U.S. DEPARTMENT OF ENERGY

PNNL-23192

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

AEP Ohio gridSMART® Demonstration Project Real-Time Pricing Demonstration Analysis

SE Widergren
K Subbarao
JC Fuller
DP Chassin
A Somani
C Marinovici
JL Hammerstrom

February 2014

Pacific Northwest NATIONAL LABORATORY

Proudly Operated by Battelle Since 1965

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

Use of the Terms Partner and Partnership

The use of the words “partner” and “partnership” herein is intended to indicate a mutually beneficial collegial relationship, rather than a partnership as defined under state law.

PNNL-23192

AEP Ohio gridSMART® Demonstration Project Real-time Pricing Demonstration Analysis

SE Widergren
K Subbarao
JC Fuller
DP Chassin
A Somani
C Marinovici
JL Hammerstrom

February 2014

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

Pacific Northwest National Laboratory
Richland, Washington 99352

Summary

This report contributes initial findings from an analysis of significant aspects of the American Electric Power, Ohio (AEP Ohio) gridSMART® Real-Time Pricing – Double Auction (RTP_{da}) demonstration project (the Project). Over the course of four years, Pacific Northwest National Laboratory (PNNL) worked with Ohio Power Company (the surviving company of a merger with Columbus Southern Power Company), doing business as AEP Ohio, and Battelle Memorial Institute to design, build, and operate an innovative system to engage residential consumers and their end-use resources in a participatory approach to electric system operations, an incentive-based approach that has the promise of providing greater efficiency under normal operating conditions and greater flexibility to react under situations of system stress. The material contained in this report supplements the findings documented by AEP Ohio in the main body of the gridSMART report. It delves into three main areas: impacts on system operations, impacts on households, and observations about the sensitivity of load to price changes.

The RTP_{da} system operated from December 2011 through the fall of 2013. An adequate population of households for system experiments was achieved in the late spring of 2013. As air conditioning equipment was the only type of load under RTP_{da} control, and as the great majority of this equipment only operated in cooling mode, the period of analysis was set from June 1 to September 30, 2013. The system was designed to collect a large amount of operational data, including the status of the enhanced programmable communicating thermostat (ePCT) parameters, indoor temperature, household energy consumption, and the RTP_{da} market data (such as household and supply price and quantity bids, market cleared price, and total distribution feeder[1] load). This was supplemented by data from the meter data management system, the billing system, weather data, and demographic data about the households (such as square footage and type of construction).

As with any operating system, the data are incomplete and testing behavior can pose challenges. Gaps in data from communications errors, equipment failures, and the like offered challenges to the analysis and add a level of uncertainty to the findings. In addition, the investigation found that the Project's plans to frequently exercise the RTP_{da} system with congestion events imposed on the households to observe their response under different circumstances (for example, days of week, times of day, and temperature conditions). A congestion event occurs when the load level of a distribution circuit (otherwise known as a feeder) exceeds the capacity limit of the feeder. Operators can impose a congestion event by setting the capacity limit below the present load level. This causes the market clearing process to drive prices higher. The investigation found that by frequently imposing congestion events, the resulting high prices desensitized the response of the equipment to normal market fluctuations when not in a congestion event. Once problems such as these were uncovered, the analysis attempted to compensate for their impacts, and thus come closer to a more accurate picture for addressing the questions under investigation. Simulations of the RTP_{da} system were also performed to help address some of these challenges and to scale the household resources to a size that allows for the investigation of system impacts.

[1] The term "distribution feeder" refers to the electric line that feeds a community of houses and terminates in a distribution substation. This is also known as a distribution circuit, as used elsewhere in the AEP Ohio gridSMART Demonstration Project Report. For the sake of brevity it is referred to simply as a feeder in this document.

The findings confirm the basic premise correlating reduction of short-term energy use with price increases and conversely, increase in energy use with price decreases. From a system impact point of view, simulations show that with a 35% penetration of RTP_{da} households, a load reduction of about 5% can be obtained for a 3.5-hour system peak event. For a 2-hour local, feeder peak event, a nearly 8% load reduction can be obtained. Regarding the impact on 5-minute wholesale energy purchases, the field data analysis indicates that, if there were no congestion events, overall energy consumption by the average RTP_{da} household could be reduced by over 5% and wholesale costs could similarly be reduced by 5% compared with the average non-responsive control group household. Simulations of the same wholesale impacts report an average of 1.2% reduction in energy consumption per household and 2.5% reduction in wholesale energy costs.

Consumer impacts studied include household bills, their thermostat statistics, and the actual energy use of the air conditioning equipment. When the RTP_{da} households' bills are computed using the RTP_{da} tariff versus the standard tariff, the study shows that there is good dispersion of relatively minor increases and decreases across all household energy use levels. Average monthly bills decrease slightly using the RTP_{da} tariff, thanks largely to the incentive savings. When investigating the average RTP_{da} bill compared to a calculation of the average bill of the non-responsive control group on the standard tariff, the analysis indicates about 5% reduction in the average RTP_{da} household bill, with a slight increase in overall energy usage. The components that appear to contribute to the average bill reduction are the incentive payments from the frequent congestion events and the flexibility to alter energy use in response to market price fluctuations. The energy usage is not reduced as reported in the wholesale energy purchases above because the congestion events are not excluded in this analysis as they were for the wholesale purchases analysis. Simulations indicate a roughly 4% savings in RTP_{da} bills versus the same households on the standard tariff that are not responding to price signals and incentives.

A study of thermostat settings shows a wide variety of settings by consumers with some indications of clusters, such as those who prefer more comfort and those who balanced comfort and economy more. To study consumer learning patterns, their behavior would need to be monitored for a longer period of time that included multiple seasons. The congestion events indicate that only 4% of the consumers overrode their thermostat setting at some point during the 2-hour events, whereas 10% of the consumers overrode their thermostat setting at some point during the 4-hour events. This provides some verification of consumer fatigue that would need careful attention in operating such programs. Lastly, the amount of energy bid in the market for the air conditioning units appears to have been underestimated from the observed energy draw on these units. The amount of energy bid into the real-time market should be more accurate in a full-scale deployment.

To analyze the sensitivity of load to price changes, the energy data measured at 5-minute intervals for each household was correlated with the corresponding 5-minute wholesale market information. Though the distribution of individual household responses is quite scattered, a filtering of the information corroborates the expectation that energy use decreases when price increases. This is particularly pronounced during hot periods when there is a great deal of air conditioning load operating in the presence of high, but fluctuating, energy prices. In addition, an analysis of the RTP_{da} household response to congestion events (resulting in high market prices) shows a strong dependence on outside temperature and the timing of the events (for example, peak versus off-peak periods). These factors affect the amount of energy curtailment initially available from the population of RTP_{da} resources, as well as the subsequent response of these resources to maintain, degrade, or enhance curtailment levels over the duration of the event. The findings contained in this report are termed *initial* because they only begin to address some of the questions about the operation of the RTP_{da} system. Due to the complex nature of interactions between consumers and the electricity system, and the complexity of electric system operations in general, many more questions arise about the performance and potential benefits of this approach. The data gathered as a result of this project will be of significant value for further research.

Acronyms and Abbreviations

AEP Ohio	American Electric Power, Ohio
CAISO	California Independent System Operator
CPP	critical peak pricing
DOE	U.S. Department of Energy
ePCT	enhanced programmable communicating thermostat
ERCOT	Electric Reliability Council of Texas
HEM	home energy manager
HVAC	heating, ventilation, and air conditioning
kW	kilowatt(s)
kWh	kilowatt-hour(s)
LMP	locational marginal price
MDM	meter data management
MSE	mean-squared error
MWH	megawatt-hours
P_{base}	base price of the supply curve
P_{cap}	price cap
P_{clear}	cleared price
PDF	probability distribution function
PJM	PJM Interconnection, LLC, AEP Ohio's Regional Transmission Organization
PNNL	Pacific Northwest National Laboratory
Q_{clear}	cleared load/quantity
RTP	real-time pricing, also used as a shortened version of $\mathrm{RTP}_{\mathrm{da}}$
$\mathrm{RTP}_{\mathrm{da}}$	real-time pricing, double auction
SMART Shift Plus^{SM}	a form of critical peak pricing implemented in the GridSMART Project
SRMCP	synchronized reserve market cleared price
$T_{desired}$	desired temperature
T_{max}	maximum temperature
T_{min}	minimum temperature

Contents

Figures

Tables

1.0 Introduction

The gridSMART® Real-Time Pricing– Double Auction (RTP_{da}) Demonstration Project (the Project) is a part of the consumer-oriented projects within the overall American Electric Power, Ohio (AEP Ohio) gridSMART program. This project engages residential households to adapt their electricity use in response to a fluctuating 5-minute price signal. In particular, heating, ventilation, and air conditioning (HVAC) units are managed by intelligent software in the home that interacts with a real-time electricity market. The electricity supply price is function of the PJM wholesale market price of electricity as described in the real-time tariff approved by the Public Utility Commission of Ohio (Schedule RS-RTP, 2012).

Significant effort went into the specification, design, development, and deployment of the RTP_{da} demonstration so that AEP Ohio, the U.S. Department of Energy (DOE), and the Project partners could learn from the experience of this innovative approach to engaging end-use systems to the benefit of the consumer and the service provider. The analysis of the RTP_{da} demonstration addresses the question, "What did we learn from the RTP_{da} experiment?" The topics for analysis were developed by the AEP Ohio and Pacific Northwest National Laboratory (PNNL) Project team and each organization was given responsibility for a portion of the topics. This report covers analysis topics assigned to PNNL. Other analysis topics related to the RTP_{da} demonstration, such as customer satisfaction, are covered elsewhere in the gridSMART Project report.

While this analysis report provides insights into the behavior of the RTP_{da} system and its implications for service providers and consumers, it represents only a step on a path to discovering the characteristics and capabilities of end-use systems to participate in system operations and how to best engage them for consumer, local, and regional system objectives. Where appropriate, the report provides perspective for the results and lists additional issues that still need to be addressed.

1.1 Analysis Objectives

The RTP_{da} demonstration represents the first time that a real-time electricity market with an approved regulatory tariff has operated in a realistic situation of approximately 200 households. These households are supplied by four distribution feeders and represent a small fraction of the roughly 2000 total number of households on these feeders. While the measurements on the HVAC systems in these households provide good data to help quantify their price-responsive behavior, the penetration level is too low to address other analysis questions that require significant penetration levels of RTP_{da} households. For this reason, simulations of a higher penetration of RTP_{da} households are needed. Once calibrated to behave similarly to actual household loads, simulations can be configured to provide insights into questions that would be difficult and costly to address in the demonstration.

This analysis report investigates the following areas:

- the potential benefits of RTP_{da} for system-capacity and feeder-capacity issues
- the potential benefits of improving wholesale purchases in the real-time (5-minute) market and participation in a spinning reserve market

- the impacts of RTP_{da} from the consumer's perspective, including consumer bills and consumer configuration of the thermostat set point and adjustments of it over time
- a characterization of the sensitivity of the RTP_{da} loads to price fluctuations and their behavior when called upon for system events.

This analysis report explains the approach taken for the investigation, the source of the information, and the results obtained. Other areas of analysis, such as the implications of RTP_{da} in overall energy consumption or customer satisfaction with the program offering, were done by AEP Ohio. This document supplements that other analysis.

The analysis also includes the results of simulations of RTP_{da} households. While the measurements on these HVAC systems in these households provide good data to help quantify their price-responsive behavior, the penetration level is too low to address other analysis questions that require significant penetration levels of RTP_{da} households. For this reason, simulations of a higher penetration of RTP_{da} households are needed. Once calibrated to behave similarly to actual household loads, simulations were used to provide insights into questions that would be difficult and costly to address in the demonstration.

1.2 RTP_{da} Theory of Operations

The following sections describe the way in which the RTP_{da} system operates. This provides a context for better understanding the analysis results described in this report. The theory of operations starts with a description of how the distribution feeder market works. This is followed by a high-level description of the RTP_{da} dispatch system and how the end-use devices interact with this this system.

1.2.1 Market Operations

The RTP_{da} system follows a transactive-control approach to coordinate household equipment participation in system operations. The term "transactive control" refers to a distributed decision-making approach that allows suppliers and consumers of energy to arrive at a coordinated solution for how each participant will operate based upon a trade-off of the value they place on electricity for a specified time. In this case, an energy market is used to resolve which HVAC loads will run in the next operating interval. The design combines

- a 5-minute retail RTP_{da} reflecting PJM wholesale locational marginal price (LMP) and capacity values
- an RTP_{da} tariff designed to be revenue neutral for the average consumer prior to any load shifting induced by the rate, and with the intent to robustly protect the consumer and the utility from long-term fluctuations in market prices
- a retail double-auction market design that directly manages congestion (that is, limits that constrain the amount of load served) at the distribution feeder level
- a retail market design capable of managing a share of congestion occurring at levels in the grid above a distribution feeder (for example, transmission), allocated to responsive load served by that feeder

- an economically rational heating/cooling thermostat design that balances a consumer's desire to save on their electric bill in exchange for their willingness to be flexible, and that bids the price at which the load it controls will operate (or not), plus the quantity of that load
- a price-normalization scheme that eliminates the need for a consumer to understand or specify price levels as (for example) high, medium, and low, and that adapts to both short-term (days) and long-term (years) changes in price.

The following sections present the operational objectives driving the Project and the design incorporating the elements listed above.

1.2.1.1 RTP$_{da}$ Market – Uncongested Conditions

A double-auction market implements a mechanism to determine the price at which supply and demand match at a given time. Bids are collected for a specified period of time from market opening to market closing, after which the market is cleared. The market clears every 5 minutes, a period that approximately matches the typical air conditioning load cycle.

After the market clears, the cleared price (P_{clear}) becomes the new prevailing retail RTP$_{da}$ and the cleared load (Q_{clear}) varies with the demand curve. When the cleared price is published, devices can respond appropriately based on internal price-response logic. The auction itself does not provide any bookkeeping or enforcement of the price-response logic. It simply provides a central facility for buyers and sellers to deliver their price and quantity response information and obtain the prevailing RTP. The following figure shows the feeder supply curve, the ordered demand curve of bids for energy from the RTP$_{da}$ households, and the market clearing at the intersection of these two curves.

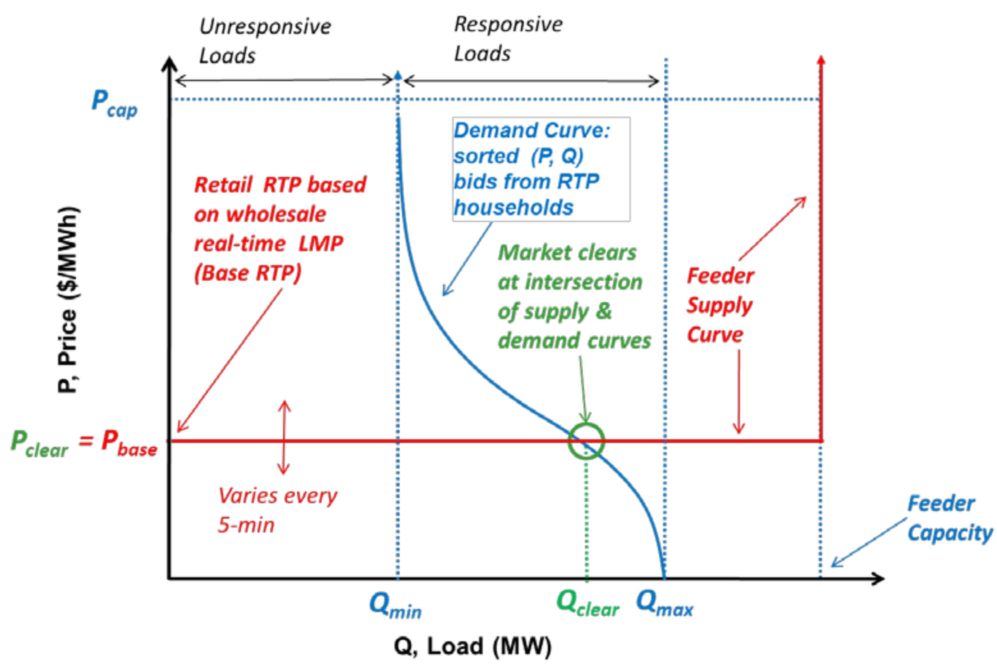

Figure 1.1. RTP$_{da}$ Market Clearing – Uncongested Condition

1.2.1.2 RTP$_{da}$ Market – Distribution Congestion

Congestion reflects feeder capacity limit constraints or system-wide operational constraints whose resolution could benefit from load reduction. Congestion can be addressed by allocating load reduction at the distribution feeder level in proportion to household bids on that feeder. By participating in a reoccurring market mechanism to negotiate energy need by willingness to pay, participants' actions can dynamically mitigate congestion limits.

During congestion, the cleared price (P_{clear}) is greater than the 5-minute, price of supply (P_{base}), and the cleared quantity (Q_{clear}) equals the feeder capacity, as shown in the Figure 1.2. Every time period (5 minutes), P_{clear} varies in order to try to keep load at the feeder capacity. The market auction proceeds as follows:

- The cleared price (P_{clear}) is set to clear the total load (Q_{clear}) at feeder capacity.
- When congested, $P_{clear} > P_{base}$.
- P_{clear} varies every 5 minutes to try to keep load at feeder capacity.
- If there is an inadequate amount of responsive load to hold the feeder capacity, the market will clear at its limit, that is, the price cap, P_{cap}.
- As shown in Figure 1.2, the total load on the feeder can theoretically vary between a minimum (Q_{min}) and maximum (Q_{max}).

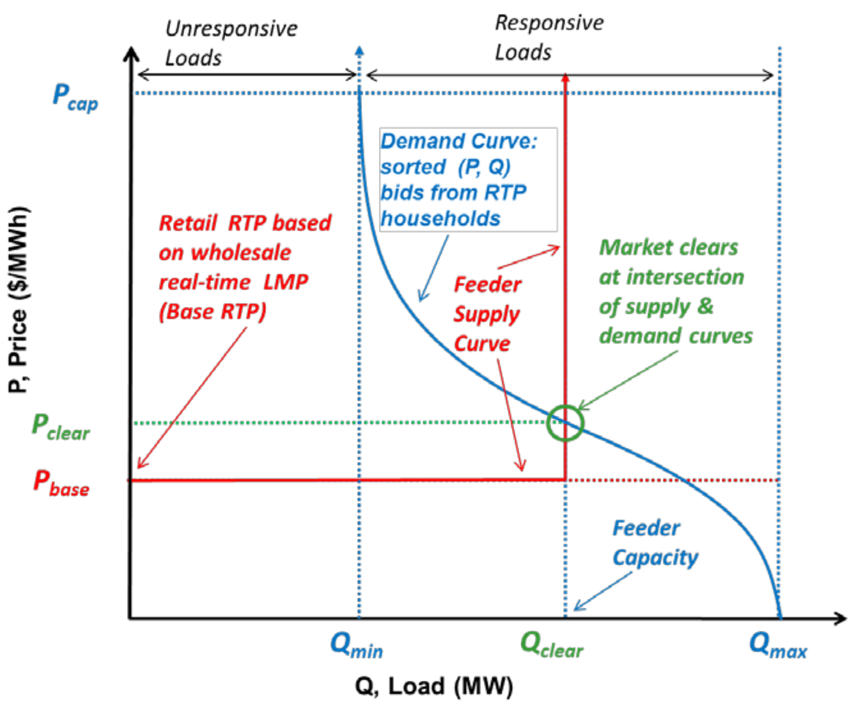

Figure 1.2. RTP$_{da}$ Market Clearing – Congested Condition

The RTP$_{da}$mechanism can also be used to address system congestion issues, that is, issues that the service provider may have with system capacity constraints. In that case, an overall system load reduction target is desired. To accomplish this, a share of the system-wide load reduction target can be allocated to each distribution feeder's households in proportion to their price and quantity bids. The number of RTP_{da} households per feeder and the feeder load itself play important roles in keeping Q_{clear} below the feeder capacity.

1.2.1.3 RTP_{da} Market – Rebate and Incentive Mechanisms

The market clearing at a higher price during congestion events encourages the bidding equipment to curtail operations. This allows consumers to avoid paying a high market price for energy; however, their exposure to higher prices was done to benefit the system, not to make the price-responsive consumer pay more for energy than a flat-rate consumer. The excess payment of the RTP_{da} consumers due to the higher cleared price than the base supply price (P_{base}) is indicated in Figure 1.3 as the congestion surplus (in the figure, only the households in the RTP_{da} program are represented). This surplus is either rebated back to the consumers at the end of the month, or equivalently, the consumers are only charged the P_{base} price even though the market cleared above this price. As the service provider did not experience added costs from associated wholesale price increases, and was able to avoid other, higher priced solutions (such as purchasing generation), the congestion surplus represents revenue for the service provider. Returning the congestion surplus back to the consumer removes the unfair burden of charging price-responsive consumers more, when in fact they are helping the service provider to avoid more costly alternatives.

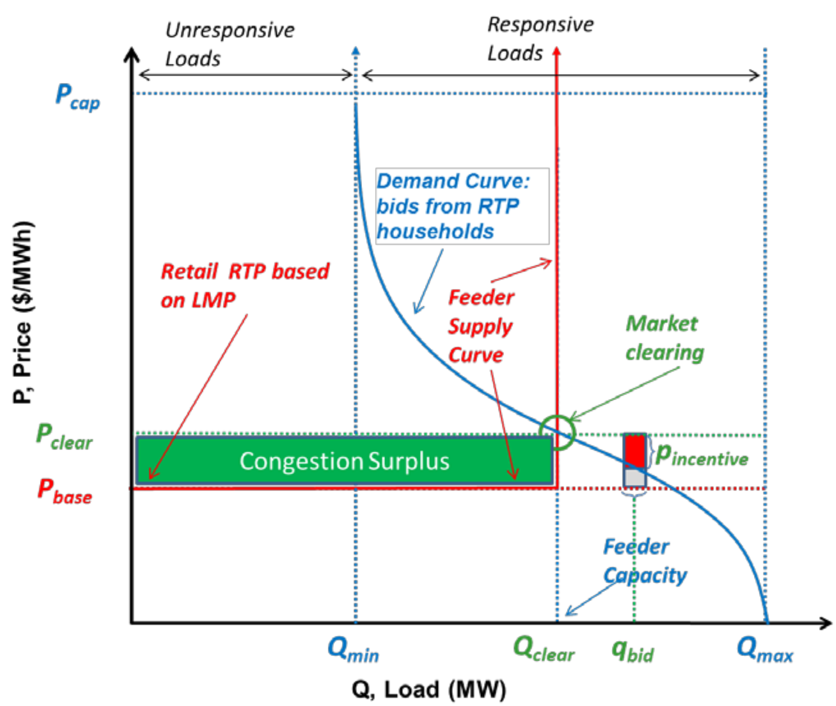

Figure 1.3. RTP_{da} Congested Condition – Congestion Surplus Rebate and Incentive

In fact, the rebate of the congestion surplus does not include any percentage of the added, long-term benefit that system operation achieves by reducing or moving a peak load condition (for example, deferring distribution system infrastructure upgrades). The value of this long-term benefit can be shared with the RTP_{da} households that actually reduced their load by using an incentive mechanism. Several alternatives were reviewed by the Project team as to how an incentive could be calculated. The team chose an algorithm meant to reward consumers who are the most flexible to price changes. Figure 1.3 shows an example of the incentive provided to a consumer with the bid (p_{bid}, q_{bid}). The incentive is computed as the quantity of energy consumed (that is, the product of the bid and the 5-minute time interval) times a function of the difference between cleared price and RTP_{da} base price.

If $P_{base} \leq p_{bid} \leq P_{clear}$,

then $p_{incentive} = q_{bid} \times \mathrm{F}(P_{clear} - P_{base})$.

If not, then no incentive is applied.

1.2.2 Dispatch System

The RTP_{da} system runs an electricity market on a distribution feeder-by-feeder basis. For the demonstration, four markets are running simultaneously, one for each of four feeders that supply the participating households. A simplified drawing of one of the markets is depicted in Figure 1.4.

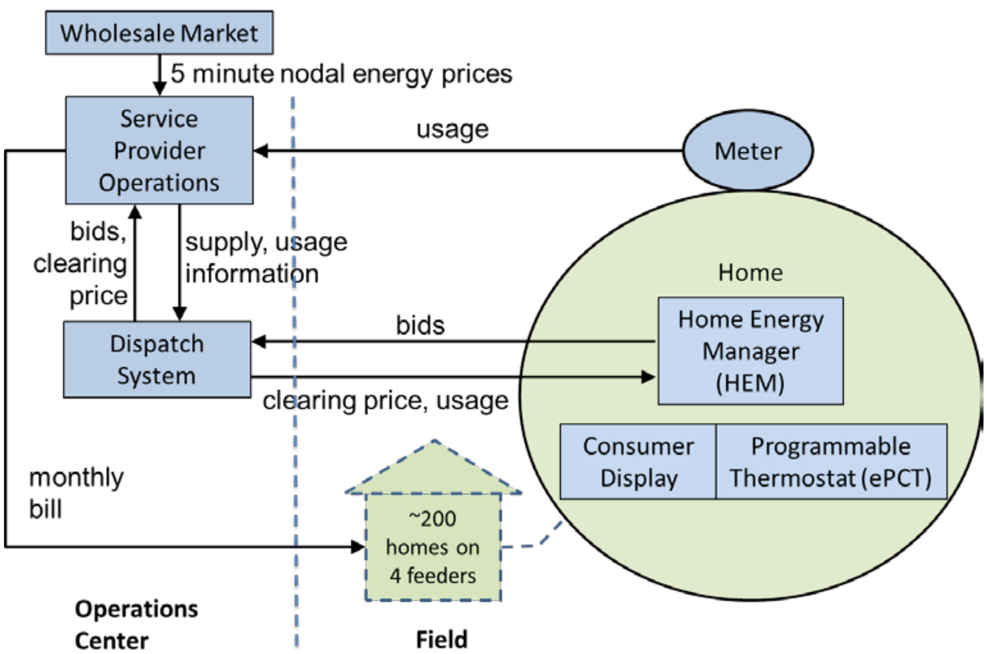

Figure 1.4. RTP_{da} System Overview

Within the home is an electronic program-controlled thermostat (ePCT) communicating with an HVAC unit and a home energy manager (HEM). The HEM hosts a software agent that monitors the market price of electricity and converts the residents' desired temperature set point, the current deviation from that set point, and their preference setting for relative comfort and savings into an amount it is willing to bid for the next 5 minutes of electricity. The HEM takes this price, along with the amount of electricity needed to run the HVAC unit, and assembles all bids in the home (in this case there is only one, representing the HVAC ePCT) and communicates the bid information via a cellular connection to the dispatch system located in the operations center.

The dispatch system assembles the bids from all households on the feeder along with the market price for supplying electricity as determined by the RTP_{da} tariff and based on the LMP for electricity in the feeder's service area. The dispatch system clears the market of supply and demand bids where the two curves intersect, creating a cleared price (as shown in Figure 1.1). The cleared price is broadcast to all homes' HEMs and sent to the service provider's operations system for billing. The billing system exchanges information with the advanced metering infrastructure smart meter at the home to obtain the energy used during the 5-minute interval so the bill can be calculated. The HEM communicates the results of the auction to the ePCT, which sends the appropriate operating signal to HVAC unit. A consumer display is built into the ePCT; it displays the estimated billing price for energy so the consumer can participate with other energy saving actions, should they be monitoring the system.

This transactive-control approach results in very simple message exchange. In general, the approach is sensitive to data-exchange privacy concerns because the transacting parties only need to share what they are willing to pay for a quantity of electricity. What is returned to all participants on the feeder is the market cleared price. For experimental purposes, additional information is collected to understand the performance of the RTP_{da} system. For example, the observed temperature in the home is recorded, as is the deviation of the temperature from the desired set point. In addition, the configuration of the ePCT is monitored, including the residents' preference for savings or comfort and any system overrides, so that consumer behavior can be studied.

1.2.2.1 Thermostat Agent

The smart thermostat agent is configured by the consumer to address their preference for comfort versus economy. For each daily period of operation (for example, "Home," "Away," or "Night"), the homeowner specifies their desired temperature ($T_{desired}$), and influences their minimum and maximum temperatures ($T_{min,}$ T_{max}), through a five-level setting for their preference for more comfort (tighter temperature control) or more savings (more flexible temperature control), as represented by the slope (k) in Figure 1.5. To simplify the discussion, only cooling mode is described below.

The thermostat agent's price-responsive controller is programmed to account for two market phenomena: price trends and price variability. In the case of price trends, the agent needs to determine whether a price is expensive or inexpensive. What may seem like a high price today may seem like a low price tomorrow, and vice versa. We see this in the fluctuation of the price of gasoline, where today's price may seem low compared with the price paid several months ago. In the case of price variability, we look at the volatility of short-term changes in price. Although the average price over a period may be relatively constant, the variability of the actual price above and below the average can change.

The volatility (standard deviation divided by mean) of wholesale and retail prices varies over time. Because the price-responsive controllers are designed to attenuate their response in the presence of more volatile prices, the determination of volatility is essential to the operation of the overall RTP_{da} system. In the case of this demonstration, the time window for the calculation of price volatility is the most recent 24 hours. The effect of this implementation is to attenuate the responses of the thermostat agents during the 24 hours that follow a period of significantly increased price volatility. The longer the duration of increased volatility, or the greater the volatility, the more the thermostat agents' responses are attenuated. For this reason it is typical to see diminished response to LMP fluctuations during the 24 hours that follow a feeder constraint event.

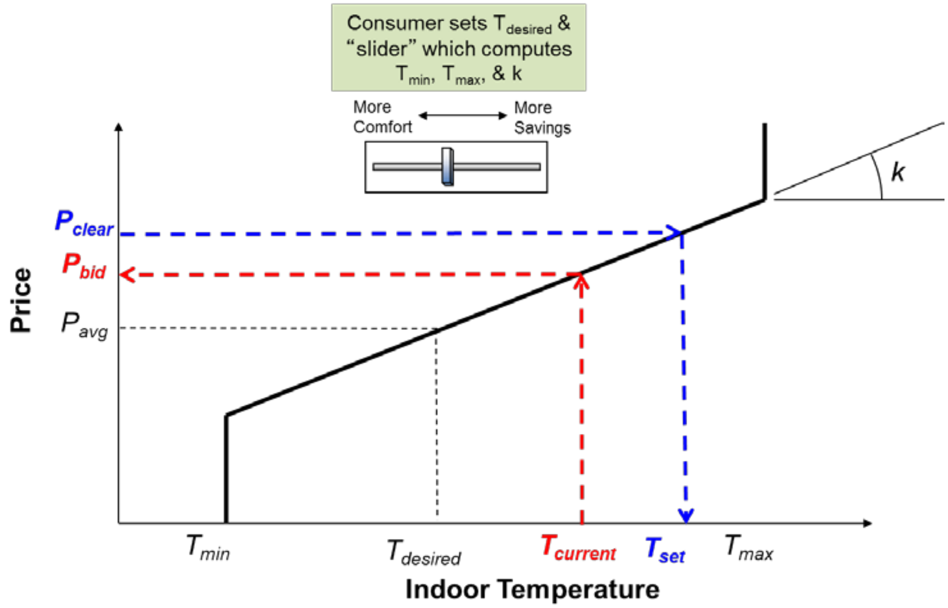

Figure 1.5. HVAC Thermostat Agent Price-Response Curve in Cooling Mode[1]

Bids are submitted every 5 minutes up until 60 seconds before the market is cleared. The bid price P_{bid} in the figure above is computed by each thermostat agent as follows:

$$P_{bid} = P_{avg} + \frac{(T_{current} - T_{desired}) \times k \times P_{dev}}{T_{max} - T_{min}} \tag{1.1}$$

where

where	P_{avg}	=	the average price over the last 24 hours
	P_{dev}	=	the standard deviation of the price over the last 24 hours
	$T_{current}$	=	the current indoor air temperature
	$T_{desired}$	=	the desired indoor air temperature
	k	=	the responsiveness desired by the consumer
	T_{max}	=	the maximum temperature limit
	T_{min}	=	the minimum temperature limit.

