A unified framework for coordination of thermostatically controlled loads [⋆]

Austin Coffman [a], Ana Bušić [b], Prabir Barooah [c]

[a] *University of Florida, Gainesville, FL, USA*

[b] *Inria, Paris, France*

[c] *University of Florida, Gainesville, FL, USA*

Abstract

A collection of thermostatically controlled loads (TCLs) – such as air conditioners and water heaters – can vary their power consumption within limits to help the balancing authority of a power grid maintain demand supply balance. Doing so requires loads to coordinate their on/off decisions so that the aggregate power consumption profile tracks a grid-supplied reference. At the same time, each consumer's quality of service (QoS) must be maintained. While there is a large body of work on TCL coordination, there are several limitations. One is that they do not provide guarantees on the reference tracking performance and QoS maintenance. A second limitation of past work is that they do not provide a means to compute a suitable reference signal for power demand of a collection of TCLs. In this work we provide a framework that addresses these weaknesses. The framework enables coordination of an arbitrary number of TCLs that: (i) is computationally efficient, (ii) is implementable at the TCLs with local feedback and low communication, and (iii) enables reference tracking by the collection while ensuring that temperature and cycling constraints are satisfied at every TCL at all times. The framework is based on a Markov model obtained by discretizing a pair of Fokker-Planck equations derived in earlier work by Malhame and Chong [21]. We then use this model to design randomized policies for TCLs. The balancing authority broadcasts the same policy to all TCLs, and each TCL implements this policy which requires only local measurement to make on/off decisions. Simulation results are provided to support these claims.

Key words: Distributed control, Grid support, Randomized control, Thermostatically controlled loads.

Contents

1   Introduction

Many loads are flexible in their power demand: they can vary their demand around a baseline without adversely affecting consumers' quality of service (QoS). The flexibility can be used by a balancing authority (BA) to balance supply and demand in a power grid. The baseline demand refers to the power demand under normal operation, when each load operates only to meet its consumer's QoS without any interference from the BA. Since the rated power of each load is small, it is necessary to use a collection of loads. To provide grid support, the collection has to vary its demand from its baseline. It is envisioned that the BA would supply a reference signal for power demand and the actions of the loads in a collection would be coordinated so that their total demand tracks this reference.

Thermostatically controlled loads (TCLs) - such as residential air conditioners, heat pumps, and water heaters - are recognized to be valuable sources of flexible demand [4,6,23,18]. For an air conditioner or a heat pump, baseline demand is largely dictated by ambient weather conditions. There are at least two QoS requirements: the indoor temperature must be maintained within a prespecified range and compressor short-cycling must be avoided, meaning, once the compressor turns on it cannot turn off until a prespecified time period elapses, and vice versa. Coordination of TCLs involves two conflicting requirements: (i) the TCLs collectively need to track the reference power demand signal, and (ii) every TCL's QoS need to be maintained.

The actuation at each TCL is discrete: it can either be on or off. Direct load control [7], in which a centralized controller at the BA directly commands on/off status of each TCL is not scalable to large populations. A more scalable idea, that subsequent works on TCL coordination use, is for the BA to broadcast a low dimensional control command to all TCLs, which is translated by each TCL into its actuation command with a local policy. To avoid confusion between the decision making at the BA and a TCL, we use the word "policy" to mean the algorithm at a TCL that makes on/off decisions. The literature on decentralized coordination of TCLs differ in their choice of the broadcast signal (i.e., BA's control command) and the policy at the TCL that translates this broadcast to on/off decisions. Coordination architectures can be divided into two broad categories based on these choices: (i) thermostat set point change [4,1] and (ii) probabilistic control [23,20,6,9]. These are discussed in more detail in Section 1.1.

A framework for coordinating TCLs needs two parts. The coordination scheme is one part. The other part is reference computation: the framework must provide the BA with a method to determine a suitable reference signal for the TCLs. That is, the reference must be such that the TCLs can collectively track the signal while each TCL maintains its QoS. Otherwise, even the best coordination scheme will fail to meet either the BA's need, which is reference tracking, or the consumers' need, which is maintaining indoor temperature etc., or both.

This work presents a unified framework for coordination of a collection of TCLs for providing grid support services. The framework enables both of the above mentioned components, i.e., (i) planning a suitable reference for a collection of TCLs and (ii) designing a randomized policy for coordination of the individual TCLs, so that both the BA's requirement and consumers' QoS are satisfied. In the proposed framework, the BA computes randomized control policies for the TCLs and broadcasts them to all the TCLs. Each TCL receives the same policy and implements it using locally measurable information. The framework is computationally tractable for an arbitrary number of TCLs. The communication burden is low: only a few numbers need to be broadcast by the BA at every sampling instant. Feedback from TCLs to the BA can be infrequent.

Underlying the framework is: (i) a Markov chain model that is derived from partial differential equations developed in the early work of Malhame and Chong [21], (ii) state augmentation to incorporate cycling constraints,

and (iii) convexification of the non-convex problem that appears in the design of the randomized control policy for the individual TCL. Additionally, we show that the assumption made about the effect of weather in earlier work [3] on randomized control, under certain conditions, is in fact true.

1.1 Literature review and contribution

Before reviewing coordination methods, we discuss two interrelated modeling approaches that underpin many of the ideas in the TCL control architectures. These are the Markov chain and partial differential equation (PDE) models [16,21,26,17,28,24], which stem from the early work of Malhame and Chong [21]. In [21] a pair of coupled Fokker-Planck equations are developed to model a collection of TCLs under thermostat control. The Fokker-Planck equations are PDEs that describe the time evolution of a certain probability density functions (pdf) over the state space of temperature and on/off mode. The PDEs can be used to model the entire collection or a single TCL: the probability that a single TCL is "on" is approximately the fraction of TCLs that are "on". Discretizing the PDE yields a Markov chain model, though some works have obtained Markov models without using the PDEs. Hence, *one* set of PDEs can model a collection of TCLs. Thus, methods that base control design on the PDE or Markov chain framework scales well with the number of TCLs.

Due to the lack of scalability of direct load control, we limit our attention to the two broad classes mentioned earlier: (i) thermostat set point changes, (ii) probabilistic policy. There are many forms of probabilistic policy, which can be roughly subdivided into two sub categories: (ii-A) "bin switching" and (ii-B) "randomized policy". We discuss these in detail below.

In the thermostat setpoint change coordination architecture, a time-varying thermostat set point is broadcast to all TCLs, and each TCL makes on/off decisions based on this new setpoint [4,1]. This approach may ask for an extremely small change in thermostat setpoint, far below the resolution of the temperature sensor at each TCL. Or it may ask for large changes in thermostat setpoint which will violate occupant comfort.

In a probabilistic policy architecture, the TCL policy - the mapping from BA's broadcast command to a TCL's on/off decision - is a non-deterministic mapping. Works in this category typically first model the population of TCLs under thermostat control, which is a deterministic policy, as a Markov chain. The continuous temperature range is divided into a number of discrete bins. A finite dimensional state vector, a probability mass function, is then defined. Each entry of the state vector represents "the fraction of TCLs that are on (or off) and has temperature in a certain range."

Since the basic Markov model is derived for the thermostat policy, introduction of the BA's control to manipulate TCLs' on/off state is somewhat ad-hoc. In the the bin switching literature, the control command from the BA is chosen so as to affect the fraction of TCLs in the temperature bins directly. In [23], the BA's control command is chosen to be another vector, whose i^{th} entry represents "the fraction of TCLs in bin i to increase/decrease". A policy is then proposed to translate this command to on/off action at each TCL, which requires knowledge of the state of the Markov model. In [20], BA's control command is chosen to be a scalar. The probability of a TCL turning on or off is proportional to this scalar. Subsequent works have proposed various refinements, such as BA's command affecting the rate of fractions to switch instead of fraction to switch [26]. Providing performance guarantees with bin switching architecture has proved challenging, either on reference tracking or on QoS maintenance for individual TCLs.

An alternative to bin switching that still uses probabilistic on/off decision making is randomized policy [3,6]. A randomized policy is a specification of the conditional probability of turning on or off given the current state. On/off decisions are computed with the help of a random number generator and the policy. In this architecture it is envisioned that the thermostat policy at the TCL is replaced with a randomized policy. In [3,6], the policy is parameterized by a scalar $\zeta(t)$. Coordination of the population is then achieved by appropriate design of $\zeta(t)$, which is computed and broadcast by the BA. This architecture also uses a Markov model of the evolution of binned temperature, but assumes a certain factorization: the next values of the temperature and mode are conditionally independent given the current joint pair of temperature and mode values under the effects of the randomized policy and exogenous disturbances, especially weather. That is, the transition matrix of the state process is a point wise product of two controlled transition matrices. In an optimal control setting, computation of the BA's control command, $\zeta(t)$, for reference tracking is a non-convex optimization problem [11]. The probability of turning on when temperature exceeds the upper limit, or off when temperature dips below the lower limit, is set to 1 by design. This will ensure the temperature constraint is maintained. Attempts have been made to maintain the cycling constraint [9]. But a formal design method to incorporate the cycling constraint has been lacking.

A complete framework for coordination of TCL needs not only a control algorithm to make decisions at TCLs, but also a method to compute a *feasible* reference signal for the collection's power demand. Feasible means that no TCL needs to violate local constraints in order for the collection to track the reference. The topic is sometimes described as "flexibility capacity" and has been examined in many recent works, with various defi-

nitions of flexibility [15,25,13,10]. A unified treatment of reference design and coordination algorithm design that would provide a complete framework is lacking.

In short, existing work on TCL coordination has a number of scattered disadvantages. Direct load control suffers from scalability/privacy issues and thermostat setpoint methods have implementation issues. Bin switching does not provide guarantees on reference tracking and often requires solving a challenging state estimation problem. Prior work on randomized control requires non-convex optimization and is based on an assumed conditional independence. Finally, there is a lack of unified treatment of the reference design and policy design problems.

In this work we develop a unified framework for coordination of TCLs that addresses the weaknesses of prior work described above. Our major contributions are as follows.

(1) We provide a complete framework that allows the BA to compute (a) an optimal reference signal that is feasible for the collection and (b) optimal randomized policies for the TCLs. When the TCLs implement these policies, their total power demand collectively tracks the reference signal and the policies guarantee that temperature and cycling QoS requirements at each TCL are satisfied. Optimal reference means it is closest to what the BA wants while being feasible for the TCLs. Implementation of the policy at a TCL is easy; it requires only local measurements. The communication burden for coordination is also low. At each sampling time, a randomized control policy - parameterized by a few numbers - is broadcast to all TCLs. Feedback from TCLs to the BA can be infrequent.

(2) Our framework is based on a careful discretization of the partial differential equation (PDE) model described in [21]. This discretization shows that a certain "conditional independence" that was assumed in [3] indeed holds. The conditional independence separates the effects of the policy at the TCL (control) and weather (disturbance) on the transition matrix, and greatly facilitates computation of policies.

(3) Numerical experiments are provided to illustrate the efficacy of the framework. Simulations show that TCLs are able to track the optimal reference collectively while each TCL is able to maintain both temperature and cycling constraints. Matlab implementation is made publicly available at [8].

Figure 1 illustrates the two parts of the proposed framework.

The Markov model obtained by discretizing a PDE was presented in [12]. For completeness, we include the discretization in this paper as an Appendix.
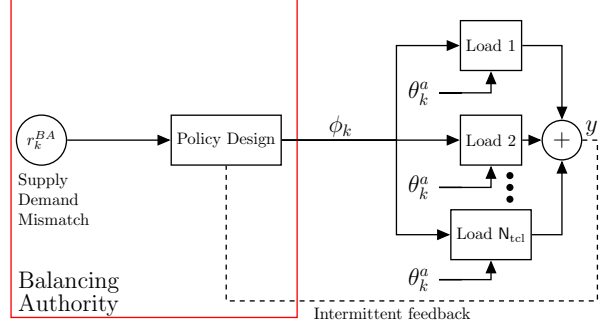

Fig. 1. Coordination architecture with the proposed framework.

1.2 Notation

The symbol $\mathbb{1}$ denotes the vector of all ones, $\mathbf{e}_i$ denotes the i^{th} canonical basis vector, and $\mathbf{0}$ denotes the zero matrix or vector, all of appropriate dimension. For a vector v, $\mathrm{diag}(v)$ denotes the diagonal matrix with entries of v, i.e., $\mathrm{diag}(v)\mathbb{1} = v$. Further, $\otimes$ denotes matrix Kronecker product and $\mathbf{I}_A(\cdot)$ the indicator function of the set A.

2 Modeling: Individual TCL

A thermostatically controlled load (TCL) is an on/off device that ensures the temperature of a given environment remains within a specified region, e.g., an air conditioner. During its operation, the TCL must adhere to certain operational requirements (QoS constraints). We consider two: the temperature constraint and the cycling constraint. The temperature constraint is that the TCL's temperature must remain within a prespecified deadband, $[\lambda^{\min}, \lambda^{\max}]$. This is achieved by switching the TCL on or off when it is too hot or cold. The cycling constraint is that the TCL can only change from "on" to "off" or vice versa once every τ (discrete) time instants, where τ is a prespecified constant. The cycling constraint is to ensure the mechanical hardware is not damaged. In both cases, ensuring the two constraints amounts to appropriately deciding when to switch the TCL on or off.

2.1 Temperature dynamics of TCLs

The typical model for the TCL's temperature $\theta(t)$ in the literature is the following ordinary differential equation (ODE),

$$\frac{d}{dt}\theta(t) = f_m(\theta, t), \quad \text{with}$$
$$f_m(\theta, t) = -\frac{1}{RC}(\theta - \theta^a(t)) - m(t)\frac{\eta P_0}{C}. \quad (1)$$

The rated electrical power consumption is denoted P_0 with coefficient of performance (COP) η. The parame-

ters R and C denote thermal resistance and capacitance, respectively. The signal $\theta^a(t)$ is the ambient temperature. The quantity $m(t)$ is the on/off mode, and in the following we identify $m(t) = 1$ and $m(t) = $ on, as well as $m(t) = 0$ and $m(t) = $ off. We denote arbitrary temperature values through the variable λ, and the thermostat setpoint as λ^{set}. The values λ^{max} and λ^{min} set the upper and lower limit for the temperature deadband.

