

NIST Technical Note
NIST TN 2268r2

Flexible Resource Controller

*A Price and Voltage-Responsive Device Controller
for GridLAB-D Co-Simulations*

David Holmberg
Thomas Roth

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.TN.2268r2>

NIST Technical Note
NIST TN 2268r2

Flexible Resource Controller

*A Price and Voltage-Responsive Device Controller
for GridLAB-D Co-Simulations*

David Holmberg
*Building Energy and Environment Division
Engineering Laboratory*

Thomas Roth
*Wireless Networks Division
Communications Technology Laboratory*

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.TN.2268r2>

September 2026



U.S. Department of Commerce
Howard Lutnick, Secretary

National Institute of Standards and Technology
Arvind Raman, NIST Director and Under Secretary of Commerce for Standards and Technology

Certain equipment, instruments, software, or materials, commercial or non-commercial, are identified in this paper in order to specify the experimental procedure adequately. Such identification does not imply recommendation or endorsement of any product or service by NIST, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

Publication History

Approved by the NIST Editorial Review Board on 2026-09-24

Supersedes NIST TN 2268r1 (Jan 2025) <https://doi.org/10.6028/NIST.TN.2268r1>

How to Cite this NIST Technical Series Publication:

Holmberg D, Roth T (2026) Flexible Resource Controller: A Price and Voltage-Responsive Device Controller for GridLAB-D Co-Simulations. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Technical Note (TN) NIST TN 2268r2. <https://doi.org/10.6028/NIST.TN.2268r2>.

Author ORCID iDs

David Holmberg: 0000-0002-1738-8190

Thomas Roth: 0000-0002-9986-7784

Abstract

The Flexible Resource Controller (FRC) is a multiple-objective controller that manages residential heat pumps (HP), electric domestic water heaters (DHW), electric vehicle (EV) charging, and batteries. The first objective is to manage energy use to save customers money—charging storage and EVs when electricity prices are low and avoiding energy use (or discharging the battery) when prices are high, while maintaining thermal comfort. A second objective is to implement reactive power control at the battery inverter based on local voltage levels (Volt-VAR) to minimize voltage excursions outside of acceptable limits.

The FRC has been developed for use in transactive energy co-simulations using the GridLAB-D distribution grid simulator to study the impact of price-responsive controls on voltage and other building and grid metrics. It takes as inputs the house thermal properties, weather, and electricity price data, and generates control variables (e.g., hot water temperature setpoint, battery real and reactive power settings) that are sent to the GridLAB-D simulation engine during its execution. The FRC makes use of hourly day-ahead prices (DAP) in addition to optional adjustment of control parameters in response to 5 min real-time prices (RTP). Alternatively, the FRC can adjust control parameters based on a congestion dynamic price (CDP) generated using service transformer power flows. Finally, the FRC can work in a pass-through mode, taking power levels for each device from the transactive energy market (TEM) and passing those on to GridLAB-D.

The FRC currently implements simple patterns that shift power usage based on prices. The FRC HP controller implements precooling prior to the evening price peak and a raised house setpoint temperature during the price peak while maintaining thermal comfort. The DHW controller lowers the tank setpoint after noon to reduce heating on peak while still providing hot water. The EV is charged at the lowest priced intervals within the owner's scheduled charging sessions. The storage battery is set to charge at night when rates are lowest and then discharge during price peaks. Battery inverter Volt-VAR voltage control is implemented independently of price response, operating at all hours of the day when activated.

This document describes the control algorithms that were implemented in the Java programming language and provides an overview of a co-simulation that connects the controller to a GridLAB D distribution grid simulation.

Keywords

Congestion dynamic price; distributed energy resource; device controller; dynamic price; flexibility; price response; voltage control; Volt-VAR.

Table of Contents

1. Introduction.....	1
2. FRC Design	3
2.1. Calculating CDP	4
2.2. Transactive Energy Market Integration	5
2.3. HP Controller.....	6
2.3.1. Precooling Based on DAP	7
2.3.2. Adjustment Based on RTP	8
2.3.3. Response to CDP.....	10
2.3.4. Heat Pump Operation in TEM Pass-through Mode.....	10
2.4. DHW Control	10
2.5. Battery Control.....	11
2.5.1. Battery DAP-Responsive Charge and Discharge of Real Power	12
2.5.2. Battery RTP Adjust.....	12
2.5.3. Battery CDP Response	12
2.5.4. Battery Reactive Power Control	14
2.5.5. Battery Control in TE Market Pass-through Mode	15
2.6. EV Charge Control	15
2.6.1. EV Charging Patterns	16
2.6.2. DAP Responsive Charging.....	17
2.6.3. CDP Responsive Charging	18
2.6.4. EV Charge Control in TE Market Pass-through Mode	19
2.7. Implementation Details	19
2.7.1. Supporting Federates	21
2.7.2. Flexible Resource Controller Federate	21
3. Conclusion	24
References.....	25
Appendix A. Baseline EV Charge Profile Development.....	26

1. Introduction

The Flexible Resource Controller (FRC) was developed in the context of transactive energy (TE) simulation studies to understand the effect of different TE approaches (different markets, dynamic price tariffs, and controllers) on customer economics, customer comfort, and grid power quality. Power quality metrics include voltage volatility and acceptable voltage levels as well as power factor, losses, and transformer loadings. The FRC was developed to study the impact of price-responsive flexibility on distribution voltage and the value of voltage-responsive reactive power control (inverter Volt-VAr) relative to other voltage-control measures. The best sources of price-responsive flexibility include the largest loads in the house, electric heating and cooling equipment (e.g., electric heat pump) and electric water heater, in addition to an EV or stationary battery [1]. Each of these devices has thermal storage (i.e., building thermal mass) or electrical storage (battery) to enable load shifting.

Every house in the R4-12.47-1 test grid model (the same GridLAB-D model used in [2]) has a heat pump, electric water heater, and stationary battery, and some have an EV (represented as charge-only batteries). The capacities of the HPs and DHWs vary from house to house, as do the thermal properties of each house. The FRC manages the HP and the DHW in each house to maximize customer financial savings while staying within allowable comfort/performance boundaries. The battery is managed for energy arbitrage—charging during low price times and discharging during high price times. The purpose of the Volt-VAr control is to use inverter-based reactive power control to move voltage closer to the nominal 120 V level.

In previous experiments where devices responded to 5 min wholesale market real-time prices (RTP), the resulting power flow and voltage were volatile, due to synchronized response to price changes [2]. The FRC avoids this for the HP by adjusting the cooling load profile to precool (store energy in house thermal mass) and then raise temperature on peak; this profile is different for each house according to its thermal constants. For the DHW and battery, some randomization across houses is used to avoid synchronized response. The randomization may be similar to what occurs when different controllers are present with different control logic, responding in slightly different ways to the same price signal.

The FRC controls primarily respond to the wholesale market hourly day-ahead prices (DAP). Based on previous research [2], response to DAP allows planning for storage use (including load shifting) and avoids the larger power swings that occur when all devices react to a volatile RTP signal. On the other hand, the FRC also has logic that, if enabled, causes the controller to avoid using power during RTP price spikes, as may be applicable to a tariff where part of a customer's energy use is paid at the RTP while a baseline amount is paid at DAP. In addition, the FRC can generate and respond to a congestion dynamic price (CDP) which is adjusted based on local transformer congestion. In the case of raised CDP, devices may shift load to reduce cost to the customer.

The Volt-VAr control has been implemented to reduce the voltage range of fluctuations on the grid and reduce the count of voltage excursions outside of ANSI C84.1 Range A limits [3]. As

voltage levels rise, inverter reactive power consumption increases to lower voltage. Likewise, as voltage drops, reactive power injection increases to raise voltage.

