

RESEARCH ARTICLE

Optimal Power Consumption for Demand Response of Thermostatically Controlled Loads

Abhishek Halder*¹ | Xinbo Geng² | Fernando A.C.C. Fontes³ | P.R. Kumar² | Le Xie²

¹Department of Applied Mathematics and Statistics, University of California, Santa Cruz, CA, USA

²Department of Electrical and Computer Engineering, Texas A&M University, TX, USA

³ISR-Porto and Faculdade de Engenharia, Universidade do Porto, Porto, Portugal

Correspondence

*Abhishek Halder, Department of Applied Mathematics and Statistics, University of California, Santa Cruz, CA 95064. Email: ahalder@ucsc.edu

Summary

We consider the problem of determining the optimal aggregate power consumption of a population of thermostatically controlled loads such as air conditioners. This is motivated by the need to synthesize the demand response for a load serving entity (LSE) catering a population of such customers. We show how the LSE can opportunistically design the aggregate reference consumption to minimize its energy procurement cost, given day-ahead price, load and ambient temperature forecasts, while respecting each individual load's comfort range constraints. The resulting synthesis problem is intractable when posed as a direct optimization problem after Euler discretization of the dynamics, since it results in a mixed integer linear programming problem with number of variables typically of the order of millions. In contrast, in this paper we show that the problem is amenable to continuous-time optimal control techniques. Numerical simulations elucidate how the LSE can use the optimal aggregate power consumption trajectory thus computed, for the purpose of demand response.

KEYWORDS:

thermostatically controlled loads, Pontryagin's maximum principle, day-ahead price, demand response

1 | INTRODUCTION

Motivated by the goal of sustainable electricity generation, renewables such as solar and wind are of increasing interest as electric energy resources. Concomitantly, the inherent time variability of such renewable generation is shifting modern power system operation from the traditional "supply follows demand" paradigm to the one where "demand adapts to supply". This new operational paradigm, called "demand response"^{1,2}, can leverage demand side flexibility to offset variability in generation. Of particular interest in this context are thermostatically controlled loads (TCLs) such as air conditioners. In this paper we examine how an "aggregator", also known as a "load serving entity" (LSE), can employ a population of its customers' TCLs to shape the aggregate power consumption, while adhering to each load's comfort constraints. We consider an LSE buying energy from the day-ahead market for a population of TCLs. We address the question of designing the optimal aggregate power consumption trajectory for this population, given a forecast of day-ahead price trajectory.

Related Work

Modeling the dynamics of a population of TCLs has been investigated in several papers^{3,4,5,6,7,8,9,10}, with the aim of deriving control-oriented models that can accurately predict the aggregate power trajectory. Once such a model is obtained, the predominant focus in these and other papers^{11,12,13} is to design a model-based setpoint controller to enable the TCL population track a given reference aggregate power trajectory in real-time, thereby compensating for the possible mismatch between the real-time and forecasted ambient temperatures. In contrast to these papers, where the availability of a reference power trajectory is assumed for real-time control design, we focus on the case where the LSE determines this reference to minimize its energy procurement cost while guaranteeing that the reference trajectory can indeed be tracked by the aggregate dynamics without violating individual comfort range constraints.

In Paccagnan et al.¹⁴, the range of feasible reference power trajectories was studied. Reulens et al.¹⁵ adopted a model-free approach to schedule a cluster of electric water heaters using batch reinforcement learning. Also relevant to our work is the paper by Mathieu et al.¹⁶ where minimizing TCL energy consumption cost subject to end users' comfort zones was considered (see equation (5) in Section III of that reference). The perspective and results of our paper differ significantly from the aforesaid formulation in that we allow a target total energy budget constraint for the LSE, which in turn prohibits transcribing the overall (discrete version of the) optimization problem into a set of decoupled mixed integer linear programs (MILPs), as was the case in Mathieu et al.¹⁶ In fact, it is precisely the dynamic coupling that makes the non-convex optimal control problem difficult to solve by a direct "discretize-then-optimize" approach, as we explain further in Section 3.3.

Contributions of This Paper

For an LSE managing a finite population of TCLs, we formulate and solve the optimal aggregate power consumption design as a finite horizon deterministic optimal control problem. The contribution of the present paper beyond our previous work^{17,18} is that herein, we analytically solve the continuous-time optimal control problem (Section 4), thereby revealing qualitative insights on how the LSE can use the knowledge of day-ahead price forecast, load forecast, and ambient temperature forecast, for the purpose of energy procurement at least cost. Furthermore, when there is additional constraint on minimum thermostatic switching period, we provide an algorithm (Section 5 and 6) to recover the optimal binary controls from the corresponding convexified optimal control solutions.

In the presence of state inequality constraints arising from comfort range contracts between the LSE and individual TCLs, the optimal controls are shown to depend on both the shape of the day-ahead price trajectory[†], and on the minimum switching period (also known as "lockout constraint"⁹) at the upper and lower comfort boundaries allowable by the thermostats. Specifically, the application of Pontryagin's maximum principle (PMP) reveals that the optimal policy is a function of certain "threshold price" to be computed from the day-ahead price forecast. The resulting optimal indoor temperature trajectories are described in terms of the so-called "two-sided Skorokhod maps"^{19,20} parameterized by individual TCL's upper and lower comfort boundaries.

This paper is organized as follows. In Section 2, we describe the mathematical models. In Section 3, we formulate the design of power consumption as an optimal control problem. Sections 4 and 5 present the solution of the power consumption design problem. Specifically, the results of Section 4 use PMP to analytically derive non-convex optimal controls but assume zero-amplitude chattering; the results of Section 5 analytically maps the (numerically computed) convexified optimal controls with prescribed chattering constraints to the corresponding non-convex optimal controls respecting those constraints. The purpose of the first is to gain qualitative insights on the structure of the optimal controls, while the latter is used for realistic numerical simulation accounting device constraints. In Section 6, numerical results based on the day-head price forecast data from Electric Reliability Council of Texas (ERCOT) and the ambient temperature forecast data from a weather station in Houston, Texas are reported, to illustrate how the LSE can use the optimal power consumption trajectory computed via the proposed framework, for the purpose of demand response. Section VII concludes the paper.

[†]We clarify here that the analytical results in Section 4 are derived assuming continuous and differentiable price forecast trajectory which helps to reveal the structure of the solution. For the purpose of numerical simulation, we use the day-ahead price forecast data from ERCOT, which is available piecewise-constant for each hour, and show how our results from Section 5 can be applied to this case. In particular, the framework and results of this paper are not limited to the context of the day-ahead price, and can be applied to other price forecasts such as those from intra-day real-time market¹⁷.




FIGURE 1 The dynamic behavior of a TCL with fixed setpoint s_0 , is illustrated for a time-varying ambient temperature $\theta_a(t)$. The horizontal axis plots time t and the vertical axis plots temperature. The comfort temperature interval for the TCL is $[L_0, U_0]$ with range $U_0 - L_0 = 2\Delta$. The indoor temperature trajectory $\theta(t) \in [L_0, U_0]$ consists of alternating OFF (blue, up-going) and ON (red, down-going) segments, with the boundaries L_0 and U_0 acting as reflecting barriers for the ON and OFF segments, respectively.

Notation

We use the symbols $\mathbb{1}_X$ and $|X|$ to respectively denote the indicator function, and the Lebesgue measure of set X . The set of integers, reals and positive reals are denoted by \mathbb{Z}, \mathbb{R} and \mathbb{R}^+ , respectively. We recall that càdlàg functions are defined to be everywhere right-continuous functions having left limits everywhere, and use $\mathbb{D}(Y)$ to denote the space of càdlàg functions whose range is set Y . For $a, b \in \mathbb{R}$, we use the notations $a \vee b := \max(a, b)$, $a \wedge b := \min(a, b)$, $[a]^+ := 0 \vee a$, and $[a] := \min\{n \in \mathbb{Z} : n \geq a\}$. The symbol \circ denotes the composition operator, and $\text{spt}(\cdot)$ denotes the support of a function.

