



**NIST Technical Note
NIST TN 2241**

Comparison of Two Approaches for Managing Building Flexibility Using Energy Market Prices

Impact on Cost, Comfort and Power Quality

David Holmberg
Farhad Omar
Thomas Roth

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.TN.2241>

**NIST Technical Note
NIST TN 2241**

Comparison of Two Approaches for Managing Building Flexibility Using Energy Market Prices

Impact on Cost, Comfort and Power Quality

David Holmberg
Farhad Omar

*Building Energy and Environment Division
Engineering Laboratory*

Thomas Roth

Communications Technology Laboratory

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.TN.2241>

October 2022



U.S. Department of Commerce
Gina M. Raimondo, Secretary

National Institute of Standards and Technology
Laurie E. Locascio, NIST Director and Under Secretary of Commerce for Standards and Technology

NIST TN 2241

October 2022

Certain commercial entities, equipment, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

NIST Technical Series Policies

[Copyright, Fair Use, and Licensing Statements](#)

[NIST Technical Series Publication Identifier Syntax](#)

Publication History

Approved by the NIST Editorial Review Board on 2022-10-17

How to Cite this NIST Technical Series Publication

Holmberg D, Omar F, Roth T (2022) Comparison of Two Approaches for Managing Building Flexibility Using Energy Market Prices. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Technical Note (TN) NIST TN 2241. <https://doi.org/10.6028/NIST.TN.2241>

NIST Author ORCID iDs

David Holmberg: 0000-0002-1738-8190

Farhad Omar: 0000-0002-0618-0060

Thomas Roth: 0000-0002-9986-7784

Abstract

Transactive energy may provide a way to support integration of customer distributed energy resources (DER) into the electric grid and to help manage DER flexibility to support distribution system and bulk grid needs. Use of DER flexibility includes peak shifting in response to forward market prices, real-time regulation in response to 5-minute prices, and, potentially, voltage regulation in response to distribution nodal prices. The work presented here compares two approaches for managing heat pump cooling to provide grid services. The first approach implements a rule-based real-time controller that manages power consumption based on 5-minute real-time prices. The second controller uses a model-based approach to optimize load based on day-ahead hourly prices. These approaches are evaluated on two reference grids using GridLAB-D with the same weather and price information. The real-time controller reacts to a volatile price signal to reduce homeowner energy use and energy costs. The model-based controller minimizes energy use by pre-cooling before the price peak and allowing indoor temperature to drift up during peak price times. The model-based controller strategy results in cost savings when temperatures are allowed to move farther from the desired setpoint. Because this controller does not respond to real-time price volatility, there are no cost savings during sub-hourly real-time price spikes.

Simulation results show both controllers adjusting temperature setpoints in response to price movement, reducing customer cost while staying within pre-defined comfort boundaries. These temperature adjustments resulted in significant power flow volatility at the substation. The average power flow change for the real-time controller, time step to time step, was approximately 10 % of the peak load. Some observed power flow changes were equivalent to simultaneously turning off or on all price-responsive loads on the grid. This power flow volatility resulted in voltage volatility at customer meters and greatly increased voltage regulator and capacitor bank actions on one of the tested grids. Analysis suggests that power flow volatility may be expected when many devices are price-responsive and controllers have the goal of reducing cost without considering voltage. Care must be taken to reduce the volatility of the price signal and boost hosting capacity to support synchronized price response by a large percentage of load.

Keywords

co-simulation; distribution grid; dynamic pricing; energy flexibility; load forecasting; grid-edge systems; GridLAB-D; heat pump controller; hosting capacity; power quality; price-responsive controller; real-time market; real-time price; transactive energy; volatility; voltage violations

Table of Contents

1. Introduction	1
1.1. Flexibility and Demand Response.....	1
1.2. Transactive Energy Management.....	2
1.3. Impact of DER on Voltage.....	3
2. Experiment.....	4
2.1. Co-simulation Environment.....	4
2.2. Grid Federate.....	4
2.3. Heat Pump Transactive Controllers	6
2.3.1. Real-Time Controller	6
2.3.2. Load Forecasting Tool for NIST Market Interactions.....	7
2.4. Experiment Design.....	9
3. Results.....	12
3.1. Comparing Thermal Comfort.....	12
3.2. Energy Forecasts	15
3.3. Cost Performance	15
3.4. 8500 Grid Power and Voltage Dynamics.....	21
3.4.1. Voltage Violations.....	21
3.4.2. Substation Power Flows	22
3.4.3. Voltage Regulator and Capacitor Bank Action.....	25
3.4.4. Correlation of Voltage and Price Movement	25
3.5. R4-1 Grid Power Quality	30
3.6. Discussion	33
4. Conclusion	36
5. References.....	37

Nomenclature

CC	Comfort coefficient (unitless)
DAP	Day-ahead market price
deadband	Range of allowed temperature variation around T_{set} without cooling ($^{\circ}\text{C}$)
DER	Distributed energy resources
HVAC	Heating, ventilation, and air-conditioning
k	PNNL real-time controller comfort parameter (unitless)
LA	Learning Algorithm
LFA	Load Forecasting Algorithm
LFC	NIST Load Forecasting Controller
NIST	National Institute of Standards and Technology
<i>offset_limit</i>	Extent that adjusted setpoint can move above or below T_{set} ($^{\circ}\text{C}$)
P_{avg}	Average price over the last 24 hours (\$/kWh)
P_{bid}	Bid price submitted to the double auction market (\$/kWh)
P_{DA}	Day-ahead price of electricity (\$/kWh)
RTC	PNNL real-time controller
RTP	Real-time market price
RTPavg	Hourly-averaged RTP
T_{csp}	Adjusted cooling setpoint temperature ($^{\circ}\text{C}$)
TE	Transactive Energy
T_{max}	Maximum allowed customer setpoint temperature ($T_{csp} + offset_limit$) ($^{\circ}\text{C}$)
T_{set}	Base setpoint (customer desired) temperature ($^{\circ}\text{C}$)
UCEF	Universal Cyber-Physical Systems Environment for Federation
λ	LFC comfort parameter (unitless)
σ	Price standard deviation (\$/kWh)

1. Introduction

The electric grid is undergoing significant changes due to the proliferation of distributed energy resources (DER) and accompanying load growth. DER include photovoltaics (PV), electric vehicles (EV), batteries, and loads capable of providing demand response. The increasing load and the challenges of variable generation can strain existing distribution grids. Successfully adapting to this changing environment requires careful research and planning to avoid overburdening existing infrastructure while enabling the integration of new DER technologies and advanced automation and controls.

Buildings consume 74 % of total electricity sold in the U.S. [1], but in the new electric grid paradigm, buildings are not solely electricity consumers. They also contain storage (both thermal and electrical) and generation that can offer flexibility for grid services such as voltage regulation and frequency response. Heating, ventilation, and air conditioning (HVAC), and hot water heating are significant sources of electricity consumption in buildings, while battery charging for vehicles and stationary applications is growing. Intelligent controls can manage these loads to provide flexibility to the grid and support the integration of intermittent renewable generation.

Transactive energy (TE) may offer a solution to address some of these challenges. TE can support energy conservation, efficiency, and integration of renewable energy at the distribution levels. This report examines the cost, comfort, and grid power quality impacts of two different heat pump control strategies. The strategies use different price signals, but both have a goal of reducing cost while maintain comfort within acceptable limits.

1.1. Flexibility and Demand Response

Demand flexibility is the capability of DER, including generation, storage, and loads, to adjust a building's demand profile across different timescales [2, 3]. It may be viewed as the amount of energy that a customer facility/DER can provide to the grid in response to grid needs. Grid needs are communicated by demand response (DR) event signals, market signals, or automatic sensing (e.g., voltage and frequency). In response to providing grid services, customers receive payments or other incentives that are provided by a market, DR program, or regulatory requirement.

Two common approaches used to access building energy resource flexibility are event-based DR and dynamic price tariffs [3]. DR programs have been implemented widely across the U.S. and typically provide an incentive payment to a customer for allowing a utility or an aggregator to adjust the customer's air conditioner setpoint temperature or alter water heater operation in response to grid conditions and electricity market prices. According to the 2020 Federal Energy Regulatory Commission (FERC) assessment [4], there are 32 GW of DR available as a resource for peak load shedding. This is approximately 5 % of the U.S. peak demand.

Despite this success in implementing DR programs, event-based DR has drawbacks: (1) it does not provide a regular signal that would continuously access the value of the building flexibility, and (2) customers are paid for the power that they do not use relative to an

artificial calculated baseline¹ rather than only paying a given market price for the power that they do use. The importance of the first drawback is that event-based DR relegates customer generation and storage resources to reserves or emergency status rather than making them full-fledged 24/7 assets for grid stability. The importance of the second drawback is that the DR event approach is sub-optimal. It exists because of earlier lack of interval meter data and the related operations of the wholesale markets. Utilities forecast load and bid that into the market, but retail customers do not see market prices and cannot adjust energy consumption based on those prices.

In contrast, dynamic price tariffs have been shown to be effective for incentivizing both manual and automated responses [5]. The most simple and common tariff is the time-of-use (TOU) tariff with fixed-price levels on a fixed schedule. A more dynamic tariff is the real-time price (RTP) tariff, where customers are charged for their energy use based on the wholesale real-time market prices plus a distribution fee. While TOU tariffs are common in the U.S., they are not typically mandatory; less than 10 % of Americans are enrolled [5]. The number of RTP tariffs is smaller but there is a growing number in use [6, 7]. More complex proposed dynamic price approaches might incorporate distribution system modeling to find the equivalent of a distribution nodal price (akin to the transmission system locational marginal price). These nodal prices could vary across the distribution grid to address localized conditions, potentially even addressing both real and reactive power for voltage control [9, 10].

The research presented here investigates the use of dynamic price tariffs for transactive energy management.

1.2. Transactive Energy Management

Transactive energy uses market signals (cleared prices or market tenders) to engage customer DER as market participants. Load controllers manage energy use in response to prices.

The experiments reported here compare two transactive energy management approaches that work in fundamentally different ways with the intent of applying building flexibility, using home thermal properties and customer preference, to reduce feeder load on the bulk grid when the grid is stressed. The first approach responds to a RTP signal and enables fast response to real-time grid needs. The second approach responds to a day-ahead price signal to enable effective load planning. These approaches may be combined to provide effective use of building flexibility to provide grid frequency support.

The first approach, adopted from Pacific Northwest National Laboratory (PNNL) and introduced in Section 2.3.1, responds to fluctuating 5-minute real-time wholesale market prices. The second approach, implemented in a controller developed at the National Institute of Standard and Technology (NIST) and covered in Section 2.3.2, responds to next-day wholesale market prices and plans energy use for the coming day.

¹ A counterfactual baseline depends on averaged historical energy use and assumption of infrequent events in order to estimate what power a customer would have used on a given day [40], rather than allowing the customer to determine the value of energy given all the real-time information at hand, including schedules and other private customer information.

The day-ahead wholesale energy market price (DAP) provides an hourly price of power for tomorrow. The wholesale RTP represents the marginal price of generation to the bulk system at different transmission system nodes and is more volatile than the DAP. Typically, 95 % of energy transactions in the wholesale energy markets are scheduled in the day-ahead market, with the rest scheduled in real-time [10].²

While wholesale RTP provides a signal for incentivizing real-time demand adjustment to support the electric grid and can save money for customers who react quickly, it also has some weaknesses. A real-time signal is not useful for planning purposes, except by reference to its historical behavior. The real-time market is responding to minute-to-minute imbalances in supply and demand and therefore the RTP can change in unpredictable ways. In the context of distribution system flexibility, RTP induces DER to adjust power consumption. When price increases, controllers are likely to reduce load and the feeder aggregate power flow will drop in proportion to the price increase³. Wholesale RTP is a marginal price indicating bulk grid supply needs; it is not intended to have a connection to local voltage levels. However, aggregate response of loads to RTP can induce local voltage changes.

DER management based on next-day hourly prices can enable optimization of load across the next day, balancing cost and comfort. Unlike the response to RTP, it does not provide cost savings via cutting load temporarily during short-term price spikes. Nonetheless, next-day prices are more stable, and energy purchased day-ahead may benefit resource planning both for the building owner and for the grid operator. One approach that has been tested in a utility pilot utilizes DAP to provide a forward planning signal which is combined with RTP for real-time adjustments [12, 13]. This essentially implements a day-ahead retail market to bring customer demand flexibility into the wholesale market cycle, resulting in wholesale clearing prices that consider demand flexibility.

1.3. Impact of DER on Voltage

A number of research reports in recent years have focused on DER-related voltage challenges [14–17], specifically tied to PV systems. Inverters may trip offline in response to a grid voltage dip, which can have the undesired effect of losing generation when the grid needs it. PV generation will produce power when the sun shines and may produce more power at solar peak than the local grid requires, which may result in backflow of power to the substation as well as elevated voltage levels. After the sun sets, heavy demand and/or localized large loads (e.g., from EV charging) may result in under-voltage conditions.

In response to these issues, the IEEE 1547 standard has been revised to enable smart inverters to ride through voltage dips [18] and to use reactive power to correct overvoltage due to PV over-production [17, 19]. Demand response may also be used to increase consumption to reduce voltage. And in the case of low voltage at peak loads, demand response may be used to reduce load and raise voltage [16]. In this study, the high penetration of PV systems does result in over-voltage conditions at certain times and

² Not all energy purchased by wholesale customers is procured through the short-term energy markets; a large percentage may be purchased via long-term bilateral contracts and many utilities own a significant percentage of their own generation resources.

³ Controller and device heterogeneity will help to minimize this effect, but if every controller has as a goal to reduce cost with reference to the same volatile price signal, then power flow volatility will result.

locations, while peak load produces undervoltage conditions in a few locations on one of the distribution feeders studied.

This research demonstrates the impact price volatility can have on voltage stability due to rapid shifts in power flows when heat pump controllers react to a volatile RTP signal in an unconstrained manner. The algorithms described in Section 2.3.1 and Section 2.3.2 were neither constrained by the grid conditions nor optimized to consider voltage stability. This study highlights the impact of controllers sensitive only to price that could lead to creating an undesired emergent behavior on voltage stability. This effect has anecdotal support but seems unreported in the literature.

2. Experiment

2.1. Co-simulation Environment

In this study, we used the U.S. National Institute of Standards and Technology (NIST) co-simulation tool called the Universal Cyber-Physical Systems Environment for Federation (UCEF) [20]. The UCEF co-simulation environment supports integration of different simulation tools that are written in various software languages using the IEEE 1516-2010 High Level Architecture (HLA) [21]. In HLA, an individual simulator is called a federate and combination of simulators forms a federation.

A transactive energy co-simulation requires several key simulation components: the electric grid, the loads and generators connected to the grid, the local and supervisory controllers that control the loads and generators, and the market that coordinates supply and demand [22]. In UCEF, each of these components could be modeled as separate simulation federates or a combination of components as one single federate. GridLAB-D combines the loads and generators together with the grid itself for grid simulation.

In this research, key simulation components consist of the GridLAB-D grid solver working with two grid models, along with two different heat pump controllers operating based on market prices.

2.2. Grid Federate

In the co-simulation experiments, two separate grid models were used. The first model is a larger feeder based on the IEEE 8500 reference grid (hereafter called “8500 grid”), and the second is a smaller grid feeder model, the R4-12.47-1 reference grid (hereafter “R4-1 grid”).

The 8500 grid is described in the NIST Transactive Energy Challenge Phase II report [23] and shown in Fig. 1. This residential feeder model is populated with 1977 single family homes, each with controllable air conditioning systems (air-source heat pumps) and uncontrolled plug loads. For these homes, 90 % have PV systems and 50 % have electric resistance hot water heaters. The 8500 grid was modeled in GridLAB-D [14]. The houses have different base temperature setpoints that are unchanged throughout the day and uniformly distributed over the range $21.1\text{ }^{\circ}\text{C} - 26.7\text{ }^{\circ}\text{C}$ ($70\text{ }^{\circ}\text{F} - 80\text{ }^{\circ}\text{F}$), with different randomized temperature deadbands (uniformly distributed over the range $1.11\text{ }^{\circ}\text{C} - 1.67\text{ }^{\circ}\text{C}$ ($2.0\text{ }^{\circ}\text{F} - 3.0\text{ }^{\circ}\text{F}$)) and offset limits ($1.67\text{ }^{\circ}\text{C} - 2.78\text{ }^{\circ}\text{C}$ ($3.0\text{ }^{\circ}\text{F} - 5.0\text{ }^{\circ}\text{F}$)). The offset limit represents the maximum temperature deviation from the base setpoint that a homeowner

considers acceptable in response to price variation. Houses vary in their floor and window area, wall thermal capacity and insulation. Uncontrolled loads (hot water heater, appliances) follow typical residential usage with randomized sizes and schedules across houses.

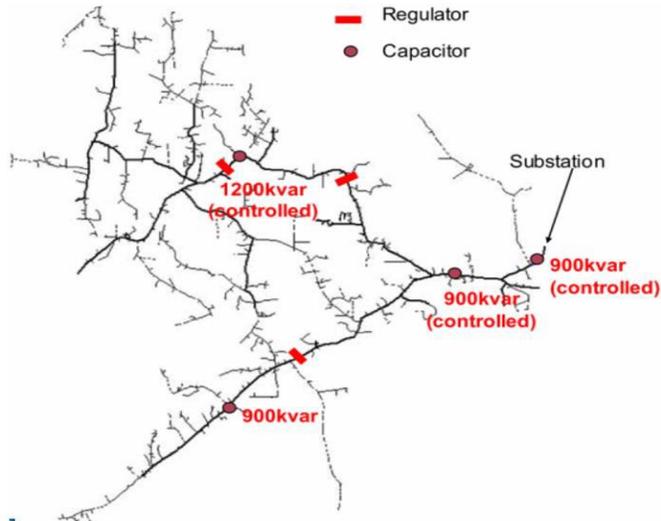


Fig. 1 IEEE 8500 reference grid [25] schematic.

The R4-1 grid, Fig. 2, contains 523 residential homes, both single family as well as several multi-family dwellings. Approximately 30 % of the load on the grid can be attributed to small commercial and industrial facilities spread across the feeder. All homes on the grid have rooftop PVs and price-responsive heat pumps while half the homes have electric resistance hot water heaters. The R4-1 grid peak load is 3 MW compared to the 8500 grid's larger 8 MW peak. In this experiment, setpoint temperatures, deadbands, and offset limits were set equal to the values used for the 8500 grid. The R4-1 grid is a smaller feeder but has the same wire diameters and voltage (12.47 kV) on the main trunk as the 8500 grid. These grid parameters result in a more robust grid with fewer voltage problems compared to the 8500 grid.

