PNNL-24116

Pacific Northwest NATIONAL LABORATORY

Proudly Operated by Battelle Since 1965

Transactive Control and Coordination of Distributed Assets for Ancillary Services: Controls, Markets and Simulations

February 2015

K Subbarao
K Kalsi
E Mayhorn
J Fuller
J Lian

U.S. DEPARTMENT OF ENERGY

Prepared for the U.S. Department of Energy under Contract DE-AC05-76RL01830

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062;
ph: (865) 576-8401
fax: (865) 576-5728
email: reports@adonis.osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312
ph: (800) 553-NTIS (6847)
email: orders@ntis.gov <http://www.ntis.gov/about/form.aspx>
Online ordering: http://www.ntis.gov

This document was printed on recycled paper.
(8/2010)

PNNL-24116

Transactive Control and Coordination of Distributed Assets for Ancillary Services: Controls, Markets and Simulations

K Subbarao JC Fuller
K Kalsi J Lian
E Mayhorn

February 2015

Prepared for
the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99352

Executive Summary

Today ancillary services are supplied by power plants that are not producing power at their most efficient output levels, tying up expensive capital investment, wasting fuel and increasing wear and tear from continually adjusting their output in response to the immediate balancing needs of the grid. The need for ancillary services is projected to increase substantially as renewable generation penetrates to greater levels of power plant capacity and beyond. Distributed energy resources (DERs), such as distributed generation, storage and responsive load, can provide equivalent services by adjusting their power demand and output rather than relying on centralized power plant output adjustments.

A transactive control approach allows DERs, which are largely owned and operated by consumers and third-parties, to be seamlessly integrated with the operation of traditional grid infrastructure and coordinated to produce a smooth, stable, predictable response as required by grid operators. As compared to traditional direct load control (DLC) approaches, transactive control solutions emphasize a decentralized approach, where consumer and third-party decision-making is kept local and private, to manage such large-scale deployments with (eventually) hundreds of millions of individual assets while maintaining free will on the part of consumers. While transactive approaches have been successfully demonstrated in engaging demand response (DR) and DERs for peak reduction and reducing wholesale energy costs (GridWise® Olympic Peninsula Demonstration Project and gridSMART® Demonstration Project), the potential for transactive control and coordination is far greater. Ancillary service markets are one potential avenue for increasing revenue generated by demand-side resources, increasing the overall value of DERs.

The purpose of this project is to develop a plausible transactive framework for DER participation in a regulation market. This document focuses on the methodology for creating a transactive-based regulation market, using one class of end-use devices as an example. The system contains two parts, one for acquiring resources at a longer timescale and a second for controlling the devices in a distributed manner at much shorter timescales. The first is based on a formal double-auction market where every five minutes each device bids the amount of resource it is able to provide and the minimum price that it would accept to provide that resource. The bid price is determined by the current state of the device and the willingness of the consumer to participate. The market system collects and orders the bids by price, and then determines a cleared price to meet the level of regulation needed. It broadcasts the cleared price to the devices, which results in contracting the services of the least cost resources. By contract, the devices that cleared the market are now engaged for the next five-minute market period. They are part of a distributed control system that allows them to respond at four-second intervals to a broadcasted regulation signal. The approach also limits the number of times devices can cycle between states (say on to off) in a given amount of time to protect the equipment life.

GridLAB-D, an open-source distribution system simulator, is used to determine the efficacy of the method. An existing retail double-auction system was adapted to meet the needs of the designed control system. A control agent was also created to interface with a programmable thermostat inside of house model, using local information to formulate and transmit a bid to the auction, determine whether it is cleared to participate in the ancillary service call, and apply the control system algorithm when cleared to do so. A diverse 5,000-home model was used to evaluate the performance of the distributed control system under a variety of constraints.

The devices were tested against Regulation A (low frequency) and Regulation D (high frequency) signals using PJM's historical data [25], as shown in Figure i. Green lines represent the desired load level of the controlled devices to properly track the regulation signal; blue lines represent the actual load level of the devices participating. In both cases, the aggregate group of HVAC systems (sometimes totaling greater than 600 homes to less than 25) is capable of tracking the desired load level requested by the regulation signals. Note that in cases where the blue line is not visible, it is behind the green line.

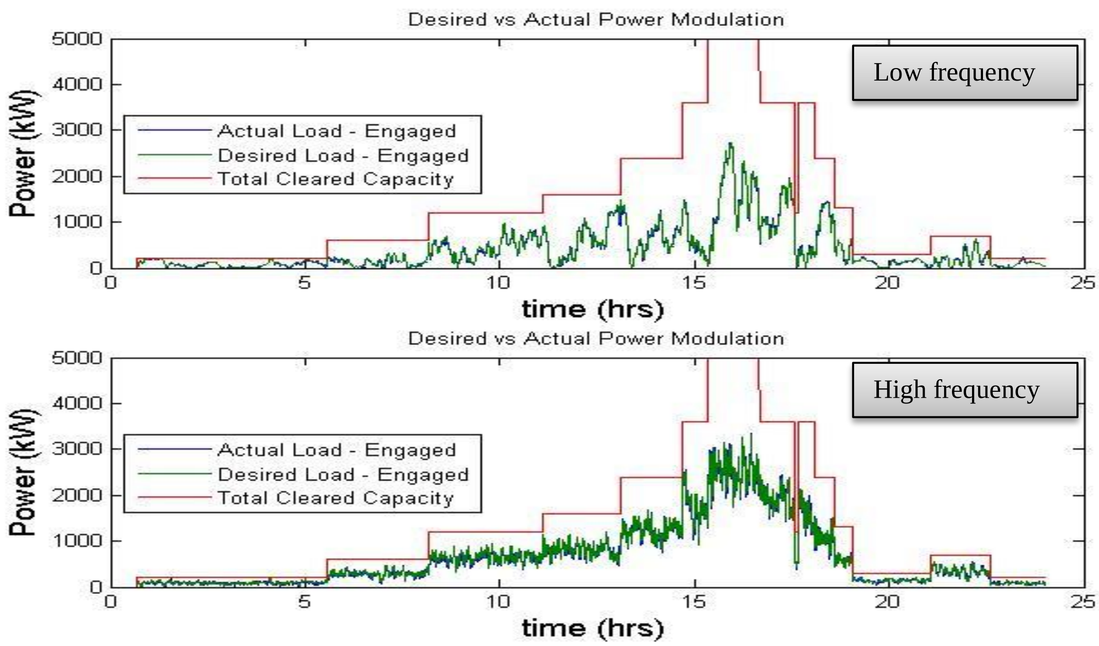

Figure i: Comparison of the desired versus actual load using different frequency components.

Other variants were tested to evaluate the system in different operating conditions. This included the evaluation of the impact of market periods – shorter market periods equated to better control and greater resource availability, but there is a trade-off in terms of communication and data requirements. Population size was investigated – populations as low as 500 homes were able to adequately participate in this system on a warm day (high resource availability), but greater numbers of devices equated to a stronger correlation with the regulation signal. "Forced" and "unforced" control modes were also tested to understand how much integration with thermostats would be needed – could the controller be an external application that only interacted with the thermostat (e.g., a Home Energy Manager), or did it require embedding it in the thermostat itself. Forced control means that the device's natural state changes are overridden to enforce the device controller's desired state while unforced *assumes* the device behaves according to its natural state changes, operating normally except to the controller's request for a state change. In Figure ii, it can be seen that forced control tracks the signal significantly better that unforced control, especially when trying to track the low frequency signal. This implies certain design decisions and deployment strategies.

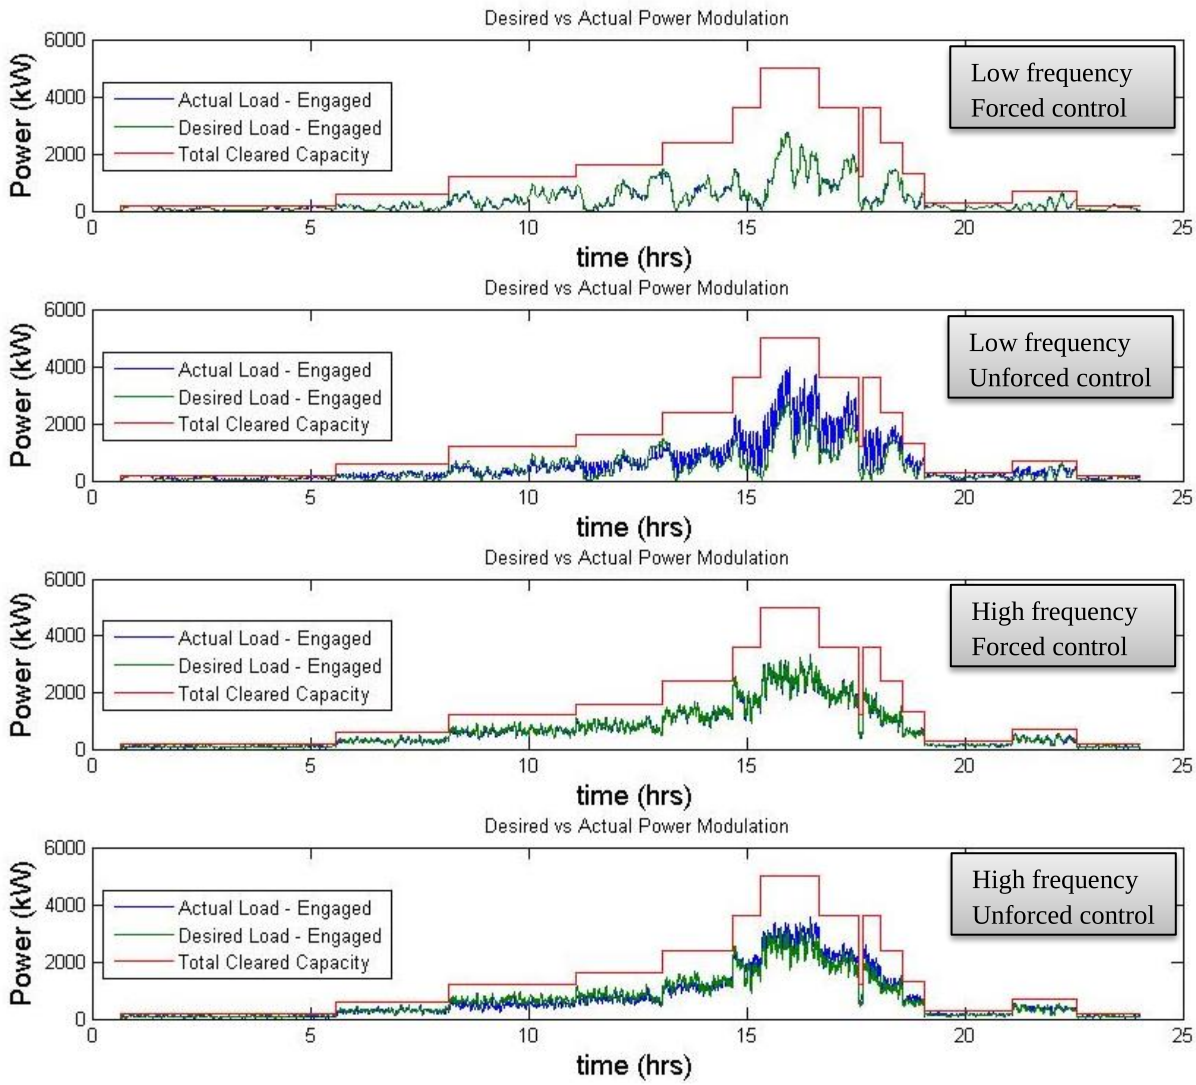

Figure ii: Comparison of the desired versus the actual load using forced and unforced control.

Overall, the simulations show that the hierarchical market and control system could be used to engage and reward DERs in ancillary service markets at longer timescales while applying distributed controls for tracking the regulation signal at much shorter timescales.

Measurement and Verification systems are widely used with traditional resources (e.g., industrial loads or generators) to validate their performance in respect to the requested service, however, few examples exist for addressing the performance of an aggregation of small devices (such as residential HVACs). A system operator might want to assess the performance of an aggregator, or an aggregator may want to assess the performance of a residence, an individual device, or a subsystem. An example M&V system is created at the device level using device state information, ensuring that individual devices are performing as desired by the aggregator control system. The required state changes of the devices are stochastic in nature, so the devices must only track the regulation system in a probabilistic manner over a given amount of time (e.g., days). An aggregator M&V method, also using the device states in the subsequent market period, is used to detect whether the population of devices is delivering as much "mileage" as the regulation system requires. This does not determine whether the population tracks the four-second signal, but rather if it produced the correct amount of up- and down-movement within a five-minute period. Finally, an example utility-level M&V system, using substation four-second SCADA measurements, is used to determine whether the load is properly moving up and down in response to the regulation signal. This technique uses relatively simple signal processing to detect a relatively small change in load behavior within a relatively noisy measurement, as other loads on the system are continuing to change states in an uncontrolled manner. All methods showed some success, but future work is required to improve on the M&V system.

Future work includes development of bidding and control methods for devices capable of more than one state change per market cycle (e.g., clothes dryer) and those that have continuous setpoints (e.g., variable speed fan), and the requisite aggregator functionality to support this. Work is needed to understand the aggregator interaction with wholesale markets, including methods for how the aggregator constructs its bid into a wholesale market, how much risk to assign to the bid, and how much extra resource should be acquired to offset the risk. New M&V methods should be explored, particularly when including a diverse set of devices and stakeholders. Impacts of communication failures should be investigated and used to determine whether the control mechanism should be more robust or what the minimum communication requirements are for the system. Finally, as this system is designed to be flexible to new entrants, standards and interoperability needs should be addressed.

Table of Contents

Table of Figures

Acronyms and Abbreviations

ACE	Area Control Error
AEP	American Electric Power
ANSI	American National Standards Institute
CAISO	California Independent System Operator
CEA	Consumer Electronics Association
DLC	Direct Load Control
DOE-OE	U.S. Department of Energy Office of Electricity Delivery and Energy Reliability
DR	Demand Response
DSO	Distribution System Operator
EPRI	Electric Power Research Institute
ERCOT	Electric Reliability Coordinator of Texas, Inc.
HVAC	Heating, Ventilation, and Air Conditioning
ISO	Independent System Operator
ISO-NE	Independent System Operator – New England
LAAR	Load as a Resource
M&V	Measurement and Verification
MISO	Midcontinent Independent System Operator
NYISO	New York Independent System Operator
PJM	Pennsylvania-New Jersey-Maryland interconnection
QoS	Quality of Service
RTO	Regional Transmission Operator
SCADA	Supervisory Control and Data Acquisition
TCL	Thermostatically Controlled Loads

1 Introduction

Today ancillary services are supplied by power plants that are not producing power at their most efficient output level, tying up expensive capital investment, wasting fuel and increasing wear and tear from continually adjusting their output in response to the immediate balancing needs of the grid. Further, the need for ancillary services is projected to increase substantially as renewable generation penetrates to 20% of power plant capacity and beyond. DERs, such as distributed generation, storage and responsive loads, can provide equivalent services by adjusting their power demand and output rather than relying on centralized power plant output adjustments. To engage DERs in a meaningful manner, a unified approach to control, coordination and market mechanisms is need to enable large-scale penetration of controllable devices.

Several DLC approaches have been proposed [1]-[4] to provide regulation services using thermostatically controlled loads (TCLs). However, the control is typically coarse, limited to specific applications (e.g., peak shaving or energy shifting) and/or number of uses per year. Continuous (and potentially high speed) price-based mechanisms coupled with automated systems that respect user inputs for flexibility and comfort have the potential to provide more granular and smooth response when coordinated across a large group of devices. Price-based mechanisms allow consumers to determine their own willingness to participate, in near real-time if they choose to do so, and reward that participation through incentives or reduced bills. Several such approaches have been proposed [5]-[6]. However, a significant concern with price-based approaches is confidence in the ability to achieve a predictable and stable system response.

A transactive control approach allows DERs, which are largely owned and operated by consumers and third parties, to be seamlessly integrated with the operation of traditional grid infrastructure and coordinated to produce a smooth, stable, predictable response as required by grid operators. As compared to traditional DLC approaches, transactive control solutions emphasize a decentralized approach, where consumer and third-party decision-making is kept local and private, to manage such large-scale deployments with (eventually) hundreds of millions of individual assets while maintaining free will on the part of consumers.

PNNL pioneered the use of transactive controls in distribution systems during the GridWise® Olympic Peninsula Demonstration Project (funded by DOE-OE, Bonneville Power Administration, Pacificorp, and Portland General Electric), engaging consumer loads (HVAC and water heaters), commercial generation units, and large municipal water pumps via a transactive retail energy market to reduce peak demand and manage wholesale prices. The experiment was a success, and prompted the development of GridLAB-D as a simulator capable of helping to understand the applicability to different climate regions and to explore additional control designs [7]. Many of the principles explored in the Olympic Peninsula Demonstration [8] are applied and refined in the AEP gridSMART® [9] and Pacific Northwest Smart Grid Demonstration [10] (both funded by DOE-OE and partner utilities through the American Recovery and Reinvestment Act). The potential for transactive control and coordination is far greater than demand response for peak load and wholesale price management. Ancillary service markets are one potential option. Engaging responsive load resources (and potentially distributed generation and storage) in ancillary service markets may provide additional revenue for both the consumer and the aggregator, increasing the value of distributed assets. Different end-use devices may lend themselves better to different service markets, depending on the alignment of time scales, opening the potential of engaging a wide variety of devices, for a range of different services.

There are a number of different ancillary services, each having different characteristics. Regulation requirements are determined every few seconds, e.g., 4 seconds. Ramping and spinning reserves are relatively slower acting, ranging from minutes to hours. Additionally, devices have significantly different availability during different times of the day (e.g., day versus night, summer versus winter). The aforementioned demonstrations have shown that a retail market with a period of five minutes is appropriate for engaging residential HVAC systems, due to alignment with the mechanical transition behavior of the HVAC units and with most real-time energy markets. By extension and with some modifications, the same five-minute system could be used to deliver ramping and spinning reserve services, as the periods are greater than five minutes. On the other hand, regulation services occur on a much short timescale (2-4 seconds). To extend a similar transactive system to a regulation market, the market period would change from minutes to seconds, requiring two-way communication every 2-4 seconds. This would undoubtedly place undue burden on today’s equipment and communication systems. This shines doubt on the efficacy of transactive mechanisms when high-speed control is required.

The purpose of this project is to develop a plausible transactive framework for DR and DER participation in a regulation market. This document will focus on the methodology for creating a transactive-based regulation market, using one class of end-use devices as an example.

The system contains two parts, one for acquiring resources at a longer timescale and a second for controlling the devices in a distributed manner at much shorter timescales. The first is based on a formal double-auction market where every five minutes each device bids the amount of resource it is able to provide and the minimum price that it would accept to provide that resource. The bid price is determined by the current state of the device and the willingness of the consumer to participate. The market system collects and orders the bids by price, and then determines a cleared price to meet the level of regulation needed. It broadcasts the cleared price to the devices, which results in contracting the services of the least cost resources. By contract, the devices that cleared the market are now engaged for the next five-minute market period. They are part of a distributed control system that allows them to respond at four-second intervals to a broadcasted regulation signal. The approach also limits the number of times devices can cycle between states (say on to off) in a given amount of time to protect the equipment life.

Details of the both methods are described in Section 3. Models of the market and control systems were developed in GridLAB-D [7]. Simulations were performed to test the system operations for delivering regulation services, evaluating both the effectiveness and the limits of the system. The details of these simulation results are discussed in Section 4. Potential “Measurement and Verification” mechanisms, used to determine the effectiveness of the system, are discussed in Section 0.

2 Background

Historically, generators have been major players in ancillary services and the markets and control systems have been developed accordingly. Participation by DR and DERs is in a nascent stage and is largely limited to industrial plants and large commercial buildings. In this project, the focus is on a large number of small devices providing such services. Participation of such DR in the energy markets has been the focus of the previously mentioned transactive control projects – the Olympic Peninsula and gridSMART® projects. A good understanding of the availability of DR and DERs when the services are needed is an important first step, as discussed in a previous report [11]. The general conclusion of this and many other studies referenced in the previous report was that DR resources are capable of adequately providing ancillary services. Further understanding of the interplay between energy and ancillary markets, and tools for DR aggregators to acquire resources and deploy them are needed; this is important regardless of the specific implementation.

Participation of DR and DERs, especially a large number of small devices, in the energy and ancillary services market are, as noted before, in a nascent stage. Of the ancillary services considered, spinning reserve, ramping and regulation, the first two share certain important characteristics with energy. They all target the responsive energy consumption to be held to a certain level over a period of time on the order of several minutes, whereas regulation requires responsive energy consumption to move up or down every few seconds in reaction to a system signal. Although one can envision separate transactive markets for energy, spinning reserve and ramping, implementation at the device level, and participation in more than one market, becomes quite complex. If these are priced differently at the wholesale level, the aggregator could respond to a portfolio of wholesale market positions by developing appropriate prices and incentives at the retail level. Regulation service requires response by the devices every few seconds (two seconds in the case of PJM) in response to the regulation signals derived from the area control error (ACE) that has a frequency component and a tie error component [12]. With today's technology, it is not practical to impose device-level transactive bidding every two seconds, so a bid period of five minutes was considered with the requirement of equipment responding to regulation signals every two seconds during a bid period.