P_{bid} and q_{bid} are sent by the HEM to the market system, where they are assembled with the other bids and the market is cleared. The cleared price (P_{clear}) is then published to all the RTP_{da} HEMs on the feeder, which pass it on to the thermostat agents where the price-response curve is used to define the temperature set point for the next 5 minutes of operation (T_{set}, see Figure 1.5). In the cooling mode case in Figure 1.5, the fact that $P_{clear} > P_{bid}$ results in T_{set} being set higher than $T_{current}$ and less than T_{max} so the HVAC unit will not run. A higher comfort setting would result in higher prices bid as the indoor temperature deviates from $T_{desired}$.

[1] Hammerstrom, D. J., et al, "Pacific Northwest GridWise® Testbed Demonstration Projects, Part I. Olympic Peninsula Project," Pacific Northwest National Laboratory, PNNL-17167, October 2007.

In cooling mode, if the current temperature is above the maximum temperature ($T_{current} > T_{max}$), then a bid at the price cap with zero quantity is submitted (that is, the consumer is fully unsatisfied). If the current temperature is below the minimum temperature ($T_{current} < T_{min}$) then a bid of zero price and zero quantity is submitted as a programming convention to represent that the consumer is fully satisfied.

The bid quantity is provided by the RTP_{da} equipment installer based on the estimated nominal power demand of the heating/air-conditioning unit. The bid state is determined by the operating mode of the heating/air-conditioning unit—for example, "Off," "Cool," or "Heat."

If a previously submitted bid is invalidated by a change in ePCT state (for example, "Off" to "Cool"), or if there was a disruption of service, then a new bid is computed and submitted to replace the previous bid. All bids received are recorded in the system database, but only the last bid received is used to clear the market.

1.2.2.2 HVAC Operating States

To better understand the operating status of the HVAC equipment and its interplay with the market bidding system, the following states are considered in the summer cooling scenario of the analysis. A diagram of the HVAC states and their possible transitions over time is depicted in Figure 1.6. The flag, "Included in Market," indicates that the HEM successfully communicated with the RTP_{da} dispatch system so that its bid can be included in the next auction.

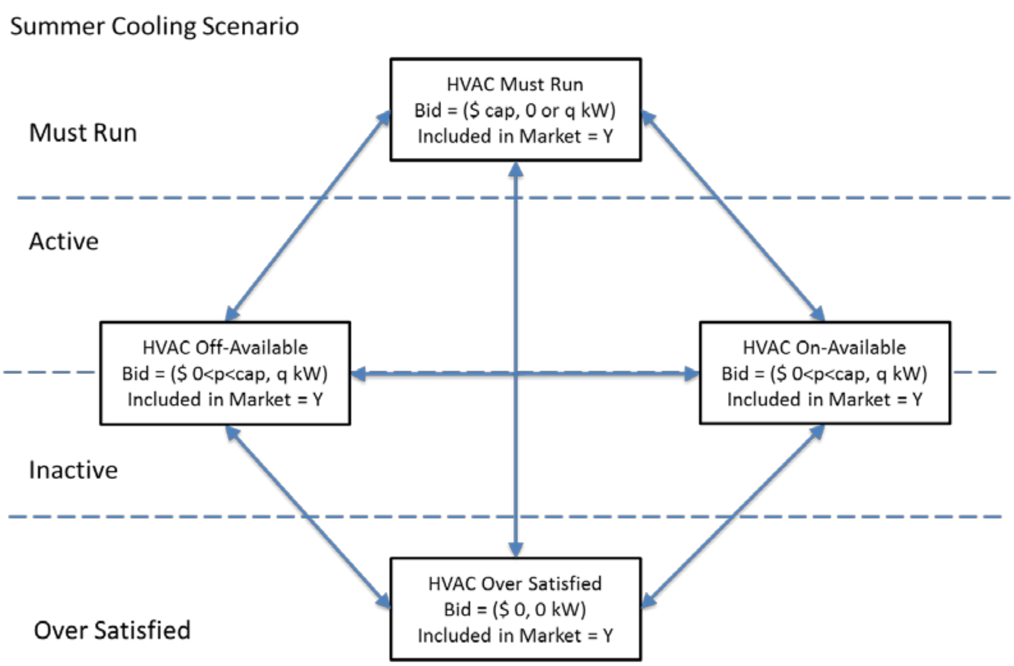

Figure 1.6. HVAC State Diagram

- Must Run: This state is reached in two situations. First, if the household temperature and comfort settings in the ePCT result in a bid at the highest price allowed in the market, P_{cap}, then the unit will automatically clear the market and will be expected to be in the "On" state during the next 5-minute period. Second, in the case where $T_{current} > T_{max}$ (that is, the household temperature is over the maximum temperature set point), then the thermostat agent is programmed to bid price = \$0 and quantity = q_{bid}, and it is expected to be in the "On" state no matter what the cleared market price. This distinction is made so that the HVAC unit is counted as unresponsive load. This convention allows the dispatch system to more easily recognize it as unresponsive.

- Over-Satisfied: If $T_{current} < T_{min}$ (that is, the household temperature is below the minimum temperature set point), then the HVAC unit is in the Over-Satisfied state and will bid price = \$0 and quantity = 0.
- Active: Represents the state where the HVAC bid was cleared (p_{bid} higher than P_{clear}) to run in the market period. The HVAC unit can either be Off and available to remain Off or turn On, or On and available to remain On or turn Off.
- Inactive: Represents the state where the HVAC bid was not cleared (p_{bid} lower than P_{clear}) to run in the market period. As with the Active state, the HVAC unit can either be On or Off and available to switch states.

In moving from one auction to the next, it is possible that an HVAC unit may stay in the same state or move to any other state. As the internal temperature is increasing in conditions of steady market supply price, one would expect an HVAC unit to move from Inactive to Active, and possibly to Must Run, if it could not keep up with the temperature increases. However, under volatile market conditions and congestion events, state changes could be more dramatic.

Considering the entire RTP_{da} household load under control (RTP_{da} Load), a more detailed market clearing illustration is presented in Figure 1.7 to reflect the different states of the HVAC units in a

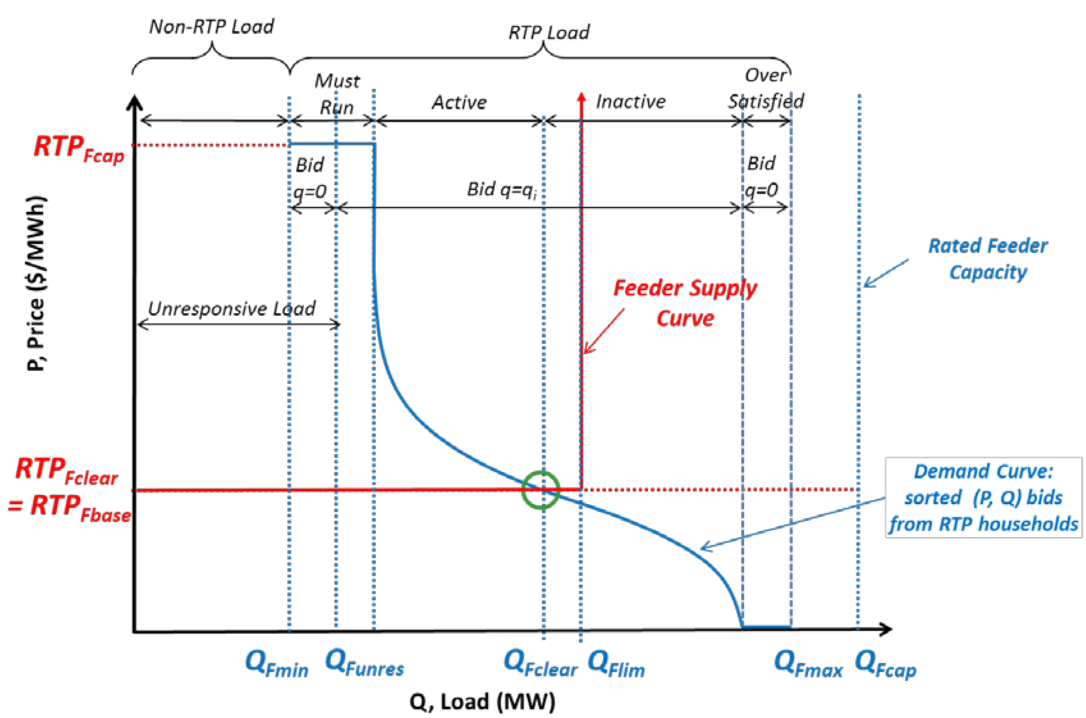

Figure 1.7. RTP_{da} Load Bidding Classifications – Non-Congested Case

particular 5-minute auction. The HVAC units to the left in the figure are in the Must Run state, followed by the Active units, then the Inactive units, and lastly the Over-Satisfied units. The remaining load on the feeder is referred to as Non-RTP_{da} Load. This is slightly different from the definition of the Unresponsive Load, which also includes the Must Run devices that submit a zero bid. The subscript F in the figure refers to feeder-based variables, with new variable Q_{Funres} being the quantity of unresponsive load on the feeder, Q_{Flim} being the congestion limit placed on the feeder, and Q_{Fcap} being the rated capacity of the feeder.

1.3 RTP_{da} Experiment Setup

To run the RTP_{da} demonstration, households were recruited to participate under the RTP_{da} tariff and they were outfitted with the ePCTs and HEMs. The RTP_{da} dispatch system was commissioned and communication was enabled between the various components of the system. An operations experiment plan was developed for testing the RTP_{da} system and performing congestion experiments.

1.3.1 RTP_{da} Households

The RTP_{da} households were selected from a pool that already had smart meters installed. These meters provided data to the meter data management (MDM) system and were also read by the HEMs, which returned metered data with their market bids, the status of the ePCT, and the indoor temperature. Based on consumer recruitment into the RTP_{da} program, changes that occurred with the households, and the eventual decommissioning, the number of participants grew over the spring and summer of 2013 and diminished in the fall months as their equipment was removed. The household equipment was configured at the time of installation, and the consumer was trained on how to enter their desired thermostat settings and change them to reflect their preferences over time. Any changes were recorded and sent back to the RTP_{da} dispatch system from the HEM.

The RTP_{da} dispatch system and the HEMs were designed to handle problems with communications. For example, default values were used for the PJM LMP price if there were delays in getting that from PJM. If a HEM's bid came in too late for the 5-minute market auction, then it did not participate in that auction, but it could participate in the next auction in which a successful bid was submitted and received. To properly analyze the behavior of the system, missing or bad data need to be detected and removed. An understanding of the default or backup settings is needed, as their appearance in the data collection can become regular, potentially skewing analysis results and observations.

The RTP_{da} system operated from December 2011 through the fall of 2013, but a sufficient population of households for conducting the experiments was not installed and operational until June, 2013. As most of the HVAC resources only operated in cooling mode, there was little heating HVAC market interaction after the beginning of October. For this reason, we limit the bulk of RTP_{da} analysis to the period from 1 June 2013 through 30 September 2013.

1.3.2 Operations Experiments and Data Collection

To identify and quantify various value streams, and to fully characterize the behavior of RTP_{da} resources, various operating scenarios were designed for the congestion experiments. The operating scenarios involved changing feeder congestion limits for varying durations to engage the RTP_{da} resources.

The RTP_{da} experiments were conducted to test the response of RTP_{da} resources based on parameters such as time of day (peak/off-peak), day of week, and weather conditions (temperature, wind, etc.). Operating scenarios were also designed to test the response of RTP_{da} resources during the critical peak pricing (CPP) events called by AEP Ohio. Finally, fatigue experiments were designed to test the extent to which the RTP_{da} households continued responding to high clearing prices, by letting the indoor temperature rise, before manually overriding the thermostat settings.

Feeder Limit Setting for RTP_{da} Congestion Experiments

First, the process of inducing feeder congestion to conduct an RTP_{da} experiment will be described. Figure 1.8 presents a conceptual view of how congestion limits were set to engage RTP_{da} resources during experiments. The dispatch system allows the operator to enter a percentage ($C_{\%}$) of the feeder's rated capacity (Q_{FCap}) to define the feeder congestion limit (Q_{Flim}). The initial plan was to conduct the congestion experiments by setting the feeder congestion limit in a manner that would engage 10–25% (α) of the total RTP_{da} responsive load on the feeder, using the following formula:

$$C_{\%} = Q_{Flim} / Q_{Fcap} \times 100\% \tag{1.2}$$

$$C_{\%} = (Q_{Ftotal} - \alpha Q_{res}) / Q_{Fcap} \times 100\% \tag{1.3}$$

where
$\alpha < 1$ = portion of Q_{res} to engage
Q_{Flim} = feeder congestion limit
Q_{FCap} = feeder rated capacity
Q_{Ftotal} = total feeder load
Q_{res} = responsive feeder load
$C_{\%}$ = percent of the feeder rated capacity.

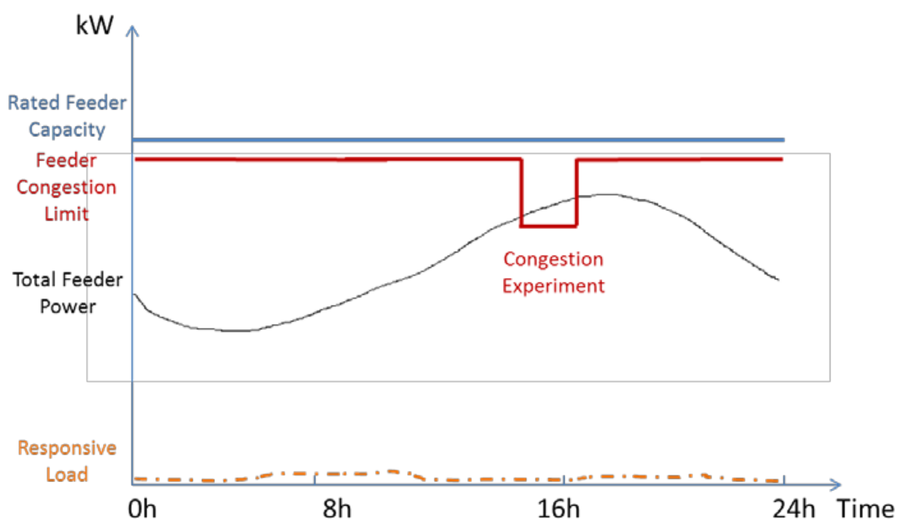

Figure 1.8. Engaging Responsive Load on a Feeder by Varying the Congestion Limit

However, the responsive loads were low compared to the total feeder load levels to the extent that the normal unresponsive load variations were greater than the total responsive load; thus, the reduction of the feeder capacity limit as a percentage of the responsive load did not always lead to congestion on the feeder. To be sure that the RTP_{da} resources would be engaged, the feeder congestion limit (Q_{Flim}) in the congestion experiments was set at 10% below the total (responsive plus unresponsive) prevailing feeder load, instead of the total responsive load. As an example of setting the feeder congestion limit, consider a total load level of 3 MW on a 10 MW feeder with only 100 kW of total responsive load. To impose congestion during an experiment, the feeder limit would be set at 10% below 3 MW, i.e., 2.7 MW (Q_{Flim}) or at 27% ($C_{\%}$) of the rated feeder capacity of 10 MW (Q_{FCap}).

Period of Study

The initial two-week period in the beginning of June was deemed a practice period. Congestion experiments were scheduled during this period on a limited basis to help shed light on the systemic behavior of RTP_{da} resources under different operating conditions. The information gathered during this period was instructive in setting up more extensive experiments later.

Practice Period

RTP_{da} resources were initially engaged for 60 minutes by setting the feeder capacity limit at 10% below the prevailing feeder load. If the RTP_{da} resources were not exhausted during the test period, the length of time to impose congestion was increased by 30 minutes during the experiment, while keeping the same congestion limit. Congestion experiments during the practice period were conducted under constant supervision of staff members at both PNNL and AEP Ohio.

Normal Operation

After the initial practice period, congestion experiments were scheduled daily during the last two weeks of June. The experiments were initially conducted under constant supervision of PNNL and AEP Ohio staff members. However, the experiments were later scheduled to run without constant supervision, once it was determined that market conditions were not being violated and that the RTP_{da} resources were not being exhausted during the course of the experiment. Table 1.1 and Table 1.2 present a breakdown of congestion experiments scheduled over different hours of a day, as well as weekend versus weekday experiments. As can be seen in Table 1.1 and Table 1.2, respectively, majority of the congestion experiments were conducted during peak periods on weekdays. The feeder capacity limits were set at 10% below the prevailing feeder load at the start of the experiment to ensure that the RTP_{da} resources were engaged.

Table 1.1. Breakdown of Congestion Experiments by Hour of Day

5:00–10:00	10	10.42%
10:00–14:00	25	26.04%
14:00–22:00	61	63.54%
Total	**96**	100.00%

Table 1.2. Breakdown of Congestion Experiments by Day of Week

Weekend	25	26.04%
Weekday	71	73.96%
Total	**96**	**100.00%**

Table 1.3 presents the breakdown of congestion experiments based on the experiment durations. As mentioned earlier, 4-hour and 6-hour experiments were conducted to test consumer fatigue, as measured by the number of manual adjustments to thermostat controls.

Table 1.3. Breakdown of Congestion Experiments by Experiment Duration

2 Hours	70	72.94%
4 Hours	25	26.04%
6 Hours	1	1.04%
Total	**96**	**100.00%**

SMART Shift Plus Events and Feeder Constraints

To study the response of RTP_{da} resources during SMART Shift Plus^{SM} events called by AEP Ohio, experiments were scheduled to coincide with and span the duration of the SMART Shift Plus events. SMART Shift Plus events were typically called for 4 hours; these also served as consumer fatigue tests. Table 1.4 below shows the experiments scheduled on the SMART Shift Plus event days, when congestion experiments were scheduled to coincide with the SMART Shift Plus events.

Table 1.4. Congestion Experiments Scheduled during SMART Shift Plus Events

SMART Shift Plus Day	SMART Shift Plus Date	Start Time (Eastern)	Duration (Hours)
Tue	7/16/2013	13:00	4
Wed	7/17/2013	15:00	4
Thu	7/18/2013	15:00	4
Thu	8/22/2013	15:00	4
Tue	8/27/2013	14:00	4
Thu	8/29/2013	14:00	4
Fri	8/30/2013	15:00	4
Tue	9/10/2013	15:00	4
Wed	9/11/2013	15:00	4

1.4 Control Groups

A number of control households were selected that were expected to have characteristics similar to the RTP_{da} households, but that remained under the standard residential tariff. Changes in behavior of the RTP_{da} group can be estimated by comparing RTP_{da} results against those of the control group. A pool of thousands of households who did not participate in the customer-oriented projects was established from which control group households could be chosen. From this pool, PNNL developed a control group of 272 households for comparison in several of the analyses in this report. Note that this RTP_{da} control group is different from other control groups used in other parts of the gridSMART report.

This section describes the way in which the control group was selected from 2010 metered data and how the control group data from 2013 were corrected for use in comparisons with RTP_{da}2013 metered data. A set of definitions for the groups of households used for the analyses follows.

1. RTP13: This group represents the 192 households who were technology-enabled and participated in the RTP_{da} market some or all the time in 2013. Often their energy use in an interval such as 5 minutes will also be referred to as RTP13.
2. RTP10: This is a group of 272 households who were recruited in 2013 as potential RTP_{da} participants and were in the RTP_{da} system database. The 15-minute energy use data for 2010 was obtained for these households. This group includes the 192 RTP13 households that participated in the RTP_{da} market during the demonstration period; however, the selection was done prior to the analysis of how many RTP_{da} households actually participated in the market.
3. Ctrl10: From the pool of households who did not participate in the gridSMART program, a set of 272 households were identified as close to RTP10 in their 15-minute energy use profiles. These are referred to as Ctrl10.
4. Ctrl13: The energy use by the same set of households as Ctrl10 in 2013 is referred to as Ctrl13.
5. RTPnr10: The average energy used by a household in Ctrl10 was then adjusted to improve the comparison with the RTP_{da} households in 2010 (RTP10). This adjusted control group is referred to as RTPnr10 (“nr” meaning non-responsive).
6. RTPnr13: The average energy used by a household in Ctrl13 was then adjusted to improve the comparison with the RTP_{da} households in 2013 (RTP13). This adjusted control group is referred to as RTPnr13.

1.4.1 Control Group Member Selection Process

The following describes how the control group was initially selected to create a set similar to RTP_{da} households, but not in the program. This involved acquiring data for candidate households to be in the control group, processing the data, (including handling bad or missing values in the acquired data), and employing data filtering mechanisms to help match load shapes to select control group members from the candidates that could represent non-responsive RTP_{da} households.

Data Acquisition: The 2010 15-minute MDM data for approximately 11,800 homes were used to identify the control group. The analysis interval was from June 1, 2010 through September 30, 2010.

Bad and Missing Values: 15-minute values that exceeded 40 kWh (an unusually high and suspect value) were treated as bad data and removed by setting the large value to zero so that they would be handled as missing data during subsequent processing. Missing values in the household meter data were replaced using a zero-order hold before any selection filter was applied. Note that the missing values were reset to zero after the filter (see below) was applied to avoid affecting the match with an RTP_{da} household itself.

Filtering: Load data collected at sub-hourly intervals can exhibit large fluctuations in the average energy (that is, power) measurement due to the cycling behavior of large loads. The quantity of interest is the duty cycle, but this quantity cannot be directly observed from interval energy data. However, for time intervals longer than the cycling time of the loads the average load, P_{avg}, is related to the duty cycle, *D*, as

$$D = \frac{P_{average}}{P_{on}} \tag{1.4}$$

where P_{on} is power measured when the equipment is On.

This property is used to estimate the total load at various time intervals by filtering the load data using a sliding-window filter based on the $1/p$-state binomial probability distribution function (PDF). The binomial PDF is defined as

$$\text{Prob}\{t|n,p\} = \binom{n}{t} p^t (1-p)^{(n-t)} \tag{1.5}$$

where $p = 2$ and describes the probability that the true state (On or Off) at the time $t = n/2$ is described by the observed state at the time t. The choice of the window size n was based on the load cycling time relative to the sampling time Δt_{sample}; for example,

$$n = \frac{\Delta t_{on} + \Delta t_{off}}{\Delta t_{sample}} \tag{1.6}$$

Note that the filtered data has zero lag (it is shifted back by a half window). In addition, the last sample is held for an additional half window to provide a smooth end to the filtered data. The result of applying such a filter with window size of 8 15-minute periods is shown in Figure 1.9.

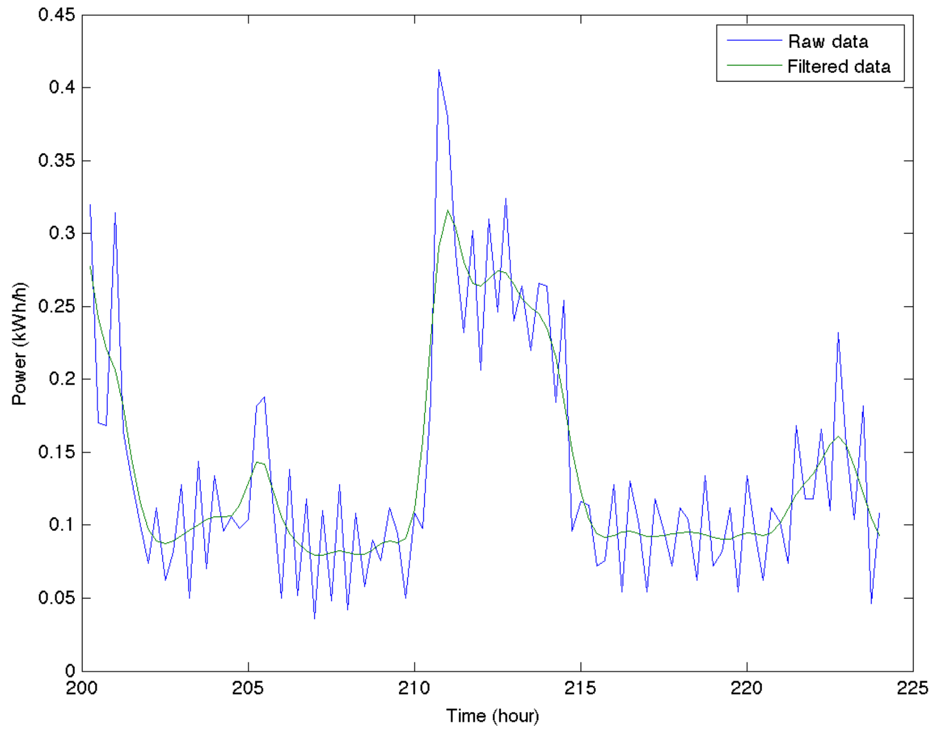

Figure 1.9. Raw and Filtered 15-Minute Electric Meter Data with a 2-Hour Binomial Window Size

Load Shape Matching: The simplest method for load shape matching is based on minimizing the total mean-squared error (MSE) between candidate load shapes

$$MSE = \sum_{t=1}^{N} (x_t^+ - y_t^+)^2 \quad (1.7)$$

where N is the number of samples in the time series, and x^+ and y^+ are the non-zero values from the load shape time-series vectors. An example of a match is shown in Figure 1.10. The corresponding figures, known as heat maps because they show the high (hot) and low (cold) areas, are shown in Figure 1.11. Each RTP_{da} home was assigned a single control home. However, some control home choices were matched to more than one RTP_{da} home. In such cases, the next-best match that was not already selected for the control group was chosen.

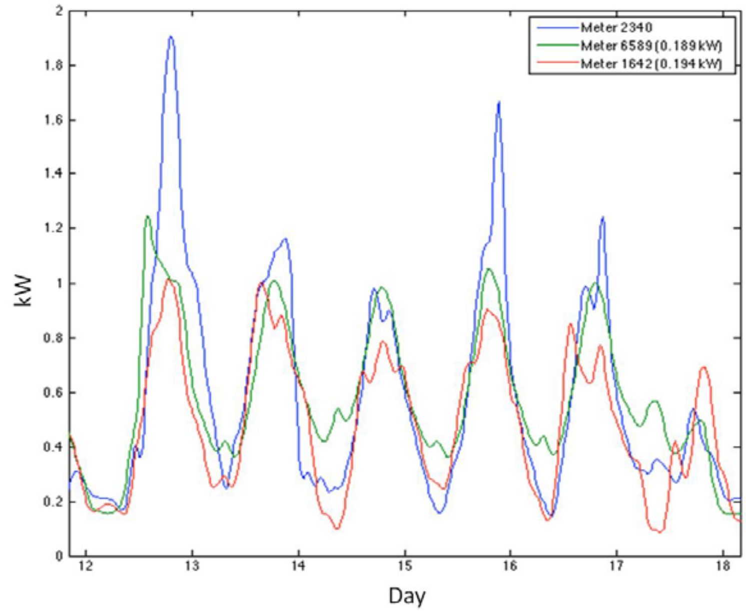

Figure 1.10. Illustration of Match of Reference Load (blue) to Best Fit (green) and Second-Best Fit (red)

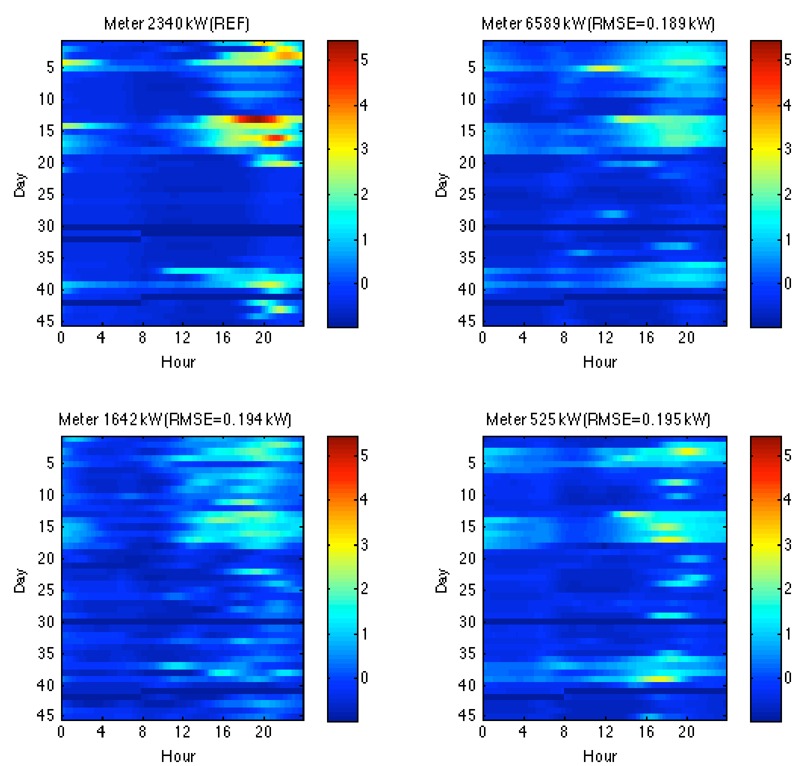

Figure 1.11. Load Shape Maps for Reference (upper left), Best Match (upper right), Second-Best Match (lower left), and Third-Best Match (lower right)

1.4.2 Control Group Adjustment for RTP_{da} Group Comparison

Recall that the RTP_{da} load data from the 2010 data are referred to as RTP10 and the RTP_{da} data measured during the course of the demonstration period in 2013 are referred to as RTP13. Similarly, the control group data for 2010 and 2013 are referred to as Ctrl10 and Ctrl13, respectively. If Ctrl13 accurately represents the behavior of RTP13 had they not been price responsive, then it is a simple matter to subtract the interval data for the average of the two groups to determine price response. Despite the optimal search for a best fit of Ctrl10 with RTP10, there were substantial differences. A typical 7-day profile of energy use by the two groups in 2010 is shown in Figure 1.12. The variations, especially of peak loads, are of concern. While the selection process emphasized load shape matching, it did not match peak energy use. An adjustment was made to make Ctrl10 closer to RTP10. The correction procedure is explained below, and the resulting group is RTPnr10.

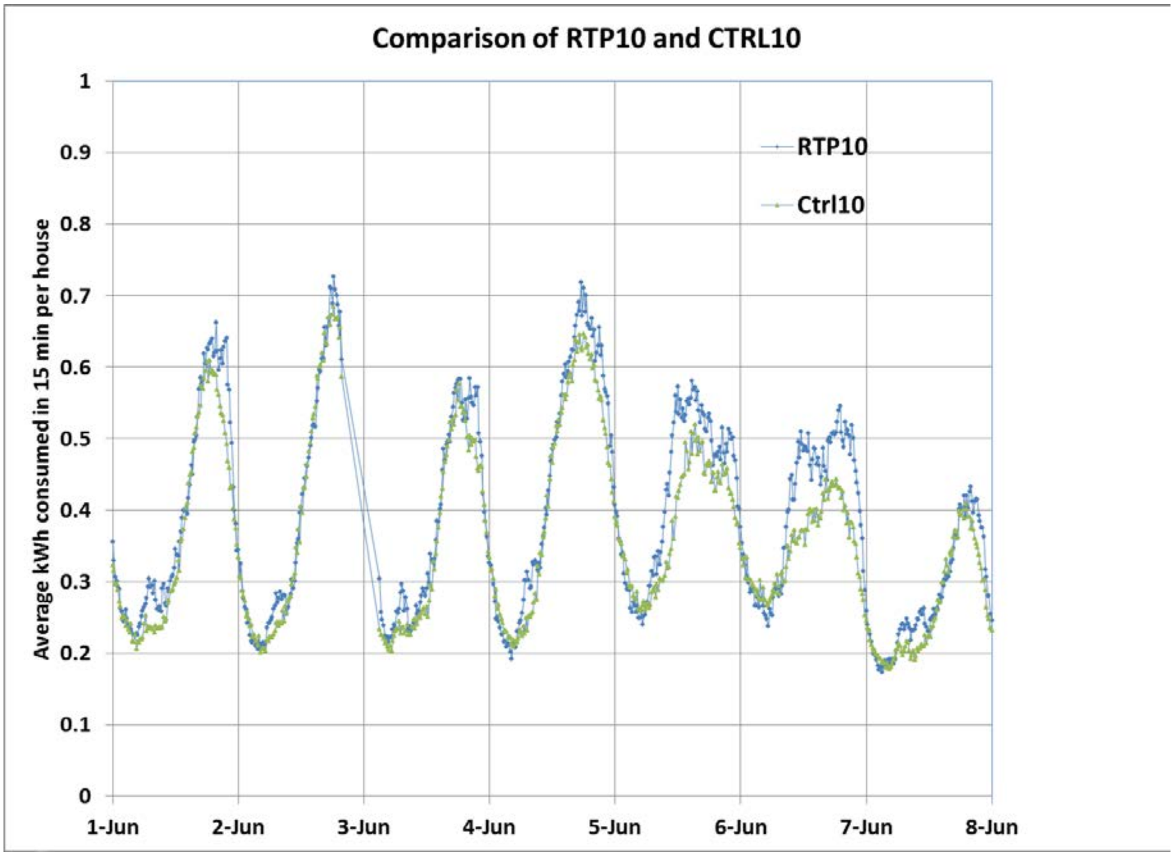

Figure 1.12. Comparison of RTP10 and Ctrl10 Profiles for a Typical 7-Day Period

Consider a relationship between RTP10(t) at time t and Ctrl10(t) of the form

$$RTP10(t) \sim f(Ctrl10(t), Time\ of\ day) \tag{1.8}$$

As RTP10 and Ctrl10 are experiencing the same outdoor temperature, the time of day (actual date is not relevant) turned out to be a good proxy for many of the un-modeled variables, including outdoor temperature. Another approach is to consider the difference between RTP10 and Ctrl10 as a function of Ctrl10 and time-of-day. A Ctrl10 value of 0.6 at 3 pm on a day in July (this would happen on a relatively cool July day) gets the same correction as a Ctrl10 value of 0.6 at 3 pm on a day in September (this would happen on a relatively warm September day). Thus, the Ctrl10 value and time-of-day act as proxies for temperature. So *RTP10 - Ctrl10 ~ f(Ctrl10, time-of-day)* is our non-parametric model. And *RTP10 ~ Ctrl10 + f(Ctrl10, time-of-day)* is just another *f(Ctrl10, time-of-day*). If the function *f* is parameterized in some fashion, the parameters can be estimated by a method such as the least-squares method. However, because the functional form as well as the parameters and their interpretation are not of much interest, a non-parametric method was used. The method of choice was LOcal regrESSion (LOESS).[2] This was implemented in MATLAB®.