A model for the temperature state that accounts for modeling errors in (1) and will be crucial in developing the content in Section 2.2 is the following Itô stochastic differential equation (SDE),

$$d\theta(t) = f_m(\theta, t)dt + \sigma dB(t). \tag{2}$$

The term $B(t)$ is Brownian motion with parameter $\sigma > 0$, and the quantity $\sigma dB(t)$ captures modeling errors in (1). In either model, the baseline power for the TCL is the value of P so that $f_1(\lambda^{\text{set}}, t) = 0$, solving yields:

$$\text{Baseline Power:} \quad \bar{P}^{\text{ind}}(t) = \frac{\theta^a(t) - \lambda^{\text{set}}}{\eta R}. \tag{3}$$

For N_{tcl} TCLs the baseline power $\bar{P}(t)$ and maximum power P_{agg} are,

$$\bar{P}(t) \triangleq \mathsf{N}_{\text{tcl}} \bar{P}^{\text{ind}}(t), \quad \text{and} \quad P_{\text{agg}} \triangleq \mathsf{N}_{\text{tcl}} P_0. \tag{4}$$

The total electrical power consumption of the collection, whether with thermostat policy or some other policy, is denoted by $y(t)$:

$$y(t) \triangleq P_0 \sum_{\ell=1}^{\mathsf{N}_{\text{tcl}}} m^\ell(t) \tag{5}$$

where $m^\ell(t)$ is the on/off state of the ℓ-th TCL.

2.1.1 Policy (at the TCL)

The mode state of a TCL evolves according to a policy. The following policy, which we denote as the *thermostat policy*, ensures the temperature constraint:

$$\lim_{\epsilon \to 0} m(t + \epsilon) = \begin{cases} 1, & \theta(t) \geq \lambda^{\text{max}}. \\ 0, & \theta(t) \leq \lambda^{\text{min}}. \\ m(t), & \text{o.w.} \end{cases} \tag{6}$$

We add the following set of assumptions about the individual TCL discussed so far.

A.1 The thermostat policy does not violate the cycling constraint.
A.2 For all $t \geq 0$ and $\theta \in [\lambda^{\text{min}}, \lambda^{\text{max}}]$, $f_{\text{on}}(\theta, t) \leq 0$ and $f_{\text{off}}(\theta, t) \geq 0$.

A.3 The TCL's cycling and temperature constraint are both simultaneously feasible.

The sizing/design of the TCL is most likely to ensure that **A.1** holds. With **A.1** , we depart from discussing the cycling constraint until Section 5 since up to that point the mode state is assumed to evolve according to (6).

Assumption **A.2** states that when the TCL is on, the temperature does not increase and when the TCL is off the temperature does not decrease. All prior works focusing on cooling TCLs (e.g., air conditioners) implicitly make this assumption. Every result that is to follow is also valid for heating TCLs (e.g., a water heater or a heat pump) with a sign reversal.

Like **A.2**, assumption **A.3** is also implicit in any work that considers both the TCLs temperature and cycling constraint.

2.2 PDE model

We now describe a PDE model of a TCL's temperature with thermostat policy originally derived in [21]. Consider the following marginal pdfs $\mu_{\text{on}}, \mu_{\text{off}}$:

$$\mu_{\text{on}}(\lambda, t)d\lambda = \mathsf{P}\left((\lambda < \theta(t) \leq \lambda + d\lambda),\ m(t) = \text{on}\right), \tag{7}$$

$$\mu_{\text{off}}(\lambda, t)d\lambda = \mathsf{P}\left((\lambda < \theta(t) \leq \lambda + d\lambda),\ m(t) = \text{off}\right), \tag{8}$$

where $\mathsf{P}(\cdot)$ denotes probability, $\theta(t)$ evolves according to (2) and for now $m(t)$ evolves according to (6). It was shown in [21] that the densities μ_{on} and μ_{off} satisfy the Fokker-Planck equations,

$$\frac{\partial}{\partial t}\mu_{\text{on}}(\lambda, t) = \frac{\sigma^2}{2}\nabla^2_\lambda \mu_{\text{on}}(\lambda, t) - \nabla_\lambda\left(f_{\text{on}}(\lambda, t)\mu_{\text{on}}(\lambda, t)\right) \tag{9}$$

$$\frac{\partial}{\partial t}\mu_{\text{off}}(\lambda, t) = \frac{\sigma^2}{2}\nabla^2_\lambda \mu_{\text{off}}(\lambda, t) - \nabla_\lambda\left(f_{\text{off}}(\lambda, t)\mu_{\text{off}}(\lambda, t)\right) \tag{10}$$

that are coupled through their boundary conditions [21]. The boundary conditions are listed in Appendix B.2.

Remark 1. The coupled equations (9)-(10) can be used to model either: (i) a *single* TCL or (ii) a *collection* of TCLs. For (i) the quantities (7)-(8) represent the *probability* that a single TCLs temperature and on/off mode reside in the respective region. For (ii) the quantities (7)-(8) represent the *fraction* of TCLs whose temperature and on/off mode reside in the respective region. How the equations (9)-(10) (specifically their discretized form) can be used to model an ensemble is discussed further in Section 5.2.
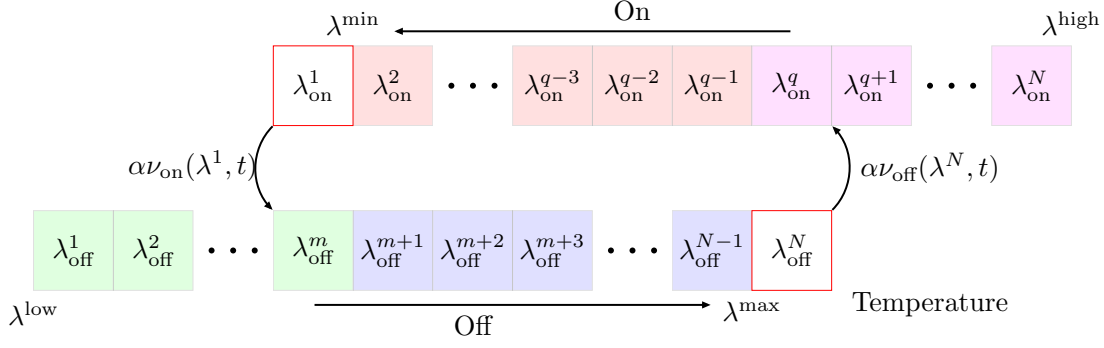
Fig. 2. The control volumes (CVs). The colors correspond to the colors found in Figure 3. The values in each CV represent the nodal temperature for the CV. The arrows describe the sign of the convection of the TCL through the CVs. The values are such that $N = m + q$. The terms involving α model rate of transfer between the corresponding CVs due to the thermostat policy, where $\alpha = \gamma + \frac{\sigma^2}{(\Delta\lambda)^2}$. The parameter $\gamma > 0$ is a design parameter; see Remark 4.

3 Markov model from PDE Discretization

We use the finite volume method (FVM) to discretize the PDEs (9) and (10). The discretization of (9) and (10) yields a finite dimensional probabilistic model for a single TCL (equation (16)). We discretize the PDEs (9) and (10) in a way that a control input for the BA can then be identified. More on this point will be discussed in Section 4, however the discretization here will play a role.

3.1 Spatial discretization

The FVM bins the continuous temperature into N control volumes (CV). The layout of the CVs is shown in Figure 2. The N CVs for both the on and off mode state, as shown in Figure 2, are defined through the nodal temperature values (λ_{on} and λ_{off}) and their boundaries (λ_{on}^+ and λ_{off}^+) and (λ_{on}^- and λ_{off}^-):

$$\lambda_{\text{on}} = (\lambda_{\text{on}}^i)_{i=1}^N, \quad \lambda_{\text{on}}^+ = \lambda_{\text{on}} + \frac{\Delta\lambda}{2}, \quad \lambda_{\text{on}}^- = \lambda_{\text{on}} - \frac{\Delta\lambda}{2},$$
$$\lambda_{\text{off}} = (\lambda_{\text{off}}^i)_{i=1}^N, \quad \lambda_{\text{off}}^+ = \lambda_{\text{off}} + \frac{\Delta\lambda}{2}, \quad \lambda_{\text{off}}^- = \lambda_{\text{off}} - \frac{\Delta\lambda}{2},$$

where $\Delta\lambda$ is the CV width. All intermediate values of λ_{on} and λ_{off} are separated from each other by $\Delta\lambda$. The values in λ_{on}^+ (respectively, λ_{off}^+) are the right edges of the CVs and the values λ_{on}^- (respectively, λ_{off}^-) are the left edges of the CVs, for example, $\lambda_{\text{off}}^{1,-} = \lambda^{\text{low}}$. The quantities λ^{min} and λ^{max} specify the thermostat deadband, and are *different* from the quantities λ^{high} and λ^{low} (see Figure 2).

The steps taken to obtain the spatially discretized PDEs is detailed in Appendix B. To give an overview, the discretization is done in two parts: (i) for the internal CV's (Appendix B.1) and (ii) for the boundary CV's
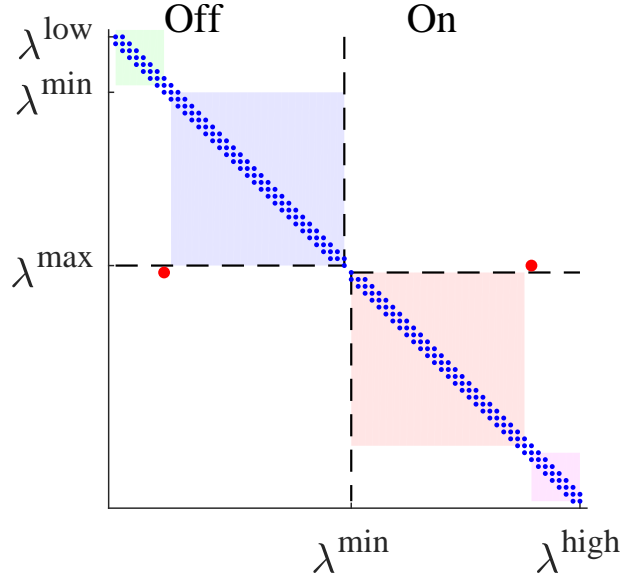
Fig. 3. Sparsity pattern of the matrix $A(t)$ for $N = 51$ CVs for both the on and off state. The colors correspond to the colors found in Figure 2.

(Appendix B.2). We describe here the end result of the derivation in Appendix B. First, define the following quantities

$$\nu_{\text{off}}(\lambda^i, t) \triangleq \mu_{\text{off}}(\lambda^i, t)\Delta\lambda, \quad \text{and} \quad (11)$$
$$\nu_{\text{on}}(\lambda^i, t) \triangleq \mu_{\text{on}}(\lambda^i, t)\Delta\lambda, \quad (12)$$

then construct the row vector, $\nu(t) = [\nu_{\text{off}}(t), \nu_{\text{on}}(t)]$. with

$$\nu_{\text{off}}(t) \triangleq [\nu_{\text{off}}(\lambda^1, t), \dots, \nu_{\text{off}}(\lambda^N, t)], \quad \text{and} \quad (13)$$
$$\nu_{\text{on}}(t) \triangleq [\nu_{\text{on}}(\lambda^1, t), \dots, \nu_{\text{on}}(\lambda^N, t)]. \quad (14)$$

By combining all the ordinary differential equations (ODEs) for the $\nu_{\text{off}}(\lambda^i, t), \nu_{\text{on}}(\lambda^i, t)$ for all the i's, we

obtain the linear time varying system

$$\frac{d}{dt}\nu(t) = \nu(t)A(t). \qquad (15)$$

The sparsity pattern of $A(t)$ is shown in Figure 3. The system (15) is the spatially discretized version of the PDEs (9)-(10). The matrix $A(t)$ also satisfies the properties of a transition rate matrix, described in the following lemma.

Lemma 1. For all t, the matrix $A(t)$ is a transition rate matrix. That is, for all t

(i): $A(t)\mathbb{1} = \mathbf{0}$.
(ii): for all i, $A_{i,i}(t) \leq 0$, and for all $j \neq i$ $A_{i,j}(t) \geq 0$.

Proof. See Appendix A.1. □

Remark 2. The choice of the FVM and how we discretize the convection and diffusion terms appearing in (9)-(10) is important for $A(t)$ to satisfy the conditions in Lemma 1. This issue is well known in the CFD literature, and also recognized in the related work [2]. If a finite difference method had been used with central differences for both diffusion and convection terms, the resulting $A(t)$ would require restrictive conditions on both σ^2 and $\Delta\lambda$ to satisfy the properties in Lemma 1 [27].

3.2 Temporal discretization

To temporally integrate the dynamics (15) we use a first order Euler approximation with time step $\Delta t > 0$. Making the identifications $\nu_k \triangleq \nu(t_k)$ and $A_k \triangleq A(t_k)$ we have

$$\nu_{k+1} = \nu_k P_k, \quad \text{with} \quad P_k = I + \Delta t A_k. \qquad (16)$$

In the continuous time setting elements of the vector $\nu(t)$ were referred to as, for example, $\nu_{\text{on}}(\lambda^i, t)$. The counterpart to this, in the discrete time setting, is referring to elements of ν_k as, for example, $\nu_{\text{on}}[\lambda^i, k]$. We further have the following.

Lemma 2. The matrix P_k is a Markov transition probability matrix if

$$\forall i, \text{ and } \forall k, \quad 0 < \Delta t \leq |[A_k]_{i,i}|^{-1}.$$

where $[A_k]_{i,i}$ is the i^{th} diagonal element of the matrix A_k.

Proof. From Lemma 1 we have that $P_k\mathbb{1} = I\mathbb{1} + \Delta t A_k\mathbb{1} = \mathbb{1}$ since $A_k\mathbb{1} = 0$. Also from Lemma 1, every element of A_k is non-negative, save for the diagonal elements. Under the hypothesis on A_k, then every diagonal element of $I + \Delta t A_k$ will be in $[0,1]$. □

Remark 3. The bound on the time step Δt given in Lemma 2 is $O(\Delta\lambda)$, which follows from the PDE discretization; see Appendix B . Since $\Delta\lambda = \frac{\lambda^{\text{high}} - \lambda^{\text{min}}}{N}$, as the temperature resolution $\Delta\lambda$ becomes finer the time resolution Δt must also become finer at the same rate. See also Remark 4 for a related comment.

4 Discrete space model of a TCL: structure and grid friendly policies

Recall that the dynamics (16) derived in the previous section was for the thermostat policy. We now delve into the structure of these dynamics so to introduce a BA control input. We first formalize a discrete state space for the dynamics (16). We will then show that the transition matrix in (16) can be written as $P_k = \Phi G_k$ where Φ depends on the thermostat policy and G_k on the TCL temperature dynamics and weather. The isolation of the policy then indicates how a BA could introduce grid friendly policies in place of the thermostat policy Φ.

