# ThermalThrift: Cost Effective Thermal Energy Storage for Load Shifting with Water Heaters

Erin Griffiths
University of Virginia
erg3wb@virginia.edu

Kamin Whitehouse
University of Virginia
whitehouse@virginia.edu

## Abstract

Thermal energy storage (TES) is among the cheapest forms of energy storage and is used to efficiently shift thermal loads such as HVAC. Recent work has begun to explore using water heaters for TES but does not yet address several key challenges, including making TES cost effective for all consumers. This paper presents *ThermalThrift*, a system that learns the water consumption patterns of each household in order to better manage TES standby loss to cost effectively perform thermal load shifting for individual households. We evaluate ThermalThrift by collecting 78 days of *in-situ* hot water usage data from 6 pairs of participants and 12 different time-of-use (TOU) pricing schedules from utility companies. Results indicate that after only 7 days of learning data for a given household, ThermalThrift is able to achieve a 47-62% peak load reduction while reducing cost for consumers up to 25% and never increasing cost for an individual. These results indicate that TES with water heaters can be made cost effective for each household by learning and applying a model of individual water usage patterns.

## Categories and Subject Descriptors

C.3 [**Special-Purpose and Applications-Based Systems**]: Real-time and embedded systems

## General Terms

Design, Experimentation

*Keywords*

Thermal Energy Storage, Standby Loss, Water Heating

## 1 Introduction

Thermal energy storage (TES) is the practice of storing energy in a thermal mass for later use. Of the four main types of energy storage (thermal, mechanical, electrical, and chemical), TES has some of the lowest capital costs ($60/kWh), even compared to Pb-acid batteries ($400/kWh) [7]. However, converting the stored heat back into electricity results in a loss of 40-70% of the energy [7], and so TES is most commonly used for heating and cooling (HVAC) so that the thermal energy is used directly and conversion to electricity is avoided. To achieve this, TES technologies such as stored ice, chilled water, or heat bricks are used to store or remove heat during off-peak hours to meet HVAC needs during peak hours, thereby reducing a building's energy bill and reducing peak load on the energy grid [22, 18, 25]. Recent work has begun to look beyond HVAC and target water heaters for TES due to their similar thermal end use during peak hours in the form of showers, dishwashing, and washing machine use in the morning and afternoon. However, current approaches do not address how to make water heater TES cost effective for consumers given the highly individual nature of hot water usage. Without cost effective TES, consumers may not be incentivized to adopt the technology and the potential 7-19 GWh of storage in the nearly 100 million homes in Europe with hot water tanks remains unused [6, 3]. In comparison, the combined capacity from thermal, battery, compressed air, and flywheel storage currently available totals about 12 GWh worldwide [32]. If it could be made cost effective for consumers, the existing culture of water heaters has the potential to provide large scale grid storage at very low cost.

The main goal in water heater TES is to store enough energy in the tank to facilitate turning off the heating elements for a period during peak hours and using stored thermal energy to service demand [1, 11, 26, 4, 19]. Current studies focus on two main techniques to increase this off period: increasing the size of the tank (or adding additional tanks) and raising the temperature setpoint of the tank [10, 13]. However, there are two key challenges that prevent these approaches from being cost effective. The first challenge is *usage dynamics*: water usage patterns vary greatly from one household to the next. The second challenge is *standby loss*: the loss of heat through the walls of the tank due to imperfect thermal insulation. All TES systems suffer from some standby loss, typically discharging 0.5-1% of their stored energy per day [7], but standby loss is especially problematic for water heaters, which can discharge 11% or more (even the most expensive tanks discharge at least 3%). To make storage cost effective, standby loss must therefore be carefully managed. However, doing so is challenging because of water usage dynamics. For example, storing too much

International Conference on Embedded Wireless Systems and Networks (EWSN) 2017
20–22 February, Uppsala, Sweden

ISBN: 978-0-9949886-1-4

energy on a day with little peak water usage can actually increase a household's energy bill by creating excess standby losses. Similarly, more energy is required to perform load shifting for a hot shower at the end of peak hours than at the beginning because heat must be stored for a longer period. Existing solutions focus exclusively on load shifting, as opposed to consumer costs, and use a static tank size or static setpoint temperature during off-peak hours without explicitly addressing usage dynamics and standby loss [10, 13]. Their results indicate these approaches reduce peak load and total cost on average across many homes, but they do not achieve optimal performance for individual households and can even increase some households' energy costs. In this paper, we extend this work by analyzing whether water heaters can be cost effective TES devices for the individual consumer by dynamically managing standby loss based on individual water usage profiles and real time-of-use pricing schemes.

