GreenThrift: Optimizing Carbon and Cost for Flexible Residential Loads

PHUTHIPONG BOVORNKEERATIROJ, University of Massachusetts Amherst, Amherst, United States
WALID A. HANAFY, University of Massachusetts Amherst, Amherst, United States
DAVID IRWIN, University of Massachusetts Amherst, Amherst, United States
PRASHANT SHENOY, University of Massachusetts Amherst, Amherst, United States

Reducing buildings' carbon emissions is an important sustainability challenge. While scheduling flexible building loads has been previously used for a variety of grid and energy optimizations, carbon footprint reduction using such flexible loads poses new challenges since such methods need to balance both energy and carbon costs while also reducing user inconvenience from delaying such loads. This article highlights the potential conflict between electricity prices and carbon emissions and the resulting tradeoffs in carbon-aware and cost-aware load scheduling. To address this tradeoff, we propose GreenThrift, a home automation system that leverages the scheduling capabilities of smart appliances and knowledge of future carbon intensity and cost to reduce both the carbon emissions and costs of flexible energy loads. At the heart of GreenThrift is an optimization technique that automatically computes schedules based on user configurations and preferences. We evaluate the effectiveness of GreenThrift using real-world carbon intensity data, electricity prices, and load traces from multiple locations and across different scenarios and objectives. Our results show that GreenThrift can replicate the offline optimal and retains 97% of the savings when optimizing the carbon emissions. Moreover, we show how GreenThrift can balance the conflict between carbon and cost and retain 95.3% and 85.5% of the potential carbon and cost savings, respectively.

CCS Concepts: • **Social and professional topics** → **Sustainability;**

Additional Key Words and Phrases: Load scheduling, carbon optimization, cost optimization, smart appliances, demand response, residential energy management, sustainability, and temporal load shifting

ACM Reference Format:
Phuthipong Bovornkeeratiroj, Walid A. Hanafy, David Irwin, and Prashant Shenoy. 2025. GreenThrift: Optimizing Carbon and Cost for Flexible Residential Loads. *ACM J. Comput. Sustain. Soc.* 3, 2, Article 14 (May 2025), 21 pages. https://doi.org/10.1145/3724408

This research is supported by NSF grants 2213636, 2105494, 2211302, 2211888, 2325956 and U.S. Department of Energy Award DE-EE0010143.

Authors' Contact Information: Phuthipong Bovornkeeratiroj, University of Massachusetts Amherst, Amherst, Massachusetts, United States; e-mail: phuthipong@cs.umass.edu; Walid A. Hanafy, University of Massachusetts Amherst, Amherst, Massachusetts, United States; e-mail: whanafy@cs.umass.edu; David Irwin, University of Massachusetts Amherst, Amherst, Massachusetts, United States; e-mail: irwin@ecs.umass.edu; Prashant Shenoy, University of Massachusetts Amherst, Amherst, Massachusetts, United States; e-mail: shenoy@cs.umass.edu.

ACM 2834-5533/2025/05-ART14
https://doi.org/10.1145/3724408

1 Introduction

Buildings account for 30% of global energy consumption and 26% of global energy-related emissions, as per the IEA [15]. Consequently, decarbonizing the building sector has emerged as a critical challenge in our society's transition to a low-carbon future. Traditional methods to reduce a building's carbon footprint have focused on increasing the penetration of renewable energy sources in the electric grid or installing distributed renewable systems, such as rooftop solar, directly on buildings. Other efforts have focused on electrification of gas-based building heating systems or the use of distributed energy storage systems [16, 24, 27]. While these approaches have proven effective, they come with substantial infrastructure investment costs and only partially address the broader decarbonization challenge [6].

In contrast to these supply-side methods, a complementary approach is demand-side carbon footprint optimization, where a building modulates its energy (and carbon) demand over time to optimize its overall carbon footprint. Since the carbon intensity of electric supply is known to vary over time—for example, due to intermittent generation from renewables—such demand-side techniques can schedule flexible building loads or time-shift them to periods of low carbon intensity, thereby performing the same work at a lower carbon footprint. While carbon-aware load scheduling in buildings is a relatively new problem, building load scheduling is well-studied in other contexts.

Scheduling of flexible loads via time shifting has been well studied in other contexts. For example, prior efforts have studied load scheduling techniques to address problems such as peak load shaving and cost optimizations in the presence of variable electricity pricing [8, 19, 20, 23, 33]. Automated demand-response optimization has also explored delaying or time-shifting loads to reduce energy demand during periods of grid stress [2, 6, 34]. More recently, researchers have begun to explore load scheduling for optimizing carbon footprint of buildings or grid loads [24]. While prior load scheduling approaches can provide inspiration for optimizing the carbon footprint of buildings, they cannot be applied directly for two reasons.

First, carbon reduction techniques solely focus on variations in energy's carbon intensity, measured in g·CO_2eq/kWh; however, it does not account for variations in electricity prices, increasing the total electricity cost incurred while reducing carbon emissions. While users may want to reduce their buildings' carbon emissions, they may be unwilling to incur higher electricity bills or may even want to reduce them. This introduces carbon and cost tradeoffs that have not been addressed in prior work.

Second, scheduling techniques that rely on time-shifting flexible, low carbon, or low electricity cost periods can increase user inconvenience since loads, such as laundry cycles or EV charging, take longer to complete. Reducing user inconvenience by mapping user preferences to delays that are tolerable is an important aspect of the usability of such techniques. Such carbon-user convenience tradeoffs have also not been explored in prior work.

Motivated by the above challenges, this article presents GreenThrift, a carbon-aware scheduler for flexible home loads, that utilizes real-time carbon intensity signals from the grid and variable pricing signals to reduce carbon usage while optimizing cost and meeting user constraints. Specifically, scheduling of flexible building loads in GreenThrift considers a three-way tradeoff between carbon, energy cost and user constraints. Through careful scheduling in the presence of real-world constraints, our GreenThrift approach demonstrates that it is possible to achieve meaningful carbon reductions in residential buildings while also optimizing electricity costs and reducing user inconvenience from such time-shifting methods. In designing, implementing, and evaluating GreenThrift, our article makes the following contributions:

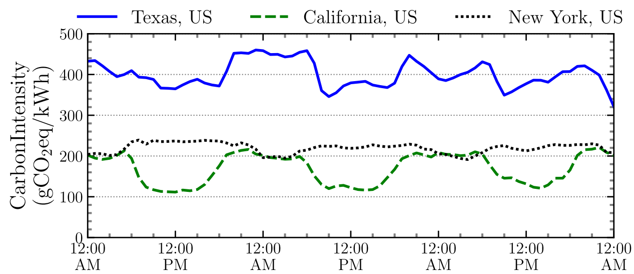

Fig. 1. Carbon intensity across three locations for July 5-7, 2023.

- We present an analysis of the potential conflict between energy's carbon intensity and prices to demonstrate scenarios where reducing carbon emissions may come at the expense of increasing monthly electricity costs.
- We present GreenThrift Algorithm, a joint optimization approach that can optimize carbon usage and electric cost while respecting user constraints. We discuss how to embed this optimization into an online scheduling algorithm that can dynamically time-shift flexible loads in a building while leaving inflexible loads untouched.
- We evaluate the efficacy of GreenThrift using real-world carbon intensity and variable electricity data from different regions of the United States. Our results show that GreenThrift can replicate the offline optimal behavior by retaining 97% of the savings when optimizing the carbon emissions. Lastly, we show how GreenThrift can balance the conflict between carbon and cost and retain 95.3% and 85.5% of the potential carbon and cost savings, respectively.

2 Background

This section provides background on energy's carbon intensity and pricing models for residential consumers. It then explains the role of load shifting in reducing carbon emissions and electricity bills.

2.1 Carbon Intensity

Carbon intensity refers to the amount of greenhouse gas emissions, measured in $g{\cdot}CO_2eq/kWh$, per unit of energy produced. At the grid level, carbon intensity represents a weighted average of the energy sources. For instance, coal-fired power plants typically have high carbon intensity, whereas renewable sources like wind and solar have near-zero carbon emissions. The intermittent nature of renewable energy sources introduces variability in the grid carbon intensity. For example, during the day, solar energy is abundant, which decreases the contributions of fossil-based sources, while at night, utilities may rely more on fossil fuel-based generation, increasing the carbon intensity of the grid.

