

Aggregation of Thermostatically Controlled Loads by Formal Abstractions

Sadegh Esmaeil Zadeh Soudjani and Alessandro Abate

Abstract—This work discusses a two-step procedure based on formal abstractions, which generates a finite stochastic dynamical model as an aggregation of the continuous temperature dynamics of a population of Thermostatically Controlled Loads (TCL). The temperature of each single TCL is described by a stochastic difference equation and the TCL status (ON, OFF) by a deterministic switching mechanism. The procedure is formal as it allows the exact quantification of the error introduced by the abstraction. The work discusses extensions to the case of controlled TCL. The procedure is tested on a case study and benchmarked against an alternative approach in the literature.

I. INTRODUCTION

Thermostatically Controlled Loads (TCL) have shown potential for the use in load balancing and regulation. Recent studies have focused on developing reliable models for aggregated populations of TCL. The goal of the work in [1] has been that of finding an aggregated model under homogeneity assumptions over the population, meaning that all TCL are assumed to have the same dynamical description. [1] puts forward a simple Linear Time Invariant (LTI) model for a population characterized by an input as the temperature set-point, and an output as the total consumed power. The parameters of the LTI model are estimated based on data and the model is used to track fluctuations in electricity generation from wind.

The work in [2] proposes an approach, based on the partitioning of the TCL temperature range, to obtain an aggregate state-space model for a population of TCL units that are heterogeneous in their thermal capacitance. The full information of the state variables of the model is used to synthesize a control strategy for output (total power) tracking via a deterministic Model Predictive Control scheme. The contribution in [3], [4], extends the results in [2] by considering a population of TCL that are heterogeneous in all their parameters. Furthermore, the Extended Kalman Filter is used to estimate the states and to identify the state transition matrix. The control of the population is performed by switching ON/OFF a portion of the TCL. Although the control strategy in [1] appears to be implementable over the current infrastructure with negligible costs, the model parameters are not directly related to the dynamics of the TCL population. On the other hand, the control methods proposed in [2], [3] may be impaired by the practical

implementation costs. Furthermore, the state-space models in [2], [3] are valid under two unrealistic assumptions: first, the temperature evolution is assumed to be deterministic, leading to a deterministic state space model; second, after partitioning the temperature range in separate bins, the TCL temperatures are assumed to be uniformly distributed within each state bin. Moreover, from a practical standpoint there seems to be no clear connection between the number of state-bins (width of temperature intervals) and the quality of the aggregated model: increasing the number of state bins (refining the model) does not necessarily improve the performance.

This article proposes a two-step abstraction procedure to generate a finite stochastic dynamical model as an aggregation of the dynamics of a population of TCL. The approach relaxes the assumptions in [2], [3] by providing a model based on the probabilistic evolution of the TCL temperature. The abstraction is made up of two separate parts: (1) going from continuous-space models for a TCL to finite state space models, which obtains a population of Markov chains; and (2) taking the cross product of the Markov chain models for the population and lumping the obtained model by finding its coarsest probabilistically bisimilar Markov chain [5]: as such the reduced-order Markov chain is an exact representation of the larger model. The approach is fully developed for the case of a homogeneous population of TCL, providing guaranteed error bounds for the abstraction.

The article also describes a dynamical model for the evolution of the abstraction, and shows convergence results as the population size grows: increasing the number of state bins improves the accuracy of the aggregated model, leading to a convergence of the error to zero. This result is aligned with the work [6] on the aggregation of continuous-time deterministic thermostatic loads. The analytic relation between model and population parameters enables the development of a set-point control strategy for reference tracking over the total load power. A modified version of the Kalman Filter is employed to estimate the states and the total power consumption of the population is regulated by a simple Model Predictive Control approach. The procedure is tested on a case study and benchmarked against the approach from [2], [3]. The promising outcomes encourage exploring extensions to the case of a population of heterogeneous TCL.

