

Lawrence Berkeley National Laboratory

LBL Publications

Title

COHORT: Coordination of Heterogeneous Thermostatically Controlled Loads for Demand Flexibility

Permalink

<https://escholarship.org/uc/item/2tk6997t>

ISBN

9781450380614

Authors

Chen, B
Francis, J
Pritoni, M
et al.

Publication Date

2020-11-18

DOI

10.1145/3408308.3427980

Peer reviewed

COHORT: Coordination of Heterogeneous Thermostatically Controlled Loads for Demand Flexibility

Bingqing Chen
Carnegie Mellon University
Pittsburgh, PA, USA
bingqinc@andrew.cmu.edu

Jonathan Francis
Carnegie Mellon University
Pittsburgh, PA, USA
jmf1@cs.cmu.edu

Marco Pritoni
Lawrence Berkeley National Lab
Berkeley, CA, USA
mpritoni@lbl.gov

Soumya Kar
Carnegie Mellon University
Pittsburgh, PA, USA
soumyyak@andrew.cmu.edu

Mario Bergés
Carnegie Mellon University
Pittsburgh, PA, USA
mberges@andrew.cmu.edu

ABSTRACT

Demand flexibility is increasingly important for power grids. Careful coordination of thermostatically controlled loads (TCLs) can modulate energy demand, decrease operating costs, and increase grid resiliency. We propose a novel distributed control framework for the Coordination Of Heterogeneous Residential Thermostatically controlled loads (COHORT). COHORT is a practical, scalable, and versatile solution that coordinates a population of TCLs to jointly optimize a grid-level objective, while satisfying each TCL's end-use requirements and operational constraints. To achieve that, we decompose the grid-scale problem into subproblems and coordinate their solutions to find the global optimum using the alternating direction method of multipliers (ADMM). The TCLs' local problems are distributed to and computed in parallel at each TCL, making COHORT highly scalable and privacy-preserving. While each TCL poses combinatorial and non-convex constraints, we characterize these constraints as a convex set through relaxation, thereby making COHORT computationally viable over long planning horizons. After coordination, each TCL is responsible for its own control and tracks the agreed-upon power trajectory with its preferred strategy. In this work, we translate continuous power back to discrete *on/off* actuation, using pulse width modulation. COHORT is generalizable to a wide range of grid objectives, which we demonstrate through three distinct use cases: generation following, minimizing ramping, and peak load curtailment. In a notable experiment, we validated our approach through a hardware-in-the-loop simulation, including a real-world air conditioner (AC) controlled via a smart thermostat, and simulated instances of ACs modeled after real-world data traces. During the 15-day experimental period, COHORT reduced daily peak loads by an average of 12.5% and maintained comfortable temperatures.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

BuildSys '20, November 16-19, 2020, Yokohama, Japan

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8061-4/20/11.

<https://doi.org/10.1145/3408308.3427980>

CCS CONCEPTS

• **Hardware** → **Smart grid**; *Energy generation and storage*; • **Computing methodologies** → **Distributed algorithms**; Control methods; Planning and scheduling.

KEYWORDS

distributed control; thermostatically controlled loads

ACM Reference Format:

Bingqing Chen, Jonathan Francis, Marco Pritoni, Soumya Kar, and Mario Bergés. 2020. COHORT: Coordination of Heterogeneous Thermostatically Controlled Loads for Demand Flexibility. In *The 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys '20)*, November 16-19, 2020, Yokohama, Japan. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3408308.3427980>

1 INTRODUCTION

Growing peak demand in some regions and increasing penetration of renewable generation in others are presenting challenges for grid operators to balance supply and demand [13, 20]. Traditionally, demand-side load is viewed as uncontrollable, while supply-side resources manage power generation to match it. An emerging paradigm is to tap into the flexibility of demand-side resources to reduce, shift, or modulate their loads in response to price or control signals [26]. Such demand flexibility can be utilized to provide grid services, improve grid resiliency [26], and reduce operating costs [4].

In this work, we focus on load control¹[4] of residential thermostatically controlled loads (TCLs), such as air conditioners (ACs), refrigerators, and electric water heaters, which account for about 20% of the electricity consumption in the United States (US) [16]. Due to their inherent flexibility through thermal inertia, they can provide grid services without compromising their end uses. However, there are two challenges to utilizing TCL flexibility. Firstly, TCLs must be aggregated across a population to be a meaningful resource at the grid-level [26], which results in a control problem with a large state-action space. Secondly, the constraints posed by each TCL are combinatorial and thus non-convex [2], due to the fact that a TCL operates in discrete action space, i.e., *on* or *off*.

We present a novel framework (Figure 1) for the Coordination Of Heterogeneous Residential Thermostatically controlled loads

¹As opposed to methods based on economic incentives. We refer interested readers to [4] for a comparison of load control vs. price-based methods.

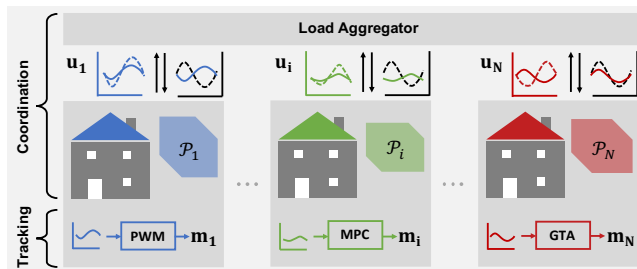


Figure 1: Framework. The load aggregator coordinates a population of TCLs to jointly optimize a grid-level objective, while satisfying each TCL’s constraints, characterized by the set \mathcal{P}_i . The load aggregator and each TCL coordinates at the level of its power trajectory, u_i , until a consensus is reached among the population. Each TCL is responsible for its own control and tracks u_i locally with its preferred strategy.

(COHORT) to jointly optimize a grid-level objective, while satisfying each TCL’s end-use requirements and operational constraints. To effectively handle the large state-action space, we adopt a distributed control architecture, where each TCL is responsible for its own control and coordinates with others to find a grid-level solution. Similar to [2], we decompose the grid-level problem into smaller subproblems and coordinate their solutions using the alternating direction method of multipliers (ADMM) [1]. The advantages of the distributed architecture compared with centralized and decentralized approaches are elaborated in Section 2. To address the second challenge, we characterize each TCL’s flexibility, i.e., the set of all admissible power profiles [36], as a convex set through relaxation. As a result, COHORT is computationally viable for tasks with long planning horizons (e.g., 24 hours), thereby addressing a limitation of [2]—namely that its computational cost grows exponentially with the planning horizon. After coordination, we use low-frequency pulse width modulation (PWM), inspired by [3], to translate the solution of the convex-relaxed problem back to *on/off* activation. Since the coordination process makes no assumption on each TCL’s control scheme, the TCL may opt for alternative strategies with reference-tracking capability, such as model predictive control (MPC) and global thermostat adjustment (GTA).

COHORT can incorporate detailed, system-specific dynamics and constraints of individual TCLs. At the same time, its computational cost scales well with both population size and planning horizon. As a result, COHORT is generalizable to a wider variety of grid objectives, compared to existing methods, which we demonstrate through three distinct use cases: generation following, minimizing ramping, and peak load curtailment. As a proof of concept, we evaluated COHORT in simulation, targeting challenges arising from increasing penetration of renewable generation [20] under, generation following and minimizing ramping, based on load profiles from California Independent System Operator (CAISO) [19].

Then, we validated that COHORT is practical for real-world TCLs. To do that, we developed a hardware-in-the-loop (HIL) simulation, including a real-world residential AC controlled via a smart

thermostat, and simulated instances of ACs modeled after real-world data traces. In a 15-day experimental period, COHORT reduced daily peak loads by 12.5% on average based on load profiles from Pennsylvania-New Jersey-Maryland (PJM) Interconnection [29], while maintaining comfort in the real-world testbed.

The rest of the paper is organized as follows. Section 2 reviews prior work on TCL control. Section 3 provides technical background, and Section 4 formally introduces COHORT. We then describe experiments and results both in the simulation study (Section 5) and in the HIL simulation (Section 6), and present conclusions and possible future work in Section 7.