The TEM pass-through mode enables a local market to determine prices and resulting load profile. The TEM runs separately from the FRC as a day-ahead market. The TEM is used to balance load across the day in order to minimize prices and maximize device flexibility, shifting EV and battery charging to the lowest cost intervals while still providing the power requested during the charging windows. The final result is a power profile (hourly load or generation) for each hour of the coming day, and this is passed to FRC which, in turn, feeds to profiles to GridLAB-D (GLD) each 5 minutes throughout the following 24 hours.

2. FRC Design

The FRC serves as a simple algorithmic controller that manages the HP, DHW, battery and EV charging operations based on price. The FRC receives both DAP and RTP. The FRC uses a simple approach to manage load based on price: reducing load or discharging when the price is high and charging when the price is low. The FRC also implements standard Volt-VAR control to inject or consume reactive power based on the local voltage to move voltage closer to nominal.

The FRC has a configuration file that enables or disables control of different devices so that the flexibility of that device in its different forms can be examined independently or in combination with other forms of flexibility. The operation of the FRC is shown in Fig. 1.




Fig. 1 Flexible Resource Controller architecture within UCEF along with the data elements moving between UCEF components: Experiment Manager, FRC, and GridLAB-D (GLD).

The FRC was implemented using the NIST Universal CPS Environment for Federation (UCEF) co-simulation development environment [4]. UCEF was used to generate the Experiment Manager, FRC, and GLD federates (software components) that participate in the co-simulation (Fig. 1). GLD simulates distribution grid power flows and the energy consumption of various loads in residential homes, taking control inputs from the FRC and generating power flow and house and device parameters as outputs. The FRC has been developed in Java and is hosted on GitHub [5].

As shown in Fig. 1, the FRC consumes prices, inverter Volt-VAR constants, house thermal constants and occupant comfort parameters, and configuration options (i.e., which controls to use for each device). These are used to generate house temperature cooling setpoints (T_{csp}) to drive HP operation, DHW tank setpoints to manage when the DHW heats water, and the charge/discharge cycle (P_{out}) for each battery. Additionally, the FRC reads in the local voltage

measured at each house meter at each time step to determine the reactive power setting (Q_{out}) and real power setting (P_{out}) to pass to the battery inverter for Volt-VAr control. To calculate the CDP, the FRC also reads the transformer real power flows ($power_{in}$). The specific design of each FRC component is provided below.

2.1. Calculating CDP

Figure 1 contains a “CDP Calculator” at the top of the FRC box. The FRC can use a congestion-aware CDP that is calculated based on the current grid state. The CDP Calculator takes a transformer’s real power flow and divides it by transformer capacity to get a multiplier (M_{flow}). The congestion price will increase as M_{flow} increases, and flexible devices may respond to increasing CDP by reducing consumption if possible and if configured to be price-responsive. The battery (and EV) will reduce charging if CDP rises (assuming the house occupant has some desire to save money), the HP may raise the set point if CDP rises significantly, and the DHW may reduce the temperature set point when CDP is too high.

CDP is calculated as a function of DAP and M_{flow} . The value of M_{flow} is mapped to a price multiplier, M_{price} , to arrive at this CDP definition:

$$CDP = M_{price} * DAP \quad (1)$$

Where:

- M_{price} is as defined in Eqn. 2,
- $M_{flow} = P_{trans,avg} / Capacity_{trans}$,
- $P_{trans,avg}$ = real power flowing through the transformer (average over past n_{avg} min),
- $Capacity_{trans}$ = transformer rated power flow capacity.

M_{price} increases as a function of transformer loading from $1.0 * DAP$ up to $5.0 * DAP$ as M_{flow} increases above 0.75. The price curve used is shown in Fig. 2. This exponential increase makes the use of electricity by any house behind the transformer more and more expensive as the demand (M_{flow}) increases above the transformer capacity. In real life, a transformer can operate temporarily at elevated flow rates above rated capacity. There is no existing public utility commission (PUC)-approved tariff using this approach, however, there have been two California pilots (RATES and CALFUSE) that similarly used load congestion mapped to a quadratic cost recovery curve, raising the price as the load increases [6].




Fig. 2: Transformer congestion pricing curve showing the influence of power flow through the transformer, relative to transformer capacity, on the resulting congestion dynamic price.

The curve in Fig. 2 can be represented by:

$$M_{price} = 0 \text{ for } M_{flow} < -1.25 \quad (2)$$

$$M_{price} = 2.5 - 2 M_{flow} \text{ for } -1.25 \leq M_{flow} < -0.75$$

$$M_{price} = 1.0 \text{ for } -0.75 \leq M_{flow} \leq 0.75$$

$$M_{price} = M_{flow}^{3.074} / 2 + 0.7935 \text{ for } 0.75 < M_{flow} \leq 2$$

$$M_{price} = 5 \text{ for } M_{flow} > 2$$

2.2. Transactive Energy Market Integration

The TE Market (TEM) [7] is a key federate for transactive energy co-simulations, working in tandem with the FRC. This section gives a short introduction of how it works with the FRC.

The TEM provides an hourly market with DAP as the initial market price. Just like CDP, the TEM market agent tracks congestion in the local transformer, and the market price will rise as load increases for subsequent power purchase transactions for any given hourly interval.

The TEM brings CDP into a forward market framework. Instead of reacting to a real-time price (5 min CDP) where the only action is to back off and delay use until later (if flexibility is available), the market allows surveying electricity price over the next 24 hours. The battery and EV can see where the lowest price is and plan to charge then. Other devices may be able to plan to shift some load to lower priced intervals.




Fig. 3 TE Market and FRC internal components and data values communicated among federates. GLD is the GridLAB-D grid simulator.

The market works with separate user agents (TEUA, Fig. 3) for the EV, battery, and house loads. Each TEUA bids in the market for the power that it needs (or expects to need). There is no adjustment, as initially implemented, based on customer comfort criteria (λ , see Table 1). There is also no RTP adjustment (2.3.2).

The TEM runs once per day in the afternoon to settle the power commitments for the next day (the market covers 24 hourly intervals starting at 5:00 p.m.). At market close (after multiple rounds), each TEUA has bought (or sold) the needed amount of power resulting in a power profile for the coming 24 hours. As shown in Fig. 3, the TEM optimized power profiles are then passed to GLD. Effectively, this replaces the FRC algorithms for EV and battery scheduling that are discussed here in the FRC report. For initial TEM implementations, the TEM does not adjust the HP and DHW power profiles (no power shifting away from the request based on DAP profile). Thus, FRC only passes on the same temperature setpoints to GLD as for the DAP control profiles.

2.3. HP Controller

The primary goal of the FRC HP controller is to provide cost savings to the consumer within thermal comfort requirements by avoiding running the HP during high-price times. There are two components: precooling based on DAP and adjustment based on either the RTP or CDP. The FRC logic initially only manages the summer cooling season. Control variables are given in Table 1.

