Aggregation and Control of Populations of Thermostatically Controlled Loads by Formal Abstractions

Sadegh Esmaeil Zadeh Soudjani and Alessandro Abate

***Abstract*—This work discusses a two-step procedure, based on the use of formal abstractions, to generate a finite-space stochastic dynamical model as an aggregation of the continuous temperature dynamics of a homogeneous population of Thermostatically Controlled Loads (TCLs). The temperature of a TCL is described by a stochastic difference equation and the TCL status (ON, OFF) by a deterministic switching mechanism. The procedure is deemed to be formal, as it allows the quantification of the error introduced by the abstraction. As such, it builds and improves on a known, earlier approximation technique used in the literature. Further, the contribution discusses the extension to the instance of heterogeneous populations of TCLs by means of two approaches. It moreover investigates the problem of global (population-level) power reference tracking and load balancing for TCLs that are explicitly dependent on a control input. The procedure is tested on a case study and benchmarked against the mentioned existing approach in the literature.**

I. Introduction

Thermostatically Controlled Loads (TCLs) have shown potential to be engaged in power system services such as load shifting, peak shaving, and demand response programs. The regulation of the total power consumption of large populations of TCLs, with the goal of smoothing the uncertain demand over the grid or of tracking the uncertain power production, while abiding by strict requirements on users comfort, can lead to economically relevant repercussions for an energy provider. The modeling of TCLs in view of their application to load control has a rich history, which can be traced back to the work [12], where the model of a TCL is used to describe the evolution of the thermostat state. A diffusion approximation framework is introduced in [23] and a discrete-time simulation model is developed in [27]. Building on these foundations, recent studies have focused on the development of practically usable models for aggregated populations of TCLs. In particular, [10] provides an approximate analytical solution to the coupled Fokker–Planck equations originally developed in [23] for a population of homogeneous TCLs (meaning that TCLs are assumed to have the same dynamics and parameters), and puts forward a Linear Time-Invariant (LTI) population model, where the coefficients of its transfer function are estimated by means of system identification techniques. The contribution in [5] develops a bilinear PDE model and designs a Lyapunov stable controller. The work in [22] proposes a new technique, based on the partitioning of the TCL temperature range, to obtain an aggregate state-space model for a population of TCLs that is now heterogeneous over its thermal capacitances. The full information of the state variables of the model is used to synthesize a control strategy over the model output (namely, the total power consumption), in order to attain tracking via a (deterministic) Model Predictive Control (MPC) scheme. The contributions in [24], [26] extend the results in [22] by considering a population of TCLs that are heterogeneous over all their parameters: this general setup requires the use of the extended Kalman filter to estimate the states of the model and to identify its characteristic transition matrix. The control of the population is performed by switching ON/OFF a portion of the TCLs in the population. Additional recent contributions have targeted the application of the approaches in [22], [24], [26], towards higher-order dynamics [29], [30] and energy arbitrage [25].

Matrices and parameters of the state-space aggregate model can be computed analytically or via system identification techniques [5], [10], [22], [24], [26]. The only available analytical derivation of the state-space aggregate model [22] works in discrete time and is based, under two rather restrictive assumptions, on the underlying model of the single TCL: first, the TCL temperature evolution is assumed to be deterministic, thus leading to a deterministic state-space model; second, after partitioning the temperature range in separate intervals, the temperatures of the TCLs within each interval are assumed to be uniformly distributed. Moreover, from a practical standpoint there seems to be no clear connection between the precision of the aggregation and the performance of the aggregated model: more specifically, an increase in the number of state bins (that is, a decrease in the width of the introduced temperature intervals) does not necessarily improve the performance of the aggregated model. The approach based on system identification [10], [24], [26] on the other hand estimates the parameters of an LTI aggregate model from data.

This article proposes a new, formal two-step abstraction procedure to generate a finite stochastic dynamical model as the aggregation of the dynamics of a population of TCLs. The approach relaxes the limiting assumptions employed in [22] by providing a model based on the native probabilistic evolution of the single TCL temperature. The abstraction comprises two separate parts: (1) translating a continuous-space model of a TCL to a finite state-space model, which obtains a Markov chain; and over a population of TCLs (2) taking the cross product of the single Markov chains and lumping the obtained model, by finding its coarsest probabilistically bisimilar Markov chain [4]. The approach is developed for the case of a homogeneous population of TCLs, and extended to a heterogeneous population – in the latter case the aggregation (second step) employs an *approximate* probabilistic bisimulation relation, which as we shall see introduces an error. In both the homogeneous and the heterogeneous case, it is possible to quantify the abstraction error of the first step, and furthermore in the homogeneous instance the error of the overall abstraction procedure can be quantified – this is unlike the approach based on approximations in [22] and the approach based on system identification in [10], [24], [26].

This article also describes a dynamical model for the time evolution of the abstraction and provides asymptotic results as the population size grows. Moreover, it shows that increasing the number of state bins always improves the accuracy, leading

A. Abate is with the Department of Computer Science, University of Oxford, United Kingdom. Both authors are affiliated with the Delft Center for Systems and Control, TU Delft – Delft University of Technology, The Netherlands. {S.EsmaeilZadehSoudjani,A.Abate}@tudelft.nl, aabate@cs.ox.ac.uk The work has been supported by the European Commission STREP project MoVeS 257005, by the European Commission IAPP project AMBI 324432, and by the NWO VENI grant 016.103.020. This article generalizes and completes the results presented in [17], [21].

to a convergence of the introduced abstraction error to zero. This result is aligned with the work in [5], [6] on the aggregation of continuous-time deterministic thermostatic loads. The explicit relationship between aggregate model and population parameters enables the development of a set-point control strategy aimed at reference tracking over the population total power consumption (cf. Figure 2): a conditional Kalman filter [11] is employed to estimate the state of the model, which is then used to regulate the power consumption of the population via a simple one-step prediction approach. As such, the control architecture does not require knowledge of the single TCL states, but directly leverages the measurement of the total power consumption. Alternatively, a stochastic model predictive control scheme is proposed. Both procedures are tested on a case study and the abstraction technique is benchmarked against the analytical approach proposed in [22].

The article is organized as follows. Section II, after introducing the model of the TCL dynamics, describes its abstraction as a Markov Chain, and further discusses the aggregation of a homogeneous population of TCLs – the errors introduced by both steps are quantified. Section III focuses on heterogeneous populations of TCLs and elucidates two techniques to aggregate their dynamics: one based on averaging, and a second based on clustering the uncertain parameters. The latter approach allows for a general quantification of the error. Section IV discusses TCL models endowed with a control input, and the synthesis of global (acting at the population level – cf. Figure 2) controllers to achieve regulation of the total consumed power – this is achieved by two alternative schemes. Finally, all the discussed techniques are tested on a case study described in Section V.

II. Formal Abstraction of a Homogeneous Population of TCLs

A. Continuous Model of the Temperature of a TCL

Throughout this article we use the notation $\mathbb{N}$ to denote the set of positive integers, $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$, $\mathbb{N}_n = \{1, 2, 3, \dots, n\}$, and $\mathbb{Z}_n = \mathbb{N}_n \cup \{0\}$. We denote vectors with bold typeset and use corresponding indexed letters with normal typeset for their elements.

The evolution of the temperature in a TCL can be characterized by the following stochastic difference equation [10], [23]:

$$\theta(t+1) = a\,\theta(t) + (1-a)(\theta_a \pm m(t) R P_{rate}) + w(t), \quad (1)$$

where θ_a is the ambient temperature, P_{rate} is the energy transfer rate of the TCL, C and R indicate the thermal capacitance and resistance respectively, and $a = e^{-h/RC}$ with a discretization step h. The process noise $w(t), t \in \mathbb{N}_0$, is made up by i.i.d. random variables characterized by a density function $t_{\mathfrak{w}}(\cdot)$. We denote with $m(t) = 0$ a TCL in the OFF mode at time t, and with $m(t) = 1$ a TCL in the ON mode. In equation (1) the symbol $\pm$ signifies the following: a plus sign is used for a heating TCL, whereas a minus sign for a cooling TCL. In this work we focus on a population of cooling TCLs, with the understanding that the case of heating TCLs can be similarly obtained. The distribution of the initial temperature and mode is denoted by $\pi_0(m, \theta)$. The temperature of the cooling TCL is regulated by a control signal $m(t+1) = f(m(t), \theta(t))$ based on discrete switching as

$$f(m,\theta) = \begin{cases} 0, & \theta < \theta_s - \delta/2 \doteq \theta_- \\ 1, & \theta > \theta_s + \delta/2 \doteq \theta_+ \\ m, & \text{else,} \end{cases} \quad (2)$$

where θ_s denotes a temperature set-point and δ a dead-band, and together characterize an operating temperature range. The power consumption of the TCL at time t is equal to $\frac{1}{\eta} m(t) P_{rate}$, which is equal to zero in the OFF mode and positive in the ON mode, and where the parameter η is the coefficient of performance. The constant $\frac{1}{\eta} P_{rate}$, namely the power consumed by the TCL when it is in the ON mode, will be shortened as P_{ON} in the sequel.

B. Finite Abstraction of a TCL by State-Space Partitioning

The composition of the dynamical equation in (1) with the algebraic relation in (2) allows one to consider a TCL as a Stochastic Hybrid System [2], namely as a discrete-time Markov process evolving over a hybrid (that is, discrete/continuous) state space. The hybrid state space is characterized by a variable $s = (m, \theta) \in \mathbb{Z}_1 \times \mathbb{R}$ with two components, a discrete (m) and a continuous (θ) one. The one-step transition density function of the stochastic process, $t_{\mathfrak{s}}(\cdot|s)$, made up of the dynamical equations in (1), (2), and conditional on point s, can be computed as

$$\begin{aligned} t_{\mathfrak{s}}\left((m', \theta')|(m, \theta)\right) =& \delta[m' - f(m,\theta)] \times \\ & t_{\mathfrak{w}}(\theta' - a\,\theta - (1-a)(\theta_a - m R P_{rate})), \end{aligned}$$

where $\delta[\cdot]$ denotes the discrete unit impulse function. This interpretation allows leveraging an abstraction technique, proposed in [1] and extended in [15], [16], [18], aimed at reducing a discrete-time, uncountable state-space Markov process into a (discrete-time) finite-state Markov chain. This abstraction is based on a state-space partitioning procedure as follows. Consider an arbitrary, finite partition of the continuous domain $\mathbb{R} = \cup_{i=1}^{n} \Theta_i$, and arbitrary representative points within the partitioning regions denoted by $\{\bar{\theta}_i \in \Theta_i, i \in \mathbb{N}_n\}$. Introduce a finite-state Markov chain $\mathcal{M}$, characterized by $2n$ states $s_{im} = (m, \bar{\theta}_i), m \in \mathbb{Z}_1, i \in \mathbb{N}_n$. The transition probability matrix related to $\mathcal{M}$ is made up of the following elements

$$\mathbb{P}(s_{im}, s_{i'm'}) = \int_{\Theta_{i'}} t_{\mathfrak{s}}\left((m', \theta')|m, \bar{\theta}_i\right) d\theta', \quad (3)$$

for all $m, m' \in \mathbb{Z}_1$, $i, i' \in \mathbb{N}_n$. The initial probability mass for $\mathcal{M}$ is obtained as $p_0(s_{im}) = \int_{\Theta_i} \pi_0(m, \theta) d\theta$. For simplicity of notation we rename the states of $\mathcal{M}$ by the bijective map $\ell(s_{im}) = mn + i, m \in \mathbb{Z}_1, i \in \mathbb{N}_n$, and accordingly we introduce the new notation

$$P_{ij} = \mathbb{P}(\ell^{-1}(i), \ell^{-1}(j)), \quad p_{0i} = p_0(\ell^{-1}(i)), \quad \forall i, j \in \mathbb{N}_{2n}.$$

Notice that the conditional density function of the stochastic system capturing the dynamics of a TCL is discontinuous, due to the presence of equation (2). This can be emphasized by the following alternative representation of the discontinuity in the discrete conditional distribution, for all $m, m' \in \mathbb{Z}_1, \theta \in \mathbb{R}$:

$$\begin{aligned} \delta\left[m' - f(m,\theta)\right] =\ & m' \mathbb{I}_{(\theta_+, \infty)}(\theta) + (1-m') \mathbb{I}_{(-\infty, \theta_-)}(\theta) \\ & + (1 - |m - m'|) \mathbb{I}_{[\theta_-, \theta_+]}(\theta), \end{aligned}$$

where $\mathbb{I}_{\mathcal{A}}(\cdot)$ denotes the indicator function of a general set $\mathcal{A}$. The selection of the partitioning sets then requires special attention: a convenient way to obtain that is to select a partition for the dead-band $[\theta_-, \theta_+]$, thereafter extending it to a partition covering the whole real line $\mathbb{R}$ (cf. Figure 1). Let us select two constants $\mathsf{l}, \mathsf{m} \in \mathbb{N}$, $\mathsf{l} < \mathsf{m}$, compute the partition size $\upsilon = \delta/2\mathsf{l}$

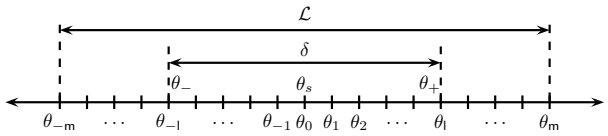


Fig. 1. Partitioning of the temperature axis for the abstraction of the dynamics of a TCL.

and quantity $\mathcal{L} = 2\mathsf{m}\upsilon$. Now construct the boundary points of the partition sets $\{\theta_i\}_{i=-\mathsf{m}}^{i=\mathsf{m}}$ for the temperature axis as follows:

$$\begin{array}{l}\theta_{\pm\mathsf{l}} = \theta_s \pm \delta/2, \quad \theta_{\pm\mathsf{m}} = \theta_s \pm \mathcal{L}/2, \quad \theta_{i+1} = \theta_i + \upsilon, \\ \mathbb{R} = \cup_{i=1}^{n}\Theta_i, \quad \Theta_1 = (-\infty, \theta_{-\mathsf{m}}), \quad \Theta_n = [\theta_{\mathsf{m}}, \infty), \\ \Theta_{i+1} = [\theta_{-\mathsf{m}+i-1}, \theta_{-\mathsf{m}+i}), \quad i \in \mathbb{N}_{n-2}, \quad n = 2\mathsf{m}+2,\end{array} \tag{4}$$

and let us render the Markov states of the infinite-length intervals Θ_1, Θ_n reflecting. Let us emphasize that the discontinuity in the discrete transition kernel $\delta\left[m' - f(m,\theta)\right]$ and the above partition induce the following structure on the transition probability matrix of the chain $\mathcal{M}$:

$$P = \left[\begin{array}{cccc} Q_{11} & 0 & Q_{31} & 0 \\ 0 & Q_{22} & 0 & Q_{42} \end{array}\right]^T, \tag{5}$$

where $Q_{11}, Q_{42} \in \mathbb{R}^{n\times(\mathsf{m}+\mathsf{l}+1)}$, whereas $Q_{22}, Q_{31} \in \mathbb{R}^{n\times(\mathsf{m}-\mathsf{l}+1)}$, which leads to $P \in \mathbb{R}^{2n\times 2n}$.

