2020 59th IEEE Conference on Decision and Control (CDC)
Jeju Island, Republic of Korea, December 14-18, 2020

A Privacy Preserving Model-Free Optimization and Control Framework for Demand Response from Residential Thermal Loads

S. Sivaranjani, P. R. Kumar and Le Xie

Abstract— **We consider the problem of optimizing the cost of procuring electricity for a large collection of homes managed by a load serving entity, by pre-cooling or pre-heating the thermal inertial loads in the homes to avoid procuring power during periods of peak electricity pricing. We would like to accomplish this objective in a completely privacy-preserving and model-free manner, that is, without direct access to the state variables (temperatures or power consumption) or the dynamical models (thermal characteristics) of individual homes, while guaranteeing personal comfort constraints of the consumers. We propose a two-stage optimization and control framework to address this problem. In the first stage, we use a long short-term memory (LSTM) network to predict hourly electricity prices, based on historical pricing data and weather forecasts. Given the hourly price forecast and thermal models of the homes, the problem of designing an optimal power consumption trajectory that minimizes the total electricity procurement cost for the collection of thermal loads can be formulated as a large-scale integer program (with millions of variables) due to the on-off cyclical dynamics of such loads. We provide a simple heuristic relaxation to make this large-scale optimization problem model-free and computationally tractable. In the second stage, we translate the results of this optimization problem into distributed open-loop control laws that can be implemented at individual homes without measuring or estimating their state variables, while simultaneously ensuring consumer comfort constraints. We demonstrate the performance of this approach on a large-scale test case comprising of 500 homes in the Houston area and benchmark its performance against a direct model-based optimization and control solution.**

I. INTRODUCTION

In traditional power grids, uncertainties typically arise in the demand-side and are countered by an increase or decrease in the generation of power using operating reserves. However, the large-scale integration of renewables has introduced additional uncertainties into the supply-side, due to the variability of renewable energy resources. Since generation from renewable energy resources cannot be directly controlled, this new uncertainty in the supply-side will need to be offset by tuning the demand via controllable loads [1]-[3]. This approach, known as demand response, is a rapidly emerging operational paradigm in the modern power grid, wherein an aggregator or load serving entity (LSE) manages a collection of controllable loads that function as a new type of operating reserve, albeit one that is now on the demand-side [4].

Thermal inertial loads such as air conditioners (ACs), heaters and refrigerators comprise nearly half of the residential demand in the United States [5], and are attractive candidates for demand response due to their ability to store energy and alter (delay or advance) consumption without causing significant discomfort to the consumer [6][7]. This demand response potential can be exploited by LSEs to provide ancillary services to the grid, while simultaneously reducing energy costs for individual consumers [8]-[10]. Early instances of demand response from thermal inertial loads typically employed coarse models of the duty cycles of the loads to compute pre-defined trajectories for load curtailment during periods of peak pricing [11]-[13]. More recent approaches involve estimating the models and states of the loads, and utilizing this information to design and track a desired power trajectory that minimizes costs or provides operational support to the grid [14]-[16].

In this context, it is desirable to develop model-free privacy-preserving approaches for thermal inertial load management, for three reasons. First, thermal models can be used to infer information about the size, layout and construction of the consumers' homes, which may constitute a violation of consumer privacy. Second, it is challenging to obtain such models for demand response programs involving large-scale participation from thousands of homes, even with intrusive measurement and monitoring. Finally, for privacy reasons, it is not desirable to measure the temperatures or power consumption of individual homes. Recently, learning-based model-free approaches for the optimization and control of thermal loads have been proposed [17][18]; however, these approaches are typically not privacy-preserving in that they still involve measuring the internal temperatures and power consumption profiles of homes. Alternatively, privacy-preserving approaches to thermal inertial load management, wherein the power consumption of individual homes is not directly measured have been proposed [19]-[21]. However, all of these approaches still utilize thermal models of homes to compute and implement optimal control actions for electricity cost minimization. The aim of this paper is to bridge this gap by proposing a model-free privacy-preserving approach for the management of thermal inertial loads.

Specifically, we consider the problem of minimizing the cost of procuring electricity for a large collection of homes managed by an LSE. The objective is to pre-cool (or pre-heat) homes by controlling residential thermal loads, in order to avoid procuring power during periods of peak electricity pricing. Further, we would like to accomplish this objective in a completely privacy-preserving and model-free manner, that is, without direct access to the state variables (temperatures and power consumption) or models (thermal characteristics) of individual homes.

The authors are with the Department of Electrical and Computer Engineering, Texas A&M University, {sivaranjani,prk,le.xie}@tamu.edu. This work was supported in part by NSF grants CCF-1934904, ECCS-2038963, ECCS-1839616, NSF Science Technology Center Grant CCF-0939370, and Department of Energy Contract DE-EE0009031.