Table 1 Heat Pump Controller Variables

Variable Name	Description
λ	Customer comfort parameter, between 0 and 1; 0 for low comfort requirements and 1 for max comfort. This abstract parameter serves for temperature setpoint adjustment as well as willingness to reduce EV charging with high price in equations in following sections.
H_{pp}	Number of hours of precool preceding start of peak period (width of precool period, in hours)
P_{peakda}	Peak price for that day, using DAP (\$)
P_{rtp}	RTP 5 min price (\$)
R_{rtp}	P_{rtp}/P_{peakda} , ratio of current RTP to DAP peak price
R_{cdp}	CDP/P_{peakda} , ratio of current CDP to DAP peak price
τ	House thermal mass time constant (h)
T_{csp}	Final house cooling setpoint temperature sent to GLD, including any adjustment based on RTP (F) or CDP (F)
$T_{max\lambda}$	$T_{set} + offset_limit \times (1-\lambda)$, where $offset_limit$ is a maximum house temperature offset in response to price and λ .
$T_{min\lambda}$	$T_{set} - offset_limit \times (1-\lambda)$
t_{peak_ctr}	Center of the DAP peak price hour (time)
t_{peak_start}	Start time for the peak period (time)
t_{peak_wid}	Width of the peak time period (h)
$t_{precool_start}$	Start time for the precool
T_{set}	Baseline setpoint temp for each house before price adjustment (F)
T_{set_dap}	Temperature before RTP Adjust (could be T_{set} , T_{set_pre} , or T_{set_peak})
T_{set_peak}	Setpoint temp during peak time period. When responding to DAP, this temperature will be set equal to $T_{max\lambda}$ (F).
T_{set_pre}	Setpoint temp during precooling (F)

2.3.1. Precooling Based on DAP

The air conditioning in a typical house can comfortably be off for 1 to 3 hours¹ while house temperatures rise during a peak price period, depending on house thermal parameters and customer comfort parameters. The temperature setpoint during this peak price period (T_{peak_wid}) can be set to some $T_{max\lambda}$ adjusted for customer comfort requirements. Likewise, before this peak period there can be a precool time to lower the house temperature (thermal mass temperature) that enables a larger and longer temperature rise during the peak period. The setpoint during precooling can be set to $T_{min\lambda}$. The FRC determines the length of a peak window and a precooling window. The peak $T_{max\lambda}$ window is centered around the DAP peak price hour. Window sizes are determined based on house thermal constants. Determining T_{csp} follows this progression.

¹ Based on simulation results tracking time for temperature to rise from T_{min} to T_{max} for different houses and average λ values.

1. Find t_{peak_ctr} . For the Tucson July data, the peak price is in the 7 p.m. to 8 p.m. hour and t_{peak_ctr} is 7:30 p.m.
2. Find t_{peak_wid} . This width varies with the thermal time constant of the home, just as the precool period varies.
 - a. $t_{peak_wid} = 0.75 H_{pp}$, and
 - b. $H_{pp} = 2 + \tau/6$ (see discussion below)
3. Calculate $t_{peak_start} = t_{peak_ctr} - t_{peak_wid}/2$. For the Tucson July data, t_{peak_start} will vary from 5:00 p.m. to 6:30 p.m. approximately.
4. Calculate the precool period start time such that the precooling period ends at t_{peak_start}
 - a. $t_{precool_start} = t_{peak_start} - H_{pp}$
5. Set temperature during the precooling
 - a. $T_{csp} = T_{set_dap} = T_{set_pre} = T_{min\lambda}$
6. Set temperature during the peak time period
 - a. $T_{csp} = T_{set_dap} = T_{set_peak} = T_{max\lambda}$
7. Set temperature outside precool and peak periods at
 - a. $T_{csp} = T_{set_dap} = T_{set}$

The FRC experiments on the R4-12.57-1 grid [2] used Tucson weather and CA Independent System Operator (CAISO) prices from June 23 to July 7, 2017. The DAP peak is similar across all days with the peak price in the 7 p.m. to 8 p.m. hour and the highest prices from 6 p.m. to 9 p.m. Only the magnitude of the peak changes. The ideal H_{pp} (precooling time) allows the house to cool to $T_{min\lambda}$, and this amount of time varies with the thermal time constant of the house. GLD stores a thermal time constant, τ^2 , for each house. The R4-12.47-1 grid houses' average τ is around 14 h, with a minimum of 4 h and a maximum of 29 h. There is a trade-off between increased energy use to maintain T_{set_pre} at $T_{min\lambda}$ and the amount of cooling of thermal mass that is achieved. The first few hours of cooling are the most valuable. The value of H_{pp} as a function of τ was set to $H_{pp} = 2 + \tau/6$. This relationship results in H_{pp} values from about 3 h for low τ ($\tau \approx 6$ h) to H_{pp} values of around 6 h for higher values of τ ($\tau \approx 24$ h). The varied peak width and precooling times also serve to stagger the timing of heat pumps shutting off and turning on to minimize shock on the power flow.

2.3.2. Adjustment Based on RTP

The availability of day-ahead prices allows effective use of storage, both thermal and electrical, to optimize load shedding during peak price times. At the same time, a tariff may charge the customer for power consumption based in part on RTP. In this case, it will be to the customer's advantage to avoid using power when prices are high and particularly during any price spike. For the CAISO RTP, the RTP signal is a more volatile version of the DAP, and the peak in the RTP signal aligns generally with the DAP peak, Fig. 3. However, RTP has price spikes. The FRC's RTP Adjust component raises the HP setpoint temperature during these price spikes.

² τ represents the time required for the thermal mass temperature of the house to change to a new value when T_{csp} is changed, specifically the time to reach 63 % of the temp change from the old to the new value.




Fig. 4 Real-time and day-ahead prices for test days in [2]. The FRC controls storage based on DAP peaks and valleys, but also allows response to price spikes as seen on July 7.

This approach assumes that it is in the financial interest of the customer to minimize power use during the DAP peak period. If the RTP Adjust algorithm responds to RTP at every instance, then it will no longer be responding to DAP. For this reason, the RTP Adjust algorithm does not change the temperature setpoint unless RTP is greater than the peak DAP, that is, T_{csp} is raised only when $Prtp > Ppeakda$ ($Rrtp > 1$). A ramp in T_{csp} is implemented from the existing T_{csp} up to a maximum of $Tmax\lambda + Tbump$, where $Tbump$ represents a small additional temperature increase (set to 1 °F (0.5 °C) by default³), as follows.

- If $Rrtp \leq 1$, then $T_{csp} = Tset-dap$
- If $1 < Rrtp < 2$, then $T_{csp} = Tset-dap + (Rrtp - 1) (Tmax\lambda - Tset-dap + Tbump)$
- If $Rrtp \geq 2$, then $T_{csp} = Tmax\lambda + Tbump$

The decision to add $Tbump$ above $Tmax\lambda$ is arbitrary, but the assumption is that people will not want to run their air conditioner during a price spike and would be willing to slightly raise the temperature to avoid the extra marginal cost. The customer still has the option not to enable RTP Adjust.

³ Note that GridLAB-D uses Fahrenheit units.

2.3.3. Response to CDP

The operation of the HP in response to CDP is the same as for RTP, following the equations in 2.3.2, except that R_{rtp} is replaced with R_{cdp} . Therefore, as congestion increases and the CDP rises above the DAP peak price, then the HP T_{csp} will start to rise as well.

Note that if CDP response is enabled, then there is no independent RTP Adjust. RTP Adjust only applies to the DAP control.

2.3.4. Heat Pump Operation in TEM Pass-through Mode

The TEM House TEUA (the user agent representing the HP to buy power in the market) seeks to buy the power required for the HP using the FRC DAP control-generated power profile which includes precooling. Specifically, the FRC generates DAP temperature setpoints for HP that are input to GLD and the house power consumption that results (output from GLD) becomes the “house load profile” input to TEM from the “House load forecaster” in Fig. 3.

The TEM is used only to purchase the house load profile, with no time shifting. Thus, during TEM co-simulation, the HP and DHW temperature setpoints that FRC sends to GLD are identical to the DAP profiles unchanged by TEM. Likewise, plug loads and photovoltaic (PV) generation are unchanged. The energy consumption of all of the House TEUA components (PV, HP, DHW, plug) are determined by GLD internal models with the only input from FRC being the HP and DHW temperature set points. More information about the operation of the TEM is available in [7].

