Preprints of the 19th World Congress
The International Federation of Automatic Control
Cape Town, South Africa. August 24-29, 2014

Demand Response of Thermostatic Loads by Optimized Switching-Fraction Broadcast ⋆

Luminita C. Totu, Rafael Wisniewski

Department of Electronic Systems, Faculty of Engineering and Science
Aalborg University, 9000 Aalborg Denmark (e-mail: lct,raf@es.aau.dk)

Abstract: Demand response is an important Smart Grid concept that aims at facilitating the integration of volatile energy resources into the electricity grid. This paper considers the problem of managing large populations of thermostat-based devices with on/off operation. The objective is to enable demand response capabilities within the intrinsic flexibility of the population. A temperature distribution model based on Fokker-Planck partial differential equations is used to capture the behavior of the population. To ensure probability conservation and high accuracy of the numerical solution, Finite Volume Method is used to spatially discretize these equations. Next, a broadcast strategy with two switching-fraction signals is proposed for actuating the population. This is applied in an open-loop scenario for tracking a power reference by running an optimization with a multilinear objective.

Keywords: Smart grids, thermostatic loads, stochastic hybrid systems, optimization

1. INTRODUCTION

We consider a large group of devices of the same type, where each unit has on/off power consumption controlled by an internal thermostat. These devices are of special interest for Smart Grid scenarios since they have the potential to deliver demand response services in an automated way. The focus of this work is on aggregating a very large numbers of units. The challenge is twofold. Firstly, communication flows must be carefully designed to be feasible under cost, security and privacy criteria because of the large scale and large geographic spread of the system. Secondly, computational complexity must be kept in check.

The study of large groups of thermostatic loads started in the power system and in the control literature in the '80s with the works of Ihara and Schweppe (1981) and Malhame and Chong (1985). The interest was on modeling the oscillatory effects in the power consumption after a planned (direct load control) or unplanned interruption (black-out). After three decades, the problem got a resurgence motivated by the advancement of demand response concepts, which in turn are motivated by new to challenges in power system operations, mainly the integration of intermittent generation.

The modeling approach that we are focused on is based on physical principles and stochastics. The thermostatic unit is described by a stochastic hybrid lumped-state dynamical model, while a thermostatic population can be described by a distribution function over the hybrid state-space.

If the distribution function is taken in its continuous form (a density), an infinite dimensional state-space description is obtained where dynamics are given in the form of a Partial Differential Equation (PDE) system and boundary conditions. This form was first derived in Malhame and Chong (1985), and has been subsequently used in e.g. Callaway (2009), Bashash and Fathy (2013), Moura et al. (2013). Another approach is to divide the hybrid state-space into a finite number of partitions, and work with a discrete distribution i.e. probability mass. Careful considerations based on the unit model can be made to derive transition probabilities from one partition to another, and thus the dynamics can be expressed in form of a Markov chain, e.g. in Koch et al. (2011), Zhang et al. (2013), Soudjani and Abate (2013). These two forms are essentially equivalent, leading to a standard, linear state-space description.

However, in terms of numerical performance not all methods all equivalent. For example, in Bashash and Fathy (2013) a finite-difference numerical scheme is used to obtain the Markov chain representation from the PDE description. This method does not inherently conserve probability and the solution might require periodic rescaling. In Koch et al. (2011) the transition probabilities of the Markov chain are derived based on an uniformity approximation over each partition. In this article, we propose the Finite Volume Method (FVM) for obtaining the Markov chain transition matrix from the PDE system. The main advantage is that the overall probability is guaranteed to be conserved and the resulting dynamic matrix has a proper rate transition form, i.e., columns that sum to zero. Furthermore, numerical results have a high degree of accuracy also over long horizons, since the PDE form is exact and the only errors are related to the size of the state-space partitions. Finally, there is an advantage in using a well-developed framework with recognized robustness and performance as opposed to custom solutions.

⋆ This work has been carried out as part of the "Smart & Cool" project financed by the Southern Denmark Growth Forum and the European Regional Development Fund.

Other important aspects for population modeling are heterogeneity and minimum on/off times. Heterogeneity is a difficult problem because exact descriptions suffer from the "curse of dimensionality". We add a few remarks, but do not address directly the modeling of heterogeneity. Nonetheless, results show that the used control appears to have good robustness to heterogeneity. The effects of enforcing minimum on/off times at the unit level are modeled using a technique essentially equivalent to Zhang et al. (2013).

