Data-driven Network-aware Control of Thermostatically Controlled Loads with Unknown Dynamics

Sunho Jang, Johanna L. Mathieu, Necmiye Ozay

***Abstract*— This paper presents a novel data-driven control algorithm that coordinates a large aggregation of heterogeneous thermostatically controlled loads (TCLs) with unknown temperature dynamics and disturbance distributions to provide balancing services to the grid. Each TCL is subject to local constraints on their temperatures due to user comfort and global constraints on the number of TCLs in a mode, called mode-counting constraints, for the safe operation of the distribution network. We first develop a data-driven method to compute sets of modes at each discretized interval of the temperature state space in which the probability of moving outside of the temperature dead-band at the next step is bounded at a certain confidence level. Subsequently, we design a model predictive control (MPC) algorithm that instructs every TCL to switch to the modes within the obtained set while satisfying the mode-counting constraints. A case study compares the proposed control algorithm with a benchmark and verifies its effectiveness in maintaining end-user comfort and the safe operation of the distribution network.**

I. Introduction

Aggregations of distributed energy resources (DERs) can provide significant flexibility to the power grid, helping balance intermittent and uncertain renewable energy generation. DERs include distributed generation, distributed storage, and flexible loads. These resources could be coordinated by an electric utility company, or a third-party (non-utility) aggregator that participates in electricity markets. In the U.S., aggregators have recently been further enabled by FERC Order 2222 [1]. A challenge with third-party aggregators is that they do not have detailed knowledge of the distribution networks within which the DERs they coordinate operate.

Here, we consider a particular type of DER – Thermostatically Controlled Loads (TCLs) such as residential air conditioners, which switch on/off to maintain a temperature within a dead-band. TCLs have thermal energy storage capacity and so the timing of their power consumption can be slightly shifted, making them a great resource for energy-neutral grid balancing services such as frequency regulation [2]. While an aggregator is coordinating TCLs, its actions should not be disruptive to end-users. In particular, end-users choose a temperature set-point, which defines a temperature dead-band (a small temperature range around the set-point), and the aggregator should control each TCL's on/off mode such that its temperature stays within its dead-band. Furthermore, TCL coordination should not cause any issues in the distribution network such as voltage violations or current overloading, which is possible according to [3].

The aggregator's job is difficult if it does not have full information about the TCLs and the distribution network. In general, a utility cannot share full details of the network with the aggregator for security reasons [4]. FERC has noted that utility-aggregator coordination is a key challenge to DER integration [5]. Recent research has proposed a variety of strategies for this type of coordination [6]–[10], but the "best" strategy is still up for debate, and likely location- and context-dependent. In addition to not having full network information, the aggregator also has only partial information about the TCLs' temperature dynamics and disturbance distributions. It is difficult to obtain an accurate TCL model. The widely used first-order model fails to accurately represent both transient and steady-state dynamics [11]. While the second-order model in [12] is more accurate, some parameters are unidentifiable from input/output data [13]. Overall, these gaps in knowledge make it difficult for the aggregator to evaluate how its actions would affect the TCLs' temperature and the safety of the distribution network.

In this work, we propose a data-driven control algorithm for an aggregation of heterogeneous TCLs with unknown temperature dynamics and disturbance distributions. The approach is non-disruptive to the end-users and ensures the safe operation of the distribution network. We propose a data-driven approach to identify modes in which temperatures do not go outside of their dead-bands. Leveraging the obtained sets, we develop a model predictive control (MPC) algorithm that selects and limits the TCLs' aggregate power within network-safe power bounds obtained from the utility, thereby guaranteeing both the end-users' comfort and the safe operation of the distribution network.

Some prior work has proposed TCL coordination methods that are non-disruptive to the users [11], [14]–[17], as well as cycling constraints to prevent damage to TCL compressors [18]–[20]. While these papers only consider local constraints on temperatures of individual TCLs, [6] additionally considered global constraints on the TCLs' aggregate power for distribution network safety. Specifically, it proposed a heuristic TCL coordination algorithm to keep the number of turned-ON TCLs below a certain bound to promote safety. In addition, [10] proposed a TCL coordination framework that satisfies a chance constraint on voltage violations, ensuring distribution network safety with high probability. Ref. [20], [21] utilized formal abstractions of subsystems to design algorithms that ensure deterministic local constraints and mode-counting constraints for aggregate homogeneous

This work was supported by U.S. NSF Award CNS-1837680.
The authors are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI, USA {sunhoj,jlmath,necmiye}@umich.edu.

switched systems. All the works cited above assume precise knowledge of TCL temperature dynamics, though some of these works conduct simulation-based testing on aggregations of "noisy" TCLs, e.g., [6], [14]. Ref. [22] employed distributionally robust optimization to coordinate an aggregation of TCLs with unknown temperature dynamics. However, this paper focuses on homogeneous TCLs (though it does propose a method for extension to heterogeneous TCLs) and does not address distribution network safety. To the best of our knowledge, past work has not developed a network-safe aggregate TCL control approach that explicitly takes into account unknown temperature dynamics and disturbance distributions.

In contrast to previous work, the main contributions of this paper are three-fold. First, we introduce a network-safe control approach for aggregations of heterogeneous TCLs with unknown dynamics and disturbance distributions to provide grid balancing services with probabilistic guarantees on end-user comfort. Using realistic assumptions about TCL temperature dynamics, our data-driven method computes a set of on/off modes at each discretized interval of the temperature space in which the probability of TCLs remaining within the specified dead-band exceeds a certain threshold with a certain confidence level. Second, we develop an MPC algorithm that leverages these sets of modes to ensure both end-user comfort and distribution network safety. The proposed MPC algorithm is capable of managing large aggregations of TCLs since it does not scale with the number of TCLs. Third, our case study demonstrates that the proposed MPC algorithm effectively maintains temperatures within a specified range, while satisfying bounds on the aggregate power, which ensures network safety.

The organization of the paper is as follows. Section II introduces the problem formulation. Section III proposes a data-driven approach that constructs a set of admissible modes for local temperature constraint satisfaction. Section IV proposes a novel MPC algorithm, and Section V demonstrates the proposed controller through a case study. The appendix includes proofs of two theorems and one proposition.

Notation: Set $\mathbb{N}$ refers to the set of all natural numbers, $\mathbb{N}_0$ refers to $\{0\} \cup \mathbb{N}$, and $[N]$ refers to the set of natural numbers $\{1, \ldots, N\}$.