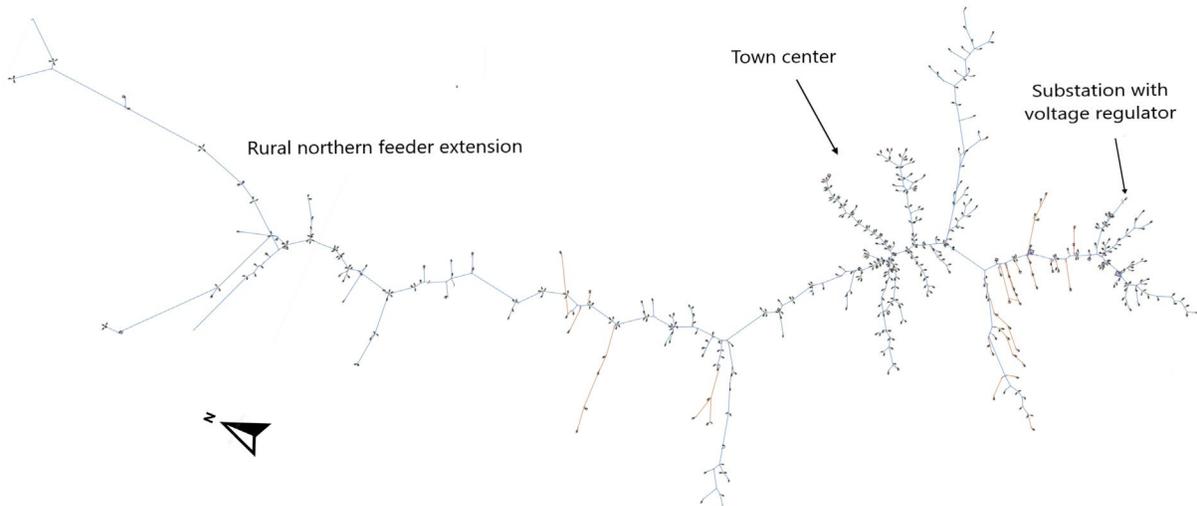


Fig. 2 R4-12.47-1 reference grid [26–28] schematic.

2.3. Heat Pump Transactive Controllers

The experiment made use of two separate controllers: the PNNL Real-Time Controller (RTC), and the NIST Load Forecasting Controller (LFC).

2.3.1. Real-Time Controller

The RTC uses rule-based control that determines house setpoint temperatures based on the real-time price for the next time interval. The RTC used in this experiment has been implemented in a transactive energy pilot program [29]. It adjusts the indoor setpoint temperatures at every 5-min interval, moved up or down in proportion to the changes in RTP.

The RTC federate includes both a thermostat setpoint controller and a double auction market as shown in Fig. 3. The RTC interacts with GridLAB-D, sending new house setpoint temperatures every 5 min and receiving back the current house temperatures. The blue oval inside the GridLAB-D Federate represents a GridLAB-D implementation of the grid topologies described above, including Loads and Generators components, as key actors in a TE co-simulation.

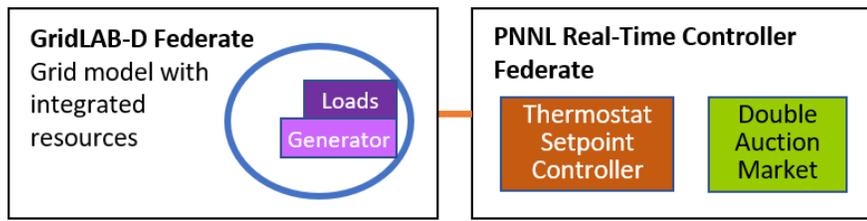


Fig. 3 The PNNL Real-Time Controller with a GridLAB-D distribution grid model.

The RTC market accepts bids representing the amount per kWh that each house is willing to pay for energy required to run the heat pump during the next 5-minute time interval. The operation of the double auction market is documented in Section 1.2.1 of [30] and in [31]. The thermostat setpoint controller submits a bid to the market for each house. By design, the RTC market clearing price is set by the wholesale market real-time price plus distribution costs and congestion limits (that may constrain the cleared load). In the current implementation, distribution congestion (voltage) is not used to adjust the price signal. The clearing price for the next five minutes is returned to the controller which translates this back to setpoint temperatures for each house. The setpoint temperatures are the control actions that GridLAB-D uses to maintain the desired thermal comfort.

The specific algorithm for price-based control, implemented by the RTC, generates bid prices using the following equation (Eq. 1.1 from [30]).

$$P_{bid} = P_{avg} + \frac{k\sigma(T-T_{set})}{2 \cdot offset_limit}, \quad (1)$$

where

- P_{bid} = bid price submitted to the double auction market (\$/kWh)
- P_{avg} = average price over the last 24 hours (\$/kWh)
- σ = standard deviation of the price over the last 24 hours (\$/kWh)

- T = current indoor air temperature (°C)
- T_{set} = desired indoor air temperature (°C)
- k = responsiveness desired by the consumer (unitless)
- $offset_limit$ = extent that adjusted setpoint can move above or below T_{set} (°C).

P_{bid} is the price that a customer is willing to pay to maintain thermal comfort closer to the desired setpoint temperature. In this research, where line congestion and distribution voltage are not accounted for, and where distribution costs are not considered, the clearing price is equal to RTP. The thermostat controller calculates the cooling setpoint temperatures (T_{csp}) for each house as:

$$T_{csp} = T_{set} + \frac{2 \cdot offset_limit}{k \cdot \sigma} (price - P_{avg}) \quad (2)$$

where $price$ is the clearing price from the RTC double auction market, which equals the RTP for the next 5-minute interval. T_{csp} values are sent to the GridLAB-D model every five minutes to update the house temperature control setpoints. Equation 2 shows that setpoint temperature increases linearly with price for each house, and the slope of the increase is proportional to the $offset_limit$ and inversely to comfort parameter, k , and price standard deviation. That is, a customer with higher comfort parameter will see less temperature increase, and likewise there will be less temperature increase when the range of prices (standard deviation) increases. Reference to yesterday's prices serves well when the range of prices is consistent day to day but can lead to over or under reaction by the RTC to prices if there is a shift in level of volatility from one day to another.

2.3.2. Load Forecasting Tool for NIST Market Interactions

The NIST Load Forecasting Controller (LFC) conceptual design, shown in Fig. 4, is based on the NIST TE Market Controller [32]. The LFC is a reduced version of the full TE Market Controller architecture; in this experiment it runs once per day to generate next-day house setpoint temperatures. The Learning Algorithm (LA) learns effective thermal parameters for each house for use with a simple house model. The Load Forecasting Algorithm (LFA) then uses the trained model parameters and next-day prices to optimize heat pump operation across the test day with the goal of minimizing daily energy cost while maintaining temperatures within occupant specified comfort limits.

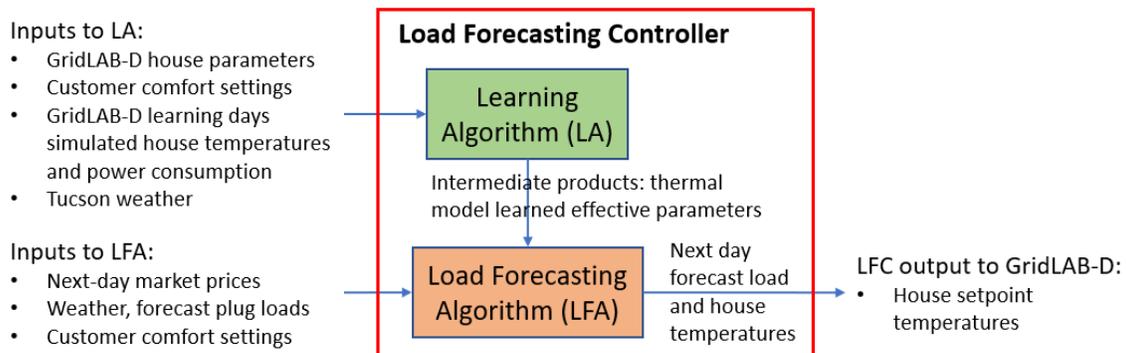


Fig. 4 Schematic representation of Load Forecasting Controller with information inputs, intermediate products, and outputs to GridLAB-D.

Realizing the objective of the LFC requires intelligent learning and control algorithms to manage cost and comfort while adapting to changing weather conditions, the thermal characteristics of homes, and the price of electricity. The intended application of the LFC is to control house resources including flexible loads such as heat pumps, water heaters, and electrical energy storage, based on next-day energy market transactions. For the work presented in this paper, only the heat pump operation is managed by sending adjusted setpoint temperatures to the GridLAB-D simulation engine. The water heater and other potential flexible loads are not adjusted.

The LFC is comprised of two main components, a Learning Algorithm (LA) and a Load Forecasting Algorithm (LFA), as shown in Fig. 4. The LA [33] uses parameter optimization to learn key thermal parameters of a first-order lumped capacitance model from historical measured or simulated data for each house. It optimizes these thermal parameters by minimizing the error between simulated or measured indoor temperature and the output of the lumped capacitance model. The lumped capacitance model forecasts the indoor temperatures from estimates of heat pump energy, solar heat gain, heat gain from plug-loads, and outdoor air temperature. The output of the LA is input to the LFA. GridLAB-D simulations of three preceding days provide training data for the LA. The objective function of the LFA is formulated in such a way that it minimizes cost while maintaining thermal comfort. The mathematical representation of the objective function is given in Eq. 3 (Eq. 1.12 from [32]), describing the multi-objective optimization problem

$$\min_{k \in [2, n]} \lambda \cdot |T_k - T_{SP}| + (1 - \lambda) \cdot (u_{k-1} P_{hpe_{k-1}} \cdot w_{k-1} \cdot x_{k-1}), \quad (3)$$

where:

k represents the discrete simulation time steps [min];

n represents the forecast horizon 1440 [min]. For speed and stability of the optimization solver, the forecast horizon has been divided into 144 bins. Simulation data in each bin represents 10 min of the forecast horizon;

u represents the binary decision variable [dimensionless], and at each simulation time step, it is defined as

$$u \in \begin{cases} 1, & \text{if the heat pump operating} \\ 0, & \text{otherwise;} \end{cases}$$

P_{hpe} represents the electrical power associated with the heat pump operation [W];

w represents heat pump power normalization factor [$^{\circ}\text{C}/\text{W}$]. The normalization factor is defined as $1^{\circ}\text{C}/\max(P_{hpe})$;

λ is a value between 0 and 1, representing the relative dominance between comfort and cost [dimensionless, varies between house models];

T_k represents the predicted indoor temperature at each simulation time step k [$^{\circ}\text{C}$];

T_{sp} represents the setpoint temperature [$^{\circ}\text{C}$]. Setpoint temperatures are obtained from GridLAB-D;

x represents the vector of normalized values of the price of electricity [dimensionless] and is given by

$$x = P_{DA} / (\bar{P}_{DA}), \forall P_{DA} \in [k, n], \text{ where } P_{DA} \text{ is the day-ahead price of electricity, and } \bar{P}_{DA} \text{ is}$$

the average value of day-head price of electricity in [\$/kWh].

Heat pump power consumption (P_{hpe}) and price (x) were normalized to keep the terms of the objective function from dominating the solution. The P_{DA} and \bar{P}_{DA} values for electricity were obtained from two weeks (June 23rd to July 7th, 2017), using the day-ahead price of electricity. Detailed description of the optimization procedure for the LFA is documented in [32].

The LFA is responsible for predicting the hourly energy consumption of a residential house, using market prices, weather data, forecast plug-loads, and customer thermal comfort requirements. Customers' thermal comfort requirements are expressed using the λ parameter that is analogous to the k parameter used in the RTC approach. Weather conditions were obtained from historical data for Tucson Arizona.⁴ A forecast of plug-loads was obtained from simulation of the 8500 grid in GridLAB-D. The LFA utilizes a multi-objective optimization model [34] and constraints to forecast heat pump control actions. The optimization method attempts to find an optimal heat pump schedule that balances cost and comfort, outputting an optimal or an integer feasible solution. The LFA is terminated after it finds an optimal policy or exceeds a pre-defined stopping criterion of 3 min maximum timeout limit⁵.

2.4. Experiment Design

The UCEF co-simulation environment was used to simulate grid operations and house performance for July 6 and July 7, 2017. The experiment used California Independent System Operator (CAISO) DAP and RTP at the Tucson, AZ node. Prices for the test days and two weeks preceding are shown in Fig. 5 (preceding days are shown for comparison purposes). Day-ahead prices (in orange) are seen to generally peak around (0.05 - 0.07) \$/kWh and dip down to (0.01 – 0.02) \$/kWh at night. The RTP has some narrow price spikes and even a negative price spike on June 23. The average price for RTP and DAP is the same (0.032 \$/kWh) across this two-week period. More detail on July 6 and 7 prices are provided below in this section. Fig. 6 provides the outdoor temperature and solar irradiance for the same days. These are key parameters for estimating heat pump load.

⁴ If using forecast weather data, one might expect the next-day heat pump energy forecast to be less accurate, resulting in actual day-of energy use greater or less than forecast. This would result in a required additional purchase or sell back of energy at RTP. For example, if forecast temperatures were lower than actual then additional energy would need to be purchased at RTP to compensate, or vice versa.

⁵ The LFA was implemented in a desktop computer with Intel® Xeon® CPU E5-1630 v3 3.7 GHz processor and 16 GB of RAM, and the data time step was 1 min.

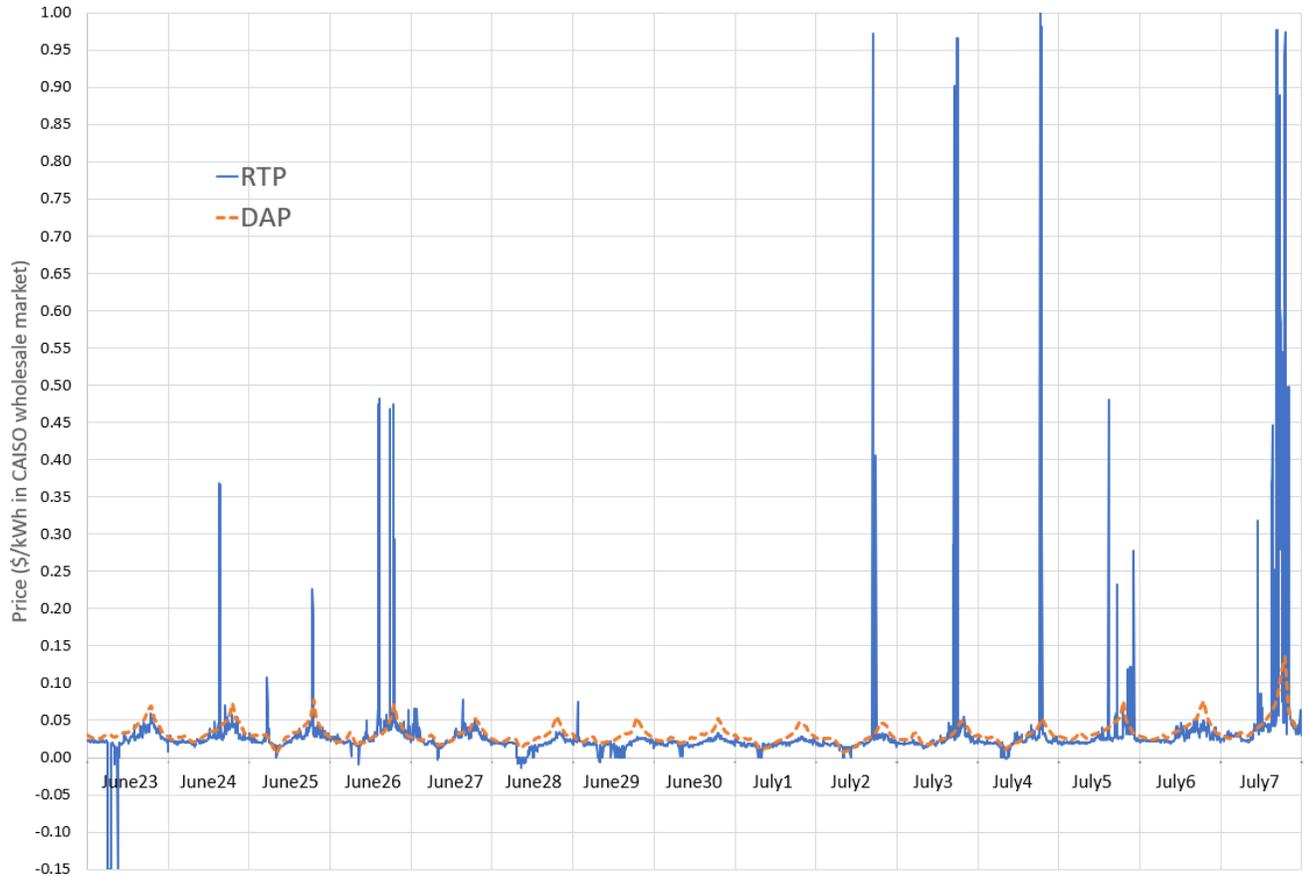


Fig. 5 CAISO RTP and DAP prices for June 23, 2017 to July 7, 2017.

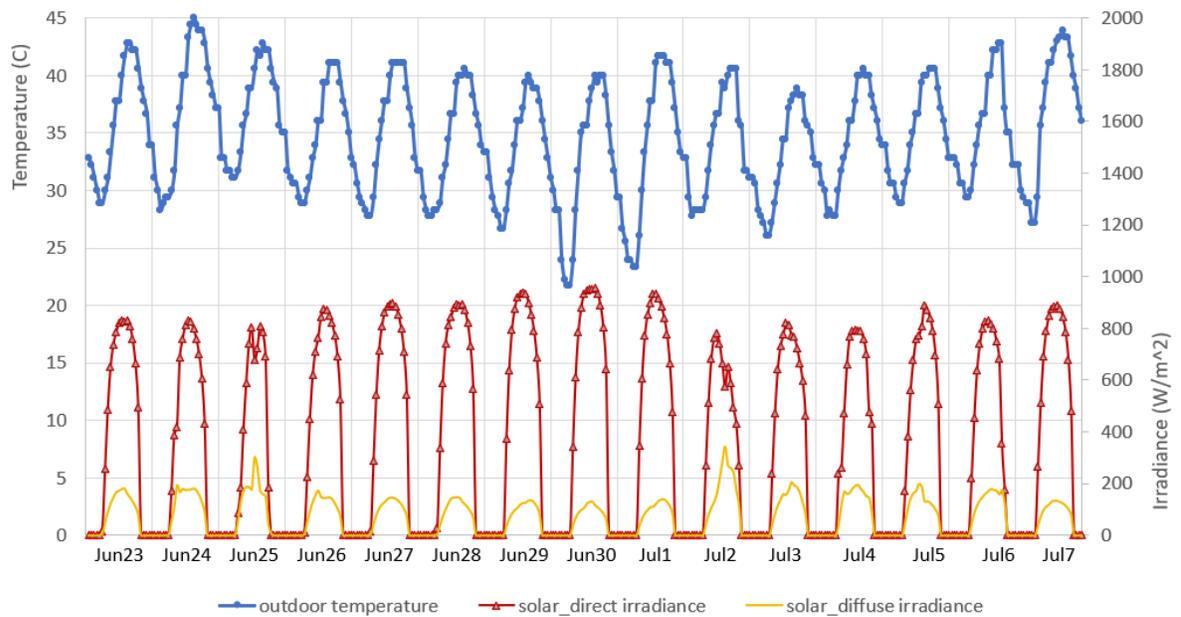


Fig. 6 Measured outdoor temperature and solar irradiance (right axis) for Tucson, Arizona between June 23, 2017, and July 7, 2017.