For control, our focus is on broadcasting strategies since we believe that these have an implementation advantage. In particular, we are interested in a particular form of toggling control (Koch et al., 2011) that involves the broadcast of two switching fractions. This is a randomized method of actuation, and has been introduced in Zhang et al. (2013) and Totu et al. (2013) in a closed-loop form where just one switching fraction is used at a time. In this work, we set-up and analyze an open-loop configuration based on a non-convex, predictive horizon optimization.

Section 2 presents the modeling used for the individual Thermostatically Controlled Load (TCL) and for the population, Section 3 introduces the randomized broadcast actuation with the switching fractions, Section 4 sets up the optimization formulation and presents numerical results, and Section 5 concludes.

2. TCL MODELING

For this work, we consider cooling units, in particular domestic refrigerators. As under realistic conditions, the TCLs are independent of each other, do not communicate nor share states. A main object of interest is the aggregated power consumption, which is simply the sum of the individual power consumptions.

2.1 TCL unit model

Stochastic Hybrid Unit Model The basic model for a TCL is a hybrid dynamical system with two modes, corresponding to the "on" or "off" state of the vapor-compression cooling cycle. When the TCL is "on", it is consuming power and the temperature in the food storage compartment is lowering. When the TCL is "off", it is not consuming power and the temperature in the food storage compartment is rising due to ambient conditions. The heating and cooling processes are modeled with a first-order dynamic. This is similar to Ihara and Schweppe (1981), Callaway (2009), Bashash and Fathy (2013) and others. Although more complex, second-order dynamics should be studied for air-conditioning or heat-pumps TCLs (Zhang et al., 2013), this is not considered necessary in the case of domestic refrigerators, because there is no dominant secondary temperature mass. Additionally, random fluctuation are introduced in the dynamics as a white noise term.

Other possible random disturbances not pursued at this time are jump processes that would correspond to "door-opening" events. Furthermore, power consumption is considered to be constant when the mode is "on". This assumption might need to be revisited in future work.

Summing up, the model of a TCL unit is a stochastic hybrid system of the following form:

$$\begin{cases} dT(t) &= -\dfrac{\mathrm{UA}}{C}\left(T(t) - T_a + m(t)\dfrac{\eta W}{\mathrm{UA}}\right)dt + \sigma dw \\ &= \Big(aT(t) + b + m(t)c\Big)dt + \sigma dw \\ y(t) &= dm(t) \end{cases} \tag{1}$$

where $T(t) \in \mathbb{R}$ is the continuously-valued temperature state, $m(t) \in \{0, 1\}$ is the discrete-valued state corresponding to the "off" and "on" modes respectively, $y \in \mathbb{R}_+$ is the power consumption viewed here as model output, and a, b, c, d and σ are time invariant coefficients.

The dynamics of the discrete-valued state, m_i, are given by

$$m(t^+) = \begin{cases} 1, & T(t) \geq T_{\max} \\ m(t^-), & T(t) \in (T_{\min}, T_{\max}) \\ 0, & T(t) \leq T_{\min} \end{cases}, \tag{2}$$

a standard thermostat mechanism with boundaries at $T_{\min}$ and $T_{\max}$. This is a deterministic state-dependent switching.

Probabilistic unit model Because of the random influences introduced in the continuous dynamics, the hybrid state $\big(T(t), m(t)\big) \in \mathbb{R} \times \{0, 1\}$ is a stochastic variable. It can be characterized by a probability density function (pdf) defined in the following way:

$$f_y(x, t) \triangleq \Pr[T(t) \in [x, x + \delta x) \wedge m(t) = y] \tag{3}$$

The Fokker-Planck equation The temperature dynamics corresponding to each of the two modes (on and off) are continuous-time continuous-state Markov processes, and in particular diffusions. If we look at any of the modes in isolation and do not take the switches into consideration, given the probability distribution of the temperature state at time t_0 we can determine the probability distribution at any future time $t \geq t_0$ using the Fokker-Planck (Forward Kolmogorov) equation, which can be seen as a transport and conservation law for probability.

