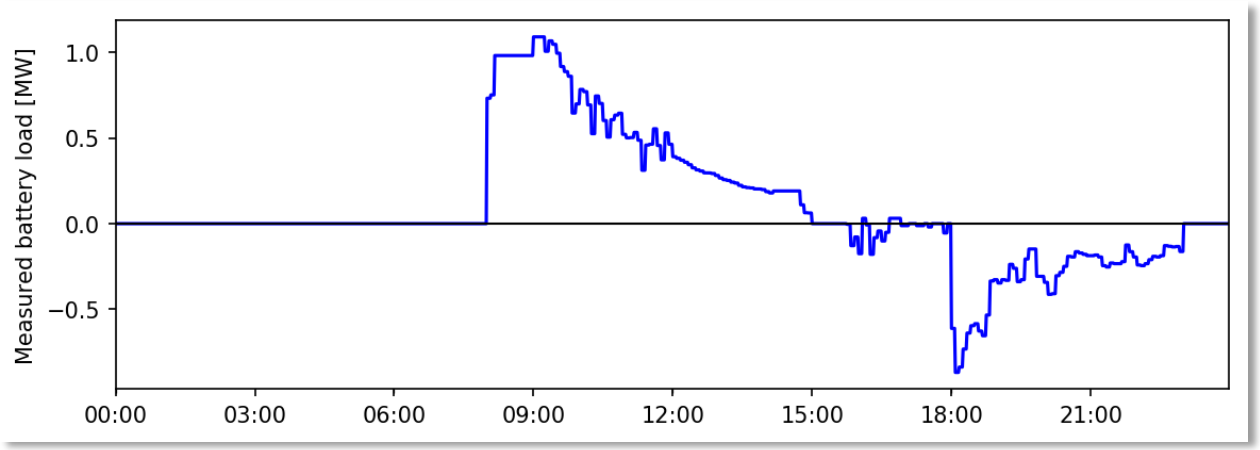

Source: SLAC National Accelerator Laboratory

Local Energy Markets with Electric Vehicles

To include the impact of charging electric vehicle batteries in the system, two types of charging units are considered: large charging stations, and commercial or domestic units. Data from a large U.S. charging operation was used to estimate arrivals, departures, and charging rates. To analyze the impact of fast charging, four types of cars are listed in Table 6.

Table 6: EV Models Included in the Simulation

Type	Storage Volume [kWh]	Level 3 Charging Rate [kW]
BMW i3	22	24
Nissan Leaf	32	50
Chevrolet Bolt	60	50
Tesla S 85	90	120

Source: SLAC National Accelerator Laboratory

Figure 52 and Figure 53 depict the overall electric vehicle charging load under 25 percent penetration for both types of charging units during a week in July.

Figure 52: Measured EV Load for Fast Charging

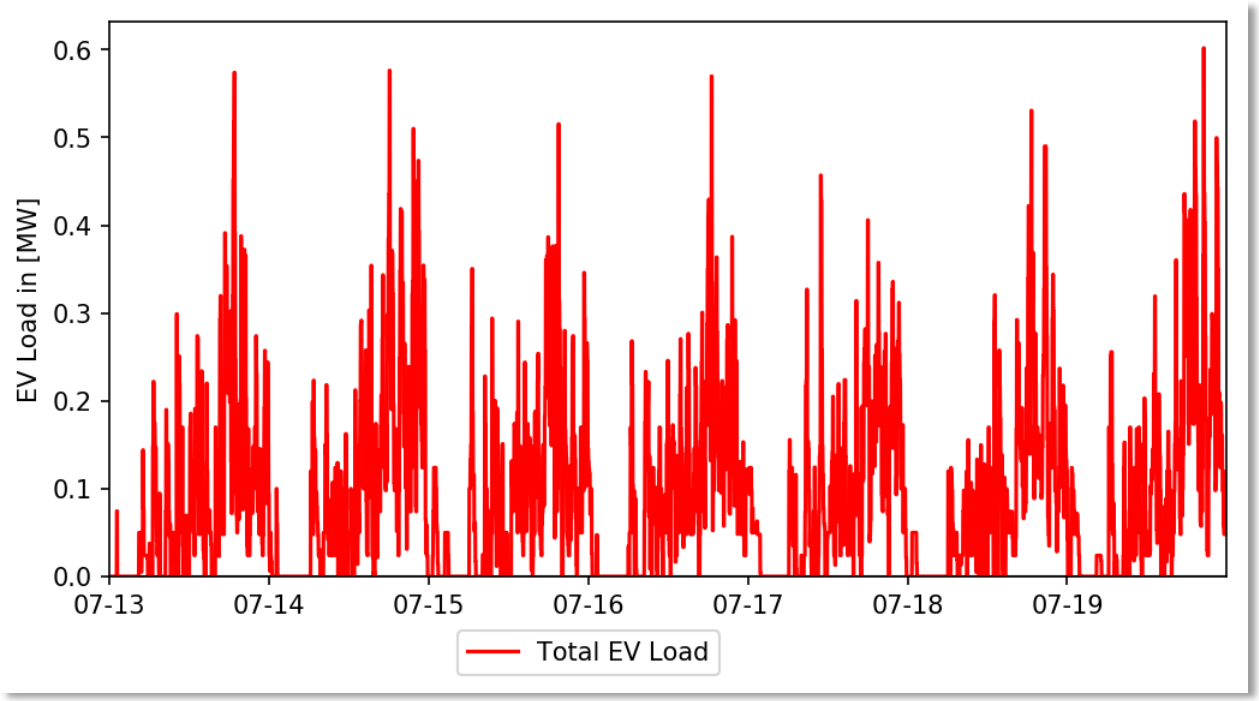

EV load for 25% penetration with arrivals, departures, and initial state of charge estimated from real data. Charging rates and maximum state of charge assigned randomly from existing fast-charging technical specifications.

Source: SLAC National Accelerator Laboratory

Figure 53: Measured EV Load for Domestic Units

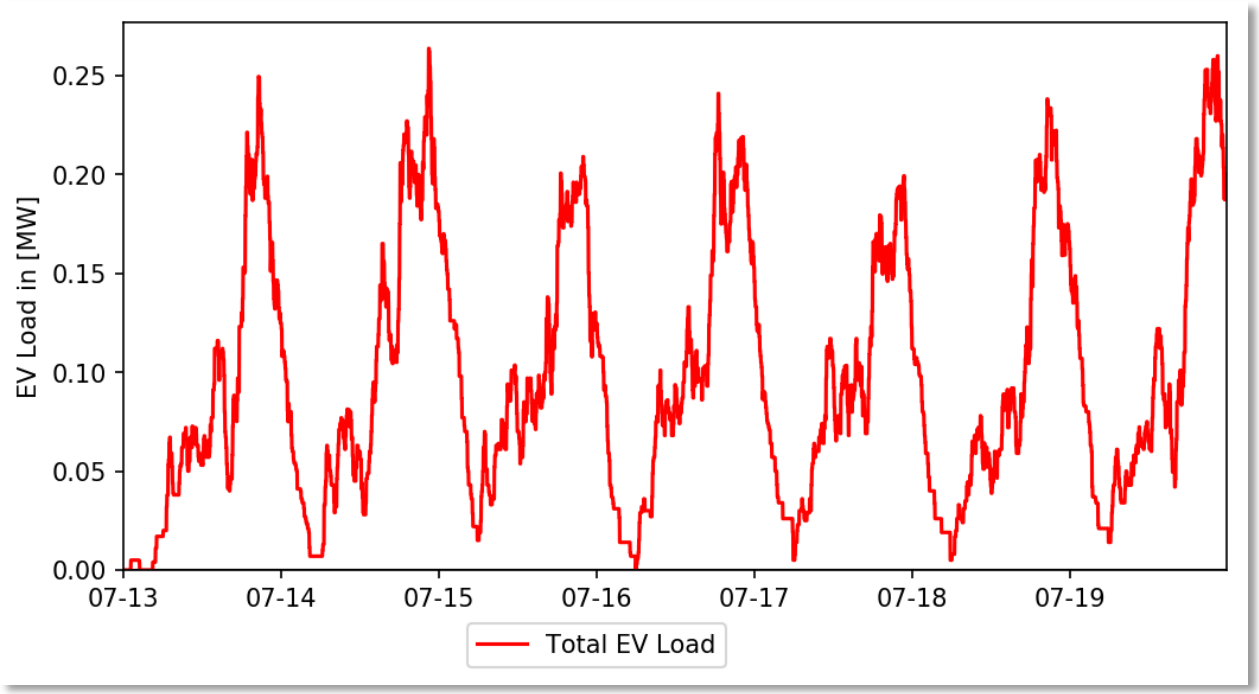

EV load for 25% penetration with arrivals, departures, initial state of charge, and charging rates estimated from real data

Source: SLAC National Accelerator Laboratory

It is possible to observe that drivers usually arrive in the evening, start charging the vehicle battery, and leave during the morning. The EV load therefore experiences peak during the night when wholesale market prices are actually highest but local system load is low. Figure 54 plots the aggregate demand for the distribution system, including electric-vehicle battery charging. The charging load predominantly fills the valleys of low demand during the night.

Figure 54: Measured Real Power With 25 Percent of EV Batteries

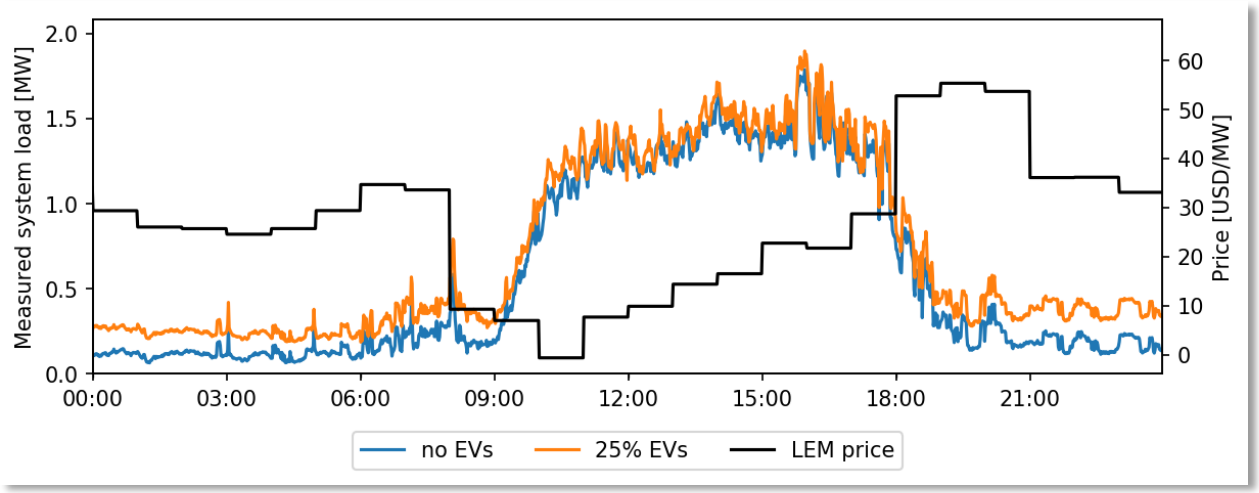

Source: SLAC National Accelerator Laboratory

Full System

The following section presents the performance of the full system with 100 percent flexible HVAC systems, 25 percent PV, 10 percent batteries, and 10 percent electric vehicles, for an effectively unconstrained 3.6 MW and an effectively constrained 1.0 MW system. Percentage numbers refer to the number of households equipped with such devices. Batteries and electric vehicles are only found in households with installed PV systems.

Figure 55 displays the real measured power at the slack bus as well as wholesale market and LEM prices over July 16 for the unconstrained system.

Figure 55: Measured Real Power and Prices

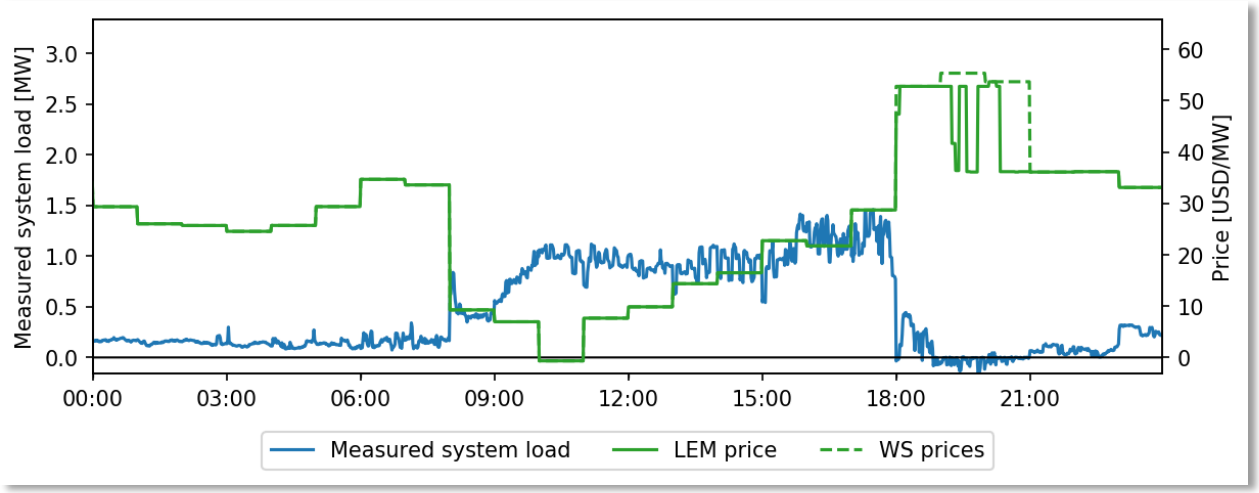

Real measured power at Slack bus for 100% flexible HVAC systems, 25% penetration of PV, 10% batteries, and 10% EVs, for a non-binding capacity constraint of 5.0 MW.

Source: SLAC National Accelerator Laboratory

Measured real power increases in the morning, is relatively flat during the middle of the day due to PV infeed, and reaches its minimum in the early evening hours. Wholesale market prices show inverse behaviors. The constraint of 3.6 MW is not binding in this example, which is why the LEM price is mostly identical to the wholesale market price; however, the LEM price is even lower than the wholesale market price in the early evening when local supply injects power into the distribution network. It is important to note that this is the case because local generation is not allowed to export to the wholesale market at the higher prices, but only at the LEM price.

Figure 56 disaggregates supply for an unconstrained system. The colored area below the curve shows the contribution to overall supply. It shows that the wholesale market supplies 100 percent of the load during the early morning, while during the day about a third of the load can be satisfied by solar generation. In the early evening, however, when wholesale-market prices increase, batteries pursue arbitrage profits and feed into the grid. Since the local load also goes down, the local energy system during some market intervals exports energy to the overlying energy system. As Figure 55 illustrates, this export is only paid for at the local LEM price, which decreases below the wholesale price.

Figure 56: Disaggregated Supply in an Unconstrained System

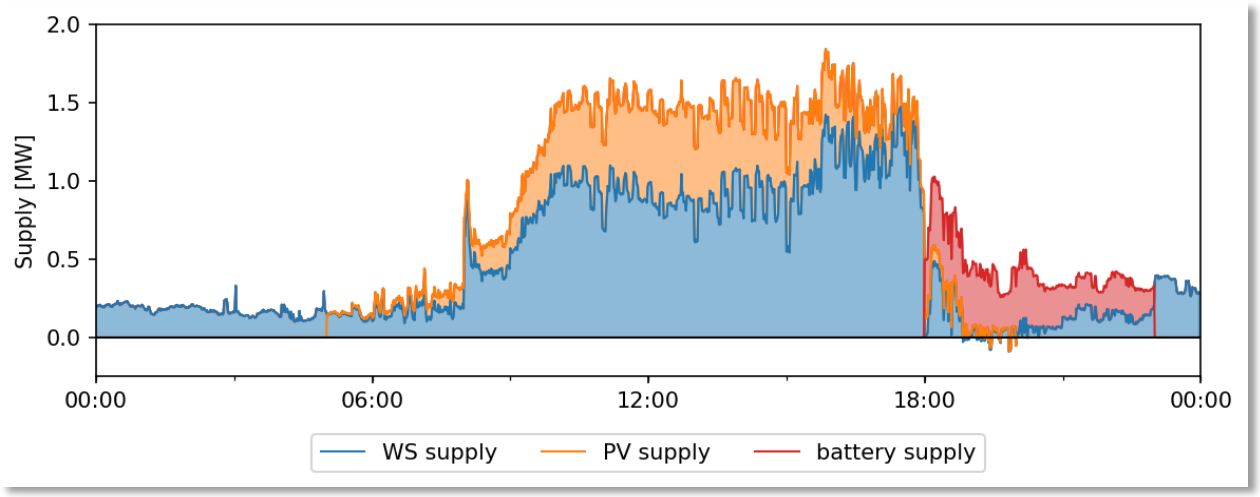

Supply at Slack bus for 100% flexible HVAC systems, 25% penetration of PV, 10% batteries, and 10% EVs on 07/16/2015, for a non-binding capacity constraint of 5.0 MW.

Source: SLAC National Accelerator Laboratory

Similarly, Figure 57 shows disaggregated consumption. Blue indicates the inflexible base load (non-HVAC), which is relatively constant over the day and increases in the late afternoon and early evening. The biggest part of the load is driven by HVAC systems, which explains the high-load demand between 9 a.m. and 6 p.m. Electric vehicle charging is mostly happening in the evening and explains approximately half of that load. Batteries are charged between 8 a.m. and 1 p.m., exceeding the base load during that time, by size.

Figure 57: Disaggregated Load in an Unconstrained System

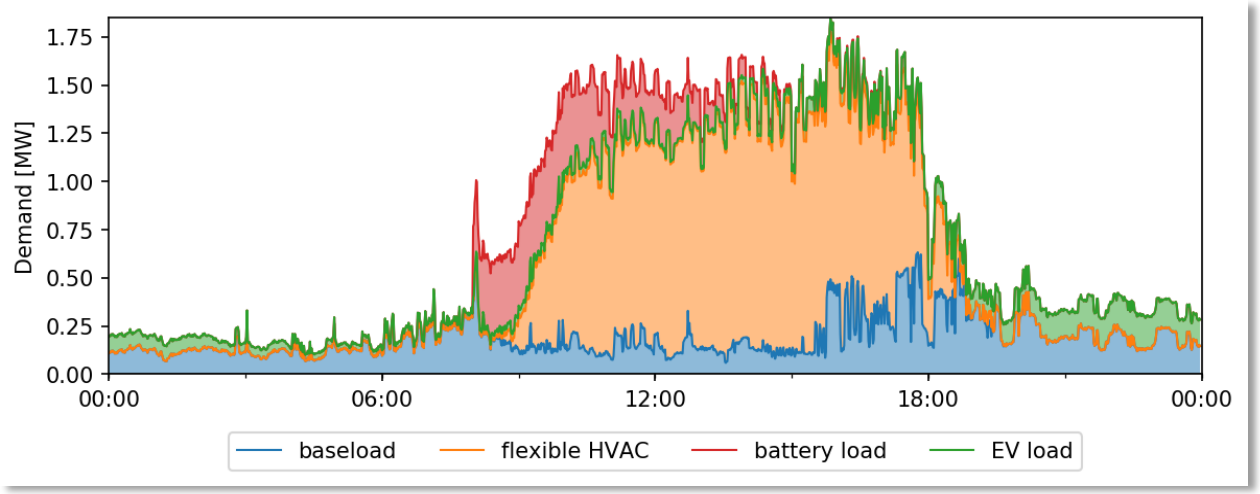

Load at Slack bus for 100% flexible HVAC systems, 25% penetration of PV, 10% batteries, and 10% EVs on 07/16/2015, for a non-binding capacity constraint of 5.0 MW.

Source: SLAC National Accelerator Laboratory

In general, the aggregate load profile mostly follows the physical necessities of the system since the capacity constraint is not binding. Only the batteries that are driven by arbitrage opportunities in the wholesale market significantly change the aggregate load shape by

plateauing aggregate load in the morning hours and even causing reverse flows in the early evening hours.

The impact of a capacity constraint in the system and its behavior is analyzed in Figure 58 and Figure 59. In this case, the topology of the system and the penetration of renewable sources are the same as before though a maximum power of 1MW can be transmitted through the interconnection point between the distribution network and the main grid. Figure 57 shows two main differences in the disaggregated supply. The power at the interconnection point is limited to 1MW as expected and the battery system is providing energy to the distribution network after 3 p.m.

Figure 58: Disaggregated Supply in a Constrained System

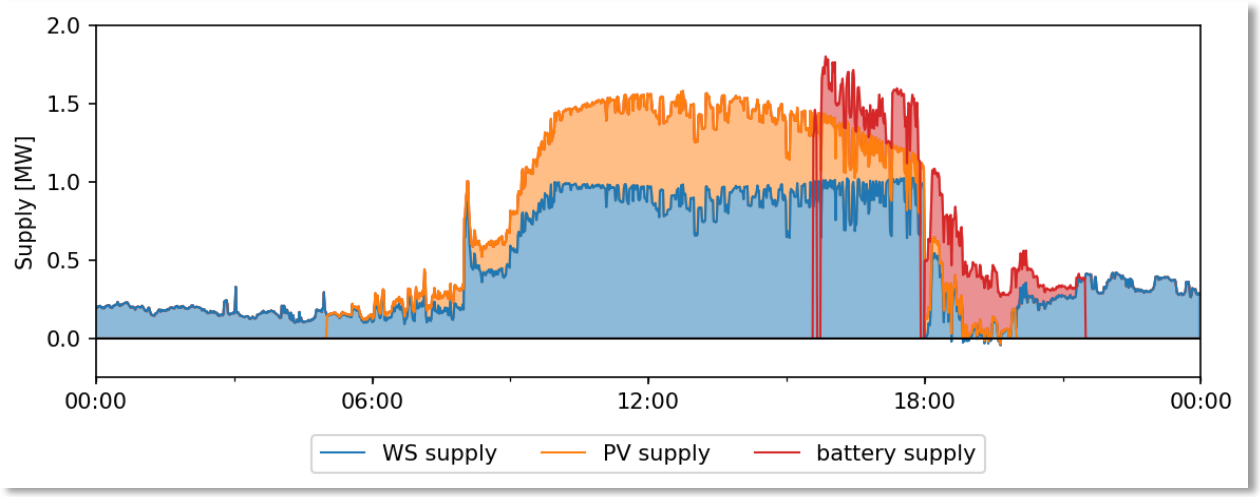

Supply at Slack bus for 100% flexible HVAC systems, 25% penetration of PV, 10% batteries, and 10% EVs on 07/16/2015, for a binding capacity constraint of 1 MW.

