Coordination between an Aggregator and Distribution Operator to Achieve Network-Aware Load Control

Stephanie C. Ross, Necmiye Ozay, and Johanna L. Mathieu
Department of Electrical Engineering and Computer Science
University of Michigan
Ann Arbor, MI, USA
{sjcrock, necmiye, jlmath}@umich.edu

***Abstract*—Operational issues could arise on a distribution network if a third-party aggregator controls a large number of the network's loads without visibility of network states or topology. This paper investigates mechanisms by which an aggregator could coordinate with a distribution network operator such that aggregate load actions do not negatively impact the network. We propose two possible frameworks for coordination and develop a specific coordination scheme in which the operator can partially block the aggregator's control inputs to loads. We design a control strategy for the aggregator assuming it has little to no information about how the operator is blocking its control. The proposed controller estimates the number of loads that are not receiving the aggregator's commands and compensates accordingly. In a simulation study, the proposed controller consistently outperforms a benchmark controller in terms of average tracking error.**

***Index Terms*—Estimation, Thermostatically controlled loads, Aggregator, Distribution network, Frequency regulation**

I. Introduction

Frequency regulation and energy balancing services will be increasingly important for power system operators as the penetration of intermittent renewable power generation continues to grow. Aggregators can provide these services by controlling thousands of flexible loads such that their total power consumption tracks a desired signal. However, large-scale participation of loads in these transmission-level services could result in local operational issues on distribution networks, such as violations of power flow and voltage constraints [1], [2].

Transmission operators, aggregators, and distribution operators will need to coordinate to prevent negative impacts on distribution networks [2], [3]. How this coordination should be structured is an open question. The U.S. Federal Energy Regulatory Commission has recently requested comments on the following questions: What, if any, real-time information do distribution operators need about aggregations or individual resources within an aggregation? Should distribution operators be able to override real-time dispatch of aggregations to resolve local reliability issues? [4].

Traditionally, load-control algorithms for frequency regulation services have not taken distribution network constraints into account. Instead, it has been assumed that the effect on network operation will be negligible if the percentage of loads participating in the aggregation is relatively small. To the best of our knowledge, only a few papers [5], [6] have proposed real-time algorithms for network-aware load control at the time scale required by frequency regulation services (i.e., seconds). However, in both [5] and [6], the distribution operator also acts as the load aggregator – an additional role that some operators may not want. Moreover, in some locations, operators may be required to allow independent aggregators to participate on their networks.

The purpose of this paper is to achieve network-aware load control through aggregator-operator coordination. In Section II we propose two different frameworks for real-time coordination. In Section III, we propose a specific coordination scheme in which an operator blocks an aggregator's commands if they will cause network issues, and we describe the resulting control problem for an aggregator of thermostatically controlled loads (TCLs). In Section IV, we propose a control strategy for the aggregator. Finally, in Section V, we test the proposed controller against a benchmark controller in a simulation study.

The contributions of this paper are: 1) development of two frameworks for real-time coordination between an independent aggregator and a distribution operator; 2) development of a controller for an aggregator whose commands can be blocked by an operator (specifically, design of a Kalman Filter that estimates the portion of TCLs not receiving commands and compensates within the controller accordingly); 3) a comprehensive simulation study that compares the proposed controller to a benchmark controller across multiple scenarios.

II. Operator-Aggregator Coordination

A. Objectives of Aggregator and Operator

The aggregator's objective is to non-disruptively control hundreds to thousands of loads such that their total power consumption accurately tracks a frequency regulation signal. Load control is considered non-disruptive if the end-use service delivered by loads (e.g., refrigeration) is not disrupted [7].

The distribution operator's objective is to provide sufficient quality power to consumers, while also ensuring safe operation of network components. Voltage magnitudes should be maintained between 0.95 and 1.05 p.u. and unbalance should be less than 3% [8]. Network components, such as lines and transformers, should not be loaded beyond their ratings.

This work was supported by the Rackham Predoctoral Fellowship and U.S. NSF Grant No. CNS-1837680.

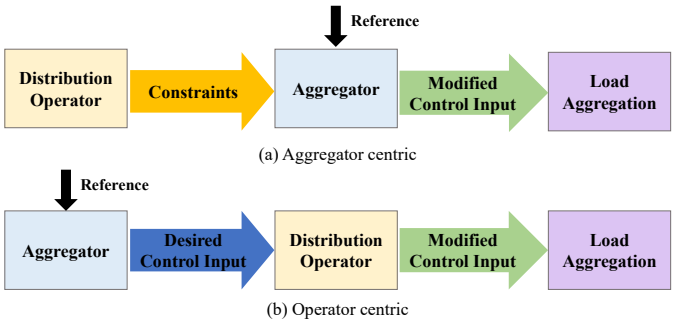

Fig. 1. Proposed coordination frameworks. The distribution operator provides constraints on load actions, the aggregator provides its desired control input, and one entity ((a) aggregator or (b) operator) combines this information into a modified control input that is sent to the loads.