II. FORMAL ABSTRACTION OF A HOMOGENEOUS POPULATION OF TCL

Throughout this article we use the notation \mathbb{N} for natural numbers, $\mathbb{Z} = \mathbb{N} \cup \{0\}$, $\mathbb{N}_n = \{1, 2, 3, \dots, n\}$, and $\mathbb{Z}_n =$

This work is supported by the European Commission MoVeS project FP7-ICT-2009-257005, by the European Commission Marie Curie grant MANTRAS 249295, by the European Commission IAPP project AMBI 324432, and by the NWO VENI grant 016.103.020.

The authors are with the Delft Center for Systems and Control, TU Delft – Delft University of Technology, The Netherlands - {S.EsmaeilZadehSoudjani,A.Abate}@tudelft.nl

$\mathbb{N}_n \cup \{0\}$. We denote vectors with bold typeset and with a letter corresponding to that of its elements.

The evolution of the temperature in a single TCL can be characterized by the following stochastic difference equation [1], [7]

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

where θ_a is the ambient temperature, C and R indicate the thermal capacitance and resistance respectively, P_{rate} is the rate of energy transfer, and $a = e^{-h/RC}$, with a discretization step h . The process noises $w(t), t \in \mathbb{Z}$, are i.i.d. and characterized by a density function $t_w(\cdot)$. In equation (1) a $+$ sign is used for a heating TCL, whereas a $-$ sign for a cooling TCL. 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. The distribution of the initial mode and temperature is denoted by $\pi_0(m, \theta)$. The temperature dynamics is regulated by the discrete switching control $m(t+1) = f(m(t), \theta(t))$, where

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

where θ_s denotes the temperature set-point and δ its dead-band, and together characterize the temperature range.

The power consumption of the TCL at time t is equal to $\frac{1}{\eta}m(t)P_{rate}$, which is zero in the OFF mode and positive in the ON mode. The parameter η is the Coefficient Of Performance (COP). The constant $\frac{1}{\eta}P_{rate}$ is the power consumption of the TCL in the ON mode, which is indicated in the sequel by $P_{rate, ON}$.

The composition of the dynamical equation in (1) with the algebraic relation in (2) allows us to consider a single TCL as a Stochastic Hybrid System (SHS) [8], namely as a discrete-time Markov process evolving over a hybrid state space characterized by a variable $s = (m, \theta) \in \mathbb{Z}_1 \times \mathbb{R}$.

The interpretation as a SHS leads to leverage an abstraction technique first proposed in [9], aimed at reducing a discrete-time, uncountable state-space Markov process into a (discrete-time) finite-state Markov chain. This abstraction is based on state-space partitioning 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

$$P(s_{im}, s_{i'm'}) = \delta[m' - f(m, \bar{\theta}_i)] \cdot \int_{\Theta_{i'}} t_w(\bar{\theta} - a\bar{\theta}_i - (1-a)(\theta_a \pm mRP_{rate})) d\bar{\theta},$$

where again $m' \in \mathbb{Z}_1, i' \in \mathbb{N}_n$, and where $\delta[\cdot]$ denotes the discrete unit impulse function. The initial probability mass for \mathcal{M} is obtained as

$$p_0(s_{im}) = \int_{\Theta_i} \pi_0(m, \theta) d\theta.$$




Fig. 1. Partition of the temperature axis for the abstraction of a TCL.

For the ease 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 introduce the new notation

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

Notice that the conditional density function of the stochastic system describing the dynamics of a single TCL is discontinuous: this is due to the presence of equation (2), which can be alternatively expressed via the following discrete conditional distribution, for all $m, \bar{m} \in \mathbb{Z}_1, \theta \in \mathbb{R}$:

$$t(\bar{m}|\theta, m) = m(1 - \bar{m})\mathbb{I}_{(-\infty, \theta_-)} + m\bar{m}\mathbb{I}_{(\theta_-, \infty)} + (1 - m)(1 - \bar{m})\mathbb{I}_{(-\infty, \theta_+]} + (1 - m)\bar{m}\mathbb{I}_{(\theta_+, \infty)},$$