We propose a two-stage optimization and control framework to address this problem. In the first stage, we use a long-short term memory (LSTM) based recurrent neural network architecture to forecast hourly electricity prices from historical price data and weather forecasts. Given the hourly price forecast and the thermal models of the homes, the problem of designing an optimal power consumption trajectory that minimizes the total electricity procurement cost can be formulated as a large-scale integer program (with millions of variables) due to the on-off cyclical dynamics of such loads. This integer program has typically been solved using linear relaxations or dynamic programming [19][22], with explicit closed-form solutions available in special cases where prices are assumed to be monotone [23]. In this paper, we propose a simple heuristic relaxation to convert this large-scale optimization problem into a model-free optimization problem that can be solved in an explicit and computationally tractable manner. In the second stage, we translate the results of this optimization problem into distributed open-loop control laws that can be implemented at the individual homes without measuring or estimating their state variables, while respecting consumer comfort constraints. We demonstrate the performance of this approach on a large-scale test case comprising of 500 homes in the Houston area, with pricing data from the Electric Reliability Council of Texas (ERCOT), and benchmark the performance of the proposed approach by comparing it with the direct model-based approach in [19].

Notation: $\mathbb{R}$, $\mathbb{R}_+$, and $\mathbb{R}^n$ denote the sets of real numbers, positive real numbers including zero, and n-dimensional real vectors respectively. Given $a, b \in \mathbb{R}$, $a \wedge b = \begin{cases} a, & a < b \\ b, & a > b \end{cases}$ and $a \vee b = \begin{cases} a, & a > b \\ b, & a < b \end{cases}$. Given two sets A and B, $A \backslash B$ represents the set of all elements of A that are not in B. We denote the the Laplace density function with zero mean and scale parameter $a \in \mathbb{R}_+ \backslash \{0\}$ by $\text{Lap}(a)$. The gamma density function with parameters $a, b \in \mathbb{R}_+ \backslash \{0\}$ is denoted by $\Gamma(a, b)$ and the exponential density function with rate $\lambda \in \mathbb{R}_+ \backslash \{0\}$ is denoted by $\text{Exp}(\lambda)$. We denote by $\mathcal{N}(\mu, \sigma, a, b)$ a truncated univariate Gaussian density function with mean μ, standard deviation σ and support $[a, b]$.

II. PROBLEM FORMULATION

We begin by describing the model of a collection of residential thermal loads, and formulate the problem of minimizing the electricity procurement cost. For simplicity, we assume that all the loads are air conditioners (ACs). Note that the same analysis can be carried out for heaters, with the objective of pre-heating, rather than pre-cooling homes.

A. System Model

Consider a population of N homes with controllable ACs managed by a load serving entity (LSE). Assume that each home has a temperature set point that is private to the consumer, denoted by s_i, and a comfort range Δ_i, $i \in \{1, 2, \dots, N\}$, which denotes the deviation from the set point that each consumer is willing to tolerate. Therefore, the temperature of the i-th home at any time $t \in \mathbb{R}_+$, denoted by $\theta_i(t)$, must lie in the comfort band $[L_{i0}, U_{i0}] = [s_i - \Delta_i, s_i + \Delta_i]$. The flexibility of the i-th consumer, $i \in \{1, 2, \dots, N\}$, can be quantified by the range of the consumer's comfort band, that is, $2\Delta_i$. The temperature dynamics of the i-th home, $i \in \{1, 2, \dots, N\}$, is governed by

$$\dot{\theta}_i(t) = -\alpha_i(\theta_i(t) - \theta_a(t)) - \beta_i P_i \sigma_i(t), \tag{1}$$

where $\theta_a(t)$ represents the ambient temperature at time $t \in \mathbb{R}_+$, P_i represents the power consumption of the i-th AC, α_i and β_i represent the heating time constant (h^{-1}) and thermal conductivity ($^\circ C/kWh$) of the i-th home, and $\sigma_i(t) \in \{0, 1\}$ denotes the ON/OFF state of the i-th AC at time $t \in \mathbb{R}_+$, where $\sigma_i(t) = 1$ indicates that the AC is ON and $\sigma_i(t) = 0$ indicates that the AC is OFF. When the AC is OFF, the temperature of the home rises until it reaches the upper bound of the consumer's comfort band U_{i0}, at which point the AC turns ON. Similarly, when the temperature reaches the lower bound of the comfort band, L_{i0}, the AC turns OFF. Therefore, the switching behavior of the i-th AC, $i \in \{1, 2, \dots, N\}$, can be defined as

$$\sigma_i(t) = \begin{cases} 1, & \theta_i(t) = U_{i0} \\ 0, & \theta_i(t) = L_{i0} \\ \sigma_i(t^-), & \text{otherwise.} \end{cases} \tag{2}$$

The total electrical power consumed by the population of ACs is given by $P_{total} = \sum_{i=1}^{N} P_i/\eta_i$, where η_i is the coefficient of performance of the i-th AC.

