

Strategies for Network-Safe Load Control with a Third-Party Aggregator and a Distribution Operator

Stephanie C. Ross and Johanna L. Mathieu

Abstract

When providing bulk power system services, a third-party aggregator could inadvertently cause operational issues at the distribution level. We propose a coordination architecture in which an aggregator and distribution operator coordinate to avoid distribution network constraint violations, while preserving private information. The aggregator controls thermostatic loads to provide frequency regulation, while the distribution operator overrides the aggregator's control actions when necessary to ensure safe network operation. Using this architecture, we propose two control strategies, which differ in terms of measurement and communication requirements, as well as model complexity and scalability. The first uses an aggregate model and blocking controller, while the second uses individual load models and a mode-count controller. Both outperform a benchmark strategy in terms of tracking accuracy. Furthermore, the second strategy performs better than the first, with only 0.10% average RMS error (compared to 0.70%). The second is also able to maintain safe operation of the distribution network while overriding less than 1% of the aggregator's control actions (compared to approximately 15% by the first strategy). However, the second strategy has significantly more measurement, communication, and computational requirements, and therefore would be more complex and expensive to implement than the first strategy.

I. INTRODUCTION

High penetrations of aggregator-controlled distributed energy resources (DERs) pose a challenge to safe operation of distribution networks. In particular, DER aggregations participating in wholesale markets may cause unexpected power flows at the distribution level that result in reliability issues, as was noted by the U.S. Federal Energy Regulatory Commission (FERC) in a recent notice [1]. This is of particular concern with third-party aggregators, which are separate from distribution operators and do not necessarily have awareness of the local impacts of their control actions. In the U.S., third-party load aggregators are active in ancillary service markets as a result of FERC Order No. 719 [2]. As an example of the problem, consider a third-party load aggregator participating in frequency regulation. To achieve a maximum setpoint, the aggregator may need 80% of its loads to draw power at the same time—effectively turning a distribution network's diversified load into coincident load—which could cause violations of lower limits on voltage magnitudes [3].

In this paper, we propose new strategies for controlling a particular type of load aggregation—an aggregation of residential, thermostatically controlled loads (TCLs). Examples of TCLs include air conditioners (ACs), heat pumps, and water heaters. The term “TCL” refers to both the thermal space (e.g., house) and appliance (e.g., AC). Every TCL has some thermal energy storage capacity, which allows for slight shifts in when electrical power is consumed without affecting temperature regulation. By controlling thousands of TCLs in this way, an aggregator can provide balancing services in the wholesale market [4]. TCL aggregations are particularly well suited for services that are (or can be made) energy neutral, such as frequency regulation (i.e., secondary frequency control) [5]. Most research on TCL control for fast balancing services has focused on the development of control strategies that 1) do not have a large impact on the TCLs' thermal regulation performance and 2) have low measurement and communication requirements (e.g., [6]–[14]). Notably, all of the above-mentioned strategies have neglected distribution network operation and constraints.

The objective of this paper is to develop TCL control strategies that provide frequency regulation without causing network constraint violations, all while preserving the privacy of the operator's and third-party aggregator's private information. The distribution operator's network model and measurements are assumed private to ensure system security and consumer privacy, and the aggregator's control algorithm for providing frequency regulation is assumed private to maintain a competitive advantage commercially. Given these privacy constraints, we propose a coordination architecture in which the aggregator controls a TCL aggregation to provide frequency regulation, but, when necessary for network safety, the operator overrides the aggregator's commands to specific TCLs.

Our approach contrasts with that of prior work on network-safe load control, which has not focused on maintaining privacy and separation between the operator and aggregator. In [15], the aggregator and operator are treated as the same entity, and a centralized AC optimal power flow (OPF) is solved. In [16]–[18], distributed strategies are proposed, but the aggregator executes a prescribed (i.e., not private) algorithm that is coupled to the operator's algorithm through Lagrange multipliers. In [19], a method for certifying if a distribution network will operate safely given any possible load control action is proposed; however, a control solution is still needed for networks that cannot be certified as safe, which are the types of networks we consider in this paper. Nazir and Almassalkhi are also investigating grid-aware approaches to dispatching demand [20] and coordination of distributed energy resources [21]; though the setting is not exactly the same as ours.

S. C. Ross and J. L. Mathieu are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 USA (e-mail: sjcrock@umich.edu; jlmath@umich.edu).

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

TABLE I
DISTRIBUTION NETWORK CONSTRAINTS

Component	Variable	Lower Limit	Upper Limit
Line	Current magnitude	–	100% of rating
Transformer	Avg. apparent power	–	100% of rating
Service node	Voltage magnitude	0.95 p.u.	1.05 p.u.

TABLE II
AIR CONDITIONER PARAMETERS

Parameter	Values	Unit
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
Setpoint temperature (θ_{set})	18-27	°C
Width of temperature range (δ)	0.25-1	°C
Coefficient of performance (ζ)	2.5	–

The main contributions of this paper are as follows. First, we propose a privacy-preserving coordination architecture that enables an aggregator and operator to coordinate to achieve network-safe TCL control. There are two variants of the architecture: with and without direct communication between operator and aggregator. Second, we propose two control strategies that use this coordination architecture. The strategies differ in terms of ease of implementation and accuracy, and thus provide options for operators and aggregators with different capabilities and preferences. Third, we evaluate the performance of the proposed strategies in simulations of a realistic distribution network model with a high penetration of aggregated TCLs.

Certain aspects of the proposed control strategies relate to our prior work in [22], [23]. In [22], we proposed an alternative strategy for controlling a TCL aggregation to provide frequency regulation in a network-safe manner. This paper’s Strategy I is similar to the strategy in [22] in that they both use blocking to ensure network safety; the two strategies are different in all other respects, including coordination architecture as well as the aggregator’s control algorithm. In [23], we considered a different objective: providing local services to the distribution network by controlling small groups of TCLs. In this paper’s Strategy II, the operator uses the “mode-count control algorithm” from [23] to relieve voltage and line constraints. Lastly, [22] and [23] did not include a full network model; in this paper, we test the strategies on a 613-node network model.

II. PROBLEM DESCRIPTION

The *aggregator’s objective* is to control the total power consumption of its TCL aggregation such that it tracks a frequency regulation signal with sufficient accuracy. We assume that the aggregator’s control actions should be non-disruptive [24] to the TCL’s end-user. Thus, we restrict an aggregator’s control actions to switching TCLs on/off only when the TCLs are within their user-set temperature range like in [6]–[8], [10]. The aggregator’s control is subject to the individual dynamics and constraints of each TCL.

A *distribution operator’s objective* is to reliably deliver power of sufficient quality to consumers. In Table I, we list the set of constraints for power quality and reliability that are considered in this paper. For transformers, apparent power is averaged over an hour because a transformer’s thermal mass enables short-term violations of its power rating without causing overheating [25]. Note we will use the term “network safety” to refer to a distribution network’s power quality and reliability.