The time series

$$RTPnr10(t) = f(Ctrl10(t), Time\ of\ day) \tag{1.9}$$

is a significantly better approximation of RTP10 than Ctrl10. Weekdays and weekends were treated separately. The resulting fit for the same 7-day period as in Figure 1.12 is shown in Figure 1.13.

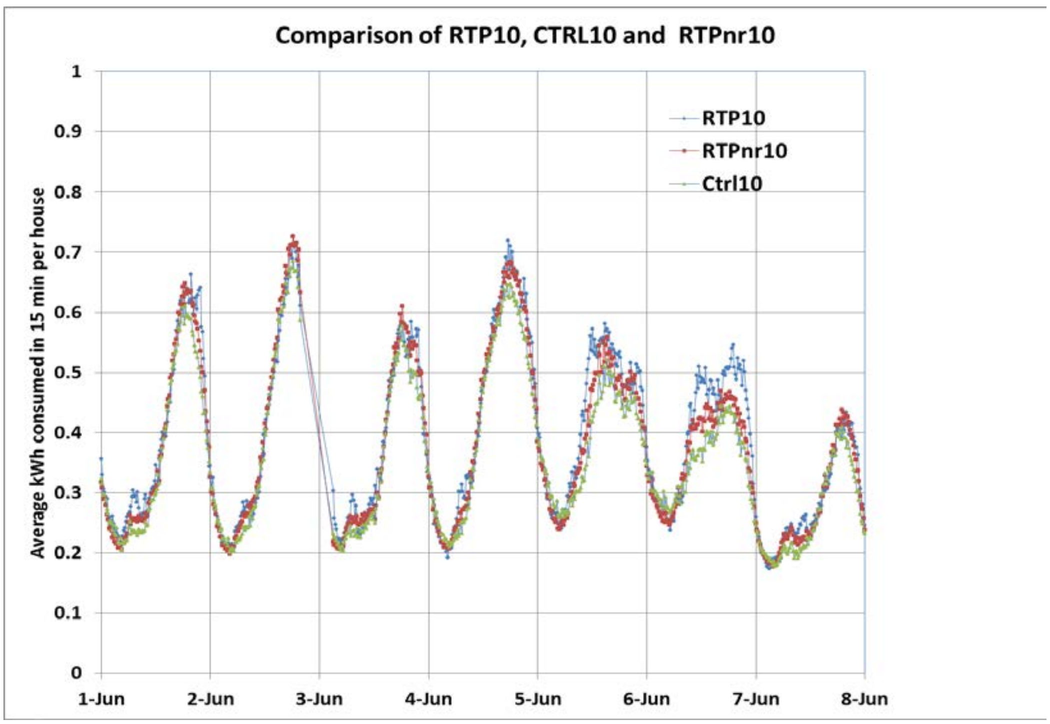

Figure 1.13. Comparison of RTP10, RTPnr10, and Ctrl10 Profiles for the Same 7-day Period as in Figure 1.12

The fit for the entire 122-day period is shown in Figure 1.14.

[2] http://www.mathworks.com/products/datasheets/pdf/curve-fitting-toolbox.pdf

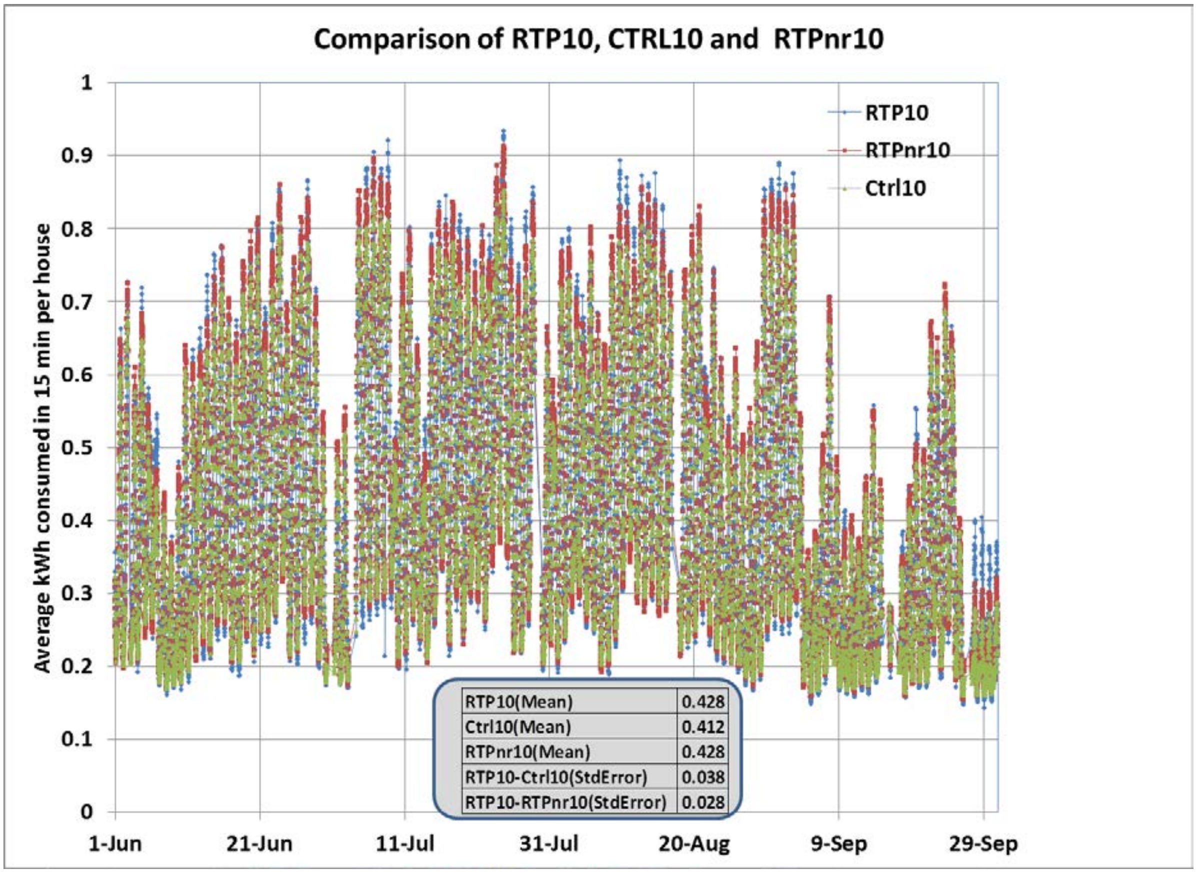

Figure 1.14. Comparison of RTP10, RTPnr10, and Ctrl10 Profiles for the Period June to September 2010

Although Figure 1.14 is crowded, one can discern that RTPnr10 is closer to RTP10 than Ctrl10 is. This is quantified by the comparison of means and standard errors in the inset in Figure 1.14. The performance of LOESS in matching peak loads can be assessed by comparing the top 5% of RTP10 loads with the coincident Ctrl10 and RTPnr loads. This is shown in Figure 1.15. It is clear that, for our purposes, RTPnr is a much more accurate representation of RTP_{da} than Ctrl.

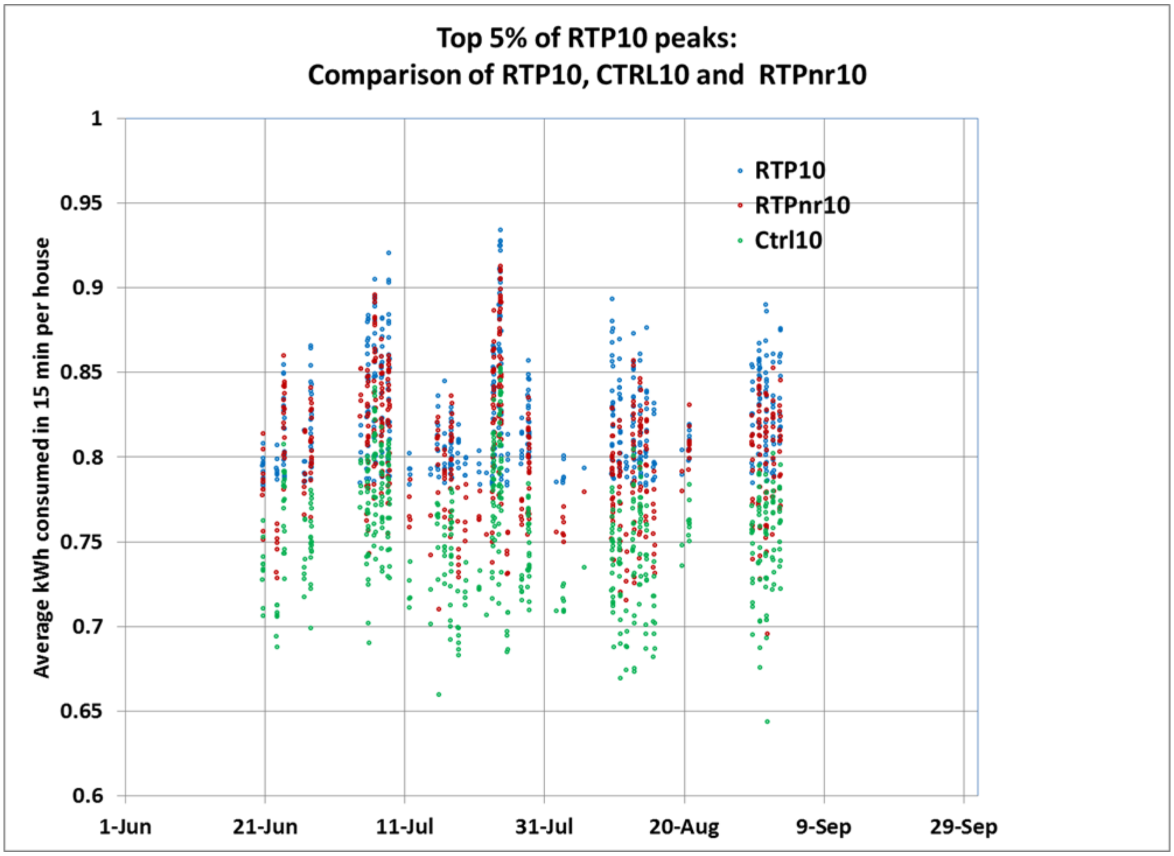

Figure 1.15. Comparison of the Top 5% of RTP10 Loads with the Coincident Ctrl10 and RTPnr Loads

We can now use the LOESS model to generate RTPnr13 from Ctrl13. From this point on, we will use RTPnr13 as the corrected control group with which RTP13 is to be compared. RTP13 is available in 5-minute intervals, whereas RTPnr13 (generated from 15-minute Ctrl13) is in 15-minute intervals. A linear interpolation method was used to generate 5-minute data from the 15-minute data. (This process is much cleaner than rolling up 5-minute RTP13 data into 15-minute data, for which missing data creates a number of special issues.)

Consistent handling of missing data is essential. Although Ctrl10 is an aggregate over a maximum of 272 households, there are many periods with fewer households. For obtaining statistically good aggregate data, only data to which >80% of the 272 contributed were retained. Similar processing was done for RTP10. The missing data time stamps for the two need not be the same. Only time stamps for which both RTP10 and Ctrl10 data are present are retained. Similar processing was done for RTP13 and RTPnr13.

1.5 Document Structure

This report supplements the RTP_{da} system analysis done by AEP Ohio in the main body of the gridSMART Project report. It covers three major areas:

- an analysis of the impacts of the RTP_{da} approach to engage end-use resources for system operations, including its application to address system capacity concerns, wholesale purchases, and spinning reserves,
- an analysis of impacts related to the consumer, including household bills, the consumers' interactions with the thermostats, and a comparison of the amount of energy bid into the market for running the HVAC units and the actual consumption of those units,
- and an analysis of the sensitivity of the of the RTP_{da} load to the fluctuating price of energy, including the observed response of the RTP_{da} resources to the congestion experiments.

2.0 System Impacts

The following sections describe the results of an analysis of impacts that affect system operations of the service provider. These include system- and feeder-capacity issues, wholesale power purchases, and the potential of applying RTP_{da} resources to spinning reserve markets.

2.1 Capacity

This analysis measures the reduction in capacity expansion requirements due to a price-induced shift in household peak load. The benefit of this analysis will be presented in terms of kW/household reduction in peak load.

Evaluating the capacity reduction is a complex problem that can be difficult to observe and characterize under real-world conditions, especially when the penetration level of RTP_{da} households is relatively low compared to other groups. A series of experiments were developed on the end-use resources to characterize their behavior and limitations. The results of these experiments were used to better calibrate the parameters of the simulation models. The simulation models were then used to quantify the potential for capacity reduction at various penetration levels.

2.1.1 Results of Analysis

The simulation models of the RTP_{da} system were evaluated on three peak days in July to determine the greatest sustainable capacity reduction that was achievable. On these days (July 16–18), the temperature was greater than 90^{o} F on five successive days. The evaluation was performed by (1) lowering the capacity limit until the cleared price reached the price cap during peak system load hours and (2) lowering the capacity limit until the price cap was reached during the projected peak feeder hours. The simulations in (2) were also run over the four-month test period to verify that the capacity could be maintained at a lower level throughout the four-month period. The simulations were performed at 15%, 25%, 35%, and 50% RTP_{da} penetration levels. The models were "tuned" to be responsive only to peak conditions, and not wholesale price fluctuations; this is similar to a day following a high-price event, when the controllers are desensitized to small fluctuations in the wholesale price.

Figure 2.1 shows a representative simulation during a 3.5-hour peak system load event. Figure 2.2 shows a representative simulation on the same day, but focusing on feeder peak reduction. All measurements are 15-minute average demand and are translated into a kW/household basis. Notice that the feeder peak is near the end of the event, highlighting that system and feeder peak demands do not necessarily align; hence the need to look at the availability of the resource in different periods. In Figure 2.1, notice that after the event is triggered, the two lines approach each other after approximately 2.5 hours, indicating that the resource is no longer able to hold a load reduction and the households begin to become less responsive, while in Figure 2.2 the resource begins to reduce sooner (approximately 2 hours into the event). The time the reduction is called for this peak load event affects the overall availability of the resource.

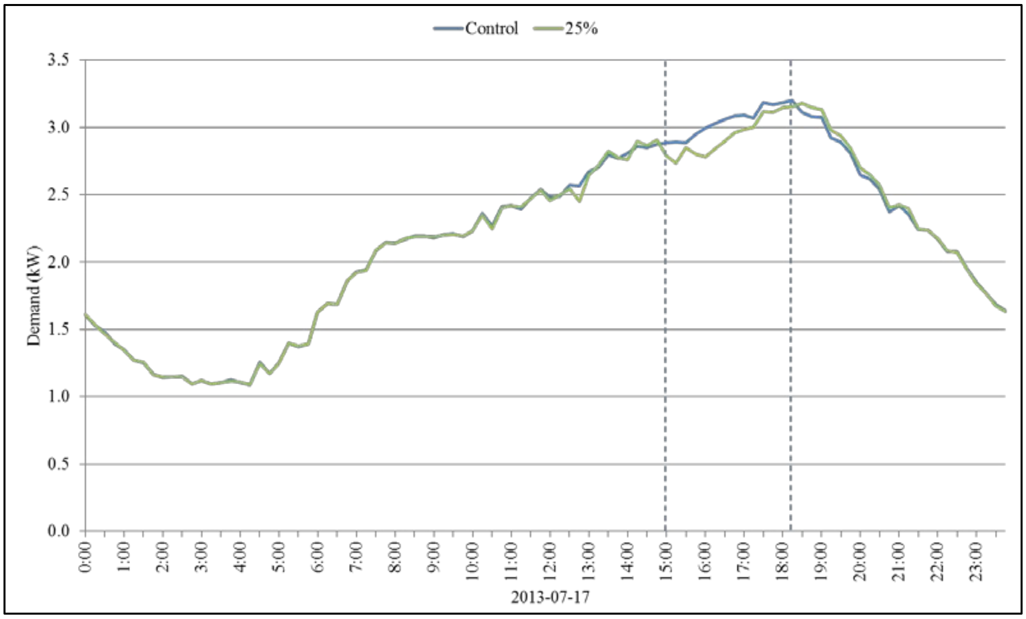

Figure 2.1. Time-Series Simulation during Peak System Load Event with 25% Penetration of RTP_{da} Households

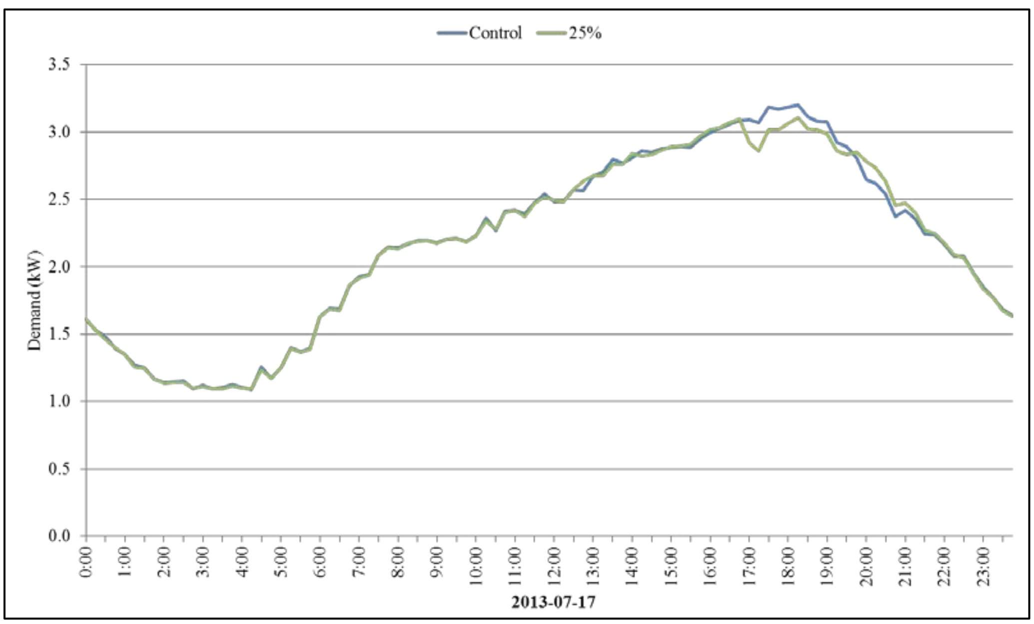

Figure 2.2. Time-Series Simulation of Feeder Peak Reduction with 25% Penetration of RTP_{da} Households

Figure 2.3 and Figure 2.4 show the peak demand reduction (i.e., the difference between the greatest demand before the capacity constraint was applied and the greatest demand after it was applied) as a function of RTP_{da} penetration levels for peak load events and feeder peak reduction, respectively. Additionally, a linear trend line has been added to the figures for clarification. Notice that the ability to reduce the peak during the feeder peak situation is much greater than during a peak system load event. This is for a number of reasons. The first is that the peak system load event lasted longer than the feeder peak event, meaning the resources are spread out over a longer period. The second is that the availability of resources for reduction is lower in non-feeder peak periods. If the values are extrapolated to 100% penetration (or the average response of an RTP_{da} household), it is seen that the RTP_{da} households provide a 13% load reduction during a peak system load event and a 22% reduction during feeder peak events.

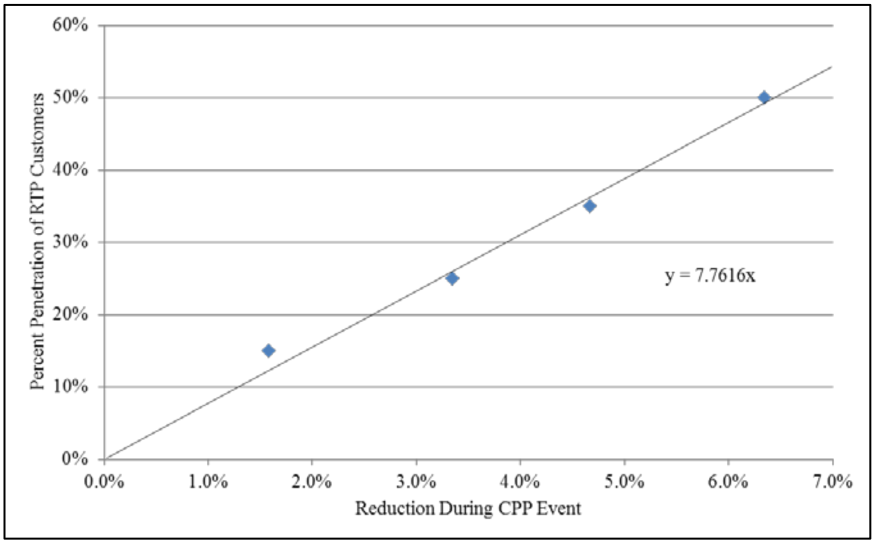

Figure 2.3. Comparison of Peak Reduction during a Peak System Load Event at Different RTP_{da} Penetration Levels

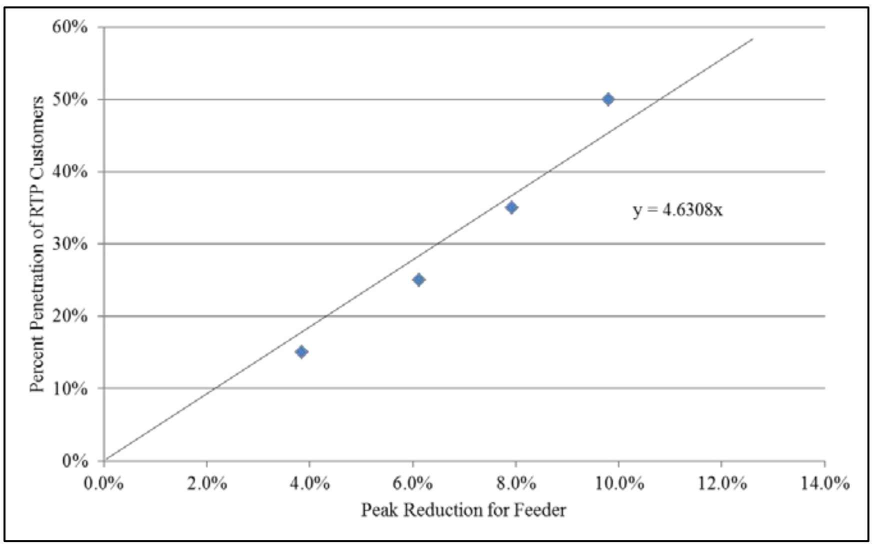

Figure 2.4. Comparison of Feeder Peak Reduction at Different RTP_{da} Penetration Levels

Note that these values represent a specific case for RTP_{da} household response, and in some ways, the "best case." In all simulations, it was known ahead of time when system peak events would occur and for how long, and what the load would be during a peak system load event. In an actual system, this will not be well known and the determination of the capacity limit may overuse resources (leading to early decay of the response) or underuse resources (leaving unused capacity from this resource). Even in simulation, the reduction did not provide a perfectly flat load (see Figure 2.2), as the market lags behind the changing load of the non-RTP_{da} households. Incorporation of short-term load prediction may improve this aspect of RTP_{da} system performance.

Additionally, the length of the time the resource is needed affects the amount of reduction available. For example, a 1-hour peak system load event is able to sustain deeper reductions than a 6-hour peak system load event. Figure 2.5 shows the reduction of demand as a function of the length of a peak system load event using the same simulation and day shown previously. This is shown with 100% penetration of RTP_{da} households. Notice that after 4 hours, the load has effectively returned to a new baseline state, with a minor reduction coming from the thermostat setback. Also, note the magnitude of the rebound after releasing the peak system load event. While significant rebounds occur in the peak periods, if the end of the event is timed correctly after the control peak, the rebound is relatively minor and much lower than during the peak period. The recovery period for all events is such that most devices do not return to normal operation until 22:00 hours, seven hours after the start of the event.

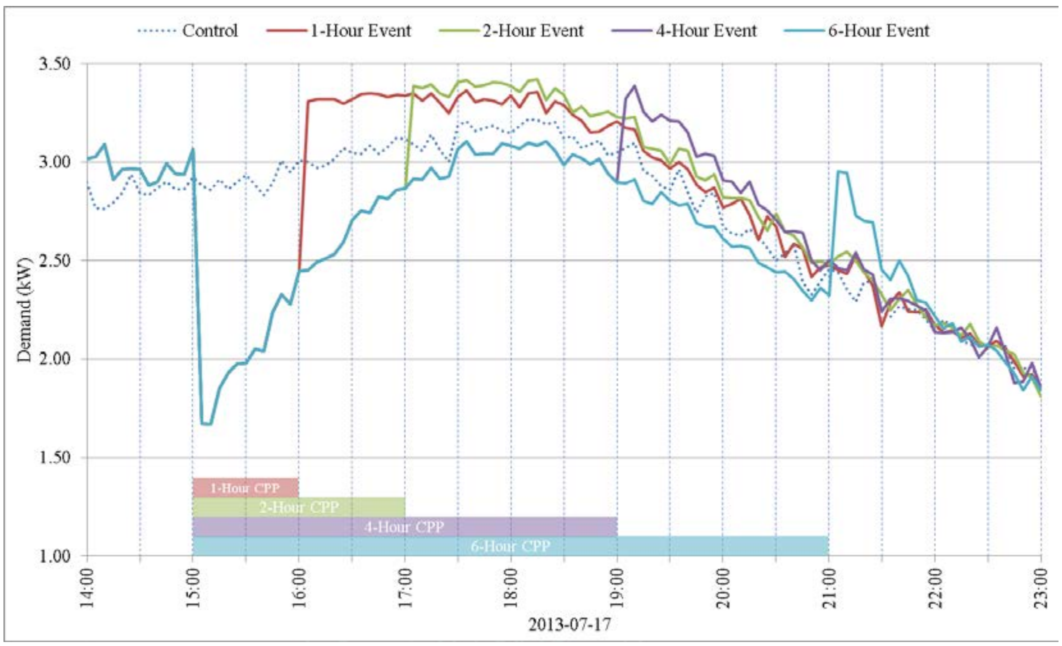

Figure 2.5. Comparison of Average Household Demand during a 1- to 6-Hour Peak System Load Event

2.2 Wholesale Purchases

Price-responsive loads alter their load shape in response to the retail energy prices. If the retail prices are determined in real time by wholesale market LMPs, then demand response to prices should result in decreased cost of wholesale energy purchases. The purpose of this analysis is to examine the impact on wholesale purchases using the data captured related to energy use by the HVAC systems in response to the market signal.

The approach is to compare the energy use in response to RTP_{da} every 5 minutes by the RTP_{da} households against the energy use by the control group. The difference is attributed to price response. Knowing the LMP, the difference in wholesale purchase costs can be calculated.

In this section, only the aggregate response by the participants is considered. That is, data from individual participants were aggregated. The way missing data were handled was considered in Section 1.4.2. The definitions of the household groups were listed in Section 1.4.

2.2.1 Response to Prices

It is instructive to examine RTP13 and RTPnr13 for typical periods. In addition, LMP data also were acquired. Figure 2.6 shows a comparison between RTP13 and RTPnr13 in the top panel and the difference between RTP13 and RTPnr13 and LMP in the bottom panel, which uses the same x-axis day intervals.

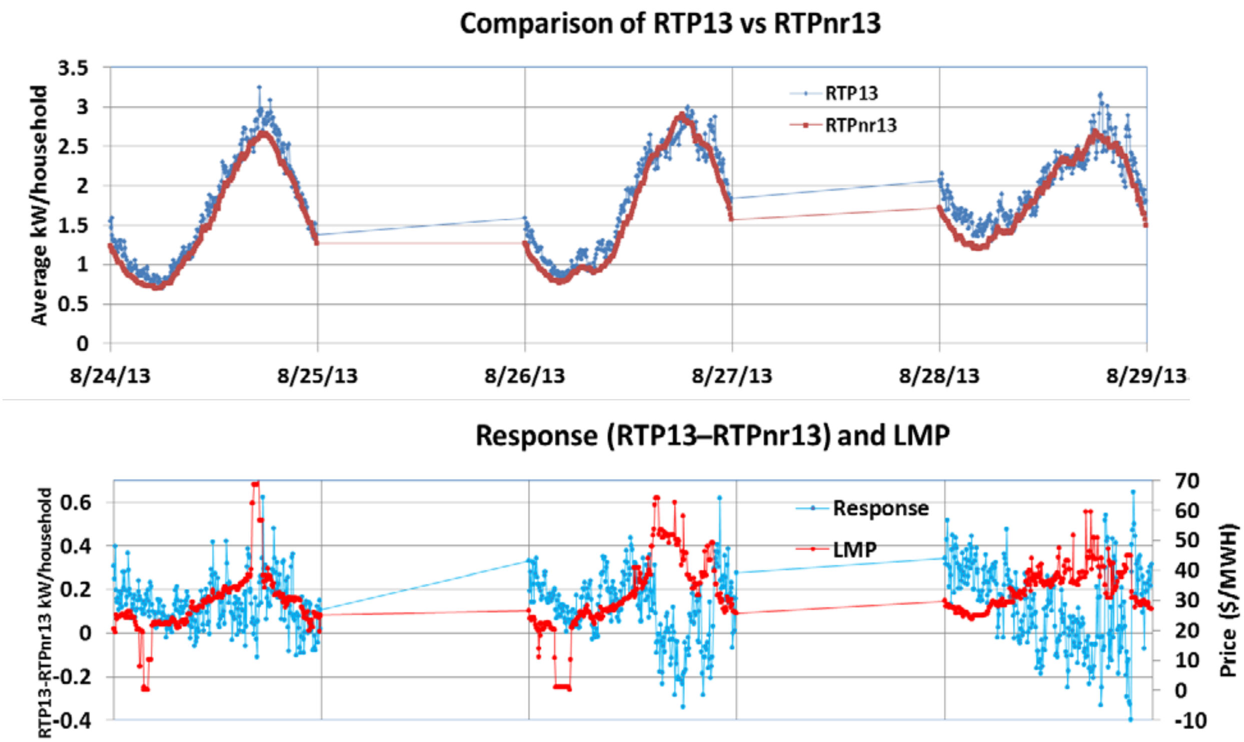

Figure 2.6. A Comparison between RTP13 and RTPnr13 (top), and between RTP13-RTPnr13 and LMP (bottom)

No data are plotted for August 25 and August 27. On those days, feeder congestion experiments were done, so on those days, the system is responding to real-time prices generated by simulated feeder congestion and not to wholesale LMP-generated prices. For this reason, feeder congestion experiment days were excluded from the wholesale purchase analysis. August 26 shows a very discernable load response to prices.