2 RELATED WORK

We review existing work on TCL control, and include work on other flexible loads if the methodology is relevant.

Architectures for TCL control. The primary challenge for jointly controlling a large number of TCLs is the large state-action space. To address this challenge, a popular approach is to develop an aggregate model for the population and control the population in a centralized manner. Examples of such aggregate model include the state bin transition model [23, 35] and the virtual battery model [16, 36]. However, these aggregate models depend on the assumptions that each system may be characterized by a 1st- (or 2nd- [35]) order linear model, and that all systems in the population share the same model structure and control scheme. These aggregate models have low fidelity and do not capture system-specific dynamics. Specifically, aggregate TCL modeling is ill-suited for predicting long-term responses—a pre-requisite for tasks with long planning horizons, such as load shifting [2]. Alternatively, one can jointly control building loads as a centralized MPC problem [33], but, while this approach allows for incorporation of detailed building models and system-specific constraints, it is computationally expensive and also raises privacy concerns, as it requires each building to share an excessive amount of information with the load aggregator, including: thermal models, system specifications, control logic, and occupants’ usage pattern and comfort preferences.

Aside from the centralized architecture, decentralized control [32] and distributed control [2, 8, 24] approaches have also been proposed in the literature. Taking advantage of the fact that system frequency is a universally-available indicator for supply-demand imbalance, [32] determines the action of each TCL with a power response model based on locally-available information. The key advantage of a decentralized control approach is that each system can be controlled based on local information, without any communication. However, this characteristic also constrains the applications of decentralized control to frequency or voltage regulation and real-time load shaping [2].

In a distributed architecture, which we adopt in this work, each system is responsible for its own control, and it coordinates with others to jointly achieve a grid-level objective. In [24], the distributed MPC scheme allocated the aggregate load to TCL clusters following a time-invariant weight. This allocation scheme does not account for the fact that the flexibility available at each building is time-varying, and thus does not fully utilize the aggregate flexibility [33]. [5] used a distributed MPC approach similar to [24], but learned the allocation scheme with an evolutionary strategy. Most

similar to this work is [2], which also used ADMM for distributed optimization. A major advantage is that the computation is distributed to and parallelized at each TCL. A significant limitation of [2] the computational cost grows exponentially with the planning horizon. Specifically, it represented each TCL as a set of feasible state-action trajectories, the size of which is N_a^T , where N_a is the number of alternative actions considered at each time-step, and T is the number of time-steps in the planning horizon. Given that the trajectories depends on initial state and future disturbances, the trajectories need to be unrolled at each time-step. Another limitation is that the solution is only optimal with respect to the set of trajectories under consideration.

Experimental Validation. The majority of works on this topic validated their approaches in simulation. In particular, works such as [2, 16, 23, 24, 32, 36] validated their approach on population of TCLs simulated with 1st-order linear thermal model, using model parameters sampled from assumed distributions. It is unclear how well such validation reflects performance on real-world systems. [34] demonstrated 1st- and 2nd-order models failed to accurately capture the thermodynamics of an individual electric water heater. Furthermore, there is a large difference between performance reported in simulation and in real-world testbeds. For instance, [16] reported a maximum absolute percentage error of less than 1% tracking a 4s frequency regulation signal in a simulation study. In comparison, in a real-world experiment on 300 residential ACs, [25] reported a median absolute percentage error of 6.7% executing 1-hour demand response (DR) events, which is arguably a much simpler task. Such discrepancy calls for more realistic evaluation. Other attempts at realistic evaluation include [6] which developed linear models based on configurations of real households, and [27], which used a co-simulation environment with EnergyPlus models.

Optimization Objectives. A myriad of grid-level objectives have been discussed in the literature, such as: energy cost minimization [6], DR events [25, 27, 32, 35], frequency regulation [16, 36], generation following [2, 24], reference tracking [23], and peak load reduction [8]. However, these works generally formulate their approaches based on their specific use case, without discussing their generalizability to other applications.

3 BACKGROUND

We now present background technical concepts used by COHORT, including TCL modeling (Section 3.1) and ADMM (Section 3.2).

3.1 TCL Model and Flexibility

Here, we introduce the modeling of an individual TCL and define its flexibility. The contents is largely inspired by [36], from which we made modifications based on our problem.

3.1.1 System Dynamics. The temperature dynamics of an individual TCL is commonly modeled with Eq. 1a [16, 23, 36], where T_t is the TCL temperature, $T_{a,t}$ is the ambient temperature, and $m_t \in \{0, 1\}$ is the binary control variable representing the operating state, i.e., *on* or *off*, at time t . The negative sign associated with the control action assumes the TCL is operating in cooling mode, which could be changed to a positive sign to reflect heating. P_m is the rated power. Denoting the thermal resistance and capacitance

of the TCL as R and C respectively, the model parameters can be calculated as: $a = \exp\{-\Delta T/(RC)\}$ and $b_t = \eta_t R$, where ΔT is the time-step and η_t is the time-varying coefficient of performance (COP). While the thermodynamics is linear, it is difficult to analyze the system dynamics in Eq. 1a due to the discrete control variable m_t . A common approach is to apply convex relaxation to the discrete control variable, which results in a linear system approximation (Eq. 1b) [16, 23, 36]. The new control variable $u_t \in [0, P_m]$, i.e., power consumption of the TCL, is continuous instead of discrete. This approach is justified as the aggregate behavior of a TCL population can be approximated accurately by Eq. 1b [36].

$$T_{t+1} = aT_t + (1-a)(T_{a,t} - b_t m_t P_m) \quad (1a)$$

$$T_{t+1} = aT_t + (1-a)(T_{a,t} - b_t u_t) \quad (1b)$$

The TCL dynamics over a planning horizon, $t : t + T - 1$, is thus characterized by Eq. 2a (or more concisely Eq. 2b), where $\mathbf{B}_u = \text{diag}(-(1-a)b_t, \dots, -(1-a)b_{t+T-1})$. Throughout this work, boldface lower-case letters, e.g., \mathbf{x} , are vectors, and boldface upper-case letters, e.g., \mathbf{A} , are matrices; \mathbf{x} , \mathbf{x}_0 , $\mathbf{u} \in \mathbb{R}^T$ and \mathbf{A} , $\mathbf{B}_u \in \mathbb{R}^{T \times T}$. Denoting the number of disturbance terms as l , we have $\mathbf{D} \in \mathbb{R}^{T \times l}$ and $\mathbf{b}_d \in \mathbb{R}^l$. In this case, the disturbance term only includes the ambient temperature, $l = 1$.

$$\underbrace{\begin{bmatrix} 1 & & & & \\ -a & 1 & & & \\ & & \ddots & & \\ & & & -a & 1 \end{bmatrix}}_{\mathbf{A}} \underbrace{\begin{bmatrix} T_{t+1} \\ T_{t+2} \\ \vdots \\ T_{t+T} \end{bmatrix}}_{\mathbf{x}} = \underbrace{\begin{bmatrix} aT_t \\ 0 \\ \vdots \\ 0 \end{bmatrix}}_{\mathbf{x}_0} + \underbrace{\mathbf{B}_u}_{\mathbf{B}_u} \underbrace{\begin{bmatrix} u_t \\ u_{t+1} \\ \vdots \\ u_{t+T-1} \end{bmatrix}}_{\mathbf{u}} + \underbrace{\begin{bmatrix} T_{a,t} \\ T_{a,t+1} \\ \vdots \\ T_{a,t+T-1} \end{bmatrix}}_{\mathbf{D}} \underbrace{\begin{bmatrix} 1-a \\ \vdots \\ 1-a \end{bmatrix}}_{\mathbf{b}_d} \quad (2a)$$

$$\mathbf{A}\mathbf{x} = \mathbf{x}_0 + \mathbf{B}_u\mathbf{u} + \mathbf{D}\mathbf{b}_d \quad (2b)$$

