



A Human-Centered Power Conservation Framework Based on Reverse Auction Theory and Machine Learning

ENRICO CASELLA, University of Wisconsin-Madison, Madison, Wisconsin, USA

SIMONE SILVESTRI, University of Kentucky, Lexington, Kentucky, USA

D. A. BAKER and SAJAL K. DAS, Missouri University of Science & Technology, Rolla, Missouri, USA

Extreme outside temperatures resulting from heat waves, winter storms, and similar weather-related events trigger the Heating Ventilation and Air Conditioning (HVAC) systems, resulting in challenging, and potentially catastrophic, peak loads. As a consequence, such extreme outside temperatures put a strain on power grids and may thus lead to blackouts. To avoid the financial and personal repercussions of peak loads, demand response and power conservation represent promising solutions. Despite numerous efforts, it has been shown that the current state-of-the-art fails to consider (1) the complexity of human behavior when interacting with power conservation systems and (2) realistic home-level power dynamics. As a consequence, this leads to approaches that are (1) ineffective due to poor long-term user engagement and (2) too abstract to be used in real-world settings. In this article, we propose an auction theory-based power conservation framework for HVAC designed to address such individual human component through a three-fold approach: *personalized preferences* of power conservation, *models of realistic user behavior*, and *realistic home-level power dynamics*. In our framework, the System Operator sends Load Serving Entities (LSEs) the required power saving to tackle peak loads at the residential distribution feeder. Each LSE then prompts its users to provide *bids*, i.e., *personalized preferences* of thermostat temperature adjustments, along with corresponding financial compensations. We employ *models of realistic user behavior* by means of online surveys to gather user bids and evaluate user interaction with such system. *Realistic home-level power dynamics* are implemented by our machine learning-based Power Saving Predictions (PSP) algorithm, calculating the individual power savings in each user's home resulting from such bids. A machine learning-based PSPs algorithm is executed by the users' Smart Energy Management System (SEMS). PSP translates temperature adjustments into the corresponding power savings. Then, the SEMS sends bids back to the LSE, which selects the auction winners through an optimization problem called POWER Conservation Optimization (POCO). We prove that POCO is NP-hard, and thus provide two approaches to solve this problem. One approach is an optimal pseudo-polynomial algorithm called DYnamic programming Power Saving (DYPS), while the second is a heuristic polynomial time algorithm called Greedy Ranking AllocationN (GRAN). EnergyPlus, the high-fidelity and gold-standard energy simulator funded by the U.S. Department of Energy, was used to validate our experiments, as well as to collect data to train PSP. We further evaluate the results of the auctions across several scenarios, showing that, as expected, DYPS finds the optimal solution, while GRAN outperforms recent state-of-the-art approaches.

This work was partially supported by the NSF grants EPCN-1936131, CPS-1545037, and CNS-2008878; and the NSF CAREER grant CPS-1943035.

Authors' Contact Information: Enrico Casella (Corresponding author), Department of Animal & Dairy Sciences, University of Wisconsin-Madison, Madison, Wisconsin, USA; e-mail: ecasella@wisc.edu; Simone Silvestri, Department of Computer Science, University of Kentucky, Lexington, Kentucky, USA; e-mail: simone.silvestri@uky.edu; D. A. Baker, Department of Psychological Science, Missouri University of Science & Technology, Rolla, Missouri, USA; e-mails: denise.baker@cdc.hhs.gov, bakerden@mst.edu; Sajal K. Das, Department of Psychological Science, Missouri University of Science & Technology, Rolla, Missouri, USA; e-mail: sdas@mst.edu.

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 2378-9638/2024/7-ART28

<https://doi.org/10.1145/3656348>

CCS Concepts: • **Human-centered computing** → *Ambient intelligence*; **Interaction design theory, concepts and paradigms**;

Additional Key Words and Phrases: HVAC power conservation, machine learning power saving predictions, reverse auctions, cyber-physical pervasive computing, human-centered cyber-physical systems

ACM Reference format:

Enrico Casella, Simone Silvestri, D. A. Baker, and Sajal K. Das. 2024. A Human-Centered Power Conservation Framework Based on Reverse Auction Theory and Machine Learning. *ACM Trans. Cyber-Phys. Syst.* 8, 3, Article 28 (July 2024), 26 pages.
<https://doi.org/10.1145/3656348>

1 Introduction

Residential power consumption has increased significantly in the last years, due to factors such as growing urbanization [37], and it is expected to be rising further as a consequence of economic development and limited progress toward energy efficiency [52]. Besides, more and more frequent weather events, such as winter storms [24] and heat waves [5], lead **Heating Ventilation and Air Conditioning (HVAC)** systems to cause *peak loads*, due to the extreme outside temperatures [2]. These loads saturate the power grid capacity and are economically costly, due to the exponential growth of the price–demand ratio [48], and may lead to even greater financial consequences when they cause blackouts. As a matter of fact, in 2021, Texas witnessed a historical winter peak demand record of 69,150 MW [24], leading to customer bills of \$5,000 and a wholesale energy increase of 17,900% [14]. In addition, blackouts resulting from peak loads not only bring discomfort to users, but they also represent a threat to the users’ health [42]. In fact, Texas peak demand record registered a wide number of users left with no electricity in freezing temperatures [30].

Efforts to tackle peak loads have been actively investigated by the research and industry communities for several years. Early attempts include **Price-Based Demand Response (PBDR)** [56, 58], according to which higher tariffs are set to discourage users from consuming power during periods of high demand. However, since users do not keep up with, and often ignore, such dynamic tariffs, this approach has been proved to be ineffective in the long term [4, 33]. To reduce the impact of peak loads on the power grid, several countries, such as Iran [34], South Africa [45], and the United States [5, 21], are focusing on alternative methodologies. For instance, some utility companies urge users, through various communication and media outlets, to reduce their power consumption when a peak load is expected [40]. Since users often do not comply with these requests, utility companies resort to *rotating outages* to prevent large-scale blackouts. However, these outages often last much longer than originally planned and cause great discomfort to the users [5]. Rotating outages, as well as the recent extreme weather events and consequent blackouts, recently occurred in Texas [24] and California [5] bear witness to the reality that peak loads remain a present issue without a concrete solution. Moreover, this issue is rapidly worsening, since the number of blackouts from weather-related events has grown exponentially since 2000, while blackouts from non-weather-related events has stayed more or less constant since 1984 [31, 38].

The diffusion of **Internet-of-Things (IoT)** devices in power systems, such as smart thermostats (e.g., Nest [25]), energy management systems, and the Advanced Metering Infrastructure [9, 10], enables a wide range of approaches to monitor power consumption and realize previously impossible fine-grained energy management solutions. One of such solutions is *power conservation*, where power consumption is reduced on the user side by changing the settings, shutting down, or delaying the use of certain appliances. As a matter of fact, a recent study [1] has shown that power conservation has the potential to greatly reduce peak loads, by temporarily reducing the




Fig. 1. A schematic overview of the proposed framework.

power consumption by more than 50% on a global scale. Several studies [15, 44] have focused on IoT-based power conservation techniques to address peak loads. Despite these numerous efforts, many solutions lack effectiveness due to the unsuccessful long-term user engagement. In fact, these solutions require users to excessively interact with both their devices and the utility company, thus leading them to experience response fatigue, disengagement, and potentially even abandonment of such systems [33, 47]. Hence, it is important to provide a more flexible program with *personalized preferences* that prioritizes user well-being and achieves long-term engagement. As a consequence, more recent efforts have focused on developing *models of realistic user behavior* [12, 20, 32] that study user interaction of users with appliances and with the power conservation program. However, such approaches often consider abstract appliances with oversimplified (constant) power profiles, resulting in impractical power conservation strategies. Thus, to design accurate and well-targeted power conservation programs, it has been shown that it is important to include *realistic home-level power dynamics* [50] that consider easy-to-obtain information to support large-scale deployment. In summary, a successful power conservation technique need to explicitly address the individuality of user needs and behaviors. Such individuality consists of three main elements, namely,

- (1) providing *personalized preferences* of power conservation,
- (2) adopting *models of realistic user behavior*, and
- (3) performing *realistic home-level power dynamics*.

To the best of our knowledge, our work is the first to design a comprehensive framework for power conservation that addresses all these challenges simultaneously, while fulfilling the utility company's goal to perform power conservation, and thus reduce the peak load.

In this work, we design a comprehensive framework for power conservation that, to the best of our knowledge, is the first to address the above-mentioned challenges. Specifically, as depicted in Figure 1, we envision a **System Operator (SO)**, which handles power generation and transmission and monitors the load of the power grid. When a peak load is anticipated, the SO engages with **Load Serving Entities (LSEs)** to perform bus-wise power conservation exploiting reverse auction theory. Each LSE prompts their users to provide a set of *personalized preferences*, consisting of a set of HVAC temperature adjustments along with the corresponding expected financial compensations. Such

preferences are called *bids*, in auction terminology, and can be submitted automatically based on a pre-defined profile. Since the power saving corresponding to a given temperature adjustment strictly depends on the home-specific factors, we develop a machine learning-based model, called **Power Saving Prediction (PSP)** that captures such *home-level power dynamics*. To enable large-scale deployment, the PSP uses a black-box approach relying on easy-to-obtain information. Furthermore, different from previous approaches, the PSP focuses on the prediction of HVAC power savings during a *transient* state of the power trend due to short-term temperature adjustments.

