2022 IEEE International Conference on Pervasive Computing and Communications (PerCom) | 978-1-6654-1643-6/22/$31.00 ©2022 IEEE | DOI: 10.1109/PerCom53586.2022.9762402

HVAC Power Conservation through Reverse Auctions and Machine Learning

Enrico Casella*, Atieh R. Khamesi[§], Simone Silvestri*, D. A. Baker[†], and Sajal K. Das[‡]
*Department of Computer Science, University of Kentucky, Lexington, KY, USA
[§]Ericsson, Santa Clara, CA, USA
[†]Department of Psychological Science, Missouri University of Science & Technology, Rolla, MO, USA
[‡]Department of Computer Science, Missouri University of Science & Technology, Rolla, MO, USA
Email: {enrico.casella, atieh.khamesi, simone.silvestri}@uky.edu, {bakerden, sdas}@mst.edu

***Abstract*—Prolonged rotating outages and exorbitant energy bills, recently experienced in California and Texas, have exposed the limitations and need for modernizing electric power systems. The occurrence of such events is a consequence of *peak loads*, often due to *extreme outside temperatures* that simultaneously trigger Heating Ventilation Air Conditioning (HVAC) systems. Leveraging pervasive computing technologies, such as smart meters and smart thermostats, this paper introduces a comprehensive approach to perform residential HVAC power conservation and prevent these catastrophic events. Differently from previous solutions, our approach models realistic user behavior and HVAC dynamics of individual homes. Specifically, we formulate a novel reverse auction-based problem, called *POwer Conservation Optimization* (POCO). The goal is to perform power conservation by motivating users to temporarily adjust their HVAC thermostat settings in exchange for financial rewards. We prove that POCO ensures *truthfulness* and *individual rationality* of the auction mechanism, although it is an *NP-hard* problem. Therefore, we propose an efficient heuristic, called *Greedy Ranking AllocatioN* (GRAN), which we prove ensures the same formal properties, while incurring only a polynomial complexity. To predict power savings resulting from an HVAC thermostat adjustments, we propose a novel machine learning-based technique called *Power Saving Prediction* (PSP). In addition, we conduct an *online survey* to study the willingness to adopt the proposed system and to model realistic user behavior. Survey results show willingness of adoption above 79% and a highly heterogeneous and non-linear user behavior. We perform extensive experiments using high-fidelity simulator *EnergyPlus*. Results show that PSP outperforms a state-of-the-art solution obtaining 85% predictions within a 5% error margin. Furthermore, GRAN achieves near-optimal performance, outperforming a recent state-of-the-art approach obtaining results between 58% and 68% closer to the optimum.**

***Index Terms*—HVAC power conservation, machine learning power saving predictions, reverse auctions, smart homes, cyber-physical pervasive computing**

I. Introduction

Motivation: In recent years, increased urbanization has led to a significant rise of electric power consumption in residential buildings, contributing to the overall power consumption [1]. Simultaneously, increased occurrences of extreme temperature [2] have resulted in episodes of extreme weather such as winter storms [3], heat waves [4], and wildfires [5]. Such events have exacerbated *peak loads* in smart grids owing to spikes in user demand because *extreme outside temperatures* often trigger Heating Ventilation and Air Conditioning (HVAC) systems [6]. For example, the February 2021 winter storm in Texas led to a historical winter peak demand record of $69,150$MW [3]. Consequently, many Texas residential customers received energy bills higher than $\$5,000$ and the wholesale energy saw a $17,900\%$ increase [7].

Utility companies struggle to deal with peak loads and high demands that put the generation, transmission, and distribution systems under enormous stress, increasing the risk of blackouts [8]. For example, in August 2020, heat waves in California led to multiple hours of outage due to poor planning that included rotating outages for several hours [4]. Overall, the United States have witnessed a 67% increase in major power outages from weather-related events since 2000 [9].

Although smart grids are experiencing a larger penetration of renewable energy sources and energy storage devices [10] to help reduce the impact of peak loads, as of today, they are still insufficient to deal with high peaks during extreme weather conditions. Hence, power systems in several countries (e.g., Texas [11]–[13] and California [4] in the US, Iran [14], and South Africa [15]) are also focusing on alternative techniques, such as *power conservation*, by reducing the demand through *rotating outages* [11]. Unfortunately, such outages severely impact the user comfort and often last longer than the originally planned few hours [4] [3].

An alternative approach to power conservation is to *directly ask* users (via social media, local newspapers, and television or radio advertisements) to reduce their electricity consumption during extreme hot and cold days. This practice has been implemented in 2021 by the Electric Reliability Council of Texas (ERCOT) [12] [13]. For such efforts, HVAC has been specifically targeted since, as stated by ERCOT, "every degree of cooling increases your energy use by six to eight percent" [13]. According to the U.S. Energy Information Administration, in 2015, 51% of total power consumption from residential buildings was related to space heating and air conditioning [16], a trend that has been increasing since then. Moreover, high peak demand is usually associated with extreme outside temperatures [6], making HVAC the most suitable appliance for power conservation.

State of the Art: Recent years have seen a diffusion of *pervasive computing* devices in smart grids realized through Internet of Things (IoT) devices, such as smart meters in the Advanced

Metering Infrastructure (AMI) [17], smart thermostats [18], and home energy management systems [19]. Increasing availability of such devices enables the design and implementation of novel, more effective, and automated approaches to achieve power conservation [20]. Smart grid operators are focusing on these approaches, as recently stated by the president and CEO of California Independent System Operator, calling for a "*modernized and well-integrated resource adequacy framework*" [21] that includes flexible power conservation [22].

Some existing works in the literature have focused on power conservation that exploits IoT-enabled home devices. For instance, in [23] and [24], IoT appliances are scheduled to reduce the electricity load during peak hours. Similarly, in [25], auction theory is exploited to encourage users to reduce power consumption in exchange for a financial reward. However, these schemes mostly consider *abstract appliances* and do not take into account either user behavioral models or the complex dynamics of different appliances on power consumption, thus limiting their applicability and effectiveness in practice. A few works focus on the impact of HVAC during periods of high load [6], [26]. However, the authors of [6] only provide preliminary results to support the use of HVAC, while the auhors of [26] proposes a basic flat-rate framework that does not incentivize participation due to monthly commitments.

Contributions and key novelties: To the best of our knowledge, our work is the first to design a comprehensive framework for power conservation that simultaneously addresses user engagement and bidding behaviors, specific dynamics of HVAC related to individual homes, and the system operator's overall objective.

Specifically, we develop an IoT-enabled framework to achieve HVAC-based power conservation under extreme temperatures. Our approach exploits *reverse auctions* to realize an Incentive-Based Power Conservation (IBPC) scheme. As depicted in Fig. 1, our system includes a utility company and a set of residential homes (referred to as *"users"* in this paper) equipped with *Smart Energy Management Systems* (SEMSs). When the utility company anticipates a peak load, it asks the user SEMSs to submit their preferences, called *bids* in auction terminology. Here, bids represent monetary rewards corresponding to the potential temperature changes (i.e., thermostat setting) that the users would be willing to set for a certain period of time, say an hour or more. A SEMS may directly inquire its user, through a smartphone app, or automatically submit bids using a pre-defined profile.

A change in the HVAC thermostat setting may result in non-linear energy savings depending on the outside weather (temperature, wind, humidity, etc.) and house characteristics (size, insulation, windows U-factor, etc.). Since some of these factors are often unknown to the SEMS and exhibit complex interactions, we develop a machine learning-based algorithm, called *Power Saving Prediction* (PSP), to predict power savings from thermostat changes. The PSP is independently executed in each user SEMS.

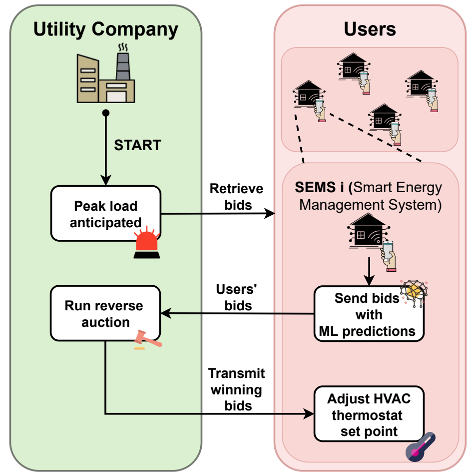

Fig. 1: A schematic overview of the proposed framework

Subsequently, different SEMSs send their bids and the corresponding power saving (predicted by PSP) to the utility company that solves a truthful and individually rational *POwer Conservation Optimization* (POCO) problem to select the winners of the reverse auction. The set of winners are communicated to the corresponding SEMSs which adjust the thermostat settings accordingly. We prove that POCO is an NP-hard problem, and propose a Greedy Ranking AllocatioN (GRAN) algorithm to find a sub-optimal solution efficiently. We prove that GRAN is also *truthful* and *individually rational*.

