Formal Synthesis and Validation of Inhomogeneous Thermostatically Controlled Loads

Sadegh Esmaeil Zadeh Soudjani[1], Sebastian Gerwinn[2], Christian Ellen[2], Martin Fränzle[2], and Alessandro Abate[3,1]

[1] Delft Center for Systems and Control, TU Delft, The Netherlands
[2] OFFIS, Universität Oldenburg, Germany
[3] Department of Computer Science, University of Oxford, United Kingdom
S.EsmaeilZadehSoudjani@tudelft.nl

Abstract. This work discusses the construction of a finite-space stochastic dynamical model as the aggregation of the continuous temperature dynamics of an inhomogeneous population of thermostatically controlled loads (TCLs). The temperature dynamics of a TCL is characterized by a differential equation in which the TCL status (ON, OFF) is controlled by a thresholding mechanism, and which displays inhomogeneity as its thermal resistance changes in time according to a Poisson process. In the aggregation procedure, each TCL model in the population is formally abstracted as a Markov chain, and the cross product of these Markov chains is lumped into its coarsest (exact) probabilistic bisimulation. Quite importantly, the abstraction procedure allows for the quantification of the induced error. Assuming that the TCLs explicitly depend on a control input, the contribution investigates the problem of population-level power reference tracking and load balancing. Furthermore, for the corresponding closed-loop control scheme we show how the worst case performance can be lower bounded statistically, thereby guaranteeing robustness versus power-tracking when the underlying assumption on the inhomogeneity term is relaxed.

Keywords: Thermostatically controlled load, Markov chain, Poisson process, Formal abstraction, Probabilistic bisimulation, Stochastic optimal control, Noisy optimization.

1 Introduction and Background

Household appliances such as water boilers/heaters, air conditioners and electric heaters – all referred to as thermostatically controlled loads (TCLs) – can store energy due to their thermal mass. These appliances generally operate within a dead-band around a temperature set-point. Control of the aggregate power consumption of a population of TCLs can provide a variety of benefits to the electricity grid. First, ancillary service requests can be partially addressed locally, which reduces the need for additional transmission line capacity. Second,

G. Norman and W. Sanders (Eds.): QEST 2014, LNCS 8657, pp. 57–73, 2014.

controlling a large population of TCLs may improve robustness, since even if a few TCLs fail to provide the required service the consequence on the population as a whole would be small. Such benefits highlight the importance of precise modeling and quantitative control of TCL populations.

Modeling efforts over populations of TCLs and applications to load control arguably initiate with the work in [9]. A discrete-time stochastic model for a TCL is studied in [24], where a simulation model is developed based on a Markov chain approximation of the discrete-time dynamics. A diffusion approximation framework is introduced in [20] to model the dynamics of the electric demand of large aggregates of TCLs by a system of coupled ordinary and partial differential equations. These equations are further studied in [8], where a linear time-invariant dynamical model is derived for the population.

A range of recent contributions [4,18,21,23] employ a partitioning of the TCL temperature range to obtain an aggregate state-space model for the TCL population. Matrices and parameters of the aggregate model are computed either analytically or via system identification techniques. Additional recent efforts have targeted the application of this approach towards higher-order dynamical models [28,29] and the problem of energy arbitrage [22]. The main limitation of these approaches is the lack of a quantitative measure on the accuracy of the constructed aggregated model. Motivated by this drawback, [14,17] have looked at the problem from the perspective of formal abstractions: in contrast to all related approaches in the literature, stymied by the lack of control on the introduced aggregation error, [14,17] have introduced a formal abstraction procedure that provides an upper bound on the error, which can be precisely tuned to match a desired level before computing the actual aggregation.

The purpose of this work is to focus on a new inhomogeneous model for the TCL dynamics, where the inhomogeneity enters the model through a thermal resistance capturing the effect of the opening/closing of windows, of people entering/leaving the room, and so on. The inhomogeneity enters randomly via a Poisson process with a fixed arrival rate, which changes the value of the thermal resistance within a given finite set. We show that the TCL dynamics can be equivalently represented by a discrete-time Markov process, of which we explicitly compute its stochastic kernel. Next, we employ the mentioned abstraction techniques in [1,12] to formally approximate it with a Markov chain. Thus, the aggregated behavior of *a population* of Markov chains can be modeled as a stochastic difference equation [13,14], which is then used for state estimation and closed-loop control of the total power consumption for tracking a load profile.

A crucial assumption in our work underpinning the construction of the stochastic abstraction and the synthesis of the corresponding control scheme is the homogeneity in the *parameters* of the population of TCLs. This assumption might be practically violated. Additionally, some variables such as the initial state of the population and the desired load profile to be tracked are potentially not known in advance. Nevertheless, we would like to guarantee a desired performance of a control scheme. For example, a desirable property could be that the control scheme is able to follow a load profile up to a given deviation while keeping the temperature of the individual households within a certain range. Once a property is specified, guaranteeing such performance amounts to solving a stochastic satisfiability problem. More precisely, one has to choose initial states and load profiles pessimistically to minimize the expect value of the associated cost function. Mathematically this problem can be formulated using stochastic satisfiability modulo theory [16]. Unfortunately, most tools for solving this kind of satisfiability problems are not suited to handle continuous non-determinism. However, there is a tight connection between noisy optimization and stochastic satisfiability [10], rendering the problem suitable for dedicated methods such as [19,27]. We employ these techniques to investigate the robustness of the overall control scheme.

The article is organized as follows. Section 2 introduces the dynamics of a TCL in continuous time, along with the inhomogeneity injected via the Poisson process. Modeling of the TCL dynamics as a discrete-time Markov process and computation of the associated stochastic kernel are presented in Section 3. Abstraction and aggregate modeling of a population of TCLs are then discussed in Section 4. Power reference tracking through closed-loop control of the population is described in Section 5. Finally in Section 6, robustness of the performance of the synthesized controller is validated a-posteriori, against violations on the assumptions on the model.

Throughout this article we use the following notation: $\mathbb{N} = \{1, 2, \ldots\}$ for the natural numbers, $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$, and $\mathbb{N}_n = \{1, 2, \ldots, n\}$ for $n \in \mathbb{N}$.