Source: SLAC National Accelerator Laboratory

Figure 59 depicts the disaggregated load in the case where the system is operating with limited capacity.

Figure 59: Disaggregated Load in a Constrained System

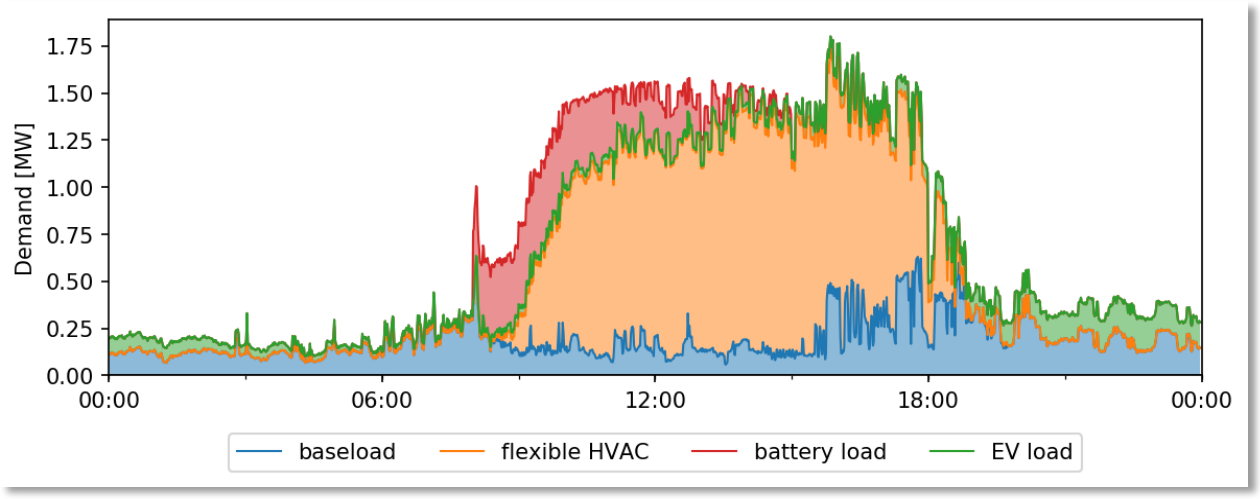

Load at Slack bus for 100% flexible HVAC systems, 25% penetration of PV, 10% batteries, and 10% EVs on 07/16/2015, for a binding capacity constraint of 1 MW.

Source: SLAC National Accelerator Laboratory

Figure 60 and Figure 61 highlight demand and supply changes between unconstrained and constrained systems. For demand, no persistent changes (except for the shift of a couple of single-period load peaks) were observed. This is very different for the supply side, where it shows that the batteries anticipate discharging partially in the late hours of the afternoon when LEM prices are driven up by binding capacity constraints.

Figure 60: Change in Supply

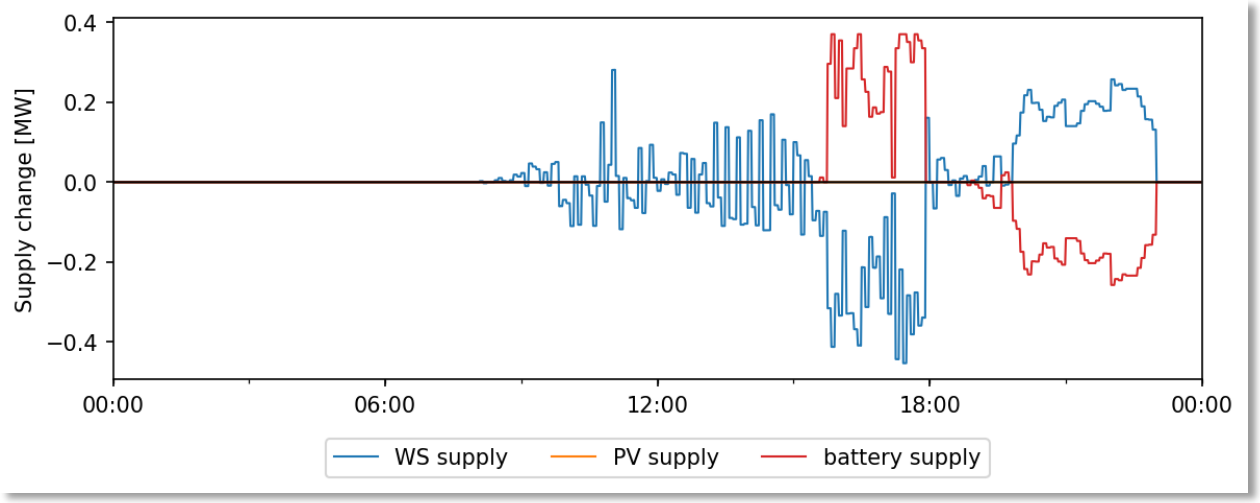

Switching from an unconstrained (5.0 MW) to a constrained system (1.0 MW)

Source: SLAC National Accelerator Laboratory

Figure 61: Change in Load

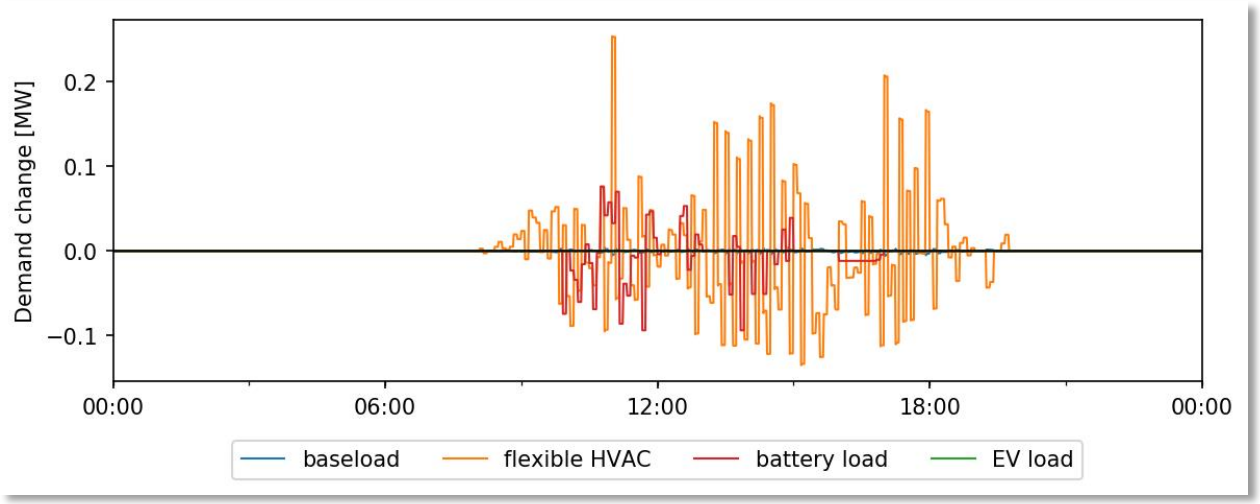

Switching from an unconstrained (5.0 MW) to a constrained system (1.0 MW)

Source: SLAC National Accelerator Laboratory

The comparison of the two examples shows that a transactive system or local energy market can provide effective incentives for flexible loads like EVs and batteries to shift load and provide flexible supply to avoid violations of capacity constraints.

In conclusion, LEM prices are generally higher in constrained systems, which reflect the scarcity of grid capacity through the wholesale market. For the sample day of 07/16/2015, in the unconstrained system, the LEM price averages 25.50 USD/MWh and is 0.89 USD/MWh lower than the wholesale market price. In the constrained system, the LEM price averages 27.74 USD/MWh and is 1.36 USD/MWh higher than the wholesale market price.

Regarding the sources of supply and demand, in both systems residential electricity consumption throughout 7/16/2015 was 14.0 MWh, out of which 9.5 MWh were consumed by the HVAC system and 1.6 MWh by electric vehicles. PV systems were producing 4.5 MWh of solar energy. Batteries shifted 1.5 MWh (after losses).

Welfare Considerations

The system has five main stakeholders: consumers, owners of generation capacity, retailers, grid operators, and market operators. These stakeholders may coincide in one institution, for example under a CCA program. The impact of a transactive system on each stakeholder group will be described and partially quantified. An overall analysis, however, depends on the regulatory framework and on the specifications of fixed retail tariffs, regulated utility income, as well as on general system conditions.

The primary concern here is to analyze the welfare changes to the stakeholders involved. In general, a transactive system should only be implemented if it provides additional value for the overall system. Furthermore, implementation details of a transactive system, such as appliances involved or the choice of the market operator, can have a decisive effect on overall welfare and the distribution between stakeholders. They should therefore be evaluated by a welfare analysis and compared with other design alternatives. Finally, general acceptance by all stakeholders is key to their participation in a transactive system. If, for any participant

class, it is evident that a null participation set is likely, redesign of the system must be considered as well as its suitability for that class of stakeholder.

Consumers

First of all, the consumer bill will change. Instead of the fixed retail tariff, the consumer now faces time-variable energy prices, at least for some appliances. Presumably, higher-than-retail tariff prices in high-price times will be experienced, and lower ones during low-price times. In principle, that could result in an increase or a decrease in utility costs, depending on the ability and the willingness of the customer to respond to price changes and shift load in lower price times.

An analysis was pursued over one week for the months of January, July, and October to analyze customer bill changes. In this scenario, all customers participated in the transactive system with their HVAC systems. The reference price used for bidding was calculated on the average over the last 24 hours. The changes to costs for customers, compared with a scenario with no market and a fixed retail tariff, were calculated. The fixed retail tariff was determined by dividing the procurement cost for energy on the wholesale market by the energy consumed. Alternatively, a risk premium of 2.5 percent was added to the retail tariff to account for the fact that, with a fixed retail tariff, a retailer is entirely exposed to the price risk.

Figure 62: Histogram of Cost Changes

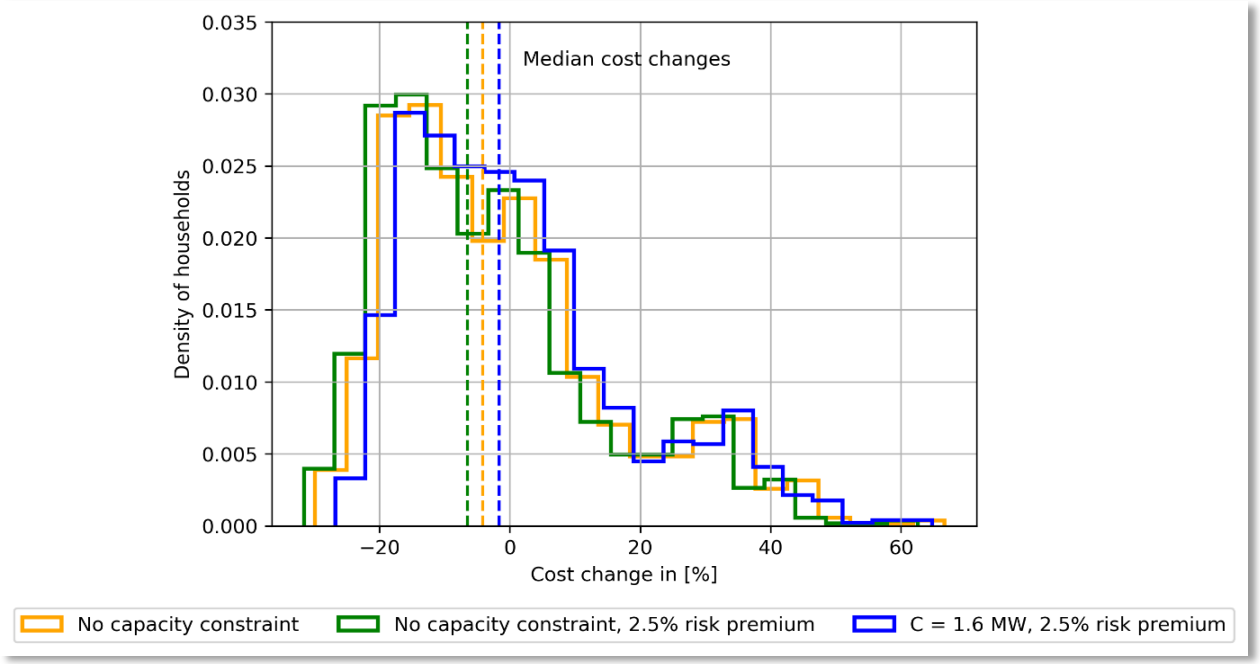

Cost changes for houses with flexible HVAC systems participating in a transactive system compared with a fixed retail tariff

Source: SLAC National Accelerator Laboratory

Figure 63 shows these results. In a scenario without a capacity constraint and without a risk premium, more than 50 percent of customers profited from energy bill savings. The yellow dashed line indicates the median household that profited from a bill decrease of 4.2 percent. If a conservative risk premium of 2.5 percent was included and the retail tariff was respectively 2.5 percent higher, the median household's bill decreased by 6.5 percent (green line). If,

however, capacity was tight and prices generally increased for participants of the transactive system, the median experienced a bill increase of 0.8 percent (not indicated) or savings of 1.6 percent if a risk premium was considered. Saving are expected to be much higher if a control considering forward markets were implemented.

The numbers mentioned, however, do not consider the value that flexible customers bring to the system with regard to capacity. Specifically, their participation in the transactive system gives the system operator indirect control rights which enable active control for peaks, which potentially avoid curtailment and capacity expansion. This should be considered by a reduction of capacity components of the bill, for example setting an additional incentive for consumers to choose the transactive system.

If participation in the transactive system is voluntary, the customer will supposedly only join the transactive system if it is favorable or opt out to stay on the fixed retail tariff if it is not. Apart from the energy-procurement effect, customers may also profit from the positive effects of the capacity usage of the system, in particular the shaving of peaks and potential savings in infrastructure investment or costs of system operations. Those savings could and should be re-distributed back to customers and offer an additional reward. The system can only succeed in the long term if customers in a transactive system save money by switching to the new tariff.

Second, comfort levels could change. Indicators for changes in welfare could be significant deviations of temperature from set points in high-priced times, or non-fully charged EVs at departure. Because the comfort settings are controlled by the consumer, the researchers presume that these changes are voluntary and that discomfort is transient during the learning period. The system can only succeed in the long term if a majority of customers is not significantly adversely impacted by comfort loss and eventually switches to a tariff that impacts their comfort less.

Third, the customer might face a price risk, resulting from volatility in the monthly bill, the risk of non-optimal control of flexible appliances during high-priced times, or the risk of non-well-defined settings in the HH (for example, too low a price sensitivity or too high a maximum price). Here, too, exposure to price volatility is voluntary and any welfare impact is transient during the learning period, if any. The system can only succeed in the long term if the majority of customers with sufficient demand elasticity are willing to accept exposure to high-price volatility.

Owners of Generation Capacity

If the local energy market positively deviates from the wholesale market price, local generation is able to collect an additional profit since it provides additional value to the system by providing local supply when it is scarce. This is true in terms of energy procurement. As flexible demand resources, they additionally provide positive capacity value by shaving net peak loads. If distributed generation owners collected excessive scarcity rents, the resulting loss of revenue may be too great for the grid operator to remain financially viable, and the grid operator would be unlikely to participate in the system.

On the other hand, the transactive system might also assist in quantifying the true value of local generation. In particular, net metering for PV customers has been criticized for over-incentivizing the installation of solar power and transferring wealth from low-income to high income customers. Then, switching to a transactive system where PV owners collect the real-

time price can mean a loss in generation income and net bill increases. This is especially true in systems where the price is negatively correlated with PV production as it is in the example chosen (San Diego). In this example, illustrated in Figure 63, an estimated increase of 45 USD during the year in the bill is experienced by the median PV owner while around 18 USD is saved by the household without PV. This illustrates the subsidization of PV owners by non-owners, which occurs through net-metering.

A transactive system can only succeed if generation owners do not experience overly high increases in their bills compared with the outside option. This might require the termination of net-metering.

Figure 63: Histogram of Cost Changes for Customers With and Without PV Systems

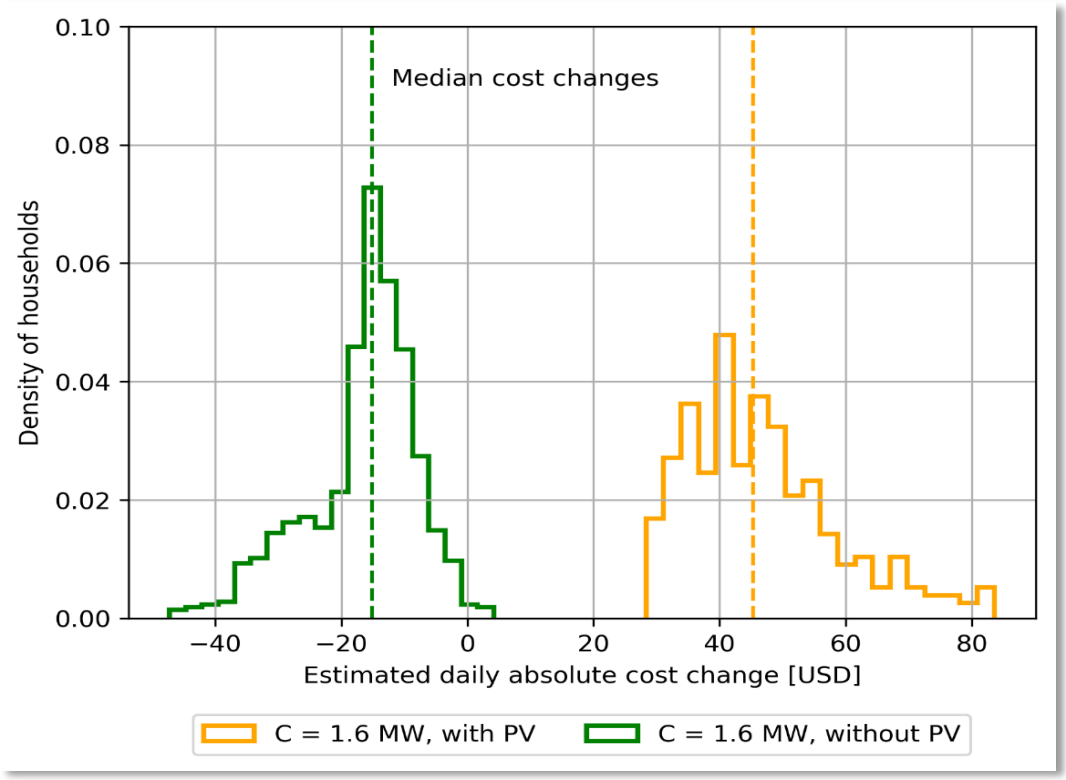

Cost changes for houses with and without PV systems participating in a transactive system compared with a fixed retail tariff

Source: SLAC National Accelerator Laboratory

Owners of Distributed Electricity Resources

DERs include not only solar systems but also batteries and electric vehicles. It is assumed that in the future those appliances will be much more widespread and become a significant part of electric systems. A transactive system can only be successful if it is able to integrate relevant resources and set positive incentives for an increasingly flexible electric system.

With transactive systems, the energy component of the fixed retail tariff increases from 26.17 USD/MWh to 33.40 USD/MWh for inflexible loads (procurement cost for inflexible load over the amount of energy purchased). Under this assumption, median households with flexible HVAC systems but no DERs will experience estimated yearly savings of 2.07 USD. Without a transactive system, the median household had an energy bill of 79.05 USD. For households

with PV systems, the energy portion of the bill increases by 33.30 USD. While before, under net metering, the median household earned 36.72 USD, this income reduces to 5.97 USD. Finally, for households with PV, batteries, and EVs, savings are highest. While their bill was 152.35 USD (batteries not active under a constant retail tariff), the median household even receives a net payment of 9.74 USD. A transactive system can only be successful if it is able to incentivize investment in DERs.

Figure 64: Histogram of Cost Changes for Customers with Different Distributed Energy Resources

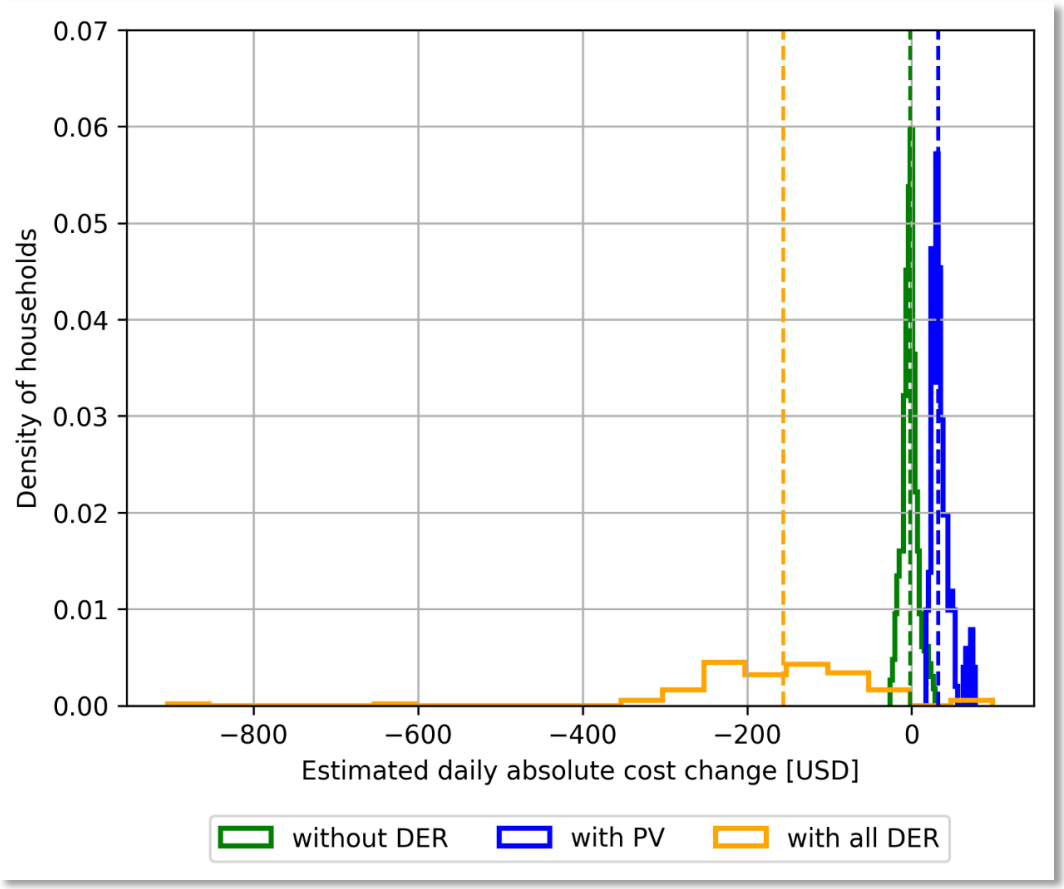

Cost changes for houses without DER, with PV systems, and with PV and battery systems as well as electric vehicles participating in a transactive system compared to a fixed retail tariff.