2 | MODEL

2.1 | Dynamics of Individual Thermostatically Controlled Load

The dynamic behavior of an individual TCL is shown in Fig. 1. At time t , let us denote the indoor temperature by $\theta(t)$, and the ambient temperature by $\theta_a(t)$. At $t = 0$, an occupant privately sets a temperature $s_0 := s(0)$, called setpoint, close to which the indoor temperature $\theta(t)$ must lie at all times. If the occupant is willing to tolerate at most $\pm\Delta$ temperature deviation from s_0 , then we define its “temperature comfort range” as $[L_0, U_0] := [s_0 - \Delta, s_0 + \Delta]$. If the setpoint s does not change with time, then $s(t) \equiv s_0$, and consequently, the comfort boundaries L_0 and U_0 remain fixed over time. For specificity we consider the problem of controlling air-conditioning rather than heating, though the theory developed in the sequel applies to both.

The rate of change of $\theta(t)$ is governed by Newton’s law of heating/cooling given by the ordinary differential equation (ODE)

$$\dot{\theta}(t) = -\alpha (\theta(t) - \theta_a(t)) - \beta P \sigma(t), \quad (1)$$

where $\sigma(t)$ is the ON/OFF mode indicator variable of the air-conditioner, given by

$$\sigma(t) := \begin{cases} 1 & \text{if } \theta(t) = U_0, \\ 0 & \text{if } \theta(t) = L_0, \\ \sigma(t^-) & \text{otherwise.} \end{cases} \quad (2)$$

In other words, $\sigma(t) = 1(0)$ indicates that the TCL is in ON (OFF) mode. In (1), the parameters $\alpha, \beta, P > 0$ respectively denote the heating time constant, thermal conductivity, and amount of thermal power drawn by the TCL in ON mode. A parameter $\eta > 0$ called load efficiency, relates the thermal power drawn P , with the electrical power drawn P_e , via the formula $P_e = \frac{P}{\eta}$. The state of a TCL at time t is the tuple $\{s(t), \theta(t), \sigma(t)\} \in \mathbb{R}^2 \times \{0, 1\}$.

As shown in Fig. 1, starting from an initial condition $(s_0, \theta_0, \sigma_0 = 0)$, the indoor temperature $\theta(t)$ rises exponentially until it hits the upper boundary U_0 , at which time an OFF→ON mode transition occurs, and subsequently $\theta(t)$ decreases exponentially until it hits the lower boundary L_0 , at which time an ON→OFF transition takes place, and so on. Thus, the dynamics of a TCL




FIGURE 2 A typical day-ahead price forecast trajectory $\hat{\pi}(t)$ (\$/MWh versus time t in hours) shown as a piecewise constant curve. As shown, $\hat{\pi}(t)$ is typically increasing till the late afternoon, and decreasing thereafter. The shown trajectory is for Houston on August 10, 2015, available from the day-ahead energy market of Electric Reliability Council of Texas (ERCOT)²¹.

is hysteretic in the sense that if $\theta_0 \in [L_0, U_0]$, then $\theta(t) \in [L_0, U_0]$ for all $t > 0$. While the qualitative behavior shown in Fig. 1 is true for any $\theta_a(t) > U_0$, temporal variations of $\theta_a(t)$ engender time-varying heating/cooling rates for $\theta(t)$.

2.2 | Day Ahead Price Forecasts

We suppose that the LSE is exposed to a price forecast $\hat{\pi}(t)$ and ambient temperature forecast $\hat{\theta}_a(t)$, over a time horizon $[0, T]$. For example, if $T = 24$ hours, and the forecast is made on the previous day, then $\hat{\pi}(t)$ is the forecasted price from the day-ahead energy market. We also allow the LSE to have a target for the total energy E to be consumed over $[0, T]$ by the population of N TCL customers managed by that LSE. The choice of E may be restricted by the parameters of the TCL population, an issue we address in Section 3.2. To minimize its energy procurement cost, the LSE would like to schedule purchase of energy when $\hat{\pi}(t)$ is low, and defer purchase when $\hat{\pi}(t)$ is high, while satisfying the total energy budget (E), as well as maintaining the comfort constraints described below specified by each of the TCLs.

A typical day-ahead price forecast trajectory $\hat{\pi}(t)$ is shown in Fig. 2. Since the aggregate power consumption around late afternoon is expected to be higher than at other times of the day, the day-ahead price is typically forecasted to be increasing till late afternoon and decreasing thereafter.

2.3 | Comfort Range Contracts

Each of the N TCLs managed by an LSE, may have different comfort ranges $[L_{i0}, U_{i0}]$ with different tolerances Δ_i , $i = 1, \dots, N$. Let $\theta_i(t)$ be the indoor temperature of the i^{th} home at time t . The LSE is obligated to maintain the indoor temperatures of its customer TCLs within their specified comfort ranges. That is, for each $i = 1, \dots, N$, the LSE must ensure that $\theta_i(t) \in [L_{i0}, U_{i0}]$ for all $t \geq 0$. Such an agreement constitutes a *contract* between the LSE and an individual TCL.

An important part of this agreement is the *flexibility* of a load, captured by its range 2Δ . The LSE can utilize this flexibility to optimally time its purchase of power. The LSE's business model essentially consists of sharing part of the realized savings with the customers in terms of serving their needs for energy at low cost. Naturally, a customer with a greater flexibility 2Δ is more valuable to the LSE and such customers can obtain better contracts from the LSE.

2.4 | Assumptions

For the purpose of theoretical analysis, we make the following assumptions.

- The ambient temperature forecast $\hat{\theta}_a(t)$, and price forecast $\hat{\pi}(t) > 0$, are continuous and differentiable functions of time t . (See the footnote at p. 2 regarding the assumption on $\hat{\pi}(t)$.)
- All TCLs are cooling, i.e., for all $t \in [0, T]$, we have $\hat{\theta}_a(t) > \max_{i=1, \dots, N} U_{i0}$.
- Each TCL in the population, when ON, draws the same thermal power P . Further, each TCL is assumed to have same load efficiency η .
- Without loss of generality, the initial indoor temperatures $\theta_{i0} := \theta_i(0) \in [L_{i0}, U_{i0}]$, for all $i = 1, \dots, N$.

3 | PROBLEM FORMULATION

Consider an LSE managing N TCLs with thermal coefficients $\{\alpha_i, \beta_i\}_{i=1}^N$, initial conditions $\{(s_{i0}, \theta_{i0}, \sigma_{i0})\}_{i=1}^N$, and comfort tolerances $\{\Delta_i\}_{i=1}^N$. We now formulate the optimal power consumption design problem. We suppose that the LSE has available estimates^{22,23} of the parameters and initial conditions at the beginning of the time horizon.

3.1 | The Load Serving Entity's Objective

Denoting by $n_{\text{ON}}(t)$ the number of ON TCLs at time t , the aggregate electrical power drawn by the TCL population at time t is

$$P_e n_{\text{ON}}(t) = \frac{P}{\eta} \sum_{i=1}^N u_i(t).$$

We take the switching trajectories $\sigma_i(t)$ as decision variables $u_i(t) \in \{0, 1\}$, and introduce an extended state vector

$$\mathbf{x}(t) := \left(\theta_1(t), \dots, \theta_N(t), t, \int_0^t \sum_{i=1}^N u_i(\zeta) d\zeta \right)^\top$$

of size $(N + 2) \times 1$. At time t , the components of $\mathbf{x}(t)$ are $x_i(t) := \theta_i(t)$ for $i = 1, \dots, N$, $x_{N+1}(t) := t$, and $x_{N+2}(t) := \int_0^t \sum_{i=1}^N u_i(\zeta) d\zeta$. Then, to minimize the procurement cost for total energy consumption over $[0, T]$, the LSE needs to

$$\underset{u_1(t), \dots, u_N(t) \in \{0, 1\}^N}{\text{minimize}} \quad \frac{P}{\eta} \int_0^T \hat{\pi}(t) \sum_{i=1}^N u_i(t) dt, \quad (3)$$

subject to the constraints:

C1. (Indoor temperature dynamics)

$$\dot{x}_i(t) = -\alpha_i(x_i(t) - \hat{\theta}_a(t)) - \beta_i P u_i(t), \quad x_i(0) = \theta_{i0}, \quad i = 1, \dots, N, \quad (4a)$$

$$\dot{x}_{N+1}(t) = 1, \quad x_{N+1}(0) = 0, \quad x_{N+1}(T) = T, \quad (4b)$$

C2. (Energy/isoperimetric constraint)

$$\dot{x}_{N+2}(t) = \sum_{i=1}^N u_i(t), \quad x_{N+2}(0) = 0, \quad x_{N+2}(T) = \frac{\eta E}{NP}, \quad (5)$$

C3. (Contractual comfort/state inequality constraint)

$$L_{i0} \leq x_i(t) \leq U_{i0}, \quad (6)$$

where $[L_{i0}, U_{i0}] := [s_{i0} - \Delta_i, s_{i0} + \Delta_i]$ for $i = 1, \dots, N$.