To evaluate the performance of our proposed framework, we use the high-fidelity gold-standard energy simulator *EnergyPlus*, a software funded by the U.S. Department of Energy [27], and tested according to ASHRAE Standard 140 methodology[1]. Additionally, we conduct an *online survey* involving 200 subjects to verify the willingness of using such auction-based system and model the realistic bidding behavior with respect to different potential temperature changes. Survey results demonstrate that more than 79% of users are willing to adopt such system and bidding preferences are highly variable and non-linear. Extensive simulation experiments demonstrate that PSP outperforms a recent state-of-the-art approach and provides 85% predictions within a 5% error margin, while requiring few training samples. Moreover, GRAN achieves near-optimal performance and outperforms a recently proposed auction-based approach by obtaining results between 58% and 68% closer to the optimum in different scenarios.

In summary, the main contributions of this work are:

- We propose a comprehensive reverse auction-based approach for power conservation that simultaneously addresses user engagement and bidding behaviors, specific dynamics of HVAC, and the system operator's objective;
- We propose optimal and heuristic solutions for selecting auction winners and prove their truthfulness and individual rationality;
- We show through realistic experiments that our solutions are superior with respect to state-of-the-art approaches.

[1] According to U.S. Department of Energy, *"The ASHRAE Standard 140 is the framework for establishing confidence in energy modeling engines"* [28].

The remainder of the paper is organized as follows. The related work is analyzed in Sec. II. Sec. III introduces the POCO problem. Sec. IV discusses the heuristic GRAN and proves its formal properties. In Sec. V we describe the power saving prediction algorithm PSP, while in Sec. VI we present the survey. Results and performance evaluation are presented in Sec. VII. A discussion of the most relevant findings, requirements, and limitations is presented in Sec. VIII, while Sec. IX concludes the paper.

II. Related Work

Peak load reduction, also known as peak load shaving, has been previously studied in power grids. Price-Based Demand Response (PBDR) has been one of the most investigated approaches to address this problem. According to PBDR, the price of electricity varies throughout the day, with the purpose of shifting user habits and reduce the likelihood of peaks [29]. Tariffs can be static [30], or change in real-time [31]. A comprehensive survey of PBDR solutions is presented in [29]. However, it has been shown that PBDR results in low user engagement, since users quickly stop keeping track of different tariffs and consumption trends [32], [33]. As a result, PBDR has low effectiveness in reducing consumption peaks. Moreover, variable rates may cause extremely high energy bills under conditions of high demand [3].

To address these shortcomings, Incentive-Based Power Conservation (IBPC) approaches have been introduced, with the goal of engaging users effectively through monetary incentives. Although IBPC incurs in an additional cost for the utility company, a recent study has shown that increasing 100MW generation unit can be as costly as providing monetary rewards up to a period of 36.2 years [26]. Therefore, utility companies are highly incentivized and have wide margins of profitability in using an IBPC approach. Early attempts of IBPC include direct load control [34] and curtail-able load programs [35] [36]. However, while the former is invasive and does not solve the engagement problem, the latter focuses on appliances which may not always be running during extreme events of peak load.

A more recent approach for incentive-based approaches consists in *demand-side bidding*, which is the main focus of this paper. In this approach, users actively submit their bids to the utility company, declaring how much money they want in exchange of a load reduction. Demand-side bidding is supported by recent studies that show how user engagement can be enhanced if users are actively involved in decision making [37]. The authors of [38] implement an auction mechanism that uses electric vehicles' storage during peak loads. However, vehicles' prices represent a strong limitation for the pervasive deployment of such solution [39]. The authors of [23] and [24] propose the use of auctions to schedule or turn off appliances in order to avoid peak loads. However, both these works only consider abstract appliances, overlooking the specific dynamics that impact power consumption and user comfort [40] [41]. Furthermore, introducing a large amount of appliances in the power conservation scheme may result in excessive user effort, which can potentially lead to *response fatigue* [32], [37]. This affects long-term user engagement and, as a consequence, successful power conservation. In order to cope with events like winter storms and heat waves, it is important to develop an approach that is easy and intuitive to the user, while addressing appliances that incur high consumption, especially during weather events, such as HVAC.

Few works focus on HVAC for power conservation. For example, the authors in [6] provide a preliminary study to motivate the use of HVAC for peak loads, i.e., its highest share of annual energy consumption, and the correlation of its high usage during peak load periods. However, the authors do not include a comprehensive framework to realize power conservation. A more comprehensive approach based on HVAC is proposed in [26]. However, the authors propose a solution with a monthly flat-rate and a single option of thermostat change. Compared to our approach, this limits the choice of thermostat change and preferred price, while constricting to month long commitments rather than daily. An approach related to power conservation in data centers during periods of peak loads has been recently proposed in [25]. Similarly to our paper, the authors propose a reverse auction-based solution. Due to these similarities, we use this approach for performance comparison as described in Section VII-C1. However, since the focus of [25] is on data centers, the dynamics of power savings do not address HVAC specifically, and are not explicitly described. Furthermore, the authors do not model user behavior and only consider one bid per user. Conversely, in this paper we address HVAC dynamics by proposing the PSP algorithm to predict power saving resulting from a thermostat change. We also consider realistic user behavior modeling by means of an online survey, and design an approach that considers multiple bids submitted by the same user.

Overall, this paper advances the state-of-the-art by considering a comprehensive auction-based framework for power conservation that simultaneously addresses user engagement, realistic bidding behaviors, specific dynamics of HVAC related to individual homes through power saving predictions, and the system operator's overall objective.

III. Problem Formulation

We consider a set of N home users served by a utility company. Users are equipped with an Internet-connected Smart Energy Management System (SEMS) that monitors, learns, analyzes, and controls the HVAC system to set temperature changes. The SEMS also interacts with the utility company to implement the auction framework. When the utility company predicts a peak load with an expected total power consumption P_T, it calculates the *power cap* $P_C = \alpha \times P_T$, where $\alpha \in [0, 1)$, according to the system's characteristics, such as generation capacity, cost of generation, capacity of transmission/distribution lines, etc.[2] Therefore, the required *power saving* is $P_S = P_T - P_C$.

[2]We assume that the utility company predicts the peak load and its duration. The proposed framework is supposed to be executed for the duration predicted by the utility company.

The utility company alerts the user SEMSs that the power conservation auction is activated, requesting for bids. A SEMS asks its user directly, or submits the bids based on the pre-defined profile. In the following formulation, we assume that all users participate in the auction. Such formulation can be easily extended to consider only a portion of participating users. As a result, the SEMS of user i submits to the utility company a set $\mathbf{B_i} = \{B_{ij} = (\Delta P_{ij}, \Delta T_{ij}, C_{ij}) : 1 \leq j \leq M_i\}$ of M_i bids, where ΔP_{ij}, ΔT_{ij}, C_{ij} represent respectively the power saving, temperature change, and monetary compensation for user i and bid j. Note that, in the following, we use the terms "cost" and "monetary compensation" to represent financial rewards from the utility company and user perspective, respectively. We discuss in Section V how the power savings ΔP_{ij}, corresponding to the temperature change ΔT_{ij}, are predicted by the PSP algorithm.

After receiving bids from the users, the utility company performs the auctioneer tasks, i.e., selects the winners and computes the payments. The winners are a subset of N users, and the utility company only selects one bid per winner. The winner selection strategy is formulated as an Integer Linear Programming (ILP) optimization problem that aims to minimize the costs in terms of the paid compensations, while satisfying the power cap constraint. As shown in Section VI, there is a correlation between the cost C_{ij} and the temperature change ΔT_{ij} of a bid. Intuitively, a user bids higher for higher temperature changes due to higher discomfort. As a result, minimizing the cost has also the implicit effect of minimizing the discomfort of the user. We refer to this as the *POwer Conservation Optimization* (POCO) problem defined as:

$$\min \sum_{i=1}^{N} \sum_{j=1}^{M_i} C_{ij} w_{ij} \tag{1a}$$

subject to

$$w_{ij} \in \{0, 1\}, \quad i = 1, \dots, N, \quad j = 1, \dots, M_i \tag{1b}$$

$$\sum_{j=1}^{M_i} w_{ij} \leq 1, \quad i = 1, \dots, N \tag{1c}$$

$$\sum_{i=1}^{N} \sum_{j=1}^{M_i} \Delta P_{ij} w_{ij} \geq P_S \tag{1d}$$

Expression (1a) defines the goal of minimizing the total cost. Constraint (1b) defines the decision variable w_{ij}, which is equal to 1 when user i is selected as a winner in the j-th bid, and 0 otherwise. Constraint (1c) ensures that no more than one bid is selected for each user. Finally, inequality (1d) guarantees that the power cap constraint is met.