Operators may need new tools to ensure network safety when third-party aggregators are active on their networks. Traditional network management tools, such as network reconfiguration and voltage regulation with regulators, have not been designed for the short but recurring fluctuations in power that can occur when aggregated loads provide frequency regulation. In this paper, we give the operator the ability to override the aggregator’s control actions when necessary to avoid network constraint violations. As with the aggregator, we assume that the operator’s TCL control actions must be non-disruptive to the end-user.

The *control plant*, for both the aggregator and operator, is an aggregation of thousands of heterogeneous TCLs. Each TCL in the plant is modeled separately to capture its individual temperature dynamics and constraints. We use an individual TCL model that was developed in [26] and is commonly found in the literature (e.g., [6], [10]). For an aggregation of cooling TCLs, the i th TCL’s temperature at time step $t+1$ is modeled as

$$\theta_{t+1}^i = \begin{cases} a^i \theta_t^i + (1 - a^i)(\theta_{a,t} - r^i p_\theta^i) & \text{if } m_t^i = 1, \\ a^i \theta_t^i + (1 - a^i)\theta_{a,t} & \text{if } m_t^i = 0, \end{cases} \quad (1)$$

where m represents the TCL’s power status (1 for *on* and 0 for *off*) and θ_a is the ambient temperature outside of the TCL. Parameter a^i is calculated as $a^i = \exp(-h/(c^i r^i))$ where h is the duration of the discrete model’s time step. All other parameters are defined in Table II. The table also lists parameter values from [6], [27] for residential ACs.




Fig. 1. Coordination Architecture. The operator controls a portion of TCLs when necessary to prevent network violations. The aggregator controls all other TCLs to provide tracking. The dashed arrow indicates a variant in which the aggregator and operator directly communicate in real-time.

To ensure non-disruptive control, we give the TCL's internal thermostat controller priority over external controllers. A TCL's power status is switched on/off by its internal controller if it reaches its upper or lower temperature limit. This condition is given by

$$m_{t+1}^i = \begin{cases} 1 & \text{if } \theta_t^i \geq \bar{\theta}^i, \\ 0 & \text{if } \theta_t^i \leq \underline{\theta}^i, \end{cases} \quad (2)$$

where $\underline{\theta}^i = \theta_{\text{set}}^i - \delta^i$ and $\bar{\theta}^i = \theta_{\text{set}}^i + \delta^i$ (see Table II for parameter definitions).

TCLs with compressors (e.g., ACs) can also have a lockout constraint that prevents compressors from cycling too frequently. After a TCL's compressor is switched on/off, it is "locked" for a period of time before it can be switched again. Given a lockout period of τ_L time steps, a TCL's power status is constrained by

$$m_t^i = m_{t-1}^i \quad \text{if} \quad \sum_{k=t-\tau_L}^{t-1} \mathbb{1}_{\{m_{t-1}^i\}}(m_k^i) < \tau_L, \quad (3)$$

where the indicator function $\mathbb{1}_{\{f\}}(g)$ equals one if $g = f$ and zero otherwise. In this paper, we assume the lockout constraint is designed to prevent excessive switching by external controllers and that it does not apply to internal control.

Finally, the total power of an aggregation of N TCLs is given by $P_{\text{total},t} = \sum_{i=1}^N p^i m_t^i$, where p^i is a TCL's electric power rating with $p^i = p_{\theta}^i / \zeta^i$.

III. COORDINATION ARCHITECTURE

We propose a coordination architecture by which the aggregator and operator can both achieve their control objectives while protecting their private information. Figure 1 shows the architecture: on the left side, the aggregator computes its next tracking command given TCL measurements and the regulation signal but does not have access to the operator's private information (i.e., network model and measurements); on the right side, the operator computes its next safety command based on network measurements (e.g., currents and voltages) and TCL measurements but does not have access to the aggregator's private information (i.e., its control algorithm and next control actions). It should be noted that optimal control could be achieved by sharing all information between the aggregator and operator and making centralized control decisions, but this would not achieve the privacy objective.

We propose two control strategies—Strategy I and II—both of which use the proposed architecture. Strategy II is higher accuracy but harder to implement (with higher modeling, communication, and measurement requirements) than Strategy I; Table III directly compares the strategies. In Strategy I, the aggregator uses an aggregate-model based control strategy to provide frequency regulation, while the operator blocks particular TCLs from receiving the aggregator's commands when the network is at risk. Since control partially synchronizes TCLs [3], the aggregate power consumption of blocked TCLs is generally less variable than that of controlled TCLs. In Strategy II both the aggregator's and operator's control algorithms rely on individual models of each TCL; in contrast to Strategy I, the operator *actively* controls a subset of TCLs to reduce the variability of their aggregate power consumption to ensure network safety. Because Strategy II relies on individual models of TCLs, it is less scalable: Strategy II's models increase linearly with TCL-population size, while Strategy I's aggregate model is independent of population size. Strategy I is also more scalable in terms of communication requirements: the aggregator receives aggregated measurements and sends aggregated commands that are independent of population size.

TABLE III
CONTROL STRATEGIES (AGG. = AGGREGATOR, OP. = OPERATOR)

	Entity	Controller	Measurements	Commands to TCLs
Strategy I	Agg.	Aggregate-Model Based Control	Aggregate power	Probabilistic, one value for each bin
	Op.	Blocking Control	Power of each TCL	Block or unblock to each TCL
Strategy II	Agg.	Individual-Model Based Control	Power & temp. of each TCL	On/off to each TCL
	Op.	Mode-Count Control	Power & temp. of each TCL	On/off to each TCL

IV. METHODS: STRATEGY I

A. Blocking Control

In Strategy I, we propose that operators use blocking control to ensure safe network operation. When a TCL is blocked, it is unresponsive to aggregator commands, and it returns to its regular on/off cycles governed by its internal thermostat. To protect a particular network constraint, we block a group of TCLs whose coincident demand would otherwise contribute to the violation of the constraint. Blocking a group of TCLs reduces their on/off synchronization—and the resulting peaks and valleys in coincident demand—that would otherwise be caused by the aggregator’s commands.

To protect a network in real-time, we must identify the set of constraints that are at risk due to aggregator-control and the appropriate TCLs to block to reduce this risk. This problem is non-trivial: it needs to be solvable in real-time while accounting for 3-phase unbalanced power flow and network constraints, TCLs’ discrete operating states (on and off), uncertainty in the aggregator’s future control actions, and uncertainty in the uncontrolled load consumption at each bus. To the best of our knowledge, solution methods for this type of problem do not yet exist in the literature. Work on a related problem [19] shows promise but relies on simplifying assumptions (e.g., balanced power flow enabling consideration of a single-phase equivalent network). Our preliminary work on this problem [28] develops an OPF-type approach to constrain the norm of the vector of nodal power deviations caused by an aggregator, but it is conservative (in a deterministic setting) and does not consider uncertainty. Ref. [20] poses the idea of constraining dispatchable demand as a hosting capacity problem. It uses convex inner approximations of network-feasible power flows, but again assumes a deterministic setting.