PDE system for the hybrid dynamic The result in Malhame and Chong (1985) gives the pdf dynamics in the particular case of the TCL hybrid dynamic (1),(2). It consists of Fokker-Planck equations written for each mode (4), and a set of boundary conditions (5).

Before stating the result, some preliminaries need to be addressed. The temperature domain needs to be divided into three subsets: the thermostat range $[T_{\min}, T_{\max}]$, the left hand side $[-\infty, T_{\min})$ and the right hand-side $(T_{\max}, \infty)$. This is a natural partition with respect to the operation of the TCL unit, and is necessary because boundary conditions apply in the points $T_{\min}$ and $T_{\max}$, and also because the pdf is not x-differentiable here. It is convenient to denote the three subsets with the letters b, a and c respectively. It is also important to note that in normal operation the pdf corresponding to the off mode, $f_0(x, t)$, is zero-valued on the c domain, because if the temperature becomes greater than $T_{\max}$ the thermostat mechanism ensures that the mode can not remain "off". The c domain accounts only for units in state "off", whose temperature becomes greater than $T_{\max}$ due to diffusion

effects. Similarly, the pdf corresponding to the on-mode, $f_1(x,t)$ is zero-valued on the a-domain. In numerical work, the infinity domains limits can be cut short since it is realistic to assume that the temperature inside a working refrigerator cannot drop below some T_{minmin} value and cannot rise above some T_{maxmax} value.

$$\begin{cases} \dfrac{\partial f_{0j}}{\partial t} + \dfrac{\partial}{\partial x}\Big((ax+b)f_{0j}\Big) = \dfrac{\sigma^2}{2}\dfrac{\partial^2 f_{0j}}{\partial^2 x}, j \in \{a,b\} \\ \dfrac{\partial f_{1j}}{\partial t} + \dfrac{\partial}{\partial x}\Big((ax+b+c)f_{1j}\Big) = \dfrac{\sigma^2}{2}\dfrac{\partial^2 f_{1j}}{\partial^2 x}, j \in \{b,c\} \end{cases} \tag{4}$$

$$\begin{cases} f_{1b}(T_{\min},t) = 0, f_{0b}(T_{\max},t) = 0 \\ f_{0a}(T_{\text{minmin}},t) = 0, f_{1c}(T_{\text{maxmax}},t) = 0 \\ f_{0b}(T_{\min},t) = f_{0a}(T_{\min},t), f_{1b}(T_{\max},t) = f_{1c}(T_{\max},t) \\ \dfrac{\partial}{\partial x} f_{1b}(T_{\min},t) = \dfrac{\partial}{\partial x} f_{0b}(T_{\min},t) + \dfrac{\partial}{\partial x} f_{0a}(T_{\min},t) \\ \dfrac{\partial}{\partial x} f_{1b}(T_{\max},t) = \dfrac{\partial}{\partial x} f_{0b}(T_{\max},t) - \dfrac{\partial}{\partial x} f_{1c}(T_{\max},t) \end{cases} \tag{5}$$

Fig. 1(a) shows the temperature domains, and the stationary shape of the pdfs obtained after a long operation time of the TCL unit. In this case, the pdf is almost uniform across the thermostat temperature range. An important quantity of interest, the probability that the TCL is on, is given by the area under the pdf associated with the on-mode,

$$\Pr[m(t)=1] = \int_{T_{\min}}^{T_{\text{maxmax}}} f_1(T,t)dT. \tag{6}$$

Furthermore, the area under both pdfs equals to 1, as it represents total probability.

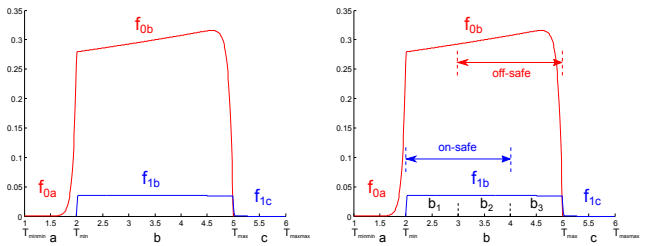

(a) Prob. density functions (b) Actuation ranges

Fig. 1. Temperature domains and sketch of the temperature distributions at equilibrium

A PDE system could be recovered using the Fokker-Planck framework also for extensions of the continuous dynamics, such as multidimensionality (e.g. second order temperature dynamics) and jump-noises (modeling user interaction /door opening events).