We present a system that *learns* the water usage patterns of each household in order to create cost effective TES. We call the system *ThermalThrift*. It first builds a statistical model of the household's historical water usage over time to make predictions about future hot water demand. Then, it combines these predictions with a time-of-use (TOU) pricing scheme and a thermal model of the water tank to decide whether and how much thermal energy to store in advance of peak hours. If raising the temperature is expected to increase total cost for the consumer, ThermalThrift does not store energy and effectively reverts to conventional setpoint water heater. ThermalThrift can create cost effective storage in two ways: a *consumer-facing* variant that minimizes the total energy cost for the consumer, favoring storage TOU savings that exceed standby loss costs and a *utility-facing* variant that minimizes peak load such that the cost of energy for the consumer does not exceed a conventional, non-TES water heater. These two variants represent the two main stakeholders: consumers, who may install such a system to respond to TOU prices and save money, and the utility, who wants consumers to install such a system to reduce peak load.

To evaluate ThermalThrift, we collect water usage data from 6 pairs of participants over periods of 2-3 weeks each, for 78 total days. The collection of this empirical data was essential to our evaluation because the goal of ThermalThrift is to provide customized performance for each household, and it therefore cannot be evaluated in simulation using the ASHRAE domestic hot water consumption profile [23] used by other studies [10, 13]. Additionally, we collect 12 different TOU pricing schedules from existing electric utilities. Using these 6 usage datasets and 12 pricing schemes, we analyze ThermalThrift's use of TES for cost effective load shifting. Results indicate that after only 7 days of training data for a given household, ThermalThrift is able to achieve within 5% and 36% of the optimal cost and peak load reduction. The consumer-facing ThermalThrift reduces consumer cost by 25% and peak load by 47% on average, using individualized storage temperatures ranging from 51°C to 93°C. Additionally, the utility-facing ThermalThrift reduces peak load 62% for water heating without increasing energy costs for consumers. For utility companies looking to reduce peak load, ThermalThrift's savings translate to a potential collective 1GWh of peak load reduction for approximately 500,000 homes. Additionally, ThermalThrift never increases costs for the studied households – indicating that individual consumers would be able to use TES cost effectively.

## 2 Background and Related Work

Work related to cost effective TES in water heaters can be roughly categorized into three areas: cost effective TES approaches in HVAC and other applications, leveraging existing stored energy in water heaters, and current TES approaches to load shifting with water heaters.

The foundation of cost effective TES is the design of efficient storage materials that minimize energy loss regardless of when energy is generated or consumed. Such materials are often analyzed with respect to variable renewable energy generation, such as solar or combined heating, cooling and power plants, and physical environment of the buildings or spaces where the storage operates [12, 30, 29, 33]. When the consumption of energy can be controlled–a standard occurrence in buildings with HVAC–techniques such as optimization and model predictive control (MPC) have been shown to save energy, shift peak load, and reduce building operating costs by choosing when to store energy and managing thermal leakage [18, 8, 28, 17]. For example, Oldewurtel *et al.* combined model predictive control with weather predictions to improve energy efficiency in climate control by storing thermal energy in advance of weather changes and Ma *et al.* uses MPC to minimize costs of precooling tank water for a water cooled A/C system subject to TOU pricing on a college campus [21, 16]. However, space heating and cooling is quite different than water heating because the latter is more dependent on occupant behaviors, which are different for each water heater. Effectively using MPC to apply TES for HVAC is very dependent on building an accurate model of the building and having weather information. Occupant usage patterns affect heating and cooling demand, but only by 15-30% at most [15, 5] and therefore existing studies do not account for the effects of dynamic occupancy when shifting peak heating or cooling load. Applying MPC to water heaters also requires a thermal model, but the models aren't as unique as buildings and their HVAC systems. In contrast, the key to success is modeling the occupants' hot water usage patterns. Unlike weather predictions, which can be found online, occupant usage patterns are different in every building and must be learned. In this paper, we evaluate how quickly ThermalThrift can learn occupancy patterns to make TES in water heaters cost effective for consumers.

Several studies have explored the use of thermal energy stored during normal operation of a tank (49-60°C) for load shifting. In these approaches heating is simply turned off during a period of peak hours, without first charging the tank with extra heat, and the tank coasts on whatever energy is already in the tank [1, 11, 26, 4, 19]. This approach is currently used by many utility companies today. However, this approach allows the temperatures to drop below user setpoints when coasting through peak periods. Therefore, it can and often does affect user comfort, e.g. when using the shower during peak hours. Research in this area focuses on minimizing or adjusting the resulting payback period when water

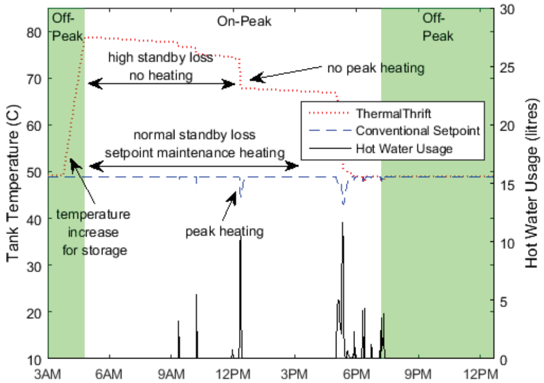


Figure 1: ThermalThrift stores thermal energy in off-peak hours (solid, edges) in preparation for usage during peak hours (white, center). ThermalThrift prevents the purchase of energy during peak hours while higher temperature stored energy is available. A conventional water heater purchases energy during the entire peak period.

heaters turn on at the end of peak hours [2]. Some work uses a fuzzy logic controller to decide when to turn off the heating elements, while Kepplinger et. al. used linear optimization to control water heaters based on stock exchange prices [4, 1, 11]. Since no additional energy is used for TES, cost effective TES is not considered in these approaches.