II. Problem Formulation

A. Overview

In this work, we address the problem of coordinating a large aggregation of heterogeneous TCLs to provide grid-balancing services, such as frequency regulation. We consider a control architecture previously explored in [20], as shown in Fig. 1. The framework includes three entities: aggregator, utility, and Independent System Operator (ISO). The aggregator is responsible for offering balancing capacity to the ISO, monitoring the states of the TCLs (temperatures and on/off modes), and controlling the TCLs' on/off modes. The aggregator also gathers historical data from each TCL, specifically temperature and mode trajectories, which are

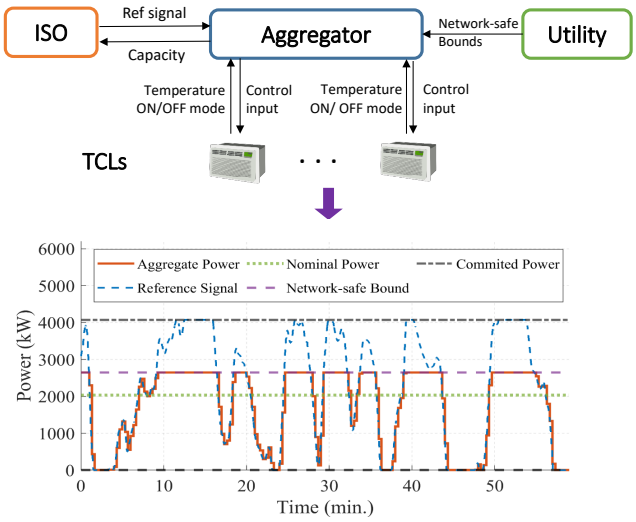

Fig. 1. Illustration of the coordination framework and plot of the aggregate power consumption of the TCLs given a reference signal from the ISO, scaled to be within the committed power (total procured power capacity of participating TCLs), and given network-safe power bounds from the utility. The nominal power is the average aggregate power when TCLs are not providing balancing services.

leveraged in controller construction. The ISO publishes a normalized reference signal for the balancing service, and the aggregator scales it by the committed capacity and shifts it by the nominal TCL aggregate power consumption. The aggregator should adjust the TCL aggregate power consumption to track the reference signal as closely as possible.

The utility, which has detailed distribution network information (e.g., network topology and parameters), is responsible for maintaining the network's safe operation. It uses network state forecasts to compute bounds on TCL aggregate power consumption that prevent constraint violations within the distribution network. Some previous studies, such as [23], have explored methods for calculating these bounds. If the aggregator's committed capacity remains within these bounds, no action is required from the utility.

However, when actual network states significantly deviate from their forecasts, the network-safe power bounds may shrink causing the aggregator's committed capacity to become larger than the bounds. In this case, the utility sends the aggregate power bounds to the aggregator and the aggregator should adjust the TCLs' aggregate power to lie within these bounds, while still tracking the reference signal as well as possible. The plot in Fig. 1 illustrates how the TCLs' aggregate power consumption should look given the reference signal from the ISO and the bounds from the utility.

B. Mathematical Formulation

We next provide the details of the problem formulation. We consider N heterogeneous TCLs, indexed by $i \in [N]$. Each TCL has a compact state space $\mathcal{T}^i := [\underline{T}^i, \overline{T}^i] \subseteq \mathbb{R}$ with two modes $\mathcal{M} := \{0, 1\}$, where the state corresponds to the TCL's internal temperature, i.e., for an air conditioner, the indoor air temperature, and modes 0 and 1 indicate OFF

and ON modes, respectively. We denote the state of the ith TCL by $T^i(t)$ and its mode by $m^i(t)$. The state of each TCL evolves according to

$$T^i(t+1) = f^i_{m^i(t)}\left(T^i(t)\right) + W^i_{m^i(t)}(t), \tag{1}$$

where $f^i_m : \mathcal{T}^i \to \mathcal{T}^i$ is the state evolution function of the ith TCL in mode m, and $W^i_m(t) \sim \mathcal{W}^i_m$ is the i.i.d. disturbance in mode m with probability distribution $\mathcal{W}^i_m$. It should be noted that the i.i.d. assumption may not fully capture real-world scenarios as the disturbances could exhibit temporal correlations as well as dependencies on temperature. Despite this limitation, this assumption is a practical simplification for TCL modeling.

While models of TCL dynamics exist, in practice these models are often inadequate at capturing the dynamical behavior of real TCLs [11], [12]. Thus, we assume that both f^i_m and $\mathcal{W}^i_m$ for all $m \in \mathcal{M}$ and $i \in [N]$ are unknown to the aggregator. In this paper, we assume that all TCLs are cooling TCLs (air conditioners) and the following assumption holds.

Assumption 1. *The state evolution functions f^i_m for all $i \in [N]$ satisfy the following three conditions.*

1) ***Monotonicity**: Each function f^i_m is monotonically increasing, i.e., $f^i_m(T_1) \le f^i_m(T_2)$ for all $T_1 \le T_2$.*
2) ***Lipschitz Continuity**: There exists a common constant L^f_m for any $m \in \mathcal{M}$:*

$$\left|f^i_m(T_1) - f^i_m(T_2)\right| \le L^f_m \cdot |T_1 - T_2|. \tag{2}$$

Assumption 1 is justifiable for real TCLs. The monotonicity condition captures the normal behavior of TCLs, specifically, if the initial temperature is lower than another temperature, the temperature at the next time step will not be higher than that associated with the other temperature under the same mode and disturbance. The Lipschitz continuity condition reflects the bounded rate of TCL temperature change. Assumption 1 holds for the standard affine temperature dynamics introduced in [12] with Lipschitz constant $L^f_m = 1$, which will be used for the case studies in Section V.

The state space of each TCL is composed of a *comfort set* $\mathcal{T}^{\mathrm{s},i} = [\underline{T}^{\mathrm{s},i}, \overline{T}^{\mathrm{s},i}]$ within which the temperature of the TCL should stay, and a *discomfort set* $\mathcal{T}^{\mathrm{d},i} := \mathcal{T}^i \setminus \mathcal{T}^{\mathrm{s},i} = [\underline{T}^i, \underline{T}^{\mathrm{s},i}) \cup (\overline{T}^{\mathrm{s},i}, \overline{T}^i]$ that every TCL should avoid. We define the *comfort condition*, which requires the temperature of a TCL to remain within $\mathcal{T}^{\mathrm{s},i}$. TCLs must turn OFF to heat up when the temperature falls below $\underline{T}^{\mathrm{s},i}$, and they must turn ON to cool down when the temperature goes above $\overline{T}^{\mathrm{s},i}$ to return to $\mathcal{T}^{\mathrm{s},i}$. Therefore, while the TCLs within $\mathcal{T}^{\mathrm{s},i}$ can be either in ON or OFF mode, the TCLs within $[\underline{T}^i, \underline{T}^{\mathrm{s},i})$ must be OFF and those within $(\overline{T}^{\mathrm{s},i}, \overline{T}^i]$ must be ON.

We denote the power consumption of each TCL in the ON mode by p^i and the upper and lower network-safe TCL aggregate power bounds (sent by the utility) by $\underline{p}$ and $\overline{p}$, respectively:

$$\underline{p} \le \sum_{i=1}^{N} p^i m^i(t) \le \overline{p} \qquad \forall t \in \mathbb{N}_0. \tag{3}$$

While the aggregator does not know the power consumption p^i of each TCL, we assume it has an estimate of the average power consumption p_{avg} of the TCL aggregation obtained from historical data. To satisfy (3), we impose the following *mode-counting constraints* on the number of ON TCLs:

$$\underline{N} = \left\lceil \frac{\underline{p}}{p_{\text{avg}}} \right\rceil \le \sum_{i=1}^{N} m^i(t) \le \overline{N} = \left\lfloor \frac{\overline{p}}{p_{\text{avg}}} \right\rfloor. \tag{4}$$

Note that, for heterogeneous TCLs, the mode-counting constraint serves as only an approximation of the actual constraints on the aggregate power (3). Consequently, the satisfaction of the mode-counting constraint does not necessarily ensure the satisfaction of (3). However, as the number of TCLs N increases, the approximation of the TCL aggregate power $p_{\text{avg}} \sum_{i=1}^{N} m^i(t)$ becomes more accurate due to the law of large numbers, leading to smaller violations of (3).