2.2.2 Totals for the 4-Month Period

The aggregation of RTP13 and RTPnr13 will now be examined to compare total loads with and without price response, and RTP13 × LMP and RTPnr13 × LMP will be examined to compare wholesale purchase costs with and without price response. Out of the 122 days in June, July, August, and September, all days when congestion experiments were performed were excluded. Furthermore, only aggregate data that received contributions from >80% of the maximum number of contributors was included. This resulted in 50 days of usable data—31 weekdays and 19 weekend days. Not all 50 days had data for every one of the 288 5-minute periods. Because this affected both RTP13 and RTPnr13 similarly, this was not considered sufficient reason to exclude a day. Table 2.1 shows a summary for the 50 days.

Table 2.1. Summary of Energy and Wholesale Costs for July–September Before Adjustments

Energy		
RTP13 (kWh/day/house)	36.21	
RTPnr13 (kWh/day/house)	35.55	
RTP13 is	1.9%	higher than RTPnr13
Wholesale cost		
RTP13 ($/day/house)	$1.432	
RTPnr13 ($/day/house)	$1.42	
RTP13 is	0.7%	higher than RTPnr13

Feeder congestion days were excluded, but they affected the behavior of the HVAC units the following day. The high prices (~$1000/MWH) experienced during the congestion period made the prices expected by HEMs high for 24 hours following the conclusion of the feeder congestion experiment. This resulted in the normal prices appearing low, and the HEMs responded by lowering the house temperatures. This can be seen in Figure 2.7, which shows that the daily average observed temperature for the non-congestion days was generally substantially lower than the desired set points. The moving average in the figure is computed over 5 points.

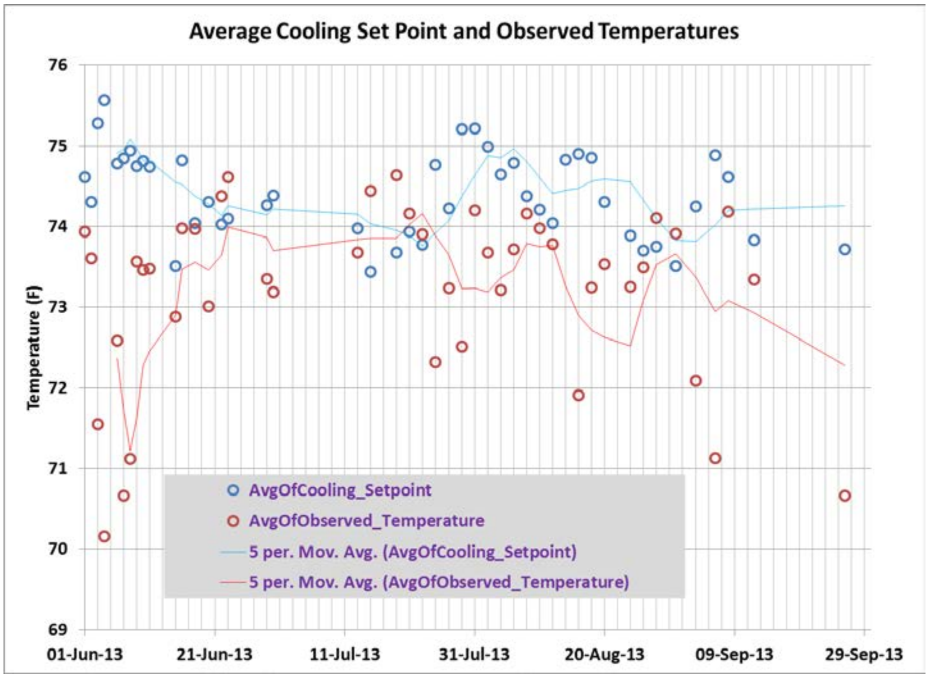

Figure 2.7. Daily Average Cooling Set Points and Observed Temperatures during Non-Congestion Days

The undesirable situation of temperatures below desired set points necessitated adjustments for the additional cooling energy use. This was done as follows. From a plot of daily average outside temperature obtained for the Columbus, Ohio, weather station, a regression of daily energy use versus daily average outside temperature was performed. Separate regressions were performed for weekdays and weekends. The results are shown in Figure 2.8 and Figure 2.9.

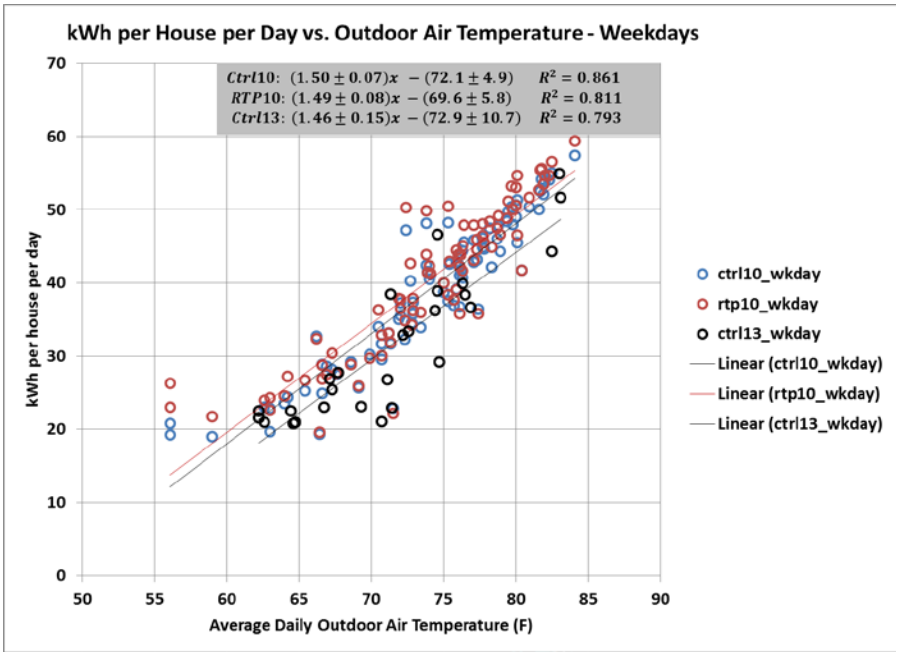

Figure 2.8. Plot of Daily Average Energy Use per House versus Average Outdoor Temperature for Weekdays

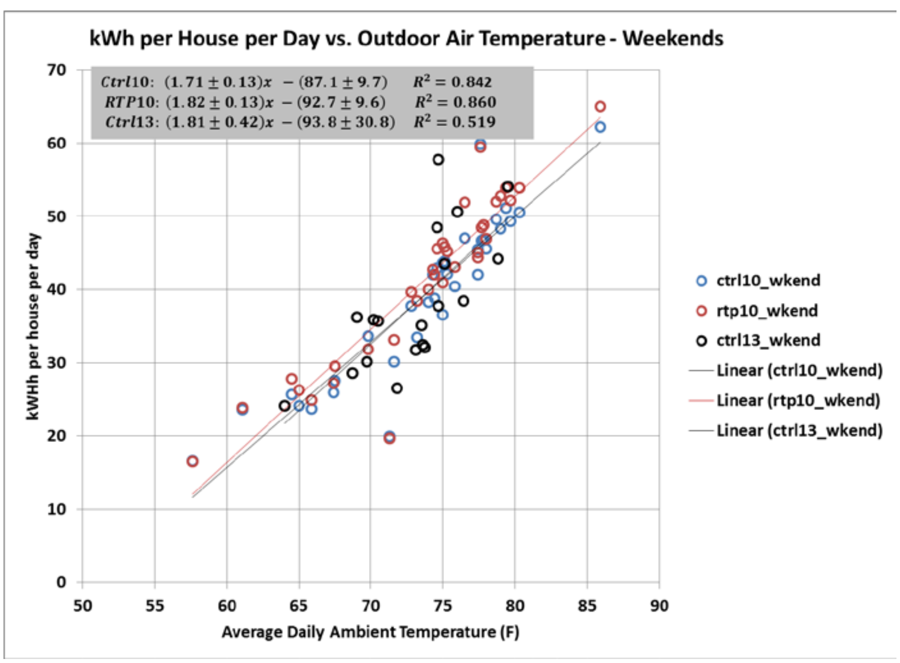

Figure 2.9. Plot of Daily Average Energy Use per House versus Average Outdoor Temperature for Weekend Days

From the above regressions, a value of 1.5 kWh/day/house/°F was derived for weekdays and 1.8 kWh/day/house/°F for weekend days. The behavior of Ctrl13 is statistically indistinguishable from that of Ctrl10. This means that on a weekday, if the average outdoor temperature increases by 1°F, the energy use for the day increases by 1.5 kWh/house. Similar considerations apply for the weekend days. Increasing the outside temperature by 1°F is, to a very good approximation for energy use, equivalent to decreasing the inside temperature by 1°F. If it is now assumed that, in the absence of feeder experiments, the average observed temperature would have been equal to the average desired temperature set point (rather than the observed temperature resulting from the RTP_{da}-driven set point), a compensation term for the daily energy use and the resulting impact to wholesale cost can be applied. The compensation term for the wholesale cost can be calculated as the average LMP for the day times the change in kWh per household for the day. The results of applying such compensation are shown in Table 2.2.

Table 2.2. Summary of Energy and Costs for July–September after Compensating for Congestion Experiments

After compensating for the impact of congestion experiments on non-congestion days		
Energy		
RTP13_Compensated (kWh/day/house)	33.66	
RTPnr13 (kWh/day/house)	35.55	
RTP13_Compensated is	5.3%	lower than RTPnr13
Wholesale cost		
RTP13_Compensated ($/day/house)	$1.351	
RTPnr13 ($/day/house)	$1.423	
RTP13_Compensated is	5.0%	lower than RTPnr13

The kWh usage is reduced by 5.3% and the wholesale costs by 5.0%. Comparing this table with Table 2.1, it can be seen that the effect of the temperature compensation on cost reduction was not commensurate with the effect of temperature compensation on kWh reduction. This is due to the fact that the compensation was effective largely during periods of low prices, as seen in Figure 2.10.

Some sources of error and their impacts follow. These include:

- There is a difference in the demographics of RTP_{da} households over the years 2010 to 2013, and similarly in the control group.
- The RTP10 group had 272 households, whereas RTP13 had 192 households. This is not necessarily a problem, but further review that the pool characteristics match would add confidence.
- The compensation method for the excessive cooling due to congestion experiments in the preceding day deserves further investigation.
- Good data representing only 50 days of operation survived the various filters.
- Even on these 50 days of good data, the number of houses contributing was variable.

No attempt was made to quantify the errors arising from these sources in this report; however, it remains a good topic for future investigation.

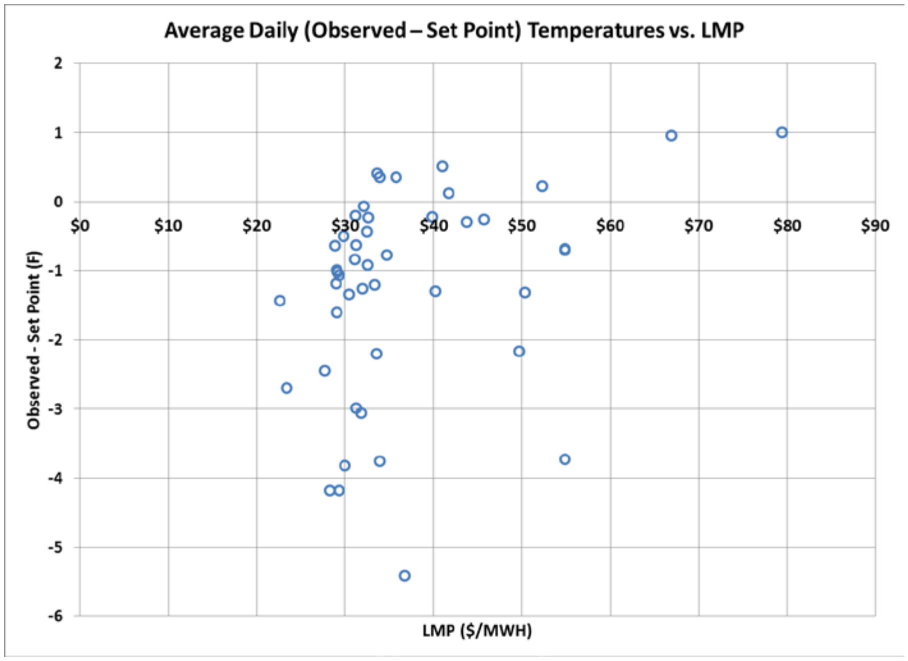

Figure 2.10. The Depression of Observed Temperature Below Set Point versus Average LMP for the Day

2.2.3 Simulated Results

The RTP_{da} household information was used to calibrate GridLAB-D[1] simulated models. These household models were used within GridLAB-D to represent 25% penetration of RTP_{da} households; 300 households were "experimental" while 900 were operated similarly in each simulation and did not respond to variations in price. The experimental households were modeled using three different scenarios:

1. Control – the households were simulated using the standard pricing tariff.
2. RTP_{da} – the households were simulated using the residential RTP_{da} service tariff and responded to wholesale price fluctuations in a manner similar to those observed in the pricing experiments (for example, thermostat slider and temperature settings, internal air temperature decay rates, etc.).
3. RTP_{da} Congested – the households were simulated using the residential RTP_{da} service tariff, responded to wholesale price fluctuations, and responded to capacity limits placed on the feeder aligned with the actual experiments (96 experiments in four months).

The simulations were run for four months, collecting energy consumed at 5-minute intervals by each of the 300 RTP_{da} households. An "average" RTP_{da} household was constructed from the resulting information by summing all 300 loads and dividing by 300 in each interval. The total wholesale energy cost for the average household is shown in Table 2.3. Table 2.4 shows these same values in terms of percent of total energy costs, showing an average RTP_{da} household savings of 2.5% for wholesale energy costs on a per RTP_{da} household basis. When accounting for the effects of the congestion experiments,

[1] www.gridlabd.org

this number is reduced to 1.5%. This is expected, as the effect of the congestion experiments is to reduce sensitivity to wholesale price fluctuations. Note that the energy costs per day match closely with the experiment results shown in Table 2.1 above.

Table 2.3. Comparison of Monthly Wholesale Energy Costs for an Average Household ($)

	Monthly Wholesale Energy Cost Per Household ($)					
	June	July	August	September	Average	Per Day
Control	$43.93	$57.57	$42.58	$37.92	$45.50	$1.492
RTP_{da}	$42.80	$55.68	$41.69	$37.17	$44.34	$1.454
RTP_{da} Congested	$43.10	$56.55	$42.05	$37.44	$44.79	$1.470

Table 2.4. Reduction of Wholesale Energy Costs for an Average Household (%)

	Change in Consumer Wholesale Energy Cost (% Savings)				
	June	July	August	September	Average
RTP_{da}	2.6%	3.3%	2.1%	2.0%	2.5%
RTP_{da} Congested	1.9%	1.8%	1.2%	1.3%	1.5%

Looking at the impact on energy consumption, Table 2.5 shows the average energy reduction for each of the cases (a positive number indicates reduced energy consumption). The average reduction in energy consumption is 1.2%, decreasing to 0.9% during the congestion experiments due to the effects of precooling.

Table 2.5. Reduction of Energy Consumption for an Average Household (%)

	Change in Consumer Energy (% Reduced)				
	June	July	August	September	Average
RTP_{da}	1.3%	1.3%	1.4%	0.7%	1.2%
RTP_{da} Congested	1.1%	0.9%	1.0%	0.7%	0.9%

While these values do not perfectly align with the estimated results in Table 2.2 above, they are very similar, and provide additional veracity to the method described in Section 2.2.2 for adjusting the load shapes.

2.3 Spinning Reserves

This analysis investigates the spinning reserve capacity that can be achieved at any given time due to the demand response capability of the loads in the demonstrations. “Spinning reserve” is the extra generating or demand response capacity that is available to the system operator within a short interval of time to meet demand in case a generator goes down or there is another disruption to the supply. Most system operators require the spinning reserve capacity to be available to compensate for the loss of the largest power plant (plus a fraction of the peak load) within a preset amount of time (typically 10 minutes), and be available to respond continuously for a preset amount of time (typically 30 minutes). The spinning reserve capacity is a short-term capability and can therefore be measured in kW/household/hour or kWh/household/hour. It can also be measured in annual monetary terms as the $/household earned in the spinning reserve market. This analysis evaluates the capability of RTP_{da} households to participate in spinning reserve markets. While this analysis is not comprehensive, it is used to determine a "best case" for RTP_{da} households participating in spinning reserve markets, using average market prices from various independent system operators.

2.3.1 Results of Analysis

Spinning reserve is an ancillary service that can be bid into the ancillary services market. Within PJM, the Synchronized Reserve and Regulation Market decides the Synchronized Reserve market cleared price (SRMCP). As load becomes price responsive, like the HVAC loads in the RTP_{da} demonstration, the load can be considered as a spinning reserve for specific durations in the day and can be bid into ancillary services markets. Knowing the amount of load that can be safely bid into the spinning reserve market is important not only to make a bid, but to also be assured that the required reserve requirement can be successfully satisfied if called upon to provide the service. The auction collects all RTP_{da} customer resource availability and the bid curve can be used to determine the amount of resource available at any given market period, as shown in Figure 2.11. The highlighted area represents the households that are willing to reduce their demand given the proper incentive (via the RTP). The service provider may engage only a fraction of the total available resource to participate in the spinning reserves market by setting the congestion limit, and hence the cleared price, appropriately. The utility would weigh the cost of acquiring the resource (incentive payment to displaced consumers) against the benefit from provision of spinning reserve capacity (revenue from spinning- reserve markets, or avoidance of self-scheduling cost).

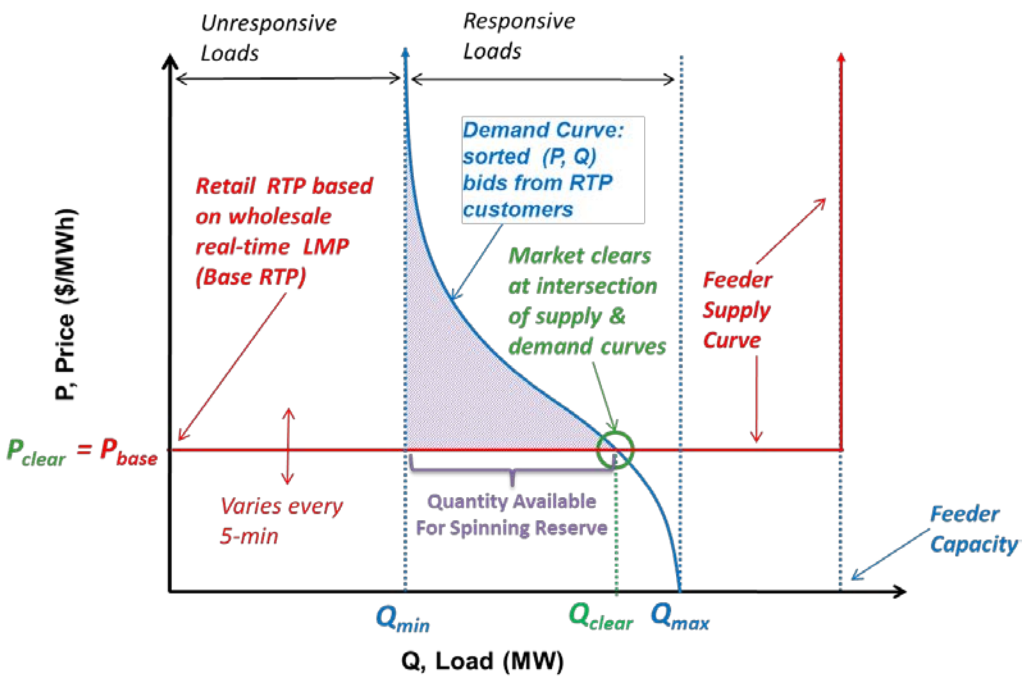

Figure 2.11. Resource Available for Spinning Reserve During Each Market Period

In each market period, the total quantity (kW) of responsive load less than the load at the cleared price is available for participation in the spinning reserve call. In this study, the quantity is calculated for each 5-minute interval, then averaged over the hour to determine the amount of load available for a one-hour spinning reserve call (although most spinning reserve calls last for periods much shorter than one hour).

The overall spinning reserve capabilities of the system during the study period are difficult to determine because the frequent congestion experiments held over the operational period distort a more natural behavior of the bidding system. To better estimate the benefits, PNNL instead used the simulated models to determine the average amount of resource available at every hour of the summer. By then comparing that value to the average PJM market price at each hour, the maximum amount of possible revenue generation in the spinning reserve market can be estimated. Note that this includes every hour in the summer and assumes that no spinning reserve calls were made to affect future household behavior. It also assumes that there is enough resource in the RTP_{da} system to participate in the market. These assumptions are used to determine the “best case” scenario for capturing spinning reserve revenue.

Table 2.6 shows the results of this study in the form of total revenue generated per RTP_{da} household per month. This is calculated by determining the total amount available (as shown in Figure 2.11), then dividing by the total number of RTP_{da} households. The spinning reserve prices were exceptionally low in the PJM market for the summer evaluated, averaging $0.49/MWh between June and October of 2013. A number of additional historical spinning reserve markets were evaluated for comparison. While the potential in PJM’s market is extremely small, in other markets, where spinning reserve resources are in higher demand, differing amounts of revenue can be generated with relatively little impact on the consumer (assuming the resource is called relatively infrequently). The results of three other markets are shown for comparison. This analysis does not address the impact on the spinning reserve market itself or the reduction of overall prices as additional demand- response resources participate, but rather highlights the potential uses of this system.

Table 2.6. Spinning Reserve Markets and the Maximum Amount of Revenue Available to RTP_{da}

	PJM 2013[(a)]	CAISO 2013[(b)]	ERCOT 2013[(b)]	ERCOT 2008[(a)]
Total Revenue Per Household Per Month	$0.08	$1.78	$5.79	$13.64

(a) Based on average hourly prices
(b) Based on average monthly prices

3.0 Household Impacts

This chapter analyzes the impacts of the RTP_{da} approach on consumers and their residential equipment. These include household electricity bills, consumer interactions with their thermostats, and the quantity of HVAC energy bid into the market versus the amount observed from the metering data.

3.1 Household Bill Impacts

This section analyzes the impact on RTP_{da} household bills (per tariff Schedule RS-RTP, 2012) for the months of June through September 2013. The RTP_{da} bills are divided into the following components:

$$B_{RTPtot} = B_{RTPeng} - B_{RTPinc} + B_{RTPfixed} \tag{3.1}$$

where B_{RTPtot} = RTP_{da} household total bill per period of interest
B_{RTPeng} = RTP_{da} household energy-sensitive component of the bill per period of interest
B_{RTPinc} = RTP_{da} household incentive savings component of the bill per period of interest
$B_{RTPfixed}$ = RTP_{da} household fixed, non-energy-sensitive component of the bill per period of interest.

The incentive savings is calculated as explained in Section 1.2.1.3. In any one month, the incentive savings (B_{RTPinc}) is not allowed to exceed the RTP_{da} market-based energy component (a portion of B_{RTPeng}) of the RTP_{da} bill; however, it is possible that $B_{RTPinc} > B_{RTPeng}$ on a daily or hourly basis. As the monthly billing periods for the households are staggered throughout a month, and 5-minute energy usage information is available from the meters, the non-energy portion of the monthly bill is spread evenly over 5-minute intervals to obtain $B_{RTPfixed}$ and the 5-minute energy data are used with the RTP_{da} tariff to obtain 5-minute portions of B_{RTPeng}. The 5-minute household market bidding data are used to calculate a 5-minute B_{RTPinc} component, and Equation (3.1) is used to calculate B_{RTPtot}. This 5-minute data forms the basis for calculating average hourly bills. The average bill for any one hour is calculated based on the population of households that are participating during that hour. This is done for every hour (for which good data exist) over the four-month period. These hours are also analyzed in subsets of peak and off-peak, and hot and mild temperature periods.

Because the bills in this analysis are calculated based on the energy use data captured by the RTP_{da} system about and the RTP_{da} tariff, there will be discrepancies with the actual bills calculated by the AEP Ohio billing system. That system must handle various complicating situations with regard to metering and household changes, and make appropriate adjustments to the final bill that are not replicated here.

The average \$/hr billing information over the months of June through September for all households in the RTP_{da} group is presented in Table 3.1 below. The total bill and the contributions to it are represented according to averages of the totals as well as the averages for the top and bottom 25% of households by bill component area. In addition, the average \$/hr of the bills for households in off-peak (22:00–14:00) and peak (14:00–22:00) periods is also listed. Both off-peak and peak periods are further filtered for mild (outdoor temperature ≤ 80°F) and hot (outdoor temperature > 80°F) weather. These same quantities are repeated in Table 3.2, but reflect percentages based on the average total RTP13 bill.

One can see from the tables that the incentive savings were dispersed to households of all levels of energy usage by amounts on the order of 3–4%. When looking at all the households according to incentive savings component, the average of the top 25% bills of the households received a savings of 12% of the average total bill of all households. This counteracts a 2.5% increase in the energy-sensitive portion of the bill, for a total bill savings of about 5% from the average $/hr over the entire population of households. Those in the bottom 25% have no incentive savings contribution and show a higher average total bill of about 10.5% compared with the average of all households.

Table 3.1. RTP$_{da}$ Bill $/hr Averages for All Households

Metric		RTP13				
		Total Bill	kWh	Energy	Incentive	Fixed
Average $/hr Total Bill		0.1916	1.4812	0.1909	0.0081	0.0087
	Top 25%	0.2563	1.9742	0.2539	0.0064	0.0087
	Bottom 25%	0.1289	0.9987	0.1277	0.0076	0.0087
Average $/hr Energy Bill						
	Top 25%	0.2531	1.9706	0.2529	0.0084	0.0087
	Bottom 25%	0.1290	0.9805	0.1253	0.0050	0.0087
Average $/hr Incentives						
	Top 25%	0.1820	1.5486	0.1963	0.0230	0.0087
	Bottom 25%	0.2116	1.5643	0.2029	0.0000	0.0087
Average $/hr Off-Peak		0.1506	1.2089	0.1457	0.0039	0.0087
	Off-Peak Mild	0.1204	1.0101	0.1148	0.0031	0.0087
	Off-Peak Hot	0.1754	1.3727	0.1712	0.0045	0.0087
Average $/hr Peak		0.2736	2.0257	0.2812	0.0164	0.0087
	Peak Mild	0.1891	1.5225	0.1884	0.0081	0.0087
	Peak Hot	0.3405	2.4247	0.3547	0.0230	0.0087

Table 3.2. RTP$_{da}$ Bill $/hr Averages: Percent of Average Total Bill for All Households

Metric		RTP13				
		Total Bill	kWh	Energy	Incentive	Fixed
Average $/hr Total Bill		100.0%	100.0%	99.7%	4.2%	4.6%
	Top 25%	133.8%	133.3%	132.6%	3.3%	4.6%
	Bottom 25%	67.3%	67.4%	66.7%	3.9%	4.6%
Average $/hr Energy Bill						
	Top 25%	132.1%	133.0%	132.0%	4.4%	4.6%
	Bottom 25%	67.3%	66.2%	65.4%	2.6%	4.6%
Average $/hr Incentives						
	Top 25%	95.0%	104.6%	102.5%	12.0%	4.6%
	Bottom 25%	110.5%	105.6%	105.9%	0.0%	4.6%
Average $/hr Off-Peak		78.6%	81.6%	76.1%	2.0%	4.6%
	Off-Peak Mild	62.8%	68.2%	59.9%	1.6%	4.6%
	Off-Peak Hot	91.6%	92.7%	89.4%	2.4%	4.6%
Average $/hr Peak		142.8%	136.8%	146.8%	8.5%	4.6%
	Peak Mild	98.7%	102.8%	98.4%	4.2%	4.6%
	Peak Hot	177.8%	163.7%	185.2%	12.0%	4.6%

The bill components can also be viewed across the peak and off-peak periods with hot and mild weather. This ranges from off-peak+mild incentive savings of 1.6% of the average $/hr total bill to peak+hot incentive savings of 12%. The impact on the hourly rate of the bills over these periods is about 63% for off-peak+mild to about 180% for peak+hot periods when compared to the average $/hr for all time periods. This range is more dramatic than for the kWh consumed in those periods because energy prices for the RTP$_{da}$ households are generally greater during peak+hot periods than during off-peak+mild periods.

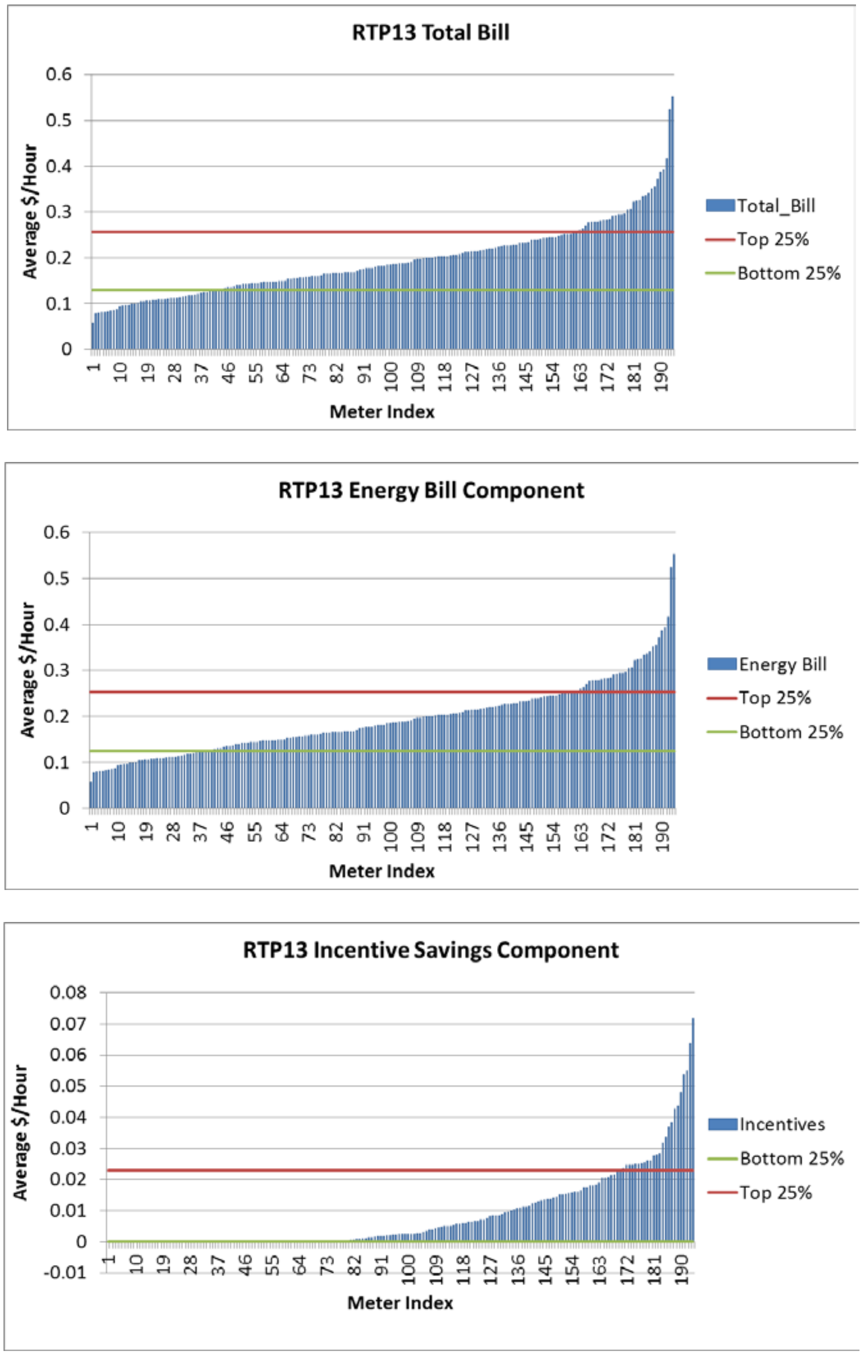

Figure 3.1. RTP$_{da}$ Average Hourly Household Bills with Total, Energy, and Incentive Savings Components

The graphs in Figure 3.1 show the distribution of the hourly rates for the population of households as ordered from lowest to highest household for each bill component; the households (indicated by meter index in the graphs) are reordered in each graph from lowest to highest household contribution. Each figure includes the average of the top and bottom 25% for the population being displayed.