B. Coordination Frameworks

We propose two general frameworks for real-time coordination between the aggregator and operator, as shown in Fig. 1. In the aggregator-centric framework, the operator sends the aggregator constraints on load actions (e.g., maximum and minimum power consumption of groups of loads), and the aggregator computes a control input that will satisfy these constraints. In the operator-centric framework, the operator receives the aggregator's desired control input (e.g., on/off commands to each load) and modifies the input if necessary to satisfy network constraints. The aggregator uses feedback from the loads to compute its control inputs, and the operator uses feedback from the network to estimate three-phase power flows and determine constraints on load actions.

Next we compare the frameworks.

1) Tracking Accuracy: In the aggregator-centric framework, the aggregator knows the constraints on its load actions and can adjust its control to mitigate the effect on tracking accuracy. In the operator-centric framework, the aggregator does not have explicit knowledge of the constraints and will have a harder time compensating for them.

2) Network reliability: In the operator-centric framework, the operator has full control over its network and can ensure its reliability; however, this reliability could come at the expense of the aggregator's objective. In the aggregator-centric framework, the operator may only be able to ensure network reliability by providing overly conservative constraints on load actions, which could result in poor tracking accuracy for the aggregator.

3) Information privacy: In the aggreagtor-centric framework, the aggregator's information is kept private, but the aggregator gains information about the operator and may be able to learn the network's parameters and configuration. In the operator-centric framework, the operator's information is kept private, but the operator gains information about the aggregator and may be able to learn the aggregator's control algorithm.

The proposed frameworks are only two of the many options for operator-aggregator coordination. Other options include a blend of these frameworks and a "load-centric" framework in which loads receive direct commands from both the operator and aggregator.

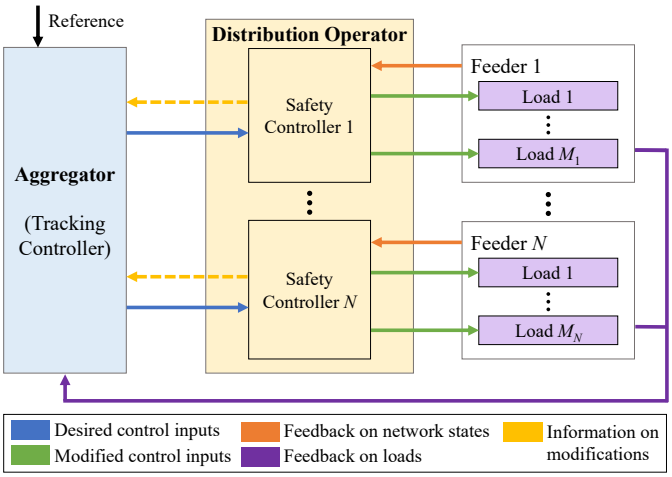

Fig. 2. Specific coordination scheme within the operator-centric framework. The aggregator calculates its desired control input based on the reference signal and feedback from loads. The operator estimates the feeders' states based on measurements and modifies the control inputs if necessary. Information on modifications may or may not be shared with the aggregator.

TABLE I
LEVELS OF FEEDBACK FOR AGGREGATOR

	Full	Moderate	Minimal
Aggregation's total power consumption	✓	✓	✓
Percent of loads blocked	✓	✓	
Blocked state of each load	✓		
On/Off state of each load	✓		

III. Problem Description

This section describes a specific coordination scheme within the operator-centric framework and the resulting control problem for the aggregator.

A. Specific Coordination Scheme

A block diagram of the coordination scheme is shown in Fig. 2. As indicated by the figure, the aggregator's loads may be located on more than one of the operator's feeders, and each feeder has a safety controller to ensure the feeder's constraints are satisfied. When a safety controller receives the aggregator's desired control input, it uses its current estimate of the feeder's states to determine if the input needs to be modified to ensure constraint satisfaction.

Although the modification of the aggregator's input could take many forms, in this problem we assume the operator "blocks" control inputs to certain loads. Unblocked inputs are sent unmodified to the appropriate loads; blocked inputs are simply not sent. Note, we refer to a load as "being blocked" if its control input is blocked. Finally, we define $\beta(k)$ to be the percentage of loads that are blocked at time step k.

We consider three levels of feedback for the aggregator: "full", "moderate", and "minimal". Each level of feedback is defined in Table I. The key difference between the levels is the amount of information that the aggregator receives about the operator's blocking actions. Note that all three levels of feedback include a measurement of the total power consumed by the load aggregation.

B. Aggregator's Control Problem

The aggregator's objective is to minimize the error between the reference power P^*_{total} and the load aggregation's actual power consumption P_{total} in each time step. In this problem, the load aggregation consists of TCLs, such as water heaters, air conditioners, and refrigerators. We use the term "TCL" to refer to both the conditioning device (e.g., heat pump) and the system being conditioned (e.g., water tank or house).

