# Modeling Options for Demand Side Participation of Thermostatically Controlled Loads

Maryam Kamgarpour[*], Christian Ellen[†], Sadegh Esmaeil Zadeh Soudjani[‡],
Sebastian Gerwinn[†], Johanna L. Mathieu[§], Nils Müllner[†],
Alessandro Abate[‡], Duncan S. Callaway[¶], Martin Fränzle[†], John Lygeros[*]
[*] Automatic Control Laboratory, Swiss Federal Institute of Technology, Switzerland
[†] OFFIS, Universität Oldenburg, Germany
[‡] Delft Center for Systems and Control, TU Delft, The Netherlands
[§] Power System Laboratory, Swiss Federal Institute of Technology, Switzerland
[¶] Energy and Resources Group, University of California, Berkeley, USA
E-mail: {mkamgar,lygeros}@control.ee.ethz.ch

## Abstract

Residential thermostatically controlled loads (TCLs) have potential for participation in electricity markets. This is because we can control a large group of these loads to achieve aggregate system behavior such as providing frequency reserves while ensuring the control actions are non-disruptive to the end users. A main challenge in controlling aggregations of TCLs is developing dynamical system models that are simple enough for optimization and control, but rich enough to capture the behavior of the loads. In this work, we propose three classes of models that approximate aggregate TCL dynamics. We analyze these models in terms of their accuracy and computational tractability. The models demonstrate a progression from models that help us analyze and predict TCL population behavior to those that help us develop large-scale automatic control strategies. Specifically, we demonstrate how formal methods from computer science and optimal control can be used to derive bounds on model error, guarantees for trajectory tracking, and algorithms for price arbitrage. We find that the accuracy of the analytic results decreases as TCL parameter heterogeneity is introduced. Thus, we motivate further development of analytical tools and modeling approaches to investigate realistic TCL behavior in power systems.

## I. Introduction

Household appliances such as water boilers/heaters, electric heaters, and air-conditioners, referred to as thermostatically controlled loads (TCLs) generally operate within a dead-band around a temperature set point. These appliances can store energy due to their thermal mass. Thus, we can control them by turning them on/off prematurely or slightly adjusting their temperature set point, and still ensure that they can provide the service expected by the user. In this way, one can hope to control a large population of TCLs so that their aggregate power consumption tracks a signal or minimizes a cost. For example, the aggregate power can follow a desired trajectory determined based on the power fluctuations of renewable energy sources or the demand on the grid [1].

Control of the aggregate power output of a TCL population can provide a variety of benefits to the electricity grid. A key change to the grid is the increase in renewable energy sources, such as wind and solar, which are hard to accurately predict. This leads to the need for additional ancillary services [2], [3] such as control reserves which are today mainly provided by conventional generators. With Smart grid sensing and communication infrastructure, population of TCLs can provide additional means for providing control reserves [4]. There are several advantages to using TCLs for this task [5]. First, ancillary service needs can be partially addressed locally, which reduces the need for additional transmission line capacity. Second, using a large population of TCLs may improve robustness, since if a few TCLs fail to provide the required service, the consequence in the large population would be small. Third, the resource potential is large [6].

To integrate TCL aggregations into the power systems, several challenges in terms of modeling, control, and communication need to be addressed. First, although accurate models for individual TCLs exist, developing a computationally tractable and accurate aggregate model for a population of TCLs is challenging. Next, control schemes must be developed so that the population tracks a given power trajectory with a required accuracy or can minimize a cost function. The control must be robust with respect to uncertainties in the model and exogenous inputs such as weather, uncertainty in renewable energy forecasts, and uncertainty in price forecasts. Since an aggregator will perform the estimation and control of the TCL population, there must be a communication link between each TCL and the aggregator. The design of such communication infrastructure and trade-offs between its cost and accuracy in estimation and control is the subject of investigation [7].

The purpose of this paper is to propose and examine several classes of models for analysis and control of TCL populations. Each proposed model makes certain assumptions about the TCL parameters and the information available to a central aggregator and has certain computational and deductive capabilities about the population performance. In particular, the objective of the first part of the paper is to derive an abstraction of the TCL population dynamics which analytically

characterizes the process noise in the abstraction and derives bounds on the error between the power output of the TCL population and that predicted by the model. Partial results of this approach appeared in our recent work [8]. The additional contributions here include the consideration of a heterogeneous TCL population. The second section explores formulation of a new method for modeling the uncertainties in TCL population. The objective of this section is to determine whether it is feasible to track a power signal within a required accuracy, given observations of the power consumption of the TCL population. The third section examines the potential of a TCL population to participate in energy markets through energy arbitrage. This section extends our recent results in [9] by finding upper bounds on energy savings, computed assuming each TCL optimizes its own energy costs. Each section works with a different model and all three models are simulated based on realistic parameter data for the TCLs. The tradeoffs in the accuracy, complexity, and application of the models are analyzed.

The paper is organized as follows. In Section II, we provide some background on TCL population modeling. We also include the individual TCL model and parameters that are used as the ground truth in all sections. In Section III, we develop a Markov chain abstraction of population of TCLs with accurate analysis of stochasticity in the population dynamics and error bounds of the model. In Section IV, we develop a framework in which, with only historic measurements, deductive quantitative conclusions about the performance of the TCL population can be determined. In Section V, motivated by optimization of long-run TCL operation, we develop an aggregate model suitable for optimization and apply the model from Section III for optimal control around the optimized power trajectory. In Section VI, we make comparisons of the models and recommendations of feasibility of each model in different scenarios. Finally, we conclude in Section VII.

## II. Background on TCL Modeling

### II.A. Individual TCL Model

The starting point of all three models is the following discrete time Stochastic Hybrid System (SHS) [10], which describes the evolution of the temperature of a single cooling TCL [1], [11], [12], [13], [14]:

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

where $\theta$ is the temperature of the load, $\theta_a$ 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 noise $w(t) \in \mathbb{R}$ is independent identically distributed (i.i.d.) and characterized by a probability density function $p_w$. The temperature dynamics are regulated by the discrete switching variable $m(t) \in \{0, 1\}$.

The switching dynamics for a cooling TCL are

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

where $\theta_s$ denotes the temperature set-point and $\delta$ the dead-band width. We define $\theta_- = \theta_s - \delta/2$ and $\theta_+ = \theta_s + \delta/2$ as the lower and upper boundary of the temperature range. The power consumption of a TCL at time $t$ is equal to $\frac{1}{\eta}m(t)P_{rate}$, where the parameter $\eta$ is the Coefficient Of Performance (COP). For notation simplicity, we define

$$\bar{P}_{rate} = \frac{1}{\eta}P_{rate}.$$

In the discrete time model, the changes in the thermostat state occur only at discrete time steps.

### II.B. Probability Distribution Modeling

One of the early works on TCL population models was [14]. Here, the authors focused on characterizing the probability distribution of homogenous populations of TCLs in temperature space. They derived partial differential equations (PDEs) for the evolution of the probability density function (PDF) of the TCLs in the On and Off mode. The PDF characterization suffers from a few drawbacks including computation of the solutions, extensions to non-homogeneous populations, and its use for control synthesis. An exact solution of the PDE was developed in [1]; however, it was shown that the aggregate power output predicted by the PDE does not approximate the entire temperature distribution accurately in the case of heterogenous loads.

In an effort to better understand TCL aggregate behavior, input-output models [1], state queuing models [15], [16] and Markov chain models have been proposed. Instead of using PDFs, the Markov chain models approximate the evolution of TCL populations over discretized temperature intervals and on/off modes, and at discrete time points [17], [18], [19], [20], [7]. The Markov chain models are in the class of linear time varying system

$$\mathbf{X}(t+1) = P^T\mathbf{X}(t) + \mathbf{W}(t), \quad (3)$$

where $\mathbf{X}$ denotes a vector containing the probability mass of TCLs within each discretized temperature state and On/Off mode, $P$ is the transition probability matrix, $(.)^T$ denotes the transpose of the matrix, and $\mathbf{W}$ denotes a noise term. This modeling approach was used for tracking control over short horizons and estimation of state. More detailed individual TCL models based on three state hybrid systems were used in [21] and a Markov chain abstraction of the population's corresponding probability density function was also addressed.