B. Optimization Problem

Define the indicator variable $u_i(t) : \mathbb{R}_+ \to \{0, 1\}$, $\forall i \in \{1, 2, \dots, N\}$ where $u_i(t) = 1$ if the i-th AC is ON at time $t \in \mathbb{R}_+$, and $u_i(t) = 0$ otherwise. We also denote the total number of ACs that are ON at any time $t \in \mathbb{R}_+$ by $n_{ON}(t)$. For simplicity, we assume without loss of generality that all the ACs have an identical power consumption and coefficients of performance, that is, $P_i = P$ and $\eta_i = \eta$, $\forall i \in \{1, 2, \dots, N\}$. Let the electricity price forecast and ambient temperature forecast at time $t \in \mathbb{R}_+$ be denoted by $\hat{\pi}(t) : \mathbb{R}_+ \to \mathbb{R}_+$ and $\hat{\theta}_a(t) : \mathbb{R}_+ \to \mathbb{R}$ respectively. If these forecasts are known over a T-hour horizon, that is, $\forall t \in [0, T], T \in \mathbb{R}_+ \backslash \{0\}$, then, the problem of minimizing the total cost of procuring electricity by the LSE for the collection of ACs over the time horizon $[0, T]$ can be formulated as

$$\begin{aligned} \mathcal{P} : & \min_{u_1(t), \dots, u_N(t) \in \{0,1\}^N} \frac{P}{\eta} \int_0^T \widehat{\pi}(t) \sum_{i=1}^{N} u_i(t) \mathrm{d}t \\ \text{s.t.} \quad & \dot{\theta}_i(t) = -\alpha_i \left(\theta_i(t) - \widehat{\theta}_a(t) \right) - \beta_i P u_i(t) \\ & \frac{P}{\eta} \int_0^T \sum_{i=1}^{N} u_i(t) \mathrm{d}t \leq E \\ & L_{i0} \leq \theta_i(t) \leq U_{i0}, \end{aligned} \tag{3}$$

where $E > 0$ is the maximum energy budget of the LSE for the time horizon $[0, T]$.

Assumption 1: We make the following assumptions pertaining to the feasibility of the optimization problem $\mathcal{P}$.

- Without loss of generality, we assume that the initial temperatures are within the user's comfort constraints, that is, $\theta_i(0) \in [L_{i0}, U_{i0}]$.
- For every $i \in \{1, 2, \ldots, N\}$, when the states are at the upper or lower bound of the comfort band $[L_{i0}, U_{i0}]$, there exists a control policy that can maintain the state inside the comfort band. In other words, the dynamics (1) are such that for all possible $\hat{\theta}_i(t)$, the temperature $\theta_i(t)$ increases with $\sigma_i(t) = 0$, and decreases with $\sigma_i = 1$, or $\forall t \in \mathbb{R}_+$ and $i \in \{1, 2, \ldots, N\}$,

$$-\alpha_i(L_{i0} - \hat{\theta}_i(t)) > 0, \quad -\alpha_i(U_{i0} - \hat{\theta}_i(t)) - \beta_i < 0.$$

Note that the control inputs to maintain the temperature at the upper or lower comfort bounds are given by $u_i^{UP}(t) = \frac{\alpha_i}{\beta_i}(\theta_i(t) - U_{i0})$ and $u_i^{DOWN}(t) = \frac{\alpha_i}{\beta_i}(\theta_i(t) - L_{i0})$ respectively.

C. Model-based Solution for Benchmarking

We now outline a model-based approach to obtain a power reference trajectory for the collection of loads by solving the optimization problem (3), and design a privacy-preserving control law to track this power reference trajectory [19]. We will later use this approach as a benchmark against which we will validate our proposed model-free solution.

If the dynamics of the thermal loads (1) are known, the optimization problem $\mathcal{P}$ can be discretized in the time variable t and directly solved as a mixed-integer linear program (MILP). However, for N homes with a discretization time step of 1 minute, the MILP would involve n $N \times 2 \times 24 \times 60$ variables, which would be prohibitively large (of the order of millions of variables) for hundreds or thousands of homes. Hence, the typical approach to solving this MILP involves a linear programming (LP) relaxation, where the integer variable $u_i(t)$ is allowed to vary continuously in the interval $[0, 1]$, that is $u_i(t) : \mathbb{R}_+ \to [0, 1], i \in \{1, 2, \ldots, N\}$. Then, $u_i(t)$ can be interpreted as the fraction of time that the i-th AC is ON during each discretization time interval. Let $\{u_i^*(t)\}_{i=1}^N$ be the solution to the optimization problem $\mathcal{P}$. Then, the optimal power reference trajectory can be computed as $P_{total}^{ref}(t) = \frac{P}{\eta}\sum_{i=1}^N u_i^*(t)$.