Our proposed approach assumes that each user is equipped with a **Smart Energy Management System (SEMS)** capable of monitoring the HVAC power consumption and executing a machine learning prediction algorithm, such as PSP. Thus, the user SEMSs communicate bids and the corresponding power savings to the LSE, which in turn exploits the *reverse auction theory* to perform power conservation. The auction mechanism is composed by an optimization problem named **Power Conservation Optimization (POCO)** and a payment rule. POCO minimizes the financial rewards that the utility company pays the users, thus intuitively minimizing their discomfort, while guaranteeing the required power conservation. We prove that the mechanism is *truthful* and *individually rational*. However, we also prove that POCO is NP-hard. Therefore, we propose two truthful and individually rational approaches, namely, a **Dynamic programming Power Saving (DYPS)** and a **Greedy Ranking Allocation (GRAN)** heuristic. DYPS is a pseudo-polynomial algorithm that optimally solves POCO, while GRAN provides a heuristic solution in polynomial time.

We carry out realistic experiments to evaluate the performance of our framework using data collected from the high-fidelity energy simulator EnergyPlus. EnergyPlus is funded by the U.S. Department of Energy [55], and tested according to the **American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE)** Standard 140 methodology [54], thus representing the gold standard of energy data simulations. Furthermore, our online survey shows that 79% of participants are willing to join such a power conservation program, but it also reveals that bids are non-linear and highly variable. Results show that PSP is able to predict power savings at different time frames with over 95% of samples within only a 5% error. Additionally, our results demonstrate that DYPS finds the optimal solution, while GRAN is within 30% of such solution. Conversely, recent state-of-the-art approaches are within 140% and 860% of the optimal solution.

2 Related Work

Peak loads have been investigated for quite a long time. Early studies focused on techniques known as PBDR [53], where fixed time-of-use tariffs [58] or real-time time-of-use tariffs [16] were intended to set higher prices during peak load periods [4, 33] to discourage users to heavily consume power. Due to PBDR's poor engagement of users, who would simply not keep up with such dynamic tariffs, **Incentive-Based Demand Response (IBDR)** has been proposed. IBDR involves users by providing a financial reward in exchange of a certain adjustment in their power consumption. Efforts in IBDR include curtailable load programs [18, 56] and direct load control [22]. However, curtailable load programs do not represent a robust solution due to the fact that during a peak load certain appliances may already be off or idle. On the other hand, while direct load control has the potential to curtail loads, it is perceived as too invasive, hence resulting in low engagement and high abandonment rate.

The need for better user engagement in end-user power conservation is a core element that has come to light in the current state of the art. One way to further engage users in the long term is to allow users to submit their own set of *personalized preferences* for power conservation, in exchange of a financial reward. This idea is at the basis of an approach known as *demand side bidding*. For

instance, [15] and [44] ask users to turn off certain appliances during a peak load period in exchange of a payment. Similarly, in [59], the authors exploit the storage of electric vehicles during times of peak loads, in the context of an auction mechanism, where users provide their preferred options of power conservation. A recent paper that provides power conservation for data center networks, following an auction-based approach similar to ours, is [17]. We use this as a comparison approach given the similarities of the considered problem and solution. Despite the above efforts at engaging users by providing the possibility to submit personalized preferences, [47] has shown that, if the system requires frequent interactions with users to adjust a set of appliances, they might experience *response fatigue*. According to this phenomenon, users stop interacting with the system because they become tired of engaging with it.

To further understand how users can be incentivized, while avoiding the response fatigue phenomenon, researchers have proposed to include *models of realistic user behavior* in the design of power conservation strategies. These models help understand and emulate how users interact with a power conservation program. For instance, [20] and [32] strictly focus on models of realistic user behavior in terms of the perceived utility of each appliance. However, the former does not provide a way to submit personalized preferences of power consumption, while the latter only considers abstract appliances with a power consumption represented by fixed constant values. This significantly oversimplifies the impact of the operation of such appliances on the home power consumption and on the overall demand response approach. As stated in [50], this represents a considerable limitation in the design of power conservation systems.

Another important aspect to consider in the design of effective power conservation strategies is to include *realistic home-level power dynamics* [49]. Additionally, HVAC has been shown to be a promising appliance for power conservation, due to its energy consumption being highly correlated with peak loads and the rapid adoption of smart thermostat that can perform temperature adjustments automatically [2, 12]. An HVAC power conservation mechanism has been proposed in [2]. This approach provides equal rewards to all users, irrespective of their preferences. Additionally, the authors rely on a *white-box approach*, which requires a profound knowledge of the chemical and structural properties of the house materials, as well as layout. Thus, it is not only unfeasible for the large-scale deployment needed to achieve sufficient power conservation, but it is also impractical to gather such information about a single house. Since this work shows similarities to ours, we use it as additional comparison approach in the experimental section.

In summary, as shown in Table 1, most existing works lack at least one of the necessary elements of successful power conservation related to the individuality of users. Conversely, this article proposes an effective power conservation framework that aims at long-term user engagement by providing

- (1) *personalized preferences* for each user,
- (2) *models of realistic user behavior* to study how users engage with such system,
- (3) *realistic home-level power dynamics*, based on easily available information, that allow large-scale deployment of the system.

Finally, note that other important aspects often addressed in demand response papers were not included in our table. For instance, although we mention the concept of comfort, we do not explicitly talk about the related topic of occupancy. Occupancy is a parameter that is often leveraged to achieve further power savings in this context [7], while also addressing user comfort by adopting more or less aggressive power conservation strategies [36]. Regarding comfort, in our work, we aim at minimizing compensation, which indirectly maximizes comfort. Furthermore, our work expects users to submit the options of power savings that they are comfortable with. Thus, if a user is not home and is willing to set higher temperature changes, they would be able to do so.

Table 1. Related Work Summary

	Personalized preferences	Models of realistic user behavior	Realistic home-level power dynamics
Baldi et al. [7]	×	×	×
Korkas et al. [36]	×	×	×
Korkas et al. [35]	×	×	×
Chen et al. [17]	✓	×	×
Khamesi et al. [44]	✓	×	×
Chapman et al. [15]	✓	×	×
Zhou et al. [59]	✓	×	✓
Shafie-Khah et al. [47]	×	✓	✓
Ciavarella et al. [18]	×	✓	×
Khamesi et al. [32]	×	✓	✓
Dolce et al. [20]	×	✓	✓
Kim et al. [33]	×	×	×
Asadinejad et al. [4]	×	×	×
Ali et al. [2]	×	×	✓
Wang et al. [56]	×	×	✓
Shi et al. [49]	×	×	✓
Ericson [22]	×	×	✓
Our approach	✓	✓	✓

While we believe that not relying on occupancy information is an advantage of our approach, as no additional sensor is required, such information could be integrated to, for example, automate the bidding process. Ultimately, within the topic of large-scale deployment, we also choose to only rely on implementations that are based not only on widely available weather information but also on widely available and affordable smart thermostats, such as Nest. For this same reason, we also do not consider the adoption of additional energy storage units, as in [35], since they are not widely available and are very expensive. However, if the utility company had storage units and renewable energy could be stored in them, our implementation could be simply extended by including these availabilities in the power cap set by the LSE. Similarly, if certain users have local storage units, our framework could be extended to consider the additional power saving that the user would obtain from such storage within the auction period.

A preliminary version of this work appeared in [12]. This article significantly extends the conference version by proposing DYPS, a new truthful and individually rational solution based on dynamic programming. Furthermore, a more comprehensive and realistic experimental validation is considered, with more realistic parameters, finer time resolution, and new performance metrics. The experiments show the superiority of DYPS with respect to the state of the art, including the algorithm originally proposed in [12] and an additional recent work proposed in [49].

3 Problem Formulation

We consider a SO controlling a power grid with a set of L buses managed by LSEs. Each bus $l \in L$ serves a set of N_l residential users equipped with an Internet-connected SEMS which monitors, learns, analyzes, and controls the HVAC system to set temperature changes. The SEMS of user i at bus l , for $1 \leq i \leq N_l$, also interacts with the LSE to implement the power conservation at

the residential distribution feeder. We assume the SO is able to predict the expected aggregated power consumption, for example a day ahead. Note that, this is a relatively easy problem to solve, widely studied in the literature [3], since deviations from the historical average of individual user behaviors are canceled out through aggregation [13]. When the SO predicts a peak load at bus l with an expected total power consumption $P_T^{(l)}$, it calculates the *power cap* $P_C^{(l)} = \alpha^{(l)} \times P_T^{(l)}$, where $\alpha^{(l)} \in [0, 1)$, according to the system's characteristics, such as the generation capacity, cost of generation, capacity of transmission/distribution lines, and so forth.¹ Therefore, the required *power saving* is $P_S^{(l)} = P_T^{(l)} - P_C^{(l)}$. Since the power saving is provided by the SO to each LSE, which conserves the power independently through its users, in the following we formulate the problem for a single LSE. Thus, without loss of generality, we drop the superscript l from the problem formulation.