2.4. DHW Control

The water heaters in the GLD model have a distribution of sizes, water consumption per day, water temperature, and times of water consumption. In aggregate, there is a morning peak and an afternoon peak in energy consumption. At the same time, the price peaks in the late afternoon. In general, observing the amount of water in the tanks and the amount of water drawn from these tanks in the afternoon, most customers would not need to heat water at all in the afternoon if the tank was fully heated after the morning peak. That is, a fully heated tank would serve the afternoon hot water demand of most customers without reheating before midnight.

Based on this information, a simple control strategy is used for the DAP, CDP and TEM approaches. The tank temperature is set to T_{max} in the morning, and to T_{min} in the afternoon. More specifically,

- $T_{setpoint} = T_{max}$ from 2:00 a.m. to noon, and
- $T_{setpoint} = T_{min}$ from noon to 2:00 a.m.

The R4-12.47-1 grid model uses the “MULTILAYER” water heater from GridLAB-D with both “lower_tank_setpoint” and “upper_tank_setpoint”. For the FRC implementation, these values were set equal to each other and to the tank_setpoint, T_{set} . The grid model DHW setpoints are

uniformly distributed between 125 °F (51.7 °C) and 135 °F (57.2 °C). The setpoints are adjusted as follows.

- $T_{max} = T_{set} + T_{max-offset}$
- $T_{min} = T_{set} - (1-\lambda) * T_{min-offset}$

$T_{max-offset}$ default is set to 5 °F (2.8 °C). $T_{min-offset}$ default is set to 20 °F (11.1 °C). The 2:00 a.m. restart time for each house is delayed by a random 0 min to 120 min to avoid all water heaters turning on at the same time.

RTP Adjust for the DHW works as follows. A large draw of water at any time of day will reduce the DHW tank temperature and generally cause the heating element to turn on. To reduce the likelihood of the heating element operating during a price spike, the DHW T_{set} is reduced to the GLD minimum allowed value, $T_{min-gld} = 90$ °F (32 °C).

- If $Prtp > 2 * P_{peakda}$, then $T_{set} = T_{min-gld}$

CDP response is identical, replacing RTP with CDP.

- If $CDP > 2 * P_{peakda}$, then $T_{set} = T_{min-gld}$

2.5. Battery Control

All batteries are the same in the R4-12.47-1 grid model, as implemented by the Pacific Northwest National Laboratory (PNNL) feeder generator script 2022 [8] and are based on a commercial battery with 13.5 kWh capacity, max 5 kW charge and discharge rate, with 86 % round trip efficiency.

The FRC battery control uses available GLD battery modes to enable both Volt-VAr (managing reactive power for voltage control) as well as charge and discharge of real power based on price. Since the goal is to manage the battery based on both price and voltage, the battery inverters use the GLD CONSTANT_PQ mode. In this mode, both the real power (P_{out}) and reactive power (Q_{out}) values are set in GLD by the FRC battery controller. P_{out} is used to indicate charge (negative value), discharge (positive value), or no power flow. Q_{out} directs the amount of reactive power to inject or consume based on the local voltage. For this action, the FRC battery controller reads the GLD house meter's real voltage output and determines Q_{out} based on the Volt-VAr settings. If Q_{out} is too high, such that the apparent power, $SQRT(P_{out}^2 + Q_{out}^2)$, is greater than 5 kW, then P_{out} is reduced to keep apparent power at or below 5 kW. The resulting power values are passed back into GLD to direct the battery inverter.

The actual charge and discharge of the battery may be different than expected (i.e., realized P_{out} and Q_{out} may differ from FRC control values) for a variety of reasons. For example, the battery may be fully charged, so charging drops to zero even when the FRC battery controller is commanding continued charging because the GLD battery model will not allow overcharging. Similarly, discharge may stop when the state of charge drops to a minimum.

2.5.1. Battery DAP-Responsive Charge and Discharge of Real Power

DAP-responsive control charges the battery during the low-price period at night and discharges during the peak-price period. The nighttime charge, P_{charge} , is activated during the low-price period from (2:00 to 8:00) a.m. Charging follows a 4.5 h triangular (ramp up, ramp down) pattern. The ramps focus discharge at the peak price and avoid a step change in load that would affect voltage. The charge pattern is as follows:

- P_{charge} ramps up from 0 kW to 4.8 kW over 0.5 hours then ramps back down to 0 kW over 4 hours with linear ramps.
- Start time is randomized for each house, uniformly distributed from (1:00 to 3:00) a.m.
- $P_{out} = -P_{charge}$

The late afternoon discharge tracks the DAP peak hours.

- $P_{discharge}$ ramps up for 3 h from 0 kW to 3.6 kW peak at the middle of the peak hour, then back down to 0 kW over the last three hours.
- $P_{out} = P_{discharge}$

For the Tucson CAISO data, this algorithm results in ramp up from (4:30 - 7:30) p.m. and ramp down from (7:30 - 10:30) p.m., with a total discharge of 10.8 kWh. The SOC is ignored because the GLD battery model will stop charging/discharging when full/empty.

2.5.2. Battery RTP Adjust

The RTP Adjust can be enabled independently of the DAP response. The battery will change state to discharge mode when the price spikes. This change will save money for the owner but have a sudden impact on power flow resulting in voltage disturbance. The discharge power, P_{out} , is unchanged (set according to the charge/discharge plan based on DAP) until RTP increases above the P_{peakda} price. The discharge power flow will then start to increase up to full discharge power ($P_{max} = 5$ kW) when RTP is at or above twice the peak DAP. The discharge logic is as follows.

- $P_{out_rtpadj} = P_{out}$, for $Prtp/P_{peakda} \leq 1$
- $P_{out_rtpadj} = P_{out} + (Prtp/P_{peakda} - 1)(P_{max} - P_{out})$, for $1 < Prtp/P_{peakda} < 2$
- $P_{out_rtpadj} = P_{max}$, for $Prtp/P_{peakda} \geq 2$

2.5.3. Battery CDP Response

The CDP implementation for battery charging starts by following the same ramp pattern used for DAP response, including the variable start times. However, as competition for power (indicated by rising congestion) starts, or the absolute price of power (indicated by CDP) gets too high, then the battery charging will start to reduce.

Charge rate reductions are dependent on both (a) M_{price} as a measure of congestion, and (b) the comparison of CDP to the DAP peak price.

- (a) If there is competition for power with a rising CDP, then it is desirable to have the battery at a lower priority than other loads. This is accomplished by backing off on charging as soon as M_{price} is greater than 1.0. The battery is the most flexible asset and can charge outside the EV charge windows (2.6.1) and when other loads are not operating.
- (b) If CDP is too high relative to the peak DAP (where the battery power is sold), then it makes no sense to charge. This presumes that the battery is benefiting the homeowner by buying low-priced off-peak power and selling during the high-priced peak. There is some acceptable price difference that accounts for round trip efficiency losses, capital cost repayment, and operation and maintenance costs, while providing some profit.

While the DAP response has a set ramp up and ramp down at low-price times, the CDP controller also monitors congestion. Once M_{price} rises or CDP is too high relative to peak DAP, then the charge rate will be reduced and the charge pattern will no longer follow the DAP response pattern. Once M_{price} and CDP are again within acceptable limits, the controller will attempt to raise the charge rate up to P_{max} (5 kW) and hold that until the battery is full. The battery charging algorithm is as follows.