While the CAISO day-ahead market and real-time market have the same average price over the two weeks of price data used for this study, the standard deviation of the two price signals is very different. Fig. 7 shows the RTP for July 6 and July 7. Fig. 8 shows the DAP along with the hourly-averaged RTP (RTP_{avg}). The mean and standard deviation vary markedly both day-to-day and between markets comparing July 6 and July 7 test days. The July 7 RTP is exceptional in the magnitude and number of price spikes while July 6 has no price spikes. Across the two weeks, RTP and DAP have the same average, but RTP has significant sub-hourly volatility.

If the two controllers make use of different market prices (LFC using DAP, RTC using RTP), it is difficult to compare the performance of the two control approaches on cost. For any given day, the variations in RTP (peak magnitude, peak hour, unusual price spikes such as we see July 7 morning due to random events) make it difficult to compare the underlying LFC planning effectiveness versus the RTC real-time management. In order to address this variability, the decision was made to perform LFC simulations using the hourly-average of the RTP (RTP_{avg}) in place of the DAP. This aligns the peaks and hourly magnitude across the day while removing sub-hourly volatility. Use of the RTP_{avg} enables a direct comparison of cost performance between the two approaches for these two days.

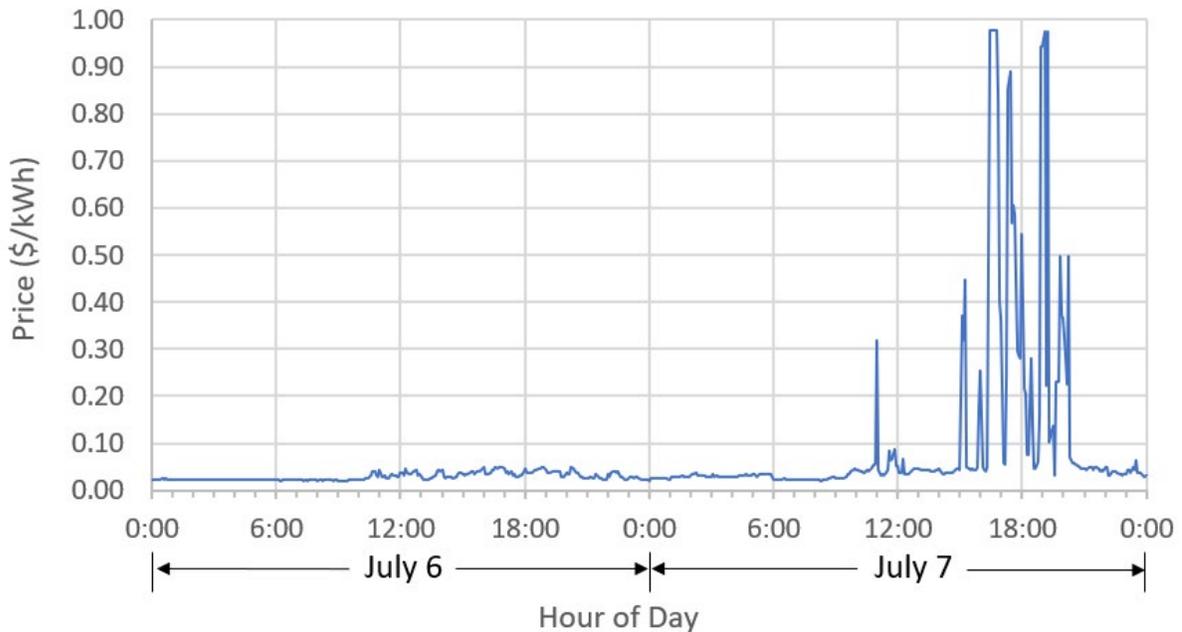


Fig. 7 CAISO Tucson Real-Time Market five-minute locational marginal prices for July 6, 2017, and July 7, 2017.

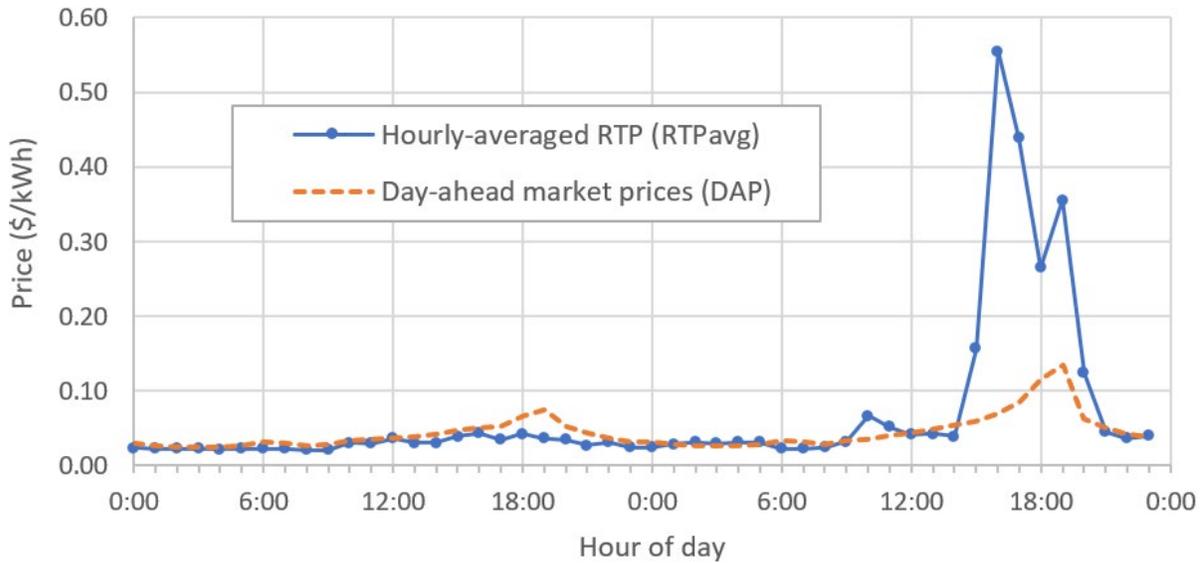


Fig. 8 Tucson day-ahead market hourly prices and hourly-averaged real-time market prices for July 6, 2017, and July 7, 2017.

For this study, GridLAB-D was simulated one day at a time using a one-minute time step. LFC and RTC temperature controllers provided adjusted setpoint temperatures on a 5-min schedule to follow the 5-min real-time market signal. Tucson outdoor temperature and solar irradiance data were provided to the grid simulator at 5-min resolution. GridLAB-D was configured to output a variety of temperature and voltage data for post-processing.

3. Results

In this section we compare the results of the two algorithms for thermal comfort, energy forecasts, cost of using electrical energy, and their impact on voltage fluctuation on the grid.

3.1. Comparing Thermal Comfort

To compare the economic performance of the LFC and RTC algorithms, there is an underlying assumption that the two algorithms weight comfort versus cost in a comparable manner. In this section we evaluate that assumption to better interpret results.

Both algorithms have a slider for comfort, enabling customers to express their desire for cost savings or thermal comfort. The RTC comfort parameter k represents the number of standard deviations that the price must diverge from the average to move T_{csp} to T_{max} (or T_{min} for low prices). The LFC comfort parameter λ is a weighting factor in the cost versus comfort optimization for the LFC. For both approaches, house temperature setpoint is always maintained within the allowed occupant temperature limits, and thus “comfortable.” The key idea is that some customers are willing to pay more money to maintain house temperature closer to their optimal setpoint while others are willing to save money and allow precooling in the morning and higher temperatures during the afternoon when it is hot.

The RTC algorithm’s adjusted setpoint (T_{csp}) will reach T_{max} when the price rises $k*\sigma$ above the average price over the past 24 hours, P_{avg} , where σ is the price standard deviation during

that time. For example, consider $P_{avg} = 0.10$ \$/kWh, with $\sigma = 0.05$ \$/kWh. If $k = 1$ and the price rises to 0.15 \$/kWh, then the house cooling setpoint will reach T_{max} . This approach uses yesterday's prices as a gauge for "high" prices and expected volatility, which may not serve well if there is a significant shift in average price and volatility day to day.

The LFC uses a different approach. It neither considers the previous day's prices, nor uses rule-based control that ties next step temperature to the next step price. Rather, it is looking at a set of hourly prices for the next day and using a house model to explore different temperature profiles throughout the next day to optimize cost while maintaining thermal comfort. The LFC approach results in a profile that includes some precooling, with the amount depending on the thermal storage capacity and time constant of each home, rather than simply lowering T_{csp} whenever the price is low. The LFC, as implemented in this study, uses hourly prices and model steps of ten minutes, not looking at five-minute RTP.

Even though the two approaches work in fundamentally different ways, they allow for temperature deviation from T_{set} as a function of price and comfort parameter. This dependency on price and comfort parameter enables us to perform a basic comparison of the two approaches. For this comparison, a comfort coefficient (CC) was defined to quantify the average normalized temperature deviation from T_{set} for a given house for a given day.

$$CC = \frac{1}{ts_f} \sum_{ts=0}^{ts_f} \frac{|T - T_{set}|}{offset_limit} \quad (4)$$

where *offset_limit* is equal to $T_{max} - T_{set}$, T is the house temperature at each time interval, and ts is the simulation time step. For a single day simulation and 5-min time steps, $ts_f = 288$.

Fig. 9 and Fig. 10 show the comparison of the CC metric as a function of k and λ for July 6 and July 7 test days, respectively. In both plots, each data point represents the CC metric of a house. The values of k associated with each house were preset as part of the IEEE 8500 grid market definition used for the TE Challenge Phase II [13]. The values of λ were set using a simple linear mapping of k to λ , such that k from 0.5 to 3 maps to values of λ from 0 to 1. A value of $\lambda = 0$ indicates a customer most willing to allow the controller to adjust temperature to save money, while $\lambda = 1$ indicates a customer most interested in maintaining temperature close to the setpoint.

As seen in Fig. 9 and Fig. 10, the two approaches do not have identical CC values over the range of k and λ for the two days with different shapes on July 6 versus July 7 due to the very different price curves on those days. Fig. 9 shows that the RTC has a CC value that drops gradually from approximately 0.4 at $k = 0.5$ to the noise floor⁶ of $CC = 0.2$ at $k = 1.5$. The CC curve for LFC has the same general shape, but with a sharper drop at low λ , indicating that LFC will allow more temperature variation at very low λ .

Fig. 10 shows CC values higher on July 7 than seen on July 6 for all values of k . This is due to the strong spikes and much higher prices compared to the previous day (Fig. 7). The higher CC values at large k indicate that the RTC algorithm is raising temperatures for even

⁶ The noise floor here is due to temperature variation within the deadband.

the most comfort-conscious customers due to the price spikes and increased P_{avg} and σ relative to July 6 values of the same.

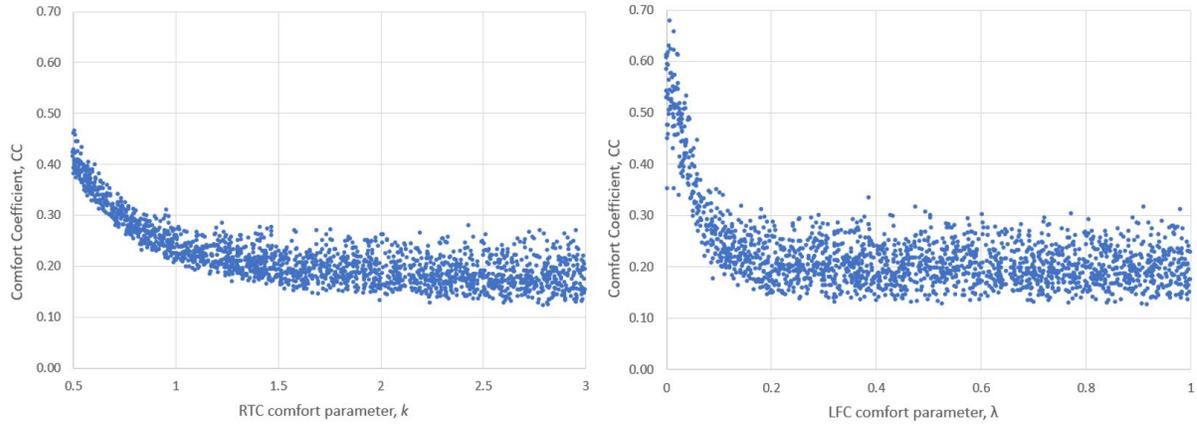


Fig. 9 Comfort Coefficient comparison across the comfort range for RTC (left) and LFC (right) for July 6. Each point represents the CC metric for a house in the 8500 grid.

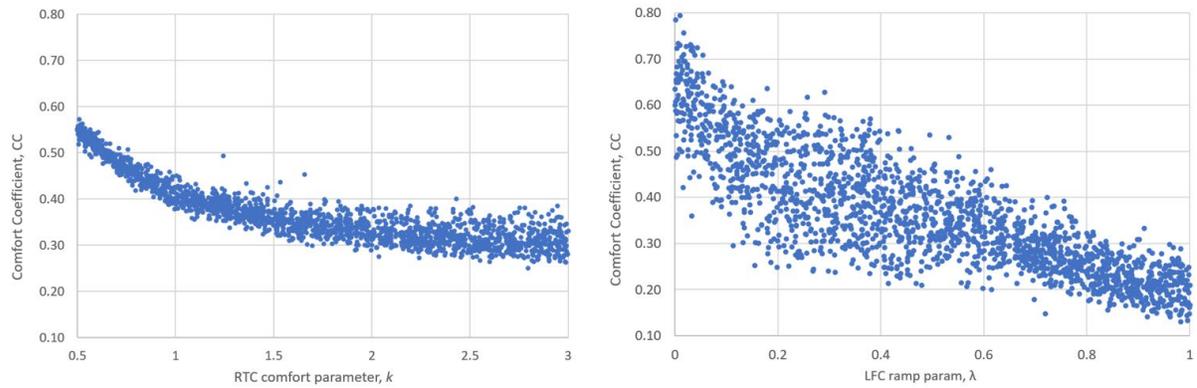


Fig. 10 Comfort Coefficient comparison across the comfort range for RTC (left) and LFC (right) for July 7. Each point represents the CC metric for a house in the 8500 grid.

The July 7 LFC CC curve shows more scatter in CC across the range of λ . This may be due to the LFC adjusting the amount of precooling as a function of house thermal mass and insulation. Note the difference between LFC at $\lambda = 1$ and RTC at $k = 3$. The LFC places all weight on the comfort term and zero weight on the cost term and thus the comfort is at the noise floor ($T_{csp} = T_{set}$). The RTC still allows some raising of T_{set} at $k = 3$ due to the higher prices on July 7 compared to July 6.

The results show that we should expect little difference in cost and comfort between the two algorithms on July 6 except at the lowest values of k and λ . Most homes will see a T_{csp} that remains near T_{set} . In contrast, on July 7, the range of T_{csp} variation increases. Over the ranges of k and λ , the LFC tends to have less T_{csp} variation at high λ compared to RTC, but more at the low end ($\lambda < 0.4$).

3.2. Energy Forecasts

The NIST LFC estimates tomorrow’s total hourly energy consumption by adding the predicted heat pump energy consumption and plug-loads for each house. The estimated hourly total energy consumption can be compared to the simulated energy use of each house in GridLAB-D. Fig. 11 shows LFC’s hourly forecast energy consumption across the day for one house and the resulting GridLAB-D’s simulated energy consumption required to maintain the house temperature at the LFC-prescribed T_{csp} (input to GridLAB-D). The results shown in Fig. 11 are representative of the observed differences between LFC generated energy forecasts and the simulated GridLAB-D energy consumption.

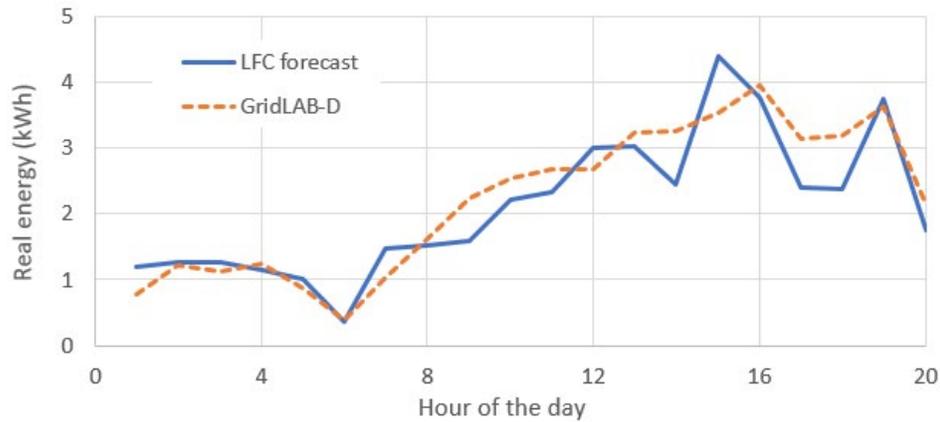


Fig. 11 Comparison of LFC forecast hourly energy consumption versus the GridLAB-D simulated energy consumption for a representative house on July 7.

In comparison, the RTC makes no forecast. It only observes the deviation of the current house temperature from the base setpoint (T_{set}) and provides a bid price above which it is willing to shut off the heat pump (or keep it off), given customer comfort requirements.