The aggregator's control problem is to track the reference power by switching TCLs on or off, in a manner that is non-disruptive to the end-user. Thus, the control method must respect the TCLs' user-set temperature constraints:

$$\theta^i_{\text{set}} - \delta^i/2 < \theta^i(k) < \theta^i_{\text{set}} + \delta^i/2, \tag{1}$$

where θ is the TCL's temperature, θ_{set} is the setpoint, and δ is the width of the allowed range. A TCL's temperature dynamics can be described by the following model (developed in [9] and frequently used in the literature [10], [11]):

$$\theta^i(k{+}1) = \begin{cases} a^i\theta_i(k) + \big(1-a^i\big)\big(\theta_{\text{a}}(k) - r^i p^i_\theta\big) & \text{if } m^i(k){=}1, \\ a^i\theta_i(k) + \big(1-a^i\big)\theta_{\text{a}}(k) & \text{if } m^i(k){=}0, \end{cases}$$

where $m^i = 1$ indicates the TCL is *on* (i.e., actively transferring thermal energy) and $m^i = 0$ indicates the TCL is *off*. Variable θ_{a} is the ambient, external temperature, and $a^i = \exp(-h/(c^i\, r^i))$, where h is the length of the model's discrete time step. All other parameters are defined in Table II. Ranges of parameter values for residential air conditioning systems, sourced from [12], [13], are shown in Table II.

We assume that the TCLs' temperature constraints and blocking conditions take precedence over the aggregator's switching commands. Let b^i indicate whether the ith TCL is blocked ($b^i = 1$) or not blocked ($b^i = 0$). Let s^i be the aggregator's switching command to the ith TCL with $s^i = 1$ indicating "switch *on*", $s^i = -1$ indicating "switch *off*", and $s^i = 0$ otherwise. Thus, when the ith TCL is under aggregator control, its *on*/*off* state is determined by

$$m^i(k+1) = \begin{cases} 1 & \text{if } \ \theta^i(k) \geq \theta^i_{\text{set}} + \delta^i/2, \\ 0 & \text{if } \ \theta^i(k) \leq \theta^i_{\text{set}} - \delta^i/2, \\ 1 & \text{if } \ s^i(k) = 1,\ b^i(k) = 0, \text{ and (1)}, \\ 0 & \text{if } \ s^i(k) = -1,\ b^i(k) = 0, \text{ and (1)}, \\ m^i(k) & \text{otherwise.} \end{cases} \tag{2}$$

Lastly, if there are N TCLs in the aggregation, their total power consumption is given by $P_{\text{total}}(k) = \sum_{i=1}^{N} m^i(k)p^i$, where $p^i = p^i_\theta/\zeta^i$ is the rated electrical power consumption of the ith TCL and ζ^i is the coefficient of performance.

IV. Controller Design

A. Probabilistic Control

We use probabilistic commands to switch TCLs *on*/*off*, as in [12], because of the light communication requirements. If the operator does not modify the aggregator's desired control inputs, then only two numbers need to be broadcast to all TCLs: u_{off} the probability with which *on* TCLs should switch *off* and u_{on} the probability with which *off* TCLs should switch *on*. Each TCL determines its individual switching command s_i by drawing a random number $z_i(k)$ from the uniform distribution between 0 and 1, and

$$s^i(k) = \begin{cases} 1 & \text{if } m^i(k) = 0 \text{ and } z^i(k) < u_{\text{on}}(k), \\ -1 & \text{if } m^i(k) = 1 \text{ and } z^i(k) < u_{\text{off}}(k), \\ 0 & \text{otherwise.} \end{cases}$$

Once s^i is calculated, a TCL switches according to (2).

Probabilities to switch are calculated based on the predicted error between $P^*_{\text{total}}(k+1)$ the desired total power in the next time step and $P_{\text{total}}(k+1|k)$ the predicted total power in the next time step. If $P^*_{\text{total}}(k{+}1) \geq P_{\text{total}}(k{+}1|k)$, then $u_{\text{off}}(k) = 0$ and

$$u_{\text{on}}(k) = K\frac{|P^*_{\text{total}}(k+1) - P_{\text{total}}(k+1|k)|}{\hat{P}_{\text{off}}(k)}. \tag{3}$$

If $P^*_{\text{total}}(k+1) < P_{\text{total}}(k+1|k)$, then $u_{\text{on}}(k) = 0$ and

$$u_{\text{off}}(k) = K\frac{|P^*_{\text{total}}(k+1) - P_{\text{total}}(k+1|k)|}{\hat{P}_{\text{on}}(k)}. \tag{4}$$

In (3) and (4), K is a proportional control gain, and $\hat{P}_{\text{on}}$ and $\hat{P}_{\text{off}}$ are estimates of the total capacity of all *on* and unblocked TCLs and all *off* and unblocked TCLs, respectively. Lastly, we constrain switching probabilities to within the range $[0, 1]$.