1. While the battery is charging, check every n_{check} minutes whether $M_{price} > 1$.
 - a. If true, then back off the charging rate according to

$$P_{charge,t} = (1 - b)P_{charge,t-1},$$

$$b = n_{backoff} \left(1 - \left(\frac{1}{M_{price}} \right) \right),$$

Where b is the backoff coefficient and $n_{backoff}$ allows adjustment to the rate of backoff on charge rate. For initial experiments, the values of $n_{avg} = n_{check} = 5$ were used such that M_{price} is evaluated every 5 minutes based on an average of the past 5 minutes of M_{flow} values. The optimal value of $n_{backoff}$ was found to be in the range of 0.5 to 1.0. P_{charge} will continue to decrease as long as M_{price} is above 1.0, and the higher M_{price} is, the faster P_{charge} is reduced.

- b. If false, and if $P_{charge,t-1} < P_{charge,t-2}$, then hold steady ($P_{charge,t} = P_{charge,t-1}$). This instruction forces a pause between reducing P_{charge} and subsequently raising it, a one-step delay to avoid oscillatory behavior.
 - c. If false (after being true), and if $P_{charge,t-1} < 5$ kW, and $P_{charge,t-1} \geq P_{charge,t-2}$, then:

$$P_{charge,t} = \min(P_{charge,t-1} + n_{check} * pmi, P_{max}),$$

where pmi is the desired “per minute increase” charge rate increase. The existing battery and EV ramp rates are in the range of $pmi = 0.04$ up to 0.1 kW.

2. While the battery is charging, check the price differential, $P_{dif} = P_{peakda} - CDP$, every n_{check} minutes. The charge rate will be fully curtailed when CDP rises above $P_{peakda} - P_{dif-econ}$, where $P_{dif-econ}$ equals the price difference representing an acceptable return on investment for this

battery buying low and selling high. Additionally, the battery charge rate (P_{charge}) will be reduced across a small range, $P_{dif-range}$ below the cut-off price. The CDP check takes the $P_{charge,t}$ above as input and adjusts it according to:

$$\begin{aligned}
 P_{charge,t,adj} &= P_{charge,t} && \text{for } P_{dif} \geq P_{dif-econ} + P_{dif-range} \\
 P_{charge,t,adj} &= P_{charge,t} \left((P_{dif} - P_{dif-econ}) / P_{dif-range} \right) && \text{for } P_{dif-econ} + P_{dif-range} > P_{dif} > P_{dif-econ} \\
 P_{charge,t,adj} &= 0 && \text{for } P_{dif} \leq P_{dif-econ}
 \end{aligned}$$

3. Continue charging until full or until noon, whichever comes first.

The battery discharging algorithm is the same as for DAP, following the same ramp pattern during the evening peak DAP hours. It is theoretically possible to have such a large power flow out to the grid that M_{price} starts to reduce below 1.0 (see left side of Fig. 2), and CDP would be reduced. This result would reduce income and call into question the amount of battery storage installed behind the transformer. No CDP adjustment algorithm has been implemented for the battery discharge.

It is not desirable to charge the battery if CDP is too high, that is, when the price spikes up due to congestion during the expected “low price” charging window. The battery algorithm reduces charging as CDP increases according to the following equations (similar to 2.4.2 RTP Adjust).

- $P_{out_cdp} = P_{out}$, for $CDP/P_{peakda} \leq 1$
- $P_{out_cdp} = P_{out} + (CDP/P_{peakda} - 1)(P_{max} - P_{out})$, for $1 < CDP/P_{peakda} < 2$
- $P_{out_cdp} = P_{max}$, for $CDP/P_{peakda} \geq 2$

2.5.4. Battery Reactive Power Control

The reactive power control (Volt-VAr) is a separate option that can be toggled on or off independent of price-responsive real power control. Volt-VAr control adjusts the phase angle of voltage and current at the inverter which has the effect to change the measured voltage. The Volt-VAr dimensionless settings for IEEE 1547 Cat.B [9] (high-DER) default are as follows: $Q_{set}=0.44$, $V_{min}=0.92$, $V_{lo}=0.98$, $V_{hi}=1.02$, $V_{max}=1.08$. The IEEE standard Q_{set} value of 0.44 ensures that real power generation from photovoltaic inverters is not overly impacted by reactive power provision. $Q_{set}=0.44$ corresponds to power factor near 0.9. The control logic is as follows.

1. Read the house meter voltage, V , from GLD
 - a. $V_{pu} = V/V_{nom}$, with $V_{nom} = 120$ V for triplex house meters
2. Set Q_{out} based on voltage, where $Q_{out} = Q_{pu} \times Q_{max}$ and $Q_{max} = 5000$ VAR
 - a. If $V_{lo} \leq V_{pu} \leq V_{hi}$, $Q_{pu} = 0$
 - b. If $V_{min} \leq V_{pu} < V_{lo}$, then $Q_{pu} = Q_{set} \left(1 - \frac{V_{pu} - V_{min}}{V_{lo} - V_{min}} \right)$
 - c. If $V_{hi} < V_{pu} \leq V_{max}$, then $Q_{pu} = -Q_{set} \left(1 - \frac{V_{max} - V_{pu}}{V_{max} - V_{hi}} \right)$

- d. If $V_{pu} < V_{min}$, then $Q_{pu} = Q_{set}$
- e. If $V_{pu} > V_{max}$, then $Q_{pu} = -Q_{set}$
3. Check that $\sqrt{P_{out}^2 + Q_{out}^2} \leq Q_{max}$
4. If $> Q_{max}$, then reduce P_{out} according to $P_{out} = \sqrt{Q_{max}^2 - Q_{out}^2} \times \text{SGN}(P_{out})$

2.5.5. Battery Control in TE Market Pass-through Mode

When co-simulations use the TEM, the battery controller in FRC is not used. The TEM battery TEUA searches for the lowest priced intervals after the evening peak to charge and buys power during those intervals in the market. For discharge, the battery TEUA looks for the highest priced intervals and sells power to the market in those intervals. The TEM battery control approach is explained in detail in [7]. The TEM does not implement any reactive power control, nor use the RTP Adjust.

2.6. EV Charge Control

EVs were not included in the original PNNL grid model [10], so they were added to the GLD model, combining a new EV meter with an inverter and battery. The EV meter (“mev”) is placed alongside the PV (“msol”), house (“mhse”), and battery (“mbat”) meters that together sit behind the house main meter (“mtr”), Fig. 4.




Fig. 5 GridLAB-D model power flow chart from substation transformer to service transformer 123 to house main meter and then 4 sub-meters below that.

In the current simulations, EV batteries are charged by the grid and do not discharge back to the grid. The EV battery was sized large enough to absorb 2 days of charging to enable a 2-day simulation experiment with no need to model discharge during driving. A longer duration simulation would need logic to reset the state of charge (SOC) each day based on expected driving patterns and energy use on the road.

2.6.1. EV Charging Patterns

Baseline charging amounts and times were roughly modeled after the Austin, TX Pecan Street EV charging data [11], for which there was a flat price tariff. The charging times vary, with approximately 60 % of EV charging in the evenings (5:00 – 11:00) p.m., 25 % during the day (7:00 a.m. to 5:00 p.m.) and 15 % at night. The typical Pecan St. EV charge is less than 10 kWh daily. The distribution of the charge amount (Pev) is approximately modeled with a log-normal curve, with charge amount distribution per Fig. 6. One can see that most charges are small, but there is a long tail where some EV charge amounts are greater than 20 kWh. In other published sources, daily driving is greater (especially in larger cities with longer commutes, such as is the case in the Washington, DC, area where NIST is located [12]) and the resulting Pev may be higher. For this reason, the average EV charge amount was raised 50 % from 6 kWh to 9 kWh⁴.