4.1 Discrete state space

When the conditions of Lemma 2 are met P_k is a transition matrix and hence each ν_k is a marginal pmf if ν_0 is a pmf. The structure of this marginal is given from (7) for the on state (a similar interpretation holds for the off state) as,

$$\nu_{\text{on}}[\lambda^i, k] = \mathsf{P}\left(\theta(t_k) \in \text{CV}(i), \ m(t_k) = \text{on}\right), \qquad (17)$$

where $\theta(t_k)$ is the temperature. Now denote, $\theta_k \triangleq \theta(t_k)$, $m_k \triangleq m(t_k)$, and

$$I_k \triangleq \sum_{i=1}^{N} i \mathbf{I}_{\text{CV}(i)}(\theta_k, m_k). \qquad (18)$$

The quantity I_k indicates which CV the TCLs temperature resides in at time k. It also is a function of m_k since the CV index for the on mode is different from the index for the off mode. We then define the following discrete state space:

$$\mathsf{Z} \triangleq \{m \in \{\text{on}, \text{off}\}, \ I \in \{1, \ldots, N\}\}, \qquad (19)$$

with cardinality $|\mathsf{Z}| = 2N$. Using the newly defined quantity I_k we rewrite the marginals $\nu_{\text{on}}[\lambda^i, k]$ and $\nu_{\text{off}}[\lambda^i, k]$ as functions on Z,

$$\nu_{\text{on}}[\lambda^i, k] = \mathsf{P}\left(I_k = i, \ m_k = \text{on}\right), \quad \text{and} \qquad (20)$$
$$\nu_{\text{off}}[\lambda^i, k] = \mathsf{P}\left(I_k = i, \ m_k = \text{off}\right). \qquad (21)$$

From the above, the matrix P_k (with the conditions of Lemma 2 satisfied) is the transition matrix for the joint

process (I_k, m_k) on the state space Z. The dynamic equation $\nu_{k+1} = \nu_k P_k$ is then a probabilistic model for a TCL with state space Z and operating under the thermostat policy.

4.2 Conditional independence in P_k

In the following, we refer to the values of I_k with i and j and the values of m_k with u and v. We introduce the following notation to refer to the elements of the transition matrix P_k:

$$P_k((i, u), (j, v)) \triangleq \tag{22}$$
$$\mathsf{P}\left(I_{k+1} = j, \ m_{k+1} = v \mid I_k = i, \ m_k = u, \ \theta_k^a = w_k\right).$$

Recall, the matrix P_k is derived for the thermostat policy. We will now show that the matrix P_k can be written as the product of two matrices. One depends only on the thermostat policy (control) and the other depends only on weather and TCL temperature dynamics. That is, we show that each entry of P_k factors as

$$P_k((i, u), (j, v)) = \phi_u^{\mathrm{TS}}(v \mid i) P_k^u(i, j) \tag{23}$$

where, for each given values of θ_k^a, $P_k^u(i, j)$ is a *controlled transition matrix* on Z:

$$P_k^u(i, j) \triangleq \mathsf{P}\left(I_{k+1} = j \mid I_k = i, \ m_k = u, \ \theta_k^a = w_k\right) \tag{24}$$

and $\phi_u^{\mathrm{TS}}(v \mid i)$ is an instance of a *randomized policy* $\phi_u(v \mid i)$ on Z:

$$\phi_u(v \mid i) \triangleq \mathsf{P}\left(m_{k+1} = v \mid I_k = i, \ m_k = u\right). \tag{25}$$

We show the factorization (23) through construction next.

4.2.1 Constructing the factorization

The quantity $\phi_u^{\mathrm{TS}}(v \mid i)$ in (25) is the thermostat policy on Z, which is formally defined as follows.

Definition 1. The thermostat policy on Z is specified by the two vectors, $\phi_{\mathrm{off}}^{\mathrm{TS}}, \phi_{\mathrm{on}}^{\mathrm{TS}} \in \mathbb{R}^N$, where $\phi_{\mathrm{off}}^{\mathrm{TS}} \triangleq \phi_{\mathrm{off}}^{\mathrm{TS}}(\mathrm{on} \mid \cdot) = \mathbf{e}_N$, $\phi_{\mathrm{on}}^{\mathrm{TS}} \triangleq \phi_{\mathrm{on}}^{\mathrm{TS}}(\mathrm{off} \mid \cdot) = \mathbf{e}_1$, and $\phi_{\mathrm{off}}^{\mathrm{TS}}(\mathrm{off} \mid \cdot) \triangleq 1 - \phi_{\mathrm{off}}^{\mathrm{TS}}$, $\phi_{\mathrm{on}}^{\mathrm{TS}}(\mathrm{on} \mid \cdot) \triangleq 1 - \phi_{\mathrm{on}}^{\mathrm{TS}}$.

The quantity $P_k^u(i, j)$ in (24) represents the policy-free (open loop) evolution of the TCL on Z. That is, it describes how the TCLs temperature evolves under a fixed mode. We define matrices with entries $P_k^u(i, j)$ next.

Definition 2. Let $P_k^{\mathrm{off}}, P_k^{\mathrm{on}} \in \mathbb{R}^{N \times N}$ have (i, j) entries given by,

$$P_k^{\mathrm{off}}(i, j) = P_{w_k}((i, \mathrm{off}), (j, \mathrm{off})), \quad i \neq N \text{ and } j \neq N,$$
$$P_k^{\mathrm{on}}(i, j) = P_{w_k}((i, \mathrm{on}), (j, \mathrm{on})), \quad i \neq 1 \text{ and } j \neq 1,$$

with $P_k^{\mathrm{off}}(N, N) = 1$ and $P_k^{\mathrm{on}}(1, 1) = 1$.

The quantities defined in Definition 1 and 2 correspond to entries of P_k. To construct the promised factorization, from these definitions, the idea is to construct its four sub-matrices that correspond to all possible combinations of $u, v \in \{\mathrm{on}, \mathrm{off}\}$ (see Figure 3). For example, the $\mathrm{off} - \mathrm{off}$ quadrant of P_k is given by the matrix product

$$\left(I - \mathrm{diag}(\phi_{\mathrm{off}}^{\mathrm{TS}})\right) P_k^{\mathrm{off}}.$$

However, since the temperature associated with the i^{th} CV for the on mode is not the same temperature associated with the i^{th} CV for the off state (see Figure 2) it is *not* true that the $\mathrm{off} - \mathrm{on}$ quadrant of P_k is given as $\mathrm{diag}(\phi_{\mathrm{off}}^{\mathrm{TS}}) P_k^{\mathrm{off}}$. The entries of the matrix P_k^{off} need to be re-arranged so to correctly account for the difference in CV index between the on/off mode. We define such correctly re-arranged matrices next.

Definition 3. Let $I^{\mathrm{off}} = \{m, \ldots, N\}$, $I^{\mathrm{on}} = \{1, \ldots, q\}$, $m^- = m - 1$, and $S_k^{\mathrm{off}}, S_k^{\mathrm{on}} \in \mathbb{R}^{N \times N}$ with (i, j) entries

$$S_k^{\mathrm{off}}(i, j - m^-) = \begin{cases} P_k^{\mathrm{off}}(i, j) & i, j \in I^{\mathrm{off}} \\ 0 & \text{otherwise.} \end{cases} \tag{26}$$

$$S_k^{\mathrm{on}}(i, j + m^-) = \begin{cases} P_k^{\mathrm{on}}(i, j) & i, j \in I^{\mathrm{on}} \\ 0 & \text{otherwise.} \end{cases} \tag{27}$$

The above definition is based on the construction that $N = q + m$. The quantities in Definition 3 let us construct, e.g., the $\mathrm{off} - \mathrm{on}$ quadrant of P_k as $\mathrm{diag}(\phi_{\mathrm{off}}^{\mathrm{TS}}) S_k^{\mathrm{off}}$.

The next result shows that $P_k = \Phi^{\mathrm{TS}} G_k$ under certain conditions and for appropriate choices of the matrices Φ^{TS} and G_k.

Lemma 3. Let the time discretization period Δt and the parameter α that appears as a design choice in discretizing the PDEs to ODEs be chosen to satisfy $\alpha = (\Delta t)^{-1}$. Let $\Phi_{\mathrm{off}}^{\mathrm{TS}} = \mathrm{diag}(\Phi_{\mathrm{off}}^{\mathrm{TS}})$ and $\Phi_{\mathrm{on}}^{\mathrm{TS}} = \mathrm{diag}(\Phi_{\mathrm{on}}^{\mathrm{TS}})$, and

$$\Phi^{\mathrm{TS}} \triangleq \begin{bmatrix} I - \Phi_{\mathrm{off}}^{\mathrm{TS}} & \Phi_{\mathrm{off}}^{\mathrm{TS}} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \Phi_{\mathrm{on}}^{\mathrm{TS}} & I - \Phi_{\mathrm{on}}^{\mathrm{TS}} \end{bmatrix} \quad \text{and} \tag{28}$$

$$G_k \triangleq \begin{bmatrix} \mathbf{0} & S_k^{\mathrm{off}} & \mathbf{0} & P_k^{\mathrm{on}} \\ P_k^{\mathrm{off}} & \mathbf{0} & S_k^{\mathrm{on}} & \mathbf{0} \end{bmatrix}^T, \tag{29}$$

then

$$P_k = \Phi^{\text{TS}} G_k. \qquad (30)$$

Proof. See Appendix A.2. □

Remark 4. The condition $\alpha = 1/\Delta t$ can be satisfied as long as time and temperature discretization intervals are chosen to satisfy $\Delta t < (\Delta \lambda)^2/\sigma^2$. To understand how, recall that in the discretizing the PDE to the coupled ODEs, a design parameter $\gamma > 0$ appears: some rate of density is transferred out of the control volume λ_{off}^N and into the CV λ_{on}^q (as depicted in Figure 2) due to thermostatic control. The rate of the density transfer is then given as $-\gamma \nu_{\text{off}}(\lambda^N, t)$, where $\gamma > 0$ is a modeling choice and a constant of appropriate units that describes the discharge rate. We then define $\alpha \triangleq D + \gamma$ where $D = \frac{\sigma^2}{(\Delta \lambda)^2}$. Recall that σ^2 is the variance in the Fokker-Planck equation (9)-(10) and $\Delta \lambda$ is the temperature discretization interval. Thus, as long as $1/\Delta t > D$, a positive γ can be chosen while meeting the condition $\alpha = 1/\Delta t$. The condition $1/\Delta t > D$ is equivalent to $\Delta t < (\Delta \lambda)^2/\sigma^2$.

Remark 5. The conditional independence factorization (30) has been a useful assumption in the design of algorithms in [3]. In the present it is a byproduct of our spatial and temporal discretization of the PDEs (9)-(10). There are other works [2,1,25] that develop Markov models for TCLs through discretization of PDEs. However, to our knowledge, our work is the first to uncover this factorization.

Lemma 3 informs us how to define the dynamics of the marginals (20) under a different policy than the thermostat policy, which is described next.

4.3 BA control command = policy

In light of the previous section, an arbitrary randomized policy can replace the thermostat policy to control the state process on Z. That is equivalent to replacing Φ^{TS} in (30) with a new matrix Φ that corresponds to a policy designed for grid support. From the viewpoint of the BA this randomized policy *is* the control input that it must design and broadcast to a TCL. The TCL now implements this policy to make on/off decisions instead of using the thermostat policy. As we shall soon see, if the BA appropriately designs and sends the randomized policy to multiple TCLs it can achieve coordination of the TCLs for grid support.

To distinguish from thermostat policy $\phi_{\text{off}}^{\text{TS}}$ and $\phi_{\text{on}}^{\text{TS}}$ in the prior section that only maintains temperature, we denote the newly introduced policies for providing grid

support with the superscript 'GS'. We require the policies, $\phi_{\text{on}}^{\text{GS}}$ and $\phi_{\text{off}}^{\text{GS}}$, to have the following structure

$$\phi_{\text{off}}^{\text{GS}}(\text{on} \mid j) = \begin{cases} \kappa_j^{\text{on}}, & (m+1) \leq j \leq (N-1). \\ 1, & j = N. \\ 0, & \text{o.w.} \end{cases} \qquad (31)$$

$$\phi_{\text{on}}^{\text{GS}}(\text{off} \mid j) = \begin{cases} \kappa_j^{\text{off}}, & 2 \leq j \leq (q-1). \\ 1, & j = 1. \\ 0, & \text{o.w.} \end{cases} \qquad (32)$$

with $\phi_{\text{off}}^{\text{GS}}(\text{off} \mid \cdot) = 1 - \phi_{\text{off}}^{\text{GS}}(\text{on} \mid \cdot)$ and $\phi_{\text{on}}^{\text{GS}}(\text{on} \mid \cdot) = 1 - \phi_{\text{on}}^{\text{GS}}(\text{off} \mid \cdot)$ and $\kappa_j^{\text{on}}, \kappa_j^{\text{off}} \in [0,1]$ for all j. The policies could also be time varying, for example: $\kappa_j^{\text{off}}[k]$ and $\kappa_j^{\text{on}}[k]$. The dependence of the policies on time is denoted as $\phi_{\text{off}}^{\text{GS}}[k]$ and $\phi_{\text{on}}^{\text{GS}}[k]$. Designing the grid support control policies is then equivalent to choosing the values of $\kappa_j^{\text{on}}[k]$ and $\kappa_j^{\text{off}}[k]$ for all j and k.

We have required $\phi_{\text{off}}^{\text{GS}}(\text{on} \mid j) = 0$ for $1 \leq j \leq m$ since the temperatures corresponding to these indices are below the permitted deadband temperature, λ^{min}. Hence, turning on at these temperature does not make physical sense. The arguments for the zero elements in $\phi_{\text{on}}^{\text{GS}}$ are symmetric.

Remark 6. From the individual TCL's perspective, implementing grid support randomized policies of the form (31)-(32) is straightforward: (i) the TCL measures its current temperature and on/off status, (ii) the TCL "bins" this temperature value according to (18) and (iii) the TCL flips a coin to decide its next on/off state according to the probabilities given in (31)-(32). Note that the thermostat policy is a special case of the grid support control policy, and both policies enforce the temperature constraint.

5 Proposed framework

We are now in a position to present our unified framework for coordination of TCLs. We first expand the state of the model (16) so to incorporate cycling, following [19,26]. We then shift the viewpoint from a single TCL to that of a collection of TCLs (recall Remark 1) to develop our control oriented aggregate model. Using this model we develop a method for designing both reference and policy through convex optimization.

5.1 Individual TCL model with cycling

We now augment the model for a TCL's temperature evolution with cycling dynamics. Recall the cycling constraint: as soon as a TCL switches its mode, the TCL becomes stuck in that mode for τ time instances. This

constraint can be represented as the evolution of a state, specifically, a counter variable. First defining the binary variable s_k as $s_k = 1$ if the TCL is stuck in the current mode at time k and 0 if it is not stuck. The counter variable is defined as follows

$$\mathsf{L}_{k+1} \triangleq \begin{cases} \mathsf{L}_k + 1, & s_k = 1. \\ 0, & s_k = 0. \end{cases} \tag{33}$$

This variable denotes the time spent in the "stuck" mode ($s_k = 1$). A TCL has flexibility to help the grid only when $\mathsf{L}_k = 0$, which means it is not stuck in either the on or off mode. If $\mathsf{L}_k > 0$, it is stuck in either the on or off mode, and switching the mode to help the grid will violate the cycling constraint.