Notice that the constraints (4) and (6) are decoupled, while the cost function (3) and constraint (5) are coupled. Denoting the solution of the open loop deterministic optimal control problem (3)–(6) by $\{u_i^*(t)\}_{i=1}^N$, the optimal power consumption trajectory is given by $P_{\text{total}}^{\text{ref}}(t) = \frac{P}{\eta} \sum_{i=1}^N u_i^*(t)$.

Remark 1. The optimal control problem (3)–(6) is non-autonomous since there are explicit dependences on time in the cost function (via $\hat{\pi}(t)$) and in the dynamics (via $\hat{\theta}_a(t)$). This motivates the inclusion of time t as a component of the extended state vector $\mathbf{x}(t)$.

3.2 | Feasibility

Let $\tau := \frac{\eta E}{P}$, and notice that the constraint (5) imposes a necessary condition for feasibility,

$$0 \leq \bar{\tau} := \frac{\tau}{NT} = \frac{\eta E}{NPT} = \frac{x_{N+2}(T)}{T} \leq 1. \quad (7)$$

Given an ambient temperature forecast $\hat{\theta}_a(t)$ and parameters of the TCL population, further restriction of $\bar{\tau}$ is needed to include the possibility of zero dynamics on (meaning the temperature trajectory chatters along) the boundaries L_{i0} and U_{i0} . Such a restriction is of the form

$$0 \leq \bar{\tau}_\ell \leq \bar{\tau} \leq \bar{\tau}_u \leq 1, \quad (8)$$

where $\bar{\tau}_\ell := \frac{\tau_\ell}{NT} = \frac{\eta E_\ell}{NPT}$, and likewise for $\bar{\tau}_u = \frac{\eta E_u}{NPT}$. Here E_ℓ (resp. E_u) is the aggregate energy consumed if the entire population were to be maintained at their private upper (lower) setpoint boundaries, thus resulting in the lowest (highest) total energy consumption while respecting (6). In other words, $E_\ell = \frac{P}{\eta} \sum_{i=1}^N \int_0^T u_i(t) dt$, where the zero dynamics controls are $u_i(t) = \frac{\alpha_i}{\beta_i P} (\hat{\theta}_a(t) - U_{i0})$, and hence

$$\bar{\tau}_\ell = \frac{1}{NP} \left(\sum_{i=1}^N \frac{\alpha_i}{\beta_i} (\langle \hat{\theta}_a \rangle - U_{i0}) \right), \quad (9)$$

where $\langle \hat{\theta}_a \rangle := \frac{1}{T} \int_0^T \hat{\theta}_a(t) dt$. Similar calculation yields

$$\bar{\tau}_u = \frac{1}{NP} \left(\sum_{i=1}^N \frac{\alpha_i}{\beta_i} (\langle \hat{\theta}_a \rangle - L_{i0}) \right). \quad (10)$$

Thus, (8) characterizes the necessary and sufficient conditions for feasibility of the optimal control problem (3)–(6).

Remark 2. Notice that constraint **C2** expresses a total energy budget $\frac{P}{\eta} \int_0^T \sum_{i=1}^N u_i(t) dt = E$. In the absence of **C2**, the optimal controls $\{u_i^*(t)\}_{i=1}^N$ that minimize (3) subject to **C1** and **C3**, satisfy $\frac{P}{\eta} \int_0^T \sum_{i=1}^N u_i^*(t) dt = E_\ell$, where from (9), we have

$$E_\ell = \frac{T}{\eta} \left(\sum_{i=1}^N \frac{\alpha_i}{\beta_i} (\langle \hat{\theta}_a \rangle - U_{i0}) \right). \quad (11)$$

Remark 3. Notice also that condition (8) is equivalent to the energy inequality $E_{\min} \leq E_\ell \leq E \leq E_u \leq E_{\max}$, where $E_{\min} := 0$, $E_{\max} := \frac{NPT}{\eta}$, E_ℓ is given by (11), and from (10) we have $E_u = \frac{1}{NP} (\sum_{i=1}^N \frac{\alpha_i}{\beta_i} (\langle \hat{\theta}_a \rangle - L_{i0}))$. As a consequence, the feasible energy budget E must belong to an interval $[E_\ell, E_u]$ with length $E_u - E_\ell = \frac{2T}{\eta} \sum_{i=1}^N \frac{\alpha_i}{\beta_i} \Delta_i$.

3.3 | Difficulty in Direct Numerical Simulation

A direct numerical approach converts the optimal control problem (3)–(6) to an optimization problem via time discretization. Such a “discretize-then-optimize” strategy leads to a mixed integer linear program (MILP), since $\theta_i(t) \in \mathbb{R}$ and $u_i(t) \in \{0, 1\}$, for $i = 1, \dots, N$. Typically, the day-ahead price forecast $\hat{\pi}(t)$ is available as a function that is piecewise constant for each hour, and so taking the Euler discretization for dynamics (4a) with 1 minute time resolution results in an MILP with $24 \times 60 \times 2 \times N$ variables. In our experience, solving the MILP even for $N = 2$ homes for the day-ahead price, is computationally expensive (with CPU runtime over 24 hours) in Gurobi²⁴. On the other hand, a linear program (LP) relaxation of the MILP, resulting from the control convexification $u_i \in \{0, 1\} \mapsto \tilde{u}_i \in [0, 1]$, has much faster runtime and was reported in our earlier work¹⁷. Furthermore, the optimal solution of the LP relaxation has the physical meaning of average ON duration over a discretization interval (see Fig. 3).

However, the MILP equality constraint does not satisfy total unimodularity (see for example, Ch. 5 in Schrijver²⁵); consequently the optimal solution of the LP relaxation is not the optimal solution of the MILP. In the following Sections IV and V, we




FIGURE 3 A comparison of the optimal indoor temperatures $\theta^*(t)$ and optimal controls $u^*(t)$, computed from the discretized MILP and its LP relaxation for $N = 1$ home. Top: A monotone increasing day-ahead price forecast $\hat{\pi}_{DA}(t)$ that is piecewise constant for each hour, is shown with $\tau = \frac{\eta E}{P}$. Bottom: The corresponding solution of the optimal power consumption problem (3)–(6) via discretized MILP (from Gurobi) and LP (from MATLAB linprog) for 1 home with $[L_0, U_0] = [20^\circ\text{C}, 30^\circ\text{C}]$, and constant ambient forecast $\hat{\theta}_a(t) = 32^\circ\text{C}$. The parameter values used in this simulation are from Table 1, p. 1392 in Callaway⁵.

solve the continuous time optimal control problem (3)–(6) using Pontryagin’s maximum principle (PMP)²⁶, thereby obtaining qualitative insights into the optimal solution, which are otherwise difficult to gauge from the direct numerical solution of the discretized LP relaxation.

4 | SOLUTION OF THE OPTIMAL CONTROL PROBLEM

For the optimal control problem (3)–(6), if we remove constraints **C2** and **C3**, then the optimal control is trivial: $u_i^*(t) = 0$ for all $i = 1, \dots, N$, for all $t \in [0, T]$. In Section 4.1, we first discuss the non-trivial case of solving (3) subject to **C1** and **C2**, i.e., in the absence of the inequality constraints **C3**. This is followed up with the solution for general case in Section 4.2, with all constraints **C1**–**C3** active.

4.1 | Solution with Constraint C3 Inactive

The following Theorem summarizes our results for this case. In particular, it reveals key structural properties of the optimal solution, viz. (i) the optimal controls are synchronizing across the TCL population, (ii) the optimal policy is of threshold type, i.e., there is a unique threshold price π^* to be determined from the given day-ahead price trajectory $\hat{\pi}(t)$ such that the optimal control is ON (resp. OFF) whenever $\hat{\pi}(t)$ falls below (resp. exceeds) that threshold, (iii) the computation of π^* amounts to performing monotone rearrangement (p. 276, Ch. 10 in Hardy et al.²⁷) of the trajectory $\hat{\pi}(t)$ and allocating the requisite ON time $\frac{\tau}{N}$ in the interval $[0, T]$ such a way that corresponds to the least price segment (this will be further elaborated in Remark 4 following the proof).