Finite Volume Method To solve the PDE, or equivalently propagate the pdf in time starting from an initial condition, numerical methods are employed and special care must be taken not to introduce unnecessary errors. Because of the conservation form of the Fokker-Planck equations, we choose FVM. This assures that the invariant of the system, probability, is conserved.

The idea is to partition the temperature domain, consisting of the sub-domains a, b, c, into a finite number of cells N_a, N_b, N_c. For convenience, we choose cells of equal sizes ΔT_a, ΔT_b and respectively ΔT_c for each domain. The continuous pdf function is replaced by a finite number of approximations points $F = [F_0^a, F_0^b, F_1^b, F_1^c]^T \in \mathbb{R}^{(N_a+2N_b+N_c)\times 1}$, each representing the average density value over a cell. This means that $F_{0a}(i)$ is designed to approximate the value $\int_{\text{Cell}_i} f_{0a}(T)dT$. Because of the linear form of the equations (the PDEs (4) are linear in the unknown function f) and of the boundary conditions, a linear finite-dimensional approximation dynamic can be obtained,

$$\dot{F}(t) = AF(t). \tag{7}$$

The resulting matrix A has the form of a transition rate matrix, i.e., has the property that the columns sum to 0. This is consistent with the fact that the space discretized description is a Markov chain representation of the stochastic hybrid TCL model. Furthermore, as a consequence of the transition rate form, matrix A has stable eigen values except one, which is exactly zero.

Lastly, the expected power consumption output of the TCL unit can be written as

$$y(t) = CF(t), \tag{8}$$

with $C = d \cdot [0_{1\times N_a}\ 0_{1\times N_b}\ \Delta T_b 1_{1\times N_b}\ \Delta T_c 1_{1\times N_c}]$.

2.2 Population model

Homogeneous Population The state distribution model has so far been developed for a single TCL. If we consider a population of N *identical* units, the dynamic model (7) holds, where the vector F simply changes meaning from probabilities to fractions [1] of the population. The expected power consumption of the population is given by

$$y(t) = NCF(t). \tag{9}$$

Heterogeneous Population "Small" heterogeneities of the TCL population should not cause severe modeling errors, but "large" heterogeneities will cause a significant departure from the homogeneous case. If we consider heterogeneity given in the form of parameter distribution, an exact modeling approach is to augment the TCL model and add parameters as states with dynamic zero. This is sketched below for a single branch of the hybrid dynamic,

$$\begin{cases} dT = \Big(aT(t) + b + m(t)c\Big)dt + \sigma dw \\ \dot{a} = 0, \dot{b} = 0, \dot{c} = 0, \dot{\sigma} = 0. \end{cases} \tag{10}$$

The corresponding multidimensional Fokker-Planck equation in the pdf $f(\mathbf{x},t) = f([T,a,c,b,\sigma],t)$ resolves to

$$\frac{\partial f(\mathbf{x},t)}{\partial t} + \frac{\partial}{\partial x_1}\Big((x_2x_1 + x_3)f(\mathbf{x},t)\Big) = \frac{x_5^2}{2}\frac{\partial^2 f(\mathbf{x},t)}{\partial^2 x_1}. \tag{11}$$

Unfortunately the approach suffers from the curse of dimensionality, since the space-variable $\mathbf{x} \in \mathbb{R}^5$ and a fine meshing of the five dimensional space is required to accurately recover the dynamics. Another observation is that, since the dynamics of the parameter-state are not connected to the dynamics of the temperature state, this formulation essentially leads to a clustering strategy, as used in Zhang et al. (2013). The clustering strategy corresponds to a rough partitioning of the parameter space.

[1] The state vector $F \in \mathbb{R}^{N_a+2N_b+N_c}$ is defined to represent average density values. It can be scaled by the cell sizes to yield probability quantities, which are equivalent to fractions in the population case.

Another approach to modeling heterogeneous dynamics is by adding an extra term in the PDE (4) that can be fitted to account for the empirically observed damping/dissipation effect. This approach is taken in Moura et al. (2013) using an increase of the already present diffusion term $\frac{\partial^2}{\partial x^2}$. Other operators could be consider for obtaining a better fit, perhaps taking inspiration from mechanical modeling.