2 Model of a Thermostatically Controlled Load

Consider the temperature $\theta(t)$ of a TCL evolving in continuous time according to the equation

$$d\theta(t) = \frac{dt}{R(t)C}(\theta_a \pm m(t)R(t)P_{rate} - \theta(t)) \ , \tag{1}$$

where θ_a is the ambient temperature, C indicates the thermal capacitance, and P_{rate} is the rate of energy transfer. In equation (1) a $+$ sign is used for a heating TCL, whereas a $-$ sign is used for a cooling TCL. Inhomogeneity enters via the thermal resistance $R(\cdot) : \mathbb{R}^{\geq 0} \to \{R_0, R_1\}$, which is a function of time and switches between two different values (R_0, R_1), where the switching times are distributed according to the (homogeneous) Poisson process $N(\cdot)$, namely

$$R(t) = \begin{cases} R_0 & \text{if } N(t) \equiv 0 \pmod 2 \\ R_1 & \text{if } N(t) \equiv 1 \pmod 2 \ . \end{cases} \tag{2}$$

The Poisson process accounts for the number of switches and their occurrence time within a given time interval. $N(\cdot)$ is characterized by a specified rate parameter λ, so that the number of switches within the time interval $(t, t + \tau]$ follows a Poisson distribution with

$$\mathbb{P}\{N(t + \tau) - N(t) = n\} = \frac{e^{-\lambda\tau}(\lambda\tau)^n}{n!} \quad \forall n \in \mathbb{N}_0 \ .$$

In equation (1), the quantity $m(\cdot)$ represents the status of the thermostat, namely $m : \mathbb{R}^{\geq 0} \to \{0, 1\}$, where $m(t) = 1$ represents the ON mode and $m(t) = 0$ the OFF mode. For the sake of simplicity, we assume that the Poisson process $N(\cdot)$ is initialized probabilistically according to

$$N(0) = \begin{cases} 0 & \text{with probability } q = 1 - p \\ 1 & \text{with probability } p \ . \end{cases}$$

Moreover, we select $p = q = 1/2$, which means $\mathbb{P}\{R(t) = R_0\} = \mathbb{P}\{R(t) = R_1\} = 1/2$, for all $t \in \mathbb{R}^{\geq 0}$. This simplifying assumption on the initialization of the Poisson process and on the special selection of parameter p can be easily relaxed by including the thermal resistance within the discrete state of the TCL. Without loss of generality, we further assume that $R_0 > R_1$.

With focus on a cooling TCL ($-$ sign in equation (1)), the temperature of the load is regulated by a digital controller $m(t + \tau) = f(m(t), \theta(t))$ that is based on a binary switching mechanism, as follows:

$$f(m, \theta) = \begin{cases} 0 & \text{if } \theta < \theta_s - \delta_d/2 \doteq \theta_- \\ 1 & \text{if } \theta > \theta_s + \delta_d/2 \doteq \theta_+ \\ m & \text{else} \ , \end{cases} \tag{3}$$

where θ_s and δ_d denote the temperature set-point and the dead-band width, respectively, and together characterize the operating temperature range. Note that the switching control signal is applied only at discrete time instants $\{k\tau, k \in \mathbb{N}_0\}$: the mode $m(t)$ may change only at these times, and is fixed in between any two time instants $k\tau$ and $(k+1)\tau$, during which the temperature evolves based on equations (1)-(2) with a fixed $m(k\tau)$. In other words, the operational frequency of the digital controller is $\frac{1}{\tau}$.

The power consumption of the single TCL at time t is equal to $\frac{1}{\eta} m(t) P_{rate}$, which is then equal to zero in the OFF mode and is positive in the ON mode, and where the parameter η is the coefficient of performance (COF). The constant $\frac{1}{\eta} P_{rate}$, namely the power consumed by a single TCL when it is in the ON mode, will be shortened as P_{ON} in the sequel.

In the next section we show that the power consumption of the TCL can be modeled as a Markov process in discrete time, of which we compute the stochastic kernel.

3 Discrete-Time Markov Process Associated to the TCL

We consider a discrete-time Markov process (dtMP) $\{s_k, k \in \mathbb{N}_0\}$, defined over a general (e.g., continuous or hybrid) state space [2]. The model is denoted by the pair $\mathfrak{S} = (\mathcal{S}, T_{\mathfrak{s}})$ in which $\mathcal{S}$ is an uncountable state space. We denote by $\mathcal{B}(\mathcal{S})$ the associated sigma algebra and refer the reader to [2,5] for details on measurability and topological considerations. The stochastic kernel $T_{\mathfrak{s}} : \mathcal{B}(\mathcal{S}) \times \mathcal{S} \to [0, 1]$ assigns to each state $s \in \mathcal{S}$ a probability measure $T_{\mathfrak{s}}(\cdot|s)$, so that for any set

$A \in \mathcal{B}(\mathcal{S}), k \in \mathbb{N}_0, \mathbb{P}\{s_{k+1} \in A | s_k = s\} = T_{\mathfrak{s}}(A|s)$. We assume that the stochastic kernel $T_{\mathfrak{s}}$ admits a representation by its conditional density function $t_{\mathfrak{s}} : \mathcal{S} \times \mathcal{S} \to \mathbb{R}^{\geq 0}$, namely $T_{\mathfrak{s}}(d\bar{s}|s) = t_{\mathfrak{s}}(\bar{s}|s)d\bar{s}$, for any $s, \bar{s} \in \mathcal{S}$.

The digital controller of Section 2 ensures that the power consumption of the TCL is a piecewise-constant signal, where the jumps can happen only at the sampling times $\{k\tau, k \in \mathbb{N}_0\}$. We define $\theta_k = \theta(k\tau), m_k = m(k\tau)$ as the values of the random processes $\theta(\cdot), m(\cdot)$ at the sampling time $k\tau$. Despite the fact that the temperature evolves stochastically in between sampling times, we show that the dynamics of the temperature and of the mode at the sampling times (that is, in discrete time with sampling constant τ) can be modeled as a dtMP and we compute the corresponding conditional density function. In other words, the goal of this section is to compute the density function of the process state at the next discrete time step, conditioned on the state at the current time step.

Define the hybrid state of the dtMP $\mathfrak{S}$ as $s_k = (m_k, \theta_k) \in \mathcal{S} \doteq \{0, 1\} \times \mathbb{R}$. The evolution of the mode is given by the deterministic equation $m_{k+1} = f(m_k, \theta_k)$, while the temperature evolves stochastically and depends on the conditional density function $t_\theta(\theta_{k+1}|\theta_k, m_k)$. Then for all $s_k = (m_k, \theta_k), s_{k+1} = (m_{k+1}, \theta_{k+1}) \in \mathcal{S}$,

$$t_{\mathfrak{s}}(s_{k+1}|s_k) = \boldsymbol{\delta}\left[m_{k+1} - f(m_k, \theta_k)\right] t_\theta(\theta_{k+1}|\theta_k, m_k) \ ,$$

where $\boldsymbol{\delta}[\cdot]$ is the discrete Kronecker delta function. The rest of this section is dedicated to the computation of t_θ, and to the study of its dependency on characteristic parameters. We focus on the explicit computation of t_θ for the case where the mode of the current state is OFF, namely $t_\theta(\theta_{k+1}|\theta_k, m_k = 0)$. The case of $t_\theta(\theta_{k+1}|\theta_k, m_k = 1)$ is similar and thus discussed at the end of this section.