One can see from these figures that there are a few households with relatively high total average hourly energy bills, with the remainder distributed with a relatively flat slope. When one looks at the incentive savings component, about 40% of the households received little or no incentive savings, while about 10% of households had significant average \$/hr incentive savings. Note that the incentive saving is allowed to exceed the RTP_{da} market-based contribution to the energy portion of the bill on an hourly basis, but not on a monthly basis.

The graphs in Figure 3.2 depict the distribution of the average \$/hr energy-sensitive components of the bills for households in off-peak and peak periods. Both off-peak and peak periods are further filtered for mild and hot weather as defined above. For each graph, the households are reordered from lowest to highest contribution for the population displayed.

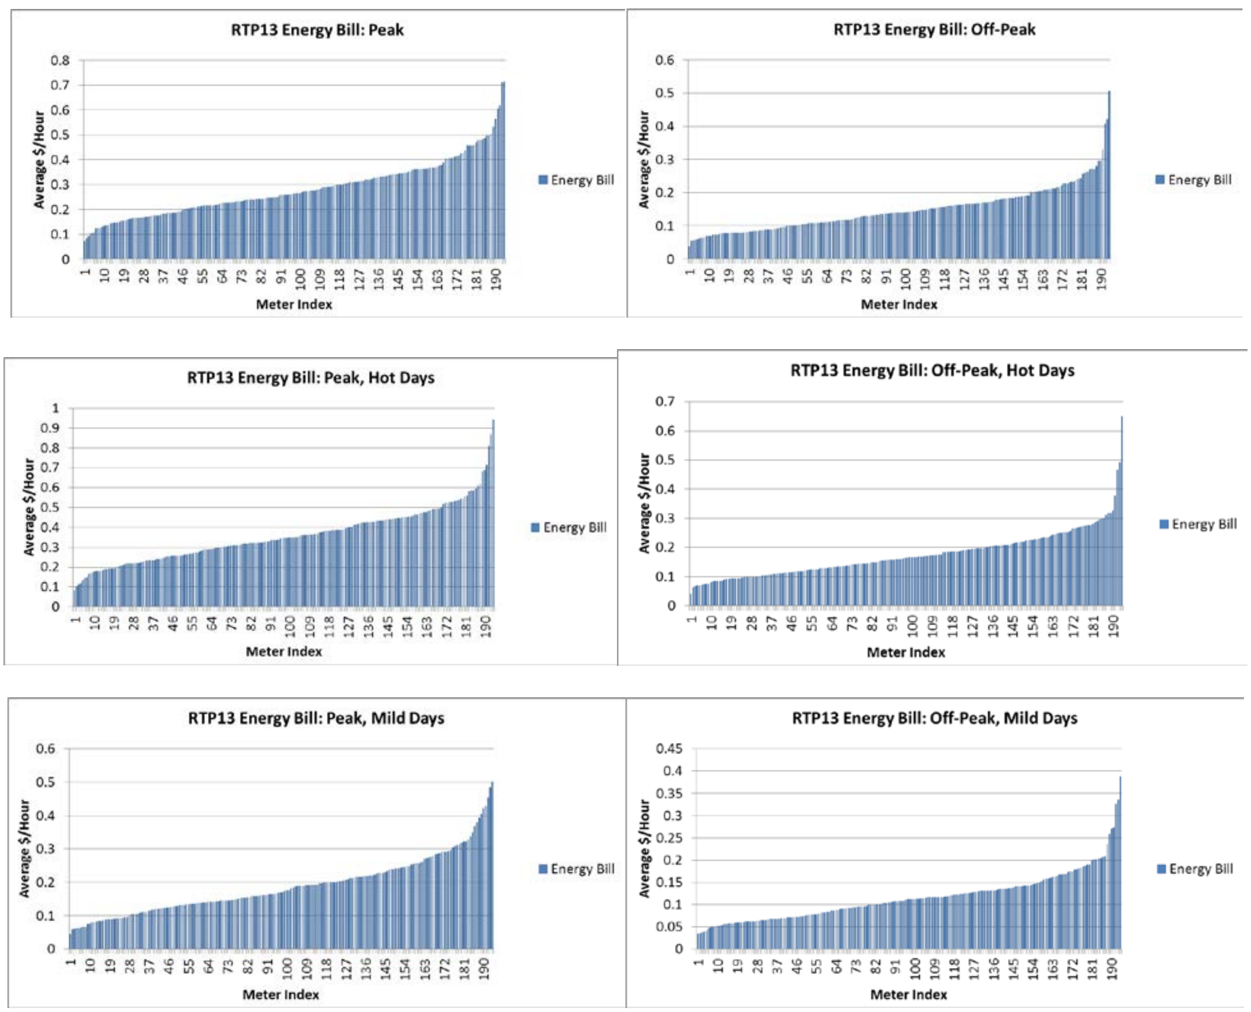

Figure 3.2. RTP_{da} Average Hourly Household Energy-Sensitive Bill Component for Peak and Off-Peak Periods

From these graphs one can see the relatively small number of households at the extremes of the billing range and the significant differences in peak+hot and peak+mild versus off-peak conditions. Similarly, the incentive savings contributions to the RTP_{da} bills are shown in Figure 3.3.

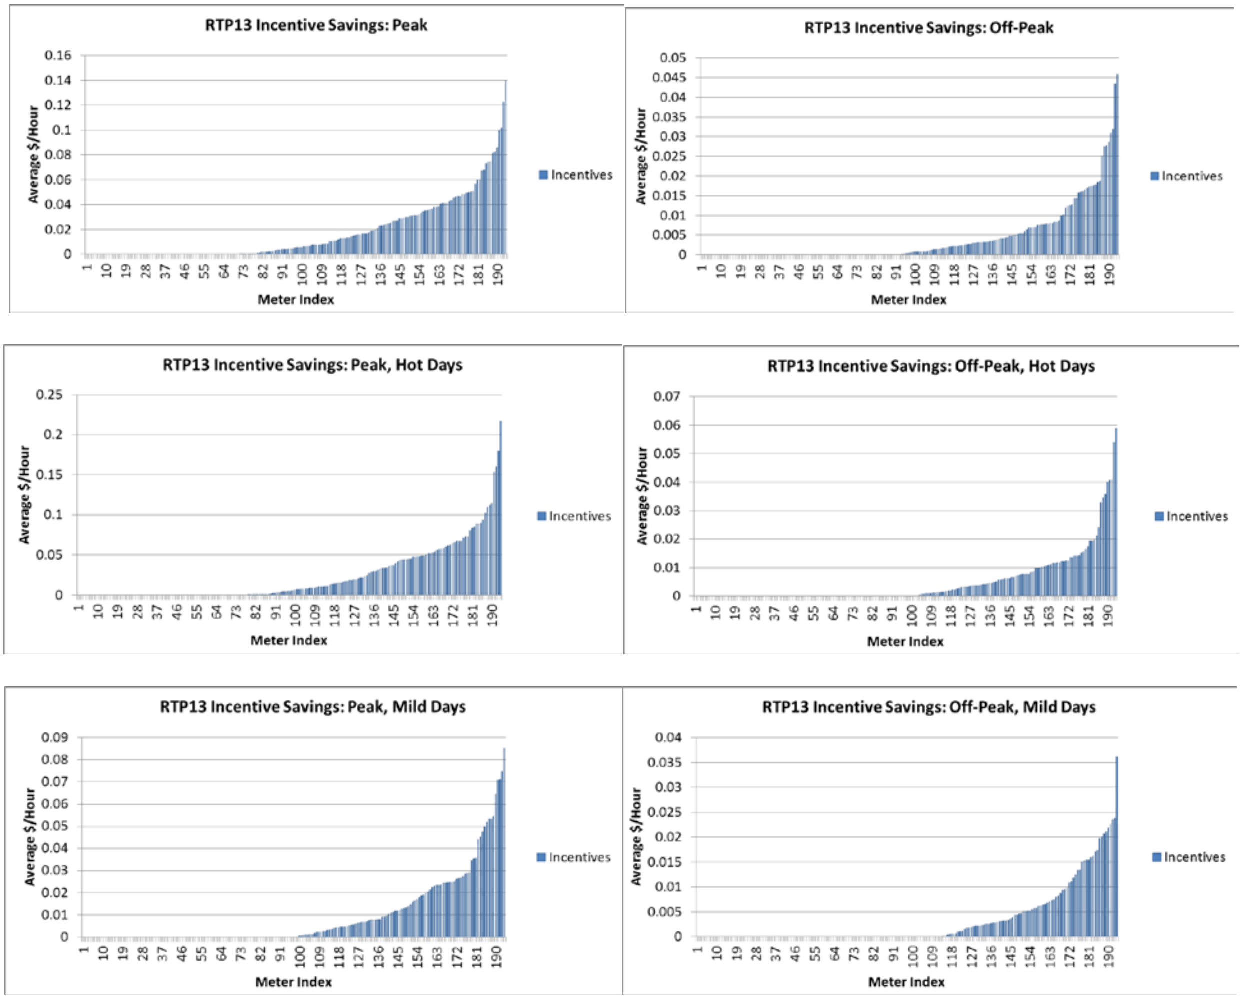

Figure 3.3. RTP_{da} Average Hourly Household Incentive Savings for Peak and Off-Peak Periods

A review of the incentive savings indicates that the peak periods have significantly more savings than the off-peak periods. This is to be expected, as there is more HVAC resource available for market-based curtailment in the peak periods. The difference is also likely increased because more congestion experiments were run during peak periods (see Section 1.3.2). The hot days also have significantly more incentive savings than the mild days, again likely because there is more HVAC resource available to participate in the market.

The following subsections describe comparisons of the RTP_{da} bill with bills for the same households subjected to the standard tariff (RTPstd). Since the RTP_{da} congestion experiments called upon the households more frequently than would be expected in typical operations, an additional section is provided that compares the bills of RTP_{da} to non-RTP_{da} households.

3.1.1 RTP_{da} Household Bills Compared with Standard Tariff

Figure 3.4 presents the total RTP_{da} bill with the total standard tariff bill. Each "index number" is one household in one month (which may be June, July, August, or September); some households will be listed four times (June-September), while others may only appear once. Household bills that had less than 80% data acquisition through the RTP_{da} dispatch system were removed from this analysis, due to large variability and errors in the bill calculation. Figure 3.4 orders the household bills by kWh consumed over the four-month period and compares the households' RTP_{da} bills against the same households' bills as if they had been charged the standard tariff. One can see that there is a wide spread of relatively small bill increases and decreases at all levels of consumption; this is more clearly shown in Figure 3.5.

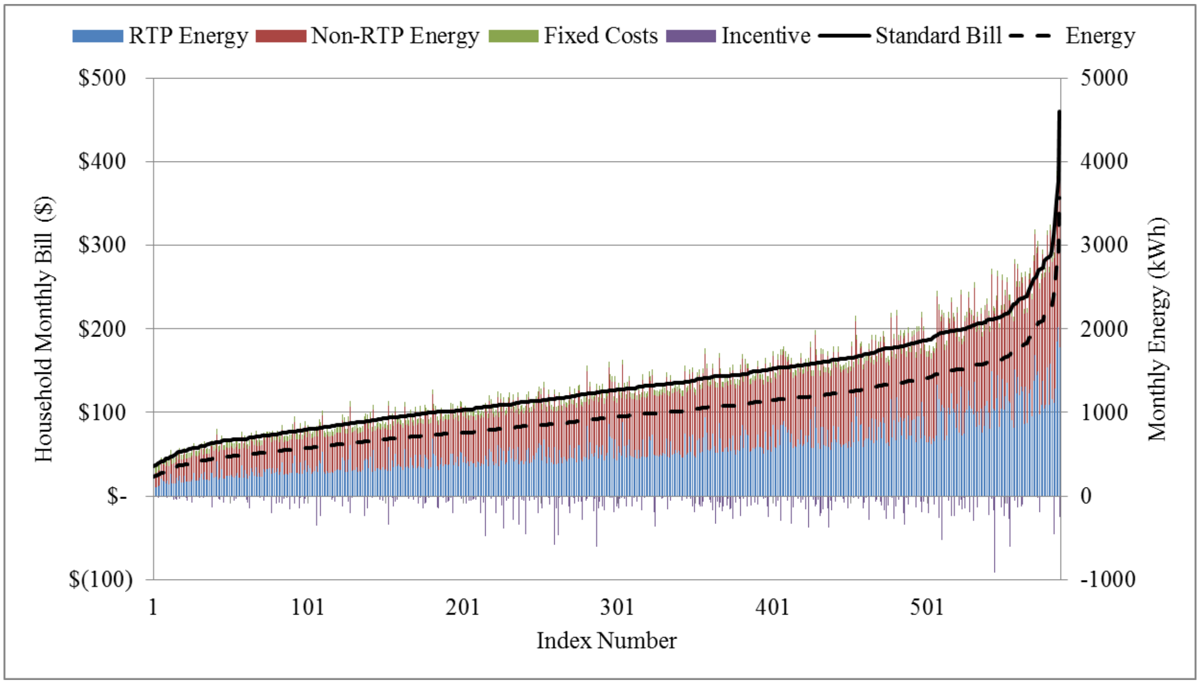

Figure 3.4. RTP_{da} Bill Comparisons with Standard Tariff Applied to the Same Households

The distribution of the difference between the RTP_{da} bills' without incentive savings compared to the same energy consumption calculated with the standard tariff is plotted in Figure 3.5. Again, this is sorted by monthly energy consumption. Figure 3.6 re-sorts the data by the change in overall bill. The two figures indicate that slightly more than half of the households were paying less under the new tariff; however, the households paying more did so by a greater amount, and on average, household bills were increased by $3.68, evenly spread across all household sizes. However, Figure 3.7 and Figure 3.8, in which the incentive is included in the bill calculation, show that many of the households were saving by switching to the new tariff; on average, household bills were decreased by $1.99, again spread across the sizes of the households. This indicates that the large number of congestion experiments may have made the prices appear lower (causing precooling) and disrupted the revenue neutrality calculations. However, when including the incentive, a majority of households were saving, indicating that by responding to the congestion events households are able to see savings.

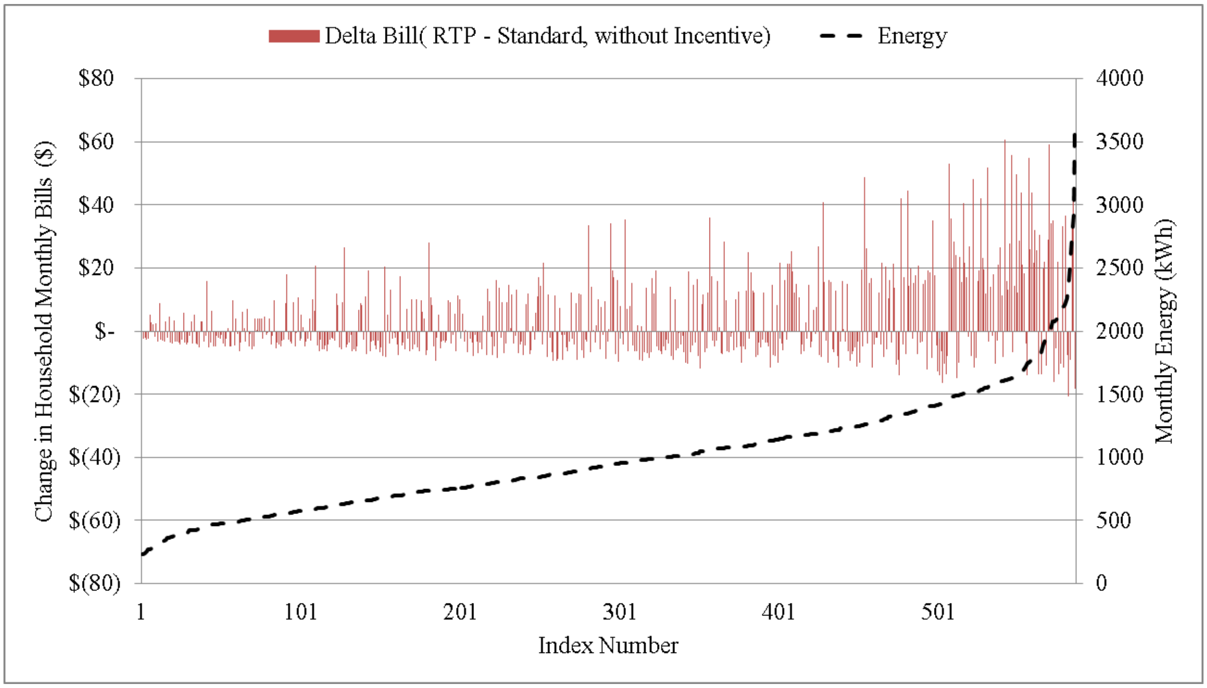

Figure 3.5. RTP$_{da}$ Bill (Without Incentive Component) Minus Standard Tariff Bill

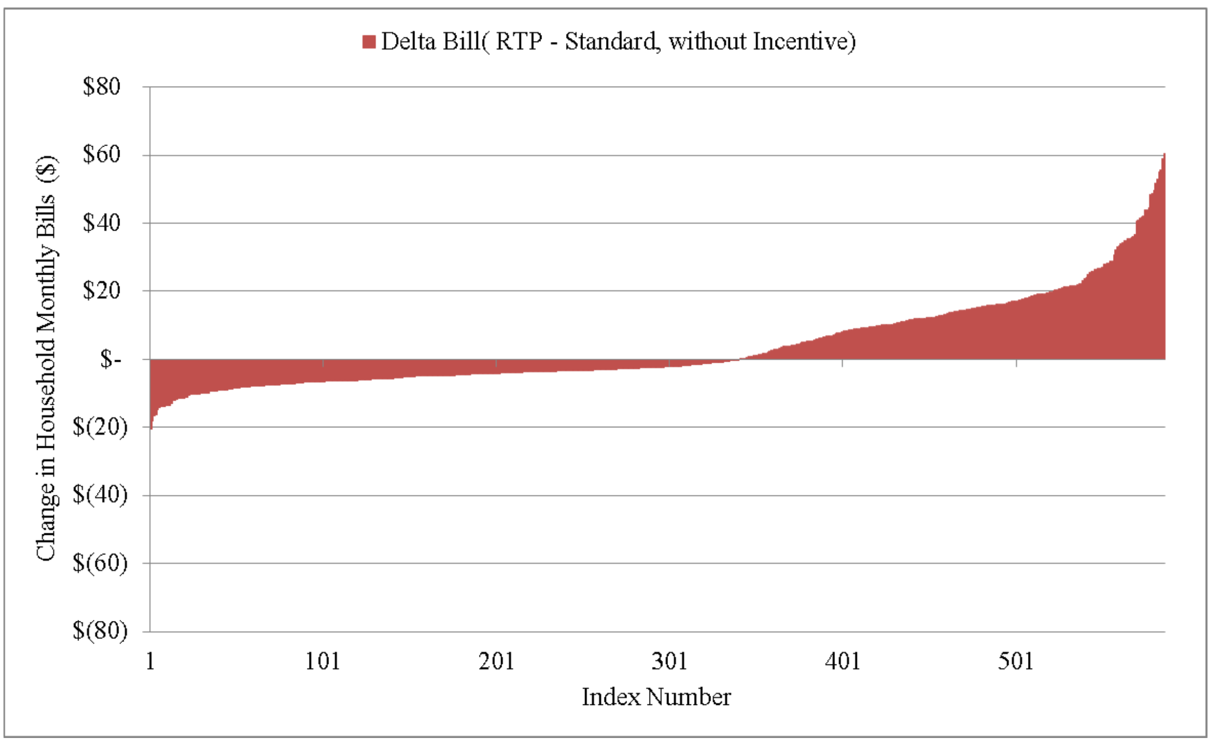

Figure 3.6. RTP$_{da}$ Bill (Without Incentive Component) Minus Standard Tariff Bill, Sorted by Change in Bill

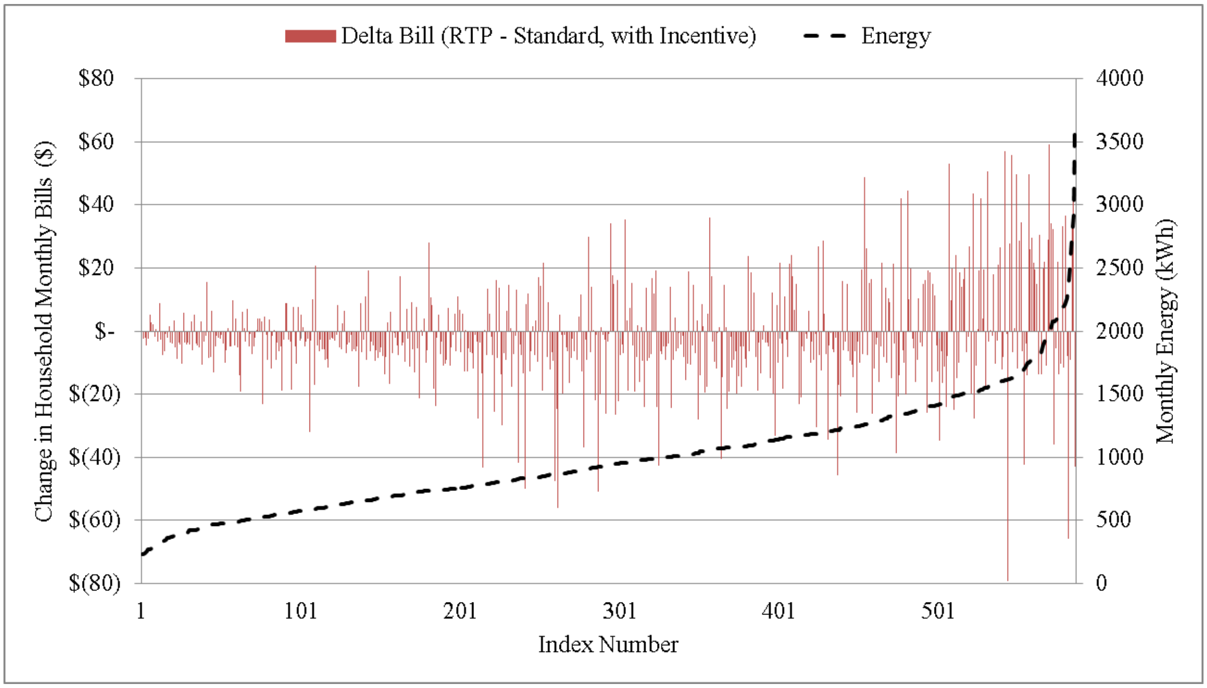

Figure 3.7. RTP_{da} Bill (With Incentive Component) Minus Standard Tariff Bill

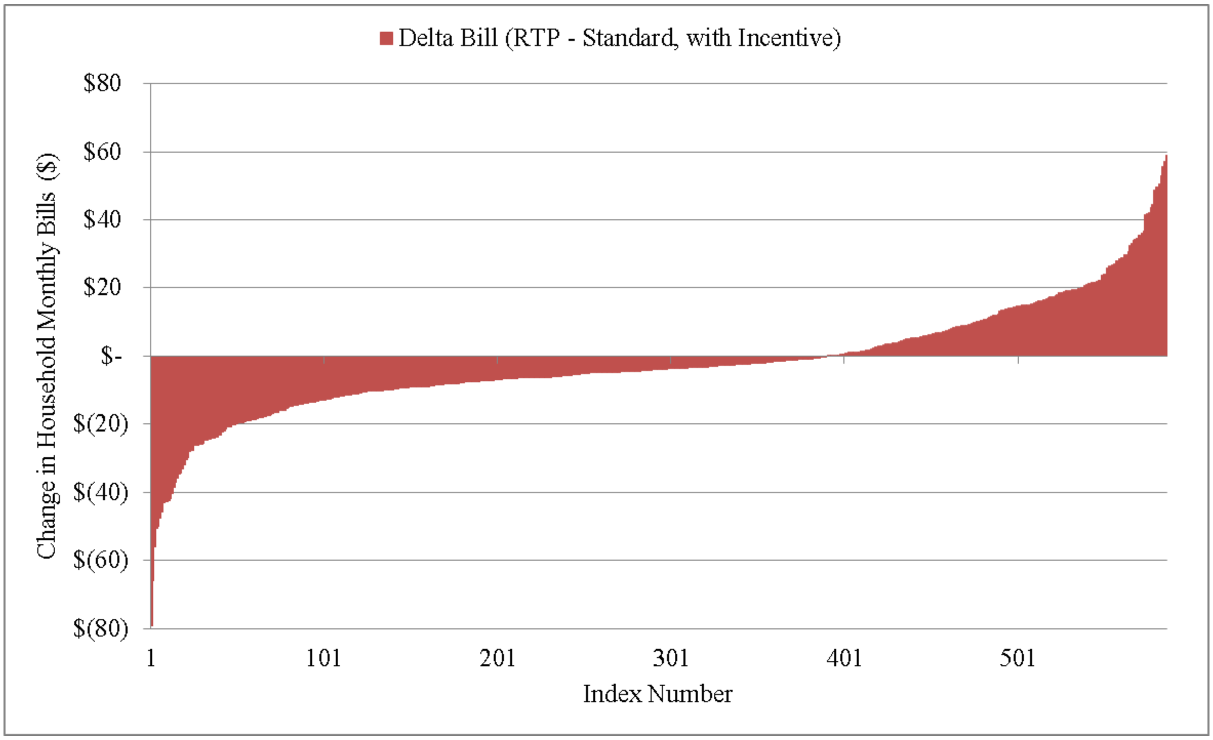

Figure 3.8. RTP_{da} Bill (With Incentive Component) Minus Standard Tariff Bill, Sorted by Change in Bill

3.1.2 RTP_{da} Bills Compared to Control Group

The average $/hr billing information over the months of June through September for all households in the RTP_{da} group and the RTPnr control group, and their percentages with respect to the average of all households' RTP_{da} bills, are presented in Table 3.3 and Table 3.4. The fixed-cost portion of the bills is the same for both RTP_{da} and RTPnr bills; it is listed in Table 3.1 and not shown here, so the bill components will not sum precisely to the total bill. The difference represented in the Delta Energy column only shows the energy component and does not include the incentive savings component of the bills.

The RTP_{da} billing information is calculated as reported in Section 3.1. The RTPnr billing information is calculated from the 15-minute metered data obtained for the households in the control group, then averaged and adjusted as explained in Section 1.4.2 to create an average RTPnr household energy use for each hour. The standard tariff is then applied to this average energy use to calculate the total bill and the energy component. As with the RTP_{da} bill calculation, discrepancies may exist between the RTPnr average billing calculation and the actual bills for the control group.

Table 3.3. $/hr Averages for RTP_{da} Bill Households Compared with Control Group (RTPnr)

Metric	RTP13				RTPnr13			Delta		
	Total Bill	kWh	Energy	Incentive	Total Bill	kWh	Energy	Total Bill	kWh	Energy
Average $/hr Total	0.1916	1.4812	0.1909	0.0081	0.2016	1.4516	0.1928	-0.0101	0.0296	-0.0019
Average $/hr Off-Peak	0.1506	1.2089	0.1457	0.0039	0.1666	1.1878	0.1578	-0.0160	0.0211	-0.0120
Off-Peak Mild	0.1204	1.0101	0.1148	0.0031	0.1381	0.9736	0.1293	-0.0178	0.0365	-0.0145
Off-Peak Hot	0.1754	1.3727	0.1712	0.0045	0.1939	1.3934	0.1851	-0.0185	-0.0207	-0.0139
Average $/hr Peak	0.2736	2.0257	0.2812	0.0164	0.2717	1.9792	0.2629	0.0019	0.0465	0.0183
Peak Mild	0.1891	1.5225	0.1884	0.0081	0.2056	1.4817	0.1968	-0.0165	0.0408	-0.0084
Peak Hot	0.3405	2.4247	0.3547	0.0230	0.3283	2.4055	0.3195	0.0122	0.0192	0.0352

Table 3.4. $/hr Average Percentages of Average Total RTP_{da} Household Bill Compared with Control Group (RTPnr)

Metric	RTP13				RTPnr13			Delta		
	Total Bill	kWh	Energy	Incentive	Total Bill	kWh	Energy	Total Bill	kWh	Energy
Average $/hr Total	100.0%	100.0%	99.7%	4.2%	105.3%	98.0%	100.7%	-5.3%	2.0%	-1.0%
Average $/hr Off-Peak	78.6%	81.6%	76.1%	2.0%	87.0%	80.2%	82.4%	-8.4%	1.4%	-6.3%
Off-Peak Mild	62.8%	68.2%	59.9%	1.6%	72.1%	65.7%	67.5%	-9.3%	2.5%	-7.6%
Off-Peak Hot	91.6%	92.7%	89.4%	2.4%	101.2%	94.1%	96.6%	-9.6%	-1.4%	-7.2%
Average $/hr Peak	142.8%	136.8%	146.8%	8.5%	141.8%	133.6%	137.2%	1.0%	3.1%	9.6%
Peak Mild	98.7%	102.8%	98.4%	4.2%	107.3%	100.0%	102.7%	-8.6%	2.8%	-4.4%
Peak Hot	177.8%	163.7%	185.2%	12.0%	171.4%	162.4%	166.8%	6.4%	1.3%	18.4%

Overall, the average $/hr savings in the bills of all households in this analysis is about 5% in RTP_{da} versus the RTPnr control group; however, the overall energy consumption is about 2% higher. When one looks at the sensitivity of the bills to the different types of operating periods, further insights can be gained. For the off-peak periods, the RTP_{da} bills show a slightly greater savings compared to the control group even though their energy usage is slightly higher.

A potential reason for this is the ability of the RTP_{da} households' HVAC units to respond to market price fluctuations in the off-peak periods. As explained in Section 2.2, the many congestion experiments performed had the effect of desensitizing the thermostat controllers to high prices. This had the effect of making prices that were not near the market cap appear to be bargains for a significant period of time after a congestion experiment. This could have resulted in overcooling. When market prices remained high after the congestion experiment, the effect was to use more energy during a normally high-price period. This phenomenon appeared to be emphasized when looking at the average $/hr during the peak+hot period. In this case, the energy consumption was 1.3% higher for the RTP_{da} group than for the control group; however, the energy component of the bill was 18.4% higher and the total bill was 6.4% higher. This is likely because, on average, the additional energy was being purchased at high market prices relative to the standard tariff. The effect of the incentive savings during these periods was to significantly reduce the impact of the large energy component on the overall average total bills.

When looking at peak+mild days, a greater amount of energy was used by the RTP_{da} group on average; however, the total bills were reduced by 8.6% compared to the control group, likely because the mild weather suppressed the market prices. Similar savings were seen in all off-peak figures, likely due to the lower market prices during the off-peak periods.

To summarize, the bill comparison between the RTP_{da} households and the RTPnr control group indicated bill savings in the summer months. More-detailed examination of the behavior of the RTP_{da} group in different periods of operation and the changes in the bill components revealed a variety of differences between the two bills. The low penetration of RTP_{da} households on each feeder and the frequent congestion experiments had a large impact on the behavior of the RTP_{da} resources and their interaction with the market. In addition, the accuracy of the representation of the control group as a "non-responsive" reflection of the RTP_{da} households deserves further scrutiny. More investigation is needed to fully understand these impacts.

One approach to isolate the impact of the congestion experiments as well as to compare results with a "perfect" control group is to model the RTP_{da} system using the GridLAB-D simulator. Section 3.1.3 reports the results of the use of simulation to both increase the penetration of RTP_{da} households and independently look at the performance of the RTP_{da} system with and without congestion experiments.