A regulation signal can be decomposed into low and high frequency components. This separation into two frequency ranges is similar to PJM's Regulation A and D [12]. PJM refers to "RegA" as traditional resources with limited ramp rates but no limit on duration and "RegD" as fast-moving resources that can respond quickly, but are unable to sustain that level for long periods. It should be noted that this decomposition is not a requirement for implementing the specific transactive control approach discussed in this report, but instead serves as an illustrative example. One of the motivations for PJM's separation was to reward faster acting regulation devices in a separate manner than slower ones. In the control algorithms that follow, regulation refers to the high frequency component if such a separation is done or the full regulation, if such separation is not performed – Section 4 will discuss the control system's ability to track low and high frequency signals.

Markets for secondary frequency control, or regulation reserves, i.e., generation and demand response capacity, exist in all energy regions of the country. These ancillary services markets pay reserve capacity to be available to restore frequency levels and power interchanges with other systems to their nominal levels following an imbalance. Table 1 presents the terminology used by various ISO/RTOs for their regulation reserve markets. Two of the six ISO/RTOs operate separate regulation up and regulation down markets, while the others make no distinction between reserve capacities that are used to provide either up or down regulation. ERCOT defines regulation down as capacity to respond within five seconds between a generator's base point and the lowest sustainable limit, and regulation up between a generator's base point and highest sustainable limit [13].

Table 1: Regulation reserve markets in various ISOs

	CAISO	ERCOT	ISO-NE	MISO	PJM	NYISO
Product Name	Regulation Reserve Regulation up Regulation down	Regulation Services Reg Service–Up Reg Service–Down	Regulation	Regulation Reserve	Regulation	Regulation
Minimum Offer Capacity (MW)		1			0.1	
Minimum Ramp Rate (MW/min)			1			
Max Time to Deliver(min)	10-30	10	5	5	5	5

Separate markets for regulation up and down reserves helps to better reflect system conditions; for instance, regulation down may be valued more highly than regulation up reserves at night times when generation is high and load levels are low. In such situations, conventional units typically operate at their minimum generation level (or close to it), having greater regulation up rather than regulation down capacity. Greater supply of regulation up reserves implies that this capacity should be priced lower than regulation down reserves, and having separate markets allows differential prices to be made available to market participants.

While not directly addressed in the evaluation of this transaction-based control system, it will be clear that communications could potentially influence the behavior of the system. Communication needs for demand response and regulation is an active area of research. A discussion of history and recent advances can be found in [14]-[15]. Communication standards such as ANSI/CEA 2045 [15] that can accommodate multiple technologies such as Wi-Fi, cellular and radio frequency are evolving; devices conforming to such standards may even be installable by the homeowners. The Electric Power Research Institute tracks advances in information and communication technologies [16]. These activities lead to the communication technologies necessary for DERs to provide transactive ancillary services. Initial simulation studies were performed to understand the impact of communication latency on a gridSMART-like transactive mechanism [17]; *extreme* communication delays affected the behavior of the market clearing, eventually leading to an "unsolvable" market that erroneously settled to the price cap. This can have a direct impact on consumers and the control of the system. The control system described in this report requires not only two-way five-minute market information but also one-way 4-second broadcasts. To evaluate the potential of this system, and others like it, additional work is needed. One possible evaluation technique is to co-simulate the behavior of the power system, the loads, the control system, the market, and the communication network [18]. This method allows users to explore a wide range of possible communication impacts on the reliability of the control system, such as latency, lost information, congestion on shared network resources, and security. The simulation environment can also be used to determine the minimum requirements of the communication network, including whether it is sufficient to layer multiple control systems on the same network, or whether a stand-alone communication network is required.

While minimizing the amount of communication required was considered as part of the development of this control system, refactoring may be necessary to create a system more robust to communication failures. Future works should evaluate the communication network as part of the system of study.

3 Market and Control Framework

This section describes the overall transactive control framework as applied to regulation services. The system contains two parts, one for acquiring resources transactively at a longer timescale and a second for controlling the devices in a distributed manner at much shorter timescales. The hierarchical system is shown pictorially in Figure 1.

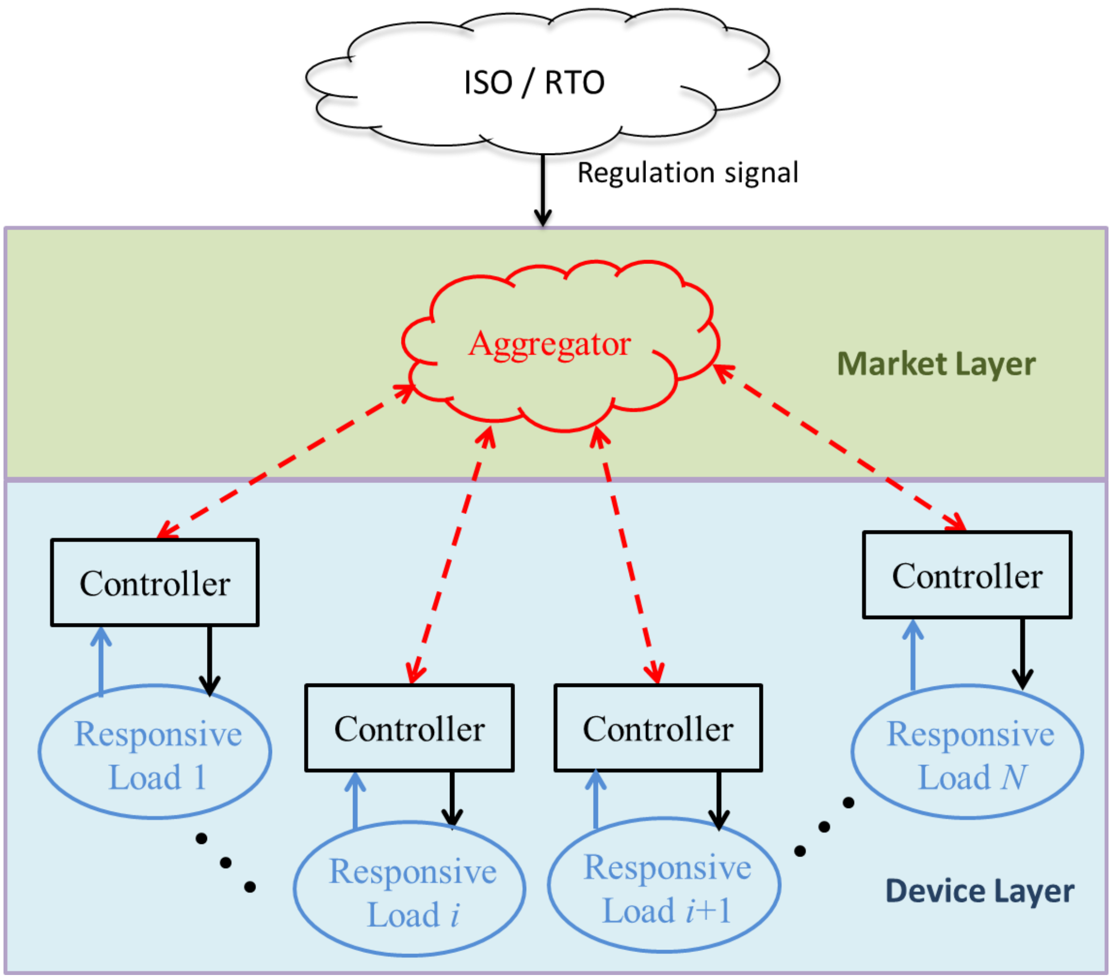

Figure 1: Market-based control strategy framework

An asset that is engaged by the aggregator (i.e., the price was what the aggregator was willing to pay) means that the control algorithms they use to provide regulation services are available to be activated. With a "winning" bid (at or below the market clearing price), the control system is activated within the device during the five-minute market period. Once activated, the device uses a distributed control approach to respond to the regulation signal broadcasted by the aggregator. Through this system, the device has the potential to change states (on to off or off to on) every four seconds via a Markov-chain device layer model that will be described in Section 3.2. To avoid equipment damage, devices can determine how many times they are willing to change states within a five-minute period, relying on the aggregator to provide enough resources to create a smooth response at the system level. Assets with the winning bids will be and rewarded for 1) their availability to be controlled and/or 2) their performance in delivering as promised.

More detail on the operation of the two subsystems is discussed in the following sections.

3.1 Resource Acquisition

This section describes the transactive market mechanism, both the aggregator agent and the device-bidding agent, developed to acquire regulation resources from end-use devices. The next sections will describe this by 1) the device bidding mechanism(s), including information shared with the central clearinghouse, 2) the central clearinghouse mechanism, including information broadcasted to the devices, and 3) the way information is then used by the control mechanism described in Section 3.2.

Resources are acquired on a five-minute basis via a transactive market mechanism. The market mechanism determines the number of devices available in each state (on versus off) at the beginning of the market period, acquires the needed amount of resource by calculating the market clearing price, and determines the transition probability matrix at that time instant. While other mechanisms might be used, to show the feasibility of this system, *two* coordinated double-auction markets are used to determine the price and availability of resources available to either 1) increase or 2) decrease their individual consumption. For simplicity, we will call these the off-to-on (currently off and available to turn on) and on-to-off auctions (currently on and available to turn off). Devices construct a bid for each five-minute market cycle that reflects the amount of resource they can provide, the number of times they can provide it, and the price at which they will engage. For the purposes of this study, each device will be limited to one allowable state change per five-minute market,

Markov chain

A Markov chain is a “memoryless” random process that describes the transition probabilities from one state space to another. It is often used to describe statistical models of real-world processes where the next state only depends on the current state, and the process occurs at discrete sets of time. In the case of residential appliances, states can be as simple as ON or OFF. The Markov-chain model describes the probability of a set of similar devices transitioning from ON to OFF and OFF to ON. In the example below, there is a 25% chance a device that is ON will turn OFF in the next time step, a 35% chance an OFF device will turn ON, and a 40% chance an OFF device will re-enter the OFF state. This is described mathematically in a transition matrix.

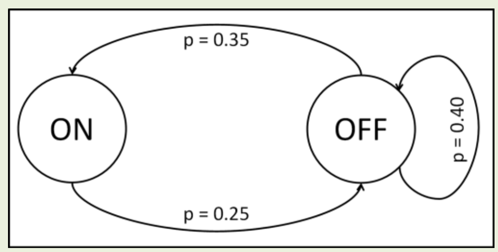

Transitioning From \ Transitioning To	ON	OFF
ON	0%	25%
OFF	35%	40%

Transition Matrix

$$\begin{bmatrix} 0.00 & 0.25 \\ 0.35 & 0.40 \end{bmatrix}$$

In a power system application, the load on the system at the next time step due to this class of devices simply becomes the number of devices in the ON state times their individual demand(s). However, more complex states can also exist, such as HIGH, LOW, SATISFIED, LOCKED, or a variety of others. If a finite number of states can be used to adequately describe the behavior of a population of devices, well-known control methodologies can be applied to control the system.

limiting the stress on the devices. This work will focus on how to construct the system from the aggregator down to the device level. Additional investigation is needed to determine the best methods for an aggregator to participate in the wholesale ancillary service market.

3.1.1 Device Bidding

During each five-minute bidding cycle, individual devices make a decision about 1) whether they can participate in a regulation market (whether able to increase or decrease load), 2) how much load (kW) is available to participate, and 3) what flexibility they have to participate in the market ($/kW). This information is delivered to the market clearinghouse in the form of a {price, quantity} or {*P*, *Q*}[1] pair. Additional information may be required during this bidding scenario (such as "quality of service" or how many times it is available to change states), but will be limited to a {price, quantity} pair for the purposes of this study. Future work may be required to understand how an aggregator might mix non-homogenous elements (i.e., that are able to switch states more often than once per market cycle) in an equitable manner.

Determining whether the device can participate, and which market to participate in, is driven by its current state and whether it has the ability to change states in the next market cycle. For example, if a refrigerator compressor has been running for 15 minutes in normal operations (i.e., the internal air temperature is within safety ranges and is not recovering from a defrost cycle or extended door opening), it has the ability to turn off before the completion of its normal cycle. In fact, it may already have a "desire" to change state as it approaches its lower deadband. However, the device may have recently turned on, and to protect the equipment or ensure efficient operations, it may make a determination that it is not available to change states. Prior to the beginning of each market cycle, the device control agent makes a decision about whether it is available to participate in the on-to-off or off-to-on auction.

If the device is able to participate, it must also determine how willing it is to participate. This is the essence of previous transactive approaches, as its "desire" to participate will be reflected in its bid price. The device operates as a supplier of services, and as such, offers a supply bid (hence, a lower price will reflect a greater willingness to participate). To keep the bidding strategy relatively simple, an adaptation of the methodology in [8] and [9] was used here.

In the experiments discussed in Section 4, the internal air temperature of the home relative to its desired set point was translated via a linear function to relative prices. This effectively translates a consumer's comfort requirements (or flexibility) into a monetary value, which allows the agent acting on behalf of the consumer to participate in the market. An example of this curve is shown in Figure 2. The desired setpoint anchors the curve to the relative prices at the average price, while the current air temperature is mapped onto the y-axis as the desire to run (a demand bid, in this case, where a higher price reflects a greater desire to run). Relative prices are calculated as a function of a 24-hour rolling

[1] There is often confusion when discussing *P* and *Q* variables in the context of markets and power systems. In the context of markets, *P* refers to price and *Q* to quantity. In the context of power systems, *P* is real power while *Q* is reactive power. For the purposes of this discussion, we will use the former definition.

window of the cleared price, where P_{avg} is the average price of the last 24 hours while σ_{avg} is the standard deviation of the last 24 hours. Refer to [8] and [9] for more details on the operation of this system.

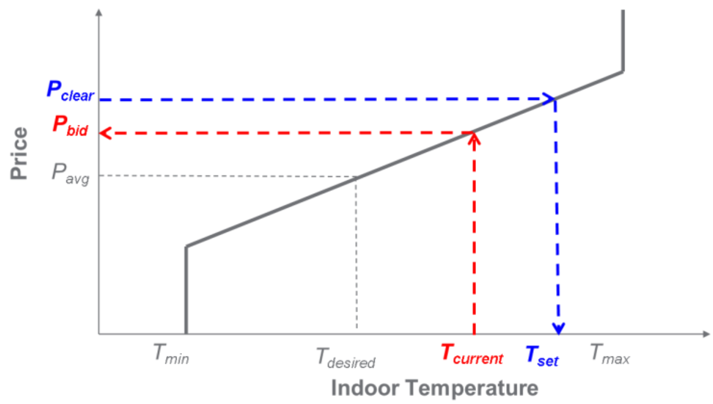

Figure 2: Thermostat bidding agent used in previous transactive demonstrations

To specify this system and make it more applicable to the proposed aggregation auctions, a few modifications are needed depending on the appliance in question. For a refrigerator, HVAC, or other TCLs, this is a simple translation. If the device is on, and willing to turn off, it means that it is within its safe operating range, specified by $T_{\min}$ and $T_{\max}$. The bidding strategy will then be defined by the linear curve shown in Figure 3.a, where the current air temperature (measured) of the cavity in comparison to the desired set point is used to represent the willingness of the device to participate in the market. Figure 3.b shows a similar curve, except in the situation where the device is off and willing to turn on. Note that there is no requirement on the linearity of the curve or dependency on a continuous curve between plus and minus three standard deviations. This formulation expresses one possible method that may be employed. Future methods may employ curves that are more in line with a manufacturer's recommendations to discourage short-cycling or may act as a more strategic bidding strategy (e.g., predictive). This formulation is merely used to represent a basic case that can be used to test the ability of appliances to participate. At this time, the length of time to calculate the average and standard deviation of price is left as an open question. It is assumed that this value will be far shorter than the 24 hours used in the wholesale energy case – on the order of 4 hours, possibly – but the actual value may depend on the period that the regulation signal remains generally un-biased or how much participation is required by the aggregator.

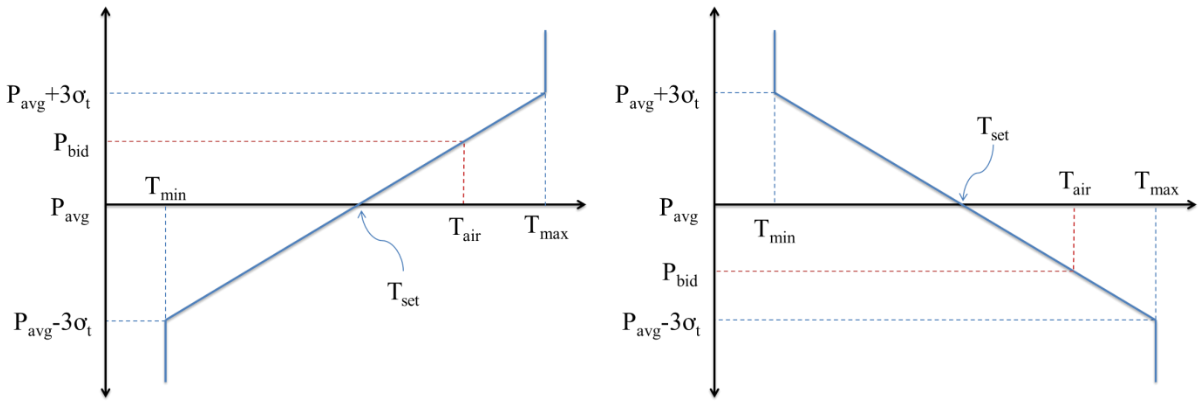

(a) (b)

Figure 3: Bidding curve for a refrigerator that is able to (a) turn off and (b) turn on

The quantity (*Q*) may be obtained from manufacturers' specifications, (e.g., the rated load of the compressor for a refrigerator or HVAC system). A reasonable value for a refrigerator might be 0.15 kW. While there is some variation in actual demand and the rated demand, we will assume the difference is negligible for this experiment. This quantity is not the overall resource available to participate in the regulatory market, as the refrigerator cannot continuously provide this level of service, but rather indicates the available resource in a single state change. To understand the resource available to the aggregator as a whole, a "quality of service" (*QoS*) value will also be required. To help sort which market the devices are participating in at any given time, the quantity will be a signed value. If the device is able to reduce demand by the given amount, the quantity will be negative. If it is able to increase demand by the given amount, the quantity will be positive. This will be purely for bookkeeping purposes. Other thermostatically controlled devices that have access to internal state measurements (i.e., temperature) could use a method similar to this. For example, a heat pump water heater, which typically is able to digitally measure an effective water temperature, could use the same method (as could a thermostat for a residential HVAC system). For devices that operate differently, different methods will need to be designed for calculating P_{bid}, however quantity and quality of service will remain similar. An interesting example would be a variable speed fan, which has a continuous output from 0%-100% of demand. However, not all output levels are equal. If the variable speed fan sits at 60% of its rated output, it could easily bid 40% (60% output +/- 40%) of its total load with a *QoS* of 1.0 and have little to no effect on its efficiency (see Figure 4). However, the efficiency is greatly affected by where it sits on the output curve, so it would make sense to bid the first 20% at a much greater value (region 2 in Figure 4) than the latter 80% of the load (region 1). The participation of other devices will be left for a later discussion, but this example clearly shows that strategic bidding will play an important role in devices capturing the most value.

The central market clearinghouse must receive this bid prior to the closing of the market. Bids received after the market closure will be considered invalid and may not participate in the market (nor be paid for participation). Note that while this is out-of-scope for this work, this may have an impact on communication requirements.

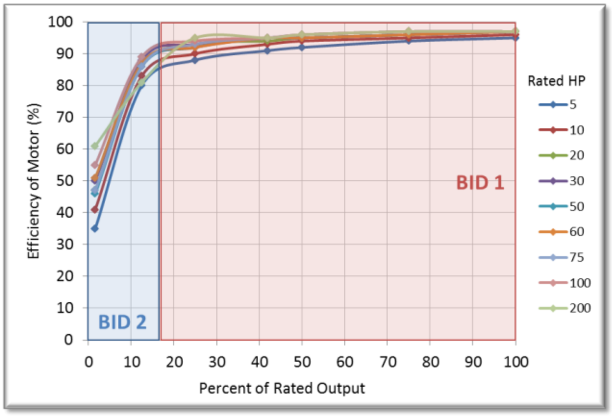

Figure 4: Variable speed drive part-load efficiency curves; data extracted from [19]

Quality of Service

QoS is borrowed from the computer network industry, where it refers to the ability to provide or guarantee a certain level of service to data flows, often with differentiating priorities to different applications. It is sometimes referred to as a quality measure. In this context, we use it to indicate the number of times the device can perform the required state change during the market period. This indicates how well suited the device is to providing regulation services. This also indicates the amount of resource available from an aggregated group of similar appliances. By using the *QoS* value, the aggregator can mix heterogeneous devices into a single auction; heterogeneous sets of devices with different QoS's and quantities may be mixed in this system, as long as the bid can be discretized into a function of price, quantity and QoS, allowing a single market to serve large sets of dissimilar devices.