B. Estimator

To implement the controller in (3)-(4), the aggregator needs to estimate the total capacities available to switch, $\hat{P}_{\text{on}}$ and $\hat{P}_{\text{off}}$. We assume the aggregator has little to no information about the operator's blocking actions. Thus, our approach is to estimate the number of TCLs that are *on* and not blocked N_{on} and the number that are *off* and not blocked N_{off}, and then estimate the capacities as

$$\hat{P}_{\text{off}}(k) = \overline{p}_{\text{off}}\hat{N}_{\text{off}}(k) \ \text{ and } \ \hat{P}_{\text{on}}(k) = \overline{p}_{\text{on}}\hat{N}_{\text{on}}(k).$$

Here $\hat{}$ indicates an estimate, and $\overline{p}_{\text{off}}$ and $\overline{p}_{\text{on}}$ are the average power ratings of an *off* TCL and an *on* TCL, respectively. For simplicity, we model $\overline{p}_{\text{off}}$ and $\overline{p}_{\text{on}}$ as constant parameters; in practice, they are time-varying for a heterogeneous aggregation.

We formulate a linear time varying model that represents aggregate TCL dynamics and for which a Kalman Filter can be designed. The aggregate model is

$$\begin{aligned} \boldsymbol{x}(k+1) &= \boldsymbol{A}(k)\boldsymbol{x}(k) + \boldsymbol{w}(k) \\ \boldsymbol{y}(k) &= \boldsymbol{C}\boldsymbol{x}(k) + \boldsymbol{v}(k), \end{aligned} \tag{5}$$

TABLE II
Air Conditioner Parameters

Parameter	Values	Unit
Setpoint temperature (θ_{set})	18-27	°C
Width of temperature range (δ)	0.25-1	°C
Thermal resistance (r)	1.2-2.5	°C/kW
Thermal capacitance (c)	1.5-2.5	kWh/°C
Thermal energy transfer rate (p_θ)	10-18	kW
Coefficient of performance (ζ)	2.5	–

where $\boldsymbol{w}$ and $\boldsymbol{v}$ are process noise and measurement noise, respectively. The state $\boldsymbol{x}$ is defined as

$$\boldsymbol{x}(k) = \begin{bmatrix} N_{\text{on}}(k) & N_{\text{off}}(k) & N_{[\text{on}]}(k) & N_{[\text{off}]}(k) \end{bmatrix}^{\text{T}},$$

where $N_{[\text{on}]}$ the number *on* and blocked, and $N_{[\text{off}]}$ the number *off* and blocked.

We model the state dynamics as a Markov chain, in a similar manner to [12]. However, in this formulation the Markov chain is time-varying because we include the time-varying commands u_{on} and u_{off} within the transition probabilities, as in [14]. The Markov chain can be represented by the transition matrix

$$\boldsymbol{A}(k) = \begin{bmatrix} 1-f-u_{\text{off}}(k) & g+u_{\text{on}}(k) & 0 & 0 \\ f+u_{\text{off}}(k) & 1-g-u_{\text{on}}(k) & 0 & 0 \\ 0 & 0 & 1-f & g \\ 0 & 0 & f & 1-g \end{bmatrix}.$$

The upper block in $\boldsymbol{A}$ consists of the transition probabilities for TCLs that are not blocked; the lower block consists of transition probabilities for TCLs that are blocked. The model does not include transition probabilities between blocked and unblocked states because the operator determines these transitions with its blocking actions; thus, from the perspective of the aggregator, these transitions are an unknown disturbance to the state. Scalars f and g represent the probability of transitions caused by TCLs' internal thermostat control, i.e., from "internally switching". Scalar f is the probability that an *on* TCL is switched *off* internally, and scalar g is the probability that an *off* TCL is switched *on* internally. In the upper block, the probability of transitioning to a new state is equal to the sum of the probabilities of switching internally and switching externally.

The output $\boldsymbol{y}$ and corresponding $\boldsymbol{C}$ matrix are defined as

$$\boldsymbol{y}(k) = \begin{bmatrix} P_{\text{total,meas}}(k) \\ N \end{bmatrix} \quad \boldsymbol{C} = \begin{bmatrix} \overline{p}_{\text{on}} & 0 & \overline{p}_{\text{on}} & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}.$$

Here the first output equation relates the number of TCLs that are *on* to the measured power consumption of the aggregation. The second output equation is an equality constraint on the states: the sum of the states must be equal to the number of TCLs in the aggregation. To incorporate this constraint into the Kalman Filter, we assume the second output is a perfect measurement with zero measurement noise [15].