After selecting the winners by solving the above problem, we propose the payment rule as follows. Let the objective function in (1a) be denoted as $f(\cdot)$. The payment E_k to the user k, who is a winner of the reverse auction, is given by

$$E_k = f(\mathbf{w}^{(-\mathbf{k})*}) - f(\mathbf{w}^*) + \sum_{j=1}^{M_k} C_{kj} w_{kj}^* \tag{2}$$

where $\mathbf{w}^*$ is the optimal solution of POCO, $\mathbf{w}^{(-\mathbf{k})*}$ is the optimal solution when user k does not participate, and $\sum_{j=1}^{M_k} C_{kj} w_{kj}^*$ corresponds to the winning bid of user k. Each winning user k gains a non-negative utility, i.e., revenue, defined as $U_k = E_k - C_k$. In the following, we prove *truthfulness* and *individual rationality* of POCO in order to ensure an effective power conservation program [42]. Truthfulness prevents users from under-bidding or over-bidding, as it would lead to a reduced utility U_k, thus preventing potential unhealthy behaviors. Individual rationality guarantees that each winning user is paid an amount that ensures non-negative utility ($U_k \geq 0$). We provide a sketch of the proof of truthfulness due to space limitations.

Theorem 1. *The reverse auction mechanism, as defined by the POCO problem and the payment rule in Eq.* (2)*, is truthful.*

Proof Sketch. Given the payment rule E_k for user k as defined in Eq. (2), we consider the utility of that user as $U_k = E_k - \sum_{j=1}^{M_k} V_{kj} w_{kj}$, where V_{kj} is the so called *true valuation* following the auction terminology. The property of truthfulness guarantees that a user does not gain by bidding higher or lower than its true valuation. Let U_k and U'_k be respectively the utility of the true valuation ($C_k = V_k$) and untrue valuation ($C_k \neq V_k$). To prove truthfulness, we prove that $U_k - U'_k \geq 0$. After several steps, we obtain: $U_k - U'_k = -\sum_{i=1}^{N} \sum_{j=1}^{M} C_{ij} w_{ij}^* + \sum_{i=1}^{N} \sum_{j=1}^{M} C_{ij} w_{ij}^{*'}$. Since w_{ij}^* is the optimal solution when the user k declares a truthful compensation (i.e., $V_k = C_k$), the first sum is always less than or equal to the second sum, because the optimal solution of a minimization problem holds the smallest objective value than any other solution. Therefore $U_k - U'_k \geq 0$, and the equality holds when $w_{ij}^* = w_{ij}^{*'}$. □

Next, we prove that POCO is individually rational, i.e., the revenue U_k of each user is non-negative.

Theorem 2. *POCO is individually rational.*

Proof. In order to prove non-negative compensation, let us consider the payment rule in Eq. (2). Assuming $\sum_{j=1}^{M_k} C_{kj} w_{kj}^* \geq 0$, since users can not ask for a negative compensation, we need to prove that $f(w_{ij}^{(-k)*}) \geq f(w_{ij}^*)$. This is always true. In fact, since $f(\cdot)$ is a minimization problem, the first term can not improve the solution found in the second term, since it has fewer elements to pick from. □

Finally, we prove the NP-hardness of POCO, motivating the need for an efficient heuristic.

Theorem 3. *POCO is an NP-hard problem.*

Sketch of Proof. The NP-hardness can be proven as a reduction from the minimum 0-1 knapsack problem (minKP) [43]. The minKP looks for the set of items with minimum weight and a cumulative value larger than or equal to a target value. We can translate any instance of minKP into an instance of POCO by considering one bid per user. Given a minKP item, we create a bid for a user with cost equal to the weight of the item, and the power saving equal to its value. We also set the power saving P_S equal to the target value of minKP.

Solving the POCO problem is to select the set of winning bids with minimum cost and the power saving larger than or equal to P_S. This corresponds to the set of elements with minimum weight and satisfying the target value constraint. In other words, POCO is at least as hard as minKP, hence it is NP-Hard. $\square$

IV. The GRAN Mechanism

In this section, we propose a heuristic called Greedy Ranking AllocatioN (GRAN) to find an efficient solution for POCO, while guaranteeing truthfulness and individual rationality of the auction mechanism.

A. Winner Selection and Payment Rule

The basic idea of GRAN is to prioritize bids with a better ratio of cost over the amount of power saved. This ratio is used to calculate a *ranking criterion* sorted in non-decreasing order. Winners are selected by picking their best bid according to the ranking criterion, until the desired power saving P_C is reached. The pseudo-code of GRAN is provided in Algorithm 1.

Algorithm 1: GRAN: Greedy Ranking AllocatioN

Input : P_T, α, and $\mathbf{B}_i$ $i = 1, \ldots, N$
Output: List of Winners $\mathbf{W}$
1 $\mathbf{W} \leftarrow \emptyset, P_{CS} = 0$ \\Variables initialization
2 $P_C = \alpha \cdot P_T$ \\Power cap
3 $P_S = P_T - P_C$
4 $\mathbf{R} \leftarrow \{R_{ij} = \frac{C_{ij}}{\Delta P_{ij}} \quad i = 1, \ldots, N, j = 1, \ldots, M_i\}$
5 Sort elements of list $\mathbf{R}$ in a non-decreasing order
6 **while** $P_{CS} < P_S$ *and* $\mathbf{R} \neq \emptyset$ **do**
7 Let $R_{\hat{i}\hat{j}}$ be the first element in $\mathbf{R}$ and
8 Let $B_{\hat{i}\hat{j}} = (\Delta P_{\hat{i}\hat{j}}, \Delta T_{\hat{i}\hat{j}}, C_{\hat{i}\hat{j}})$ be the bid corresponding to $R_{\hat{i}\hat{j}}$ \\ *Update values*
9 $P_{CS} = P_{CS} + \Delta P_{\hat{i}\hat{j}}$
10 $\mathbf{W} \leftarrow B_{\hat{i}\hat{j}}$
11 Remove all bids of user $\hat{i}$ from $\mathbf{R}$
12 **end**
13 **return** $\mathbf{W}$

In line 1 of Algorithm 1, we initialize the list of the auction winners $\mathbf{W}$ and the variable storing the cumulative power saving P_{CS}. We then calculate the power cap P_C and the amount of power saving P_S that represents the power constraint in the Inequality (1d) (lines $2-3$). Since our goal is to minimize the objective function in (1a), GRAN uses a ranking criterion which gives precedence to the bids with low cost and large power saving. GRAN uses a list $\mathbf{R}$ that stores the values of ranking criterion in non-decreasing order (lines $4-5$).

In the while loop (lines 6-12), we go through the list until the power cap constraint is satisfied, i.e., the cumulative power saving P_{CS} is greater than or equal to the required power saving P_S. At each iteration, we pick the bid $B_{\hat{i}\hat{j}}$ with the smallest ranking criterion $R_{\hat{i}\hat{j}}$ in $\mathbf{R}$ (line $7-8$). Therefore, we increase P_{CS} by the corresponding power saved (line 9) and we add the winning bid $B_{\hat{i}\hat{j}}$ to the list of winners $\mathbf{W}$ (line 10). Finally, we remove all other elements from user $\hat{i}$ in $\mathbf{R}$ (line 11), since only one bid per winner should be selected.

GRAN terminates as soon as the power saving is met, i.e., $P_{CS} \geq P_S$. Subsequently, the new thermostat settings of the winners are sent to the corresponding SEMSs, and the utility company pays the winners. For this purpose, we propose a truthful payment rule for GRAN as described in Algorithm 2. It may be possible that GRAN is unable to meet the power cap and terminates the while loop because $\mathbf{R} = \emptyset$. In this case, the utility company may increase the power cap, thus reducing the required power saving, by supplementing the auction mechanism with other approaches for power conservation. Nevertheless, in all our experiments, we use a power cap that far exceeds similar power reductions [26], and this situation never occurred.