Developing an online solution to this problem is important future work, but a significant undertaking on a different type of problem than the one considered here, and so considered outside the scope of this paper. Instead, we determine the set of TCLs to block ensuring network safety for the time horizon of the problem (e.g., an hour) using iterative offline three-phase unbalanced power flow simulations in which perfect knowledge of the full problem, including the aggregator’s control actions over the time horizon, is assumed. Details of the method are given in Appendix A. Therefore, our results provide an upper-bound on performance; in practice, the distribution operator would not be able to select exactly the right set of TCLs to block (maximizing TCL control capacity while ensuring satisfaction of all network constraints) because of uncertainty.

B. Aggregate-Model Based Tracking Control

Strategy I’s proposed tracking controller is based on an aggregate model of TCL state progression. In Section VI, we benchmark the tracking performance of the proposed, aggregate-model based controller against that of a model-less, PI controller. We expect the proposed controller to have better performance because it uses model-based prediction to compensate for the internal control actions of TCLs and uses state estimation to prioritize switching TCLs that are about to internally switch.

1) *Aggregate TCL Model:* The proposed model takes the form of a population transition model, which has been commonly used in the literature on TCL control (e.g., [6], [7], [13], [29], [30]). In this type of model, the state space of all TCLs is discretized into the same number of “bins” and the progression of TCLs from bin to bin is modeled with transition probabilities.

In this paper, we extend the population transition model to include TCL lockout dynamics. Figure 2 depicts the proposed model. Each bin is represented by a circle and is defined by a locked/unlocked status, on/off status, and one of N_I temperature intervals. In Fig. 2, horizontal arrows represent transitions from one temperature interval to the next due to TCLs’ temperature dynamics. Diagonal arrows represent transitions due to thermostat actions (i.e., internal control). Vertical arrows represent TCLs becoming unlocked after their lockout times have elapsed. Transitions due to aggregator-control are not shown but would be downward diagonals in each temperature interval.

We formulate the aggregate model as a linear time-varying system:

$$\begin{aligned} \mathbf{x}_{t+1} &= \mathbf{A}_t \mathbf{x}_t \\ \mathbf{y}_t &= \mathbf{C} \mathbf{x}_t, \end{aligned} \tag{4}$$

where $\mathbf{x}$ is the state vector, $\mathbf{A}$ the state transition matrix, $\mathbf{y}$ the output, and $\mathbf{C}$ the output matrix. The state $\mathbf{x}$ is the distribution of the TCL population across the bins. Given the different types of bins, the state vector can be divided into four categories:




Fig. 2. State Transition Diagram with Locked States. Arrows indicate the most likely state transitions for an uncontrolled TCL with a lockout period.

$\mathbf{x} = [\underbrace{x^1 \dots x^{N_I}}_{\text{off, unlocked}} \underbrace{x^{N_I+1} \dots x^{2N_I}}_{\text{on, unlocked}} \underbrace{x^{2N_I+1} \dots x^{3N_I}}_{\text{off, locked}} \underbrace{x^{3N_I+1} \dots x^{4N_I}}_{\text{on, locked}}]^T$. The entries of the state transition matrix $\mathbf{A}_t$ are the probability that a TCL will transition from bin i to bin j in the next time step. For example, the entry in the i th row and j th column is the probability that a TCL in bin j will transition to bin i in the next time step.

We model transitions due to the aggregator's probabilistic commands within the $\mathbf{A}$ matrix. This is similar to the Markov Decision Process approach in [31]. We formulate $\mathbf{A}_t$ as the product of two transition matrices: $\mathbf{A}_t = \mathbf{A}_{u,t} \mathbf{A}_s$. Matrix $\mathbf{A}_{u,t}$ models "external transitions" that occur due to the TCLs' response to the aggregator's command $\mathbf{u}_t$. Matrix $\mathbf{A}_s$ models "internal transitions" that occur due to the TCLs' autonomous dynamics, including temperature dynamics, lockout dynamics, and internal control. For the purpose of modeling, we have assumed that, in a given time step, internal transitions occur before external transitions. The matrix $\mathbf{A}_s$ can be identified for a particular outdoor temperature by counting bin transitions when aggregator-control is inactive. Here, for simplicity, we assume a constant outdoor temperature, which makes $\mathbf{A}_s$ time-invariant and is not an unreasonable assumption over the duration of a frequency regulation market period, e.g., 1 hour.

Transition probabilities in $\mathbf{A}_{u,t}$ depend on the aggregator's command $\mathbf{u}_t$. We design the aggregator's command to be a vector that can be broadcast to all TCLs and thereby reduce communication requirements. The i th entry of $\mathbf{u}_t$ is the probability with which TCLs in bin i should switch at time step t . The entries only correspond to bins that are unlocked, so $\mathbf{u}_t$ has length $2N_I$. As an example, consider a probabilistic command in which $u_5 = 0.4$, i.e., "TCLs in bin 5 switch with 40% probability". Each TCL receives this command, checks its current state, and, if in bin 5, decides whether to switch by running a Bernoulli random trial with 40% chance of success (i.e., of switching). Given this definition of $\mathbf{u}_t$, the matrix $\mathbf{A}_{u,t}$ takes the form

$$\begin{bmatrix} \mathbf{I} - \text{diag}(u_t^1, \dots, u_t^{N_I}) & \mathbf{0} \\ \mathbf{0} & \mathbf{I} - \text{diag}(u_t^{N_I+1}, \dots, u_t^{2N_I}) \\ \mathbf{0} & \text{adiag}(u_t^{2N_I}, \dots, u_t^{N_I+1}) \\ \text{adiag}(u_t^{N_I}, \dots, u_t^1) & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \\ \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix},$$

where $\mathbf{I}$ is the identity matrix, $\text{diag}(\cdot)$ maps the input vector to a diagonal matrix, and $\text{adiag}(\cdot)$ maps the input vector to an anti-diagonal matrix, where the diagonal runs from the upper right corner to the lower left corner. The left side of $\mathbf{A}_{u,t}$ models transitions due to aggregator control: the diagonal sub-matrices model transitions out of unlocked bins, and the anti-diagonal sub-matrices model transitions into locked bins.

The two outputs of the aggregate model are: 1) the total power consumption of the aggregation and 2) the fraction of TCLs in the aggregation, which should always equal one. The output matrix is given by $\mathbf{C} = [\mathbf{C}_s \ \mathbf{C}_s]$, where

$$\mathbf{C}_s = \begin{bmatrix} 0 & \dots & 0 & \bar{p}_{\text{on}} N & \dots & \bar{p}_{\text{on}} N \\ 1 & \dots & 1 & 1 & \dots & 1 \end{bmatrix}$$

and $\bar{p}_{\text{on}}$ is the average power consumption of an *on* TCL.

2) *State Estimation*: We use a time-varying Kalman filter [32] to estimate the aggregate state in (4). The output of the Kalman filter in each time step is $\hat{\mathbf{x}}_t$. We account for process noise and measurement noise by adding terms $\mathbf{w}$ and $\mathbf{v}$ to the state and output equations in (4), respectively. We treat the covariances of $\mathbf{w}$ and $\mathbf{v}$, denoted as $\mathbf{Q}$ and $\mathbf{R}$, as tuning parameters. For output measurements, we use $\mathbf{y}_{\text{meas},t} = [P_{\text{total},t} \ 1]^T$. We assume the second output measurement is perfect (i.e., has zero measurement noise) and set the corresponding entry in $\mathbf{R}$ to zero. In this way, the second output equation acts as an equality constraint within the Kalman filter [33].