Fig. 6 Austin Pecan St. EV charge amount distribution and log-normal curve fit used in the FRC.

The code implementation first determines Pev as a random value taken from the inverse log-normal fit to the Pecan St. data with a multiplier to increase the charge per the above discussion. Second, the charge window, whether in the day, evening, or night, is determined. Third, a specific start time is generated. The charging pattern is then constructed with a defined ramp up, level region (with maximum charge rate $P_{max-ev} = 7.2$ kW), and ramp down.

The actual details of the ramp up, duration and magnitude of a flat charging period (if any), and the ramp down (if any), depend on different conditions, including time window, Pev, start time, and price signal. Each time window (day, evening, night) has a constrained amount of time that can be used for charging. The daytime window was set from 10:00 a.m. to 5:00 p.m. The evening window was set from 5:00 p.m. to 11:00 p.m., although the charge was allowed to extend into the night for large Pev. The night window was 11:00 p.m. to 7:00 a.m. The detailed baseline charge profile development path used in the FRC code is presented in Appendix A.

The configuration file allows a seed value to be used as input to the randomization, and the same randomization value is used to generate the charge amount, charge window, and start time. This approach allows for a repeatable set of charge patterns that can be used to compare between baseline, DAP response, CDP response, and any TE market designs.

⁴ 9 kWh translates to 30 – 40 miles or 50 – 65 km, depending on the efficiency of the EV.

The baseline usage, as with the Pecan St. data, is what may be expected when customers have a flat tariff with uniform prices across the day. Customers typically return home and plug in their EV to recharge in the evening. The FRC is focused on at-home charging and its impact on grid power flows and does not consider workplace or public charging.

2.6.2. DAP Responsive Charging

The EV DAP responsive charging makes two changes to the baseline. First, evening charging is pushed later to move charging to the lower night prices. Second, all charging is spread out across the day or night charge windows to avoid overloading the transformer. All day charge events are shifted up to start at 9:00 a.m., while every evening and night charge event is moved to start at 10:00 p.m. The day charge is spread out over the window of 9:00 a.m. to 4:00 p.m., while the night charges are spread out over the 10:00 p.m. to 7:00 a.m. window, Fig. 7. The day charge pattern has a one-hour ramp up to a flat charge interval followed by a one-hour ramp down, and the charge rate during the flat interval equals $P_{ev}/6$. The night charge has a three-hour ramp up to a flat charge rate interval with a one-hour ramp down. The charge rate during the flat interval equals $P_{ev}/7$.

Since the day charge window is only seven hours long, and the maximum Level 2 charge rate is 7.2 kW, the limit to the total possible charge is 43.2 kWh. Thus, P_{ev} is truncated to this value if the originally calculated charge quantity is greater. In a similar way, the longer night charge window has a maximum quantity equal to 50.4 kWh.




Fig. 7 DAP responsive EV charging profiles, day (top) and night (bottom).

This approach to flattening out the charge profiles reduces overloading on the transformers. On the other hand, it does not necessarily eliminate overloading and may not match the needs of individuals who require the charging to be done earlier. Likewise, it doesn't guarantee the lowest price, although it avoids the highest price times. A lower cost option for the consumer, if charging during the day, would be to charge at full power in order to finish earlier, since DAP prices are generally rising from morning lows up to afternoon peaks. Also, given that the DAP price profile can change with location and season, it would be better to have a more intelligent algorithm that examines the DAP profile for the next day, compares that to the requested power, and finds the lowest-priced time window during the day or night that meets the charging request. These enhancements have not been implemented.

2.6.3. CDP Responsive Charging

The EV CDP charging follows the DAP response charge profile until M_{price} rises above the threshold for a given house, after which the EV charging rate starts to reduce. When M_{price} reduces to within acceptable range, then charging will begin to ramp up to P_{max-ev} . There is no final down ramp and charging continues until the full P_{ev} charge quantity is added to the EV battery or until the end of the charging window (7:00 a.m. for the night charge; the day charge window is extended to 5:00 p.m. to enable a longer charge if needed). Note that this algorithm could (but does not) limit charging based on CDP, given some customer-defined max price.

The EV charging algorithm is as follows.

1. While the EV battery is charging, check every n_{check} minutes whether $M_{price} > 1 + \lambda/2$.
 - a. If true, then back off charging rate according to

$$P_{charge,t} = (1 - b)P_{charge,t-1},$$

$$b = n_{backoff} \left(1 - \left(\frac{1 + \frac{\lambda}{2}}{M_{price}} \right) \right),$$

P_{charge} will continue to decrease as long as M_{price} is above the threshold value, and the higher M_{price} is, the faster P_{charge} is reduced. For $n_{backoff} = 1$ and $M_{price} = 1.25$, the $\lambda = 0$ house will back off to 80 % of the previous P_{charge} rate, but the $\lambda = 1$ house will not start to back off until $M_{price} > 1.5$. At $M_{price} = 2$, the $\lambda = 0$ house will back off to 50 % of the previous P_{charge} (that is, charging power is reduced to 50 % of the value at the last adjustment) while the $\lambda = 1$ house backs off to 75 %. This can result in the low- λ houses reducing charging close to zero before the high- λ houses start to reduce at all, and thus an undersized transformer may force a customer to raise their comfort setting or face frustration when they can't get their EV charged.

- b. If false, and if $P_{charge,t-1} < P_{charge,t-2}$, then hold steady.
 - c. If false, and if $P_{charge,t-1} < P_{max-ev}$, and $P_{charge,t-1} \geq P_{charge,t-2}$, then:

$$P_{charge,t} = \min(P_{charge,t-1} + n_{check} * pmi, P_{max-ev}).$$

2. Continue charging until Pev kWh has been added to the battery charge or until the end of the charging window, whichever comes first.

Note that RTP Adjust functionality has not been implemented for the EV controller.

2.6.4. EV Charge Control in TE Market Pass-through Mode

When co-simulations use the TEM, the FRC is still used to generate the charge profiles. Only the charging window (day, evening or night) and the requested charge amount are passed to the TEM for each EV. As done for DAP control, any baseline EV charge requests in the evening charge window are pushed later. The EV TEUA searches for the lowest priced intervals in the day or night charge windows and buys power in those market intervals. This TEM EV TEUA operation is explained in detail in [7].

2.7. Implementation Details

As was shown in Fig. 3, the system has four main components: a GridLAB-D distribution grid simulation that models (among other structures) residential houses and their energy resources, the FRC that controls these resources as described above, a transactive energy market that provides current energy prices that could influence the controller behavior, and an experiment manager to handle configuration and provide historic price data.

GridLAB-D was run in its server mode, which allows external processes to access and modify model properties during the simulation execution. The FRC was implemented as a separate process using server mode to control the GridLAB-D model. For scalability, the FRC controls multiple houses where each house has different control parameters, and one instance of the FRC controlled all the houses in GridLAB-D. The communication between the FRC and GridLAB-D was handled by a co-simulation standard called the High-Level Architecture (HLA [13]).




Fig. 8 Model of the High-Level Architecture (HLA) federation for the FRC modeled in WebGME, including messages communicated between federates.

Fig. 8 shows the different processes (called federates) and their data exchanges in the Web-based Generic Modeling Environment (WebGME [14]). The federates (called a federation) include the GridLAB-D *GridModel*, the *FlexibleResourceController* implemented in Java, the transactive energy market *TEM*, and two support federates called *SimulationTime* and *PriceReader* that represent the experiment manager.