Source: SLAC National Accelerator Laboratory

Retailer

The gains of the retailer depend on the institutional framework chosen. Currently, a retailer acts on the wholesale market on behalf of the customer, buying energy and selling it at a fixed tariff. Because the procurement risk is carried by the retailer, shielding the customer from wholesale market fluctuations, any profits consist of the service of procurement as well as the risk premium or insurance premium that can be collected from the customer. It needs to be determined if the retailer keeps this role of procuring energy or if the retailer will be entirely substituted by the market operator. Although retailers are not required (or in some jurisdiction even allowed) to do so, they provide an important alternative to an integrated utility. If a

retailer cannot make a business case for participating in the system, it is not clear then whether the system is actually viable, even if operated under the monopoly utility model.

Grid Operator

A grid operator can save costs by postponing or substituting investments in the local infrastructure, for example in a transformer to the overlaying grid. The net benefit depends on the regulatory details. In fact, if regulation focuses solely on long-term capital expenditures (CAPEX), a cost-saving investment into a transactive system architecture might be profit-decreasing if it decreases capital investments (which in many regulations build the basis for calculating allowable profits). In that case, this might be an advert incentive and lead to an omission of otherwise welfare-improving investments. The grid operator is a necessary participant, and as the sole participant has effective veto power over the entire system. If the grid operator cannot be made whole, then it seems unlikely the system will be able to operate at all.

Market Operator

The market operator collects the congestion rents, or the profit from selling electricity locally (at a higher price during congestion) and buying electricity on the wholesale market at a lower price that does not include congestion costs. Sensitive regulation needs to ensure that those benefits get re-distributed to customers through investing in network capacity enhancements. If the market operator cannot realize profits by arbitraging price differences between wholesale and retail operations, then it seems unlikely that the system can be operated successfully in the long term because there will be no wholesale benefits to share with retail participants.

Furthermore, if the retailer, the market, and the grid operator are united in one institution, this might increase the adverse incentive of non-investing.

Lessons for Market Design

The results show that, in principle, a local energy market is able to manage capacity constraints, inform the system operator about economic curtailment, and provide long-term investment incentives into more flexible loads.

However, for a local energy market to be efficient and effective, it must be well designed and operated. In particular, good forecasting tools regarding unresponsive load are valuable to make the best use of available grid resources. If the forecast is too low, the technical constraints will be exceeded. Otherwise, if the forecast is too high, the available grid capacity will be under-utilized and customers that could be part of the allocation do not get dispatched.

Having a power-only market with specified market intervals (here at five minutes) limits the ability of the system or market operator to adjust dispatch resources when problems arise. Smaller market intervals might be promising but do not align with the duty cycles of many appliances and their must-run requirements. Market designs other than power-only markets or capacity-only markets, for example, ramping or storage, as well as the introduction of efficient forward markets or flexible time-horizons could be promising.

Another important factor that limits the effectiveness of the market design as described this project is the observation that the actual price response of devices does not necessarily

conform with the bids posted to the market. The reason is that actual power consumption depends on the internal state of the appliance and its environment (for example, ambient temperature or down time), which may be unavailable, or uncertainty when the initial bid is submitted. Consumers are not necessarily correctly incentivized to respond in accordance with their original bids, but they also do not have full control over the response even if they wished to respond correctly. In current wholesale market designs, retailers are required to balance procurement and demand for the customers they supply. If they fail to do so, they face penalties. Various balancing concepts and "strategy-proof" mechanisms for local energy markets that engage individual customers and devices have been proposed in the literature, but they have yet to be tested in the field and should be the subject of future research.

CHAPTER 7:
Conclusions and Future Research

Conclusions

This project demonstrated a cloud-based DER management system that considered integration with legacy systems for monitoring and controlling all functions within energy-smart communities (including smart inverters and smart meter functions that handle renewable intermittency issues). Powernet managed demand-response and controlled renewable energy, storage, and loads through an auction market.

A large-scale simulation platform based on GridLAB-D was created to simulate and analyze results for different scenarios, markets, and levels of penetration of renewable energy resources. The platform allowed integration of a distribution system's physical and economic layers to conduct these analyses. The goal was to expand market simulation capabilities to better understand these systems on a large scale and examine a local energy market's ability to control load flexibility on a distribution level.

A mathematical model of the feedback system controlled by the market was developed to define the stability and performance of the system. It allowed designing and quantifying the market controller and evaluating its impact in the system of the resident's preference through the bidding to the market. The advantage of conducting this analysis through analytical models gives the option of quantifying stability and performance limits that are cumbersome to obtain running multiple time-domain simulations for different scenarios and operation conditions. The model provided valuable insight into market design at the distribution level.

The technology developed was installed in a small number of houses in the City of Fremont, California. It tested communication between the different devices and the Cloud, and the algorithms developed, using the simulation platform and the laboratories. The results in the communication were consistent with previous test results by PNNL. Communication times between the Cloud and the devices were negligible with respect to the 5-minute clearing time used to operate the market.

An analysis of the system for different operation conditions and penetrations of renewable energy resources was performed. The evaluation of the system performance was conducted when different renewable energy and storage units were participating in the market with differing resident preferences. A welfare analysis was performed for different types of renewable and storage energy units and for flexible loads and degrees of penetration in the market. The analysis showed that a transactive system (or local energy market) was able to provide effective incentives for flexible loads like electric vehicles and batteries to both shift loads and provide flexible supplies to avoid violations of the capacity constraint. A local energy market could also increase welfare by achieving efficiency in the short-term, the long term, and in the efficient use of supply.

Lessons Learned

The results showed that, in principle, a local energy market is able of managing capacity constraints, informing the system operator of economic curtailments, and providing long-term investment incentives into flexible loads. However, for a local energy market to be both efficient and effective, it must be well designed and operated. For a transactive system to be beneficial, the welfare analysis should also be positive.

Different stakeholders and customer groups were differently affected. For example, PV customers in the benchmark case profited from net metering and had bill increases, while customers with batteries had substantial bill decreases. In the long run, this dynamic provided additional investment incentives for storage, though this vision is far from being fully realized—in part because of hardware costs and performance limitations, in part because of a lack of financial incentives for continuous participation. Part of this issue will be addressed in a future follow-up project.

Suggestions for Future Research

A second project is being planned based on this project. The United States Department of Energy has funded the initial deployment of the Transactive Energy Services System (TESS) platform, based in part on the experience and technical capabilities developed in this project.

- The initial version of TESS is expected to include a retail market system to discover the real-time price of electricity for customers with controllable loads, batteries, rooftop PV, and electric vehicles.
- TESS will demonstrate a platform for utilities that addresses the principal phases of adopting a transactive energy system: program development, resource deployment, and system operation. TESS will apply a whole-business, life-cycle approach to transactive energy services by providing an end-to-end framework that a majority of North American utilities can use to plan, finance, design, deploy and operate a transactive energy system. A goal is to ease the challenges associated with the high penetration of renewable resources on the electric grid.
- Evaluate the system when operating under emergency conditions or islanding. To maximize system flexibility, the transactive-control strategy must be combined with other control strategies to be able to reconfigure the distribution network in case of emergencies.
- Evaluate other products in a market (for example, storage energy) for the transaction, in addition to the capacity and power consumed.
- Conduct a larger deployment after all previous points have been addressed to fully demonstrate all the platform's features.

GLOSSARY AND LIST OF ACRONYMS

Term	Definition
ADC	Analog to Digital Converter
Agent	Element capable of producing an effect
API	Application Program Interface
California ISO	California Independent System Operator
CC	Cloud Coordinator
CCA	Community Choice Aggregator
CosPhi	Cosine of Phi, where Phi is the angle between voltage and current
DERs	Distributed Energy Resources
Docker	Provides containers for standardized software code development conditions across different machines
EV	Electric Vehicle
Grid Edge Resources	Technology, solutions, business models advancing the transition toward a decentralized, distributed, transactive electric grid
HH	Home Hub
HIL	Hardware in the Loop
HVAC	Heater, Ventilation, Air Conditioning
LEM	Local Energy Market
MIMO	Multi-Input Multi-Output
PEID	Power Electronics Interface Device
PV	Photovoltaic
Reduced Model	Simplified model of a system where the order of the dynamics is reduced (e.g.: Dynamic system modeled by N differential equations, *Reduced Model* is modeled by M diff. equations, where N>M – See Appendix A)
SOC	State of Charge
VSC	Variable Structure Control
Welfare	State of doing well respect to fortune or prosperity
WS	Wholesale

REFERENCES

Borenstein. 2005. "The Long-Run Efficiency of Real-Time Electricity Pricing." *The Energy Journal* 26 (3): 93-116.

Caramanis, Schweppe, and Tabors. 1983. *Sopt Pricing and Its Relation to Other Management Methods.* MIT-EL 83-001, MIT Energy Laboratory.

Cezar, Gustavo, Elizabeth Buechler, Nikola Milivojevic, Abbas El-Gamal, and Ram Rajagopal. 2020. "Powernet in Farms: A could-based approach to manage electrical loads in livestock farming." *TechConnect World Innovation Conference.* Accepted.

Chassin, David, and GridLAB-D team PNNL. 2008. *GridLAB-D: Residential module user's guide.* http://gridlab-d.shoutwiki.com/wiki/Residential_module_user%27s_guide.

Chassin, David, Jason Fuller, and Ned Djilali. 2014. "GridLAB-D: An agent-based simulation framework for smart grids." *Journal of Applied Mathematics* 1-12.

Chassin, David, K. Schneider, and C. Gerkensmeyer. 2008. "GridLAB-D: An open-source power systems modeling and simulation environment." *IEEE-PES Transmission and Distribution Conf. and Exp.* Chicago, IL. 1-5.

DeCarlo, R.A., S.H. Zak, and G.P. Matthews. 1988. "Variable Structure Control of Nonlinear Multivariable Systems: A Tutorial." *Proc. IEEE* 76 (3): 212-232.

EIA. 2009. *Residential Energy Consumption Survey (RECS).* https://www.eia.gov/consumption/residential.

ENERGY STAR. 2011. *Home Sealing Specification.* https://www.energystar.gov/ia/home_improvement/home_sealing/ES_HS_Spec_v1_0b.pdf.

Filippov, A.F. 1988. *Differential Equations with Discontinuous Righthand Sides.* Dordrecht, The Nederlands: Kluwer Academic.

Gao, Weibing, Yufu Wang, and Abdollah Homaifa. 1995. "Discrete-Time Variable Structure Control Systems." *IEEE Trans. on Industrial Electronics* 42 (2): 117-122.

Hammerstron, D.J., R. Ambrosio, J. Brous, T.A. Carlon, D.P. Chassin, J.G. DeSteese, R.T. Guttromson, et al. 2007. *Pacific Northwest GridWise Testbed Demonstration Projects - Part I. Olympic Peninsula Project.* PNNL-17167, Richland, Washington : Pacific Northwest National Laboratory.

Hu, Junjie, Guangya Yang, Koen Kok, Yusheng Xue, and Henrik Bindner. 2017. "Transactive control: a framework for operating power systems characterized by high penetration of distributed energy resources." *J. Modern Power System Clean Energy* 451-464.

Hung, John, Weibing Gao, and James Hung. 1993. "Variable Structure Control: A Survey." *IEEE Trans. on Industrial Electronics* 40 (1): 2-22.

IEEE, Distribution System Analysis Subcommittee. 2010. *IEEE 123 Node Test Feeder.* http://sites.ieee.org/pes-testfeeders/resources/.

Khalil, Hassan. 1996. "Ch 9: Singular Perturbations." In *Nonlinear Systems*, 351-398. Upper Saddle River, New Jersey - USA.: Prentice-Hall, Inc.

Kokotovic, P.V., H.K. Khalil, and O'Reilly. 1986. *Singular Perturbation Methods in Control: Analysis and Design.* London: Academic Press, Inc.

NREL. 2018. *NREL Data.* https://sam.nrel.gov https://pvwatts.nrel.gov/pvwatts.php.

Pecan Street, Co. 2015. *Pecan Street Data.* https://www.pecanstreet.org/dataport/about/.

PNNL. 2012. *Building America Solution Centers.* https://basc.pnnl.gov/images/iecc-climate-zone-map.

Radavanovic, Ana, Ram Rajagopal, and Sila Kiliccote. 2016. "Powernet for Distributed Energy Resources Networks." *PESGM.*

Rasch, K. 2013. *Smart Assistants for Smart Homes.* Doctoral Dissertation, Stockholm - Sweden: Royal Institute of Technology.

Roozbehani, Mardavij, Munther Dahleh, and Sanjoy Mitter. 2010. "On the Stability of Wholesale Electricity Market under Real-Time Pricing." *49th IEEE Conf. on Decision and Control.* Atlanta, GA, USA. 1911-1918.

Tesfatsion, Leigh, and Swathi Bettula. 2019. *Notes on the GridLAB-D Household Equivalent Thermal Parameter Model.* May 28. https://lib.dr.iastate.edu/econ_workingpapers/60.

Utkin, V. 1978. Sliding Modes and Their Application in Variable Structure Control. Moscow: MIR.

_____. 1992. Sliding Modes in Control and Optimization. Berlin: Springer.

APPENDIX A:
Control of Distributed Resources

Two Time-Scale Dynamics

For such systems where there is a large ratio between the dynamics, it is possible to analyze and simplify them using singular perturbation theory (Khalil 1996). The dynamics of the system can be expressed generically as (singular perturbation form)

$$\dot{x} = f(t, x, z, \epsilon),$$

$$\epsilon \dot{z} = g(t, x, z, \epsilon)$$

where $x \in R^n, z \in R^m$, are the state variables, $t \in R^+$, is the time and the functions $f(.), g(.)$ are smooth. The state variables $x(t), z(t)$ are function of the time, in this work the explicit time dependence can be omitted if it is clear in the context. The parameter $\epsilon \to 0$ could be defined as a ratio among slow and fast time constants of the system. In particular in the limit, setting $\epsilon = 0$ causes a fundamental and abrupt change in the dynamic properties of the system as the differential equation $\epsilon \dot{z} = g$ degenerates into the algebraic equation $g(t, x, z, 0) = 0$. In that case, the overall dimension of the state equation reduces from $n + m$ to n. If $g(t, x, z, 0) = 0$ has multiple isolated real roots $z = h_i(t, x)$, with $i = 1, 2, \dots, k, \dots$, the model is in the standard form. This assumption ensures that a well-defined *n*-dimensional reduced model will correspond to each root $z = h_i(t, x)$. To obtain the n^{th}-reduced model for $\epsilon = 0$

$$\dot{x} = f(t, x, h_i(t, x), 0)$$

This model is sometimes called a quasi-steady-state model or slow model.

Singular perturbation causes a multitime-scale behavior of dynamic system characterized by the presence of slow and fast transients in the system's response to external stimuli. To evaluate the time-scale properties of the standard model, the initial value problem can be solved

$$\dot{x} = f(t, x, z, \epsilon), \quad x(t_o) = \psi(\epsilon)$$

$$\epsilon \dot{z} = g(t, x, z, \epsilon), \quad z(t_o) = \eta(\epsilon)$$

where $\psi(\epsilon)$ and $\eta(\epsilon)$ depend smoothly on ϵ and $t_o \in [0, t_1)$. Let $x(t, \epsilon)$ and $z(t, \epsilon)$ denote the solutions of the full differential equation. When the corresponding problem is defined for the reduced model, it is possible only to specify n initial conditions since the model is n^{th}order. The reduced problem is

$$\dot{x} = f(t, x, h(t, x), 0), \quad x(t_o) = \psi_o = \psi(0)$$

Since the variable the variable $z(t)$ has been excluded from the reduced model, the only information that it is possible to obtain about $z(t)$ is to compute $\bar{z}(t) = h(t, x(t))$, which describes the quasi-steady-state behavior of $z(t)$. The original variable $z(t)$ starts at t_o from $z(t_o) = \eta(\epsilon)$, while the quasi-steady-state variable is not free to start from a prescribed value

and it is forced to start at $\bar{z}(t_o) = h\big(t_o, x(t_o)\big) = h(t_o, \psi(0))$. It may be a large discrepancy between both initial values $z(t_o)$ and $\bar{z}(t_o)$.

If for $[t_b, t_1]$, with $t_o < t_b < t_1$, the error $z(t) - \bar{z}(t) \to 0$, it must be true that during the initial interval $[t_o, t_b]$ the variable $z(t)$ approaches $\bar{z}(t)$. This transient is fast because the speed of $z(t)$ is large since $\dot{z}(t) = g/\epsilon$. It is more convenient in the analysis of this fast dynamics to perform a change of variables $y(t) = z(t) - h(t, x(t))$ that shifts the quasi-steady state of $z(t)$ to the origin. In the new variables (x, y), the full problem is

$$\dot{x} = f(t, x, y + h(t, x), \epsilon), \quad x(t_o) = \psi(\epsilon)$$

$$\epsilon\dot{y} = g(t, x, y + h(t, x), \epsilon) - \epsilon\frac{\partial h}{\partial t}$$

$$-\epsilon\frac{\partial h}{\partial x} f(t, x, y + h(t, x), \epsilon), \quad y(t_o) = \eta(\epsilon) - h(t_o, \psi(\epsilon))$$

To analyze the fast dynamics, it is important to note that $\epsilon\dot{y}$ may remain finite even when ϵ tends to zero and $\dot{y}$ tends to infinity. It can be set

$$\epsilon\frac{dy}{dt} = \frac{dy}{d\tau}; \text{hence}, \frac{d\tau}{dt} = \frac{1}{\epsilon}$$

The new time variable $\tau = (t - t_o)/\epsilon$ is stretched, that is if ϵ tends to zero, τ tends to infinity even for finite t only slightly larger than t_o by a fixed difference. In the τ time scale, the fast dynamics is represented by

$$\frac{dy}{d\tau} = g(t, x, y + h(t, x), \epsilon) - \epsilon\frac{\partial h}{\partial t}$$

$$-\epsilon\frac{\partial h}{\partial x} f(t, x, y + h(t, x), \epsilon), \quad y(0) = \eta(\epsilon) - h(t_o, \psi(\epsilon))$$

The variables t and x in the foregoing equation will be slowly varying since $t = t_o + \epsilon\tau$ and $x = x(t_o + \epsilon\tau, \epsilon)$. Setting $\epsilon = 0$ freezes these variables at $t = t_o$ and $x = x(t_o) = \psi(0)$ and reduces the fast dynamics to the autonomous system

$$\frac{dy}{d\tau} = g(t_o, \psi(0), y + h(t_o, \psi(0)), 0), \quad y(0) = \eta(0) - h(t_o, \psi(0))$$

This model is also called the boundary-layer system. In summary, the original system in R^{n+m} represented by

$$\dot{x} = f(t, x, z, \epsilon), \quad x(t_o) = \psi(\epsilon)$$

$$\epsilon\dot{z} = g(t, x, z, \epsilon), \quad z(t_o) = \eta(\epsilon)$$

can be approximated by the reduced model (slow dynamics) in R^n:

$$\dot{x} = f(t, x, h(t, x), 0), \quad x(t_o) = \psi(0)$$

and the boundary-layer system (fast dynamics) in R^m:

$$\frac{dy}{d\tau} = g(t, x, y + h(t, x), 0), \quad y(0) = \eta(0) - h(t_o, \psi(0))$$

where t and x are slow variables. For some values of ϵ in the interval [0, ϵ^*), the original system can be analyzed using this simplified two-scale separation. In particular, the results on stability analysis of the simplified model can be extended to the original model for some range of the parameter ϵ.

The stability analysis of the original system based on the decomposition of the reduced and boundary-layer system has be presented in detail in (Khalil 1996, Kokotovic, Khalil and O'Reilly 1986). To conclude, a result of that analysis that is applicable to systems used in this project is summarized as follows:

Given the singularly perturbated system

$$\dot{x} = f(t, x, z, \epsilon), \quad x(t_o) = \psi(\epsilon)$$
$$\epsilon \dot{z} = g(t, x, z, \epsilon), \quad z(t_o) = \eta(\epsilon)$$

with the origin $(x, z) = (0,0)$ as equilibrium point. Assuming it satisfices all the conditions presented above to define the reduce and boundary-layer systems, if the origin $x = 0$ of the reduced system $\dot{x} = f(t, x, h(t, x),0)$ is exponentially stable and the origin $(x, z) = (0,0)$ of the boundary-layer system $\frac{dy}{d\tau} = g(t, x, y + h(t, x), 0)$ is exponentially stable, uniformly in (t, x), then there exists $\epsilon^* > 0$ such that for all $\epsilon < \epsilon^*$, the origin $(x, z) = (0,0)$ of the original singularly perturbated system is exponentially stable.

Applying this theorem to the system analyzed in the project, for some parameters where the conditions of scale separation are valid, if one assumes that the fast dynamics is exponentially stable and uniformly in (t, x), where x are the state variables of the reduced system, the overall system is stable if the design of the transactive energy control renders stable the slow dynamics. From a practical point of view, if the interaction between the power converters interfacing the energy resources to the grid and the grid components that define the fast dynamics is exponentially stable for all time and operation points (uniformity in (t, x)), then if the transactive energy control designed to distribute the energy resources in the grid is exponentially stable then the operation point of the overall system will be stable.