3) *Control Policy*: We propose a control policy based on one-step model prediction. The policy selects a control input such that the model's predicted output in the next time step matches the desired output. Appendix B provides a derivation of the policy. We note that the proposed policy is similar to that of [6]; slight differences emerge due to differences between our aggregate model and that of [6].




Fig. 3. Demonstration of Mode-Count Control. The policy maintains the on-count ≤ 3 by switching *off* TCLs *on*, as soon as possible after entering the upper margin, and if needed, switching an *on* TCL *off* to compensate. TCLs are indicated by numbered circles; boxes indicate locked status.

Here we describe the policy for when an increase in power is needed. First, the choice of $\mathbf{u}_t$ must satisfy the equation:

$$\sum_{n=1}^{N_I} u_t^n x_{\text{pred},t}^n = K \frac{P_{\text{total},t+1}^* - \mathbf{C}^1 \mathbf{A}_s \hat{\mathbf{x}}_t}{\bar{p}_{\text{on}} N}, \quad (5)$$

where $\mathbf{x}_{\text{pred},t} = \mathbf{A}_s \hat{\mathbf{x}}_t$, K is a proportional gain, P_{total}^* is the desired power in the next time step, and $\mathbf{C}^1$ is the first row of $\mathbf{C}$. Because (5) is under-determined when $N_I > 1$, we have the freedom to choose $\mathbf{u}$ as long as (5) is satisfied. We design a rule for selecting $\mathbf{u}$ that prioritizes switching TCLs that are closest to being switched by internal control. When switching TCLs *on*, bin N_I has first priority, bin $N_I - 1$ has second priority, and so on. Let b be the minimum index for which $\sum_{k=0}^b x_{\text{pred},t}^{N_I-k}$ is greater than the right hand side of (5). Then we set the entries of $\mathbf{u}$ equal to 1 for all bins of higher priority than $N_I - b$ and equal to zero for all bins of lower priority. Finally, the value for bin $N_I - b$ is chosen such that (5) is satisfied.

The control policy is very similar when a decrease in power is needed. In this case, $\mathbf{u}$ must satisfy (9) in Appendix B, and bins $2N_I$ to $N_I + 1$ are switched out of, in that order.

Finally, when there is direct communication in Strategy I, the operator informs the aggregator of the number of TCLs it is blocking in the current time step. To incorporate this information into the aggregator's control policy, we use a lookup table of proportional gains for each of five blocking levels: 0%, 10%, 20%, 30%, and 40%. When the number blocked changes, we update K in (5) by linearly interpolating between the nearest values in the lookup table.

V. METHODS: STRATEGY II

A. Mode-Count Control

For the operator in Strategy II, we propose the use of mode-count control to ensure network safety. Mode-count control was originally developed for TCLs in [34] and was extended to account for TCLs with lockout in [23]. In this paper, we propose two applications of mode-count control: 1) prevention of under-voltages or over-currents by reducing a TCL group's maximum demand, and 2) prevention of over-voltages by increasing a TCL group's minimum demand. For brevity, we present only the first application; the methods for the second application follow directly.

For a group of co-located TCLs, the policy determines which TCLs to switch *on/off* to constrain the group's "on-count", and thereby the group's demand. A group's on-count is the number of TCLs in the group that are *on* at a particular time and is denoted H_t^j for group j at time t . We constrain a group's on-count such that $H_t^j \leq \bar{H}^j$, where $\bar{H}^j$ is a low (but feasible) upper bound.

Figure 3 provides an example of the control policy. Consider a group of 7 TCLs with upper bound $\bar{H} = 3$. The central idea of the policy is to switch TCLs *on* as soon as possible once they have entered the "upper margin" of their temperature range. At the time step shown in Fig. 3, the policy switches TCL #4 *on* because it has just entered its upper margin, and switches TCL #1 *off* to satisfy the counting constraint. A main goal of the policy is to avoid having too many TCLs locked *on* at the same time, which could happen if we waited to switch TCLs at their upper temperature limit.

We summarize the steps of the control policy in Algorithm 1 and define three terms used within the algorithm as follows. (For a more detailed treatment of the policy, we refer the reader to [23].) First, a TCL's "time-to-upper-limit" is the time it would take a TCL to progress from its current temperature θ_t^i to its upper limit $\bar{\theta}^i$ in the *off* direction. This time can be explicitly calculated from the individual TCL model, and for the i th TCL is given by $t_{\text{UL}}^i = r^i c^i \ln((\theta_a - \theta_t^i)/(\theta_a - \bar{\theta}^i))$. Second, a TCL's "upper margin temperature" $\bar{\theta}_L^i$ is the lesser of two temperatures: 1) the temperature from which it takes the TCL τ_L time steps to reach $\bar{\theta}^i$ when *off*, and 2) the temperature reached by the TCL τ_L time steps after leaving $\bar{\theta}^i$ when *on*. Third, a TCL "counter-switches" with another TCL by switching at the same time but in the opposite direction (i.e., to the opposite power status).

To prevent the violation of a particular network constraint, an operator must first identify the group of TCLs to control and then set a lower bound on the group's on-count. As with blocking, we use an iterative offline method to identify the groups of TCLs to

Algorithm 1 Mode-Count Control for Upper Bound

In each time step,

- 1) Initialize $\Delta H = \bar{H} - H_t$, and find the *off* unit in the upper margin with shortest time-to-upper-limit. Let this unit's index be g .
 - 2) If unit g is *unlocked*:
 - a) If $\Delta H > 0$, then switch unit g *on* and decrement ΔH .
 - b) Otherwise, find a TCL that is available to counter-switch; this TCL must be *on* and *unlocked*, must have $\theta_t^i < \bar{\theta}_L^i$, and must not be reserved to counter-switch with a different TCL. Let s be the index of the available TCL with the longest time-to-upper-limit. Switch unit g *on* and unit s *off*.
 - 3) If unit g is *locked*:
 - a) If possible, find a TCL that is available to counter-switch. Reserve this TCL for counter-switching with g in the future.
 - b) Otherwise if $\Delta H > 0$, decrement ΔH so that unit g will be able to switch in the future, as soon as it is unlocked.
 - 4) Find the *off* unit with the next shortest time-to-upper limit. Let this unit's index be g .
 - 5) Repeat 2)-4) until either there are no more *off*, *unlocked* units in the upper margin, or $\Delta H = 0$ and there are no more *on* TCLs available for counter-switching.
-

Algorithm 2 Individual-Model Based Tracking Control

In each time step,

- 1) Calculate $\Delta P_{\text{track},t}$, the change in power needed to track $P_{\text{total},t+1}^*$, taking into account changes in power due to internal switching and, if there is coordination, safety control.
 - 2) Update priority stacks $\mathbb{S}_{\text{on},t}$ and $\mathbb{S}_{\text{off},t}$.
 - 3) If $\Delta P_{\text{track},t} \geq P_{\text{small}}$
 - a) Find index $j^* \in \mathbb{S}_{\text{off},t}$ that minimizes $|\sum_{i=1}^j p^i - \Delta P_{\text{track},t}|$
 - b) Switch *on* unit j^* and all units of higher priority in $\mathbb{S}_{\text{off},t}$
 - 4) Otherwise, if $\Delta P_{\text{track},t} \leq -P_{\text{small}}$
 - a) Find index $j^* \in \mathbb{S}_{\text{on},t}$ that minimizes $|\sum_{i=1}^j p^i + \Delta P_{\text{track},t}|$
 - b) Switch *off* unit j^* and all units of higher priority in $\mathbb{S}_{\text{on},t}$
-

control for network safety. Once identified, we set each group's upper bound $\bar{H}^j$ to its lowest feasible value; our prior work [23] provides a conjecture on this value.