2.3 Numerical verification

For a population composed of identical TCLs with known parameters, the only sources of inaccuracies are the fact that number of units in the population is finite (this error is small for large populations) and the spatial discretization of the PDE (this error is also small in the FVM case). A more concerning source of errors is that a real population would not be composed of identical units.

The graphs in Fig. 2 compare the power output of the model (7), (9) with Monte Carlo simulations of populations composed of $N = 10000$ units and different levels of heterogeneity. The scenario is that of a free response to synchronized initial condition where all units start from the same state $(T_{\max}, 0)$. This type of synchronization generates a well known oscillatory behavior.

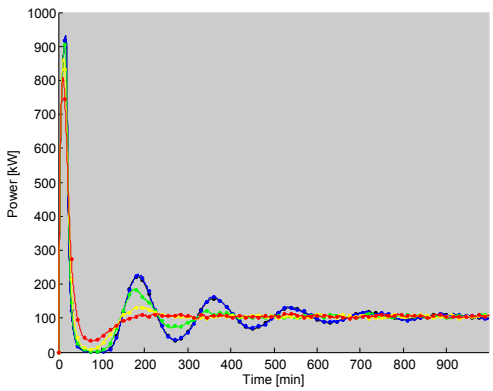

Fig. 2. Free response of the power consumption. The model is shown in black, and Monte Carlo simulations in blue (identical units), green (10% heterogeneity), yellow (20% heterogeneity) and red (30% heterogeneity).

The heterogeneous populations are composed of units where the key Equivalent Thermal Parameters of the TCL model (C, UA, W) and σ are each distributed according to a Gaussian profile with standard deviations 10%, 20% and 30% respectively, and truncated at $\pm 3\sigma$. It can be seen that the model matches well the power consumption of the population with identical units, but starts to degrade in performance with the increase of heterogeneity.

The parameters used for the TCL unit model are given in Table 1, resulting in a duty cycle with an on-time of about 18 minutes and off-time of 160 minutes.

Table 1. TCL unit parameters

C (J/K)	UA (W/K)	T_a ($^\circ C$)	W (W)
93920	1.432	24	100
η	σ ($^\circ C/s$)	$T_{\min}(^\circ C)$	$T_{\max}(^\circ C)$
2.8	0.0065	2	5

3. SWITCHING-FRACTION MODELING

In order to externally influence the power consumption of the TCL population, a switching-fraction signal is introduced. This is composed of two rational numbers that have the meaning of percentages, $\epsilon = \{\epsilon_0, \epsilon_1\} \in \mathbb{Q}^2$, and is broadcast every T_c seconds. The signal triggers a percentage ϵ_0 of the units "on", which are also in the "on-safe" temperature range (to be defined next), to "switch-off", and a percentage ϵ_1 of the units "off", which are also in the "off-safe" temperature range, to "switch-on". The "on-safe" range is $[T_{\min}, T_{\max} - \Delta T_0]$, and the "off-safe" range is $[T_{\min} + \Delta T_1, T_{\max}]$.

3.1 Actuation at the unit-level

While the switching-fractions are given at the population level, the actual switch is decided at the unit level. The TCLs have to meet the requested fractions without communicating with each other. To do so, an individual TCL that is in the target group switches based on the result of a binomial trial with success rate equal to the broadcast fraction value. If the target group is large enough, by the law of large numbers, the response of the population will be close to the requested one.

Referring to the stochastic hybrid model from Sec. 2.1, this represents a change in the discrete dynamic (2). The thermostat mechanism that generates deterministic state-dependent switches remains in place, and in addition, a stochastic component is introduced. This can generate spontaneous switches in response to the external signal. The temperature safe-zones are specially defined to ensure that switches do not occur close to a relevant boundary, since they would quickly be reversed by the thermostat. Furthermore, minimum on/off guards can be introduced to block the generation of an external switch, again to avoid the undesirable fast switching behavior.