The control performance at each time step is measured by the difference between the aggregate power consumption $p_{\text{agg}}(t)$ of the TCLs and the scaled/shifted reference signal $r(t)$, i.e., $|p_{\text{agg}}(t) - r(t)|$. The objective of this work is to develop the aggregator's control algorithm to manipulate the ON/OFF modes of a large number of TCLs each represented by (1) to maximize control performance while trying to satisfy the comfort condition and satisfying the mode-counting constraints (4) to ensure network safety.

Remark 1. *The approach proposed in this paper, though specifically considering the coordination of an aggregation of TCLs, is applicable to a wide range of collections of switched subsystems that satisfy the monotonicity and Lipschitz continuity conditions in Assumption 1. It can even be extended to collections of multi-dimensional subsystems that exhibit monotonicity and Lipschitz continuity across every dimension.*

III. Data-driven identification of admissible modes ensuring the comfort condition

In this section, we develop a data-driven strategy to select the modes of the TCLs to satisfy the comfort condition. Specifically, we discretize the state spaces $\mathcal{T}^i$ and derive a set of modes for each discretized interval in which the probability of moving outside of the temperature dead-band in the next step is bounded at a certain confidence level.

We first uniformly discretize the comfort set $\mathcal{T}^{\mathrm{s},i}$ for all $i \in [N]$ into K^{s} intervals with length $\delta^i = (\overline{T}^{\mathrm{s},i} - \underline{T}^{\mathrm{s},i})/K^{\mathrm{s}}$ and denote the kth interval by $\mathcal{T}^i_k$:

$$\mathcal{T}^i_k = \begin{cases} [\underline{T}^{\mathrm{s},i} + (k-1)\,\delta^i, \underline{T}^{\mathrm{s},i} + k\delta^i) & \text{if } k \ne K^{\mathrm{s}} \\ [\underline{T}^{\mathrm{s},i} + (k-1)\,\delta^i, \underline{T}^{\mathrm{s},i} + k\delta^i] & \text{if } k = K^{\mathrm{s}}. \end{cases} \tag{5}$$

Furthermore, we discretize $\mathcal{T}^{\mathrm{d},i}$ with the same length δ^i and index them in order. Then, we denote by K the total number of intervals in the state space $\mathcal{T}^i$. Fig. 2 shows how the discretization is conducted and the intervals are indexed.

Suppose that the aggregator has a set of historical state transitions $\{T^i_{m,l}, T^{i+}_{m,l}\}^{D^i_m, N}_{l=1,i=1}$ from all TCLs in each mode m, where l indexes each data point, D^i_m is the number of data

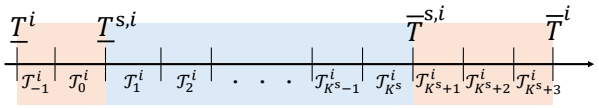

Fig. 2. Illustration of state space discretization. The blue area corresponds to $\mathcal{T}^{s,i}$ and the orange area corresponds to $\mathcal{T}^{d,i}$. Here, K is equal to K^s+5.

points obtained from TCL i in mode m, and $T^{i+}_{m,l}$ denotes the state of the ith TCL in mode m after its transition from $T^i_{m,l}$. We denote by $w^i_{m,l}$ the disturbance under which the transition from $T^i_{m,l}$ to $T^{i+}_{m,l}$ happens: $w^i_{m,l} := T^{i+}_{m,l} - f^i_m(T^i_{m,l})$.

Based on the data, we identify modes for each discretized interval $\mathcal{T}^i_k$ where the probability of comfort condition's satisfaction at the next time step is assured. In particular, in Lemma 1, for each mode and each disturbance $w^i_{m,l}$, we derive lower and upper bounds on the possible temperatures at the next time step given the starting temperature is within $\mathcal{T}^i_k$. Subsequently, Theorem 1 demonstrates that if, for a given mode, these bounds lie within the comfort set $\mathcal{T}^{s,i}$ for all $l \in [D^i_m]$, the probability of satisfaction of the comfort condition when starting from $\mathcal{T}^i_k$ with that mode is guaranteed with a certain confidence level.

Lemma 1. *Define* $\underline{T}^{i+}_{m,k,l}$ *and* $\overline{T}^{i+}_{m,k,l}$ *as*

$$\begin{aligned}\underline{T}^{i+}_{m,k,l} &:= T^{i+}_{m,l} - (L^f_m(T^i_{m,l} - \underline{T}^i_k))^+ \\ \overline{T}^{i+}_{m,k,l} &:= T^{i+}_{m,l} + (L^f_m(\overline{T}^i_k - T^i_{m,l}))^+.\end{aligned} \tag{6}$$

Then, for any starting temperature $T \in \mathcal{T}^i_k$*, the temperature of TCL* i *at the next time step under disturbance* $w^i_{m,l}$ *is within* $[\underline{T}^{i+}_{m,k,l}, \overline{T}^{i+}_{m,k,l}]$.

Proof. Suppose that $\tilde{T}^{i+}_l := f^i_m(T) + w^i_{m,l}$ is the state at the next time step given starting temperature T, mode m, and disturbance sample $w^i_{m,l}$. Then, we find

$$\begin{aligned}\tilde{T}^{i+}_l &= \left(f^i_m\left(T^i_{m,l}\right) + w^i_{m,l}\right) + f^i_m(T) - f^i_m\left(T^i_{m,l}\right) \\ &= T^{i+}_{m,l} + f^i_m\left(T\right) - f^i_m\left(T^i_{m,l}\right).\end{aligned} \tag{7}$$

In addition, we obtain the following inequality from the monotonicity and Lipschitz continuity conditions in Assumption 1:

$$f^i_m(T) - f^i_m\left(T^i_{m,l}\right) \geq \begin{cases} -L^f_m\left(T^i_{m,l} - T\right) & \text{if } T \leq T^i_{m,l} \\ 0 & \text{otherwise.}\end{cases}$$

Thus, we obtain the following for any $T \in \mathcal{T}^i_k$:

$$\begin{aligned}f^i_m(T) - f^i_m(T^i_{m,l}) &\geq -\left(L^f_m\left(T^i_{m,l} - T\right)\right)^+ \\ &\geq -\left(L^f_m\left(T^i_{m,l} - \underline{T}^i_k\right)\right)^+.\end{aligned} \tag{8}$$

Similarly, we obtain

$$\left(L^f_m\left(\overline{T}^i_k - T^i_{m,l}\right)\right)^+ \geq f^i_m(T) - f^i_m\left(T^i_{m,l}\right). \tag{9}$$