The system (5) is observable in the time interval $\tau = [k, (k+3)]$ if the observability matrix $\boldsymbol{O}(\tau) = \big[\boldsymbol{C}; \boldsymbol{C}\boldsymbol{A}(k); \boldsymbol{C}\boldsymbol{A}(k+1)\boldsymbol{A}(k); \boldsymbol{C}\boldsymbol{A}(k+2)\boldsymbol{A}(k+1)\boldsymbol{A}(k)\big]$ has rank 4 [16]. We find sufficient conditions for observability by finding conditions under which the matrix has four linearly independent rows. Let the matrix $\boldsymbol{O}_{\text{sub}}$ be made up of rows $\{1, 2, 3, 5\}$ of $\boldsymbol{O}$. Then $\boldsymbol{O}$ has four linearly independent rows if $\boldsymbol{O}_{\text{sub}}$'s determinant is non-zero. The determinant is given by

$$\det \boldsymbol{O}_{\text{sub}}(\tau) = \overline{p}_{\text{on}}^{3}\Big(u_{\text{on}}(k)\big(u_{\text{off}}(k)(1-f) - f u_{\text{on}}(k+1)\big) + u_{\text{off}}(k)\big(g u_{\text{off}}(k+1) + u_{\text{on}}(k+1)(g-1)\big)\Big).$$

Recall that in each time step either $u_{\text{off}} = 0$ or $u_{\text{on}} = 0$. Thus, conditions for a non-zero determinant are as follows:

1) $u_{\text{off}}(k+1) \neq 0$ & $u_{\text{on}}(k) \neq 0$, if $u_{\text{on}}(k+1)=0$ & $u_{\text{off}}(k)=0$;
2) $u_{\text{on}}(k+1) \neq 0$ & $u_{\text{off}}(k) \neq 0$, if $u_{\text{off}}(k+1)=0$ & $u_{\text{on}}(k)=0$;
3) $u_{\text{off}}(k+1) \neq 0$ & $u_{\text{off}}(k) \neq 0$, if $u_{\text{on}}(k+1)=0$ & $u_{\text{on}}(k)=0$;
4) $u_{\text{on}}(k+1) \neq 0$ & $u_{\text{on}}(k) \neq 0$, if $u_{\text{off}}(k+1)=0$ & $u_{\text{off}}(k)=0$.

In summary, a sufficient condition for system observability is for either u_{on} or u_{off} to be non-zero in each time step. Note, the internal switching probabilities f and g cannot cause the determinant to be zero because we assume TCLs are cycling through their temperature range, which implies f and g are not equal to 0 or 1 by definition.

We use a time-varying Kalman Filter [17] to estimate the states of the stochastic model described by (5). Process noise $\boldsymbol{w}$ and measurement noise $\boldsymbol{v}$ represent plant-model mismatch and unknown disturbances. We assume that $\boldsymbol{w}$ and $\boldsymbol{v}$ are zero-mean, Gaussian, white noise processes with covariance $\boldsymbol{Q}$ and $\boldsymbol{R}$, respectively.

C. Implementation

In the case of full feedback, the proposed controller does not use all of the information available to it in every time step. Instead, the Kalman Filter is used to estimate the state, except for time steps in which the percentage blocked changes, i.e., $\beta(k) \neq \beta(k-1)$. At these time steps, the Kalman Filter is bypassed and the state estimate is updated according to $\hat{\boldsymbol{x}}(k) = \big[\sum m^i(k)\big(1-b^i(k)\big), \sum(1-m^i(k))\big(1-b^i(k)\big), \sum m^i(k)b^i(k), \sum(1-m^i(k))b^i(k)\big]$, where all summations are from $i = 1$ to N.

In the case of moderate feedback, the proposed controller uses information on the percentage blocked to improve its state estimate. At time steps when $\beta(k) \neq \beta(k-1)$, the Kalman filter produces the state estimate $\hat{\boldsymbol{x}}(k)$ as usual, and then we update the estimate such that the percentage of TCLs estimated to be blocked equals $\beta(k)$. The update is given by $\hat{\boldsymbol{x}}(k) = \big[\big(1-\beta(k)\big)\big(\hat{x}_1(k)+\hat{x}_3(k)\big), \big(1-\beta(k)\big)\big(\hat{x}_2(k)+\hat{x}_4(k)\big), \beta(k)\big(\hat{x}_1(k)+\hat{x}_3(k)\big), \beta(k)\big(\hat{x}_2(k)+\hat{x}_4(k)\big)\big]$.

D. Benchmark Controller

The benchmark controller computes switching probabilities with (3)-(4) but does not use a Kalman Filter to estimate $\hat{P}_{\text{on}}$ and $\hat{P}_{\text{off}}$. Instead, with full feedback, the benchmark controller estimates are given by $\hat{P}_{\text{on}}(k) = \sum_{i=1}^{N} p^i m^i(k)\big(1-b^i(k)\big)$ and $\hat{P}_{\text{off}}(k) = \sum_{i=1}^{N} p^i\big(1-m^i(k)\big)\big(1-b^i(k)\big)$. With moderate feedback, the benchmark controller estimates are given by $\hat{P}_{\text{on}}(k) = \big(1-\beta(k)\big)P_{\text{total}}(k)$ and $\hat{P}_{\text{off}}(k) = \big(1-\beta(k)\big)\big(\sum_{i=1}^{N} p^i - P_{\text{total}}(k)\big)$. Finally, with minimal feedback, the controller uses the same equations as with moderate feedback, but $\beta(k)$ is set to 0 because the aggregator does not have any information about blocking.