Figure 1 depicts the three-day carbon intensity in July in Texas, California, and New York. As shown, the carbon intensity highly varies across locations and time of day as the energy source changes. For example, California's carbon intensity is typically low around noon due to its high dependency on solar energy, while Texas has no noticeable pattern due to its high reliance on wind energy. In contrast, New York highly depends on Gas and, hence, has a more stable carbon intensity. In addition, the figure shows locations where users would typically prefer to run their load to optimize their carbon emissions and the expected benefits of these actions. For instance, in

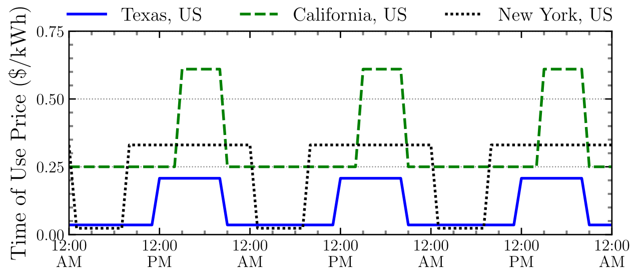

Fig. 2. TOU price across three locations for July 5-7, 2023.

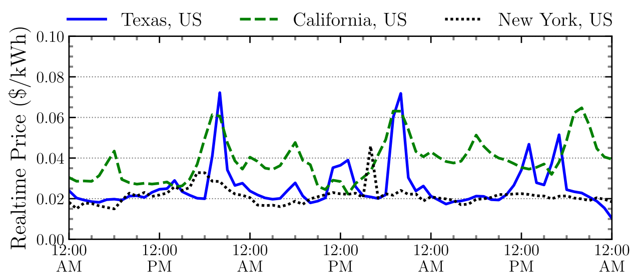

Fig. 3. Real-time price across three locations for July 5-7, 2023.

California, instead of running appliances early in the morning or later in the evening, users can shift their loads to noon and reduce their emissions by almost 2×.

2.2 Electricity Prices

Electricity pricing models for residential consumers typically fall into two categories: flat-rate pricing and dynamic pricing. Flat-rate pricing charges consumers a constant rate per **kilowatt-hour** (**kWh**) regardless of the time of use. While this model is simple and predictable, it does not reflect the true cost of electricity generation, which varies throughout the day. Dynamic pricing models, on the other hand, adjust the cost of electricity based on real-time demand and supply conditions.

Time-of-Use (**TOU**) pricing charges higher rates during peak demand periods (on-peak prices) and lower rates during off-peak periods (off-peak prices). Figure 2 depicts the TOU prices for Texas, California, and New York [7, 11, 28]. As shown, the price differs greatly between times of day. For instance, in New York, peak prices are applied between 8 AM and 12 AM, while in California and Texas, peak prices are for a few hours. Nonetheless, in all cases, the differences between on-peak and off-peak highly encourage users to shift their loads, which can result in up to 5.8×, 2.44×, and 14.8× cost savings for Texas, California, and New York, respectively.

Although TOU pricing is the most common approach for residential homes, another approach is **Real-Time Pricing** (**RTP**), which takes this a step further by varying the price of electricity on an hourly basis, reflecting the real-time cost of generating and delivering electricity as per the energy market [10]. Although this pricing model can significantly reduce total energy costs, it is often very dynamic, requiring agile demand-response strategies and higher flexibility. Figure 3 shows the real-time energy prices across Texas, California, and New York based on their respective **Independent System Operator** (**ISO**), ERCOT, CAISO, and NYISO. As shown, energy prices highly fluctuate,

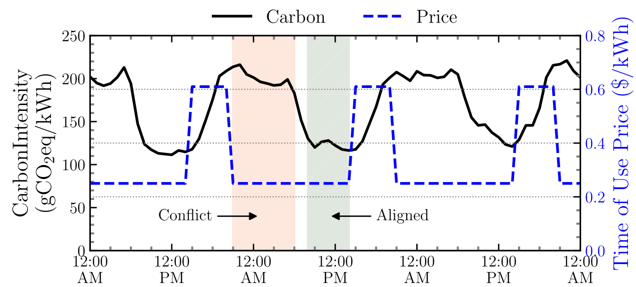

Fig. 4. Relation between carbon intensity and TOU prices during July 5-7, 2023, in US, California. The graph highlights examples of carbon-price conflict and alignments.

highlighting a cost reduction potential of 7×, 3×, and 3× for Texas, California, and New York, respectively.

2.3 Temporal Load Shifting

To exploit variations in energy prices and prices, researchers have proposed multiple temporal load techniques that move the consumption away from time slots with high energy prices [8, 19, 20, 23, 33]. Similarly, users have explored the potential of temporal shifting to reduce the carbon emissions of residential loads [3]. Shifting consumption can be implemented through various mechanisms, such as scheduling flexible appliances, such as dishwashers or dryers, to run during off-peak hours. Another approach is to utilize energy storage systems like batteries by charging the batteries when energy is cheaper or greener.

2.4 Problem Formulation

Our work considers a house with multiple electricity loads, which we categorize as flexible and inflexible loads. Flexible loads have temporal flexibility and can be delayed to later times. Examples of these loads are dishwashers, washing machines, EV charging, and other loads that users do not typically directly interact with. On the other hand, inflexible loads, such as lights, cookers, HVAC systems, and refrigerators, cannot be shifted as they require direct user interaction. We aim at optimizing carbon emissions and costs by scheduling flexible smart appliances to shift their energy usage to different times by utilizing carbon intensity and cost variations. We assume that users provide a deadline or a waiting limit that GreenThrift can use in load scheduling. To our knowledge, we are the first work to consider the tradeoff in carbon and cost optimizations in residential workloads. In Section 3, we highlight the tradeoffs in implementing carbon and cost-aware load shifting.

3 Motivation and Carbon-Cost Tradeoffs

While exploring variations in the electricity grid's carbon intensity by shifting appliance usage from high to low carbon periods can effectively reduce emissions, previous research did not consider the correlations between the grid's carbon intensity and energy prices. In this section, we utilize real-world carbon intensity and pricing traces to quantify the carbon-cost tradeoffs of temporal load shifting. We note that although we focus on the relation between carbon intensity and TOU prices, the conclusions also apply to real-time prices.

Figure 4 highlights the relation between energy's carbon intensity and TOU prices in California during July 5-7, 2023. The figure highlights that there is no clear correlation between energy's carbon intensity and prices, and they may even strictly oppose each other (see highlighted

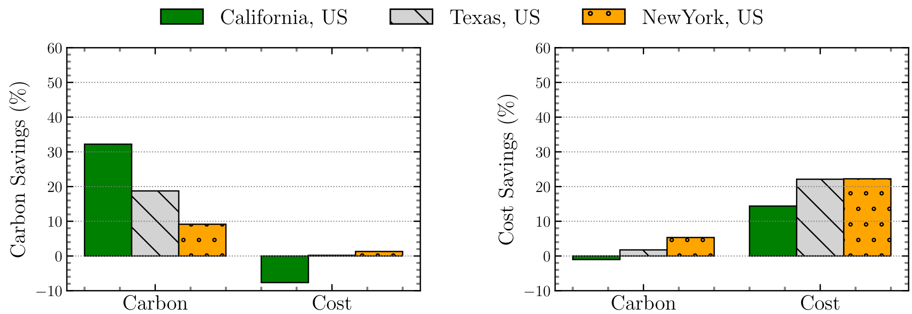

(a) Carbon Savings (%) (b) Cost Savings (%)

Fig. 5. Demonstrating the conflict when scheduling a washing machine (1 kWh load) with a 24 hrs deadline across regions and objectives using TOU prices.

periods). Moreover, this relationship is consistent across different electricity grids, with year-long correlation coefficients between carbon intensity and TOU prices of 0.17, 0.01, and 0.49 for Texas, California, and New York, respectively. Although analyzing the reasons behind this conflict is beyond the scope of this article, one possible reason is that TOU pricing schemes only focus on demand patterns, are used to limit grid-level peak usage, and do not consider the energy source.

We devise a simple example to demonstrate the breadth of the conflict between carbon-aware and cost-aware temporal shifting. In this example, we consider a washing machine that consumes 1kWh of energy and must run for an hour. We repeat this experiment for every hour of the year and schedule the load within the following 24 hours. We utilize two shifting policies: a carbon-aware policy—which selects the time slot with the lowest carbon intensity—and a cost-aware policy—which selects the time slot with the lowest price.

Figure 5 presents the carbon savings (Figure 5(a)) and cost savings (Figure 5(b)) of the carbon-aware and cost-aware policies compared to starting the load immediately across three regions while considering the TOU prices. As shown, carbon-aware temporal shifting can result in more than 50% carbon savings. In contrast, cost-aware temporal shifting can result in more than 20% cost savings, depending on the variations in energy's carbon and price. More importantly, the figure demonstrates how strictly following a cost-aware policy increases carbon emissions, and likewise, following a carbon-aware policy increases the cost. Moreover, the figure highlights how the nuances of the relationship between energy's prices and carbon intensity affect the conflict. For instance, following a carbon-aware scheduler in New York still introduces cost savings, while in California, carbon-aware scheduling increases the total cost.