In relation to past work, Section III of this paper provides insight into the Markov chain abstraction method (referred to as Model 1) by characterizing the process noise term $\mathbf{W}(t)$ of

| Parameter | Interpretation | Model 1 & 2 | Model 3 |
|-----------|----------------|-------------|---------|
| $\theta_s$ | set-point | 20°C | 18–27°C |
| $\delta$ | dead-band width | 0.5°C | 0.25–1°C |
| $\theta_a$ | ambient temperature | 32°C | varying |
| $R$ | thermal resistance | 2°C/kW | 1.5–2.5°C/kW |
| $C$ | thermal capacitance | 10kWh/°C | 1.5–2.5 kWh/°C |
| $P_{rate}$ | power | 14 kW | 10–18 kW |
| $\eta$ | COP | 2.5 | 2.5 |
| $h$ | time step | 10 sec | 10 sec |

Table I: Parameters for simulation of air conditioner populations from [1] and [20].

## II.C. Model Parameters

In this paper, we consider populations of air conditioners. Equations (1) and (2) are used as the plant. The parameters in Sections III and IV are set according to Table I, column 3 [1]. In the third modeling approach, in Section V, a heterogenous population is considered and the parameters are uniformly distributed within the bounds given in Table I, column 4 [20].

## III. Model 1 - A Markov Chain Abstraction with Analytic Error Bounds

We develop a novel two-step abstraction procedure to derive a linear stochastic dynamical model for the TCL population. In the first step, a population of discrete time Markov chains is generated based on the probabilistic evolution of the continuous state model of each TCL temperature; in the second step, a Markov chain of reduced order is derived which is an exact representation of the population model, that is, it is probabilistically bisimilar to the original model. The approach is analytically developed for the case of a homogeneous population of TCLs, and extended to a heterogeneous population. While the resulting model is similar to those proposed in the literature [17], [18], [19], [20], [7], the analytic derivation allows us to quantify the abstraction error and thus the difference between the original and the modeled population. This section extends our work in [8].

Throughout this section we use the notation $\mathbb{N}$ for natural numbers, $\mathbb{Z} = \mathbb{N} \cup \{0\}$, $\mathbb{N}_n = \{1, 2, 3, \cdots, n\}$, and $\mathbb{Z}_n = \mathbb{N}_n \cup \{0\}$. We denote vectors with bold typeset and with a letter corresponding to that of its elements.

## III.A. Abstraction of a Single TCL

The interpretation of (1)-(2) as a SHS enables the use of an abstraction technique first proposed in [22], aimed at
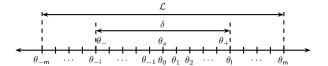


Figure 1: Partition of the temperature axis for abstraction of a TCL.

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 $\{\bar{\theta}_i \in \Theta_i, i \in \mathbb{N}_n\}$. The hybrid state space is characterized by a variable $s = (m, \theta) \in \mathbb{Z}_1 \times \mathbb{R}$. 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 $\mathcal{M}$ consists of the elements

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

where $m' \in \mathbb{Z}_1, i' \in \mathbb{N}_n$ and $\delta_d[m' - f(m, \bar{\theta}_i)]$ is the Dirac-delta function. For 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

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

Due to the presence of switching dynamics, the conditional density function of the stochastic system describing the dynamics of a single TCL is discontinuous. 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 $\mathbb{R}$ as shown in Figure 1. Let us select two constants $\mathsf{l}, \mathsf{m} \in \mathbb{N}, \mathsf{l} < \mathsf{m}$, compute the partition size $\tau = \delta/2\mathsf{l}$ and quantity $\mathcal{L} = 2\mathsf{m}\tau$. Now construct the boundary points of the partition sets $\{\theta_i\}_{i=-\mathsf{m}}^{i=\mathsf{m}}$ for the temperature axis as

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

We render the Markov states of the infinite length intervals $\Theta_1, \Theta_n$ absorbing, that is, once the temperature reaches one of these intervals, it remains there forever.

## III.B. Abstraction of a Homogeneous Population of TCLs

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 $\theta_s, \delta, \theta_a, R, C, P_{rate}, \eta, h$, and the distribution $p_w$ of $w(t)$. Each TCL can be abstracted

as a Markov chain with the transition probability matrix $P = [P_{ij}]$, where $i,j \in \mathbb{N}_{2n}$. This abstraction leads to $n_p$ identical Markov chains $\mathcal{M}$.

The homogeneous TCL population can be represented by a single Markov chain $\Xi$, built as the cross product of the $n_p$ Markov chains. The state of the Markov chain $\Xi$ is

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

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

The Markov chain $\Xi$ has $(2n)^{n_p}$ states, which in general can be very large. As the second step of the abstraction procedure, we are interested in further aggregating this model. The motivation for this approach stems from the fact that for studying aggregate power consumption, it is sufficient to know the number of TCLs in each discrete state. Formally, the aggregation is achieved through the notion of (exact) probabilistic bisimulation [23].

Let us introduce $AP$ as a constrained vector with a dimension corresponding to the number of states of Markov chain $\mathcal{M}$:

$$AP = \left\{ \mathbf{x} = [x_1, x_2 \cdots, x_{2n}]^T \in \mathbb{Z}_{n_p}^{2n} \bigg| \sum_{i=1}^{2n} x_i = n_p \right\}.$$

A function $L : \mathcal{Z} \to AP$, referred to as a labeling function, associates to a configuration $\mathbf{z}$ of $\Xi$ a vector $\mathbf{x} = L(\mathbf{z})$, the elements $x_i \in \mathbb{Z}_{n_p}$ of which count the number of TCLs 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 $\Xi$. Define an equivalence relation $\mathcal{R}$ on the state space of $\mathcal{Z}$, such that

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

This 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. The equivalence relation $\mathcal{R}$ is an exact probabilistic bisimulation relation on $\Xi$ [23], which means for any set $T \in \mathcal{Z}/\mathcal{R}$

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

where $P_\Xi(\mathbf{z}, T) = \sum_{\mathbf{z}_1 \in T} P_\Xi(\mathbf{z}, \mathbf{z}_1)$. Given an observation $\mathbf{x}(t) \in AP$ at time $t$ over the Markov chain $\Xi$, it is of interest to compute the probability mass function of the conditional random variable $(x_i(t+1)|\mathbf{x}(t))$ as $\mathsf{P}(x_i(t+1) = k|\mathbf{x}(t))$, for any $k \in \mathbb{Z}_{n_p}, i \in \mathbb{N}_{2n}$. Using the law of total probability we get the following result.

*Theorem 1:* 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 [24]. Their mean, variance, and covariance are characterized by

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

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

Theorem 1 indicates that the distribution of the conditional random variable $(\mathbf{x}(t+1)|\mathbf{x}(t))$ is independent of the underlying state $\mathbf{z}$ of $\Xi$ in which $L(\mathbf{z}) = \mathbf{x}$.

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 $\mathbf{X}$. 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 which allows expressing the following dynamical model for the variable $\mathbf{X}$:

$$\mathbf{X}(t+1) = P^T \mathbf{X}(t) + \mathbf{W}(t), \qquad (5)$$

where the distribution of $\mathbf{W}(t)$ depends only on the state $\mathbf{X}(t)$. We use the Lyapunov central limit theorem [24] 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

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

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

We have modeled the evolution of the TCL population with an abstract model based on linear stochastic difference equations (5). The approach in derivation of the stochastic model above is different than that of [7] in that the noise covariance is derived analytically in the current approach, in contrast to estimated via simulation in the latter approach.