3.3. Cost Performance

The LFC model produces a house temperature setpoint profile that should ideally meet customer comfort requirements at lowest cost, if energy is paid for based on the same day-ahead prices. The temperature setpoint profile for each house is provided to GridLAB-D as input, and GridLAB-D returns simulated indoor temperatures and energy consumption for each house. While the RTC approach charges all energy use during each 5-min interval at the RTP for that interval, the cost for the LFC controller is computed per Equation (5). For LFC, day-ahead forecast energy is paid for at the day-ahead price. Any energy consumption difference from the forecast (more or less) is paid for at RTP.

$$C_{LFC} = \frac{E_{hf}}{12} P_{DA} + \left(E - \frac{E_{hf}}{12} \right) P_{RTP} \quad (5)$$

where

- C_{LFC} = cost per 5-min time interval for LFC controller (\$)
- E_{hf} = LFC hourly forecast energy consumption for next day (kWh)
- P_{DA} = next day hourly price (\$/kWh)

E = energy consumption per 5-min time interval (kWh)

P_{RTP} = RTP price for the time interval (\$/kWh)

Fig. 12 presents cumulative energy cost for RTC, LFC, and a no-control baseline (temperatures held constant at T_{set}) for a single house that has a low comfort (more cost-conscious) setting on July 7. As shown in Fig. 12, prior to 10 a.m., the RTC and LFC track the baseline. After 10 a.m., the RTC reacts to the rising prices by allowing T_{csp} to drift upwards (Fig. 13), saving some energy and reducing cost. The LFC performs some pre-cooling that keeps the cost higher than the RTC until the peak pricing period when the cost of the LFC drops below the cost of the RTC.

Note that GridLAB-D had power flow convergence issues that seem to be tied to price fluctuations and resulting power flow fluctuations on the 8500 grid⁷. The 24-hour baseline simulations ran for the full day. LFC hourly price fluctuations induced GridLAB-D solver convergence failure typically around peak load. With RTC and its stronger price fluctuations, GridLAB-D encountered convergence issues earlier, as seen in Fig. 12. Due to these convergence issues, the analyses of the results on cost and power flow volatility are only valid until simulation stopped.

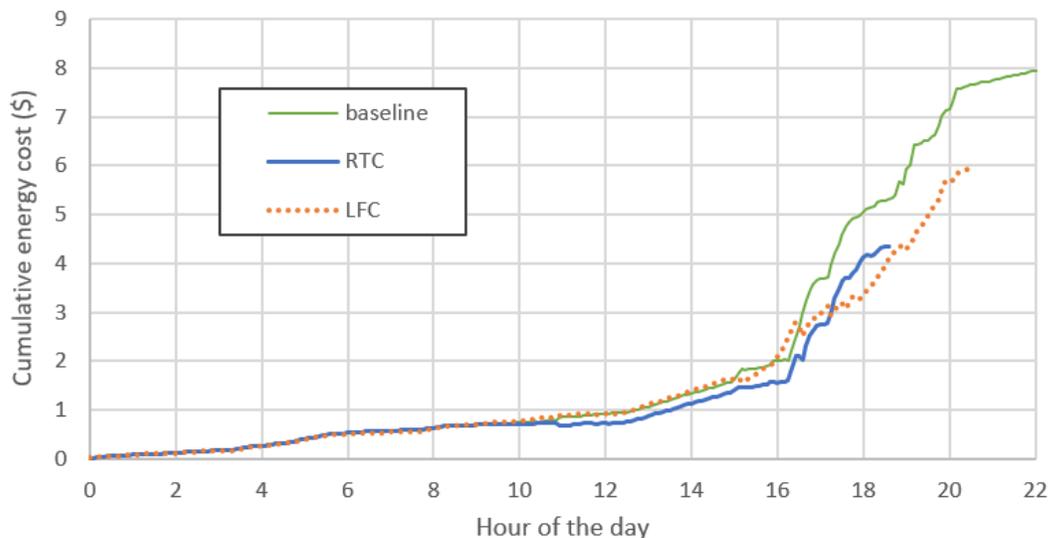


Fig. 12 Cumulative electricity cost for a typical house with low-comfort parameter $\lambda = 0.087$), July 7.

Fig. 13 shows the T_{csp} profiles for the two controllers along with the RTP profile on the secondary axis. The RTC algorithm raises T_{csp} in the morning due to the already high (compared to July 6 average) morning prices. The LFC, on the other hand, is performing pre-cooling (from 8:00 a.m. to 9:30 a.m. and 12:00 p.m. to 3:00 p.m.) in anticipation of even higher afternoon prices. LFC attempts to reduce cost by letting the indoor temperature float during the peak price periods.

⁷ Similar issues with this grid model were encountered in previous publications.

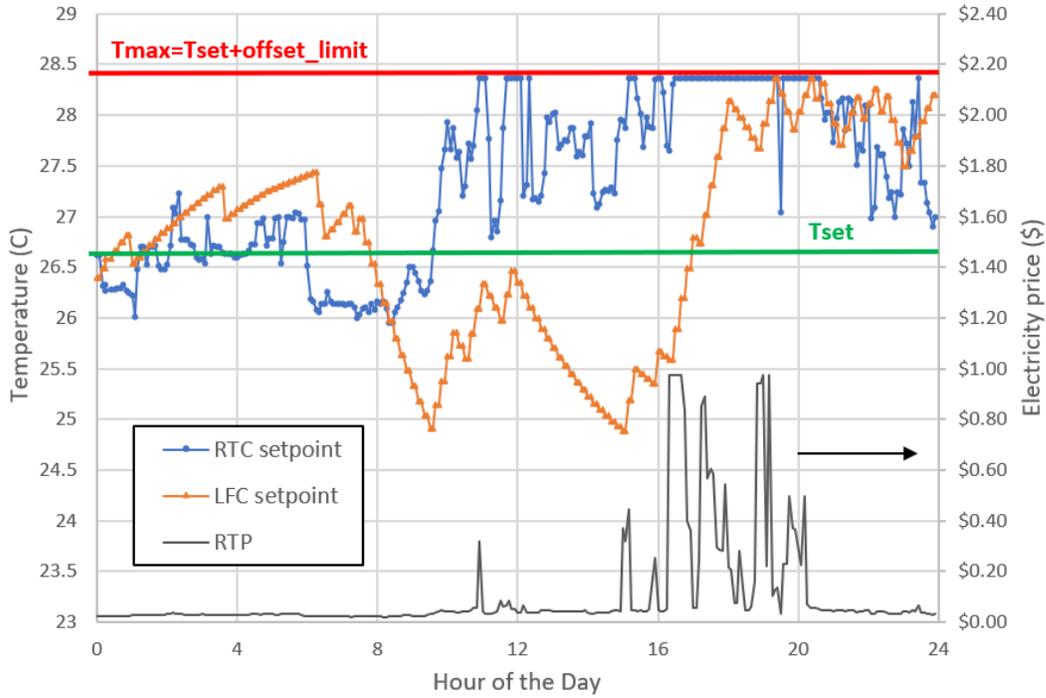


Fig. 13 Cooling setpoint temperatures from RTC and LFC controllers for the same low-comfort house as in Fig. 12, July 7, 8500 grid. The RTP profile is shown in black on secondary vertical axis.

The house indoor temperature bounces around T_{csp} within the deadband, as shown in Fig. 14.

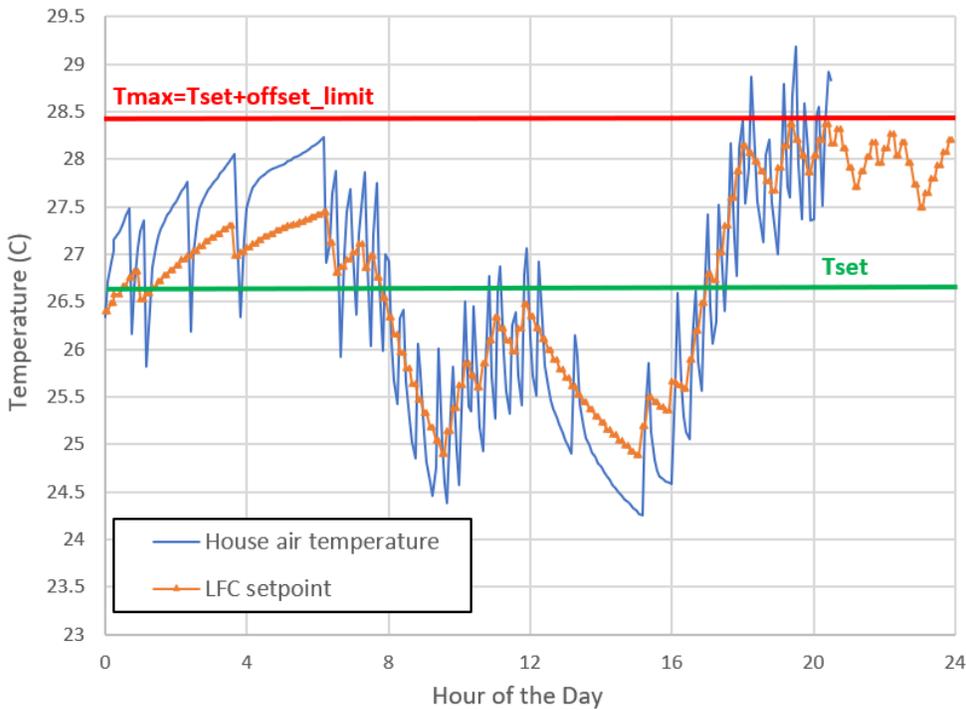


Fig. 14 Comparing house air temperature to LFC T_{csp} for the house seen in Fig. 13.

Fig. 15 compares controllers' actions for a house with a high-comfort parameter. In this case, we expect the controller to maintain the temperature closer to T_{set} , given the customer preference for comfort over cost savings, as shown in Fig. 15 versus Fig. 13. The RTC also reduces the amount of temperature excursion away from T_{set} , although setpoints still hit T_{max} .

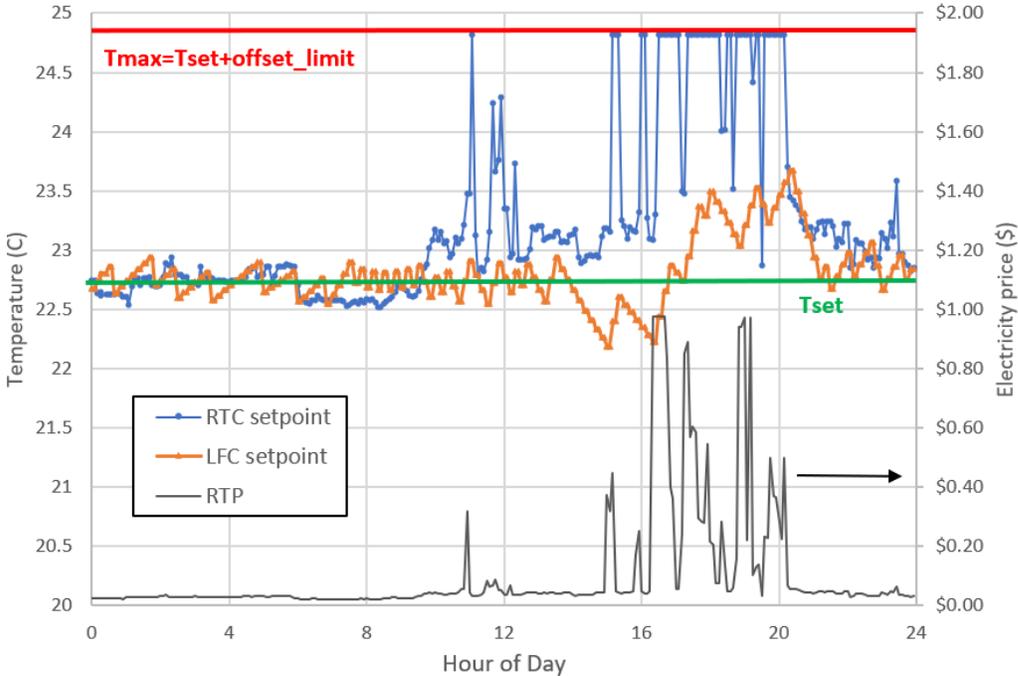


Fig. 15 Cooling setpoint temperatures for RTC and LFC controllers for a house with a high-comfort setting ($\lambda = 0.82$), July 7, 8500 grid.

Considering Fig. 15, the RTC algorithm reacts to the high prices on July 7 (compared to July 6 when peak prices were less than \$0.10) and responds by raising T_{csp} in the morning. The RTC algorithm judges whether the current price is high or low based on prices from the day before and volatility of those prices. On July 7th, RTC reacts strongly to the much higher prices, raising T_{csp} to T_{max} during the price spike in the morning. RTC has no foresight to expect even higher prices in the afternoon. On the other hand, the LFC attempts to balance comfort and cost across the day without reference to July 6 temperatures. T_{csp} is only raised about 1 °C during the peak price period.

Fig. 16 shows the cumulative cost impact for the same house. The LFC cost tracks closer to the baseline, since T_{csp} is kept close to T_{set} . The RTC cost is less due to raising T_{csp} up well above T_{set} as shown in Fig. 15. After 4 p.m. the RTC cost drops close to zero. This is due to solar PV energy production combined with an energy efficient house and the strong RTC response to price spikes. Under these conditions, the PV is generating excess energy that is flowing back to the grid. The excess generation is being reimbursed at the high RTP, resulting in reduced cost to the homeowner. In contrast, the LFC does not react to RTP price spikes; however, it is minimizing energy use generally during the hours with high average RTP cost such that we see some decline in the cost of energy between 16:00 and 18:00 relative to the baseline.

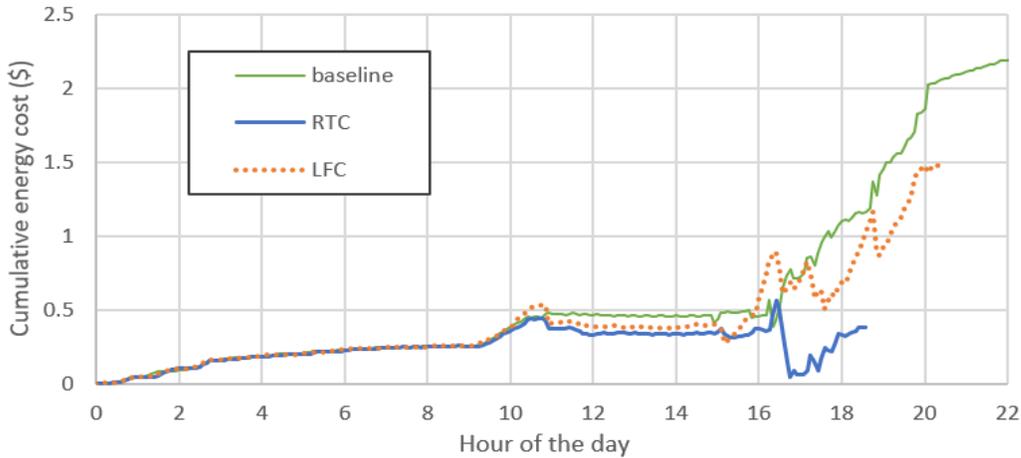


Fig. 16 Cumulative electricity cost for the same high-comfort house shown in Fig. 15. The energy cost drops when PV generation is greater than house consumption and the owner is paid for power flowing out to the grid.

Finally, consider a similar high-comfort house with no PV on the roof as shown in Fig. 17. In this case, the RTC saved money by raising T_{csp} in response to price spikes, but there is no drop in cost due to PV net metering. The RTC algorithm saves money by driving the setpoint temperature higher (Fig. 15) while the LFC keeps the adjusted setpoint closer to the base setpoint and thus the cost also tracks close to the baseline.

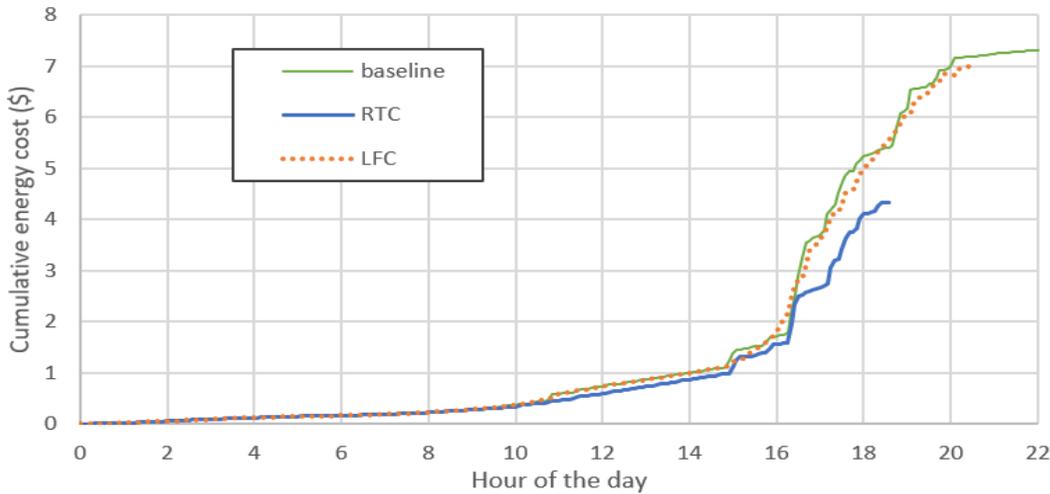


Fig. 17 Cumulative electricity cost, comparing controller performance to baseline for a high-comfort house with no PV generation ($\lambda = 0.88$).

All the cost results presented so far are from July 7. On July 6, the RTP price peak and volatility is lower compared to July 7. For reference, consider the temperature setpoint results (Fig. 18) and cost results (Fig. 19) on July 6 for the same low-comfort house shown in Fig. 12 and Fig. 13. In Fig. 18, the RTC T_{csp} follows the rise and fall of the RTP, rising up toward T_{max} . The LFC allows the T_{csp} to drift up to T_{max} during the highest price periods. In Fig. 19, despite this house having a low-comfort setting ($\lambda = 0.087$), the cost data shows only small deviations from baseline for both the RTC and LFC algorithms. Fig. 19 shows that GridLAB-D still had convergence issues, even with the lower price volatility. However, examination of

the voltage and power flow fluctuations in Section 3.4 show that the July 6 prices still induced synchronized temperature adjustments resulting in strong power flow fluctuations.