Algorithm 2: GRAN payment rule

Input : List of Winners $\mathbf{W}$, $GRAN$ Algorithm
Output: Payment Vector $\mathbf{E}$
1 **foreach** $B_{ij} \in \mathbf{W}$ **do**
2 $\mathbf{W}_{-i} = GRAN(\mathbf{B_{-i}})$; // $\mathbf{B_{-i}} = \bigcup_{k=1}^{N} \mathbf{B}_k \setminus \{\mathbf{B}_i\}$
3 Let $B_{\bar{i}\bar{j}}$ be the last element added to $\mathbf{W}_{-i}$
4 $E_i = R_{\bar{i}\bar{j}} \Delta P_{ij}$
5 **end**
6 **return** $\mathbf{E}$

To define a truthful *payment rule*, we guarantee that each user i is paid the *critical value* E_i, which is defined as follows with respect to the *critical bid* $B_{\bar{i}\bar{j}}$. If user i submits a compensation $C_{ij} > E_i$, it loses; otherwise, it wins. In Algorithm 2, we obtain the critical bid as follows. We find the solution $\mathbf{W}_{-i}$ of GRAN when user i is not participating in the auction (line 2). Then, we select the critical bid $B_{\bar{i}\bar{j}}$ as the last bid added to the solution set (line 3). Finally, in line 4, we define the critical value $E_i = R_{\bar{i}\bar{j}} \Delta P_{ij}$. In the following subsection, we will prove that this payment rule, paired with the winner selection algorithm, guarantees truthfulness of the GRAN mechanism.

B. GRAN Properties

To prove that the $GRAN$ mechanism is truthful, we follow the approach similar to [44]. More precisely, we first prove that the winner selection algorithm (Algorithm 1) is monotonic, and then that the payment rule (Algorithm 2) pays the critical value.

Definition 1 (Monotonicity)**.** *An algorithm is monotonic if, by substituting any winning bid $B_{ij} = (\Delta P_{ij}, \Delta T_{ij}, C_{ij})$ with $\tilde{B_{ij}} = (\Delta P_{ij}, \Delta T_{ij}, C_{ij} - \delta)$, $\tilde{B_{ij}}$ is selected as a winner.*

Theorem 4. *Algorithm 1 is monotonic.*

Proof. Suppose the bid B_{ij} wins in the q^{th} iteration. If we substitute B_{ij} with $\tilde{B_{ij}} = (\Delta P_{ij}, \Delta T_{ij}, C_{ij} - \delta)$, $\delta > 0$, and execute Algorithm 1 with such new input, $\tilde{B_{ij}}$ would appear in the ranking criterion $\mathbf{R}$ before the position of B_{ij} in the original execution. As a result, $\tilde{B_{ij}}$ would be selected on or before the q^{th} iteration. $\square$

Theorem 5. *Each winning bid is paid the critical value.*

Proof. Our goal is to prove that the payment rule we defined in Algorithm 2 pays the critical value, as defined in line 4. More specifically, paying the critical value is equivalent to proving that if user i submitted a compensation $C_{ij} > E_i$, then it will lose; otherwise (i.e., if user i submitted a compensation $C_{ij} \leq E_i$), then it will win.

Consider a winning bid $B_{ij} = (\Delta P_{ij}, \Delta T_{ij}, C_{ij})$ selected by Algorithm 1, and consider the critical bid $B_{\bar{i}\bar{j}}$ in line 3 of Algorithm 2, i.e., the last selected winning bid when B_{ij} is not participating in the auction.

[Case 1]: if $C_{ij} > E_i$, then B_{ij} is a losing bid. The inequality $C_{ij} > E_i$ can be rewritten as $C_{ij} > R_{\bar{i}\bar{j}} \Delta P_{ij}$. Dividing both members by ΔP_{ij} yields $R_{ij} > R_{\bar{i}\bar{j}}$. Because Algorithm 1 sorts values of ranking criterion ascendingly, R_{ij} would be placed after $R_{\bar{i}\bar{j}}$ in the list $\mathbf{R}$. Therefore, B_{ij} will be a losing bid.

[Case 2]:if $C_{ij} \leq E_i$, bid B_{ij} is a winning bid. Similarly to Case 1, $C_{ij} \leq E_i$ yields $R_{ij} \leq R_{\bar{i}\bar{j}}$, which makes B_{ij} a winning bid, since R_{ij} would be placed before $R_{\bar{i}\bar{j}}$ in $\mathbf{R}$. □

Theorem 6. *The GRAN mechanism is truthful.*

Proof. Following [44, Theorem 9.36], the proof of this theorem follows from Theorems 4 and 5 proved above. □

Theorem 7. *The GRAN mechanism holds the property of individual rationality.*

Proof. To prove individual rationality, we need to show that the GRAN payment rule defined in line 4 of Algorithm 2 is non-negative, i.e., $E_i = R_{\bar{i}\bar{j}} \Delta P_{ij} \geq 0$. This is straightforward given the ranking criterion R_{ij} and the power saving ΔP_{ij} are non-negative. □

Let us know analyze the computational complexity of the GRAN mechanism.

Theorem 8. *The time complexity of the GRAN mechanism is $O(N^2 M_{\max} \log(N M_{\max}))$, where $M_{max} = \max_{i=1,\ldots,N} |B_i|$ is the maximum number of bids submitted by a user.*

Proof. We analyze the time complexity of the winner selection (Algorithm 1) and the payment rule (Algorithm 2) separately.

[Algorithm 1]: In line 4 of Algorithm 1, we generate the list $\mathbf{R}$ and sort it in line 5. The list size $O(NM_{\max})$, where $M_{max} = \max_i |B_i|$. Thus the overall complexity is $O(NM_{\max} \log(NM_{\max}))$. The "while" loop in lines $6-12$ is executed at most N times, since each iteration selects a user and all other bids of that user are removed from $\mathbf{R}$. The cost of each iteration is dominated by the cost of removing bids for the selected user from $\mathbf{R}$ in line 11. By using a hash list to store the pointers to the bids, this operation can be done in $O(M_{\max})$ time, implying the while loop requires $O(NM_{\max})$ time. Therefore, the time complexity of Algorithm 1 is $O(NM_{\max} \log(NM_{\max}))$.

[Algorithm 2]: The "for" loop in line 1 makes at most N iterations, since the maximum number of winners is N. At each iteration, we execute Algorithm 1. Therefore, it requires $O(N^2 M_{\max} \log(NM_{\max}))$ time.

Overall, the time complexity of the GRAN mechanism is $O(N^2 M_{\max} \log(NM_{\max}))$, dominated by Algorithm 2. □

V. Power Saving Predictions

In order to effectively select the winners of the auction and meet the power cap constraint, it is necessary to know the power saving corresponding to each bid. Predicting the power consumption for a given thermostat setting is a complex task that depends on a plethora of parameters, such as weather, house size, solar gain, physical and chemical characteristics of the house materials, etc. [45]. It is even more challenging to predict the power saving resulting from a sudden and short-time change in the thermostat setting.

A. Background on Power Prediction

In the literature, there exist two different approaches to predict power consumption, namely the white box and black box approaches. In the *white box* approach, a physical model consisting of equations that formulate the physical and chemical characteristics of the house, layout, occupancy, and materials [26]. However, this approach is often impractical mainly due to two reasons: (i) most of these parameters are unknown in practice [45]; and (ii) a different model would be required for each house. In contrast, the *black box* approaches rely on the historical data of power consumption. The goal is to train a machine learning model capable of predicting the future time series of power consumption under the *current* environmental (e.g., weather information) and house conditions [46].

Most previous works adopting a black box approach focus on predicting the *steady-state* energy consumption of a given house at a specified thermostat setting [46], [47]. However, in our proposed work in this paper, we are interested in predicting the power saving during a *transient state*, i.e., after a sudden and short-term change of the thermostat setting. In most circumstances, the peak load period is not long enough to allow the power consumption to reach the steady state [48], making the prediction problem extremely challenging. To the best of our knowledge, no other work strictly focuses on the *transient state* predictions of residential power consumption.