Complete Thermal Model of the System

The characteristic of the thermal system is that the energy supplied to the home by the HVAC unit is in discrete events (State ON: heating/cooling power, State OFF: zero power/idling power) at different times setting an intrinsic combination of continuous and discrete subsystems. The reduced thermal model for the home was presented in Chapter 3, and it can be represented by (Chassin and PNNL 2008) (Tesfatsion and Bettula 2019)

$$C_A \frac{dT_A}{dt} = -U_A(T_A - T_O) - H_M(T_A - T_M) + Q_A$$
$$C_M \frac{dT_M}{dt} = -H_M(T_M - T_A) + Q_M$$

where Q_M, represents the heat from appliances and solar radiation directly delivered to mass of the house, and Q_A the energy delivered to the air from the heating/cooling system and also a fraction of heat from appliances and solar radiation coupled through the air. Q_A can be expressed by $Q_A = \gamma Q_M + p_d$, where the first term models the fraction of the heat energy coupled to the air and p_d the power delivered by the HVAC unit. This equation can be represented in a compact form for the j^{th} home by

$$\frac{dT_j}{dt} = A_j T_j + B_{1_j} Q_{EXT_j} + B_{2_j}\, u_j$$

where $T_j = \left[T_{A_j}\, T_{M_j}\right]^T$ represents the temperature vector for the j^{th} house including the air and mass temperatures, $Q_{EXT_j} = \left[T_{o_j}\, Q_{M_j}\right]^T$ a vector that includes the perturbations given by the external temperature T_{o_j} and the energy Q_{M_j}, defined by the heat from appliances and solar radiation, $u_j = p_{d_j}$, and the matrices

$$A_j = \begin{bmatrix} -\frac{U_A + H_M}{C_A} & \frac{H_M}{C_A} \\ \frac{H_M}{C_M} & -\frac{H_M}{C_M} \end{bmatrix}, \quad B_{1_j} = \begin{bmatrix} \frac{U_A}{C_A} & \frac{\gamma}{C_A} \\ 0 & \frac{1}{C_M} \end{bmatrix}, \quad B_{2_j} = \begin{bmatrix} \frac{1}{C_A} \\ 0 \end{bmatrix}.$$

The output of the thermal system per house is defined by $y_j = T_{A_j}$, then $y_j = [1\ 0]\, T_j = C_j T_j$. Assuming the resident relates his/her comfort with the price of the energy willing to pay, it is possible to assume that the bidding willingness to pay in the market for this unit, pr_j, is a function of the air temperature T_{A_j}. Then

$$pr_j = f_j\left(T_{A_j}\right) = f_j(y_j)$$

where $f_j(.): R \to R^+$. The full model of the thermal system for one home including the preference of the resident is

$$\frac{dT_j}{dt} = A_j T_j + B_{1_j} Q_{EXT_j} + B_{2_j}\, u_j, \quad y_j = C_j\, T_j, \ \ pr_j = f_j(y_j),$$

Extending this model to multiple homes $\in D = \{1, \ldots., j, \ldots. n_d\}$, the thermal model for the full system is

$$\frac{dT}{dt} = A_{th} T + B_1 Q_{EXT} + B_2\, u, \quad y = C\, T, \ \ pr_d = f(y),$$

where

$$T = [T_1 \quad \cdots \quad T_{n_d}]^T, Q_{EXT} = \left[Q_{EXT_1} \quad \cdots \quad Q_{EXT_{n_d}}\right]^T, u = [u_1 \quad \cdots \quad u_{n_d}]^T, y = [y_1 \quad \cdots \quad y_{n_d}]^T, pr_d = [pr_1 \quad \cdots\ pr_j \cdots \quad pr_{n_d}]^T$$

and

$$A_{th} = \begin{bmatrix} A_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & A_{n_d} \end{bmatrix}, B_1 = \begin{bmatrix} B_{1_1} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & B_{1_{n_d}} \end{bmatrix}, B_2 = \begin{bmatrix} B_{2_1} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & B_{2_{n_d}} \end{bmatrix}, C = \begin{bmatrix} C_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & C_{n_d} \end{bmatrix}$$

It is important to observe that because both the thermal system of the houses and the comfort preference of each resident are independent each other, each house is controlled independently. That is evident from the structure of the aggregate system where the matrices A_{th}, B_1, B_2, C and $f(y)$ are block diagonals.

For each thermal system, the control signal u_j can take two values

$$u_j = p_{d_j} = \begin{cases} d_j & \rightarrow Equipment\ ON \\ 0 & \rightarrow Equipment\ OFF \end{cases}$$

then, the thermal system for a single residence has two dynamical structures which are defined by the differential equations:

$$\frac{dT_j}{dt} = A_j T_j + B_{1_j} Q_{M_j} + B_{2_j} d_j \quad for\ u_j = d_j$$

$$\frac{dT_j}{dt} = A_j T_j + B_{1_j} Q_{M_j} \qquad for\ u_j = 0$$

It defines an intrinsic variable structure system (VSS) defined by two continuous structures that are selected based on the discrete magnitude taken by the control signal u_j.

Sliding Mode Solution

Sliding mode control systems, also known as variable structure control systems (VSC), are characterized by a discontinuous feedback control law which switches the structure of the system during the evolution of the state vector to maintain the state trajectories in a predefined subspace. The state-feedback control law is not a continuous function of time, it can switch from one continuous structure to another based on the current states in the state-space (DeCarlo, Zak and Matthews 1988) (Utkin 1978) (Utkin 1992) (Hung, Gao and Hung 1993)

The multiple control structures are designed so that trajectories always move toward an adjacent region with a different control structure, and so the ultimate trajectory will not exist entirely within one control structure. Instead, it will slide along the boundaries of the control structures. The trajectory of the system as it slides along these boundaries is called sliding mode and the geometrical locus consisting of the boundaries is called sliding hypersurface. The purpose of the switching control law is to steer the system's state trajectory onto a prespecified surface in the state-space and to maintain the plant's state trajectory on this surface for all subsequent time. The surface is called a switching surface or sliding surface. The plant dynamics restricted to this surface represent the controlled system's behavior. Figure A-1 depicts as an example the state trajectories for a system comprised by two structures defined on the surfaces $\mathcal{H}_1$ and $\mathcal{H}_2$. The sliding surface is defined by the intersection between these surfaces $\mathcal{H} = \mathcal{H}_1 \cap \mathcal{H}_2$. In the plot, it can be seen trajectories moving in the surfaces $\mathcal{H}_1$ and $\mathcal{H}_2$ and steered toward the sliding surface. Once they reach the sliding or switching function, the control law forces to maintain the trajectories on the surface $\mathcal{H}$.

Figure A-1: State Space for a Variable Structure System

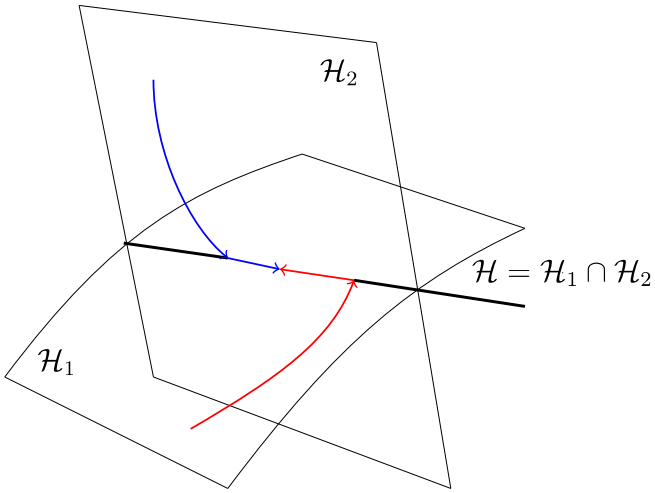

Source: SLAC National Accelerator Laboratory

The first critical phase of a VSC design is to define properly a switching surface so that the plant, restricted to the surface, has desired dynamics, such as stability to the origin, tracking, regulation, etc. The VSC control design breaks down into two phases, the first one corresponds to choose a switching surface so that the plant state restricted to the surface has desired dynamics. The second phase is to design a switched control that will drive the plant state to the switching surface and maintain it on the surface upon interception.

Consider as a generic plant, the thermal system of the j^{th} house represented by the reduced model

$$\frac{dT_j}{dt} = A_j T_j + B_{1_j} Q_{EXT_j} + B_{2_j} u_j, \quad y_j = C_j T_j$$

where the temperature $T_j = \left[T_{A_j}\, T_{M_j}\right]^T \in R^n \equiv R^2$ represents the state variables of the system and

$$u_j \in R^m \equiv R^1;\ u_j = p_{d_j} = \begin{cases} d_j & \rightarrow Equipment\ ON \\ 0 & \rightarrow Equipment\ OFF \end{cases}$$

the discrete control signal.

The objective of the sliding mode control strategy is to drive the states of the system into a $(n-m) = (2-1)$ dimensional subspace $\mathcal{H} \subset R^n \equiv R^2$ and to maintain the subsequent motion of the state trajectories on this $(n-m)$ dimensional manifold $\mathcal{H}$ for time $t > t_1 > 0$. The sliding manifold $\mathcal{H}$ can be defined by

$$\mathcal{H} = \{(T,t) \in R^{2+1} \mid \sigma(T,t) = 0\}$$

As examples, if a control is designed based on a thermostat unit sensing the air temperature at one point of the house (internal temperature control), the sliding manifold is

$$\mathcal{H} = \left\{(T,t) \in R^3 \mid \sigma(T,t) = T_{ref} - T_{A_j} = 0\right\}$$

assuming the reference temperature T_{ref} set by the resident is constant. In case the resident uses the thermal unit to participate local market bidding $pr_j(T_{A_j})$, the sliding manifold is

$$\mathcal{H} = \left\{(T,t) \in R^3 \mid \sigma(T,t) = pr_j(T_{A_j}) - \lambda^*(T_A,t) = 0\right\}$$

where the market price $\lambda^*(T_A,t)$ is a function of the temperature of all of the houses bidding in the market and can change in the time. The bidding price for the j^{th} house, $pr_j(T_{A_j})$, is a function only of the state T_{A_j}.

The control law is not continuous and has the ability to drive trajectories to the sliding manifold $\mathcal{H}$ in finite time. However, once the trajectories reach the sliding surface, the system takes on the character of sliding mode and may converge asymptotically to the equilibrium point in this surface. The sliding manifold $\mathcal{H}$ can be seen as a distance that the state T_{A_j} is away from the surface $\mathcal{H}$. If the state T_{A_j} is outside the surface $\mathcal{H}$, $\sigma(T,t) \neq 0$, while if T_{A_j} is on the sliding surface, $\sigma(T,t) = 0$. The sliding control law switches from one state to another based on the sign of this distance. The variable structure control is

$$u_j = \begin{cases} u_j^+(T,t) & \text{when } \sigma(T,t) > 0 \\ u_j^-(T,t) & \text{when } \sigma(T,t) < 0 \end{cases}$$

such that the state T_{A_j} reaches $\sigma(T,t) = 0$ in finite time. This reaching condition can be evaluated by choosing a Lyapunov function candidate

$$V(T,t) = \sigma(T,t)^T \sigma(T,t) = \|\sigma(T,t)\|^2 > 0 \quad when \ \sigma(T,t) \neq 0$$

The global reaching condition is given by

$$\dot{V}(T,t) < 0 \quad when \ \sigma(T,t) \neq 0$$

It derivates in

$$\dot{V}(T,t) = 2\,\sigma(T,t)\,\dot{\sigma}(T,t) < 0 \text{ or } sign\,\sigma(T,t) \neq sign\,\dot{\sigma}(T,t)$$

Note that

$$\dot{\sigma}(T,t) = \frac{\partial \sigma}{\partial T}\frac{\partial T}{\partial t} + \frac{\partial \sigma}{\partial t} = \frac{\partial \sigma}{\partial T}(A_j T_j + B_{1_j} Q_{EXT_j} + B_{2_j} u_j) + \frac{\partial \sigma}{\partial t},$$

then control signal u_j has to define the $sign\,\dot{\sigma}(T,t)$ to make the trajectories to converge toward $\sigma(T,t) = 0$ and remain in the sliding surface $\mathcal{H}$.

Applying to the thermal model of the house controlled by a thermostat, the surface $\mathcal{H}$ was defined by

$$\mathcal{H} = \left\{(T,t) \in R^3 \mid \sigma(T,t) = T_{ref} - T_{A_j} = 0\right\},$$

then $\frac{\partial \sigma}{\partial T} = \left[\frac{\partial \sigma}{\partial T_{A_j}}, \frac{\partial \sigma}{\partial T_{M_j}}\right] + \frac{\partial \sigma}{\partial t} = [-1\ 0] + 0$, and if $sign\ \dot{\sigma}(T,t)|_{u_j^-} \neq sign\ \dot{\sigma}(T,t)|_{u_j^+}$, the reaching condition is satisfied if

$$[-1\ 0](A_j T_j + B_{1_j} Q_{EXT_j} + B_{2_j}\ u_j^+) = [-1\ 0](A_j T_j + B_{1_j} Q_{EXT_j} + B_{2_j}\ d_j) < 0 \quad if \quad \sigma(T,t) = T_{ref} - T_{A_j} > 0$$

and

$$[-1\ 0](A_j T_j + B_{1_j} Q_{EXT_j} + B_{2_j}\ u_j^-) = [-1\ 0](A_j T_j + B_{1_j} Q_{EXT_j}) > 0 \quad if \quad \sigma(T,t) = T_{ref} - T_{A_j} < 0$$

Then $u_j = d_j\ sign\big(\sigma(T,t)\big) = d_j\ sign\left(T_{ref} - T_{A_j}\right)$ is the switching control law to control the temperature of the house assuming a heating system ($d_j > 0$). Following similar considerations, the switching control law for a cooling system is $u_j = d_j\ sign\big(\sigma(T,t)\big) = d_j\ sign\left(T_{ref} - T_{A_j}\right)$. Replacing the control law into the differential equation, it becomes:

$$\frac{dT_j}{dt} = A_j T_j + B_{1_j} Q_{EXT_j} + B_{2_j}\ d_j\ sign\big(T_{ref} - C_j\ T_j\big), \quad y_j = C_j\ T_j$$

This equation has unique solution if $sign\big(T_{ref} - C_j\ T_j\big) = 1$ for $T_{ref} > C_j\ T_j$ or $sign\big(T_{ref} - C_j\ T_j\big) = 0$ for $T_{ref} < C_j\ T_j$, and corresponds to the reaching condition analyzed previously. When the trajectory reaches the sliding surface $\mathcal{H}: \sigma(T,t) = 0$, the $sign\big(T_{ref} - C_j\ T_j\big)$ is not defined and the differential equation has not unique solution because it has discontinuous right-hand side. Thus, they fail to satisfy conventional existence and uniqueness results of differential equation theory (Lipschitz continuous). Nevertheless, an important aspect of VSC is the presumption that the plant behaves uniquely when restricted to $\sigma(T,t) = 0$. One of the earliest approaches addressing existence and uniqueness is the Filippov method (Filippov 1988). Another solutions are presented ideally interpreting the $sign(.)$ function as a relay operating at infinitely high frequency between the limits 0 and 1 when $\sigma(T,t) = 0$. Under this condition, the sliding mode motion follows the invariance properties such that $\sigma(T,t) = 0$ and $\dot{\sigma}(T,t) = 0$. The last condition is necessary for the state trajectory to stay on the switching surface $\sigma(T,t) = 0$. It is possible to define an equivalent control signal u_{EQ_j} for the thermal system such that the state trajectory stays on the switching surface $\sigma(T,t) = 0$ (Utkin 1978). This signal u_{EQ_j} as the averaged signal of the control input u_j when it is switching between 0 and 1 at high frequency to keep the trajectory invariant on $\sigma(T,t) = 0$. Using the invariance condition, the equivalent control signal u_{EQ_j} can be calculated as

$$\dot{\sigma}(T,t) = \frac{\partial \sigma}{\partial T}(A_j T_j + B_{1_j} Q_{EXT_j} + B_{2_j}\ u_{EQ_j}) + \frac{\partial \sigma}{\partial t} = 0$$

For the example analyzed, $\sigma(T,t) = T_{ref} - T_{A_j}$, $\frac{\partial \sigma}{\partial t} = 0$ and $\frac{\partial \sigma}{\partial T} = [-1\ 0]$ then

$$u_{EQ_j} = -\left(\frac{\partial \sigma}{\partial T} B_{2_j}\right)^{-1} \left(\frac{\partial \sigma}{\partial T}(A_j T_j + B_{1_j} Q_{EXT_j})\right),$$

assuming that $\left(\frac{\partial \sigma}{\partial T} B_{2_j}\right)^{-1}$ exists. The equivalent control exists and forces the state trajectory to stay on $\sigma(T,t) = 0$ if u_{e_j} is within the range of $u_j, 0 < u_{EQ_j} < d_j$.

The mathematical analysis presented uses the limit condition that the relay ideally switches at infinitely high frequency. This assumption is not practical because the heating/cooling system cannot switch between the ON-OFF states at infinite frequency. Given the dynamics associated with the thermal system in the houses, one option commonly implemented in commercial units is to transform the relay action by a hysteresis such that the equipment does not switch between ON-OFF states at high frequency. In this case, the sliding manifold $\mathcal{H}$ is transformed from $\sigma(T,t) = 0$ to $|\sigma(T,t)| \leq h$, where $2h$ is the hysteresis width. Forced by this control law, state trajectories converge to the manifold $\mathcal{H}$ and the equipment switches between ON-OFF state to control the temperature within the range

$$|\sigma(T,t)| \leq h \;\rightarrow\; \left|T_{ref} - T_{A_j}\right| \leq h \;\rightarrow\; T_{ref} - h \leq T_{A_j} \leq T_{ref} + h$$

When the temperature T_{A_j} is within that range, the equipment operates switching ON-OFF at a non-fixed frequency defined by the thermal parameters of the house and the external perturbations. Another practical option is to implement the relay action but defining a fixed time interval when the decision is taken. The system, called discrete-time VSC, is forced periodically to commute at the switching event $t = k\,T_s$ with $k = 0, 1, \ldots, \infty$ and T_s is the minimum switching period (Gao, Wang and Homaifa 1995). Then, the control strategy is

$$u_j(T,k) = \begin{cases} d_j \rightarrow \sigma(T,k) = T_{ref} - T_{A_j}(k) > 0 \\ 0 \rightarrow \sigma(T,k) = T_{ref} - T_{A_j}(k) < 0 \end{cases} \quad \text{for } k = 0, 1, \ldots, \infty.$$

Figure A-2 depicts the responses of air and mass temperatures of the j^{th} house when the two practical controls described above are used. In the case of the hysteresis (Fig. A-2-a), the trajectory of the air temperature T_{A_j} starts at $T_{A_j}(0) = 60^o F$ outside of the sliding manifold $\mathcal{H}$: $\left|T_{ref} - T_{A_j}\right| \leq h$ and converges to it in finite time. When T_{A_j} is within $T_{ref} - h \leq T_{A_j} \leq T_{ref} + h$, with $T_{ref} = 70^o F$ and $h = 1^o F$, the heating system switches ON-OFF to keep the air temperature around the reference temperature rejecting external perturbations like the energy due to the sun irradiance and the outside temperature. Similarly, Fig. A-2-b shows the case where the thermal system in the j^{th} house is controlled by the discrete-time variable structure control strategy assuming that the minimum switching period $T_s = 5\,min.$

Figure A-2: Response of the Thermal System of a House

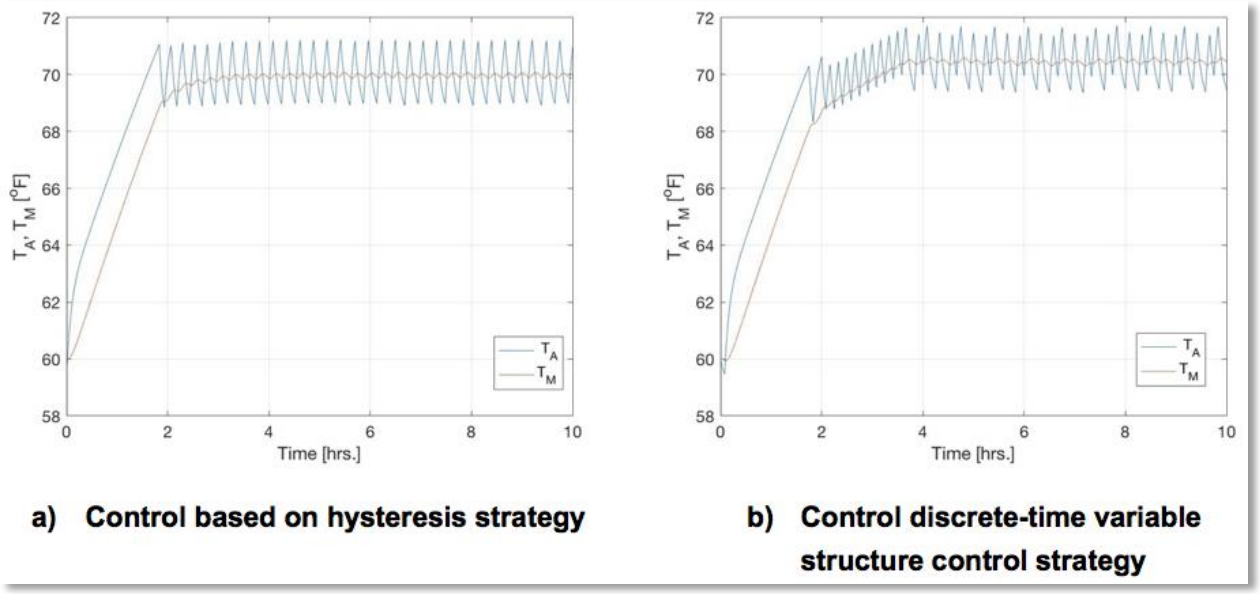

a) Control based on hysteresis strategy

b) Control discrete-time variable structure control strategy

Source: SLAC National Accelerator Laboratory

The example presented above can be extended to cover the aggregated thermal system comprised by multiple distributed houses. The complete system's model was presented in detail in this appendix as

$$\frac{dT}{dt} = A_{th}T + B_1 Q_{EXT} + B_2\, u, \quad y = C\ T,\ \ pr_d = f(y),$$

where all the matrices have a block-diagonal structure defining dynamic interdependence among the houses. The control for each unit will follow similar considerations and strategies to the one presented for the j^{th} house.

Associating a n_d number of houses in the aggregated distribution, the order of the systems is $2n_d$, corresponding to the temperatures in the air and the envelope of each house. To design the control, it is possible to define a $(2n_d + 1)$-dimensional switching surface or equilibrium manifold for the case corresponding to houses participating with the equipment in the market. It is given by,

$$\mathcal{H} = \{(T_A, t) \in R^{2n_d+1} \mid \sigma(T_A, t) = 0\}$$

where

$$\sigma(T_A, t) = [\sigma_1(T_A, t), \quad \dots,\ \sigma_j(T_A, t), \dots, \quad \sigma_{n_d}(T_A, t)]^T$$
$$= \left[pr_1(T_{A_1}, t) - \lambda^*(T_A, t), \quad \dots, \quad pr_{n_d}(T_{A_{n_d}}, t) - \lambda^*(T_A, t)\right]^T$$

for the case the agent is participating with the thermal equipment in a local market. These surfaces are designed so that the system state trajectory, restricted to $\sigma(T_A, t) = 0$, has a desired behavior such as stability or tracking. After proper design of the surface, a switched controller $u(T, t) = \left[u_1(T, t), \dots, u_{n_d}(T, t)\right]^T$, where the control signals $u_j(T, t)$ with $j = 1, \dots, n_d$ are defined by

$$u_j \in R^1;\ u_j = p_{d_j} = \begin{cases} d_j & \rightarrow Equipment\ ON \\ 0 & \rightarrow Equipment\ OFF \end{cases}$$

In this control strategy, each surface $\sigma_j(T_A, t)$ is associated with a control signal u_j. The control law is very similar to the one presented for a single house given the independence that exist among the units. As such,

$$u_j = \begin{cases} u_j^+(T,t) & \text{when } \sigma_j(T,t) > 0 \\ u_j^-(T,t) & \text{when } \sigma_j(T,t) < 0 \end{cases}$$

then the control for each house is $u_j = d_j\ sign\left(\sigma_j(T,t)\right)$, assuming that the reaching conditions for each house are satisfied.