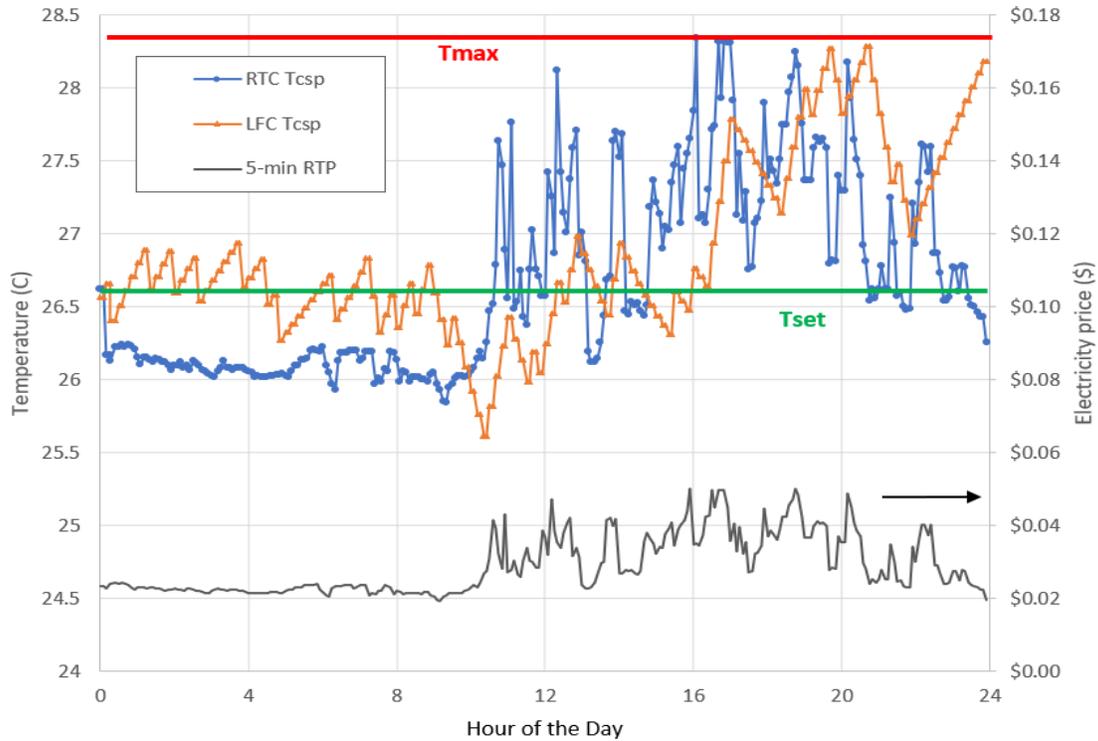


Fig. 18 Cooling setpoint temperatures for RTC and LFC, July 6, 8500 grid. Compare to July 7 data for the same low-comfort house in Fig. 13.

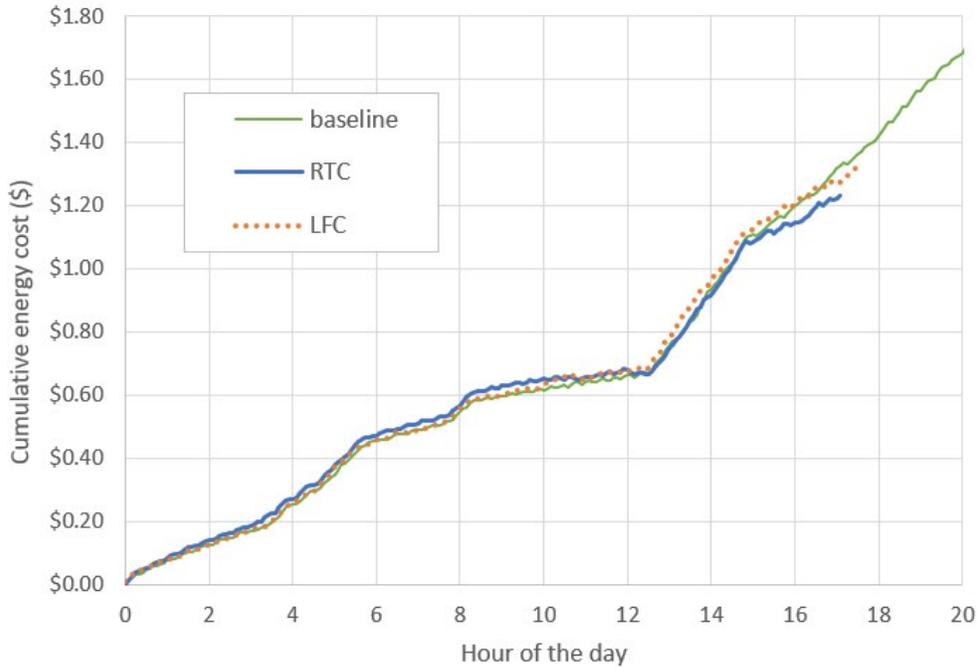


Fig. 19 Cumulative cost of energy for July 6 for low-comfort house shown in Fig. 12.

3.4. 8500 Grid Power and Voltage Dynamics

GridLAB-D simulations provided minutely results for substation power flows, house meter voltages, as well as voltage regulator and capacitor bank actions (mechanical switching). Results suggest that the RTC and LFC controllers can have significantly different impact on the distribution power quality and hardware action.

3.4.1. Voltage Violations

The normal acceptable voltage range for service below 600 V as defined by ANSI C84.1 [22] is nominal voltage +/- 5 %, or a service voltage of 114 V up to 126 V. This is termed Range A. The ANSI standard states that “The occurrence of service voltage variation outside this range should be infrequent.” Therefore, the number of occurrences of voltage outside this range, and duration of time outside this range is a measure of distribution grid power quality and can be used to judge the impact of the control algorithms.

The number of PV systems on the 8500 grid used in this experiment [14] is above what a utility would normally permit, given that overvoltage conditions are present when the sun is shining. The simulation results in Table 1 show the presence of voltage violations. In the case of the baseline simulations, voltage violations can be attributed to over-voltage conditions that are almost entirely limited to the morning when there is significant solar generation while the heat pump load is still low.

In contrast, the over-voltage counts for LFC and RTC occur throughout the day and are primarily tied to price changes. The LFC and RTC controllers adjust T_{csp} in response to prices, leading to synchronized power flow changes and voltage fluctuations. In the case of RTC, T_{csp} is changed in proportion to RTP every 5 min. In the case of LFC, each house is separately optimized based on hourly day-ahead RTP_{avg} but temperatures move up and down at each LFC optimization time step (10 min), impacting heat pump operation. For reference, compare the T_{csp} variations for baseline (green T_{set} line), LFC and RTC in Fig. 13 and Fig. 15.

Table 1 Voltage violations outside Range A levels and average time step to time step power flow changes for the 8500 grid.

	July 6 Voltage violations ^a		July 7 Voltage violations ^a		July 6 average change in power flow ^b (kW)	July 7 average change in power flow ^b (kW)
	Count (per h)	Duration (min/h)	Count (per h)	Duration (min/h)		
baseline	0.29	0.80	0.22	0.65	101	101
RTC	1.80	7.2	2.16	9.2	745	870
LFC-RTP _{avg}	1.11	3.3	1.16	3.7	423	447

a. average per house

b. timestep-to-timestep change, absolute value

The right-hand columns of Table 1 give a measure of the power flow volatility, showing the average change in power flow from one time step to the next. RTC produces stronger fluctuations than LFC and LFC stronger than baseline. This power flow volatility leads to

voltage volatility, with voltage violation counts and durations as shown in the left columns of Table 1. However, despite the stronger price fluctuations on July 7 (Fig. 7) the voltage violations and power flow fluctuations on July 6 are only slightly lower than what is observed on July 7. The reason for this can be tied to the operation of the two controllers producing similarly volatile T_{csp} profiles on July 6th versus July 7th (Fig. 18 and Fig. 13). The relationship between temperature and voltage is studied more in section 3.4.4.

3.4.2. Substation Power Flows

Another way to look at voltage volatility is to examine power flows through the substation for the 8500 grid. Fig. 20 shows the baseline substation power flow, a classic duck curve with decrease in load in the morning followed by an increase in load up to when the sun sets at 8 p.m. In addition, power fluctuations are minimal (refer back to Table 1). In contrast, the fluctuations in power flow at the substation are greater for the LFC (Fig. 21) and even stronger for the RTC (Fig. 22). There are also significant negative power excursions in response to the RTC indicating power flow back to the transmission grid. The substation power flow and voltage fluctuations are not incorporated as constraints in the LFC and RTC algorithms. However, this analysis shows the indirect impact of grid condition-unaware controllers on grid stability.

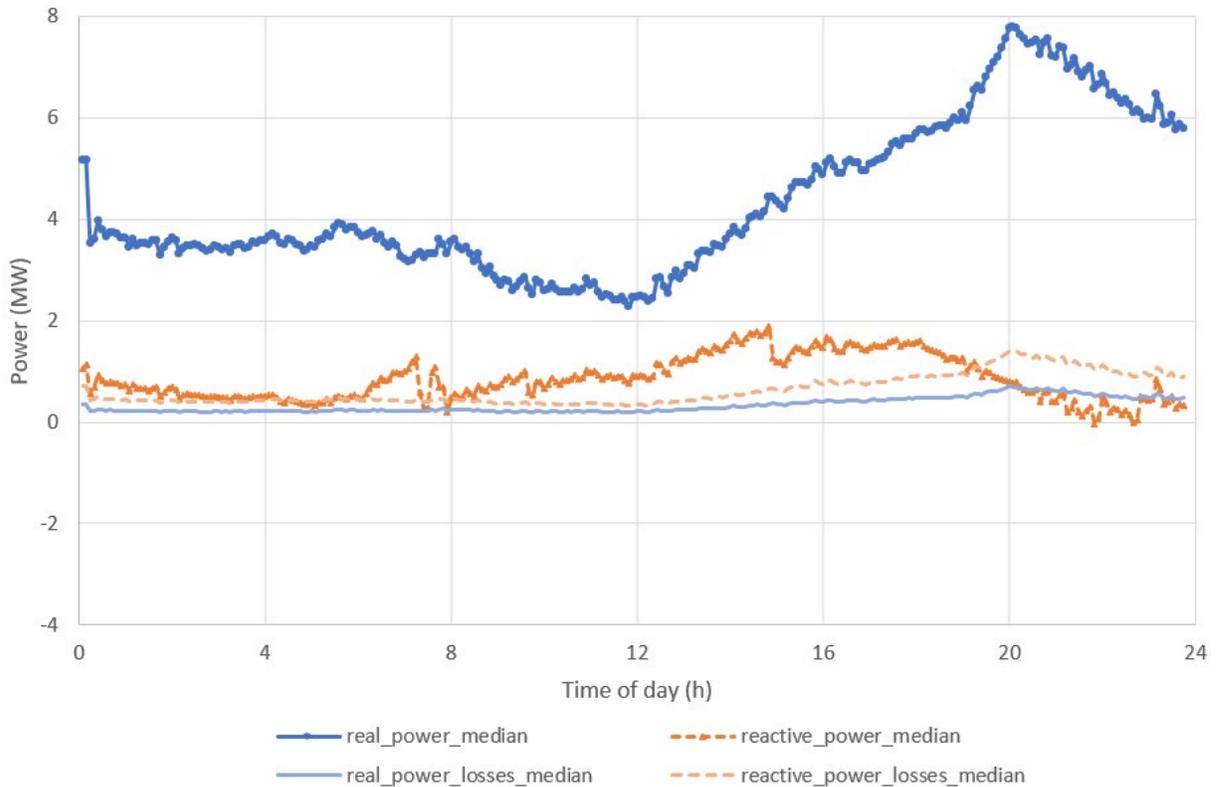


Fig. 20 Substation power flows for the baseline case, July 7.

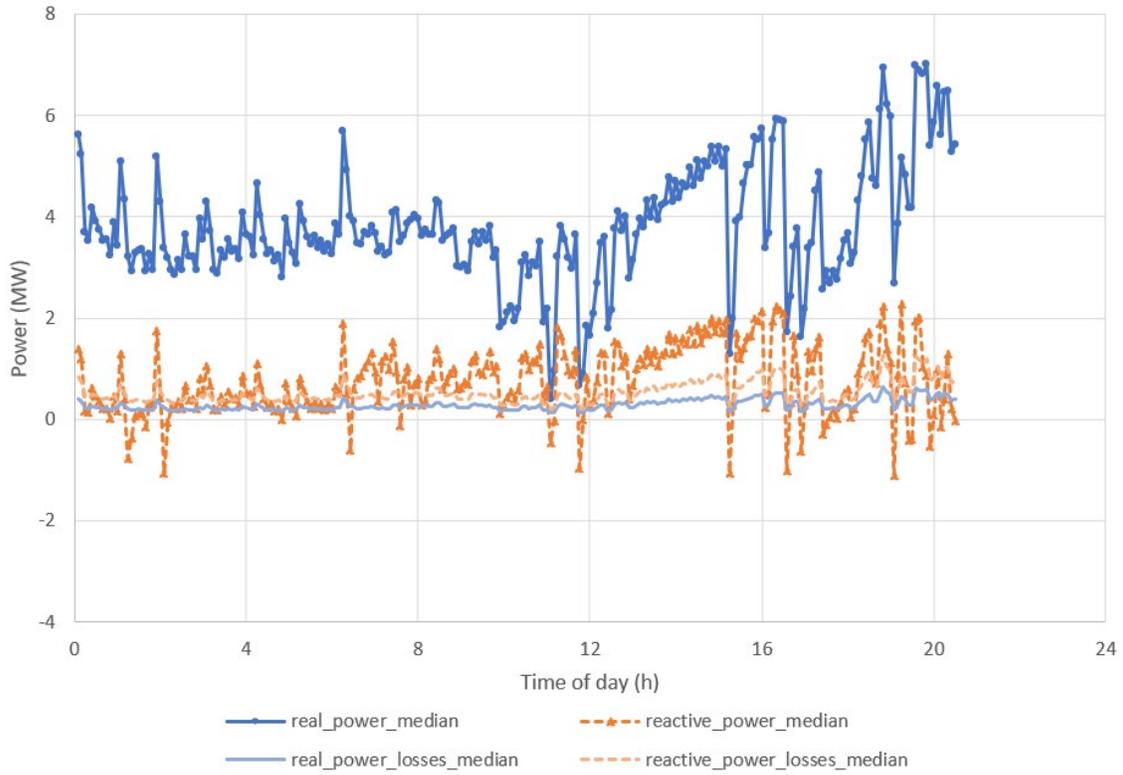


Fig. 21 Substation power flows for the LFC, July 7.

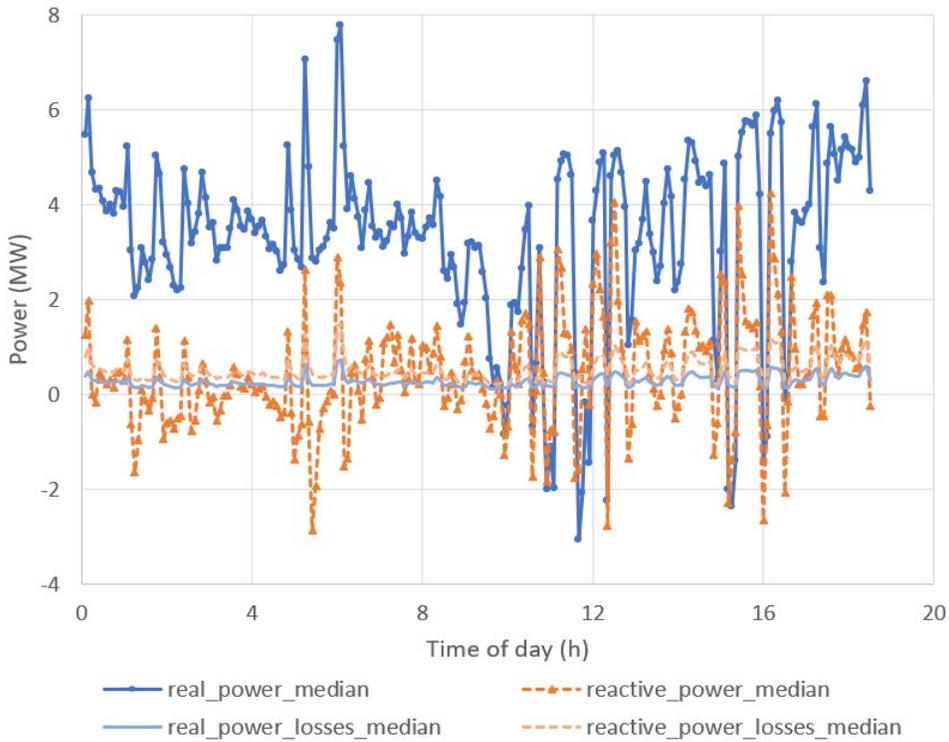


Fig. 22 Substation power flows for the RTC, July 7.

Note that the power flow figures above present the feeder aggregate flows at the substation represented by the average of five recorded minutely values during each 5 min time interval. There is actually tremendous variation of power flows among the minutely data within each 5 min time interval. The *average* max-min difference seen in the minutely data is 1507 kW, approximately twice the magnitude of the 5 min average power flow changes shown in Table 1. This implies that there are even greater power flow changes than the ones indicated in these power flow plots.

These power flow changes result in voltage fluctuations as seen in Fig. 23 (RTC July 7), Fig. 24 (RTC July 6), and in Fig. 25 (LFC July 7). The voltages presented here are the average voltages across all 1977 house meters on the 8500 grid. July 6 fluctuations are less than July 7, and LFC fluctuations less than RTC, as seen in the voltage violation and power flow data given in Table 1.

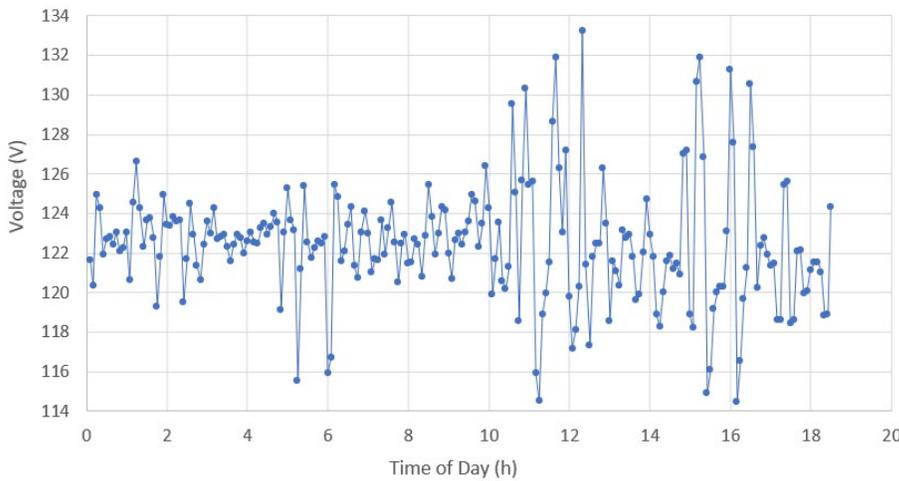


Fig. 23 Average hourly voltage across house meters, 8500 grid using RTC on July 7.