B. The PSP Algorithm

The proposed Power Saving Prediction (PSP) algorithm is based on a regression technique that predicts cumulative power saving resulting from a thermostat change. We assume that the SEMS of each user keeps track of thermostat setting adjustments that occur over time during non-peak load periods, and the resulting power consumption. These changes may be due to sporadic manual adjustments, or automatic event-based adjustments supported by modern thermostats [49]. Note that the user may or may not be at home when such changes take place. For each of these adjustments, the SEMS records the power saving at different time scales (e.g., multiple of

15 minutes), representing the potential duration of a peak load. A different model is trained for each of these durations. Training is performed with a set of features easily available to the SEMS. Therefore, potentially useful but hard to obtain information, such as the window *U-factor* [45], is purposely omitted. Specifically, the PSP algorithm is based on the following features:

- *Weather information*: outside temperature, wind speed, humidity at the beginning of the peak load period;
- *House information*: default thermostat set point, new thermostat set point, inside temperature; and
- *Time*: hour of the day.[3]

The above features could be used to train several types of machine learning models. However, since the data collected are from individual homes, and thus *limited*, models that require large training sets (e.g., deep neural networks) would not be practical [50]. As a result, the PSP algorithm exploits a *regression* technique that allows us to learn the correlation between the features given as input, and the power saved during the peak load period, with limited training data and at a very fast rate, as shown in the experimental study (Section VII). We evaluated the performance of several regression algorithms (Artificial Neural Networks, Random Forest, ElasticNet, Support Vector Machines, Nearest Neighbors Regression, and Naive Bayes), and tested the values of various parameters with a grid search. We found *Random Forest regression* (RF) [51] to provide the best performance. In our experiments, we set the parameters as follows: (i) the criterion to measure split quality to the *mean squared error* (MSE), (ii) the maximum depth of the tree to 1000, (iii) the number of estimators to 150, (iv) the minimum number of samples needed to split a node to 2, (v) the maximum number of features while deciding the best split equal to the total number of features (7 in this case), and (vi) the minimum number of samples required to be at a leaf node to 1.

VI. Online Survey

We conduct an online survey involving 200 subjects to assess the bidding behavior and willingness to participate in the proposed Incentive-Based Power Conservation (IBPC) program. The study was approved by the Institutional Review Board at the University of Missouri System (#IRB-2025242-ST). This section discusses the survey and the results.

A. Overview of the Survey

The participants are recruited using Amazon Mechanical Turk are pre-screened to include only Florida residents who use an adjustable thermostat in their homes, receive an energy bill each month based on the energy usage, and review their bill every month or most months. We focused on a specific geographic area for a more uniform perception of the system. Eligible participants were informed of their rights and compensation before completing the survey. In our online study, the mean time to complete the survey was just under 10 minutes and the participants were compensated for $1.75. This translates into the rate of $10.60/hr which was above the federal minimum hourly wage of $7.75/hr at the time the study was conducted, and above the top 4% earning rate of $7.50/hr for M-Turk workers [52].

The survey began by asking the participants to first indicate their typical thermostat setting on a hot summer day. They then read a brief description of peak load and power conservation to ensure that each participant had an understanding of the context. This was followed by a description of the proposed system that would ostensibly help reduce the energy consumption during peak times by compensating the customers via an automated system to temporarily adjust their thermostat setting.

The participants were asked to imagine that they were participating in such a program and setting up their smart thermostat temperature. This was completed in two steps. First, the participants were reminded of their response for their typical thermostat setting on a hot summer day. Then, from a list of options, they were asked to select the highest thermostat setting to which they would be willing to occasionally adjust for a maximum of one hour per day. This list of options was customized for each participant to include 8 degrees of change above their typical setting (for example, if their typical setting was 70^{o}F, their range of options was 71^{o}F-78^{o}F). Next, for each thermostat temperature setting within the selected range, the participants were asked to use a slider to indicate the minimum compensation they would like to be paid in order to allow the thermostat adjusted to that setting. For uniformity of the results, we asked everyone to imagine the following scenario when they provided their bids:

"Imagine it is daytime on a hot summer day, you are at home, and you are doing low to moderate effort activities (for example, sleeping, sitting, or light chores) and imagine that the maximum duration of the change would be 1 hour, at which point the thermostat then returns to the previous setting."

The slider range was $0.00 to $5.00 and could be moved in increments of .01. This dollar range was proposed in [26]. The participants were told a compensation of $0.00 implied they would make the adjustment for free. Finally, they were asked whether they would participate in such a system if it existed. The outcome of the survey is reported in the following.

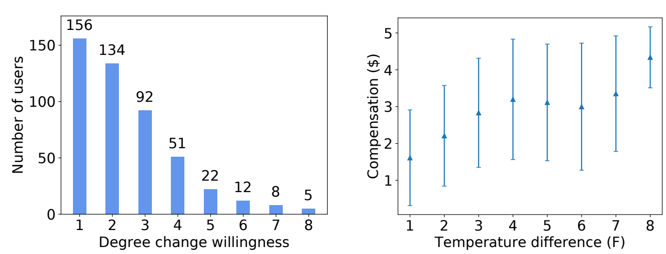

(a) User willingness (b) Average compensation

Fig. 2: Summary of online survey results per degree change

[3]The *season* could be another important factor. Different models could be trained for different seasons to take into account this aspect.

B. Survey Results

The results do not include 44 users who failed to correctly answer the attention check questions. Overall, more than 79% users answered that they would be willing to use this system in their homes. Fig. 2a shows, for a given change in the temperature value (measured in degrees) on the X-axis, the number of survey participants who agreed to change their thermostat to at most that value. The results are clearly non-linear: most users are comfortable with small temperature changes, and become less comfortable as the change increases.

Fig. 2b shows the mean and standard deviation of user compensations. The plot shows a monotonic trend, suggesting that higher temperature changes require higher monetary incentives. Nevertheless, the users show significant heterogeneity in the requested amount for a given temperature change. This, coupled with the non-linear willingness to adjust the temperature setting, results in an interesting and non-trivial optimization scenario for our proposed approach.

Overall, the survey results support the feasibility of the proposed IBPC auction framework. We use these results to model a realistic user behavior in engaging with the power conservation framework. Specifically, we follow the survey results to determine how many degrees a user is willing to change and the corresponding compensation.

VII. Performance Evaluation

This section presents the experimental setup and the performance of our method in comparison with other approaches.

A. Experimental Setup

We adopt *EnergyPlus* and integrate it with Python scripts implementing our solutions as well as other approaches used for comparison. EnergyPlus is a simulator funded by the U.S. Department of energy, and tested according to ASHRAE Standard 140 methodology [27] which makes it the gold standard of power data simulation. It is a high-fidelity tool that allows for modeling of very low-level parameters of residential buildings, with the goal of producing extremely accurate power consumption data [27].

In order to consider a variety of realistic houses, we employed the *EnergyPlus residential prototype building models* provided by the U.S. department of Energy in collaboration with the Pacific Northwest National Laboratory [53]. The models have 4 foundation types (slab, crawlspace, heated basement, and unheated basement) and 2 cooling system types (central air conditioning cooling and heat pump cooling). The combination of these characteristics gives us a total of 8 considerably different houses and therefore different utility loads. Furthermore, EnergyPlus allows low-level control of many house details. Hence, we exploit this functionality by varying the *window U-factor*, a parameter that greatly impacts the thermal resistance of a residential building. For each one of the 8 models previously mentioned, we generate 5 additional models by changing the *U-factor* within $[2, 4] W/(m^2K)$ range [54]. As a result, we obtain a total of 40 heterogeneous models that capture a wide spectrum of thermal resistance of a house. We used each model twice for a total of 80 houses. Note that, further increase in the number of houses by using additional copies of these models would result in more homogeneous, and thus less realistic, scenarios. Since the total power consumption P_T and the power cap P_C scale linearly with the number of houses, we expect the trends observed in our results to hold in larger deployments of the system.

B. Performance of the PSP Algorithm

In this section, we study the performance of the power saving prediction (PSP) algorithm.

1) Comparison Approach Sha-SVR: We compare PSP to a recent state-of-the-art approach for power prediction proposed in [45], which we refer to as Sha-SVR. We select this approach because it is designed to work in a specific building setting where, similar to our framework, data are limited and the features need to be easily available.

The authors of Sha-SVR adopt the Support Vector Regression (SVR) as the prediction model. To select the features set, the Pearson correlation coefficients between a vast array of meteorological parameters and the HVAC power data are analyzed. This allows to considerably reduce the size of the feature set. The authors conclude that the *dry-bulb temperature* has the highest impact, with a correlation coefficient of 0.91 on a summer day, which is the season considered in our experiments. Besides the dry-bulb temperature, the authors also consider the *balance point temperature*, T_c. The dry-bulb temperature is transformed into Cooling Degree-Day CDD), a simple but effective method for building energy analysis [55]. Here $CDD = \max\{(T_{\max} - T_{\min})/2 - T_c, 0\}$, where $T_{\max}$ and $T_{\min}$ are the maximum and minimum hourly temperature in a day, and $T_c = 59°F(15°C)$ is the standard temperature value they intuitively set for their experiments. Finally, the authors add two features to describe the behavioral pattern of users, by adding the month type and the day type.