Clearly, the abstraction of the dynamics in (1)-(2) over this partition of the state space leads to a discretization error: in Section II-E we formally derive bounds on this error as a function of the partition size υ and of the quantity $\mathcal{L}$. This guarantees the convergence (in expected value) of the power consumption of the abstracted model to that of the entire population [1], [15], [16].

C. Aggregation of a Population of TCLs by Bisimulation Relation

Consider now a population of n_p homogeneous TCLs, that is a population of TCLs which, after possible rescaling of (1)-(2), share the same set of parameters $\theta_s, \delta, \theta_a, C, R, P_{rate}, P_{ON}$ (and thus η), h, and noise terms $t_{\mathfrak{w}}(\cdot)$. Each TCL can then be abstracted as a Markov chain $\mathcal{M}$ with the same transition probability matrix $P = [P_{ij}]_{i,j}$, where $i, j \in \mathbb{N}_{2n}$, which leads to a population of n_p homogeneous Markov chains. The initial probability mass vector $p_0 = [p_{0i}]_i$ might still vary over the TCL population.

The homogeneous population of TCLs can be represented by a single Markov chain Ξ, built as the cross product of the n_p homogeneous Markov chains, and endowed with the state

$$\mathbf{z} = [z_1, z_2, \ldots, z_{n_p}]^T \in \mathcal{Z} = \mathbb{N}_{2n}^{n_p},$$

where $z_j \in \mathbb{N}_{2n}$ represents the state of the j^{th} Markov chain. We denote by P_Ξ the transition probability matrix of Ξ.

It is understood that Ξ, having exactly $(2n)^{n_p}$ states, can in general be quite large, and thus cumbersome to manipulate computationally. As the second step of the abstraction procedure, we are interested in aggregating this model and we employ the notion of probabilistic bisimulation to achieve this [4]. Let us introduce a finite set of atomic propositions[1] as a constrained vector with a dimension corresponding to the number of bins of the TCL $\mathcal{M}$:

$$AP = \left\{ \mathbf{x} = [x_1, x_2, \ldots, x_{2n}]^T \in \mathbb{Z}_{n_p}^{2n} \middle| \sum_{r=1}^{2n} x_r = n_p \right\}.$$

The labeling function $L : \mathcal{Z} \to AP$ associates to a configuration $\mathbf{z}$ of Ξ a vector $\mathbf{x} = L(\mathbf{z})$, the elements $x_i \in \mathbb{Z}_{n_p}$ of which count the number of thermostats in bin $i, i \in \mathbb{N}_{2n}$. Notice that the set AP is finite with cardinality $|AP| = (n_p + 2n - 1)!/(n_p!(2n-1)!)$, which for $n_p \geq 2$ is (much) less than the cardinality $(2n)^{n_p}$ of Ξ.

Let us define an equivalence relation $\mathcal{R}$ [4] on the state space of $\mathcal{Z}$, such that

$$(\mathbf{z}, \mathbf{z}') \in \mathcal{R} \Leftrightarrow L(\mathbf{z}) = L(\mathbf{z}').$$

A pair of elements of $\mathcal{Z}$ (each of them a vector representing a state of Ξ) is in the relation whenever the corresponding number of TCLs in any of the introduced bins is the same (recall that the TCL population is assumed to be homogeneous). Such an equivalence relation provides a partition of the state space of $\mathcal{Z}$ into equivalence classes belonging to the quotient set $\mathcal{Z}/\mathcal{R}$, where each class is uniquely specified by the label associated to its elements. We plan to show that $\mathcal{R}$ is an exact probabilistic bisimulation relation on Ξ [4], which requires proving that, for any set $\mathscr{T} \in \mathcal{Z}/\mathcal{R}$ and any pair $(\mathbf{z}, \mathbf{z}') \in \mathcal{R}$

$$\mathbb{P}_\Xi(\mathbf{z}, \mathscr{T}) = \mathbb{P}_\Xi(\mathbf{z}', \mathscr{T}), \tag{6}$$

This is achieved by Corollary 1 in the next Section. We now focus on the stochastic properties of Ξ, which we study via its quotient Markov chain obtained with $\mathcal{R}$.

D. Properties of the Aggregated Quotient Markov Chain

We study the one-step probability mass function associated to the codomain of the labeling function (that is, to the labels set), conditional on the state of the chain Ξ.

Lemma 1. *The conditional random variable $(x_i(t+1)|\mathbf{z}(t))$, $i \in \mathbb{N}_{2n}$, has a Poisson-binomial distribution over the sample space $\mathbb{Z}_{n_p}$, with the following mean and variance:*

$$\begin{array}{l}\mathbb{E}[x_i(t+1)|\mathbf{z}(t)] = \sum_{r=1}^{n_p} P_{z_r(t)i}, \\ var(x_i(t+1)|\mathbf{z}(t)) = \sum_{r=1}^{n_p} P_{z_r(t)i}(1 - P_{z_r(t)i}).\end{array} \tag{7}$$

Proof. Recall that the states of all the Markov chains at time t are known; Markov chain r jumps to the state i with probability $P_{z_r(t)i}$ and fails to jump to the state i with probability $(1 - P_{z_r(t)i})$. The definition of the variable x_i implies that the conditional random variable $(x_i(t+1)|\mathbf{z}(t))$ is the sum of n_p independent Bernoulli trials with different success probabilities $P_{z_r(t)i}$, which is a Poisson-binomial distribution with the mean and variance as in (7). □

Conditional on an observation $\mathbf{x} = [x_1, x_2, \ldots, x_{2n}]^T$ at time t over the Markov chain Ξ, it is of interest to compute the probability mass function of the conditional random variable $(x_i(t+1)|\mathbf{x}(t))$ as $\mathbb{P}(x_i(t+1) = j|\mathbf{x}(t))$, for any $j \in \mathbb{Z}_{n_p}$ — notice the difference with the quantity discussed in (7) where the conditioning is over variable $\mathbf{z}(t)$. For any label $\mathbf{x} = [x_1, \ldots, x_{2n}]^T$ there are exactly $n_p!/(x_1!x_2!\cdots x_{2n}!)$ states of Ξ such that $L(\mathbf{z}) = \mathbf{x}$. We use the notation $\mathbf{z} \hookrightarrow \mathbf{x}$ to indicate the states in Ξ associated to label $\mathbf{x}$, that is

[1] An atomic proposition is a proposition whose truth or falsity does not depend on that of other propositions. Sets of atomic propositions are used to assign labels to the states of a Markov chain [4].

$\mathbf{z} : L(\mathbf{z}) = \mathbf{x}$. Based on the law of total probability for conditional probabilities, we can write

$$\begin{aligned}\mathbb{P}(x_i(t+1) &= j|\mathbf{x}(t)) \\ &= \frac{\sum_{\mathbf{z}(t)\hookrightarrow\mathbf{x}(t)} \mathbb{P}(x_i(t+1)=j|\mathbf{z}(t))\mathbb{P}(\mathbf{z}(t))}{\mathbb{P}(\mathbf{x}(t))} \\ &= \mathbb{P}(x_i(t+1)=j|\mathbf{z}(t))\frac{\sum_{\mathbf{z}(t)\hookrightarrow\mathbf{x}(t)} \mathbb{P}(\mathbf{z}(t))}{\mathbb{P}(\mathbf{x}(t))} \\ &= \mathbb{P}(x_i(t+1)=j|\mathbf{z}(t)),\end{aligned} \tag{8}$$

where the sum is over all states $\mathbf{z}(t)$ of Ξ such that $L(\mathbf{z}(t)) = \mathbf{x}(t)$: in these states we have $x_1(t)$ Markov chains in state 1 with probability P_{1i}, $x_2(t)$ Markov chains in state 2 with probability P_{2i}, and so on. The simplification above is legitimate since the probability of having a label $\mathbf{x} = (x_1, x_2, \ldots, x_{2n})$ is exactly the sum of the probabilities associated to the states $\mathbf{z}$ generating such a label. This further allows expressing the quantities in (7) as

$$\mathbb{E}[x_i(t+1)|\mathbf{z}(t)] = \sum_{r=1}^{n_p} P_{z_r(t)i} = \sum_{r=1}^{2n} x_r(t)P_{ri}.$$

The generalization of the previous results to vector labels leads to the following statement.

Theorem 1. *The conditional random variables $(x_i(t+1)|\mathbf{x}(t))$ are characterized by Poisson-binomial distributions, whereas the conditional random vector $(\mathbf{x}(t+1)|\mathbf{x}(t))$ by a generalized multinomial distribution. Mean, variance, and covariance are described, $\forall i, j \in \mathbb{N}_{2n}, i \neq j$, by*

$$\begin{aligned}&\mathbb{E}[x_i(t+1)|\mathbf{x}(t)] = \sum_{r=1}^{2n} x_r(t)P_{ri}, \\ &var(x_i(t+1)|\mathbf{x}(t)) = \sum_{r=1}^{2n} x_r(t)P_{ri}(1-P_{ri}), \\ &cov(x_i(t+1), x_j(t+1)|\mathbf{x}(t)) = -\sum_{r=1}^{2n} x_r(t)P_{ri}P_{rj}.\end{aligned}$$

Theorem 1 indicates that the distribution of the conditional random variable $(\mathbf{x}(t+1)|\mathbf{x}(t))$ is independent of the underlying state $\mathbf{z}(t) \hookrightarrow \mathbf{x}(t)$ of Ξ. With focus on equation (6), this result allows us to claim the following.

Corollary 1. *The equivalence relation $\mathcal{R}$ is an exact probabilistic bisimulation over the Markov chain Ξ. The resulting quotient Markov chain is the coarsest probabilistic bisimulation of Ξ.*

Without loss of generality, let us normalize the values of the labels $\mathbf{x}$ by the total population size n_p, thus obtaining a new variable $\mathbf{X}$. The conditional variable $(\mathbf{X}(t+1)|\mathbf{X}(t))$ is characterized by the following parameters, for all $i, j \in \mathbb{N}_{2n}, i \neq j$:

$$\begin{aligned}&\mathbb{E}[X_i(t+1)|\mathbf{X}(t)] = \sum_{r=1}^{2n} X_r(t)P_{ri}, \\ &var(X_i(t+1)|\mathbf{X}(t)) = \frac{1}{n_p}\sum_{r=1}^{2n} X_r(t)P_{ri}(1-P_{ri}), \\ &cov(X_i(t+1), X_j(t+1)|\mathbf{X}(t)) = -\frac{1}{n_p}\sum_{r=1}^{2n} X_r(t)P_{ri}P_{rj}.\end{aligned} \tag{9}$$

Based on the expression of the first two moments of $(\mathbf{X}(t+1)|\mathbf{X}(t))$, we apply a translation (shift) on this conditional random vector as $\mathbf{W}(t) = \mathbf{X}(t+1) - \mathbb{E}[\mathbf{X}(t+1)|\mathbf{X}(t)]$, where $\mathbf{W}(t) = [\omega_1(t), \ldots, \omega_{2n}(t)]^T$ and $\omega_i(t)$ are guaranteed to be (dependent) random variables with zero mean and covariance described by (9) and dependent on the state $\mathbf{X}(t)$. Such a translation allows expressing the following dynamical model for the variable $\mathbf{X}$:

$$\mathbf{X}(t+1) = P^T\mathbf{X}(t) + \mathbf{W}(t). \tag{10}$$

Remark 1. *We have modeled the evolution of the TCL population with the abstract aggregated model* (10)*, characterized by a stochastic difference equation. The dynamics in* (10) *represent a direct generalization of the model abstraction provided in [22]. A closer look on the covariance terms in* (9) *suggests that they are bounded by the quantity $1/n_p$, and thus decrease to zero as n_p grows, regardless of the number of bins n. This result not only relates the model in* (10) *to that in [22], but also allows generalizing the results of [8], [9], where convergence to a* deterministic *difference equation for large populations of Markov chains is investigated in the context of mean field limits. Notice further that the state process noise $\mathbf{W}(t)$ is obtained because of the use of a finite population size, which can be seen as a specific form of modeling mismatch, as indeed qualitatively claimed in [24], [26] and now formally explained in this work.*

Above we have characterized the random variable $(X_i(t+1)|\mathbf{X}(t))$ with a Poisson-binomial distribution. We employ the Lyapunov central limit theorem [7] to show that this distribution converges to a Gaussian one.