Lemma 1. *Suppose the TCL is in the OFF mode, $m(t) = 0$, during the interval $t \in [t_1, t_2]$. The value of the temperature at the end of the interval solely depends on the relative time the temperature evolves with either of the two resistance values $\{R_0, R_1\}$, and is independent of the actual order or number of occurrence of the two values.*

Lemma 1 states that the distribution of the temperature at the next time step θ_{k+1}, conditioned on the current temperature value θ_k, depends exclusively on the relative time duration that the temperature evolves with any of the two resistances, within the interval: the number or the order of switchings between the resistances is not important, thus the corresponding time can be simply accumulated. Since the resistance changes value based on the jumps of a Poisson process, we define the sum of the length of the sub-intervals of $[k\tau, (k+1)\tau]$ in which the temperature evolves with R_0 (resp. R_1) as the random variable w_0 (resp. w_1). Let us now compute the density functions of these two random variables. Despite the fact that ω_0, ω_1 are defined with respect to the particularly chosen interval $[k\tau, (k+1)\tau]$, next lemma shows that their density functions are independent of k and solely depend on the length of the interval τ.

Lemma 2. *The density function of ω_0 can be expressed as*

$$f_{\omega_0}(x) = e^{-\lambda\tau}\boldsymbol{\delta}(\tau - x) + \lambda e^{-\lambda\tau}\left[I_0\left(2\lambda\sqrt{x(\tau - x)}\right) + \sqrt{\tfrac{x}{\tau - x}}I_1\left(2\lambda\sqrt{x(\tau - x)}\right)\right], \quad (4)$$

which is parametrized by λ, τ, *and is independent of* k. $I_0(\cdot), I_1(\cdot)$ *are modified Bessel functions of the first kind [3, Chapter 9] and* $\boldsymbol{\delta}(\cdot)$ *is the Dirac delta function. The density function of* ω_1 *is* $f_{\omega_1}(x) = f_{\omega_0}(\tau - x)$.

Lemma 2 indicates that the random variable ω_0 (resp. ω_1) is of mixed type, including a probability density function with the interval $[0, \tau]$ as its support, and a probability mass at $x = \tau$ (resp. $x = 0$). Once we know the distributions of ω_0, ω_1, we can compute $t_\theta(\theta_{k+1}|\theta_k, m_k)$ in the OFF mode.

Theorem 1. *The conditional density function* $t_\theta(\theta_{k+1}|\theta_k, m_k = 0)$ *is of the form*

$$t_\theta(\theta_{k+1}|\theta_k, m_k = 0) = \frac{1}{2} t_0(\theta_{k+1}|\theta_k, R_0) + \frac{1}{2} t_1(\theta_{k+1}|\theta_k, R_1) \ , \tag{5}$$

where the functions t_0, t_1 *are computed based on the density functions of* ω_0, ω_1*:*

$$t_i(\theta_{k+1}|\theta_k, R_i) = \frac{1}{|(\theta_{k+1} - \theta_a)\gamma|} f_{\omega_i}\left(\frac{1}{\gamma}\left[\ln\frac{\theta_{k+1} - \theta_a}{\theta_k - \theta_a} + \tau_1\tau\right]\right) \quad i \in \{0, 1\} \ , \tag{6}$$

and where $\tau_0 = \frac{1}{R_0 C}$, $\tau_1 = \frac{1}{R_1 C}$, *and* $\gamma = \tau_1 - \tau_0 > 0$.

We emphasize that the conditional density function t_θ is independent of k, which results in a time-homogeneous dtMP. It has the support $[\theta_{min}, \theta_{max}]$ with

$$\theta_{min} \doteq \theta_a + (\theta_k - \theta_a)e^{-\tau_0\tau} \ , \quad \theta_{max} \doteq \theta_a + (\theta_k - \theta_a)e^{-\tau_1\tau} \ ,$$

and includes two Dirac delta functions at the boundaries of its support

$$\frac{1}{2} e^{-\lambda\tau} \left[\boldsymbol{\delta}\left(\theta_{k+1} - \theta_{min}\right) + \boldsymbol{\delta}\left(\theta_{k+1} - \theta_{max}\right)\right] \ .$$

Moreover it is discontinuous at the boundaries of its support with the following discontinuities:

$$t_\theta(\theta_{min}|\theta_k, m_k = 0) = \frac{\lambda e^{-\lambda\tau}(2+\lambda\tau)}{2|(\theta_k - \theta_a)\gamma|e^{-\tau_0\tau}} \ , \ t_\theta(\theta_{max}|\theta_k, m_k = 0) = \frac{\lambda e^{-\lambda\tau}(2+\lambda\tau)}{2|(\theta_k - \theta_a)\gamma|e^{-\tau_1\tau}} \ . \tag{7}$$

All the above derivations for the density function t_θ conditioned in the OFF mode can be likewise obtained for that in the ON mode: $t_\theta(\theta_{k+1}|\theta_k, m_k = 1)$ is formulated exactly as $t_\theta(\theta_{k+1}|\theta_k, m_k = 0)$ where the quantity θ_a is replaced by the steady-state value of the temperature in the ON mode. The only required assumption is that the temperature trajectories are steered toward the same steady-state value regardless of the thermal resistance. Such a steady-state value is $\theta_\infty = \theta_a - m(t)R(t)P_{rate}$, which is a function of $R(t)$ in the ON mode. This assumption technically allows us to swap the order of the intervals in which the temperature evolves with different resistances and leads to being able to simplify the computations and to obtain the conditional density functions.