All the federates have a shared representation of scenario time⁵, which they advance in lockstep with a configurable 60 s step size. This arrangement means that, from the perspective of scenario time, the FRC sends and receives data to its controlled devices every minute and does not interact at the timescale of seconds. The step size can be configured and was selected

⁵ The federation simulates a range of dates based on the available price and weather data (for instance, it could simulate from July 5, 2017 at 17:00 to July 7, 2017 at 17:00). The scenario time refers to the current time being simulated within this range.

for improved performance given that price data is provided at the 5 min scale and the federation does not model the grid transient response. Both the data exchange between federates, and the synchronization of their execution, are managed by the HLA standard.

2.7.1. Supporting Federates

The *SimulationTime* federate specifies the scenario time reference using the *SimTime* message:

- *timeScale*: time step size between synchronization points (in floating-point seconds)
- *timeZone*: default time zone for the federation represented in Olson format
- *timeZonePosix*: default time zone for the federation represented in POSIX format
- *unixTimeStart*: a unix time stamp that specifies experiment scenario start time
- *unixStopTime*: a unix time stamp that specifies experiment scenario end time

The *PriceReader* federate provides price data to the federation. The price data is configured using two CSV files, one for real-time price and one for day-ahead price, where each row contains an ISO 8601 timestamp (including date, time, and UTC offset) paired with the price in \$/kWh. Because the federates only interact at multiples of the *timeScale* specified in the *SimTime* message, the time scale should be selected based on the price data to ensure the federation synchronizes as new price data is available. Real-time prices are provided each time step when new price data is available from the CSV file. Day-ahead (hourly) prices are provided for the next day at midnight. In addition, during initialization of the federation, the *PriceReader* federate also provides the day-ahead prices for the day represented by *unixTimeStart*.

The *TEM* federate is an abstraction of multiple federates that comprise the market described in detail in [7]. Each controllable resource participates as a party in the market to buy and sell power for the next 24 hours, resulting in a market position for each resource containing their market transactions for each hour. The *TEM* federate provides this market position using the *Commitment* message with these components:

- *partyId*: unique identifier for the resource (house, battery, or electric vehicle)
- *interval*: market interval (hour of the day) described by the message
- *totalCost*: total transaction costs over the interval (negative values are profit)
- *totalQuantity*: total power transacted for the interval (negative values are generation)

2.7.2. Flexible Resource Controller Federate

The FRC federate implements the heat pump control, water heater control, battery control, and electric vehicle charging methods described in the previous sections. Controllers for each resource can be enabled or disabled through a configuration file, and the real and reactive power controls for batteries can be enabled or disabled independently of each other. In addition, the real-time price adjustment can be enabled or disabled per controller. The congestion dynamic price, if enabled, is global (either all controllers will use CDP, or none will use CDP) and takes precedence over the RTP adjustment (if CDP is enabled, RTP is ignored). Likewise, using the transactive energy market is global and takes precedence over the normal

battery and electric vehicle controls. Because the FRC federate controls all houses in the GridLAB-D simulation together, a CSV file configures the individual house parameters. These parameters include:

- house: the unique name of the house in GridLAB-D,
- Tcsp: the cooling setpoint for the heat pump controller,
- offset_limit: the maximum allowed deviation from cooling setpoint for the heat pump,
- tank_setpoint: the setpoint for the water heater controller,
- lambda: the comfort parameter for the house,
- tau: the thermal time constant for the house,
- minute_delay: a random (0 to 120) min time offset to stagger control changes, and
- has_electric_vehicle: a Boolean value to indicate whether EV charging is applicable.

The heat pump control, each day, computes the times for precooling start, peak start, and peak end for each house based on its thermal time constant. The controller also computes three different setpoints for normal operation, precooling, and peak operation based on the house comfort parameter and offset limit. At each time step, the current scenario time is compared against these time values to determine the state (normal, precooling, peak) of each house. The corresponding heat pump setpoint is selected, and a congestion dynamic price or real-time price adjustment is applied if enabled.

The water heater control, each day, computes a time interval for each house when the water heater operates at a higher setpoint. The default time range for all houses is from 2 a.m. to 12 p.m. The minute delay value configured for each house is added to the start time of this interval to stagger the water heater operation to prevent a sudden spike in power consumption on the grid at 2 a.m. This results in individual intervals for each house that start in the range from 2 a.m. to 4 a.m. and end at 12 p.m. During this interval, the water heater setpoint operates at the specified tank setpoint plus 5 °F (2.8 °C). Outside of this interval, it operates at a setpoint calculated based on the house comfort parameter. If a congestion dynamic price or real-time price adjustment is necessary, regardless of the interval, the setpoint is set to 90 °F (32.2 °C)⁶.

The battery real power controller, each day, computes time intervals for each house to charge and discharge the batteries. The default charge interval starts at 1:00 a.m. with a duration of 270 min. The minute delay value configured for each house acts as a delay to the start time of this interval to spread out power consumption. This procedure results in charging intervals for houses that range from [1:00 a.m., 5:30 a.m.] to [3:00 a.m., 7:30 a.m.]. The discharge interval is the same for all houses and is the 6 h interval centered on the predicted peak hour⁷. If the peak hour begins at 7:00 p.m., then this interval will be calculated as [4:30 p.m., 10:30 p.m.] which is centered around the time 7:30 p.m. When the battery is not charging and the RTP exceeds a threshold (as given earlier), regardless of the discharge interval, the control will override the battery real power output based on the RTP adjustment. The batteries are charged and

⁶ 90 °F is not an arbitrary value. In reality, the controller should turn off when the real-time price is too high, but there is no easy mechanism to disable the water heater control in GridLAB-D. Instead, while the prices are too high, the setpoint is set to a very low value (compared to the current water temperature in the tank) to prevent the water heater from operating. However, GridLAB-D does not allow for water heater setpoints less than 90 °F, and so the minimum allowed value for GridLAB-D was selected.

⁷ The peak hour is computed each day, at midnight, based on the largest value from the day-ahead price data.

discharged according to this schedule or alternatively according to the RTP Adjust mode, as described in the previous section. The CDP implementation is discussed at the end of this section.

The battery reactive power control for each house monitors the voltage at the house meter to set the reactive power output. This control is not time-dependent and is implemented as described in the previous section. However, the batteries have a maximum charge and discharge rate of 5 kVA (apparent power) that depends on both the real and reactive parts of the battery output. To prevent the FRC from exceeding this 5 kVA limit, the reactive battery control calculates the magnitude of the complex power output after it computes the reactive power output. If the magnitude exceeds the limit, then the real power output is reduced to the maximum possible value that fits under the limit. To support this implementation, the reactive battery control runs sequentially after the real battery control to ensure the latest real power value is used in the magnitude calculation.

The electric vehicle control, each day, generates charge profiles at 8:00 a.m. A charge profile consists of a duration (in minutes) and a ramp rate (in kW/h) for both the ramp up and ramp down periods, a duration (in minutes) and a maximum charge rate (in kW) for the maximum charge period, and the charge start time. The profiles are generated as described in Section 2.5, with duration values being set to 0 when certain charge periods, like ramp down, should be skipped. The FRC uses these charge profiles to update the real power output to the battery that represents the electric vehicle each time step (minute) when using day-ahead price control.

CDP control, for both the batteries and the electric vehicles, relies on additional code to monitor the states of both the transformers and the batteries. For each battery (or electric vehicle), the FRC tracks both the total amount charged and the previous real power output to implement the CDP responsive charging discussed in previous sections. The total amount charged is reset per resource when it completes charging. For each transformer, the FRC uses data received from GridLAB-D to store both its capacity and a list of recently received real power flows through the transformer. These values are used to compute, for each transformer, the CDP using Eqn. 1. The FRC associates each house with a transformer based on the naming scheme used in the GridLAB-D model (both houses and their transformers share a common string prefix). Then, during runtime, the batteries are controlled as described in prior sections using the CDP and M_{flow} values of their associated transformer.