The LSE alerts the user SEMs that the power conservation auction is activated, requesting for bids. An SEM asks its user directly or submits the bids based on a 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 the participating users. As a result, to provide *personalized preferences*, each user i submits to the LSE, through their SEM, a set $B_i = \{B_{ij} = (\Delta P_{ij}, \Delta T_{ij}, C_{ij}) : 1 \leq i \leq N, 1 \leq j \leq M_i\}$ of M_i bids. Here ΔP_{ij} , ΔT_{ij} , C_{ij} , respectively, represent the power saving, temperature change, and monetary compensation² for user i and bid j . We discuss in Section 6 how *realistic home-level power dynamics* are implemented through the PSP algorithm, which predicts power saving ΔP_{ij} corresponding to the temperature change ΔT_{ij} .

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. Note that, we assume that users will keep their thermostat setting during the power conservation period, if their bid is not selected as a winning bid. This way, the PSP algorithm can predict the energy saving associated to each user bid. Furthermore, we assume that if a user is selected as a winner, the proposed thermostat change is enforced by the SEM. Therefore, the user is not able to change the setting at their discretion. The participating users agree on a contract to participate in the power conservation framework, implying the loss of complete control of the HVAC setting during the power conservation period.

The winner selection strategy is formulated as an Integer Linear Programming optimization problem that aims to minimize the costs in terms of the paid compensations, while satisfying the power cap constraint. As shown in Section 7, 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 POCO problem defined as:

$$\begin{aligned} & \min \sum_{i=1}^N \sum_{j=1}^{M_i} C_{ij} w_{ij} \\ & \text{subject to} \end{aligned} \tag{1a}$$

¹We assume the SO predicts the peak load and its duration. The proposed framework is supposed to be executed for the predicted duration.

²In this article, we use the terms "cost" and "monetary compensation" to represent financial rewards from the LSE and user perspective, respectively.

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

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

$$\sum_{i=1}^N \sum_{j=1}^{M_i} \Delta P_{ij} w_{ij} \geq P_S. \quad (1d)$$

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 Expression (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 Expression (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

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

where $\mathbf{w}^*$ is the optimal solution of POCO, $\mathbf{w}^{(-k)*}$ 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., a revenue, defined as $U_k = E_k - C_k$. In the following, we prove *truthfulness* and *individual rationality* of POCO to ensure an effective power conservation program [57]. Truthfulness prevents potential unhealthy bidding behavior, by providing reduced utility U_k when users bid differently than the true valuation. Individual rationality guarantees that each winning user is paid an amount that ensures non-negative utility ($U_k \geq 0$). We provide proof of truthfulness in the following.

THEOREM 3.1. *The reverse auction mechanism, as defined by the POCO problem and the payment rule in Equation (2), is truthful.*

PROOF. The proof is inspired by the proof of truthfulness of the well-known Vickrey–Clarke–Groves auction presented in [6]. As stated in Equation (2), the payment rule E_k to a selected winning user k is given by $f(\mathbf{w}^{(-k)*}) - f(\mathbf{w}^*) + \sum_{j=1}^{M_k} C_{kj} w_{kj}^*$. Therefore, the utility of this user is defined as $U_k = E_k - \sum_{j=1}^{M_k} C_{kj} w_{kj}$. Specifically, the utility is intended as a revenue with respect to the truthful compensation V_{kj} . Thus, we can write $U_k = E_k - \sum_{j=1}^{M_k} V_{kj} w_{kj}$. To prove truthfulness, we show that the utility U_k of a user k in case that their declared compensation is equal to the true valuation ($C_k = V_k$) is greater than the utility U'_k of the same user in case the declared compensation is not equal to the true valuation ($C_k \neq V_k$), i.e., $U_k - U'_k > 0$, while all other bids are unchanged and exactly the same. Let us define U_k and U'_k as

$$U_k = f(\mathbf{w}^{(-k)*}) - f(\mathbf{w}^*) \quad (3)$$

$$U'_k = f(\mathbf{w}^{(-k)*}) - f(\mathbf{w}^*) + \sum_{j=1}^{M_k} (C_{kj} - V_{kj}) w_{kj}, \quad (4)$$

where $f(\mathbf{w}^*) = \sum_{i=1}^N \sum_{j=1}^{M_i} C_{ij} w_{ij}^*$. Note that in calculating $U_k - U'_k$, the terms $f(\mathbf{w}^{(-k)*})$ and $f(\mathbf{w}^{(-k)*})$ are equal, since the optimal solution without the user k is the same in both cases, truthful and untruthful bidding of user k

$$\begin{aligned}
U_k - U'_k &= - \sum_{i=1}^N \sum_{j=1}^{M_i} C_{ij} w_{ij}^* + \sum_{i=1}^N \sum_{j=1}^{M_i} C_{ij} w_{ij}^{*'} - \sum_{j=1}^{M_k} (C_{kj} - V_{kj}) w_{kj}^{*'} \\
&= - \sum_{i=1}^N \sum_{j=1}^{M_i} C_{ij} w_{ij}^* + \sum_{i \neq k}^N \sum_{j=1}^{M_i} C_{ij} w_{ij}^{*'} + \sum_{j=1}^{M_k} C_{kj} w_{kj}^{*'} - \sum_{j=1}^{M_k} (C_{kj} - V_{kj}) w_{kj}^{*'} \\
&= - \sum_{i=1}^N \sum_{j=1}^{M_i} C_{ij} w_{ij}^* + \sum_{i \neq k}^N \sum_{j=1}^{M_i} C_{ij} w_{ij}^{*'} + \sum_{j=1}^{M_k} V_{kj} w_{kj}^{*'}.
\end{aligned}$$

Now we combine the last two terms into one summation by considering that V_{kj} is the user compensation in the truthful scenario ($V_k = C_k$), therefore, we have

$$U_k - U'_k = - \sum_{i=1}^N \sum_{j=1}^{M_i} C_{ij} w_{ij}^* + \sum_{i=1}^N \sum_{j=1}^{M_i} C_{ij} w_{ij}^{*'}.$$

Since w_{ij}^* is the optimal solution (of a minimization problem), given a truthful compensation ($V_k = C_k$) by user k in the first nested sum, then this value is always less than or equal to the second nested sum, and therefore, $U_k - U'_k \geq 0$ is always true. The equality holds when $w_{ij}^* = w_{ij}'$. $\square$

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

THEOREM 3.2. *POCO is individually rational.*

PROOF. Consider payment rule defined in Equation (2) $E_k = f(\mathbf{w}^{(-k)*}) - f(\mathbf{w}^*) + \sum_{j=1}^{M_k} C_{kj} w_{kj}^*$ and utility $U_k = f(\mathbf{w}^{(-k)*}) - f(\mathbf{w}^*)$. Since we want to prove that $U_k \geq 0$, then we want to prove that $f(\mathbf{w}^{(-k)*}) \geq f(\mathbf{w}^*)$.

By contradiction, it would be impossible to have $f(\mathbf{w}^{(-k)*}) < f(\mathbf{w}^*)$. In fact, if $\mathbf{w}^{(-k)*}$ is the solution when user k does not participate in the auction, then the optimal solution $\mathbf{w}^*$ can only improve, and thus we have $f(\mathbf{w}^*) \leq f(\mathbf{w}^{(-k)*})$, which proves that $U_k \geq 0$. $\square$

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

THEOREM 3.3. *POCO is an NP-hard problem.*

PROOF. The NP-hardness can be proven as a reduction from the minimum 0-1 knapsack problem (minKP) [19]. The minKP looks for the set of items with minimum weight and a cumulative value larger than or equal to a target value. A reduction of POCO can be solved by minKP. Thus, we set $M_i = 1$, according to which each user in POCO can only bid one offer. Therefore, the decision variable w_{ij} can be written as a simple w_i which is equal to 1 when the user is a winner of the auction, and 0 otherwise, thus obtaining the following formulation:

$$\min \sum_{i=1}^N C_i w_i \quad (5a)$$

$$\text{subject to: } w_i \in \{0, 1\} \quad (5b)$$

$$\sum_{i=1}^N \Delta P_i w_i \geq P_S. \quad (5c)$$

The optimal solution to this instance of our problem is a formulation of minKP where values are represented in terms of compensations, while weights and maximum capacity are represented in

terms of power consumption values and the required powered saving. Because this instance of POCO, formulated in the form of a minKP, provides a solution that is also a solution to the knapsack problem, we can say that solving POCO is at least as difficult as solving KP, and therefore, POCO is NP-hard. $\square$

4 The DYPS Algorithm

In this section, we present a pseudo-polynomial algorithm to solve POCO called DYPS. DYPS is based on dynamic programming and it is composed of two phases. The first phase divides the original problem into sub-problems, and it exploits a *recursive relation* that provides the solution of bigger sub-problems by exploiting the solutions of smaller sub-problems. The output of the first phase is the *value* of the optimal solution of POCO. The second phase uses a recursive algorithm to find the set of winning bids that provide the optimal solution of POCO by back-tracking the first-phase decisions. Note that, since DYPS solves POCO optimally, we pair it with the same payment rule in Equation (2), and thus it inherits the property of individual rationality and truthfulness.