Constraints. Each TCL needs to satisfy the end-use requirements and respect the operational constraints. In this case, we require the TCL temperature to be within the deadband, i.e., $T_t \in [T_{sp} - \Delta, T_{sp} + \Delta]$, where T_{sp} is the setpoint and Δ is half of the deadband. At the same time, the system needs to be operating within its power limits, i.e., $P_t \in [0, P_m]$. Combining the system dynamics given in Eq. 2b, the aforementioned constraints can be written as Eq. 3, where $\underline{\mathbf{u}} = [0]$, $\bar{\mathbf{u}} = [P_m]$, $\underline{\mathbf{x}} = [T_{sp} - \Delta]$, and $\bar{\mathbf{x}} = [T_{sp} + \Delta]$.

$$\underline{\mathbf{u}} \leq \mathbf{u} \leq \bar{\mathbf{u}}; \quad \underline{\mathbf{x}} \leq \mathbf{A}^{-1}(\mathbf{x}_0 + \mathbf{B}_u\mathbf{u} + \mathbf{D}\mathbf{b}_d) \leq \bar{\mathbf{x}}; \quad (3)$$

3.1.2 Flexibility. We adopt the definition of flexibility proposed in [36], where the flexibility of a system is the set of all admissible power profiles. Our formulation of Eq. 4 is generalized from that in [36], to incorporate non-linear systems. \mathcal{P} denotes the flexibility of the system, $\mathcal{T}(x_k, u_k)$ denotes the state transition function, and \underline{x}_k , \bar{x}_k , \underline{u}_k , and \bar{u}_k are the lower and upper bounds for state and action of the given system at time k . An important intuition is that the flexibility is coupled over time through the thermodynamics [36]:

$$\mathcal{P} = \left\{ \begin{bmatrix} u_{t:t+T-1} \end{bmatrix} \left| \begin{array}{l} x_{k+1} = \mathcal{T}(x_k, u_k); \\ \underline{u}_k \leq u_k \leq \bar{u}_k; \quad \forall k \in \{t, \dots, t+T-1\} \\ \underline{x}_{k+1} \leq x_{k+1} \leq \bar{x}_{k+1}; \end{array} \right. \right\} \quad (4)$$

For the specific case of TCL, which follows the linear dynamics in Eq. 1b, the flexibility \mathcal{P} of a TCL can be expressed as Eq. 5.

$$\mathcal{P} = \mathcal{U} \cap \mathcal{X} \quad (5)$$

where $\mathcal{U} = \{\mathbf{u} | \underline{\mathbf{u}} \leq \mathbf{u} \leq \bar{\mathbf{u}}\}$ and $\mathcal{X} = \{\mathbf{x} | \underline{\mathbf{x}} \leq \mathbf{A}^{-1}(\mathbf{x}_0 + \mathbf{B}_u \mathbf{u} + \mathbf{D} \mathbf{b}_d) \leq \bar{\mathbf{x}}\}$, as derived in Eq. 3. Note that \mathcal{P} boils down to a set of linear inequalities, which is geometrically interpreted as a polytope² [36].

3.2 ADMM

ADMM is a well-established distributed convex optimization algorithm, which decomposes a large problem into smaller subproblems and coordinates the solutions to find a global optimum [1]. Generally, ADMM solves problems in the form of Eq. 6.

$$\begin{aligned} \min_{u,v} \quad & f(u) + g(v) \\ \text{s.t.} \quad & Au + Bv = c \end{aligned} \quad (6)$$

Specifically, we introduce the application of ADMM to a canonical problem: the *sharing problem*, as given in Eq. 7, where f_i is a local objective for agent i , and g is the global objective—defined as a function of the aggregate of all decision variables from the agents:

$$\min_{u_i} \sum_i^N f_i(u_i) + g\left(\sum_i^N u_i\right) \quad (7)$$

By introducing a copy of the decision variable u_i as v_i , the sharing problem can be written in a ADMM-compatible form (Eq. 8):

$$\begin{aligned} \min_{u,v} \quad & \sum_i^N f_i(u_i) + g\left(\sum_i^N v_i\right) \\ \text{s.t.} \quad & u_i - v_i = 0, \quad i = 1, \dots, N \end{aligned} \quad (8)$$

The update rules for solving the problem are given in Eq. 9, where w is the dual variable, ρ is a hyperparameter, and the superscript (k) denotes the value of a variable at the k^{th} iteration. We elaborate on the intuition behind these update rules here. The sharing problem can be interpreted as the agents coordinating their decisions so as to strike a balance between the local and the global objectives. Hence, u_i and v_i are each agent's solutions to its local problem and the global objective, respectively. The dual variable w_i , as calculated in Eq. 9c, is the cumulative disagreement between u_i and v_i . Thus, w_i , which we also call the incentive variable, communicates how to adjust each agent's solutions such that they would agree, i.e., $u_i = v_i$. Thus, in the *u-update* step (Eq. 9a), each agent solves its local problem, while mindful of its solution to the global problem. Similarly, in the *v-update* step (Eq. 9b), the agents jointly optimize the global objective, while ensuring their decisions are close to those of their local problems.

$$u_i^{(k+1)} = \arg \min_{u_i} f_i(u_i) + \frac{\rho}{2} \|u_i - v_i^{(k)} + w_i^{(k)}\|_2^2 \quad (9a)$$

$$v_i^{(k+1)} = \arg \min_{v_i} g\left(\sum_i^N v_i\right) + \frac{N\rho}{2} \|u_i^{(k+1)} - v_i + w_i^{(k)}\|_2^2 \quad (9b)$$

$$w_i^{(k+1)} = w_i^{(k)} + u_i^{(k+1)} - v_i^{(k+1)} \quad (9c)$$

A downside of the updates rules given in Eq. 9 is that it requires a copy of the decision variable for each agent, i.e., v_i . Intuitively, the global objective only depends on the aggregate behavior of the population, and thus a more efficient algorithm (Eq. 10) is possible using the mean of the variables, denoted as \bar{u} , \bar{v} , and \bar{w} respectively.

²A polytope can be characterized as a set $\mathcal{S} = \{x \in \mathbb{R}^n | Ax \leq b\}$.

$$u_i^{(k+1)} = \arg \min_{u_i} f_i(u_i) + \frac{\rho}{2} \|u_i - u_i^{(k)} + \bar{u}^{(k)} - \bar{v}^{(k)} + \bar{w}^{(k)}\|_2^2 \quad (10a)$$

$$\bar{v}^{(k+1)} = \arg \min_{\bar{v}} g(N\bar{v}) + \frac{N\rho}{2} \|\bar{u}^{(k+1)} - \bar{v} + \bar{w}^{(k)}\|_2^2 \quad (10b)$$

$$\bar{w}^{(k+1)} = \bar{w}^{(k)} + \bar{u}^{(k+1)} - \bar{v}^{(k+1)} \quad (10c)$$

An equality such as $v_i^{(k)} = \bar{v}^{(k)} + (u_i^{(k)} + w_i^{(k)}) - (\bar{u}^{(k)} + \bar{w}^{(k)})$, to show Eq. 9 and Eq. 10 are equivalent, can be derived from the stationarity condition [2]. For more details on the algorithm, we refer interested readers to [1].

4 APPROACH

We first formulate the problem and elaborate on the optimization procedure, with focus on the coordination between the load aggregator and the TCLs (Section 4.1). We then describe a PWM-based strategy for TCL-level control (Section 4.2).