B. Individual-Model Based Tracking Control

For the aggregator in Strategy II, we propose a control policy based on individual TCL models. The general principal of the control policy is to switch TCLs that are closest—in terms of time—to being switched internally. The proposed controller is similar to the priority stack method of [10], but we use time to determine priority instead of temperature. We summarize the steps of the control policy in Algorithm 2 and discuss each step in the remainder of the section.

In step 1, we calculate the change in power that is needed to achieve tracking in the next time step. When there is no direct communication between aggregator and operator, this change in power is given by

$$\Delta P_{\text{track},t} = P_{\text{total},t+1}^* - (P_{\text{total},t} + \Delta P_{\text{internal}}); \quad (6)$$

see the end of this section for the case in which there is direct communication. Variable $\Delta P_{\text{internal}}$ is the change in power that will occur in the next time step due to internal switching. We determine $\Delta P_{\text{internal}}$ by predicting which TCLs will be switched by their internal controllers in the next time step. For each TCL, we calculate its time-to-upper-limit t_{UL} (as defined in Section V-A) and its time-to-lower-limit t_{LL} , defined as the time it would take a TCL to progress from its current temperature to its lower limit in the *on* direction. Both metrics can be solved for using (1). A TCL is predicted to switch in the next time step if it is *off* and has $t_{\text{UL}} < h$, or if it is *on* and has $t_{\text{LL}} < h$.

In step 2, we update the priority stacks. Priority stack $\mathbb{S}_{\text{on}}$ is composed of *on* and *unlocked* TCLs and is sorted by time-to-lower-limit; priority stack $\mathbb{S}_{\text{off}}$ is composed of *off* and *unlocked* TCLs and is sorted by time-to-upper-limit. In each time step, we update the priority stacks according to the switching actions from the last time step and newly calculated values for t_{UL} and t_{LL} . The set of TCLs under control by the operator are excluded from the priority stacks; we assume the aggregator is able to identify this set given its TCL measurements.

In step 3, we select which TCLs to switch. To prevent excessive switching, we only switch TCLs if the desired change in power has magnitude greater than a threshold, P_{small} . Parameter P_{small} can be thought of as a tuning parameter; here we set it to 25% of the smallest, individual power rating in the aggregation. If the threshold is reached, then in step 3a we calculate how much of the stack should switch to minimize the tracking error. In this calculation, we take the sum of the TCLs' power ratings in order of priority. In step 3b, the selected TCLs are switched.

Step 4 is similar to step 3, except ΔP_{track} is negative, so units need to be switched *off* rather than *on*.

Finally, when there is direct communication in Strategy II, the operator sends the aggregator ΔP_{safety} , the change in power its safety control actions will cause in the next time step. The aggregator compensates for the operator's actions by subtracting ΔP_{safety} from the right-hand side of (6). All other aspects of the control are the same.

TABLE IV
CASE STUDY: AGGREGATION AND NETWORK DETAILS

# of Primary Nodes	# of Houses	Baseline AC Load	Non-AC Load
613	2265	4.42 MW	2.01 MW

TABLE V
TUNING PARAMETERS

Strategy I:	$K = 1$	$\mathbf{Q} = \text{diag}(I, 10^2 I)$	$\mathbf{R} = 10^9 \text{diag}(1, 0)$
Benchmark:	$T_1 = 100$	$K_P = 3.4 \times 10^{-4}$	

$\text{diag}(\cdot)$ maps the input matrices to a block diagonal matrix

VI. CASE STUDY

A. Benchmark Strategy

In the case study, we compare the proposed control strategies to a benchmark strategy. The benchmark is identical to Strategy I, except the aggregator's tracking controller is a model-free proportional-integral (PI) controller. Therefore, comparison to the benchmark allows us to determine the benefit of the model-based controller and state estimator in Strategy I. The PI controller computes a scalar switching command, whose magnitude is the probability with which TCLs should switch and whose sign is the direction with which TCLs should switch: positive (*on*) or negative (*off*). We use a discrete-time PI control algorithm with anti-windup (see pp. 311-312 of [35]). The tuning parameters are a proportional gain K_P and an integral gain K_P/T_1 , where T_1 is the integral time constant. We set the anti-windup time constant equal to that of the integral gain. When there is direct communication, we use a look-up table for the gains.

B. Setup

We test the control strategies in simulations of a distribution network with a high-penetration of aggregator-controlled residential ACs. We assume 100% of houses on the network have ACs, and we use an outdoor temperature of 32°C in order to capture a peak-load hour. We compare the strategies' performances across 10 1-hour trials, where each trial uses a different segment of the RegD regulation signal from the PJM Interconnection [36]. We scale the signal such that its amplitude is +/-33% of the ACs' baseline power consumption. Good control performance is indicated by high accuracy in tracking and low prevalence of network violations.

The simulation model has two main parts: 1) the network and 2) the AC aggregation and controllers. We use a realistic network model and GridLAB-D [37] to solve the network's three-phase, unbalanced power flow. The network model is of an actual system: R1-12-47-1 in Pacific Northwest National Lab's (PNNL's) database [38], [39]. The network model includes dynamic voltage regulation in the form of one voltage regulator and two capacitor banks. During simulations, we measure current flow through lines, apparent power flow through transformers, and voltages at residential meters. Appendix C lists a few modifications made to the network model.

We model the AC aggregation and controllers in MATLAB. We determine the number of houses served by each residential transformer using a disaggregation method developed by PNNL [40]. Each house is composed of one AC model and one constant-power baseload. We use (1)-(3) to model each AC and generate a heterogeneous population by randomly selecting parameter values from the uniform distributions described in Table II. The lockout period for each AC is set to 2 minutes. A power factor of 0.97 is used for ACs and 0.95 for baseload. Table IV reports the AC aggregation's baseline power consumption, as well as other network and load metrics.

For Strategy I's aggregate load model, we use five temperature intervals, i.e., $N_1 = 5$, and identify the matrix $\mathbf{A}_s$ by counting bin transitions within historical data of TCLs' temperatures and on/off statuses. For Strategy II's individual TCL models, we assume perfect models (i.e., the models are identical to those in the simulated plant); in future work, we plan to identify the individual TCL models from historical measurements.