Theorem 2. *The random variable $(X_i(t+1)|\mathbf{X}(t))$ can be explicitly expressed as*

$$X_i(t+1) = \sum_{r=1}^{2n} X_r(t)P_{ri} + \omega_i(t), \tag{11}$$

where the random vector $\mathbf{W}(t) = [\omega_1(t), \ldots, \omega_{2n}(t)]^T$ has a covariance matrix $\Sigma(\mathbf{X}(t))$ as in (9)*, and converges (in distribution) to a multivariate Gaussian vector $\mathcal{N}(0, \Sigma(\mathbf{X}(t)))$, as $n_p \uparrow \infty$.*

Proof. In order to prove that the random vector $(\mathbf{X}(t+1)|\mathbf{X}(t))$ converges to a multivariate normal random variable, we show that every linear combination of its components converges to a normal random variable. Consider any arbitrary vector $\boldsymbol{\nu} = [\nu_1, \nu_2, \ldots, \nu_{2n}]^T \in \mathbb{R}^{2n}$. The random variable $(\boldsymbol{\nu}^T\mathbf{X}(t+1)|\mathbf{X}(t))$ can be seen as the sum of n_p independent (normalized) categorical random variables y_j [20] over the sample space $\{\frac{\nu_1}{n_p}, \frac{\nu_2}{n_p}, \ldots, \frac{\nu_{2n}}{n_p}\}$, where $x_1(t)$ of them have success probability $\mathbf{p}_1 = [P_{11}, P_{12}, \ldots, P_{12n}]$, $x_2(t)$ have success probability $\mathbf{p}_2 = [P_{21}, P_{22} \ldots, P_{22n}]$, and so on. Then

$$\begin{aligned}&\lim_{n_p\to\infty} \frac{1}{s_{n_p}^3}\sum_{j=1}^{n_p}\mathbb{E}\left[|y_j-\mu_j|^3\right] \\ &= \lim_{n_p\to\infty} \frac{n_p^{3/2}}{n_p^2}\frac{\sum_{r=1}^{2n}\sum_{i=1}^{2n}\mathbf{X}_r(t)|\nu_i - \mathbf{p}_r\boldsymbol{\nu}|^3}{\left(\sum_{r=1}^{2n}\sum_{i=1}^{2n}\mathbf{X}_r(t)|\nu_i-\mathbf{p}_r\boldsymbol{\nu}|^2\right)^{3/2}} = 0,\end{aligned}$$

since both numerator and denominator of the second fraction are constant and thus independent of n_p. On the other hand, the mean and variance can be obtained based on the direct definition of $\boldsymbol{\nu}^T\mathbf{X}(t+1)$ and the relation in (9). Based on the Lyapunov central limit theorem [7] we are able to conclude that (the convergence is in distribution)

$$\left(\frac{\boldsymbol{\nu}^T\mathbf{X}(t+1) - \boldsymbol{\nu}^T P^T\mathbf{X}(t)}{\sqrt{\boldsymbol{\nu}^T\Sigma(\mathbf{X}(t))\boldsymbol{\nu}}}\middle|\mathbf{X}(t)\right) \xrightarrow{d} \mathcal{N}(0,1).$$

Defining the variable $\omega_i(t)$ through equation (11) leads to $\mathbf{W}(t) \xrightarrow{d} \mathcal{N}(0, \Sigma(\mathbf{X}(t)))$. □

Theorem 2 practically states that the conditional distribution of the random vector $\mathbf{W}(t)$ for a relatively large population size can be effectively replaced by a multivariate Gaussian distribution with known moments.[2] We shall exploit this result in the state estimation step, where we employ a conditional Kalman filter, as discussed in Section IV. Notice that this conclusion holds for any population of homogeneous TCLs that are characterized by Markov chains having the same transition probability matrix, and initial distributions that can instead vary.

In the previous theorem we have developed a model for the evolution of $X_i(t)$, which in the limit includes a Gaussian noise $\omega_i(t)$. As discussed in (9), these Gaussian random variables depend on $\mathbf{X}(t)$. The covariance matrix in (9) is guaranteed to be positive semi-definite for all $X_r(t) \in \{0, \frac{1}{n_p}, \frac{2}{n_p}, \dots, \frac{n_p-1}{n_p}, 1\}$, provided that $\sum_{r=1}^{2n} X_r(t) = 1$. In order to enable a more general use in (10), we next show that the covariance matrix remains positive semi-definite when the model is extended over the variables $X_r \in [0, 1]$.

Proposition 1. *The covariance matrix $\Sigma(\mathbf{X})$ is positive semi-definite for all $X_r \geq 0$. The entries of the random vector $\mathbf{W}$ are dependent on each other, since $\sum_{r=1}^{2n} \omega_r = 0$ whenever $\sum_{r=1}^{2n} X_r = 1$. Condition $\sum_{r=1}^{2n} X_r(0) = 1$ implies that $\sum_{r=1}^{2n} X_r(t) = 1$, for all $t \in \mathbb{N}$.*

Proof. The matrix $\Sigma(\mathbf{X})$ can be written as $\frac{1}{n_p} \sum_{r=1}^{2n} X_r \Phi_r$, where

$$\Phi_r = \begin{bmatrix} P_{r1}(1-P_{r1}) & -P_{r1}P_{r2} & \cdots & -P_{r1}P_{r2n} \\ -P_{r2}P_{r1} & P_{r2}(1-P_{r2}) & \cdots & -P_{r2}P_{r2n} \\ \vdots & \vdots & \vdots & \vdots \\ -P_{r2n}P_{r1} & -P_{r2n}P_{r2} & \cdots & P_{r2n}(1-P_{r2n}) \end{bmatrix}$$

The above structure of matrix Φ_r allows us to compute, for all $\boldsymbol{\nu} \in \mathbb{R}^{2n}$,

$$\boldsymbol{\nu}^T \Phi_r \boldsymbol{\nu} = \sum_{i=1}^{2n} P_{ri} \nu_i^2 - \left(\sum_{i=1}^{2n} P_{ri} \nu_i \right)^2 .$$

The Cauchy-Schwartz inequality implies that $\boldsymbol{\nu}^T \Phi_r \boldsymbol{\nu} \geq 0$, where the equality holds at least for the vectors $\boldsymbol{\nu} = c \mathbb{1}_{2n}^T$ with any arbitrary constant c. The positive semi-definiteness of all Φ_r implies the positive semi-definiteness of $\Sigma(\mathbf{X})$, for all $X_r \geq 0$.

The random variable $\omega = \mathbb{1}_{2n} \mathbf{W} = \sum_{r=1}^{2n} \omega_r$ has expected value $\mathbb{E}[\omega] = \mathbb{1}_{2n} \mathbb{E}[\mathbf{W}] = 0$ and variance $\sigma^2(\omega) = \mathbb{E}[\omega \omega^T] = \mathbb{1}_{2n} \Sigma(X) \mathbb{1}_{2n}^T = 0$. Then the random variable ω is deterministic: $\omega = 0$. The last part of the theorem is proven by taking the sum of all the equations of the dynamical system and noticing that matrix P is a stochastic matrix:

$$\sum_{r=1}^{2n} X_r(t+1) = \sum_{r=1}^{2n} X_r(t) + \sum_{r=1}^{2n} \omega(t) = \sum_{r=1}^{2n} X_r(t).$$

□

[2]The interested reader may refer to [7] for a discussion on the rate of convergence in the central limit theorem and on the quality of finite approximations. In practice any choice of a finite population size ought to be with respect to the stochastic matrix P and to the conditional state vector $\mathbf{X}(t)$.

E. Explicit Quantification of the Errors of the Abstraction and of the Aggregation Procedures

Let us now quantify the power consumption of the aggregate model, as an extension of the quantity discussed after equation (2). The total power consumption obtained from the aggregation of the original models in (1)-(2), with variables $(m_j, \theta_j)(t), j \in \mathbb{N}_{n_p}$, is

$$y(t) = \sum_{j=1}^{n_p} m_j(t) P_{ON}. \tag{12}$$

With focus on the abstract model (with the normalized variable $\mathbf{X}$), the power consumption is

$$y_{abs}(t) = H \mathbf{X}(t), \quad H = n_p P_{ON} [0_n, \mathbb{1}_n], \tag{13}$$

where 0_n and $\mathbb{1}_n$ are row vectors with all the entries equal to zero and one, respectively. For the quantification of the error we consider a homogeneous population of TCLs with dynamics affected by Gaussian process noise $w(\cdot) \sim \mathcal{N}(0, \sigma^2)$, and the abstracted model constructed based on the partition introduced in (4). The main result of this section hinges on two features of the Gaussian density function, its continuity and its decay at infinity. In order to keep the discussion focused we continue considering Gaussian distributions, however the results can be extended to any distribution with these two features. Since the covariance matrix in (9) is small for large population sizes, the first moment of the random variable $y(t)$ provides sufficient information on its behavior over a finite time horizon. The total power consumption in (12) is the sum of n_p independent Bernoulli trials over the sample space $\{0, P_{ON}\}$, each with different success probability. Then for the quantification of the modeling error we study the error produced by the abstraction over the expected value of the binary TCL mode (ON, OFF).

Consider a TCL, with initial state $s_0 = (m_0, \theta_0)$. Also select a desired final time T_d and discrete time horizon $N = T_d/h$, where h is the discretization step. Based on the evolution equation of the discrete mode (2), the TCL is in the ON mode at time step N, namely $m(N) = 1$, if and only if $s(N-1) \in \mathcal{A}$, where $\mathcal{A} = \{1\} \times [\theta_-, +\infty) \cup \{0\} \times [\theta_+, +\infty)$. Then the expected value of its mode at time N, $m(N)$, can be computed as

$$\begin{aligned} \mathbb{E}[m(N)|m_0, \theta_0] &= \mathbb{P}\left(m(N) = 1 | m_0, \theta_0\right) \\ &= \mathbb{P}\left(s(N-1) \in \mathcal{A} | m_0, \theta_0\right). \end{aligned} \tag{14}$$

This quantity can be characterized via value functions $\mathcal{V}_k : \mathcal{S} \to [0, 1]$, $k \in \mathbb{N}_N$, which are computed recursively as follows:

$$\mathcal{V}_k(s_k) = \int_{\mathcal{S}} \mathcal{V}_{k+1}(s_{k+1}) t_{\mathfrak{s}}(s_{k+1}|s_k) ds_{k+1}, \tag{15}$$

for all $k \in \mathbb{N}_{N-1}$, initialized with $\mathcal{V}_N(s) = \mathbb{I}_{\mathcal{A}}(s)$. Knowing these value functions, we have that $\mathbb{E}[m(N)|m_0, \theta_0] = \mathcal{V}_1(m_0, \theta_0)$. Computationally, the calculation of these quantities can leverage the results in [1], [14], [16], which however require extensions 1) to conditional density functions of the process that are discontinuous, and 2) to an unbounded state space. The first issue is addressed by the following lemma.

Lemma 2. *The density function $t_{\mathfrak{s}}(s'|\cdot)$ is piecewise-continuous within the continuity regions*

$$\begin{aligned} &\{0\} \times (-\infty, \theta_+], \quad \{0\} \times (\theta_+, +\infty), \\ &\{1\} \times (-\infty, \theta_-), \quad \{1\} \times [\theta_-, +\infty). \end{aligned}$$

The value functions $\mathcal{V}_k(s)$ are piecewise-Lipschitz continuous, namely:

$$|\mathcal{V}_k(m,\theta) - \mathcal{V}_k(m,\theta')| \leq \frac{2a}{\sigma\sqrt{2\pi}}|\theta - \theta'|,$$

where a, σ represent respectively the TCL parameter and the variance of the process noise $w(t)$ as in equation (1)*, and where $(m,\theta),(m,\theta')$ is any pair of points belonging to one of the four continuity regions of the function $t_{\mathfrak{s}}(s'|\cdot)$.*

Proof. The density function of the temperature process noise $w(\cdot)$ is $t_{\mathfrak{w}}(w) = \frac{1}{\sigma}\phi(\frac{w}{\sigma})$, where $\phi(\cdot)$ is the standard Gaussian density function $\phi(\mathsf{x}) = \exp(-\mathsf{x}^2/2)/\sqrt{2\pi}$. We prove the statement for one of the continuity regions, namely $m = 0$ and $\theta, \theta' \in (-\infty, \theta_+]$, the other regions being treated in the same way, via the following chain of inequalities:

$$\begin{aligned}
&|\mathcal{V}_k(m,\theta) - \mathcal{V}_k(m,\theta')| \\
&\leq \textstyle\int_{\mathbb{R}} \mathcal{V}_{k+1}(0,\bar{\theta})|t_{\mathfrak{w}}(\bar{\theta} - a\theta - (1-a)\theta_a) \\
&\qquad\qquad -t_{\mathfrak{w}}(\bar{\theta} - a\theta' - (1-a)\theta_a)|d\bar{\theta} \\
&\leq \textstyle\int_{\mathbb{R}} \left|t_{\mathfrak{w}}(\bar{\theta} - a\theta - (1-a)\theta_a) - t_{\mathfrak{w}}(\bar{\theta} - a\theta' - (1-a)\theta_a)\right| d\bar{\theta} \\
&= \textstyle\frac{1}{\sigma}\int_{\mathbb{R}} \left|\phi\left(\frac{\bar{\theta} - a\theta - (1-a)\theta_a}{\sigma}\right) - \phi\left(\frac{\bar{\theta} - a\theta' - (1-a)\theta_a}{\sigma}\right)\right| d\bar{\theta} \\
&= 2\textstyle\int_0^{\infty}\left[\phi\left(u - \frac{a|\theta-\theta'|}{2\sigma}\right) - \phi\left(u + \frac{a|\theta-\theta'|}{2\sigma}\right)\right] du \\
&= 2\textstyle\int_{-a|\theta-\theta'|/2\sigma}^{a|\theta-\theta'|/2\sigma} \phi(v)dv \leq \frac{2a}{\sigma\sqrt{2\pi}}|\theta - \theta'|.
\end{aligned}$$

□

In order to cope with the second issue, we study the limiting behavior of the value functions at infinity, apply a truncation over the state space, as proposed in Section II-B, and properly select the value of the functions outside of this region. Lemma 3 shows that $\lim_{\theta\to-\infty}\mathcal{V}_k(m,\theta) = 0$, $\lim_{\theta\to+\infty}\mathcal{V}_k(m,\theta) = 1$, and provides an upper bound on the distance between $\mathcal{V}_k(m,\theta)$ and its limiting values, which hinges on the parameter $\mathcal{L}$. This parameter represents the length of the truncated part of the temperature range $[\theta_{-\mathsf{m}}, \theta_{\mathsf{m}}]$, which is then further partitioned to construct the abstract Markov chain. An upper bound on the error of the value functions produced by state-space truncation and partitioning is further quantified in Theorem 3.