where $\mathbb{I}_{\mathcal{A}}$ denotes the indicator function of the set \mathcal{A} . The selection of the partitioning sets then requires special attention: it is convenient to select a partition for the dead-band $[\theta_-, \theta_+]$, thereafter extending it to a partition over the whole line \mathbb{R} (cf. Figure 1). Let us select two constants $l, m \in \mathbb{N}, l < m$, compute the partition size $\tau = \delta/2l$ and quantity $\mathcal{L} = 2m\tau$. Now construct the boundary points of the partition sets $\{\theta_i\}_{i=-m}^m$ for the temperature axis as

$$\begin{aligned} \theta_{\pm 1} &= \theta_s \pm \delta/2, & \theta_{\pm m} &= \theta_s \pm \mathcal{L}/2, & \theta_{i+1} &= \theta_i + \tau, \\ \mathbb{R} &= \cup_{i=1}^n \Theta_i, & \Theta_1 &= (-\infty, \theta_{-m}), & \Theta_n &= [\theta_m, \infty), \\ \Theta_{i+1} &= [\theta_{-m+i-1}, \theta_{-m+i}), & i &\in \mathbb{N}_{n-2}, & n &= 2m + 2, \end{aligned} \quad (3)$$

and let us render the Markov states of the infinite-length intervals Θ_1, Θ_n absorbing.

In the next section we assess the discretization error of the abstraction of the dynamics in (1)-(2) introduced by the partition of the state space as a function of the partition size τ and of the quantity \mathcal{L} . This guarantees the convergence (in expected value) of the error between the power consumption of the model and the actual one for the entire population [9].

Consider now a population of n_p homogeneous TCL, that is a population of TCL which, after possible rescaling of (1)-(2), share the same set of parameters $\theta_s, \delta, \theta_a, C, R, P_{rate}, \eta, h$, and the noise term $t_w(\cdot)$. We focus on a population of cooling TCL, with the understanding that the case of heating TCL can be similarly obtained. Each TCL can be abstracted as a Markov chain 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 \mathcal{M} . Still, notice that the initial probability mass vector $p_0 = [p_{0i}], i \in \mathbb{N}_{2n}$, may vary over the population.

The homogeneous TCL population can be represented by a Markov chain Ξ , built as the cross product of the n_p Markov chains. The Markov chain Ξ has state

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

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. As the second step of the abstraction procedure, we are interested in further aggregating this model: we employ the notion of (exact) probabilistic bisimulation to achieve this [5]. Let us introduce a finite set of atomic propositions as a constrained vector with a dimension corresponding to the number of states of the single \mathcal{M} :

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

The labeling function $L : \mathcal{Z} \rightarrow AP$ 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} on the state space of \mathcal{Z} [5], such that

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

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 of its elements. We plan to show that \mathcal{R} is an exact probabilistic bisimulation relation on Ξ [5], which requires proving that, for any set $T \in \mathcal{Z}/\mathcal{R}$ and any pair $(\mathbf{z}, \mathbf{z}') \in \mathcal{R}$

$$P_{\Xi}(\mathbf{z}, T) = P_{\Xi}(\mathbf{z}', T), \quad (4)$$

This is achieved by Corollary 1 in the next Section. We now focus on the stochastic properties of \mathcal{R} , and in particular of Ξ and of its quotient Markov chain.