We tune the tracking controllers of Strategy I and the benchmark strategy for the best average tracking performance over a set of tuning trials. When there is no direct communication, we use 50 tuning trials consisting of 25 different hours of the RegD signal and two different blocking levels: 0% and 20%. Table V lists the resulting tuned parameters. When there is direct communication, we tune the controllers for each blocking level independently. The tuned parameters have the following ranges of values: $K \in [1, 1.15]$, $K_P \in [3 \times 10^4, 5 \times 10^5]$, and $T_1 \in [80, 800]$, where the lower and upper values correspond to 0% and 40% blocked, respectively.

C. Results

Table VI shows the results of the case study, with values averaged over the 10 test trials. For all strategies, the safety controllers are successful—*no network violations occur*. Without safety control, there are an average of 2.2 over-voltage nodes and 1.2 overloaded lines, where the averages are taken over all trials and strategies. The number of ACs under safety control varies

TABLE VI
CASE STUDY RESULTS: AVERAGES OVER 10 TRIALS

Strategy	ACs under Safety Control		Tracking Error (RMS)	
	No Comm.	Comm.	No Comm.	Comm.
Benchmark	13.8%	10.3%	1.37%	1.28%
Strategy I	15.3%	15.2%	0.72%	0.70%
Strategy II	0.4%	0.4%	0.10%	0.10%

“Comm.” indicates direct communication between aggregator and operator

considerably, depending on the strategy used. In the benchmark strategy, 13.8% of ACs are blocked, on average, in order to avoid network violations; this number falls slightly to 10.3% when there is direct communication. Similarly, in Strategy I, 15.3% of ACs are blocked, and 15.2% with direct communication. In contrast, Strategy II’s mode-count controller requires only 0.4% of the aggregation to prevent network violations, with and without direct communication.

In terms of tracking error, Strategy II outperforms Strategy I, and both outperform the benchmark strategy. In Table VI, root-mean-square (RMS) error in tracking is reported as a percentage of the baseline power of the aggregation. Comparing against the benchmark, we find that the tracking error of Strategy I (0.72%, and 0.70% with direct communication) is almost 50% less than that of the benchmark (1.37%, and 1.28% with direct communication), and the tracking error of Strategy II (0.10% with and without direct communication) is more than an order of magnitude less than that of the benchmark. In all strategies, tracking performance improves given direct communication; we note that for Strategy II this improvement is not shown by Table VI due to rounding. To put these numbers in context, we can compare them to those reported in the literature. For example, [6] reported RMS errors of 0.2–9% depending upon the communication scenario and [41] reported RMS errors of 0.5–28% depending upon the number of TCLs participating, though neither study considered distribution network constraints.

Although Strategy II outperforms Strategy I in terms of tracking, Strategy I’s advantage is in ease of implementation. Strategy I is more computationally scalable than Strategy II because Strategy I’s aggregate model is independent of aggregation size whereas Strategy II uses individual TCL models and therefore the model increases linearly with aggregation size. In the case study, we expect the benchmark strategy to be the fastest since it is model-free, Strategy I to be the next fastest with its 20-state model, and Strategy II to be the slowest with its 4530-state model. All strategies were run in MATLAB using a 4-core machine (3.30 GHz) with 32 GB of RAM. Over the 10 trials without coordination, the average time to compute commands for the next time step was 0.045×10^{-4} s for the benchmark strategy, 1.5×10^{-4} s for Strategy I, and 7.9×10^{-4} s for Strategy II. All of the strategies are very fast given that the frequency regulation signal changes every two seconds. We expect Strategy I’s computational advantages over Strategy II to be more pronounced with larger sized aggregations. Recall that Strategy I also requires substantially lower measurement and communication requirements than Strategy II (see Table III).

In Fig. 4, we demonstrate a few of Strategy I’s key features using data from a selected trial. In the upper plot, Strategy I’s aggregate-model based controller closely tracks the regulation signal, despite a large percentage—in this trial, 38%—of ACs being blocked. However, the tracking is not perfect: there is noticeable error in the 6th minute, when the aggregation has saturated (i.e., not enough ACs are available to switch *off* due to blocking or lockout). The middle plot of Fig. 4 demonstrates an attribute of the state estimator: for bins that frequently receive external switching commands (i.e., bins N_1 and $2N_1$), the state estimate corresponds to the number of ACs that are available to switch (i.e., unblocked ACs), rather than the total number of ACs in the bin. This helps the controller compensate for blocking. The lower plot of Fig. 4 demonstrates the effects of blocking. Without blocking, phase C of overhead line 301 experiences an over-current in minutes 7 and 8; blocking a portion of the ACs on phase C effectively reduces this over-current.

Fig. 4 also highlights the main drawback of using blocking for safety control: a relatively large percentage of ACs must be blocked to protect just a few network constraints. This is because, unlike mode-count control, blocking does not actively constrain a group’s maximum demand. Instead, blocking returns a group to its normal variations in demand and level of load diversity, where there is more load diversity with larger groups.

In Fig. 5, we demonstrate a few of Strategy II’s key features using the same selected trial. The top plot shows the high-accuracy tracking of Strategy II’s individual-model based controller. Compared to Strategy I, Strategy II has very few ACs under safety control—only 0.4% in this trial. Strategy II also results in fewer switching actions than Strategy I, and therefore, fewer locked ACs. Both of these features make Strategy II less prone to saturation, as can be seen by its accurate tracking in the 6th minute. The middle and lower plots of Fig. 5 demonstrate Strategy II’s mode-count controller. Without mode-count control, an under-voltage occurs at meter 469 between minutes 15–18. With mode-count control, the three AC’s connected to meter 469 are constrained to have an on-count ≤ 2 , and this prevents the under-voltage from occurring.

In Fig. 6, we show the benchmark strategy’s tracking accuracy for the same trial as Figs. 4 and 5. The plot shows noticeable tracking error in minutes 6, 8, 29, and 32, much more than what occurs in Strategy I or Strategy II. Similar to Strategy I, a large percentage of ACs must be blocked in this trial to ensure network constraints are satisfied.

Overall, we find that tracking performance improves as a control strategy uses more information. In this study, additional information comes in several forms: estimation, modeling, and communication. Strategy I uses state estimation and modeling and outperforms the benchmark strategy. Strategy II uses more detailed models and more communication than the other strategies,




Fig. 4. Strategy I during a selected trial. The plots shows the strategy's high tracking accuracy (top plot) and state-estimation performance for bin 5 (middle plot). The bottom plot shows that blocking a large portion of TCLs can prevent over-currents, here on overhead line # 301.




Fig. 5. Strategy II during a selected trial. Plot (a) shows the strategy's very high tracking accuracy. The plots in (b) shows that mode-count control can maintain a group of 3 TCLs' on-count to ≤ 2 (top plot) and thereby prevent under-voltages (bottom plot), here for residential meter # 469.