Recall, the discrete state space Z for a TCL included binned temperature and on/off mode. The space Z, the policies $\phi_{\text{on}}^{\text{GS}}$ and $\phi_{\text{off}}^{\text{GS}}$, the marginal pmf ν_k, and the transition matrix P_k (and consequently its factors Φ and G_k) now all need to be expanded to be defined over a state space consisting of (I_k, m_k, L_k). This expansion is described next.

We denote this newly expanded state space as the set of values: $\mathsf{X} \triangleq$

$$\left\{ m \in \{\text{on}, \text{off}\}, \ I \in \{1, \dots, N\}, \ \mathsf{L} \in \{0, \dots, \tau\} \right\}, \tag{34}$$

with cardinality $|\mathsf{X}| = 2N(\tau + 1)$. The policies on the expanded state space are:

$$\phi_{\text{off}}^{\text{E}} = \mathbf{I}_{\{0\}}(\mathsf{L})\phi_{\text{off}}^{\text{GS}} + (1 - \mathbf{I}_{\{0\}}(\mathsf{L}))\phi_{\text{off}}^{\text{TS}}, \quad \text{and} \tag{35}$$
$$\phi_{\text{on}}^{\text{E}} = \mathbf{I}_{\{0\}}(\mathsf{L})\phi_{\text{on}}^{\text{GS}} + (1 - \mathbf{I}_{\{0\}}(\mathsf{L}))\phi_{\text{on}}^{\text{TS}}.$$

To ensure that expanded policy (35) will enforce the cycling constraint, we impose the following restriction at the design stage: a TCL with $\mathsf{L}_k > 0$ will only implement the thermostat policy, and a TCL with $\mathsf{L}_k = 0$ will make on/off decisions based on the grid support policy. The construction in this way ensures a TCL will not violate its cycling and temperature constraints under the conditions in Assumption **A.2** and **A.3**.

Each entry of the expanded policy is denoted as $\phi_{\text{off}}^{\text{E}}(u \mid j, \ l)$ and $\phi_{\text{on}}^{\text{E}}(u \mid j, \ l)$. The expanded marginals are $\nu_{\text{off}}[\lambda^j, l, k]$ and $\nu_{\text{on}}[\lambda^j, l, k]$, and $\nu_{\text{off},l}$ (resp., $\nu_{\text{on},l}$) is shorthand for $\nu_{\text{off}}[\cdot, l, k]$ (resp., $\nu_{\text{on}}[\cdot, l, k]$). In vectorized form, the expanded marginal is $\nu^{\text{E}} = [\nu_{\text{off}}^{\text{E}}, \nu_{\text{on}}^{\text{E}}]$ where $\nu_{\text{off}}^{\text{E}} = [\nu_{\text{off},0}, \dots, \nu_{\text{off},\tau}]$ and $\nu_{\text{on}}^{\text{E}} = [\nu_{\text{on},0}, \dots, \nu_{\text{on},\tau}]$. Define

$$G_k^{\text{E}} \triangleq \begin{bmatrix} \mathbf{0} & D_\tau \otimes S_k^{\text{on}} & \mathbf{0} & C_\tau \otimes P_k^{\text{off}} \\ C_\tau \otimes P_k^{\text{on}} & \mathbf{0} & D_\tau \otimes S_k^{\text{off}} & \mathbf{0} \end{bmatrix}^T, \tag{36}$$
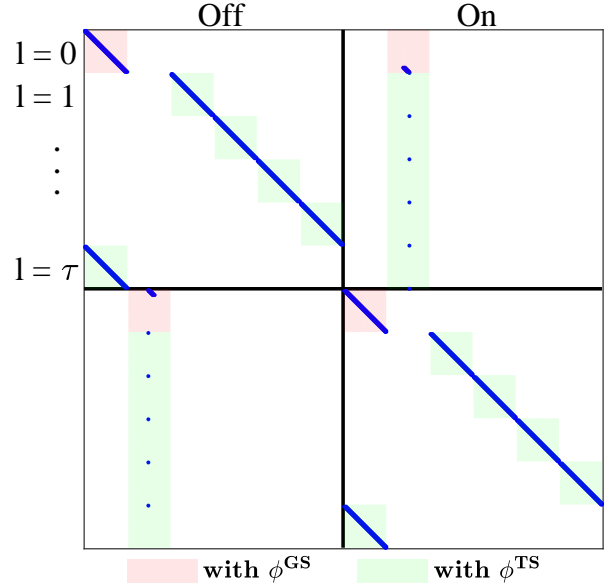

Fig. 4. The sparsity pattern of the expanded transition matrix (the dots represent non-zero entries in the matrix) with $\tau = 5$. Each shaded block is over the entire range of temperature values.

where $D_\tau \triangleq \mathbb{1}^T \otimes \mathbf{e}_2 \in \mathbb{R}^{\tau+1 \times \tau+1}$ and

$$C_\tau \triangleq \begin{bmatrix} 1 & 0 & \mathbf{0}_{\tau-1}^T \\ \mathbf{0}_{\tau-1} & \mathbf{0}_{\tau-1} & I_{\tau-1} \\ 1 & 0 & \mathbf{0}_{\tau-1}^T \end{bmatrix} \in \mathbb{R}^{(\tau+1) \times (\tau+1)}. \tag{37}$$

We define the matrix Φ_k^{E} as having the same structure as (28), but with the expanded policies $\phi_{\text{off}}^{\text{E}}$ and $\phi_{\text{on}}^{\text{E}}$, i.e.,

$$\Phi_k^{\text{E}} \triangleq \begin{bmatrix} I - \Phi_{\text{off}}^{\text{E}}[k] & \Phi_{\text{off}}^{\text{E}}[k] & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \Phi_{\text{on}}^{\text{E}}[k] & I - \Phi_{\text{on}}^{\text{E}}[k] \end{bmatrix}, \tag{38}$$

where $\Phi_{\text{off}}^{\text{E}}[k] \triangleq \text{diag}(\phi_{\text{off}}^{\text{E}}[k])$ and $\Phi_{\text{on}}^{\text{E}}[k] \triangleq \text{diag}(\phi_{\text{on}}^{\text{E}}[k])$. The model of a TCL with cycling dynamics and grid support policy becomes

$$\nu_{k+1}^{\text{E}} = \nu_k^{\text{E}} \Phi_k^{\text{E}} G_k^{\text{E}}. \tag{39}$$

The structure of the transition matrix $\Phi_k^{\text{E}} G_k^{\text{E}}$ is shown in Figure 4. For comparison, the transition matrix with policy ϕ^{GS} and *without* the cycle counter variable would simply be the four red shaded blocks appearing in their respective quadrant. In the expanded system, an on to off mode switch forces probability mass from the red shaded region ($l = 0$ and $m = $ on) to the green shaded region ($l = 1$ and $m = $ off). Mass must then transition through the chain of τ green blocks until it reaches the red block again, so to respect the cycling constraint.

5.2 Aggregate model of a collection of TCLs

We now transition from the viewpoint of a single TCL to that of a collection of N_{tcl} TCLs: $\ell = 1, \ldots, N_{tcl}$. For example, m_k^ℓ and I_k^ℓ are the mode and binned temperature of the ℓ^{th} TCL at time k. Recall Remark 1, the model (39) also describes an entire collection of TCLs. For a single TCL, we view the state ν_k^E as a marginal but for a collection of TCLs we expect the marginal pmf ν_k^E to approximate the histogram

$$h_k[u, i, l] \triangleq \frac{1}{N_{tcl}} \sum_{\ell=1}^{N_{tcl}} \left(\mathbf{I}_{\{i\}}(I_k^\ell) \mathbf{I}_{\{u\}}(m_k^\ell) \mathbf{I}_{\{l\}}(L_k^\ell) \right), \quad (40)$$

for each state $(u, i, l) \in X$ as $N_{tcl} \to \infty$. In the same regard, we define

$$\gamma_k^E \triangleq \nu_k^E C^E, \quad \text{where} \quad C^E \triangleq [\mathbf{0}^T, P_{agg}\mathbb{1}^T]^T, \quad (41)$$

where P_{agg} is the maximum possible power of the collection, defined in (4). We expect γ_k^E to approximate the total power consumption y_k of the collection of N_{tcl} TCLs:

$$y_k \triangleq P \sum_{\ell=1}^{N_{tcl}} m_k^\ell. \quad (42)$$

which is the discrete-time equivalent of $y(t)$ defined in (5). That is, we expect $\gamma_k^E \approx y_k$ for large N_{tcl}, based on a law of large numbers argument [5]. The *control oriented aggregate model of a TCL collection* is the dynamics (39) together with the output (41):

$$\nu_{k+1}^E = \nu_k^E \Phi_k^E G_k^E. \quad \text{and} \quad \gamma_k^E = \nu_k^E C^E. \quad (43)$$

5.2.1 Evaluating the aggregate model

Before proceeding to policy design with our developed model (43), we first show that it is effective in modeling a population of TCLs. We do this by comparing the state of the model to (40) and (42) obtained from a simulation of $N_{tcl} = 50,000$ air conditioning TCLs.

The comparison results are shown in Figure 5 and Figure 6. The mode state of each TCL evolves according to a control policy, where the ϕ_{off}^{GS} and ϕ_{on}^{GS} portion are shown in Figure 6 (bottom). The policy is arbitrary, designed merely to be an example of a non-thermostat policy. This policy satisfies the structure in (31) and (32) so that both temperature and cycling constraints are satisfied at each TCL. The temperature evolution evolves according to (2). We see the state ν_k^E matches the histogram h_k of the collection for the devices that are not stuck (Figure 5 (top)) and for the devices that are stuck (Figure 5 (bottom)). Additionally, the output of the aggregate model, γ_k^E, matches it's empirical counterpart y_k (shown in Figure 6 (top)).

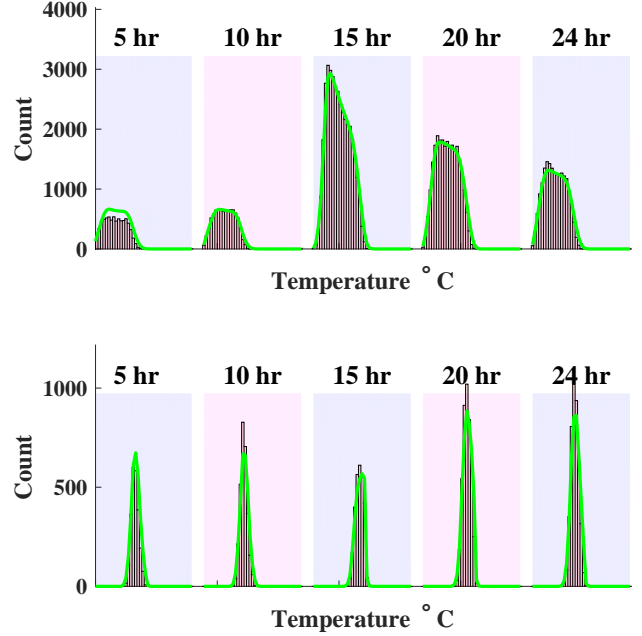Fig. 5. (Top): Histogram of the collection for the devices that are on and not stuck. (Bottom): Histogram of the collection for the devices that are on and are stuck.
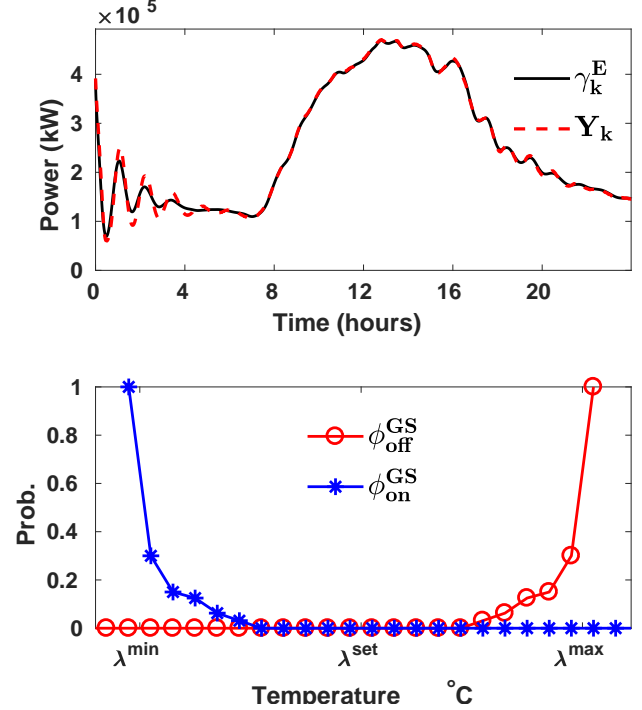

Fig. 6. (Top): Comparison of the output of the expanded aggregate model γ_k^E and the ensembles power consumption y_k. (Bottom): The policies ϕ_{off}^{GS} and ϕ_{on}^{GS} used for the numerical experiment in Section 5.2.