Lemma 3. *For the partitioning procedure in* (4) *we have that*

$$\begin{aligned}
&\mathcal{V}_k(m,\theta) \leq (N-k)\epsilon \quad \forall\theta \leq \theta_{-\mathsf{m}} = \theta_s - \mathcal{L}/2, \\
&\mathcal{V}_k(m,\theta) \geq 1 - (N-k)\epsilon \quad \forall\theta \geq \theta_{\mathsf{m}} = \theta_s + \mathcal{L}/2,
\end{aligned} \tag{16}$$

for $m \in \mathbb{Z}_1$, where $\epsilon = \dfrac{e^{-\gamma^2/2}}{\gamma\sqrt{2\pi}}$, and where

$$\begin{aligned}
\gamma &= \frac{1-a}{2\sigma}\left[\frac{a^N\mathcal{L} + \delta}{1-a^N} - \lambda\right], \\
\lambda &= RP_{rate} + |2(\theta_s - \theta_a) + RP_{rate}|.
\end{aligned} \tag{17}$$

Proof. Define the sequences $\{\beta_k\}_{k=0}^N, \{\gamma_k\}_{k=0}^N$ by the linear difference equations

$$\begin{aligned}
&\beta_k = (\beta_{k+1} - \gamma\sigma - (1-a)\theta_a)/a, \quad \beta_N = \theta_+, \\
&\gamma_k = (\gamma_{k+1} + \gamma\sigma - (1-a)(\theta_a - RP_{rate}))/a, \quad \gamma_N = \theta_-.
\end{aligned} \tag{18}$$

We prove inductively that the inequalities in (16) hold for $\theta \leq \beta_k$ and $\theta \geq \gamma_k$, respectively. Because of the similarities in the proof we present the reasoning only for $\theta \leq \beta_k$ and $m = 0$. The recursive equation (15) implies that for any $\theta \leq \beta_k$ and $m = 0$,

$$\begin{aligned}
\mathcal{V}_k(0,\theta) &= \textstyle\int_{-\infty}^{+\infty}\mathcal{V}_{k+1}(0,\bar{\theta})\frac{1}{\sigma}\phi\left(\frac{\bar{\theta} - a\theta - (1-a)\theta_a}{\sigma}\right)d\bar{\theta} \\
&\leq \textstyle\int_{-\infty}^{\beta_{k+1}}(N-k-1)\epsilon\frac{1}{\sigma}\phi\left(\frac{\bar{\theta} - a\theta - (1-a)\theta_a}{\sigma}\right)d\bar{\theta} \\
&\quad + \textstyle\int_{\beta_{k+1}}^{+\infty}\frac{1}{\sigma}\phi\left(\frac{\bar{\theta} - a\theta - (1-a)\theta_a}{\sigma}\right)d\bar{\theta} \\
&\leq (N-k-1)\epsilon + \textstyle\int_{\gamma}^{+\infty}\phi(u)\,du \\
&\leq (N-k-1)\epsilon + \epsilon = (N-k)\epsilon.
\end{aligned}$$

Since the parameter $a = e^{-h/RC} \in (0,1)$ and $[\theta_-, \theta_+] \subset [\theta_a - RP_{rate}, \theta_a]$, the sequences in (18) are monotonic. The selected value of γ in (17) ensures that $\beta_k \geq \theta_{-\mathsf{m}}$ and $\gamma_k \leq \theta_{\mathsf{m}}$ for $k \in \mathbb{N}_N$, which concludes the proof. □

Notice in particular the linear dependence of γ on $\mathcal{L}$, the temperature interval of interest.

Theorem 3. *If we abstract a TCL as a Markov chain based on the procedure of Section II-B, compute the solution of problem* (15) *over the Markov chain, and construct a piecewise-constant approximation function $\mathcal{W}_1(m,\theta)$ using the solution of* (15) *over the Markov chain, then the approximation error can be upper bounded as follows:*

$$|\mathcal{V}_1(m,\theta) - \mathcal{W}_1(m,\theta)| \leq (N-1)\left[\frac{N-2}{2}\epsilon + \frac{2a}{\sigma\sqrt{2\pi}}\upsilon\right],$$

for all $(m,\theta) \in \mathbb{Z}_1 \times [\theta_{-\mathsf{m}}, \theta_{\mathsf{m}}]$.

Notice that the error has two terms: one term accounts for the error of the approximation over infinite-length intervals ϵ, whereas the second is related to the choice of the partition size υ.

Proof. We define E_k as an upper bound for $|\mathcal{V}_k(m,\theta) - \mathcal{W}_k(m,\theta)|$ and compute it recursively. The proof of Lemma 3 indicates that the same result is valid for the functions $\mathcal{W}_1(m,\theta)$. We have $|\mathcal{V}_k(m,\theta) - \mathcal{W}_k(m,\theta)| \leq (N-k)\epsilon$, for all (m,θ) belonging to the external infinite-length intervals. Recall that the value functions $\mathcal{V}_k$ satisfy the recursion in (15). We discuss this step for $m = 0$, $\theta_+ \leq \theta \leq \theta_{\mathsf{m}}$, the other four possibilities being the same. Suppose that $\theta \in \Theta_i$ with representative point $\bar{\theta}_i$, then:

$$\begin{aligned}
&|\mathcal{V}_k(0,\theta) - \mathcal{W}_k(0,\theta)| \\
&\qquad \leq |\mathcal{V}_k(0,\theta) - \mathcal{V}_k(0,\bar{\theta}_i)| + |\mathcal{V}_k(0,\bar{\theta}_i) - \mathcal{W}_k(0,\bar{\theta}_i)| \\
&\qquad \leq \tfrac{2a}{\sigma\sqrt{2\pi}}|\theta - \bar{\theta}_i| + |\mathcal{V}_k(0,\bar{\theta}_i) - \mathcal{W}_k(0,\bar{\theta}_i)|.
\end{aligned}$$

The second term is upper bounded as follows

$$\begin{aligned}
&|\mathcal{V}_k(0,\bar{\theta}_i) - \mathcal{W}_k(0,\bar{\theta}_i)| \\
&\leq \textstyle\int_{-\infty}^{\infty}|\mathcal{V}_{k+1}(1,\bar{\theta}) - \mathcal{W}_{k+1}(1,\bar{\theta})|t_{\mathfrak{w}}(\bar{\theta} - a\bar{\theta}_i - (1-a)\theta_a)d\bar{\theta} \\
&\leq (N-k-1)\epsilon + E_{k+1}\textstyle\int_{\theta_{-\mathsf{m}}}^{\theta_{\mathsf{m}}} t_{\mathfrak{w}}(\bar{\theta} - a\bar{\theta}_i - (1-a)\theta_a)d\bar{\theta} \\
&\leq (N-k-1)\epsilon + E_{k+1} \\
&\Rightarrow E_k = \tfrac{2a}{\sigma\sqrt{2\pi}}\upsilon + (N-k-1)\epsilon + E_{k+1}, \quad E_N = 0, \\
&\Rightarrow E_1 = (N-1)(N-2)\tfrac{\epsilon}{2} + (N-1)\tfrac{2a}{\sigma\sqrt{2\pi}}\upsilon.
\end{aligned}$$

□

Collecting the results above, the following corollary quantifies an upper bound on the abstraction error over the total power consumption.

Corollary 2. *The difference in the expected value of the total power consumption of the population $y(N)$ in* (12)*, and that*

of the abstracted model $y_{abs}(N)$ in (13)*, both conditional on the corresponding initial conditions, is upper bounded by*

$$\begin{aligned}\left|\mathbb{E}[y(N)|\mathbf{s}_0] - \mathbb{E}[y_{abs}(N)|\mathbf{X}_0]\right| \\ \leq n_p P_{ON}(N-1)\left[\frac{(N-2)}{2}\epsilon + \frac{2a}{\sigma\sqrt{2\pi}}\upsilon\right],\end{aligned} \tag{19}$$

for all $\mathbf{s}_0 \in (\mathbb{Z}_1 \times [\theta_{-\mathsf{m}}, \theta_{\mathsf{m}}])^{n_p}$. The initial state $\mathbf{X}_0$ is obtained as a function of the initial states in the TCL population $\mathbf{s}_0$, as can be evinced from the definition of the state vector $\mathbf{X}$.

Proof. Suppose all the TCLs are initialized at $s_{0j} = (m_{0j}, \theta_{0j})$, $j \in \mathbb{N}_{n_p}$. Recall (12) for the total power consumption and the equality $\mathbb{E}[m_j(N)|s_{0j}] = \mathcal{V}_1(s_{0j})$ for any TCL. Then we have

$$\mathbb{E}[y(N)|\mathbf{s}_0] = P_{ON}\sum_{j=1}^{n_p}\mathbb{E}[m_j(N)|s_{0j}] = P_{ON}\sum_{j=1}^{n_p}\mathcal{V}_1(s_{0j}). \tag{20}$$

Recall (13) for $y_{abs}(t)$ and the difference equation (10),

$$\begin{aligned}\mathbb{E}[y_{abs}(N)|\mathbf{X}_0] &= H(P^T)^N\mathbf{X}_0 \\ &= n_p P_{ON}\sum_{m=0}^{1}\sum_{i=1}^{n}\mathbb{E}[m(N)|s_{im}]X_{0i},\end{aligned}$$

where $s_{im} = (m, \bar{\theta}_i)$ is the initial discrete state of the Markov chain. Since $\mathcal{W}_k(\cdot)$ is constructed based on the solution of (15) for the Markov chain, which is a piecewise-constant function over the selected partition, we have

$$\begin{aligned}\mathbb{E}[y_{abs}(N)|\mathbf{X}_0] &= P_{ON}\sum_{m=0}^{1}\sum_{i=1}^{n}\mathcal{W}_1(s_{im})n_p X_{0i} \\ &= P_{ON}\sum_{j=1}^{n_p}\mathcal{W}_1(s_{0j}).\end{aligned} \tag{21}$$

By subtracting (21) from (20) we obtain

$$\begin{aligned}|\mathbb{E}[y(N)|\mathbf{s}_0] - \mathbb{E}[y_{abs}(N)|\mathbf{X}_0]| \\ \leq P_{ON}\sum_{j=1}^{n_p}|\mathcal{V}(s_{0j}) - \mathcal{W}(s_{0j})| \leq P_{ON}n_p E_1.\end{aligned}$$

□

The result in Corollary 2 allows tuning the error over the total power consumption of the population made with the abstraction procedure. In practice, it can be reduced to a desired level by increasing the abstraction precision: this can be achieved by increasing γ and the number of state bins (by decreasing their size). This results in a larger-dimensional model in (10). To address this issue *post facto*, model-order reduction techniques like *balanced realization and truncation* or *Hankel singular values* [3] can be employed to obtain a lower-dimensional model describing the dynamics of the population power consumption. These known techniques follow the observation that the dynamics of the model are mostly determined by the largest eigenvalues of the transition probability matrix. These model reduction techniques are not exclusively applicable to homogeneous populations of TCLs, but can as well be employed for the heterogeneous populations case discussed in the following section.

III. Formal Abstraction of a Heterogeneous Population of TCLs

Consider a heterogeneous population of n_p TCLs, where the heterogeneity resides on a parameter α that takes n_p possibly different values as $\{\alpha_1, \alpha_2, ..., \alpha_{n_p}\}$. Each instance of α specifies a set of model parameters $(\theta_s, \delta, \theta_a, C, R, \sigma, P_{rate}, P_{ON})$ for the dynamics of a single TCL: notice that all these parameters influence the temperature evolution, except P_{ON}, which affects exclusively the output equation. Each dynamical model for a TCL can be abstracted as a Markov chain $\mathcal{M}_\alpha$ with a transition matrix $P_\alpha = [P_{ij}(\alpha)]_{i,j}$, according to the procedure explained in Section II. As expected, the transition probability matrix P_α obtained for a TCL depends on its own set of parameters α.

With the objective of an aggregated Markov chain model for a population of n_p TCLs, the goal is again that of abstracting it as a reduced-order (lumped) model. The apparent difficulty is that the heterogeneity in the transition probability matrix P_α of the single TCL renders the quantity $\mathbb{P}(x_i(t+1) = j|\mathbf{z}(t))$ dependent not only on the label $\mathbf{x}(t) = L(\mathbf{z}(t))$, but effectively on the current state $\mathbf{z}(t)$, namely the present distribution of temperatures of each TCL. This leads to the impossibility of simplifying equation (8), as done in the homogeneous case. Recall that computations on $\mathbb{P}(\mathbf{z}(t))$ require manipulations over the large-dimensional matrix P_Ξ, which can become practically infeasible.