Following the case of participation in a local market, when all the trajectories reach the sliding surface $\mathcal{H}: \sigma(T,t) = 0$, all the houses are setting the comfort level or internal temperature based on the common price $\lambda^*(T,t)$ defined by the market. This result is important because it means that in steady state the overall system converges to the state where the individual prices are equal to the market price. It is

$$\sigma(T_A, t) = 0\ \rightarrow \sigma_j(T_A, t) = 0,\ \ \forall\, j = 1, \dots., n_d$$

$$pr_j\left(T_{A_j}, t\right) -\ \lambda^*(T_A, t) = 0\ \rightarrow pr_j\left(T_{A_j}, t\right) =\ \lambda^*,$$

and

$$T_{A_j} =\ f_j^{-1}(\lambda^*),\ \ \forall\, j = 1, \dots., n_d$$

assuming that $f_j^{-1}(.)$ exists.

Robustness to perturbations: Previous definition of the u_{EQ} with $u_{EQ_j} \in (0, d_j)$, forces the trajectory to slide in the switching manifold satisfying , $\sigma(T(t), t) = 0$, and $\dot{\sigma}(T(t), t) = 0$. Then, the impact of the external perturbation Q_M does not affect the motion of the trajectory on the sliding surface if each component of u_{EQ} is within the boundaries $0 < u_{EQ_j} <\ d_j$. Thus, if u_{EQ} is within those limits and $\sigma(T(t), t) = 0$, then $pr_j = \lambda^*\ \forall\ i = 1, \dots, n_d$, independent of the perturbation Q_M. Then the perturbation Q_M does not affect pr_j or the individual house temperature $T_{A_j} =\ f_j^{-1}(\lambda^*)$.

APPENDIX B:
Simulation Platform: Physical layer

Physical Layer – Details

The physical layer of the simulation platform is implemented using GridLAB-D. Multiple residences including appliances, HVAC system, solar photovoltaic units and storage devices are connected to distribution feeders.

Standard distribution system models have been suggested by IEEE PES Distribution System Analysis- Subcommittees Distribution Test Feeder Working Group to tests system via simulations. For this project's scope, the IEEE 123 network model has been selected as a base network (IEEE 2010). It represents a typical residential distribution grid and follows a radial topology operating at 4.16KV and a total power of about 3.6MW. The feeder is connected to the overlaying voltage level by a single transformer where congestion potentially happens. Other studies and tests, with similar structures to the network used in this project, have used the IEEE-123 as distribution feeder in the representation of the area model developed in GridLAB-D (Chassin, Schneider and Gerkensmeyer 2008) (Chassin, Fuller and Djilali, GridLAB-D: An agent-based simulation framework for smart grids 2014). This network includes 123 active nodes where the loads are connected to the distribution system. Loads models in this normalized network are represented by passive elements (resistor, capacitor, inductor) or constant power loads. In this project, these loads have to be replaced by the models of the houses including the details of their appliances and devices.

The IEEE 123 network model was populated with house models that are statistically representative of a typical residential housing stock. Six different basic home types were defined based on the number of stories and the characteristics of the heating and cooling system, as shown in Table B-1.

Table B-1: Basic Home Types for Generating Representative Models

Number of Stories	HVAC System Type
1	Heat pump
2	Heat pump
1	Electric cooling and electric resistance heating
2	Electric cooling and electric resistance heating
1	Electric cooling and natural gas furnace heating
2	Electric cooling and natural gas furnace heating

Source: SLAC National Accelerator Laboratory

Based on these home types, house parameters such as floor area, air leakage rate, and cooling/heating setpoints were generated from probability distributions. These probability distributions were fit from data from the 2009 EIA Residential Energy Consumption Survey (RECS) (EIA 2009) and LBNL Residential Diagnostic Database, and are defined by geographic region and IECC climate zone (ENERGY STAR 2011).The goal is to probabilistically account for

the diversity of residential housing characteristics in the generated house models. Examples of probability distributions of the parameters of the homes are shown in Figure B-1 for IECC climate zone 3C (coastal California) and assuming equal penetration of the six different home types. The floor area is described by a truncated normal distribution, ranging from 1000 to 4000 ft^2. The mean cooling setpoint for this group of generated home models was approximately 71°F.

Figure B-1: Statistical Model to Define the Parameters of the Houses

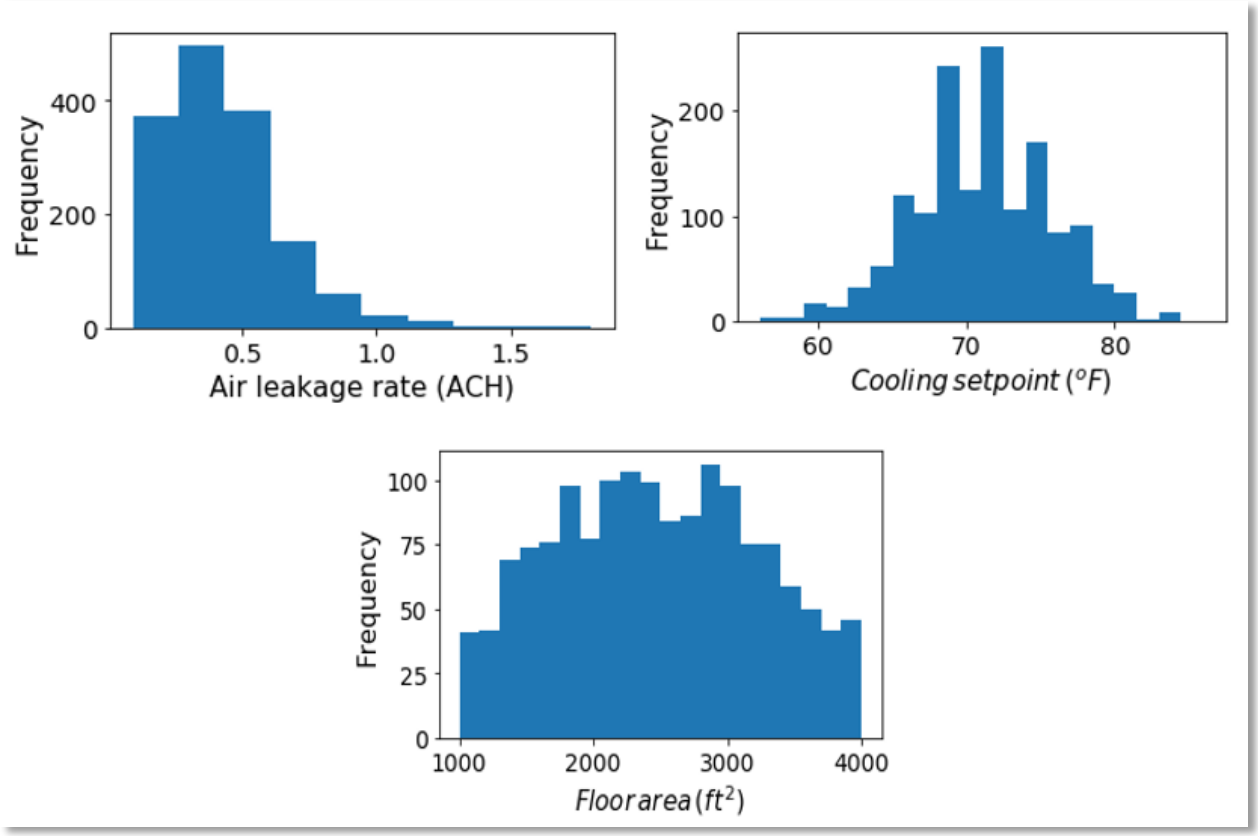

Examples of probability distributions used to define the parameters of the house models in the GridLAB-D simulation. These distributions were generated specifically for IECC climate zone 3C (coastal California).

Source: SLAC National Accelerator Laboratory

Additionally, to the pool of houses created based on statistics, a set of house's models representing the SLAC/Stanford Labs and the Fremont area houses is included in the simulation. These models are included to have a real comparison point between the simulation results and the measurements conducted in the field.

The thermal heating and cooling load of each home was modeled using the equivalent thermal parameters (ETP) model included in GridLAB-D (Chassin, Fuller and Djilali, GridLAB-D: An agent-based simulation framework for smart grids 2014) (Chassin, Schneider and Gerkensmeyer 2008).The ETP model simulates the thermal response of the home and assumes single lumped parameters for the conductance of the building envelope, the conductance of the thermal mass of the home and indoor air, and external and internal heat gains. The aggregate unresponsive load for each simulated home was modelled in GridLAB-D as a static impedance-current-power (ZIP) load using ZIP parameters typical for residential

loads and driven by time-series using historical appliance-level load data from 2015 for homes in San Diego from the Pecan Street database (Pecan Street 2015).

The house dissemination in the network is set by balancing the power consumption of the houses with the power capacity of the distribution lines. The equivalent power consumed by the set of houses connected per node has to be equal to the nominal power defined per node in the normalized network. The power consumption per house is variable along the day/week/month. To allocate the houses per distribution network node the *After Diversity Maximum Demand* (ADMD) *criterion* is used. ADMD is a metric frequently used to determine the maximum power requirements for distribution system designs, and is defined as the coincident peak load of the total consumption present in the network divided by the number of households.

To use this criterion in the simulation platform, a pool of 2000 houses with diverse features and appliances is created and the total power consumption of the houses is evaluated based on different weather conditions and real data for the daily use of appliances. This simulation is run to evaluate months of power consumption and aggregate the diversity of multiple houses to estimate the peak power. This peak power divided by the number of houses included in the simulation defines the ADMD used to allocate the residences in the distribution network. Figure B-2 shows an example of the aggregate demand for a pool of 100 simulated homes for the month of July. The ADMD per home for this calibration period would be equal to the peak power divided by the total number of simulated homes.

Figure B-2: Time Domain Representation of the Total Power Consumption

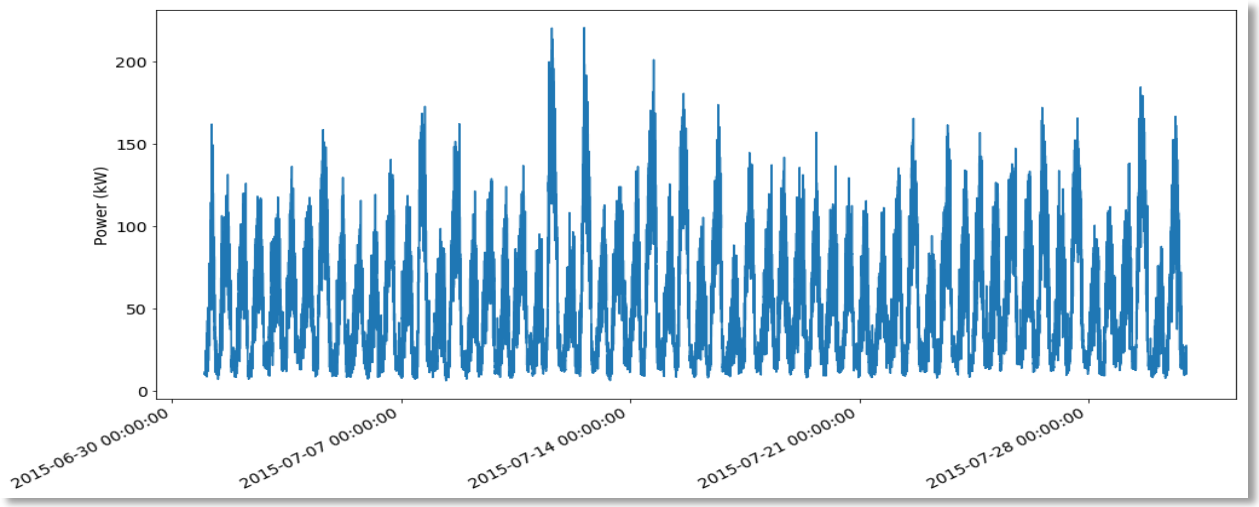

Aggregate demand for a pool of 100 simulated homes for the month of July.

Source: SLAC National Accelerator Laboratory

A part of the studies conducted in this project are related with the impact on the system and the control of renewable sources and storage units. Batteries and solar photovoltaic (PV) cells are placed in the houses using a variable penetration index ranging from 20% to 80-100%. For both devices, the models included in GridLAB-D are used. PV models were defined based on the characteristics of typical residential solar installations (e.g. single-crystalline silicon panels, 14% efficiency, south-west, fixed-axis, 27 degree tilt, etc.) and validated using NREL data from similar panels (NREL 2018). The capacity of each PV system was set during model calibration to be related with the house characteristics where the units is installed. This

magnitude corresponds to the minimum of two values: (1) the maximum capacity given the roof area of the home and (2) the calculated capacity such that the mean daily PV output is approximately equal to the mean daily load or consumption of the home. The same calibration period used for the ADMD calculation was used for PV sizing calculations.

As an example of the modeling and integration of distributed energy resources in the simulation platform, Figure B-3 shows an example of the aggregate net power consumed at bus 4 of the IEEE 123 bus network. In this case, the static loads are replaced by houses with time varying loads and PV generation. Results are shown over a 24 hour period for each different level of PV cell penetration. In the original IEEE 123 bus model, this bus has connected a load of S = 40 KW + j20 KVAR. For this specific day of simulation, there was a wide fluctuation in ambient outdoor temperature, resulting in a heating load in the early morning and a cooling load in the afternoon. In Figure B-3, the blue trace shows the case of the net power consumed by all the houses connected to node 4 when there are no PV cells included in the residences. It represents the baseline power consumed by the appliances and HVAC systems. The power in node 4 for increasing penetration of PV cells in the residences connected to that node are depicted by the brown, yellow and purple traces.

Figure B-3: Power Consumption at IEEE 123 Node 4

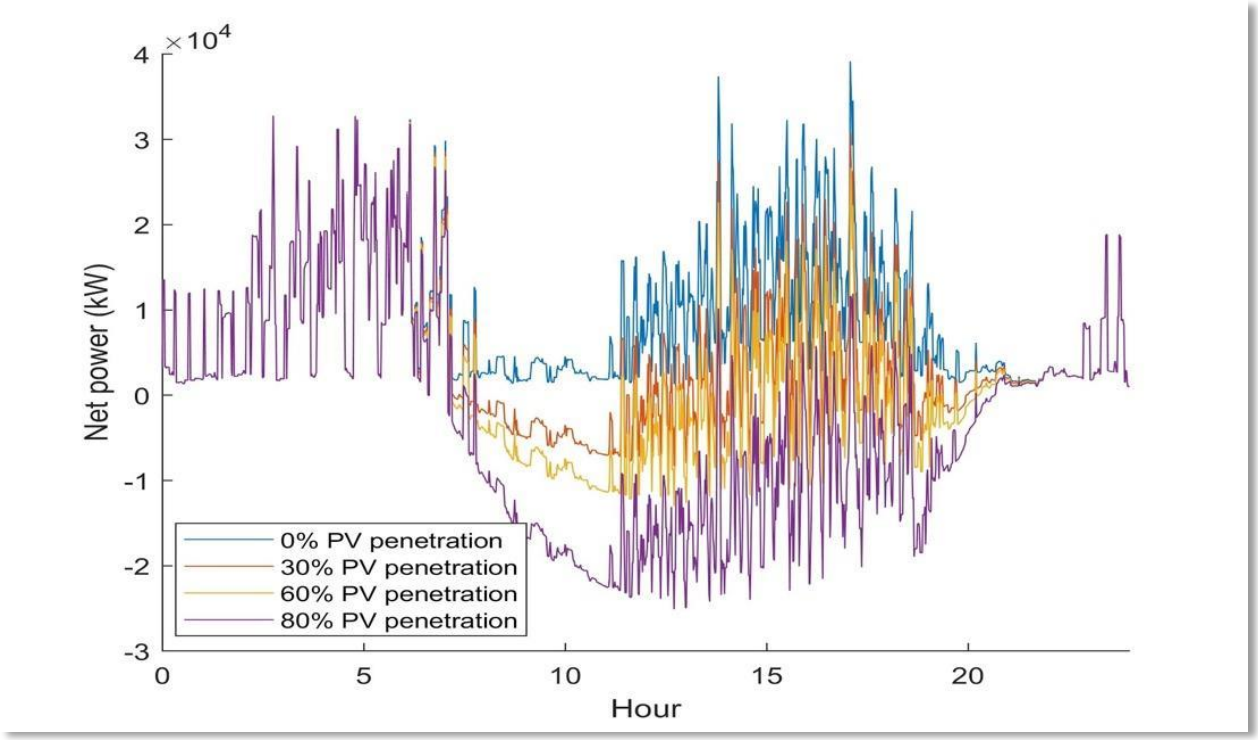

Net aggregate power consumption at bus 4 for four different levels of PV penetration: 0%, 30%, 60%, and 80%.

Source: SLAC National Accelerator Laboratory

Battery models were defined assuming standard characteristics of commercially available lithium-ion batteries. Storage units are only installed in houses with PV systems. The size of the battery is designed according to the PV system installed and such that the capacity is approximately equal to the energy produced by the PV system in one day. Based on this initial value, the final capacity is defined following the commercial characteristics of Sonnen Inc. batteries where the battery size is scaled in steps of 2 kWh.

Two types of vehicle charging stations are distinguished: commercial charging stations and workplace or home charging. The stationary charging stations are modeled by batteries that are connected to the physical system when an electric vehicle (EV) connects to the charger. Real data of residential customers in California is used to identify connection and charging times and estimate the average charging rate and the state of charge (SOC). The data available does not cover throughout the year but the researchers bundle event series for the days of interest.

The data available reports for a charging station when a car has arrived, and a new charging session started. It also includes, for every 15 minutes, the average charging rate and the energy charged, as well as the end of the charging session and time of departure (which might differ if full charge has been achieved before departure). As mentioned, the data does not cover a whole year but, depending on the charging station, covers periods from only a couple of days to a whole year. For that reason, seven consecutive days are randomly picked in the time series available to schedule EV events for a representative week.

Simulation Process

The process of preparing and running the simulation using the platform is depicted in Figure B-4. Before defining the set-up for simulating the system, proceeding with the studies and evaluating the results, it is necessary to design the distributed electrical network and the residences connected to it.

Figure B-4: Simulation Process

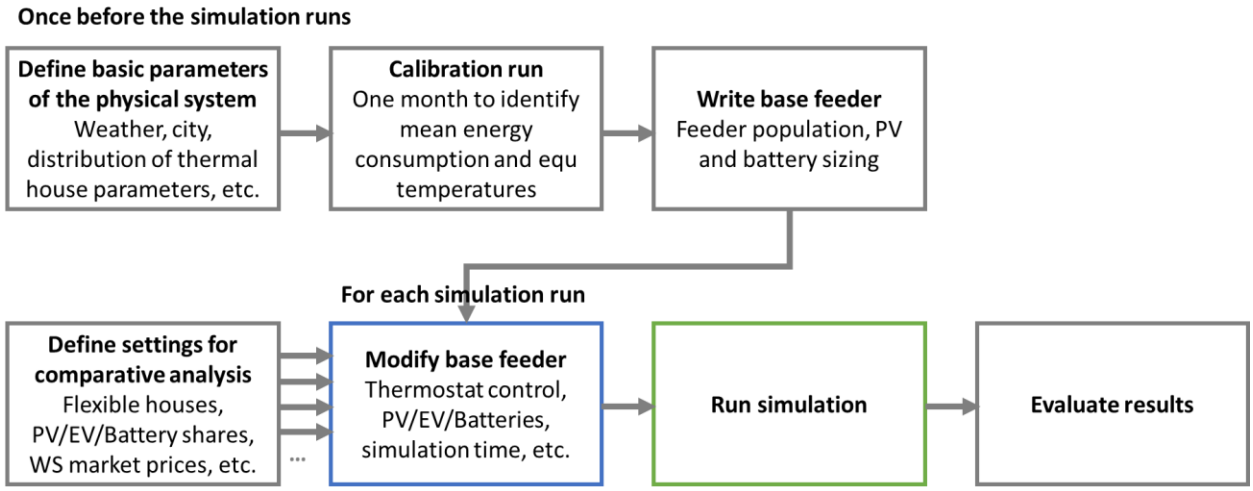

Source: SLAC National Accelerator Laboratory

The first point is to define the area or zone where the study is going to be conducted, that allows to set the information about the weather pattern, conditions to define the house parameters, etc. All this information is available through time-series data for the weather and the use of appliances or statistical information for the house parameters. Based on this data, a routine creates a pool of 2000 diverse houses based on the statistical information and assigns to each residence the internal temperature conditions and the use of appliances. Under these conditions, the operation of this set of houses is simulated for long periods (i.e. a few months) recording the power consumed by each house.

The estimation of the power consumed per house allows aggregating the total power consumption and estimate the *After Diversity Maximum Demand* (ADMD). Based on this

criterion the houses can be allocated in the distribution network and the size of the solar panels and batteries can be assigned as a function of the penetration of renewables and storage units studied. Electric vehicles per home and charging stations are allocated randomly in the set-up. The resulting feeder is used for the following simulations and, depending on the scenario definition (i.e. level of penetration of renewable-storage energy), it is equipped with additional PV modules, batteries, and EVs. The technical parametrizations of the houses as well as the sequence of where additional flexible resources and loads are placed are fixed to ensure comparability of results.

The market layer is initialized by setting the preferences of the residents in the home hub. The market is mainly defined by the wholesale market price and the participation of controllable loads and sources. CAISO real-time price data is used for the node KETTNER_6_N004 (airport in San Diego) and as provided by OASIS. As 2015 data is not available online and requests to CAISO were unanswered, the project team relied on 2016 data.

For each setting, the physical layer file is modified by adding GridLAB-D objects to the base file depending upon the penetration of renewable-storage agents, run, and the results eventually evaluated. Figure B-5 sums up the inputs and outputs of the simulation platform.