Key takeaways: *The conflict between energy's carbon intensity and prices motivates multi-criteria temporal shifting strategies that consider variations in carbon intensity and pricing structures.*

4 GreenThrift Design

In this section, we outline the design of GreenThrift, our carbon- and cost-aware home automation software, and highlight key components needed for its functions. Then, we present our offline problem formulation and a scheduling algorithm that considers realistic knowledge assumptions. Lastly, we list an illustrative example of the typical behavior of GreenThrift across scenarios.

4.1 System Architecture

Figure 6 shows the system architecture of GreenThrift where flexible appliances (washing machines, dishwashers, etc.) are scheduled in a carbon- and cost-aware manner. GreenThrift

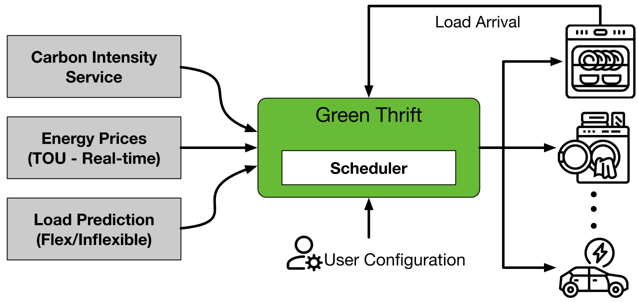

Fig. 6. GreenThrift system architecture.

scheduling decisions depend on many factors, such as carbon intensity, cost, and user preferences, and can be integrated easily into current home automation systems. To schedule flexible home appliances in a carbon- and cost-aware manner, GreenThrift utilizes the following components:

Carbon Information Service. GreenThrift bases its decision on average carbon intensity, which can be realized through third-party integrations with carbon services such as Electricity Maps [9], and CarbonCast [18]. These services provide fine-grained real-time and carbon intensity forecasts at an hour-granularity.

Price Information Service. GreenThrift also relies on electricity prices (e.g., TOU or real-time prices). TOU prices are fixed as part of the contracts with the utility company and usually remain constant; hence, it is straightforward to consider in GreenThrift. On the other hand, real-time prices are variable and require monitoring and predicting local energy markets. Because TOU is the most commonly used pricing scheme, our evaluations will mainly focus on TOU. However, we demonstrate the applicability of our approach and tradeoffs when considering real-time prices in Section 5.6.

Load Prediction. To make safe scheduling decisions, GreenThrift needs information about the expected electricity demand for the upcoming hours. This includes knowledge of the inflexible power load to ensure that scheduling decisions adhere to the maximum capacity of the circuit breaker. Additionally, GreenThrift needs to know the periodicity of the flexible loads to ensure that they are completed before the arrival of the new load.

User Interfaces. As in typical home automation systems, in GreenThrift, the user configures his preferences (e.g., the allowed delay per appliance). GreenThrift interface highlights the possible benefits of different scheduling decisions, such as allowed shifting periods and delays. The user then selects the appliance-specific flexibility based on their preferences.

Scheduling Policy. Lastly, GreenThrift combines the forecasts, user configurations, and current flexible loads to compute a carbon- and cost-aware schedule. In this case, the schedule computes the start time of each appliance, where we assume that loads must run to completion. Although some loads (e.g., batteries or EVs) can be interrupted, we limit ourselves to such use cases.

The following section presents the problem formulation and a scheduling algorithm that acknowledges the impracticalities of the offline formulation.

4.2 Problem Formulation

This section formalizes the problem of carbon- and cost-aware load scheduling. We consider a house with N flexible appliances, L_i loads for appliance $i : i \in N$. We divide the scheduling period

Table 1. GreenThrift Parameters and Decision Variables

Notation	Description
N	$N = \{0, 1, \ldots, n\}$ is the set of appliances.
L_i	L_i is the set of loads for appliance $i : i \in N$.
$e_{i,j}$	Energy consumption per slot of the jth load of appliance i.
$p_{i,j}$	Peak power consumption of the jth load of appliance i.
$l_{i,j}$	Length of the jth load of appliance i.
$a_{i,j}$	Arrival time of the jth load of appliance i.
$d_{i,j}$	Deadline for the jth load of appliance i.[1]
$\mathcal{P}_t$	Price of electricity at time t.
C_t	Carbon intensity at time t.
α	Carbon weight parameter.
β	Cost weight parameter.
$\mathcal{I}_t$	Inflexible load power consumption at time t.
$\mathcal{B}_{\max}$	Breaker peak load.
$s_{i,j,t}$	load j of appliance i starts at time t.
$x_{i,j,t}$	load j of appliance i is running at time t.

where, $i \in N$, $j \in L_i$, and $t \in [1, T]$.

into T discrete intervals of equal length (e.g., one hour) from 1 to T. Table 1 describes the system inputs and utilized decision variables. We formulate the problem as a minimization **Integer linear programming** (**ILP**) problem as follows:

$$\min \sum_{i \in N} \sum_{j \in L_i} \sum_{t \in T} e_{i,j} \times x_{i,j,t} \times (\alpha\, C_t + \beta\, \mathcal{P}_t) \tag{1}$$

s.t.

$$\sum_{t=a_{i,j}}^{d_{i,j}} s_{i,j,t} = 1 \quad \forall i, \forall j \tag{2}$$

$$\sum_{t} x_{i,j,t} = l_{i,j} \quad \forall i, \forall j \tag{3}$$

$$\sum_{j \in L_i} x_{i,j,t} \leq 1 \quad \forall i, \forall j, \forall t \tag{4}$$

$$x_{i,j,t} \leq \sum_{t'=t-l_{i,j}+1}^{t} s_{i,j,t'} \quad \forall i, \forall j, \forall t \tag{5}$$

$$\sum_{i} \sum_{j} p_{i,j} x_{i,j,t} + I_t \leq \mathcal{B}_{\max} \quad \forall t \tag{6}$$

$$s, x \in \{0, 1\} \quad . \tag{7}$$

As shown, the optimization tries to schedule workloads to optimize a parameterized cost function where $\alpha \in [0, 1]$ is the weight of carbon emissions and $\beta \in [0, 1]$ is the weight of electric prices, and is subject to multiple constraints. Equation (2) guarantees that the job only starts once

[1] Note that the deadline is computed per load, depending on the load's arrival time, the next load on the same appliance, load length, and the users' configuration.

ALGORITHM 1: GreenThrift Algorithm

Input: Appliances N, Loads L, New Load $\mathcal{L}$, Inflex Load $\hat{\mathcal{I}}$, Carbon Intensity $\hat{\mathcal{C}}$, Energy Price $\hat{\mathcal{P}}$, Breaker peak load $\mathcal{B}_{\max}$, α, β.
Output: Load start times S.

if $L_{\mathcal{L}.i} != \phi$ **then** ▹ A load exist for this appliance.
  $d_{i,1} \leftarrow l_{i,1}$ ▹ Schedule for now.
end if
$\mathcal{L}.d \leftarrow \min(\mathcal{L}.\lambda, \mathcal{L}.w) + \mathcal{L}.l$
$L.append(\mathcal{L})$
$S \leftarrow$ Solve Optimization$(N, L, \hat{\mathcal{I}}, \hat{\mathcal{C}}, \hat{\mathcal{P}}, \mathcal{B}_{\max}, \alpha, \beta)$
return S

and within the allowed time, i.e., respect the deadline. Equation (3) is the load length constraint, while Equation (4) guarantees that only one load is utilizing the appliances. Equation (5) is the non-interruptibility constraint, where loads must run to completion once started. The constraint in Equation (6) ensures that the peak power consumption of flexible and inflexible loads is within the allowed circuit breaker capacity. Finally, Equation (7) states that the decision variables are binary.

4.3 GreenThrift Algorithm

The problem formulation above relies on detailed knowledge of the electric loads and system capacities. Although some details are known, such as the circuit breaker capacity, or can easily be forecasted, such as energy's carbon intensity [9, 32], some information is burdensome or cannot be known. For instance, in contrast to the carbon intensity that typically follows a diurnal pattern, the real-time electricity prices are much more variable, resulting in higher prediction errors. Another example is inflexible load prediction, which is usually dynamic and can change abruptly. Lastly, in cases where the carbon intensity or future loads are not predictable, GreenThrift can utilize threshold-based approaches [17]; however, evaluating the effectiveness of such approaches is beyond the scope of this article.

To address these challenges, GreenThrift monitors for the arrival of new loads and the quality of its predictors at each time step. GreenThrift then follows an iterative scheduling behavior that dynamically schedules electricity loads when a new load is scheduled or whenever GreenThrift detects an abrupt change in prediction accuracy. In addition, new load arrivals may cause GreenThrift Algorithm to violate some of the scheduling constraints, e.g., a new load may arrive before the current load finishes. Our GreenThrift algorithm makes iterative and incremental modifications to the computed schedule to address such issues. Our experimental evaluation quantifies the impact of these decisions in Section 5.6.