V. Simulation Study

A. Setup

We test controller performance using 1-hour simulations of 1000 TCLs controlled to tracking a frequency regulation signal. A constant set of TCLs is blocked between minutes 20 and 40. We test each controller in 12 different scenarios, where a scenario is defined by a combination of blocking level

TABLE III
TUNING PARAMETERS

Controller	Feedback	K	$\boldsymbol{Q}$	$\boldsymbol{R}$
Benchmark	Full, Moderate	0.96	–	–
Benchmark	Minimal	1.03	–	–
Proposed	Full, Moderate	0.96	$2.5\,\mathrm{diag}(3,\ 3,\ 1,\ 1)$	$\mathrm{diag}(1,\ 0)$
Proposed	Minimal	1	$9.3\,\mathrm{diag}(1,\ 7,\ 1,\ 7)$	$\mathrm{diag}(1,\ 0)$

$\mathrm{diag}(\cdot)$ maps an n-tuple to the corresponding nxn diagonal matrix.

(0%, 20%, 40%, or 60% of the population) and feedback level (see Table I).

We run 24 randomized trials for each scenario. Each trial has: a different random instance of TCL parameters, different random numbers generated during probabilistic dispatch, and a different frequency regulation signal. TCL parameters are drawn from the uniform distributions described in Table II. For the frequency regulation signal, Trial 1 uses the first 1-hour segment of PJM's RegD signal [18] from June 3, 2018, Trial 2 uses the second segment, and so on. In all trials, the ambient temperature is 32°C, and the TCL aggregation's regulation capacity is set to +/-20% of baseline power consumption.

For these simulations, we use a persistence model to predict the total power of the aggregation in the next time step (i.e., $P_{\text{total}}(k+1|k) = P_{\text{total,meas}}(k)$ in (3)-(4)) because using the aggregate model (5) for prediction would result in poor accuracy. In addition, we do not add noise to the measurement of P_{total}, i.e., $P_{\text{total,meas}}(k) = P_{\text{total}}(k)$.

Values for tuning parameters are listed in Table III. The parameters for each combination of controller and feedback level are tuned separately, although sometimes are equal. For the proposed controller, we iteratively tune $\boldsymbol{Q}$ and $\boldsymbol{R}$ with $K = 1$. Then, with $\boldsymbol{Q}$ and $\boldsymbol{R}$ fixed for the proposed controller, we sweep through values for K from 0.95 to 1.05 for both controllers. We select the parameter values that yield the lowest sum of root-mean-square errors (RMSE) in tracking, across Trials 1-4 and all blocking levels.

Prior to the test hour, we run the simulation for an hour to ensure steady-state conditions at the start of the test. We also use this pre-test hour to calculate the aggregation's baseline power consumption, the $\boldsymbol{A}$ matrix's transition probabilities, and the parameter $\overline{p}_{\text{on}}$. For simplicity, we set $\overline{p}_{\text{off}} = \overline{p}_{\text{on}}$.

B. Results

1) Estimation Performance: The proposed controller's Kalman Filter is able to estimate the number of TCLs that are blocked, even when no explicit blocking information is available. Performance of the Kalman Filter during Trial 4 is shown in Fig. 3. State estimates are more accurate when feedback on percent blocked is available (left column) than when no feedback on blocking is available (right column), as would be expected. Step changes in the percentage blocked occur at minutes 20 and 40; in the case of minimal feedback, these step changes cause a lag in the state estimates.

2) Tracking Performance: Figure 4 shows the proposed and benchmark controllers tracking the frequency regulation signal while 60% of TCLs are blocked in Trial 2. We evaluate

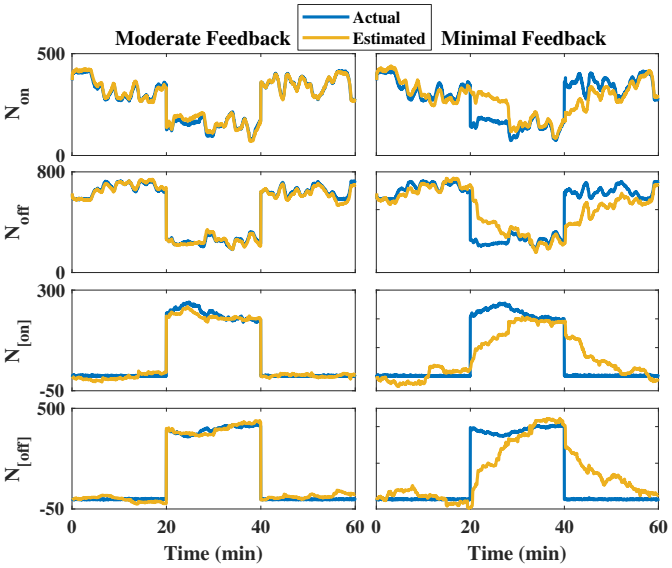

Fig. 3. Estimated states and actual states with the proposed controller when feedback on percentage blocked is available (left) and when no feedback on blocking is available (right). (Data is from Trial 4.)