III. PROPERTIES OF THE QUOTIENT MARKOV CHAIN

Let us recall the definition of a few known discrete random variables that are used to describe quantities in the abstraction procedure. The sum of independent Bernoulli trials characterized the same success probability follows a *binomial distribution* [10]. If a random variable Y is instead defined the sum of n_p independent Bernoulli trials with success probabilities p_1, p_2, \dots, p_{n_p} , then Y has a *Poisson-binomial distribution* [11] with the sample space \mathbb{Z}_{n_p} and the following mean and variance: $\mathbb{E}[Y] = \sum_{r=1}^{n_p} p_r$, $\text{var}(Y) = \sum_{r=1}^{n_p} p_r(1 - p_r)$. As a generalization of the Bernoulli trials, a *categorical distribution* describes the result of a random event that takes on one of $n > 2$ possible outcomes. Its sample space is taken to be \mathbb{N}_n and its probability mass function $\mathbf{p} = [p_1, p_2, \dots, p_n]$, such that $\sum_{i=1}^n p_i = 1$. A *multinomial distribution* is a generalization of the binomial distribution as the sum of categorical random variables with the same parameters. The sum of categorical random variables with different parameters follows instead the *generalized multinomial distribution*, defined as follows. Consider n_p independent categorical random variables defined over the same sample space \mathbb{N}_n but with different

outcome probabilities $\mathbf{p}_r = [p_{r1}, p_{r2}, \dots, p_{rn}]$, $r \in \mathbb{N}_{n_p}$. Let the random variable Y_i indicate the number of times the i^{th} outcome is observed over n_p samples. Then vector $\mathbf{Y} = [Y_1, \dots, Y_n]^T$ has a generalized multinomial distribution characterized by

$$\begin{aligned} \mathbb{E}[Y_i] &= \sum_{r=1}^{n_p} p_{ri}, & \text{var}(Y_i) &= \sum_{r=1}^{n_p} p_{ri}(1 - p_{ri}), \\ \text{cov}(Y_i, Y_j) &= - \sum_{r=1}^{n_p} p_{ri}p_{rj} \quad (i \neq j). \end{aligned}$$

Back to the abstraction procedure, we now study the one-step probability mass function for one of the labels, conditional on the state of Ξ .

Theorem 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:

$$\begin{aligned} \mathbb{E}[x_i(t + 1)|\mathbf{z}(t)] &= \sum_{r=1}^{n_p} P_{z_r(t)i}, \\ \text{var}(x_i(t + 1)|\mathbf{z}(t)) &= \sum_{r=1}^{n_p} P_{z_r(t)i}(1 - P_{z_r(t)i}). \end{aligned} \quad (5)$$

Conditional on an observation $\mathbf{x} = [x_1, x_2, \dots, x_{2n}]^T$ 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 $P(x_i(t + 1) = j|\mathbf{x}(t))$, for any $j \in \mathbb{Z}_{n_p}$. Notice that for a label $\mathbf{x} = [x_1, \dots, x_{2n}]^T$ there are exactly $n_p! / (x_1!x_2! \dots x_{2n}!)$ states of Ξ such that $L(\mathbf{z}) = \mathbf{x}$. We use the notation $\mathbf{z} \rightarrow \mathbf{x}$ to indicate the states in Ξ associated to label \mathbf{x} , that is $\mathbf{z} : L(\mathbf{z}) = \mathbf{x}$.

Based on the law of total probability for conditional probabilities, we can write

$$\begin{aligned} P(x_i(t + 1) = j|\mathbf{x}(t)) &= \frac{\sum_{\mathbf{z}(t) \rightarrow \mathbf{x}(t)} P(x_i(t + 1) = j|\mathbf{z}(t))P(\mathbf{z}(t))}{P(\mathbf{x}(t))} \\ &= P(x_i(t + 1) = j|\mathbf{z}(t)) \frac{\sum_{\mathbf{z}(t) \rightarrow \mathbf{x}(t)} P(\mathbf{z}(t))}{P(\mathbf{x}(t))} \\ &= P(x_i(t + 1) = j|\mathbf{z}(t)), \end{aligned} \quad (6)$$

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 has been possible since the probability of having a label $\mathbf{x} = (x_1, x_2, \dots, x_{2n})$ is exactly the sum of the probabilities of the states \mathbf{z} generating such a label. This further allows expressing the quantities in (5) 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 2: The conditional random variables $(x_i(t + 1)|\mathbf{x}(t))$ have Poisson-binomial distributions, whereas the conditional random vector $(\mathbf{x}(t + 1)|\mathbf{x}(t))$ has a generalized

multinomial distribution. Their mean, variance, and covariance are characterized by

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

for all $i, j \in \mathbb{N}_{2n}, i \neq j$, ■

Theorem 2 indicates that the distribution of the conditional random variable $(\mathbf{x}(t+1)|\mathbf{x}(t))$ is independent of the underlying state $\mathbf{z}(t) \rightarrow \mathbf{x}(t)$ of \mathcal{E} . With focus on equation (4), this result allows to claim the following.