## III.C. Quantification of the Abstraction Error

The total power consumption obtained from the aggregation of the individual models in (1)-(2), with variables $(m_i, \theta_i)(t), i \in \mathbb{N}_{n_p}$, denoting TCL $i$ is given as

$$y_{total}(t) = \sum_{i=1}^{n_p} m_i(t)\bar{P}_{rate}.$$

Focusing on the abstract model, described in terms of the normalized variable $\mathbf{X}$, the power consumption is equal to

$$y_a(t) = H\mathbf{X}(t), \quad H = \bar{P}_{rate}[\mathbf{0}_n, \mathbf{1}_n],$$

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

*Theorem 3:* Consider a homogeneous population of TCLs with a Gaussian process noise $w(\cdot) \sim \mathcal{N}(0, \sigma^2)$, and the abstracted model constructed based on the partitions defined in (4). 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

$$\begin{aligned} &\left| \mathbb{E}[y_{total}(t)|\mathbf{s}_0] - \mathbb{E}[y_a(t)|\mathbf{X}_0] \right| \\ &\quad \leq n_p(t-1)\bar{P}_{rate}\left[ \frac{(t-2)}{2}\epsilon + \frac{2a}{\sigma\sqrt{2\pi}}\tau \right], \end{aligned} \quad (6)$$

where the constants above are given as

$$\begin{aligned} \epsilon &= \frac{e^{-\gamma^2/2}}{\gamma\sqrt{2\pi}}, \\ \gamma &= \frac{1-a}{2\sigma}\left[ \frac{\mathcal{L}a^t + \delta}{1 - a^t} - RP_{rate} - |2(\theta_s - \theta_a) + RP_{rate}| \right], \end{aligned}$$

for all $\mathbf{s}_0 \in (\mathbb{Z}_1 \times [\theta_{-m}, \theta_m])^{n_p}$. The initial state $\mathbf{X}_0$ is a function of the initial states in the population of TCLs $\mathbf{s}_0$, according to the definition of the state vector $\mathbf{X}$.

The importance of this theorem is that it allows us to tune the error in estimating the total power consumption of the population from the abstraction. Notice that the constant $\gamma$ is an affine function of $\mathcal{L}$. The constant $\epsilon$, and consequently the first term of the error bound (6), is reduced by selecting a larger interval around the dead-band to be partitioned. The second term of the error bound (6) is decreased by a smaller partition diameter $\tau$. The error bound depends linearly on the population size since the total power consumption is the sum of power consumption of $n_p$ single TCL. The error bound depends quadratically on time. Currently, we are exploring improvements to this bound.

## III.D. Extension to a Heterogeneous Population of TCLs

Consider a heterogeneous population of $n_p$ TCLs, where heterogeneity is characterized by a parameter $\alpha$ that takes $n_p$ values. Each instance of $\alpha$ specifies a set of model parameters $(\theta_s, \delta, \theta_a, R, C, \sigma)$ for a single TCL. The dynamical model can be abstracted as a Markov chain $\mathcal{M}_\alpha$ with a transition matrix $P_\alpha = [P_{ij}(\alpha)]_{i,j}$. This transition probability matrix obtained for a TCL depends on its own set of parameters specified by $\alpha$. The apparent difficulty is that the heterogeneity in the transition probability matrices renders the quantity $P(x_i(t+1) = j|\mathbf{z}(t))$ dependent not only on the label $\mathbf{x}(t) = L(\mathbf{z}(t))$, but also on the current state $\mathbf{z}(t)$.

In contrast to the homogeneous case, which allows us to compute probabilities $P(x_i(t+1) = j|\mathbf{x}(t))$ by constructing an exact probabilistic bisimulation of $\Xi$, in the heterogeneous case we have to leverage approximate probabilistic bisimulation of the Markov chain $\Xi$. In this case, the approximation error can only be quantified empirically using the matrix

$P(\mathbf{z}(t+1)|\mathbf{z}(t))$, which in practice can be unfeasible. We encompass the population heterogeneity by constructing an empirical probability distribution $f_\alpha(\cdot)$ from the finite set of values for parameter $\alpha$.

*Theorem 4:* If the TCL population heterogeneity is characterized by a parameter $\alpha$ with empirical distribution $f_\alpha(\cdot)$, using the approximate probabilistic bisimulation of the Markov chain $\Xi$, the random vector $(\mathbf{X}(t+1)|\mathbf{X}(t))$ has the expected value $\mathbf{M}(\mathbf{X}(t)) = \bar{P}^T\mathbf{X}(t)$ and covariance matrix $\Sigma(\mathbf{X}(t))$, where for all $i, j \in \mathbb{N}_{2n}, i \neq j$,

$$\begin{aligned} \Sigma_{ii}(\mathbf{X}) &= \frac{1}{n_p}\sum_{r=1}^{2n} X_r\overline{P_{ri}(1 - P_{ri})} \\ &\quad + \frac{1}{n_p - 1}\left( \sum_{r=1}^{2n} X_r\overline{P_{ri}} \right)^2 - \frac{1}{n_p - 1}\sum_{r=1}^{2n} X_r\overline{P_{ri}}^2, \end{aligned}$$

$$\begin{aligned} \Sigma_{ij}(\mathbf{X}) &= \frac{1}{n_p - 1}\left( \sum_{r=1}^{2n} X_r\overline{P_{ri}} \right)\left( \sum_{s=1}^{2n} X_s\overline{P_{sj}} \right) \\ &\quad - \frac{1}{n_p - 1}\sum_{r=1}^{2n} X_r\overline{P_{ri}P_{rj}} - \frac{1}{n_p}\sum_{r=1}^{2n} X_r\overline{P_{ri}P_{rj}}. \end{aligned}$$

The bar notation indicates the expected value with respect to the parameters set $\alpha$, for instance,

$$\overline{P_{ri}P_{rj}} = \mathbb{E}_\alpha[P_{ri}(\alpha)P_{rj}(\alpha)] = \int P_{ri}(v)P_{rj}(v)f_\alpha(v)dv.$$

Theorem 4 enables us to use model (5) as an approximation of the dynamics of the heterogeneous population. The transition matrix $P$ and the covariance of $\mathbf{W}(t)$ must be computed with respect to the set of parameters, according to the Theorem 4.

## III.E. Numerical Benchmark

In this section, we compare the performance of our formal abstraction with a deterministic abstraction. A TCL population size of $n_p = 500$ is considered for all the simulations. Each TCL is characterized by parameters that take values in Table I. All TCLs are initialized in the Off mode ($m(0) = 0$) and with a temperature at the set-point ($\theta(0) = \theta_s$). We assume the process noise has a Gaussian distribution with a standard deviation $\sigma = 0.01\sqrt{h} = 0.032$.

For the formal abstraction proposed in this work, we construct a partition as per (4) with $l = 7, m = 35$, which leads to $2n = 144$ abstract states. We then generate the abstracted system trajectory using Equation (5), where the covariance matrix of the noise term $W(t)$ is that given by the limiting covariance in Theorem (2). 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.

For comparison, we also perform a deterministic abstraction which does not consider the analytically derived covariance
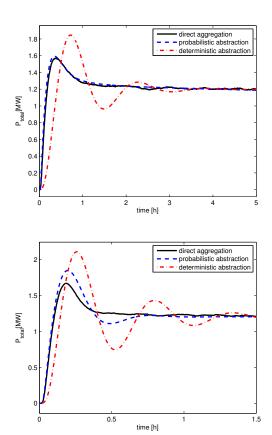
Figure 2: Comparison of a deterministic abstraction with the formal stochastic abstraction in this work for a homogeneous (top) and a heterogeneous population (bottom).



Figure 3: Tracking a piecewise constant reference signal (top) by set-point control (bottom) in a homogeneous population of TCLs.

matrix for the process noise $W(t)$. We select a $n_d = 5$ for the number of bins for this abstraction, which leads to 10 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 [19]. Figure 2 (top) presents the results of the experiment. As can be observed, generating the stochastic model based on (5) results in more accuracy than a deterministic abstraction.

To observe the performance of the proposed model for the heterogeneous population, let us assume that heterogeneity enters the TCL population in only the thermal capacitance $C$ of each single TCL, which is taken to be $C \sim \mathcal{U}([2, 18])$, that is a uniform distribution over a compact interval. We perform 50 Monte Carlo simulations with a noise level $\sigma = 0.032$, and select discretization parameters $n_d = 7$, $l = 10$, and $m = 50$. The outcome is presented in Figure 2 (bottom).

## III.F. Control of TCL Population

Among the different strategies for controlling the total power consumption of a population of TCLs, we consider the case in which the control input is the set-point $\theta_s$ of the TCL [1]. We intend to apply the same control input to all TCLs since
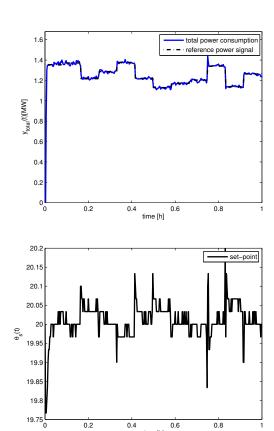
this requires no prior knowledge of the state of the single TCL. Given the model parameters, we use online measurement of the total power consumption of the TCL population, to estimate the states in $\mathbf{X}(t)$ and we use the set-point $\theta_s$ to track any reference signal based on a one-step output prediction.