4 Abstraction of a Population of Inhomogeneous TCLs

4.1 Abstraction of a TCL as a Markov Chain

The interpretation of the Poisson-driven TCL model as a dtMP allows leveraging an abstraction technique, proposed in [1] and extended in [11,12], aimed at reducing an uncountable state-space dtMP 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 $s_{im} = (m, \bar{\theta}_i), m \in \{0,1\}, i \in \mathbb{N}_n$. The transition probability matrix of $\mathcal{M}$ is made up of the entries

$$\mathbb{P}(s_{im}, s_{i'm'}) = \int_{\Theta_{i'}} t_{\mathfrak{s}}\left(m', \theta' | m, \bar{\theta}_i\right) d\theta' \quad \forall m, m' \in \{0,1\},\, i, i' \in \mathbb{N}_n \ . \tag{8}$$

The initial probability mass of $\mathcal{M}$ is obtained as $p_0(s_{im}) = \int_{\Theta_i} \pi_0(m,\theta) d\theta$, where $\pi_0 : \mathcal{S} \to \mathbb{R}^{\geq 0}$ is the density function of the initial hybrid state of the TCL. For simplicity of notation we rename the states of $\mathcal{M}$ by the bijective map $\ell(s_{im}) = mn + i, m \in \{0,1\}, i \in \mathbb{N}_n$, 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 (3). Further, the density function $t_\theta(\theta_{k+1}|\theta_k, m_k)$ is the summation of two Dirac delta functions and of a piecewise-Lipschitz continuous part. The existence of Dirac delta functions produces technical difficulties in the analysis of properties of interest [15]. Despite these irregularities, we show in the next section that we can compute an upper bound on the error of Markov chain abstraction, based on [14,17]. The abstraction error is composed of three terms related to the Dirac delta functions, the discontinuity at the boundaries, and the state-space discretization.

4.2 Error Computation of the Markov Chain Abstraction

We compute the abstraction error based on [12, pp. 933-934], which gives an upper bound for the abstraction error via the constant $\mathcal{H}$ satisfying the inequality

$$\int_{\mathbb{R}} |t_\theta(\theta_{k+1}|\theta_k, m_k) - t_\theta(\theta_{k+1}|\theta'_k, m_k)| d\theta_{k+1} \leq \mathcal{H}|\theta_k - \theta'_k| \quad \forall \theta_k, \theta'_k, m_k \ . \tag{9}$$

Notice that the integration in the left-hand side is with respect to the next state θ_{k+1}, while the changes are applied to the current state θ_k. In order to compute the constant $\mathcal{H}$, we first establish the Lipschitz continuity of the continuous part of t_θ in Theorem 2 which is founded on the Lipschitz continuity of the bounded part of the functions $f_{\omega_0}(\cdot)$, $f_{\omega_1}(\cdot)$ presented in Lemma 3.

Lemma 3. *Functions* $g_{\omega_0}(x) \doteq f_{\omega_0}(x) - e^{-\lambda\tau}\boldsymbol{\delta}(\tau - x)$ *and* $g_{\omega_1}(x) \doteq f_{\omega_1}(x) - e^{-\lambda\tau}\boldsymbol{\delta}(x)$*, representing the bounded part of functions* $f_{\omega_0}, f_{\omega_1}$*, satisfy the Lipschitz condition*

$$|g_{\omega_i}(x) - g_{\omega_i}(x')| \leq h|x - x'| \quad \forall x, x' \in [0, \tau],\, i \in \{0, 1\} \ ,$$

where $h = \lambda^2 e^{-\lambda\tau} M(2\lambda\tau)$ *and* $M(a) = \max_{u\in[0,1]} |\zeta(u, a)|$*. The function* ζ *is defined as*

$$\zeta(u, a) = I_0\left(a\sqrt{u(1-u)}\right) + \tfrac{1-2u}{\sqrt{u(1-u)}} I_1\left(a\sqrt{u(1-u)}\right) - \tfrac{u}{1-u} I_2\left(a\sqrt{u(1-u)}\right) \ ,$$

where $I_0(\cdot), I_1(\cdot), I_2(\cdot)$ *are modified Bessel functions of the first kind [3].*

Theorem 2. *The density function* $t_\theta(\theta_{k+1}|\theta_k, m_k)$ *satisfies the Lipschitz condition*

$$|t_\theta(\theta_{k+1}|\theta_k, m_k) - t_\theta(\theta_{k+1}|\theta'_k, m_k)| \leq \kappa|\theta_k - \theta'_k| \quad \forall \theta_k, \theta'_k \in [\theta_-, \theta_+],\, m_k \in \{0, 1\} \ ,$$

for all θ_{k+1} *in the intersection of the supports of* $t_\theta(\cdot|\theta_k, m_k), t_\theta(\cdot|\theta'_k, m_k)$*, with* $\kappa = \dfrac{h}{\varrho\gamma^2}$*, and the constant* $\varrho = \min\left\{(\theta - \theta_\infty)^2, \theta \in [\theta_-, \theta_+]\right\}$*.*

Recall that the density function t_θ is discontinuous at the boundaries of its support, with jumps quantified in (7). The value of these jumps appear directly in the left-hand side of inequality (9). Then we have to establish the Lipschitz continuity of the jumps of t_θ with respect to the current state, which is done in the next Lemma.

Lemma 4. *The jumps of the density functions* $t_\theta(\theta_{min}|\theta_k, m_k)$*,* $t_\theta(\theta_{max}|\theta_k, m_k)$*, as in (7), are Lipschitz continuous with respect to the current temperature* θ_k*, with the following constants:*

$$\kappa_0 = \frac{\lambda}{2\varrho|\gamma|}(2 + \lambda\tau)e^{-\lambda\tau}e^{\tau_0\tau} \ , \quad \kappa_1 = \frac{\lambda}{2\varrho|\gamma|}(2 + \lambda\tau)e^{-\lambda\tau}e^{\tau_1\tau} \ .$$

Finally, the transition probabilities of the Markov chain are computed by eliminating the Dirac delta functions and integrating over the partition sets. The total abstraction error is formulated in Theorem 3 using the constants of Theorem 2 and Lemma 4, and leveraging results of [12].

Theorem 3. *If we partition the temperature range with the diameter* δ_p*, the one-step abstraction error is* $\varepsilon = 2(e^{-\lambda\tau} + \kappa_0 + \kappa_1 + \kappa\mathcal{L})\delta_p$*, where* $\mathcal{L}$ *is the Lebesgue measure of the temperature range.*

The error computed in Theorem 3 is useful towards two different purposes. First, it provides a measure on the distance between the power consumption of the model in (1)-(3) and that of the abstracted model [14], namely

$$|\mathbb{E}\left[m(k\tau) - m_{abs}(k\tau)|\,\theta(0), m(0)\right]| \leq k\varepsilon \ .$$

Second, it can be used to check Bounded Linear Temporal Logic (BLTL) specifications over the abstracted Markov chain, providing a guarantee on the specification for the original population model. The error caused by the abstraction for checking any BLTL specification is $N\varepsilon$, where N is the horizon of the specification [25]. Notice that the error can be tuned by proper selection of the partition diameter, time-step, and arrival rate. For instance the error is $\varepsilon = 1.3 \times 10^7 \delta_p$ for the physical parameters of Table 1 (left) that are widely used in the literature, and for the values of Table 1 (right) specifically selected for the model in this study. The large constant in the expression of ε is mainly due to the Lipschitz constant of the density function, which can be reduced by selection of the discretization time step τ.