3.2 Actuation in the population model

The switching-fractions broadcast actuation has been introduced as a discrete-time strategy. We will therefore introduce it into a discrete-time state space form of (7), specifically

$$F(k+1) = A_d F(k), \quad A_d = e^{A \cdot T_c}. \tag{12}$$

Fig.1(b) shows the temperature ranges of the switching-fraction signal. A switch-off fraction ϵ_0 will instantaneously transport probability from the temperature zone $[T_{\min}, T_{\max} - \Delta T_0]$ of f_{1b} to f_{0b}. Similarly, a switch-on fraction ϵ_1 will transport probability from the temperature zone $[T_{\min} - \Delta T_1, T_{\max}]$ of f_{0b} to f_{1b}. Notations b_1, b_2 and b_3 have been introduced in Fig.1(b) to define the safe-range partitions of the thermostat domain.

At the time step k^+, right after the broadcast, the pdfs over the $[T_{\min}, T_{\max}]$ domain change in the following way,

$$\begin{cases} F_{0b_1}(k^+) &= F_{0b_1}(k) + \Delta F_{1b_1}(k) \\ F_{1b_1}(k^+) &= F_{1b_1}(k) - \Delta F_{1b_1}(k) \\ F_{0b_2}(k^+) &= F_{0b_2}(k) - \Delta F_{0b_2}(k) + \Delta F_{1b_2}(k) \\ F_{1b_2}(k^+) &= F_{1b_2}(k) + \Delta F_{0b_2}(k) - \Delta F_{1b_2}(k) \\ F_{0b_3}(k^+) &= F_{0b_3}(k) - \Delta F_{0b_3}(k) \\ F_{1b_3}(k^+) &= F_{1b_3}(k) + \Delta F_{0b_3}(k). \end{cases} \tag{13}$$

If the minimum on-off time effects are not considered, then

$$\begin{cases}\Delta F_{1b_1} = \epsilon_0 F_{1b_1}; \;\; \Delta F_{1b_2} = \epsilon_0 F_{1b_2} \\ \Delta F_{0b_2} = \epsilon_1 F_{0b_2}; \;\; \Delta F_{0b_3} = \epsilon_1 F_{0b_3}, \end{cases} \tag{14}$$

leading to the following bilinear form:

$$F(k+1) = E(k+1)A_d F(k) \tag{15}$$

where $E(k)$ is a matrix depending on $\epsilon_0(k)$ and $\epsilon_1(k)$.

To include the minimum on-off effects, it is required to account for the units that are "locked", meaning that they have switched recently and cannot do so again. This reduces the amount of TCL units that are responsive to the broadcast signal. The population modeling needs to be extended to include tracking of the locked units. The following is similar to Zhang et al. (2013).

Let us assume that the on-lock duration (minimum on time) is 300[s], the same for all units, and equivalent to $l_1 = 5$ control steps for $T_c = 60$[s]. Similarly, off-lock duration (minimum off time) is associated with l_0. Initially, after a period with no control, all units are available for external switching. The first broadcast actuation causes a change in the distribution state F as described in (15). After the broadcast, the units that have switched become locked. We have direct information about the distribution of the locked states, this is $[\Delta F_{0b_2}(k), \Delta F_{0b_3}(k)]$ (just switched on), and $[\Delta F_{1b_1}(k), \Delta F_{1b_2}(k)]$ (just switched off). These proportions of the distribution are locked for a number l_1 and l_0 respectively of time steps, and we have to also take into account that these distribution will evolve in the temperature space during this time.

The following additional states are added to the model, representing the distribution of on-locked units and off-locked units respectively, at different locking stages.

$$F^l_{1L} = [F^l_{1Lb_1}, F^l_{1Lb_2}, F^l_{1Lb_3}, F^l_{1Lc}]^T, l \in \{0, \cdots, l_1 - 1\} \tag{16}$$

$$F^l_{0L} = [F^l_{0La}, F^l_{0Lb_1}, F^l_{0Lb_2}, F^l_{0Lb_3}]^T, l \in \{0, \cdots, l_0 - 1\} \tag{17}$$

These states propagate in the following way,

$$F^{l+1}_{1L}(k+1) = A_{d1} F^l_{1L}(k) \tag{18}$$

$$F^{l+1}_{0L}(k+1) = A_{d0} F^l_{0L}(k) \tag{19}$$

where matrices A_{d1} and A_{d0} represent temperature dynamics without switches.