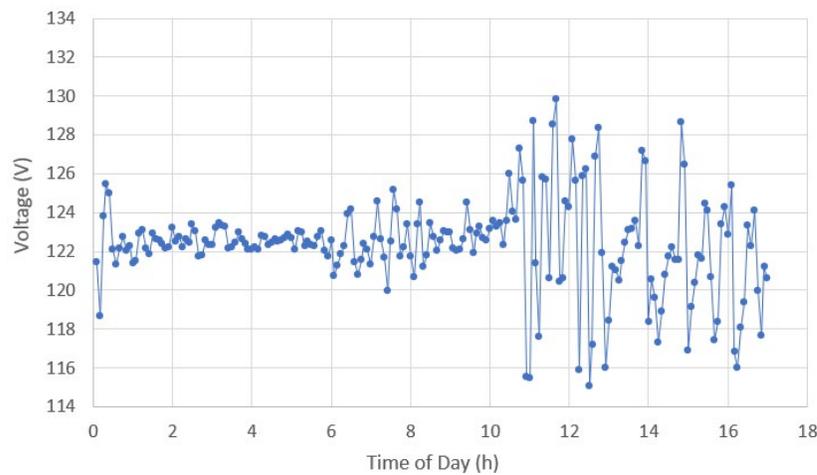


Fig. 24 Average hourly voltage across house meters, 8500 grid using RTC on July 6.

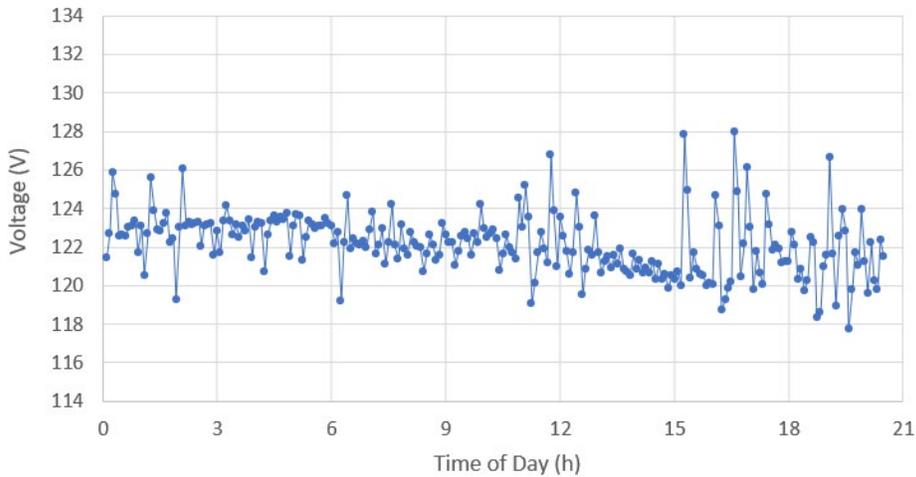


Fig. 25 Average hourly voltage across house meters, 8500 grid using LFC on July 7.

3.4.3. Voltage Regulator and Capacitor Bank Action

These changes in power and voltage cause voltage regulators and capacitor banks to adjust their operation as well. There are four voltage regulators and capacitor banks on the IEEE 8500 grid as shown in Fig. 1. As can be seen in Fig. 26 and Fig. 27, the July 7 results show that voltage regulators and capacitor banks see significantly more switching action in response to RTC control than to the LFC control. Both RTC and LFC cause increased switching action compared to the baseline controller. This increase in switching action leads to more wear and tear on the grid hardware.

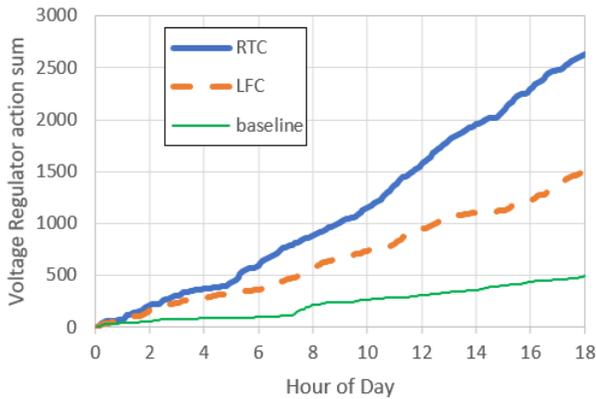


Fig. 26 Sum of voltage regulator actions for July 7.

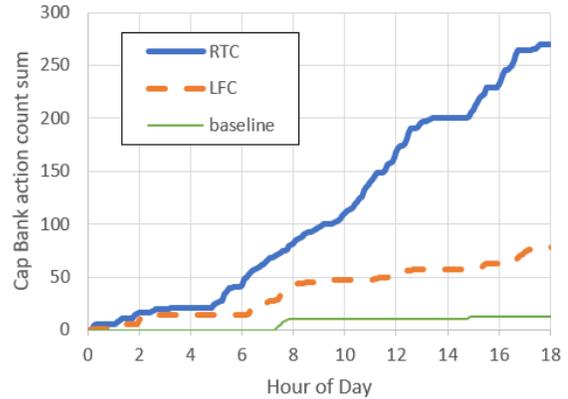


Fig. 27 Sum of capacitor bank control actions for July 7.

3.4.4. Correlation of Voltage and Price Movement

Consider the T_{csp} profiles (baseline, LFC, RTC) for the three houses shown in Fig. 13, Fig. 15, and Fig. 18. For the baseline case, each home has a flat temperature across the day (T_{set}), which produces the stable power flow seen in Fig. 20. For RTC, the adjusted T_{csp} moves in proportion to the RTP (except when the temperature is clipped at T_{max}). Raising T_{csp} leads to heat pumps shutting off while lowering T_{csp} leads to heat pumps turning on. A large

movement in price of electricity results in a large movement in T_{csp} . If a price jump follows a large price drop, there would be a change from all heat pumps ON to all heat pumps OFF in a single action with a resulting drop of approximately 6 MW load. This can be observed in several instances in Fig. 22. The end result is synchronization of heat pumps (or more generally, all price-responsive loads) on the grid.

Fig. 28 overlays the average RTC T_{csp} profile of all houses with the substation average real power flow. Fig. 29 shows the correlation of changes in temperature at each 5-min step with the changes in power flow at each step along with a fit line. The resulting correlation coefficient is $R = -0.83$.

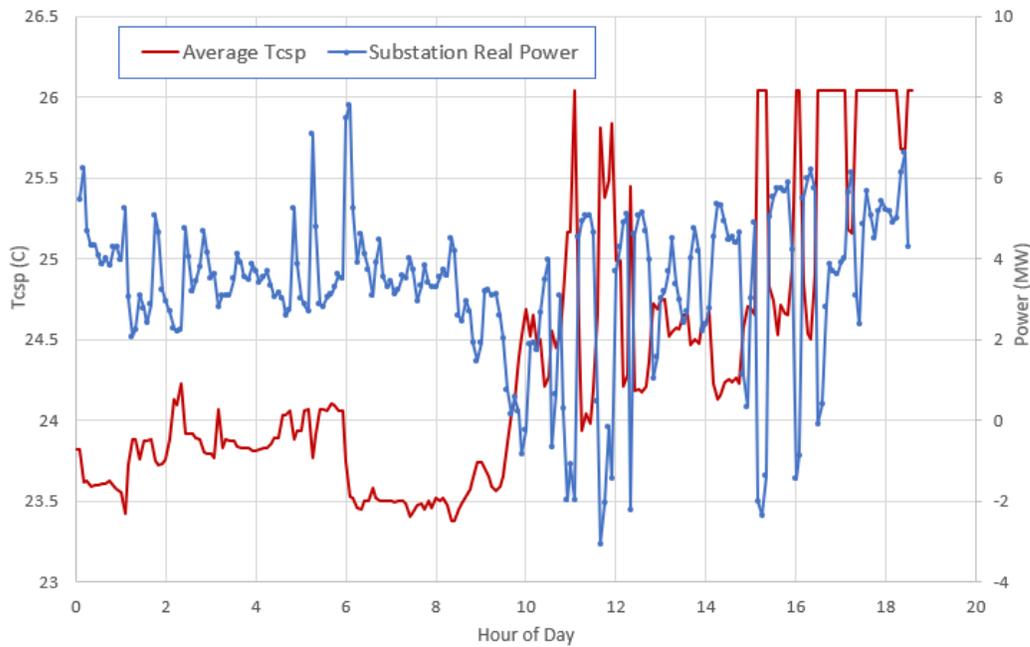


Fig. 28 RTC average T_{csp} across all 8500 grid houses compared to the resulting substation real power flow at each 5-min interval, July 7.

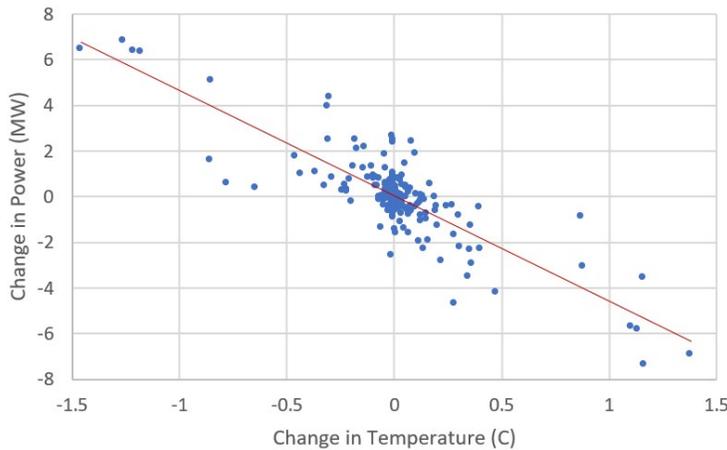


Fig. 29 Correlation of stepwise change in substation power flow with change in average house T_{csp} for each 5-min time interval, July 7, RTC, 8500 grid. $R = -0.83$. Power flow and T_{csp} as shown in Fig. 28.

The T_{csp} is itself driven by price. The relationship of RTC T_{csp} to price is linear, as can be seen from Eq. 2. However, because T_{csp} is clipped at T_{max} for high prices, the relationship is not linear in implementation, as can be seen in Fig. 15 where RTP and T_{csp} are both shown. The correlation of change in T_{csp} to change in RTP has a coefficient of $R = 0.51$, as shown in Fig. 30 with the fit line. However, 85 % of the data points fall on the more vertical line of points clustered near the origin. All the points off that line are due to clipping of T_{csp} .

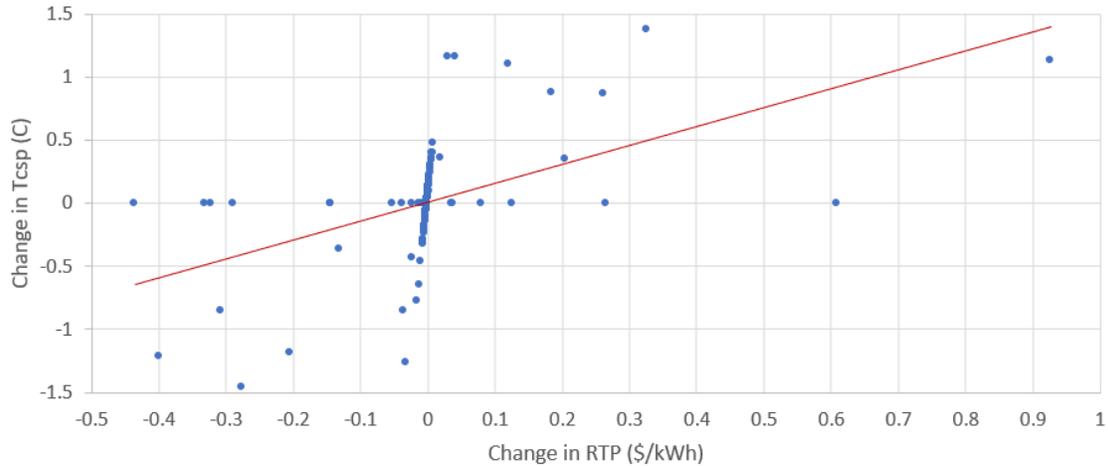


Fig. 30 Correlation of change in average house T_{csp} with change in RTP for each 5-min time interval, July 7, RTC, 8500 grid, with $R = 0.51$.

We see that price changes induce T_{csp} changes which in turn cause power flow changes. And voltage changes are strongly correlated with these power flow changes (Fig. 31), with a correlation coefficient $R = -0.91$.

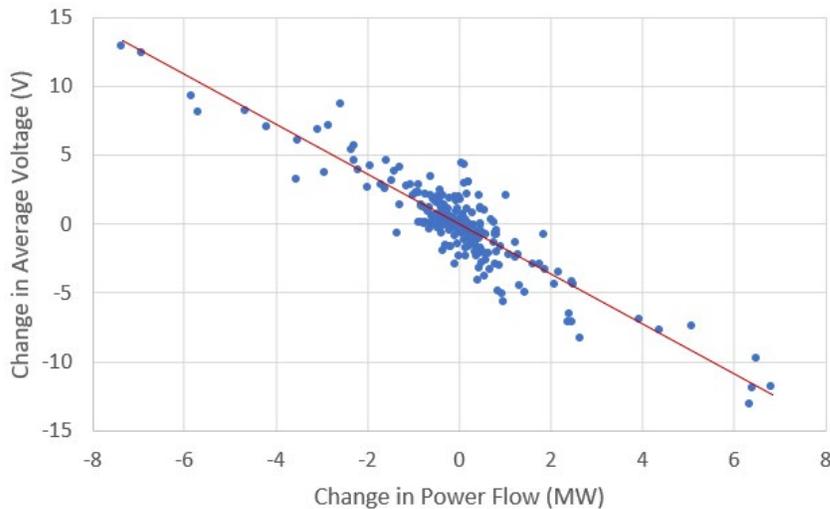


Fig. 31 Correlation of stepwise change in average house meter voltage with change in substation power flow, for each 5-min time interval, July 7, RTC, 8500 grid.

Finally, we can look at the correlation of voltage to price, which combines the correlations given in Fig. 29 to Fig. 31. This correlation is shown in Fig. 32 with a relatively low correlation coefficient of $R = 0.36$. The large majority of price changes are on the order of

0.01 \$/kWh. For the larger price changes, it is notable that positive changes in price almost always result in positive changes in voltage except for a handful of data points, and the same result is seen looking at negative price changes resulting in negative voltage changes.

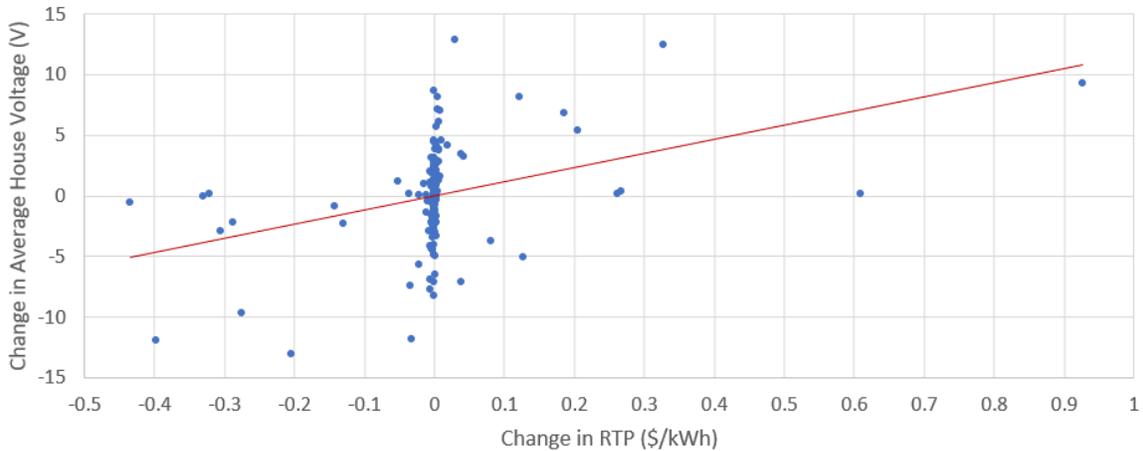


Fig. 32 Correlation of stepwise change in average house meter voltage with change in RTP, for each 5-min time interval, July 7, RTC, 8500 grid.

In the case of the LFC controller, T_{csp} is determined for each house independently by optimizing the temperature profile across the next day using the RTP_{avg} . Although each house profile is unique, there is still a pattern driven by the price signal. The T_{csp} , averaged across all houses, is shown in Fig. 33 along with the RTP_{avg} price. There is no strong correlation between RTP_{avg} and T_{csp} , while we see the LFC controller precooling prior to the peak price period and then allowing the house temperature to rise during the peak prices.

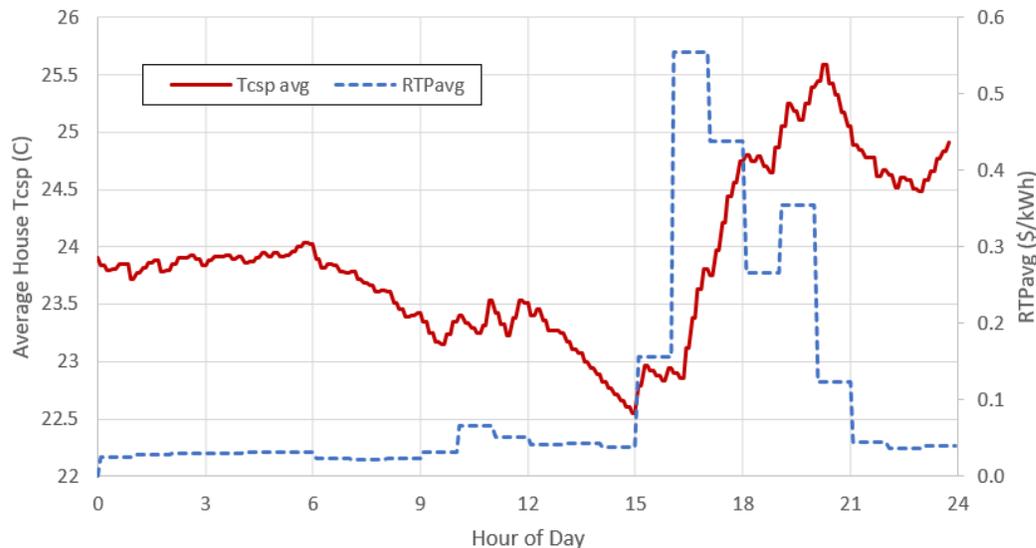


Fig. 33 LFC average house T_{csp} with corresponding RTP_{avg} price, July 7, 8500 grid.