4.1 Problem Formulation and Optimization

The problem we address can be formulated as Eq. 11: we want to simultaneously optimize a grid-level objective $g(\cdot)$, which is a function of aggregate power consumption, and make sure the actions of each TCL are admissible based on operational constraints and end-use requirements, characterized by the set \mathcal{P}_i . By applying convex relaxation to discrete actions, and characterizing the flexibility as a convex set \mathcal{P}_i , our approach is computationally viable for tasks with long planning horizons. Recall the definition in Section 3.1, $\mathbf{u}_i \in \mathbb{R}^T$ is the power consumption of a TCL over a planning horizon. The subscript i denotes the i^{th} TCL. While we did not include any TCL-level objective other than satisfying its constraints, it is possible to incorporate such objectives.

$$\begin{aligned} \min_{\mathbf{u}_i} \quad & g\left(\sum_i^N \mathbf{u}_i\right) \\ \text{s.t.} \quad & \mathbf{u}_i \in \mathcal{P}_i, \quad \forall i \end{aligned} \quad (11)$$

The problem in Eq. 11 can be written in a ADMM-compatible form (Eq. 12) by (i) introducing a copy of the variable \mathbf{u}_i as \mathbf{v}_i , and (ii) representing the constraints set \mathcal{P}_i with the indicator function $\mathbb{I}_{\mathcal{P}_i}$. By definition, $\mathbb{I}_{\mathcal{P}_i}(\mathbf{u}_i) = 0$, if $\mathbf{u}_i \in \mathcal{P}_i$, else $\mathbb{I}_{\mathcal{P}_i}(\mathbf{u}_i) = \infty$ [31]. The large penalty for an inadmissible \mathbf{u}_i forces the solver to find \mathbf{u}_i that satisfies the constraints.

$$\begin{aligned} \min_{\mathbf{u}_i, \mathbf{v}_i} \quad & \sum_i \mathbb{I}_{\mathcal{P}_i}(\mathbf{u}_i) + g\left(\sum_i \mathbf{v}_i\right) \\ \text{s.t.} \quad & \mathbf{u}_i = \mathbf{v}_i \end{aligned} \quad (12)$$

Note that Eq. 12 now has the same form as Eq. 8, and thus may be solved with the update rules in Eq. 10. Figure 2 summarizes the coordination procedure between the load aggregator and each TCL. Firstly, each TCL updates its action, \mathbf{u}_i , locally. Given that $f_i = \mathbb{I}_{\mathcal{P}_i}$ and \mathcal{P}_i is a polytope for a TCL, the *u-update* step (Eq. 10a) may be implemented efficiently as a projection operation [31], also illustrated in Figure 2. To simplify notation, we denote $\mathbf{u}_i^+ = \mathbf{v}_i^{(k)} - \mathbf{w}_i^{(k)} = \mathbf{u}_i^{(k)} - \bar{\mathbf{u}}^{(k)} + \bar{\mathbf{v}}^{(k)} - \bar{\mathbf{w}}^{(k)}$; \mathbf{u}_i^+ can be interpreted as the desired power profile for TCL i at the end of the k^{th} iteration, and the projection $\text{Proj}_{\mathcal{P}_i}(\mathbf{u}_i^+)$ ensures that the coordinated power profile

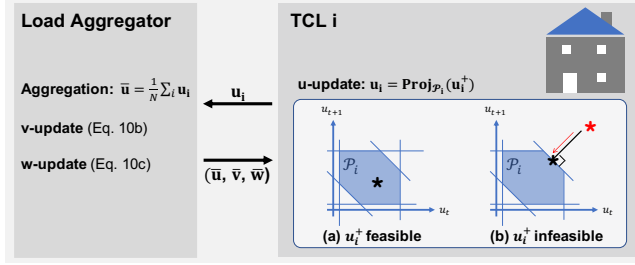


Figure 2: Coordination. The u -update step is distributed to and computed in parallel at each TCL as a projection operation. The load aggregator calculates \bar{u} , and then updates \bar{v} and \bar{w} . Finally, the load aggregator broadcasts the \bar{u} , \bar{v} , and \bar{w} to the population. This procedure repeats until convergence.

is admissible for the TCL. Secondly, the load aggregator collects the actions from all agents to find the mean, \bar{u} , and sequentially updates \bar{v} following Eq. 10b, and \bar{w} following Eq. 10c. Finally, the load aggregator broadcasts \bar{u} , \bar{v} , and \bar{w} to all TCLs, such that they can update u_i locally. This procedure repeats until convergence.

While we use ADMM in this work, we note that alternative distributed optimization algorithms exist in the literature, such as consensus-based approaches [18] that are especially relevant for distributed energy management. However, for the application in this paper, ADMM is suitable given the communication topology³ and provides scalability given the ease of solving each TCL’s local problem. There are three key advantages to the ADMM-based approach for this application. Firstly, the ADMM naturally decomposes the grid-level problem into subproblems. Thus, each TCL can ensure its local objective and constraints are satisfied, without sharing them with the load aggregator, thereby preserving privacy. Secondly, the u -update at each TCL is computed in parallel. Thus, the approach is highly scalable to large population. Finally, ADMM is guaranteed to converge to the grid-level optimum given a convex objective [1]. As will become clear through our demonstrations, a variety of grid objectives can be formulated as convex problems.

4.2 TCL-level Control

Recall that we applied convex relaxation to TCL dynamics, i.e., $m_t P_m \in \{0, P_m\}$ (Eq. 1a) to $u_t \in [0, P_m]$ (Eq. 1b), such that the grid-level objective can be optimized efficiently over long time horizons. In this section, we describe how to translate the continuous power trajectory back to on/off actuation with PWM, a method for generating quasi-continuous output from an on/off actuator. In [3], it was demonstrated that a TCL could be treated as a variable power unit via low-frequency PWM. The action normalized by rated power, $u_t/P_m \in [0, 1]$, can be interpreted as duty cycle ratio, i.e., the portion of time the TCL is *on* within the control time-step. Specifically, we implemented PWM with Sigma-Delta (Σ - Δ) modulation [28], where TCL switches between *on* and *off* when the cumulative error, $\epsilon_t = \sum_{k=0}^t (u_k/P_m - m_k)\Delta T$, exceeds the limits. Figure 3 illustrates its use in tracking a sine wave.

³i.e. there being a load aggregator that centrally collects and broadcasts information. In comparison, consensus-based methods would be more suitable for peer-to-peer communication topology.

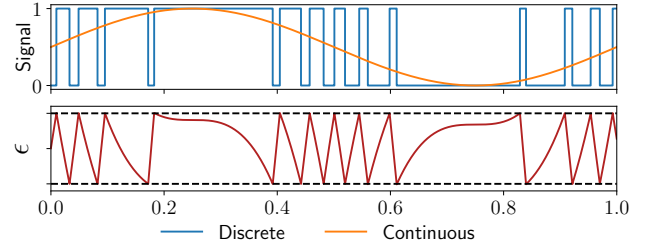


Figure 3: Sigma-Delta modulation converts a continuous signal to a discrete signal by switching states when the cumulative error, ϵ , exceeds the limits (dashed black lines).

It is clear from Figure 3 that fewer switchings are needed when a signal is close to either 0 or 1. We also observe that the solution tends to be close to a feasible initialization. The intuition may be that an individual TCL does not need to drastically change its default behavior to collaboratively achieve a grid-level objective. Given these observations, we initialize each u_i with a sequence of interlaced 0s and P_m s to encourage sparsity in the solution. Furthermore, by placing 0s and P_m s with care in the initialization, short-cycling can be reduced.

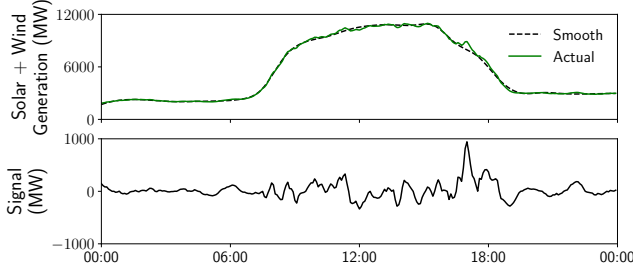
5 EXPERIMENT 1: SIMULATION STUDY

In this section, we evaluated COHORT in simulation as an initial proof of concept. Following [2, 16, 23, 24, 32, 36], we validated our approach on a population of TCLs simulated with 1st-order linear thermal models (Eq. 1a). We simulated 1000 TCLs, using parameters sampled from uniform distributions around nominal values [16, 23, 32, 36]. Specifically, we followed the same parameter distributions and values for temperature setpoint and exogenous variable as [16]. We used a deadband of $\Delta = 1^\circ\text{C}$ throughout this work. While these assumptions may not reflect realistic system dynamics and population heterogeneity, we adopted them such that the performance of our approach is directly comparable to those reported in the literature. We lifted these assumption and validated COHORT in a real-world testbed in Section 6.