Theorem 1. Consider problem (3) with constraints **C1**–**C2**. Then

- (i) the optimal controls are synchronizing, i.e., $u_1^*(t) = u_2^*(t) = \dots = u_N^*(t)$ at each $t \in [0, T]$;
- (ii) there is a unique threshold price π^* , such that $u_i^*(t) = 1(0)$ for all $i = 1, \dots, N$, iff $\hat{\pi}(t) < (\geq)\pi^*$;

(iii) let $\Phi_{\hat{\pi}}(\tilde{\pi}) := \int_0^T \mathbb{1}_{\{\hat{\pi}(t) \leq \tilde{\pi}\}} dt$, $\pi^* := \inf\{\tilde{\pi} \in \mathbb{R}^+ : \Phi_{\hat{\pi}}(\tilde{\pi}) = \frac{\tau}{N}\}$, $S := \{s \in [0, T] : \hat{\pi}(s) < \pi^*\}$, and let λ be the costate vector corresponding to the extended state vector \mathbf{x} . The optimal solution is

$$u_i^*(t) = \begin{cases} 1 & \forall t \in S \\ 0 & \text{otherwise,} \end{cases}$$

$$\lambda_i^*(t) = 0 \quad \forall i = 1, \dots, N, \quad \lambda_{N+1}^*(t) = \begin{cases} \frac{P}{\eta} N (\pi^* - \hat{\pi}(t)) & \forall t \in S, \\ 0 & \text{otherwise,} \end{cases} \quad \lambda_{N+2}^*(t) = -\frac{P}{\eta} \pi^*,$$

$$x_i^*(t) = e^{-\alpha_i t} \left(\theta_{i0} + \int_0^t e^{\alpha_i \zeta} (\alpha_i \hat{\theta}_a(\zeta) - \beta_i P u_i^*(\zeta)) d\zeta \right) \quad \forall i = 1, \dots, N, \quad x_{N+1}^*(t) = t, \quad x_{N+2}^*(t) = \begin{cases} Nt & \forall t \in S, \\ 0 & \text{otherwise.} \end{cases}$$

Proof. (i) The Hamiltonian

$$H = \sum_{i=1}^N \left[u_i(t) \left(\frac{P}{\eta} \hat{\pi}(t) - P \beta_i \lambda_i(t) + \lambda_{N+2}(t) \right) - \alpha_i \lambda_i(t) (x_i(t) - \hat{\theta}_a(t)) \right] + \lambda_{N+1}(t) \quad (12)$$

gives the first order optimality conditions

$$\dot{\lambda}_i(t) = -\frac{\partial H}{\partial x_i} = \alpha_i \lambda_i(t), \quad i = 1, 2, \dots, N, \quad (13)$$

$$\dot{\lambda}_{N+1}(t) = -\frac{\partial H}{\partial t} = -\frac{P}{\eta} \frac{\partial \hat{\pi}}{\partial t} \sum_{i=1}^N u_i(t) - \frac{\partial \hat{\theta}_a}{\partial t} \sum_{i=1}^N \alpha_i \lambda_i(t), \quad (14)$$

$$\dot{\lambda}_{N+2}(t) = -\frac{\partial H}{\partial x_{N+2}} = 0 \Rightarrow \lambda_{N+2} = \text{constant.} \quad (15)$$

The transversality condition yields

$$-\sum_{i=1}^N \lambda_i(T) dx_i(T) - \lambda_{N+1}(T) dx_{N+1}(T) - \lambda_{N+2}(T) dx_{N+2}(T) + H(T) dT = 0. \quad (16)$$

Since the terminal states $\theta_i(T)$ are free, $dx_i(T) = d\theta_i(T) \neq 0$, and hence (16) implies that $\lambda_i(T) = 0$, for all $i = 1, \dots, N$. Because T is fixed, $dx_{N+1}(T) = dT = 0$. Similarly, $dx_{N+2}(T) = \frac{d\tau}{N} = 0$. Combining $\lambda_i(T) = 0$ with (13) gives $\lambda_i(t) = 0$ for all t . Setting $\lambda_i(t) \equiv 0$ in (12), and invoking PMP yields the optimal controls as

$$\underset{(u_1(t), \dots, u_N(t)) \in \{0,1\}^N}{\text{argmin}} \left(\frac{P}{\eta} \hat{\pi}(t) + \lambda_{N+2} \right) \sum_{i=1}^N u_i(t). \quad (17)$$

Hence, if $\frac{P}{\eta} \hat{\pi}(t) + \lambda_{N+2} > (<) 0$ at any time t , then we need to minimize (maximize) $\sum_{i=1}^N u_i(t)$ over $\{0, 1\}^N$ at that time, meaning that the *optimal controls are synchronized*.

(ii) We know that $P, \eta, \hat{\pi}(t) > 0 \forall t \in [0, T]$. Thus, if $\lambda_{N+2} \geq 0$, then $\frac{P}{\eta} \hat{\pi}(t) + \lambda_{N+2} > 0$ implying $u_i^*(t) = 0 \forall i = 1, \dots, N, \forall t \in [0, T]$. This in turn leads to $\dot{x}_{N+2} = 0$ implying $x_{N+2}(t) = x_{N+2}(0) = 0 = x_{N+2}(T) = \frac{\tau}{N}$, which is impossible since $\tau \neq 0$ (given). Therefore, the constant $\lambda_{N+2} < 0$.

Notice that whether $\frac{P}{\eta} \hat{\pi}(t) + \lambda_{N+2}$ is > 0 or < 0 depends on the magnitude of the constant $\lambda_{N+2} < 0$, as well as on the magnitude of $\hat{\pi}(t) > 0$. Depending on the sign of the time-varying sum $\frac{P}{\eta} \hat{\pi}(t) + \lambda_{N+2}$, the optimal control will switch between 0 and 1.

Let us denote the optimal value of λ_{N+2} as λ_{N+2}^* , and consider a set $S \subseteq [0, T]$ given by $S := \{s \in [0, T] : \hat{\pi}(s) < -\frac{\eta}{P} \lambda_{N+2}^*\}$. Then, from PMP, $\forall i = 1, \dots, N$, we can rewrite the optimal control as $u_i^*(t) = 1 \forall t \in S$, and $= 0$ otherwise. The statement follows by letting $\pi^* := -\frac{\eta}{P} \lambda_{N+2}^* > 0$, wherein the uniqueness of π^* follows from the continuity of $\hat{\pi}(t)$ (as per assumption in Section 2.4).




FIGURE 4 Illustration of the optimal control $u_i^*(t)$, $i = 1, \dots, N$, for N TCLs corresponding to strictly monotone price forecast $\hat{\pi}(t)$. Left: Strictly increasing (red) and decreasing (green) price forecasts $\hat{\pi}(t)$, Right: the corresponding optimal control with constraint **C3** inactive is given by $u_i^*(t) = \mathbb{1}_S$, where $S = [0, \frac{\tau}{N}]$ for strictly increasing, and $S = [T - \frac{\tau}{N}, T]$ for strictly decreasing $\hat{\pi}(t)$.