Fig. 6. Benchmark strategy during a selected trial. At minutes 6, 8, 29, and 32, there is noticeable tracking error, more so than in Strategies I and II.

and has the best tracking performance. Finally, all three strategies perform better when there is direct communication between the operator and aggregator. However, information and direct communication come at a cost. Furthermore, Strategy II requires more computation than Strategy I, which requires more computation than the benchmark strategy, and computation also comes at a cost. In practice, aggregators and operators will need to analyze these trade-offs to determine which strategy best suits their performance/cost requirements.

VII. CONCLUSIONS

We have proposed two strategies for network-safe load control; in the strategies, a third-party aggregator controls TCLs to provide frequency regulation, while the operator overrides the aggregator's control when necessary to ensure network safety. The two control strategies differ markedly in terms of performance and ease of implementation. Strategy II substantially outperforms Strategy I. Its safety controller is able to prevent network violations using less than 3% of the ACs used by Strategy I. This reduction in the number of safety ACs, in combination with a more information-rich tracking controller, results in improved tracking performance for Strategy II. However, Strategy II would likely be more difficult to implement. Compared with Strategy I, it is less scalable in terms of modeling and communication requirements. Choosing between these alternative control strategies will require balancing the costs of implementing the more intensive strategy against the benefits of its improved performance.

APPENDIX A

METHOD FOR SELECTING TCLS TO BLOCK OR CONTROL FOR SAFETY

We first determine which constraints are at risk by running a preliminary simulation of the network with the TCL aggregation controlled to provide frequency regulation. Constraints that experience at least one violation during the simulation are considered at risk. Next, we select which TCLs to block to relieve these at-risk constraints; this selection process is described in detail in the next paragraph. Then we re-run the simulation with the selected TCLs blocked and check if all at-risk constraints have been relieved and if any additional violations have occurred. If there continue to be violations, additional TCLs are added to the set of blocked TCLs and the simulation is re-run. This process is repeated until no violations occur. We note that this offline method assumes advanced knowledge of the regulation signal and a perfect model of load behavior for the next hour. Thus, the method's selections (of which TCLs to block) should be interpreted as best-case selections. An online method for selecting TCLs would likely result in more TCLs being blocked and may not successfully avoid all constraint violations.

The method that we use for selecting TCLs depends on the type of constraint that is at risk. If a line or transformer are overloaded, then we incrementally block TCLs that are downstream from the component according to how much they contribute to the overload. We assume the network is radial and therefore a TCL's contribution can be approximated as the sum of its demand and associated line losses. We assume a TCL's associated line losses are proportional to the length of line between the TCL and the overloaded component. Thus, we block TCLs that are farthest downstream of the component first. Note that differences in line resistances and differences in TCL power ratings are neglected. If a service node has a voltage violation, then we incrementally block TCLs according to how much they contribute to the voltage drop between the constrained node and the substation. For loads located along the line segment between the constrained node and the substation, the loads' contributions to the voltage drop can be approximated by their distance along the line segment, with loads located farther from the substation contributing more (see Ch. 3, Section 4.1 of [42]). Thus, TCLs located at the constrained node, or downstream of the node, contribute the most to the voltage drop and are blocked first. If the violation persists, then we incrementally block TCLs along the line segment starting at the constrained node and moving toward the substation.

APPENDIX B

CONTROL POLICY DERIVATION

We derive the policy by solving the aggregate model (4) for $\mathbf{u}_t$ given state estimate $\hat{\mathbf{x}}_t$. To achieve the desired output $\mathbf{y}_{t+1}^*$, we choose $\mathbf{u}_t$ such that $\mathbf{y}_{t+1}^* = \mathbf{C}\mathbf{A}_{u,t}\mathbf{A}_s\hat{\mathbf{x}}_t$. The matrix $\mathbf{A}_{u,t}$ can be decomposed into two matrices: $\mathbf{A}_{u,t} = \tilde{\mathbf{A}}_{u,t} + \mathbf{I}$, and therefore $\mathbf{y}_{t+1}^* = \mathbf{C}\tilde{\mathbf{A}}_{u,t}\mathbf{A}_s\hat{\mathbf{x}}_t + \mathbf{C}\mathbf{A}_s\hat{\mathbf{x}}_t$. Given the structure of $\mathbf{A}_{u,t}$, the second row of this equation is satisfied by any choice of $\mathbf{u}_t$. The first row of this equation is

$$P_{\text{total},t+1}^* = \bar{p}_{\text{on}}N \left(\sum_{n=1}^{N_1} u_t^n x_{\text{pred},t}^n + \sum_{n=N_1+1}^{2N_1} -u_t^n x_{\text{pred},t}^n \right) + \mathbf{C}^1 \mathbf{A}_s \hat{\mathbf{x}}_t, \quad (7)$$

where $P_{\text{total},t+1}^*$ is the first entry of $\mathbf{y}_{t+1}^*$ and $\mathbf{x}_{\text{pred},t} = \mathbf{A}_s \hat{\mathbf{x}}_t$.

We design the control policy such that, in each time step, TCLs are switched in only one direction (*on* or *off*). The entries of $\mathbf{u}_t$ are probabilities and must be between 1 and 0. Given these restrictions, when a positive change in power is needed, (7) can be simplified and rearranged such that

$$\sum_{n=1}^{N_1} u_t^n x_{\text{pred},t}^n = K \frac{P_{\text{total},t+1}^* - \mathbf{C}^1 \mathbf{A}_s \hat{\mathbf{x}}_t}{\bar{p}_{\text{on}}N}, \quad (8)$$

where K has been added as a proportional gain. Similarly, when a negative change in power is needed, the last N_I entries of $\mathbf{u}_t$ must satisfy

$$\sum_{n=N_I+1}^{2N_I} -u_t^n x_{\text{pred},t}^n = K \frac{P_{\text{total},t+1}^* - \mathbf{C}^1 \mathbf{A}_s \hat{\mathbf{x}}_t}{\bar{p}_{\text{on}} N}, \quad (9)$$

and all other entries of $\mathbf{u}_t$ are set to zero.

APPENDIX C NETWORK MODIFICATIONS

We make the following five modifications to improve the accuracy of the network model: 1) we increase a distribution transformer's rating if its average load over an hour is higher than both its original planning load and its original rating; 2) we increase the size of triplex lines if their maximum current is larger than their rating; 3) we shift the regulation range of capacitor banks so that voltages stay within 0.95-1.05 p.u. during nominal operation; 4) we adjust the capacitor bank settings such that phases are individually controlled; 5) we set the initial conditions of capacitors and regulators so that they are at steady state condition when the test trials begin.