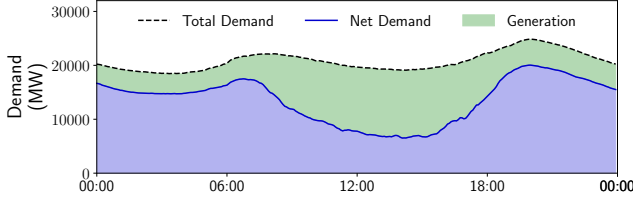
We applied COHORT to address challenges arising from increasing penetration of renewable generation. Firstly, we used the inherent flexibility in the TCL population to absorb the variations in renewable generation in a generation following use case (Section 5.1). Secondly, we shifted the TCL load to alleviate the need to quickly ramp up / down energy generation in areas of high renewable penetration (Section 5.2). The load curves used for both use cases (Figure 4) are from CAISO [19] on 31 March 2020.

5.1 Use Case 1: Generation Following

The generation following signal was produced following the same procedure in [2], as shown in Figure 4a. The TCL population tracked a scaled version of the generation following signal around its baseline power consumption. The objective function is the mean squared error (MSE) between the reference signal, denoted by \bar{u} , and actual aggregate energy consumption (Eq. 13). Other tasks such as DR events [32, 35] and frequency regulation [16, 36] boil down to



(a) Use Case 1: The 5-min generation following signal is produced by taking the difference between the actual renewable generation and its smooth spline fit, following [2].



(b) Use Case 2: The duck curve (i.e. the net demand) exemplifies the need for generators to quickly ramp up energy production when the sun sets in areas of high renewable penetration [20].

Figure 4: Load Profiles from CAISO, 2020/03/31

tracking a given reference signal, and thus may be addressed with the same objective function.

All the optimization problems in this work⁴ were solved using CVXPY [9] with hyperparameter $\rho = 10$. In this use case, we used a 5-min control time-step, and planned for the next time-step, i.e., $T = 1$. For tracking with PWM, it is necessary to use a smaller time-step. Throughout this work, we used a tracking time-step that is 1/15 of the control time-step. The error limit used for Σ - Δ modulation is 0.1kWh in the simulation study.

$$g_{\text{track}}\left(\sum_i \mathbf{u}_i\right) = \frac{1}{T} \|\bar{\mathbf{u}} - \sum_i \mathbf{u}_i\|_2^2 \quad (13)$$

The behavior of the TCL population is shown in Figure 5. COHORT tracked the reference signal with a small error, while maintaining the temperature of the population within the deadband (dashed green line). Note that the discrepancy between the reference signal and actual power consumption came solely from discretization error. The performance of our approach with comparison to seminal works on TCL control is summarized in Table 1. Our tracking performance is comparable to that in [23], reported in normalized⁵ root MSE (RMSE), and is not as good as [16], reported in mean absolute percentage error (MAPE). While [16] may be more suited for reference tracking tasks, it is not applicable to any planning-based task, e.g. load shifting. Compared to the baseline scenario where the TCL population only maintains temperature, our approach increased the switching frequency by 158.7%, which is

⁴The code is available at <https://github.com/INFERLab/COHORT>.

⁵by average aggregate power

Table 1: Performance Comparison for Reference Tracking

	Tracking		Switching
	Norm. RMSE (%)	MAPE (%)	Increase (%)
[23]	0.8-2.27	-	170-300
[16]	-	<1	116.7
COHORT	2.04	1.53	158.7

similar to [16, 23]. By adjusting the error limits in Σ - Δ modulation, one can trade-off tracking performance and switching frequency.

Figure 6 show actions of individual TCLs. We initialized each action with either 0 or P_m based on its previous action, and switched if the temperature was close to the edge of the deadband. Given this initialization scheme, the majority of continuous actions u_t are close to either 0 or P_m , making conversion to on/off actuation possible with a reasonable number of switchings.

In this use case, our approach took an average of 5.4 iterations to reach consensus. Interestingly, the number of iterations till convergence is almost independent of the population size, which is also observed in [15]. Since the u -update step is distributed to and computed in parallel at each TCL, the overall computation time and the computation cost at the load aggregator scale very well with the population size.

5.2 Use Case 2: Minimize Ramping

While we first evaluated our approach in a reference tracking use case, a major advantage of COHORT compared with existing methods is its ability to coordinate TCLs over long time horizons. Thus, we applied our approach to flatten the *duck curve* by shifting load over a day. Specifically, we used a 15-min control time-step, and plan for an entire day ahead, i.e., $T = 96$. An example of the *duck curve*, named after its resemblance to a duck [20], is given in Figure 4b. We scaled down the load curve such that the TCL demand accounts for 20% of the total demand [16]. We formulate the objective as minimizing total ramping, i.e., the difference in net demand between consecutive time-steps (Eq. 14). $\mathbf{P}_{\text{total}}$, \mathbf{P}_{net} , and \mathbf{P}_{gen} are the total demand, net demand, and renewable generation, respectively. This problem is also known as total variation minimization [7]. While this objective is trickier to optimize, convergence to global optimum is still guaranteed, because g_{tv} is a convex function.

$$g_{\text{tv}}\left(\sum_i \mathbf{u}_i\right) = \sum_{k=t+1}^{t+T} |P_{\text{net},k} - P_{\text{net},k-1}| \quad (14)$$

where, $\mathbf{P}_{\text{net}} = \mathbf{P}_{\text{total}} - \mathbf{P}_{\text{gen}}$

$$\mathbf{P}_{\text{total}} = \mathbf{P}_{\text{non-shiftable}} + \sum_i \mathbf{u}_i$$

The behavior of the TCL population for this use case is also shown in Figure 3. The TCLs systematically shifted their load and reduced ramping by 23.1% compared to the baseline scenario, where the TCL population was operated by on-off control. The temperature of the TCLs shifted accordingly within the deadband. Since the TCLs are operating in cooling mode, reduced energy consumption results in higher temperature, and vice versa. Note that the 23.1% reduction in ramping is based on the assumption that 20% of total

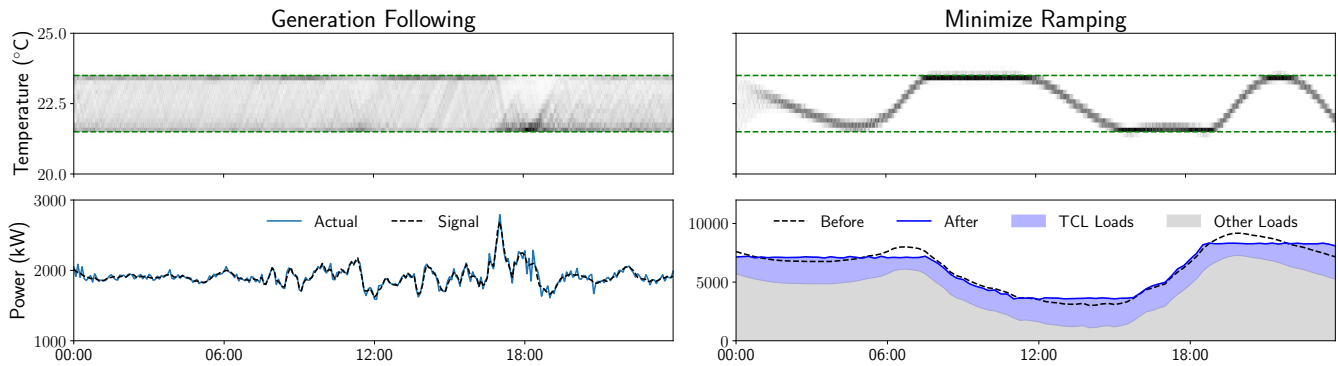


Figure 5: Simulation Study. (Top) the temperature distribution and (Bottom) the aggregate power of the population

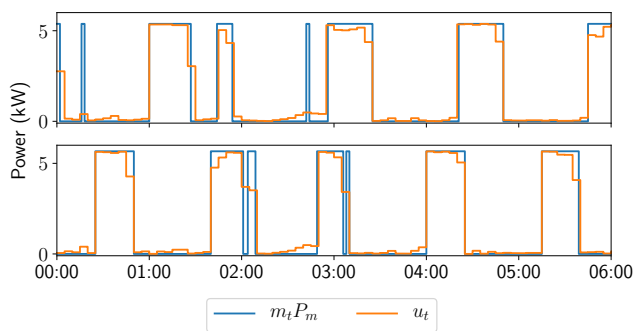


Figure 6: Actions of Individual TCLs

demand is flexible. Further reduction is possible by extending our approach to other flexible loads.