In the end, the population model with switching-fraction actuation and minimum on/off times can be written as an augmented form of (15),

$$\bar{F}(k+1) = \bar{E}(k+1)\bar{A}_d F(k). \tag{20}$$

The bilinear form of the actuation in (15) and (20) can be seen as intrinsic to the TCL problem. It also appears in the case of a thermostat set-point actuation (Bashash and Fathy, 2013). This is because when using a physically based modeling approach, the actuation does not represent an external input to the system, but rather an internal transformation/change. Although the work in (Koch et al., 2011) proposes a linear formulation of the control, by letting the decision variables be the $\Delta F(k)$ quantities from (13), this has some disadvantages. In this case, there are as many independent control channels as bins in the relevant b_1, b_2 and b_3 temperature zones. The control is therefore dependent on a particular spatial discretization, and multiple switching fractions need to be broadcast at every time step.

4. OPEN-LOOP CONTROL

The objective of this section is to use models (15) and (20) for controling the aggregate power consumption of the TCL population. We define an optimization to generate an actuation sequence consisting of switching fractions $\epsilon_0(k)$ and $\epsilon_1(k)$. This actuation sequence is applied in open loop to drive the population to consume power in a manner that closely matches a given external reference.

4.1 Optimization problem

Given a power reference over a time horizon with T steps, and the initial state of the TCL population as the distribution F_0, an optimization for minimizing the power consumption tracking error can be written as

$$\begin{aligned} \underset{\epsilon_1(k),\epsilon_0(k)}{\text{minimize}} \quad & f(\epsilon) = \sum_{k=1}^{T} \Big(CF(k) - r(k) \Big)^2 \\ \text{subject to} \quad & \\ & 0 \le \epsilon_0(k) \le 1 \\ & 0 \le \epsilon_1(k) \le 1 \end{aligned} \tag{21}$$

where, using model (15),

$$F(k) = E(k)A_d E(k-1)A_d \cdots E(1)A_d F_0. \tag{22}$$

Gradient of the objective function The first order optimality conditions on (21) lead to,

$$\begin{aligned} \frac{\partial f(\epsilon)}{\partial \epsilon_x(l)} &= 2\Big(\sum_{k=l}^{T} [CF(k) - r(k)] C^*_{l,k} \Big) E_x(l) A d F(l-1) \\ &= 0 \end{aligned} \tag{23}$$

$$\text{with } C^*_{l,k} = C \cdot \prod_{l+1}^{i=k\downarrow} E(i)A_d, \forall k \ge l, \tag{24}$$

matrix $E_x(l) = \partial E(l)/\partial \epsilon_x(l)$, and the subscript $()_x$ standing in for either 0 or 1. Since the optimization problem is not convex (the objective expression is multilinear in the decision variables), it is important to specify this gradient information to the numerical solver to improve computational time and performance. Furthermore, the line-vector $C^*_{l,k}$ can be computed recursively

$$C^*_{l,k} = C^*_{l+1,k} E(l) A_d, \tag{25}$$

and the sparse structure of the matrices involved, especially E and A_d, can be used to reduce computation time.

The optimization has the same form when considering the minimum on/off effects, requiring only to replace F, C, A_d, E with the augmented versions, $\bar{F}, \bar{C}, \bar{A}_d, \bar{E}$.

4.2 Numerical examples

This section presents numerical results for tracking two references composed of step segments over a time horizon of two hours. The control time-step is $T_c = 60$[s] and the initial state of the population is close to the equilibrium distribution shown in Fig. 1. The optimization is implemented numerically in MATLAB with a generic interior point algorithm. Figure 3 shows the results, and it can be

seen that model without the locking mechanism can be driven to consume closely to the desired reference. The power flexibility of the model with the locking mechanism is somewhat reduced. Although this is a good solution, because of the non-convex formulation we cannot conclude that these results represent the optimum.