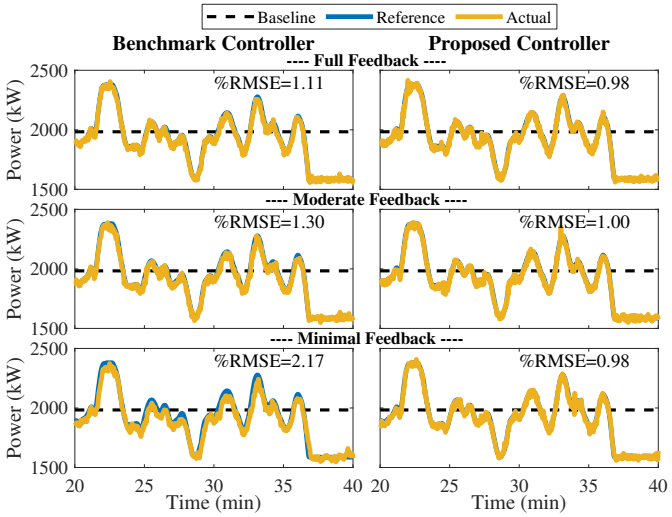

Fig. 4. Benchmark and proposed controllers tracking the frequency regulation signal during the blocking period of Trial 2 with 60% blocked. The proposed controller improves upon the benchmark controller at all levels of feedback.

performance in terms of percent RMSE, which is the error in total power consumption as a percentage of the aggregation's average baseline power consumption. In this trial, the proposed controller improves upon the benchmark controller at every level of feedback. The decrease in RMSE is largest in the case of minimal feedback.

Table IV reports the average RMSE in tracking for each scenario, with the average taken across the 24 trials. In every scenario, the proposed controller's average RMSE is less than or equal to that of the benchmark controller. On average, error increases as the percentage blocked increases and as the level of feedback decreases. The proposed controller's improvement over the benchmark is not very large in most scenarios, with a decrease in average percent RMSE of $\leq 0.03\%$ in 10 of 12 scenarios. The largest improvements occur in the remaining

TABLE IV
CONTROLLER PERFORMANCE: AVERAGE % RMSE IN TRACKING

Percent Blocked	Full Feedback		Moderate Feedback		Minimal Feedback	
	BM	PP	BM	PP	BM	PP
0%	0.76	0.75	0.76	0.75	0.77	0.75
20%	0.77	0.75	0.77	0.76	0.78	0.75
40%	0.78	0.76	0.79	0.77	0.86	0.77
60%	0.86	0.86	0.89	0.86	1.22	0.93

BM = Benchmark controller and PP = Proposed controller

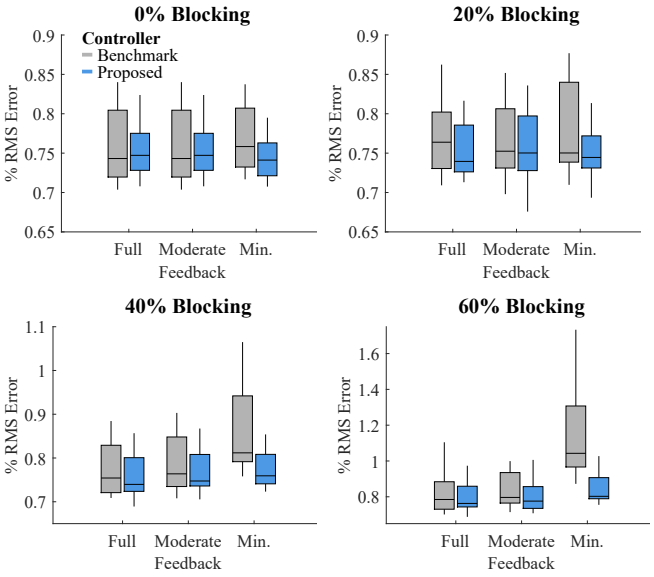

Fig. 5. Box plots of tracking errors in 24 trials for each blocking level and feedback level. Each box plot indicates the trials' median (horizontal line), interquartile range (box), and maximum and minimum points (whiskers). (Note: outliers are not plotted, and not all y-axis scales are equal.)

two scenarios when there is minimal feedback and blocking is at 40% and 60%. In these scenarios, average percent RMSE decreases by 0.09% and 0.29%, respectively.