4.1 DYPS: Recursive Relation

The core of DYPS is a recursive relation that allows us to solve larger sub-problems from the solution of smaller sub-problems. Given an instance of POCO with a set of bids $\mathbf{B}$ and a power saving P_S , we define a sub-problem as an instance of POCO with a reduced input, i.e., considering a subset of bids in $\mathbf{B}$ or a smaller value of P_S . Starting from base cases that are straightforward to solve, the size of the input is gradually increased to find the solution to the original problem, i.e., the value of the optimal solution. This is achieved by exploiting a table T of size $|\mathbf{B}| \times P_S$, which stores the solutions of the sub-problems. The element $T[x, y]$ is the solution to a sub-problem that considers the first x bids in $\mathbf{B}$ and a power saving $P_S = y$. To define such recursive relation, we first define the *base cases* and subsequently the *recursive cases*.

4.1.1 Base Cases. The base cases occur when $x = 0$ or $y = 0$. For $x = 0$, there are no bids, i.e., $\mathbf{B} = \emptyset$. Therefore, it is impossible to satisfy the power saving. As a result, we set $T[0, y] = \infty$ for all values y of power saving. Conversely, when $y = 0$ no power reduction is needed ($P_S = 0$). Consequently, there are no auction winners and the value of the optimal solution is zero. Thus, we set $T[x, 0] = 0$ for all x . Note that, when $x = 0$ and $y = 0$, we set $T[0, 0] = 0$.

4.1.2 Recursive Cases. For $x > 0$ and $y > 0$, we identify two *recursive cases*. Let x correspond to the bid $B_{ij} = (\Delta P_{ij}, \Delta T_{ij}, C_{ij})$. For $x > 0$, the bid B_{ij} is sufficient to satisfy the required entire power saving, i.e., $\Delta P_{ij} \geq y$. In this case, the optimal solution of POCO is either the cost C_{ij} (i.e., B_{ij} is the only winner), or it is the optimal solution without B_{ij} (i.e., using $x - 1$ bids). Between these two options, we choose the solution with the minimum cost, i.e., $T[x, y] = \min(C_{ij}, T[x - 1, y])$.

Now, for $y > 0$, the bid B_{ij} is not sufficient to fulfill the entire power saving, i.e., $\Delta P_{ij} < y$. Thus, the optimal solution of POCO may not include B_{ij} . If it does (i.e., B_{ij} is a winning bid), no other bid of user i can be a winning bid. Therefore, the optimal solution is composed of B_{ij} plus the bids of the optimal solution of the sub-problem with $x - j$ bids³ and a power saving of $y - \Delta P_{ij}$. Conversely, if B_{ij} is not included in the solution, the optimal solution is that of the sub-problem with $x - 1$ bids and power saving y . Among these two options, we choose the solution with the minimum cost, i.e., $T[x, y] = \min(C_{ij} + T[x - j, y - \Delta P_{ij}], T[x - 1, y])$.

³We assume that the bids in $\mathbf{B}$ are numbered such that the bids of the same user are adjacent. Since we are currently considering B_{ij} , thus $x - j$ refers to the last bid of user $i - 1$.

Algorithm 1: DYPS Recursive Relation Algorithm

```

Input :  $B, P_S$ 
Output: Table  $T$ 
/* Base cases */
1  $T[0, y] = \infty \quad \forall y$ 
2  $T[x, 0] = 0 \quad \forall x$ 
/* Recursive cases */
3 for  $x \leftarrow 1$  to  $|B|$  do
4   for  $y \leftarrow 1$  to  $P_S$  do
5      $B_{ij} = (\Delta P_{ij}, \Delta T_{ij}, C_{ij}) = x\_to\_bid(x)$ 
6     if  $\Delta P_{ij} \geq y$  then
7       // Case 1) Bid alone fulfills power cap  $y$ 
7        $T[x, y] = \min(C_{ij}, T[x - j, y])$ 
8     else
9       // Case 2) Bid alone does not fulfill power cap  $y$ 
9        $T[x, y] = \min(C_{ij} + T[x - j, y - \Delta P_{ij}], T[x - 1, y])$ 
10    end
11  end
12 end
13 return  $T$ 

```

In summary, the recursive equation is the following:

$$T[x, y] = \begin{cases} 0, & \text{if } y = 0 \\ \infty, & \text{if } x = 0 \\ \min(C_{ij}, T[x - 1, y]), & \text{if } \Delta P_{ij} \geq y \\ \min(C_{ij} + T[x - j, y - \Delta P_{ij}], T[x - 1, y]) & \text{if } \Delta P_{ij} < y \end{cases} \quad (6)$$

4.1.3 Pseudo-Code. Algorithm 1 shows the pseudo-code of DYPS algorithm that takes the set of bids B and the power saving P_S as input. The base cases are initially addressed in lines 1 and 2. Subsequently, two nested loops iterate x over the set of bids and y from 1 to P_S . Each iteration considers the sub-problem $T[x, y]$. We make use of an auxiliary function $x_to_bid()$ that returns the bid B_{ij} corresponding to x . Lines 6 and 7 correspond to the first recursive case, in which the current bid provides sufficient power saving to meet the power cap ($\Delta P[i, j] \geq y$). Similarly, Lines 8 and 9 provide the recursive solution for the case of $\Delta P[i, j] < y$.

The algorithm returns the table T , which contains in position $T[|B|, P_S]$ the value of the optimal solution of POCO. In the following algorithm, we describe how to obtain from T the actual solution, i.e., the set of winning bids.

4.2 DYPS: Finding the Solution

Given the table T returned by the previous phase of DYPS, we now provide a recursive algorithm, called DYPS-Sol, to find the actual solution, i.e., the set of winning bids. DYPS-Sol starts from $T[|B|, P_S]$, which contains the value of the optimal solution of POCO, and recursively navigates the table T to back track the decision of the algorithm during the first phase until a base case is reached. A set S is updated by adding the winning bids encountered during the recursive iterations.

Algorithm 2: DYPS Solution Algorithm

```

1 DYPS-Sol (T, B, x, y, S)
2 if x == 0 ∨ y == 0 then
3   return S
4 end
5  $B_{ij} = (\Delta P_{ij}, \Delta T_{ij}, C_{ij}) = x\_to\_bid(x)$ 
6 if  $T[x, y] == C_{ij} + T[x - j, y - \Delta P_{ij}]$  then
7    $S \leftarrow S \cup \{B_{ij}\}$ 
8   return DYPS-Sol (T, B, x - j, y -  $\Delta P_{ij}$ , S)
9 end
10 if  $T[x, y] == C_{ij}$  then
11    $S \leftarrow S \cup \{B_{ij}\}$ 
12   return S
13 end
14 return DYPS-Sol (T, B, x - 1, y, S)

```

The pseudo-code of DYPS-Sol is shown in Algorithm 2. The algorithm takes as input the table T , the set of bids B , the current sub-problem input x and y , and the current solution S . The algorithm is initially called as $DYPS-Sol(T, B, -B-, P_S, S)$, where S is initially empty.

During a generic iteration, the algorithm first checks if we are in a base case (line 2), i.e., if x or y are zero. In that case, the current solution S is returned and the algorithm terminates. If we are not in a base case, the algorithm extracts the current bid $B_{ij} = (\Delta P_{ij}, \Delta T_{ij}, C_{ij})$ in line 5. Subsequently, in line 6, the algorithm checks if the previous phase picked the current bid and this bid was not able to fulfill the entire power saving, i.e., if $T[x, y] == C_{ij} + T[x - j, y - \Delta P_{ij}]$. In this case, the bid is added to the current solution and the algorithm is recursively called on $x - j$ and $y - \Delta P_{ij}$. Then, the algorithm checks if B_{ij} was picked as a single bid able to fulfill the entire power saving y . This can be verified by checking if $T[x, y] == C_{ij}$. If this is the case, the algorithm adds B_{ij} to the current solution and terminates the recursive calls (no other winning bid can be added). Finally, if all previous cases are not true, the current bid has not been selected by the first phase. As a result, the algorithm calls itself recursively by excluding this bid, i.e., bid $x - 1$ and power saving y .

4.3 DYPS: Complexity

The following theorem shows the pseudo-polynomial complexity of DYPS.

THEOREM 4.1. *The time complexity of DYPS mechanism is $O(|B|P_S)$.*

PROOF. The time complexity is dominated by the payment rule in Equation (2), which is calculated for each winner of the auction, that is $|\mathbf{w}^*|$ times. Each calculation requires to execute DYPS twice, to find $\mathbf{w}^{(-k)*}$ and $\mathbf{w}^*$. The complexity of DYPS is dominated by the two for loops on lines 3 and 4, resulting in the time complexity of $|B| \times P_S$. As a result, the overall time complexity of DYPS is $O(|\mathbf{w}^*| \times |B| \times P_S)$. $\square$

5 The GRAN Mechanism

In this section, since POCO is NP-hard, and DYPS has pseudo-polynomial time complexity, we propose a heuristic called GRAN to find an efficient solution for POCO, while guaranteeing truthfulness and individual rationality of the auction mechanism.