6 EXPERIMENT 2: HARDWARE-IN-THE-LOOP SIMULATION

In this section, we validated COHORT in a HIL simulation, with primary focus on its performance on a real-world testbed. Specifically, we controlled a residential AC via a smart thermostat. More details on the real-world testbed is presented in Section 6.1. We augmented the testbed with simulated instances of residential ACs, modeled after real-world data traces (Section 6.2). We integrated the real-world system and the simulated instances as a HIL simulation, and showcased it in a peak load curtailment use case, based on load profiles from PJM [29], as elaborated in Section 6.3. The HIL simulation was executed from 11-25 July 2020, a 15-day period, and the results are summarized in Section 6.4.

6.1 Real-World Testbed

The testbed is first author’s apartment located in Pittsburgh, PA, USA (Figure 7). The 1632 ft² apartment has three regular occupants and was occupied most of the time during the experimental period. The AC unit is controlled via an ecobee smart thermostat, installed at a location shaded from direct solar radiation, and energy consumption of the AC is monitored for verification only via

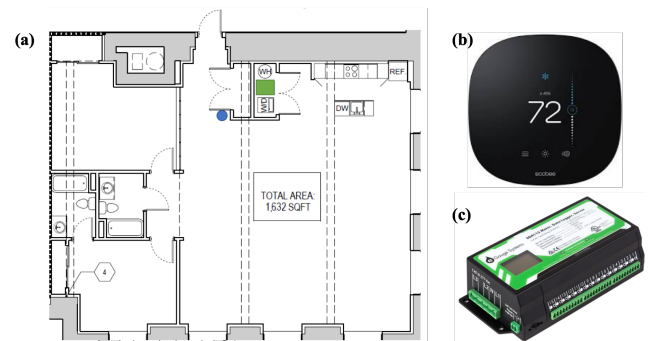


Figure 7: The floor plan of the apartment is shown in (a), with the circle and the rectangle marking the location of the smart thermostat and the indoor AC unit. The apartment is instrumented with (b) an ecobee smart thermostat and (c) an eGauge energy metering unit.

eGauge energy metering [12], with sampling rate up to 1Hz. Power measurements are not needed for control.

We monitored zone temperature and sent commands to the smart thermostat via ecobee API [11]. To get the effect of on/off commands, we sent a low temperature setpoint (70 °F) when we want the AC to be *on* and a high temperature setpoint (80 °F) when we want the AC to be *off*. This simple strategy worked surprisingly well. Figure 8 shows a comparison of the commands vs. the power draw on the AC circuit (at 1Hz) for an on/off event. The response of the AC to *on* is almost instantaneous, and the response to *off* is delayed by a few seconds, but negligible compared to the control time-step. Such response was observed throughout the experiment. The ease of integration with a smart thermostat implies COHORT could be scaled to 11% of households in the US already equipped with smart thermostats [22], with minimal effort and no retrofit.

6.2 Modeling of Residential ACs

To have a population of simulated ACs with realistic thermodynamics and population heterogeneity, we took advantage of ecobee’s *Donate Your Data* dataset [10]. For a comprehensive description of

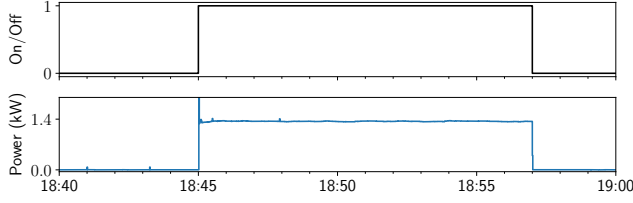


Figure 8: Comparison of the on/off commands vs. the actual power draw by the AC measured at 1Hz

the dataset, we refer interested readers to [17]. We selected households in the same municipal area as the real-world testbed using data from 2019. We only used households with single-stage cooling, no less than 60 cooling days⁶, and less than 10% missing data. These criteria left us with 106 households. While the simulated population size is relatively small, the primary focus of the HIL simulation is on the real-world testbed.

The raw data came in 5-min intervals and we down-sampled it to 15-min based on the control time-step. We used the same model form as [21] (Eq. 15), where T_t is the control temperature, u_t is the duty cycle ratio, $T_{a,t}$ is the outdoor ambient temperature, and $d_{o,t}$ and $d_{s,t}$ are binary flags for occupancy sensor activation and scheduled sleep time. These variables would be explained shortly. Both the model orders, i.e., p and q , and the disturbance terms were selected based on the Akaike Information Criterion (AIC). The median of selected model orders are $p = 5$ and $q = 2$.

$$T_{t+1} = \sum_{i=0}^{p-1} a_i T_{t-i} + \sum_{i=0}^{q-1} b_{u,i} u_{t-i} + b_a T_{a,t} + b_o d_{o,t} + b_s d_{s,t} \quad (15)$$

The state variable, the *control temperature*, is what an ecobee uses for operating the AC with respect to the setpoint. It is a weighted average of the temperature measurements at the main thermostat and remote sensors [17]. Note that the control temperature in the dataset came in 1°F resolution. The control action, *duty cycle ratio*, is the equipment run-time normalized between 0 and 1. Similar to [17], we assumed that the house was occupied if any of the motion sensors were triggered or during scheduled sleep time. Since the raw data from occupancy sensors were sharp spikes, we passed the data through a low-pass filter. We used a separate variable for scheduled sleep time. Despite having different form from Eq. 1b, Eq. 15 can also be written in the form of Eq. 2b as a linear system, and thus the same formulation of flexibility (Eq. 5) still applies.

We evaluated the modeling performance with mean absolute error (MAE) over 1 to 6 hour prediction horizons. Figure 9 shows the distribution of MAE over 106 households. The train set and the test set were the first 2/3 and last 1/3 of cooling days, respectively. The weather was checked to have a similar distribution over the train-test split. The majority of prediction error is less than 1°F even at a 6-hour prediction horizon. This result is comparable to that of [21]. Bear in mind, in interpreting the results, that the control temperature came in a low resolution of 1°F.

⁶Cooling days are defined as days the AC is operating exclusively in cooling mode, following [17].

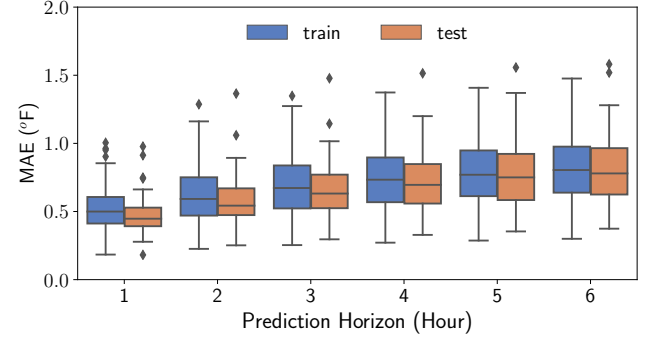


Figure 9: Distribution of model prediction error for 1 to 6 hour prediction horizons over 106 households

6.3 Experimental Setup

Growing peak demand decreases the average utilization of generators [13] and increases the need to build and operate high marginal cost peaking generation [4]. Thus, we showcased our approach on a peak load curtailment use case, the objective of which is given in Eq. 16. As an infinity norm⁷ minimization problem, the solver minimizes the maximum total load within the planning horizon. We used a 15-min control time-step and found a 16-hour planning horizon to be sufficient. We re-planned at each hour to avoid compounding modeling error. Similar to the experiment in Section 5.2, we scaled down the PJM load profile with the assumption that TCL loads account for about 20% of the total load.

$$g_{\text{peak}} \left(\sum_i \mathbf{u}_i \right) = \|\mathbf{P}_{\text{total}}\|_{\infty} \quad (16)$$

where, $\mathbf{P}_{\text{total}} = \mathbf{P}_{\text{non-shiftable}} + \sum_i \mathbf{u}_i$

We integrated the real-world testbed with simulated instances of residential ACs to form an HIL simulation. The HIL simulation is also integrated with day-ahead weather forecast via Dark Sky API [30] and total load profile from PJM Data Miner 2 [29].