Privacy-preserving Implementation: In order to track this power reference trajectory in a privacy-preserving manner, assume that the LSE does not have access to states of the home including its set point s_i, temperature $\theta_i(t)$, and the state of its AC, $\sigma_i(t)$. First, the LSE estimates the total demand $P_{\text{total}}(t)$ in a privacy-preserving manner as follows. The i-th home, $i \in \{1, 2, \ldots, N\}$, reports, with probability $p \in [0, 1]$, a corrupted power consumption $\hat{P}_i = P_i + n_i$, where $n_i \sim \Gamma\left(\frac{1}{pN}, \frac{\epsilon}{P_e}\right)$ is chosen independently and distributed identically among homes. With this setup, it can be shown that the total power can be estimated in a differentially private manner as $P_{\text{total}}(t) = \frac{N}{\hat{N}}\sum_{i=1}^{\hat{N}} P_i + n$, where $n \sim \operatorname{Lap}\left(\frac{P_e}{\epsilon}\right)$, and $pN = \hat{N}$, where $\hat{N}$ is the number of homes that report their noise-corrupted power consumption. Next, the LSE measures the deviation of the total power consumption of the homes, $P_{total}(t)$ from the optimal power reference trajectory $P_{\text{total}}^{\text{ref}}(t)$ and uses a simple PID controller with proportional, integral, and derivative gains k_p, k_i, and k_d respectively, to compute a velocity control signal $v(t) = k_p e(t) + k_i \int_0^t e(s)\mathrm{d}s + k_d \frac{\mathrm{d}e}{\mathrm{d}t}, e(t) = P_{\text{total}}(t) - P_{\text{total}}^{\text{ref}}(t)$, which is broadcast to all homes. The i -th AC, $i \in \{1, 2, \ldots, N\}$ then locally computes its new set point as $s_i(t) = \Delta_i v(t)$, and adjusts its comfort band as $[L_{it}, U_{it}] \subseteq [L_{i0}, U_{i0}]$, where

$$\begin{aligned} L_{it} &= \min\left(U_{i0}, \max\left(L_{i0}, s_i(t) - \Delta_i\right)\right) \\ U_{it} &= \max\left(L_{i0}, \min\left(U_{i0}, s_i(t) + \Delta_i\right)\right). \end{aligned} \tag{4}$$

In this manner, the temperatures of individual homes can be locally regulated in a privacy-preserving manner such that their aggregate power consumption tracks the optimal reference trajectory.

D. Problem Statement

We now state the problem addressed in this paper.

Problem: Given historical hourly data of electricity prices and ambient temperatures, and the ambient temperature forecast $\hat{\theta}_a(t)$ over a time horizon $[0, T]$, the aim of this paper is to (i) solve optimization problem $\mathcal{P}$ without explicit knowledge of the values of the thermal parameters α_i and β_i, $i \in \{1, 2, \ldots, N\}$ in (3), and (ii) design $\sigma_i(t), i \in \{1, 2, \ldots, N\}$ that results in the optimal power consumption determined by the solution of (3) when implemented locally at each AC $i \in \{1, 2, \ldots, N\}$, without access to the state variables $\theta_i(t)$ or $\sigma_i(t)$ and power consumption $P_i(t)$ or $P_{total}(t)$ by the LSE.

III. MODEL-FREE PRIVACY-PRESERVING OPTIMIZATION AND CONTROL FRAMEWORK

In this section, we present a two-stage approach to solve the problem considered in Section II-D. In the first stage, we begin by forecasting hourly electricity prices based on historical price data and ambient temperature forecasts. We then propose a heuristic relaxation to solve the optimization problem $\mathcal{P}$ in a model-free manner. In the second stage, we discuss control laws for the implementation of this solution.

A. Stage 1: Optimization

We begin by describing how the price forecast $\hat{\pi}(t)$ can be obtained from historical data.

1) LSTM-based Price Forecasting: Given the ambient temperature forecast $\hat{\theta}_a(t)$ over the horizon $t \in [0, T]$, we begin by using a long short-term memory (LSTM) neural network to forecast the hourly electricity price $\hat{\pi}(t)$, $t \in [0, T]$. We choose to use an LSTM-based prediction, since its memory structure allows us to capture features like seasonal and daily variations in prices. Real-time electricity prices vary rapidly on a minute-by-minute basis. However, significant variations are typically observed at the hourly level, and most procurement by the LSE is also carried out at this time scale. Therefore, we begin by averaging intra-hourly historical data to obtain hourly electricity price data on each day. Similarly, we obtain historical temperature data on an hourly time scale. These hourly price and temperature datasets serve as the inputs to the LSTM.

Remark 1: The window of prediction for the LSTM is chosen based on two considerations. First, in our simulations, we determined that highly accurate price predictions can be made in short time windows of less than four hours. Second, we require that the prediction window is larger than the sum of two time windows T_{ON} and T_{OFF}, defined as follows:

- T_{ON}: the average time required to cool a home from its upper comfort bound to its lower comfort bound, that is, the average over all $i \in \{1, 2, \ldots, N\}$ of the smallest time $T_{ON,i}$ such that $\theta_i(0) = U_{i0}$ and $\theta_i(T_{ON,i}) = L_{i0}$ with $\sigma_i(t) = 1, \forall t \in [0, T_{ON,i}]$, and
- T_{OFF}: the average 'duty cycle' of the residential thermal loads, that is, the average over all $i \in \{1, 2, \ldots, N\}$ of the smallest amount of time $T_{OFF,i}$, such that $\theta_i(T_{OFF}) = U_{i0}$, given that $\theta_i(0) = L_{i0}$ and $\sigma_i(t) = 0, \forall t \in [0, T_{OFF,i}]$.

This is to account for the fact that decisions to pre-cool a home will need to be taken at least $(T_{ON} + T_{OFF})$ amount of time before price peaks for a feasible implementation.