Note that Sha-SVR has been designed for the prediction during a steady-state, rather than transient-state. Hence, we adapt the algorithm as follows. In order to calculate the power saving resulting from a transition from set point T_{old} to a set point T_{new}, we first add the temperature set point as a new feature. Then, we use Sha-SVR to predict the steady-state power consumption $P_{(T_{old})}$ and $P_{(T_{new})}$ separately, and calculate the power saving $\Delta P = P_{(T_{old})} - P_{(T_{new})}$. For more details on Sha-SVR, refer to [45].

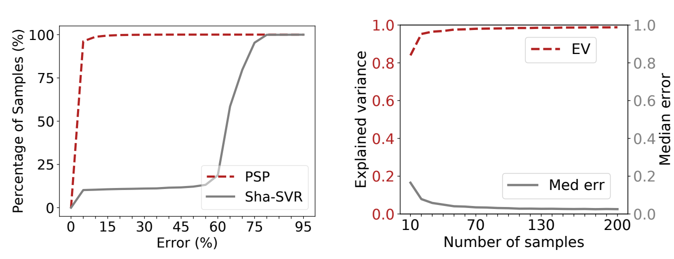

(a) Error distribution. (b) Learning rate of PSP.

Fig. 3: Performance of the prediction models

2) Results: To train and compare the PSP algorithm with Sha-SVR, we use the weather information from Miami provided with the EnergyPlus residential prototype building models [53]. Miami has been chosen since it experiences very hot summer days, and it is the area where the online survey was conducted. Because the focus is on the hottest days and hours, we consider a time range from July to September, between 1 PM and 6 PM. For each of the houses, we consider 8 thermostat change options, each representing a 1°F (approx. 0.55°C) degree difference and takes place for 1 hour. We use EnergyPlus to collect the resulting power consumption data and pair it with the features required by each algorithm. The data is then shuffled before forming the testing and training set, with 50% samples each.

In the experiments of Fig. 3, we analyze the performance of PSP in terms of the error percentage, median error, and explained variance[4]. Specifically, Fig. 3a shows the fraction of testing samples that are predicted within a certain error percentage. Observe that Sha-SVR achieves poor performance. This is due to its focus on the steady-state prediction, which prevents this approach from considering the dynamics that occur after a sudden change of the thermostat for a short period of time. As a numerical example, more than 85% of Sha-SVR's predictions incur higher than 60% error. Conversely, PSP achieves a very high accuracy, with more than 85% of our testing samples within 5% error.

Next, we study the learning rate of PSP. This is particularly important since the training data is generated by each home individually, and hence it is limited. For this purpose, we adopt a widely used metric called *Explained Variance* (EV) regression score. EV is a statistical measure used to evaluate the quality of a regression prediction, based on the variance of the real value and the error [56]. $EV \in [0, 1]$ and a higher value (i.e., close to 1) represents more accurate predictions.

Our goal is to analyze the value of EV and the median prediction error, by progressively increasing the size of the testing set. We performed tests individually for each home and averaged the results. The testing samples for a home are randomly selected. As Fig. 3b shows, 20-30 samples are sufficient to obtain a very high EV and very low median error. These results show the ability of PSP to provide accurate prediction requiring only few samples for training. Recall that these samples do not need to be collected during the peak loads, but can instead be gathered by the SEMS, during manual or automatic adjustments that are possible with modern thermostats [49].

C. Power Conservation

1) Comparison Approach MEDR: A recent paper [25] proposes a truthful auction-based IBPC approach for peak load reduction in data centers called Mechanism for Emergency Demand Response (MEDR). Similar to our scenario, in the event of a peak load, N tenants are required to reduce their power consumption below a power cap. Each tenant may submit *one bid* consisting of a power reduction and monetary compensation. This paper defines an NP-Hard problem to select winners of the auction that, similarly to POCO, aims at minimizing the overall cost. Since the users in our settings may submit multiple bids, for each user i we randomly pick a bid in the set $\mathbf{B_i}$. We implement the NP-Hard optimization problem to select winners as well as the truthful payment rule defined in [25]. This implementation gives an advantage to MEDR, since the solution of the NP-Hard problem guarantees the minimization of the objective function, at the cost of a higher complexity. For more details, refer to [25].

[4]These are typical metrics for regression algorithms, comparable to the accuracy, F-score, etc., for classification algorithms.

2) Results: In the following, we compare the performance of GRAN, MEDR, and the optimal solution of POCO, referred to as OPT, obtained with the Gurobi optimizer [57]. Experiments are run during the hot summer days in July and August 2009 with an average temperature of 89.06°F (31.7°C). Similar to the previous works, we consider a peak load period of 1 hour [25], [26]. We ran experiments for different peak lengths observing similar trends. Moreover, in all experiments, the user bidding behavior is selected from the results of the online survey. Finally, we provide each approach with two predictions of power saving resulting from thermostat adjustments, namely the case of perfect prediction (perfect knowledge), and the case of the energy prediction provided by PSP. The results are averaged over several runs to obtain reliable results.

We explore two experimental scenarios. One where we vary the percentage of participants, and another where we vary the percentage of reduction required by the utility company. In the first experimental scenario, we increase the percentage of users participating in the auction from 40% to 100% out of the total $N = 80$ users. In this setting, the total consumption calculated with EnergyPlus is $P_T = 261.95kW$. We set $\alpha = 0.95$, thus the power cap is $P_C = 0.95 \times P_T = 248.859kW$. Note that, the non-participating users contribute to P_T, but refuse to participate in the IBPC program.

Figures 4a and 4b respectively show the value of the POCO objective function and the payment, by increasing the number of participants under perfect predictions. Both objective values and payments decrease as the user participation increases, for all approaches. This is due to the availability of more bids with higher user participation, which enables to find

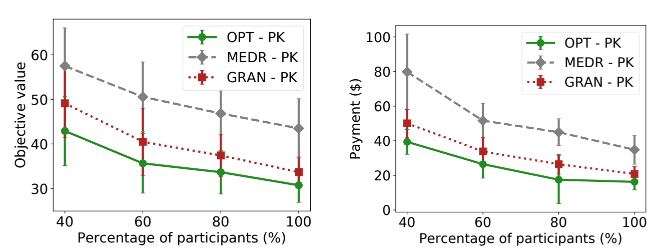

(a) Objective value (b) Payment rule

Fig. 4: Percentage of participants with perfect knowledge

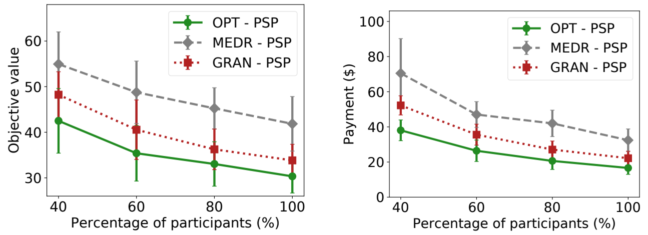

(a) Objective value

(b) Payment rule

Fig. 5: Varying percentage of participants with PSP predictions

better solutions. MEDR suffers from its inability to optimize over multiple bids per user. Conversely, GRAN significantly outperforms MEDR across all values of user participation by exploiting the ranking criterion while guaranteeing truthfulness. This is remarkable considering that MEDR solves an NP-Hard problem for each auction. Overall, GRAN achieves a value of the objective function which is on average 68.3% closer to the optimum than MEDR, and incurs a payment which is 71.65% lower.

Figures 5a and 5b show the value of the objective function and payment when the power saving is predicted by PSP. Also in this case, GRAN significantly outperforms MEDR. The high prediction accuracy of PSP allows to achieve results comparable to the case of perfect knowledge. It is worth noting that, even with a participation well below the 79% expressed by the users of the survey, our system is able to achieve the desired power conservation.

In the second experimental scenario, we study the impact of different settings of the power cap on the performance of the considered approaches. For this purpose, we increase the power reduction α from 3% to 9% of P_T, and fix the percentage of participating users to 60%. We only show results of power saving predicted by PSP; we observe similar trends with perfect knowledge. The results of the objective function and payment are shown in Figures 6a and 6b. We observe an increase in the objective function and payments for all approaches. Intuitively, increasing the power reduction requires more winners to be selected, with higher temperature changes and compensation as well. Moreover, in this case, GRAN shows superior performance than MEDR, being able to successfully exploit the available bids to find better solutions. Overall, GRAN is 58.75% closer to the optimum solution on an average and achieves payments that are 62.1% lower.