We modeled the real-world testbed using the same procedures as described in Section 6.2. Note that the temperature time series pulled from ecobee API comes in 0.1°F resolution, and thus allows for more accurate modeling. To identify the system accurately, the testbed was excited by square wave signals, one-hour on followed by one-hour off, on 10 July 2020. During the experiment, the model was updated regularly based on new data. We used $P_m=1.4\text{kW}$ for the real-world system based on actual power measurements. As discussed in Section 4.2, we initialize \mathbf{u}_i with sequences of interlaced 0s and P_m s to encourage sparsity in the solution. For the AC unit, we initialized the solution for each hour with $[P_m, 0, 0, 0]$ and used an error limit of 0.075kWh.

For the simulated instances, we selected models that performed no worse than 75th percentile on any of the prediction horizon, which left us with 72 households. The dataset does not contain information on the rated power of the AC units, and thus we sized

⁷Infinity norm is defined as $\|\mathbf{x}\|_{\infty} := \max(|x_j|)$.

the AC based on the floor area, following Energy Star recommendations [14]. We used the same occupancy data from the same time last year. The temperature setpoint for the entire population, including the real-world testbed, was assigned to be 75°F throughout the experiment.

6.4 Results

The performance of COHORT compared with the baseline scenario, where the TCL population was operated by on-off control, is exemplified in Figure 10, using time series from 24 July 2020. The baseline TCL load was simulated using a model of the real-world testbed, along with the rest of the population, based on the same weather data. The baseline total load was a scaled down version of the actual PJM load profile, under the assumption that TCLs account for about 20% of the total load. At the aggregate level (Figure 10a), the TCL peak approximately coincides with the utility peak in the baseline scenario (dashed lines). COHORT systematically shifted the TCL load away from peak hours to early morning, thereby reducing the utility peak by 15%. The temperature of the population shifted accordingly. By pre-cooling the households when the demand was low, the TCLs could reduce energy consumption during the peak hours and let temperature slowly float up, without violating comfort constraints. The same behavior is observed in the real-world testbed (Figure 10b).

Also in the real-world testbed, the planned vs. actual duty cycle ratio at each 15-min control time-step shows good overall agreement. By initializing the actions at each hour with $[P_m, 0, 0, 0]$, the *on* events were spaced apart, thereby reducing the risk of short cycling. The temperature, as logged by ecobee at 5-min intervals, was maintained within the deadband. To compare with the baseline scenario, we also include in the same plot the temperature time-series in the real-world testbed on 27 July 2020, when the ecobee maintained temperature at 75°F setpoint using its default strategy. The slightly wider temperature range and the gradual temperature shift were barely perceptible by the occupants. All occupants were happy with their thermal comfort during the experiment. In single blind tests, occupants not involved in the experiment were not able to tell if the AC was operated by COHORT or on-off control.

Over the 15-day experimental period, COHORT reduced daily peak loads by an average of 12.5% (with a 2.9% standard deviation), when the TCL accounts for about 20% of the total load. In comparison, the peak load was reduced by 8.8% in the best-performing case in [6], when the TCL peak account for 74% of the utility peak in the baseline scenario. While the experimental setups are different, the difference in performance likely comes from the fact that a mere 1-hour planning horizon was used in [6]. Instead, we planned ahead for 16-hour and managed to level the total load from about 6:00–21:00 in the case shown in Figure 10a. This affirms the advantage of being able to plan over long time horizons.

7 CONCLUSIONS

We proposed COHORT, a novel distributed control solution for coordinating TCLs to jointly optimize a grid-level objective, while satisfying each TCL’s end-use requirements and operational constraints. Our approach decomposes the grid-level problem into subproblems and coordinates their solutions to find the grid-level

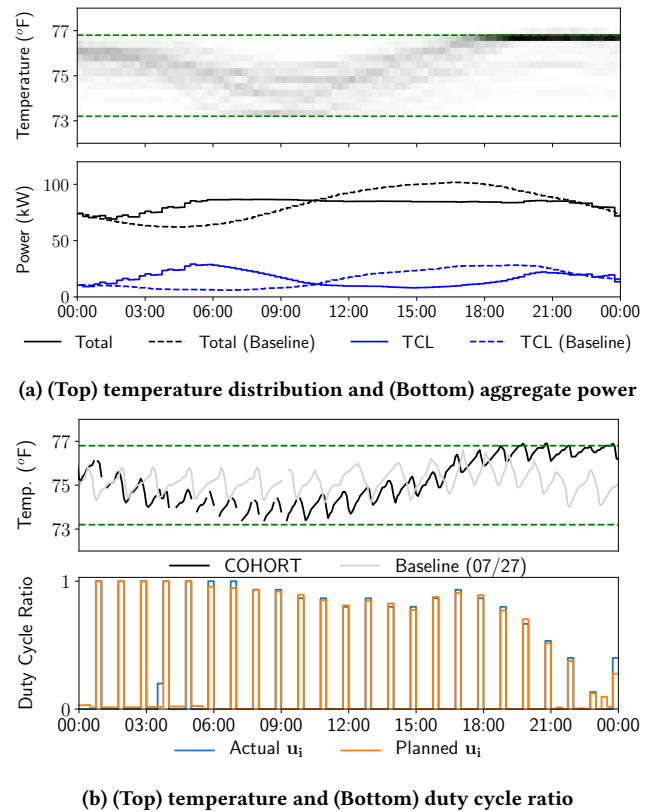


Figure 10: HIL Simulation. Behaviour of (a) the population and (b) the real-world testbed on 2020/07/24

optimum. To be computationally viable over long planning horizons, we apply convex relaxation to the discrete action space and characterize each TCL’s flexibility as a convex set. After coordination, each TCL tracks the agreed-upon power trajectory locally with a PWM-based strategy, which translates the continuous power back to on/off actuation. Since each TCL is responsible for its own control, it can incorporate detailed and system-specific dynamics and constraints, which is difficult to accomplish in a centralized architecture. Furthermore, the coordination process is independent of each TCL’s dynamics and control scheme, making COHORT extensible to other flexible loads.

We validated COHORT in simulation studies and a HIL simulation to address challenges arising from growing peak demand and increasing penetration of renewable generation. In the generation following use case, our approach showed comparable performance to prior work. A major advantage of our approach compared with existing work is its ability to coordinate TCLs over long planning horizons. Thus, we applied it to load shifting use cases, assuming TCLs account for 20% of the total load [16]. Firstly, we used it to smooth out the duck curve, based on actual load profile from CAISO, and reduced ramping by 23.1% compared to the baseline scenario. Secondly, we applied it to curtail peak load based on load profile from PJM. The experiment was conducted on a HIL simulation, including a real-world testbed. Over the 15-day experiment

period, COHORT reduced daily peak loads by an average of 12.5%. Furthermore, the occupants living in the real-world testbed, i.e., the first author's apartment, reported no discomfort and could not distinguish whether the AC was operated by on-off control or COHORT. Finally, COHORT is easily integrated with a commercial smart thermostat, which provides readily-available control and communication infrastructure. Thus, our approach is scalable to households already equipped with smart thermostats with minimal effort and no retrofit.

In summary, COHORT is shown to be a practical, scalable, and versatile solution for coordinating TCLs to provide grid services. Currently, quite a few iterations are required to reach consensus for problems with long planning horizons, which may raise concern about the communication cost. This was discussed in [15], with the conclusion that the network latency and message size between the aggregator and the end users during coordination are not bottlenecks for modern networks⁸. Regardless, the number of iterations could be reduced by initializing COHORT with an approximate solution from imitation learning, to be incorporated as future work. Another research direction would be to extend the current work to other flexible loads, with tracking strategies tailored to those loads.