2) Model-free Optimization: In order to solve the optimization problem $\mathcal{P}$ without knowledge of the dynamics of individual homes, we begin by making an assumption about the price forecast $\hat{\pi}(t), t \in [0, T]$.

Assumption 2: We assume that the price forecast $\hat{\pi}(t)$ is *unimodal* over $t \in [0, T]$, that is, there exists $t_{PEAK} \in [0, T]$, such that $\hat{\pi}(t)$ is monotonically increasing $\forall t \leq t_{PEAK}$, and monotonically decreasing $\forall t > t_{PEAK}$.

This assumption is not unreasonable since historical data indicates a strong unimodality property in hourly electricity prices, typically correlated with hourly variations in temperature and load profiles over the day, thus allowing for the electricity price forecast $\hat{\pi}(t)$ to be closely approximated by a unimodal function as illustrated in Fig. 1. We now propose a simple heuristic relaxation to the optimization problem $\mathcal{P}$, based on Assumption 2. If $\hat{\pi}(t)$ is unimodal, then, an explicit solution to (3) can be written down as follows. Intuitively, the optimal solution to (3) involves designing $u_i(t)$ such that the LSE purchases most of its power during the period when the price is low, and uses this energy to pre-cool homes to their lower comfort bound L_{i0}, allowing for the ACs to be switched off during the peak pricing period until the temperature reaches the upper comfort band U_{i0}. For this pre-

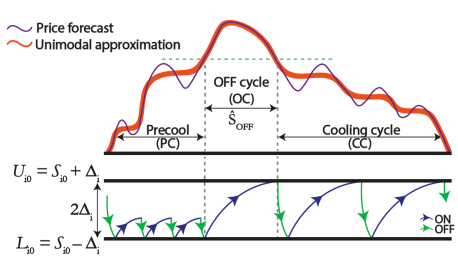

Fig. 1. Schematic of the optimization and control framework, indicating periods of pre-cooling (PC), OFF time (OC), and normal cyclical cooling operation (CC).

cooling operation, we consider the monotonically increasing portion of the the unimodal price function, that is $\hat{\pi}(t)$ such that $t \in [0, t_{PEAK}]$. Additionally, we relax the energy budget constraint by assuming $E = \infty$ (An explicit model-based solution to (3) incorporating this constraint and the switching dynamics of the loads can be provided along the lines of [23]). We have the following result on the solution to the optimal control problem $\mathcal{P}$ for the period where the price is monotonically increasing.

Theorem 1: If $\hat{\pi}(t),\ t \in [0, t_{PEAK}]$ is monotonically increasing, then there exists $t^* < t_{PEAK}$, such that the optimal solution to (3) is given by

$$u_i^*(t) = \begin{cases} 1, & t < t^*, \quad \theta_i \in (L_{i0}, U_{i0}] \\ u_i^{DOWN}(t), & t < t^*, \quad \theta_i = L_{i0} \\ u_i^{UP}(t), & t \geq t^*, \quad \theta_i = U_{i0} \\ 0, & t \geq t^*, \quad \theta_i \in [L_{i0}, U_{i0}), \end{cases} \tag{5}$$

where u_i^{UP} and u_i^{DOWN} are as defined in Assumption 1.

In order to apply the result of Theorem 1 to solving (3) with a unimodal price forecast $\hat{\pi}(t), t \in [0, t_{PEAK}]$ satisfying Assumption 2, it is first necessary to determine the pre-cooling period, denoted by PC $= [0, t^*]$ as shown in Fig. 1 such that $u_i(t^*) = u_i^{DOWN}(t^*)$ and $\theta_a(t^*) = L_{i0}$. We begin by noting that we would like to maintain $u_i = 0$ for as long as possible around the peak pricing period, without violating consumers' comfort bounds. We denote this period where $u_i = 0$ as the OFF cycle (OC) with duration $\hat{S}_{OFF}$. The longest period for which the OFF cycle can be maintained is the average duty cycle T_{OFF} as defined in Remark 1, that is $\hat{S}_{OFF} = T_{OFF}$. Working backwards, we can approximate $t^* \approx t_{PEAK} - T_{OFF}/2$. We then have the following result on the solution to the optimization problem (3) during the pre-cooling period and the OFF-cycle.

Corollary 1: If $\hat{\pi}(t)$ is monotonically increasing for $t \in [0, t_{PEAK}]$ and monotonically decreasing for $t \in [t_{PEAK}, T]$ then the solution to (3) for $t \in [0, t_{PEAK} + T_{OFF}/2]$ is given by (5) with $t^* \approx t_{PEAK} - T_{OFF}/2$.

After the OFF cycle, the price $\hat{\pi}(t), t \in [t_{PEAK} + T_{OFF}/2, T]$ is assumed to be monotonically decreasing according to Assumption 2. During this period, two types of control actions are possible as follows:

- Option 1: Maintain $\theta_i(t) = U_i(t)$ for $t \in [t_{PEAK} + T_{OFF}/2, T]$, or
- Option 2 (cooling cycle or CC): Allow the collection of ACs to evolve according to their natural dynamics (1) with control action (2).