Algorithm 1 lists the GreenThrift Algorithm, where the input contains current loads L, new load $\mathcal{L}$, and predictions such as carbon intensity and inflexible loads. The algorithm also takes the system parameters and configurations. The GreenThrift Algorithm is also executed when the current inflexible load power demand changes beyond the expected range, even if there is no new load. In this case, GreenThrift Algorithm recomputes the schedules for the current appliances. First, when a new load arrives, the algorithm forces the scheduler to start all currently scheduled loads for the corresponding appliance that have not yet begun. The algorithm then computes the deadline of the new load by considering the expected duration between loads λ, the maximum allowed waiting time w, which the users configure, and the load length l.[2] The

[2]Note that we assume that appliances' power consumption and duration are known, as they are typically static.

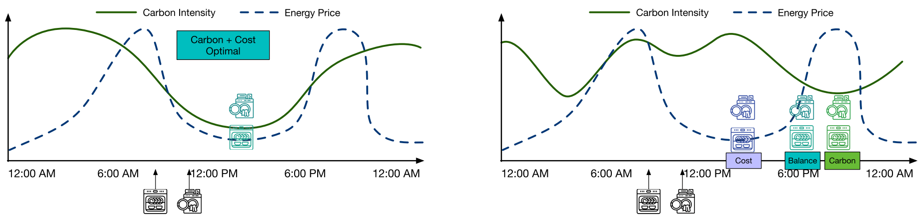

(a) Correlated carbon and price (b) Uncorrelated carbon and price

Fig. 7. An illustrative example of GreenThrift behavior across carbon-prices scenarios.

algorithm then computes the schedule using the optimization approach from Equations (1)–(7) based on the forecasted inputs. Lastly, the algorithm returns the schedule that should be followed by GreenThrift. Thus, by optimizing the schedule as new information becomes available, our algorithm effectively balances the goals of minimizing carbon emissions, reducing costs, and adhering to peak demand constraints, even in a dynamic and uncertain environment.

4.4 GreenThrift in Action

To illustrate the behavior of GreenThrift, we construct two examples of the relation between the electricity grid's time-varying carbon intensity and price, which is typically dynamic, as explained earlier. In both examples, we consider two appliances: A dishwasher and a washing machine, ready at 9 AM and 11 AM, respectively. Figure 7 lists the two cases where the first case (see Figure 7(a)) shows the availability of a slot where carbon and price are low (around noon). In this situation, GreenThrift is not subjected to the carbon-price tradeoff, and it simply can select the time slot where both are low, achieving maximum carbon and cost savings. In contrast, in the second scenario (see Figure 7(b)), carbon intensity and price are not aligned, where users must choose between optimizing for cost or carbon. However, we show that it may be possible to balance the system by finding a compromise between reducing costs and carbon emissions by selecting values for α and β. In the next section, we evaluate the tradeoffs between carbon and cost and show the effect of the weight factors on the results.

5 Evaluation

In this section, we evaluate the performance of GreenThrift in terms of carbon and cost savings. We start by evaluating the performance of the proposed approach across different objectives and scenarios. Next, we illustrate how GreenThrift under different carbon/price dynamics and traces. We then perform a sensitivity analysis of the user's settings. Lastly, we discuss the generalizability and limitations of our approach.

5.1 Experimental Setup

This section describes the real-world traces used to simulate realistic scenarios and assess GreenThrift's performance.

Carbon Intensity Traces. We utilize carbon intensity traces from ElectricityMaps [9]. The traces provide hourly average carbon intensity information, measured in grams of carbon dioxide equivalent per kilowatt-hour (g·CO_2eq/kWh). We utilize carbon intensity for California, Texas, and New York. To emulate the forecast errors, we introduce a uniformly random error to carbon intensity data, denoted as $\hat{C}_{err}$, where err is the mean of added percentage error.

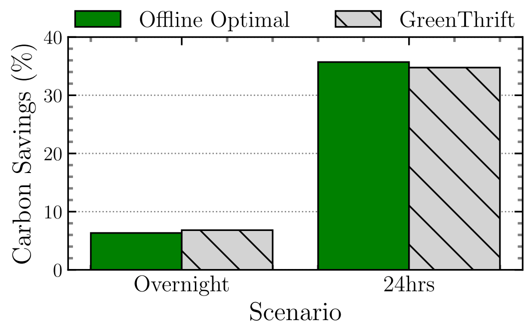

Fig. 8. Carbon savings for different scenarios and assumptions in California, USA.

Electricity Prices Traces. In addition to carbon intensity traces, we utilized electricity price traces for California, Texas, and New York. We utilize TOU and RTP data from utility providers [7, 11, 28], and real-time energy prices from EnergyOnline, a service that provides historical real-time energy prices [10]. We base most of our experiments on the TOU data, which is typically the most common pricing model for houses in the US. Nonetheless, since some states allow residential homes to directly participate in the energy market [26], in Section 5.6, we show the performance of GreenThrift when using real-time prices. Similar to carbon forecasts, we introduce a uniformly random error to real-time price data, denoted as $\hat{\mathcal{P}}_{err}$, where err is the mean of added percentage error.

Residential Load Traces. We use energy consumption traces from households across multiple open datasets(e.g., UMass Smart* dataset [1]). These traces contain various types of loads, including both flexible (e.g., washing machines and dishwashers) and inflexible (e.g., lights, refrigerators, and HVAC) loads, allowing us to simulate different load-shifting policies.

Experimental Settings. We utilize three main variations of GreenThrift by setting the values of $alpha$ and $beta$: a ***carbon-aware policy*** ($alpha = 1, beta = 0$), a ***cost-aware policy*** ($alpha = 0, beta = 1$), and a ***balance policy*** ($alpha = 0.001, beta = 1$). In addition, we evaluate the performance of GreenThrift in two load shifting scenarios, an ***overnight*** load shifting and ***24 hrs*** load shifting. The ***overnight*** is a minimum disruption approach that only shifts overnight workloads (loads are submitted after 6 PM) and ensures they finish before 8:00 AM. In contrast, the ***24 hrs*** scheduling approach allows all workloads to be moved for 24 hrs in the future. Unless otherwise stated, we use carbon intensity traces without errors and report carbon and cost savings from flexible loads. We set the peak constraints for individual houses equal to the peak reported in the original trace. Finally, we implement GreenThrift load shifting policies using Google OR-Tools [25] across different scenarios and policies and use 1hr time step t.

5.2 Carbon-aware Load Shifting

This section evaluates the carbon savings achieved by GreenThrift over a year using different knowledge assumptions and scenarios. Figure 8 compares the carbon savings of GreenThrift to the offline optimal under the 24 hrs and overnight experimental scenarios. The figure shows that the carbon savings potential for the 24 hrs approach is much higher than that of the overnight approach, achieving up to 35.7% and 6.3% for the two scenarios, respectively. This is because the carbon intensity in California is typically lowest during the day, a period that cannot be utilized when workloads are only shifted overnight.

Moreover, the figure highlights the performance similarly of GreenThrift, which depends on realistic knowledge assumptions, with the offline optimal where it achieves carbon savings of 34.7%

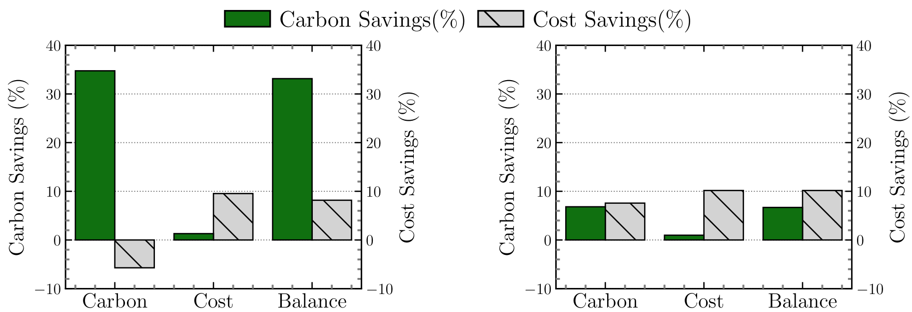

(a) 24hrs (b) Overnight

Fig. 9. Average savings when employing carbon-aware, cost-aware, and balance scheduling across scheduling scenarios in California, USA.

and 6.8% for the 24 hrs and overnight scheduling scenarios, respectively. Note that the heuristic achieves slightly higher carbon savings in the overnight scenario, as it frequently violates the deadline constraints by not finishing the loads before the new load arrives, forcing early morning loads to be slightly shifted.