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

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 with the following parameters, for all $i, j \in \mathbb{N}_{2n}, i \neq j$:

$$\begin{aligned}\mathbb{E}[X_i(t+1)|\mathbf{X}(t)] &= \sum_{r=1}^{2n} X_r(t)P_{ri}, \\ \text{var}(X_i(t+1)|\mathbf{X}(t)) &= \frac{1}{n_p} \sum_{r=1}^{2n} X_r(t)P_{ri}(1-P_{ri}), \\ \text{cov}(X_i(t+1), X_j(t+1)|\mathbf{X}(t)) &= -\frac{1}{n_p} \sum_{r=1}^{2n} X_r(t)P_{ri}P_{rj}.\end{aligned}\quad (7)$$

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

$$\begin{cases} \omega_1(t) = X_1(t+1) - \sum_{r=1}^{2n} X_r(t)P_{r1} \\ \omega_2(t) = X_2(t+1) - \sum_{r=1}^{2n} X_r(t)P_{r2} \\ \vdots \\ \omega_{2n}(t) = X_{2n}(t+1) - \sum_{r=1}^{2n} X_r(t)P_{r2n}, \end{cases}$$

where $\omega_i(t)$ are guaranteed to be (dependent) random variables with zero mean and covariance described by the matrix with elements in (7). Such a translation allows expressing the following dynamical model for the variable \mathbf{X} :

$$\mathbf{X}(t+1) = P^T \mathbf{X}(t) + \mathbf{W}(t), \quad (8)$$

where the distribution of $\mathbf{W}(t)$ depends only on the state $\mathbf{X}(t)$.

Remark 1: We have modeled the evolution of the TCL population with an abstract model (8) based on linear stochastic difference equations. The dynamics in (8) represent a direct generalization of the model abstraction provided in [2], [3], which is deterministic since its transitions are based on the trajectories of a deterministic version of (1). ■

Based on Theorem 2 we have characterized the random variable $(X_i(t+1)|\mathbf{X}(t))$ with a Poisson-binomial distribution. We use the Lyapunov central limit theorem [10] to show that this distribution converges to a Gaussian one.

Theorem 3: The random variable $(X_i(t+1)|\mathbf{X}(t))$ can be explicitly expressed as

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

where the random variables $\omega_i(t)$ converge (in distribution) as $n_p \uparrow \infty$ to the Gaussian random variables $\omega_i(t) \sim \mathcal{N}(0, \sigma_i^2(\mathbf{X}(t)))$, where $\sigma_i^2(\mathbf{X}) = \frac{1}{n_p} \sum_{r=1}^{2n} X_r P_{ri}(1-P_{ri})$. ■

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 $(m_i, \theta_i)(t), i \in \mathbb{N}_{n_p}$, is

$$y_{total}(t) = \sum_{i=1}^{n_p} m_i(t)P_{rate, ON}.$$

With focus on the abstract model (described in terms of the normalized variable \mathbf{X}), the power consumption is equal to

$$y_a(t) = H\mathbf{X}(t), \quad H = n_p P_{rate, ON} [0_n, \mathbf{1}_n],$$

where $0_n, \mathbf{1}_n$ are row vectors with all the entries equal to zero and one, respectively. The following theorem quantifies the abstraction error over the total power consumption.