- (iii) To determine $u_i^*(t)$, all that remains is to determine λ_{N+2}^* , or equivalently $-\frac{\eta}{P}\lambda_{N+2}^*$. The choice of $\lambda_{N+2} < 0$, or equivalently $-\frac{\eta}{P}\lambda_{N+2} > 0$, is constrained by the terminal condition

$$x_{N+2}(T) = \frac{\tau}{N} \Leftrightarrow \int_0^T \sum_{i=1}^N u_i(t) dt = \frac{\tau}{N} \stackrel{\text{PMP}}{\Rightarrow} \int_0^T \sum_{i=1}^N \mathbb{1}_{\{\hat{\pi}(t) < -\frac{\eta}{P}\lambda_{N+2}\}} dt = \frac{\tau}{N},$$

and hence feasible values of $-\frac{\eta}{P}\lambda_{N+2}$ comprise the set $\{-\frac{\eta}{P}\lambda_{N+2} \in \mathbb{R}^+ : \int_0^T \mathbb{1}_{\{\hat{\pi}(t) < -\frac{\eta}{P}\lambda_{N+2}\}} dt = \frac{\tau}{N}\}$. The optimal λ_{N+2}^* , that minimizes the “cost-to-go” $\frac{P}{\eta} \int_0^T \hat{\pi}(t) \sum_{i=1}^N \mathbb{1}_{\{\hat{\pi}(t) < -\frac{\eta}{P}\lambda_{N+2}\}} dt$, is given by

$$-\frac{\eta}{P}\lambda_{N+2}^* = \inf \left\{ -\frac{\eta}{P}\lambda_{N+2} \in \mathbb{R}^+ : \int_0^T \mathbb{1}_{\{\hat{\pi}(t) < -\frac{\eta}{P}\lambda_{N+2}\}} dt = \frac{\tau}{N} \right\}.$$

To determine $\lambda_{N+1}^*(t)$, combining (14) with $\lambda_i(t) = 0$ results in

$$\dot{\lambda}_{N+1} = -\frac{P}{\eta} N \mathbb{1}_{\{\hat{\pi}(t) < -\frac{\eta}{P}\lambda_{N+2}^*\}} \frac{\partial \hat{\pi}}{\partial t}.$$

Thus, $\lambda_{N+1}|_{u_i^*(t)=0} = H|_{u_i^*(t)=0} = \text{constant}$, which we enforce to be zero. On the other hand, $\lambda_{N+1}|_{u_i^*(t)=1} = -\frac{P}{\eta} N \hat{\pi}(t) + k$, where the integration constant k needs to be determined. Since the Hamiltonian evaluated at $u_i^*(t) = 1$, is

$$H|_{u_i^*(t)=1} = N \left(\frac{P}{\eta} \hat{\pi}(t) + \lambda_{N+2}^* \right) + \lambda_{N+1}|_{u_i^*(t)=1} = N \lambda_{N+2}^* + k = \text{constant},$$

which, as before, we enforce to be zero, we obtain

$$k = -N \lambda_{N+2}^* \Rightarrow \lambda_{N+1}|_{u_i^*(t)=1} = -\frac{P}{\eta} N \hat{\pi}(t) - N \lambda_{N+2}^* = \frac{P}{\eta} N (\pi^* - \hat{\pi}(t)) > 0$$

as $\pi^* > \hat{\pi}(t) \forall t \in S$.

To derive $x_i^*(t)$ for all $i = 1, \dots, N$, we simply substitute the optimal control $u_i^*(t)$ into (4a), and then integrate the resulting first-order linear non-homogeneous ODE using the method of integrating factor, yielding the desired expression. \square

Remark 4. Monotone rearrangement of $\hat{\pi}(t)$: The main insight behind the optimal control derived in Theorem 1 can be obtained by looking at strictly monotone price forecasts, as shown in Fig. 4. In these cases, it is intuitive that the optimal ON periods of the TCLs lie at either end of the interval $[0, T]$. Theorem 1 tells us that the same insight can be extended to non-monotone $\hat{\pi}(t)$, by first computing its monotone rearrangement (p. 276, Ch. 10 in Hardy et al.²⁷) $\hat{\pi}^\dagger(t)$, and then computing




FIGURE 5 Illustration of the optimal control $u_i^*(t)$, $i = 1, \dots, N$, for N TCLs corresponding to non-monotone price forecast $\hat{\pi}(t)$. Left: A non-monotone price forecast $\hat{\pi}(t)$ with $S = [T - \frac{\tau}{N}, T]$ being the subset of times in $[0, T]$, with measure $\frac{\tau}{N}$, that corresponds to minimum price. Middle: $\hat{\pi}^\uparrow(t)$ is the increasing rearrangement of the function $\hat{\pi}(t)$, plotted against the corresponding re-arranged time t^\uparrow . Right: With constraint **C3** inactive, the optimal control $u_i^*(t)$ and optimal indoor temperature trajectories $\theta_{12}^{*(i)}(t)$ are shown for $i = 1, 2$ TCLs. The subscript 12 in $\theta_{12}^{*(i)}(t)$ denotes that constraints **C1** and **C2** are active.




FIGURE 6 Some illustrative examples showing how the uniqueness of optimal control depends on the shape of the price forecast curve $\hat{\pi}(t)$. (a) For this non-monotone price forecast, the set $S = S_1 \cup S_2$ is unique, where $|S| = \frac{\tau}{N}$, and hence $u_i^*(t)$ is unique too. (b) Here also $S = S_1 \cup S_2$ such that $|S| = \frac{\tau}{N}$, however, $\frac{\tau}{N}$ and hence the threshold π^* are such that the set $S_2 \subset V$, where V is the total interval of constancy. Thus, subsets of S_2 can be interchanged with equi-measure subsets of $V \setminus S_2$ without affecting the cost. Therefore, the set S and optimal control $u_i^*(t)$ are not unique in this case. Since these sets are intervals, the number of such interchanges, and hence the number of optimal controls, is uncountable.

the threshold π^* and the ON time set S from this monotone rearrangement as a function of $\frac{\tau}{N}$, as illustrated in Fig. 5. This is especially relevant noting that the typical $\hat{\pi}(t)$ is non-monotone and looks as in Fig. 2.

Remark 5. Non-uniqueness of optimal control: From Theorem 1, the synchronized optimal control $\{u_i^*(t)\}_{i=1}^N$, is unique iff the set S is unique, where S is the pre-image of π^* . Notice that although π^* is unique for any continuous $\hat{\pi}(t)$, uniqueness of S depends on whether there exist time intervals of constancy in price forecast $\hat{\pi}(t)$. For example, in Fig. 5, there is no such interval of constancy, and hence the pre-image set S , and the optimal control $u_i^*(t)$, are unique. This remains true even when $\frac{\tau}{N}$ is large (see Fig. 6(a)). Non-uniqueness, however, can arise if there exist an interval of constancy V , and $\frac{\tau}{N}$ is large enough that S contains at least a subset of V (see Fig. 6(b)). The uncountable number of non-unique solutions arising from such a situation




FIGURE 7 A graphical illustration of the two-sided Skorokhod map $\Psi_{L,U}$ introduced in Definition 1, where the horizontal axis plots the independent variable t and the vertical axis spans $(-\infty, \infty)$. In this example, the boundaries $[L, U] = [20, 25]$ are shown as *thick solid horizontal lines*, and the Skorokhod map $\Psi_{L,U}$ maps the piecewise linear function $y(t)$ (shown above as *dashed blue*) to another piecewise linear function $z(t)$ (shown above as *solid red*), i.e., $z(t) = \Psi_{L,U}(y(t))$. In general, $y(t)$ can be any càdlàg function, i.e., $y(t) \in \mathbb{D}((-\infty, \infty))$, and consequently $z(t) \in \mathbb{D}([L, U])$.

can be resolved by fixing the convention of choosing the optimal control with minimum number of switchings. Finally, notice that a typical day-ahead price forecast is piecewise constant for each hour (see Fig. 2 and the footnote at p. 2); therefore the non-uniqueness issue is rare in practice, and in case it happens, can always be resolved by choosing the optimal control with minimum number of switchings, as before.

4.2 | Solution with Zero Amplitude Chattering

Now we focus on solving (3) subject to constraints **(C1)–(C3)**, under the assumption that the indoor temperature trajectories $\theta_i(t)$ can slide along the boundaries L_{i0} and U_{i0} , which can be thought of as the limits of small amplitude chattering. We assume that sliding along L_{i0} holds the ON mode, while the same along U_{i0} holds the OFF mode. Our objective is to obtain qualitative insight into the solution structure under these simplifying assumptions. In Section 5, we consider the practical case of finite amplitude chattering.

To describe the optimal solution, we next define the two-sided Skorokhod map^{19,20}, which generalizes the one-sided version originally introduced by Skorokhod²⁸. A graphical illustration of the two-sided Skorokhod map is shown in Fig. 7.

Definition 1. Two-sided Skorokhod Map:

Given $0 < L < U < \infty$, and scalar trajectory $y(\cdot) \in \mathbb{D}((-\infty, \infty))$, the two-sided Skorokhod map $\Psi_{L,U} : \mathbb{D}((-\infty, \infty)) \mapsto \mathbb{D}([L, U])$ is defined as $z(t) = \Psi_{L,U}(y(t)) := \Lambda_{L,U} \circ \Psi_{L,\infty}(y(t))$, where

$$\Lambda_{L,U}(\phi)(t) := \phi(t) - \sup_{0 \leq s \leq t} \left([\phi(s) - U]^+ \wedge \inf_{s \leq r \leq t} (\phi(r) - L) \right),$$

$$\Psi_{L,\infty}(y(t)) := y(t) + \sup_{0 \leq s \leq t} [L - y(s)]^+.$$

The following Theorem summarizes the solution for problem (3) with constraints **C1–C3**, under the simplifying assumption of zero amplitude chattering. The optimal controls are shown to be identical to those in Theorem 1. Interestingly, it is shown that the optimal indoor temperature trajectories in this case can be obtained by applying the two-sided Skorokhod maps on the optimal indoor temperature trajectories obtained from Theorem 1. Here, the Skorokhod maps are parameterized by the upper and lower comfort boundaries of the individual TCLs.