Box plots of the 24 trials across all scenarios are shown in Fig. 5. The proposed controller generally reduces the maximum and median RMSE across the 24 trials. In all scenarios, the maximum error decreases with the proposed controller; in all but two scenarios, the median error also decreases. As with the mean error results, the box plots show that the proposed controller yields the biggest improvements when 40% or 60% of TCLs are blocked and minimal feedback is available.

It is somewhat surprising that the proposed controller outperforms the benchmark controller when full feedback is available. This result may be because the benchmark controller lacks information on how many TCLs are outside of their set temperature range and thus unavailable to switch. In contrast, the proposed controller has some information of this nature because its Kalman Filter estimates how many TCLs are generally unavailable, whether it is due to blocking or out-of-range temperature.

VI. CONCLUSION

We have proposed two coordination frameworks – "operator centric" and "aggregator centric" – that enable network-aware load control while maintaining the independence of the aggregator and the distribution operator. Working within the operator-centric framework, we proposed a specific scheme in which the operator has the authority to block aggregator commands in order to protect network reliability. We then designed a controller for the aggregator that successfully estimates and compensates for the number of loads that are blocked. In future work, we plan to develop a strategy for the operator that it is maximally permissive of the aggregator's control and includes a real-time assessment of network reliability.

REFERENCES

[1] S. C. Ross, G. Vuylsteke, and J. L. Mathieu, "Effects of load-based frequency regulation on distribution network operation," *IEEE Trans. on Power Systems*, 2018.

[2] More than Smart, "Coordination of transmission and distribution operations in a high distributed energy resource electric grid," Tech. Rep., Jun. 2017. [Online]. Available: https://www.caiso.com/Documents/MoreThanSmartReport-CoordinatingTransmission_DistributionGridOperations.pdf

[3] H. Gerard, E. I. Rivero Puente, and D. Six, "Coordination between transmission and distribution system operators in the electricity sector: A conceptual framework," *Utilities Policy*, vol. 50, pp. 40–48, 2018.

[4] Federal Energy Regulatory Commission, "Notice Inviting Post-Technical Conference Comments," Tech. Rep. Docket No. RM18-9-000, Apr. 2018. [Online]. Available: https://www.ferc.gov/CalendarFiles/20180427135034-notice-for-comments.pdf

[5] E. Vrettos and G. Andersson, "Combined load frequency control and active distribution network management with thermostatically controlled loads," in *IEEE International Conf. on Smart Grid Comm.*, Oct. 2013.

[6] E. Dall'Anese, S. S. Guggilam, A. Simonetto, Y. C. Chen, and S. V. Dhople, "Optimal regulation of virtual power plants," *IEEE Trans. on Power Systems*, vol. 33, no. 2, pp. 1868–1881, Mar. 2018.

[7] D. S. Callaway and I. A. Hiskens, "Achieving controllability of electric loads," *Proc. of the IEEE*, vol. 99, no. 1, pp. 184–199, Jan. 2011.

[8] National Electrical Manufacturers Association, "American national standard for electric power systems and equipment – voltage ratings (60 Hertz)," Tech. Rep. ANSI C84.1-2006, Dec. 2006.

[9] R. Sonderegger, "Dynamic models of house heating based on equivalent thermal parameters," Ph.D. dissertation, Princeton Univ., N.J., 1978.

[10] D. Callaway, "Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy," *Energy Conversion and Management*, vol. 50, pp. 1389–1400, 2009.

[11] P. Du and N. Lu, "Appliance commitment for household load scheduling," *IEEE Trans. on Smart Grid*, vol. 2, no. 2, pp. 411–419, 2011.

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

[13] J. Mathieu, M. Dyson, and D. Callaway, "Resource and revenue potential of California residential load participation in ancillary services," *Energy Policy*, vol. 80, pp. 76–87, 2015.

[14] Y. Chen, A. Bušić, and S. P. Meyn, "State Estimation for the Individual and the Population in Mean Field Control with Application to Demand Dispatch," *IEEE Trans. on Automatic Control*, vol. 62, no. 3, pp. 1138–1149, 2017.

[15] D. Simon, "Kalman filtering with state constraints: a survey of linear and nonlinear algorithms," *IET Control Theory & Applications*, vol. 4, no. 8, pp. 1303–1318, 2010.

[16] J. Bay, *Fundamentals of Linear State Space Systems*. Electrical and Computer Engineering Faculty Scholarship, 1999. [Online]. Available: https://orb.binghamton.edu/electrical_fac/3/

[17] MATLAB. (2018) Kalman filtering. [Online]. Available: https://www.mathworks.com/help/control/ug/kalman-filtering.html

[18] PJM ancillary services. [Online]. Available: http://www.pjm.com/markets-and-operations/ancillary-services.aspx