Key takeaways: *Carbon-aware load shifting can reduce over 35% of carbon emissions from flexible residential workloads. GreenThrift can replicate the offline optimal behavior and retain 97% of the carbon savings.*

5.3 Balancing the Tradeoffs

As highlighted in Section 3, strictly following the carbon- or a cost-aware schedule often yields undesirable effects. Figure 9 shows a single house's carbon and cost savings when employing different scheduling objectives across different scenarios in California. Figure 9(a) highlights the carbon and cost savings and their tradeoffs. For instance, following a carbon-aware schedule ($\alpha = 1, \beta = 0$) can yield up to 34.7% carbon savings. However, it increases the cost of running the flexible load by 5.7%. On the other hand, following the cost-aware schedule ($\alpha = 0, \beta = 1$) yields up 9.5% cost savings while losing 96.2% of the possible carbon savings. Finally, the figure shows that using the balance policy can allow us to co-optimize the carbon emissions and cost and retain 95.3% and 85.5% of the potential carbon and cost savings, respectively. Figure 9(b) shows the behavior of the three policies, where the balance policy can retain 98% and 100% of the possible gains. Nonetheless, in contrast to Figure 9(a), the carbon-aware policy does not yield cost increases, as pushing the workloads to later periods always pushes it away from on-peak price periods.

Figure 10 provides additional insights into how these load-shifting policies impact load distribution throughout the day in a single home from Figure 9(a). Figure 10(a) shows the original demand and highlights the difference between flexible and inflexible load. As shown in the original schedule, most flexible loads (e.g., dishwashers) occur later in the evening, when the cost is low, which limits the possibilities of cost savings (see Figure 9). Figure 10(b) shows how the carbon-aware scheduler moves workloads around noon, when carbon intensity is typically the lowest, while Figure 10(c) pushes the demand away from the peak period, which normally moves the load later in the evening, without introducing many changes. Lastly, Figure 10(d) shows how the balance policy seeks a middle ground, distributing loads across lower carbon and cost periods. Although the result resembles carbon-aware, it avoids the on-peak TOU period at 5 PM. Finally, we note that

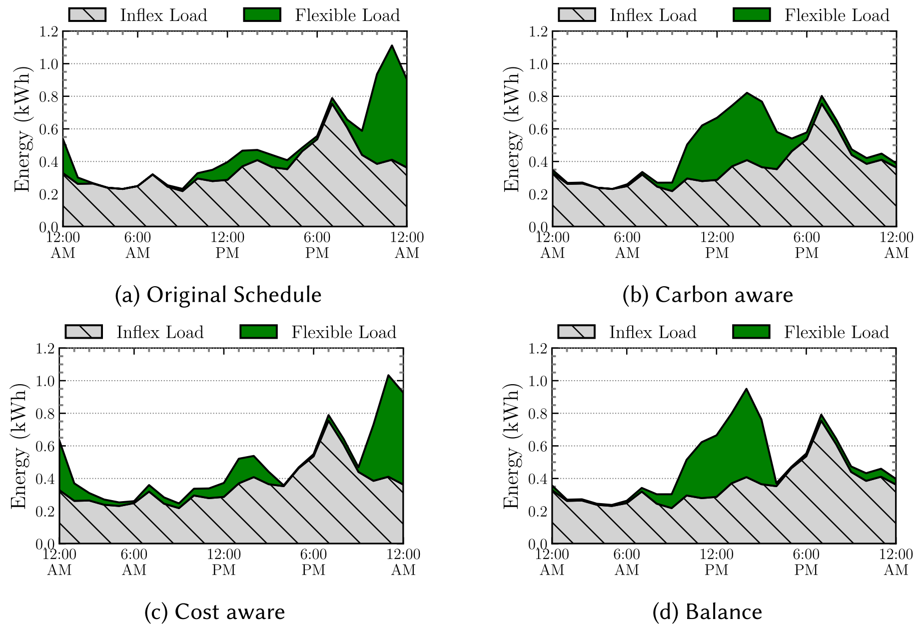

(a) Original Schedule

(b) Carbon aware

(c) Cost aware

(d) Balance

Fig. 10. Average Energy consumption across policies from Figure 9(a).

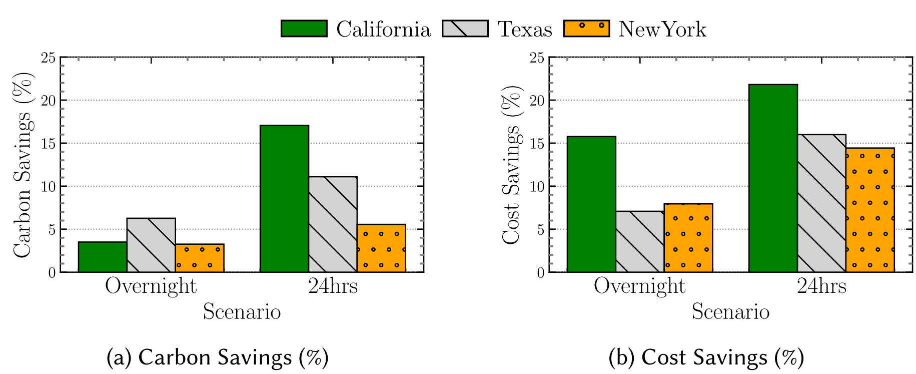

(a) Carbon Savings (%)

(b) Cost Savings (%)

Fig. 11. Average carbon savings and cost savings using balance settings across scenarios in three different regions.

load shifting does not affect expected peak demand, as inflexible loads dominate the energy consumption. In Section 5.5, we show that this conclusion applies when considering multiple houses.

Key takeaways: *GreenThrift can balance the conflict between carbon and cost and retain 95.3% and 85.5% of the potential carbon and cost savings, respectively. Since inflexible loads dominate energy consumption, shifting flexible loads does not increase peak demand.*

5.4 Carbon-Price Dynamics

In this section, we evaluate GreenThrift across multiple carbon-price dynamics. Figure 11 shows the carbon and cost savings achieved using the balance policy across scheduling scenarios in California, Texas, and New York. The results highlight the impact of the electricity pricing structures

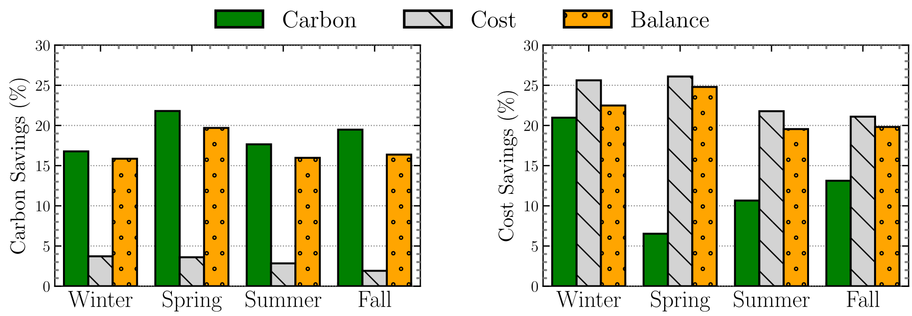

(a) Carbon Savings (%)

(b) Cost Savings (%)

Fig. 12. Carbon and cost savings across seasons using balance settings in the 24 hrs scenario, in California.

and carbon intensity profiles on the potential savings and the effectiveness of load-shifting policies. For instance, GreenThrift can achieve 17%, 11%, and 5% for California, Texas, and New York, respectively, when considering the 24 hrs scenario, which is correlated with the variability of each trace [14]. However, the high reliance on wind energy in Texas, which is typically more available at night, allows the 6 AM to 8 PM policy to retain more than 50% of the carbon savings. At the same time, GreenThrift can achieve 22%, 14%, and 16% cost savings for California, Texas, and New York, respectively, when considering the 24 hrs scenario. The figure also highlights the role of the variability in carbon and cost savings. For example, New York has a gap in the price difference with 25 ¢/kWh off-peak price and 61 ¢/kWh on-peak price during summer, which impacts its cost savings compared to other regions.

Figure 12 further analyzes the behavior of GreenThrift and highlights the carbon and cost savings for different seasons across policies in California. As shown, both the potential savings and the conflict differ across seasons. For instance, in Spring, where carbon intensity is typically more variable in California, the GreenThrift reduces carbon emissions by up 21.8% and cost by up 26.5%. Nonetheless, the figure highlights that Spring has a bigger carbon-cost conflict than other seasons.

Key takeaways: *Carbon intensity and price dynamics, which change across seasons, dictate the possible savings and conflict.*

5.5 Demand Dynamics

Next, we depict the performance of GreenThrift across different houses, representing customers' behaviors. Figure 13 illustrates the carbon and cost savings across multiple houses in California, New York, and Austin. The data shows significant variability in savings within each region, highlighting how differences in household loads and appliance usage patterns can lead to diverse outcomes. For instance, in California, houses can achieve up to 33.1% carbon savings and up to 37% cost savings. At the same time, some houses only achieve 6.6% carbon savings or only 8.1% cost savings.

To understand the reason behind variations, Figures 14 and 15 show the demand pattern of two houses in California before and after load shifting. As shown in the first house (see Figure 14(a)), the demand is typically higher at on-peak and high carbon intensity periods. Thus, load shifts highly influence the demand pattern (see Figure 14(b)), resulting in carbon and cost savings of 25.6% and 36.9%, respectively. On the other hand, Figure 15 demonstrates a scenario where the majority of the load already occurs during off-peak and low-carbon periods. This results in minimal changes in the demand pattern, leading to only 7.6% and 15.6% reductions in carbon and cost, respectively.