Table 1. Physical parameters of a residential air conditioner as a TCL [8] (left) and selected parameters for the model in (1)-(3) (right)

Parameter	Interpretation	Value
θ_s	set-point	20 [$^\circ C$]
δ_d	dead-band width	0.5 [$^\circ C$]
θ_a	ambient temperature	32 [$^\circ C$]
C	thermal capacitance	10 [$kWh/^\circ C$]
P_{rate}	power	14 [kW]
η	COF	2.5

Parameter	Interpretation	Value
R_0	thermal resistance	1.5 [$^\circ C/kW$]
R_1	thermal resistance	2.5 [$^\circ C/kW$]
τ	time step	10 [sec]
λ	arrival rate	1 [sec^{-1}]

4.3 Aggregate Model of a Population of TCLs

In the previous section we described how a TCL model is formally abstracted as a Markov chain. In this section we develop a *stochastic* difference equation (SDE) as the aggregate model of the TCL population, which is to be later employed in the state estimation and closed-loop control. Our modeling approach generalizes the results of [7,6] for our setting in which convergence to a *deterministic* difference equation for large populations of Markov chains is investigated in the context of mean field limits. To construct the SDE model, we first take the cross product of all the abstracted Markov chains, to obtain a (admittedly large) Markov chain $\mathcal{Z}$ with finitely many states characterizing the behavior of the population. As a second stage we assign labels $X \in \mathbb{R}^{2n}$ to the states of $\mathcal{Z}$ in which the i^{th} entry of $X(X(i), i \in \mathbb{N}_{2n})$ indicates the proportion of individual Markov chains within the i^{th} state. Given such a labeling, we define an equivalence relation over the labeled Markov chain $\mathcal{Z}$, which relates all states with the same label. This equivalence relation is in fact an exact probabilistic bisimulation [14] and makes it possible to consider equivalence classes as lumped states and thus reduce the state space of $\mathcal{Z}$. In order for this to be an exact probabilistic bisimulation, however, all transition matrices have to be the same, which reduces to TCL characterized by equal parameters, unlike a different ambient temperature (see Section 6). Let us remark that in practice the lumped chain can be obtained directly, with no need to go through the construction of $\mathcal{Z}$. The dynamics over the labels in the reduced Markov chain can be modeled by the following SDE

$$X_{k+1} = P^T X_k + W_k \ , \tag{10}$$

where $X_k \in \mathbb{R}^{2n}$ is the value of the label of the reduced Markov chain at sample time $k\tau$, and its i^{th} entry represents the portion of TCLs with mode and temperature inside the partition set associated to the representative point $\ell^{-1}(i)$ (cf. Section 4.1). Matrix $P = [P_{ij}]_{i,j\in\mathbb{N}_{2n}}$ is the transition probability matrix of the Markov chain $\mathcal{M}$, and W_k is a state-dependent process noise with $\mathbb{E}[W_k] = 0$ and $Cov(W_k) = [\Sigma_{ij}(X_k)]_{i,j\in\mathbb{N}_{2n}} \in \mathbb{R}^{2n\times 2n}$, where

$$\Sigma_{ii}(X_k) = \frac{1}{n_p}\sum_{r=1}^{2n} X_k(r)P_{ri}(1-P_{ri}) \ , \quad \Sigma_{ij}(X_k) = -\frac{1}{n_p}\sum_{r=1}^{2n} X_k(r)P_{ri}P_{rj} \ ,$$

for all $i,j \in \mathbb{N}_{2n}$, $i \neq j$. The process noise W_k converges in distribution to a multivariate normal random vector for large population sizes n_p [14]. Moreover, the transition probability matrix P in (10) depends on the set-point θ_s, which is utilized in the next section for power tracking.

The total power consumption obtained from the aggregation of the original models in (1)-(3), with variables $(m^j,\theta^j)(t), j \in \mathbb{N}_{n_p}$, is

$$y(t) = \sum_{j=1}^{n_p} m^j(t)P_{ON} \ ,$$

which is piecewise-constant in time due to the presence of the digital controller updating the modes of the TCLs. Then the total power consumption can be represented as

$$y(t) = \sum_{k=0}^{\infty} y_k \mathbb{I}_{[k\tau,(k+1)\tau)}(t) \ , \quad y_k = \sum_{j=1}^{n_p} m_k^j P_{ON} \ , \tag{11}$$

where $\mathbb{I}_A(\cdot)$ is the indicator function of a given set A, i.e. $\mathbb{I}_A(x) = 1$ for $x \in A$ and zero otherwise. With focus on the abstract model (10), the power consumption of the model is also piecewise-constant in time with the representation

$$y_{abs}(t) = \sum_{k=0}^{\infty} y_k^{abs}\mathbb{I}_{[k\tau,(k+1)\tau)}(t) \ , \quad y_k^{abs} = HX_k \ , \quad H = n_pP_{ON}[0_n, \mathbb{1}_n] \ , \tag{12}$$

where 0_n and $\mathbb{1}_n$ are n-dimensional row vectors with all the entries equal to zero and one, respectively.

The performance and precision of the aggregated model in (10) is displayed in Figure 1 by comparing its normalized output with the normalized power consumption of the TCL population (namely, by comparing $y_{abs}(t), y(t)$ divided by n_pP_{ON}). In these simulations, the initial temperatures of the TCLs are distributed uniformly over the dead-band and the modes are obtained as samples of Bernoulli trials with a success probability of 0.5. Two different population sizes $n_p = 100$ (top panel) and $n_p = 1000$ (bottom panel) are considered. The oscillations ranges are $[0.35, 0.58]$ and $[0.40, 0.52]$, respectively, which indicates that the oscillations amplitude depends on the population size: it decreases as the population size increases. The number of introduced abstract states in both cases is $2n = 860$. The transition probability matrix P is computed within 8.3 seconds and is sparse with 0.3% non-zero entries.