Finally, from (7)-(9), we obtain

$$\tilde{T}^{i+}_l \geq T^{i+}_{m,l} - (L^f_m(T^i_{m,l} - \underline{T}^i_k))^+ \tag{10}$$

$$\tilde{T}^{i+}_l \leq T^{i+}_{m,l} + (L^f_m(\overline{T}^i_k - T^i_{m,l}))^+. \tag{11}$$

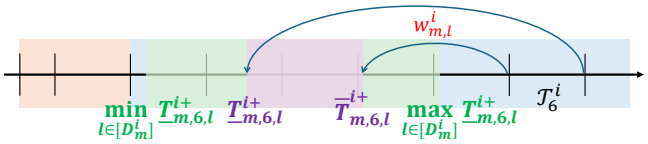

Fig. 3. The purple region represents the range of the temperatures at the next time step from $\mathcal{T}^i_6$ in mode 1 under disturbance $w^i_{m,l}$. The green regions represent the union of all purple regions generated for all $l \in [D^i_m]$. Given that the green region is within $\mathcal{T}^{s,i}$ (blue region), any state from $\mathcal{T}^i_6$ cannot move to $\mathcal{T}^{d,i}$ in mode 1 under any disturbance among $w^i_{m,1}, \ldots, w^i_{m,D^i_m}$.

Thus, the lemma is proven. □

Now, we provide a guarantee on the probability of the satisfaction of the comfort condition if a TCL chooses a mode m in which the range $[\underline{T}^{i+}_{m,k,l}, \overline{T}^{i+}_{m,k,l}]$ is within the comfort set $\mathcal{T}^{s,i}$ for all data points $l \in [D^i_m]$. Let $P^{d,i}_m(T)$ be the probability of the ith TCL moving to $\mathcal{T}^{d,i}$ from a starting temperature T in mode m, that is:

$$P^{d,i}_m(T) := \Pr\left(f^i_m(T) + W^i_m \in \mathcal{T}^{d,i}\right).$$

Then, we can obtain the following theorem.

Theorem 1. *Assume that the following inequalities hold:*

$$\underline{T}^{s,i} \leq \min_{l\in[D^i_m]} \underline{T}^{i+}_{m,k,l} \tag{12}$$

$$\overline{T}^{s,i} \geq \max_{l\in[D^i_m]} \overline{T}^{i+}_{m,k,l}, \tag{13}$$

where $\underline{T}^{i+}_{m,k,l}$ *and* $\overline{T}^{i+}_{m,k,l}$ *are defined in* (6)*. Then, for any* $\epsilon \in [0,1]$ *and* $T \in \mathcal{T}^i_k$*, the probability* $P^{d,i}_m(T)$ *that TCL* i *moves to the discomfort set* $\mathcal{T}^{d,i}$ *at the next time step is lower than* ϵ *with a confidence level over* $1 - 2(1-\epsilon/2)^{D^i_m}$.

Proof. By Lemma 1, if (12) and (13) hold, the range of possible temperatures at the next time step from $\mathcal{T}^i_k$ under any disturbance among $w^i_{m,1}, \ldots, w^i_{m,D^i_m}$ is within the comfort set $\mathcal{T}^{s,i}$; this range is illustrated as the green region in Fig. 3. Hence, the minimum/maximum disturbance among $w^i_{m,1}, \ldots, w^i_{m,D^i_m}$ is above/below the disturbance required for any temperature $T \in \mathcal{T}^i_k$ to reach the lower/upper bound $\underline{T}^{s,i}/\overline{T}^{s,i}$ of the dead-band at the next time step in mode m, that is:

$$\min_{l\in[D^i_m]} w^i_{m,l} \geq \underline{T}^{s,i} - f^i_m\left(T\right) \tag{14}$$

$$\max_{l\in[D^i_m]} w^i_{m,l} \leq \overline{T}^{s,i} - f^i_m\left(T\right). \tag{15}$$

From (14) and (15), we can derive an expression for the probability of moving to the discomfort set

$$\begin{aligned}P^{d,i}_m(T) &= \Pr\left(f^i_m(T) + W^i_m \in \mathcal{T}^{d,i}\right) \\ &= \Pr\left(W^i_m \leq \underline{T}^{s,i} - f^i_m(T) \bigvee W^i_m \geq \overline{T}^{s,i} - f^i_m\left(T\right)\right) \\ &\leq \Pr\left(W^i_m \leq \min_{l\in[D^i_m]} w^i_{m,l} \bigvee W^i_m \geq \max_{l\in[D^i_m]} w^i_{m,l}\right) \\ &\leq \Pr\left(W^i_m \leq \min_{l\in[D^i_m]} w^i_{m,l}\right) + \Pr\left(W^i_m \geq \max_{l\in[D^i_m]} w^i_{m,l}\right)\end{aligned}$$

$$= F_m^i\left(\min_{i\in[D_m^i]} w_{m,l}^i\right) + \left(1 - F^i\left(\max_{l\in[D_m^i]} w_{m,l}^i\right)\right), \quad (16)$$

where F_m^i is the cdf of W_m^i.

Suppose that $W_{m,1}^i, \ldots, W_{m,D_m^i}^i$ are D_m^i i.i.d. samples of W_m^i and $\underline{W}_m^i$ and $\overline{W}^i{}_m$ are defined as the minimum and maximum of these samples

$$\underline{W}_m^i := \min_{l\in[D_m^i]} W_{m,l}^i, \quad \overline{W}^i{}_m := \max_{l\in[D_m^i]} W_{m,l}^i. \quad (17)$$

Then, for any $\epsilon_1 \in [0,1]$, the probability that the cdf evaluated at the minimum of the samples $F_m^i(\underline{W}_m^i)$ is upper bounded by ϵ_1 is given by

$$\begin{aligned}
&\Pr\left(F_m^i(\underline{W}_m^i) \le \epsilon_1\right) \\
&\quad = 1 - \Pr\left(\min_{l\in[D_m^i]} W_{m,l}^i \ge \left(F_m^i\right)^{-1}(\epsilon_1)\right) \\
&\quad = 1 - \Pi_{l=1}^{D_m^i} \Pr\left(W_{m,l}^i \ge \left(F_m^i\right)^{-1}(\epsilon_1)\right) \\
&\quad = 1 - (1-\epsilon_1)^{D_m^i}.
\end{aligned} \quad (18)$$

Similarly, for any $\epsilon_2 \in [0,1]$, the probability that the cdf evaluated at the maximum of the samples $F_m^i(\overline{W}^i{}_m)$ is lower bounded by $1-\epsilon_2$ is given by

$$\begin{aligned}
&\Pr\left(F_m^i(\overline{W}^i{}_m) \ge 1 - \epsilon_2\right) \\
&\quad = 1 - \Pr\left(\max_{l\in[D_m^i]} W_{m,l}^i \le \left(F_m^i\right)^{-1}(1-\epsilon_2)\right) \\
&\quad = 1 - \Pi_{l=1}^{D_m^i} \Pr\left(W_{m,l}^i \le \left(F_m^i\right)^{-1}(1-\epsilon_2)\right) \\
&\quad = 1 - (1-\epsilon_2)^{D_m^i}.
\end{aligned} \quad (19)$$