In contrast to the homogeneous case, which allows one to quantify the probabilities $\mathbb{P}(x_i(t+1) = j|\mathbf{x}(t))$ over a Markov chain obtained as an exact probabilistic bisimulation of the product chain Ξ, in the heterogeneous case we resort to an *approximate* probabilistic bisimulation [13] of the Markov chain Ξ. The approximation enters in equation (8), with the replacement of the weighted average in the expression of the law of total probability with a normalized (equally weighted) average, as follows:

$$\mathbb{P}(x_i(t+1) = j|\mathbf{x}(t)) = \frac{\sum_{\mathbf{z}(t)\hookrightarrow\mathbf{x}(t)}\mathbb{P}(x_i(t+1) = j|\mathbf{z}(t))}{\#\{\mathbf{z}(t)\hookrightarrow\mathbf{x}(t)\}}, \tag{22}$$

where $\#A$ indicates the cardinality of a given finite set A. In other words we have assumed that the probability for the Markov chain Ξ to be in each labeled state is the same. Similarly, the average of the random variables $x_i(t+1)$, conditioned over $\mathbf{x}(t)$, can be obtained from (22) as $\mathbb{E}[x_i(t+1)|\mathbf{x}(t)] = \frac{\sum_{\mathbf{z}(t)\hookrightarrow\mathbf{x}(t)}\mathbb{E}[x_i(t+1)|\mathbf{z}(t)]}{\#\{\mathbf{z}(t)\hookrightarrow\mathbf{x}(t)\}}$. Unlike in the exact bisimulation instance, the error introduced by the approximate probabilistic bisimulation relation can only be quantified empirically over matrix P_Ξ.

Next, we put forward two alternative approaches to characterize the properties of the aggregated TCL population model: by an averaging argument in Section III-A, and by a clustering assumption in Section III-B. A technique related to the clustering one has also been proposed in [29], [30], however without formal quantification of the associated error and no indication on how to select the number of clusters.

A. Abstraction of a Heterogeneous Population of TCLs via Averaging

We characterize quantitatively the population heterogeneity by constructing an interpolated density function $f_\alpha(\cdot)$ from the finite set of values $\{\alpha_1, \alpha_2, ..., \alpha_{n_p}\}$ taken by parameter α. This leads to the characterization of the statistics of the conditional variable $(\mathbf{X}(t+1)|\mathbf{X}(t))$ (recall that $\mathbf{X}$ is a normalized version of $\mathbf{x}$), as follows.

Theorem 4. *Consider a TCL population with heterogeneity that is encompassed by a parameter α with empirical density function $f_\alpha(\cdot)$. Introducing an approximate probabilistic*

bisimulation of Ξ as in (22)*, the conditional random variable* $(\mathbf{X}(t+1)|\mathbf{X}(t))$ *has the following statistics:*

$$\begin{aligned}
&\mathbb{E}[X_i(t+1)|\mathbf{X}(t)] = \sum_{r=1}^{2n} X_r(t)\overline{P_{ri}},\\
&var(X_i(t+1)|\mathbf{X}(t)) = \frac{1}{n_p}\sum_{r=1}^{2n} X_r\overline{P_{ri}(1-P_{ri})}\\
&\qquad +\frac{1}{n_p-1}\left(\sum_{r=1}^{2n} X_r\overline{P_{ri}}\right)^2 - \frac{1}{n_p-1}\sum_{r=1}^{2n} X_r\overline{P_{ri}}^2,\\
&cov\left(X_i(t+1), X_j(t+1)|\mathbf{X}(t)\right)\\
&\qquad = \frac{1}{n_p-1}\left(\sum_{r=1}^{2n} X_r\overline{P_{ri}}\right)\left(\sum_{s=1}^{2n} X_s\overline{P_{sj}}\right)\\
&\qquad -\frac{1}{n_p-1}\sum_{r=1}^{2n} X_r\overline{P_{ri}P_{rj}} - \frac{1}{n_p}\sum_{r=1}^{2n} X_r\overline{P_{ri}P_{rj}},
\end{aligned}$$

where the barred quantities indicate an expected value over the parameters set α, for instance $\overline{P_{ri}P_{rj}} = \mathbb{E}_\alpha[P_{ri}(\alpha)P_{rj}(\alpha)] = \int P_{ri}(v)P_{rj}(v)f_\alpha(v)dv$.

Proof. The proof is directly obtained by computing the expected value and the covariance of $(x_i(t+1)|\mathbf{z}(t))$, and by taking an average over all the states $\mathbf{z}(t)$ that generate label $\mathbf{x}(t)$. □

Let us remark that the asymptotic properties obtained as the population size grows, as discussed in Section II-D, still hold as long as the distribution of the parameters set $f_\alpha(\cdot)$ is given and fixed. With focus on the heterogeneity in the output equation, we can similarly replace the ensemble of instances of parameter P_{ON} by its average quantity $\bar{P}_{ON}$, namely the mean rated power consumption of TCLs when all of them are in the ON mode, which is computed as the expected value of P_{ON} with respect to the parameter set $\bar{P}_{ON} = \mathbb{E}_\alpha\left[P_{ON}(\alpha)\right]$, and is time-independent. While (as discussed above) we cannot analytically quantify the error introduced by the approximate bisimulation used for the abstraction of the temperature evolution in the population, we can still quantify the error related to the heterogeneity in the output equation: this will be done shortly, in Theorem 5.

B. Abstraction of a Heterogeneous Population of TCLs via Clustering

We now propose an alternative method to reduce a heterogeneous population of TCLs into a finite number of homogeneous populations. While more elaborate than the preceding approach, it allows for the quantification of the error under the following Assumption.

Assumption 1. *Assume that the heterogeneity parameter $\alpha = (\theta_s, \delta, \theta_a, C, R, \sigma, P_{rate}, P_{ON})$ belongs to a bounded set $\Gamma_{\mathfrak{a}}$, and that the parametrized transition probability matrix P_α satisfies the following inequality expressing a condition on its continuity with respect to α:*

$$\|P_\alpha - P_{\alpha'}\|_\infty \leq h_{\mathfrak{a}}\|\alpha - \alpha'\| \quad \forall \alpha, \alpha' \in \Gamma_{\mathfrak{a}}. \tag{23}$$

Consider the range for a given parameter: we partition this uncertainty range and "cluster together" the TCLs in the given population, according to the partition they belong to; we further consider the TCLs to be homogeneous within their cluster. More precisely, select a finite partition of the set $\Gamma_{\mathfrak{a}} = \cup_i \Gamma_i$, characterized by a diameter $\upsilon_{\mathfrak{a}}$, namely $\|\alpha - \alpha'\| \leq \upsilon_{\mathfrak{a}}, \forall \alpha, \alpha' \in \Gamma_i, \forall i$. Associate arbitrary representative points $\alpha_i \in \Gamma_i$ to the partition sets. Finally, replace the transition matrix P_α and P_{ON} by $\sum_i P_{\alpha_i}\mathbb{I}_{\Gamma_i}(\alpha)$ and $\sum_i P_{ON}(\alpha_i)\mathbb{I}_{\Gamma_i}(\alpha)$, respectively. The error made by this procedure is quantified in the following statement.

Theorem 5. *Given a heterogeneous population of TCLs, suppose we cluster the heterogeneity parameter $\alpha \in \Gamma_{\mathfrak{a}} = \cup_i \Gamma_i$, assume homogeneity within the introduced clusters, and model each cluster based on the results of Section II with outputs $y_{abs,i}(N)$. Let us define the approximate power consumption of the heterogeneous population as the sum of outputs of the clusters, as follows: $y_{abs}(N) = \sum_i y_{abs,i}(N)$. The abstraction error can be upper bounded by*

$$\begin{aligned}
&\left|\mathbb{E}[y(N)|\mathbf{s}_0] - \mathbb{E}[y_{abs}(N)]\right|\\
&\quad \leq \max_\alpha n_p(N-1)P_{ON}(\alpha)\left[\tfrac{(N-2)}{2}\epsilon(\alpha) + \tfrac{2a(\alpha)}{\sigma(\alpha)\sqrt{2\pi}}\upsilon\right]\\
&\quad +n_p\left[\bar{P}_{ON}(N-1)h_{\mathfrak{a}} + 1\right]\upsilon_{\mathfrak{a}},
\end{aligned} \tag{24}$$

for all $\mathbf{s}_0 \in (\mathbb{Z}_1 \times [\theta_{-\mathfrak{m}}, \theta_{\mathfrak{m}}])^{n_p}$. The parameters $\epsilon(\cdot), \gamma(\cdot)$, and $\lambda(\cdot)$ are computed as in Corollary 2 and depend on the value of α. Finally, the quantity $\bar{P}_{ON} = \sum_i \frac{n_i}{n_p}P_{ON}(\alpha_i) = \mathbb{E}_\alpha\left[P_{ON}(\alpha)\right]$, where n_i is the population size of the i^{th} cluster, so that $\sum_i n_i = n_p$.

Proof. The equations in (20), (21) and Theorem 3 indicate that the first part of the error in (24) is an upper bound for the sum of abstraction error of each TCL.

The second part of the error is proved by studying the sensitivity of the solution of the problem (14) against parameter α. As we discussed before, the solution of this problem for the Markov chain over the time horizon N is obtained by the recursion $\bar{\mathcal{V}}_k(\alpha) = P(\alpha)\bar{\mathcal{V}}_{k+1}(\alpha)$, where $\bar{\mathcal{V}}_N(\alpha)$ is the indicator vector of the reach set, hence independent of α. Then we have

$$\begin{aligned}
\|\bar{\mathcal{V}}_k(\alpha) - \bar{\mathcal{V}}_k(\alpha')\|_\infty &= \|P(\alpha)\bar{\mathcal{V}}_{k+1}(\alpha) - P(\alpha')\bar{\mathcal{V}}_{k+1}(\alpha')\|_\infty\\
&\leq \|\left(P(\alpha) - P(\alpha')\right)\bar{\mathcal{V}}_{k+1}(\alpha)\|_\infty\\
&\quad +\|P(\alpha')\left(\bar{\mathcal{V}}_{k+1}(\alpha) - \bar{\mathcal{V}}_{k+1}(\alpha')\right)\|_\infty\\
&\leq h_{\mathfrak{a}}\|\alpha - \alpha'\| + \|\bar{\mathcal{V}}_{k+1}(\alpha) - \bar{\mathcal{V}}_{k+1}(\alpha')\|_\infty,
\end{aligned}$$

which results in the inequality

$$\|\bar{\mathcal{V}}_1(\alpha) - \bar{\mathcal{V}}_1(\alpha')\|_\infty \leq (N-1)h_{\mathfrak{a}}\|\alpha - \alpha'\|.$$

Define function $\xi(\cdot)$ that assigns to each α the representative parameter of its cluster. Then

$$\begin{aligned}
&\left|\sum_{\alpha\in\Gamma_{\mathfrak{a}}} P_{ON}(\alpha)\bar{\mathcal{V}}_1(\alpha) - \sum_i n_i P_{ON}(\alpha_i)\bar{\mathcal{V}}_1(\alpha_i)\right|\\
&\quad \leq \sum_{\alpha\in\Gamma_{\mathfrak{a}}}\left|P_{ON}(\alpha)\bar{\mathcal{V}}_1(\alpha) - P_{ON}(\xi(\alpha))\bar{\mathcal{V}}_1(\xi(\alpha))\right|\\
&\quad \leq \sum_{\alpha\in\Gamma_{\mathfrak{a}}}\left|P_{ON}(\alpha) - P_{ON}(\xi(\alpha))\right|\bar{\mathcal{V}}_1(\alpha)\\
&\quad + \sum_{\alpha\in\Gamma_{\mathfrak{a}}} P_{ON}(\xi(\alpha))\left|\bar{\mathcal{V}}_1(\alpha) - \bar{\mathcal{V}}_1(\xi(\alpha))\right|\\
&\quad \leq n_p\upsilon_{\mathfrak{a}} + (N-1)h_{\mathfrak{a}}\upsilon_{\mathfrak{a}}\sum_{\alpha\in\Gamma_{\mathfrak{a}}} n_i P_{ON}(\alpha_i).
\end{aligned}$$

□

Notice that the first part of the error in (24) is due to the abstraction of a single TCL by state-space partitioning, while the second part is related to the clustering procedure described above. Further, notice that all terms in the bound above can be reduced by selecting finer temperature partitions (smaller bins) or smaller clusters diameter for the parameter sets.

The second part of the error in (24) is computed based on the Lipschitz continuity of the transition probability matrix P_α, as per Assumption 1. This can be evaluated over the

transition probability matrices obtained by abstracting the heterogeneous TCL dynamics (characterized by the conditional density functions $t_{\mathfrak{s}}$) as Markov chains. Alternatively, we could formulate this error bound based on the Lipschitz continuity of the conditional density function $t_{\mathfrak{s}}(s'|s)$ with respect to the parameters set α by using the explicit relation (3) for the transition probabilities. Then the constant $h_{\mathfrak{a}}$ is computable as a function of the Lipschitz constant of the conditional density function of the process. As an example, the constant $h_{\mathfrak{a}}$ for the case of a Gaussian process noise and heterogeneity term residing exclusively in thermal capacitance (that is, in the parameter a) is computed as follows: $h_{\mathfrak{a}} = \dfrac{\mathcal{L}+\lambda}{\sigma\sqrt{2\pi}}$.

Remark 2. *Notice that we have presented a clustering approach that hinges on the similarity of the transition probability matrices of the associated Markov chains, measured as $\|P_\alpha - P_{\alpha'}\|_\infty$. There is no difference between heterogeneity in one or any other parameter: as long as the approximated Markov chains are similar, we can assign them to the same cluster. We can perform an abstraction by partitioning the space of each parameter, as presented in Theorem 5, however this approach clearly suffers from the curse of dimensionality. The alternative would be to first compute the transition matrices, then do clustering based on the distance measure $\|P_\alpha - P_{\alpha'}\|_\infty$, which is likely to provide a smaller number of clusters.*

As a final note, the result in Theorem 5 is applicable to the setup in Section III-A when the heterogeneity lies in the parameter P_{ON}, by considering a single cluster.

IV. Abstraction and Control of a Population of Non-Autonomous TCLs

One can imagine a number of different strategies for controlling the total power consumption of a population of TCLs. With focus on the dynamics of a single TCL, one strategy could be to vary the rate of the energy transfer P_{rate}, for instance by circulating cold/hot water through the load with higher or lower speed. Another approach could be to act on the thermal resistance R, for instance opening or closing doors and windows at the load end. Yet another strategy could be to apply changes to the set-point θ_s, as suggested in [10].