Suppose we have a homogeneous population of TCLs with known parameters. Based on Equation (5), we set up the model

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

where $\theta_s(t)$, the set-point value at time $t$, is the control input for the model. We assume that the control input is discrete and take values from a set:

$$\theta_s(t) \in \{\theta_{-l}, \theta_{-l+1}, \cdots, \theta_{l-1}, \theta_l\}, \quad \forall t \in \mathbb{Z}.$$

This assumption makes it possible to use the partitions defined in Equation set (4) at all time steps. The process noise $\mathbf{W}(t)$ is normal with zero mean and its state-dependent covariance matrix is obtained from Theorem 1. 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 which depends 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.

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

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

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

This scheme is implemented on a homogeneous population of $n_p = 500$ TCLs, for tracking a randomly generated piecewise constant reference signal. We set the discretization parameters to $l = 8$, $m = 40$, where the standard deviation of the measurement noise is $\sqrt{R_v} = 0.005$. Figure 3 displays the tracking outcome (top), as well as the required set-point signal (bottom) synthesized from the above optimization problem.

# IV. Model 2 - Constraint Satisfaction Formulated via Satisfiability Modulo Theory

Here, we explore a different method for capturing the uncertainty in an individual TCL's temperature evolution and in bounding the aggregate power consumption of the TCL over a future time horizon given an input sequence. We quantify this bound by solving a feasibility problem whose constraints are determined from the uncertain dynamics of the TCLs and the historical data on population power consumption. Our results are preliminary. We show that the approach works well for deterministic systems but more work is needed to handle stochastic systems and heterogenous parameters.

The problem of interest here is to construct a controller that comes with guaranteed performance. In particular, given a sequence of observed power consumption and a known sequence of control inputs, we seek to choose an input to apply in the next time frame, such that the resulting overall power will be guaranteed to lie in an interval $[P^* - \epsilon, P^* + \epsilon]$ around a desired power $P^*$. Instead of developing an abstraction model and then quantifying its resulting error, we use observations of the population to formulate meaningful constraints on the future trajectory of the population given a control input.

We derive a model that is a continuous time abstraction, with the temperature interval of the TCL discretized into bins. We then track the movement of the upper and lower boundaries of each bin. The continuous (temperature) and discrete (On/Off mode) states of the bin boundary provide constraints on the continuous and discrete state of a TCL whose initial temperature is within the bin. To cope with the continuous dynamics and the uncertainty in the states of

individual TCLs in a given bin, we use the formalism of the satisfiability modulo theory (SMT). An SMT instance is a formula in first-order logic, and the problem is determining whether such a formula is satisfiable. SMT is a widely used method in computer science verification, and solvers such as iSAT [25] have been developed to automatically verify and analyze model properties given a set of initial conditions or control strategies. Using an SMT solver, we provide upper and lower bounds on the performances of a set of control strategies.

## IV.A. Model Description

We use the following continuous time dynamics of the TCL, which can be derived by letting $h \to 0$ from (1) and by removing the noise term $w$:

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

$$dm(t) = \begin{cases} -1, & \theta(t) < \theta_- + u(t) \\ 1, & \theta(t) > \theta_+ + u(t) \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

$$R(t) = \begin{cases} R_0, & N(t) \equiv 0 \bmod 2 \\ R_1, & N(t) \equiv 1 \bmod 2 \end{cases}$$

In the above, $u(t)$ is the control input, $R(t) \in \{R_0, R_1\}$ capture two different thermal resistance values (reflecting, for example, closed vs. open windows); and the switching times between the resistance values are distributed according to the homogeneous Poisson process $N_t$ with a specified rate $\lambda$. The corresponding dynamics are illustrated for a single household by a sample path in the left panel of Figure 4. Although we develop the approach for thermal resistances taking two potential values, the approach could be generalized to multiple values of thermal resistances, to account for various discrete changes in the room such as occupancy or opening/closing of entrances.

The reasoning for modeling stochasticity in this way is twofold. First, changes in thermal resistance due to changes in room occupancy, opening and closing of doors/windows may be more accurate than fluctuations in the temperature captured by the noise term used in Equation (1). Second, modeling the random influences as jumps in the thermal resistances allows for an event based simulation. Although the underlying model is formulated in continuous time, only discrete events have to be considered. Also, between two random events, the dynamics are deterministic. Therefore, we can draw exact samples from the continuous time model.

## IV.B. Abstraction Approach

Given that some of the parameters and the control strategy are not known beforehand, obtaining guaranteed bounds on the power consumption of the TCL population in future time

steps requires solving a hard optimization problem. To reduce the computational complexity we aim at a safe abstraction of the model in (7) as follows:

1) We divide the temperature interval into bins and count the number of TCLs within a bin. The bins are allowed to move along the temperature axis over time, in contrast to the previous approaches.

2) By introducing non-determinism to capture the unknown temperature of individual TCLs within a bin and uncertainty in switching dynamics of the thermal resistance $R(t)$ we safely over-approximate the effects of these variables.

Here, we are only interested in properties of the power consumption of the population. Thus, we adopt an event based time resolution within the abstraction approach by considering points in time at which the power can potentially change.

*a) Abstraction dynamics:* Mathematically, the dynamics of the model can be formulated as follows: Each bin $i$ is characterized by a tuple $(\theta_i^u, \theta_i^l, m_i^u, m_i^l, R_i)$ denoting the temperature of its upper boundary $\theta_i^u$, its lower boundary $\theta_i^l$ (indicated in Fig. 4, right panel, with blue lines). Each has an associated On/Off state $m_i^u, m_i^l$ and a resistance $R_i \in \{R_0, R_1\}$. A bin is defined for every combination of temperature range, discrete state $m^l = m^u \in \{0, 1\}$ and $R \in \{R_0, R_1\}$. The total number of bins is therefore given by $4 \cdot n_d$, where $n_d$ denotes the number of bins along the temperature axis. Initially, all TCLs within bin $i$ have a temperature between $\theta_i^u$ and $\theta_i^l$, start with a thermal resistance $R_i$ and are all in the same On/Off state, $m_i^l = m_i^u$.

The temporal evolution of the bounds of all bins can be calculated using the noiseless version of (7) and hence we can compute the first time any of the temperature bins hits one of the dead-band boundaries. These switching thresholds, $\theta_-$ and $\theta_+$ are denoted in Fig. 4 by black horizontal lines. Once one of the bin boundaries hits the threshold, the discrete state associated with the upper temperature bound $m_i^u$, and the one for the lower temperature bound $m_i^l$, start to differ indicating that not all TCLs within this bin need to have the same On/Off state. To indicate the temporal dependence of the $m$-values, we write $m_i^u(t_j), m_i^l(t_j)$ and $\theta_i^u(t_j), \theta_i^l(t_j)$ respectively.

The bin description above gives an over-approximation of the dynamical model in (7). That is, one can construct a sequence of On/Off states $\tilde{m}(t_j)$ which always fulfill $\min(m^u(t_j), m^l(t_j)) \leq \tilde{m}(t_j) \leq \max(m^u(t_j), m^l(t_j))$ but for which one cannot find a temperature sequence counterpart $\tilde{\theta}(t_j), \tilde{m}(t_j)$ which fulfills the dynamical constraints at the same time. On the other hand, each trajectory $\tilde{\theta}(t_j), \tilde{m}(t_j)$ fulfilling the dynamical constraints, will also fulfill $\min(m^u(t_j), m^l(t_j)) \leq \tilde{m}(t_j) \leq \max(m^u(t_j), m^l(t_j))$ for some $i$. Therefore, the abstraction gives a safe over-approximation of the system in (7).

If there are more than one bins for which $m^l(t_j) \neq m^u(t_j)$, we cannot determine exactly how many TCLs fall into these bins based on observations of the overall power consumption. Nevertheless, these variables define upper and lower bound constraints on the number of TCLs within each bin. Within a SAT-based approach this non-determinism has to be resolved by a solver which can decide how many TCLs to put in a bin to fulfill all constraints.