The *QoS* value will be defined as $QoS = ft_p$ where f is the maximum frequency at which the device can change states (state changes per second) and t_p is the time between regulation signal updates (e.g., 4 seconds). *QoS* will be capped in its effective range when the device can change states at least once per regulation signal update, or $f > \frac{1}{t_p}$, such that $QoS = 1$. When $f < \frac{1}{t_m}$, where t_m is the market period, then $QoS = \frac{t_p}{t_m}$, as it is assumed that the device has already agreed to participate and therefore will be called upon.

For the refrigerator example with a 4-second regulation signal, we assume it can only participate once per 5-minute market cycle to prevent fast-cycling, so this value would be $QoS = 4/300 = 0.01333$. When aggregated to the market level, this indicates that 10,000 identical refrigerators at 0.15 kW will have an overall resource availability of $10000 * 0.15 * 0.01333 = 20$ kW. Note, this means that to guarantee 20 kW of regulation services, 10,000 refrigerators would need to be engaged, assuming the regulation signal moved from maximum to minimum within the market period.

The *QoS* value could potentially be used to determine a particular device's payment as a percentage of the overall resource it provided, as it effectively describes its participation level in the control system – devices with a higher *QoS* could be better rewarded than those at a lower *QoS* level.

3.1.2 Central Market Clearinghouse

The clearinghouse is formed through two double-auction markets. When the market closes, all valid bids are sorted into off-to-on and on-to-off regulation groups to form two independent markets. *QoS* and *abs(Q)* are multiplied to get the effective amount of available resource in kW from each device, forming two lists of quantity / price pairs. In each market, the bid list is sorted from lowest price to highest and the quantity is cumulatively summed to represent the supply side curve. For the purposes of this experiment, the demand curve will be represented as a boundary condition, either as a fixed price (e.g., distributed load is purely responsive to price fluctuations) or a fixed quantity (e.g., aggregator bids a quantity into the ancillary market and is procuring the necessary resource) received from the wholesale market. These two curves are represented in Figure 5. Future applications may use the demand bids to form an optimization mechanism for the aggregator to mitigate certain risk factors, but these simpler curves are used for initial experimentation. Note that heterogeneous sets of devices with different *QoS*'s and quantities may be mixed in this system, as long as the bid can be discretized into a function of price, quantity and *QoS*, allowing a single market to serve large sets of dissimilar devices.

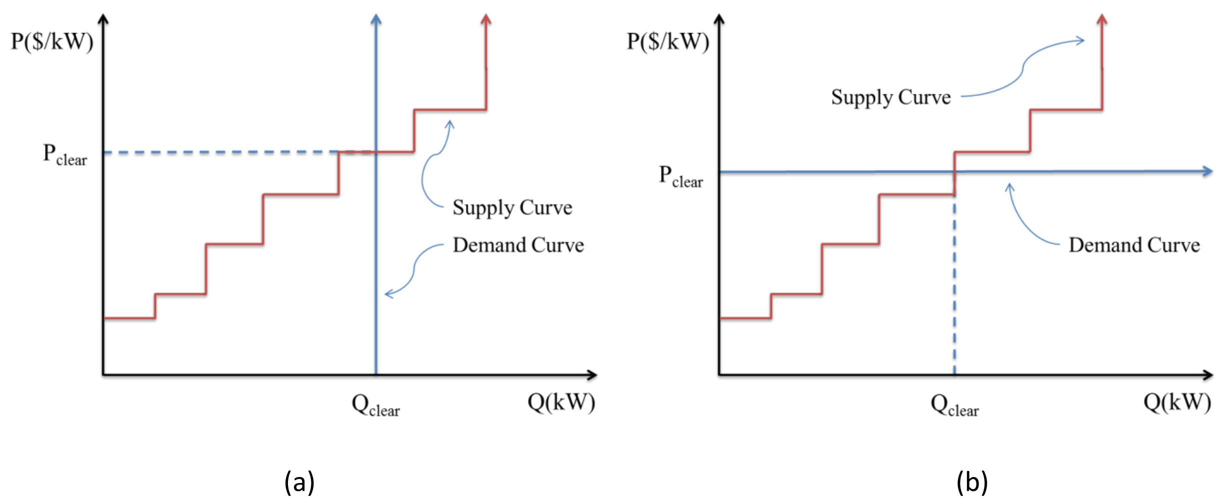

Figure 5: Example clearing of (a) fixed quantity and (b) fixed price market

The market is cleared by determining a cleared price and cleared quantity at the intersection of the two curves. Marginal clearing, exceeded capacity, and zero capacity markets will be cleared similar to previous mechanisms [8] and [9]. This process is performed every market cycle. Upon clearing, the cleared price is broadcasted to all of the participating devices. The cleared quantity may also be broadcasted, but is not used by the current mechanism; however, marginal fractions are broadcasted for use when large numbers of devices resolve to the same price as can occur during periods of low diversity.

The scope of this research did not determine how an aggregator would choose a cleared capacity to extract the most value out of the system. A discussion about the impacts of clearing quantity choices can be found in Appendix A.

3.1.3 Control Systems Interaction

Upon receiving the clearing price during each market cycle, if the device "won" the bid (i.e., the bid price is less than or equal to the clearing price) the device will be engaged by the control system, and will participate as described in Section 3.2. If the device "lost" the bid (i.e., the bid price is greater than the clearing price) the device will not participate in the control system, and will continue to operate normally. Either way, the device will have the option to participate in the following market, if it so desires and is able to do so. Note that this system depends on adequate communications being available. In previous work [8], it was required that devices bid a minimum of 30 seconds prior to market closing to ensure that the bid was received prior to market closing. High-latency communication may lead to devices participating in the control system that should not be involved, leading to a greater rate of error. If the cleared price is broadcast and received later, then the device's response may be delayed until the cleared price is received, again increasing the rate of error. These are factors that will need to be considered in future tests, but are ignored in this experiment.

3.2 Device Layer Control Design

In this section, the modeling and controls for a specific class of devices, thermostatically controlled loads (TCLs) such as HVACs, water heaters, refrigerators and clothes dryers, is discussed. Due to the inherent thermal energy storage, these loads can often be switched on-to-off or off-to-on for short periods of time (seconds to minutes) while minimally affecting the end-use performance. The power consumption of the TCLs is assumed to be zero in the OFF state, and is a non-zero constant when in the ON state. The total controllable power consumption at a bus is the sum of the powers of all the controllable TCLs in the ON state on this bus. Therefore, the objective for each controllable load is to calculate its own probability of turning ON or OFF, such that, the change in power consumption of the aggregation of the controllable loads at a given bus is equal to the regulation signal received from the aggregator. To address this problem, Markov-chain models, which are distributed and maintained at the device level, are used to estimate the aggregate load response of controllable loads following a similar approach first proposed in [20] and [21]. A Markov-chain representation implies that the transition probabilities at a given time depend only on the current state of the system and not the history of states; this assumption is easily made due to the bidding mechanism described in Section 3.1 which limits the number of state transitions within the given market period. Markov-chain models are developed for HVACs and water heaters, but a similar approach can be adopted for other types of thermostatically controlled end-use loads such as refrigerators and dryers, or storage-based devices, such as electric vehicles.

A state transition model describing TCL operation could be simply defined as two states: ON and OFF. However, some TCLs, such as HVAC systems, cannot or should not be cycled quickly between states (due to a variety of reasons, such as compressor cycle relaxation and decreased efficiency). With some simplifying assumptions, the HVAC model can be expanded to include a LOCKED state, which prohibits additional state changes for a given amount of time. Therefore, HVACs can be modeled using four operating states: ON, OFF, ON-LOCKED, and OFF-LOCKED. It is assumed that if the load switches to the locked states (ON-LOCKED or OFF-LOCKED), then it will stay in the locked states for the rest of the control period.

For the purpose of this discussion, two terms will be defined. "Unforced" state switching is defined as the state change that would occur if the control system were removed and the device returned to normal operation, most similar to its natural behavior. A "forced" state change is defined as a transition to a new state initiated by the controller that would not have occurred through normal operations; the forced state is held until the end of the control period. When a device is engaged by the market-based control system, the control algorithm ignores unforced state switching within each control period and considers only forced switching. The forced transition between groups of different operating states is depicted in Figure 6, where μ_1 denotes the forced switching probability from ON to OFF-LOCKED, and μ_0 denotes the forced switching probability from OFF to ON-LOCKED.

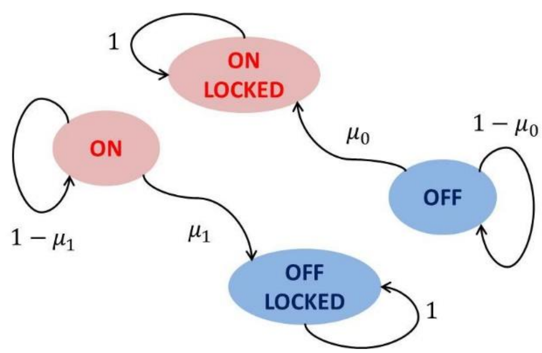

Figure 6: State transition diagram for a population of HVACs

Water heaters, however, can be modeled using two operating states: ON and OFF. Once again, the unforced state switching within each control period is ignored and only forced switching is considered. The forced transition between groups of different operating states is depicted in Figure 7.

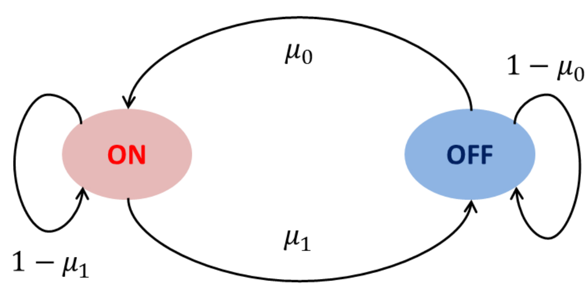

Figure 7: State transition diagram for a population of water heaters

For HVACs and water heaters, the forced switching probabilities $\mu_0(t_k)$ and $\mu_1(t_k)$ are calculated to match the required regulation service requested by the aggregator. For more details on this, the reader is referred to Section 2.4 in [20].

In the above discussion, the control design assumes that when a device is in a given state, it remains in that state unless told to change by the control mechanism. Ideally, the unforced rates of state change are much slower than the forced rates, so that even without enforcement of this behavior, the amount of error introduced is negligible. However, systems like HVACs tend to have a rapid unforced rate of state change (relative to the market period), especially when the HVAC systems are oversized (rapid cooling) or during extremely hot days (rapid warming). If allowed to operate using its normal control in addition to the regulation control function, this may lead to the device either (1) changing state prior to the controller initiated state change or (2) returning back to its prior state after a controller initiated state change. This has the potential to lead to inconsistencies between the modeled states within the controller and the actual distribution of states. This is especially true when considering that the market will tend to choose devices that are most likely to change states during the market period anyway, as their desire to change states will be reflected by a lower bid price. Devices such as water heaters, with larger amounts of thermal mass, are less likely to suffer from this potential problem, as the unforced rate of change of states may be much slower. Appendix B describes a similar process for extracting the state diagram model (and control) for other thermostatically controlled loads; in this case, a refrigerator.

4 GridLAB-D Simulations and Results

4.1 Simulation and System Setup

The control system described in Section 3.2 and the market mechanisms described in Section 3.1 were implemented in the GridLAB-D™ simulation environment to verify the behavior of the system. GridLAB-D [7] is a U.S. Department of Energy (DOE) developed open source simulation tool that incorporates advanced modeling techniques to deliver the latest in end-use load modeling technology. This agent-based system allows users to examine the interplay between the elements of a distribution system without the use of reduced-order or aggregate models. Rather, it relies on advanced physical models of individual elements and their simulated interaction through time [22].

Existing retail double-auctions within the simulator were adapted to fit the requirements discussed in Section 3.1. The performance of these auctions had been previously evaluated as part of the AEP gridSMART Demonstration project [9]. New market clearing mechanisms were added to support periods of stagnant bidding (i.e., no bids for days at a time), a function for supporting QoS as part of the bid was added (but not used in this analysis), and new statistic collection tools were created (also to support periods of stagnant or non-existent bidding).

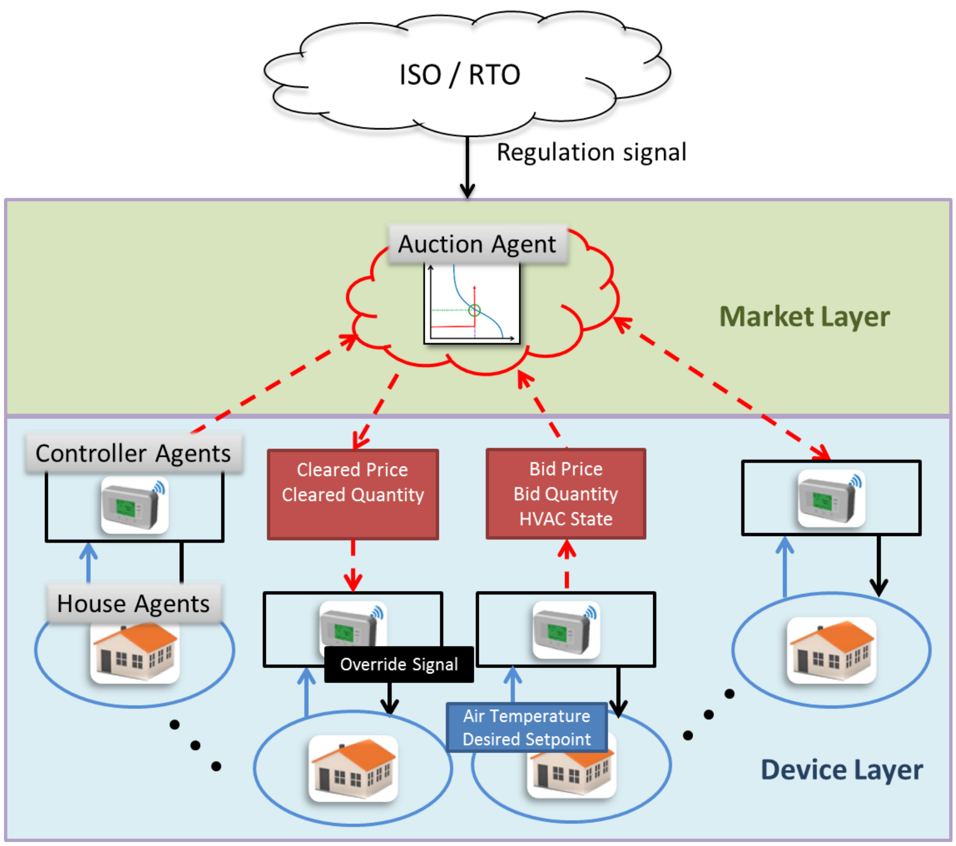

Figure 8: GridLAB-D implementation of framework.

Most of the development occurred within the *controller* agent, or an interface between the *auction* agent and the various load agents (such as a *house*). The controller has two main functions. First, it uses local information to formulate and transmit a bid to the auction (e.g., current air temperature is used to calculate a price at which the device would be willing to turn off). Once the market clears, it compares the broadcasted cleared market price with its bid price to determine whether it is cleared to participate in providing ancillary services during the corresponding market period. In this implementation, the controller also needs to track the current state of the device it is controlling to bid into the correct auction depending on the current state. The controller releases control back to the thermostat after the control action is completed, or directly overrides the unforced behavior of the device when a control action is necessary. This is represented graphically in Figure 8. The studies in this section focus on devices that are only allowed (by the controller design) to change states once per market cycle. While QoS was defined in previous sections, it was not directly tested beyond a value of once per market cycle.

To evaluate the performance of the control system, a 5,000-home model was created that included a power distribution system and the various auctions and controllers, as needed. Each home was individually specified and included air conditioning, gas or heat pump heating, a physical water heater representation, and a time series load shape to represent "other" loads within the home (e.g., lights, computers,

What is an aggregator?

An aggregator is an entity that engages services from a wide variety of distributed resources (e.g., distributed generation, distributed storage, or demand response) and represents them as a single resource to the market and system operators. The aggregator acts as a broker between individual consumers and the larger electric system. The consumers are typically rewarded financially for providing a given service, which allows the aggregator to control the device to provide a service to the larger system.

For the purposes of this discussion, an aggregator can be the distribution utility itself, a third-party aggregator contracted by the utility, or an independent aggregator. A common embodiment may also be that of a Distribution System Operator (DSO). The purpose is to describe the functional role of an aggregator, not the entity that provides the service.

A large number of customers today are self-generating electricity through rooftop PV and other distributed generation systems. As even greater numbers of customers decide to move in that direction, system operations may be impacted adversely, unless the DER operations are managed and coordinated properly. Additionally, utilities' financial models will be affected, further affecting their ability to operate the system in a reliable manner. Transactive systems hold the potential to coordinate, manage, and leverage the flexibility from DERs (including DR) to reduce the uncertainty from system operations. Hence, transactive systems can also reduce the risk of uncertain retail rates, and maintain overall quality of service for end-customers. The aggregators or other service providers who contract with end-customers assume some of the risk associated with the management of DERs (diversity of resources helps reduce some of the operational risk). In the longer-term planning horizon, given the modular and incremental nature of TE resources, the risk associated with changes to environment policy and regulation is also reduced, because the TE resources can be deployed in smaller increments than other conventional resources, allowing opportunities for course correction or selecting alternate paths.

etc.). The energy characteristics of the homes, such as insulation, square footage, thermostat set points, etc., were randomly selected in a method similar to that described in [23] and [24]. This resulted in a population of homes that roughly represents a suburban neighborhood in the Midwest of the U.S. with high penetrations of air conditioning and electric water heaters. In addition, each air conditioner's device level controller and was modeled with a 1-min lockout period. Differing weather conditions were studied, representing both near constant conditions and a peak summer day with a cool night, over the 24-hour simulation runs.

The market was constructed using historic pricing data from PJM's web site [25]-[26] so that it aligned with the weather data used in the model and the regulation/ACE signal used to control the devices. The quantity procured by the market at a given time interval was not driven by historical data (as this type of market has not been used before nor is there currently a connection to a transmission-level market), but rather was designed to test the system in various states. In the "base case" simulation, the quantity was chosen through multiple simulation runs to maximize the total amount of resource available to the aggregator. The same quantity value was used in subsequent cases. As will be seen, there are times when too much quantity was requested by the aggregator; at these times, the market is unable to clear due to uneven amounts of 'on' devices clearing versus 'off' devices clearing. Note, this does not indicate a failure of the power system, but rather a failure to resolve to a market price that allows the devices (and therefore the aggregator) to participate in the ancillary service market. While this may represent a loss of revenue, and indicates that there is room for improvement in the system, it does not represent a failure to deliver electricity to the consumers. Additionally, while not observable from the data presented in the following sections, there are market periods when more load could have been engaged, but was not. Future work could likely improve these mechanisms. This includes evaluating appropriate bidding strategies for the aggregator, in terms of when and how to use the available resources on the feeder to provide the maximum benefit, which should lead to greater benefits to the consumer, as well.

4.2 Analysis of Simulation Results

Through simulations, many of the factors affecting the ability of a population to provide a desirable response were investigated. The goal is to determine the effectiveness of the system in tracking a regulation signal, understand the limitations of the control system due to the physical constraints of the devices being controlled, and explore some of the external parameters of the control system. In particular, several studies were performed to understand the sensitivity of the resource population's performance to the following:

1) Compare low versus high frequency regulation signals,
2) Direct override of the natural behavior of individual devices regardless of indoor air temperature and lockout times versus override of the natural behavior only when air temperature is within temperature deadband limits,
3) Duration of market period,
4) Minimal device population size, and
5) Impact of time-varying weather.

4.2.1 Low versus High Frequency Regulation Signal

To compare the performance of the population of devices providing frequency regulation, 24-hour high and low frequency signals from the same day were selected from 2009 historic PJM ACE data. The signals are normalized to values between {-1, 1}, so that a single number can be broadcasted to all devices (rather than individual settings for each device). The tie-line interchange, or the exchange of power between adjacent control areas, can often lead to a bias in the amount of energy consumed over many market cycles (low frequency), while the frequency regulation is generally unbiased over relatively few market cycles (high frequency). Figure 9 shows an example of normalized regulation signal for a 24-hour period; note, that the green line roughly averages zero, while the blue line is heavily biased towards negative.