Theorem 4: Consider a homogeneous population of TCL with Gaussian process noise $w(\cdot) \sim \mathcal{N}(0, \sigma^2)$ and the abstracted model constructed based on the partition in (3). The difference in the expected value of the total power consumption of the population $y_{total}(t)$, and that of the abstracted model $y_a(t)$, both conditional on the corresponding initial conditions, is upper bounded by

$$\begin{aligned}|\mathbb{E}[y_{total}(t)|\mathbf{s}_0] - \mathbb{E}[y_a(t)|\mathbf{X}_0]| \\ \leq n_p(t-1)P_{rate, ON} \left[\frac{(t-2)}{2} \epsilon + \frac{2a}{\sigma\sqrt{2\pi}} \tau \right],\end{aligned}$$

for all $\mathbf{s}_0 \in (\mathbb{Z}_1 \times [\theta_{-m}, \theta_m])^{n_p}$, and where $\epsilon = \frac{e^{-\gamma^2/2}}{\gamma\sqrt{2\pi}}$, $\gamma = \frac{1-a}{2\sigma} \left[\frac{\mathcal{L}a^t + \delta}{1-a^t} - RP_{rate} - |2(\theta_s - \theta_a) + RP_{rate}| \right]$.

The initial state \mathbf{X}_0 is a function of the initial states in the population of TCL \mathbf{s}_0 , as from the definition of the state vector \mathbf{X} . ■

Notice that this result allows tuning the error made in estimating the total power consumption of the population from the abstraction.

IV. NUMERICAL BENCHMARKS

In this section we compare the performance of the presented abstraction with that developed in [2], which as discussed obtains an aggregated model with dynamics that are deterministic, and in fact shown to be a special (limiting) case of the model obtained in this work.

A TCL population size of $n_p = 500$ is considered for all the simulations (though the performance of our model can be tuned for arbitrary values of the population size). Each TCL is characterized by parameters that take values in Table I [1]. All TCL are initialized in the OFF mode ($m(0) = 0$) and with a temperature at the set-point ($\theta(0) = \theta_s$).

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]

TABLE I
PARAMETERS FOR SIMULATION OF THE TCL POPULATION [1].




Fig. 2. Sample trajectories of the TCL population for two different values of the standard deviation of the process noise ($\sigma = 0.0032$ and $\sigma = 0.032$).

Unlike the deterministic dynamics considered in [2], the model in (1) encompasses process noise: we select a very small standard deviation $\sigma = 0.001\sqrt{h} = 0.0032$.

The abstraction in [2] is obtained by partitioning the dead-band and shifting the probability masses outside of this interval to the next bin in the opposite mode. Notice that in the approach put forward in this work because of the presence of noisy dynamics we need to provide a partition not only over the dead-band but for the entire temperature range. The abstraction in [2] depends on a parameter n_d for the number of bins: we select a $n_d = 70$, which leads to 140 states. This selection is based on empirical tuning targeted toward optimal performance – however, there seems to be no clear correspondence between the choice of n_d and the overall precision in [2].

For the formal abstraction proposed in this work, we construct a partition as per (3) with $l = 70$, $m = 350$, which leads to $2n = 1404$ abstract states. The presence of a small standard deviation σ for the process noise requires a smaller partition size to resolve the probability of jumps between different bins.

We run 50 Monte Carlo simulations for the TCL population based on the explicitly aggregated dynamics in (1)-(2) and compute the average total power consumption. Sample trajectories of the TCL population are presented in Figure 2: the second plot, obtained for a larger value of noise level, confirms that we need to partition the whole temperature range, rather than exclusively the dead-band.

The results obtained for a noise level $\sigma = 0.0032$ are presented in Figure 3. The aggregate power consumption has an oscillatory decay because all thermostats are started in a single state bin. This outcome matches that presented in [2]: the deterministic abstraction in [2] produces precise results for the first few (2-3) oscillations, after which the disagreement between models over aggregate power increases.




Fig. 3. Comparison of the deterministic abstraction from [2] with the formal stochastic abstraction in this work, for a small process noise $\sigma = 0.0032$ and two different time scales.