Finally, replacing ϵ_1 and ϵ_2 in (18) and (19) with$\epsilon/2$, the probability that $F_m^i(\underline{W}_m^i) + 1 - F_m^i(\overline{W}^i{}_m)$ is upper bounded by an arbitrary $\epsilon \in [0,1]$ is

$$\begin{aligned}
&\Pr\left(F_m^i(\underline{W}_m^i) + 1 - F_m^i\left(\overline{W}^i{}_m\right) \le \epsilon\right) \\
&\ge \Pr\left(F_m^i\left(\underline{W}_m^i\right) \le \frac{\epsilon}{2}\right) + \Pr\left(1 - F_m^i\left(\overline{W}^i{}_m\right) \le \frac{\epsilon}{2}\right) \\
&\ge 1 - 2\left(1 - \frac{\epsilon}{2}\right)^{D_m^i}.
\end{aligned} \quad (20)$$

Note that $\min_{l\in[D_m^i]} w_{m,l}^i$ and $\max_{l\in[D_m^i]} w_{m,l}^i$ are realizations of $\underline{W}_m^i$ and $\overline{W}^i{}_m$, respectively. Hence, from (16) and (20), $P_m^{\mathrm{d},i}(T) \le \epsilon$ for any $T \in \mathcal{T}_k^i$ with a confidence level over $1 - 2\left(1-\epsilon/2\right)^{D_m^i}$, which proves the Theorem. □

Theorem 1 uses the fact that the comfort condition is ensured in mode m under a set of D_m^i i.i.d. disturbance samples to derive a guarantee on the probability of moving to the discomfort set $\mathcal{T}^{\mathrm{d},i}$. Note that the confidence level in Theorem 1 exponentially converges to 1 as the number of data points D_m^i grows so we can achieve a high confidence level even with a small number of data points; for $\epsilon = 0.01$, the confidence level is 0.987 with only $D_m^i = 1000$ data points.

Now, we define sets of *admissible modes* $\mathcal{M}_k^{\mathrm{s}}$ for each index $k \in [K^{\mathrm{s}}]$ as

$$\mathcal{M}_k^{\mathrm{s}} = \{m \mid m \text{ satisfies (12) and (13) for all } i \in [N]\},$$

and sets of *blocked modes* $\mathcal{M}_k^{\mathrm{d}} := \mathcal{M} \backslash \mathcal{M}_k^{\mathrm{s}}$. By Theorem 1, a TCL's probability of moving to the discomfort set is bounded by ϵ with confidence level specified in Theorem 1 if it (or the aggregator) selects its mode $m^i(t)$ from $\mathcal{M}_k^{\mathrm{s}}$. Therefore, we can ensure the probability of satisfying the comfort condition at a high confidence level by blocking modes in $\mathcal{M}_k^{\mathrm{d}}$. This strategy is adopted in the MPC algorithm we propose in the next section.

IV. Mode-Blocking MPC

We next develop an MPC algorithm, referred to as Mode-Blocking MPC, that is designed to ensure the satisfaction of both the comfort condition and the mode-counting constraints (4). We use a model of the temperature dynamics of a TCL aggregation, which has been developed/used in many prior works, e.g., [14], [20]. Subsequently, we propose an MPC algorithm that incorporates the sets of blocked modes $\mathcal{M}_k^{\mathrm{d}}$ to ensure the comfort condition at a certain probability with a certain confidence level while also satisfying the mode-counting constraint (4).

To create the model, suppose that the aggregator has a set of historical state trajectories $\{T_m^i(t)\}_{t=0,i=1}^{T_{\mathrm{D}},N}$ from all TCLs in each mode, where T_{D} is the length of the state trajectory horizon. Note that the same historical data could be used for these state trajectories and the historical state transitions used in Section III. The empirical transition probability from $\mathcal{T}_j^i$ to $\mathcal{T}_k^i$ is $\hat{P}_{kj}^m := D_{kj}^m / D_j^m$, where D_j^m is the number of data points such that $T_m^i(t) \in \mathcal{T}_j^i$ and D_{kj}^m is the number of the data points such that $T_m^i(t+1) \in \mathcal{T}_k^i$.

We leverage an approximation of the evolution of the TCLs' state distributions over the discretized intervals using $\hat{P}_{kj}^m$. Define state $x_{m,k}(t)$ as the number of TCLs at $\mathcal{T}_k^i$ in mode m, and input $u_{m,m',k}(t)$ as the number of TCLs at $\mathcal{T}_k^i$ switching mode from m to m' at time step t. We denote by $\boldsymbol{x}(t)$ and $\boldsymbol{u}(t)$ the vectors composed of $x_{m,k}(t)$ and $u_{m,m',k}(t)$, respectively.

Subsequently, we construct the aggregate model from the empirical transition probabilities to approximate the evolution of the TCLs' distribution among the discretized intervals

$$\boldsymbol{x}(t+1) = A\boldsymbol{x}(t) + B\boldsymbol{u}(t), \quad (21)$$

where A and B are matrices composed of the estimated transition probabilities $\hat{P}_{kj}^m$. Note that (21) does not precisely represent the actual state distribution of the TCLs.

Subsequently, the optimization problem solved in each iteration of Mode-Blocking MPC is as follows:

$$\min_{\boldsymbol{x}^h, \boldsymbol{u}^h} \sum_{h=t}^{t+H-1} c_h\left(\boldsymbol{x}^h, \boldsymbol{u}^h\right) + c_{t+H}^{\mathrm{f}}\left(\boldsymbol{x}^{t+H}\right) \quad (22a)$$

$$\text{s.t. } \forall h \in \{t, \ldots, t+H-1\}$$

$$\boldsymbol{x}^{h+1} = A\boldsymbol{x}^h + B\boldsymbol{u}^h \quad (22b)$$

$$\boldsymbol{x}^t = \boldsymbol{x}(t) \quad (22c)$$

$$0 \le u^h_{m,m',k} \le x^h_{m,k} \qquad \forall m \in \mathcal{M}, m' \neq m \tag{22d}$$

$$u^h_{m',m,k} = 0 \qquad \forall m \in \mathcal{M}^{\mathrm{d}}_k, m' \neq m \tag{22e}$$

$$x^h_{m,k} = \sum_{m' \in \mathcal{M}^{\mathrm{s}}_k} u^h_{m,m',k} \qquad \forall m \in \mathcal{M}^{\mathrm{d}}_k \tag{22f}$$

$$\underline{N} \le \sum_{k=1}^{K} x^{t+H}_{1,k} \le \overline{N} \tag{22g}$$

where $\boldsymbol{x}^h$ and $\boldsymbol{u}^h$ are the state and input at time step h, respectively, c_h is the cost function at time step h, H is the length of the prediction horizon, c^{f}_{t+H} is the cost function for the last step of the horizon, (22d) ensures the number of mode-switching TCLs do not exceed the number of total TCLs in each interval; (22e) prevents TCLs at $\mathcal{T}^i_k$ from switching to any mode within $\mathcal{M}^{\mathrm{d}}_k$; (22f) ensures that all TCLs in any mode within $\mathcal{M}^{\mathrm{d}}_k$ switch to a mode within $\mathcal{M}^{\mathrm{s}}_k$; and (22g) forces the number of turned-ON TCLs to be within the bounds $\underline{N}$ and $\overline{N}$, aligning with (4).