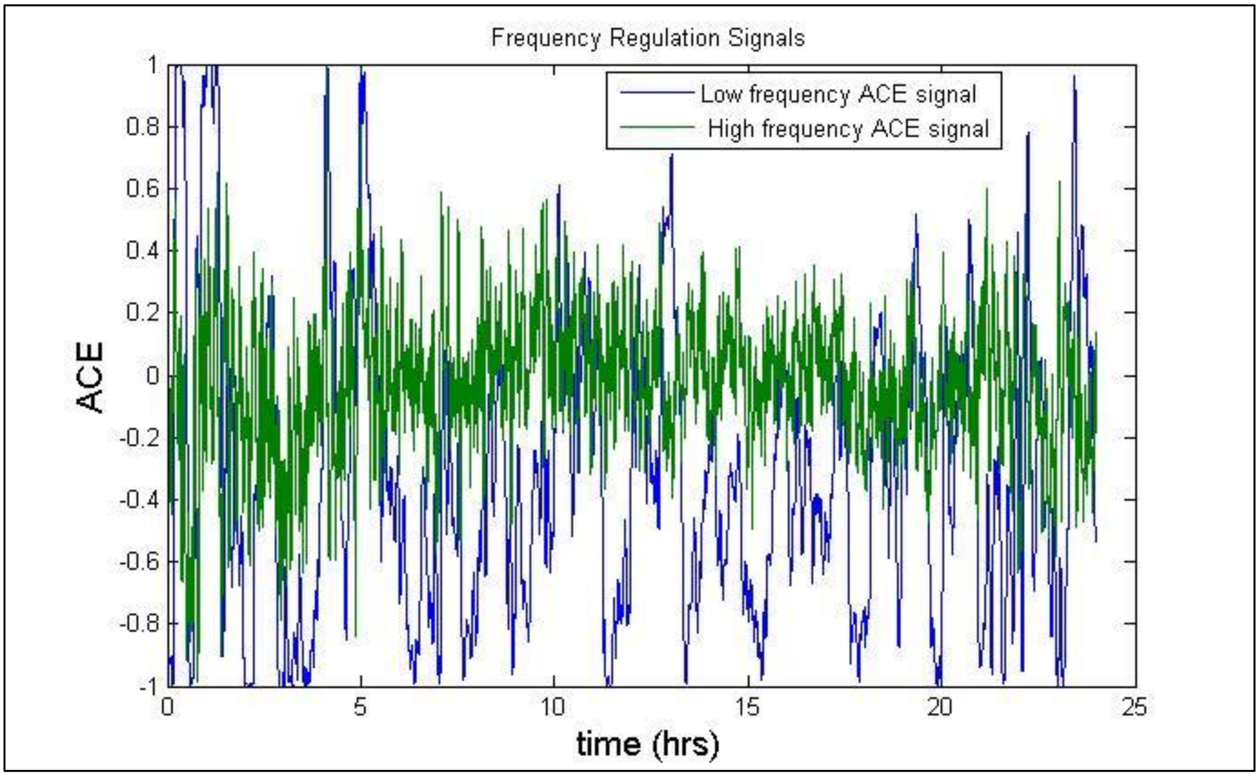

Figure 9: Example of low and high frequency ACE signal from PJM 2009.

Figure 10 shows the actual and desired response of the engaged population based on the low and high frequency signals; note that where the blue line (actual load of the appliances) is not visible, it is covered by the green line (load desired by the control system to match the regulation signal). In the current studies, it is not representing either up- or down-regulation, but rather operation around a central operating point – the market mechanism is easily adjusted to provide up- or down-regulation resources. For this reason, it is assumed that equal amounts of on-to-off and off-to-on resources are needed and the procured capacities of both markets are equal. The procured capacity can vary over time, as often as each market cycle, depending on the needs of the aggregator. For this study, as the mechanisms for aggregator participation in wholesale markets were not addressed, the capacity was selected heuristically to highlight operation. Both cases, i.e., low and high frequency cases shown in Figure 10, allow the controller to override natural behavior of engaged HVACs, which is referred to as forced control. Effectively, if a device controller makes a decision to be in a given state during the market period, it enforces that state on the HVAC system until the end of the market period – the succeeding section will explore the impact of not enforcing the forced state behavior. As a result, the simulation indicates that the population is able to closely follow the desired power modulation. The mean absolute error in actual and desired load relative to the total procured capacity is approximately 3% and 4%, respectively. In addition, compared to a baseline case, there was a negligible average increase in absolute temperature deviation from the desired temperature setpoints for each home. In the low frequency regulation case, the average increase was 0.03°F, while the high frequency regulation case was 0.02°F.

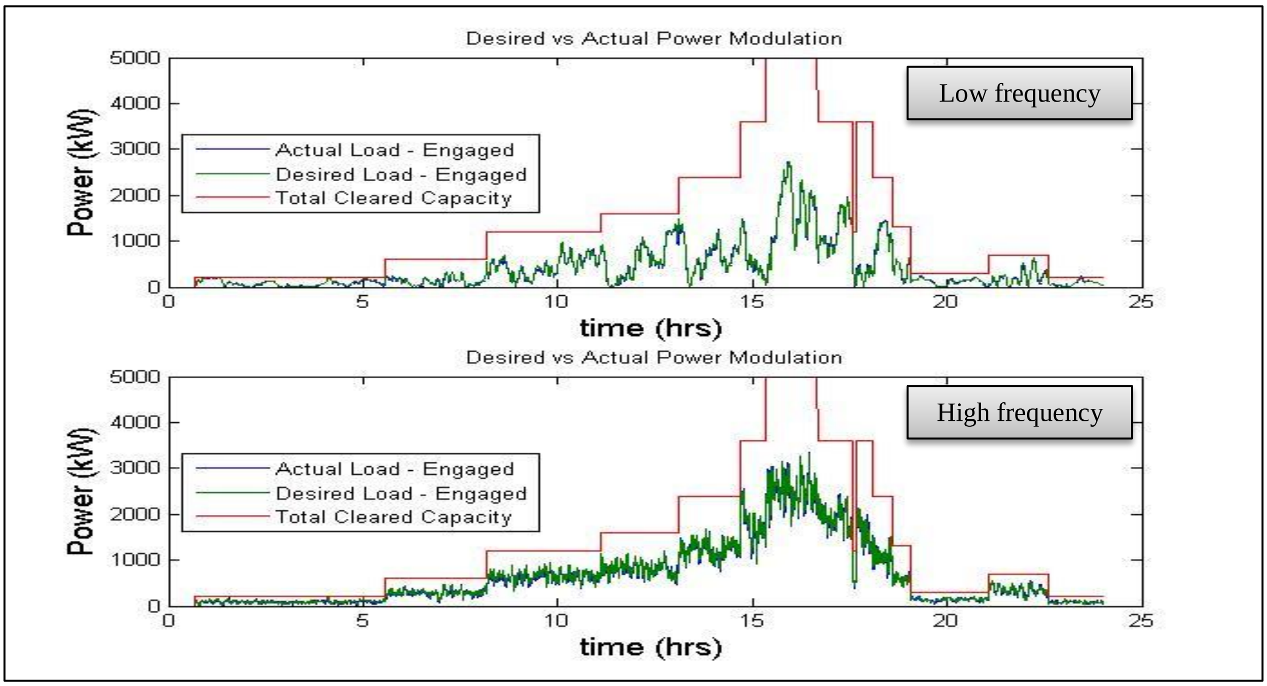

Figure 10: Comparison of the desired versus actual load with different frequency components.

4.2.2 Forced versus Unforced Control

The control design described in Section 3.2 assumes that when a device is in a given state, it remains in that state unless told to change by the control mechanism. Simulations were performed to explore the unforced case where the controller requests a state change, but the unforced device on and off cycle (in response to air temperature) are also modeled – in other words, the device will respond to the state change request by the controller AND due to temperature deviations. The unforced situation means the actual state of the device is less deterministic. Once again, the low and high frequency regulation signals are used in this study.

Figure 11 shows some of the results of these simulations. The top two figures use the low frequency component, while the bottom two figures use the high frequency component. The first and third figure show forced control, while the second and fourth show unforced control. Each figure shows the actual and desired response of the engaged population based on the regulation signals considered, as well as, the total cleared capacity of both regulation markets. The unforced control cases result in an increase in mean absolute error relative to the total cleared capacity compared to the forced control cases. The increase in error (3 to 14%) is even more pronounced in the low frequency regulation case. The error increases from 4% to approximately 9% in the high frequency regulation case.

The increased error in the unforced case is likely driven by the change in diversity of the thermostats. In normal operations, the population of devices is spread across the thermostat deadband, i.e., not all of the devices are lumped at one end of the thermostat deadband. However, as the devices are participating in the regulation service market, the diversity is affected. This is particularly true in the low frequency case, which contains a long-term bias due to the tie-line error – as can be seen in Figure 9 the low frequency component is not energy neutral. This leads to a “bunching” of the devices in the high-end of the thermostat deadband; when these devices change state and are released to operate in an unforced manner, they immediately turn ON because of high(er) temperatures within the home. This is not seen in the forced control case, as the device is held in its state until the market cycle.

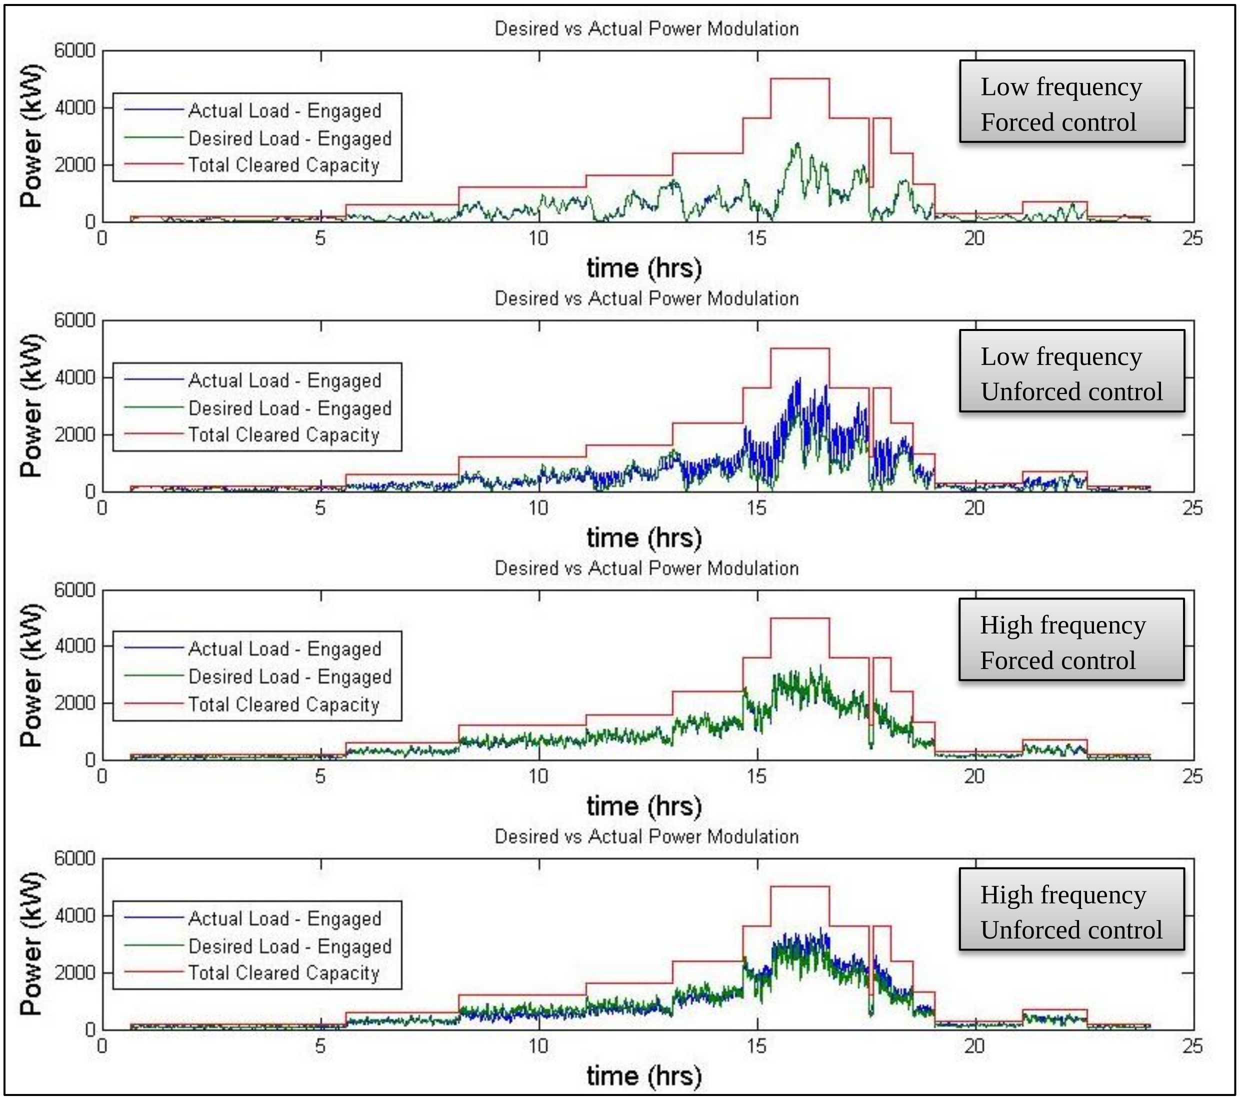

Figure 11: Comparison of the desired versus actual load using forced and unforced.

4.2.3 Market Period Impact

The market period, five minutes, was originally chosen to align with previous experiments for HVACs and real-time energy pricing. This aligns well with the duty cycle of a normally sized HVAC unit. However, it was not obviously clear what impact changing the market period might have on the operations of the control system – could the control system maintain the controllability for fifteen minutes rather than five?

The previous experiments were re-run, this time comparing a 5-minute market cycle to a 10-minute market cycle. The results are shown in Figure 12.

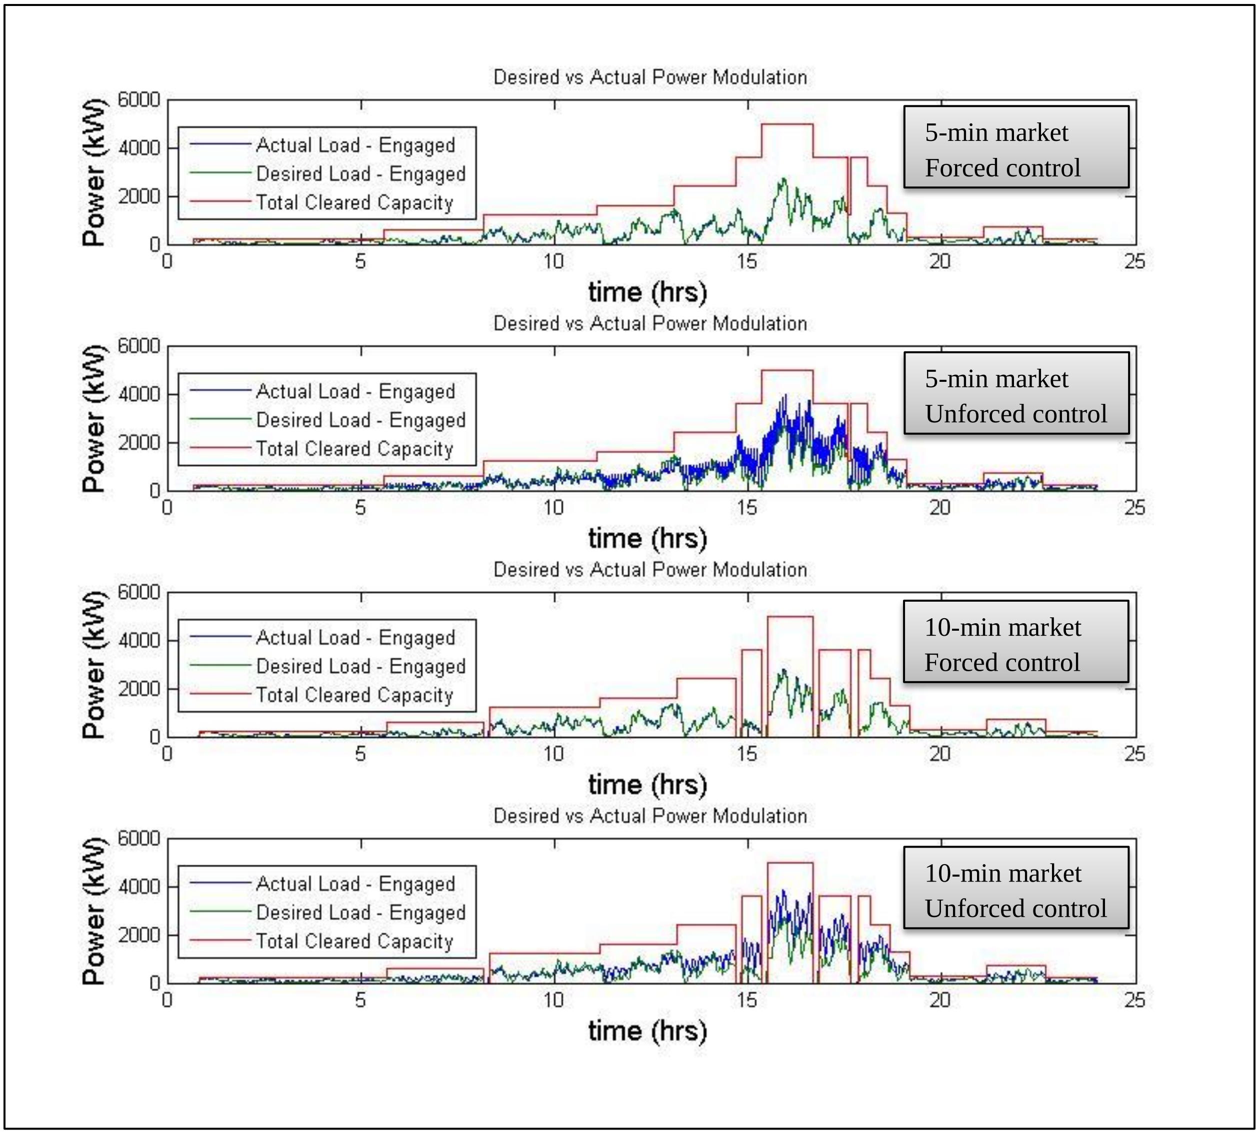

Figure 12: Comparison of the desired versus actual load with different market periods.

There is some impact to the way the market behaves. In the 10-minute market period case, there were a greater number of market periods where the auction could not clear adequate supply. In this case, the inability to clear the market is caused by lack of resources available for participation in the market. Once the market price has been established and the control engaged, however, the results are nearly identical. Mean absolute error from the desired load is the same for the forced case, while slightly decreasing from 14.0% to 13.7% in the unforced case. This indicates that for shorter market intervals, the impact on the control system is minimal.

Figure 13 shows the cleared price of the market over the course of the day. As is clearly seen, there is a difference between the cleared price in the 10-min markets (red and blue) and the 5-min markets (green and purple). It is an indicator that the length of the market period has an impact on the retail market clearing price and the distribution of the device states. As these are initial studies, it is difficult to draw conclusions about which period length is more appropriate or the impact this has on the operation of the market.

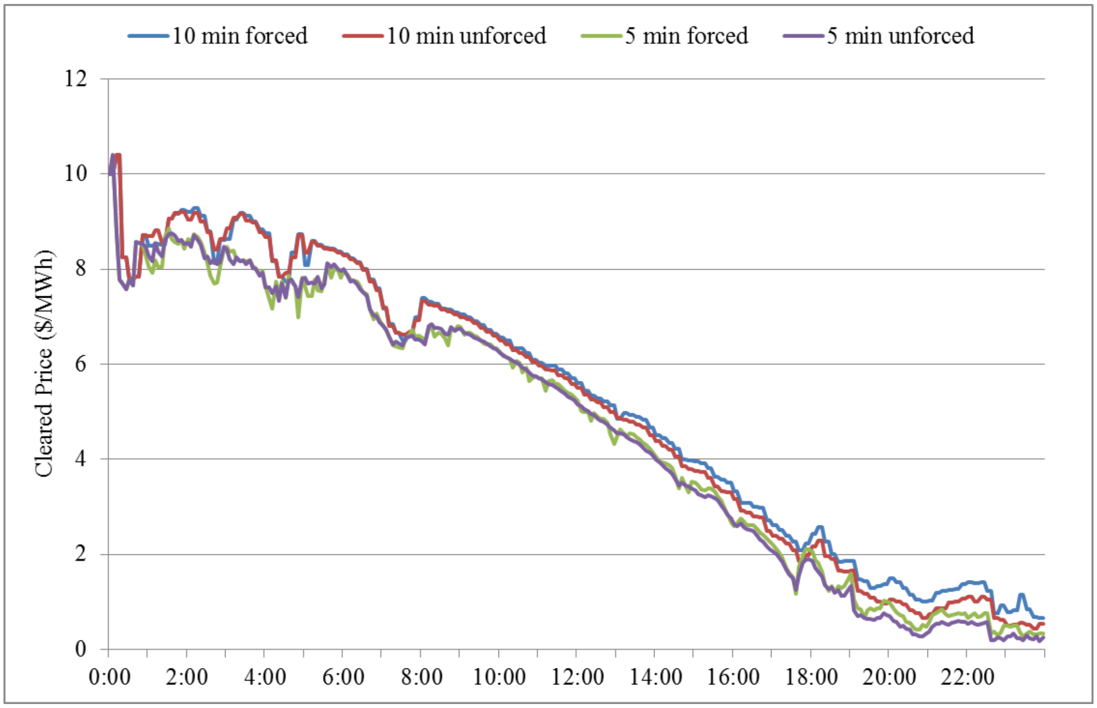

Figure 13: Comparison of cleared market prices using different operating scenarios.