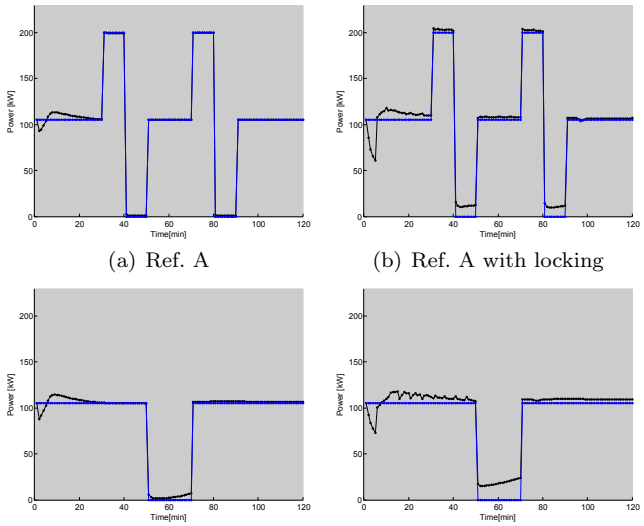

(a) Ref. A (b) Ref. A with locking

(c) Ref. B (d) Ref. B with locking

Fig. 3. Power reference (blue) and the optimized consumption (black) of the model with (20) and without locking (15).

Figure 4 shows the results of actuating Monte Carlo populations ($N = 10000$) using the optimized switching fraction signals. The performance is good for both the homogeneous and the heterogeneous populations suggesting that the switching fractions broadcast is robust to heterogeneity [2]. On the other hand, as can be seen from Fig. 4(a) and 4(c), the locking effect is significant and cannot be avoided in the modeling.

5. CONCLUSION

Demand response of TCL populations can provide a great support to the electrical grid by reducing the capacity need for fast reserves. The switching fraction broadcast is a reliable way of engaging and controlling the power output of such populations within their natural flexibility. This intrinsic flexibility is most accurately represented by a physically based modeling technique. We have introduced a switching-fraction actuation entering the system in a bilinear form. While this is not an ideal form for control, we have shown that an open-loop model predictive strategy can still be computationally tractable and robust to heterogeneity. Future work will make use of these results to complete a control architecture with feedback and state-estimation.

REFERENCES

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

[2] Results are shown here for Gaussian distributed parameters. Similar results are obtained also for uniformly distributed parameters.

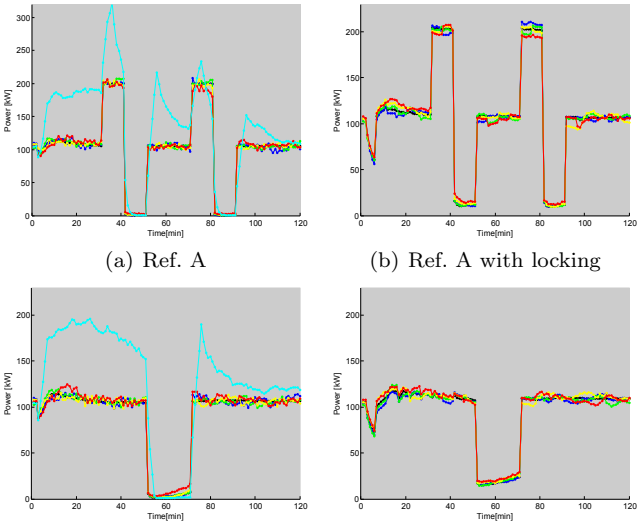

(a) Ref. A (b) Ref. A with locking

(c) Ref. B (d) Ref. B with locking

Fig. 4. TCL populations under optimized actuation. The color legend is the same as in Fig. 2. The Monte Carlo populations on the left do not have the locking mechanism implemented (except cyan), while the populations on the right all have it.

Callaway, D.S. (2009). 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.

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

Koch, S., Mathieu, J.L., and Callaway, D.S. (2011). Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services. In *Proc. PSCC*, 1–7.

Malhame, R. and Chong, C.Y. (1985). Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system. *Automatic Control, IEEE Transactions on*, 30(9), 854–860.

Moura, S., Bendtsen, J., and Ruiz, V. (2013). Modeling heterogeneous populations of thermostatically controlled loads using diffusion-advection pdes. In *Proceedings of the 2013 ASME Dynamic Systems and Control Conference, Stanford, California.*

Soudjani, S. and Abate, A. (2013). Aggregation of thermostatically controlled loads by formal abstractions. In *ECC 2013: Eropean Control Conference*, 4232–4237. Zurich, Switzerland.

Totu, L.C., Leth, J., and Wisniewski, R. (2013). Control for large scale demand response of thermostatic loads. In *American Control Conference (ACC), 2013*, 5023–5028. IEEE.

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