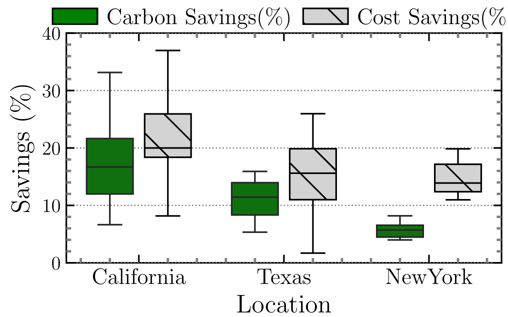

Fig. 13. Carbon savings and cost savings from multiple houses in different regions using the balance policy.

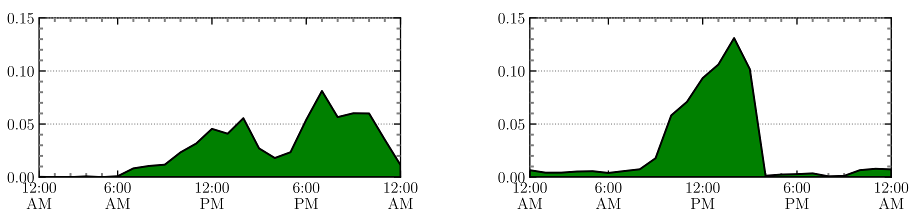

(a) Original

(b) Optimized

Fig. 14. Flexible demand in a house with high carbon and cost savings in California.

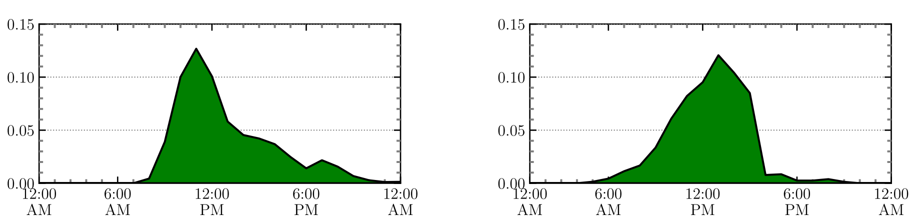

(a) Original

(b) Optimized

Fig. 15. Flexible demand in a house with low carbon and cost savings in California.

Moreover, as highlighted earlier, the savings in Texas and New York are much more limited, where houses achieve a maximum of 15.9% and 33.7% carbon and cost savings and a maximum of 10% and 26.1% carbon and cost savings, for Texas and New York, respectively. Nonetheless, houses exhibit similar variability. Lastly, it's important to note that, on average, most locations offer greater cost savings than carbon savings. This is because the ratio between low and high prices is usually higher than between low and high carbon intensity.

Finally, to analyze the impact of load shifting on overall energy demand, we combine demand from flexible and inflexible loads across households. Figures 16 and 17 depict the average daily energy consumption when combining the load of all the houses in California and New York, respectively. Figure 16(a) and (b) compares the original demand, where the peak typically happens around 6 PM, with the load after shifting. As shown, GreenThrift reduces the peak demand by 3.1%, as it shifts the load to an off-peak period. Similarly, in New York (See Figure 17(a) and (b)),

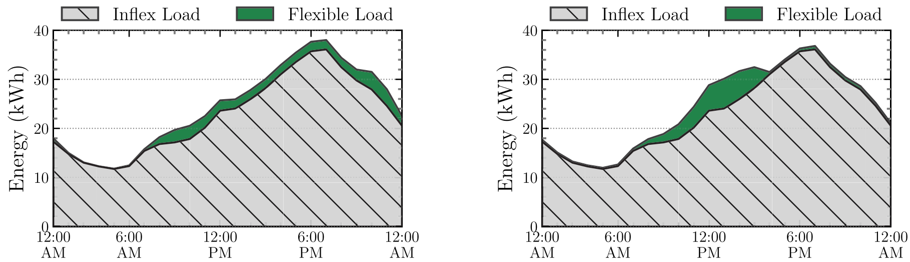

(a) Original (b) Balance Carbon-Cost

Fig. 16. Accumulated average flexible and inflexible load at the transformer level in California.

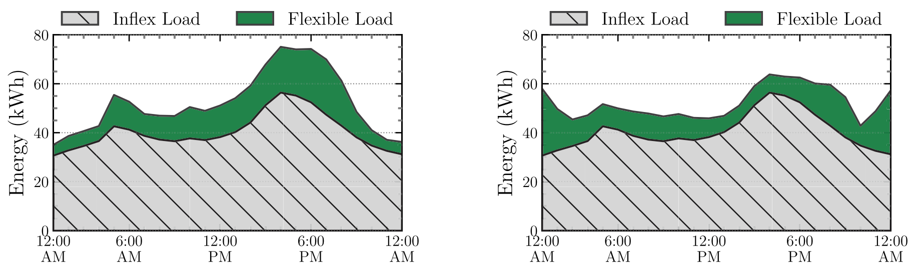

(a) Original (b) Balance Carbon-Cost

Fig. 17. Accumulated average flexible and inflexible load at the transformer level in New York.

GreenThrift is capable of reducing peak demand by 15%. As explained earlier, the reason for this is that inflexible loads dominate total consumption, and moving flexible loads is less likely to increase the peak. At the same time, since the electricity price is typically high at peak load times, shifting workloads in a cost-aware manner also shifts flexible loads away from peak demand slots.

Key takeaways: *The potential carbon and cost reductions from GreenThrift are significantly influenced by demand patterns. The benefits of GreenThrift increase when demand is not aligned with periods of low carbon intensity and prices. Additionally, GreenThrift aims to lower energy costs by shifting loads to off-peak times, thereby reducing overall peak demand.*

5.6 Sensitivity Analysis

Using Real-time Prices: Although TOU is the most common dynamic pricing scheme, some locations allow users to participate directly in the energy market and pay according to real-time prices. Figure 18 shows the behavior of GreenThrift across scheduling objectives, using the same trace as in Figure 9 while allowing loads to be shifted for 24 hrs. Similar to TOU prices, optimizing for a single objective often leads to a bias in the decisions. However, we highlight that in contrast to using TOU pricing, real-time prices allowed the house to save up 45% compared to 9.5% that was saved in Figure 9. Moreover, the figure illustrates that the tradeoff is less pronounced where all objectives result in similar carbon and cost savings.

Impact of extended deadline: To illustrate the relationship between extended deadlines and savings, we evaluate the behavior of GreenThrift under different deadlines. Figure 19 depicts this relation when using the balance policy in California. As shown in contrast to earlier work [14], extending the deadline does not always yield higher savings. For example, in California, extending

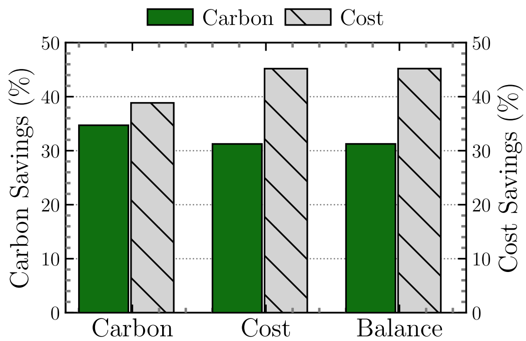

Fig. 18. Carbon and cost savings when using RTP for a single house in California across policies.

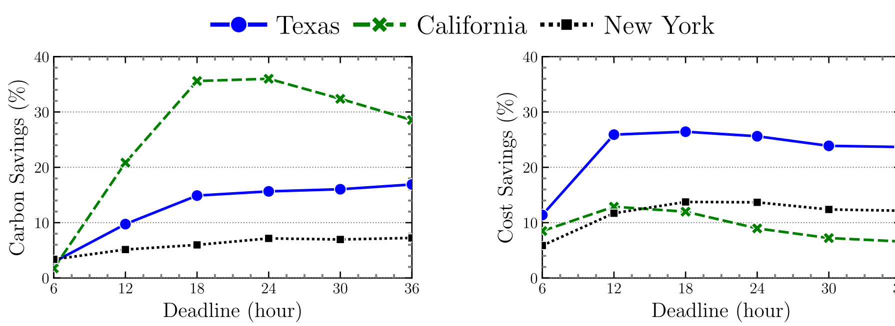

(a) Carbon Savings (%)

(b) Cost Savings (%)

Fig. 19. Carbon and cost savings when employing the balance policy for a single house in California.

the deadline beyond 24 hrs decreases the carbon savings, while extending the deadline beyond 12 hrs decreases the cost savings. This is because when a new load arrives on an appliance with a scheduled load, GreenThrift Algorithm immediately starts the scheduled load, aside from the carbon intensity and cost. The figure also highlights that this behavior is seen across locations, but the magnitude of the change depends on the load and the electricity grid's carbon and price characteristics. Lastly, it highlights how selecting the proper deadline and understanding the load dynamics is vital to maximizing the benefits of GreenThrift.