After solving (22) at time step t, we determine the input $\boldsymbol{u}(t)$ as the optimal value of $\boldsymbol{u}^t$. Note that the input $\boldsymbol{u}(t)$ just specifies how many TCLs in each interval should switch to each mode without specifying which TCLs to switch. Thus, after $\boldsymbol{u}(t)$ is determined, it must be realized by arbitrarily selecting $u_{m,m',k}(t)$ TCLs at $\mathcal{T}^i_k$ to switch from m to m'. The advantage of using aggregate dynamics over determining individual modes of all TCLs is that the number of decision variables is independent of the number of TCLs. It is only dependent on the number of the discretized intervals K, and hence the approach is scalable for a large aggregation of TCLs.

The next result follows directly from Theorem 1.

Proposition 1. *Suppose that* (22) *is feasible for all time steps* t *and the modes of the TCLs are switched according to the optimal input* $\boldsymbol{u}^t$ *from* (22)*. Then, for any* $\epsilon > 0$ *and* $i \in [N]$*, the probability* $P^{\mathrm{d},i}_m(T^i(t))$ *of* i*th TCL moving outside of the comfort set* $\mathcal{T}^{s,i}$ *is less than* ϵ *with a confidence level over* $1-2(1-\epsilon/2)^{D^i_m}$ *for all* $t \in \mathbb{N}_0$*. In addition, the mode-counting constraint* (4) *holds for all* $t \in \mathbb{N}_0$*.*

There are already MPC algorithms that leverage similar models of aggregate TCL dynamics, e.g., [14], [20]. However, the novelty of the proposed MPC algorithm is that it blocks modes in $\mathcal{M}^{\mathrm{d}}_k$ for all k to satisfy comfort conditions while also ensuring network safety. The MPC algorithm developed in [20] utilizes a similar mode-blocking strategy but is only applicable to mildly heterogeneous TCLs with known and deterministic temperature dynamics.

V. Case Study

In this section, we present a case study in which we control $N = 1,000$ heterogeneous cooling TCLs providing grid balancing services while keeping their aggregate power consumption lower than a prescribed bound to maintain distribution network safety. In particular, we consider a scenario in which the real-time network conditions differ significantly from the forecasted network conditions; the same scenario was considered in [10], though the approach in that paper did not provide any guarantee on the comfort condition.

We model each TCL using the affine temperature dynamics model described in [12]:

$$f^i_m(T) = a^i_{\mathrm{th}} T + \left(1 - a^i_{\mathrm{th}}\right)\left(\theta^i_{\mathrm{a}} + b^i_{\mathrm{th}} m\right), \tag{23}$$

where θ^i_{a} is the ambient temperature and

$$a^i_{\mathrm{th}} = \exp\left(-\frac{\Delta t}{r^i_{\mathrm{th}} c^i_{\mathrm{th}}}\right), \quad b^i_{\mathrm{th}} = r^i_{\mathrm{th}} p^i_{\mathrm{tr}},$$

where r^i_{th} is the thermal resistance and c^i_{th} is the thermal capacitance of the ith TCL. Also, p^i_{tr} is the energy transfer rate of the ith TCL, which is negative for a cooling TCL. The power consumption of each TCL in the ON mode is equal to $p^i := p^i_{\mathrm{tr}}/\zeta^i$, where ζ^i is the coefficient of performance. The temperature set-point $T^i_{\mathrm{set}} := (\underline{T}^{\mathrm{s},i} + \overline{T}^{\mathrm{s},i})/2$ is at the center of the temperature dead-band, which has width $d^i := \overline{T}^{\mathrm{s},i} - \underline{T}^{\mathrm{s},i}$. As mentioned in Section II-B, Assumption 1 holds with the Lipschitz constant $L^f_m = 1$.

The TCL parameters are determined by sampling from uniform distributions over specified ranges whose widths $\Delta_\bullet$ are set proportional to a parameter Δ that serves as a scaling factor for adjusting the widths[1]. To evaluate performance under various degrees of heterogeneity, we test the control algorithm for multiple values of Δ. The disturbance probability distributions $\mathcal{W}^i_m$ for each mode and TCL are modeled as normal distributions with mean 0 and standard deviations sampled from the range $[0.005-\Delta_w, 0.005+\Delta_w]$. The dead bands $\mathcal{T}^{\mathrm{s},i}$ are discretized into $K^{\mathrm{s}} = 50$ intervals. In addition, we assume the utility provides an upper bound on the aggregate power consumption of TCLs of $2,240$ kW to prevent under-voltages. We ran offline simulations of the TCLs to obtain state trajectory data to compute the set of admissible modes $\mathcal{M}^{\mathrm{s}}_k$, blocked modes $\mathcal{M}^{\mathrm{d}}_k$, and empirical transition probabilities $\hat{P}^m_{kj}$. The number of data points D^i_m was larger than 100 for all $i \in [N]$ and $m \in \{0, 1\}$, and so we obtain $P^{\mathrm{d},i}_m(T^i(t)) \le 0.001$ with confidence level of almost 1.

We compare Mode-Blocking MPC with a Benchmark MPC, which permits switching to all modes across each discretized interval $\mathcal{T}^i_k$ within the comfort set $\mathcal{T}^{\mathrm{s},i}$; this control algorithm has served as a benchmark in [20] as well. Through this comparison, we aim to observe the impact of mode-blocking on the comfort condition, network safety, and tracking performance. The cost functions for both of the MPC algorithms are designed to maximize tracking performance while penalizing the total number of mode switches to reduce the damage to the TCLs' compressors:

$$c_t\left(\boldsymbol{x}(t), \boldsymbol{u}(t)\right) = \underbrace{\left|\tilde{r}(t) - p_{\mathrm{avg}} \sum_{k=1}^{K} x_{1,k}(t)\right|}_{c^{\mathrm{f}}_t(\boldsymbol{x}(t))} + \alpha \left\|\boldsymbol{u}(t)\right\|_1$$

[1] $r^i_{\mathrm{th}} \in [1.8 - \Delta_{\mathrm{r}}, 1.8 + \Delta_{\mathrm{r}}]$ °C/kW, $c^i_{\mathrm{th}} \in [1.5 - \Delta_{\mathrm{c}}, 1.5 + \Delta_{\mathrm{c}}]$ kWh/°C, $\theta^i_{\mathrm{a}} \in [32 - \Delta_{\mathrm{a}}, 32 + \Delta_{\mathrm{a}}]$ °C, $p^i_{\mathrm{tr}} \in [-16 - \Delta_p, -16 + \Delta_p]$ kW, $\zeta^i \in [2.5 - \Delta_\zeta, 2.5 + \Delta_\zeta]$, $T^i_{\mathrm{set}} \in [23 - \Delta_T, 23 + \Delta_T]$ °C, $d^i = 1.0$ °C, $\Delta t = 20$ s.