Algorithm 3: GRAN: Greedy Ranking AllocationN

Input : P_T , α , and B_i $i = 1, \dots, N$
Output: List of Winners W

```

1  $W \leftarrow \emptyset, P_{CS} = 0$ ; // Variables initialization
2  $P_C = \alpha \cdot P_T$ ; // Power cap
3  $P_S = P_T - P_C$ 
4  $R \leftarrow \{R_{ij} = \frac{\Delta P_{ij}}{C_{ij}} \mid i = 1, \dots, N, j = 1, \dots, M_i\}$ 
5 Sort elements of list  $R$  in a non-ascending order
6 while  $P_{CS} < P_S$  and  $R \neq \emptyset$  do
7   Let  $R_{ij}$  be the first element in  $R$  and
8   Let  $B_{ij} = (\Delta P_{ij}, \Delta T_{ij}, C_{ij})$  be the bid corresponding to  $R_{ij}$ 
9    $P_{CS} = P_{CS} + \Delta P_{ij}$ ; // Update cumulative power saving
10   $W \leftarrow B_{ij}$ ; // Update list of winners
11  Remove all bids of user  $\hat{i}$  from  $R$ 
12 end
13 return  $W$ 

```

5.1 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 3.

In line 1 of Algorithm 3, we initialize the list of the auction winners 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 Expression (1d) (lines 2 and 3). Since our goal is to minimize the objective function in Expression (1a), GRAN uses a ranking criterion which gives precedence to the bids with low cost and large power saving. GRAN uses a list R that stores the values of ranking criterion in non-ascending order (lines 4 and 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_{ij} with the greatest ranking criterion R_{ij} in R (lines 7 and 8). Therefore, we increase P_{CS} by the corresponding power saved (line 9) and we add the winning bid B_{ij} to the list of winners W (line 10). Finally, we remove all other elements from user $\hat{i}$ in 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 4. It may be possible that GRAN is unable to meet the power cap and terminates the while loop because $R = \emptyset$. 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 [49], and this situation never occurred.

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_{ij} . If user i submits a compensation $C_{ij} > E_i$, it loses; otherwise, it wins. In Algorithm 4, we obtain the critical bid as follows. We find