The average T_{csp} is shown with the corresponding substation power flow in Fig. 34. In this case, one can see that many of the temperature drops are accompanied by power flow jumps. In fact, in most cases, each change in direction of the T_{csp} trend (from down trend to up, or up

to down) is accompanied by a change in power flow on the order of 1 MW. The correlation coefficient for change in power flow with change in temperature is $R = -0.45$, with data shown in Fig. 35.

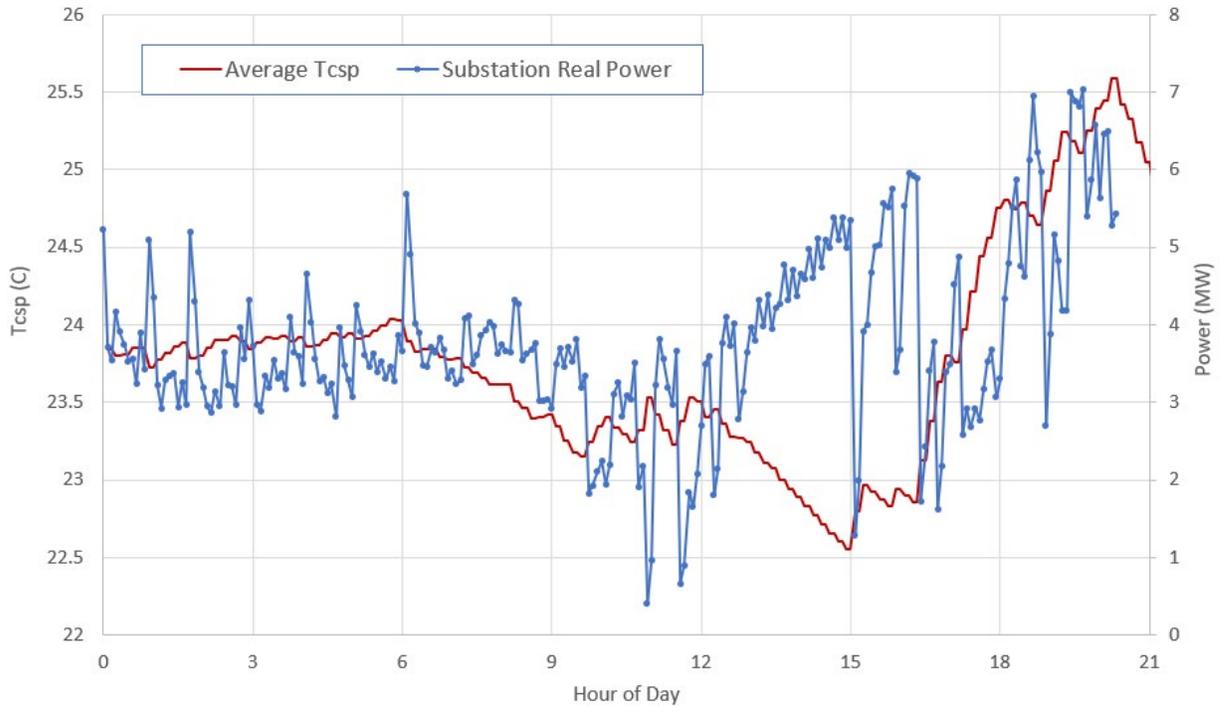


Fig. 34 LFC average T_{csp} across all 8500 grid houses together with the resulting substation real power flow at each 5-min interval

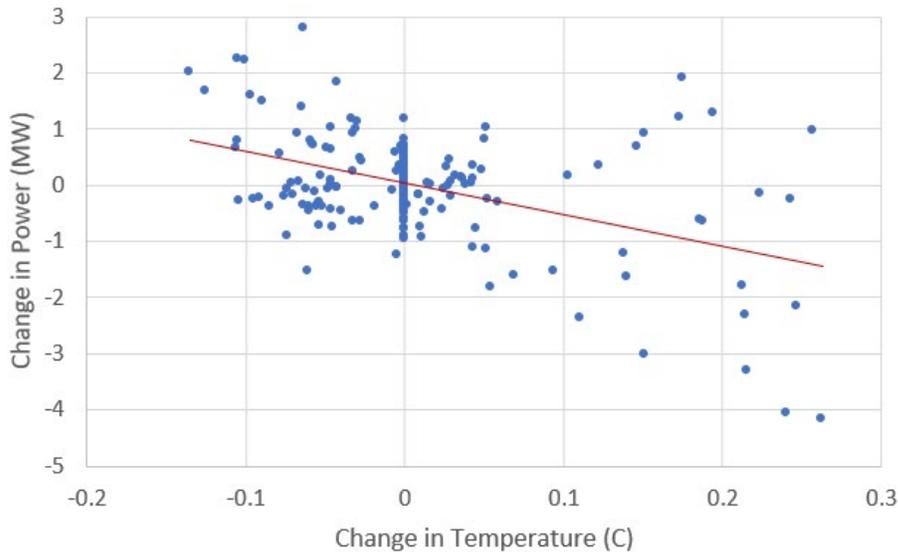


Fig. 35 Correlation of the change in substation power with the change in T_{csp} for LFC controller, July 7, 8500 grid.

The key point to be made here is that, despite the relatively gentle LFC T_{csp} movements (compared to the RTC), there are movements in power and voltage that are correlated with T_{csp} , which is itself based on an hourly price signal. Synchronization across houses is an issue

when there are changes in T_{csp} trend from cooling to warming (or *vice versa*), as evidenced by resulting significant changes in power flow. This is then reflected in the voltage with fluctuations and voltage violations (Table 1), along with voltage regulator and capacitor bank actions.

3.5. R4-1 Grid Power Quality

The R4-1 grid was included in the test plan to observe power quality impact of price changes on a different grid for comparison to the 8500 grid results. The R4-1 grid is more robust, producing less voltage fluctuation for a given power flow change. The experiments presented in previous sections for the 8500 grid were repeated for the RTC, using the same CAISO RTP and Tucson weather, similar houses and the same metrics.

Fig. 36 shows the substation power flow for the no control baseline case for the R4-1 grid on July 7, with the same duck curve profile peaking at 3 MW load at 8 p.m., as compared to 8 MW on the 8500 grid. Whereas the 8500 grid is entirely residential, the R4-1 residential load is, on average, 70 % of the total load with the remainder coming from some small commercial and industrial loads spread across the feeder.

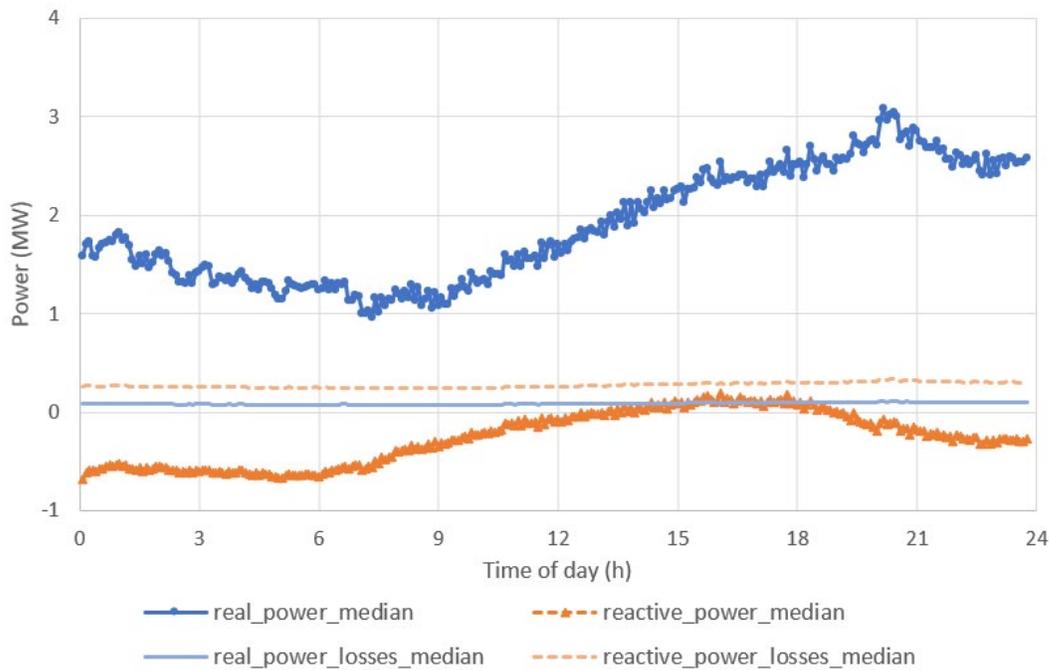


Fig. 36 July 7 R4-12.47-1 grid baseline substation power flow.

The impact of the RTC controller on the R4-1 substation power flow can be seen in Fig. 37. The real power standard deviation is close to 1 MW, which relative to the peak load is the same as what was seen for the 8500 grid (Fig. 20).

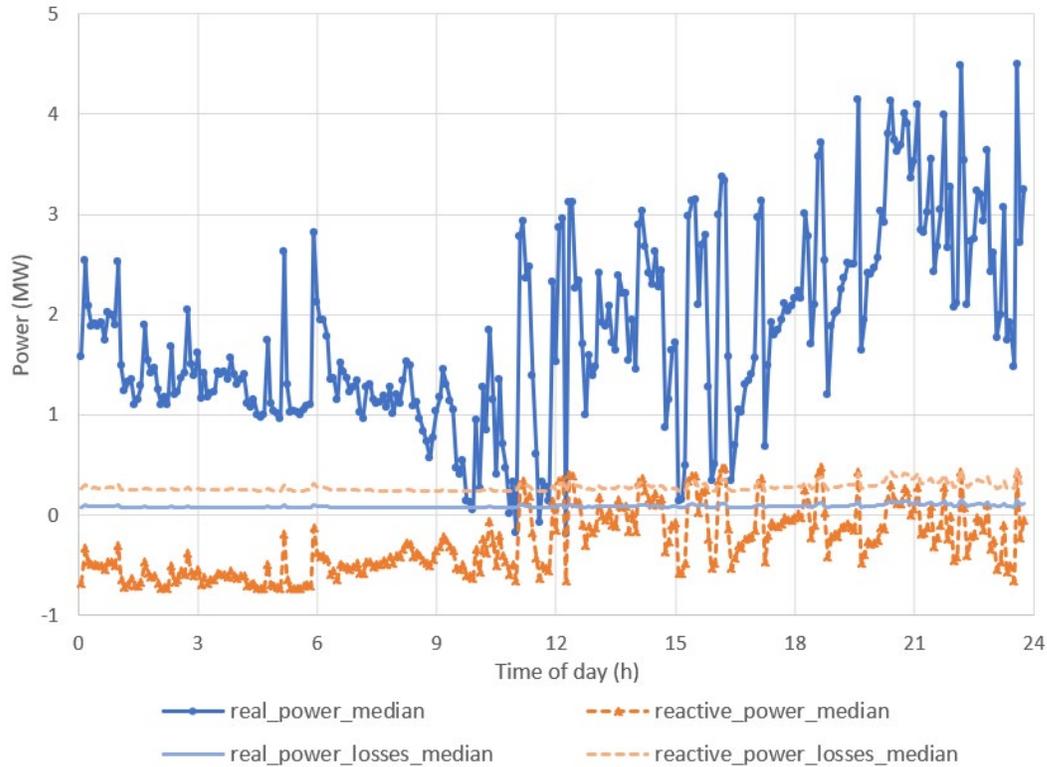


Fig. 37 July 7 R4-12.47-1 grid RTC substation power flows.

The average house meter voltage deviations are shown in Fig. 38, corresponding to the power flows in Fig. 37. The voltage fluctuations are about 20 % or less of what was seen for the corresponding 8500 grid data shown in Fig. 23. This clearly shows that the same relative power flows on the R4-1 grid produce significantly less voltage disturbance such that grid voltages could be kept within limits. In other words, smaller actual power flow fluctuations on the same size grid wires result in smaller voltage fluctuations. While one sees price-induced power flow oscillations, the more robust R4-1 grid has more stable voltages.

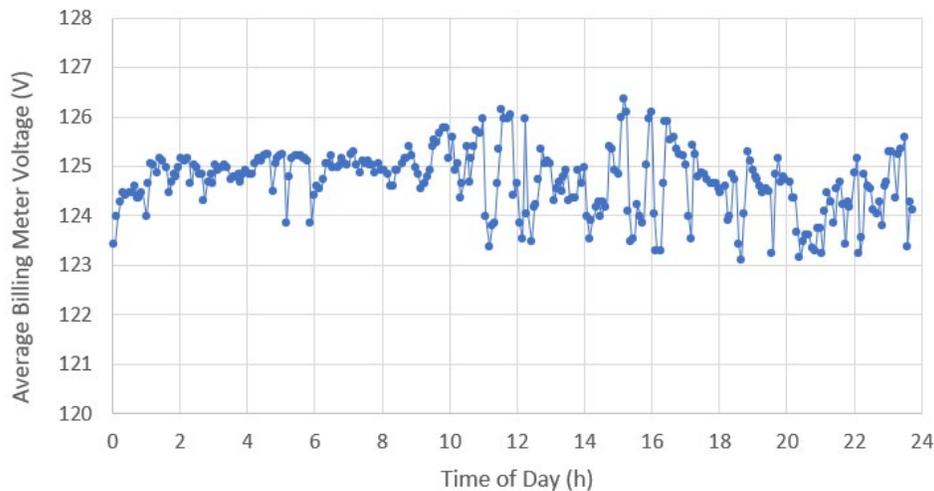


Fig. 38 July 7 R4-12.47-1 RTC average voltage across house meters.

The RTC substation power flow overlays the T_{csp} in Fig. 39. The correlation of changes in substation power flow to changes in T_{csp} is shown in Fig. 40, with $R = -0.83$. The results appear very similar to results seen in Fig. 28 and Fig. 29 for the 8500 grid.

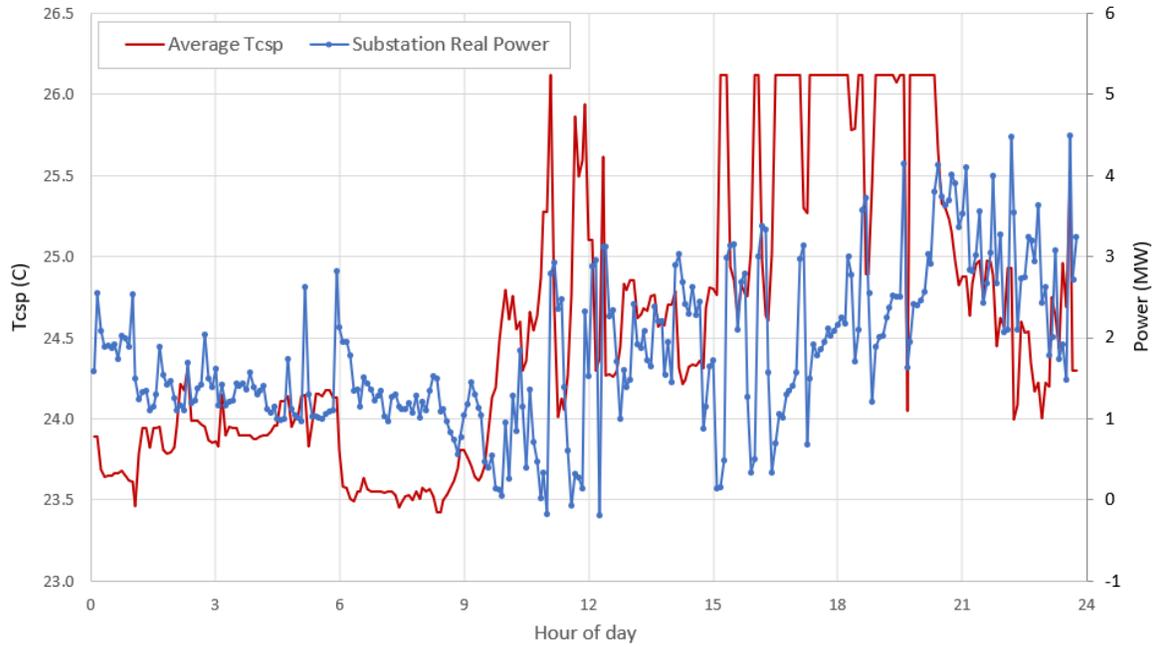


Fig. 39 RTC average T_{csp} for R4-1 grid houses compared to the resulting substation real power flow at each 5-min interval, July 7.

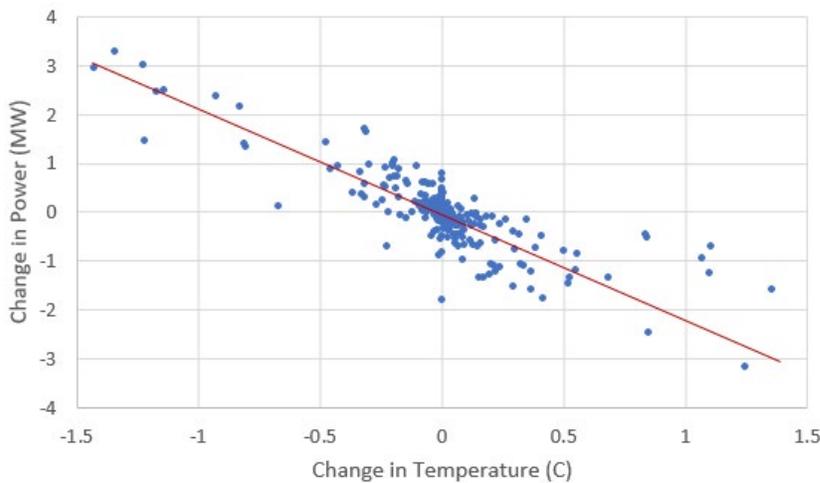


Fig. 40 Correlation of stepwise change in substation power flow with change in average house T_{csp} for each 5-min time interval, July 7, RTC, R4-1grid, $R = -0.83$.

The correlation of change in voltage to change in power is $R = -0.96$, and the final correlation of voltage to price has a correlation coefficient of $R = 0.27$. These R4-1 grid results, with price-responsive load reaction to price changes inducing power flow and voltage volatility, confirm the results seen with the 8500 grid. However, we note that the R4-1 grid is more stable with respect to voltage due to lower actual power flows for the same wire diameters.

3.6. Discussion

These experiments have demonstrated the differences in energy management for two heat pump controllers managing indoor temperature based on price signals and customer comfort selections. The results provide some insights into potential power quality issues that may be encountered when deploying price-based control algorithms.