FIGURE 8 Comparison of the indoor temperature trajectories for $N = 2$ TCLs with and without constraint **C3** active for the non-monotone price forecast $\hat{\pi}(t)$ as in Fig. 5. When constraint **C3** is active, we assume that the indoor temperature trajectories can slide along the respective comfort range boundaries (zero amplitude chattering). The optimal controls $u_i^*(t)$ (top left) with constraints **C1–C3**, are same as those shown in Fig. 5. However, the optimal indoor temperatures $\theta_{123}^{*(i)}(t)$ (right bottom) are different from $\theta_{12}^{*(i)}(t)$ (left bottom, same as in Fig. 5), for $i = 1, 2$ TCLs. The subscript 123 in $\theta_{123}^{*(i)}(t)$ denotes that constraints **C1–C3** are active.

Theorem 2. Consider problem (3) with constraints **C1–C3**. The optimal controls are synchronizing, and, as in Theorem 1, based on a price forecast threshold π^* , switch between 0 and 1. Assuming zero amplitude chattering to be feasible, the open loop optimal controls $u_i^*(t)$ are identical to those in Theorem 1. For $i = 1, \dots, N$, the optimal states $x_i^*(t)$ are the two-sided Skorokhod maps parameterized by individual comfort ranges $[L_{i0}, U_{i0}]$, acting on respective optimal states from Theorem 1, i.e., $\theta_{123}^{*(i)}(t) = \Psi_{L_{i0}, U_{i0}}(\theta_{12}^{*(i)}(t))$, where $\theta_{123}^{*(i)}(t)$ is the optimal indoor temperature trajectory when constraints **C1–C3** are active, and $\theta_{12}^{*(i)}(t)$ is the same when constraints **C1–C2** are active, for the i^{th} TCL.

Proof. From the necessary conditions for optimality under state inequality constraints^{29,30}, it can be directly verified that the π^* is as in Theorem 1. Hence the monotone rearrangement argument applies as before, and $u_i^*(t)$ are synchronized as function of time. However, the optimal states have different hitting times to the respective boundaries L_{i0} and U_{i0} . For brevity, we provide below a simple graphical argument for the optimal states for strictly monotone (w.l.o.g. decreasing) $\hat{\pi}(t)$. The non-monotone $\hat{\pi}(t)$ can be dealt via the monotone rearrangement, and the non-uniqueness due to constancy can be dealt with by adopting a minimum switching convention, as earlier.

By the argument above, consider strictly decreasing $\hat{\pi}(t)$ as in Fig. 4 left, green curve, and fix the i^{th} home with initial indoor temperature θ_{i0} , and comfort boundaries L_{i0} and U_{i0} . At time $t = 0^+$, the trajectory $\theta_i(t)$ can move either exponentially upward or downward. For $t \in [0, T]$, let us call the set of all feasible trajectories $\theta_i(t) \in [L_{i0}, U_{i0}]$ for which $u_i(0^+) = 0$, as the “initially up-going family”. Similarly, define “initially down-going family” for $u_i(0^+) = 1$. Our proof consists of the following two steps.

Step 1: Finding the optimal indoor temperature trajectory among the “initially up-going family”: We notice that among the “initially up-going family”, it is optimal to hit U_{i0} , since otherwise turning the TCL ON before hitting U_{i0} strictly increases the cost, as $\hat{\pi}(t)$ is strictly decreasing. Similarly, starting from the time at which U_{i0} is hit, it is then optimal to hold till $T - \frac{\tau}{N}$ as sliding along U_{i0} , as per assumption, does not contribute to the cost. For $t \in [T - \frac{\tau}{N}, T]$, we notice that we must keep $u_i(t) = 1$ to respect the energy constraint. Further, notice that any de-synchronization among TCLs increase cost. Thus, the optimal temperature trajectory among the “initially up-going family” looks like those shown in Fig. 8 bottom right. From Definition 1, we find that the optimal indoor temperature trajectory among the “initially up-going family” is $\Psi_{L_{i0}, U_{i0}}(\theta_{12}^{*(i)}(t))$.

Step 2: Showing that any trajectory from the “initially down-going family” has cost strictly larger than the same for the optimal trajectory in Step 1: This can be easily verified by comparing the optimal from Step 1, with any trajectory from the “initially down-going family” using that $\hat{\pi}(t)$ is strictly decreasing.

Combining the above two steps, we conclude that for $\hat{\pi}(t)$ strictly decreasing, the optimal indoor temperature trajectory found in Step 1 is the optimal among all feasible indoor temperature trajectories. Similar argument applies to $\hat{\pi}(t)$ strictly increasing,

and to monotone rearranged version $\hat{\pi}^\uparrow(t)$ in case $\hat{\pi}(t)$ is non-monotone. We eschew the details and illustrate an example in Fig. 8 to help the readers follow our main argument. \square

5 | IMPLEMENTABLE SOLUTION WITH A SPECIFIED MINIMUM SWITCHING PERIOD

Since physical thermostats have a minimum chattering amplitude, or equivalently minimum ON-OFF time period $T_m > 0$, it is important to find an algorithm that explicitly accounts for this device limitation in the control design. Given this parameter T_m , the following Theorem gives an *exact* algorithm to compute the binary optimal control via convexification.

The importance of the Theorem below lies in decoupling the two technical difficulties in solving problem (3) subject to (C1)–(C3), namely the non-convexity of the control and the presence of state inequality constraints. In other words, it allows us to first solve the convexified optimal control problem with the state inequality constraints, and then use Algorithm 1 (introduced as part of the Theorem below) as a post-processing tool to recover the binary optimal controls respecting the prescribed minimum ON-OFF time period T_m .

Theorem 3. For $i = 1, \dots, N$, consider the control convexification $u_i(t) \in \{0, 1\} \mapsto v_i(t) \in [0, 1]$, and let $v_i^*(t) \in [0, 1]$ be the optimal convexified control corresponding to the non-convex optimal control problem (3)–(6) with N TCLs having thermal coefficients $\{\alpha_i, \beta_i\}_{i=1}^N$. Let $\vartheta_i^*(t)$ (resp. $\theta_i^*(t)$) be the indoor temperature trajectory realized by the optimal control $v_i^*(t)$ (resp. $u_i^*(t)$). Then, Algorithm 1 recovers the optimal control $u_i^*(t) \in \{0, 1\}$ from $v_i^*(t) \in [0, 1]$, while guaranteeing that the indoor temperatures $\vartheta_i^*(t)$ and $\theta_i^*(t)$ coincide at the end of each minimum allowable time period of length T_m . In Algorithm 1, $\{\bar{\gamma}_i, \underline{\gamma}_i\}_{i=1}^N$

Algorithm 1 Recovering $u_i^*(t) \in \{0, 1\}$ from $v_i^*(t) \in [0, 1]$

```

1: if  $v_i^*(t) = 0$  OR  $1$  then
2:    $u_i^*(t) = v_i^*(t)$ 
3: else
4:   if  $\vartheta_i^*(t)$  is at upper boundary then
5:      $u_i^*(t) = \begin{cases} 1 & \forall t \in [(j-1)T_m, (j-1)T_m + \bar{\gamma}_i) \\ 0 & \forall t \in [(j-1)T_m + \bar{\gamma}_i, jT_m) \end{cases}$ 
6:     where  $j = 1, \dots, \left\lceil \frac{\text{spt}(v_i^*(t))}{T_m} \right\rceil, t \in [0, T]$ 
7:   else  $\vartheta_i^*(t)$  is at lower boundary
8:      $u_i^*(t) = \begin{cases} 0 & \forall t \in [(j-1)T_m, jT_m - \underline{\gamma}_i) \\ 1 & \forall t \in [jT_m - \underline{\gamma}_i, jT_m) \end{cases}$ 
9:     where  $j = 1, \dots, \left\lceil \frac{\text{spt}(v_i^*(t))}{T_m} \right\rceil, t \in [0, T]$ 
10:  end if
11: end if