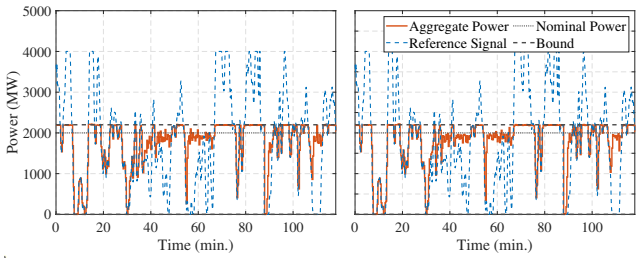

Fig. 4. Aggregate power consumption of the TCLs under the Benchmark MPC (left) and Mode-Blocking MPC (right) when $\Delta = 0.2$. The aggregate powers under both algorithms do not violate their bounds significantly.

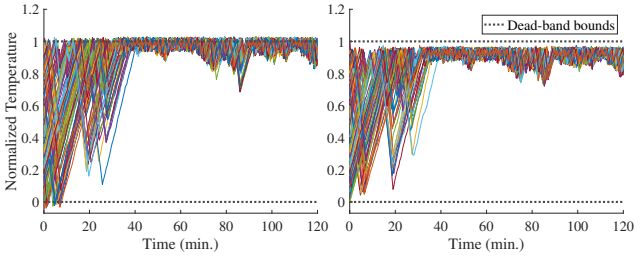

Fig. 5. Trajectories of TCL normalized temperatures under the Benchmark MPC (left) and Mode-Blocking MPC (right) when $\Delta = 0.2$. A significant number of TCLs move outside of the dead-band under the Benchmark MPC. No dead-band violations occurred under Mode-Blocking MPC.

where $\alpha = 0.01$ is penalty on the number of mode switches. We generate the reference signal $r(t)$ by shifting and scaling a 2-hour segment of the PJM RegD signal [24]. This signal is originally at a 2-second intervals, but to align with the sampling time $\Delta t = 20$ s, we employ down-sampling, resulting in a modified signal $\tilde{r}(t)$.

The length of the prediction horizon for MPC was set to $H = 5$. For each case, we analyze the tracking performance by calculating the average difference between the TCL aggregate power consumption $p_{\text{agg}}(t) := \sum_{k=1}^{K} p^i m^i(t)$ and $\tilde{r}(t)$, i.e., $\frac{1}{T_h}\sum_{t=1}^{T_h} |\tilde{r}(t) - p_{\text{agg}}(t)|$ where T_h is the total number of time steps in the experiment. We also empirically evaluate the comfort condition by computing the fraction of TCLs that move to the discomfort set $\mathcal{T}^{\text{d},i}$, and assess network safety by calculating the fraction of time steps with aggregate power bound violations and the maximum deviation from the upper bound $\overline{p}$, as a percentage of $\overline{p}$. We perform five experiments for each of eight Δ values (0, 0.05, 0.1, 0.15, 0.2, 0.25, 0.3, 0.35), keeping the TCL parameters constant for a given Δ while changing the disturbance samples across the five experiments.

Fig. 4 illustrates the TCL aggregate power consumption and reference signal $r(t)$ under each algorithm from one experiment for $\Delta = 0.2$. Fig. 5 shows the trajectories of the normalized temperatures (i.e., the TCL dead-bands are projected to $[0, 1]$) of 100 of the TCLs under each algorithm for $\Delta = 0.2$. As shown, the TCL aggregate power consumption does not accurately track the reference signal even when it is lower than $\overline{p}$ as the TCL population increasingly clusters towards the upper limit of the dead-band, which restricts TCLs from turning ON. The mode-counting constraints are always satisfied for both algorithms, though the actual aggregate power consumption of the TCLs occasionally exceeds the upper limit $\overline{p}$, as shown in Table I, since the mode-counting constraints are approximate. The maximum observed deviation from $\overline{p}$ is less than 2% of $\overline{p}$, which should have only a small effect on network voltages.

TABLE I

FRACTION OF TIME STEPS WITH VIOLATIONS AND MAX DEVIATION (%)

Δ	0.05	0.1	0.15	0.2	0.25	0.3	0.35
Fraction	0	0.015	0.038	0.015	0.050	0.122	0.213
Max Dev	0	.082	.328	.283	.544	.968	1.66

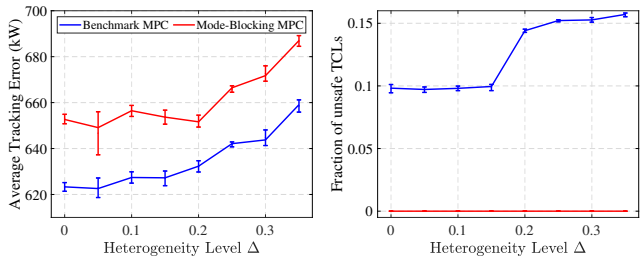

Fig. 6. Averages and ranges of reference tracking error (left) and fraction of TCLs within $\mathcal{T}^{\text{d},i}$ (right) versus heterogeneity level Δ across the five experiments. The difference in tracking error between the two controllers does not change significantly with the heterogeneity level. The fraction of TCLs outside of the dead-band increases with heterogeneity level under Benchmark MPC but stays constant at zero under Mode-Blocking MPC.

Fig. 6 shows the averages and ranges of tracking error and the fraction of TCLs in $\mathcal{T}^{\text{d},i}$ across the five experiments for different levels of heterogeneity Δ, while Table II shows the number of intervals $\mathcal{T}_k^i$ within the dead-band in which each mode is blocked. As illustrated in Fig. 6, the tracking performance of Benchmark MPC always surpasses that of Mode-Blocking MPC because it has greater flexibility from a larger number of admissible modes within the intervals in $\mathcal{T}^{\text{s},i}$. However, the average fraction of TCLs within $\mathcal{T}^{\text{d},i}$ is high (exceeding 10%) and increases with TCL heterogeneity. In contrast, Mode-Blocking MPC prevented any TCL from moving to $\mathcal{T}^{\text{d},i}$. This shows the impact of implementing mode-blocking for intervals within $\mathcal{T}^{\text{s},i}$, particularly when we wish to prioritize the comfort condition under significant TCL heterogeneity. Note that, while differences in tracking performance between the two approaches are almost constant with respect to Δ, the fraction of TCLs within the discomfort set and the difference in comfort condition outcomes between the two approaches increases with Δ in Benchmark MPC. This demonstrates that Mode-Blocking MPC successfully ensures the comfort condition without sacrificing much tracking performance. Moreover, the tracking performances of both Benchmark MPC and Mode-Blocking MPC were not necessarily proportional to the level of heterogeneity. However, beyond $\Delta = 0.2$, the tracking performance degrades as the heterogeneity increases. This can be explained by the growing fraction of TCLs with faster temperature dynamics, resulting in diminishing flexibility. Additionally, the average time to solve the Benchmark MPC is 1.28 s, and for the Mode-Blocking MPC, it is 1.15 s, both of which are shorter

TABLE II
NUMBER OF INTERVALS $\mathcal{T}_k^i$ IN WHICH EACH MODE IS BLOCKED

Δ	0	0.05	0.1	0.15	0.2	0.25	0.3	0.35
OFF Mode	3	2	3	3	2	3	3	3
ON Mode	4	4	4	4	4	4	4	4

than $\Delta t = 20$ s. Hence, both approaches can be utilized in real-time applications.