TEM control replaces the electric vehicle and battery control, including the battery reactive power control, with a simple replay of the market position. The *totalQuantity* value from the *Commitment* message is used to set the battery power output for each hour of simulation. No real-time adjustment to price (congestion-based or otherwise) was implemented when using the transactive market. The heat pump and electric water heater controls run as described above when using TEM control, as the market algorithm described in [7] treats those loads as inflexible and assumes they are not shifted in response to market price.

3. Conclusion

The FRC implements heat pump, water heater, battery, and EV charge control based on prices and other house specific parameters. The FRC control function may be given in part to the TEM for EV and battery power schedules. Each device adjusts power consumption in response to day-ahead prices, congestion prices, and, optionally, real-time price spikes. Additionally, the FRC implements battery-reactive power control based on local meter voltage. The FRC has been implemented in Java and deployed in the UCEF co-simulation environment to enable experiments examining the impact of customer device response to price signals.

The FRC has been tested in co-simulation with GridLAB-D using the R4-12.47-1 grid model to examine the impact of price response on distribution power quality. That report is in preparation.

References

- [1] Holmberg, D.G., Omar, F., “Potential for Residential Device Provision of Electric Grid Ancillary Services”, NIST Special Pub 1900-601, 2018, available at: <https://doi.org/10.6028/NIST.SP.1900-601>
- [2] Holmberg, D., Omar, F., Roth, T. (2023), “Impact of Dynamic Prices on Distribution Grid Power Quality” Technical Note 2261, National Institute of Standards and Technology, Gaithersburg, MD, [online], <https://doi.org/10.6028/NIST.TN.2261>.
- [3] 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>
- [4] 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>.
- [5] Flexible Resource Controller Java code at <https://github.com/usnistgov/FlexibleResourceController>.
- [6] Staff, California Public Utilities Commission Energy Division, “Advanced Strategies for Demand Flexibility Management and Customer DER Compensation,” 2022, available at: <https://www.cpuc.ca.gov/-/media/cpuc-website/divisions/energy-division/documents/demand-response/demand-response-workshops/advanced-der---demand-flexibility-management/ed-white-paper---advanced-strategies-for-demand-flexibility-management.pdf>
- [7] Holmberg D, Roth T (2026) NIST Transactive Energy Market: A TE Market for GridLAB-D Co-Simulations. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Technical Note (TN) NIST TN 2384. <https://doi.org/10.6028/NIST.TN.2384>
- [8] PNNL Taxonomy Feeders and Feeder Generator script, https://github.com/gridlab-d/Taxonomy_Feeders, 2008, last updated 2015.
- [9] IEEE 1547-2018 - IEEE Standard for Interconnection and Interoperability of Distributed Energy Resources with Associated Electric Power Systems Interfaces, 2018.
- [10] Feeder Taxonomy Wiki, 2024, available at http://gridlab-d.shoutwiki.com/wiki/Feeder_Taxonomy.
- [11] St. John, Jeff, “5 charts that shed new light on how people charge EVs at home,” 2022, available at: <https://www.canarymedia.com/articles/ev-charging/5-charts-that-shed-new-light-on-how-people-charge-evs-at-home>.
- [12] EV Box, “How much electricity does an electric car use?” 2022, available at: <https://evbox.com/en/ev-home-charger-electricity-usage>.
- [13] IEEE Standard for Modeling and Simulation (M&S) High Level Architecture (HLA)—Framework and Rules. IEEE Computer Society. 18 August 2010. ISBN 978-0-7381-6251-5
- [14] Maroti, M., et.al., "Next Generation (Meta)Modeling: Web- and Cloud-based Collaborative Tool Infrastructure", 2015, available at: <https://webgme.org/WebGMEWhitePaper.pdf>.

Appendix A. Baseline EV Charge Profile Development

The charge profile for the baseline case follows these steps.

1. Find kWh amount of charge, P_{ev} , from the lognormal.inv algorithm for each car.
2. Randomly assign cars to [day, evening, night] (25 % to day, 60 % to evening, 15 % to night) with day window 10:00 – 17:00, evening 17:00 – 23:00 and night 23:00 to 7:00.
3. Create a charge profile for each car, which includes the total charge duration. The charge profile has three stages: ramp up, constant charge, and ramp down. If $P_{ev} < 1.2$, then the profile only uses a truncated ramp up stage. If $1.2 \leq P_{ev} < 8.4$, then the profile uses the full ramp up stage with a truncated ramp down. If $P_{ev} \geq 8.4$, then the profile contains all three stages with the constant charge stage stretched based on the magnitude of P_{ev} .
4. If the charge duration of the charge profile exceeds the available charging time, reduce the duration of the constant charge stage until the total charge amount is equal to the max amount possible during the charge window.
 - a. A day charge has a 7-hour window to charge (10:00 to 17:00).
 - b. An evening charge has a 14-hour window to charge (17:00 to 7:00).
 - c. A night charge has an 8-hour window to charge (23:00 to 7:00).
 - d. Example: If a car assigned to charge during the day has a charge profile with a total duration of 13 hours, the profile is changed so that it fits within the 7-hour window that is available during the day. The constant charge stage of the profile is reduced by $13 - 7 = 6$ hours so the charge profile can finish by 5 pm.
5. Randomly pick a start time for the charging event during the selected charge window based on the calculated charge duration.
 - a. For day charge event with a charge duration D (hours), randomly select an event start time between 10:00 and $(17:00 - D)$ to ensure the charge completes by 17:00 or 5 p.m. Truncate if needed per above.
 - b. For an evening charge there is an exception since the charge can extend into the night charge window.
 - i. If $D \leq 9$, choose a start time between 17:00 and 22:00.
 - ii. If $D > 9$ then select the start times between 17:00 and $(07:00 - D)$.
 - iii. This ensures that the evening charge will always start in the evening charge window, but it will always finish by 7:00.
 - c. For night charge, choose a start time between 23:00 and $(07:00 - D)$.
6. Develop the charge profile.
 - a. Use a charge profile with ramp up and ramp down, with ramp up rate = 21.6 kW/h and ramp down rate of 3.6 kW/h. Thus, the ramp up requires 20 min to go from 0 to 7.2 kW, and 2 hours to ramp down. Add in constant 7.2 kW plateau between up and down ramps to get to the full charge amount (if over 8.4 kWh charge). Truncate the down ramp for $P_{ev} < 8.4$ kWh, and then truncate the up ramp for any $P_{ev} < 1.2$ kWh.

- b. For example, start time is 9:00 p.m. and $P_{ev} = 40$ kWh. In this case, the ramp up starts at 9:00 p.m. and reaches 7.2 kW in 20 min or 9:20 p.m. with charge so far of 1.2 kWh. The ramp down is 7.2 kWh. Thus, the plateau in the middle is $((40 - 8.4)/7.2) = 4$ h 23 min, and total charge length will be 6 h 43 min. In this case, the charge starts at 9:00 p.m. and ends at 3:43 a.m.
 1. Continuing from evening into night is not a problem. As long as charging is done by 17:00 (for day charge) and by 7:00 (for evening or night charges).
 2. If the given start time and P_{ev} result in a time to finish charging after 17:00 for a day charge or after 7:00 for evening/night charge, then in this case, slide the start time earlier to finish at 17:00/7:00.
7. The charge profile is converted to an array of 5-min [time, duration] pairs which will be passed into GLD for each EV battery inverter as the charging schedule.