3.1.3 RTP_{da} versus Non-RTP_{da} Bill Comparison – Simulation

Simulations of the RTP_{da} group with controls and the same households without controls have been executed in GridLAB-D. The simulated households have been configured to represent the sizes and types of housing in the RTP_{da} group. The observed RTP_{da} household thermostat statistics and energy usage information have been used to calibrate the simulated households. This section reports the comparison of their bills. Bill comparisons include the summer months without congestion events and with congestion events to better understand the impacts.

Households were simulated within GridLAB-D to represent 25% penetration of RTP_{da} households; 300 households were "experimental" while 900 were operated similarly in each simulation and did not respond to variations in price. The experimental households were run using four different scenarios:

1. Control – The households were simulated using the standard pricing tariff (Schedule R-R, 2012).
2. RTP_{da} Without Response – The households were simulated using the experimental residential real-time pricing service tariff (Schedule RS-RTP, 2012), but did not respond to price fluctuations (these households could be considered to serve a purpose equivalent to the RTPnr households described in the previous section, but with the advantage that in the simulator, they are precisely the same household models).
3. RTP_{da} – The households were simulated using the residential RTP_{da} service tariff and responded to wholesale price fluctuations in a manner similar to those observed in the pricing experiments (for example, thermostat slider and temperature settings, internal air temperature decay rates, etc.).
4. RTP_{da} Congested – The households were simulated using the residential RTP_{da} service tariff, responded to wholesale price fluctuations, and responded to capacity limits placed on the feeder aligned with the actual experiments.

The bills are calculated using all components of the tariffs, including all fixed, rider, and energy charges. The bills are presented as the average of all four months and the impact on that average monthly bill. Figure 3.9 shows the monthly billing impact when switching from the standard tariff to the RTP_{da} tariff, without changing load behavior. The left-hand axis (red bars) indicates the percentage by which the household's bill changes when switching from the standard tariff to the RTP_{da} tariff, where a negative number indicates a savings when moving from the standard tariff to the RTPda tariff. The right-hand axis (blue line) indicates the monthly energy consumption of the household, ranked from low to high (left to right). The percent differences are consistent across most household sizes with an average reduction of

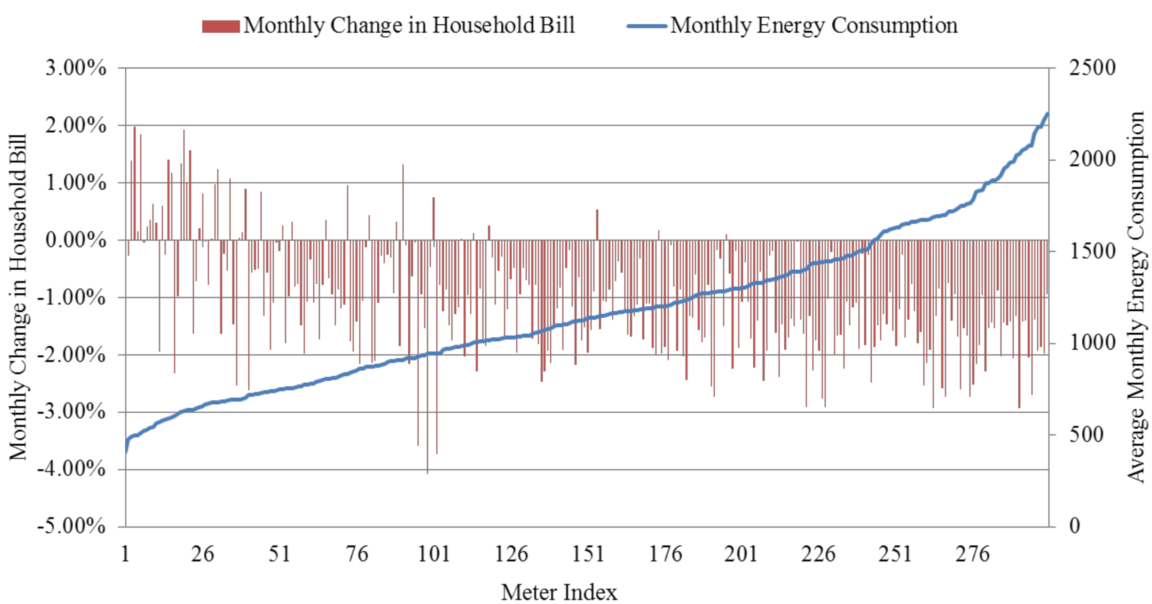

Figure 3.9. Change in Monthly Household Bills When Switching from Standard Tariffs to RTP_{da} (No Response), Without Responding to Price Fluctuations

1.1% in the bill. This indicates that during the four-month period, the RTP_{da} rate was nearly revenue neutral but slightly skewed toward decreasing the households' bills. The rate was designed to be revenue neutral over an entire year, but may show variance within any given period. The energy consumption is identical in these two cases.

Figure 3.10 shows a similar plot for households responding to wholesale price fluctuations. In this case, the difference in bills reflects moving from RTP_{da} Without Response to RTP_{da} (with response). The households are still stacked from left to right according to their energy consumption using the standard tariff (that is, Household 10 is the same Household 10 in each graph). This indicates the amount of savings seen by each household in responding to the price, with the effects of whether the rate is revenue neutral removed—in other words, the amount the household saves by allowing their thermostat to be adjusted in response to price fluctuations. The average reduction in the bill is 2.1%, with an average decrease in energy consumption of 1.2%.

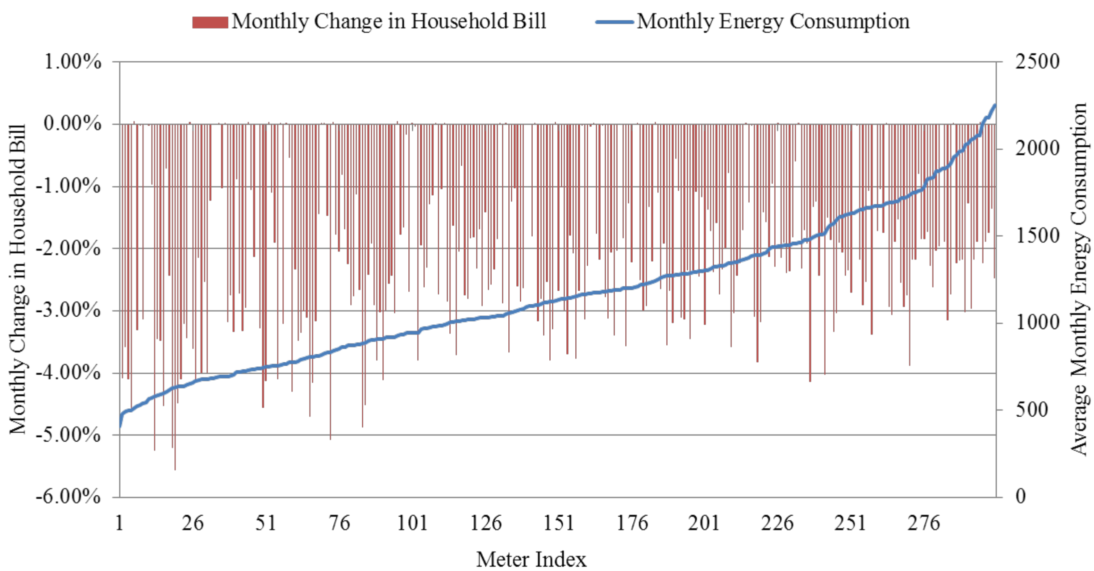

Figure 3.10. Change in Monthly Household Bills When Responding to Price Fluctuations (changing from RTP_{da} Without Response to RTP_{da} With Response)

Figure 3.11 is a combination of Figure 3.9 and Figure 3.10, moving from the standard tariff to RTP_{da} (with response). The average household reduces their bill by 3.2% and reduces energy consumption by 1.2%. Figure 3.12 shows the same information, but in terms of actual dollars saved (rather than percentage of bill). The average bill reduction is $5.11 with a maximum (average) reduction of $12.43 (one household was able to see a $22.52 reduction for the month of July). The reduction of the bill is consistent across household sizes in terms of percent reduction, with larger energy users seeing a larger decrease in proportion to their energy use.

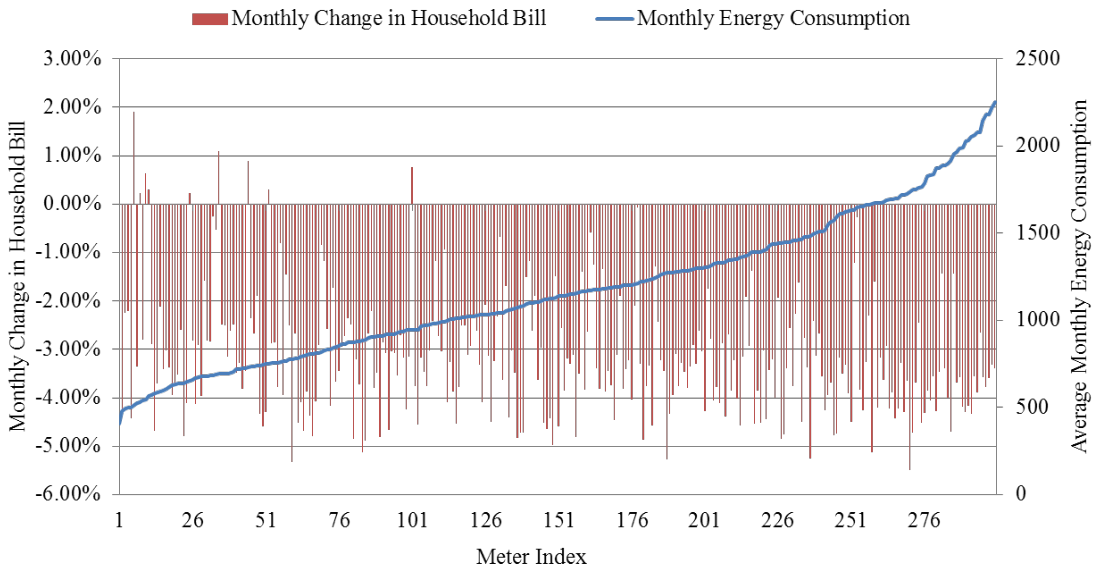

Figure 3.11. Percentage Change in Monthly Household Bills When Switching from Standard Pricing to RTP_{da} With Response

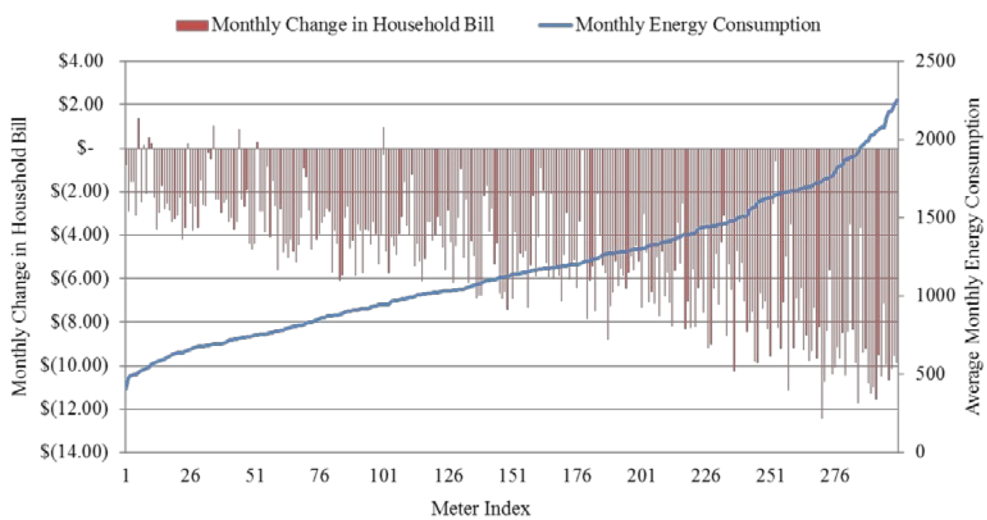

Figure 3.12. Change in Monthly Household Bills ($) When Switching from Standard Pricing to RTP_{da} With Response to Price Fluctuations

Figure 3.13 shows the impacts on the household monthly bill when moving from RTP_{da} with response to RTP_{da} Without Response with the 66 congestion experiments during the June-to-September period (out of 96 total congestion experiments done in all of 2013). The experiments caused prices to rise very high for a few hours and reduce demand, then "appear" very low for the following 24 hours as the average price was increased. This caused a number of units to lower their thermostat cooling set points and increase energy consumption (by 0.27% on average over the four-month period). For households that frequently responded to the congestion event by decreasing demand during the period, significant savings were seen (the maximum reduction was $31.16) driven by the incentive payment. Households that were not overly responsive saw a slight increase in their bills (between $0 and $3), driven by the increased demand caused by precooling. Ideally, the system would not be operated that often, so it is assumed that the level of impact and savings would be decreased.

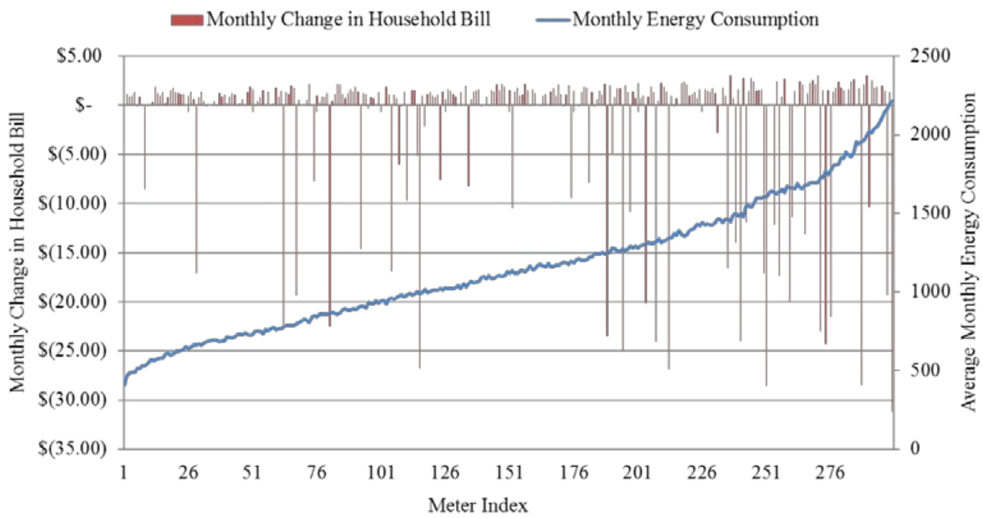

Figure 3.13. Change in Monthly Household Bills When Switching from RTP_{da} With Response to RTP_{da} Without Response, Taking into Account the Congestion Experiments Used in the Actual System

Table 3.5 through Table 3.8 show the average household bills by month for each of the cases discussed above. RTP_{da} cases include a breakdown into the base bill, rebate payments, and incentive payments. Note that the rebate payments approximately nullify the effects of increased prices due to the congestion experiments, but not quite (as indicated by the slight rise in most household bills during the experiments).

Table 3.5. Average Monthly Bill with Standard Tariff

	Average Household Bill Control Group (Standard Tariff)				
	June	July	August	September	Average
Total	**$153.14**	**$164.96**	**$161.19**	**$139.92**	**$154.80**

Table 3.6. Average Monthly Bill with RTP_{da} Tariff and No Response

	Average Household Bill RTP_{da} Without Response				
	June	July	August	September	Average
Base Bill	$148.25	$178.30	$152.40	$132.42	$152.84
Rebate	$–	$–	$–	$–	$–
Incentive	$–	$–	$–	$–	$–
Total	**$148.25**	**$178.30**	**$152.40**	**$132.42**	**$152.84**

Table 3.7. Average Monthly Bill with RTP_{da} Tariff and Response

	Average Household Bill RTP_{da} Wholesale Response				
	June	July	August	September	Average
Base Bill	$145.53	$172.97	$149.55	$130.71	$149.69
Rebate	$–	$–	$–	$–	$–
Incentive	$–	$–	$–	$–	$–
Total	**$145.53**	**$172.97**	**$149.55**	**$130.71**	**$149.69**

Table 3.8. Average Monthly Bill with RTP_{da} Tariff and 66 Congestion Experiments

	Average Household Bill RTP_{da} with Congestion Experiments				
	June	July	August	September	Average
Base Bill	$167.53	$232.78	$231.56	$199.80	$207.92
Rebate	$(21.38)	$(57.75)	$(80.80)	$(68.46)	$(57.09)
Incentive	$(1.08)	$(3.05)	$(3.35)	$(2.63)	$(2.53)
Total	**$145.08**	**$171.97**	**$147.42**	**$128.72**	**$148.30**

Table 3.9 summarizes the changes in household bills by month from the standard tariff to each of the three RTP_{da} experiments. Negative values (in parentheses) indicate a reduction in the bill. Notice the increase in July bills, even in the RTP_{da} Without Response case, indicating that wholesale prices were higher than expected in July.

Table 3.9. Comparison of Bill Reductions from Standard to RTP_{da} Tariff

	Delta Average Household Bill Control to RTP_{da}				
	June	July	August	September	Average
RTP_{da} Without Response	$(4.89)	$13.34	$(8.79)	$(7.50)	$(1.96)
RTP_{da} With Response	$(7.61)	$8.01	$(11.64)	$(9.21)	$(5.11)
RTP_{da} With Congestion	$(8.06)	$7.01	$(13.77)	$(11.20)	$(6.51)

3.2 Thermostat Statistics

This section explores the RTP_{da} consumers' interactions with their thermostat. A statistical characterization of the population of thermostat settings is presented, followed by an investigation of the thermostat override changes that occurred during congestion event periods.

3.2.1 Thermostat Settings

In the course of the Project, the consumers exercised their choice of setting the cooling and heating set points, as well as the comfort slider settings. In addition, they had the choice of overriding the system until the next scheduled period or indefinitely. A number of aspects of these choices can be studied, but the overall features will be considered first.

The period of analysis was the four-month period June 1–September 30, 2013. The occupancy status had four possibilities: "Home," "Night," "Away," and "Vacation." The day was divided into six parts of 4 hours each. Weekday and weekend differences are implicit in the occupancy status (that is, generally more hours of the day are in occupied status during the weekend), so no distinction was made for this study. The comfort setting has six possibilities: 0, 20, 40, 60, 80, and 100, with 0 being most comfort oriented and 100 being most economically oriented. The cooling set points covered a wide range: 55°F to 95°F. So the total number of bins will be (4 occupancy statuses) × (6 day periods) × (6 slider settings) × (41 cooling set points in 1°F increments), or 5904 bins. The amount of data available from a household was dependent on recruitment date, communication issues, and other matters. To normalize for this variability, each household was given one vote that could be distributed among the 5904 bins. Imagine each household receiving a sheet of paper of unit area. It can be torn into a maximum of 5904 pieces (often far fewer) in proportion to the fraction of time the house was in the state represented by a bin and placed in that bin. The areas of the pieces of paper in each bin were summed. These sums are shown in the graphs below, where a separate graph is drawn for each of the four occupancy statuses and six day periods, resulting in a possible 24 graphs. Each graph is further normalized so that the probabilities for each bin add up to 100%. "Away" status in the period midnight to 4 a.m. did not occur, so no graph is shown for that combination. The 23 graphs are shown in Figure 3.14, Figure 3.15, Figure 3.16, and Figure 3.17.

The trends seen in these graphs are generally self-explanatory. Additional studies exploring the changes during the course of the Project are possible but have not been performed. For example, a study of the default and initial comfort settings selected as part of the ePCT installation and training process could shed light on how these statistics evolved over time. As one demographic study, the impact of the size of the house on the settings was explored. Figure 3.18 shows the distribution of aggregated (overall occupancy modes and hours of day) overall occupancy statuses and day periods for the smallest 25% and the largest 25% of the houses as well as for all the households.

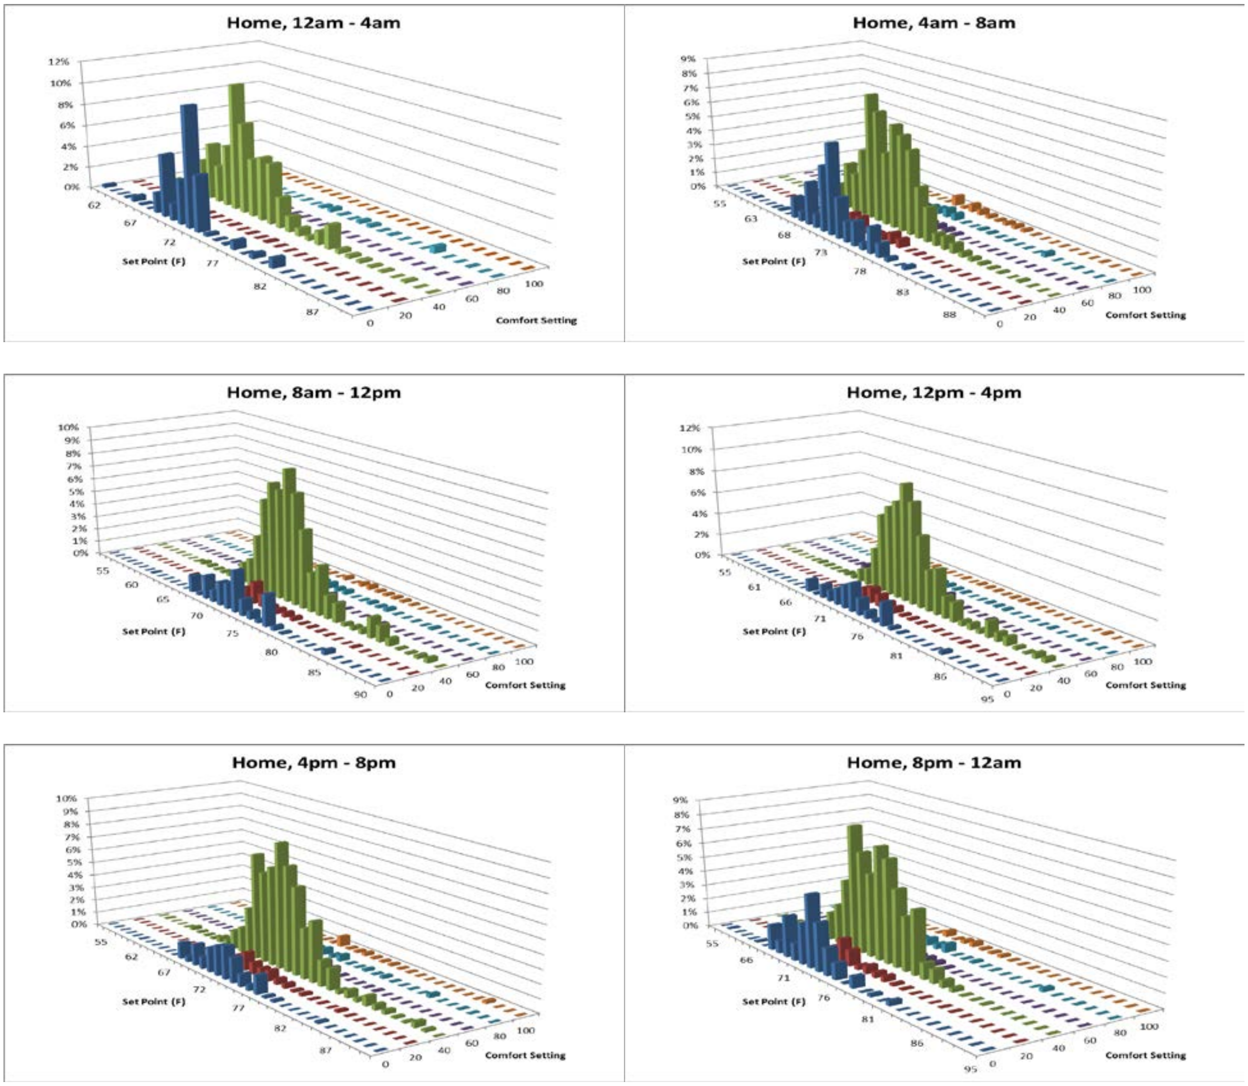

Figure 3.14. Cooling Set Point and Slider Distribution for Occupancy Status "Home"

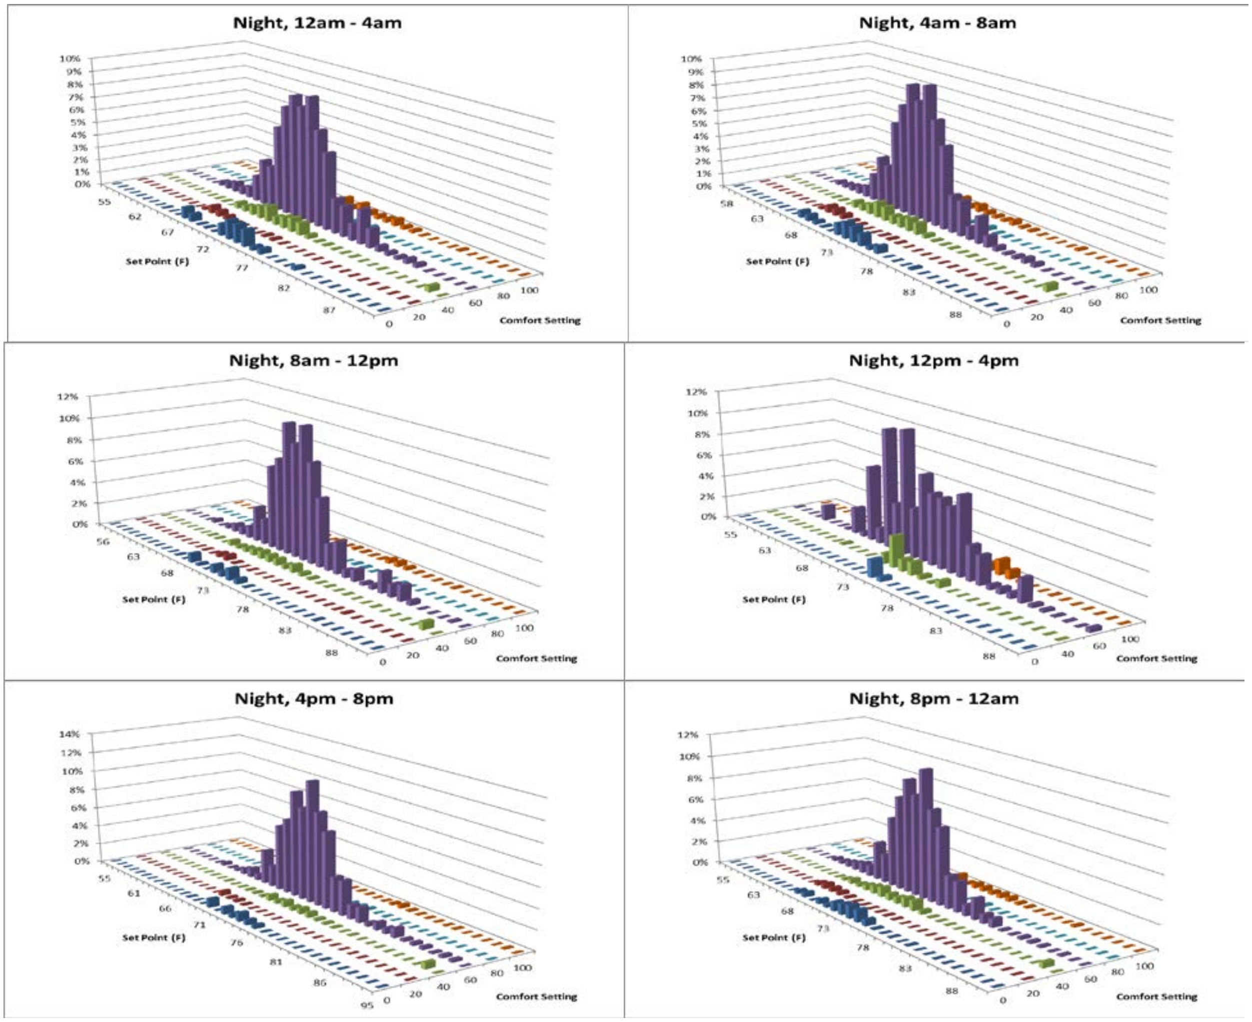

Figure 3.15. Cooling Set Point and Slider Distribution for Occupancy Status "Night"

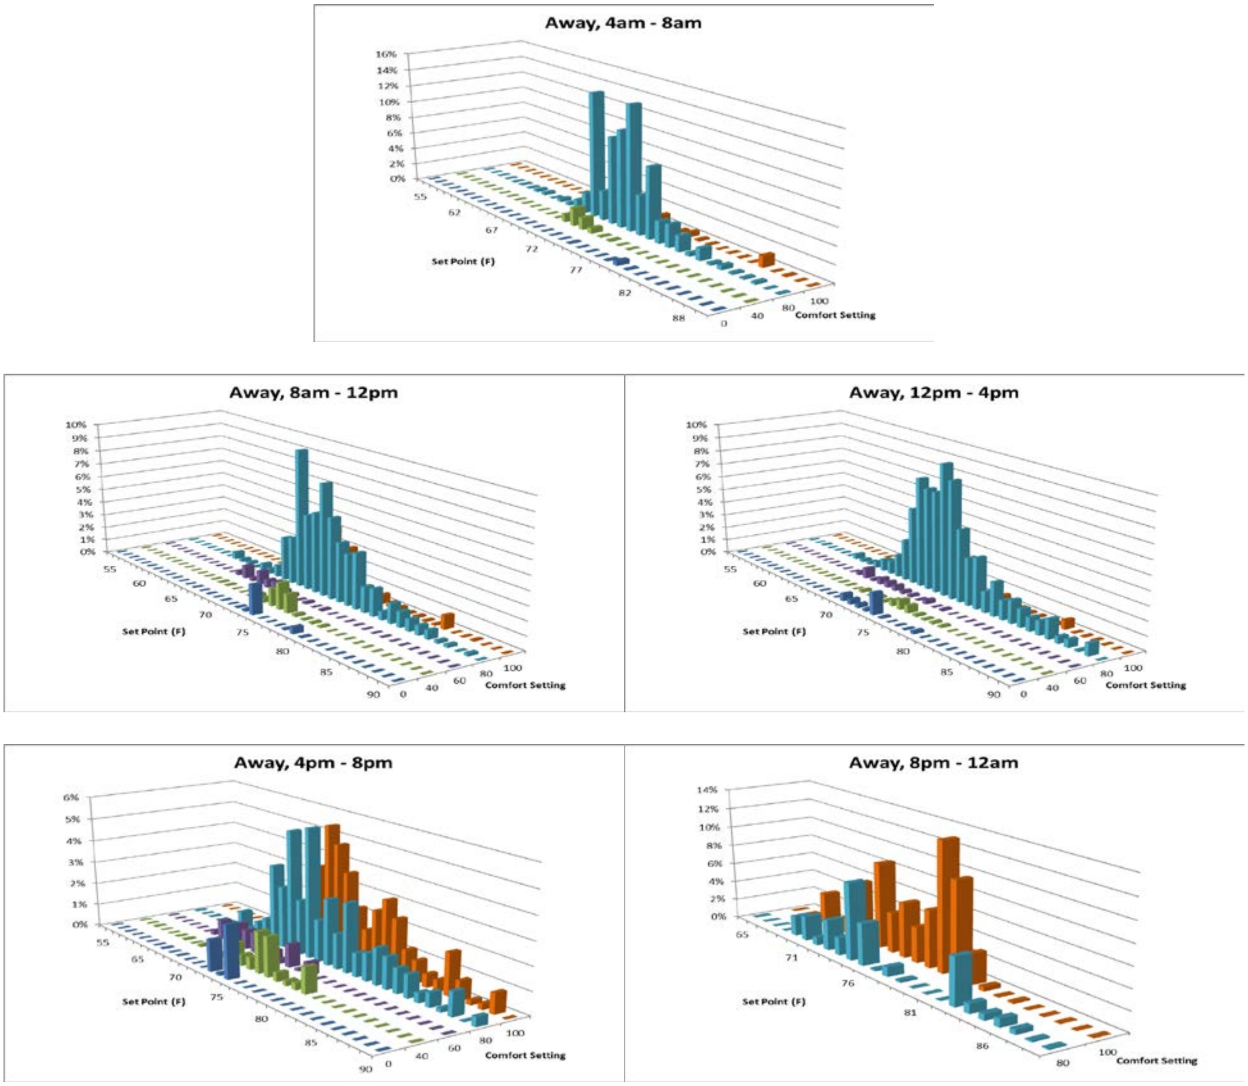

Figure 3.16. Cooling Set Point and Slider Distribution for Occupancy Status “Away”