Algorithm 4: 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    $\mathbf{W}_{-i} = \text{GRAN}(\mathbf{B}_{-i})$  ; //  $\mathbf{B}_{-i} = \bigcup_{k=1}^N \mathbf{B}_k \setminus \{\mathbf{B}_i\}$ 
3   Let  $B_{i\bar{j}}$  be the last element added to  $\mathbf{W}_{-i}$ 
4    $E_i = \frac{\Delta P_{ij}}{R_{i\bar{j}}}$ 
5 end
6 return  $\mathbf{E}$ 

```

the solution $\mathbf{W}_{-i}$ of GRAN when user i is not participating in the auction (line 2). Then, we select the critical bid $B_{i\bar{j}}$ as the last bid added to the solution set (line 3). Finally, in line 4, we define the critical value $E_i = \frac{\Delta P_{ij}}{R_{i\bar{j}}}$. In the following subsection, we will prove that this payment rule, paired with the winner selection algorithm, guarantees truthfulness of the GRAN mechanism.

5.2 GRAN Properties

To prove that the GRAN mechanism is truthful, we follow the approach similar to [11]. More precisely, we first prove that the winner selection algorithm (Algorithm 3) is monotonic, and then that the payment rule (Algorithm 4) pays the critical value.

Definition 5.1. (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 5.2. *Algorithm 3 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 3 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 5.3. *Each winning bid is paid the critical value.*

PROOF. Our goal is to prove that the payment rule we defined in Algorithm 4 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 3, and consider the critical bid $B_{i\bar{j}}$ in line 3 of Algorithm 4, 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 $C_{ij} > \frac{\Delta P_{ij}}{R_{i\bar{j}}}$. Multiplying both members by $R_{i\bar{j}}$ and then dividing by C_{ij} yields $R_{i\bar{j}} > R_{ij}$. Because Algorithm 3 sorts values of ranking criterion non-ascendingly, R_{ij} would be placed after $R_{i\bar{j}}$ 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. Similar to Case 1, $C_{ij} \leq E_i$ yields $R_{i\bar{j}} < R_{ij}$, which makes B_{ij} a winning bid, since R_{ij} would be placed before $R_{i\bar{j}}$ in $\mathbf{R}$. $\square$

THEOREM 5.4. *The GRAN mechanism is truthful.*

PROOF. Following [11, Theorem 9.36], the proof of this theorem follows from Theorems 5.2 and 5.3 proved above. $\square$

THEOREM 5.5. *The GRAN mechanism holds the property of individual rationality.*

PROOF. To prove individual rationality, we need to show that the utility $U_k = E_k - C_{kj}$ is non-negative, i.e., $\frac{\Delta P_{ij}}{R_{ij}} - C_{kj} > 0$. This is equivalent to $C_{kj} < \frac{\Delta P_{ij}}{R_{ij}}$, which is easily proven because by construction of GRAN (Algorithms 3 and 4), if we had $C_{kj} > \frac{\Delta P_{ij}}{R_{ij}}$, then user k would not be a winning bid. $\square$

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

THEOREM 5.6. *The time complexity of the GRAN mechanism is $O(N^2 M_{\max} \log(NM_{\max}))$, where $M_{\max} = \max_{i=1, \dots, N} |B_i|$ is the maximum number of bids submitted by a user.*

PROOF. We analyze the time complexity of the winner selection (Algorithm 3) and the payment rule (Algorithm 4) separately.

[Algorithm 3]: In line 4 of Algorithm 3, we generate the list $\mathbf{R}$ and sort it in line 5. The list size is $O(NM_{\max})$, 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 3 is $O(NM_{\max} \log(NM_{\max}))$.

[Algorithm 4]: 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 3. 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 4. $\square$

6 PSP

To effectively select the winners of the auction and meet the power cap constraint, it is necessary to model realistic home-level power dynamics, i.e., know the power saving corresponding to the HVAC thermostat adjustment contained in 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, and so forth. [46]. In our case, this is even more challenging, since our goal is to predict the power saving resulting from a sudden and short-time change in the thermostat setting. This implies that, 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 [43], making the prediction problem extremely challenging.

The proposed PSP algorithm allows us to model realistic home-level power dynamics by means of 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 [41]. 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* [46], 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;
- *Time*: hour of the day.⁴

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 [8]. 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 8). We evaluated the performance of several regression algorithms (Artificial Neural Networks, **Random Forest (RF)**, Elastic-Net, Support Vector Machines, Nearest Neighbors Regression, and Naive Bayes) and tested the values of various parameters with a grid search. We found *RF regression* [51] to provide the best performance. In our experiments, we set the parameters as follows: (1) the criterion to measure split quality to the **mean squared error (MSE)**, (2) the maximum depth of the tree to 1,000, (3) the number of estimators to 150, (4) the minimum number of samples needed to split a node to 2, (5) the maximum number of features while deciding the best split equal to the total number of features (7 in this case), and (6) the minimum number of samples required to be at a leaf node to 1.

7 Online Survey

We conduct an online survey involving 200 subjects to model realistic user behaviors. Specifically, the objective of the survey is to assess bidding behavior and the 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.

7.1 Overview of the Survey

The participants 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/hour which was above the federal minimum hourly wage of \$7.75/hour at the time the study was conducted and above the top 4% earning rate of \$7.50/hour for M-Turk workers [28].

The survey began by asking participants to indicate their typical thermostat setting on a hot summer day. Then, they 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

⁴The *season* could be another important factor. Different models could be trained for different seasons to take into account this aspect.




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

proposed system that would help reduce 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 1 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°F, their range of options was 71°F–78°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 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 0.01. This dollar range was proposed in [49]. 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 Figure 2.

7.2 Survey Results

A total of 200 participants took part in the survey. However, results do not include 44 users who failed to correctly answer the attention check questions. Overall, more than 79% of users answered that they would be willing to use this system in their homes. Figure 2(a) shows, for a given change in the temperature value (measured in degrees Fahrenheit) along the x -axis, the number of survey participants who agreed to change their thermostat up to that value, but not more. The results are clearly non-linear: most users are comfortable with small temperature changes and become less comfortable as the change increases.

Figure 2(b) 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, 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 define *models of 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.

8 Performance Evaluation

This section presents the experimental setup followed by a thorough performance comparison of our methods versus recent state-of-the-art solutions.

8.1 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 [55] 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 [55].

Since our auction framework is executed independently at each bus, given the power constraint provided by the SO, in the experiments we focus on a set of houses at a single bus. Nevertheless, we consider a variety of scenarios by changing parameters such as the required power constraint, the percentage of participants, and the duration of the auction.

To consider a variety of houses with *realistic home-level power dynamics*, we employed the *EnergyPlus residential prototype building models* provided by the U.S. Department of Energy in collaboration with the Pacific Northwest National Laboratory [39]. The models have four foundation types (slab, crawlspace, heated basement, and unheated basement) and two cooling system types (central air conditioning cooling and heat pump cooling). The combination of these characteristics gives us a total of eight 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 eight models previously mentioned, we generate five additional models by changing the *U-factor* within $[2, 4] \text{ W}/(\text{m}^2 \text{ K})$ range [29]. 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. However, 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.

8.2 Performance of the PSP Algorithm

In this section, we study the performance of the PSP algorithm.

8.2.1 Comparison Approach Sha-Support Vector Regression (SVR). We compare PSP to a recent state-of-the-art approach for power prediction proposed in [46], which we refer to as Sha-SVR. We select this approach because, similar to our framework, it is designed to work in specific building settings and it uses limited and easily available features to facilitate large-scale deployment of the system.

The authors of Sha-SVR adopt SVR as 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 [27]. $CDD = \max\{(T_{\max} - T_{\min})/2 - T_c, 0\}$, where $T_{\max}$ and $T_{\min}$ are the maximum and minimum hourly temperature in a day, and $T_c = 59^\circ\text{F}$ (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. We add to the feature set the current temperature set point. Then, to calculate the power saving resulting from a transition from set point T_{old} to a set point T_{new} , we use Sha-SVR to predict the steady-state power consumption $P_{(T_{old})}$ and $P_{(T_{new})}$ separately. We then calculate the power saving $\Delta P = P_{(T_{old})} - P_{(T_{new})}$. For more details on Sha-SVR, refer to [46].

8.2.2 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 [39]. 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 eight thermostat change options, each representing a 1°F (approx. 0.55°C) degree difference, and we consider three different time frames for the auction duration, namely, 15 minutes, 30 minutes, and 1 hour. The thermostat set point is altered at the beginning of the auction and restored to the original value at the end. As a result, our objective is to predict the energy saving for the auction duration. We use EnergyPlus to collect the resulting power consumption data and pair it with the features required by each algorithm. The data are then shuffled before forming training and testing set, 75% and 25% samples, respectively. We analyze the performance of PSP in terms of **Mean Average Percentage Error (MAPE)**, median error, and **explained variance (EV)**⁵. 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 [23]. $EV \in [0, 1]$ and a higher value (i.e., close to 1) represents more accurate predictions.

Figure 3 shows the percentage of testing samples, along the y-axis, that are predicted within a certain MAPE. PSP achieves very high accuracy for the vast majority of samples in all time frames. In fact, the dashed red line in Figure 3(a)–(c) indicates that 95% of testing samples are predicted within a 5% MAPE in all scenarios. On the other hand, Sha-SVR achieves poor performance. This is due to their steady-state approach, which prevents the algorithm from capturing the dynamics that occur after a sudden change of the thermostat set point. Quantitatively, across all time frames shown in each subfigure, more than 75% of Sha-SVR's predictions incur more than 60% MAPE on average, which further proves the need for a *ad hoc* transient-state approach, such as PSP.

Next, in Figure 4(a)–(c), we study the learning rate of PSP. This is particularly important in our context since, to enable large-scale deployment, and capture *realistic home-level power dynamics*, the SEMS of each user has its own PSP model, trained with only limited data. We adopt the median error as a metric, along with EV. The goal of this experiment is to analyze the value of the adopted

⁵These are typical metrics for regression algorithms, comparable to the accuracy, F-score, and so forth, for classification algorithms.




Fig. 3. Performance of prediction models based on error distribution.




Fig. 4. Performance of prediction models based on learning rate.