In our approach, we choose the latter, namely Option 2, for two reasons. First, Option 2 allows for greater comfort for residential consumers by maintaining the average temperature of the home closer to the setpoint of the consumer's choice. Second, since the ambient temperatures during this period are typically cooler, it may not be optimal to maintain the temperature at the upper comfort bound U_{i0}.

In summary, we solve the optimization problem $\mathcal{P}$ by dividing the day into three time horizons, namely, pre-cooling, OFF cycle and cooling cycle, for which the control actions $u_i^*(t)$ are determined by Corollary 1.

Remark 2: We make the following remarks about the proposed solution to the optimization problem $\mathcal{P}$.

- We note that T_{OFF} for a given ambient temperature profile can be easily inferred by observing the total load profile over a day, without any direct knowledge of the dynamics of the homes. Therefore, the solution to (3) can be constructed in a completely model-free manner.
- This solution relies on an accurate forecast of t_{PEAK}, which is obtained using the LSTM network described in Section III-A.1. Note that the actual magnitude of the peak price is not important to our approach. Therefore, while prediction errors in the price magnitude can be tolerated, it is critical that the LSTM network be tuned such that the time of peak pricing is predicted as accurately as possible.

B. Stage 2: Private Control Implementation

We now describe how the solution of the optimization problem $\mathcal{P}$ as discussed in Section III-A.2 and depicted in Fig. 1 can be implemented in a private and distributed manner at each home, without any measurement of the state (temperature and power consumption) of the home by the LSE. At any time $t \in [0, T]$, the LSE broadcasts one of the following commands to the ACs:

$$c(t) = \begin{cases} 1, & t \in [0, t_{PEAK} - T_{OFF}/2] \\ 0, & t \in [t_{PEAK} - T_{OFF}/2, t_{PEAK} + T_{OFF}/2] \\ \text{CC}, & t \in [t_{PEAK} + T_{OFF}/2, T] \end{cases} \tag{6}$$

The ACs then translate these commands into their private switching state $\sigma_i(t)$, $i \in \{1, 2, \ldots, N\}$, at each time $t \in [0, T]$ as follows:

$$\sigma_i(t) = \begin{cases} 1, & c(t) = 1, & \theta_i(t) \in (L_{i0}, U_{i0}] \\ u_i^{DOWN}(t), & c(t) = 1, & \theta_i(t) = L_{i0} \\ 0, & c(t) = 0, & \theta_i(t) \in [L_{i0}, U_{i0}) \\ u_i^{UP}(t), & c(t) = 0, & \theta_i(t) = U_{i0} \\ \sigma(t^-), & c(t) = \text{CC}, & \theta_i(t) \in (L_{i0}, U_{i0}) \\ 1, & c(t) = \text{CC}, & \theta_i(t) = U_{i0} \\ 0, & c(t) = \text{CC}, & \theta_i(t) = L_{i0} \end{cases} \tag{7}$$

In contrast to the PID-based differentially private control implementation described in Section II-C, the control actions (7) can be implemented in an extremely simple manner without any measurements being transmitted to the LSE. The only requirement is that the homes be equipped with a smart thermostat that can receive instructions broadcast by the LSE. We note that the control inputs $u_i^{DOWN}(t)$ and $u_i^{UP}(t)$ in (7) to maintain a particular temperature $\theta_i(t)$ once the home has cooled to its setpoint are also already present as an energy saving measure in most ACs, where they are implemented by turning off the compressor of the AC and do not require knowledge of the thermal parameters of the home.

IV. CASE STUDY

In this section, we demonstrate the application of the proposed optimization and control framework in Section III on a test scenario in the Houston area, and benchmark it against the model based solution described in Section II-C. We consider $N = 500$ ACs with thermal power $P = 14\text{kW}$, and efficiency $\eta = 2.5$, with thermal parameters α_i and β_i, $i = 1, 2, \ldots, N$ drawn from the truncated Gaussians $\alpha \sim \mathcal{N}(\mu_\alpha, 0.1\mu_\alpha, 0.9\mu_\alpha, 1.1\mu_\alpha)$ and $\beta \sim \mathcal{N}(\mu_\beta, 0.1\mu_\beta, 0.9\mu_\beta, 1.1\mu_\beta)$ respectively, where $\mu_\alpha = \frac{1}{RC}\text{h}^{-1}$, $\mu_\beta = \frac{1}{C}{}^\circ C/kWh$, and $R = 2°\text{C/kW}$ and $C = 10$ $\text{kWh}/°\text{C}$ represent the thermal resistance and capacitance of the ACs respectively. We assume that the comfort bands of the ACs Δ_i are uniformly distributed in the range $[1, 3]°C$.