Let us observe that the first two actions would modify the dynamics of (1), whereas the third control action would affect the relation in (2). While abstracting the TCL model as a finite-state Markov chain, a control action results in a modification of the elements of the transition probability matrix. With reference to (5), the entries of the matrices $Q_{11}, Q_{22}, Q_{31}, Q_{42}$ are computed based on (1), while the size of these matrices is determined based on (2). Since the set-point θ_s affects only equation (2), a set-point alteration affects the *structure* of the probability matrix in (5), whereas the other approaches affect the *value* of its non-zero elements. It follows that in view of the abstraction procedure the control by set-point variation has the advantage of requiring a single computation of marginals, while the other discussed methods would require this computation to be a function of the allowed control inputs.

Based on the discussion above, we consider the case where the control input is taken to be the set-point θ_s of the TCL. We intend to apply the control input to all TCLs uniformly (cf. Figure 2), which does not require differentiating among the states of different TCLs. Moreover, in order to retain the validity of the definition of state bins $\mathbf{X}(\cdot)$ regardless of the applied input signal, we discretize the domain of allowable set-points by the same parameter υ used for the partition size.

With reference to existing closed-loop control schemes in the literature, [22] assumes full knowledge of the state vector $\mathbf{X}(t)$ and employs a Model Predictive Control architecture to design the control signal. Moving forward, [24], [26] consider different scenarios for the configuration of the control architecture: states are measured completely, or known partially and a Kalman filter is used for state estimation, or both states and transition matrix are estimated by use of an Extended Kalman filter. The minimum required infrastructure for the practical implementation of the strategies in [24], [26] ranges from a TCL temperature sensor and a two-way data connection for transmitting the state information and control signal, to a one-way data connection for only sending the control signal to all TCLs without receiving the explicit information on their states. The presence of a local decision maker is essential in all the scenarios: each TCL receives a control signal at each time step, determines its current state, and generates a local control action. In contrast, the set-point control strategy in this work does not require each single TCL to know its individual state, which makes the approach applicable regardless of the thermometer precision [10].

In the following we show that the knowledge of the actual states of the TCL, or of vector $\mathbf{X}(t)$ in the aggregated model, are not necessary for the global control of the total power consumption. Given the model parameters, all that is needed is an online measurement of the total power consumption of the TCL population, which allows estimating the states in $\mathbf{X}(t)$ by use of a conditional Kalman filter [11] adapted to the state-dependent covariance of the process noise. Moreover, we attempt to mitigate the aforementioned limitations by using the set-point θ_s as the control input to track a given reference signal. The control action then comprises a simple signal for the set-point that is applied to all TCLs uniformly: no local decision maker is required.

A. State Estimation and One-Step Regulation

Suppose we consider a homogeneous population of TCLs with known parameters. As discussed earlier, we assume that the control input is discrete and takes values over a finite set, $\theta_s(t) \in \{\theta_{-\mathfrak{l}}, \theta_{-\mathfrak{l}+1}, \dots, \theta_{\mathfrak{l}-1}, \theta_{\mathfrak{l}}\}, \forall t \in \mathbb{N}_0$: the parameter $\mathfrak{l}$ is arbitrary and has been chosen to match the abstraction parameter in Figure 1 and the scheme in (4). Based on (10), we set up the discrete-time switched stochastic system $\mathbf{X}(t+1) = F_{\sigma(t)}\mathbf{X}(t) + \mathbf{W}(t)$, where by switched model we mean that the state matrix $F_{\sigma(t)}$ takes values in

$$\{P^T(\theta_{-\mathfrak{l}}), P^T(\theta_{-\mathfrak{l}+1}), \dots, P^T(\theta_{\mathfrak{l}-1}), P^T(\theta_{\mathfrak{l}})\},$$

for all $t \in \mathbb{N}_0$, (cf. (10)), and the switching signal $\sigma(\cdot) : \mathbb{N}_0 \to \mathbb{Z}_{2\mathfrak{l}}$ is a map specifying the set-point θ_s, and hence the TCL dynamics, as a function of time. The process noise $\mathbf{W}(t)$ is normal with zero mean and state-dependent covariance matrix $\Sigma(\mathbf{X}(t))$ in (9). The total power consumption of the TCL population is measured as $y_{meas}(t) = H\mathbf{X}(t) + v(t)$, where $v(t) \sim \mathcal{N}(0, R_v)$ is a measurement noise characterized by $\sqrt{R_v}$, the standard deviation of the real-time measurement in the power meter instrument.

Since the process noise $\mathbf{W}$ is state-dependent, the state of the system can be estimated by a conditional Kalman filter with the following updates:

$$\hat{\mathbf{X}}^-(t+1) = F_{\sigma(t)}\hat{\mathbf{X}}(t),$$
$$P^-(t+1) = F_{\sigma(t)}\mathbb{P}(t)F^T_{\sigma(t)} + \Sigma(\hat{\mathbf{X}}(t)),$$

and the following measurement updates:

$$K_{t+1} = P^-(t+1)H^T\left[HP^-(t+1)H^T + R_v\right]^{-1},$$
$$\mathbb{P}(t+1) = [I - K_{t+1}H]P^-(t+1),$$
$$\hat{\mathbf{X}}(t+1) = \hat{\mathbf{X}}^-(t+1) + K_{t+1}[y_{meas}(t+1) - H\hat{\mathbf{X}}^-(t+1)].$$

When the state estimates $\hat{\mathbf{X}}$ are available, we formulate the following optimization problem based on a one-step output prediction, in order to synthesize the control input at the next step:

$$\min_{\sigma(t+1)\in\mathbb{Z}_{2l}} |y_{est}(t+2) - y_{des}(t+2)|, \text{ s.t.}$$
$$\hat{\mathbf{X}}(t+2) = F_{\sigma(t+1)}\hat{\mathbf{X}}(t+1),$$
$$y_{est}(t+2) = H\hat{\mathbf{X}}(t+2),$$

where $y_{des}(\cdot)$ is a desired reference signal, whereas $\hat{\mathbf{X}}(t+1)$ is provided by the conditional Kalman filter above. The obtained optimal value for $\sigma(t+1)$ provides the set-point $\theta_s(t+1)$, which is applied to the entire TCL population at the following $(t+1)^{\text{th}}$ iteration. Figure 2 illustrates the closed-loop configuration for state estimation and one-step regulation of the power consumption.

B. Regulation via Stochastic Model Predictive Control (SMPC)

We can perform power tracking by formulating and solving the following SMPC problem [19]:

$$\min_{\sigma(\tau)} J_t = \mathbb{E}\left[\sum_{\tau=t+1}^{T} [y_{abs}(\tau) - y_{des}(\tau)]^2 + \kappa^T\mathbf{X}(T)\middle|\mathbf{X}(t)\right],$$
$$\text{such that}$$
$$\mathbf{X}(\tau+1) = F_{\sigma(\tau)}\mathbf{X}(\tau) + \mathbf{W}(\tau), \quad y_{abs}(\tau) = H\mathbf{X}(\tau),$$
$$\sigma(\tau)\in\mathbb{Z}_{2l}, \quad \forall\tau\in\{t, t+1, \dots, T-1\}. \tag{25}$$

The conditional cost-to-go function J_t comprises a running cost for tracking and a terminal cost. The terminal cost is assumed to be a linear combination (with weighting vector κ) of the model states at final time T, and practically accounts for possible penalty weights over the number of TCLs within the temperature intervals. The expectation is taken over the underlying probability space for the trajectories of the process over the time interval $[t+1, T]$. The dynamics are nonlinear due to the switching nature of the control signal. The average evolution of the states and output of the system can be expressed by the following deterministic difference equations:

$$\mathbb{E}[\mathbf{X}(\tau+1)] = F_{\sigma(\tau)}\mathbb{E}[\mathbf{X}(\tau)], \quad \mathbb{E}[y_{abs}(\tau)] = H\mathbb{E}[\mathbf{X}(\tau)].$$

The associated state transition matrix $\Phi_\sigma(T,t) = F_{\sigma(T-1)}F_{\sigma(T-2)}\dots F_{\sigma(t)}$ provides a closed form for the average evolution over the interval $[t, T]$:

$$\mathbb{E}[\mathbf{X}(T)] = \Phi_\sigma(T,t)\mathbb{E}[\mathbf{X}(t)],$$
$$\mathbb{E}[y_{abs}(T)] = H\Phi_\sigma(T,t)\mathbb{E}[\mathbf{X}(t)].$$

Thanks to the linearly state-dependent covariance matrix, we can establish the following result.

Theorem 6. *The cost function of the SMPC problem can be computed explicitly as*

$$J_t = \sum_{\tau=t+1}^{T} \left[H\Phi_\sigma(\tau,t)\mathbf{X}(t) - y_{des}(\tau)\right]^2 + \Psi_\sigma(T,t)\mathbf{X}(t), \tag{26}$$

where the matrix

$$\Psi_\sigma(T,t) = \kappa^T\Phi_\sigma(T,t) + \frac{1}{n_p}\sum_{\tau_1=t}^{T}\sum_{\tau_2=\tau_1+1}^{T}\mathscr{R}(H\Phi_\sigma(\tau_2,\tau_1+1), F_{\sigma(\tau_1)})\Phi_\sigma(\tau_1,t),$$

and where $\mathscr{R}:\mathbb{R}^{1\times 2n}\times\mathbb{R}^{2n\times 2n}\to\mathbb{R}^{1\times 2n}$ is a matrix-valued map with $\mathscr{R}(C,D) = C^{\circ 2}D - (CD)^{\circ 2}$, where the operator $\circ 2$ is the Hadamard *square of the matrix (element-wise square).*

Proof. The proof is derived by computing the summation of J_t in (25) backwards, conditioning the expected value to the intermediate states, and utilizing the equality $\boldsymbol{\nu}^T\Sigma(\mathbf{X})\boldsymbol{\nu} = \frac{1}{n_p}\mathscr{R}(\boldsymbol{\nu}^T, P^T)\mathbf{X}$, for any $\boldsymbol{\nu}\in R^{2n}$. □

The obtained explicit cost function is the sum of a quadratic cost for the deterministic average evolution of the system state and of a linear cost related to the covariance of the process noise.

Remark 3. *For both formulations of the power tracking problem, the reference signal $y_{des}(\cdot)$ is assumed to be given. This is in practice the case when the TCL population is controlled to provide ancillary services. Moreover, this holds when a power utility company (or aggregator) participates in an energy market: it can observe the profile of the energy price, solve an optimization problem at a higher level minimizing the total energy cost based on an energy storage model, and thus obtain the power reference signal [25].*

Example 1. *The SMPC formulation can accommodate problems where the population participates in the energy market to minimize the energy costs. In the real-time energy market the Locational Marginal Pricing algorithms result in the profile of energy price for time intervals of 5-minutes [28]. Given that profile, the population can save money by minimizing the total cost of its energy usage within the given time frame, i.e. consuming less energy when the price is high and more energy when the price is low, under some constraints, in the next 24 hours. Suppose the final time T is selected such that $T = 24/h$, where h is the length of the sampling time (5 minutes), and let the sequence $\{\lambda_\tau, \tau = t+1, t+2, \dots, T\}$ be the profile of the energy price provided by the energy market. The total energy consumption of the population is then $\sum_{\tau=t+1}^{T}\lambda_\tau y_{abs}(\tau)h$. The following optimization problem can be solved, given the model dynamics, in order to minimize the expected value of the energy consumption:*

$$\min_{\sigma(\tau)} \mathbb{E}\left[\sum_{\tau=t+1}^{T}\lambda_\tau y_{abs}(\tau)h\middle|\mathbf{X}(t)\right] = \min_{\sigma(\tau)}\sum_{\tau=t+1}^{T}\lambda_\tau hH\Phi_\sigma(\tau,t)\mathbf{X}(t).$$

V. Numerical Case Study and Benchmarks

In this section we compare the performance of the aggregation procedure with that developed in [22], which as discussed obtains an aggregated model with dynamics that are deterministic, and has in fact been shown to be a special (limiting) case of the model in this work (cf. Remark 1). We further present the extension to the case of heterogeneous populations (with a comparison of the two proposed approaches), and synthesize global controls over the temperature set-point to perform tracking of the total power consumption of the population.

For all simulations we consider a population of $n_p = 500$ TCLs, however recall that our abstraction is proved to perform

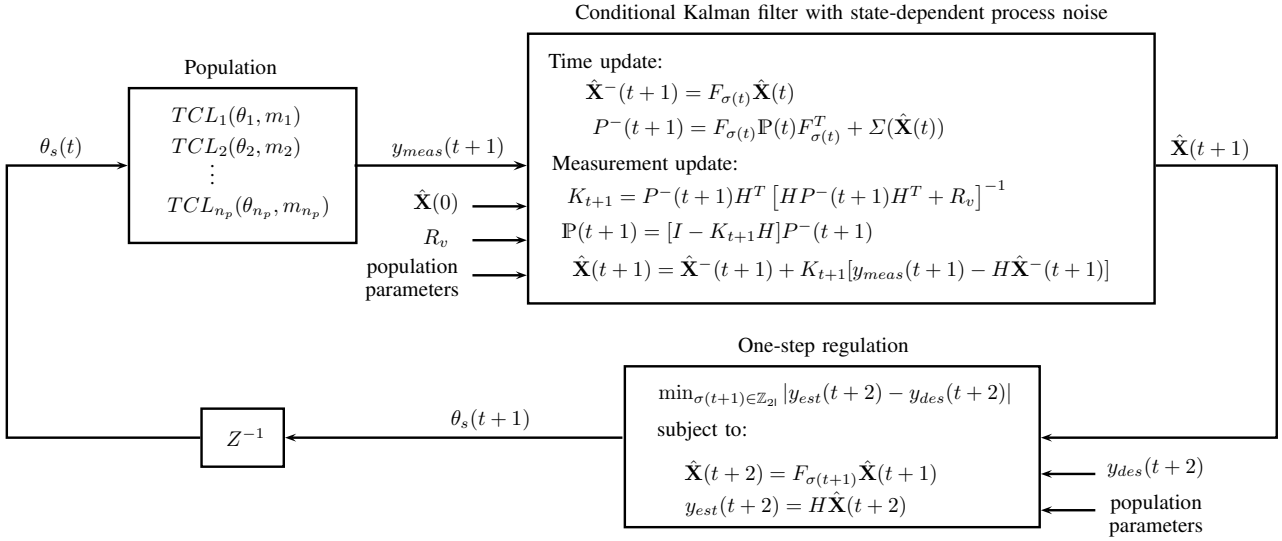


Fig. 2. State estimation and one-step regulation for the closed-loop control of the power consumption.

as desired for any value n_p of the population size. As a benchmark, we have run 50 Monte Carlo simulations for the TCL population based on the explicitly-aggregated dynamics in (1)-(2), and empirically computed the average total power consumption.