5.3 Grid support Policy design

The goal of coordinating TCLs is to help the BA balance supply and demand of electricity in the grid. We denote r_k^{BA} as the desired demand from all flexible loads and batteries that will reduce the imbalance to 0. It is unreasonable to expect any collection of TCLs to meet the entire desired demand r_k^{BA} while maintaining their QoS. Only a portion of r_k^{BA} can be supplied by TCLs, and we denote this portion by r_k. Determining r_k becomes an optimal control problem due to the time coupling produced by the TCL dynamics. We consider a planning horizon of T_{plan}. To simultaneously design grid support control policies $\phi_{\mathrm{off}}^{\mathrm{GS}}[k]$ and $\phi_{\mathrm{on}}^{\mathrm{GS}}[k]$ and determine a suitable reference signal r_k over T_{plan} the BA solves the following optimization problem,

$$\eta^* = \min_{\nu_k^{\mathrm{E}},\Phi_k^{\mathrm{E}}} \eta(\hat{\nu}) = \sum_{k\in\mathsf{T}} \left(r_k^{\mathrm{BA}} - \gamma_k^{\mathrm{E}}\right)^2 \tag{44}$$

$$\text{s.t. } \nu_{k+1}^{\mathrm{E}} = \nu_k^{\mathrm{E}}\Phi_k^{\mathrm{E}}G_k^{\mathrm{E}}, \quad \nu_{\mathsf{T}(0)}^{\mathrm{E}} = \hat{\nu}, \tag{45}$$

$$\gamma_k^{\mathrm{E}} = \nu_k^{\mathrm{E}}C^{\mathrm{E}}, \quad \nu_k^{\mathrm{E}} \in [0,1], \quad \Phi_k^{\mathrm{E}} \in \Phi. \tag{46}$$

The solution at time k is denoted $r_k \triangleq \gamma_k^{\mathrm{E},*}$, $\phi_{\mathrm{off}}^{\mathrm{GS},*}[k]$, and $\phi_{\mathrm{on}}^{\mathrm{GS},*}[k]$. We have $\mathsf{T} \triangleq \{\mathsf{T}(0),\dots,\mathsf{T}(0)+T_{\mathrm{plan}}-1\}$ is the index set of times, $\mathsf{T}(0)$ denotes the initial time index, $\hat{\nu}$ is the initial condition, and $\nu_k^{\mathrm{E}} \in [0,1]$ holds elementwise. The set Φ collects all of the constraints on the policy. This includes the equality constraints set by the structural requirements in (31)-(32) and (35) as well as the structural requirement in (38). These constraints require certain elements of the policy to be either zero or one. The policy should also be a valid conditional pmf and its elements in $[0,1]$. Hence, the set Φ is the following convex set

$$\Phi \triangleq \Big\{\Phi \in \mathbb{R}_{[0,1]}^{|\mathsf{X}|\times 2|\mathsf{X}|} \mid \Phi \text{ satisfies (38)}, \ \mathbb{1} = \Phi\mathbb{1},$$
$$\phi_{\mathrm{off}}^{\mathrm{GS}} \text{ satisfies (31)},$$
$$\phi_{\mathrm{on}}^{\mathrm{GS}} \text{ satisfies (32), and}$$
$$\phi_{\mathrm{off}}^{\mathrm{E}} \text{ and } \phi_{\mathrm{on}}^{\mathrm{E}} \text{ satisfy (35)}\Big\}. \tag{47}$$

Where, e.g., $\mathbb{R}_{[0,1]}^{|\mathsf{X}|\times|\mathsf{X}|}$ is the set of $|\mathsf{X}| \times |\mathsf{X}|$ matrices with elements in $[0,1]$.

5.3.0.1 QoS + Solution of (44)

(1) The equality constraints in Φ are present to ensure the individual TCL's QoS constraints: the structure (35) ensures the cycling constraint and the structure (31)-(32) ensures the temperature constraint. Recall that this structure guarantees QoS by requiring the policy to place zero probability on state transitions that would violate QoS.

(2) A solution to (44) yields, for $k \in \mathsf{T}$, two things: (i) the optimal randomized policies $\phi_{\mathrm{off}}^{\mathrm{GS},*}[k]$ and $\phi_{\mathrm{on}}^{\mathrm{GS},*}[k]$ and (ii) an optimal reference for the power demand of the TCL collection $r_k(=\gamma_k^{\mathrm{E},*})$. The reference is optimal in the following sense: among all power demand signals the collection can track without requiring any TCL to violate its local QoS constraints in so doing, it is the closest to the BA's desired demand r^{BA} in 2-norm. The reference is also the predicted power consumption of the TCLs whilst using the policies $\phi_{\mathrm{off}}^{\mathrm{GS},*}[k]$ and $\phi_{\mathrm{on}}^{\mathrm{GS},*}[k]$.

Remark 7. Since the reference $r_k(= \gamma_k^{\mathrm{E},*})$ from (44) is the best the TCLs can do to help the BA without any TCL having to violate its QoS, Problem (44) therefore also provides an answer to the "aggregate flexibility" question: how much can a collection of TCLs vary their demand while maintaining their local QoS constraints. This question has been investigated by many works [25,13,10,15].

5.3.1 Convex control synthesis

The problem (44) is non-convex due to the product $\nu_k^{\mathrm{E}}\Phi_k^{\mathrm{E}}$ in the constraint. A well known convexification remedy for (44) is to consider optimizing over the marginal and joint distribution instead of the marginal and the policy [22,2]. Using our identified structure from Section 4.2 we construct the following joint distribution (written in matrix form):

$$J_k = \mathrm{diag}(\nu_k^{\mathrm{E}})\Phi_k^{\mathrm{E}} \in \mathbb{R}^{|\mathsf{X}|\times 2|\mathsf{X}|}. \tag{48}$$

By construction, we have that $\nu_{k+1}^{\mathrm{E}} = \mathbb{1}^T J_k G_k^{\mathrm{E}}$ and $(\nu_k^{\mathrm{E}})^T = J_k\mathbb{1}$ since $\mathbb{1}^T\mathrm{diag}(\nu_k^{\mathrm{E}}) = \nu_k^{\mathrm{E}}$ and $\mathbb{1} = \Phi_k^{\mathrm{E}}\mathbb{1}$. It is straightforward to convert the constraint set $\Phi_k^{\mathrm{E}} \in \Phi$ to the new decision variables. For the equality constraints in Φ if we have that $\phi_{\mathrm{off}}^{\mathrm{E}}(u \mid j,l) = \kappa$, then in the decision variables J_k and ν_k^{E} we will have a linear constraint of the form

$$\mathsf{P}\,(m_{k+1} = u, \ I_k = j, \ \mathsf{L}_k = l, \ m_k = \mathrm{off})$$
$$= \kappa\nu_{\mathrm{off}}[\lambda^j, l, k], \tag{49}$$

where the LHS of the above is some element in the matrix J_k. In addition to the above equality constraints, requiring both J_k and ν_k^{E} to be within $[0,1]$ and the constraint $(\nu_k^{\mathrm{E}})^T = J_k\mathbb{1}$ will allow one to reconstruct a policy $\Phi_k^{\mathrm{E}} \in \Phi$ from J_k and ν_k^{E} (described shortly in Lemma 4). We denote the transcription of $\Phi_k^{\mathrm{E}} \in \Phi$ to the new variables as $(J_k, \nu_k^{\mathrm{E}}) \in \bar{\Phi}$. Optimizing over J_k

and ν_k^{E} yields the convex program:

$$
\begin{aligned}
\eta^* = \min_{\nu_k^{\mathrm{E}}, J_k} \ \eta(\hat{\nu}) &= \sum_{k \in \mathsf{T}} \left(r_k^{\mathrm{BA}} - \gamma_k^{\mathrm{E}} \right)^2 \\
\text{s.t.} \quad \nu_{k+1}^{\mathrm{E}} &= \mathbb{1}^T J_k G_k^{\mathrm{E}}, \quad \nu_{\mathsf{T}(0)}^{\mathrm{E}} = \hat{\nu}, \quad \gamma_k^{\mathrm{E}} = \nu_k^{\mathrm{E}} C^{\mathrm{E}}, \\
\nu_k^{\mathrm{E}}, J_k &\in [0,1], \ (\nu_k^{\mathrm{E}})^T = J_k \mathbb{1}, \ (J_k, \nu_k^{\mathrm{E}}) \in \bar{\Phi}.
\end{aligned}
\tag{50}
$$

Once the convex problem is solved, the grid support control policies need to be recovered from it by using the relation (48). If the matrix $\mathrm{diag}(\nu_k^{\mathrm{E}})$ is invertible, then the policy can be obtained trivially from inversion of $\mathrm{diag}(\nu_k^{\mathrm{E}})$. If $\mathrm{diag}(\nu_k^{\mathrm{E}})$ is not invertible, then slight care is required when reconstructing a policy from the solution of (50). We describe this in the following Lemma.

Lemma 4. Suppose for all $k \in \mathsf{T}$ that ν_k^{E} and J_k satisfy the constraints in problem (50). Then, there exists matrices $H_k = H_k(\nu_k^{\mathrm{E}})$ and $W_k = W_k(\nu_k^{\mathrm{E}})$ so that for all $k \in \mathsf{T}$ the quantity $\Phi_k^{\mathrm{E}} = H_k J_k + W_k$ satisfies (48) and $\Phi_k^{\mathrm{E}} \in \Phi$.

Proof. See Appendix A.3. □

Exact construction of H_k and W_k is given in the proof of Lemma 4. *Hence, the proof of Lemma 4 provides an algorithm for computing grid support control policies that are feasible for the problem (44) from the solutions of the convex problem (50).* Further the two problems have a certain equivalence described here in the following Theorem.

Theorem 1. Denote η_{CVX}^* the optimal cost for (50) and η_{NCVX}^* the optimal cost for (44) we have that $\eta_{\mathrm{CVX}}^* = \eta_{\mathrm{NCVX}}^*$.

Proof. See Appendix A.4. □

This result, for a similar problem setup, is also reported in [2]. While we have no guarantee on the difference of the argument minimizers (and hence the policies obtained from both), Theorem (1) says that the policies will produce the same tracking performance. Further, from Lemma 4, the policies produced from either problem are guaranteed to ensure TCL QoS.

5.3.2 Computational considerations

The dimension of the program (50) can be quite large, so that even though it is convex obtaining a solution requires care. We discuss now some practical considerations that we found necessary to consider when solving the problem (50).

Due to the structure of Φ_k^{E}, we do not need to declare every element in the matrix J_k as a decision variable since many of these elements will be zero. For instance, we see that $\mathrm{diag}(\nu_k^{\mathrm{E}}) \Phi_k^{\mathrm{E}}$ is a block matrix, where further each matrix block is diagonal. We express this as: $\mathrm{diag}(\nu_k^{\mathrm{E}}) \Phi_k^{\mathrm{E}} =$

$$
\begin{bmatrix} B_{\mathrm{off,off}}[k] & B_{\mathrm{off,on}}[k] & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & B_{\mathrm{on,off}}[k] & B_{\mathrm{on,on}}[k] \end{bmatrix}
$$
$$
\overline{\overline{\nabla}} \ \ \mathrm{sparse}(J_k),
\tag{51}
$$

where, e.g., $B_{\mathrm{off,off}}[k] = \mathrm{diag}(\nu_{\mathrm{off}}^{\mathrm{E}}[k])(I - \Phi_{\mathrm{off}}^{\mathrm{E}}[k])$. The other diagonal matrices appearing in (51) can be inferred by carrying out the matrix multiplication.

If J_k was declared directly as a decision variable the problem (50) would have $(8N^2 + 2N(\tau + 1))T_{\mathrm{plan}}$ primal variables, whereas the problem with $\mathrm{sparse}(J_k)$ as a decision variable only has $2NT_{\mathrm{plan}}(\tau + 3)$ primal variables. As an example, consider $N = 12$, $T_{\mathrm{plan}} = 360$, and $\tau = 5$, which are values used in numerical results reported later. The problem (50) without the structure exploited has ≈ 0.5 million decision variables, but only $\approx 75,000$ when the structure is exploited.

We also have found it helpful to include constraints of the form,

$$
\phi_{\mathrm{off}}^{\mathrm{GS}}(\mathrm{on} \mid j-1)\nu_{\mathrm{off}}[\lambda^{j-1}, 0, k] \le \phi_{\mathrm{off}}^{\mathrm{GS}}(\mathrm{on} \mid j)\nu_{\mathrm{off}}[\lambda^j, 0, k],
\tag{52}
$$
$$
\phi_{\mathrm{on}}^{\mathrm{GS}}(\mathrm{off} \mid j+1)\nu_{\mathrm{on}}[\lambda^{j+1}, 0, k] \le \phi_{\mathrm{on}}^{\mathrm{GS}}(\mathrm{off} \mid j)\nu_{\mathrm{on}}[\lambda^j, 0, k],
\tag{53}
$$

so to suggest that the switching on (resp., switching off) probability increases as temperature increases (resp., decreases). Adding the constraints (52)-(53) to the problem (50) is straightforward as both the LHS and RHS of the inequalities are elements in the matrix J_k.

Matlab implementation of (50) and the algorithm to extract the policies from J_k (described in the proof of Lemma 4) is available at [8].

5.3.3 Communication burden

Once solved, the policies obtained from (50) need to be sent to each individual TCL. Many of the policy state values are constrained to either zero or one, which could be pre-programmed into each TCL. At each time index, $q - 2$ (for the on to off policy) plus $N - m - 1$ (for the off to on policy) numbers are not constrained and need to be sent from the BA to each TCL. Recall that the numbers m and q are temperature bin indices (see Figure 2) and N is the number of temperature bins. For illustrative purposes, consider the values used in numerical experiments reported in the sequel: $N = 12$ with

$q = 10$ and $m = 2$ and a time discretization $\Delta t = 1$ minute. Since $N = q + m$, then the BA has to broadcast $2(q-1) = 18$ numbers every 1 minute to the TCLs. Each TCL receives the same 18 numbers.

Communication from TCLs to the BA - about their temperature and on/off state - is needed at the beginning of every planning period so that the BA can determine the initial condition $\hat{\nu}$ in (50). The frequency of this feedback is a design choice. In our numerical simulations reported later, a planning horizon of 6 hours was used, and this feedback was necessary only once in six hours. More frequent loop closure may be needed for higher robustness to uncertainty in weather prediction etc., a topic outside the scope of this paper.

6 Numerical experiments

Simulation involving coordination of $\mathsf{N}_{\text{tcl}} = 20,000$ TCLs through our proposed framework is presented here. Recall the two parts of the coordination architecture shown in Figure 1: (i) planning and (ii) real time control. Planning refers to the solution of the problem (50) at the BA to compute the following two things for the planning period T:

(1) r_k: the reference power consumption of the TCL collection, given the problem data r_k^{BA}.
(2) $\phi_{\text{off}}^{\text{GS},*}[k]$ and $\phi_{\text{on}}^{\text{GS},*}[k]$: grid support control policies for each TCL.

This computation is performed at $\mathsf{T}(0)$. Real time control is then the implementation of the grid support policies by each TCL to make on/off decisions in real time. We imagine the BA broadcasts the policies $\phi_{\text{off}}^{\text{GS},*}[k]$ and $\phi_{\text{on}}^{\text{GS},*}[k]$ at each k, though it can also broadcast all the policies, for all $k \in \mathsf{T}$, at $\mathsf{T}(0)$ and not broadcast again until the beginning of the next planning horizon.

The goal of the numerical simulations of real time control is to show the following.

(1) When each TCL uses the policies $\phi_{\text{off}}^{\text{GS},*}[k]$ and $\phi_{\text{on}}^{\text{GS},*}[k]$ to decide on/off actuation, the collection's power demand indeed tracks r_k.
(2) Every TCL's QoS constraints - both temperature and cycling - are satisfied at all times.

Temperature of each TCL is computed in these simulations with the ODE model (1).

6.1 Planning

The demand needed for demand-supply imbalance at the BA, r_k^{BA}, is chosen arbitrarily, and shown in Figure 7

Table 1
Simulation Parameters

Par.	Unit	value	Par.	Unit	value
N_{tcl}	N/A	2×10^4	η	$\frac{\text{kW-e}}{\text{kW-th.}}$	2.5
C	kWh/°C	1	P_0	kW	5.5
$\lambda^{\min}$	°C	20	$\lambda^{\max}$	°C	22
$(\Delta t)\tau$	Mins.	5	P_{agg}	MW	110
R	°C/kW	2	Δt	Mins.	1
q	N/A	10	m	N/A	2
N	N/A	12	T_{plan}	N/A	360

(top). It is infeasible for the collection: sometimes negative and sometimes far higher than the maximum power demand of the collection. This is done to simulate a realistic scenario in which many sources of demand and generation, not just TCLs, are managed by the BA.