The conservatism of Mode-Blocking MPC is influenced by the number of bins that are blocked for each mode. We observed that an increase in heterogeneity level Δ does not consistently lead to a rise in the number of blocked intervals, as shown in Table II. The tracking performance improved when the number of blocked bins was reduced (e.g., for $\Delta = 0$ to 0.05) while it significantly degraded when the number of blocked bins was increased (e.g., for $\Delta = 0.2$ to 0.25). In cases in which the number of blocked bins remained the same despite an increase in Δ, tracking performance was sometimes improved, although the general trend was a reduction in tracking performance with increased TCL heterogeneity.

VI. CONCLUSION

This paper proposed an MPC algorithm for a large aggregation of heterogeneous TCLs with unknown temperature dynamics to provide balancing services. Leveraging a data-driven method, we constructed a set of admissible modes for each discretized interval that ensures the satisfaction of the comfort condition when TCLs choose modes from these sets. We then developed a Mode-Blocking MPC that leverages these sets to concurrently satisfy a comfort condition and mode-counting constraints, which approximates network safety. A key feature of Mode-Blocking MPC is it guarantees the probability of satisfaction of the comfort condition while coordinating heterogeneous TCLs with unknown temperature dynamics. A case study demonstrated that Mode-Blocking MPC satisfies the comfort condition for a large range of TCL heterogeneity levels, while not sacrificing much in terms of tracking performance. This demonstrated the benefits of Mode-Blocking MPC versus a Benchmark MPC approach that does not use mode-blocking.

In the future, we aim to enhance Mode-Blocking MPC, e.g., by developing approaches to ensure recursive feasibility and increase robustness against error from the approximated TCL dynamics.

Acknowledgment: The authors would like to thank Haechan Jeong for his assistance in experimenting with various previous versions of the Mode-Blocking MPC algorithm.

REFERENCES

[1] Federal Energy Regulatory Commission, "FERC order no. 2222: Participation of distributed energy resource aggregations in markets operated by regional transmission organizations and independent system operators." Sep. 2020. [Online]. Available : https://www.ferc.gov/sites/default/files/2020-09/E-1_0.pdf.

[2] 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.

[3] S. C. Ross, G. Vuylsteke, and J. L. Mathieu, "Effects of load-based frequency regulation on distribution network operation," *IEEE Transactions on Power Systems*, vol. 34, no. 2, pp. 1569–1578, 2018.

[4] Y. Yan, Y. Qian, H. Sharif, and D. Tipper, "A survey on smart grid communication infrastructures: Motivations, requirements and challenges," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 1, pp. 5–20, 2012.

[5] Federal Energy Regulatory Commission, "Notice inviting post-technical conference comments," Tech. Rep. Docket No.. RM18-9-000, Apr. 2018. [Online]. Available : https://cms.ferc.gov/sites/default/files/2020-09/Notice-for-Comments-on-NOPR-RM18-9.pdf.

[6] S. Ross and J. Mathieu, "Strategies for network-safe load control with a third-party aggregator and a distribution operator," *IEEE Transactions on Power Systems*, vol. 36, no. 4, pp. 3329–3339, 2021.

[7] H. D. Nguyen, K. Dvijotham, and K. Turitsyn, "Constructing convex inner approximations of steady-state security regions," *IEEE Transactions on Power Systems*, vol. 34, no. 1, pp. 257–267, 2018.

[8] D. Lee, K. Turitsyn, D. K. Molzahn, and L. A. Roald, "Robust AC optimal power flow with robust convex restriction," *IEEE Transactions on Power Systems*, vol. 36, no. 6, pp. 4953–4966, 2021.

[9] N. Nazir and M. Almassalkhi, "Grid-aware aggregation and realtime disaggregation of distributed energy resources in radial networks," *IEEE Transactions on Power Systems*, vol. 37, no. 3, pp. 1706–1717, 2021.

[10] S. Jang, N. Ozay, and J. L. Mathieu, "Probabilistic constraint construction for network-safe load coordination," *IEEE Transactions on Power Systems*, vol. 39, no. 2, pp. 4473–4485, 2024.

[11] W. Zhang, J. Lian, C.-Y. Chang, and K. Kalsi, "Aggregated modeling and control of air conditioning loads for demand response," *IEEE Transactions on Power Systems*, vol. 28, no. 4, pp. 4655–4664, 2013.

[12] R. C. Sonderegger, *Dynamic models of house heating based on equivalent thermal parameters.* Ph.D. dissertation, Princeton University, 1978.

[13] I. M. Granitsas, I. A. Hiskens, J. L. Mathieu, and G. S. Ledva, "Parameter identifiability and estimation of thermostatically controlled loads," in *IEEE PowerTech*, Belgrade, Serbia, 2023.

[14] J. Mathieu, S. Koch, and D. Callaway, "State estimation and control of electric loads to manage real-time energy imbalance," *IEEE Transactions on Power System*, vol. 28, no. 1, pp. 430–440, 2012.

[15] S. Bashash and H. K. Fathy, "Modeling and control of aggregate air conditioning loads for robust renewable power management," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 4, pp. 1318–1327, 2012.

[16] S. H. Tindemans, V. Trovato, and G. Strbac, "Decentralized control of thermostatic loads for flexible demand response," *IEEE Transactions on Control Systems Technology*, vol. 23, no. 5, pp. 1685–1700, 2015.

[17] S. E. Z. Soudjani and A. Abate, "Aggregation and control of populations of thermostatically controlled loads by formal abstractions," *IEEE Transactions on Control Systems Technology*, vol. 23, no. 3, pp. 975–990, 2014.

[18] C. Ziras, S. You, H. W. Bindner, and E. Vrettos, "A new method for handling lockout constraints on controlled TCL aggregations," in *Power Systems Computation Conference (PSCC)*, 2018.

[19] A. Coffman, A. Bušić, and P. Barooah, "A unified framework for coordination of thermostatically controlled loads," *Automatica*, vol. 152, 2023.

[20] S. Jang, N. Ozay, and J. L. Mathieu, "An invariant set construction method, applied to safe coordination of thermostatic loads," *IEEE Transactions on Control of Network Systems*, pp. 1–12, 2024.

[21] P. Nilsson and N. Ozay, "Control synthesis for permutation-symmetric high-dimensional systems with counting constraints," *IEEE Transactions on Automatic Control*, vol. 65, no. 2, pp. 461–476, 2019.

[22] A. Hassan, R. Mieth, D. Deka, and Y. Dvorkin, "Stochastic and distributionally robust load ensemble control," *IEEE Transactions on Power Systems*, vol. 35, no. 6, pp. 4678–4688, 2020.

[23] S. C. Ross and J. L. Mathieu, "A method for ensuring a load aggregator's power deviations are safe for distribution networks," *Electric Power Systems Research*, vol. 189, p. 106781, 2020.

[24] PJM, "RTO regulation signal data for 7.19.2019 & 7.20.2019.xls," https://www.pjm.com/markets-and-operations/ancillary-services.aspx, accessed: 2019-10-22.