Carbon Intensity Error: Although researchers have shown that carbon intensity forecasts are often very accurate [18, 32], forecasting errors are still possible. Figure 20 evaluates the performance of GreenThrift when considering carbon forecast errors, which we emulate by adding uniform errors, as explained in Section 5.1. As shown, the effect of errors barely changes the performance of GreenThrift, highlighting the robustness of our scheduling approaches. For instance, even when adding 30% carbon intensity errors, the carbon savings are reduced by only 2%.

Effect of α and β. GreenThrift typically uses a weighted average of the energy's carbon intensity and price. In this subsection, we explore the sensitivity of our load-shifting policies to different values of carbon weight parameter α and cost weight parameter β in terms of carbon and cost savings. Figure 21 shows the effect of α and β on the carbon emissions and savings. As expected, the values of α and β highly influence the carbon and cost savings. For example, when $\alpha = \beta$ or $\alpha > \beta$, GreenThrift behaves as a carbon-aware system as the magnitude of the carbon intensity is

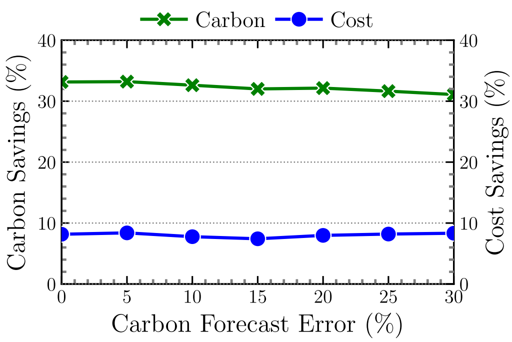

Fig. 20. Effect of Carbon Forecast Errors in California, using the balance policy.

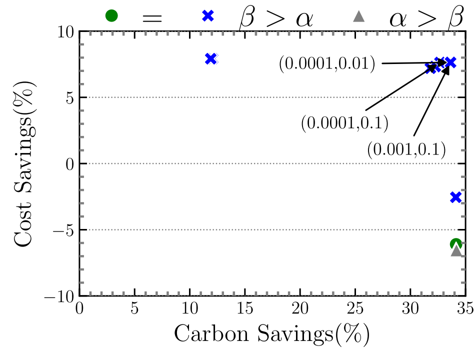

Fig. 21. Changes in carbon and cost savings across different values of α and β in a single house in California.

much higher. However, when $\beta > \alpha$, the system behaves differently based on the ratio between alpha and beta. For example, when $\beta \gg \alpha$, it behaves as a cost-aware, while if $\beta \approx \alpha$, it still behaves as a carbon-aware. Finally, we note that we found many values for α and β that can bring GreenThrift to a balance; we highlight a few of them in the figure.

Key takeaways: *Although GreenThrift configurations and quality of inputs significantly impact the possible carbon and cost savings, the configuration space is easy to navigate, and in case of forecast errors, GreenThrift can retain most benefits.*

5.7 Discussion

We have shown the benefits of GreenThrift in minimizing carbon and costs by exploiting the temporal flexibility in residential loads. Next, we highlight other benefits of the proposed methods and their limitations.

Generalizability of GreenThrift. In this work, we showed the benefits of GreenThrift in reducing carbon emissions and the cost of residential flexible loads. However, flexible loads only represent a small fraction of the total load. However, GreenThrift load-shifting policies can be leveraged to implement holistic reductions using batteries and rooftop solar panels. In this case, each house can compute its own energy carbon intensity and prices and shift load according to cheap and green energy availability. Evaluating the usage of batteries and solar energy and addressing the challenges and tradeoffs in this case is part of our future work.

Limitations. Although load shifting can significantly reduce carbon emissions and costs, GreenThrift makes assumptions that must be addressed in real-world deployments. First, in this work, we assume no dependency between loads. However, in reality, there is some causality between loads. For example, the washer must finish before the dryer, and should complete its cycle when the user is nearby to manually unload and load the machines.

Although these relations can be modeled, they require further user involvement to configure multiple loads simultaneously. Second, our experiments assume that load time and power consumption are fixed and known beforehand. Although this is true for most appliances, some smart appliances have automatic modes that change their behavior with the state. For example, some dryers stop when the load gets dry enough, which requires further profiling of household workloads. Moreover, our experiments assume knowledge of flexible loads demand, which may not be true, and require insights on the typical users' behavior. Lastly, although we show that GreenThrift typically reduces the peak, this was an artifact of our traces, where flexible loads are a minority. Although this is true for most residential buildings, houses with many flexible loads may lead to increases in their peak demand.

6 Related Work

Load shifting is a commonly used technique to manage residential electricity demand. It helps to lower costs, reduce peak consumption, and decrease overall carbon emissions [13]. Researchers have employed load shifting to cut down on energy expenses by shifting electricity usage away from peak demand periods when electricity is typically more expensive [8, 12, 19, 20, 23]. For example, the authors of [23] proposed a scheduling technique for flexible loads to minimize electricity costs while maintaining user comfort. In [20], the authors used batteries to lower all loads' energy costs by charging a battery when energy prices are low and using the stored energy when prices are high. Lastly, the authors of [12] analyzed the challenges in the broad deployment of demand-response techniques. In contrast to these single-objective approaches, GreenThrift focuses on simultaneously optimizing costs and carbon emissions.

Moreover, researchers have demonstrated that load shifting not only helps in cost optimization but also aids in reducing peak demand. For example, [34] conducted simulations to analyze the advantages of load shifting for flexible loads in lowering the peak-to-average ratio. Additionally, [4] [5] illustrated how cooperative load shifting can reduce grid-wide peaks. Furthermore, researchers have also explored the benefits of house-wide load shifting in reducing peak demands. For example, in [2], researchers have utilized load shifting for background loads (e.g., an HVAC unit) to decrease the peak demand while considering users' comfort. In contrast, in [20, 22, 31], authors explored how batteries can help reduce load peaks. Although in GreenThrift, we only focus on flexible loads, peak reduction was not a direct objective. We have demonstrated that cost considerations typically move loads away from peak demand slots, decreasing the average cost.

Researchers have also studied ways to reduce the carbon emissions produced by residential energy usage. For instance, [21, 33] have demonstrated that using local renewable energy sources can replace some of the energy obtained from the grid with carbon-free renewable energy, leading to lower emissions and costs. Furthermore, the authors of [3, 29, 30] analyzed the potential of load shifting in directly reducing carbon emissions by exploiting temporal variability of energy's carbon intensity. In GreenThrift, we consider grids where energy's cost and carbon intensity are variable, unlike other approaches that only consider variations in either.

7 Conclusion

In this article, we analyzed the potential of temporal load shifting of flexible loads to decrease carbon emissions and costs in residential houses. We proposed GreenThrift, an optimization technique that automatically computes schedules based on user configurations and preferences while considering the tradeoffs between energy's carbon intensity and prices. Our results from trace-driven simulations based on real-world traces show that our approach can replicate the offline optimal behavior by retaining 97% of the savings when optimizing carbon emissions. Moreover, we show how GreenThrift can balance the conflict between carbon and cost and retain 95.3% and 85.5% of the potential carbon and cost savings, respectively. In future work, we will analyze the applicability of our load-shifting techniques and the breadth of the carbon-cost tradeoffs in the presence of local renewables and energy storage.

Acknowledgments

We thank the anonymous reviewers and Electricity Maps for access to their carbon intensity data.

References

[1] Sean Barker, Aditya Mishra, David Irwin, Prashant Shenoy Emmanuel Cecchet, and Jeannie Albrecht. 2012. Smart*: An open dataset and tools for EnablingResearch in sustainable homes. In *Proceedings of the 2012 Workshop on Data Mining Applications in Sustainability (SustKDD'12), Beijing, China, August 2012*.

[2] Sean Barker, Aditya Mishra, David Irwin, Prashant Shenoy, and Jeannie Albrecht. 2012. SmartCap: Flattening peak electricity demand in smart homes. In *Proceedings of the 2012 IEEE International Conference on Pervasive Computing and Communications*. 67–75. DOI: https://doi.org/10.1109/PerCom.2012.6199851

[3] Phuthipong Bovornkeeratiroj, Noman Bashir, Vivek Deulkar, Bharathan Balaji, Prashant Shenoy, David Irwin, and Mohammad Hajiesmaili. 2023. Quantifying the decarbonization potential of flexible load. In *Proceedings of the 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys'23)*. 429–433. DOI: https://doi.org/10.1145/3600100.3626346

[4] Phuthipong Bovornkeeratiroj, John Wamburu, David Irwin, and Prashant Shenoy. 2021. VPeak: Exploiting volunteer energy resources for flexible peak shaving. In *Proceedings of the 8th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys'21)*. 121–130. DOI: https://doi.org/10.1145/3486611.3486667

[5] Davide Caprino, Marco L. Della Vedova, and Tullio Facchinetti. 2014. Peak shaving through real-time scheduling of household appliances. *Energy and Buildings* 75 (2014), 133–148. DOI: https://doi.org/10.1016/j.enbuild.2014.02.013

[6] Wesley J. Cole, Danny Greer, Paul Denholm, A. Will Frazier, Scott Machen, Trieu Mai, Nina Vincent, and Samuel F. Baldwin. 2021. Quantifying the challenge of reaching a 100% renewable energy power system for the United States. *Joule* 5, 7 (2021), 1732–1748.