A. Aggregation of a Homogeneous Population of TCLs

We consider a population of TCLs comprised by air conditioners of residential kind (which affects the choice of parameters). Each TCL is characterized by parameters that take value as in Figure 3 (left). The TCLs are initialized with a temperature at the set-point ($\theta(0) = \theta_s$), half of them in the OFF mode ($m(0) = 0$), and the other half in the ON model ($m(0) = 1$). Unlike the deterministic dynamics considered in [22], the model in (1) includes a process noise: we select initially a small value for its standard deviation as $\sigma = 0.001\sqrt{h} = 0.0032$.

The abstraction in [22] is obtained by partitioning exclusively the dead-band and by "moving the probability mass" outside of this interval to the nearest bin in the opposite mode. Recall that in the new approach put forward in this work we need to provide a partition not only for the dead-band but for a larger range of temperatures (cf. Fig. 1). Sample trajectories of the TCL population are presented in Figure 3 (right): the second set of trajectories, obtained for a larger value of noise level, confirms that we need to partition a wider temperature range, rather than exclusively the dead-band. The abstraction in [22] depends on a parameter n_d, denoting the number of bins: we select $n_d = 70$, which leads to a total of 140 states. The selection of n_d has been steered by empirical tuning, targeted towards optimal performance; however, in general there seems to be no clear correspondence between the choice of n_d and the overall precision of the abstraction procedure in [22].

For the formal abstraction proposed in this work, we construct the partition as in (4) with parameters $\mathsf{l} = 70, \mathsf{m} = 350$, which leads to $2n = 1404$ abstract states. Notice that the presence of a small standard deviation σ for the process noise (not included in the dynamics of [22]) requires a smaller partition size to finely resolve the jumping probability between adjacent bins. Let us emphasize again that an increase in n_d for the method in [22] does not lead to an improvement of the outcomes.

The results obtained for a small noise level $\sigma = 0.0032$ are presented in Figure 4 (top). The aggregate power consumption has an oscillatory decay, since all thermostats are started in a single state bin (they share the same initial condition) and are thus "synchronized" at the outset. This outcome matches that presented in [22]: the deterministic abstraction[3] in [22] produces precise results for the first few (2-3) oscillations, after which the disagreement increases.

Let us now select a larger standard deviation for the process, to take the value $\sigma = 0.01\sqrt{h} = 0.032$, all other parameters being the same as before. We now employ $n_d = 5$ (obtained by empirical optimal tuning as per [22]), and $\mathsf{l} = 7$, and $\mathsf{m} = 35$, which leads to 10 and 144 abstract states, respectively. Figure 4 (bottom) presents the results of the experiment. It is clear that the model abstraction in [22] is not capable to generate a valid trajectory for the aggregate power, whereas the output of the formal abstraction proposed in this work nicely matches that of the average aggregated power consumption. Let us again remark that increasing the number of bins n_d does not seem to improve the performance of the deterministic abstraction in [22], but in this case rather renders the oscillations more evident. On the contrary, our approach allows the quantification of an explicit bound on the error made: for instance, the error on the normalized power consumption with parameters $N = 2$ and $\mathsf{l} = 70$ is equal to 0.226 (absolute quantity).

In order to better elucidate and distinguish the contributions in the literature, notice that the deterministic abstraction in [22] has been developed for a deterministic model of the TCLs (i.e. neglecting $w(t)$ in equation (1)), whereas the present approach is novel in that it provides an analytical derivation of an aggregated model given stochastic TCL dynamics that [22] did not consider; secondly, the new approach can handle larger noise values than [22]. As a final remark, let us emphasize that the outputs of both abstract models converge to steady-state values that may be slightly different from those

[3]Let us again remark that by "deterministic abstraction" we mean that the aggregate model $\mathbf{X}(t+1) = P^T\mathbf{X}(t)$ obtained in [22] is a deterministic equation. Notice however that the process noise is part of the temperature evolution of the TCLs.

Parameter	Interpretation	Value
θ_s	temperature set-point	$20\,[^\circ C]$
δ	dead-band width	$0.5\,[^\circ C]$
θ_a	ambient temperature	$32\,[^\circ C]$
R	thermal resistance	$2\,[^\circ C/kW]$
C	thermal capacitance	$10\,[kWh/^\circ C]$
P_{rate}	power	$14\,[kW]$
η	coefficient of performance	2.5
h	time step	$10\,[sec]$

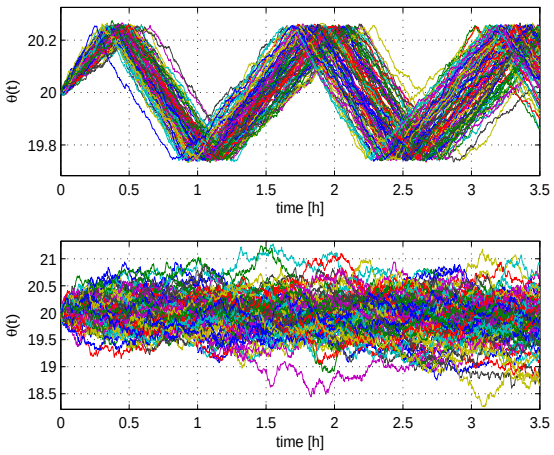


Fig. 3. Nominal values of parameters for residential air conditioners as from [10] (left), sample trajectories of the TCL population for two different values of the standard deviation of the process noise, $\sigma = 0.0032$ and $\sigma = 0.032$ (right).

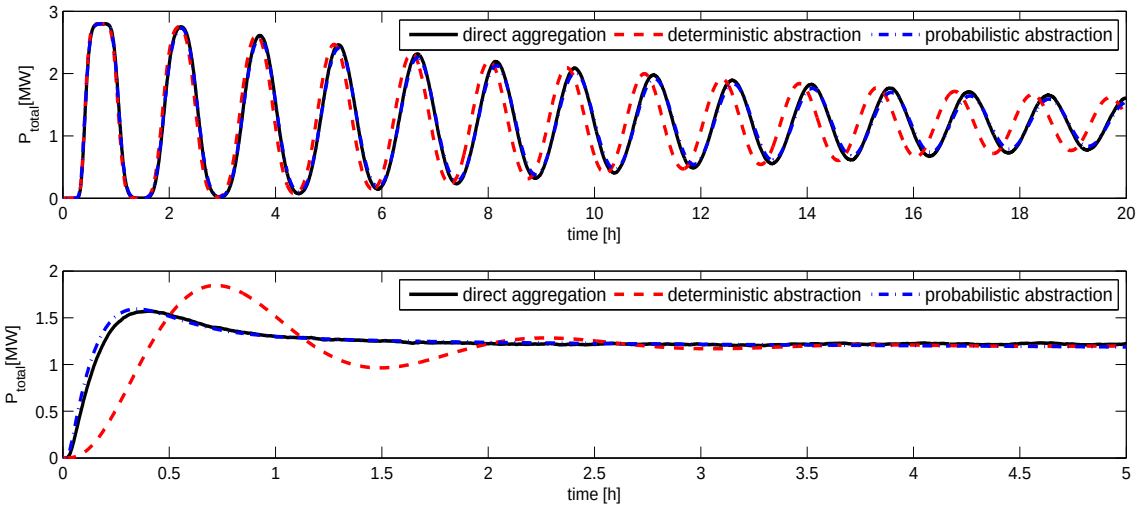


Fig. 4. Homogeneous population of TCLs. Comparison of the deterministic abstraction from [22] with the formal stochastic abstraction, for a small process noise $\sigma = 0.0032$ (top panel) and a larger value $\sigma = 0.032$ (bottom panel).

obtained as the average of the Monte Carlo simulations for the model aggregated directly. This discrepancy is due to the intrinsic errors introduced by both the abstraction procedures, which approximate a concrete continuous-space model (discontinuous stochastic difference equations) with discrete-space abstractions (finite-state Markov chains). However, whereas the abstraction in [22] does not offer an explicit quantification of the error, the formal abstraction proposed in this work does, and further allows the tuning (decrease) of such an error bound, by choice of a larger cardinality for the partitions set. However as a tradeoff, recall that increasing the number of partitions demands managing a Markov chain abstraction with a larger size.

B. Aggregation of a Heterogeneous Population of TCLs

In the rest of the simulations, unless otherwise stated we fix the standard deviation for the process noise to the larger value $\sigma = 0.032$. Let us assume that heterogeneity enters the TCL population over the thermal capacitance C of the TCLs, which is taken to be $C \sim \mathcal{U}([8, 12])$, that is described by a uniform distribution over a compact interval.

Monte Carlo simulations are performed with the discretization parameters $n_d = 6$ (deterministic abstraction), and $\mathsf{l} = 10, \mathsf{m} = 50$ (probabilistic abstraction via averaging). Figure 5 (top) compares the results of the two abstraction methods: the plots are quite similar to those for the homogeneous case, since the allowed range for the parameter is small. However, let us now increase the level of heterogeneity by enlarging the domain of definition of the thermal capacitance, so that $C \sim \mathcal{U}([2, 18])$. The (empirically) best possible deterministic abstraction is obtained by selecting $n_d = 7$, whereas we again select $\mathsf{l} = 10$, $\mathsf{m} = 50$ for the probabilistic abstraction based on averaging. The outcomes are presented in Figure 5 (bottom).

Figure 5 also compares the performance of the two abstraction approaches described respectively in Section III-A (via averaging) and in Section III-B (via clustering), for the two ranges of thermal capacitance ($[8, 12]$ and $[2, 18]$ respectively) characterizing the heterogeneity in the population. For the approach of Section III-B the population is clustered into

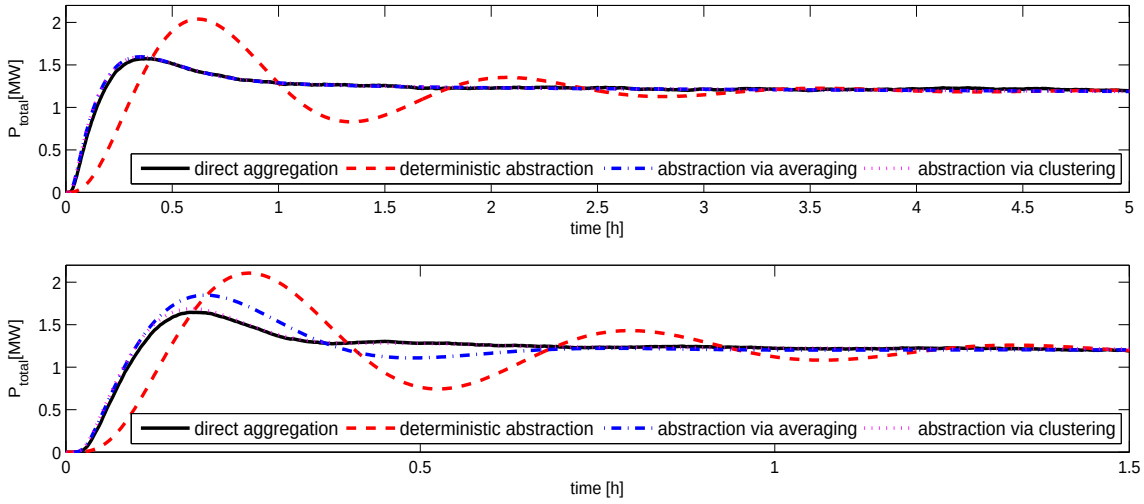


Fig. 5. Heterogeneous population of TCLs. Comparison of the deterministic abstraction from [22] with the formal probabilistic abstraction based on averaging, for two different ranges of the thermal capacitance: $[8, 12]$ (top panel) and $[2, 18]$ (bottom panel).

5 and 20 clusters, respectively. Figure 5 indicates that the performance of clustering approach surpasses that of the averaging approach. In conclusion, while the latter can be suitable for instances of small heterogeneity, the former is essential for large heterogeneity in the population.

C. Abstraction and Control of a Population of TCLs

With focus on the abstraction proposed in this work for a homogeneous population (again of $n_p = 500$ TCLs), the Kalman state estimation and one-step regulation scheme of Section IV-A is applied with the objective of tracking a randomly generated piecewise-constant reference signal. We have used discretization parameters $\mathsf{l} = 8$, $\mathsf{m} = 40$, and the standard deviation of the measurement ($\sqrt{R_v}$) has been chosen to be 0.5% of the total initial power consumption. Figure 6 displays the tracking outcome (top), as well as the required set-point signal synthesized by the above optimization problem (bottom). Notice that the set-point variation is bounded to within a small interval, which practically means that the temperature values (and as such the users) in the TCLs are unaffected by that.

A similar performance, as displayed in Figure 7, is obtained in the case of a heterogeneous population (again of 500 TCLs), where heterogeneity is characterized by the parameter $C \in \mathcal{U}([2, 18])$. The averaging approach of Section III-A and clustering approach of Section III-B with 5 clusters are employed for the abstraction of the population. While in general that the clustering approach can provides a more accurate model, its performance in this reference tracking example is quite similar to the averaging approach (their relative accumulated error is less than 10%). The total variation of the set-point signal is also the same for both approaches with the relative error less than 10%. This is because of the smaller state-space model of the averaging approach (the dimension is 5 times smaller) and its better performance in state estimation.