metrics by progressively increasing the size of the testing set along the x-axis. We performed tests individually for each home and averaged the results. The testing samples for a home are randomly selected. As shown in Figure 4, 20–30 samples are sufficient to obtain a very high *EV* and very low median error for all time frames. These results show that PSP can quickly learn and perform accurate individual home PSPs. Recall that these samples do not need to be collected during peak load periods but can instead be gathered by the SEMS during manual or automatic adjustments that are possible with modern thermostats even more than once a day [41].

8.3 Power Conservation

In this section, we study the performance of the auction framework in dealing with peak loads. We first introduce two recent comparison approaches and then discuss the results.

8.3.1 Comparison Approach Mechanism for Emergency Demand Response (MEDR). A recent paper [17] proposes a truthful auction-based IBPC approach in data centers called 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 B_i . We implement the authors' NP-hard optimization problem to select winners. 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. The authors also propose their own truthful payment rule, which we also implement for the calculation of payments. For more details, refer to [17].

8.3.2 Comparison Approach FLAT. The recent article [49] proposes a subscription-based power conservation system. Users agree to participate on a monthly contract which pays all users the same fixed amount in exchange of a fixed temperature change. Once users subscribe, the adjustment can occur up to once a day at the utility company's will. Due to such flat approach to payments with fixed temperature changes, we call this approach FLAT.

Although the context of this approach is similar to ours, in the sense that users get paid to perform a temperature change in the thermostat, its subscription-based implementation is actually quite different. Hence, for the sake of fair comparison, we slightly adapt the FLAT by performing the following changes. First, we divide the required power saving P_S by all participating users N . Then, we require that each house performs a thermostat adjustment to the closest integer temperature that ensures a power saving of P_S/N . Note that, the integer adjustment is due to the intrinsic characteristics of thermostats which allow temperature changes at 1°F steps. To process payments, we calculate the average compensation requested by all users of our survey, for each temperature change. We pay users the average compensation corresponding to required temperature change.

Note that, FLAT *does not* provide the properties of truthfulness nor individual rationality. This is an advantage for FLAT in terms of raw performance metrics, since providing such properties leads to payments that are always greater than or equal to the requested value.

8.3.3 Results. In the following, we compare the performance of different approaches, namely, the optimal solution OPT obtained with the Gurobi optimizer [26], DYPs, the heuristic solution GRAN presented in this work, our conference version⁶ solution GRAN_PC [12], and the comparison approaches MEDR [17] and FLAT [49]. Note that, as highlighted in Table 1, FLAT and MEDR cover different key properties of power conservation mechanisms, jointly providing a robust comparison for our approaches.

Experiments are run during hot summer days in July and August 2009 with an average temperature of 89.06°F (31.7°C). In all experiments, the total power consumption is $P_T = 261.95\text{ kW}$, and it is the result of the power consumption of all users. We consider power reduction values between 3% and 9% of P_T , in line with other state-of-the-art power conservation approaches [17, 49]. User bidding behavior, i.e., bids consisting of temperature change and financial compensation are selected from the *models of realistic user behavior* gathered through our online survey. To obtain and provide reliable results, we average the values of each experiment over several runs. Furthermore, we use perfect knowledge of PSP for OPT, and the PSP predictions of power saving (PSP) for the other algorithms. Finally, we explore three experimental scenarios, investigating the impact of different dimensions on the algorithms' performance.

Scenario 1: Varying the Percentage of Auction Participants. In this scenario, we consider a 1 hour auction, and we set the power cap to 95% of P_T . Thus, the required power saving is $P_S = 12.14\text{ kW}$. To achieve this saving, we increase the percentage of users participating in the auction from 40% to 100%. Figure 5 shows the reverse auction performance in terms of objective value (Figure 5(a)), payment (Figure 5(b)), average temperature change (Figure 5(c)), and number of users changing the temperature (Figure 5(d)).

Overall, a clear downward trend of both objective value and payment is reported as the number of participants increases for all approaches, except FLAT. This is due to the fact that all auction-based approaches become more efficient as more bids become available, since better winning bids can be selected. On the contrary, FLAT suffers as more participants become available. This is due to the lack of flexibility of this approach, resulting from the equal power saving imposed to all users and the step-wise 1°F thermostat adjustments. Intuitively, as we increase the percentage of participants, the total payment of FLAT increases linearly, as more and more participating users will have to apply

⁶The main differences between GRAN and GRAN_PC are the bids ranking criterion and the payment rule.




Fig. 5. Performance of reverse auction varying user participation.

a thermostat adjustment. The linear growth is briefly interrupted as soon as there are sufficient participating users to decrease the adjustment to the next lower temperature. This happens, in our experiments, at 70% of participating users, and thus leads to brief interruption of the increasing trend. However, after the thermostat adjustment reaches 1°F for all participating users (as further shown in Figure 5(c)), the total payment increases linearly with the number of participants.

On the other hand, MEDR shows better performance than FLAT, proving the benefits of auction-based approaches versus a flat subscription. However, MEDR underperforms in comparison to DYPS, even though MEDR is solving an NP-hard problem to select the winning bids. This is due to MEDR's limitation of allowing a single bid per user. Conversely, DYPS is able to efficiently make use of multiple bids per user to find a better solution. As expected, DYPS matches the performance of OPT, and outperforms GRAN, by finding the optimal solution of POCO in pseudo-polynomial time. Overall, results show that providing multiple *personalized preferences* can achieve more power conservation with lower payments. Additionally, GRAN is within 30% of the optimal solution found by DYPS, while MEDR and FLAT are within 140% and 860% of that solution, respectively. Finally, it is worth highlighting that although in these results DYPS adopts PSP, it achieves performance comparable to OPT, which instead uses perfect predictions. This is another proof of the quality of the predictions provided by PSP.

In Figure 5(c), we show the average temperature change for winning users, and in Figure 5(d), the number of users who change temperature settings (for FLAT, this refers to all participating users, while for auction-based approaches, it refers to auction winners only). The figures show that, as more participants are available, the average temperature change tends to decrease for all approaches, while the number of users who change temperature tends to increase. FLAT shows the least temperature change, which intuitively results from all participating users changing their temperature, thus requiring less change on average. However, this also results in very high cost, as previously shown. All auction-based approaches benefit from having more participating users, since this allows more flexibility in selecting bids and thus to find more efficient solutions. These approaches tend to select a relatively stable number of winners, as more users are available. This is to be expected, since the power saving required is kept constant. Nevertheless, the presence of more bids allows the selection of better bids, with lower cost. OPT and DYPS tend to select fewer bids, thus resulting in higher temperature changes for winners. GRAN shows a relatively stable temperature change with the participating users. This results from the bids' ranking that gives preference to bids that balance power savings and cost. This is a desirable secondary property, since it makes the system actions more predictable to the participating users over different scenarios.

Scenario 2: Varying the Power Reduction. In the second scenario, we fix the percentage of participants to 60% and decrease the power reduction from 9% to 3% of P_T . We consider again the objective value, payment, temperature change, and number of users changing the temperature. Figure 6 shows the results. FLAT once again performs the worst. Although performance slightly improve with lower power reductions, the approach lacks the flexibility of selecting a subset of the users




Fig. 6. Performance of reverse auction varying percentage of power reduction.




Fig. 7. Performance of reverse auction at different auction durations.

with a minimum temperature change to achieve the power reduction objective. Such inefficiency is particularly evident when the power reductions are smaller (more realistic scenario), since FLAT imposes all users to change their temperature, even when this is not needed. This of course results in a lower temperature change and a constant number of users changing their temperature with respect to the percentage of reduction.

In this scenario, we see a downward trend for all auction-based approaches with respect to all considered metrics, as the percentage of reduction decreases. This is due to lower reductions leading to a lower number of winners. Similar to the previous scenario, DYPS finds the optimal solution to the problem, while GRAN outperforms both GRAN_PC and MEDR. Even in this case, MEDR suffers from the inability of handling multiple bids. Once again, GRAN shows stability even under different percentages of reduction, further supporting the predictability, and thus the acceptance, of this approach. Numerically, GRAN is within 30% of the optimal solution found by DYPS, while MEDR and FLAT are within 140% and 860% of the optimal solution, respectively.

Scenario 3: Duration of the Reverse Auction. In the third scenario, we investigate the impact of the auction duration on the performance of the algorithms. We consider an auction duration of 15 minutes with a bus load $P_T = 60.66$ kW, 30 minutes with a bus load $P_T = 121.37$ kW, and 1 hour with a bus load $P_T = 242.81$ kW. Furthermore, we fix the percentage of participants to 60% and the power reduction to 5%.

In Figure 7, we show the impact that the auction duration has on the objective function and payment. FLAT clearly is penalized by requiring a temperature adjustment to all participating users. The performance of all other algorithms is in line with the previously discussed experiments: DYPS matches the results of OPT and outperforms GRAN, MEDR, and FLAT for all auction durations.

It is interesting to note that the objective value and payments increase non-linearly with the auction duration. This is due to the fact that the transitory nature of the temperature adjustment has non-trivial impact on the HVAC dynamics. Intuitively, during the first few minutes following a temperature adjustment, the HVAC stop cooling the house until the new temperature is reached. Many physical and weather factors affect the temperature dissipation that determines this process, highlighting the importance of modeling individual home-level power dynamics. As a result, the amount of energy to curtail for a 60 minutes auction is much higher than twice the energy of the 30 minutes auction, requiring more bids to be selected, and thus increasing objective value and payments.

9 Conclusion and Future Work

In this work, we present a human-centered HVAC-based power conservation framework based on reverse auctions. Specifically, under this framework, a reverse auction-based approach lets users submit their *personalized preferences* of power conservation, consisting of thermostat temperature adjustments along with a financial compensation for such change. We formulate an optimization problem that selects the winning users of the auction and a payment rule that guarantees truthfulness and individual rationality. We prove that the problem is NP-hard and thus propose an algorithm that finds the optimal solution in pseudo-polynomial time. We employ *models of realistic user behavior* by means of an online survey to gather user bids, while showing user willingness to participate in a similar power conservation system. Further, we are able to predict *realistic home-level power dynamics* through a machine learning prediction algorithm. By using power data from EnergyPlus, the gold standard of energy simulation, we show that our solution outperforms recent state-of-the-art approaches under multiple scenarios.