The experiments demonstrate that energy and cost reductions are possible while maintaining occupant comfort within selected ranges. For this hot summer scenario, energy is saved when the house temperature is raised above the base setpoint, and cost is lowered by reducing heat pump run time when prices are high.

Results show the sensitivity of cost savings to house thermal characteristics, customer comfort preferences, the presence of rooftop PV, and net metering policies. Only heat pump control was considered in this work, but it can be expected that an individual customer would save more or less money depending on when they operate other household devices (e.g., the clothes dryer) or whether they have price-aware controllers on other appliances such as the hot water heater.

The two control approaches are sufficiently different that combining them could produce additional cost savings. That is, the RTC algorithm might benefit from day-ahead planning to take advantage of pre-cooling and use of the resulting stored thermal energy during the highest price period. Likewise, the LFC might achieve additional cost savings by responding to short-term extreme price spikes.

Use of dynamic price control can deliver significant adjustment of feeder load to benefit the bulk grid, but the bulk market price does not consider distribution voltage. A number of serious power quality issues were observed in the simulations presented here: voltage violations, voltage volatility, and related impact on voltage regulator and capacitor bank actions. The voltage violations are a key factor determining hosting capacity, and it appears that response to dynamic prices has the potential to reduce hosting capacity. This raises several questions that should be discussed.

1. Are these power and voltage fluctuations representative of what we would see on a real distribution grid, considering that these simulated grids use identical heat pump controllers responding to the same signals?
2. Is it reasonable to base conclusions on an extreme peak day like July 7?
3. Is the bulk CAISO RTP signal a realistic signal to pass to customer DER?

In response to question 1, the two reference grids used in these experiments are based on two real-world feeders. The only significant change is that nearly all the residential homes are fitted with rooftop solar PV which contributes to over-voltage conditions for the 8500 grid, even for the baseline case. Additionally, all houses have price-responsive heat pump temperature control. Most of the observed voltage violations (Table 1) are not due to the PV, but rather to power flow fluctuations induced by the price-responsive heat pump controllers.

One may point to utility pilots and existing dynamic price tariffs and the apparent lack of any of these power flow fluctuations and resulting voltage violations. However, the authors are not aware of published power flow and voltage violation data. In addition, existing dynamic

tariffs are not mandatory, but subscribed to by only a limited number of customers. And finally, there are likely a limited number of devices responding automatically to the price signals. Thus, the penetration of price-responsive load on any existing grid is much less (as a percentage of total load) than seen in this simulation study. Yet, as utilities seek to integrate more customer resource flexibility, and as customers add more price-responsive DER (consider EVs and batteries in addition to electric hot-water heaters), it is possible that dynamic tariffs will become more common and that more loads will be price-responsive.

If the primary controller goal is to save the customer money, then the behavior of different controllers will align—use power when the price is low and avoid consumption when the price is high. This leads to the conclusion that the voltage instability seen in this study is very possible on a future grid with a higher percentage of price-responsive load despite a heterogeneity of devices and controllers. Nonetheless, there may be other drivers that reduce synchronized price response, such as an aggregator selling flexibility into a regulation market, or a vendor marketing a controller that seeks to only operate a device when renewable power is available, independent of an energy market price signal. Price signal volatility may also be reduced.

Furthermore, because the control algorithms used in this work are based solely on bulk system energy price and customer comfort preference, distribution grid voltage and power flow conditions are treated as externalities to these control schemes. It is therefore not surprising that factors kept external to the market optimization scheme will not be optimized. A primary lesson from this work is that, under very high DER penetrations, grid operations would benefit from the development and use of control strategies that internalize some of the considerations examined in this work but excluded from consideration by the implemented controllers.

In response to question 2, note that voltage violations in Table 1 were similar comparing July 6 to July 7. Even the LFC, which uses an hourly price signal and relatively gentle temperature adjustments, saw significant voltage violations on July 6 compared to baseline. With regard to these relatively small and gentle temperature adjustments impacting voltage, it was noted that a change in the T_{csp} trend direction from up to down or down to up induces synchronized demand adjustment leading to the observed voltage volatility. Per the discussion above, the heat pump controllers only look at bulk price and comfort parameters and do not consider voltage impact. A distribution utility should be able to moderate the price signal to reduce voltage impacts.

Finally, for question 3, we should understand how the CAISO Real-Time Market (RTM) works, what the signal communicates, and then consider what utilities are doing today as they roll out dynamic prices. The CAISO RTM is an imbalance market used to match generation to supply in real time. Generally, market clearing prices are higher when the grid load is higher. However, RTM prices can spike or drop in unexpected ways due to unscheduled or must-run generation or load that was not considered in forecasts [37]. In addition, volatility can be partially attributed to a small pool of available generators participating in the market at peak load conditions. One may conclude, based on [37], that RTP is a useful measure of real-time grid load that might be used to incentivize additional distribution load and generation to reduce imbalances. However, the results presented here

provide a caution—more price-responsive loads will lead to more power flow fluctuations that may be a problem on some grids.

We can observe now that some utilities and local public utility commissions recognize the importance of accessing customer DER flexibility and the value of a dynamic price signal. The bulk market RTP or DAP provide readily available signals. Some utilities are using a real-time 5-min price based on RTP [12]. Some are using a day-ahead hourly signal based on DAP [38]. Some are using an hourly signal based on RTP_{avg} [39]. All of these are in play and may become more common.

Some principles can be discerned out of this discussion.

- Price variation incentivizes flexibility to help balance the grid, but price-responsive controllers will have some impact on voltage.
- DAP and RTP communicate bulk grid capacity needs but do not consider distribution grid constraints. Some price signals may not be suitable for more-constrained grids.
- Controllers with the goal of saving money will adjust power usage to avoid consumption when price is higher and shift the load to when price is lower.
- DAP and RTP present different time scales of price variation leading to different effective strategies for using thermal storage to achieve cost savings. Controllers must adapt their strategy to the available electricity tariff(s).
- An increased number of price-responsive controllers as well as larger changes in price will both lead to bigger power flow changes and greater savings to customers.
- A more volatile price signal will tend to induce more synchronized load response with corresponding voltage volatility.
- The degree to which power flow changes impact voltage depends on grid wires and hardware. Increased wire thickness reduces voltage volatility.

There are some potential methods to reduce the impact of price response on grid voltage fluctuations. The first approach is to smooth the price signal. Price volatility may be reduced to avoid large movements in power flow and voltage. This might be accomplished with a smoothed RTP signal as well as removal of price spikes.

A second approach is to consider adjusting a price signal based on local voltage levels. This might be used to align peak prices with the time of local peak load (which may be different from time of the peak of the bulk grid). Alternatively, prices might be adjusted at a more granular level: lower prices at a PV hot spot (high voltage) or higher at an EV load pocket (low voltage). The adjustment could be made such that the average price at all points on the grid is the same across a day while the amount of price variation may be more or less in one location or another.

Another solution may be to avoid billing a customer based on 5-min RTP, which incentivizes fast response to price changes. An alternative, as used in [39], may be to communicate the RTP, but to bill based on the past hour's RTP_{avg} . In this way, smart devices will see increasing prices and try to reduce consumption across the hour but not react as strongly to every price change.

Finally, the utility may consider various hardware solutions, such as use of more capable voltage regulators, or installation of batteries for the sole purpose of reducing voltage volatility. In this case, it seems the batteries would be working against the price signal, thus indicating the price by itself is not ideal. But this might be an overall simpler solution that does not require adjusting prices throughout a system. This approach and the others just mentioned, when combined with existing tools like Volt-Var control, should help to boost the hosting capacity of distribution grids while still successfully engaging customer flexibility.

4. Conclusion

Two transactive management approaches were simulated side-by-side, using the same houses on the same grids with the same weather. The rule-based real-time price controller adjusts house cooling setpoint temperature directly in response to RTP, considering current house temperature relative to the setpoint temperature, customer comfort, and previous day price average and standard deviation. The model-based load forecasting controller determines the cooling setpoint a day ahead based on hourly prices for the next day. This enables planning for precooling and avoidance of peak price times, making optimal use of available thermal storage while considering customer comfort. The two approaches respond to different incentive signals which makes direct comparison difficult. For these tests, the 5-minute real-time market price signal was averaged over each hour to produce a comparable signal for the model-based load forecasting controller to use as a day-ahead price.

Simulation results show that both controllers reduce homeowner electricity cost below a baseline with a fixed heat pump temperature setpoint. The LFC can perform better than baseline for cost-conscious homeowners by using pre-cooling strategies to take advantage of short-term thermal storage. The RTC can save money in the presence of a volatile RTP by responding in real-time with setpoint temperature adjustments.

However, both of these price-responsive controllers induced power flow fluctuations with resulting voltage fluctuations. For the more constrained 8500 grid, the result was significant voltage violations, more so for the RTC than the LFC. For the more robust R4-1 grid, voltage fluctuations were significantly less. On the 8500 grid, LFC voltage violations (Table 1) increased on the order of 400 % above baseline. The RTC controller doubled that again for voltage deviations, voltage regulator actions and capacitor bank actions leading to reduced power quality.

Analysis of the impact of dynamic prices on voltage suggests these results are indicative of voltage problems on a future grid with a high percentage of price-responsive loads if care is not taken to manage volatility of the price signal and if the grid is near its hosting capacity. It appears that large price changes, or a shift in price trend from rising to falling (or *vice versa*) can result in synchronized load response, with devices turning on or off together. This can result in large shifts in power flow at the substation. The grid must be able to handle these jumps in power flow and resulting voltage changes.

Future research will investigate methods for further reducing the voltage volatility to support increased hosting capacity and improve power quality while still engaging customer flexibility to support a high-DER grid.

5. References

- [1] EIA - Annual Energy Outlook 2021 Available at https://www.eia.gov/outlooks/aeo/tables_ref.php
- [2] Neukomm M, Nubbe V, Fares R (2019) Grid-interactive Efficient Buildings Technical Report Series: Overview of Research Challenges and Gaps.
- [3] Villara J, Bessab R, Matosc M (2018) Flexibility products and markets: Literature review. *Electric Power Systems Research* 154.
- [4] Yan X, Ozturk Y, Hu Z, Song Y (2018) A review on price-driven residential demand response. *Renewable and Sustainable Energy Reviews* 96:411–419. <https://doi.org/10.1016/j.rser.2018.08.003>
- [5] Staff (2020) 2020 Assessment of Demand Response and Advanced Metering. Available at [https://cms.ferc.gov/sites/default/files/2020-12/2020 Assessment of Demand Response and Advanced Metering_December 2020.pdf](https://cms.ferc.gov/sites/default/files/2020-12/2020%20Assessment%20of%20Demand%20Response%20and%20Advanced%20Metering_December%202020.pdf)
- [6] Faruqui A, Aydin MG (2017) Moving Forward with Electricity Tariff Reform. *Regulation*:42. Available at <https://www.cato.org/sites/cato.org/files/serials/files/regulation/2017/9/regulation-v40n3-5.pdf>
- [7] Faruqui A (2021) *The Rate Design Imperative: Why the Status Quo is Not Viable* (Brattle). Available at https://brattlefiles.blob.core.windows.net/files/22871_the_rate_design_imperative_-_why_the_status_quo_is_not_viable_indiana_ga_slides.pdf
- [8] A. Faruqui (2020) *Time of Use and Dynamic Pricing Rates in the US* (Florence School of Regulation). Available at <https://fsr.eui.eu/time-of-use-and-dynamic-pricing-rates-in-the-us/>
- [9] Xie Q, Hui H, Ding Y, Ye C, Lin Z, Wang P, Song Y, Ji L, Chen R (2020) Use of demand response for voltage regulation in power distribution systems with flexible resources. *IET Generation, Transmission and Distribution* 14(5):883–892. <https://doi.org/10.1049/IET-GTD.2019.1170>
- [10] Nair VJ, Venkataramanan V, Haider R, Annaswamy A (2021) A Hierarchical Local Electricity Market for a DER-rich Grid Edge. Available at <https://arxiv.org/abs/2110.02358v1>
- [11] FERC Staff Report (2015) Energy Primer: A Handbook of Energy Market Basics. Available at <https://www.ferc.gov/sites/default/files/2020-05/energy-primer.pdf>
- [12] Cazalet EG, Kohanim M, Hasidim O (2020) Complete and Low-Cost Retail Automated Transactive Energy System (RATES). Available at <https://www.energy.ca.gov/publications/2020/complete-and-low-cost-retail-automated-transactive-energy-system-rates>
- [13] Bai L, Wang J, Wang C, Chen C, Li F (2018) Distribution Locational Marginal Pricing (DLMP) for Congestion Management and Voltage Support. *IEEE Transactions on Power Systems* 33(4):4061–4073. <https://doi.org/10.1109/TPWRS.2017.2767632>
- [14] Singh AK, Parida SK (2018) A review on distributed generation allocation and planning in deregulated electricity market. *Renewable and Sustainable Energy Reviews* 82(3). <https://doi.org/doi.org/10.1016/j.rser.2017.10.060>

- [15] Coster E, Myrzik J, Kruimer B, Kling W (2011) Integration Issues of Distributed Generation in Distribution Grids. *Proceedings of the IEEE, Vol. 99* Available at <https://ieeexplore.ieee.org/abstract/document/5565399>
- [16] Xie Q, Hui H, Ding Y, Ye C, Lin Z, Wang P, Song Y, Ji L, Chen R (2020) Use of demand response for voltage regulation in power distribution systems with flexible resources. *IET Generation, Transmission & Distribution* 14(5):883–892. <https://doi.org/10.1049/iet-gtd.2019.1170>
- [17] Australian Energy Market Operator (2019) Technical Integration of Distributed Energy Resources. Available at <https://aemo.com.au/-/media/Files/Electricity/NEM/DER/2019/Technical-Integration/Technical-Integration-of-DER-Report.pdf>
- [18] IEEE (2018) *IEEE 1547 2018 Standard for Interconnection and Interoperability of Distributed Energy Resources with Associated Electric Power Systems Interfaces*.
- [19] Haider R, D’Achiardi D, Venkataramanan V, Srivastava A, Bose A, Annaswamy AM (2021) Reinventing the Utility for DERs: A Proposal for a DSO-Centric Retail Electricity Market. <https://doi.org/10.1016/j.adapen.2021.100026>
- [20] Roth T, Song E, Burns M, Neema H, Emfinger W, Sztipanovits J (2017) Cyber-Physical System Development Environment for Energy Applications. *ASME 2017 11th International Conference on Energy Sustainability* (American Society of Mechanical Engineers). <https://doi.org/10.1115/ES2017-3589>
- [21] Scudder R, Saunders R, Möller B, Mores K l. (2010) IEEE Standards for M&S HLA-Framework and Rules. 41. Available at <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5953411>
- [22] Burns M, Song E, Holmberg D (2016) The Transactive Energy Abstract Component Model. <https://doi.org/10.6028/NIST.SP.1900-750>
- [23] Holmberg D, Burns M, Bushby S, Gopstein A, McDermott T, Tang Y, Huang Q, Pratt A, Ruth M, Ding F, Bichpuriya Y, Rajagopal N, Ilic M, Jaddivada R, Neema H (2019) NIST transactive energy modeling and simulation challenge phase II final report. (Gaithersburg, MD). <https://doi.org/10.6028/NIST.SP.1900-603>
- [24] GridLAB-D power distribution system simulator Available at <https://gridlabd.org>
- [25] Arritt RF, Dugan RC (2010) The IEEE 8500-node test feeder. *2010 IEEE PES Transmission and Distribution Conference and Exposition: Smart Solutions for a Changing World*. <https://doi.org/10.1109/TDC.2010.5484381>
- [26] Feeder Taxonomy Wiki Available at http://gridlab-d.shoutwiki.com/wiki/Feeder_Taxonomy
- [27] Schneider K, Chen Y, Chassin D, Pratt R, Engel D, Thompson S (2008) Modern Grid Initiative Distribution Taxonomy Final Report. Available at http://emac.berkeley.edu/gridlabd/taxonomy_graphs/
- [28] Cohen M GridLAB-D Taxonomy Feeder Graphs.
- [29] Widergren S, Somani K Subbarao C Marinovici JC Fuller JL Hammerstrom DP Chassin SA (2014) AEP Ohio gridSMART Demonstration Project Real-Time Pricing Demonstration Analysis.

- [30] Widergren Somani K Subbarao C Marinovici JC Fuller JL Hammerstrom DP Chassin SA (2014) AEP Ohio gridSMART -Demonstration Project Real-Time Pricing Demonstration Analysis. Available at https://www.pnnl.gov/main/publications/external/technical_reports/PNNL-23192.pdf
- [31] GridLAB-D WiKi Available at <http://gridlab-d.shoutwiki.com/wiki/Spec:Market>
- [32] Omar F, Holmberg D (2021) Load Forecasting Tool for NIST Transactive Energy Market.
- [33] Omar F, Bushby ST, Williams RD (2017) A self-learning algorithm for estimating solar heat gain and temperature changes in a single-Family residence. *Energy and Buildings* 150:100–110. <https://doi.org/10.1016/J.ENBUILD.2017.06.001>
- [34] Omar F, Bushby ST, Williams RD (2019) Assessing the performance of residential energy management control Algorithms: Multi-criteria decision making using the analytical hierarchy process. *Energy and Buildings* 199:537–546. <https://doi.org/10.1016/J.ENBUILD.2019.07.033>
- [35] ANSI C84.1-2020: Electric Power Systems Voltage Ratings (60 Hz) - ANSI Blog Available at <https://blog.ansi.org/2020/10/ansi-c84-1-2020-electric-voltage-ratings-60/>
- [36] Holmberg D, Omar F (2018) Characterization of residential distributed energy resource potential to provide ancillary services. (Gaithersburg, MD). <https://doi.org/10.6028/NIST.SP.1900-601>
- [37] CAISO Dept of Market Monitoring (2006) Real Time Market Performance. Available at https://www.caiso.com/Documents/Chapter3_RealTimeMarketPerformance11-Apr-06.pdf
- [38] GA Power (2020) *Real-Time Pricing Large Customer Tariff, RTP-DA-5*. Available at <https://www.georgiapower.com/content/dam/georgiapower/pdfs/business-pdfs/rates-schedules/RTP-DA-5.pdf>
- [39] Commonwealth Edison *Hourly Pricing Program*. Available at <https://hourlypricing.comed.com/>
- [40] Holmberg D, Hardin D, Koch EE Towards Demand Response Measurement and Verification Standards. *Grid Interop 2012* Available at https://www.gridwiseac.org/pdfs/forum_papers12/holmberg_hardin_koch_paper_gi12.pdf