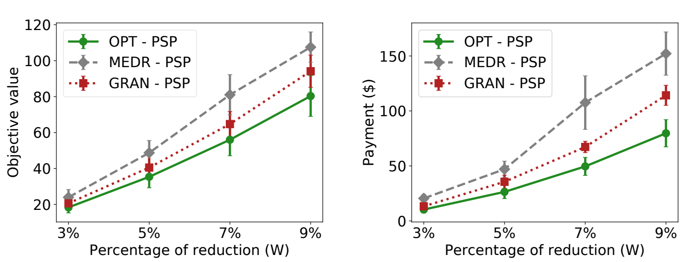

(a) Objective value

(b) Payment rule

Fig. 6: Varying percentage of reduction with PSP predictions

VIII. Discussion

The work presented in this paper is the first holistic approach to comprehensively considering 1) explicit power dynamics with realistic power data, along with 2) the corresponding machine learning-based power saving predictions due to thermostat changes, and 3) modeling of realistic user behavior by means of an online survey. Results of the survey show that 79% of the users would be interested in using such system and that the users have highly heterogeneous and non-linear bidding behavior. This percentage highly exceeds the findings of our scalability experiments, which demonstrate that 40% of user participation is able to fulfill the requested power cap, hence proving significant opportunities for our approach to successfully accomplish power conservation.

The proposed approach has the following limitations that will be addressed in our future work.

Real implementation requirements: Nowadays, the utility companies are able to monitor the residential power consumption at a fine-grain resolution, owing to the wide deployment of the Advanced Metering Infrastructure (AMI) [58], [59]. Therefore, a real life implementation of the proposed approach would additionally require each user home to be equipped with an Internet-connected Nest-like thermostat that is already capable of learning user energy patterns [18]. This thermostat needs to be equipped with additional machine learning capabilities for power saving predictions. The user should also have a smart-phone app to interact with the auction system and submit bids. The app should interface with the utility company and the thermostat to adjust the temperature setting according to the outcome of the auction.

Robustness: Robustness is an important criterion for power conservation approaches, since high demand may generate large-scale blackouts. For this purpose, the system operator sets the power cap constraint to a comfortable level for the system to operate. However, there are several uncertainties that make it challenging to satisfy such a constraint. First, power saving predictions need to be accurate. In this paper, we showed that our PSP algorithm for power saving predictions needs few samples to provide accurate results, with a median error below 10% after only 20 samples of training. Second, the users may not be willing to participate in the auction and consequently adjust their thermostat settings. We show that, even when only 40% of the users are willing to participate in the auction, we are able to meet the power cap constraint. This is particularly encouraging, since 79% of the participants in the survey expressed willingness to engage in such a system. Note that, the power system operator may add an additional layer of safety by further increasing the power cap to provide additional wiggle room to account for the aforementioned uncertainties.

Bidding competitively: Many participants do not intuitively know what a fair or competitive bid amount is. Nevertheless, each participant has a "floor" value, referred to as the true valuation in auction terms, that they are willing to go for in order to win the auction. The truthfulness of the auction helps the user converge the bids towards this value, as untruthful bids will result in reduced revenue. As a consequence, the user "learns" over time to bid their true valuation.

IX. Conclusions and Future Work

In this work, we proposed a comprehensive framework for HVAC-based power conservation. We developed a reverse auction-based approach according to which the users submit bids requesting a monetary compensation to adjust their thermostat settings. We formulated an optimization problem to select the set of winners and a payment rule that provides truthfulness and individual rationality. Proving that the winner selection problem is NP-Hard, we formulated an efficient heuristic which guarantees the same formal properties. Furthermore, we developed a machine learning module to predict power savings resulting from the thermostat adjustments of individual homes. We also performed an online survey to study user willingness to engage with such a system and their bidding behavior. Finally, we compared our approach with state-of-the-art solutions using the high-fidelity simulator EnergyPlus. Experimental results show that our approach outperforms prior solutions with near-optimum results; and our machine learning prediction algorithm provides accurate predictions with minimum training.

There are several future research directions worth investigating. As an example, energy-hungry appliances (e.g., electric vehicles and water heaters) could be included in the power conservation approach. However, such appliances will require machine learning algorithms in order to learn the dynamic impact on the power savings, and surveys to study the user behavior and perceived comfort. Although this paper focused only on the residential settings, the system operators may also benefit from curtailing the industrial energy consumption. However, industrial settings require ad-hoc solutions that take into account potentially very different business needs. Finally, while this work focuses on power conservation during hot summer days, it is worth investigating how a similar approach would perform in colder climates. This may require to tackle additional challenges. For instance, certain heating systems use back-up heating sources from natural gas when the thermostat is turned off, which could significantly impact the power grid, as well as the power saving predictions.

Acknowledgments

This work is partially supported by the National Institute for Food and Agriculture (NIFA) under the grant 2017-67008-26145; the NSF grants EPCN-1936131, CPS-1545037, and CNS-2008878; and the NSF CAREER grant CPS-1943035.

References

[1] Y. M. Lee, L. An, F. Liu, R. Horesh, Y. T. Chae, and R. Zhang, "Applying science and mathematics to big data for smarter buildings," *Annals of the NEW YORK Academy of Sciences*, vol. 1295, no. 1, pp. 18–25, 2013.

[2] Masson-Delmotte, V., P. Zhai, A. Pirani and others, "IPCC, 2021: Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change," Tech. Rep.

[3] A. Goard, ""Texas' power grid set a new winter peak demand record Sunday evening"." [Online]. Available: https://www.kxan.com/news/texas/texas-power-grid-set-a-new-winter-peak-demand-record-sunday-evening/

[4] Associated Press, Staff (ABC10). "California heat, poor planning led to August power outages". [Online]. Available: https://www.abc10.com/article/news/local/california/august-power-outages-report/103-267eb366-f08e-47ad-a387-e37d37910597

[5] BBC News, ""France wildfire: Thousands evacuated as blaze rages near Riviera"," https://www.bbc.com/news/world-europe-58243066.

[6] M. Ali, A. Safdarian, and M. Lehtonen, "Demand response potential of residential hvac loads considering users preferences," in *IEEE PES innovative smart grid technologies, Europe*. IEEE, 2014, pp. 1–6.

[7] V. Cavaliere. "Why some Texas residents are ending up with $\$5,000$ electric bills after the winter storms". [Online]. Available: https://www.businessinsider.com/why-texas-residents-hit-with-soaring-electric-bills-winter-storms-2021-2

[8] Q. Qdr, "Benefits of demand response in electricity markets and recommendations for achieving them," *US Dept. Energy, Washington, DC, USA, Tech. Rep*, vol. 5, 2006.

[9] A. Kenward and U. Raja, "Blackout: Extreme weather climate change and power outages," *Climate central*, vol. 10, pp. 1–23, 2014.

[10] G. Pulazza, N. Zhang, C. Kang, and C. A. Nucci, "Transmission planning with battery-based energy storage transportation for power systems with high penetration of renewable energy," *IEEE Transactions on Power Systems*, 2021.

[11] E. Douglas. "Texas was "seconds and minutes" away from catastrophic months long blackouts, officials say". [Online]. Available: https://www.texastribune.org/2021/02/18/texas-power-outages-ercot/

[12] The Electric Reliability Council of Texas (ERCOT), "Grid operator requests energy conservation for system reliability," http://www.ercot.com/news/releases/show/225151, Accessed: August 2021.

[13] ——, "Tight grid conditions expected due to high number of forced generation outages," Accessed: August 2021. [Online]. Available: http://www.ercot.com/news/releases/show/233037

[14] M. Kolnegari, E. Moghimi, A. A. Basiri, A. T. Qashqaei, and M. Hazrati, "A new committee to address the threats of power grids to birds in Iran," *Biodiversity*, vol. 20, no. 4, pp. 161–164, 2019.

[15] T. Schoeman and M. Saunders, "The impact of power outages on small businesses in the City of Johannesburg," in *10th International Conference on Education, Business, Humanities and Social Sciences Studies*, 2018.

[16] U.S. Energy Information Administration, "Use of energy explained. energy use in homes," 2020. [Online]. Available: https://www.eia.gov/energyexplained/use-of-energy/homes.php

[17] Landis+Gyr, "Landis+gyr manage energy better," 2021. [Online]. Available: https://www.landisgyr.com/product/maxsys-e850-elite-advanced-metering/

[18] google.com, "Nest learning thermostat," 2021. [Online]. Available: https://store.google.com/product/nest_learning_thermostat_3rd_gen?hl=en-US

[19] sense.com, "Welcome to sense. take command of your energy use with total home monitoring," 2021. [Online]. Available: https://sense.com

[20] C. Mahapatra, A. K. Moharana, and V. Leung, "Energy management in smart cities based on internet of things: Peak demand reduction and energy savings," *Sensors*, vol. 17, no. 12, p. 2812, 2017.