As described in Section III-A.1, we begin by using an LSTM to forecast the hourly price given historical price data and the ambient temperature profile as shown in Fig. 2-Top. To obtain this forecast, we consider an input dataset comprising of (i) real-time electricity price data for Houston, Texas (LZ-HOUSTON node) at 15-min intervals over a period of 7 years ranging from 2013-2019, available from the Electric Reliability Council of Texas (ERCOT) at `http://www.ercot.com/mktinfo/prices`, and (ii) hourly historical weather data, available from the National Centers for Environmental Information at `https://www.ncdc.noaa.gov/cdo-web/datatools`. We begin by averaging the 15-min prices from the ERCOT dataset to obtain the average hourly historical prices. After suitably scaling the temperature and hourly price datasets, we separate them into training and test data sets, where the training data set comprises of all price and temperature information for the years 2013-2017, and the test data set comprises of the same information for the years 2018-2019. We then implement an LSTM network comprised of one hidden layer with 5 LSTM neurons using Keras (`https://keras.io`). Based on the considerations described in Remark 1, we choose a forecast window of 3 hours. The network was found to converge in 10 epochs, with a mean absolute error (MAE) of 4.06%.

We then compare the two following approaches:

- **Private model-free control scheme:** We compute the solution to the optimal control problem (3) using the approach in Section III-A.2, compute the control commands broadcast by the LSE according to (6) and implement the corresponding switching actions (7).
- **Model-based control scheme:** We compute the solution to the optimal control problem (3) by an LP relaxation as described in Section II-C, compute the velocity control commands broadcast by the LSE using a PID controller with gains $k_p = 10^{-4}$, $k_i = 10^{-6}$ and $k_d = 10^{-4}$, and determine the control actions of the individual homes according to (4).

We simulate the response of the homes to each of these control schemes by solving (1) with switching action (7) over a horizon of $T = 24\text{h}$ by discretization using the Euler method with a step size of 1 second, and compute the total power consumption at each time step. Fig. 2-Bottom shows the temperature profiles of the homes with the model-free control scheme, clearly satisfying consumer comfort constraints. We observe that the ACs are pre-cooled from $t \in [0, 16]\text{h}$ and are turned off during the period of peak

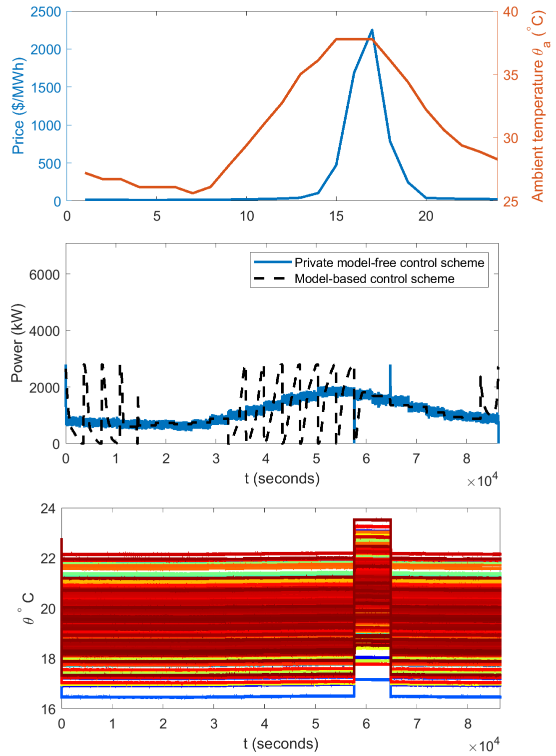

Fig. 2. Top: Hourly price and ambient temperature, Middle: Comparison: Power consumption of the proposed model-free framework vs model-based solution in Section II-C, Bottom: Temperature profiles of ACs.

pricing between $t \in [16, 17.5]$h. It can be verified that this OFF cycle aligns with the average duty cycle of the ACs computed from α_i and β_i, $\{1, 2, \ldots, N\}$.

The total power consumption under the model-free and model-based control schemes are compared in Fig. 2-Middle. We observe that the power consumption of the model-free control scheme approximately tracks the mean of the power consumption trajectory generated by the model-based scheme. The average energy consumption E_{avg} and energy cost savings E_s over the day for each control scheme are found to be as follows:

$$\begin{aligned} &\text{Uncontrolled:} && E_{avg} = 25.68\text{MWh}, && E_s = 0 \\ &\text{Model-based:} && E_{avg} = 22.4\text{MWh}, && E_s = \$3787 \\ &\text{Model-free:} && E_{avg} = 23.0\text{MWh}, && E_s = \$3597. \end{aligned}$$

Strikingly, the proposed model-free approach has almost no loss of performance as compared to the complex model-based scheme, indicating its potential.

V. Conclusion and Future Work

In this paper, we proposed a model-free framework to minimize the cost of procuring electricity for a collection of residential thermal loads by pre-cooling them to avoid purchasing power during peak pricing periods. The proposed approach is privacy-preserving in the sense that it does not require knowledge of the thermal dynamics or measurement of the states of the individual homes. Future work will involve improving the forecast of the time at which the peak price occurs in an online manner to dynamically shape the duration and frequency of the pre-cooling cycles.

References

[1] D. S. Kirschen, "Demand-side view of electricity markets," *IEEE Transactions on power systems*, vol. 18, no. 2, pp. 520–527, 2003.