The noise process can also be tackled by introducing non-determinism. To this end, we can calculate the temporal evolution of each bin given its initial condition (represented by the tuple) assuming no noise. Given two bins with different initial $R$ values but potentially the same On/Off state, one can now check if the evolutions cross or overlap each other in the temperature axis. If so, it is possible for a TCL within one bin to jump to the other bin by changing its $R$ parameter. This procedure leads to a set of bins $M_i$ to which a TCL within bin $i$ can jump to due a change in its dynamics, in this case the resistance $R$. Computing this for all time $t_j$ leads to a sequence of sets $M_i(t_j)$ which indicate the possibility of a jump within the time-frame $[t_j, t_{j+1}]$.
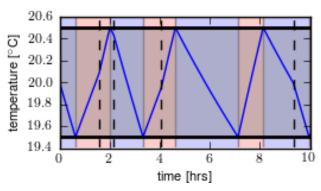
The computational load is increased by introducing non-determinism, that is, the solver has to decide how many TCLs to put in each bin based on the available constraints. However, the effort for handling multiple TCLs is drastically decreased since only the number of TCLs within a bin are determined.

## IV.C. Controller Verification Using the SMT framework

We consider a discrete time controller that acts on all TCLs by setting the same temperature set point $\theta_s$. For every given sequence of such a control signal, the quantities describing the bins can be calculated beforehand, which characterizes the behavior of a population of TCLs. The observed sequence of population power $P(t_k)$, can then be used to write down a set of constraints on the number of TCLs $n_i(t_k)$ within a bin $i$ at a given point in time $t_k$. More precisely, as we have upper and lower bounds $m_i^u(t_k), m_i^l(t_k)$, on the state of the TCLs within a given bin $i$, we have the following constraints on $n_i(t_k)$:

$$\sum_i m_i^l(t_k) n_i(t_k) \leq P(t_k) \leq \sum_i m_i^u(t_k) n_i(t_k) \qquad (8)$$

The number of TCLs within a bin changes over time as the individual TCLs may change their dynamics due to the probabilistic switches in the $R$ parameters. To capture this behavior in the constraints, we introduce $Q(i, j, t_k)$ to keep track of the number of TCLs jumping from bin $i$ to bin $j$ within the time-interval $[t_{k-1}, t_k]$. Thus, in addition to Constraint (8)
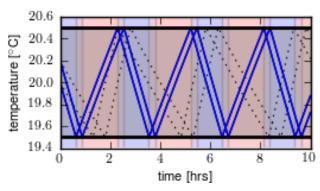
Figure 4: **Left:** In blue, a single sample path of the dynamics (7) is plotted. Dashed vertical lines indicate time points, at which the parameter $R$ changes. **Right:** The blue lines indicate the dynamics of a bin with $R = R_0$. Black dotted illustrate the dynamics with $R$ set to $R_1$. Overlapping red and blue regions indicate the time-intervals in which TCLs within this bin can be in either the On or the Off state.

we have the following two constraints:

$$n_j(t_k) = \sum_i Q(i,j,t_k) \qquad (9a)$$

$$Q(i,j,t_k) \le \sum_{i'} Q(i',i,t_{k-1})M(i,j,t_k), \qquad (9b)$$

where $M(i,j,t_k) \in \{0,1\}$ is a pre-computed deterministic quantity, which indicates whether it is possible for a TCL in bin $i$ to jump to bin $j$ within $[t_{k-1}, t_k]$ according to the noisy dynamics. Additionally, as there are $n_p$ number of TCLs distributed across the bins, we have the trivial constraint:

$$\sum_{i,j} Q(i,j,t_k) = n_p. \qquad (10)$$

All together, we can construct a set of constraints that give a safe over-approximation of the behavior of a population of TCLs, given an observed sequence of aggregate power consumption and a sequence of applied control inputs. To verify a controller, we check the feasibility of the constraints (8)-(10), with the following additional constraints capturing the performance requirements in the time step $t_k$:

$$P^u(t_k) := \sum_j m_j^u(t_k)n_j(t_k) \le P^* + \epsilon, \qquad (11a)$$

$$P^l(t_k) := \sum_j m_j^l(t_k)n_j(t_k) \ge P^* - \epsilon. \qquad (11b)$$

If all constraints are feasible, then the given control input is guaranteed to produce an aggregate power which is within $\epsilon$ bound of $P^*(t_k)$, the desired power at the next time step.

Under all control inputs which satisfy the performance guarantees, we can then either choose randomly or optimize with respect to further objectives. Note that the constraints do not consider the probability of the set of trajectories fulfilling the constraints but only characterize an over-approximation (superset) of the set of trajectories. Therefore, rare extreme situations are covered at the cost of potentially pessimistic guarantees.

## IV.D. Simulation of the Abstraction Procedure

To illustrate the approach, we simulated 200 TCLs. All parameters used for the simulation are summarized in Table I. The number of bins was set to 22. In Figure 5 upper panel the individual temperature trajectories as well as the control inputs in terms of the desired temperature set-point is shown. We observe the population power, shown in the lower panel of Figure 5, for over 18 time steps and use this information to construct a set of constraints for the next 3 minutes, as described above.

Figure 5 shows the result for the deterministic setting, that is $R = 2.0$ and is not changing. Given a set of control inputs, we determine which of them satisfy aggregate power bounds and if they satisfy a given bound $\epsilon$, what is the smallest bound $\epsilon'$ they can satisfy. We then implemented a controller that guarantees the least bound, that is least tracking error. The resulting guaranteed range for the chosen control inputs over the time horizon is shown as a gray region in the lower panel.

In principle, the approach can capture heterogeneity in parameters and stochasticity, e.g. switching $R$ values. In this case, the worst and best case evolution of the bin temperature boundaries lead to temperature ranges and discrete states of each bin upper and lower boundaries which drift apart from each other over time. As a result, the guaranteed predictions are pessimistic. Because of too much conservatism, this approach does not work well for the heterogenous or stochastic models. To address this conservatism, one can use more bins, resulting in increased computational load. In particular, the approach scales linearly in terms of the observed history length and scales quadratically in terms of the number of bins.

## V. Model 3 - TCL Population Model for Analyzing Arbitrage Potential

The objective of this section is to understand if non-disruptive direct load control of TCL aggregations could be used to
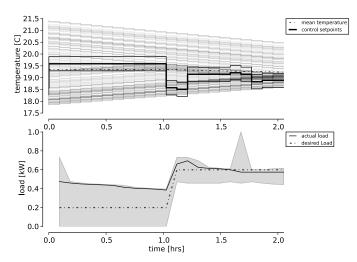
Figure 5: A set of 200 TCLs is simulated and controlled by an input with a guaranteed performance. Individual temperature dynamics are shown in the top panel, whereas the bottom panel shows the aggregate power. The interval around the desired load which can be guaranteed for the determined control inputs are plotted in the lower panel as the gray shaded region.

arbitrage intra-hour electricity market prices. Given that price variability is on the orders of minutes, to address this problem we need models that allow us to optimize the power consumption of TCL aggregations over a horizon of minutes to hours. Most models, for example, the Markov chain model, were developed for short prediction horizons during which the ambient temperature remains approximately constant. Here, we need to account for time-varying temperature and longer prediction horizons. The results in this section are an extension of the work presented in [9].

The energy arbitrage problem has been investigated by a number of researchers [26], [27], [28], [29], [30], [31]. However, past research has not taken into account the specific capabilities and constraints of TCL aggregations. Recently, we developed an aggregate model of a heterogenous TCL population and used this model to derive practically-feasible upper bounds on the amount of money that TCL aggregations could save through energy arbitrage in 5-minute energy markets [9]. We assumed that an aggregator rather than individual TCLs, arbitrages prices, and that the aggregator sends control signals to individual TCLs based on aggregate models and aggregate measurements only; he does not have access to individual TCL parameters or states. Here, we extend our analysis by comparing our previous results to those generated for the case when each individual TCLs arbitrages prices. This gives us actual upper bounds for the TCL energy arbitrage problem and helps us understand the value of information. In both cases, we assume control is via on/off switching, not temperature set point adjustment. Moreover, we assume that TCLs can only be controlled when they are within their dead-band, which ensures that our control is non-disruptive to the end users.