```

are time duration pairs such that the binary optimal trajectory $u_i^*(t)$ consists of two duty cycles: $\bar{\gamma}_i/T_m$ and $\underline{\gamma}_i/T_m$, where

$$\bar{\gamma}_i = \frac{1}{\alpha_i} \log \left(1 + \alpha_i \int_0^{T_m} e^{\alpha_i s} v_i^*(s) ds \right), \quad (18)$$

$$\underline{\gamma}_i = \frac{1}{\alpha_i} \log \left(\frac{1}{1 - \alpha_i e^{-\alpha_i T_m} \int_0^{T_m} e^{\alpha_i s} v_i^*(s) ds} \right). \quad (19)$$

Proof. Since $v_i^*(t)$ is binary iff the respective upper and lower comfort boundaries are not hit, hence $u_i^*(t) = v_i^*(t)$ when $v_i^*(t) = 0$ or 1 . We know that $v_i^*(t) \in (0, 1)$ iff $\vartheta_i^*(t)$ is at either upper or lower boundary. Clearly, at the upper (lower) boundary, such cycles should begin with an ON (OFF) segment, and end with an OFF (ON) segment. Matching indoor temperature values at

each end of these switching period means $\vartheta_i^*(T_m) = \theta_i^*(T_m)$, which gives

$$\int_0^{T_m} e^{\alpha_r s} v_i^*(s) ds = \int_0^{T_m} e^{\alpha_r s} u_i^*(s) ds. \quad (20)$$

At upper boundary, RHS of (20) equals $\int_0^{\bar{\gamma}_i} e^{\alpha_r s} ds$, which solved for $\bar{\gamma}_i$ yields (18). At lower boundary, RHS of (20) equals $\int_{T_m - \underline{\gamma}_i}^{T_m} e^{\alpha_r s} ds$, which solved for $\underline{\gamma}_i$ yields (19). \square

It should be noted that although the indoor temperature trajectories θ_i^* and ϑ_i^* (corresponding to the controls u_i^* and v_i^* , respectively) periodically coincide at the end of each time period T_m , the cost of the respective solutions are not necessarily the same. The solution corresponding to v_i^* is the theoretical, though non-implementable, optimal solution; the solution corresponding to u_i^* is a suboptimal solution that has the closest implementable trajectory.

Remark 6. Our optimal control problem can be viewed as an optimal control problem of a switched system; see e.g.³¹. The use of an approximate relaxed version of the problem with a convexified control set, as is done here, has been studied in the literature (see³² and references therein). More recently, a so-called embedding principle has been investigated in the works^{31,33,34}. In those works, a relaxed version of the optimal control problem is first solved in a convexified input set, and then a projection operator is used to obtain the input in the original discrete set. This conceptual path is also followed here in Theorem 3. Nevertheless, the goal with which the techniques are used is different in our approach. In the references mentioned, one of the main concerns is to address the limitation that the switched optimal control problem, without imposing additional assumptions related to the possibility of chattering, only has a solution when the space of controls is enlarged to the space of relaxed controls. The goal then becomes to construct approximate solutions that are consistent, in the sense that in the limit they converge to the optimal solution. In our case, the starting point is a physical limitation of the system imposing a minimum ON-OFF time period of $T_m > 0$. This physical limitation itself prevents infinite frequency chattering, and the problem becomes to construct solutions to an approximate problem such that they can be easily projected into a physical realizable solution space.

6 | NUMERICAL SIMULATION

To illustrate how the LSE can use the results derived so far for the purpose of demand response, we now provide a numerical example where the LSE computes the day-ahead minimum cost energy procurement for its $N = 500$ customers' TCLs based on ERCOT day-head price forecast ($\hat{\pi}(t)$) data as shown in Fig. 2, and the ambient temperature forecast ($\hat{\theta}_a(t)$) data for the same day (August 10, 2015) available from a weather station in Houston, Texas. These forecast data and the real-time ambient temperature ($\theta_a(t)$) data on August 11, 2015, are shown in Fig. 9.

With the initial conditions and parameters of the heterogeneous TCL population as in Section V-A in Halder et al.¹⁸, $\bar{\tau} = \frac{1}{3}$ (which was verified to be feasible using (8)), comfort tolerances $\{\Delta_i\}_{i=1}^N$ sampled randomly from a uniform distribution over $[0.1^\circ\text{C}, 1.1^\circ\text{C}]$, and for $\hat{\pi}(t)$ and $\hat{\theta}_a(t)$ as in Fig. 9, the LSE solves the optimal control problem (3)–(6) by first convexifying the controls $u_i(t) \in \{0, 1\} \mapsto v_i(t) \in [0, 1]$ for $i = 1, \dots, N$, and then recovering the optimal controls $\{u_i^*\}_{i=1}^N$ using Theorem 3. For this computation, we used 1 minute time-step for Euler discretization of dynamics (4a), and solved the resulting LP with 1 million 440 thousand decision variables (see Section 3.3) using MATLAB `linprog`. In Fig. 10, we show the *convexified* optimal controls $v_i^*(t)$ (gray curves) and corresponding indoor temperature trajectories $\vartheta_i^*(t)$ (black curves) for two representative TCLs out of the total $N = 500$ TCLs. This computation was followed by applying Algorithm 1 to evaluate the mapping $(v_i^*(t), \vartheta_i^*(t)) \mapsto (u_i^*(t), \theta_i^*(t))$ with $T_m = 1.5$ minutes. The resulting optimal aggregate power consumption trajectory $P_{\text{total}}^{\text{ref}}(t) = P_e \sum_{i=1}^N u_i^*(t)$ is shown as the black curve in Fig. 11. We emphasize again that the black curve in Fig. 11 is the optimal *planned* aggregate consumption, computed by the LSE ahead of the actual time duration under consideration (in our case, 24 hours ahead). In operation, the LSE also needs to implement *real-time* setpoint control across its customers' TCL population, so as to make the real-time aggregate consumption track the planned optimal aggregate consumption, given the mismatch between the forecasted and real-time ambient temperatures. The brick colored curve in Fig. 11 corresponds to the real-time aggregate consumption for the TCL population with same T_m , and real-time ambient temperature $\theta_a(t)$ as in the solid blue curve in Fig. 9, for a PID velocity control gain tuple (k_p, k_i, k_d) used to control the setpoint boundaries as part of a mixed centralized-decentralized control. We refer the readers to Section III.1 in Halder et al.¹⁸ for details on the real-time setpoint control. The purpose of Fig. 11 is to




FIGURE 9 The day-ahead ambient temperature forecast $\hat{\theta}_a(t)$ (dashed blue), real-time ambient temperature $\theta_a(t)$ (solid blue), ERCOT day-ahead price $\hat{\pi}(t)$ (dashed green), and ERCOT real-time price $\pi(t)$ (solid green, not used for computation in this paper) data for Aug. 11, 2015 in Houston.

highlight how the solution of the open-loop optimal control problem (3)–(6) can be used by the LSE as a reference aggregate consumption to be tracked in real-time, to elicit demand response.

7 | CONCLUDING REMARKS

In this paper, we have addressed how an aggregator or load serving entity can design an optimal aggregate power consumption trajectory for a population of thermostatically controlled loads. We have formulated this operational planning problem as a deterministic optimal control problem in terms of the day-ahead price forecast, ambient temperature forecast, and an energy budget available from the load forecast. A direct numerical approach to solve the problem is computationally hard. We use tools from optimal control theory to gain analytic insights into the solution of the problem of designing optimal power consumption while respecting individual comfort range constraints. A numerical example is worked out to illustrate how an LSE can use the optimal aggregate power consumption trajectory computed offline, as a reference signal to be tracked in real-time by its customers' TCL population for the purpose of demand response.

ACKNOWLEDGMENTS

This work is supported in part by NSF Contract 1760554, ECCS-1546682, NSF Science & Technology Center Grant CCF-0939370, and the Power Systems Engineering Research Center (PSERC).