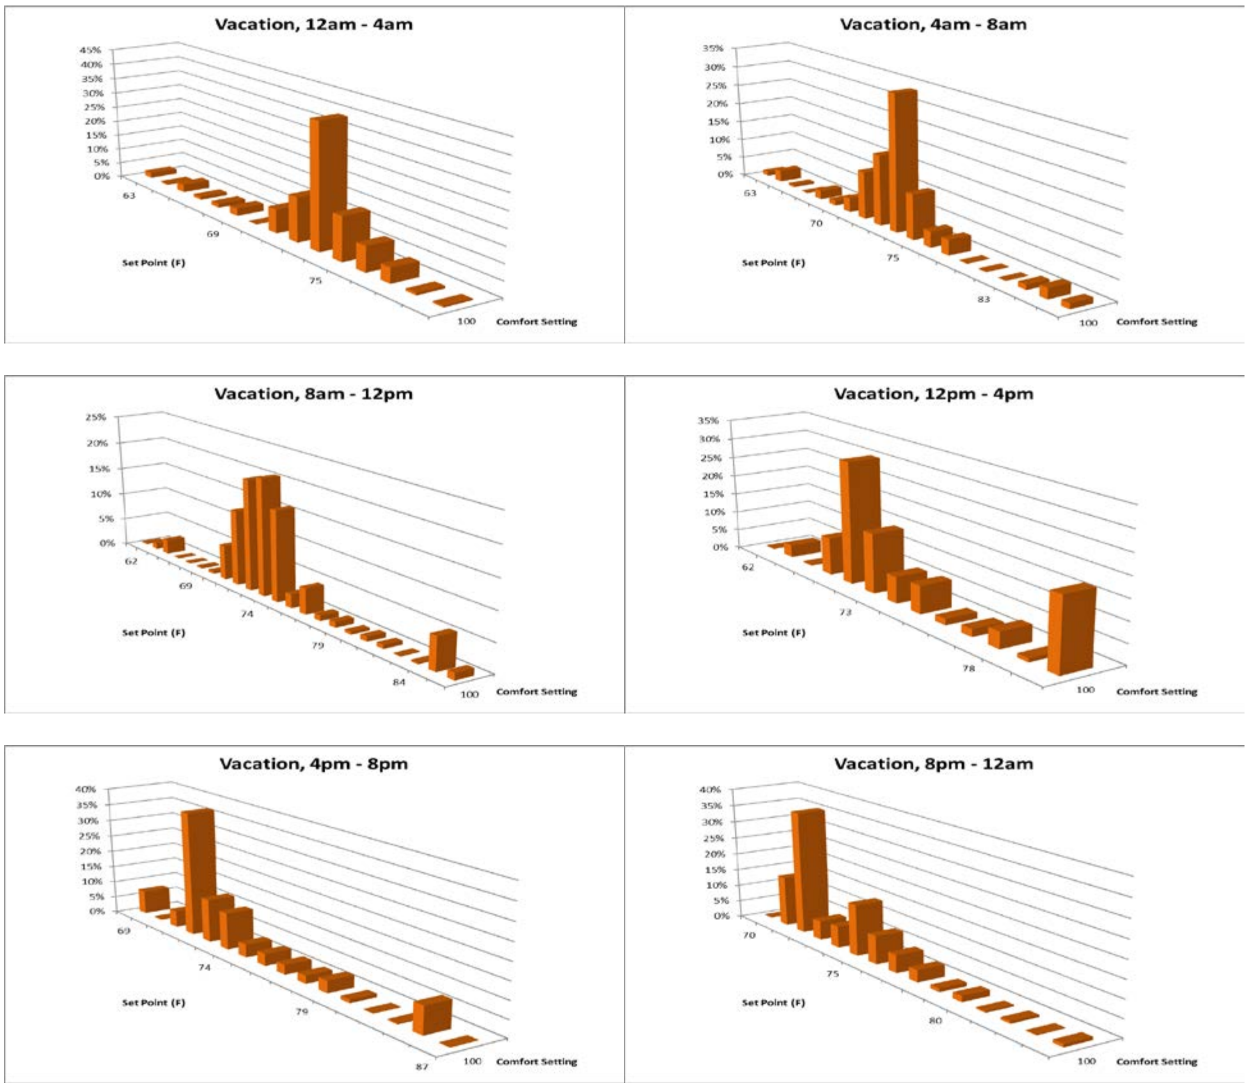

Figure 3.17. Cooling Set Point and Slider Distribution for Occupancy Status “Vacation”

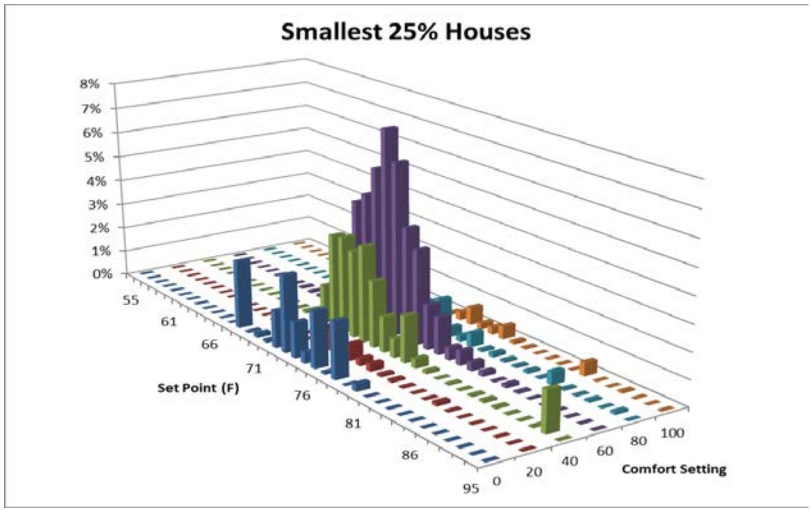

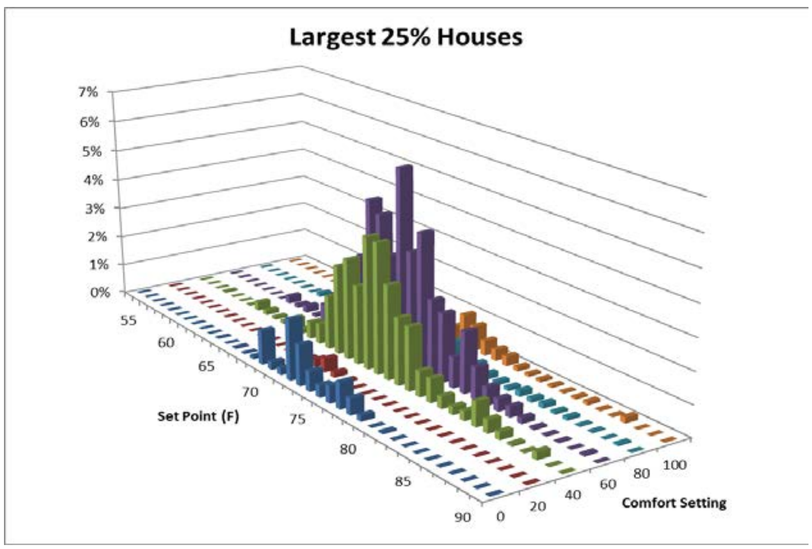

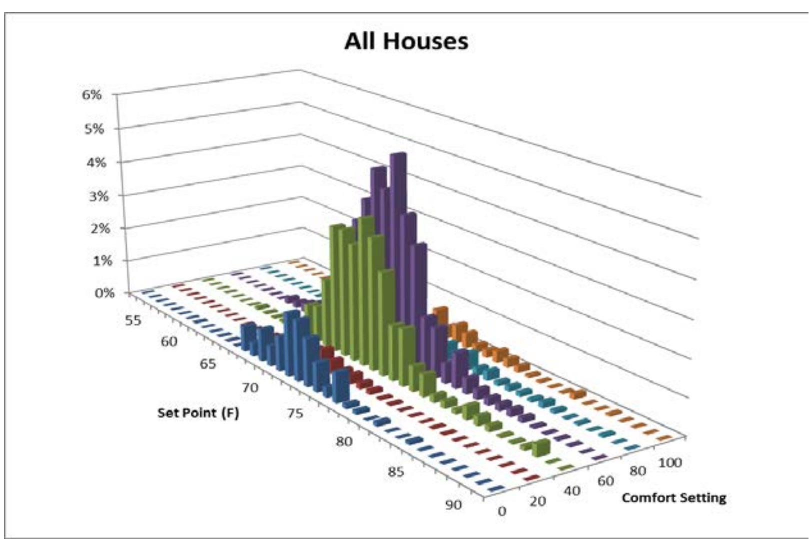

Figure 3.18. Effect of House Size on Cooling Set Points and Slider Settings Aggregated over All Occupancy Settings and Hours of Day

3.2.2 Thermostat Override Statistics

Table 3.10 shows the percentage of households that overrode their programmed thermostat settings during 2-hour and 4-hour congestion experiments. The override status was calculated as being positive for only those thermostats that were not in the override mode at the start of the experiment (that is, they were participating in the market), but at some point during the experiment were manually overridden. In 14 out of the total 69 2-hour experiments no thermostats were overridden, while only three 4-hour experiments recorded no overridden thermostats.

Table 3.10. Thermostat Override Statistics for 2-hour and 4-hour Congestion Experiments

2-Hour Experiments			4-Hour Experiments		
% Households that Overrode	Frequency	Probability	% Households that Overrode	Frequency	Probability
0%	14	20.00%	0%	3	11.54%
0–1%	13	18.57%	0–1%	3	11.54%
1–2%	25	35.71%	1–2%	4	15.38%
2–3%	10	14.29%	2–3%	3	11.54%
3–4%	5	7.14%	3–4%	2	7.69%
4–5%	3	4.29%	4–5%	3	11.54%
5–6%	0	0.00%	5–6%	3	11.54%
6–7%	0	0.00%	6–7%	3	11.54%
7–8%	0	0.00%	7–8%	0	0.00%
8–9%	0	0.00%	8–9%	1	3.85%
9–10%	0	0.00%	9–10%	1	3.85%
Total	**70**	**100%**		**26**	**100%**

Figure 3.19 compares the numbers of households that overrode their programmed thermostat settings during on-peak and off-peak 2-hour congestion experiments. It is evident that more households in override status were recorded during on-peak periods (14:00 – 22:00), as compared to off-peak period (22:00 – 14:00) experiments.

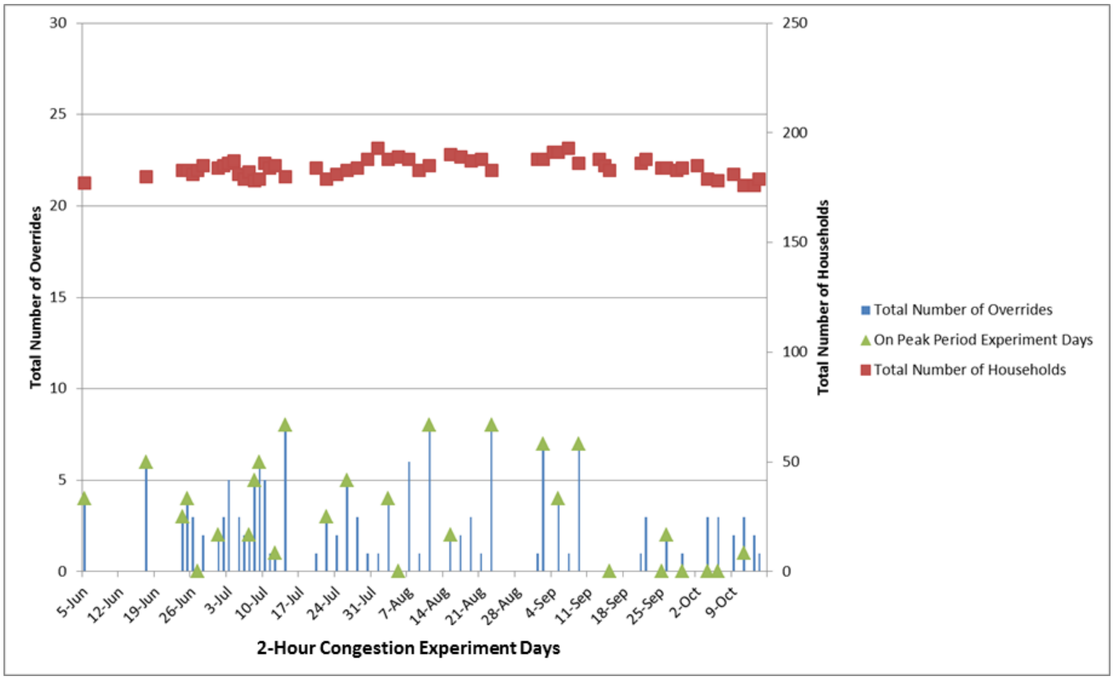

Figure 3.19. Thermostat Override Statistics during 2-hour Congestion Experiments

Figure 3.20 presents the override statistics recorded during 4-hour congestion experiments, which were conducted over the on-peak periods of the day. The figure also presents a comparison of the override statistics when the experiment was called during a SMART Shift Plus (critical peak pricing) event versus, along with the other 4-hour experiments. It is evident that a greater number of households overrode their programmed thermostat settings during SMART Shift Plus events.

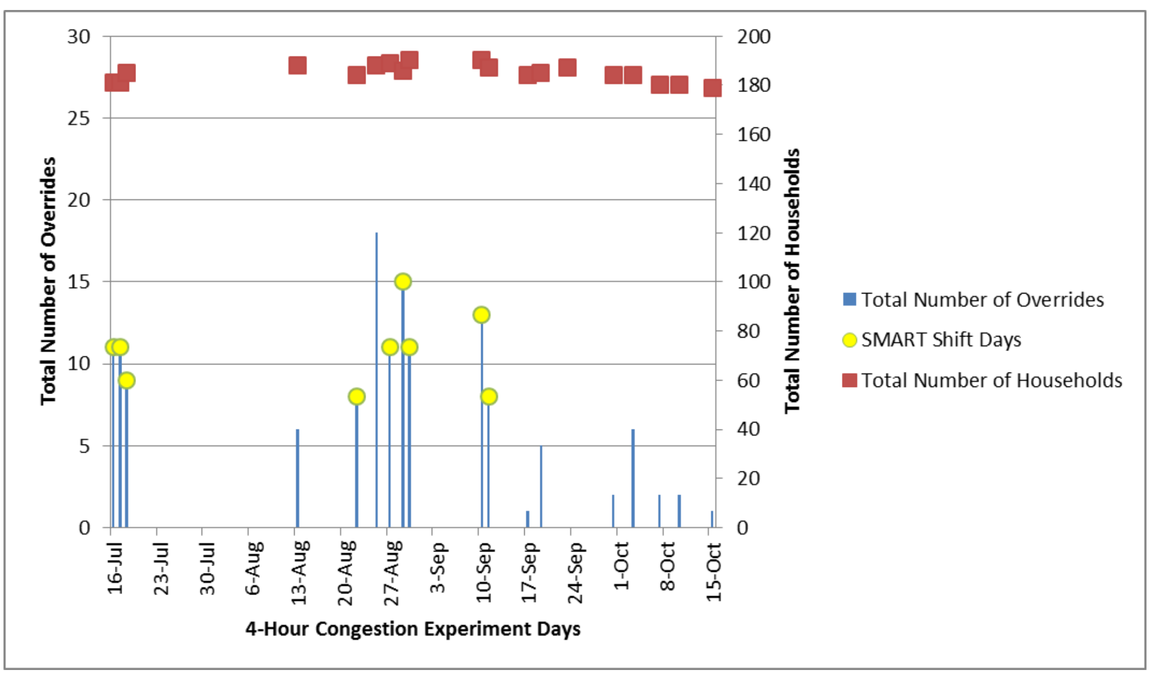

Figure 3.20. Thermostat Override Statistics during 4-hour Congestion Experiments

Figure 3.21 and Figure 3.22 present comparisons of the total numbers and the percentages of households that overrode their programmed thermostat settings during 2-hour versus 4-hour congestion experiments. The figures may be interpreted as override "duration" curves, presenting the numbers and percentages, respectively, of households in override status during 2-hour and 4-hour experiment periods. From both the figures, it is evident that a greater number of households overrode their programmed thermostat settings during 4-hour experiments than during 2-hour experiments. This may be attributed to greater discomfort due to the rising house temperatures during 4-hour experiments when HVACs stayed off for a considerably longer duration.

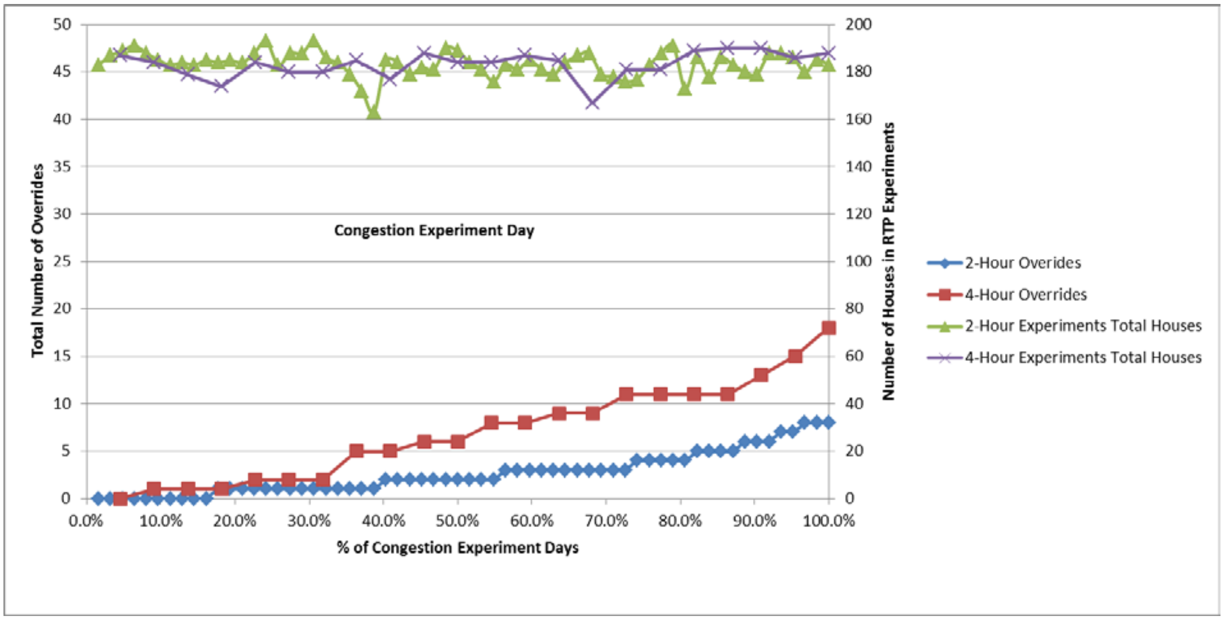

Figure 3.21. Override Duration Curves for 2-Hour and 4-Hour Congestion Experiments

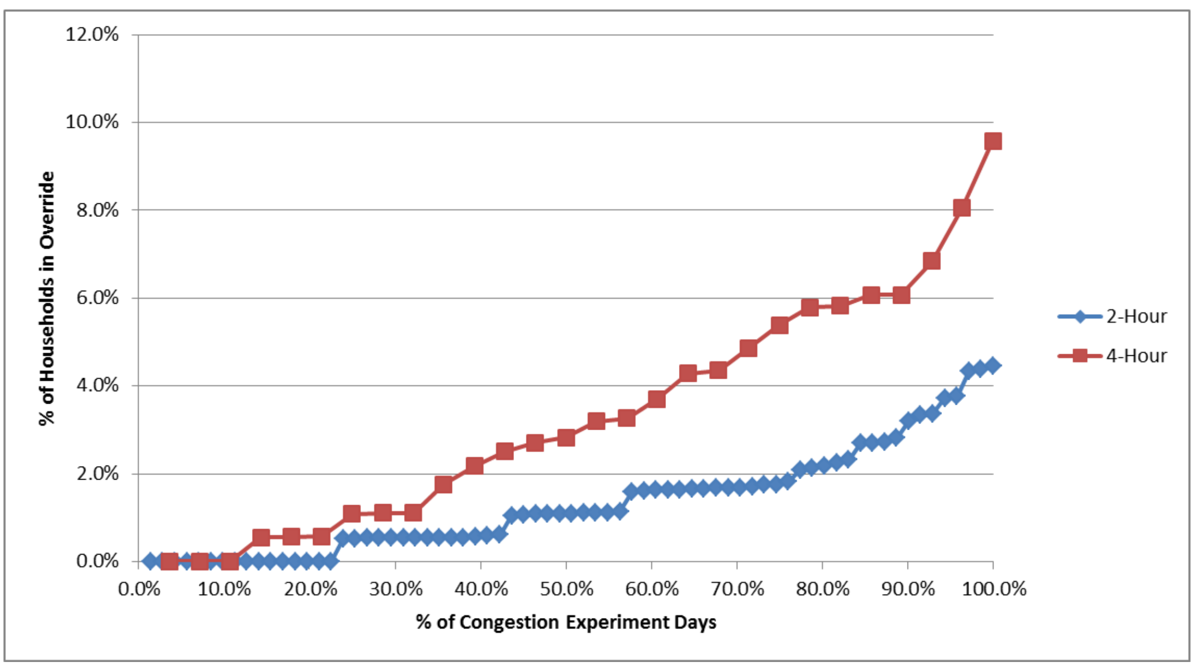

Figure 3.22. Override (% of total) Duration Curves for 2-Hour and 4-Hour Congestion Experiments

3.3 HVAC Bid Quantity versus Actual Load

When ePCT equipment was installed in a home, it was configured to store the estimate for the amount of power that the HVAC unit would draw when operating. The Project assumed that the compressors are fixed speed, which appears reasonable today, but will likely change in the future as HVAC units become more efficient. The installer estimated the power draw based on the nameplate rating and/or size of the compressor. A look-up table was provided to help convert HVAC heating/cooling size to the power draw. The estimated power draw was stored in the HEM equipment for use as the HVAC bid quantity (q_{bid}). The accuracy of this estimate when compared to the actual power draw could be important to the performance and stability of the RTP_{da} system under wide-scale deployment.

Analysis of the metered data for the RTP_{da} households was undertaken to determine the actual power drawn for each unit. The results of this analysis are presented below. Details of the methodology are left for a future publication.

3.3.1 Results of Analysis

Figure 3.23 plots the meter analysis value for each household as compared with the nominal value (q_{bid}) used in the RTP_{da} auction. If the values corresponded well, then the points in the graph would be expected to cluster closely around the dotted diagonal. Instead we see that, in general, the HVAC equipment is drawing more power than the household is bidding into the market. The distributions for 99 households of their bid power quantities and estimated power quantities are shown in Figure 3.24, indicating a significant deviation from the quantities bid.

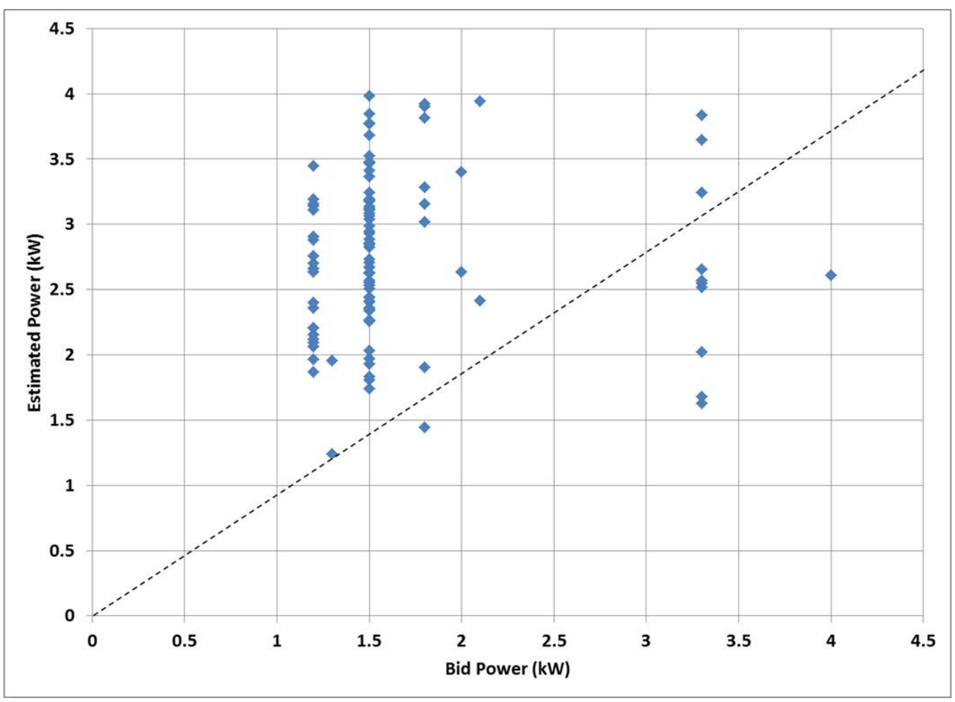

Figure 3.23. Estimated HVAC Power from Metered Data versus Bid Power for the Same Household

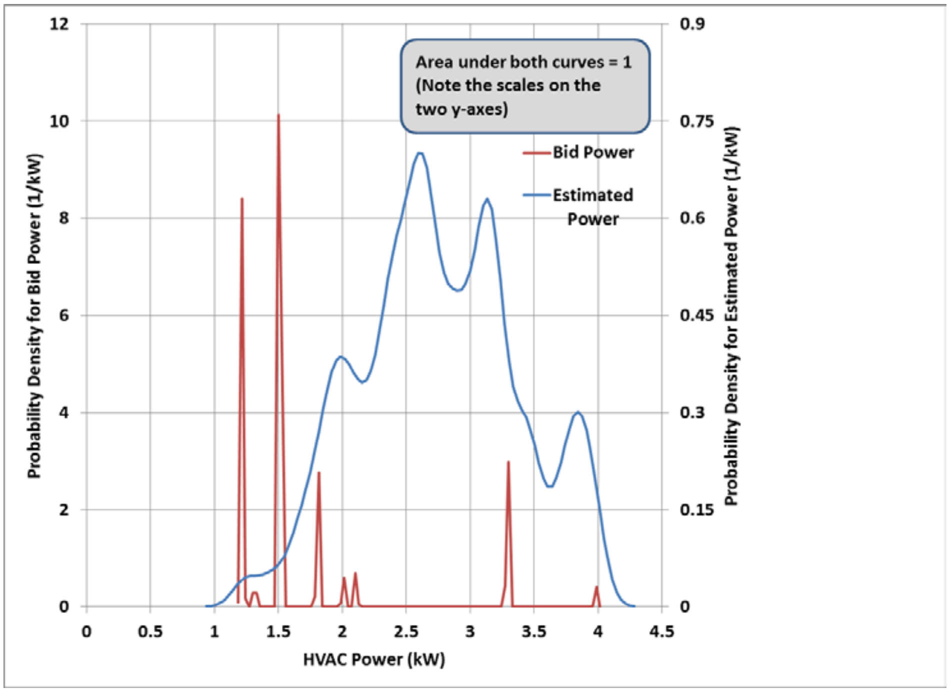

Figure 3.24. Distributions of HVAC Bid Power and Estimated Power

The linear vertical groupings of points in Figure 3.23 likely occurred because the look-up table listed a few key values: 1.2, 1.5, 1.8, 2.1, and 3.3. These values are also clearly apparent in the red line in Figure 3.24.

If the system becomes constrained either due to limited supply or due to congestion, the market cleared price is determined by price, energy bids. These bids should be accurate for proper operation. However, in this project, the responsive load was so small that the market either cleared at the feeder base price (non-congested situation) or at the feeder market cap (congested situation). In either case, the inaccuracy of the bid quantities did not affect the market clearing. As the energy-sensitive portion of the household bills were calculated from the metered quantity, this portion of the bill was not affected. However, the incentive saving calculation is dependent on the bid quantity, so inaccuracies can have an effect here.

Bid power inaccuracy would also have an impact on the analysis of load sensitivity to price when there is a high penetration of RTP_{da} resources on the feeder. In this case, a congestion situation can occur where the market clears between P_{base} and P_{cap}. The bid power inaccuracy may cause the market to clear at different energy quantity and price points (see Figure 3.25).

This situation is a topic for future analysis. A simulation of the RTP_{da} system could be set up to run two sets of scenarios using the PJM real-time market pricing information and the congestion event periods. One set of scenarios would run the households that bid q_{bid}, but size the HVAC models to match the distribution seen from the meter data analysis as in the demonstration. The other set of scenarios would adjust q_{bid} to accurately reflect the meter data analysis values. The scenarios would be scaled to show different penetrations of RTP_{da} households on the feeders. Comparisons of these runs would provide insight regarding the impact of bid power errors on the performance of the RTP_{da} system.

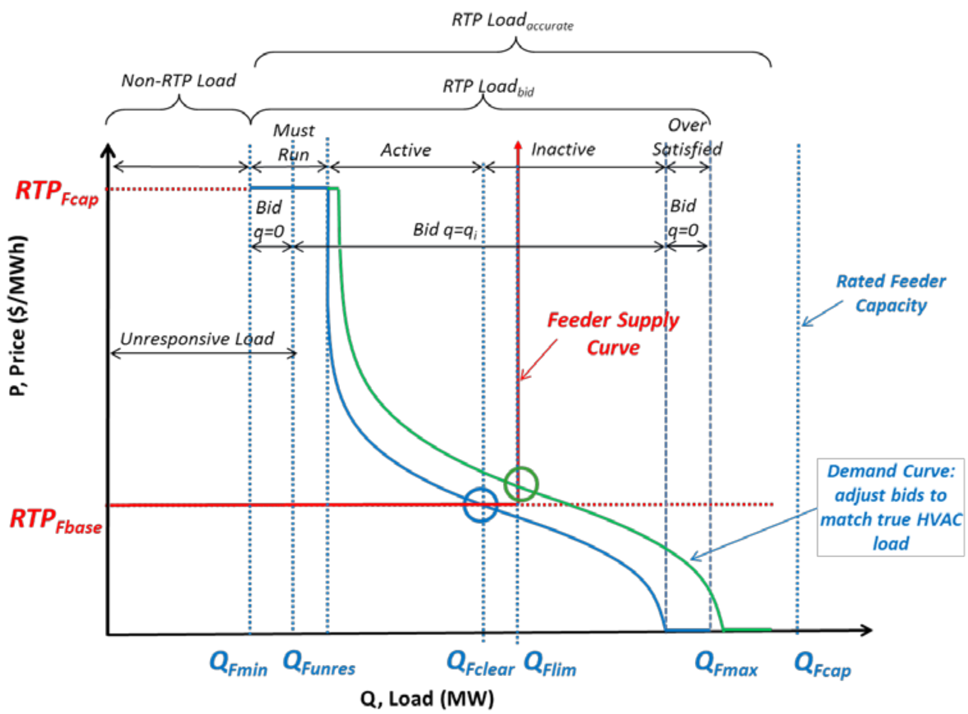

Figure 3.25. Conceptual Bid Curve Comparison of Bid versus Accurate HVAC Draw

4.0 RTP$_{da}$ Load Sensitivity to Price

This chapter investigates the price-responsive nature of the RTP$_{da}$ resources over the course of the summer. The first section analyzes data collected for each household to explore the population of RTP$_{da}$ resources' sensitivity to the 5-minute, LMP-based market price fluctuations experienced during the experiment. It is followed by a similar investigation done using the GridLAB-D simulator. The final section investigates the response of the RTP$_{da}$ resources to the imposed congestion events, where the market cleared at the price cap for the duration of the event.