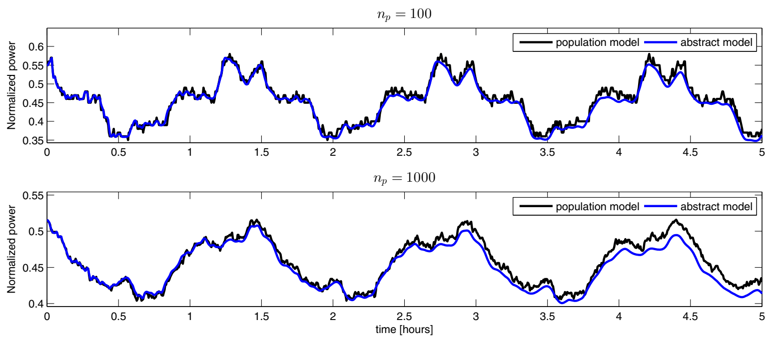

Fig. 1. Comparison of the total power consumption of the population with that of the aggregated model, with population sizes 100 (top) and 1000 (bottom) respectively. The initial conditions have been distributed uniformly.

5 Closed-Loop Control of the TCL Population

As recently investigated in related work [14,17], we consider the set-point θ_s as the control input and regulate the value of this control uniformly over all TCLs. The application of this control scheme practically leads to a change in the position of the non-zero entries of the transition matrix P derived from the dynamics of the single TCL. We discretize the domain of allowable set-point control input by the same partition diameter δ_p used for the temperature range: this allows retaining the definition of the states in X_k.

We employ the stochastic Model Predictive Control (MPC) framework used in [14] over the controlled model in order to track a reference signal for the total power consumption. More precisely, we optimize the following cost function at each time step to track the reference power signal $y_{ref}(\cdot)$:

$$\min_{\theta_s(k\tau)\in[\theta_-,\theta_+]} |\mathbb{E}[y_{abs}(k\tau+\tau)|X_k, \theta_s(k\tau)] - y_{ref}(k\tau+\tau)| \ . \tag{13}$$

A Kalman filter with state-dependent process noise is employed for state estimation when the information of states X_k is not available and only the total power consumption of the population $y(k\tau)$ is measured. Figure 2 presents the closed-loop control scheme for the power reference tracking problem.

The performance of the closed-loop control of Figure 2 is illustrated in Figure 3 (left). A population size $n_p = 100$ is selected and the number of states $2n = 1000$

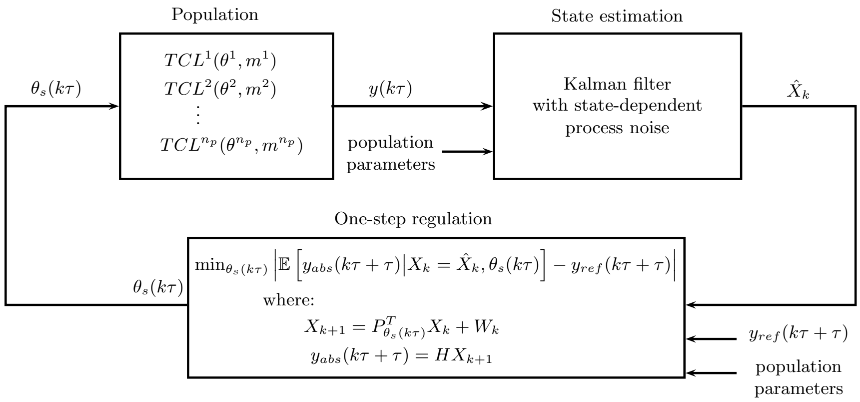

Fig. 2. State estimation and one-step regulation architecture for the closed-loop control of the total power consumption

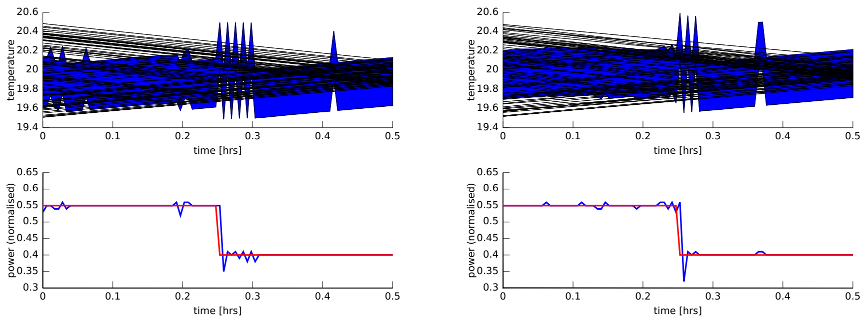

Fig. 3. Illustration of the controller based on a Kalman filter (ref. Figure 2), applied to a population of 100 TCLs, modeled with equations (1)-(3), for either a homogeneous (left) or heterogeneous (right) population. Each TCL has its own initial condition. The upper panels show the evolution of the temperatures across the population in black. The blue region indicates the applied set-point $\pm$ the δ_d dead-band. The lower panels plot in red the desired load profile. The blue lines show the actual load, as achieved when applying the control scheme. For the heterogeneous case the ambient temperature of each TCL is allowed to vary randomly with ± 2 C$^\circ$ around the average ambient temperature of the homogeneous case.

is considered. The lower panel presents the reference signal and the normalized power consumption of the population, while the upper panel shows the synthesized temperature dead-band $[\theta_s(k) - \delta_d/2, \theta_s(k) + \delta_d/2]$. The simulations indicate that the population can accurately track the desired power reference signal. In order to examine the robustness of the control scheme against heterogeneity in the parameters of the population, we run simulations for the case where the ambient temperature of each TCL is allowed to vary randomly with ± 2 C$^\circ$ around the average ambient temperature of the population. The result is presented in Figure 3 (right), which shows that the control scheme is qualitatively robust against violation of this homogeneity assumption. In the next section we quantify more formally the robustness of the closed-loop control scheme.

6 Validation of the Robustness of the Synthesized Controller

As claimed theoretically and further suggested by Figure 3, the control scheme devised on the aggregated population in the previous section can accurately perform tracking of a given load profile for the total power consumption. This requires an assumption on the parameter homogeneity of the TCL population. The simulation shown in Figure 3, however, further suggests that the presented control scheme is robust against heterogeneity, which is to be investigated in this section. Although the following approach could be optimally applied in a closed-loop setting in order to automatically validate the actions taken by the control scheme, we restrict ourselves to an a-posteriori validation due to the computational complexity that is discussed below. We have investigated how well the control scheme based on the Kalman-filter is able to track power consumption in a system that switches modes instantaneously once the temperature hits given thresholds. To apply this control in practice we would like to know how safely the control scheme performs in the worst case, that is across possible unknown variables such as the initial conditions and the desired load profiles. With safe here we mean that the achieved load is within a given range of the desired load profile. As the underlying system is stochastic our goal is to bound the probability of violating such a safety target. Mathematically, we can state such a problem as a noisy optimization problem, representing the probability as an expectation over a binary function:

$$\max_{x \in \mathcal{X}} \mathbb{E}_{y(t), T_0 < t < T}[\phi_{rob}(y)|x] \ , \tag{14}$$

where ϕ_{rob} is an indicator function characterizing whether a trajectory $y(t)$ is in close proximity of the desired load profile. Specifically, we use

$$\phi_{rob}(y) = \mathbb{I}_{|y - y_{ref}| \leq \Delta_l}(y) = \begin{cases} 1 & \text{if } \forall t \in [T_0, T] : |y(t) - y_{ref}(t)| \leq \Delta_l \\ 0 & \text{else} \ , \end{cases}$$

where Δ_l is a parameter controlling the desired degree of load tracking. The vector x represents all variables (in a given space, to be discussed shortly) over which we would like to optimize safety, including the desired load profile (y_{ref}), the initial conditions $(\theta^1(0), m^1(0), \ldots, \theta^{n_p}(0), m^{n_p}(0))$ of TCLs, and the ambient temperatures $\theta_a^1, \ldots \theta_a^{n_p}$. Parameters T_0 and T are used to select the relevant time period over which robustness is validated. Formally, equation (14) falls into the same class of problems as equation (13) and therefore could be solved similarly. As we are interested in robustness properties against heterogeneity, the corresponding optimization space ($\mathcal{X}$) comprises as many dimensions as individual TCLs (x contains all initial conditions of individual TCLs, see above). Formulating the problem as a maximization over an expectation can be interpreted as assigning a probability value to a SSMT formula, comprised of an existentially quantified variable followed by a randomized quantified variable, ϕ_{rob} being an atomic SMT formula (see [16,10] for more details). As formal approaches to solve such problems depend critically on the dimension of the state space, we follow a statistical approach instead, thereby relying only on simulations of the TCL population. To solve such a noisy optimization problem, we are adopting statistical methods, presented in [19,27] to obtain *probably approximate near optimizers*.

Suppose that $g : Y \to \mathbb{R}$, that $\mathbb{P}_Y$ is a given probability measure on Y, and that $\alpha, \epsilon > 0$ are given numbers. A number $g_0 \in \mathbb{R}$ is said to be a Type 3 near minimum of g to level α, or a probably approximate near minimum of g to accuracy ϵ and level α, if $g_0 \geq \min_{y \in Y} g(y) - \epsilon$, and in addition: $\mathbb{P}_Y\{y \in Y : g(y) < g_0 - \epsilon\} \leq \alpha$.

If the probability measure $\mathbb{P}_Y$ is chosen to be the uniform distribution, the probably approximate near minimum is equivalent to the notion of approximate domain optimizer with value imprecision ϵ and residual domain α [19]. These notions of approximate near optimizers can further be extended to hold only with a given confidence ρ, if the approximate near optimizers g_0 have an additional dependence on further random variables [19]. We present a corresponding algorithm to obtain such an optimizer based on uniform sampling.

As mentioned, we aim at a statistical solution to the noisy optimization problem of equation (14). To this end, we use the following simple algorithm [19,27]. The algorithm first samples the parameters over which we would like to optimize the probability of satisfying the robustness property. For each such parameter (in particular containing the initial conditions), the behavior of each TCL is sampled under the closed-loop control from the previous section. Using these samples the probability of satisfying the robustness property can be estimated. The necessary number of samples to achieve the desired accuracy can be calculated in advance (N and M in the above algorithm). For this algorithm it can be shown that the output is a probably approximate near minimum to accuracy ϵ and level α with probability at least ρ, see [26,19,27].

Algorithm 1. Randomized Black-Box optimization algorithm

```
function RandOpt(satProperty, α, ε, ρ)
    N ← log(2/(1−ρ)) / log(1/(1−α)); M ← (1/(2ε²)) log(4N/(1−ρ))
    x^1, ..., x^N ← SampleOptimizationParams(𝒳)        ▷ sample conditions/profiles
    for n = 1, ..., N do
        y_n^1, ..., y_n^M ← SampleTrajectories(x^n)     ▷ Sample power for condition x^n
        for m = 1 ..., M do
            g_{n,m} ← CheckRobustnessOnSample(φ_rob(y_n^m))
        end for
    end for
    ĝ_n ← Average(g_{n,m} across m)          ▷ Estimate robustness for initial conditions
    return min_n ĝ_n
end function
```

Table 2. Parameters for a posteriori verification of the closed-loop control scheme

Δ_l	0.15	Population size n_p	100
Range within ambient temp.	$\pm$ 2°	Level α	0.8
Accuracy ϵ	0.1	Confidence $1-\delta$	0.8
Duration between load switches	15 min	Time horizon T_0, T	25 min, 60 min

We are interested in estimating the worst case performance of the control scheme of Section 5. In such a setting, controllers are typically designed for performance, say to follow a predefined load profile as closely as possible (see equation (13)). Safety is usually not considered during the design phase. Having a system specification for which one can simulate an already designed controller also allows for verifying safety with a statistical procedure such as the one presented in the previous section. If we define safety for a given controller by a maximal allowed deviation Δ_l of the actual load from the desired load profile, we can quantify the number of simulations needed (in terms of precision, level and confidence) in order to guarantee the safety of a given controller for a worst case scenario. To simulate, we have to additionally assume a finite time horizon: we have chosen 4 times the period needed for updating the desired load profile, in order to cover the relevant changes between desired values (see Table 2). Using Algorithm 1 we can now verify that such a controller achieves a desired load profile robustly across a heterogeneous parameter within the ambient temperature. More precisely, using the parameters in Table 2, we could verify that the worst-case probability of resulting in a non-safe system using the closed-loop control scheme is guaranteed to be bounded by ϵ (accuracy=0.1) up to residual α level and confidence, as given in Table 2. The number of simulations for such parameters is $N = 738$ and the objective is to have the total power consumption within ± 0.15 of the desired reference load. The worst case were determined over all possible initial conditions (temperatures and modes of TCLs) as well as over a set of desired load profiles. For the set of load profiles we considered step-functions, which change the desired load every 15 minutes and have a height $\in [0.4, 0.6]$. Due to the independence between different simulations, Algorithm 1 can be parallelized efficiently. Nevertheless, the necessary number of samples quickly increases with the different parameters for accuracy using this procedure, therefore, it still needs considerable computational effort. Although such a procedure could in principle also be used to verify safety or solve the control problem in a closed-loop setting, we did not investigate such scheme due to the computational effort. To illustrate the feasibility of the approach and due to the high computational load, we set α, or residual domain to 0.8 thereby allowing 80 percent of the optimization domain to have a potentially worse robustness level.