## V.A. Arbitrage for an Individual TCL

We consider the case in which each TCL optimizes its power consumption given a forecasted price signal. Consider the individual TCL model (1). We define the control for each TCL to be $u \in \{0, 1\}$, where 0 turns a TCL off and 1 turns a TCL on. The discrete variable $m$ is now updated as follows:

$$m(t+1) = \begin{cases} 0, & \theta(t+1) < \theta_- \text{ or} \\ & u(t) = 0 \wedge \theta(t+1) \in [\theta_-, \theta_+] \\ 1, & \theta(t+1) > \theta_+ \text{ or} \\ & u(t) = 1 \wedge \theta(t+1) \in [\theta_-, \theta_+] \\ m(t), & \text{otherwise} \end{cases} \quad (12)$$

Let $l(t)$ be the cost of energy at time step $t$ and $N$ be the prediction horizon. The arbitrage problem for one TCL can be written as:

$$\min_{u \in U} h \sum_{t=t_0}^{t_0+N} l(t) m(t) \bar{P}_{rate} \quad (13)$$

$$\text{s.t. (1) and (12)},$$

where $U = \{0, 1\}^N$ and as a reminder $\bar{P}_{rate} = P_{rate}/\eta$.

The above optimization problem can be solved using Dynamic Programming (DP). We tackle the DP problem by discretizing the state space. Each state represents both a specific temperature interval within or just outside the dead-band and the On/Off state. We then precompute all possible state transitions for each input and a range of outdoor air temperatures. With this information we can compute the optimal policy given forecasts for outdoor air temperatures and electricity prices.

In general, the cost of each state at time $t$ is $m(t-1)\bar{P}_{rate}$ since a TCL switches only at the end of each time step. To formulate the DP, we need to assign a cost to each state that is not a function of time. The cost of states within the dead-band is simply $m\bar{P}_{rate}$; however, the cost of states just outside of the dead-band is a function of whether the TCL has just switched or not. Here, we assume that if a TCL is outside of the dead-band, it has just switched and its cost is $(1 - m)\bar{P}_{rate}$. We can choose the discretization step so that this is true nearly all of the time. However, sometimes this may not be the case. To solve this problem, one could introduce virtual states that capture not only temperature and on/Off state but recent switching history, but we leave this to future work.

An important consideration when picking the discretization step is that TCLs move at different speeds at different outdoor air temperatures. Therefore, a good discretization step for a high temperature may not be good at relatively low temperatures. For example, if the discretization step is too large and the outdoor air temperature is just above the dead-band then TCLs in the Off state may move so slowly that they do not switch bins in each time step. Decreasing the discretization steps alleviates this problem but can lead to others including numerical issues and divergence from the DP cost approximation described above. In sum, the DP works
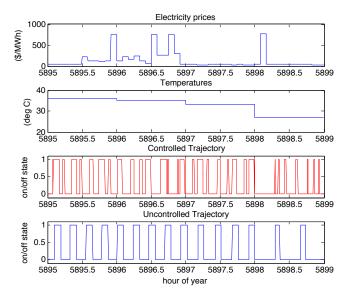
Figure 6: Example of individual TCL arbitrage results.



Figure 7: Histogram of energy cost savings for individual TCL arbitrage.

trol responses by sending broadcast control signals to TCLs.

## V.C. Aggregate Thermal Battery Model

For a population of TCLs, the previous optimization problem can be solved by formulating the cost function as the aggregate of costs of the individual TCLs: $h \sum_{t=t_0}^{t_0+N} \sum_{i=1}^{n_p} l(t) m^i(t) \bar{P}_{rate}^i$. This approach is impractical for two reasons. First, the aggregator would need to know all of the individual TCL parameters and states. Second, the aggregator would need communication links with each individual TCL to send individual on/off control signals. Here we assume that the aggregator only has access to aggregate system parameters and measurements and that he can only broadcast control vectors to TCL populations, as in [7].

To address the aggregate optimization, we initially investigated use of a modified version of the Markov chain model described in the first section; however, we found that it is unsuitable for describing the behavior of the aggregate system when it is repeatedly pushed to its constraints, as is done in arbitrage [9]. Therefore, we propose modeling the TCL population as a time-varying thermal battery, described more fully in [9]. This model keeps track of a TCL population's energy state, $S(k)$, as a function of its mean aggregate power usage, $\bar{P}_{agg}$, in each price interval, $[t_k, t_{k+1}]$, of width $\Delta T$. The energy state of the population is defined similar to a battery's state of charge; it describes how full an energy storage unit is. We can obtain a difference equation for the evolution of the energy state as well as upper and lower envelopes of achievable power and energy for a population of TCLs.

Without external control, a TCL population's time-varying power trajectory is referred to as its "baseline." Figure 8 shows an air conditioner population's mean aggregate power baseline, $\bar{P}_{agg, baseline}$, over a day. A TCL population increases it energy state when $\bar{P}_{agg}(k) > \bar{P}_{agg, baseline}(k)$, and decreases it when $\bar{P}_{agg}(k) < \bar{P}_{agg, baseline}(k)$:

$$S(k+1) = S(k) + (\bar{P}_{agg}(k) - \bar{P}_{agg, baseline}(k))\Delta T. \quad (14)$$

## V.B. Results of Individual Optimization

We consider a 10 hour period with highly volatile prices (from California ISO node MERCED_1_N001 [32]) and high outdoor air temperature (from NOAA weather station Merced 23 WSW [33]). We use a population of $n_p = 1,000$ central air conditioners parameterized with the heterogeneous parameters in Table I, and we assume $w(t) = 0 \, \forall t$ in order to get an upper bound on the savings. For each TCL, we divide its dead-band into 100 temperature intervals and so we end up with 200 within-dead-band states, and we use the same discretization step for states just outside the dead-band. We precompute all bin transitions for only integer values of temperatures and use this as a look-up table. We compute the optimal control policy based on the discretized system and perfect price and temperature forecasts, and apply that to the TCL model of (1) and (12). Figure 6 shows an example uncontrolled and controlled trajectory for one TCL. We find that the population saves about 28% of its total energy costs, while individual TCLs save -15% to over 55% of their individual energy costs (Figure 7). Differences in savings results from different thermal parameters, initial conditions, and DP model accuracy.

Since we have assumed a deterministic system and perfect forecasts, the analysis here provides upper bounds on energy cost savings through arbitrage for individual TCLs. In reality, the individual TCLs may not have direct access to time-varying price signals or may not have the local computational capabilities to do the optimization. Therefore, we consider a more realistic scenario in which a load aggregator uses aggregate models, aggregate measurements, and forecasts to compute optimal control trajectories and then coordinate con-
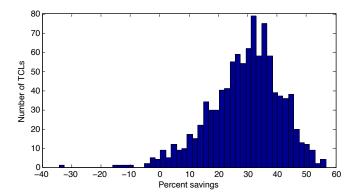
much better if the outdoor temperature is well above the dead-band temperatures.
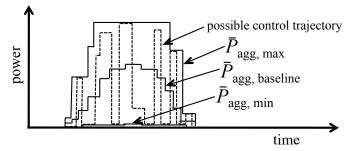
Figure 8: A TCL population's baseline and power constraints. The TCL population also has energy constraints, not shown.
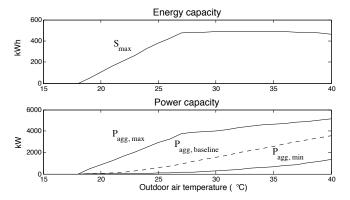


Figure 9: Energy and power capacity of a population of air conditioners as a function of outdoor air temperature.

As shown in Fig. 8, the choice of $\bar{P}_{\text{agg}}(k)$ is constrained:

$$\bar{P}_{\text{agg, min}}(k) \leq \bar{P}_{\text{agg}}(k) \leq \bar{P}_{\text{agg, max}}(k). \tag{15}$$

The energy state is also constrained at each time step:

$$0 \leq S(k) \leq S_{\text{max}}(k). \tag{16}$$

These bounds define the *power and energy capacity* of a TCL population. When $S = 0$ the thermal battery is depleted meaning all TCLs operate at one edge of the dead-band (e.g., for cooling TCLs all operate near $\theta_{\text{set}} + \delta/2$). When $S = S_{\text{max}}$, the thermal battery is full meaning all TCLs operate at the other edge of the dead-band.