The baseline demand trajectory is defined by the equation (4), which is approximately the power consumption for this collection of air conditioners under thermostat control. The ambient air temperature is time varying and is obtained from `wunderground.com` for a typical summer day in Gainesville, Florida, USA. The other parameters that affect the Markov model are shown in Table 1.

Planning computations are done with Matlab and CVX [14] using a desktop Linux machine, with $N = 12$, and for a six hour planning horizon with 1 minute discretization ($T_{\text{plan}} = 360$). The problem (50) takes about a minute to solve. The quantity r_k^{BA}, the baseline power $\bar{P}_k$, and the reference signal r_k, obtained from solving (50), are shown in Figure 7 (top). The optimal reference for the collection, r_k, is as close to r_k^{BA} as the dynamics of TCLs allows without violating their QoS constraints; recall Remark 7. Figure 7 (bottom) shows the two grid support control polices for one time instant.

6.2 Real time control

The power consumption of the collection making on/off decisions according to the obtained policies is shown in Figure 8 (top). The figure shows that the TCLs are able to collectively track the reference signal r_k. We emphasize that the computational effort at each TCL is negligible. Recall Remark 6: once a TCL receives a grid support policy (≈ 18 floating point numbers, see Section 5.3.3) it only has to measure its current state (temperature and on/off mode) and generate a uniformly distributed random number in $[0, 1]$ to implement the policy.

Verification of the grid support policies in ensuring QoS is shown in Figure 8. The bottom plots shows a histogram of the times between switches for 300 randomly chosen TCLs. The middle plot shows a histogram of
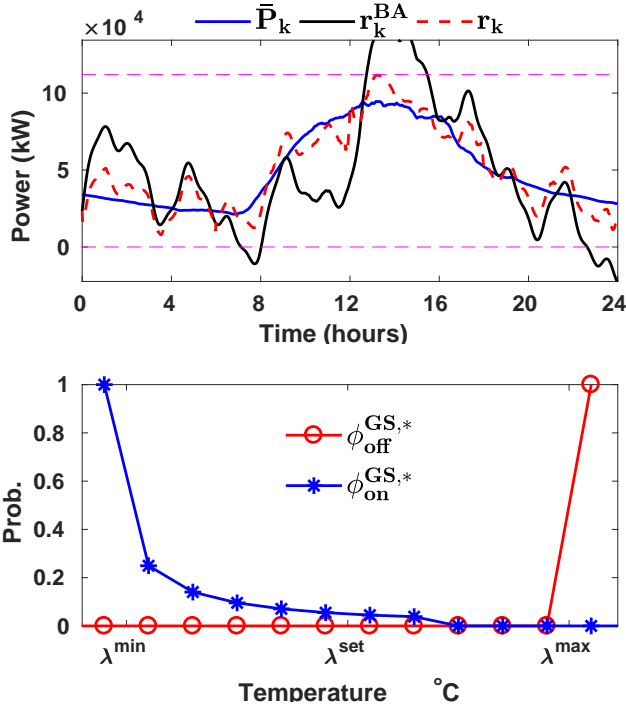
Fig. 7. (Top): The quantity r_k obtained from solving (50), the dashed horizontal lines represent all of the TCLs on (top line) and off (bottom line). (Bottom): Grid support control policies, obtained from solving (50), at one time instance.

temperature from 200 randomly chosen TCLs' temperature trajectories. The histograms show that the policies designed with (50) indeed satisfy the QoS constraints, which is specified by the vertical lines in the figures. Some TCLs do escape the temperature deadband by a little bit, which is expected and occurs also in thermostatic control: the sensor must *first* register a value outside the deadband in order decide to switch the on/off state.

7 Conclusion

In this work we present a unified framework for the distributed control of TCLs. The framework enables: (i) reference planning for a collection of TCLs and (ii) design of a randomized control policy for the individual TCLs, so that both the BA's requirement and consumers' QoS are satisfied. The resulting framework is (i) scalable to an arbitrary number of loads and is implemented through *local* feedback and minimal communication, (ii) able to guarantee both temperature and cycling constraints maintenance in each TCL, and (iii) based on convex optimization. Matlab/cvx implementation is publicly available [8].

There are several avenues for future work. The optimal control problem is solved in an open-loop fashion here. Feedback from TCLs is used only to compute an initial condition that is needed as problem data for the off- line
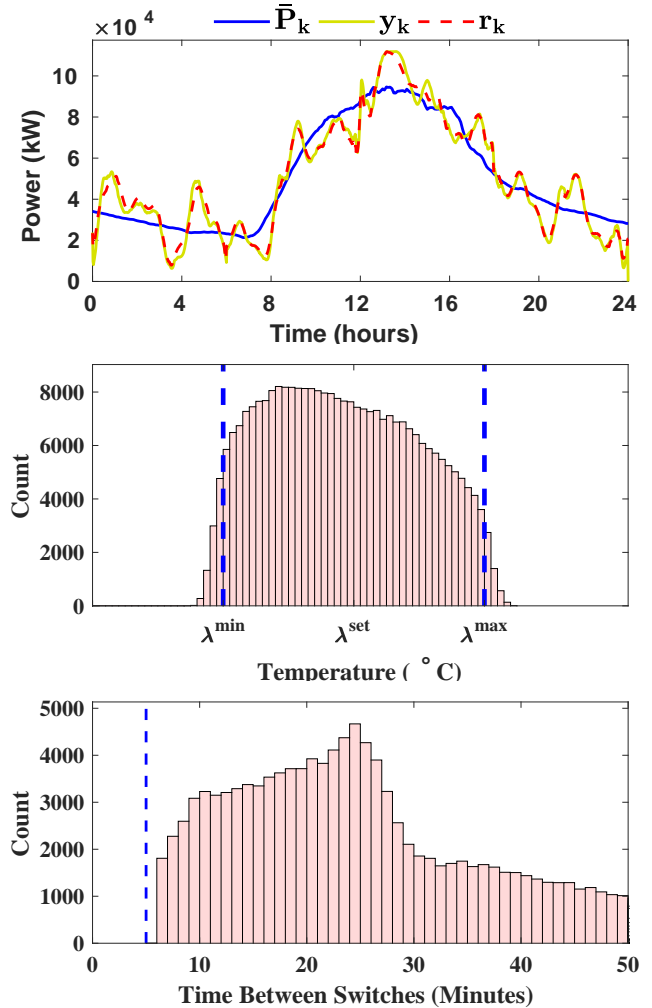

Fig. 8. (Top): Reference tracking results for the TCLs under the influence of the grid support control policies obtained by solving (50). (Middle): Histogram of the 200 TCL's temperature trajectories over the entire simulation horizon. (Bottom): Histogram of the time between switches over 3000 TCLs with the vertical line representing the minimum allowable time between switches.

planning problem. It is straightforward to close the loop between the TCL collection and the BA with greater frequency for robustness to uncertainty in weather forecast and TCL parameters. It will be of interest to identify scenarios where closing loop, say, by using Model Predictive Control, is (i) necessary, and (ii) at what frequency should information be communicated from the TCLs to the BA. Another avenue is to investigate how the problem (50) could be solved at each TCL, intermittently, instead of at the BA. Since the computational power of the processor at each TCL is lower than that of the processor at the BA, online distributed algorithms for convex optimization could play a role. The Fokker-Planck equations from [21] we used here are convenient for modeling TCL populations with a small deegree of heterogeneity. Distributed computation of optimal poli-

cies locally at each TCLs may help extend the method to a highly heterogeneous population of TCLs.

References

[1] Saeid Bashash and Hosam K Fathy. Modeling and control of aggregate air conditioning loads for robust renewable power management. *IEEE Transactions on Control Systems Technology*, 21(4):1318–1327, 2013.

[2] Emilio Benenati, Marcello Colombino, and Emiliano Dall'Anese. A tractable formulation for multi-period linearized optimal power flow in presence of thermostatically controlled loads. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 4189–4194. IEEE, 2019.

[3] Ana Bušić and Sean Meyn. Distributed randomized control for demand dispatch. In *IEEE conference on decision and control*, pages 6964–6971, 2016.

[4] D.S. Callaway and I.A. Hiskens. Achieving controllability of electric loads. *Proceedings of the IEEE*, 99(1):184–199, 2011.

[5] Yue Chen, Ana Bušić, and Sean P. Meyn. State estimation for the individual and the population in mean field control with application to demand dispatch. *IEEE Transactions on Automatic Control*, 62(3):1138–1149, 2017.

[6] Yue Chen, Md Umar Hashmi, Joel Mathias, Ana Bušić, and Sean Meyn. Distributed control design for balancing the grid using flexible loads. In *IMA Volume on the Control of Energy Markets and Grids*, pages 1–26. 2017.

[7] Chi-Min Chu and Tai-Lang Jong. A novel direct air-conditioning load control method. *IEEE transactions on power systems*, 23(3):1356–1363, 2008.

[8] Austin Coffman. Distributed control of TCLs through convex optimization. https://gitlab.com/austinrc925/distributed-control-of-tcls-through-convex-optimization/-/tree/main, 2021.

[9] Austin Coffman, Ana Bušić, and Prabir Barooah. Virtual energy storage from TCLs using QoS preserving local randomized control. In *5th ACM International Conference on Systems for Built Environments (BuildSys)*, page 10, November 2018.

[10] Austin Coffman, Neil Cammardella, Prabir Barooah, and Sean P. Meyn. Aggregate capacity of TCLs with cycling constraints. *ArXiV.org*, October 2020. arXiv:1909.11497.

[11] Austin R. Coffman, Ana Bušić, and Prabir Barooah. Aggregate capacity for TCLs providing virtual energy storage with cycling constraints. In *IEEE Conference on Decision and Control*, December 2019.

[12] Austin R. Coffman, Ana Bušić, and Prabir Barooah. Control oriented modeling of TCLs. In *American Control Conference*, May 2021. accepted, extended version in arXiv:2009.12960.

[13] Austin R. Coffman, Zhong Guo, and Prabir Barooah. Characterizing capacity of flexible loads for providing grid support. *IEEE Transactions on Power Systems*, 36:2428 – 2437, May 2021.

[14] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 1.21. http://cvxr.com/cvx, February 2011.

[15] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent. Aggregate flexibility of thermostatically controlled loads. *IEEE Transactions on Power Systems*, 30(1):189–198, Jan 2015.

[16] Emre C Kara, Mario Bergés, and Gabriela Hug. Impact of disturbances on modeling of thermostatically controlled loads for demand response. *IEEE Transactions on Smart Grid*, 6(5):2560–2568, 2015.

[17] Adil Khurram, Roland Malhamé, Luis Duffaut Espinosa, and Mads Almassalkhi. Identification of hot water end-use process of electric water heaters from energy measurements. *Electric Power Systems Research*, 189:106625, 2020.

[18] Zachary E. Lee, Qingxuan Sun, Zhao Ma, Jiangfeng Wang, Jason S. MacDonald, and K. Max Zhang. Providing Grid Services With Heat Pumps: A Review. *ASME Journal of Engineering for Sustainable Buildings and Cities*, 1(1), 01 2020. 011007.

[19] M. Liu and Y. Shi. Model predictive control of aggregated heterogeneous second-order thermostatically controlled loads for ancillary services. *IEEE Transactions on Power Systems*, 31(3):1963–1971, May 2016.

[20] M. Liu, Y. Shi, and X. Liu. Distributed MPC of aggregated heterogeneous thermostatically controlled loads in smart grid. *IEEE Transactions on Industrial Electronics*, 63(2):1120–1129, 2016.

[21] Roland Malhame and Chee-Yee Chong. Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system. *IEEE Transactions on Automatic Control*, 30(9):854–860, 1985.

[22] Alan S. Manne. Linear programming and sequential decisions. *Management Science*, 6(3):259–267, 1960.

[23] J. L. Mathieu, S. Koch, and D. S. Callaway. State estimation and control of electric loads to manage real-time energy imbalance. *IEEE Transactions on Power Systems*, 28:430–440, 2013.

[24] Md Salman Nazir and Ian Hiskens. Analysis of synchronization in load ensembles. *Electric Power Systems Research*, 190:106779, 2021.

[25] Dario Paccagnan, Maryam Kamgarpour, and John Lygeros. On the range of feasible power trajectories for a population of thermostatically controlled loads. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 5883–5888, 2015.

[26] L. C. Totu, R. Wisniewski, and J. Leth. Demand response of a TCL population using switching-rate actuation. *IEEE Transactions on Control Systems Technology*, 25(5):1537–1551, 2017.

[27] Henk Kaarle Versteeg and Weeratunge Malalasekera. *An introduction to computational fluid dynamics: the finite volume method*. Pearson education, 2007.

[28] Wei Zhang, Jianming Lian, Chin-Yao Chang, and Karanjit Kalsi. Aggregated modeling and control of air conditioning loads for demand response. *IEEE Transactions on Power Systems*, 28(4):4655–4664, 2013.

A   Proofs

A.1   Proof of Lemma 1

See Appendix B before reading this proof. Property (ii) is a consequence of the upwind difference scheme used. We see that for the internal CVs we have

$$\text{off CVs:} \quad -\left(F_{\text{off}}^{i,+} + D\right) \quad\quad \text{(A.1)}$$

$$\text{on CVs:} \quad \left(F_{\text{on}}^{i,-} - D\right). \quad\quad \text{(A.2)}$$

From Assumption **A.2** we have that $F_{\text{on}}^{i,-} \leq 0$ and $F_{\text{off}}^{i,-} \geq 0$ so that both of the above terms are negative. The upwind scheme is what ensured appropriate sign was added to the terms $F_{\text{on}}^{i,-}$ and $F_{\text{off}}^{i,-}$ so that the above coefficients are negative. Similar arguments can be applied for the off diagonal terms of the internal CVs and the boundary CVs.

To show property (i) we consider solely an internal CV for the off state as the arguments for all other CVs are identical in structure. Note that showing $A(t)\mathbb{1} = \mathbf{0}$ is equivalent to $\mathbb{1}^T \mathcal{A}(t) = \mathbf{0}^T$. Hence we need to show, for an arbitrary i that all coefficients acting on $\nu_{\text{off}}(\lambda^i, t)$ sum to 0. We collect the coefficients corresponding to $\nu_{\text{off}}(\lambda^i, t)$:

$$\begin{aligned} \text{From CV}(i): \quad & -F_{\text{off}}^{i,+}(t) - D. \\ \text{From CV}(i-1): \quad & \frac{D}{2}. \\ \text{From CV}(i+1): \quad & \frac{D}{2} + F_{\text{off}}^{i+1,-}(t). \end{aligned}$$

We then require the sum of these coefficients to be zero for all t and any index i for the internal off CVs, adding yields

$$F_{\text{off}}^{i+1,-}(t) - F_{\text{off}}^{i,+}(t) = \frac{f_{\text{off}}(\lambda^{i+1,-}, t) - f_{\text{off}}(\lambda^{i,+}, t)}{\Delta\lambda} = 0$$

since by construction $\lambda^{i+1,-} = \lambda^{i,+}$ for the off CV's. This procedure can be repeated for $\nu_{\text{off}}(\lambda^i, t)$ with $i \in \{1, m, N\}$, i.e., the boundary CVs in the off state and all of the on CVs in a similar fashion.