Figure B-5: Setup of the Cloud Coordinator – Home Hub Architecture

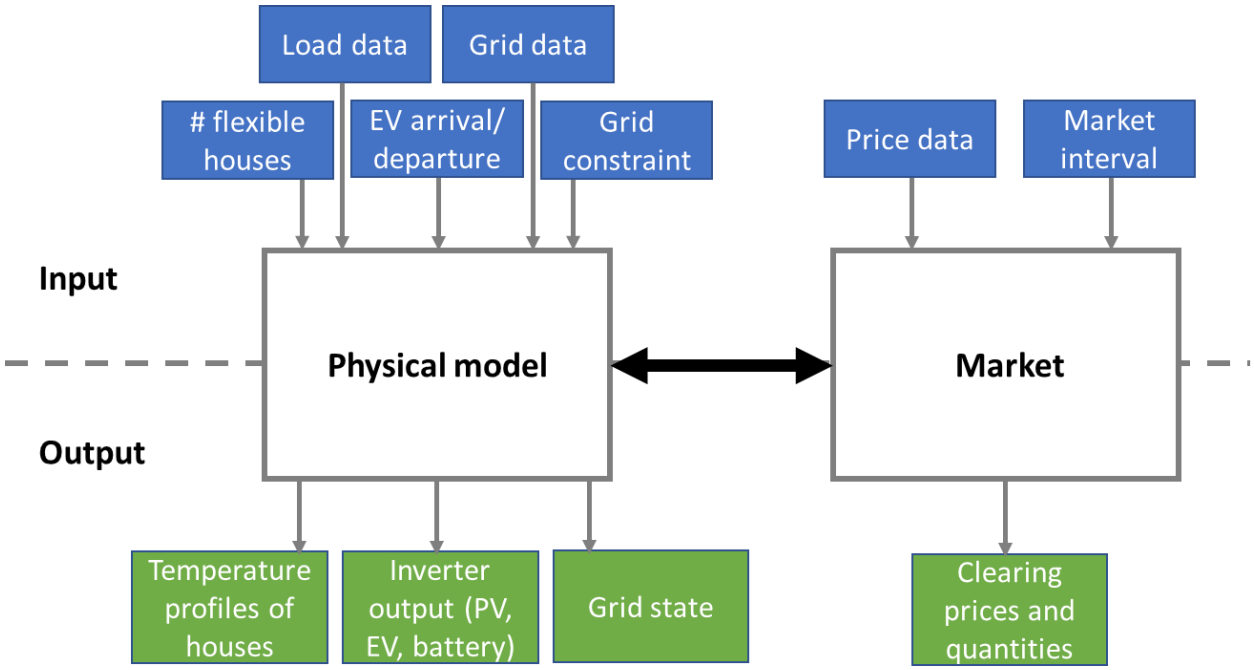

Source: SLAC National Accelerator Laboratory

APPENDIX C:
Implementation of the Home Hub and Cloud Coordinator

Introduction

The physical and platform architecture of the project was presented briefly in Chapter 2. Two main hardware/software components are:

- The **Home Hub (HH)** that aggregates resources behind-the meter and connects them to a **cloud coordinator (CC)**
- The Cloud Coordinator, that operates as the market-based controller of the feedback system.

The platform architecture is shown again in Figure C-1.

Figure C-1: Powernet with Markets System Architecture

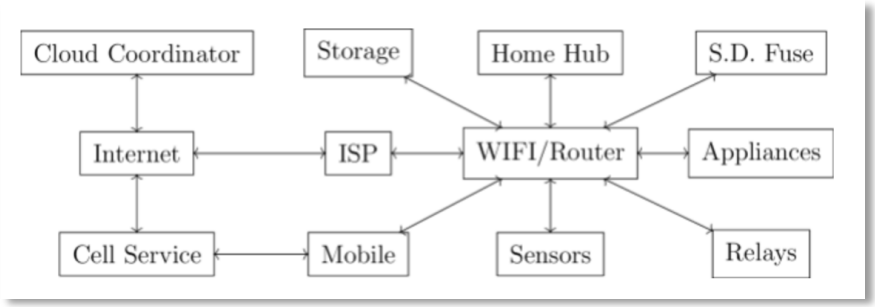

Source: SLAC National Accelerator Laboratory

This appendix presents the implementation of the main components of this system, the Home Hub and the Cloud Coordinator; giving details on the hardware, the communication strategy with other devices, the User Interface (UI) applications and the test results. These components have been deployed in the Stanford Bits and Watts Laboratory and residences in City of Fremont. In particular, in the laboratory, a full electric equipped house including multiple appliances, solar panels and battery energy storage systems has been implemented.

Home Hub

The purpose of the Home Hub (HH) is to serve as a local brain for the cloud coordinator (CC). In this project, the HH conducts real-time measurements of the main electrical variables inside the houses, as voltage, current, power at both the device level and sub-circuit level, exchanges information with the cloud coordinator, communicates the status of the appliances to the CC, accounts for user preferences and comfort and elaborates the bidding strategy for each appliance included in the home.

The project team investigated two separate approaches for the HH architecture. The first approach was having a physical hardware device with local computational power and storage to be able to control resources within the home. The second option was having all

computational power in the cloud and measurements and controls happen at the device level, leveraging the Internet of Things (IoT), capabilities. Both solutions were implemented - the physical hardware was the one chosen to run in the laboratories as it gives more flexibility for tests and fine tuning algorithms, the cloud solution is chosen for the field deployment as all equipment utilized are UL listed and do not pose any complications with the electrical code and authority having jurisdiction (AHJ).

Both solutions are presented below. In summary, the key difference is the HH's computational power, or smartness, being located at a local level — hardware — or in the cloud.

Home Hub V1.0 - Local Intelligence

The basic computing unit and communication functionalities of HH V1.0 are the following:

- Computing Unit: Device used to communicate and store information from home devices, runs HH optimization and elaborates biddings for the appliances.
- Basic Communication with CC, appliances, smart dimmer fuse, battery and solar inverter, and NEST-like devices.

Figure C-2 shows the hardware design concept for the home hub.

Figure C-2: Home Hub V1.0 Hardware Design Concept

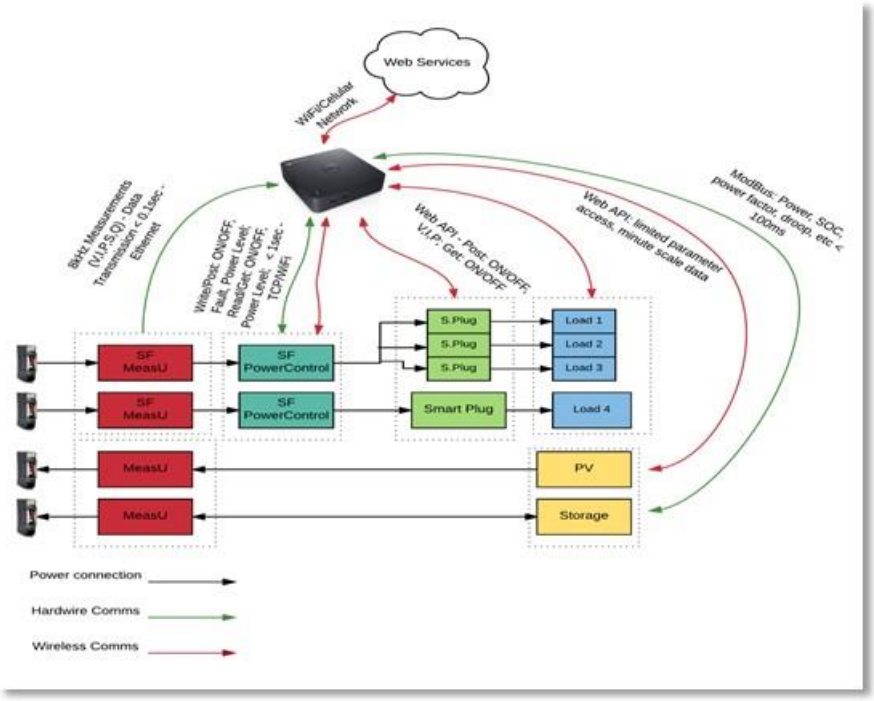

Source: SLAC National Accelerator Laboratory

Hardware

The hardware developed and assembled for the Home Hub is composed by a microprocessor running Linux operating system, which is in charge of all the processing. This microprocessor is responsible for running all the software, performing the communication with the cloud through internet, communicating with the Smart Dimmer Fuse through Ethernet and with the battery storage system through Modbus TCP, sampling appliances electrical variables and

controlling relays to turn ON/OFF appliances. This microprocessor has a built-in microcontroller that enables an integrated solution of the management capabilities and data collection and actuation.

Each unit is able to monitor up to 8 sub-circuits and turn ON/OFF 15+ devices that operates in 120V/240V and consumes up to 10A. These values are based on the relays available in the market, that are cheap and simple to interface. Increasing the amperage, the available relays in the market starts to get more expensive and harder to interface because it is needed to build specific circuitry to allow the microprocessor to interface it. However, a few important points need to be made: (1) it is not common to have a 120V residential load consuming more than 10A; (2) the 240V residential loads in the US are actually 2x 120V loads as they are, most of the time, independent (i.e. in an electric stove, one 120V "leg" is used to power the displays, lights and a few small heating coils and the other 120V "leg" is used to power the large heating coils). Therefore, controlling each individual 120V "leg" of a 240V is sufficient if the current in each leg does not go above 10A, which is the case most of the time and (3) with todays' loads that have a display to enable the user to set parameters before start the load cycle, only in a few cases make sense to control the loads individually as interrupting/enabling power to a specific load will not necessarily achieve the end goal. For example, a washing machine cannot be started, even if the relay is in the ON state unless the user goes to its display and manually set it to start. This makes the end user to still be required to turn it on therefore not enabling the full functionality and automation that the Home Hub can provide. A final point that is worth mentioning is that currently there is a misconception about the term "smart appliances". Besides the fact that they are much more expensive these appliances are limited in their "smart" capabilities as, as of now, they do not allow third-party applications to interface with them — one can only control remotely these appliances using the manufacturers proprietary application and not allowing interoperability across different platforms. However, it has been shown in the Stanford Laboratory that if the manufacturer provides an open Application Programming Interface (API), it is possible to integrate easily the appliance with the Home Hub platform and provide the full capabilities of a smart home.

The unit that is fully operational and running live in Stanford Laboratory is monitoring 8 sub-circuits and controlling 5 appliances. This unit is also leveraging the open API provided by Phillips Hue to showcase the capability of the platform to seamlessly integrate with appliances with such capabilities and thus enabling its full functionality.

Measurements for each individual appliance are taken using current transformers and being sampled at 600Hz. The data is averaged over 100 samples and sent to the Home Hub cloud hosted at Google Cloud Platform.

The storage requirement is dependent on the algorithms running locally in the Home Hub as the data should be readily available. As these algorithms are based on a longer time scale, i.e. minute level power data. The microprocessor has an internal storage of 4 Gb and the ability to be expanded with a microSD card, then the internal data storage requirements are sufficient to provide all required information for the algorithms. All the other data that is not required to be stored locally is sent to the cloud service that is hosting the Home Hub instance. This means that storage requirement is virtually limitless.

Management Capabilities

In this section, basic communication with inverter-based system and NEST-like devices are described. The latency, data storage needs and available communication capabilities to use as a benchmark were measured and reported in [1].

Communication with Inverter Based-Systems

- *Inverter types* — There are two types of inverters currently installed at Stanford University Lab: a) microinverter and b) central inverter. Microinverters are used for the solar system and they are manufactured by Enphase. Each unit has a maximum power of 250W. The central inverter is used for the storage system only and it is rated at 3.3 kW peak.
    - Each of these inverters provide the user with an account to access their information through a web portal. In addition, they also provide an Application Programming Interface (API) available to users so they are capable of integrate system's information and status to an application. Although the API provided is open, meaning that access to the system information is available, users are only allowed to read information from the system but not actively control the inverters. The microinverters, specifically, only provide the API which means that there is no option to actively set parameters and gains on the controller. The central inverter on the other hand has the capability of being interfaced through Modbus protocol. Using this protocol, reading and writing functionalities are enabled.
    - Obtaining information from the API also limits the update rate. Even if requests are made in less than one minute, the system values are only updated in 15min interval. Using Modbus instead, user can get update rates in real time.
- *Tests* – In this subsection, sample responses from devices communication capabilities are presented. For the Energy Storage and Solar systems there are multiple requests to obtain different information. A sample of one type of request for each system is presented in this subsection. These requests are then parsed, and the relevant information are saved. For further information on the type of requests that it is possible to check the manufacturer's API documentation.
    - For Enphase microinverters: https://developer.enphase.com/docs#summary
    - For Solaredge inverter: https://www.solaredge.com/sites/default/files/se_monitoring_api.pdf
- Communication Through Web API
    - Energy Storage System: The response from a HTTPS GET request to the energy storage system from SolarEdge is depicted in Figure C-3. Only real power (Watts) is reported.
    - Solar System: The response from a HTTPS GET request to the solar system from Enphase is shown in Figure C-4. Only real power (Watts) is reported.

Figure C-3: Sample Response from Energy Storage System

<Response [200]>

{u'sites': {u'count': 1, u'site': [{u'status': u'Active', u'primaryModule': {u'maximumPower': 280.0, u'modelName': u'Unsure', u'manufacturerName': u'****'}, u'peakPower': 3.3, u'lastUpdateTime': u'2017-10-29', u'notes': u'', u'uris': {u'OVERVIEW': u'/site/549572/overview', u'DETAILS': u'/site/549572/details', u'DATA_PERIOD': u'/site/549572/dataPeriod'}, u'name': u'Standford Bits & Watts 2', u'ptoDate': None, u'installationDate': u'2017-09-27', u'location': {u'city': u'Stanford', u'countryCode': u'US', u'zip': u'94305', u'country': u'United States', u'address2': u'', u'state': u'California', u'stateCode': u'CA', u'address': u'Via Ortega 473', u'timeZone': u'America/Los_Angeles'}, u'publicSettings': {u'isPublic': False}, u'type': u'Optimizers & Inverters', u'id': 549572, u'accountId': 17244}]}}

Source: SLAC National Accelerator Laboratory

Figure C-4: Sample Response from Solar System

{u'status': u'normal', u'energy_lifetime': 143544, u'current_power': 840, u'modules': 6, u'operational_at': u'2017-09-14T14:33:42-07:00', u'summary_date': u'2017-10-29T00:00:00-07:00', u'source': u'meter', u'energy_today': 1772, u'system_id': 1308698, u'last_report_at': u'2017-10-29T14:43:53-07:00', u'last_interval_end_at': u'2017-10-29T14:42:01-07:00', u'size_w': 1680}

Source: SLAC National Accelerator Laboratory

- Communication Through Modbus
    - Communication with the storage inverter through Modbus is fully operational. The type of Modbus protocol being used is the Modbus TCP. Inverter registers reading and writing capabilities are completed and integrated in the Home Hub system. For a complete information of parameters that can be read and write please refer to this documentation: https://www.solaredge.com/sites/default/files/sunspec-implementation-technical-note.pdf
    - The Powernet team decided not to follow the Sunspec Alliance documentation as it was found to have a lot of errors and "bugs" and created more problems and confusion. In addition to that, the inverter manufacturer — SolarEdge — does not provide official support for Modbus communication which complicated even further the development. Therefore, the team decided to develop its own communication protocol leveraging Python programming language and an open-source Python Modbus TCP library. The Home Hub system was able to control charging and discharging power rates, real and reactive power, phase angle etc. The latency of the communication is limited only by the speed of the Modbus and TCP protocols.
    - The requirements to successfully run the implementation is knowing the mapping between the registers address and the actual variables one wants to control, registers' size and registers possible values. With this information the full capabilities of the inverter can be accessed.

- Additionally, to the capabilities presented above, three new ones were implemented. They are, being able to read the State of Energy from the battery, read power factor (CosPhi) and changing the power factor value to any number between -1.0 and 1.0. The ability to change the power factor allows the Home Hub system to control real and reactive power, injection and/or absorption of power to/from the grid.
- With respect to the latency in the Modbus communication, theoretically, it is not possible to control Ethernet latency. However, practical solutions have been developed. Unpredictable latency occurs when there are complex routing paths and a large number of Ethernet elements. If controllers and devices are directly connected to each other, or if a private network is used, then the latency can be kept to within a specific range. A direct Ethernet link could be established between the Home Hub system and the inverter, with no other devices on the network. In a 100 Mbps Ethernet environment, the transmission latency for 100 bytes of data could be calculated as follows[1]:

$$\frac{1}{100 \times 10^6} \times 8 \frac{bits}{byte} \times (100\ bytes + 54\ bytes) = 12.32 \mu s$$

- This means the predictable latency of the Modbus communication is below $13 \mu s$.

Communication with NEST-like Devices

The category of NEST-like devices involves all equipment that is capable of communicating through the internet using an open API. This communication is mainly done through HTTPS. In the lab, there are 3 devices that leverage this communication protocol, storage and solar inverters and the Phillips Hue system.

Data Collection and Logging: Measurements from each individual appliance are taken using current transformers and being sampled at 600Hz. From each device, before data goes to the HH's ADC, it passes through a low pass filter to remove high frequency noise. Then current is sampled until a buffer of 100 samples is filled. Current RMS value is computed and assuming voltage RMS is either 120V or 240V*, depending on appliances' voltage rating real power is computed. Additionally, devices' state (ON/OFF) is also collected.

In the cases where devices provide APIs or use a different communication protocol, such as Modbus, data is collected either through HTTPS requests or leveraging communication protocols capabilities. For the specific case of the batteries, which communicate through Modbus, the variables that are being collected are: state [OFF/CHARGING/DISCHARGING], power, state of energy, CosPhi.

The HH storage management module collects data from sensors at a high resolution from the sub-circuits in the home and plugs, stores data in a database and prepares it for various uses. The HH storage management module should be able to log and store data from more than 50 data sources sampled at 1 min granularity in PostgreSQL database for local use and compressed data sent to cloud PostgreSQL database. The data types to be stored include: load (kW), status of devices and indoor air conditions (temp, humidity, etc.).

The data logging process is responsible for logging all the information into HH's internal storage and it runs in its own thread in the processor. Power values, devices states and time

when data was collected are stored in HH's internal database (sqLite). Weather information and devices that provide information through APIs are also stored in the local database.

The storage requirement is dependent on the algorithms running locally in the Home Hub as the data should be readily available. As these algorithms are based on a longer time scale, i.e. minute level power data, and the microprocessor has an internal storage of 4Gb and has the ability to be expanded with a microSD card, the internal data storage requirements are sufficient to provide all required information for the algorithms. All the other data that is not required to be stored locally is sent to the cloud service that is hosting the Home Hub instance. This means that the storage requirement is virtually limitless.

Home Hub V2.0 - Cloud Intelligence

As mentioned previously, the biggest difference between the two HH designs is where the "smartness" is located. For HH V2.0, it is located at the cloud level. One of the key reasons for this new design was the focus on the field deployments. Many counties and places require most of electrical equipment that will be monitoring power consumption, especially at the breaker panel level, to be UL listed. As this process can take several months - and even years - the project team decided to use solutions that would provide the necessary measurements. However, as most of these solutions send their measurements to their own cloud servers, the value of having a physical device in the home goes away.

Cloud

The new HH architecture is depicted in Figure C-5. All the resources within the home communicates with their own private clouds and the HH cloud gets information from them and sets values to them, if they accept control commands. This solution has many advantages over HH V1.0 but the most important one is scalability. However, the (partial) drawback is the reliance of this solution in internet communication. The reason why it is a partial drawback is that the first version communication with many of the home devices was also done through the cloud and thus, only the physically hardwired devices would be able to provide information/measurements and within this set a smaller number of devices would have control capabilities when offline.

Figure C-5: Home Hub V2.0 Hardware Design Concept

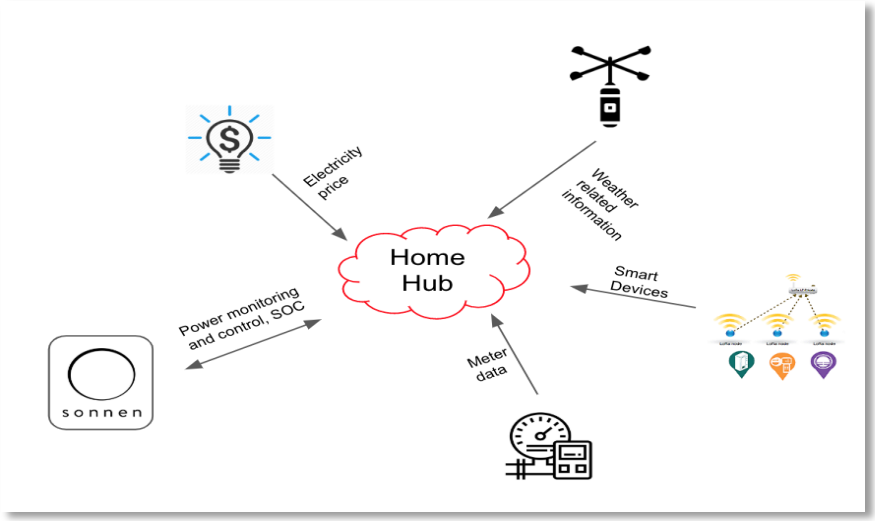

Source: SLAC National Accelerator Laboratory

Management Capabilities

As all the communication capabilities from HH V2.0 was inherited from HH V1.0 wireless communication/management, please refer to this section in HH V1.0 for more information.

Data Collection and Logging

In HH V2.0 the meter level data, at the sub-circuit level, is measured using an E-Gauge device. This device provides minute level metered graded measurements of power (current, voltage and phase angle). Measurements are uploaded to E-Gauge cloud server and is accessible to third-party solution providers given owner's permission. Figure C-6 shows the installation of the E-Gauge monitoring system.