[7] Con Edison. 2024. Time-of-Use Rates. Retrieved March 30, 2025 from https://www.coned.com/en/accounts-billing/your-bill/time-of-use

[8] Alessandro Di Giorgio and Francesco Liberati. 2014. Near real time load shifting control for residential electricity prosumers under designed and market indexed pricing models. *Applied Energy* 128 (2014), 119–132. DOI: https://doi.org/10.1016/j.apenergy.2014.04.032

[9] Electricity Maps. 2024. Retrieved March 30, 2025 from https://www.electricitymap.org/

[10] Energy Online. 2024. Real-time Energy Prices. Retrieved March 30, 2025 from http://www.energyonline.com/

[11] Entergy Texas. 2024. Residential Service Time-of-Day Rate Schedule. Retrieved September 26, 2024 from https://cdn.entergy-texas.com/userfiles/content/price/tariffs/eti_rs-tod.pdf

[12] Samuel Gyamfi, Susan Krumdieck, and Tania Urmee. 2013. Residential peak electricity demand response–Highlights of some behavioural issues. *Renewable and Sustainable Energy Reviews* 25 (2013), 71–77. DOI: https://doi.org/10.1016/j.rser.2013.04.006

[13] Haider Tarish Haider, Ong Hang See, and Wilfried Elmenreich. 2016. A review of residential demand response of smart grid. *Renewable and Sustainable Energy Reviews* 59 (2016), 166–178. DOI: https://doi.org/10.1016/j.rser.2016.01.016

[14] Walid A. Hanafy, Qianlin Liang, Noman Bashir, David Irwin, and Prashant Shenoy. 2023. CarbonScaler: Leveraging cloud workload elasticity for optimizing carbon-efficiency. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 7, 3, Article 57 (2023), 28 pages. DOI: https://doi.org/10.1145/3626788

[15] International Energy Agency (IEA). 2024. Energy System - Buildings. Retrieved September 26, 2024 from https://www.iea.org/energy-system/buildings

[16] Rishikesh Jha, Stephen Lee, Srinivasan Iyengar, Mohammad H. Hajiesmaili, David Irwin, and Prashant Shenoy. 2020. Emission-aware energy storage scheduling for a greener grid. In *Proceedings of the 11th ACM International Conference on Future Energy Systems (e-Energy'20)*. 363–373. DOI: https://doi.org/10.1145/3396851.3397755

[17] Adam Lechowicz, Nicolas Christianson, Jinhang Zuo, Noman Bashir, Mohammad Hajiesmaili, Adam Wierman, and Prashant Shenoy. 2023. The online pause and resume problem: Optimal algorithms and an application to carbon-aware load shifting. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 7, 3, Article 53 (2023), 36 pages.

[18] Diptyaroop Maji, Prashant Shenoy, and Ramesh K. Sitaraman. 2022. CarbonCast: Multi-day forecasting of grid carbon intensity. In *Proceedings of the 9th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys'22)*. 198–207. DOI: https://doi.org/10.1145/3563357.3564079

[19] Francesco Mancini, Jacopo Cimaglia, Gianluigi Lo Basso, and Sabrina Romano. 2021. Implementation and simulation of real load shifting scenarios based on a flexibility price market strategy–the Italian residential sector as a case study. *Energies* 14, 11 (2021). DOI: https://doi.org/10.3390/en14113080

[20] Aditya Mishra, David Irwin, Prashant Shenoy, Jim Kurose, and Ting Zhu. 2012. SmartCharge: Cutting the electricity bill in smart homes with energy storage. In *Proceedings of the 3rd International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet (e-Energy'12)*. Article 29, 10 pages. DOI: https://doi.org/10.1145/2208828.2208857

[21] Aditya Mishra, David Irwin, Prashant Shenoy, Jim Kurose, and Ting Zhu. 2013. GreenCharge: Managing RenewableEnergy in smart buildings. *IEEE Journal on Selected Areas in Communications* 31, 7 (2013), 1281–1293. DOI: https://doi.org/10.1109/JSAC.2013.130711

[22] Aditya Mishra, David Irwin, Prashant Shenoy, and Ting Zhu. 2013. Scaling distributed energy storage for grid peak reduction. In *Proceedings of the 4th International Conference on Future Energy Systems (e-Energy'13)*. 3–14. DOI: https://doi.org/10.1145/2487166.2487168

[23] Bruno Mota, Pedro Faria, and Zita Vale. 2022. Residential load shifting in demand response events for bill reduction using a genetic algorithm. *Energy* 260 (2022), 124978. DOI: https://doi.org/10.1016/j.energy.2022.124978

[24] Kingsley Nweye, Kathryn Kaspar, Giacomo Buscemi, Tiago Fonseca, Giuseppe Pinto, Dipanjan Ghose, Satvik Duddukuru, Pavani Pratapa, Han Li, Javad Mohammadi, et al. 2024. CityLearn v2: Energy-flexible, resilient, occupant-centric, and carbon-aware management of grid-interactive communities. *Journal of Building Performance Simulation* 0, 0 (2024), 1–22. DOI: https://doi.org/10.1080/19401493.2024.2418813

[25] Laurent Perron and Vincent Furnon. 2024. Google OR-Tools v9.10. Retrieved March 30, 2025 from https://developers.google.com/optimization/

[26] Plug-in Illinois. 2024. Utility Real Time Pricing. Retrieved July, 2024 from https://plugin.illinois.gov/your-available-choices/utility-real-time-pricing.html

[27] Ana Radovanović, Anand Ramesh, Ross Koningstein, David Fork, Wolf Weber, Sangsun Kim, Jim Schmalzried, Jyoti Sastry, Michael Dikovsky, Konstantin Bozhkov, et al. 2016. Powernet for distributed energy resource networks. In *Proceedings of the 2016 IEEE Power and Energy Society General Meeting (PESGM'16)*. 1–5. DOI: https://doi.org/10.1109/PESGM.2016.7742035

[28] Southern California Edison. August 2024. Time-Of-Use Residential Rate Plans (TOU-D-Prime). Retrieved March 30, 2025 from https://www.sce.com/residential/rates/Time-Of-Use-Residential-Rate-Plans

[29] Pia Stoll, Gargi Bag, Judith E. Y. Rossebø, Larisa Rizvanovic, and Mikael Åkerholm. 2011. Scheduling residential electric loads for green house gas reductions. In *Proceedings of the 2011 2nd IEEE PES International Conference and Exhibition on Innovative Smart Grid Technologies*. 1–8. DOI: https://doi.org/10.1109/ISGTEurope.2011.6162805

[30] Jing Wang, Prateek Munankarmi, Jeff Maguire, Chengnan Shi, Wangda Zuo, David Roberts, and Xin Jin. 2022. Carbon emission responsive building control: A case study with an all-electric residential community in a cold climate. *Applied Energy* 314 (2022), 118910. DOI: https://doi.org/10.1016/j.apenergy.2022.118910

[31] Ye Xu, David Irwin, and Prashant Shenoy. 2013. Incentivizing advanced load scheduling in smart homes. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings (BuildSys'13)*. 1–8. DOI: https://doi.org/10.1145/2528282.2528292

[32] Xiaoyang Zhang and Dan Wang. 2023. A GNN-based day ahead Carbon intensity forecasting model for cross-border power grids. In *Proceedings of the 14th ACM International Conference on Future Energy Systems (e-Energy'23)*. 361–373. DOI: https://doi.org/10.1145/3575813.3597346

[33] Ting Zhu, Aditya Mishra, David Irwin, Navin Sharma, Prashant Shenoy, and Don Towsley. 2011. The case for efficient renewable energy management in smart homes. In *Proceedings of the 3rd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys'11)*. 67–72. DOI: https://doi.org/10.1145/2434020.2434042

[34] Fatih Çakıl and İbrahim Gürsu Tekdemir. 2024. A novel probabilistic load shifting approach for demand side management of residential users. *Energy and Buildings* 323 (2024), 114751. DOI: https://doi.org/10.1016/j.enbuild.2024.114751

Received 2 December 2024; revised 2 December 2024; accepted 1 March 2025