A.2   Proof of Lemma 3

If $\alpha = (\Delta t)^{-1}$, the diagonal elements of A_k with α in them will go to zero and the non diagonal elements will go to 1. These non-diagonal elements with value 1 are the red dots in Figure 3 and encapsulate the thermostat control law. Thus the construction of Φ^{TS} with the canonical basis vectors. Now, multiplying out the matrix we have,

$$\Phi^{\text{TS}}G_k = \begin{bmatrix} (I - \Phi_{\text{off}}^{\text{TS}})P_k^{\text{off}} & \Phi_{\text{off}}^{\text{TS}}S_k^{\text{off}} \\ \Phi_{\text{on}}^{\text{TS}}S_k^{\text{on}} & (I - \Phi_{\text{on}}^{\text{TS}})P_k^{\text{on}} \end{bmatrix} \quad \text{(A.3)}$$

where $(I - \Phi_{\text{off}}^{\text{TS}})P_k^{\text{off}}$ (respectively, $(I - \Phi_{\text{on}}^{\text{TS}})P_k^{\text{on}}$) is the matrix P_k^{off} (respectively, P_k^{on}) but with the last (respectively, first) row zeroed out. The exact opposite statement is true for $\Phi_{\text{off}}^{\text{TS}}P_k^{\text{on}}$ and $\Phi_{\text{on}}^{\text{TS}}P_k^{\text{off}}$. Hence, by definition of the matrices in G_k we have $P_k = \Phi^{\text{TS}}G_k$ where each non-zero element holds the interpretation (23).

A.3   Proof of Lemma 4

We define the following transformation for $l \in \{0, \ldots, \tau\}$ and $j \in \{1, \ldots, N\}$ as

$$T(j, l) = lN + j \quad \text{(A.4)}$$

that maps the integers j and l that label the state values to the absolute index of either of the vectors $\nu_{\text{off}}^{\text{E}}$ and $\nu_{\text{on}}^{\text{E}}$. Now consider the following two sets

$$\mathcal{W}_{\text{off}} \triangleq \left\{ (u, j, l) \in \mathsf{X} \;\middle|\; \phi_{\text{off}}^{\text{E}}(u \mid j, \, l) = \beta_{\text{off}}(u, j, l) \right\} \quad \text{(A.5)}$$

$$\mathcal{W}_{\text{on}} \triangleq \left\{ (u, j, l) \in \mathsf{X} \;\middle|\; \phi_{\text{on}}^{\text{E}}(u \mid j, \, l) = \beta_{\text{on}}(u, j, l) \right\}. \quad \text{(A.6)}$$

The values β_{off} and β_{on} are chosen to ensure the structural requirements in (31)-(32) and (35). For example, for $l = 1$ and $u = \text{on}$ we have that $\beta_{\text{off}}(\text{on}, \cdot, 1) = \phi_{\text{off}}^{\text{TS}}(\text{on}|\cdot)$ (and hence $\beta_{\text{off}}(\text{off}, \cdot, 1) = 1 - \phi_{\text{off}}^{\text{TS}}(\text{on}|\cdot)$) so to enforce the structural requirement in (35). Define for each $k \in \mathsf{T}$, $u, v \in \{\text{on}, \text{off}\}$, $j \in \{1, \ldots, N\}$, and $l \in \{0, \ldots, \tau\}$ the following vectors

$$h_k^u[T(j, l)] \triangleq \begin{cases} (\nu_u^{\text{E}}[\lambda^j, l, k])^{-1} & \text{if } \nu_u^{\text{E}}[\lambda^j, l, k] > 0. \\ 0 & \text{otherwise.} \end{cases}$$

$$w_k^{u,v}[T(j, l)] \triangleq \begin{cases} \beta_v(u, j, l) & \text{if } (u, j, l) \in \mathcal{W}_v \text{ and} \\ & \nu_v^{\text{E}}[\lambda^j, l, k] = 0. \\ 0.5 & \text{if } (u, j, l) \notin \mathcal{W}_v \text{ and} \\ & \nu_v^{\text{E}}[\lambda^j, l, k] = 0. \\ 0 & \text{otherwise.} \end{cases}$$

where $T(\cdot, \cdot)$ is defined in (A.4), $\mathcal{W}_{\text{off}}$ in (A.5), and $\mathcal{W}_{\text{on}}$ in (A.6). Let $W_k^{u,v} = \text{diag}(w_k^{u,v})$ and $H_k^{u,v} = \text{diag}(h_k^{u,v})$, and construct the following matrices

$$H_k = \begin{bmatrix} H_k^{\text{off}} & \mathbf{0} \\ \mathbf{0} & H_k^{\text{on}} \end{bmatrix}, \quad \text{and} \quad \text{(A.7)}$$

$$W_k = \begin{bmatrix} W_k^{\text{off,off}} & W_k^{\text{off,on}} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & W_k^{\text{on,off}} & W_k^{\text{on,on}} \end{bmatrix}. \quad \text{(A.8)}$$

We first show that $\Phi_k^{\text{E}} = H_k J_k + W_k$ satisfies (48). Note that $\text{diag}(\nu_k^{\text{E}})W_k = \mathbf{0}$ since by construction if the i^{th} row of W_k has a non zero entry then the i^{th} diagonal entry of $\text{diag}(\nu_k^{\text{E}})$ is zero. The product $\text{diag}(\nu_k^{\text{E}})H_k$ is a diagonal matrix with with entries of either zero or one. The zero entries also correspond to the zero entries of ν_k^{E}. In this case, the respective entry in J_k is also zero so that $\text{diag}(\nu_k^{\text{E}})H_k J_k = J_k$ as desired.

We now show that $\phi_k^{\mathrm{E}} = H_k J_k + W_k \in \Phi$. First consider an arbitrary state indexed by (off, j, l) at time k, if the corresponding value in $\nu_{\mathrm{off}}[\lambda^j, l, k] > 0$ then the two policy values are defined as

$$\frac{\mathsf{P}\left(m_{k+1} = \mathrm{on}, \; I_k = j, \; \mathsf{L}_k = l, \; m_k = \mathrm{off}\right)}{\nu_{\mathrm{off}}[\lambda^j, l, k]} \quad (\mathrm{A.9})$$

$$\frac{\mathsf{P}\left(m_{k+1} = \mathrm{off}, \; I_k = j, \; \mathsf{L}_k = l, \; m_k = \mathrm{off}\right)}{\nu_{\mathrm{off}}[\lambda^j, l, k]}. \quad (\mathrm{A.10})$$

If either of the above values are fixed in the constraint set Φ, then the constraint (49) will ensure this. Further, since we have that $\nu_k^{\mathrm{E}}, J_k \in [0,1]$ and that $(\nu_k^{\mathrm{E}})^T = J_k \mathbb{1}$ this ensures that the above policy values are within $[0,1]$ and sum to 1. The above argument is valid for any pair of state values such that the corresponding value of ν_k^{E} is non-zero. If the corresponding value of $\nu_{\mathrm{off}}[\lambda^j, l, k] = 0$ and the policy (conditioned on this state value) has a constraint, the first if case in the definition of $w_k^{u,v}$ ensures this constraint. Further, the constraint values must also be chosen to ensure the respective policy values are in $[0,1]$ and sum to one. Lastly, if $\nu_{\mathrm{off}}[\lambda^j, l, k] = 0$ and there is no constraint for the policy conditioned on this state value the second if case in the definition of $w_k^{u,v}$ ensures the policy value sums to 1 and the respective elements are in $[0,1]$. Thus $\Phi_k^{\mathrm{E}} \in \Phi$ for all $k \in \mathsf{T}$.

A.4   Proof of Theorem 1

The proof structure is similar to the one in [2]. The idea is to exploit the fact that: (i) ν_k^{E} is a decision variable for both optimization problems (50) and (44) and (ii) the objective function is the same for both problems and solely a function of the marginal ν_k^{E}. We rewrite these problem compactly below,

$$\eta_{\mathrm{CVX}}^* = \min_{(\nu^{\mathrm{E}}, J) \in X} \eta(\nu^{\mathrm{E}}), \quad (\mathrm{A.11})$$

$$\eta_{\mathrm{NCVX}}^* = \min_{(\nu^{\mathrm{E}}, \Phi^{\mathrm{E}}) \in Y} \eta(\nu^{\mathrm{E}}), \quad (\mathrm{A.12})$$

where the sets X and Y collect all of the relevant constraints for the problems. The variables ν^{E}, Φ^{E}, and J are concatenated over the considered finite time horizon and hence are not sub-scripted by k. We proceed by showing that $\eta_{\mathrm{CVX}}^* \leq \eta_{\mathrm{NCVX}}^*$ and $\eta_{\mathrm{NCVX}}^* \leq \eta_{\mathrm{CVX}}^*$ to give the desired result.

A.4.1 $\eta_{CVX}^* \leq \eta_{NCVX}^*$

Pick any argument minimizer that achieves value η_{NCVX}^* and denote the pair as $(\nu_{\mathrm{NCVX}}^{\mathrm{E}}, \Phi_{\mathrm{NCVX}}^{\mathrm{E}})$. Trivially construct J through the relation (48) so that this constructed J and $\nu_{\mathrm{NCVX}}^{\mathrm{E}}$ (that is optimal for (44)) are also feasible for (50), i.e., $(\nu_{\mathrm{NCVX}}^{\mathrm{E}}, J) \in X$. This is since

$\mathbb{1}^T \mathrm{diag}(\nu_k^{\mathrm{E}}) = \nu_k^{\mathrm{E}}$ and $\mathbb{1} = \Phi_k^{\mathrm{E}} \mathbb{1}$. Hence we have that

$$\eta_{\mathrm{CVX}}^* = \min_{(\nu^{\mathrm{E}}, J) \in X} \eta(\nu^{\mathrm{E}}) \leq \eta(\nu_{\mathrm{NCVX}}^{\mathrm{E}}) = \eta_{\mathrm{NCVX}}^* \quad (\mathrm{A.13})$$

since by definition η_{CVX}^* is the minimum value over the set of feasible solutions.

A.4.2 $\eta_{NCVX}^* \leq \eta_{CVX}^*$

We take a pair $(\nu_{\mathrm{CVX}}^{\mathrm{E}}, J_{\mathrm{CVX}})$ that achieve optimal cost η_{CVX}^* and construct a feasible solution for (44), denoted $(\eta_{\mathrm{NCVX}}^{\mathrm{E}}, \Phi_{\mathrm{NCVX}}^{\mathrm{E}})$, as follows (for each k)

$$\Phi_{k,\mathrm{NCVX}}^{\mathrm{E}} = H_k J_k + W_k, \quad \text{and} \quad (\mathrm{A.14})$$

$$\nu_{k,\mathrm{NCVX}}^{\mathrm{E}} = \nu_{k,\mathrm{CVX}}^{\mathrm{E}}. \quad (\mathrm{A.15})$$

Where H_k and W_k are defined in Lemma 4. This constructed solution is then feasible for (44) as the constraint $\Phi_{\mathrm{NCVX}}^{\mathrm{E}} \in \Phi$ is part of the result in Lemma 4 and

$$\nu_{k,\mathrm{NCVX}}^{\mathrm{E}} \Phi_{k,\mathrm{NCVX}}^{\mathrm{E}} G_k^{\mathrm{E}} = \nu_{k,\mathrm{NCVX}}^{\mathrm{E}} \left(H_k J_k + W_k\right) G_k^{\mathrm{E}} \quad (\mathrm{A.16})$$

$$= \mathbb{1}^T J_k G_k^{\mathrm{E}} = \nu_{k+1,\mathrm{NCVX}}^{\mathrm{E}}. \quad (\mathrm{A.17})$$

The fact that $\nu_{k,\mathrm{NCVX}}^{\mathrm{E}} \left(H_k J_k + W_k\right) = \mathbb{1}^T J_k$ is since $\nu_{k,\mathrm{NCVX}}^{\mathrm{E}} W_k = \mathbf{0}$ and $\nu_{k,\mathrm{NCVX}}^{\mathrm{E}} H_k J_k = \mathbb{1}^T J_k$. The matrix W_k only has non zero entries for row indices where the index of the row vector $\nu_{k,\mathrm{NCVX}}^{\mathrm{E}}$ is zero so that the resulting product is the zero vector. The product $\nu_{k,\mathrm{NCVX}}^{\mathrm{E}} H_k$ is a vector of ones and zeros, specifically, if the i^{th} element of this vector is zero then the entire i^{th} column of the matrix J_k will be the zero vector. Thus the equivalence between $\nu_{k,\mathrm{NCVX}}^{\mathrm{E}} \left(H_k J_k + W_k\right)$ and $\mathbb{1}^T J_k$. Since the constructed solution is feasible we have that

$$\eta_{\mathrm{NCVX}}^* = \min_{(\nu^{\mathrm{E}}, \Phi^{\mathrm{E}}) \in Y} \eta(\nu^{\mathrm{E}}) \leq \eta(\nu_{\mathrm{CVX}}^{\mathrm{E}}) = \eta_{\mathrm{CVX}}^* \quad (\mathrm{A.18})$$

since by definition η_{NCVX}^* is the minimum value over the set of feasible solutions.

B   PDE discretization

We denote the i^{th} CV as $\mathrm{CV}(i)$ and further adopt the following notational simplifications,

$$\mu_{\mathrm{off}}(\lambda^i, t) \triangleq \mu_{\mathrm{off}}(\lambda_{\mathrm{off}}^i, t), \quad \text{and} \quad \mu_{\mathrm{on}}(\lambda^i, t) \triangleq \mu_{\mathrm{on}}(\lambda_{\mathrm{on}}^i, t).$$

Highlighted red in Figure 2 are the two control volumes to assist in enforcing boundary conditions that coincide with the thermostat policy (6). This is discussed further in Appendix B.2 when the boundary conditions CVs are discretized.