Fig. 4. Comparison of the deterministic abstraction from [2] with the formal stochastic abstraction in this work, for a larger process noise $\sigma = 0.032$.

Let us increase the standard deviation of the process noise to a larger value $\sigma = 0.01\sqrt{h} = 0.032$, all other parameters being the same. We now employ $n_d = 5$, $l = 7$, and $m = 35$, which leads to 10 and 144 abstract states, respectively. We also run 50 Monte Carlo simulations for the explicitly aggregated TCL population.

Figure 4 presents the results of the experiment. It is clear that the model abstraction from [2] is not able to generate a good trajectory for the aggregate power, whereas the output of the formal abstraction nicely matches that of the average aggregated power consumption. Let us again remark that increasing number of bins n_d does not seem to improve the performance of the deterministic abstraction, but renders the oscillations more evident.

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 that obtained from the average Monte Carlo simulations. This discrepancy is due to the intrinsic errors in the abstraction procedures, which approximate a concrete continuous model (discontinuous stochastic difference equation) with a discrete abstraction (a finite state Markov chain). However, whereas the abstraction in [2] 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 error bound, depending on the choice of the (larger) cardinality of the partitions set. As a tradeoff, increasing number of partition sets demands handling an abstraction with a larger size.

V. EXTENSION TO NON-AUTONOMOUS TCL MODELS

Among the different strategies for controlling the total power consumption of a population of TCL, we consider

the case where the control input is the set-point θ_s of the TCL [1]. Recall that we intend to apply the control input to all TCL uniformly, in order to obtain a homogeneous population of TCL since this requires no a-priori knowledge of the state of the single TCL. This is unlike [2], which consider the control signal as an external input that is applied based on the knowledge of states of the single TCL: this requires adding thermometers (with relatively high accuracy) to the TCL. More precisely, [2] assumes full knowledge of the state vector $\mathbf{X}(t)$ and employs a Model Predictive Control architecture to design the control signal. Instead [3] considers different scenarios ranging from measuring all the states to a subset of them, and implements the extended Kalman Filter to identify the model parameters, together with estimating the states of the model – the approach seems to break down when the number of states becomes large.

In the following we show that the knowledge of the actual values of the states in $\mathbf{X}(t)$ is not necessary. Given the model parameters, all is needed is an online measurement of the total power consumption of the TCL population, which allows estimating the states in $\mathbf{X}(t)$ and using the set-point θ_s to track any reference signal based on a one-step output prediction.

Suppose we have a homogeneous population of TCL with known parameters. Based on (8), we set up the model $\mathbf{X}(t+1) = F(\theta_s(t))\mathbf{X}(t) + \mathbf{W}(t)$, where $\theta_s(t)$, the set-point value at time t , is the control input for the model, and matrix $F = P^T$. We assume that the control input is discrete and take values from a set: $\theta_s(t) \in \{\theta_{-1}, \theta_{-1+1}, \dots, \theta_{l-1}, \theta_l\}$, $\forall t \in \mathbb{Z}$. This assumption makes it possible to use the partition in (3) at all time steps. The process noise $\mathbf{W}(t)$ is normal with zero mean and the state-dependent covariance matrix in (7), which is denoted by $\Sigma(\mathbf{X}(t))$. The total power consumption of the TCL population is measured as $y_m(t) = H\mathbf{X}(t) + v(t)$, where $v(t) \sim \mathcal{N}(0, R_v)$ is a measurement noise and $\sqrt{R_v}$ represents a standard deviation depending on the real-time measurements from power meters.