[2] F. C. Schweppe, R. D. Tabors, J. L. Kirtley, H. R. Outhred, F. H. Pickel, and A. J. Cox, "Homeostatic utility control," *IEEE Transactions on Power Apparatus and Systems*, no. 3, pp. 1151–1163, 1980.

[3] G. Strbac, "Demand side management: Benefits and challenges," *Energy policy*, vol. 36, no. 12, pp. 4419–4426, 2008.

[4] A. M. Carreiro, H. M. Jorge, and C. H. Antunes, "Energy management systems aggregators: A literature survey," *Renewable and Sustainable Energy Reviews*, vol. 73, pp. 1160–1172, 2017.

[5] U. E. I. Administration. (2015) Residential energy consumption survey. [Online]. Available: http://www.eia.gov/consumption/residential/index.cfm

[6] D. S. Callaway and I. A. Hiskens, "Achieving controllability of electric loads," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 184–199, 2010.

[7] W. Turner, I. Walker, and J. Roux, "Peak load reductions: Electric load shifting with mechanical pre-cooling of residential buildings with low thermal mass," *Energy*, vol. 82, pp. 1057–1067, 2015.

[8] J. Mathieu, M. Dyson, D. Callaway, and A. Rosenfeld, "Using residential electric loads for fast demand response: The potential resource and revenues, the costs, and policy recommendations," in *ACEEE Summer Study on Energy Efficiency in Buildings*. Citeseer, 2012, pp. 189–203.

[9] G. Heffner, "Loads providing ancillary services: Review of international experience," 2008.

[10] H. Hao, Y. Lin, A. S. Kowli, P. Barooah, and S. Meyn, "Ancillary service to the grid through control of fans in commercial building hvac systems," *IEEE Transactions on smart grid*, vol. 5, no. 4, pp. 2066–2074, 2014.

[11] D.-c. Wei and N. Chen, "Air conditioner direct load control by multi-pass dynamic programming," *IEEE Transactions on Power Systems*, vol. 10, no. 1, pp. 307–313, 1995.

[12] C. Kurucz, D. Brandt, and S. Sim, "A linear programming model for reducing system peak through customer load control programs," *IEEE Transactions on Power Systems*, vol. 11, no. 4, pp. 1817–1824, 1996.

[13] T.-F. Lee, M.-Y. Cho, Y.-C. Hsiao, P.-J. Chao, and F.-M. Fang, "Optimization and implementation of a load control scheduler using relaxed dynamic programming for large air conditioner loads," *IEEE Transactions on Power Systems*, vol. 23, no. 2, pp. 691–702, 2008.

[14] J. L. Mathieu, S. Koch, and D. S. Callaway, "State estimation and control of electric loads to manage real-time energy imbalance," *IEEE Transactions on Power Systems*, vol. 28, no. 1, pp. 430–440, 2012.

[15] W. Zhang, J. Lian, C.-Y. Chang, and K. Kalsi, "Aggregated modeling and control of air conditioning loads for demand response," *IEEE transactions on power systems*, vol. 28, no. 4, pp. 4655–4664, 2013.

[16] D. Docimo and H. K. Fathy, "Demand response using heterogeneous thermostatically controlled loads: Characterization of aggregate power dynamics," *Journal of Dynamic Systems, Measurement, and Control*, vol. 139, no. 10, 2017.

[17] F. Ruelens, B. J. Claessens, S. Vandael, B. De Schutter, R. Babuška, and R. Belmans, "Residential demand response of thermostatically controlled loads using batch reinforcement learning," *IEEE Transactions on Smart Grid*, vol. 8, no. 5, pp. 2149–2159, 2016.

[18] B. J. Claessens, D. Vanhoudt, J. Desmedt, and F. Ruelens, "Model-free control of thermostatically controlled loads connected to a district heating network," *Energy and Buildings*, vol. 159, pp. 1–10, 2018.

[19] A. Halder, X. Geng, P. Kumar, and L. Xie, "Architecture and algorithms for privacy preserving thermal inertial load management by a load serving entity," *IEEE Transactions on Power Systems*, vol. 32, no. 4, pp. 3275–3286, 2016.

[20] A. Halder, X. Geng, F. A. Fontes, P. Kumar, and L. Xie, "Optimal power consumption for demand response of thermostatically controlled loads," *Optimal Control Applications and Methods*, vol. 40, no. 1, pp. 68–84, 2019.

[21] E. Liu, P. You, and P. Cheng, "Optimal privacy-preserving load scheduling in smart grid," in *2016 IEEE Power and Energy Society General Meeting (PESGM)*. IEEE, 2016, pp. 1–5.

[22] S. Bhattacharya, K. Kar, and J. H. Chow, "Economic operation of thermostatic loads under time varying prices: An optimal control approach," *IEEE Transactions on Sustainable Energy*, vol. 10, no. 4, pp. 1960–1970, 2018.

[23] F. A. Fontes, A. Halder, J. Becerril, and P. Kumar, "Optimal control of thermostatic loads for planning aggregate consumption: Characterization of solution and explicit strategies," *IEEE Control Systems Letters*, vol. 3, no. 4, pp. 877–882, 2019.