Finally, we have employed the SMPC scheme described in Section IV-B, combined with the Kalman state estimator of Section IV-A, to track a piecewise-constant reference signal over a homogeneous population of TCLs. A prediction horizon of $T - t = 5$ steps has been selected. The discrete nature of the optimization variable in (25) requires us, at each time step, to compute the cost function J_t for each sequence of $\sigma(\cdot)$ and to find the optimal one. In order to reduce computational burden of the optimization we introduce the following constraint on the variation of the set-point: $\left|\frac{d\theta_s}{dt}\right| \simeq \left|\frac{\theta_s(t+1)-\theta_s(t)}{h}\right| \leq \upsilon = 0.025$. Figure 8 presents the power consumption of the population (top) and the required set-point variation (bottom). The displayed response consists of a transient phase and of a steady state. It takes 2 minutes to reach the steady-state phase because of the limitations on the max rate of set-point changes. This can be seen from the plot of the set-point control signal, which first decreases and then increases within the transient phase with a constant rate. In order to obtain a faster transient phase, the upper bound for the set-point changes may be increased.

VI. Conclusions and Future Work

This work has put forward a formal approach for the abstraction of the dynamics of a TCL and the aggregation of a population model. The approach starts by partitioning the state space and constructing Markov chains for each TCL. Given the transition probability matrix of the Markov chains, it is possible to write down the state-space model of the population and further to aggregate it. This leads to a simple dynamical model that can be later analyzed. The article has discussed approaches to deal with models heterogeneity and to perform controller synthesis over the aggregated model. The article has derived a formal error bound for autonomous populations, which can be extended to controlled populations.

Looking forward, developing alternative approaches for the heterogeneous case, synthesizing new control schemes, and improving the error bounds are directions that are research-worthy in order to render the approach further applicable in practice.

References

[1] A. Abate, J.-P. Katoen, J. Lygeros, and M. Prandini. Approximate model checking of stochastic hybrid systems. *European Journal of Control*, 6:624–641, 2010.

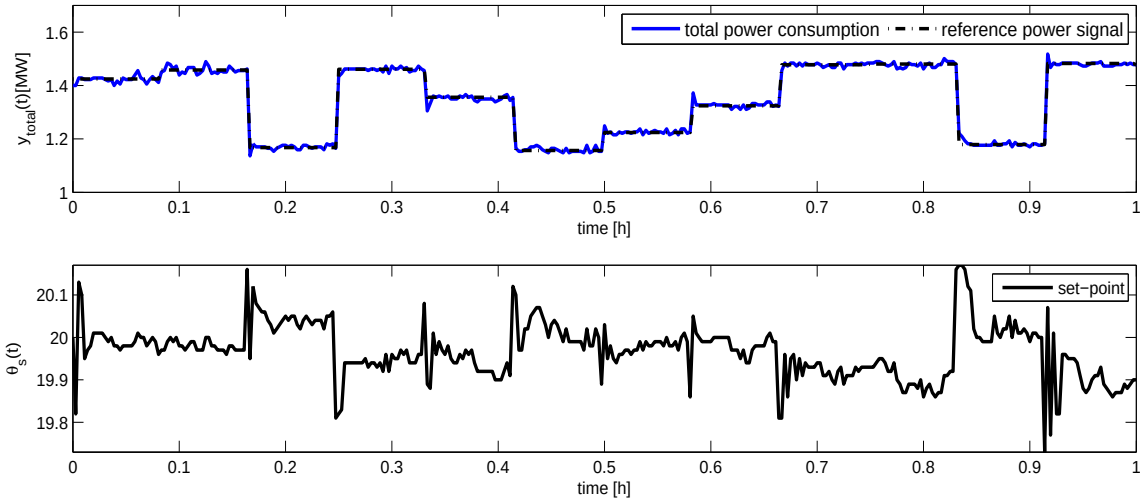


Fig. 6. Tracking of a piecewise-constant reference signal (top panel) by set-point control (bottom panel) in a homogeneous population of TCLs abstracted by the formal probabilistic approach.

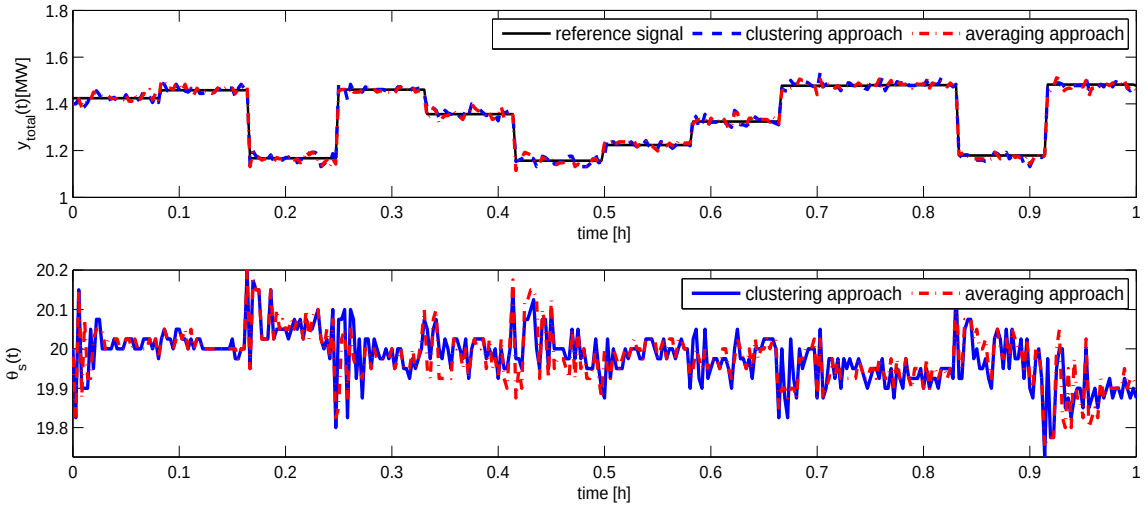


Fig. 7. Heterogeneous population of TCLs with $C \in \mathcal{U}([2, 18])$: tracking of a piecewise-constant reference signal (top) by set-point control (bottom), abstracted via clustering (5 clusters) and averaging.

[2] A. Abate, M. Prandini, J. Lygeros, and S. Sastry. Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica*, 44(11):2724–2734, November 2008.

[3] A.C. Antoulas. *Approximation of large-scale dynamical systems*. The Society for Industrial Mathematics, 2005.

[4] C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.

[5] S. Bashash and H.K. Fathy. Modeling and control insights into demand-side energy management through setpoint control of thermostatic loads. In *Proceedings of the 2011 American Control Conference*, pages 4546–4553, June 2011.

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

[7] P. Billingsley. *Probability and Measure - Third Edition*. Wiley Series in Probability and Mathematical Statistics, 1995.

[8] L. Bortolussi, J. Hillston, D. Latella, and M. Massink. Continuous approximation of collective system behaviour: A tutorial. *Performance Evaluation*, 70(5):317–349, 2013.

[9] J.-Y. Le Boudec, D. McDonald, and J. Mundinger. A generic mean field convergence result for systems of interacting objects. In *4th International Conference on the Quantitative Evaluation of SysTems (QEST)*, pages 3–18, 2007.

[10] D.S. Callaway. Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy. *Energy Conversion and Management*, 50(5):1389–1400, 2009.

[11] H.-F. Chen, P.R. Kumar, and J.H. van Schuppen. On Kalman filtering for conditionally Gaussian systems with random matrices. *Systems & Control Letters*, 13(5):397–404, 1989.

[12] C. Y. Chong and A.S. Debs. Statistical synthesis of power system functional load models. In *18th IEEE Conference on Decision and Control including the Symposium on Adaptive Processes*, volume 18, pages 264–269, 1979.

[13] J. Desharnais, F. Laviolette, and M. Tracol. Approximate analysis of probabilistic processes: logic, simulation and games. In *Proceedings of the International Conference on Quantitative Evaluation of Systems*, pages 264–273, Sept. 2008.

[14] S. Esmaeil Zadeh Soudjani and A. Abate. Adaptive gridding for abstraction and verification of stochastic hybrid systems. In *Proceedings of the 8th International Conference on Quantitative Evaluation of Systems*, pages 59–69, Sept. 2011.

[15] S. Esmaeil Zadeh Soudjani and A. Abate. Higher-Order Approximations for Verification of Stochastic Hybrid Systems. In S. Chakraborty and M. Mukund, editors, *Automated Technology for Verification and*

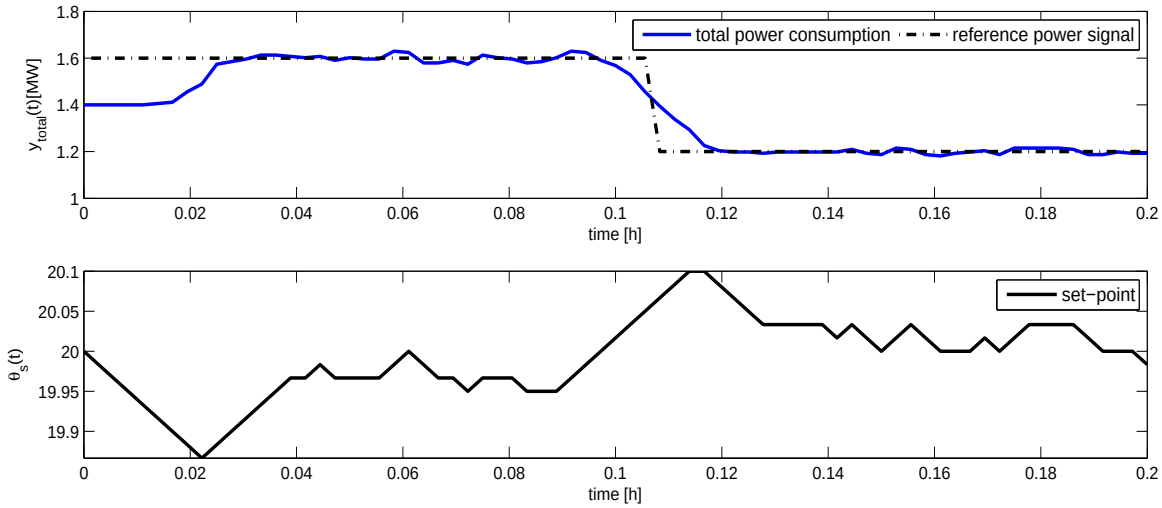


Fig. 8. Tracking of a piecewise-constant reference signal (top panel) by set-point control (bottom panel) for a homogeneous population of TCLs using the SMPC scheme.

Analysis, volume 7561 of *Lecture Notes in Computer Science*, pages 416–434. Springer Verlag, Berlin Heidelberg, 2012.

[16] S. Esmaeil Zadeh Soudjani and A. Abate. Adaptive and sequential gridding procedures for the abstraction and verification of stochastic processes. *SIAM Journal on Applied Dynamical Systems*, 12(2):921–956, 2013.

[17] S. Esmaeil Zadeh Soudjani and A. Abate. Aggregation of thermostatically controlled loads by formal abstractions. In *European Control Conference*, pages 4232–4237, Zurich, Switzerland, July 2013.

[18] S. Esmaeil Zadeh Soudjani and A. Abate. Probabilistic reach-avoid computation for partially-degenerate stochastic processes. *IEEE Transactions on Automatic Control*, 59(2):528–534, 2014.

[19] P. Hokayem, D. Chatterjee, and J. Lygeros. On stochastic receding horizon control with bounded control inputs. In *Proceedings of the 48th IEEE Conference on Decision and Control*, pages 6359–6364, Shanghai, PRC, December 2009.

[20] N. L. Johnson, S. Kotz, and N. Balakrishnan. *Discrete Multivariate Distributions*. Wiley Series in Probability and Statistics, 1997.

[21] M. Kamgarpour, C. Ellen, S. Esmaeil Zadeh Soudjani, S. Gerwinn, J.L. Mathieu, N. Mullner, A. Abate, D.S. Callaway, M. Fränzle, and J. Lygeros. Modeling options for demand side participation of thermostatically controlled loads. In *International Conference on Bulk Power System Dynamics and Control (IREP)*, pages 1–15, August 2013.

[22] S. Koch, J.L. Mathieu, and D.S. Callaway. Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services. In *17th Power Systems Computation Conference*, Stockholm, Sweden, August 2011.

[23] R. Malhame and C.-Y. 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.

[24] J.L. Mathieu and D.S. Callaway. State estimation and control of heterogeneous thermostatically controlled loads for load following. In *Hawaii International Conference on System Sciences*, pages 2002–2011, Hawaii, USA, 2012.

[25] J.L. Mathieu, M. Kamgarpour, J. Lygeros, and D.S. Callaway. Energy arbitrage with thermostatically controlled loads. In *2013 European Control Conference (ECC)*, pages 2519–2526, 2013.

[26] 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(1):430–440, 2013.

[27] R.E. Mortensen and K. P. Haggerty. A stochastic computer model for heating and cooling loads. *IEEE Transactions on Power Systems*, 3(3):1213–1219, 1988.

[28] S. Widergren, C. Marinovici, T. Berliner, and A. Graves. Real-time pricing demand response in operations. In *2012 IEEE General Meeting of the Power and Energy Society*, pages 1–5, 2012.

[29] W. Zhang, K. Kalsi, J. Fuller, M. Elizondo, and D. Chassin. Aggregate model for heterogeneous thermostatically controlled loads with demand response. In *IEEE PES General Meeting*, San Diego, CA, July 2012.

[30] W. Zhang, J. Lian, C.-Y. Chang, and K. Kalsi. Aggregated modeling and control of air conditioning loads for demand response. *IEEE Transactions on Power Systems*, 28(4):4655–4664, Nov 2013.