4.1 Results from Measured Data

The responses of the thermostat agents are not really to energy prices but to variations in the prices around the "expected" price, the expected price being the average price experienced during the previous 24 hours. Consider the relationship between response defined as (RTP13 − RTPnr13) and LMP. The LMPs are used as proxies for the real-time prices that are used for the bids. An LMP of, for example, \$40/MWH may be perceived as high at some times and as low at other times. If perceived as high, the result is a tendency for reduction of energy use compared to the control houses. If it is perceived as a low price, the tendency is toward increased energy use. Therefore for a given LMP, the response at different times can vary over a range of values. This is seen in Figure 4.1. In this figure, the 5-minute average change in energy between RTP13 and RTPnr13 households for non-congestion experiment days is plotted against the corresponding time period's 5 minute LMP. A heavily smoothed response is shown as the blue points. Also shown is a histogram of the LMPs in the data points. The series of vertical streaks

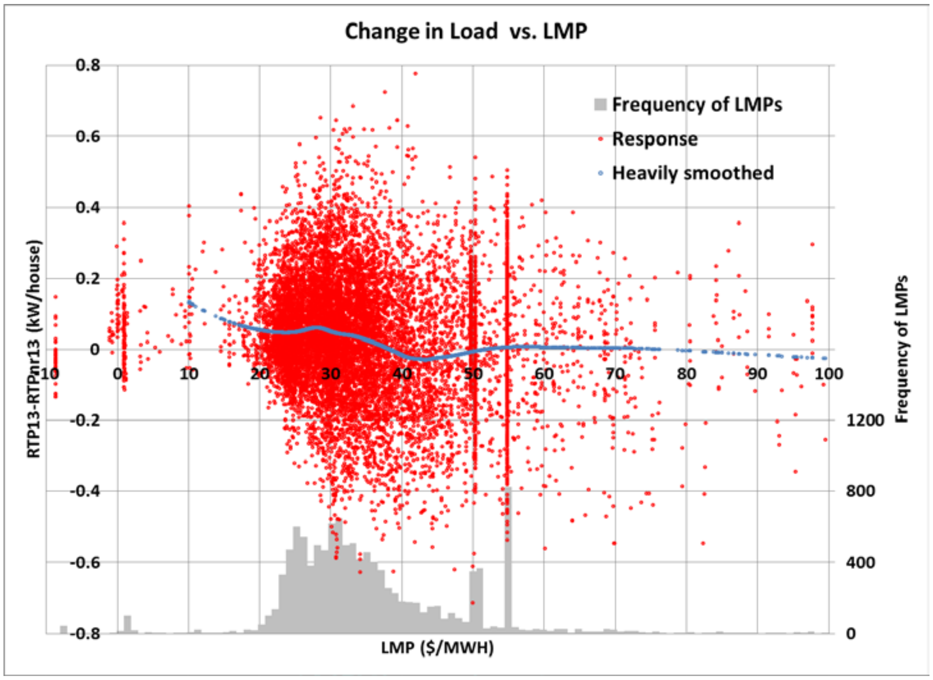

Figure 4.1. Response versus LMP for About 12,000 5-Minute Data Points Covering the Period June–September 2013

correspond to some default prices used when no LMP was transmitted. The LMPs are truncated at 100 \$/MWH because points beyond that are sparse. A histogram of the full range of LMPs from \$−8.67 to \$659.09 is shown in Figure 4.2.

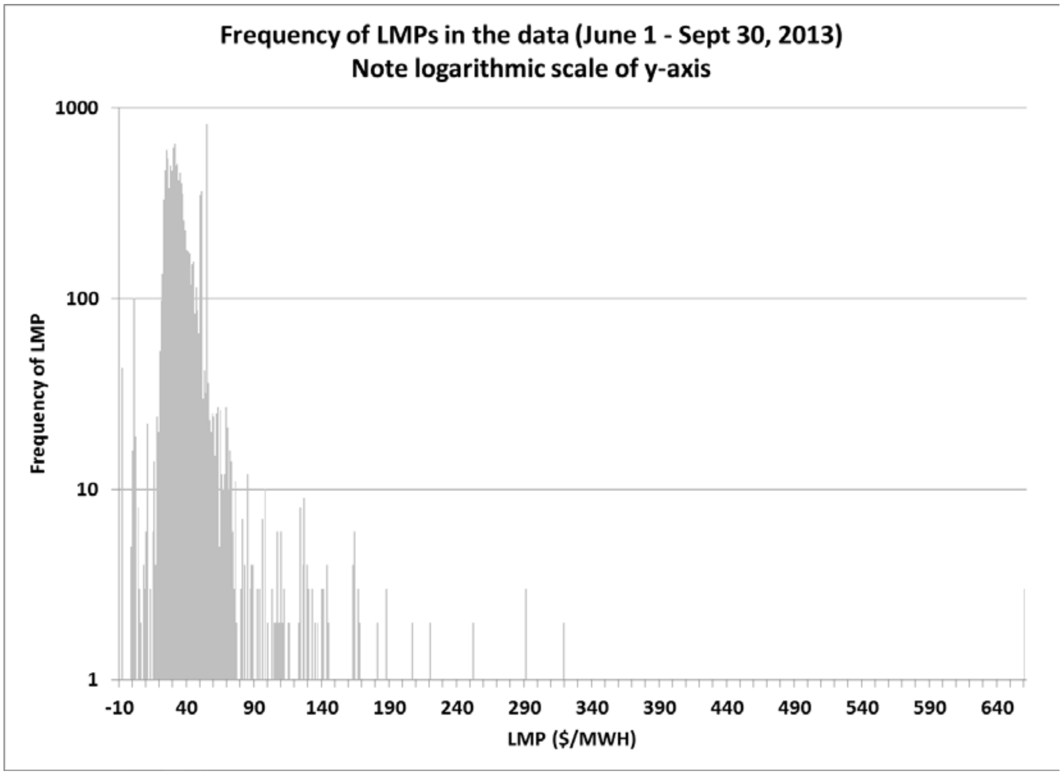

Figure 4.2. A Histogram of the Full Range of LMPs (from −\$8.67 to \$659.09) Seen During the Analysis Period

Even though the response is to variations in LMP and not to LMP itself, a heuristic expectation is that higher prices should generally result in a negative response. It is therefore of interest to determine the correlation coefficient between the response and LMP. This coefficient is −0.17, confirming heuristic expectations of the design of the RTP_{da} system.

4.2 Simulated Results

Similar to previous sections, households were simulated within GridLAB-D to represent 25% penetration of RTP_{da} households; 300 households were "experimental" while 900 were operated similarly in each simulation and did not respond to variations in price. The experimental households were modeled using three different scenarios:

1. Control – the households were simulated using the standard pricing tariff.
2. RTP_{da} – the households were simulated using the residential RTP_{da} service tariff and responded to wholesale price fluctuations in a manner similar to those observed in the pricing experiments (for example, thermostat slider and temperature settings, internal air temperature decay rates, and such).

3. RTP_{da} Congested – the households were simulated using the residential RTP_{da} service tariff, responded to wholesale price fluctuations, and responded to capacity limits placed on the feeder aligned with the actual experiments (96 experiments in four months).

The simulation results offer an additional comparison against the actual data, without the necessity of using filtering and regression techniques, as simulation provides a "perfect" control group. The graph in Figure 4.3 is similar to Figure 4.1, but uses the simulated response data (RTP_{da} minus Control) and includes every day within the four-month period (rather than a subset in the field data case where congestion days were removed). The graph is truncated as before and does not show some periods of very high price and high load response. The figure is presented to show the similarities between the simulated and actual responsive loads; particularly the area around LMPs of 30 to 60 $/MWh. A negative value on the y-axis indicates a reduction in demand when moving to the RTP_{da} group. It can be seen that when the LMP is relatively low (less than $30/MWh), there is almost no change in demand when moving to the RTP_{da} rate. Further research is needed, though this may be indicative of having very little resource available for reduction during relatively low LMP periods (e.g., early morning periods). However, as LMP becomes higher, there is a significant trend toward reduced demand. It should be noted that data points higher than $60/MWh are of much lower density and the trend line is less certain.

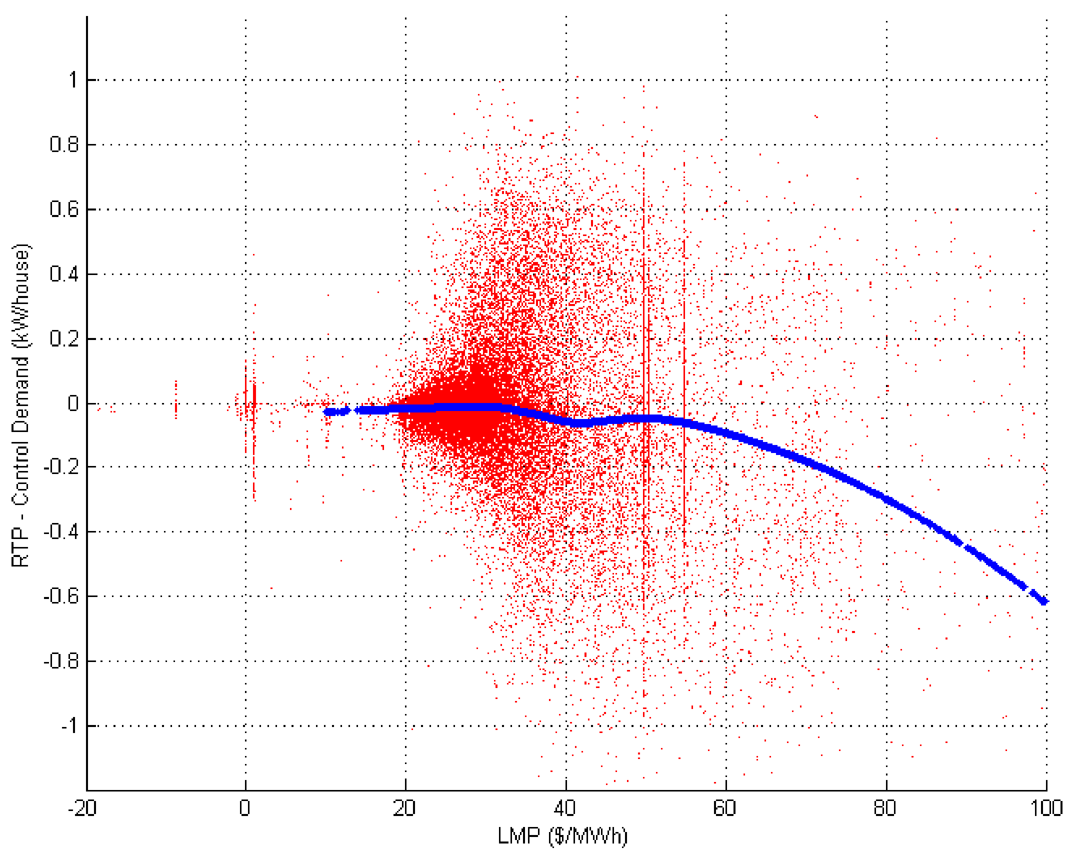

Figure 4.3. Change in Load versus LMP between the Control and RTP_{da} Groups (positive value indicates increased load)

To replicate what was seen in the deployed system, the 96 congestion events were applied to the population of devices. Figure 4.4 shows the relationship between LMP (that is, not the price directly seen by the consumer) and the change in load behavior. The application of the congestion events tended to increase demand during periods of high prices and negate some of the decreases seen during low to medium price periods. This is most likely driven by the congestion event raising the average price and making the LMP look like a better deal in the hours following the events. This may have a significant impact immediately after a congestion event, when units are trying to recover, especially if high LMPs are coincident with the congestion event. The absolute LMP is indifferent to whether the congestion event occurred; however, the thermostat agents will perceive this price as relatively lower following a congestion event, and will increase their consumption following the event. While this is greatly exaggerated by using 96 congestion events, it does suggest that for 24 hours after a high-price or congestion event, households will tend to "over consume" relative to high LMP values, as this price will appear to be a relatively low price. More investigation is warranted to better understand and quantify the simulation results and calibrate simulation models with further information that can be gleaned from the field data.

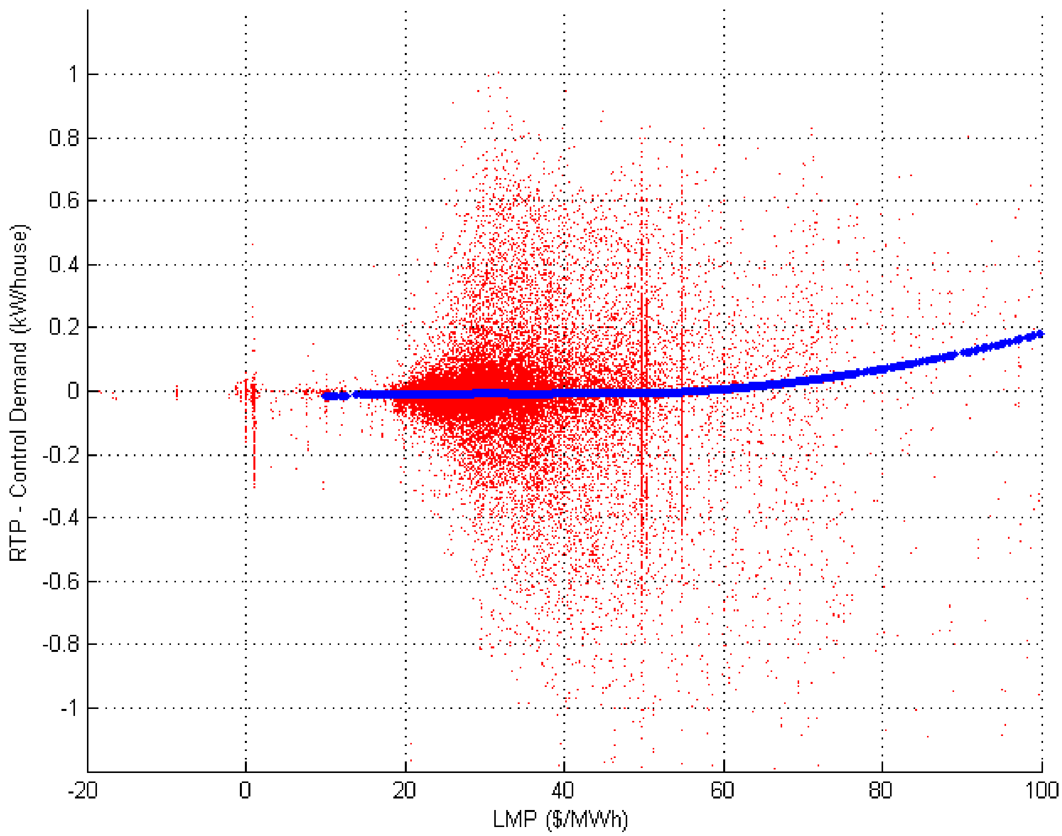

Figure 4.4. Change in Load versus LMP between the Control and RTP_{da} Groups Including Congestion Experiments

To look at this effect, Figure 4.3 and Figure 4.4 were adapted to look at the response to the supply price in terms of standard deviations from the average price. This effectively translates the price into how the thermostat controller views it; the current price is always relative to the average price from the previous 24 hours. So, the price is calculated and displayed as

$$P_{\sigma} = \frac{P_{supply\ bid} - P_{avgerage\ 24\ hours}}{P_{standard\ deviation\ 24\ hours}} \tag{4.1}$$

Figure 4.5 shows the same (non-congested) case as Figure 4.3, but with the price translated into relative prices. The patterns can be difficult to discern. In general, one would expect that as relative price increases, the amount of reduction in demand should increase. However, during any given market period, the response to price is dependent on what happened in previous markets. For example, after an extended period of +2.5 standard deviation prices, in which loads were continuously deferring, a price of +1.5 standard deviations might be a relatively attractive price due to the deferral of operation. So, at any given 5-minute period, the RTP_{da} load may increase with higher prices, but the overall trend should be strongly toward decreased demand during relatively high-price periods.

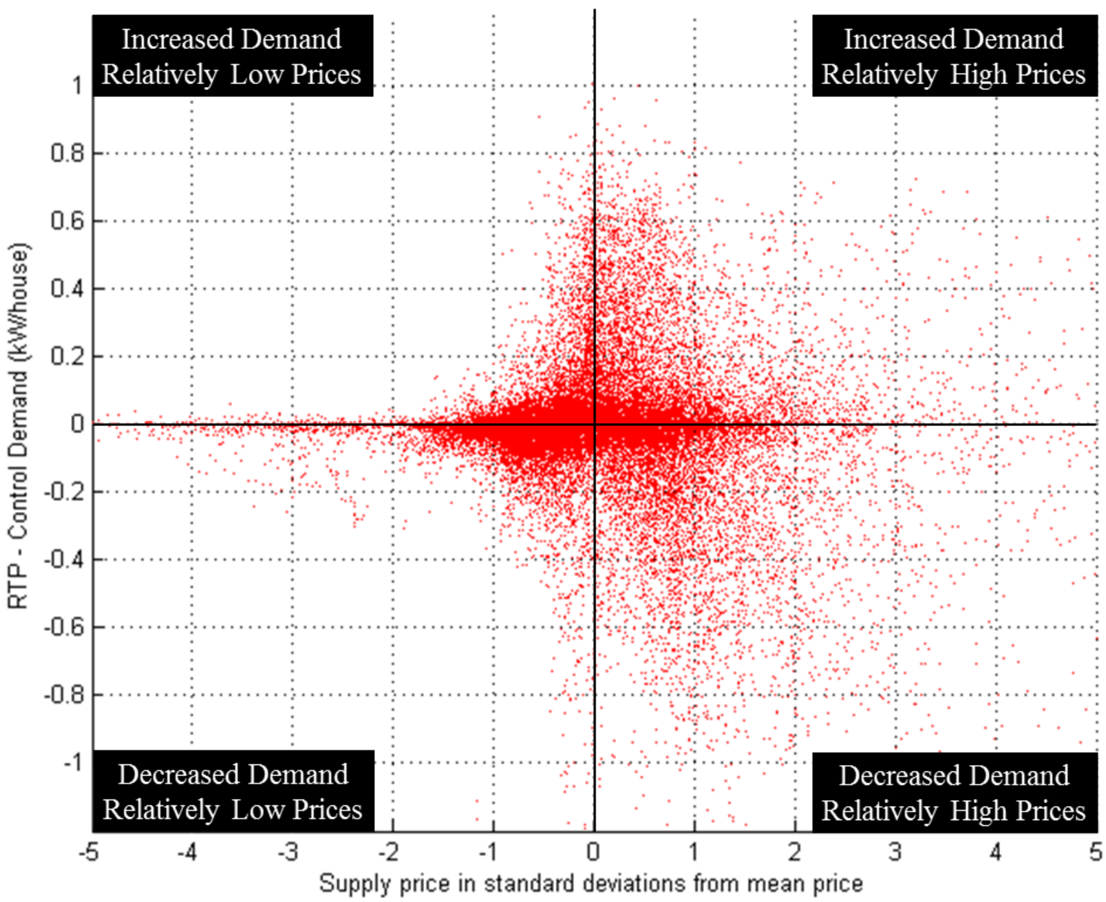

Figure 4.5. Change in Load versus Relative Price between Control and RTP_{da} Groups

Figure 4.6 breaks the same data into three temperature "bins," where the temperature bin represents the current outdoor temperature during that 5-minute market interval. Blue represents temperatures less than 70°F, green between 70 and 80°F, and red over 80°F. The black lines are trend lines determined using the same technique as in Figure 4.1. When presented this way, a number of trends are quickly identifiable. Again, the trends are important and not the individual points. In any given individual market clearing, depending on what occurred in previous market clearings, the RTP_{da} load may increase or decrease relative to the control group no matter the relative price. For example, if the RTP_{da} load has been deferred for 30 minutes, a relative price of +1 may appear quite reasonable (recovery from the deferral may cause an increase in load). However, if the price has been relatively low for the past 30 minutes, a price of +1 is too high for the current resource status (and decreases the load). Looking at the trend lines reduces this noise and determines whether the system in aggregate is behaving as designed, or in other words, decreasing demand during relatively high-price periods.

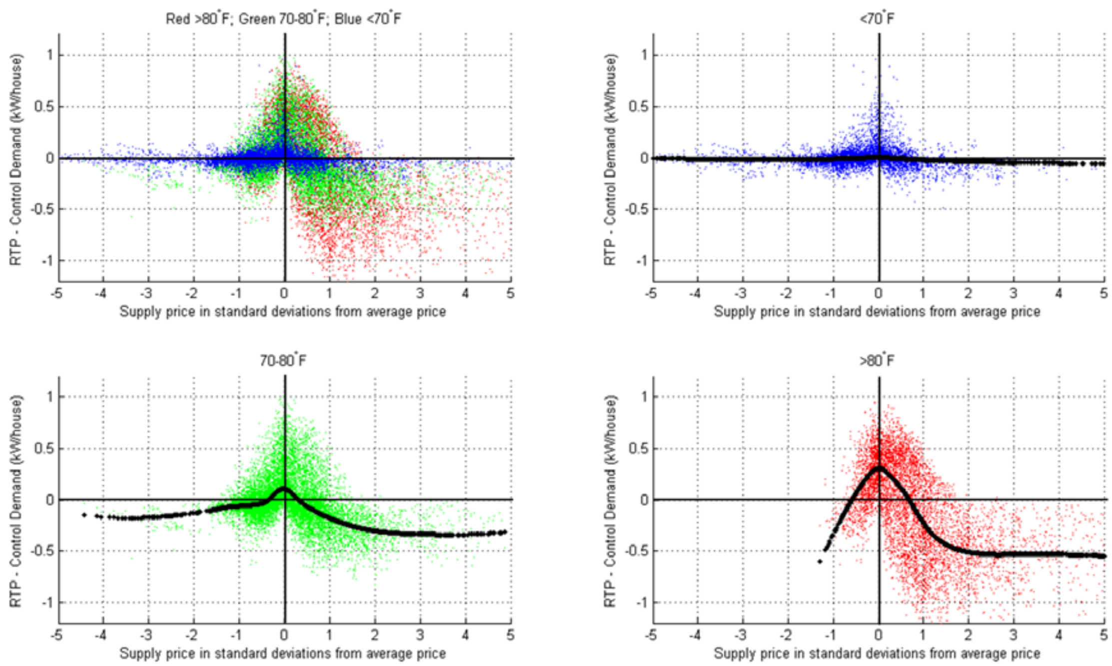

Figure 4.6. Change in Load versus Relative Price between Control and RTP_{da} Groups Broken into Outdoor Air Temperature Bins

In the “less than 70 degrees” bins (the simulations were for summer periods only), when there is minimal air conditioning, the trend is very flat with a very slight trend toward reducing demand as price increases. This is expected, as there is minimal resource during these periods (albeit some resources still available). Also, while there are some individual cases where demand increases or decreases, the overall trend is to minimally decrease load as a function of price. When temperatures are greater than 80 degrees (red), it is clear that as relative price increases, load decreases to a plateau value of approximately 0.5 kW/household. This is as expected (and desired). Of note, however, is that when price is between −0.5 and 0.5 standard deviations, the trend is actually to increase load. Most likely, this is caused by the devices “recovering” during slightly higher prices after very high-price periods that tend to occur more often during hot periods of the day. Data for the 70–80 degree time periods (green) are in between data for the other two graphs. Additionally, as temperature climbs, the relative price also climbs, indicating that a higher temperature day after a cooler temperature day tends to experience higher LMPs. In future applications, this observation could be used to better predict upcoming prices and better tune the controllers to respond to high and low price excursions, which would allow for a rough prediction of demand reduction available during any given market cycle as a function of outdoor air temperature.

In conclusion, this analysis of field and simulated data is but a start to understand and provide insight to the complex interactions at play in the RTP_{da} system. The basic trend observed of reducing load as LMPs rise is directional evidence of the desired behavior with the system; however, many more questions are raised about the strength of this correlation and its behavior under different market, weather, and temporal-related conditions. Further investigation is needed to address these questions and gain greater insight.

4.3 RTP_{da} Load Event Response

Due to the relatively low penetration of RTP_{da} households on each feeder, the impact of each congestion experiment was to engage all of the resources and drive the cleared market price to the price cap. A benefit of these experiments is that they demonstrated the maximum amount of response that could be obtained under various operating conditions. This section investigates the magnitude of the responses from the resources to the congestion events and how well the resources responded over the duration of the event.

4.3.1 Events Response Magnitude

The magnitude of the response is estimated by evaluating the difference between the fractional change in the control group load and the RTP_{da} group load relative to their averages for the 4 hours prior to the beginning of the event. The 4 hour period is chosen as it strikes a balance between prior conditions sample size and variance. The fractional response of response group x relative to control group y is defined as

$$r_t(x,y) = \frac{\sum_{n=1}^{N_x} x_{n,t}}{\sum_{n=1}^{N_y} y_{n,t}} - 1 \qquad (4.2)$$

The mean response for the 4 hours prior to the start of the event at $t = 0$, where the time interval is 5 minutes and t is in hours, is evaluated as

$$\bar{r}(x,y) = \frac{1}{16} \sum_{-4<t\leq 0} r_t(x,y) \qquad (4.3)$$

and is used to normalize all the responses thereafter. The percent response is evaluated relative to this 4-hour mean prior to the event.

The magnitude of the response after the start of the event relative to the response prior to the event is thus

$$R(x,y) = r_t(x,y) - \bar{r}(x,y) \qquad (4.4)$$

This result is shown for 2-hour and 4-hour duration events by the solid lines in Figure 4.7.

4.3.2 Event Response Uncertainty

The uncertainty of the response is estimated by first evaluating the variance of the control group and the RTP_{da} response group for various response types (for example, 2-hour event, 4-hour event, mild day, hot day, off-peak period, on-peak period). All the responses for the selected response types were grouped after being normalized with respect to the conditions prior to the event. The variance of a group load x^N (N is the number of active meters) at the time t is given by

$$v_t(x) = \frac{1}{N^2}\sum_{n=1}^{N}\left(x_{n,t} - \bar{x}_t\right)^2 \qquad (4.5)$$

To compute the uncertainty of the difference between response group x and the control group y, we must compute the covariance

$$c_t(x,y) = \frac{1}{N^2}\sum_{n=1}^{N}\left(x_{n,t} - \bar{x}_t\right)\left(y_{n,t} - \bar{y}_t\right) \qquad (4.6)$$

The 63% confidence interval for the response of the response group x relative to the control group y is

$$\sigma_t = \sqrt{v_t(x) + v_t(y) - 2c_t(x,y)} \qquad (4.7)$$

This result for 2- and 4-hour events is shown by the dotted lines in Figure 4.7.

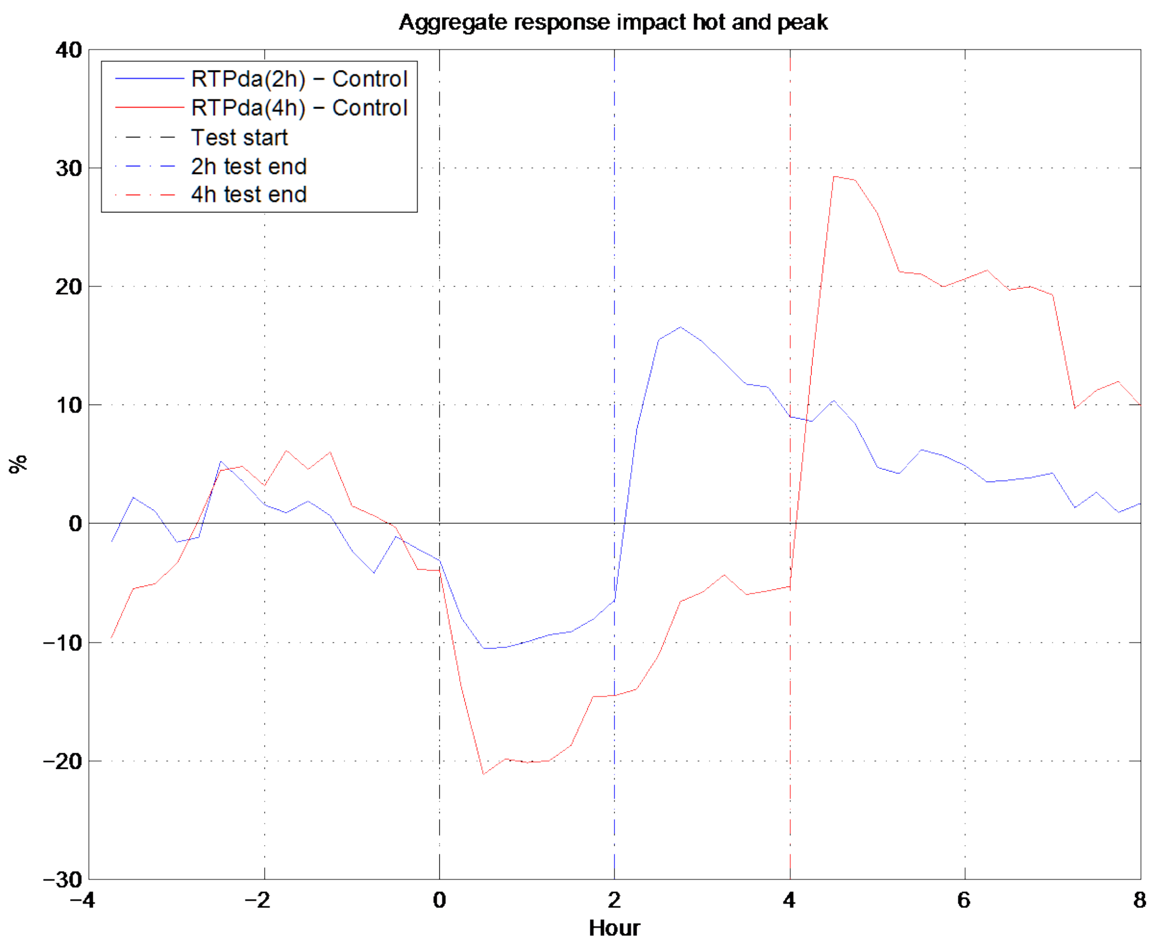

Figure 4.7. Aggregate Response to 2- and 4-hour Congestion Events

The results for the various event and day types are summarized in Table 4.1. The initial response is evaluated 15 minutes after the start the event. The response trend refers to the general trajectory of the response after the initial response to the event. It is evaluated for the duration of the event as a linear fit to the data in percent per hour with the same percent scale as initial response. The initial uncertainty is given for the 63% confidence interval on the same scale as the initial response. The uncertainty trend refers to the general trajectory of the uncertainty after the initial response to the event. It is evaluated in %/h on the same scale as the initial response for the duration of the event.

Table 4.1. Summary of Group Responses by Event and Day Type

Event Type	Initial Response (%)	Response Trend (%/h)	Initial Uncertainty (%)	Uncertainty Trend (%/h)
All events	−8.6	1.0	4.2	−0.2
All 2 h	−6.4	−0.3	5.9	−0.3
Hot+peak 2 h	−10.5	1.3	11.4	−0.6
Mild+peak 2 h	−1.4	−3.1	16.4	−0.8
All 4 h	−17.9	4.1	17.2	−0.6
Hot+peak 4 h	−22.5	4.7	23.1	−0.7
Mild+peak 4 h	−9.3	3.0	63.4	−2.0

The 2-hour events have a much shallower initial response than the 4-hour events (10.5% versus 22.5%). This is mostly driven by the timing of the events. Many of the 4-hour events occurred during high peak periods in the late afternoon on successive hot days, while the 2-hour events occurred over a greater variety of situations in both time of day and daily temperatures. Also, on mild days 4-hour events tended to start later in the day relative to 2-hour events. Thus the 2-hour events tended to start with very few RTP_{da} HVAC resources available and as more unconstrained HVAC resources began operating, the system became more responsive (hence the negative percent response trend). In contrast, 4-hour events started with more RTP_{da} resources already in operation and moved into times of the day where RTP_{da} resources became more constrained in their ability to respond. Therefore, the trend was for decreasing RTP_{da} resource response (positive response percent trend). On peak days, events usually began with more RTP_{da} resources available and the trend was generally toward fewer RTP_{da} resources as time passed. As the resources were depleted, it became harder to distinguish the experiment response from the control group.

Further analysis of the data, such as segmenting the graphs according to time of day, weekday versus weekend, and temperature can help in more fully characterizing the response of the RTP_{da} resources to these events.

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

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