Since the process noise $\mathbf{W}(t)$ is state-dependent, the state of the system can be estimated by modifying the classical Kalman Filter with the following time update:

$$\begin{aligned}\hat{\mathbf{X}}^-(t+1) &= F(\theta_s(t))\hat{\mathbf{X}}(t), \\ P^-(t+1) &= F(\theta_s(t))P(t)F(\theta_s(t))^T + \Sigma(\hat{\mathbf{X}}(t)),\end{aligned}$$

and the following measurement update:

$$\begin{aligned}K_{t+1} &= P^-(t+1)H^T [HP^-(t+1)H^T + R_v]^{-1}, \\ P(t+1) &= [I - K_{t+1}H]P^-(t+1), \\ \hat{\mathbf{X}}(t+1) &= \hat{\mathbf{X}}^-(t+1) + K_{t+1}[y_m(t+1) - H\hat{\mathbf{X}}^-(t+1)].\end{aligned}$$

When the state estimates are available, the following straightforward one-step Model Predictive Control scheme is employed to synthesize the next control input:

$$\begin{aligned}\min_{\theta_s(t+1)} & |\hat{y}(t+2) - y_d(t+2)|, \text{ s.t.} \\ \hat{\mathbf{X}}(t+2) &= F(\theta_s(t+1))\hat{\mathbf{X}}(t+1) \\ \hat{y}(t+2) &= H\hat{\mathbf{X}}(t+2) \\ \theta_s(t+1) &\in \{\theta_{-1}, \theta_{-1+1}, \dots, \theta_{l-1}, \theta_l\},\end{aligned}$$




Fig. 5. Tracking a piece-wise constant reference signal (top) by set-point control (bottom) in a homogeneous population of TCL.

where $y_d(\cdot)$ is the reference signal and $\hat{\mathbf{X}}(t+1)$ is provided by Kalman Filter. The obtained optimal value for $\theta_s(t+1)$ is applied to the TCL population at the following iteration.

The above scheme is implemented on a homogeneous population of $n_p = 500$ TCL, for tracking a randomly generated piece-wise constant reference signal. We have used a discretization parameters $l = 8$, $m = 40$, where the standard deviation of the measurement is 0.5%. Figure 5 displays the tracking outcome (top), as well as the required set-point signal synthesized from the above optimization problem.

VI. CONCLUSIONS AND FUTURE WORK

This work has studied a formal approach for the aggregation of the dynamics of a homogeneous population of TCL, leading to their control. The authors are currently focusing on the extension of the results to the case of a population of heterogeneous TCL.

REFERENCES

- [1] D. Callaway, "Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy," *Energy Conversion and Management*, vol. 50, no. 5, pp. 1389–1400, 2009.
- [2] S. Koch, J. Mathieu, and D. Callaway, "Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services," in *Proceedings of the 17th Power Systems Computation Conference*, Stockholm, Sweden, 2011.
- [3] J. Mathieu and D. Callaway, "State estimation and control of heterogeneous thermostatically controlled loads for load following," in *Hawaii International Conference on System Sciences*, Hawaii, USA, 2012, pp. 2002–2011.
- [4] J. Mathieu, S. Koch, and D. Callaway, "State estimation and control of electric loads to manage real-time energy imbalance," *IEEE Transactions on Power Systems* (accepted), 2012.
- [5] C. Baier and J.-P. Katoen, *Principles of Model Checking*. MIT Press, 2008.
- [6] S. Bashash and H. Fathy, "Modeling and control insights into demand-side energy management through setpoint control of thermostatic loads," in *Proceedings of the 2011 American Control Conference*, San Francisco, CA, June 2011, pp. 4546–4553.
- [7] 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*, vol. 30, no. 9, pp. 854–860, 1985.
- [8] A. Abate, M. Prandini, J. Lygeros, and S. Sastry, "Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems," *Automatica*, vol. 44, no. 11, pp. 2724–2734, November 2008.
- [9] A. Abate, J.-P. Katoen, J. Lygeros, and M. Prandini, "Approximate model checking of stochastic hybrid systems," *European Journal of Control*, vol. 6, pp. 624–641, 2010.
- [10] P. Billingsley, *Probability and Measure - Third Edition*. Wiley Series in Probability and Mathematical Statistics, 1995.
- [11] Y. Wang, "On the number of successes in independent trials," *Statistica Sinica*, vol. 3, no. 2, pp. 295–312, 1993.