4.2.4 Population Size Impact

This control system relies on a distribution of controller states and therefore relies on there being a statistically large enough population to contribute adequate response. To explore what the size limits are, additional simulations results are shown in Figure 14. The upper two graphics show 5,000 homes with forced and unforced response, respectively. The lower two graphics use 500 homes. Note that 500 homes are participating overall, but at any given time, as few as 25 homes are cleared in the market and participating in the control action. The mean absolute error from the desired output power increases from 2.7% to 7.4% and 14.0% to 17.7% in the forced and unforced cases, respectively. As expected, as the population size decreases, there is greater error. When the signal is this small in magnitude (some of the changes are less than 0.5 kW, which is much smaller than an average HVAC load of 2.5-3.0 kW), there is not enough resolution in the load response to provide the necessary load change. Also, note that these experiments are performed at a constant outdoor temperature that tends to drive the devices to a high duty cycle, increasing the availability. To use this system on "shoulder" days, when relatively few devices are available, the overall size of the population should be such that enough devices can be engaged continuously – this implies that a minimum clearing population will likely need to be specified for the market that depends upon the amount of ancillary service procurement. Further studies are needed to characterize this relationship and estimate appropriate values.

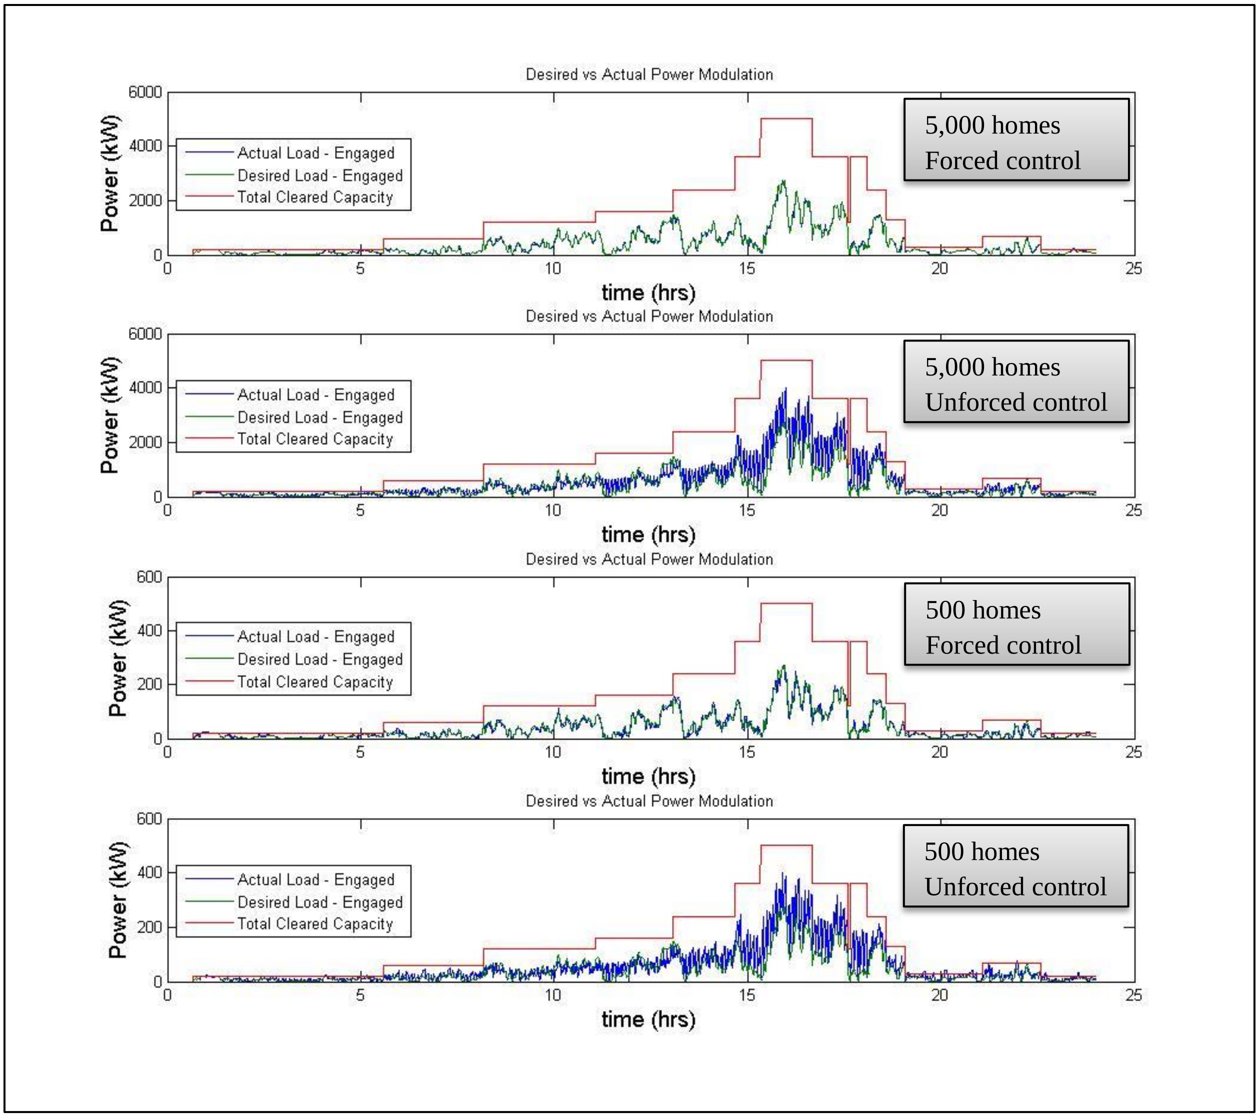

Figure 14: Comparison of the desired versus actual load with differing population size.

4.2.5 Outdoor Temperature Impact

Previous experiments kept the outside air temperature constant to maintain a higher than normal availability of the devices. Simulations were re-run using outdoor air temperatures (and solar radiation) from Columbus, OH during the same period as the ACE signal, summer 2009. As can be seen in Figure 15, the results are similar – all four figures represent the varying weather conditions. The upper two show the response to a high frequency component, forced and unforced, while the second two show the response to a low frequency component. The control system behaves comparably, with a slight decrease in overall error as compared to the fixed condition cases. However, there is a clear difference in the market behavior. The availability of the resources is decreased during the cooler hours, affecting both the amount of demand that can be procured in the market and the cleared market price. This affects how the aggregator might forecast the amount of available resource under different conditions, and how many resources it should expect to acquire to meet wholesale market expectations. This will affect the cost of procuring the resources and how often they are used to support the system needs, but it does not generally affect the control system performance. Future work is needed to understand how aggregators would aggregate this behavior as a wholesale market bid.

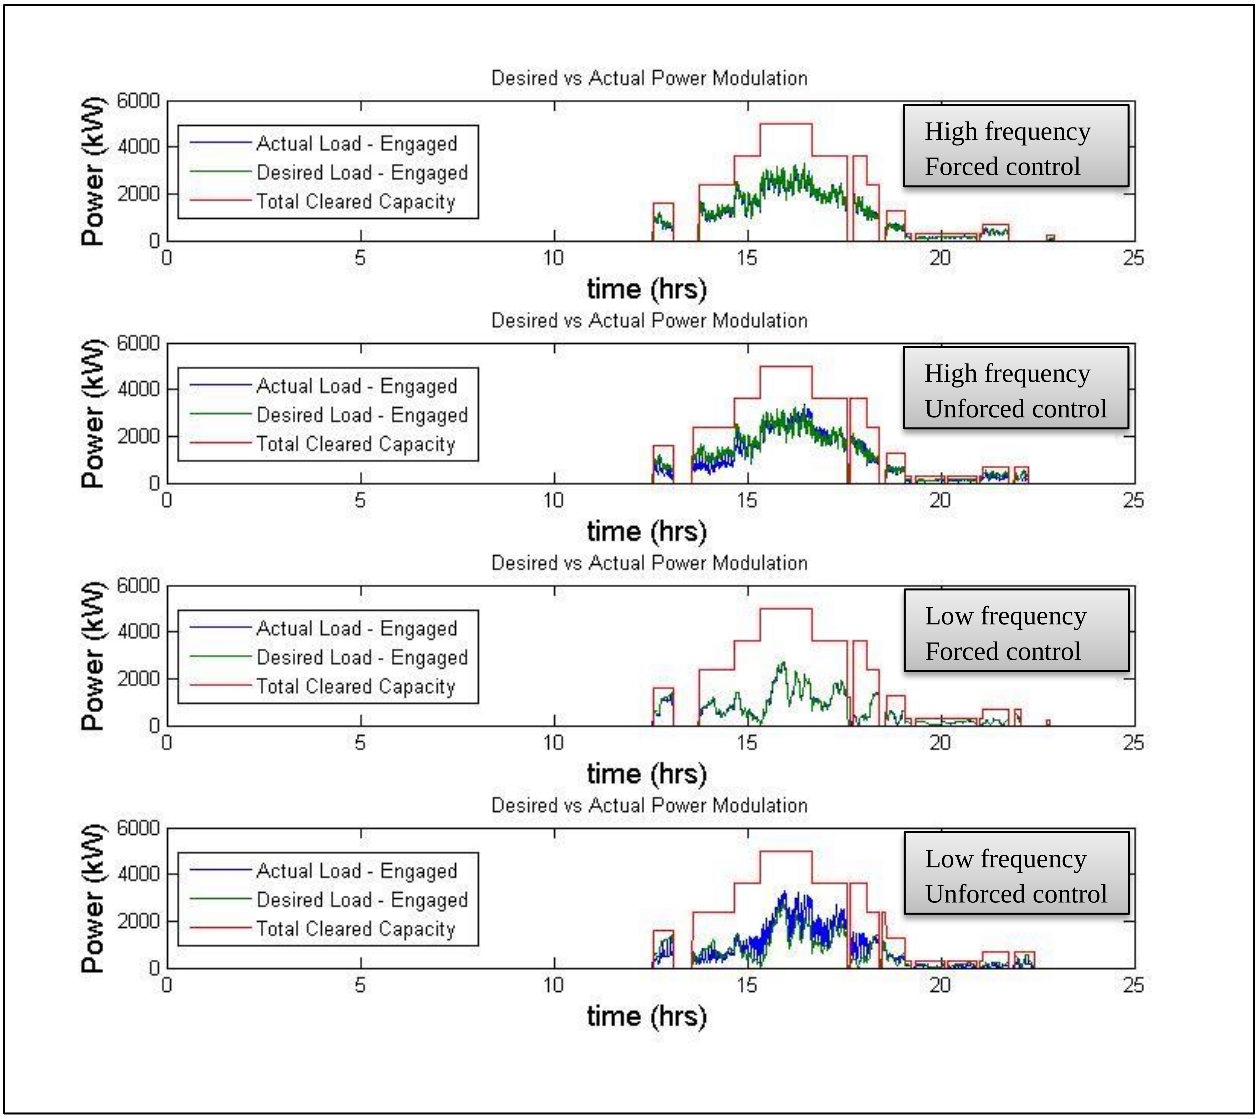

Figure 15: Comparison of the desired versus actual load using varying weather conditions.

Figure 16 shows the percent of devices (out of 5,000) engaged during each market cycle (low frequency, forced control case). At any given time, no more than 30% of the devices are required to participate in the control system, minimizing the impact to the operation of most devices.

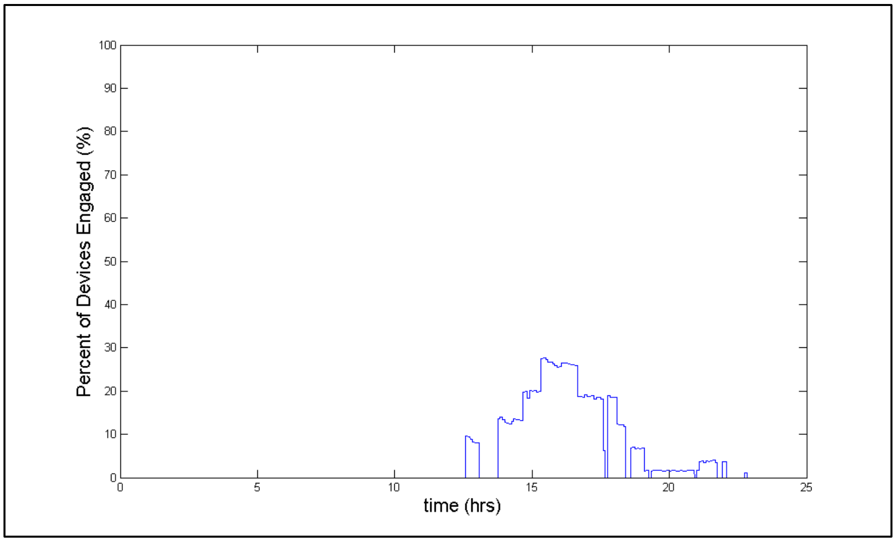

Figure 16: Percent of devices engaged in the market.

Discussion of communication requirements and timing of the control system

The communication requirements for the described regulation control system are more intensive than used in commercially available demand response systems. The system is designed to minimize the amount of two-way communications required (market layer), but still requires very fast one-way communications (device layer).

The Olympic Peninsula Demonstration Project used customer-owned Internet services to support the five-minute market, while the AEP gridSMART RTP Demonstration used cellular modems - both readily available, off-the-shelf technologies. A five-minute market is not required, so trade-offs can also be made in terms of communication costs versus accuracy of the acquisition system. Communication delays in this system may impact the aggregator's ability to secure the right resources at the right time, and may have a financial impact, but should not cause system control issues.

The device-layer controls require significantly faster communication performance, but only as a one-way system with the same signal broadcast to all participating devices. Communication latency or lost packets could cause significant performance issues, which may impact system operations. The timing of the control layer with system characteristics is a necessary feature, and is only possible with reliable communication systems and internal clocks. Most of these can be provided relatively easily through Internet protocols, while other communication technologies may need additional functions to support this. The extent of these issues have not be evaluated. All of the simulations in this study assume perfect, no latency communication channels.

One can also envision that if the devices were only to respond to frequency deviations (i.e., not using the ACE signal), then the frequency could be locally sensed and appropriate control performed. This paper on supervisory control discusses some of these concepts [21].

5 Measurement & Verification

Measurement and Verification (M&V) is the process of quantifying the quality of a service, typically used to verify that a resource is providing the agreed upon service in a quality manner. In the case of regulation services, the response of the device (whether a generator, battery, or aggregate load) should sufficiently track the desired regulation signal. The M&V should be sufficiently accurate, especially when it has contractual and financial implications, but the cost of data acquisition and analysis should be small in comparison to the value of the service. A system operator might want to assess the performance of an aggregator, or an aggregator may want to assess the performance of a residence, an individual device, or a subsystem. Both aggregator level and device level M&V are addressed in the following sections, using data that would likely already be available from the market aggregator, the advanced metering system, and/or utility SCADA systems.

M&V systems are widely used with more traditional resources (i.e., industrial loads or generators), however, few examples exist for addressing the performance of an aggregation of small devices (such as residential HVACs). Interval meter data analysis with a non-intrusive load monitoring method is an approach that has been considered by many, e.g., EPRI [27]. This method depends heavily on the accuracy of the particular hardware/software implementation.

For simplicity of nomenclature, "aggregate level M&V" will be used when the focus is at a more aggregated level rather than an individual device, e.g., at the feeder level or substation level. "Device level M&V" will be used when discussing the ability of a particular device to follow a regulation signal within an aggregated system. Recalling from Section 3.1 that a device participating in a service provides its status as part of the bid, this data is available and used for M&V. Sections 5.1 and 5.2 will discuss M&V approaches using this information. Section 5.3 will discuss a method that uses only SCADA-level substation data. Appendix C contains a discussion about methods beyond the scope of this work, focusing on M&V for energy and ramping services. 'Mileage' will be referred to as the total change in generation or load, both up and down (sum of $|\Delta MW|$), within a given market period, and tracks the total amount of movement.

In any control situation, in particular for stochastic control, it is important to have an understanding of possible control errors. There are two cases where control errors are possible at the system level even if individual devices perform as instructed:

1. The total mileage required is higher than what the recruited devices can deliver. In this case, as discussed in the Appendix A, once the resources are exhausted, no more service is delivered towards the end of the 5-minutes (or whatever the market clearing period is).
2. The smallest increment of service is provided by turning one device on or off. Typical magnitude of the service required should be large enough such that it takes several devices to change state to provide the service. Otherwise, large control errors are possible.

The number of devices in the ensemble should be large enough, and the service expected from the ensemble should be large enough to make control errors acceptable. In the case of unforced control, there is an additional source of error arising from unwanted or premature transitions. All of these can be quantified by further analysis; this is left for the future. We will simply assume that the inherent control errors are small enough, and the purpose of M&V is to assess the performance in providing such service. The unforced case, when the devices are allowed to undergo "natural" change of state, e.g., based on their thermostat settings, results in premature changes of state and/or unwanted multiple changes of state. This will result in control errors. An analysis of such errors is left for future work.

5.1 Device Level

A narrow view of M&V is to assess service on a four-second basis, i.e., measure the power output of each device on a four-second basis. Collecting and transmitting this much data is cost-prohibitive. Additionally, as each device state is binary in nature (i.e., on or off), it is not realistic for an individual device to track an analog signal (such as a regulation signal); only an aggregation of devices should be expected to track the analog signal. In the method described in Section 3.2, the devices only need to track the regulation in a probabilistic manner; in other words, the average response of a device should follow the regulation signal, but at any given four-second interval, this is not necessary. The following section will formulate M&V in a way that does not require four-second data, but still provides an assessment of overall behavior over timescales on the order of hours, days, or longer. The goal is to use data that is already available in the market mechanism (state of a device at the beginning and end of each market cycle) to determine whether a device changed state, and then over longer time intervals determine if it behaved in a statistically similar manner to the requested regulation signal.

5.1.1 Forced Control

Consider a single device in the pool of devices providing regulation during an interval α. For notational consistency, consider multiple devices and add a suffix δ to label a specific device. At the beginning of the interval it is ON. It can be shown that the probability p_α^+ that it is OFF at the end of the interval α is

$$p_\alpha^+ = \min(\,fS_\alpha^+, 1) \tag{1}$$

The quantities f and S_α^+ are defined and discussed in Appendix A.

In the experiment described in Section 4 that includes a population of similar devices, this is the same for all devices in this pool and so a suffix δ is unnecessary. Let $n_{+-\alpha}^{+\delta}$ be the number of state changes from ON to OFF in interval α by device δ. The superscript '+' denotes that the device is participating in providing regulation by turning "on to off", while '-' indicates it is participating from "off to on". For the forced control case, it can only be 0 or 1: if the actual state of the device at the end of the interval is ON, it is 0, and 1 if it is OFF. Considering the expectation values, it can be shown that, for N intervals

$$\sum_{\alpha=1}^{N} n_{+-\alpha}^{+\delta} \;\cong\; \sum_{\alpha=1}^{N} p_\alpha^+ \tag{2}$$

where the symbol ≅ denotes stochastic equality of the expectation value. Further details are given later through a discussion of standard deviation of their difference.

A few comments are in order. If a device does not participate in an interval, it needs to be excluded from the sum during that interval. A device that "wins" the bid in an interval is required to send its state at the end of the interval, whether it bids for the next interval or not. A value of $fS_\alpha^+ > 1$ implies saturation: all devices that can be turned OFF have already been turned OFF, and additional service is unavailable for the last part of the interval. The value of N should be large enough such that the sum should be relatively large compared to the standard deviation of the difference between the two sums.

It can be shown that the standard deviation of the estimate of the sum is $\sqrt{\sum_{\alpha=1}^{N} p_{\alpha}^{+}\,(1-p_{\alpha}^{+})}$. Thus, the previous equation can be modified (using a range of 1 standard deviation on either side - other ranges are possible perhaps desirable - a range of 2 standard deviations is used in some of the graphs) to give the M&V equation

$$\sum_{\alpha=1}^{N} n_{+-\alpha}^{+\delta} \;\cong\; \sum_{\alpha=1}^{N} p_{\alpha}^{+} \pm \sqrt{\sum_{\alpha=1}^{N} p_{\alpha}^{+}\,(1-p_{\alpha}^{+})} \tag{3}$$

The OFF devices can be analyzed similarly. Let $n_{-+\alpha}^{-\delta}$ be the number of state changes from off to on in interval α by device δ participating in providing down regulation. In the forced control case, this can only be 0 or 1. It can be shown that the probability p_{α}^{-} that it is on at the end of the interval α is

$$p_{\alpha}^{-} = \min(-fS_{\alpha}^{-}, 1) \tag{4}$$

and the M&V equation

$$\sum_{\alpha=1}^{N} n_{-+\alpha}^{-\delta} \;\cong\; \sum_{\alpha=1}^{N} p_{\alpha}^{-} \pm \sqrt{\sum_{\alpha=1}^{N} p_{\alpha}^{-}\,(1-p_{\alpha}^{-})} \tag{5}$$

The quantity S_{α}^{-} has been defined in Appendix A. The validity of these equations is checked for each device at the end of each M&V period, e.g., a day. Standard statistical methods can be employed to determine if the equation is satisfied within a margin of error.

Note that Equations 3 and 4 do not directly validate four-second performance of the device. This M&V at the device level is performed integrated over time of the order of a day. If the equations are closely obeyed for all days examined, it serves to validate the performance of that particular device. In this approach, a device is not penalized for not delivering more mileage than it possibly can, i.e., if the aggregator 'asks' for more mileage than the devices stated they could provide, the devices should not be penalized.