7 Conclusions

In this work the problem of aggregate modeling and control of a population of TCLs has been addressed. The temperature evolution is modeled in continuous time and combined with a digital ON/OFF switching controller. The TCLs are allowed to have different thermal resistances which change values in time based on a Poisson arrival process, and thus induce inhomogeneity in the dynamics. The power consumption of each TCL is modeled as a Markov process and formally abstracted to a Markov chain, which is then used to develop an aggregate model for the total power consumption of the population. Finally, a control scheme is proposed to track a power reference signal and its robustness is examined against violation of the homogeneity assumption by use of statistical techniques.

References

1. Abate, A., Katoen, J.-P., Lygeros, J., Prandini, M.: Approximate model checking of stochastic hybrid systems. European Journal of Control 6, 624–641 (2010)
2. Abate, A., Prandini, M., Lygeros, J., Sastry, S.: Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. Automatica 44(11), 2724–2734 (2008)
3. Abramowitz, M., Stegun, I.A.: Handbook of mathematical functions with formulas, graphs, and mathematical tables. Dover, New York (1964), ninth dover printing, tenth gpo printing edition
4. Bashash, S., Fathy, H.K.: Modeling and control insights into demand-side energy management through setpoint control of thermostatic loads. In: Proceedings of the 2011 American Control Conference, pp. 4546–4553 (June 2011)
5. Bertsekas, D.P., Shreve, S.E.: Stochastic optimal control: the discrete-time case. Athena Scientific (1996)
6. Bortolussi, L., Hillston, J., Latella, D., Massink, M.: Continuous approximation of collective system behaviour: A tutorial. Performance Evaluation 70(5), 317–349 (2013)
7. Le Boudec, J.-Y., McDonald, D., Mundinger, J.: A generic mean field convergence result for systems of interacting objects. In: 4th International Conference on the Quantitative Evaluation of SysTems (QEST), pp. 3–18 (2007)
8. Callaway, D.S.: 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)
9. Chong, C.-Y., Debs, A.S.: Statistical synthesis of power system functional load models. In: 18th IEEE Conference on Decision and Control including the Symposium on Adaptive Processes, vol. 18, pp. 264–269 (1979)
10. Ellen, C., Gerwinn, S., Fränzle, M.: Confidence bounds for statistical model checking of probabilistic hybrid systems. In: Jurdziński, M., Ničković, D. (eds.) FORMATS 2012. LNCS, vol. 7595, pp. 123–138. Springer, Heidelberg (2012)
11. Esmaeil Zadeh Soudjani, S., Abate, A.: Higher-order approximations for verification of stochastic hybrid systems. In: Chakraborty, S., Mukund, M. (eds.) ATVA 2012. LNCS, vol. 7561, pp. 416–434. Springer, Heidelberg (2012)
12. Esmaeil Zadeh Soudjani, S., Abate, A.: Adaptive and sequential gridding procedures for the abstraction and verification of stochastic processes. SIAM Journal on Applied Dynamical Systems 12(2), 921–956 (2013)
13. Esmaeil Zadeh Soudjani, S., Abate, A.: Aggregation of thermostatically controlled loads by formal abstractions. In: European Control Conference (ECC), Zurich, Switzerland, pp. 4232–4237 (July 2013)
14. Esmaeil Zadeh Soudjani, S., Abate, A.: Aggregation and control of populations of thermostatically controlled loads by formal abstractions. IEEE Transactions on Control Systems Technology (accepted, 2014)

15. Esmaeil Zadeh Soudjani, S., Abate, A.: Probabilistic reach-avoid computation for partially-degenerate stochastic processes. IEEE Transactions on Automatic Control 59(2), 528–534 (2014)
16. Fränzle, M., Hermanns, H., Teige, T.: Stochastic satisfiability modulo theory: A novel technique for the analysis of probabilistic hybrid systems. In: Egerstedt, M., Mishra, B. (eds.) HSCC 2008. LNCS, vol. 4981, pp. 172–186. Springer, Heidelberg (2008)
17. Kamgarpour, M., Ellen, C., Esmaeil Zadeh Soudjani, S., Gerwinn, S., Mathieu, J.L., Mullner, N., Abate, A., Callaway, D.S., Franzle, M., Lygeros, J.: Modeling options for demand side participation of thermostatically controlled loads. In: International Conference on Bulk Power System Dynamics and Control (IREP), pp. 1–15 (August 2013)
18. Koch, S., Mathieu, J.L., Callaway, D.S.: Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services. In: 17th Power Systems Computation Conference, Stockholm, Sweden (August 2011)
19. Lecchini Visintini, A., Lygeros, J., Maciejowski, J.: Stochastic optimization on continuous domains with finite-time guarantees by Markov chain Monte Carlo methods. IEEE Transactions on Automatic Control 55(12), 2858–2863 (2010)
20. Malhame, R., Chong, C.-Y.: 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)
21. Mathieu, J.L., Callaway, D.S.: State estimation and control of heterogeneous thermostatically controlled loads for load following. In: Hawaii International Conference on System Sciences, Hawaii, USA, pp. 2002–2011 (2012)
22. Mathieu, J.L., Kamgarpour, M., Lygeros, J., Callaway, D.S.: Energy arbitrage with thermostatically controlled loads. In: European Control Conference (ECC), pp. 2519–2526 (2013)
23. Mathieu, J.L., Koch, S., Callaway, D.S.: State estimation and control of electric loads to manage real-time energy imbalance. IEEE Transactions on Power Systems 28(1), 430–440 (2013)
24. Mortensen, R.E., Haggerty, K.P.: A stochastic computer model for heating and cooling loads. IEEE Transactions on Power Systems 3(3), 1213–1219 (1988)
25. Tkachev, I., Abate, A.: Formula-free finite abstractions for linear temporal verification of stochastic hybrid systems. In: Proceedings of the 16th International Conference on Hybrid Systems: Computation and Control, Philadelphia, PA, pp. 283–292 (April 2013)
26. Vidyasagar, M.: Learning and generalization: with applications to neural networks. Springer (1997)
27. Vidyasagar, M.: Randomized algorithms for robust controller synthesis using statistical learning theory. Automatica 37, 1–2 (2001)
28. Zhang, W., Kalsi, K., Fuller, J., Elizondo, M., Chassin, D.: Aggregate model for heterogeneous thermostatically controlled loads with demand response. In: IEEE PES General Meeting, San Diego, CA (July 2012)
29. Zhang, W., Lian, J., Chang, C.-Y., Kalsi, K.: Aggregated modeling and control of air conditioning loads for demand response. IEEE Transactions on Power Systems 28(4), 4655–4664 (2013)