Figure C-6: E-Gauge Installation in Participant Home

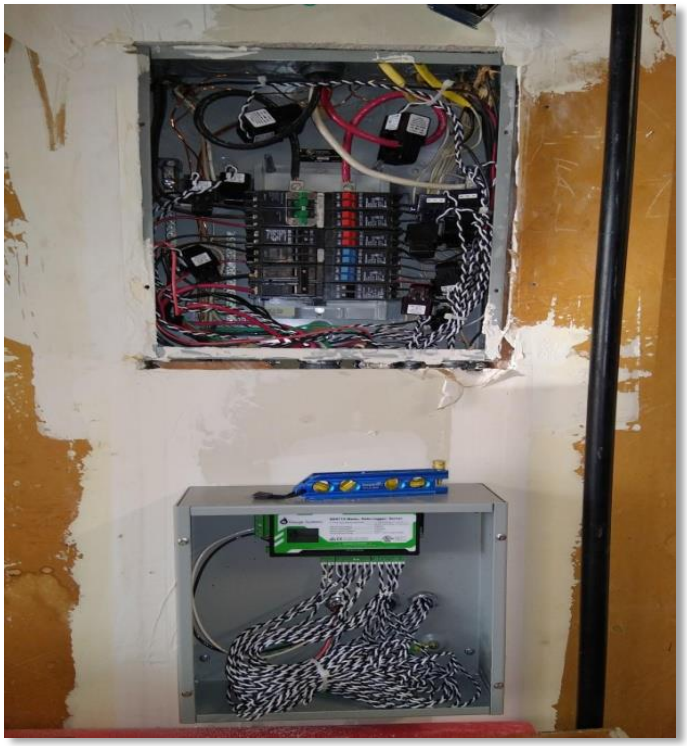

Source: SLAC National Accelerator Laboratory

Tests

Tests were conducted to show the Powernet platform capability of scaling this solution. The system was setup in such a way that simulation and real hardware connected in the labs were able to respond to signals coming from the cloud coordinator. Additionally, to the simulation, 10 lab homes with 10 devices each were actively participating in the experiment. The goal was to show the capability of the platform to operate 10 real homes, where all included HH V1.0, plus the simulated homes while responding to regulation signal. Figure C-7 shows the homes implemented in the laboratory used for the experiment.

Figure C-7: Laboratory Tests

Top image shows the main Laboratory home with big appliances and bottom image shows 9 "temporary" homes with different set of loads

Source: SLAC National Accelerator Laboratory

Figure C-8 shows the aggregated regulation signal including the simulation and the homes in the laboratory and the total tracking error was 6.365%.

Figure C-8: Aggregated Regulation Signal Tracking

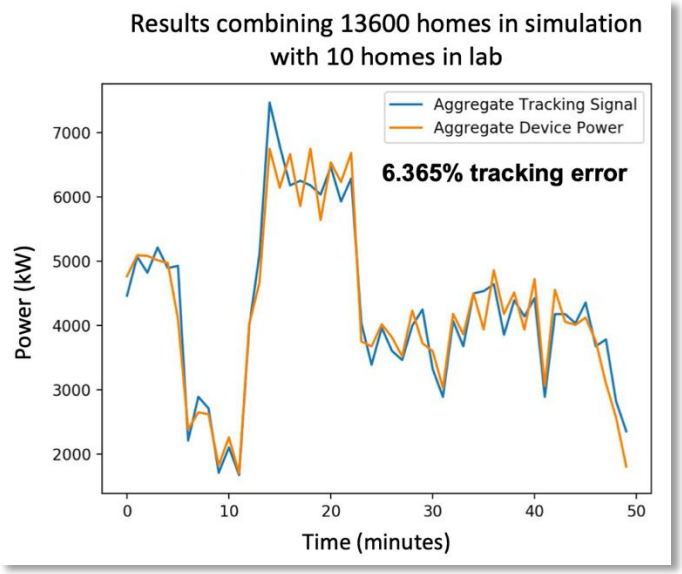

Source: SLAC National Accelerator Laboratory

Deployment in Residences in the City of Fremont

To test the feasibility of the cloud-based control strategy and validate the operation of the communication and information exchange among the DERs and the cloud, a small deployment of the hardware and the software was conducted in residences in the City of Fremont. The pool of houses selected was very diverse, including some devices such as, solar panels, electric vehicles and central air conditioning. The home hub controller for these houses was implemented in a cloud server.

Twelve homes where equipped with a monitor system to measure the current (real power) in the main sub-circuits of each home. Additionally, smart thermostats where installed in those houses to be able to control the HVAC system. Four houses where equipped with a Sonnen storage system of 10 kWh capacity. Some of the measurement results are shown in the following figures. Figure C-9 shows the battery system installed and the battery charge and discharge for a period of 24-hours. Figure C-10 shows the main currents measured for different sub-circuits in one of the residences. This measurement is important to estimate the power of the aggregated unresponsive load needed to define the demand curve in the auction-market.

Figure C-9: Installed Storage Unit and Battery State of Charge (SOC)

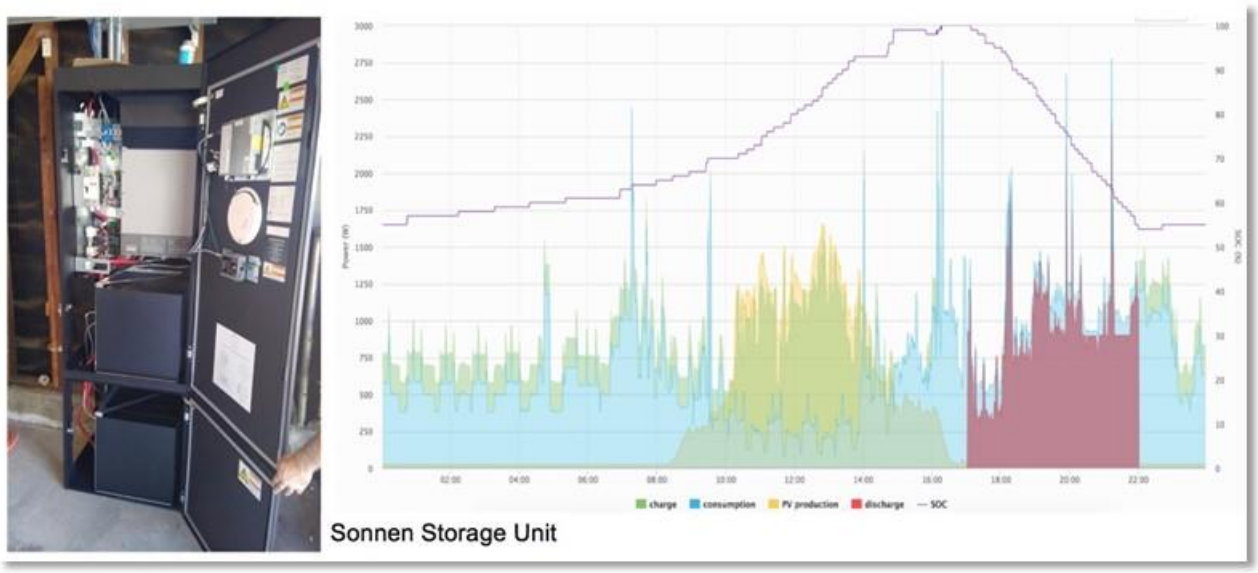

Source: SLAC National Accelerator Laboratory

Figure C-10: Installed Monitoring Unit and Sub-Circuit Currents

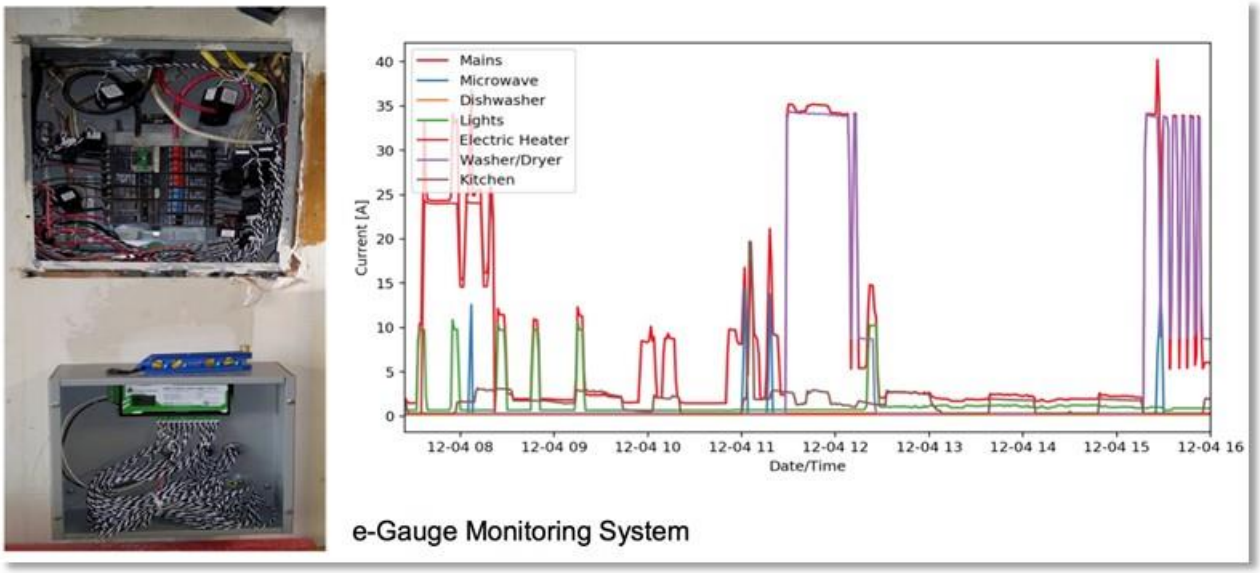

Source: SLAC National Accelerator Laboratory

Preliminary Web User Interface

The preliminary web User Interface (UI) is hosted in the Google Cloud Platform (GCP) and the user can access it and get information from all the appliances/sub-circuits being monitored, their status — whether it's ON/OFF, charging/discharging — and power consumption. In addition, there is a preliminary version of the visualization that is being used for demoing the lab capabilities. Figures C-11 – C-15 show these features.

Figure C-11: PV and Storage Status

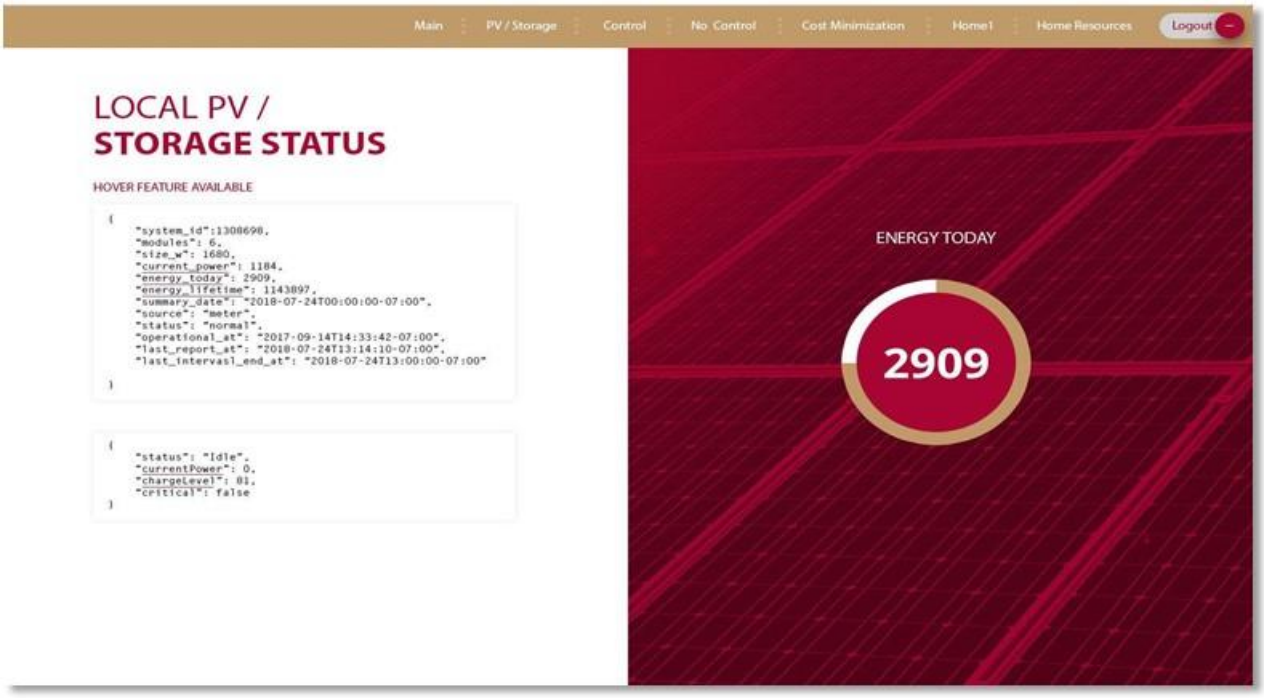

Source: SLAC National Accelerator Laboratory

Figure C-12: Changing Device Status

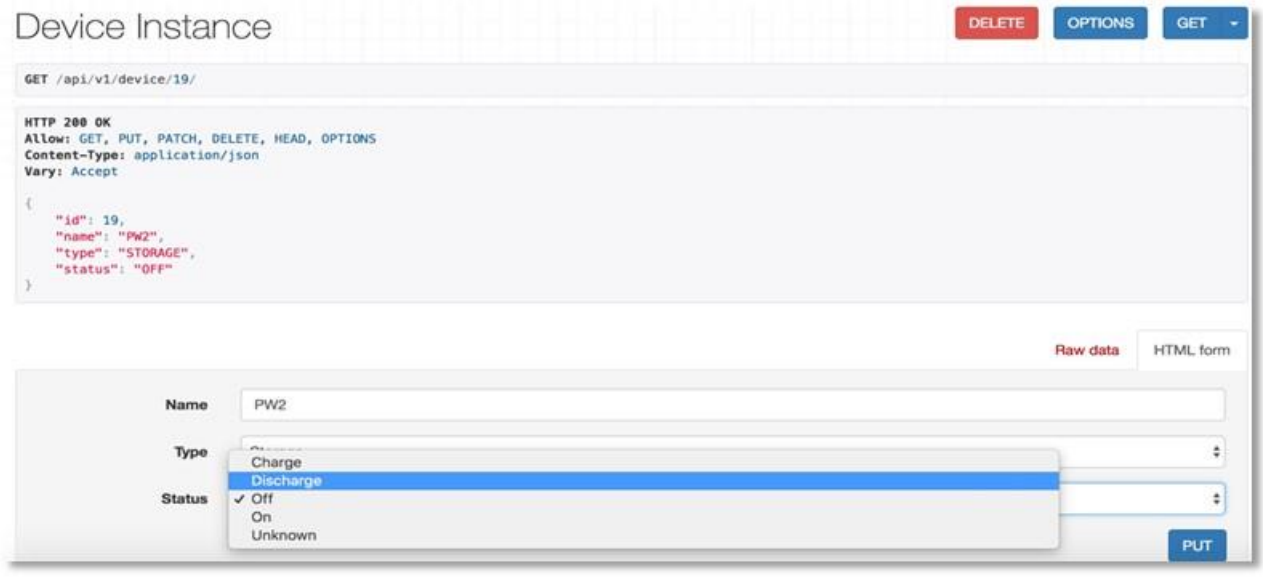

Source: SLAC National Accelerator Laboratory

Figure C-13: Device Status

Devices Being Monitored and Their Status.

Devices Current Measurements and Timestamp.

Source: SLAC National Accelerator Laboratory

Figure C-14: Home Web Visualization with Statistics

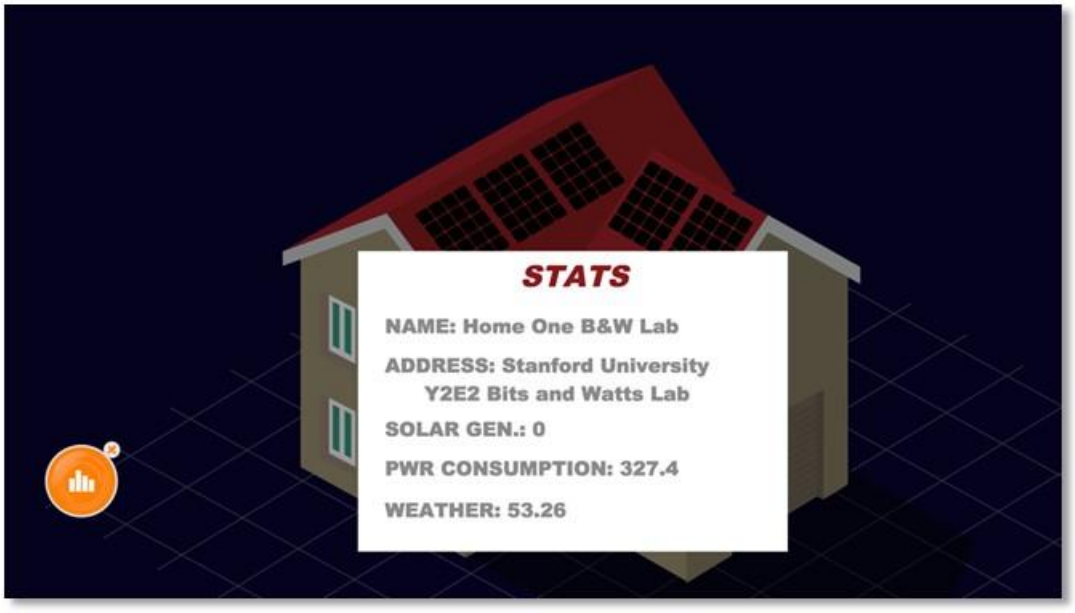

Source: SLAC National Accelerator Laboratory

Figure C-15: Home Visualization Showing Appliance Power Consumption and ON/OFF Controls

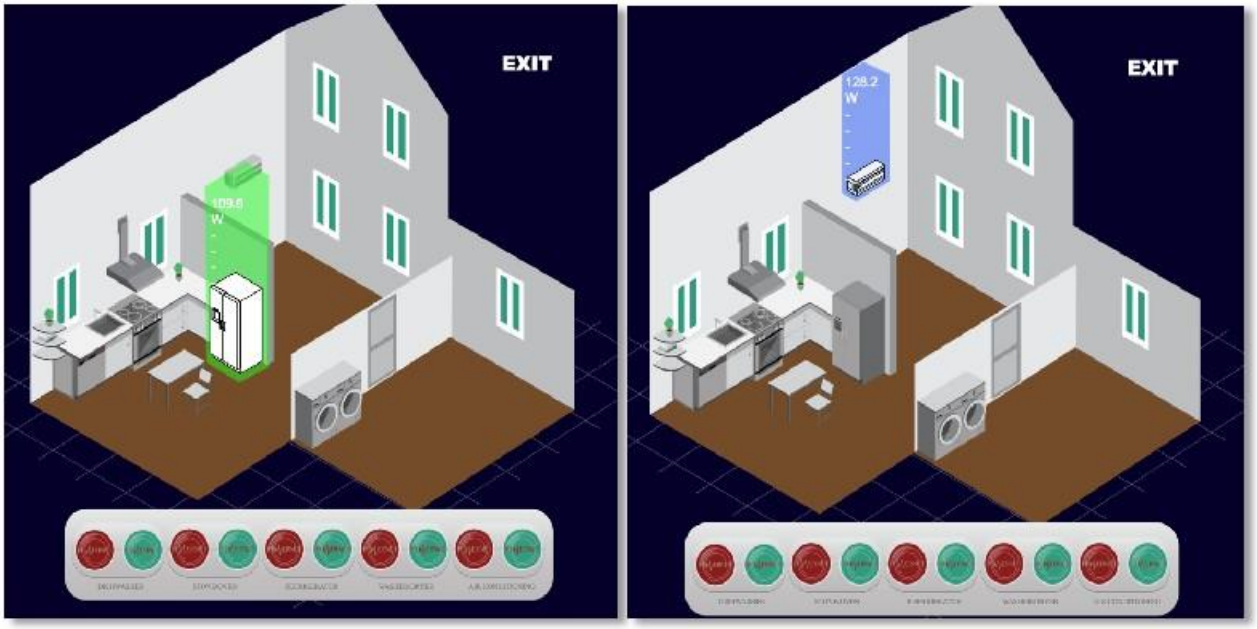

Source: SLAC National Accelerator Laboratory

Preliminary Mobile User Interface

The mobile application being developed is also intended to give the user the ability to see the same information, control the devices and set preferences as in the web application.

The Application is currently talking to the Home-hub through REST API. It fetches the operational status and power consumption data for connected smart Appliances. We collect this data for every 10 seconds. Along with fetching the data, the researchers enable users to control the appliances. This would be a single pane of glass-view for monitoring and controlling all the connected smart appliances. Some pictures demonstrating these capabilities are shown in Figures C-16 and C-17.

Figure C-16: Navigation Drawer Displaying all the Appliances.