Figure 17 shows this method applied to the 500-home case simulated in Section 4. It plots the desired number of state changes (and two standard deviations on either side of the expected due to the stochastics of control) and the actual number of state changes for each of the 500 devices in the market for the forced case using the low frequency signal. The top graph is when the devices participate in the off-to-on control and the bottom on-to-off. The desired number of state changes need not be an integer. Note the slight bias in the off-to-on device, with the devices switching more often than expected, driven by the bias of the low frequency signal – the same bias is not seen in the on-to-off devices, as the low frequency signal tended to be biased in one direction throughout the day.

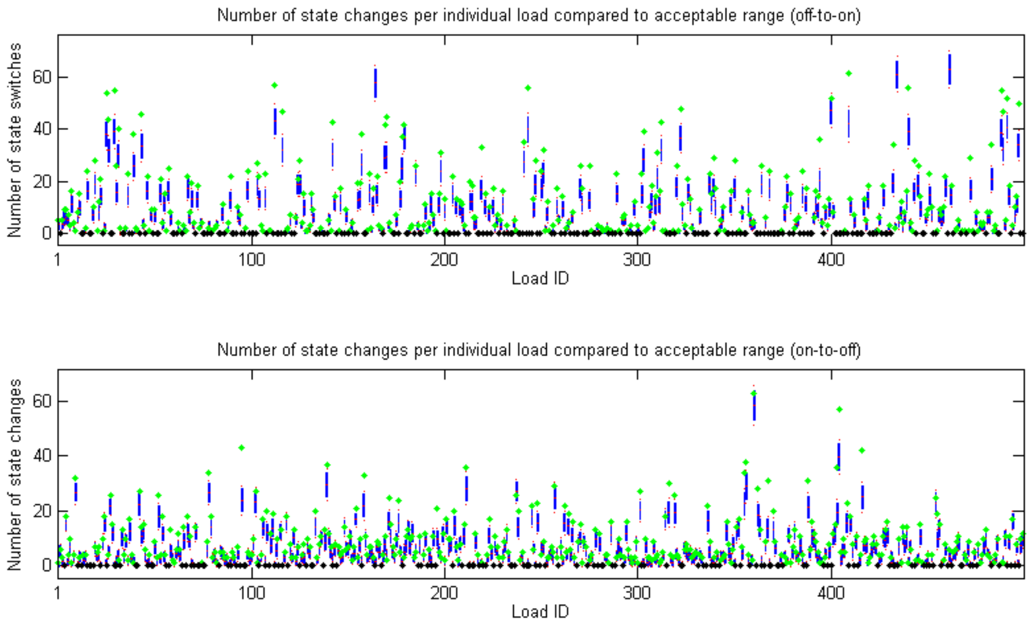

Figure 17: Desired versus expected number of state changes by device.

5.1.2 Unforced Control

In this case, the device may undergo more than one state change during an interval resulting in increased control error. The quantities $n_{+-\alpha}^{+\delta}$ and $n_{-+\alpha}^{-\delta}$ are no longer constrained to be 0 or 1, but can be a larger integer. Furthermore, we need to add additional equations, one for the OFF state

$$\sum_{\alpha=1}^{N} n_{-+\alpha}^{+\delta} \cong 0 \tag{6}$$

and another for the ON state

$$\sum_{\alpha=1}^{N} n_{+-\alpha}^{-\delta} \cong 0 \tag{7}$$

Note that the right hand side is 0 with no statistical error. Thus, the number of unwanted transitions is a measure of the control error. For simplicity, one can combine equations 3 and 6 as well as 5 and 7:

$$\sum_{\alpha=1}^{N} (n_{+-\alpha}^{+\delta} + n_{-+\alpha}^{+\delta}) \cong \sum_{\alpha=1}^{N} p_{\alpha}^{+} \pm \sqrt{\sum_{\alpha=1}^{N} p_{\alpha}^{+} (1 - p_{\alpha}^{+})} \tag{8}$$

$$\sum_{\alpha=1}^{N} (n_{-+\alpha}^{-\delta} + n_{+-\alpha}^{-\delta}) \cong \sum_{\alpha=1}^{N} p_{\alpha}^{-} \pm \sqrt{\sum_{\alpha=1}^{N} p_{\alpha}^{-} (1 - p_{\alpha}^{-})} \tag{9}$$

One can view this as comparing, in each interval, desired up mileage with the actual up mileage plus unwanted down mileage (and similarly for down mileage). The forced case is considerably simpler and more accurate than the unforced case.

5.2 Aggregated Level

Unlike the individual devices, the aggregated level M&V can be performed each bid interval. The goal of this system is to use state information from the bids at the beginning and end of each market cycle to track the overall performance of the group of devices in terms of "mileage". This method does not use power measurements, but rather the state of the devices as a proxy for power.

5.2.1 Forced Control

The expected total mileage from all the devices participating in providing regulation in interval α is

$$P_\alpha \cdot \min(\, fS_\alpha^+, 1) \tag{10}$$

The actual mileage can be determined from the state of the devices at the end of the interval (N_D^+ and N_D^- denote the number of on-to-off and off-to-on devices):

$$\sum_{\delta=1}^{N_D^+} n_{+-\alpha}^{+\delta}\, P_\alpha^\delta \tag{11}$$

P_α^δ is the power consumption when on by the device δ in interval α. By incorporating the standard deviation, two additional M&V formulas are constructed

$$\sum_{\delta=1}^{N_D^+} n_{+-\alpha}^{+\delta}\, P_\alpha^\delta \;\cong\; p_\alpha^+ P_\alpha \;\pm\; \sqrt{p_\alpha^+(1-p_\alpha^+)\;\sum_\delta^{N_D^+}(P_\alpha^\delta)^2} \tag{12}$$

Similarly for off-to-on devices

$$\sum_{\delta=1}^{N_D^-} n_{-+\alpha}^{-\delta}\, P_\alpha^\delta \;\cong\; p_\alpha^- P_\alpha \;\pm\; \sqrt{p_\alpha^-(1-p_\alpha^-)\;\sum_\delta^{N_D^-}(P_\alpha^\delta)^2} \tag{13}$$

Note that Equations 12 and 13 do not directly verify four-second performance. However, if they are obeyed within statistical errors for all the 5-minute intervals examined, it serves to validate the aggregate level performance of the system. In this formulation, the system is not penalized for not delivering more mileage than it possibly can. A simple reformulation can address imposition of such penalty.

Figure 18 shows this method applied to the 500-home case with forced control and using the low frequency signal. The plot shows the desired mileage (and two standard deviations on either side expected from the stochastic nature of the control) and actual mileage for each of the 288 5-minute market periods. The top graph is for off-to-on devices, while the bottom is on-to-off. Note the same bias in the off-to-on devices. The results of these two methods indicates that the 500-home case is likely too small and does not provide enough resolution to properly track the signal.

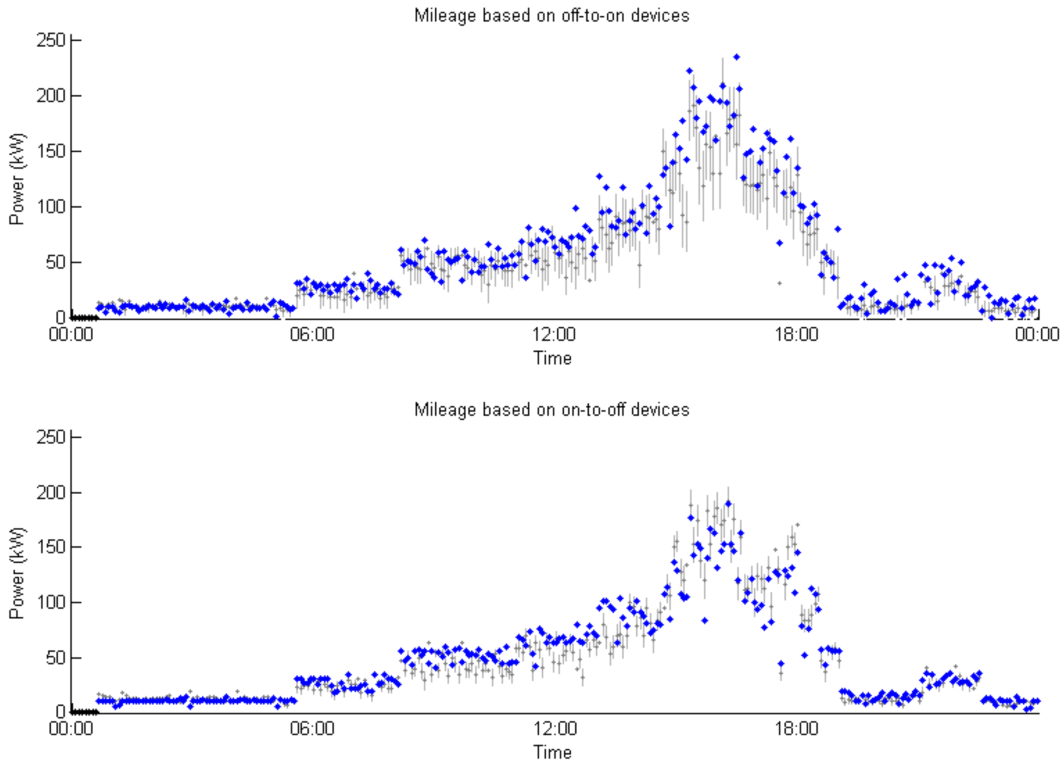

Figure 18: Desired and actual mileage for each of the 288 5-minute market periods.

5.2.2 Unforced Control

In this case, the device may undergo a premature state change or more than one state change during an interval resulting in increased control error. The up and down mileage considerations can be readily extended to this case. It is, however, important to complement this with an analysis of control errors. This is left for future work, and we will restrict ourselves to forced case in the rest of the report.

5.3 Alternate Method from Substation Measurements

The previous methods relied on state data from the individual devices. It also relied on an accurate and consistent translation of device state to device load, which may not always be the case (e.g., a constant impedance device such as a water heater on a distribution system with widely varying voltages will have a constantly changing load). As mileage, rather than four-second measurements were used, there is a possibility of an undetected phase shift in response of the load.

The objective of this approach was to determine if it would be possible to detect the presence of a group of loads following a regulation signal in substation measurements of real power. Typical SCADA systems acquire data at ten-second (or faster) intervals, although they are typically averaged and recorded at longer intervals. It may be possible that an aggregator could use four-second, real-power measurements at the substation to determine the efficacy of the controlled units to track the regulation signal. Note that this requires detecting a relatively small change in load behavior within a relatively noisy measurement, as other loads on the system are continuing to change states in an uncontrolled manner. The 5,000-home GridLAB-D simulations from Section 4 are used to demonstrate this process – in these simulations, at any given time tens to hundreds of devices were participating in the regulation control system, while thousands of devices were not.

To detect the regulation signal, the substation measurements required preprocessing. The measurements before preprocessing for the forced case are presented in Figure 19. To remove the daily trend that is apparent in Figure 19, the data was high-pass filtered with a cutoff frequency of 0.002 Hz. Because the regulation signal specifies a requested change in power, a first order derivative filter was also applied to the data. The preprocessed data is presented in Figure 20. Note that the effects of the regulation signal, plotted in Figure 21 for the forced case, are apparent in the measurement data.

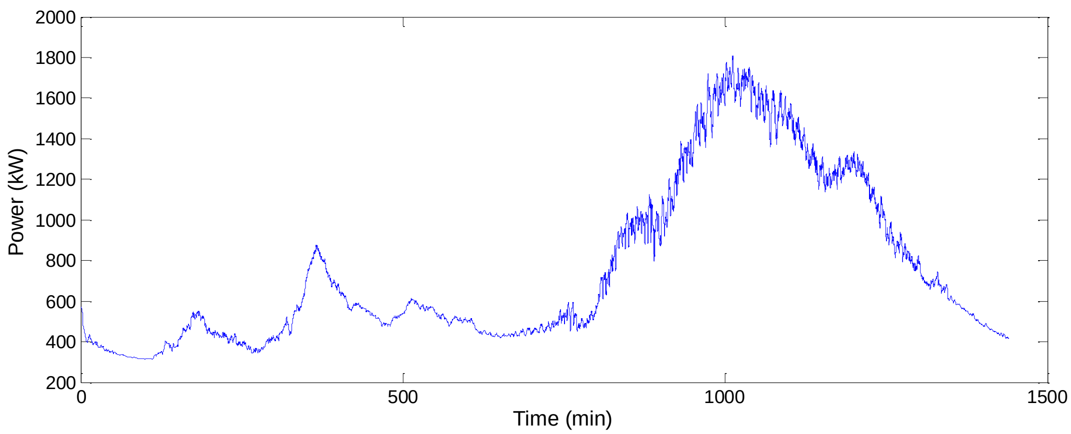

Figure 19: Measured substation power before preprocessing for the forced case.

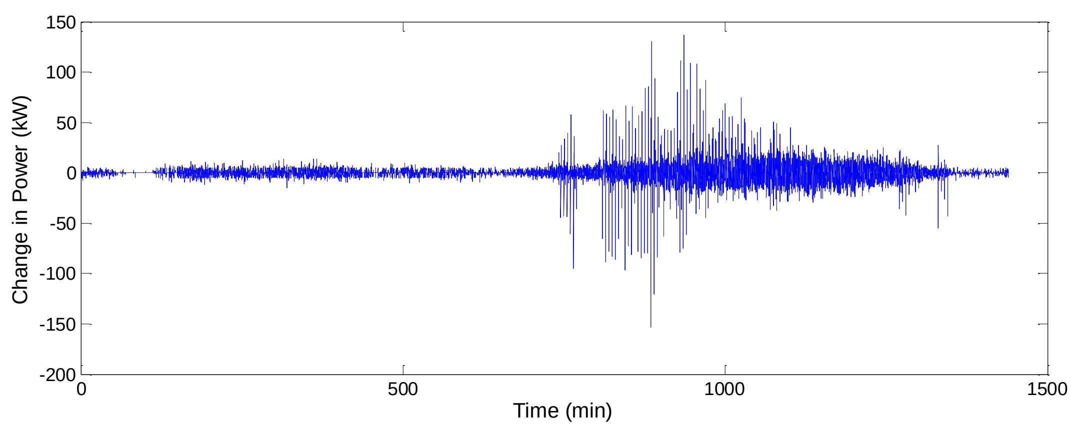

Figure 20: Measured substation power after preprocessing for the forced case.

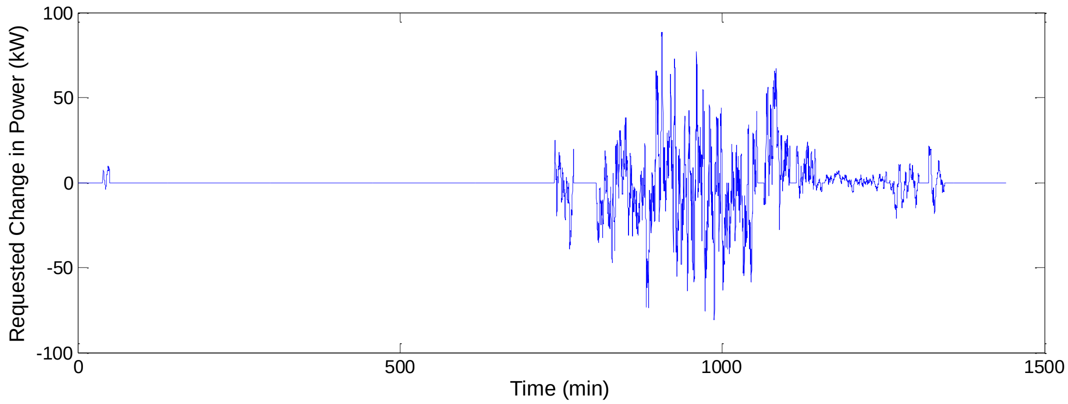

Figure 21: Regulation signal for the forced case.

Optimal detection of a known signal in Gaussian white noise is accomplished using a correlator. Essentially, the cross-correlation of the known signal with the measured data at a zero lag value is compared to a threshold. In this case, an optimal detector, along with its meaningful detection threshold, is not available because the noise is not Gaussian or white. However, the approach can still be applied. The cross-correlation between the preprocessed data and the regulation signal in Figure 21 is presented in Figure 22 for the forced case and Figure 23 for the unforced case. In both cases, the maximum value at a lag of -4 seconds indicates that the signals are correlated. The correlation is significantly stronger for the forced case.

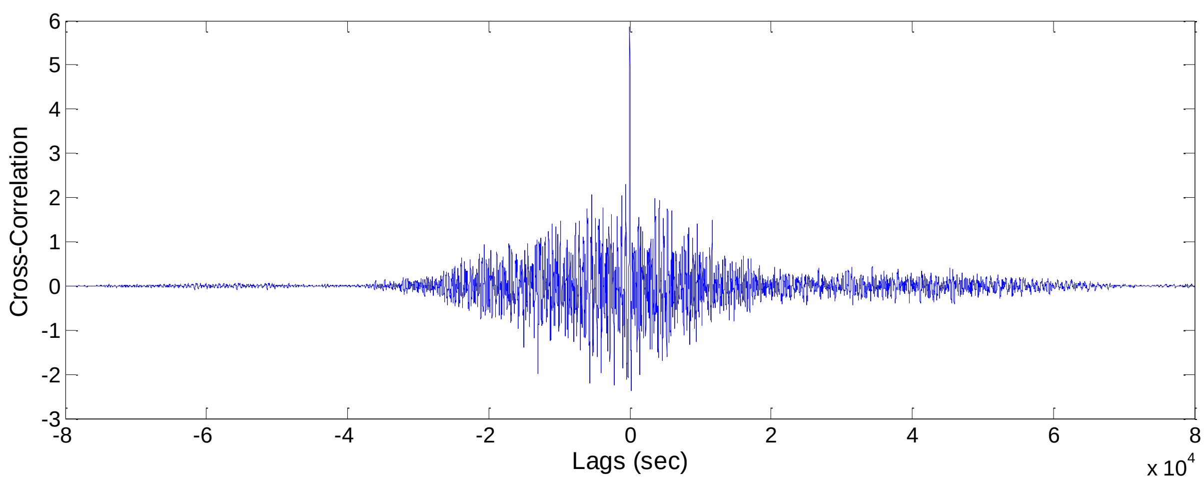

Figure 22: Cross-correlation between the measured data and regulation signal for the forced case.

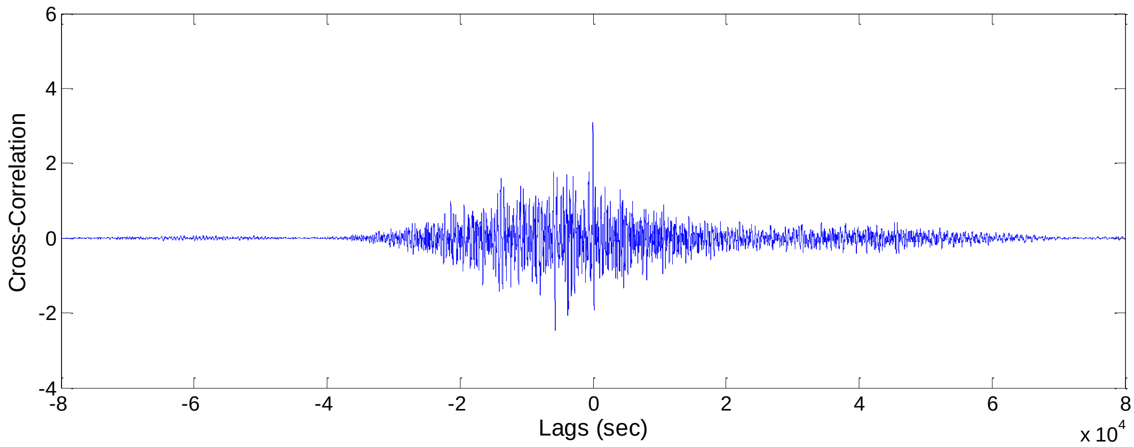

Figure 23: Cross-correlation between the measured data and regulation signal for the unforced case.

The strong correlation seen at a lag of four seconds is due to the strong “presence” of the regulation signal in the substation power data. Further work is needed to develop this into a quantitative tool for M&V, but the results indicate that this may be a promising approach.

6 Summary

In this project, a transactive framework was developed for DERs to provide regulation services. Specifically, a hierarchical market-based control framework is proposed which consists of two parts, one for acquiring resources transactively at a longer timescale (minutes) and a second for controlling the devices in a distributed manner at much shorter timescales (seconds). An aggregator engages a population of thermostatically controlled loads such as HVACs, water heaters, refrigerators and clothes dryers. A double-auction market is then used to acquire the devices every five minutes. After having acquired the contracted devices, the devices respond to a regulation signal broadcast by the aggregator every 4 seconds. Each contracted device controller uses this signal to autonomously decide the operation of the device using a stochastic algorithm. The result is that the system of devices responds to meet the desired regulation quantity.

The proposed transactive control approach was implemented in GridLAB-D and the efficacy of the approach was demonstrated through a variety of simulation studies. It was shown that the control strategy was able to respond to the regulation signal with better accuracy for high frequency cases (frequency component only) as compared to low frequency (frequency and tie-line bias components). Furthermore, allowing for natural state transitions (unforced control) led to a smaller increase in error for the high frequency case, as compared to the low frequency case.