References

1. U.S. Dept. Energy. Benefits of demand response in electricity markets and recommendations for achieving them. *US Dept. of Energy, Washington, DC, USA, Tech. Rep.* 2006.
2. Callaway DS, Hiskens IA. Achieving controllability of electric loads. *Proceedings of the IEEE.* 2011;99(1):184–199.
3. Chong CY, Debs AS. Statistical synthesis of power system functional load models. *18th IEEE Conference on Decision and Control including the Symposium on Adaptive Processes.* 1979;18:264–269.
4. Malhame R, Chong C-Y. Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system. *IEEE Transactions on Automatic Control.* 1985;30(9):854–860.




FIGURE 10 Convexified optimal control $v_i^*(t) \in [0, 1]$ (gray) and corresponding indoor temperature $\vartheta_i^*(t)$ in $^{\circ}\text{C}$ (black), $i = 1, 2$, trajectories for two out of the total population of 500 TCL customers for which the LSE designs the optimal power consumption corresponding to the day-ahead price and ambient temperature forecast shown in Fig. 9. Details of the simulation setup are described in Section 6. For the two representative TCLs shown above, their lower (resp. upper) comfort boundaries L_i (resp. U_i) are depicted as blue (resp. red) thick horizontal lines. Specifically, $[L_1, U_1] \equiv [18.40^{\circ}\text{C}, 19.48^{\circ}\text{C}]$ and $[L_2, U_2] \equiv [20.08^{\circ}\text{C}, 22.22^{\circ}\text{C}]$, as shown. The parameters $(\alpha_1, \alpha_2) = (4.4032, 4.1067) \times 10^{-3}$ seconds $^{-1}$, and $(\beta_1, \beta_2) = (8.4510, 8.5286) \times 10^{-3}$ $\frac{^{\circ}\text{C}}{\text{kW seconds}}$. The optimal control trajectories $v_i^*(t) \in [0, 1]$ (gray) are scaled between the respective comfort boundaries, i.e., $v_i^* = 1$ (resp. 0) at the upper (resp. lower) boundary. Notice that $v_i^*(t)$ takes fractional values whenever $\vartheta_i^*(t)$ lies at either upper or lower comfort boundary, as mentioned in the proof of Theorem 3.

5. Callaway DS. Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy. *Energy Conversion and Management*. 2009;50(5):1389–1400.
6. Bashash S, Fathy K. Modeling and control insights into demand-side energy management through setpoint control of thermostatic loads. *2011 American Control Conference (ACC)*. 2011;4546–4553.
7. Kundu S, Sinitsyn N, Backhaus S, Hiskens I. Modeling and control of thermostatically controlled loads. *17th Power Systems Computation Conference, Stockholm, Sweden*. 2011.
8. Mathieu JL, Koch S, Callaway DS. State estimation and control of electric loads to manage real-time energy imbalance. *IEEE Transactions on Power Systems*. 2013;28(1):430–440.
9. Zhang W, Lian J, Chang C-Y, Kalsi K. Aggregated modeling and control of air conditioning loads for demand response. *IEEE Transactions on Power Systems*. 2013;28(4):4655–4664.
10. Totu LC, Wisniewski R. Demand response of thermostatic loads by optimized switching-fraction broadcast. *IFAC Proceedings Volumes*. 2014;47(3):9956–9961.
11. Ghaffari A, Moura S, Krstic M. Analytic modeling and integral control of heterogeneous thermostatically controlled load populations. *ASME 2014 Dynamic Systems and Control Conference*. 2014;V002T22A002–V002T22A002.
12. Grammatico S, Gentile B, Parise F, Lygeros J. A mean field control approach for demand side management of large populations of thermostatically controlled loads. *2015 European Control Conference (ECC)*. 2015;3548–3553.




FIGURE 11 The *black* curve shown above is the optimal power consumption trajectory $P_{\text{total}}^{\text{ref}}(t) = P_e \sum_{i=1}^N u_i^*(t)$ computed by solving the optimal control problem (3)–(6), via Algorithm 1 with $T_m = 1.5$ minutes. The *brick* colored curve is the real-time controlled aggregate consumption $P_{\text{total}}(t)$ corresponding to control gain tuple $(k_p, k_i, k_d) = (10^{-4}, 10^{-6}, 10^{-4})$ used to move the thermostatic boundaries (see Section III.1 in¹⁸ for details). The LSE can invoke demand response by controlling the setpoints of the TCL population in such a way that their real-time aggregate consumption $P_{\text{total}}(t)$ track the reference aggregate consumption $P_{\text{total}}^{\text{ref}}(t)$. The tracking error between the two curves depend on the forecasted versus real-time ambient temperature mismatch, as well as on the thermal inertia of the TCLs in the population.

13. Meyn SP, Barooah P, Bušić A, Chen Y, Ehren J. Ancillary service to the grid using intelligent deferrable loads. *IEEE Transactions on Automatic Control*. 2015;60(11):2847–2862.
14. Paccagnan D, Kamgarpour M, Lygeros J. On the range of feasible power trajectories for a population of thermostatically controlled loads. *IEEE 54th Annual Conference on Decision and Control (CDC)*. 2015;5883–5888.
15. Ruelens F, Claessens BJ, Vandael S, Iacovella S, Vingerhoets P, Belmans R. Demand response of a heterogeneous cluster of electric water heaters using batch reinforcement learning. *2014 Power Systems Computation Conference*. 2014;1-7.
16. Mathieu JL, Kamgarpour M, Lygeros J, Andersson G, Callaway DS. Arbitraging intraday wholesale energy market prices with aggregations of thermostatic loads. *IEEE Transactions on Power Systems*. 2015;30(2):763–772.
17. Halder A, Geng X, Sharma G, Xie L, Kumar PR. A control system framework for privacy preserving demand response of thermal inertial loads. *2015 IEEE International Conference on Smart Grid Communications (SmartGridComm)*. 2015;181–186.
18. Halder A, Geng X, Kumar PR, Xie L. Architecture and algorithms for privacy preserving thermal inertial load management by a load serving entity. *IEEE Transactions on Power Systems*. 2017;32(4):3275–3286.
19. Kruk Ł, Lehoczyk J, Ramanan K, Shreve S. An explicit formula for the Skorokhod map on $[0, a]$. *The Annals of Probability*. 2007;1740–1768.
20. Kruk Ł, Lehoczyk J, Ramanan K, Shreve S. Double Skorokhod map and reneging real-time queues. *Markov Processes and Related Topics: A Festschrift for Thomas G. Kurtz*. 2008;169–193.
21. <http://tinyurl.com/z3gmvt6>. accessed on Sep. 17, 2016.
22. Moura S, Bendtsen J, Ruiz V. Observer design for boundary coupled PDEs: Application to thermostatically controlled loads in smart grids. *2013 IEEE 52nd Annual Conference on Decision and Control (CDC)*. 2013;6286–6291.
23. Moura S, Bendtsen J, Ruiz V. Parameter identification of aggregated thermostatically controlled loads for smart grids using PDE techniques. *International Journal of Control*. 2014;87(7):1373–1386.
24. Gurobi Optimization Inc.. Gurobi Optimizer Reference Manual. <http://www.gurobi.com>. 2015.

25. Schrijver A. *Combinatorial optimization: polyhedra and efficiency*. Springer Science & Business Media; 2002.
26. Pontryagin LS, Boltyanskii VG, Gamkrelidze RV, Mishchenko E. *Mathematical theory of optimal processes*. Wiley Interscience, New York; 1962.
27. Hardy GH, Littlewood JE, Pólya G. *Inequalities*. Cambridge university press; 1952.
28. Skorokhod AV. Stochastic equations for diffusion processes in a bounded region. *Theory of Probability & Its Applications*. 1961;6(3):264–274.
29. Bryson AE, Denham WF, Dreyfus SE. Optimal programming problems with inequality constraints. *AIAA journal*. 1963;1(11):2544–2550.
30. Vinter R. *Optimal control*. Springer Science & Business Media; 2010.
31. Bengea SC, DeCarlo RA. Optimal control of switching systems. *Automatica*. 2005;41(1):11–27.
32. Berkovitz L, Medhin N. *Nonlinear Optimal Control Theory*. Chapman and Hall/CRC; 2012.
33. Vasudevan R, Gonzalez H, Bajcsy R, Sastry SS. Consistent approximations for the optimal control of constrained switched systems—Part 1: a conceptual algorithm. *SIAM Journal on Control and Optimization*. 2013;51(6):4463–4483.
34. Chen H, Zhang W. On weak topology for optimal control of switched nonlinear systems. *Automatica*. 2017;81:409–415.