B.1 Internal CV's

Consider the RHS of the pde (9) integrated over CV(i):

$$\int_{\text{CV(i)}} \left(\frac{\sigma^2}{2} \frac{\partial^2}{\partial \lambda^2} \big(\mu_{\text{on}}(\lambda, t) \big) - \frac{\partial}{\partial \lambda} \big(f_{\text{on}}(\lambda, t) \mu_{\text{on}}(\lambda, t) \big) \right) d\lambda$$

$$= \left(\frac{\sigma^2}{2} \frac{\partial}{\partial \lambda} \mu_{\text{on}}(\lambda, t) - f_{\text{on}}(\lambda, t) \mu_{\text{on}}(\lambda, t) \right) \Big|_{\lambda^{i,-}}^{\lambda^{i,+}}, \quad \text{(B.1)}$$

where equality is by the divergence theorem [27]. Note, the points $\lambda^{i,-}$ and $\lambda^{i,+}$ are not control volume variables, but rather the boundaries of a single control volume. Hence, quantities in (B.1) need to be approximated in terms of the nodal points of the neighboring control volumes. The approximations for the partial derivative are,

$$\frac{\partial}{\partial \lambda} \mu_{\text{on}}(\lambda^{i,+}, t) \approx \frac{\mu_{\text{on}}(\lambda^{i+1}, t) - \mu_{\text{on}}(\lambda^i, t)}{\Delta \lambda}, \quad \text{and} \tag{B.2}$$

$$\frac{\partial}{\partial \lambda} \mu_{\text{on}}(\lambda^{i,-}, t) \approx \frac{\mu_{\text{on}}(\lambda^i, t) - \mu_{\text{on}}(\lambda^{i-1}, t)}{\Delta \lambda}. \tag{B.3}$$

For the integrated convective term, we use the so-called upwind scheme [27]. This scheme elects the FVM equivalent of a forward or backward difference based on the sign of the convective velocity $f_{\text{on}}(\lambda, t)$. By assumption **A.2**, $f_{\text{on}}(\lambda, t) \leq 0$ and the upwind scheme prescribes:

$$f_{\text{on}}(\lambda^{i,-}, t) \mu_{\text{on}}(\lambda^{i,-}, t) = f_{\text{on}}(\lambda^{i,-}, t) \mu_{\text{on}}(\lambda^i, t), \quad \text{and}$$
$$f_{\text{on}}(\lambda^{i,+}, t) \mu_{\text{on}}(\lambda^{i,+}, t) = f_{\text{on}}(\lambda^{i,+}, t) \mu_{\text{on}}(\lambda^{i+1}, t). \tag{B.4}$$

When the TCL is off $f_{\text{off}}(\lambda, t) \geq 0$ (also by Assumption **A.2**) the upwind scheme prescribes:

$$f_{\text{off}}(\lambda^{i,-}, t) \mu_{\text{off}}(\lambda^{i,-}, t) = f_{\text{off}}(\lambda^{i,-}, t) \mu_{\text{off}}(\lambda^{i-1}, t), \quad \text{and}$$
$$f_{\text{off}}(\lambda^{i,+}, t) \mu_{\text{off}}(\lambda^{i,+}, t) = f_{\text{off}}(\lambda^{i,+}, t) \mu_{\text{off}}(\lambda^i, t). \tag{B.5}$$

Now returning to the discretization of the PDE (9) over an arbitrary internal CV. We approximate the LHS of (9) integrated over the control volume as,

$$\int_{\text{CV(i)}} \frac{\partial}{\partial t} \mu_{\text{on}}(\lambda, t) d\lambda \approx \frac{d}{dt} \mu_{\text{on}}(\lambda^i, t) \Delta \lambda = \frac{d}{dt} \nu_{\text{on}}(\lambda^i, t),$$

where we have defined

$$\nu_{\text{on}}(\lambda^i, t) \triangleq \mu_{\text{on}}(\lambda^i, t) \Delta \lambda. \tag{B.6}$$

Now, denote the following

$$D \triangleq \frac{\sigma^2}{(\Delta \lambda)^2}, \quad \text{and} \quad F_{\text{on}}^i(t) \triangleq \frac{f_{\text{on}}(\lambda^i, t)}{\Delta \lambda}, \tag{B.7}$$

where the quantities $F_{\text{off}}^i(t)$, $F_{\text{on}}^{i,+}(t)/F_{\text{off}}^{i,+}(t)$, and $F_{\text{on}}^{i,-}(t)/F_{\text{off}}^{i,-}(t)$ are defined similarly to $F_{\text{on}}^i(t)$, e.g., $F_{\text{off}}^{i,+}(t) \triangleq f_{\text{off}}(\lambda^{i,+}, t)/\Delta \lambda$. Now equating the approximation of the RHS (9) with the approximation of the LHS of (9) we have,

$$\frac{d}{dt} \nu_{\text{on}}(\lambda^i, t) = \left(F_{\text{on}}^{i,-}(t) - D \right) \nu_{\text{on}}(\lambda^i, t) + \frac{D}{2} \nu_{\text{on}}(\lambda^{i-1}, t)$$
$$+ \left(\frac{D}{2} - F_{\text{on}}^{i,+}(t) \right) \nu_{\text{on}}(\lambda^{i+1}, t). \tag{B.8}$$

The spatial discretization for the PDE (10) is similar and yields,

$$\frac{d}{dt} \nu_{\text{off}}(\lambda^i, t) = \frac{D}{2} \nu_{\text{off}}(\lambda^{i+1}, t) - \left(F_{\text{off}}^{i,+}(t) + D \right) \nu_{\text{off}}(\lambda^i, t)$$
$$+ \left(\frac{D}{2} + F_{\text{off}}^{i,-}(t) \right) \nu_{\text{off}}(\lambda^{i-1}, t), \tag{B.9}$$

where $\nu_{\text{off}}(\lambda^i, t) \triangleq \mu_{\text{off}}(\lambda^i, t) \Delta \lambda$.

B.2 Boundary CV's

The boundary CVs are the CVs associated with the nodal values: λ_{on}^1, λ_{on}^q, λ_{on}^N, λ_{off}^1, λ_{off}^m, and λ_{off}^N. The superscript, for example the integer q in λ_{on}^q represents the CV index. All boundary CVs can be seen in Figure 2. Discretization of the boundary CVs requires care for atleast two reasons. First, this is typically where one introduces the BCs of the PDE into the numerical approximation. Secondly, on finite domains the endpoints present challenges, for example, there is no variable $\mu_{\text{on}}(\lambda^{N+1}, t)$ for computation of the derivative values for node λ_{on}^N.

The BC's for the coupled PDEs (9)-(10) are [21]:

Absorbing Boundaries:
$$\mu_{\text{on}}(\lambda^{\min}, t) = \mu_{\text{off}}(\lambda^{\max}, t) = 0. \tag{B.10}$$
Conditions at Infinity:
$$\mu_{\text{on}}(+\infty, t) = \mu_{\text{off}}(-\infty, t) = 0. \tag{B.11}$$
Conservation of Probability:
$$\frac{\partial}{\partial \lambda} \left[\mu_{\text{on}}(\lambda^{q,-}, t) - \mu_{\text{on}}(\lambda^{q-1,+}, t) - \mu_{\text{off}}(\lambda^{N-1,+}, t) \right] = 0. \tag{B.12}$$
$$\frac{\partial}{\partial \lambda} \left[\mu_{\text{off}}(\lambda^{m,+}, t) - \mu_{\text{on}}(\lambda^{2,-}, t) - \mu_{\text{off}}(\lambda^{m+1,-}, t) \right] = 0. \tag{B.13}$$

Continuity:
$$\mu_{\text{on}}(\lambda^{q,-}, t) = \mu_{\text{on}}(\lambda^{q-1,+}, t). \tag{B.14}$$
$$\mu_{\text{off}}(\lambda^{m,+}, t) = \mu_{\text{off}}(\lambda^{m+1,-}, t). \tag{B.15}$$

As we will see, implementation of some of the above conditions will require a bit of care. However, some are quite

trivial to enforce. For example, by default, the continuity conditions (B.14) and (B.15) are satisfied due to our choice of CV structure, since, for example, for any i we have $\lambda_{\text{off}}^{i,-} = \lambda_{\text{off}}^{i-1,+}$ and $\lambda_{\text{off}}^{i,+} = \lambda_{\text{off}}^{i+1,-}$.

Now focusing on the conditions at infinity BC (B.11), we enforce instead the following conditions:

$$\frac{\partial}{\partial \lambda} \mu_{\text{off}}(\lambda^{1,-}, t) = 0, \quad \text{and} \quad \frac{\partial}{\partial \lambda} \mu_{\text{on}}(\lambda^{N,+}, t) = 0.$$
(B.16)

Our computational domain cannot extend to infinity, where the BC (B.11) is required to hold, but the temperature values λ_{off}^1 and λ_{on}^N are quite far away from the deadband and so the density here will be near zero.

Now, consider the spatial discretization of the CVs associated with the BC at infinity. First considering the CV associated with the temperature λ_{off}^1, we have that the differential equation is

$$\frac{d}{dt} \nu_{\text{off}}(\lambda^1, t) = \left(-F_{\text{off}}^{1,+}(t) - \frac{D}{2} \right) \nu_{\text{off}}(\lambda^1, t) \quad \text{(B.17)}$$
$$+ \left(\frac{D}{2} + F_{\text{off}}^{2,-}(t) \right) \nu_{\text{off}}(\lambda^2, t).$$

Considering the CV associated with the temperature λ_{on}^N, we have

$$\frac{d}{dt} \nu_{\text{on}}(\lambda^N, t) = \left(F_{\text{on}}^{N,+}(t) - \frac{D}{2} \right) \nu_{\text{on}}(\lambda^N, t) \quad \text{(B.18)}$$
$$+ \left(\frac{D}{2} - F_{\text{on}}^{N,+}(t) \right) \nu_{\text{on}}(\lambda^{N-1}, t).$$

In the above we make the assumption that $\nu_{\text{off}}(\lambda^{1,-} - \Delta\lambda, t) = 0$ and $\nu_{\text{on}}(\lambda^{N,+} + \Delta\lambda, t) = 0$.

Now focus on the absorbing boundary (B.10) and conservation of probability (B.12)-(B.13) boundary conditions. These BCs have the following meaning. The condition (B.10) clamps the density at the end of the deadband to zero. BC (B.12) reads: the net-flux across the temperature value λ_{on}^q is equal to the flux of density going from off to on. In order to enforce both (B.12) and (B.13) we will model the flux of density due to the thermostat control policy as a source/sink. Before doing this, we mention some issues with enforcing the BC (B.10).

A TCL's temperature trajectory will not satisfy the BC (B.10) since to switch its mode the TCL's temperature sensor will have to register a value outside the deadband. Therfore, we introduce two additional CV's associated with the temperatures λ_{on}^1 and λ_{off}^N, which are shown in red in Figure 2. We then transfer the BC (B.10) to one on the added CVs, which becomes:

$$\mu_{\text{on}}(\lambda^{1,-}, t) = \mu_{\text{off}}(\lambda^{N,+}, t) = 0. \quad \text{(B.19)}$$

As mentioned, to enforce the conservation of probability BC we use a source/sink type argument, which we also enforce on the added CVs. To see what we mean by source/sink argument, consider the following: some rate of density is transferred out of the CV λ_{off}^N and into the CV λ_{on}^q (as depicted in Figure 2) due to thermostatic control. We model the sink as simply $-\nu_{\text{off}}(\lambda^N, t)$. The rate of the sink is then given as $-\gamma \nu_{\text{off}}(\lambda^N, t)$, where $\gamma > 0$ is a modeling choice and a constant of appropriate units that describes the discharge rate. We shortly give insight on how to select a value for γ. Now discretizing the CV corresponding to the nodal value λ_{off}^N subject to the BC (B.19) and the sink $-\nu_{\text{off}}(\lambda^N, t)$ we obtain,

$$\frac{d}{dt} \nu_{\text{off}}(\lambda^N, t) = \left(\frac{D}{2} + F_{\text{off}}^{N,-}(t) \right) \nu_{\text{off}}(\lambda^{N-1}, t) \quad \text{(B.20)}$$
$$- \alpha \nu_{\text{off}}(\lambda^N, t),$$

where $\alpha \triangleq (\gamma + D)$. In obtaining the above, we have made the reasonable assumption that $\nu_{\text{off}}(\lambda^{N,+} + \Delta\lambda, t) = 0$. The quantity $\alpha \nu_{\text{off}}(\lambda^N, t)$ represents the rate of change of density from the CV λ_{off}^N to the CV λ_{on}^q, as depicted in Figure 2. Consequently, to conserve probability, we must add this quantity as a source to the ode for the CV λ_{on}^q, i.e.,

$$\frac{d}{dt} \nu_{\text{on}}(\lambda^q, t) = \cdots + \alpha \nu_{\text{off}}(\lambda^N, t). \quad \text{(B.21)}$$

The dots in equation (B.21) represent the portion of the dynamics for the standard internal CV (i.e., the RHS of (B.8)) for the temperature node λ_{on}^q. A similar argument is used for the BC (B.13) with the CV's λ_{on}^1 and λ_{off}^m, and the corresponding differential equations are,

$$\frac{d}{dt} \nu_{\text{on}}(\lambda^1, t) = \left(\frac{D}{2} - F_{\text{on}}^{1,+}(t) \right) \nu_{\text{on}}(\lambda^2, t) - \alpha \nu_{\text{on}}(\lambda^1, t),$$
(B.22)

$$\frac{d}{dt} \nu_{\text{off}}(\lambda^m, t) = \cdots + \alpha \nu_{\text{on}}(\lambda^1, t). \quad \text{(B.23)}$$

To better understand the role of γ consider the following example. Electing γ in the above so that $\alpha = (\Delta t)^{-1}$, where Δt is a time increment, has the following interpretation: all mass starting in state $\nu_{\text{off}}(\lambda^N, \cdot)$ at time t is transferred out by time $t + \Delta t$ into the state $\nu_{\text{on}}(\lambda^q, \cdot)$.

B.2.1 Additional conditions

Two additional conditions are enforced, namely that once mass is transferred to the nodes λ_{off}^N or λ_{on}^1 it cannot "travel backwards." For example, mass is transferred from λ_{off}^N entirely to the corresponding on temperature bin and no mass is transferred backwards to $\lambda_{\text{off}}^{N-1}$. This corresponds to setting: (i) the coefficient on $\nu_{\text{off}}(\lambda^N, t)$

in the ode for $\nu_{\text{off}}(\lambda^{N-1}, t)$ to zero and (ii) the coefficient on $\nu_{\text{on}}(\lambda^1, t)$ in the ode for $\nu_{\text{on}}(\lambda^2, t)$ to zero.

B.3 Overall system

Now, combining the odes–(B.8) and (B.9) for all of the internal CVs and (B.17), (B.18), (B.20), (B.21), (B.22), (B.23) for the BC CVs–we obtain the linear time varying system,

$$\frac{d}{dt}\nu(t) = \nu(t)A(t). \qquad \text{(B.24)}$$