[21] California Energy Commission. "CAISO, CPUC, CEC issue final report on causes of august 2020 rotating outages". [Online]. Available: https://www.energy.ca.gov/news/2021-01/caiso-cpuc-cec-issue-final-report-causes-august-2020-rotating-outages

[22] California Independent System Operator, California Public Utilities Commission, California Energy Commission, "Root Cause Analysys - Mid-August 2020 Extreme Heat Wave," Tech. Rep., 2021.

[23] A. C. Chapman and G. Verbič, "An iterative on-line auction mechanism for aggregated demand-side participation," *IEEE Transactions on Smart Grid*, vol. 8, no. 1, pp. 158–168, 2015.

[24] A. R. Khamesi and S. Silvestri, "Reverse auction-based demand response program: A truthful mutually beneficial mechanism," in *IEEE International Conference on Mobile Ad Hoc and Sensor Systems (MASS)*. IEEE, 2020, pp. 1–10.

[25] J. Chen, D. Ye, S. Ji, Q. He, Y. Xiang, and Z. Liu, "A truthful fptas mechanism for emergency demand response in colocation data

centers," in *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*. IEEE, 2019, pp. 2557–2565.
[26] Q. Shi, C.-F. Chen, A. Mammoli, and F. Li, "Estimating the profile of incentive-based demand response (IBDR) by integrating technical models and social-behavioral factors," *IEEE Transactions on Smart Grid*, vol. 11, no. 1, pp. 171–183, 2020.
[27] U.S. Department of Energy, "Energyplus." [Online]. Available: https://www.energy.gov/eere/buildings/downloads/energyplus-0
[28] ——. "ASHRAE Standard 140 Maintenance and Development". [Online]. Available: https://www.energy.gov/eere/buildings/ashrae-standard-140-maintenance-and-development
[29] M. Uddin, M. F. Romlie, M. F. Abdullah *et al.*, "A review on peak load shaving strategies," *Renewable and Sustainable Energy Reviews*, vol. 82, pp. 3323–3332, 2018.
[30] Q. Zhang and J. Li, "Demand response in electricity markets: A review," in *IEEE international conference on the European Energy market*. IEEE, 2012, pp. 1–8.
[31] D. P. Chassin and D. Rondeau, "Aggregate modeling of fast-acting demand response and control under real-time pricing," *Applied Energy*, vol. 181, pp. 288–298, 2016.
[32] J.-H. Kim and A. Shcherbakova, "Common failures of demand response," *Energy*, vol. 36, no. 2, pp. 873–880, 2011.
[33] A. Asadinejad, M. G. Varzaneh, K. Tomsovic, C.-f. Chen, and R. Sawhney, "Residential customers elasticity estimation and clustering based on their contribution at incentive based demand response," in *2016 IEEE Power and Energy Society General Meeting (PESGM)*. IEEE, 2016, pp. 1–5.
[34] T. Ericson, "Direct load control of residential water heaters," *Energy Policy*, vol. 37, no. 9, pp. 3502–3512, 2009.
[35] J. Wang, S. Kennedy, and J. Kirtley, "A new wholesale bidding mechanism for enhanced demand response in smart grids," in *IEEE Innovative Smart Grid Technologies (ISGT)*. IEEE, 2010, pp. 1–8.
[36] S. Ciavarella, J.-Y. Joo, and S. Silvestri, "Managing contingencies in smart grids via the internet of things," *IEEE Transactions on Smart Grid*, vol. 7, no. 4, pp. 2134–2141, 2016.
[37] M. Shafie-Khah and P. Siano, "A stochastic home energy management system considering satisfaction cost and response fatigue," *IEEE Transactions on Industrial Informatics*, vol. 14, no. 2, pp. 629–638, 2017.
[38] R. Zhou, Z. Li, and C. Wu, "An online procurement auction for power demand response in storage-assisted smart grids," in *IEEE Conference on Computer Communications (INFOCOM)*. IEEE, 2015, pp. 2641–2649.
[39] S. Comello and S. Reichelstein, "The emergence of cost effective battery storage," *Nature communications*, vol. 10, no. 1, pp. 1–9, 2019.
[40] A. R. Khamesi, S. Silvestri, D. A. Baker, and A. D. Paola, "Perceived-value-driven optimization of energy consumption in smart homes," *ACM Transactions on Internet of Things*, vol. 1, no. 2, pp. 1–26, 2020.
[41] V. Dolce, C. Jackson, S. Silvestri, D. Baker, and A. De Paola, "Social-behavioral aware optimization of energy consumption in smart homes," in *2018 14th International Conference on Distributed Computing in Sensor Systems (DCOSS)*. IEEE, 2018, pp. 163–172.
[42] L. Zhang, S. Ren, C. Wu, and Z. Li, "A truthful incentive mechanism for emergency demand response in colocation data centers," in *2015 IEEE Conference on Computer Communications (INFOCOM)*. IEEE, 2015, pp. 2632–2640.
[43] J. Csirik, "Heuristics for the 0-1 min-knapsack problem," *Acta Cybernetica*, vol. 10, no. 1-2, pp. 15–20, 1991.
[44] L. Blumrosen and N. Nisan, "Algorithmic game theory," *Introduction to Mechanism Design, Cambridge University Press, New York, USA*, 2007.
[45] H. Sha, P. Xu, C. Hu, Z. Li, Y. Chen, and Z. Chen, "A simplified hvac energy prediction method based on degree-day," *Sustainable Cities and Society*, vol. 51, p. 101698, 2019.
[46] M. Cai, M. Pipattanasomporn, and S. Rahman, "Day-ahead building-level load forecasts using deep learning vs. traditional time-series techniques," *Applied Energy*, vol. 236, pp. 1078–1088, 2019.
[47] Y. Huang, N. Wang, and W. o. Gao, "Loadcnn: A low training cost deep learning model for day-ahead individual residential load forecasting," *arXiv preprint arXiv:1908.00298*, 2019.
[48] P.-O. Pineau, "Peak load problem, deregulation and reliability pricing," in *Deregulation of Electric Utilities*. Springer, 1998, pp. 275–296.
[49] T. Peffer, M. Pritoni, A. Meier, C. Aragon, and D. Perry, "How people use thermostats in homes: A review," *Building and Environment*, vol. 46, no. 12, pp. 2529–2541, 2011.
[50] J. K. Basu, D. Bhattacharyya, and T.-h. Kim, "Use of artificial neural network in pattern recognition," *International journal of software engineering and its applications*, vol. 4, no. 2, 2010.
[51] V. Svetnik, A. Liaw, C. Tong, J. C. Culberson, R. P. Sheridan, and B. P. Feuston, "Random forest: a classification and regression tool for compound classification and qsar modeling," *Journal of chemical information and computer sciences*, vol. 43, no. 6, pp. 1947–1958, 2003.
[52] K. Hara, A. Adams, K. Milland, S. Savage, C. Callison-Burch, and J. P. Bigham, "A data-driven analysis of workers' earnings on amazon mechanical turk," in *Proceedings of the 2018 CHI conference on human factors in computing systems*, 2018, pp. 1–14.
[53] U. D. of Energy. (2020) Building energy codes program. [Online]. Available: https://www.energycodes.gov/development/residential/iecc_models
[54] International Code Council, "2015 international energy conservation code - chapter 4 [re] residential energy efficiency," 2016. [Online]. Available: https://codes.iccsafe.org/content/IECC2015/chapter-4-re-residential-energy-efficiencyIECC2015_Pt02_Ch04_SecR402
[55] Handbook, ASHRAE Fundamentals and others, "American society of heating, refrigerating and air-conditioning engineers," *Inc.: Atlanta, GA, USA*, 2009.
[56] A. Gelman and I. Pardoe, "Bayesian measures of explained variance and pooling in multilevel (hierarchical) models," *Technometrics*, vol. 48, no. 2, pp. 241–251, 2006.
[57] L. Gurobi Optimization, "Gurobi optimizer reference manual," 2020. [Online]. Available: http://www.gurobi.com
[58] S. Bhattacharjee, A. Thakur, S. Silvestri, and S. K. Das, "Statistical security incident forensics against data falsification in smart grid advanced metering infrastructure," in *Proceedings of the Seventh ACM on Conference on Data and Application Security and Privacy*, 2017, pp. 35–45.
[59] S. Bhattacharjee, V. P. K. Madhavarapu, S. Silvestri, and S. K. Das, "Attack context embedded data driven trust diagnostics in smart metering infrastructure," *ACM Transactions on Privacy and Security (TOPS)*, vol. 24, no. 2, pp. 1–36, 2021.