References

- [1] International Energy Agency. 2009. *World Energy Outlook*. OECD/IEA, Paris.
- [2] Mubbashir Ali, Amir Safdarian, and Matti Lehtonen. 2014. Demand response potential of residential HVAC loads considering users preferences. In *Proceedings of the IEEE PES Innovative Smart Grid Technologies, Europe (ISGT-Europe '14)*. IEEE, 1–6.
- [3] Kadir Amasyali and Nora M. El-Gohary. 2018. A review of data-driven building energy consumption prediction studies. *Renewable and Sustainable Energy Reviews* 81 (2018), 1192–1205.
- [4] Ailin Asadinejad, Mostafa G. Varzaneh, Kevin Tomsovic, Chien-Fei Chen, and Rupy Sawhney. 2016. Residential customers elasticity estimation and clustering based on their contribution at incentive based demand response. In *Proceedings of the IEEE Power and Energy Society General Meeting (PESGM '16)*. IEEE, 1–5.
- [5] Associated Press, Staff (ABC10). 2021. “California Heat, Poor Planning Led to August Power Outages”. Retrieved from <https://www.abc10.com/article/news/local/california/august-power-outages-report/103-267eb366-f08e-47ad-a387-e37d37910597>
- [6] Lawrence M. Ausubel and Paul Milgrom. 2006. The lovely but lonely Vickrey auction. *Combinatorial Auctions* 17, 3 (2006), 22–26.
- [7] Simone Baldi, Christos D. Korkas, Maolong Lv, and Elias B. Kosmatopoulos. 2018. Automating occupant-building interaction via smart zoning of thermostatic loads: A switched self-tuning approach. *Applied Energy* 231 (2018), 1246–1258.
- [8] Jayanta K. Basu, Debnath Bhattacharyya, and Tai-Hoon Kim. 2010. Use of artificial neural network in pattern recognition. *International Journal of Software Engineering and its Applications* 4, 2 (2010): 23–34.
- [9] Shameek Bhattacharjee, Venkata P. K. Madhavarapu, Simone Silvestri, and Sajal K. Das. 2021. Attack context embedded data driven trust diagnostics in smart metering infrastructure. *ACM Transactions on Privacy and Security (TOPS)* 24, 2 (2021), 1–36.
- [10] Shameek Bhattacharjee, Aditya Thakur, Simone Silvestri, and Sajal K. Das. 2017. Statistical security incident forensics against data falsification in smart grid advanced metering infrastructure. In *Proceedings of the 7th ACM on Conference on Data and Application Security and Privacy (CODASPY '17)*. 35–45.
- [11] Liad Blumrosen and Noam Nisan. 2007. *Algorithmic Game Theory. Introduction to Mechanism Design*. Cambridge University Press, New York, NY.

- [12] Enrico Casella, Atieh R. Khamesi, Simone Silvestri, Denise A. Baker, and Sajal K. Das. 2022. HVAC power conservation through reverse auctions and machine learning. In *Proceedings of the IEEE International Conference on Pervasive Computing and Communications (PerCom '22)*. IEEE, 89–100.
- [13] Enrico Casella, Eleanor Sudduth, and Simone Silvestri. 2022. Dissecting the problem of individual home power consumption prediction using machine learning. In *Proceedings of the IEEE International Conference on Smart Computing (SMARTCOMP '22)*. IEEE, 156–158.
- [14] Victoria Cavaliere. 2021. “Why Some Texas Residents are Ending up with \$5,000 Electric Bills After the Winter Storms”. Retrieved from <https://www.businessinsider.com/why-texas-residents-hit-with-soaring-electric-bills-winter-storms-2021-2>
- [15] Archie C. Chapman and Gregor Verbic. 2015. An iterative on-line auction mechanism for aggregated demand-side participation. *IEEE Transactions on Smart Grid* 8, 1 (2015), 158–168.
- [16] David P. Chassin and Daniel Rondeau. 2016. Aggregate modeling of fast-acting demand response and control under real-time pricing. *Applied Energy* 181 (2016), 288–298.
- [17] Jianhai Chen, Deshi Ye, Shouling Ji, Qinning He, Yang Xiang, and Zhenguang Liu. 2019. A truthful FPTAS mechanism for emergency demand response in colocation data centers. In *Proceedings of the IEEE Conference on Computer Communications (INFOCOM '19)*. IEEE, 2557–2565.
- [18] Stefano Ciavarella, Jhi-Young Joo, and Simone Silvestri. 2016. Managing contingencies in smart grids via the internet of things. *IEEE Transactions on Smart Grid* 7, 4 (2016), 2134–2141.
- [19] János Csirik. 1991. Heuristics for the 0-1 min-knapsack problem. *Acta Cybernetica* 10, 1–2 (1991), 15–20.
- [20] Valeria Dolce, Courtney Jackson, Simone Silvestri, Denise Baker, and Alessandra De Paola. 2018. Social-behavioral aware optimization of energy consumption in smart homes. In *Proceedings of the 14th International Conference on Distributed Computing in Sensor Systems (DCOSS '18)*. IEEE, 163–172.
- [21] Erin Douglas. 2021. “Texas was “Seconds and Minutes” Away from Catastrophic Months Long Blackouts, Officials Say”. Retrieved from <https://www.texastribune.org/2021/02/18/texas-power-outages-ercot/>
- [22] Torgeir Ericson. 2009. Direct load control of residential water heaters. *Energy Policy* 37, 9 (2009), 3502–3512.
- [23] Andrew Gelman and Iain Pardoe. 2006. Bayesian measures of explained variance and pooling in multilevel (hierarchical) models. *Technometrics* 48, 2 (2006), 241–251.
- [24] Alyssa Goard. 2021. “Texas’ Power Grid Set a New Winter Peak Demand Record Sunday Evening”. Retrieved from <https://www.kxan.com/news/texas/texas-power-grid-set-a-new-winter-peak-demand-record-sunday-evening/>
- [25] google.com. 2021. Nest Learning Thermostat. Retrieved from https://store.google.com/product/nest_learning_thermostat_3rd_gen?hl=en-US
- [26] LLC Gurobi Optimization. 2020. Gurobi Optimizer Reference Manual. Retrieved from <http://www.gurobi.com>
- [27] Handbook, ASHRAE Fundamentals and Others. 2009. American Society of Heating, Refrigerating and Air-Conditioning Engineers. Inc., Atlanta, GA.
- [28] Kotaro Hara, Abigail Adams, Kristy Milland, Saiph Savage, Chris Callison-Burch, and Jeffrey P. Bigham. 2018. A data-driven analysis of workers’ earnings on Amazon Mechanical Turk. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI '18)*. 1–14.
- [29] International Code Council. 2016. 2015 International Energy Conservation Code—Chapter 4 [Re] Residential Energy Efficiency. Retrieved from https://codes.iccsafe.org/content/IECC2015/chapter-4-re-residential-energy-efficiency#IECC2015_Pt02_Ch04_SecR402
- [30] Benji Jones. 2021. “Texas Blackouts Explained: Arctic Weather Shut Down Power Plants as Demand for Heat Surged, and the State’s Grid is on its Own”. Retrieved from <https://www.businessinsider.com/texas-blackouts-millions-lost-power-in-storm-went-wrong-2021-2>
- [31] Alyson Kenward and Urooj Raja. 2014. Blackout: Extreme weather climate change and power outages. *Climate Central* 10 (2014), 1–23.
- [32] Atieh R. Khamesi, Simone Silvestri, Denise A. Baker, and Alessandra De Paola. 2020. Perceived-value-driven optimization of energy consumption in smart homes. *ACM Transactions on Internet of Things* 1, 2 (2020), 1–26.
- [33] Jin-Ho Kim and Anastasia Shcherbakova. 2011. Common failures of demand response. *Energy* 36, 2 (2011), 873–880.
- [34] Mahmood Kolnegari, Eteazad Moghimi, Ali A. Basiri, Ali T. Qashqaei, and Mandana Hazrati. 2019. A new committee to address the threats of power grids to birds in Iran. *Biodiversity* 20, 4 (2019), 161–164.
- [35] Christos D. Korkas, Simone Baldi, and Elias B. Kosmatopoulos. 2018. Grid-connected microgrids: Demand management via distributed control and human-in-the-loop optimization. In *Advances in Renewable Energies and Power Technologies*. Elsevier, 315–344.
- [36] Christos D. Korkas, Simone Baldi, Iakovos Michailidis, and Elias B. Kosmatopoulos. 2016. Occupancy-based demand response and thermal comfort optimization in microgrids with renewable energy sources and energy storage. *Applied Energy* 163 (2016), 93–104.

- [37] Young M. Lee, Lianjun An, Fei Liu, Raya Horesh, Young T. Chae, and Rui Zhang. 2013. Applying science and mathematics to big data for smarter buildings. *Annals of the New York Academy of Sciences* 1295, 1 (2013), 18–25.
- [38] Department of Energy. 2022. “Electric Disturbance Events (DOE-417)”. Retrieved from <https://www.oe.netl.doe.gov/oe417.aspx>
- [39] U.S. Department of Energy. 2020. Building Energy Codes Program. Retrieved from https://www.energycodes.gov/development/residential/iecc_models
- [40] The Electric Reliability Council of Texas (ERCOT). 2021. “Texas was “Seconds and Minutes” Away from Catastrophic Months Long Blackouts, Officials Say”. Retrieved from <https://www.ercot.com/news/release?id=e2b19f22-7283-3fe8-bed2-a4a979dce772>
- [41] Therese Pepper, Marco Pritoni, Alan Meier, Cecilia Aragon, and Daniel Perry. 2011. How people use thermostats in homes: A review. *Building and Environment* 46, 12 (2011), 2529–2541.
- [42] Thomas Petermann, Harald Bradke, Arne Lüllmann, Maik Poetzsch, and Ulrich Riehm. 2014. *What Happens During a Blackout: Consequences of a Prolonged and Wide-Ranging Power Outage*. Final Report. TAB—Office of Technology Assessment at the German Bundestag, Berlin, Germany.
- [43] Pierre-Olivier Pineau. 1998. Peak load problem, deregulation and reliability pricing. In *Deregulation of Electric Utilities*. Springer, 275–296.
- [44] Atieh R. Khamesi and Simone Silvestri. 2020. Reverse auction-based demand response program: A truthful mutually beneficial mechanism. In *Proceedings of the IEEE International Conference on Mobile Ad Hoc and Sensor Systems (MASS '20)*. IEEE, 1–10.
- [45] Thea Schoeman and Michael Saunders. 2018. The impact of power outages on small businesses in the City of Johannesburg. In *Proceedings of the 10th International Conference on Education, Business, Humanities and Social Sciences Studies (EBHSSS '18)*. DOI : <https://doi.org/10.17758/EARES4.EAP1118411>
- [46] Huajing Sha, Peng Xu, Chonghe Hu, Zhiling Li, Yongbao Chen, and Zhe Chen. 2019. A simplified HVAC energy prediction method based on degree-day. *Sustainable Cities and Society* 51 (2019), 101698.
- [47] Miadreza Shafie-Khah and Pierluigi Siano. 2017. A stochastic home energy management system considering satisfaction cost and response fatigue. *IEEE Transactions on Industrial Informatics* 14, 2 (2017), 629–638.
- [48] Ankit K. Sharma and Akash Saxena. 2019. A demand side management control strategy using Whale optimization algorithm. *SN Applied Sciences* 1, 8 (2019), 1–15.
- [49] Qingxin Shi, Chien-Fei Chen, Andrea Mammoli, and Fangxing Li. 2020. Estimating the profile of incentive-based demand response (IBDR) by integrating technical models and social-behavioral factors. *IEEE Transactions on Smart Grid* 11, 1 (2020), 171–183.
- [50] Qie Sun, Hailong Li, Zhanyu Ma, Chao Wang, Javier Campillo, Qi Zhang, Fredrik Wallin, and Jun Guo. 2015. A comprehensive review of smart energy meters in intelligent energy networks. *IEEE Internet of Things Journal* 3, 4 (2015), 464–479.
- [51] Vladimir Svetnik, Andy Liaw, Christopher Tong, James C. Culberson, Robert P. Sheridan, and Bradley P. Feuston. 2003. Random forest: A classification and regression tool for compound classification and QSAR modeling. *Journal of chemical information and computer sciences* 43, 6 (2003), 1947–1958.
- [52] Samuel Thomas and Jan Rosenow. 2020. Drivers of increasing energy consumption in Europe and policy implications. *Energy Policy* 137 (2020), 111108.
- [53] Moslem Uddin, Mohd F. Romlie, Mohd F. Abdullah, Syahirah A. Halim, and Tan C. Kwang. 2018. A review on peak load shaving strategies. *Renewable and Sustainable Energy Reviews* 82 (2018), 3323–3332.
- [54] U.S. Department of Energy. 2021. “ASHRAE Standard 140 Maintenance and Development”. Retrieved from <https://www.energy.gov/eere/buildings/ashrae-standard-140-maintenance-and-development>
- [55] U.S. Department of Energy. 2021. EnergyPlus. Retrieved from <https://www.energy.gov/eere/buildings/downloads/energyplus-0>
- [56] Jiankang Wang, S Kennedy, and J Kirtley. 2010. A new wholesale bidding mechanism for enhanced demand response in smart grids. In *Proceedings of the IEEE Innovative Smart Grid Technologies (ISGT '10)*. IEEE, 1–8.
- [57] Linqun Zhang, Shaolei Ren, Chuan Wu, and Zongpeng Li. 2015. A truthful incentive mechanism for emergency demand response in colocation data centers. In *Proceedings of the IEEE Conference on Computer Communications (INFOCOM '15)*. IEEE, 2632–2640.
- [58] Qin Zhang and Juan Li. 2012. Demand response in electricity markets: A review. In *Proceedings of the IEEE International Conference on the European Energy Market (EEM '12)*. IEEE, 1–8.
- [59] Ruiting Zhou, Zongpeng Li, and Chuan Wu. 2015. An online procurement auction for power demand response in storage-assisted smart grids. In *Proceedings of the IEEE Conference on Computer Communications (INFOCOM '15)*. IEEE, 2641–2649.

Received 1 May 2023; revised 14 March 2024; accepted 28 March 2024