The market period was also increased from 5 minutes to 10 minutes to better understand the impact of availability of resources on control performance. The mean absolute error from the desired load was computed to be the same for the forced case, while slightly decreasing from 14.0% to 13.7% in the unforced case. This indicates that for shorter market intervals, the impact on the control system should be minimal.

It is also noted that reducing the population size introduces larger error. However, the increase in error appears relatively small. It is concluded that there needs to be a minimum population and composition of devices to clear the market to ensure that adequate resources can be engaged continuously.

Finally, the outside temperature was varied to study the impact on the performance of devices to provide the desired regulation service. In this case, the control system behaved comparably well, with a slight decrease in overall error as compared to the fixed outside temperature condition cases. However, the availability of the resources was decreased during the cooler hours, affecting both the amount of demand that could be cleared in the market and the market clearing price.

References

[1] D.S. Callaway, “Tapping the energy storage potential in electric loadsto deliver load following and regulation, with application to wind energy,” Energy Conversion and Management, vol. 50, no. 5, pp. 1389–1400, 2009.

[2] S. Koch, J. Mathieu, and D. Callaway, “Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services,” in Proc. PSCC, 2011, pp. 1–7.

[3] J.L. Mathieu, S. Koch, and D.S. Callaway, “State estimation and control of electric loads to manage real-time energy imbalance,” IEEE Transactions on Power Systems, vol. 28, no. 1, pp. 430 –440, Feb 2013.

[4] S. Kundu, N. Sinitsyn, S. Backhaus, and I. Hiskens, “Modeling and control of thermostatically controlled loads,” in the 17-th Power Systems Computation Conference, 2011.

[5] H. Chao, “Price-responsive demand management for a smart grid world,” Electr. J., vol. 23, no. 1, pp. 7–20, 2010.

[6] H. Shu, R. Yu, and S. Rahardja, “Dynamic incentive strategy for voluntary demand response based on TDP scheme,” in Proc. 2012 IEEE Asia-Pacific Signal & Information Processing Association Annual Summit Conf. (APSIPA ASC), 2012, pp. 1–6.

[7] GridLAB-D, accessed Dec 2014, http://www.gridlabd.org.

[8] D. J. Hammerstrom, et al, “Pacific Northwest GridWise® Testbed Demonstration Projects, Part I. Olympic Peninsula Project,” Pacific Northwest National Laboratory, PNNL-17167, Oct 2007.

[9] S.E. Widergren, et al, “AEP Ohio gridSMART® Demonstration Project: Real-Time Pricing Demonstration Analysis”, Pacific Northwest National Laboratory, PNNL-23192, Feb 2014.

[10] PNWSGD, accessed Dec 2014, http://www.pnwsmartgrid.org.

[11] K. Subbarao, J.C. Fuller, K. Kalsi, R. Pratt, S. Widergren, D. Chassin, “Transactive Control and Coordination of Distributed Assets for Ancillary Services,” Pacific Northwest National Laboratory, PNNL-22942, Sept 2013.

[12] PJM Regulation Performance Senior Task Force, “Performance Based Regulation: Year One Analysis,” accessed Oct 2013, t http://www.pjm.com/~/media/documents/ferc/2013-filings/20131016-er12-1204-004 .

[13] ERCOT Glossary, online, accessed Jan 2015 at: http://www.ercot.com/glossary .

[14] OpenADR, accessed Dec 2014, http://energy.ipkeys.com/wp-content/uploads/2013/11/AEEP_PJM_OpenADR_ATR_Pilot.pdf.

[15] CEA-2045, accessed Dec 2014, http://www.ce.org/Standards/Standard-Listings/R7-8-Modular-Communication-Interface-for-Energy-Ma/CEA-2045.aspx.

[16] EPRI, accessed Dec 2014, http://mydocs.epri.com/docs/Portfolio/P2014/2014_P161.pdf.

[17] J.C. Fuller, S. Ciraci, J.A. Daily, A.R. Fisher, M. Hauer, "Communication simulations for power system applications," Modeling and Simulation of Cyber-Physical Energy Systems (MSCPES), 2013 Workshop on, pp.1,6, 20-20 May 2013.

[18] “Framework for Network Co-Simulation”, accessed Feb 2015, http://ims4.pnl.gov/winmedia/2014/FNCS/FNCS.wmv

[19] DOE Advanced Manufacturing Office, Adjustable Speed Drive Part-Load Efficiency, accessed May 2014, http://www1.eere.energy.gov/manufacturing/tech_assistance/pdfs/motor_tip_sheet11.pdf

[20] K. Kalsi, et al, "Loads as a Resource: Frequency Responsive Demand," Pacific Northwest National Laboratory, PNNL SA-23764, Sept. 2014.

[21] C. Moya, W. Zhang, J. Lian, and K. Kalsi, "A Hierarchical Framework for Demand-Side Frequency Control", American Control Conference (ACC) 2014

[22] GridLAB-D Market Specifications. http://gridlab-d.sourceforge.net/wiki/index.php/Spec:Market. Accessed Aug 2014.

[23] J.C. Fuller and G.B. Parker, "Modeling of GE Appliances: Cost Benefit Study of Smart Appliances in Wholesale Energy, Frequency Regulation, and Spinning Reserve Markets", Pacific Northwest National Laboratory, PNNL-22128, Dec. 2012.

[24] J.C. Fuller, N. Prakash Kumar, C.A. Bonebrake, "Evaluation of Representative Smart Grid Investment Grant Project Technologies: Demand Response", Pacific Northwest National Laboratory, PNNL-20772, Feb 2012.

[25] PJM, accessed Aug 2014, http://pjm.com/~/media/markets-ops/ancillary/description-of-regulation-signals.ashx

[26] PJM, accessed Aug 2014, http://www.pjm.com/markets-and-operations.aspx.

[27] EPRI, accessed Dec 2014, http://smartgrid.epri.com/doc/amidata_presentation.pdf.

[28] J.F. Ellison, L.S. Tesfatsion, V.W. Loose, and R.H. Byrne, "Project Report: A Survey of Operating Reserve Markets in U.S. ISO/RTO-managed Electric Energy Regions," SAND2012-1000, Sandia National Laboratories, Albuquerque, New Mexico, accessed October 1, 2013 at http://www.sandia.gov/ess/publications/SAND2012_1000.pdf.

[29] ENERGY STAR: Residential Refrigerators and Freezers Specification Version 5.0, [Online], Accessed 6/15/2013, Available at: https://www.energystar.gov/products/specs/residential_refrigerators_and_freezers_specification_version_5_0_pdf

[30] "Measurement and Verification for Demand Response", National Forum on the National Action Plan on Demand Response: Measurement and Verification Working Group, http://www.ferc.gov/industries/electric/indus-act/demand-response/dr-potential/napdr-mv.pdf (Accessed December 19, 2014), Miriam L. Goldberg and G. Kennedy Agnew (2013).

APPENDIX A

Regulation Amounts to be Provided

The acquired DERs are expected to provide a share of the total regulation needed by the system. Their relationship is addressed in this section. Consider the PJM normalized 4-sec regulation signal for May 4, 2014 plotted in Figure 24 (from [26]). The actual regulation needed (e.g., in MW) is obtained by multiplying the normalized signal by an appropriate factor. Suppose the share of regulation to be provided by our ensemble is -0.5 MW to +0.5 MW. Suppose we acquire resources that are on at the beginning of the market cycle consuming 0.5 MW, and resources that are off at the beginning of the market cycle that would consume 0.5 MW if they are all turned on. These would span the range 0 to 1 MW. Whether this is adequate or not depends on the characteristics of the signal during the market cycle as we see below.

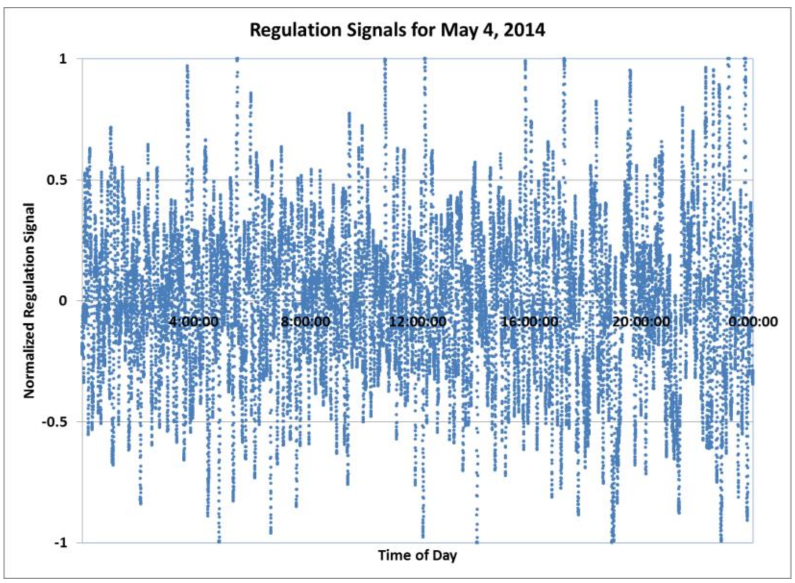

Figure 24: Normalized PJM regulation signal for a 24-hour period

The considerations are applicable to regulation signals for different ISOs. We will explicitly consider one set of signals readily available from PJM. Consider the normalized regulation signals over a day denoted by $r_{\alpha,i}$ *i*=1,…,75 denotes the value at each 4-sec mark within a 5-min interval, and α = 1,…,288 denotes the 288 5-min intervals in a day. Let P_{α}^{+} denote the power consumed by the devices that clear the “up” market. Let the power that would be consumed by the devices that cleared the “down” market, if they are all “on”, be P_{α}^{-}. It is reasonable, but not necessary, to require that the two powers be equal; this allows the combined set of devices to provide equal amount of up and down regulation during the 5-minute period. For notational simplicity, we will require them to be equal and denote the common power by P_{α}. We now need to assign a share of the total regulation to this set of devices. Let us denote this share by *f*. This means that the desired power consumption profile is given by

$$P_{\alpha}(1 + f.r_{\alpha,i}).$$

The aggregator needs to ensure that these profiles, summed over all sets of devices (of which our set is just one) participating in regulation is equal to the amount of regulation the aggregator is expected to provide to the grid.

An appropriate value of *f* should be determined. At first sight, it might seem that setting it to 1 will provide the needed regulation with capacity to spare since $r_{\alpha,i}$ are generally significantly less than 1 in most intervals. What is important is that the required up "mileage" can be provided by the devices that cleared the up market, and similarly the down market. This is analogous to hiking through a series of peaks and valleys resulting in total uphill and total downhill walks that may have little to do with net elevation change. Whenever $r_{\alpha,i+1} - r_{\alpha,i}$ is positive, certain number of "on" devices should change state to "off" and vice versa. Thus the power consumed by all the "on" devices at the end of the 5-min period is given by

$$P_\alpha . \max((1 - fS_\alpha^+), 0)$$

where

$$S_\alpha^+ = \sum_{i=1}^{75} (r_{\alpha,i+1} - r_{\alpha,i})_+ \tag{15}$$

The subscript + indicates that only positive values are kept. The quantity S_α^+ is a characteristic of the regulation signal and can depend on various factors, such as season, time of day etc. Similar considerations apply for regulation "down" case: The power consumed by all the "off" devices at the end of the 5-min period is given by

$$P_\alpha . \min(-fS_\alpha^-, 1)$$

where

$$S_\alpha^- = \sum_{i=1}^{75} (r_{\alpha,i+1} - r_{\alpha,i})_- \tag{16}$$

and a corresponding quantity S_α^- can be defined. The 288 values of S_α^+ and S_α^- are plotted in Figure 25 for May 6, 2014 with PJM data [26].

If we set *f*=1, then for the 5-minute intervals outside the -1 to +1 band in Figure 25 there are not enough devices to provide the needed regulation. This is roughly 20% on either side. A value of *f* less than one essentially moves all points closer to 0, and so reduces the number of points outside the (-1,1) band. A sufficiently small value of *f* ensures that all points are inside the band. As *f* decreases, so does the share of regulation provided by our set of devices and consequently utilization of the devices. Thus the aggregator has to optimize the value of *f*.

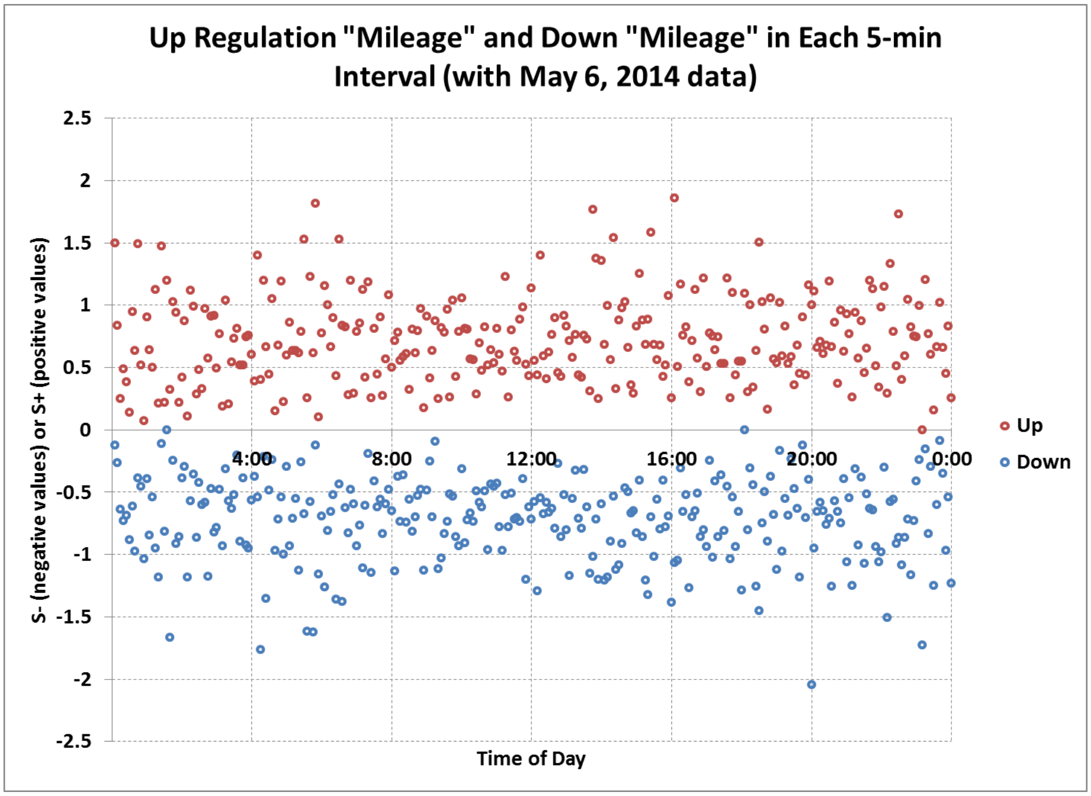

Figure 25: Up and Down Regulation Mileage over 5 minutes for a 24 hour period

Figure 26 shows the duration curves for $S^{+}$and S^{-}. Each day during the one-week period May 4-10, 2014 contributes one plot on the positive side and one on the negative side. All 7 days show qualitatively similar characteristics, except for a few outliers; these outliers can only be satisfied by setting *f* to be correspondingly low. We have not studies whether there are any seasonal trends. It is also necessary to combine the characteristics of the regulation signal with those of devices. Different classes, e.g., HVAC units and resistance water heaters may have different performance characteristics, based on which a different value of *f* may be assigned to different assert classes. Studies of this type enable an aggregator to formulate acquisition and dispatch strategies for these resources.

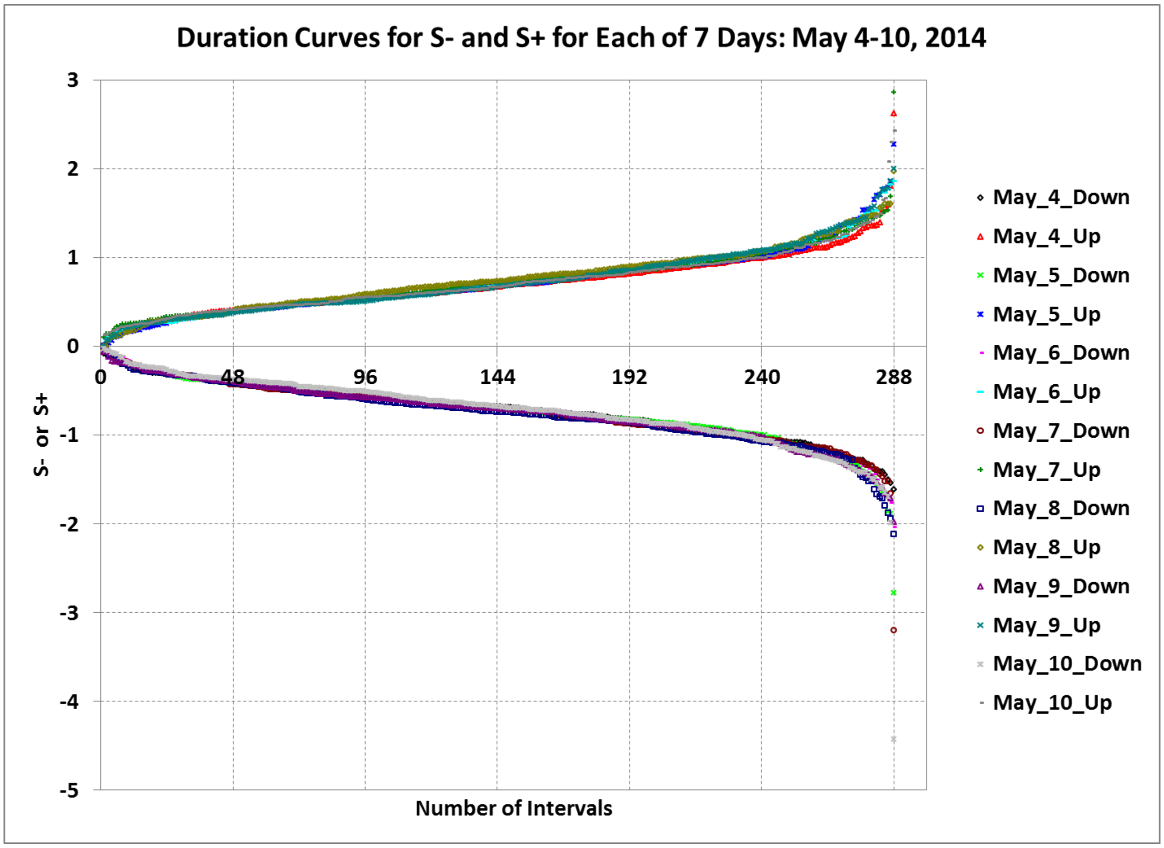

Figure 26: Duration curves for the total up and down regulation mileage in each 5-min interval, over a one week period.

In the GridLAB-D simulations discussed, the case of $f = 1$ is discussed in detail, with the expectation that during certain market cycles, the regulation needs will not be fully met. The market mechanisms and resource acquisition are general.

APPENDIX B

Detailed Description of Refrigerator State Model

In this appendix, certain characteristics of refrigerators are discussed. These are important in putting into context the approximations made in the Markov chain formulation, as well as possible future extensions of simulations to multistate systems as well as to relax the requirement of allowing only forced transitions.

The largest power and energy demand of a refrigerator is the compressor that cycles *on* and *off* to maintain the internal air temperature near a pre-set value. In some models, the compressor may be dual-speed, low for normal operations and high to pre-cool the refrigerator before a defrost cycle. In addition to the additional compressor load, the defrost cycle activates heaters around the cooling coils to melt accumulated ice on the coils. In addition to this main process, there are a number of sub-processes, depending upon model, make, and age of the equipment. These may include anti-sweat heaters for eliminating moisture from the outer shell, fans for moving air from one compartment to another, ice makers, lights, and power electronics. While very dependent on size of unit, manufacturer, efficiency rating, etc., Table 2 shows average or rough approximations of the power demand of standard processes within a refrigeration unit and average times they might be in a given state.

Table 2: Approximation of demand of individual processes within residential refrigerator.

State	Demand (W)	Avg Time On (m)
Base	~25	~45
Compressor (low) interruptible (+ Base)	~100	~60
Compressor (low) uninterruptible (+ Base)	~100	~120
Compressor (high) (+ Base)	~150	~60-75
Defrost (+ Base)	~400	~20-30
Sweat Heaters (+ Base)	~10	~10
Ice Maker (+ Base)	~100	~3

To a certain degree, consumer interaction with the appliance affects the behavior, particularly in that the duty cycle of the compressor and frequency of defrost cycle increases as consumers open the door more often or place more hot food into the cavity. This can be seen in Figure 27 as the daily load shape is roughly uniform across the time of day, with a slight increase in demand during the dinner period. This may be useful, as this means that the amount of resource available from refrigeration is somewhat independent of the time of day.