To use this model, we need to derive or identify the time-varying parameters: $\bar{P}_{\text{agg, baseline}}$, $\bar{P}_{\text{agg, min}}$, $\bar{P}_{\text{agg, max}}$, and $S_{\text{max}}$. These parameters are a function of ambient temperature dynamics, but for simplicity we assume that each is simply a function of the current ambient temperature, $\theta_{\text{a}}$. Additionally, we assume that each belongs to a finite set of values and develop a look-up table that specifies an estimate of each value as a function of $\theta_{\text{a}}$. In recent work, we described the procedure for computation and estimation of $\bar{P}_{\text{agg, baseline}}$, $S_{\text{max}}$, $\bar{P}_{\text{agg, min}}$ and $\bar{P}_{\text{agg, max}}$ through system identification [9]. The identified parameters, as a function of the current outdoor air temperature, are shown in Fig. 9.

Given the thermal battery model and the price and outdoor air temperature forecasts over a horizon, we aim to determine the optimal mean aggregate power consumption in each interval, $\bar{P}_{\text{agg}}^*$, and so we solve:

$$\min \Delta T \sum_{k=t_0}^{t_0+N} l(k)\bar{P}_{\text{agg}}(k) \tag{17}$$

$$\text{s.t.} \ (14), (15), \ \text{and} \ (16).$$

The above can be solved as a receding-horizon Linear Program (LP). We then transform $\bar{P}_{\text{agg}}^*$ into a control trajectory $p_{\text{agg}}^*$: $p_{\text{agg}}^*(t) = \bar{P}_{\text{agg}}^*(k)$ for $t = k, k+h, ..., k+\Delta T - h$, to be tracked by the TCL population. Thus, we use the less accurate time-varying thermal battery model for the purpose of optimization over a long time horizon, while we use the Markov chain model for controlling TCLs to track the power output.

In order to track the trajectory $p_{\text{agg}}^*(t)$, we calculate $u_{\text{goal}}$, the total fraction of TCLs to switch on or off in the next time step. To achieve this, we use an extension of the Markov chain model of the heterogeneous population with a predictive proportional controller (PPC) [7]. The Markov chain model is extended in order to address large ambient temperature variations over long horizons. The extensions include addition of extra bins outside the temperature dead-band and identification of the $P$ matrix in (3) for every discretized ambient temperature. Here, the stochastic noise term $\mathbf{W}$ is not explicitly considered in control synthesis.

To design the control, first, we compute:

$$u'_{\text{goal}}(t) = \frac{p_{\text{agg}^*}(t+1) - y(t+1)}{N_p \bar{P}_{\text{rate}}}, \tag{18}$$

where $y(t + 1)$ is the predicted power output of the TCL population given the Markov chain abstraction as defined in Equation (3). Then, $u_{\text{goal}}(t)$ is calculated by putting $u'_{\text{goal}}(t)$ through a saturation filter with minimum equal to the fraction of TCLs on, and maximum equal to the fraction of TCLs off. We then distribute $u_{\text{goal}}$ to the bins and, for each bin, divide the absolute fraction of TCLs to switch by the measured or estimated fraction of TCLs in the bin to determine the switch probability. Switch probabilities are broadcast to the TCLs and then TCLs switch or not based on the switch probability associated with the bin they are in. Note that $u_{\text{goal}}$ can be distributed to the bins in different ways, for example, equally or by preferentially switching TCLs that are about to switch. Here we do the latter, so that the controller would preferentially switch TCLs in bins closer to the dead-band and thus to natural switching. This helps minimize the chance of compressor short-cycling.

## V.D. Results of Aggregate TCL Optimization

We use the same population of central air conditioners as in the individual TCL arbitrage population, and consider the same 10 hour period. However, here we use a stochastic individual TCL model (1) and assume the noise standard deviation is $5 \times 10^{-4}$. For implementation of control, we use 42 bins in
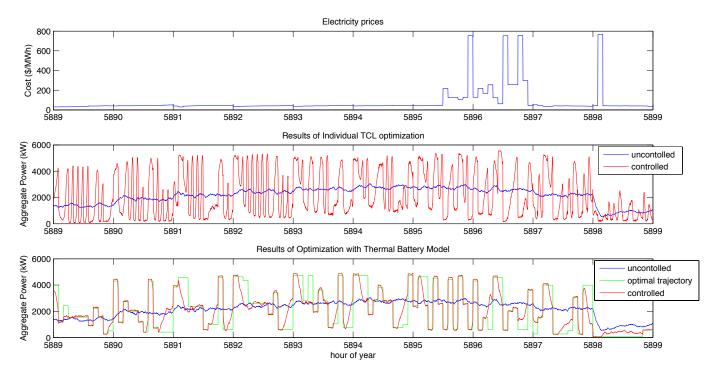
Figure 10: Comparison of optimization by individual TCLs and with the thermal battery model.

the Markov chain: 40 bins within the dead-band and 2 bins to capture temperatures just outside the dead-band into which TCLs switch. Again, we assume perfect price and temperature forecasts. We identified $\bar{P}_{\text{agg, baseline}}$, $\bar{P}_{\text{agg, min}}$, $\bar{P}_{\text{agg, max}}$, $S_{\text{max}}$, and $P$ for integer values of temperature using the system identification methods described in [9].

First, in Fig. 10, we compare the individual TCL optimization results to the aggregate population optimization results. In the individual optimization, each TCL in the population optimizes its power consumption using the DP approach described above. In the aggregate population optimization, the optimal aggregate power trajectory is found using the thermal battery model and the LP optimization. Then, the TCL population tracks this desired power output using the on/off broadcast control described in the previous section. The aggregate results indicate savings of 21% in contrast to the 28% savings given by the individual optimization. In addition to apparent inaccuracy in the aggregate model compared to individual TCL model, a reason for the decreased savings is that the battery storage model based on the 5-minute discretization step results in a piecewise constant power trajectory, which may not allow the TCL population sufficient flexibility. In fact, tracking this trajectory is difficult for the TCL population, requiring significant TCL switching. Future work will explore other methods for transforming the output of the LP into more suitable control trajectories.

To address the potential for TCL populations to arbitrage energy prices using the aggregate population optimization framework, we performed one optimization each day, i.e. $N = 24$ hours for one full year (2010) of the same data

source. Our results were discussed in detail in our recent publication [9], and here we summarize them. The results of the optimization problem predict that a population of air conditioners in Merced, CA, USA could save, at most, 17% in yearly wholesale energy cost through arbitrage in CAISO's 5-minute energy market. When we control the population to track the optimal trajectory, we find that the maximum savings are closer to 14%, specifically the uncontrolled population would have spent about \$91,500 for energy during the year while the same population doing energy arbitrage would have spent \$78,400. This translates to about \$13 in wholesale energy cost savings per TCL per year. Since this analysis assumes perfect price and weather forecasts and exogenous electricity prices, this is an upper bound on the potential practical energy costs savings in Merced, assuming future prices and price volatility are similar to those in the past.

## VI. Discussion of the Approaches

The formal Markov chain abstraction in the first section provides insight into the previously proposed models [17], [18], [19], [20], [7]. Its basis is dividing the continuous state into bins and associating a Markov chain state to each bin for each discrete on/off mode. The problem is then characterizing the evolution of the fraction of TCLs in each bin. The contribution here is to provide exact characterization of the evolution of the TCL fractions in each bin for the homogenous system. The mean of the distribution is consistent with previously derived results [7]. The method also exactly characterizes the covariance of the distribution. In addition, it shows that as the number of TCLs grow, the evolution of the distribution

approaches a linear system with additive Gaussian noise and the noise covariance is exactly characterized. This analysis leads to derivation of error bounds on the power consumption of the population and that predicted by the model. The error bound is a function of the number and length of the bins. Thus, it provides a method to tune the discretization in order to achieve a desired error behavior as well providing worst-case analysis of the error in the abstraction. The analysis is extended to a heterogeneous population of TCLs in which the model parameters for each TCL are drawn from a known distribution. The Markov chain abstraction can be used in a receding horizon approach to find the temperature set-point variation required for tracking a desired power. In addition, the linear Gaussian formulation of the dynamics leads itself naturally to Kalman filter estimation approach to estimate the state of the system given noisy measurements [7]. Given that the prediction error grows as the time horizon grows, the model although useful for optimal tracking and estimation over short time steps, has limitations for longer time steps.