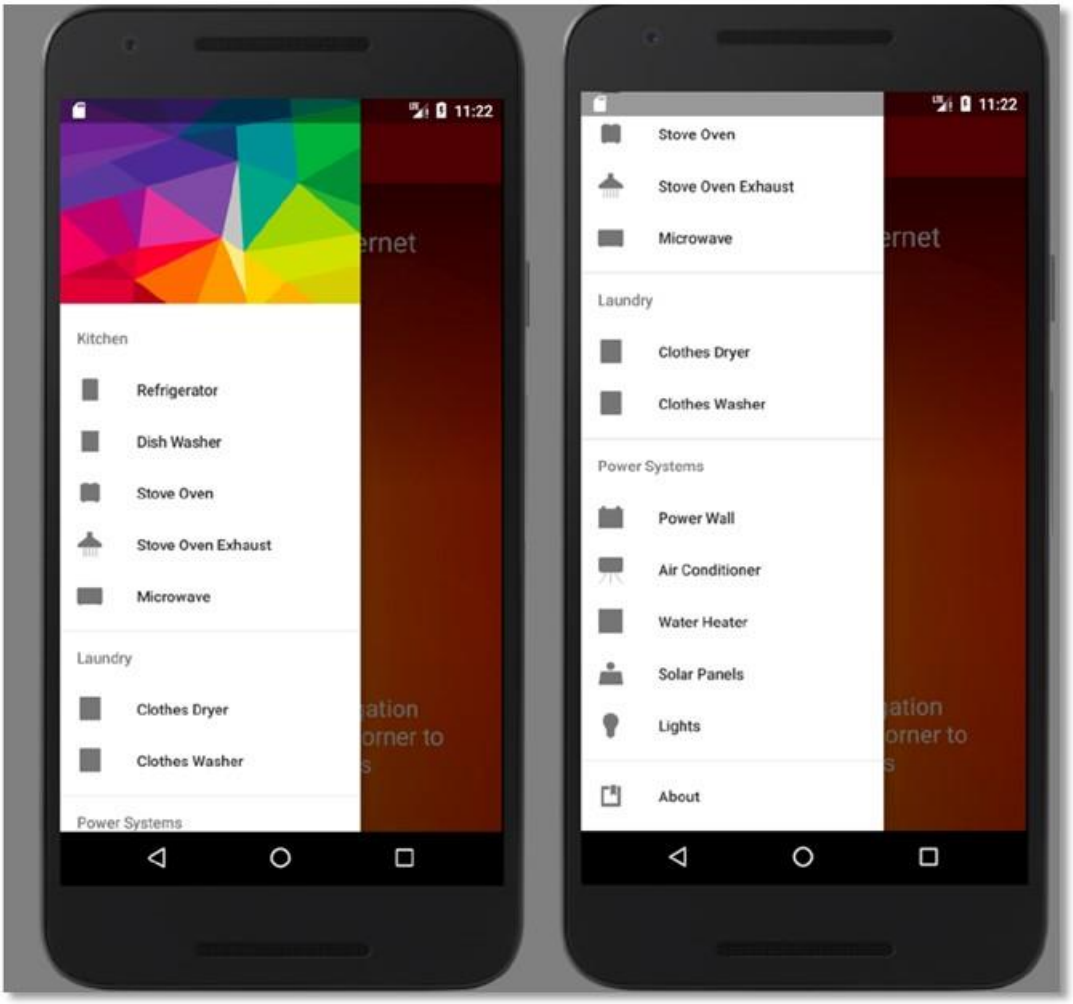

Source: SLAC National Accelerator Laboratory

Figure C-17: Display of Individual Appliances

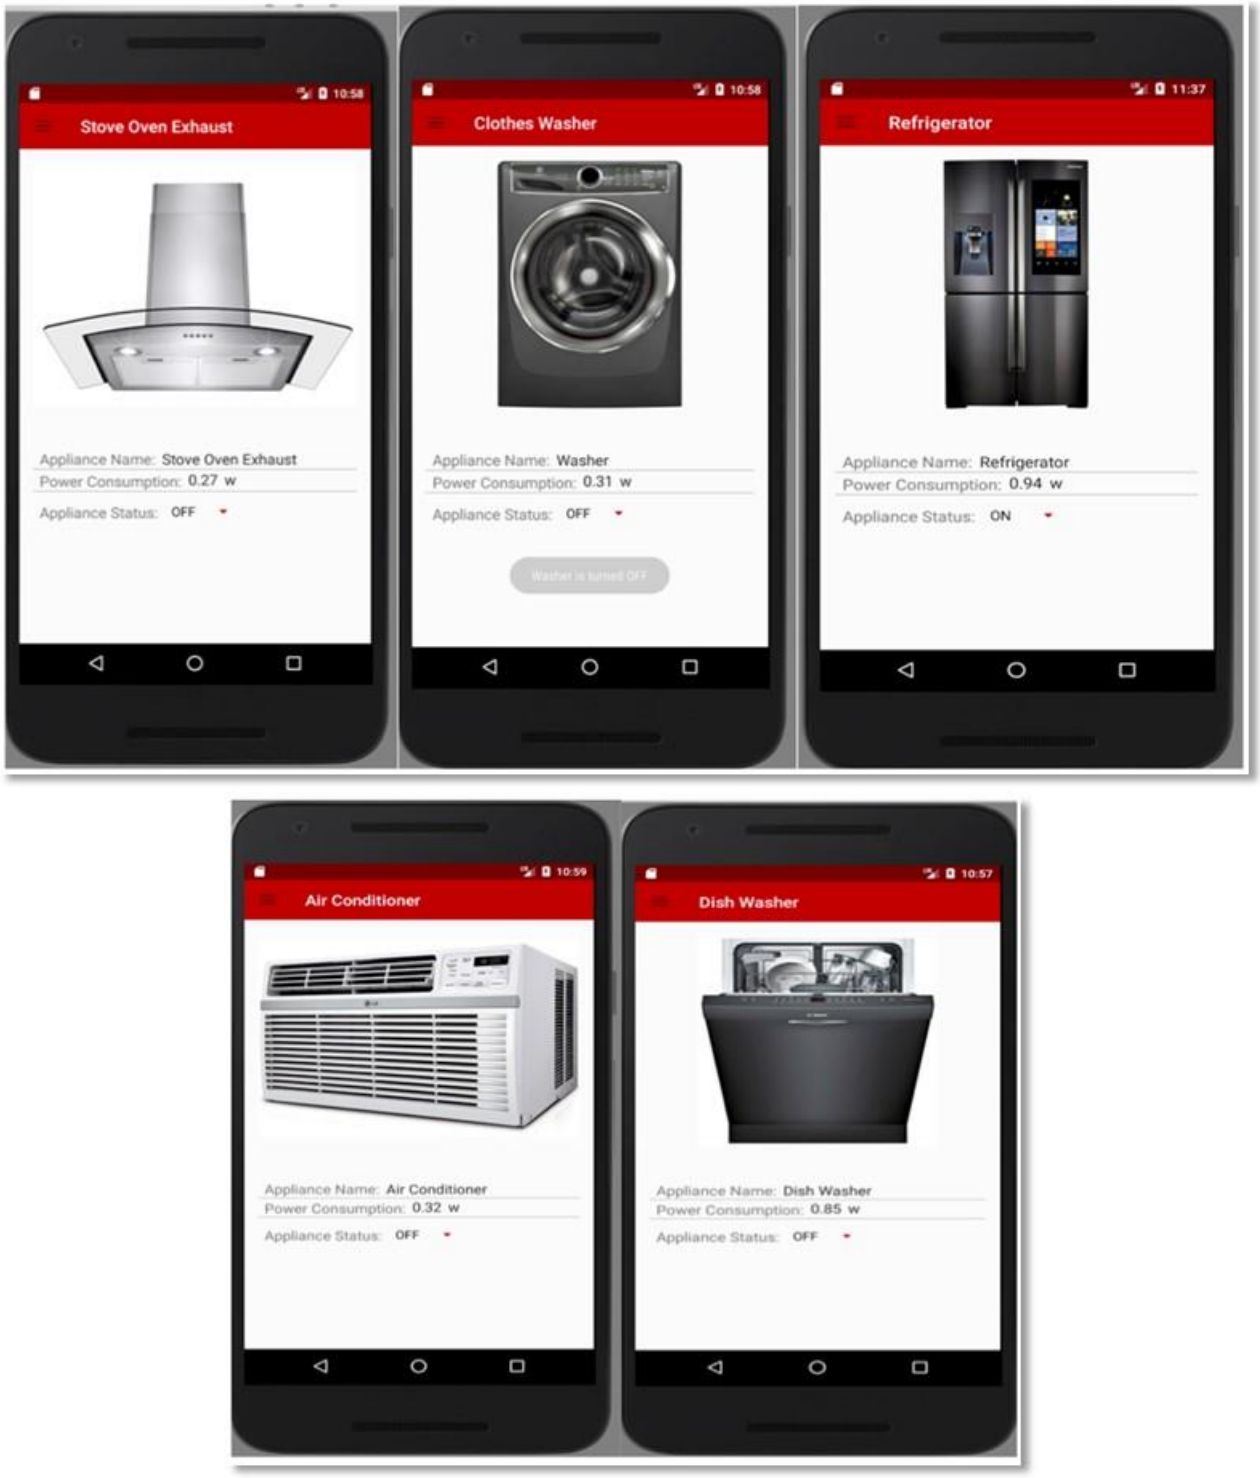

Source: SLAC National Accelerator Laboratory

Cloud Coordinator

The cloud coordinator (CC) is an aggregation platform that integrates the resources in the field/homes as well as other information such as utility rates and signals, markets and weather information. Once data is aggregated the cloud coordinator runs different algorithms based on

the tasks it needs to perform — forecasting, market, cost minimization, etc. Figure C-18 shows a schematic on how the system is designed.

As it has been mentioned before, the CC is a web platform that runs in Google Cloud Platform (GCP). Within GCP a Django web framework was developed to enable all the environment to exist. This web framework includes the application layer where the algorithms run, where user accounts get created, resources can be accessed, and user interface is managed. It also has a relational database that stores all the information collected from the field, and algorithms results.

The system as designed has a minimum of two servers running all the time to support system operation. Additionally, it is also setup in auto-scale mode so it can grow the number of instances/servers supporting the operation in case that a sudden growth in the number of requests or computational power is a requirement of the algorithms. The system also keeps a log of all the information in its logging session that can be filtered for any keywords and also send an email alert to the administrators in case errors occur. Lastly, all the code is based on a Docker image and has continuous integration and continuous delivery. This brings many benefits such as: a) whenever the code gets update on GitHub it automatically checks for inconsistencies and if it passes the code gets deployed in the production server, and b) the image can be downloadable from GitHub and deployed in any cloud service provider and the system will deploy on its own.

Figure C-18: Cloud Coordinator Architecture Design

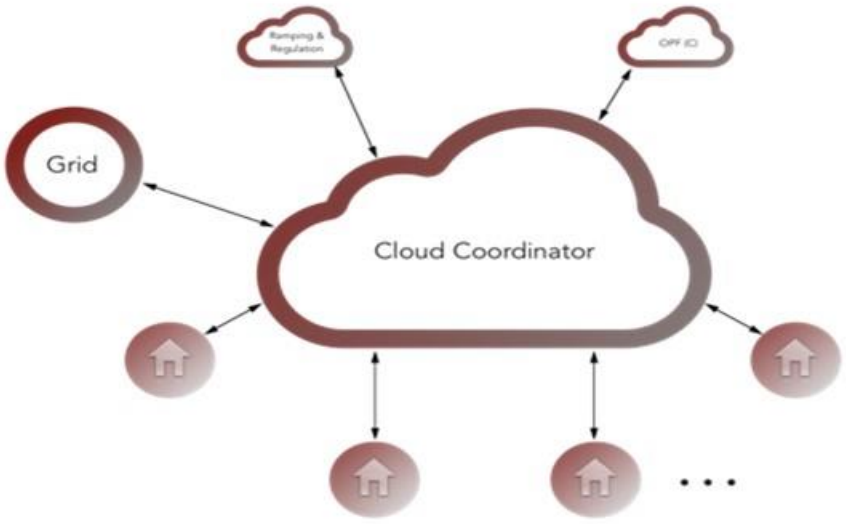

Source: SLAC National Accelerator Laboratory

All homes are connected with Django web framework and their information (data and control signals) are routed through Django Rest Framework (DRF) application programming interface, and it is transparent to the end user. Additionally, the only difference between the implementation of HHV1.0 and HHV2.0 is in the former there is a single point of communication to collect all the data and in the later there are multiple points to collect data for homes. Additional drivers were developed and standardized to make sure the

implementation is replicable. Therefore, the system can still be used for both versions of the HH without deprecation of earlier versions.

Communication Pipeline

The focus of this section is to provide information on the main communication links, starting from the system operator up to the HHs. Specifically, this will illustrate and do a breakdown of the communication links. Figure C-19 shows a breakdown of the communication links.

Figure C-19: Breakdown of the Communication Path

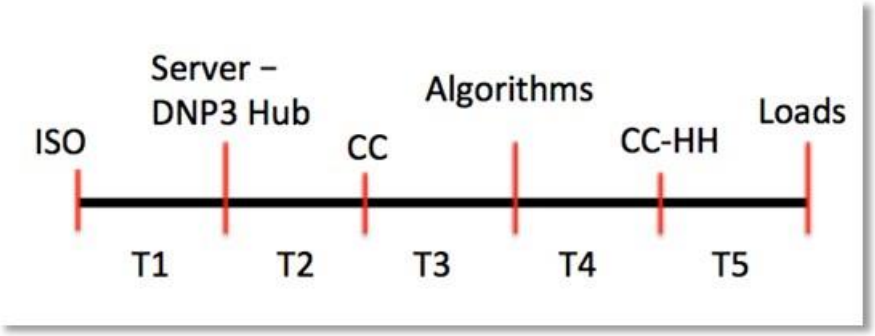

Source: SLAC National Accelerator Laboratory

T1 is the time it takes from the grid operator to "announce" the new set point (either for a ramp or regulation). This data is transmitted through DNP3 protocol to the entities participating in the ancillary service markets. These entities have a device, certified by the ISO, that takes care of this communication. This device, after reading the signal (this signal is mainly a timestamp of when it was updated with the set-point value in kW), upload it to a local server where the agents participating in these markets can pull the data from and start their algorithms to make sure they meet the targets. This time interval is *T2*. *T3* is the time it takes to run the algorithm (ramping or regulation). *T4* is the time for the CC to send this message to the HHs. Finally, *T5* is the time it takes for the HH to actually respond to those signals. These intervals have been measured using a prototype of the Home Hub and a personal computer to emulate the Cloud Coordinator.

Measurement and Evaluation of Interval T1

The intervals depicted in Figure C-17 have been measured based on the hardware and software comprising the HH. T1 is computed based on the time it takes for the entities participating in the ancillary services to retrieve the information from ISO through DNP3. The graph below shows the time T1 for 84 different calls. Results from the measurements using the Home Hub prototype are depicted in Figure C-20. The average time and variance are 0.135 seconds and 0.011 seconds, respectively.

Figure C-20: Sample of T1 times

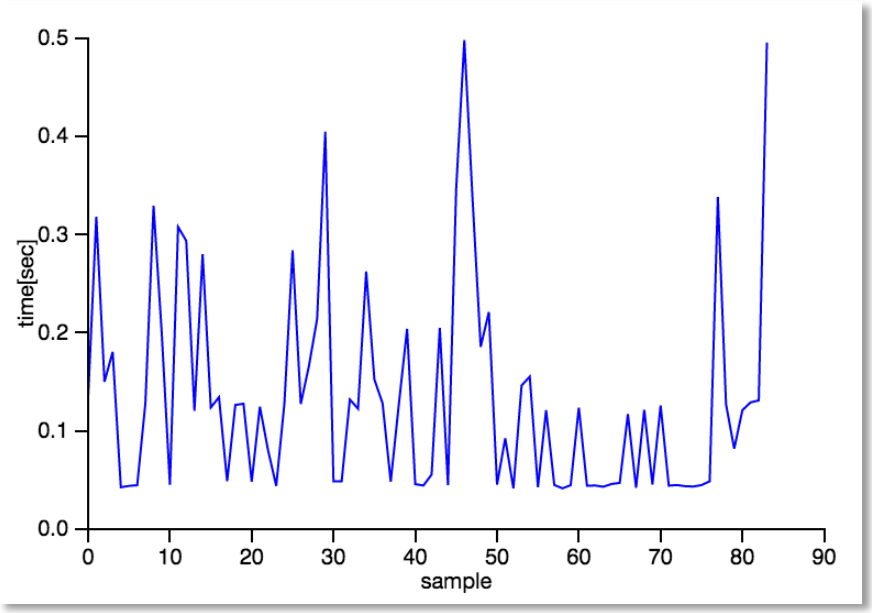

Source: SLAC National Accelerator Laboratory

Measurement and Evaluation of Interval T2

T2 is calculated based on the time it takes to perform an HTTPS request to the server hosting the grid signals. The tests were conducted to stress the communication link for the worst case possible. This happens when many entities request information from the server at virtually the same time.

The test parameters to stress system are:

- Standard load test, which performs an HTTPS GET request to a default resource on the cloud server instance.
- Response payload was standardized at 2.2 KB
- Requests grow linearly by 10 requests per second (RPS) and starts at 10.
- Total time tested ~ 10 minutes
- Total requests sent 304 K
- Total failed requests were 1.8 K
- Total successful requests handled 99.41%
- Total failed requests 0.59%

Caveat: During this load test it was deliberately decided to NOT auto scale the cloud server. The research team wanted to characterize the maximum load it could handle by a single instance. Auto scaling the cloud server and setting a fairly conservative threshold theoretically is only limited by the amount of money one wants to spend on scaling the cloud infrastructure.

Figure C-19 shows all requests during the load test plotted against its respective latency. However, there are some early outliers that are likely due to the response not being initially cached. Further improvement in response time and overall request latency can be achieved by an explicit caching policy.

Based on the data depicted in Figure C-21, the distribution of the communication latency is shown in Figure C-22, while the Table C-1 defines the percentiles. The 95% percentile for communication latency for successful HTTPS requests, is around 0.282 seconds.

Figure C-21: Per Request Distribution by Millisecond

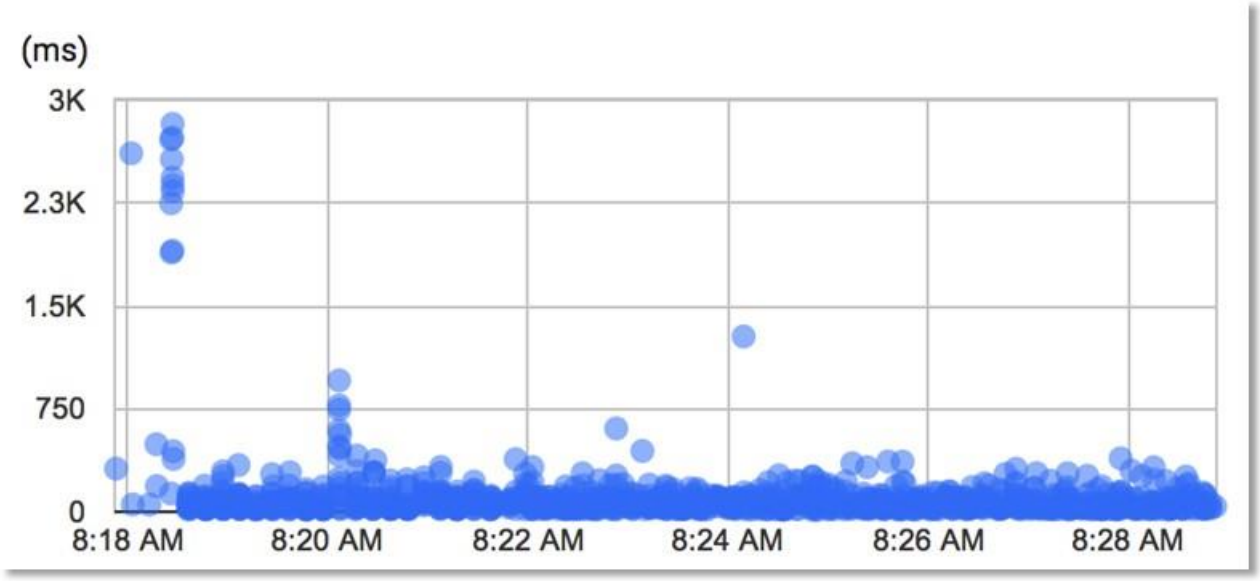

Source: SLAC National Accelerator Laboratory

Figure C-22: Latency Density Distribution

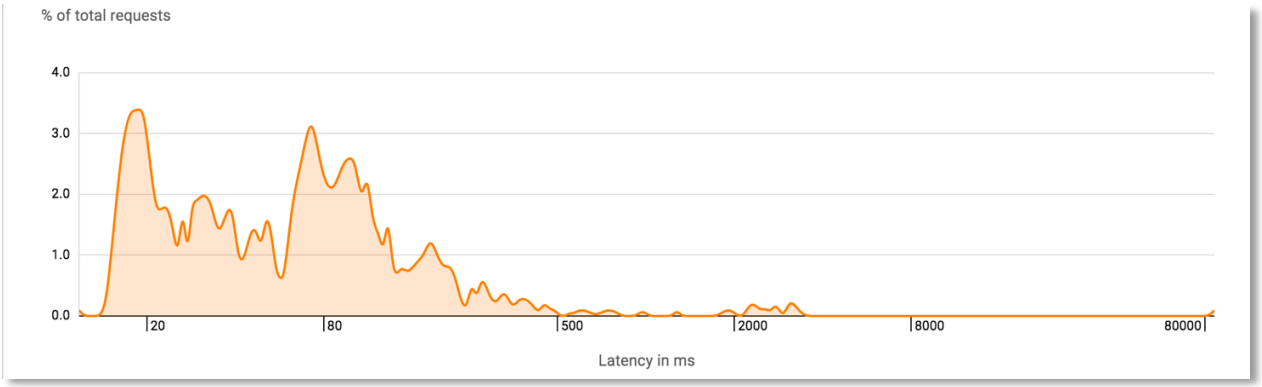

Source: SLAC National Accelerator Laboratory

The performance of the communications was tested by increasing the request-per-second (RPS) in small steps. At approximately a load of 415 RPS, the single server instance began to experience minimal (< 1%) service degradation. At this point, packets were dropped, and failure responses were sent. At approximately 600 RPS the researchers saw severe service degradation and failures began to exceed 1%.

Table C-1: Latency Percentile Breakdown

Request Percentile	Latency
25%	25 ms
50%	54 ms
90%	190 ms
95%	282 ms
98%	608 ms
99%	2,443 ms

Source: SLAC National Accelerator Laboratory

Measurement and Evaluation of Interval T3

T3 is the time it takes for the algorithms to run in the cloud, after getting the signals from the grid. Tests were conducted for the Ramping algorithm. Running on a personal computer, to implement the CC, the global control algorithm takes about 3:30 minutes to run for a 107-bus network. This global controller computes 72 times steps but it only uses the first 24. The current state of the ramping algorithm is not optimized for speed. However, the speed could be reduced by half in the same machine, if speed is optimized. Two points are worth mentioning. The first is the algorithm's speed is related to the size of the network and thus, for larger networks, the time to find a solution is going to increase assuming the same machine is used. The second point is the computer performance. All the CC algorithms are going to be run in a cloud service. Therefore, computational power can be increased or decreased by the user or automatically by the cloud servicers management system, which ultimately can increase the algorithm's speed statically or dynamically.

Measurement and Evaluation of Interval T4

T4 is the time it takes for the CC to send the set-points to HHs. This is analogous to *T2* in which the communication is made through HTTPS. Please, refer to *T2* for latency and drop packets numbers as the communication link is stressed with many requests per second.

Measurement and Evaluation of Interval T5

T5 is the time it takes for the HH to respond to the set-points received. This time is related to the algorithms running in the HH and the time it takes for the devices to respond. The algorithm running in the HH is mainly the real time load scheduler which is a sort algorithm based on load priority. These algorithms are very efficient (in general $n \log n$) and for a small number of loads, time is not an issue. Depending on the machine it is running it can be a few seconds in less powerful computers to less than a second for a standard and more powerful units. The algorithm that is going to be used for this task is still under development and thus a precise timing is going to be communicated in the future "Markets Mechanisms" report and updated in this document

Smart plug loads use HTTPS communication and thus the same time reported for *T2* and *T4* applies in this case. Batteries, specifically the Powerwall unit, have a discharge rate of 3.3 kW with communication through Modbus with a speed of 19.2 kbps. For a message with a few bytes, the communication should take no longer than 10ms.

The total time for ramping and regulation are the sum of *T1*, *T2*, $T3_{ramp/reg.}$, *T4* and *T5*.