ACKNOWLEDGMENTS

This material is partly based upon work supported by the U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy (EERE) under the Building Technologies Office Award Number DE-EE0007682. We would like to thank Eric Burger and Scott Moura, whose prior work was the primary inspiration for us. We also thank ecobee and its customers for their data. *Disclaimer: The views expressed in the article do not necessarily represent the views of the U.S. Department of Energy or the United States Government.*

REFERENCES

- [1] Stephen Boyd, Neal Parikh, and Eric Chu. 2011. *Distributed optimization and statistical learning via the alternating direction method of multipliers*. Now Publishers Inc.
- [2] Eric M Burger and Scott J Moura. 2017. Generation following with thermostatically controlled loads via alternating direction method of multipliers sharing algorithm. *Electric Power Systems Research* 146 (2017), 141–160.
- [3] William Burke and David Auslander. 2009. Low-frequency pulse width modulation design for HVAC compressors. In *ASME 2009 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. American Society of Mechanical Engineers Digital Collection, 291–297.
- [4] Duncan S Callaway and Ian A Hiskens. 2010. Achieving controllability of electric loads. *Proc. IEEE* 99, 1 (2010), 184–199.
- [5] Bingqing Chen, Weiran Yao, Jonathan Francis, and Mario Berges. 2020. Learning a Distributed Control Scheme for Demand Flexibility in Thermostatically Controlled Loads. In *2020 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm) (IEEE SmartGridComm'20)*. Tempe, USA.
- [6] Wesley J Cole, Joshua D Rhodes, William Gorman, Krystian X Perez, Michael E Webber, and Thomas F Edgar. 2014. Community-scale residential air conditioning control for effective grid management. *Applied Energy* 130 (2014), 428–436.
- [7] Laurent Condat. 2013. A direct algorithm for 1-D total variation denoising. *IEEE Signal Processing Letters* 20, 11 (2013), 1054–1057.
- [8] Giuseppe Tommaso Costanzo, Oliver Gehrke, Daniel Esteban Morales Bondy, Fabrizio Sossan, Henrik Bindner, Jacopo Parvizi, and Henrik Madsen. 2013. A coordination scheme for distributed model predictive control: Integration of flexible DERs. In *IEEE PES ISGT Europe 2013*. IEEE, 1–5.
- [9] Steven Diamond and Stephen Boyd. 2016. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research* 17, 83 (2016), 1–5.
- [10] ecobee. [n.d.]. *Donate your data*. <https://www.ecobee.com/donate-your-data/> Accessed: 2020-06-30.
- [11] ecobee. [n.d.]. *Getting Started with the ecobee API*. <https://www.ecobee.com/home/developer/api/introduction/index.shtml> Accessed: 2020-06-30.
- [12] eGauge. [n.d.]. *eGauge Core - Residential*. <https://store.egauge.net/egauge-core-residential> Accessed: 2020-05-22.
- [13] US EIA. 2014. Peak-to-average electricity demand ratio rising in New England and many other US regions. <https://www.eia.gov/todayinenergy/detail.php?id=15051> Accessed: 2020-07-30.
- [14] Energy Star. [n.d.]. *Room Air Conditioner-Buying Guidance*. https://www.energystar.gov/products/heating_cooling/air_conditioning_room Accessed: 2020-05-22.
- [15] Daniel Gebbran, Wibowo Hardjawana, Sleiman Mhanna, Branka Vucetic, Archie C Chapman, and Gregor Verbic. 2020. Practical Considerations of DER Coordination with Distributed Optimal Power Flow. In *IEEE International Conference on Smart Grids and Energy Systems*.
- [16] He Hao, Borhan M Sanandaji, Kameshwar Poolla, and Tyrone L Vincent. 2014. Aggregate flexibility of thermostatically controlled loads. *IEEE Transactions on Power Systems* 30, 1 (2014), 189–198.
- [17] Brent Huchuk, William O'Brien, and Scott Sanner. 2018. A longitudinal study of thermostat behaviors based on climate, seasonal, and energy price considerations using connected thermostat data. *Building and Environment* 139 (2018), 199–210.
- [18] Gabriela Hug, Soumyya Kar, and Chenye Wu. 2015. Consensus+ innovations approach for distributed multiagent coordination in a microgrid. *IEEE Transactions on Smart Grid* 6, 4 (2015), 1893–1903.
- [19] California ISO. [n.d.]. *Today's Outlook*. <http://www.caiso.com/TodaysOutlook/Pages/default.aspx> Accessed: 2020-06-30.
- [20] Becca Jones-Albertus. 2017. Confronting the Duck Curve: How to Address Over-Generation of Solar Energy. <https://www.energy.gov/eere/articles/confronting-duck-curve-how-address-over-generation-solar-energy> Accessed: 2020-05-15.
- [21] Emre Can Kara, Michaelangelo D Tabone, Jason S MacDonald, Duncan S Callaway, and Sila Kilicote. 2014. Quantifying flexibility of residential thermostatically controlled loads for demand response: a data-driven approach. In *Proceedings of the 1st ACM conference on embedded systems for energy-efficient buildings*. 140–147.
- [22] Jen King. 2018. Energy impacts of smart home technologies. *Report A1801* (2018).
- [23] Stephan Koch, Johanna L Mathieu, and Duncan S Callaway. 2011. Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services. In *Proc. PSCC*. Citeseer, 1–7.
- [24] Nariman Mahdavi, Julio H Braslavsky, Maria M Seron, and Samuel R West. 2017. Model predictive control of distributed air-conditioning loads to compensate fluctuations in solar power. *IEEE Transactions on Smart Grid* 8, 6 (2017), 3055–3065.
- [25] Fabian L Müller and B Jansen. 2019. Large-scale demonstration of precise demand response provided by residential heat pumps. *Applied Energy* 239 (2019), 836–845.
- [26] Monica Neukomm, Valerie Nubbe, and Robert Fares. 2019. *Grid-Interactive Efficient Buildings*. Technical Report. US Department of Energy, Navigant.
- [27] Truong X Nghiem and Colin N Jones. 2017. Data-driven demand response modeling and control of buildings with gaussian processes. In *2017 American Control Conference (ACC)*. IEEE, 2919–2924.
- [28] Sangil Park. 1999. Principles of sigma-delta modulation for analog-to-digital converters. (1999).
- [29] PJM. [n.d.]. *Data Miner 2 Seven-Day Load Forecast*. https://dataminer2.pjm.com/feed/load_frcstd_7_day Accessed: 2020-07-09.
- [30] Dark Sky. [n.d.]. *Dark Sky API*. <https://darksky.net/dev> Accessed: 2019-06-19.
- [31] Ryan Tibshirani. [n.d.]. *Proximal Gradient Descent*. <http://www.stat.cmu.edu/ryantibs/convexopt/lectures/prox-grad.pdf>. Accessed: 2019-06-19.
- [32] Simon H Tindemans, Vincenzo Trovato, and Goran Strbac. 2015. Decentralized control of thermostatic loads for flexible demand response. *IEEE Transactions on Control Systems Technology* 23, 5 (2015), 1685–1700.
- [33] Evangelos Vrettos, Frauke Oldewurtel, and Göran Andersson. 2016. Robust energy-constrained frequency reserves from aggregations of commercial buildings. *IEEE Transactions on Power Systems* 31, 6 (2016), 4272–4285.
- [34] Zhijie Xu, Ruisheng Diao, Shuai Lu, Jianming Lian, and Yu Zhang. 2014. Modeling of electric water heaters for demand response: A baseline PDE model. *IEEE Transactions on Smart Grid* 5, 5 (2014), 2203–2210.
- [35] Wei Zhang, Jianming Lian, Chin-Yao Chang, and Karanjit Kalsi. 2013. Aggregated modeling and control of air conditioning loads for demand response. *IEEE transactions on power systems* 28, 4 (2013), 4655–4664.
- [36] Lin Zhao, Wei Zhang, He Hao, and Karanjit Kalsi. 2017. A geometric approach to aggregate flexibility modeling of thermostatically controlled loads. *IEEE Transactions on Power Systems* 32, 6 (2017), 4721–4731.

⁸More concretely, the message size is less than 1KB per agent per iteration in our problem setup and the total message size grows linearly with the number of agents. Overall, the bandwidth requirement is low. As a reference, the network latency for 4G network ranges from 30 to 160ms. Even for 100 iterations, the latency is small compared to a 1-hour re-planning time-step sufficient for load shifting use cases.