The approach in Model 2 provides a fresh look at the problem of TCL population analysis. In particular, the approach is not based on developing a dynamical model of the TCL population. Rather, it uses observations of the TCL population power consumption to define bounds on the number of the TCLs within each bins, without solving for the exact values of these numbers. Then, the analysis problem of whether a given power trajectory can be tracked with a desired accuracy is cast as a feasibility problem, with the additional constraints given by the bounds on the number of TCLs in each bin defined by the historic data of population power. This particular line of analysis is important in applications in which the TCL population need to provide guarantees on the power tracking performance, such as ancillary services participation. Thus, the approach proposed could determine if there exists any feasible controller which satisfies ancillary service requirements. In theory, the model can address any form of uncertainty in the dynamics, such as uncertainty in the resistance values of the TCLs, through introducing non-determinism. The additional non-determinism maps to additional constraints in the feasibility problem under consideration. However, the limitation of the model is that as uncertainty dimension grows, the constraints lead to a very conservative approach which quickly leads to infeasibility of the problem even though in practice a feasible controller may exist. Thus, at this point, the model is not able to handle parameter heterogeneity and stochasticity.

Model 3 was developed with the aim of capturing heterogenous TCL behavior over long time horizons so that we could study the potential for TCLs to arbitrage energy prices in a realistic setting. The model needs to provide predictability of power consumption over horizons of minutes to hours in order to be able to take advantage of the temporal price differences which are of the same order of time magnitude. In order to deal with the above two issues, time-varying lower and upper bounds on achievable power and energy by a TCL population are derived. The power trajectory is designed to minimize energy costs over a prediction horizon while satisfying these bounds. The resulting optimal trajectory can be tracked by the TCLs by broadcasting control signals to the population that cause TCLs to switch on/off probabilistically. The Markov chain abstraction is used to develop the controller. We find that this approach leads to about 75% of the savings achievable with the fully optimal approach of individual TCL arbitrage. The main limitation of this model is the fact that it is difficult to characterize the power and energy capacities. Here, we approximate them through system identification. More in depth analysis of these quantities are subject of current investigation.

## VII. Conclusions

We proposed three different modeling approaches to analyze and control populations of TCLs. The development of Model 1 provides insight into the dynamics of the TCL population and model error. Model 2, on the other hand, explores the potential of verifying TCL population performance without resorting to any dynamical modeling of the population. Model 3 provides a practical approach to quantifying the potential of TCLs to participate in energy markets over long time horizons (minutes to hours). Based on this work, we propose several directions for future work. In terms of Model 1, the derivation of tight error bounds for heterogeneous populations with time-varying ambient temperature are important next steps. In terms of Model 2, formulation of constraints so that heterogeneity and stochasticity in parameters lead to not overly conservative constraints is a subject of current exploration. As for Model 3, we are working on better methods to identify the parameters of the model and characterize their uncertainty.

The potential of TCLs to serve in demand response programs or ancillary service markets is highly dependent on the ability to accurately predict and control their power consumption while ensuring users' comfort and the physical limitation of the devices such as compressor short cycling. As seen in this exploratory study, the modeling and analysis of the population dynamics is an interesting and rich problem. In addition to improvements to our proposed models, there is a potential to explore various other modeling, analysis, and control techniques for large scale stochastic hybrid systems. Given optimal control and prediction of the TCL population, ultimately, an aggregator or utility company would need to decide if arbitrage or ancillary service revenues could be sufficient to cover upfront costs including hardware, software, and installation; reoccurring costs including operations, maintenance, and incentive payments to customers; and its desired profit margin. Thus, it is desirable to provide as accurate predication and control of the population as possible using limited communication and measurement infrastructure.

## VIII. Acknowledgement

# References

[1] 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*, 50:1389–1400, 2009.

[2] Y.V. Makarov, C. Loutan, J. Ma, and P. De Mello. Operational impacts of wind generation on california power systems. *IEEE Transactions on Power Systems*, 24(2):1039–1050, 2009.

[3] D.A. Halamay, T.K.A. Brekken, A. Simmons, and S. McArthur. Reserve requirement impacts of large-scale integration of wind, solar, and ocean wave power generation. *IEEE Transactions on Sustainable Energy*, 2(3):321–328, Jul. 2011.

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

[5] D.S. Callaway. Can smaller loads be profitably engaged in power system services? In *Proceedings of IEEE Power and Energy Society General Meeting*, Detroit, MI, 2011.

[6] J.L. Mathieu, M. Dyson, and D.S. Callaway. Using residential electric loads for fast demand response: The potential resource and revenues, the costs, and policy recommendations. In *Proceedings of the ACEEE Summer Study on Buildings*, Pacific Grove, CA, August 2012.

[7] 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.

[8] S. Esmaeil Zadeh Soudjani and A. Abate. Aggregation of thermostatically controlled loads by formal abstractions. In *European Control Conference*, Zurich, Switzerland, 2013.

[9] J. L. Mathieu, M. Kamgarpour, J. Lygeros, and D. S. Callaway. Energy arbitrage with thermostatically controlled loads. In *European Control Conference*, Zurich, Switzerland, 2013.

[10] 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.

[11] S. Ihara and F.C. Schweppe. Physically based modeling of cold load pickup. *IEEE Transactions on Power Apparatus and Systems*, PAS-100(9):4142–4150, 1981.

[12] R.E. Mortensen and K.P. Haggerty. Dynamics of heating and cooling loads: models, simulation, and actual utility data. *IEEE Transactions on Power Systems*, 5(1):243–249, 1990.

[13] C. Uçak and Çağlar. The effect of load parameter dispersion and direct load control actions on aggregated load. In *Proceedings of POWERCON '98*, volume 1, pages 280–284, 1998.

[14] 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, September 1985.

[15] N. Lu and D.P. Chassin. A state queuing model of thermostatically controlled appliances. *IEEE Transactions on Power Systems*, 19(3):1666–1673, 2004.

[16] N. Lu, D.P. Chassin, and S.E. Widergren. Modeling uncertainties in aggregated thermostatically controlled loads using a state queueing model. *IEEE Transactions on Power Systems*, 20(2):725–733, 2005.

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

[18] S. Kundu, N. Sinitsyn, S. Backhaus, , and I. Hiskens. Modeling and control of thermostatically controlled loads. In *Proceedings of $17^{th}$ Power Systems Computation Conference*, Stockholm, Sweden, Aug 2011.

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

[20] J.L. Mathieu and D.S. Callaway. State estimation and control of heterogeneous thermostatically controlled loads for load following. In *Proceedings of the Hawaii International Conference on Systems Science*, pages 2002–2011, Wailea, HI, January 2012.

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

[22] 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.

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

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

[25] M. Fränzle, C. Herde, T. Teige, S. Ratschan, and T. Schubert. Efficient solving of large non-linear arithmetic constraint systems with complex boolean structure. *Journal on Satisfiability, Boolean Modeling and Computation*, 1(3-4):209–236, 2007.

[26] S.R. Vosen and J.O. Keller. Hybrid energy storage systems for stand-alone electric energy stems: optimization of system performance and cost through control strategies. *International Journal of Hydrogen Energy*, (24):1139–1156, 1999.

[27] R. Walawalkar, J. Apt, and R. Mancini. Economics of electric energy storage for energy arbitrage and regulation in new york. *Energy Policy*, (35):2558–2568, 2007.

[28] P. Denholm and R. Sioshansi. The value of compressed air energy storage with win in transmission constrained electric power systems. *Energy Policy*, (37):3149–3158, 2009.

[29] R. Sioshansi, P. Denholm, T. Jenkin, and J. Weiss. Estimating the value of electricity storage in PJM: Arbitrage and some welfare effects. *Energy Economics*, (31):269–277, 2009.

[30] V. Marano, G. Rizzo, and F.A. Tiano. Application of dynamic programming to the optimal management of a hybrid power plant with wind turbines, photovoltaic panels and compressed air energy storage. *Applied Energy (in press)*, 2012.

[31] J.A. Taylor, D.S. Callaway, and K. Poolla. Competitive energy storage in the presence of renewables. *IEEE Transactions on Power Systems*, 28(2):985–996, 2013.

[32] CAISO. Open access same-time information system site. California Independent System Operator, 2012.

[33] NOAA. National climatic data center: US climate reference network data set: Merced 23 WSW. National Oceanic and Atmospheric Administration, 2010.