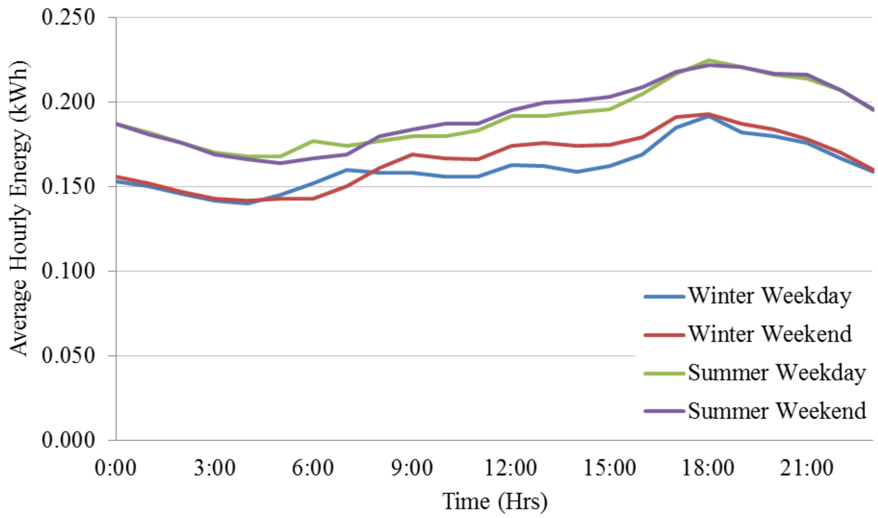

Figure 27: ELCAP daily energy consumption pattern for refrigerators, corrected for efficiency improvements.

Some components of the cycle are not interruptible. For example, to interrupt the compressor shortly after it has started on a consistent basis without taking into consideration the current runtime of the compressor may cause damage or excessive wear and tear on the units. Once a defrost cycle has started, it should not be interrupted. If the temperature of the refrigerator cavity climbs to an unsafe level, the compressor cycle cannot be interrupted. Consequently, the entire unit cannot be used for performing ancillary services by turning the unit *off* and *on*. Therefore, for the purposes of this discussion, it will be assumed that the refrigerator is a "smart" appliance with the ability to control the individual processes within the device and ancillary services cannot effectively be performed on retrofitted refrigerators. In addition, it is assumed that this technology is not driven by human interaction and will be, for the most part, automated.

ENERGY STAR Program Requirements Product Specifications for Residential Refrigerators and Freezers – Eligibility Criteria Version 5.0 discusses compliance for DR functionality [29]. In this document, it discusses a number of energy reduction methods that may be required for refrigerators to be DR compliant. All ice making must be deferrable for a minimum of four hours, although longer is acceptable, provided the ice making process has not already been started. Additionally, it is expected that pre-cooling and defrost cycles (Delayed Defrost) can also be deferred for four hours providing the cycle has not already begun. For shorter time periods (10-15 minutes), the device must be able to reduce demand by 50% from *baseline* operations (note that this does not mean *current* operations). Southern California Edison comments to these requirements indicate that response for longer events may be on the order of 10-20 W per unit and as high 100 W per unit for shorter time periods. Studies with GE Smart Appliances indicate that for spinning reserve type events, ~40% of load may be available for short-term reduction whereas ~30% may be available for a single, instantaneous regulation event [23]. The ability to increase load was not discussed in either.

The refrigerator operation is described in terms of a simplified state transition diagram in Figure 28.

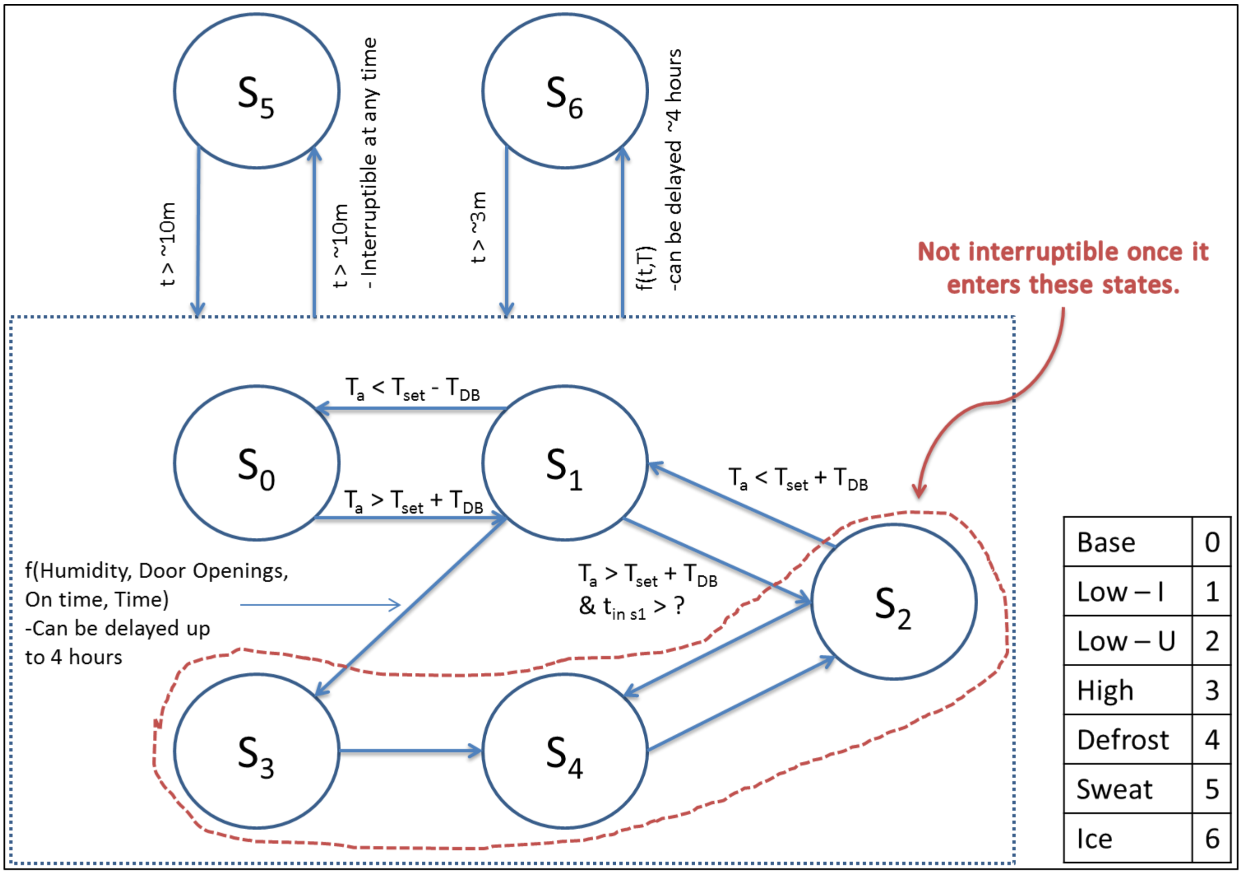

NOTE: T_a is a representation of the freezer and refrigerator air temperature. For the most part, this cycle is driven by the freezer air temperature, so for these purposes it will be the freezer temperature.

Figure 28: State transition diagram for a representative "smart" refrigerator.

However, this model is reduced into even simpler states to align with the proposed Markov Chain model combined with bidding strategies that reveal a certain subset of behavior when bidding a certain action. These states are broken into a *bidding* state and a *control* state, where the bidding state describes what processes are currently available for modification and hence defines what control actions are available. Table 3 defines the bidding states.

Table 3. Defined bidding states for a smart refrigerator.

Bidding State Number	Bidding State	Load Available for Shifting (Up or Down)
B1	Normal thermostatic	Compressor
B2	Defrost desired	Defrost cycle / High compressor*
B3	Sweat heater	Sweat heater
B4	Ice maker	Ice maker

*The defrost cycle is only available for "up" or increasing of the load.

Note that B1 and B2 are mutually exclusive; if the defrost cycle is desired then the device is deferring the demand and hence is not available for further control. This may be an oversimplification that reduces the availability of B1 at certain times, but is used as an assumption in this setup. Bidding states 3 and 4 are independent from all other states; they can occur no matter what is occurring with B1 and B2, or each other.

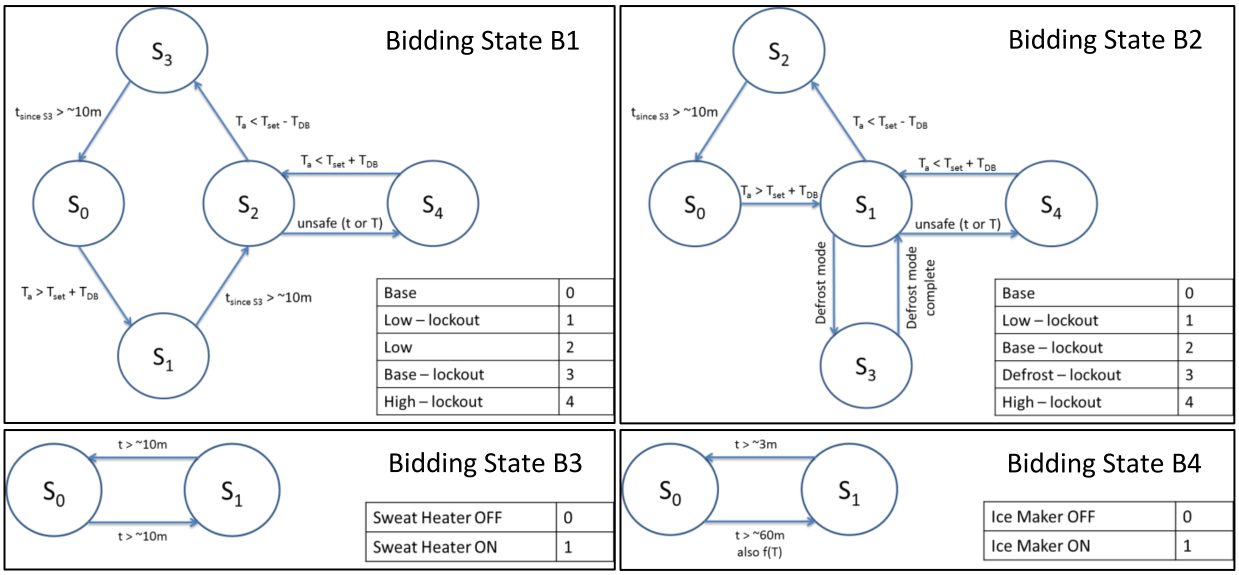

Figure 29: Simplified state diagram and bidding states for a smart refrigerator.

In B1, the refrigerator can bid either its availability to turn on the compressor (S_0 -> S_1) or its availability to turn off the compressor (S_2 -> S_3). Once it enters S_1 or S_3 it is unavailable for changing its behavior for a pre-specified amount of time to protect over-cycling of the compressor. It is assumed that this value can be determined by the manufacturer, will be a different value for each model, and will be unknown to the Markov Chain-based control system; for the purposes of this experiment, 10 minutes will be used. Note, this value is greater than the proposed 5 minute market cycle, so should not affect the operations within the 5 minute bidding window. Additionally, manufacturer safety controls will not be overridden. For example, if the refrigerator cavity is at an unsafe temperature requiring high cooling capacity, then the device will enter the "locked out" state, S_4, and cannot participate in load shifting.

Bidding state B2 is similar to B1. It will bid its availability to turn off the compressor (S_2 -> S_0) or its availability to activate the defrost cycle (S_1 -> S_3), which has a higher demand than just the compressor over an extended period. Note, entering the defrost cycle is a locked state with a series of controlled events (high cooling, defrost coils, and high cooling) that will not be interrupted once started. However, the start of the series of events can be delayed for a number of hours.

Bidding state B3 is available to bid both turning on and off the sweat heater, within certain restrictions on runtime. It is assumed that manufacturers will be able to define this in terms of minimum and maximum runtimes. For the purposes of this experiment, an approximate duty cycle of 50% with an average runtime of 10 minutes will be used.

Bidding state B4 is available to bid turning on the ice maker operation only. There is considerable flexibility in turning on the device when the operation is started. Once the cycle is started, it only takes a few minutes and is not available for load reduction.

APPENDIX C

Energy Services

Of the ancillary services considered in this report, M&V for regulation has been discussed in previous sections. M&V for the other ancillary services – ramping and spinning reserve – is considered in this section. The methods can be broadly classified as follows depending on the data requirements and analysis methods:

1. When the status of a device is known at the beginning of a bid interval and the number of transitions in the interval, and the number of desired transitions (0 or 1) is known, M&V is straightforward. Even in cases where a specific profile within the bid interval is desired (e.g., a linear ramp), the method studied in detail in this report is applicable and is in fact simpler since the desired profile is known ahead of time.
2. From interval meter data and a non-intrusive load monitoring method, M&V is possible. We will not pursue this in this report
3. From interval meter data, M&V can be done in manner similar to that for demand response. The latter has been studied extensively; a recent review can be found in [30]. A baseline is created and compared against energy use during events.
4. The method under item 3 can be employed with feeder level data for M&V of performance at feeder level. We will consider this in some detail below, by using very different baselining methods, and consider enhancements to the analysis of responses that provide a more insightful representation of the response.

Aggregate Response

Starting with energy use data at the feeder level or aggregated over of multiple premises, a baseline is constructed in one of two ways:

a. using a control group (primarily applicable for aggregation over multiple premises), and normalizing the control group to track experimental group as closely as possible during non-event periods
b. using a model (no control group needed): energy use vs outdoor temperature, with independent regressions for each period (hour, 15-mins or 5-mins) further segmented by daytype.

Using method (a), and data from a utility in the Pacific Northwest Smart Grid Demonstration Project, the kW/customer change during 173 DR events each lasting about 3 hours is plotted in Figure 30. Electric water heaters responded to these DR events. The mean change as well as standard deviation of the changes observed during an event are plotted. There were many more events during the first few months in the plot than during the rest of the period. It is more instructive to plot the same data differently: cumulative kWH from the events vs cumulative customer-hours. Customer-hours instead of just hours (representing the duration of the event) is used to account for possibly varying number of customers during each event. The resulting plot is shown in Figure 30.

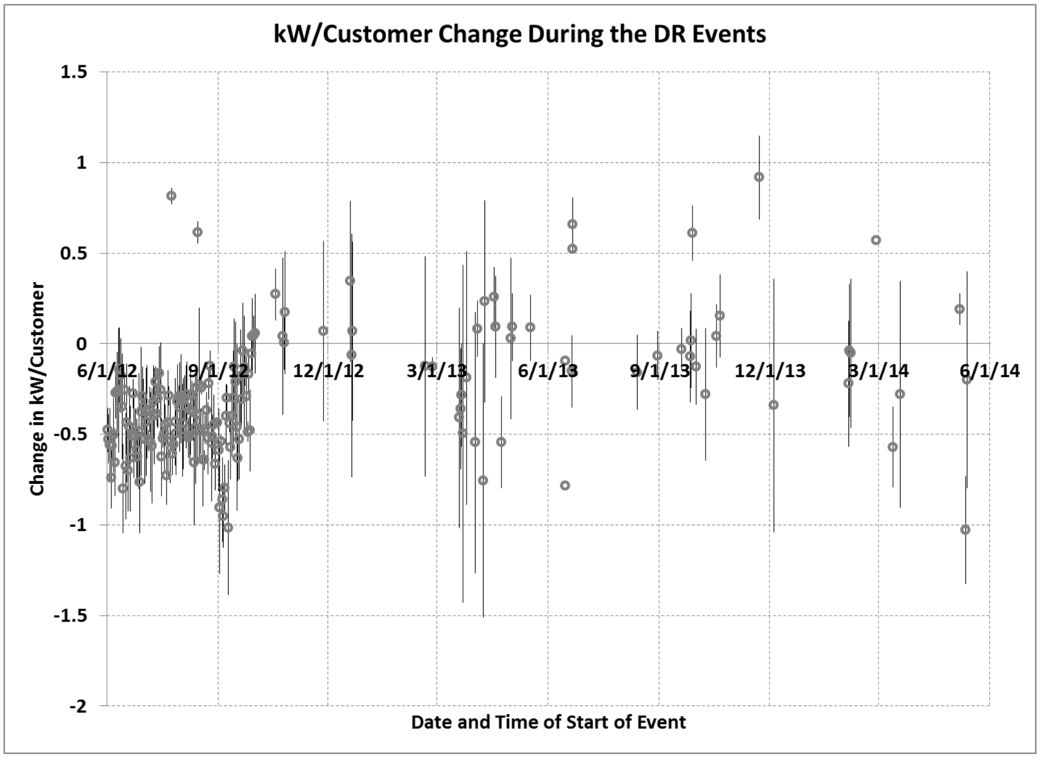

Figure 30: The change in kW/customer during each of the 173 DR events.

The slope of the curve represents kW/customer change. A relatively constant change results in a straight line. The first part of the curve shows a steady reduction of -0.421 ± 0.002 kW/customer during the event (the standard error is as reported by a linear regression over the relatively linear part of the curve). During the later events, the performance is uneven. From this, quantitative measures of performance can be developed.

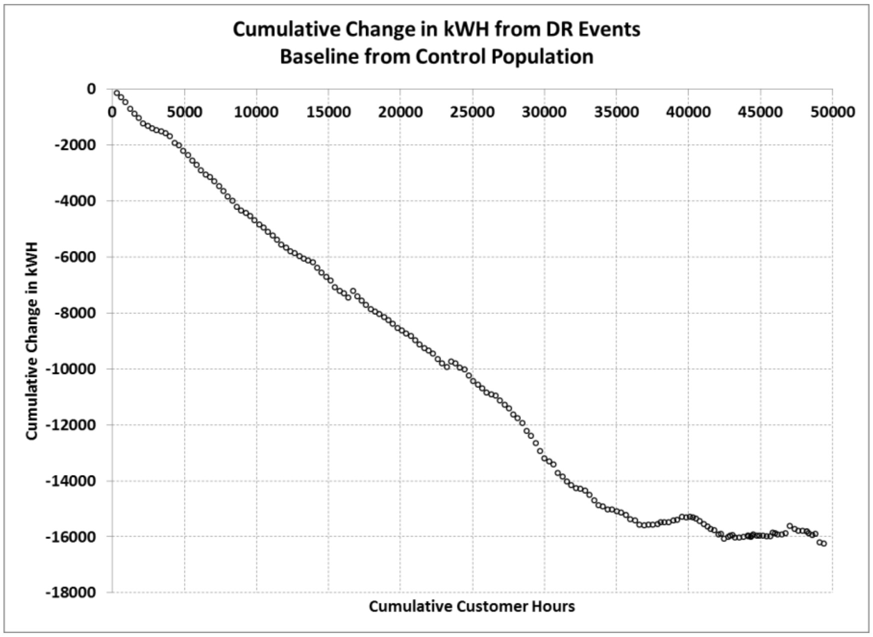

Figure 31: Cumulative change in kWh vs cumulative customer hours; baseline is from a control population.

The same responses, analyzed with temperature normalization are shown in Figure 32. The slope of the relatively linear portion of the curve is -0.576 ± 0.004 kW/customer. The two methods, although based on very different normalizations, give similar results. Because of differences in the availability of data, the numbers of points in the two curves are different, and so there is no one-to-one correspondence of points in the two graphs.

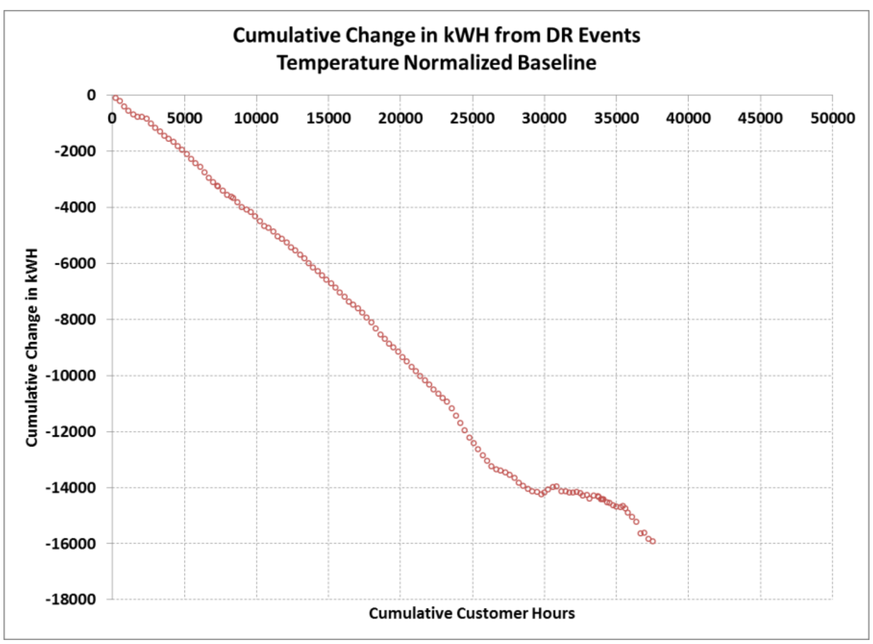

Figure 32: Cumulative change in kWh vs cumulative customer hours; baseline is through temperature normalization.

Distribution

No. of Copies		No. of Copies	
1	Dan Ton U.S. Department of Energy 1000 Independence Ave, SW Washington, DC 20585	3	Local Distribution Pacific Northwest National Laboratory Steve Widergren (K1-85) Robert Pratt (K1-85) Dale King (K9-69)

Pacific Northwest
NATIONAL LABORATORY

U.S. DEPARTMENT OF
ENERGY

Proudly Operated by **Battelle** *Since 1965*

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99352
1-888-375-PNNL (7665)

www.pnnl.gov