REFERENCES

- [1] 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>
- [2] —, "Wholesale Competition in Regions with Organized Electric Markets," Tech. Rep. Order No. 719, Oct. 2008. [Online]. Available: <https://www.federalregister.gov/documents/2008/03/07/E8-3984/wholesale-competition-in-regions-with-organized-electric-markets>
- [3] S. C. Ross, G. Vuylsteke, and J. L. Mathieu, "Effects of load-based frequency regulation on distribution network operation," *IEEE Trans. on Power Systems*, vol. 34, no. 2, pp. 1569–1578, 2019.
- [4] 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.
- [5] PJM Interconnection, "Implementation and Rationale for PJM's Conditional Neutrality Regulation Signals," Tech. Rep., Jan. 2017. [Online]. Available: <https://www.pjm.com/~media/committees-groups/task-forces/rmistf/postings/regulation-market-whitepaper.ashx>
- [6] 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.
- [7] W. Zhang, J. Lian, C.-Y. Chang, and K. Kalsi, "Aggregated Modeling and Control of Air Conditioning Loads for Demand Response," *IEEE Trans. on Power Systems*, vol. 28, pp. 4655–4664, 2013.
- [8] E. Can Kara, Z. Kolter, M. Berges, B. Krogh, G. Hug, and T. Yuksel, "A moving horizon state estimator in the control of thermostatically controlled loads for demand response," in *IEEE International Conf. on Smart Grid Comm.*, Oct. 2013.
- [9] N. Lu, "An evaluation of the HVAC load potential for providing load balancing service," *IEEE Trans. on Smart Grid*, vol. 3, no. 3, pp. 1263–1270, 2012.
- [10] H. Hao, B. Sanandaji, K. Poolla, and T. Vincent, "Aggregate flexibility of thermostatically controlled loads," *IEEE Trans. on Power Systems*, vol. 30, no. 1, pp. 189–198, Jan. 2015.
- [11] M. Almassalkhi, J. Frolik, and P. Hines, "Packetized energy management: Asynchronous and anonymous coordination of thermostatically controlled loads," in *2017 American Ctrl. Conf. (ACC)*, May 2017.
- [12] F. Ruelens, B. J. Claessens, S. Vandael, B. De Schutter, R. Babuška, and R. Belmans, "Residential demand response of thermostatically controlled loads using batch reinforcement learning," *IEEE Transactions on Smart Grid*, vol. 8, no. 5, pp. 2149–2159, 2017.
- [13] A. Radaideh, U. Vaidya, and V. Ajjarapu, "Sequential set-point control for heterogeneous thermostatically controlled loads through an extended markov chain abstraction," *IEEE Transactions on Smart Grid*, vol. 10, no. 1, pp. 116–127, 2019.
- [14] M. Liu, S. Peeters, D. S. Callaway, and B. J. Claessens, "Trajectory tracking with an aggregation of domestic hot water heaters: Combining model-based and model-free control in a commercial deployment," *IEEE Transactions on Smart Grid*, vol. 10, no. 5, pp. 5686–5695, 2019.
- [15] 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.
- [16] 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.
- [17] A. Hassan, R. Mieth, M. Chertkov, D. Deka, and Y. Dvorkin, "Optimal Load Ensemble Control in Chance-Constrained Optimal Power Flow," *IEEE Trans. on Smart Grid*, vol. 10, no. 5, pp. 5186–5195, 2018.
- [18] A. Bernstein and E. Dall'Anese, "Real-time feedback-based optimization of distribution grids: A unified approach," *IEEE Trans. on Control of Network Systems*, vol. 6, no. 3, pp. 1197–1209, 2019.
- [19] D. Molzahn and L. A. Roald, "Grid-aware versus grid-agnostic distribution system control: A method for certifying engineering constraint satisfaction," in *Proc. of the Hawaii Int. Conf. on Sys. Sci.*, Jan. 2019.
- [20] N. Nazir and M. Almassalkhi, "Convex inner approximation of the feeder hosting capacity limits on dispatchable demand," in *Proc. of the IEEE Conf. on Decision and Control (CDC)*, Dec. 2019.
- [21] —, "Grid-aware aggregation and realtime disaggregation of distributed energy resources in radial networks," *IEEE Trans. on Power Systems (in review)*, 2020.
- [22] S. C. Ross, N. Ozay, and J. L. Mathieu, "Coordination between an aggregator and distribution operator to achieve network-aware load control," in *IEEE PES PowerTech Conf.*, Jun. 2019.
- [23] S. C. Ross, P. Nilsson, N. Ozay, and J. L. Mathieu, "Managing voltage excursions on the distribution network by limiting the aggregate variability of thermostatic loads," in *American Ctrl. Conf.*, Jul. 2019.
- [24] 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.
- [25] IEEE, "IEEE Guide for Loading Mineral-Oil-Immersed Transformers," Tech. Rep. IEEE Std C57.91-1995(R2004), Dec. 2004.
- [26] R. Sonderegger, "Dynamic models of house heating based on equivalent thermal parameters," Ph.D. dissertation, Princeton Univ., N.J., 1978.
- [27] J. L. Mathieu, M. Dyson, and D. S. Callaway, "Using Residential Electric Loads for Fast Demand Response : The Potential Resource and Revenues, the Costs, and Policy Recommendations," in *Proc. of the ACEEE Summer Study on Buildings*, Aug. 2012.
- [28] S. Ross and J. Mathieu, "A method for ensuring a load aggregator's power deviations are safe for distribution networks," in *Proceedings of the Power Systems Computation Conference*, Jun. 2020.
- [29] M. Liu and Y. Shi, "Model predictive control of aggregated heterogeneous second-order thermostatically controlled loads for ancillary services," *IEEE Trans. on Power Systems*, vol. 31, no. 3, pp. 1963–1971, 2016.

- [30] Y. Ji, E. Buechler, and R. Rajagopal, "Data-driven load modeling and forecasting of residential appliances," *IEEE Transactions on Smart Grid*, vol. 11, no. 3, pp. 2652–2661, 2020.
- [31] 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.
- [32] MATLAB. (2018) Kalman filtering. [Online]. Available: <https://www.mathworks.com/help/control/ug/kalman-filtering.html>
- [33] 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.
- [34] P. Nilsson and N. Ozay, "On a class of maximal invariance inducing control strategies for large collections of switched systems," in *Proc. of the Int. Conf. on Hybrid Sys.: Comp. and Ctrl. (HSCC)*, Apr. 2017.
- [35] K. Astrom and R. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press, 2008.
- [36] PJM Interconnection. Ancillary services. [Online]. Available: <http://www.pjm.com/markets-and-operations/ancillary-services.aspx>
- [37] GridLAB-D software. [Online]. Available: <http://www.gridlabd.org/>
- [38] Pacific Northwest National Laboratory. Taxonomy_Feeders. [Online]. Available: https://github.com/gridlab-d/Taxonomy_Feeders
- [39] K. Schneider, Y. Chen, D. Chassin, R. Pratt, D. Engel, and S. Thompson, "Modern Grid Initiative Distribution Taxonomy Final Report," Pacific Northwest National Laboratory, Nov. 2008.
- [40] Pacific Northwest National Laboratory. PopulationScript. [Online]. Available: https://github.com/gridlab-d/Taxonomy_Feeders/tree/master/PopulationScript
- [41] M. M. Olama, T. Kuruganti, J. Nutaro, and J. Dong, "Coordination and control of building hvac systems to provide frequency regulation to the electric grid," *Energies*, vol. 11, no. 7, p. 1852, 2018.
- [42] W. Kersting, *Distribution System Modeling and Analysis*, 3rd ed. Boca Raton, Florida: CRC Press, 2012.