A recent white paper by EPRI and study by Lacriox were among the first to evaluate water heaters for TES by actively charging the tanks with more energy than is required for normal operation. The tank setpoints were set to arbitrarily high temperatures (92°C and 77°C respectively) during off-peak hours to achieve longer coasting periods [10, 13]. They also studied the use of extra large tanks in order to increase storage capacity [13]. However, these approaches use static tank sizes and temperatures during off-peak hours and do not actively manage standby loss or usage patterns. Because of this, an excess of thermal energy could be stored. This may increase costs for the consumer and render TES not cost effective. This paper is the first to actively create cost effective TES using water heaters by demonstrating that hot water usage patterns can be used to better manage standby loss.

## 3 Water Heater TES Approach

ThermalThrift achieves cost effective TES with water heaters by leveraging TOU pricing to buy cheaper energy during off-peak hours, store it thermally in the tank, and consume it during more expensive peak hours. ThermalThrift stores thermal energy by increasing the tank temperature above a *conventional setpoint* (49-60°C) to a *TES temperature* (60-93°C). Any hot water used during peak hours reduces the tank temperature, eliminating the need to consume additional energy until the temperature drops back to the conventional setpoint. This process is illustrated in Figure 1. This approach requires two related parameters to be defined: the *storage start time* $j_{TES}$ and the *TES temperature* $t_{TES}$. ThermalThrift must select the values of these parameters to shift peak load and be cost effective for the consumer. If $j_{TES}$ is too early or $t_{TES}$ too high, thermal storage costs due to standby loss will outweigh savings during peak hours and cost more to the consumer than conventional setpoint heating. If $j_{TES}$ is too late or $t_{TES}$ too low, the opportunity to shift peak load will not be utilized to the fullest and the energy peak may not be reduced significantly. Therefore, ThermalThrift has two variants that select these parameters to fulfill one of two objectives: minimizing costs for the consumer (*consumer-facing*) or minimizing peak load for the utility while not increasing costs for the consumer (*utility-facing*).

Table 1: The notation used to model ThermalThrift and predict usage, heating, and costs.

| | |
|---|---|
| **Tank Parameters** | |
| $g_{tank}$ | tank water size (L) |
| $s_{tank}$ | tank surface area (m$^2$) |
| $m_{tank}$ | tank water weight (g) |
| $R_{tank}$ | tank insulation value (e.g. R-1.8, m$^2$°C/W) |
| $t_{conv}$ | comfortable user temperature (°C) |
| $p_{tank}$ | heating power of the tank (W) |
| $t_a$ | ambient air temperature (°C) |
| $t_c$ | cold water temperature (°C) |
| $c_w$ | specific heat of water (4.186 J/g°C) |
| $maxheat$ | maximum possible change in temperature (°C/sec) |
| **Input Data** | |
| $U$ | historical hot water usage (L) |
| $\vec{p_{TOU}}$ | vector of electricity prices ($ /Wh) |
| **Prediction Notation** | |
| $S$ | temperature of tank at each interval (°C) |
| $H$ | change in temperature due to heating (°C) |
| $W$ | hot water drawn from the tank (L) |
| $L$ | change in temperature due to thermal losses (°C) |
| $T$ | temperature of tank after losses and mixing (°C) |
| $n$ | length of prediction horizon |
| $m$ | interval where peak hours begin, $m < n$ |
| $\tau$ | time step interval (seconds) |
| **Control Notation** | |
| $c$ | consumer-facing TES start control ({0,1}) |
| $d$ | consumer-facing TES delay control ({0,1}) |
| $t_{TES}$ | TES temperature (°C) |
| $j_{TES}$ | TES start time |

Both ThermalThrift variants optimize the predicted costs of TES, based on historical hot water usage data and a model of tank operation, to select these parameters in a process similar to model predictive control. In the consumer-facing variant, ThermalThrift implicitly finds $j_{TES}$ by evaluating the predicted costs of TES at each time step moving forward. When the predicted cost of TES falls below that of a conventional setpoint heater, storage begins. In the utility-facing variant, ThermalThrift explicitly selects $t_{TES}$ by predicting the cost of a sweep of TES temperatures and choosing the temperature that shifts the most peak load without exceeding

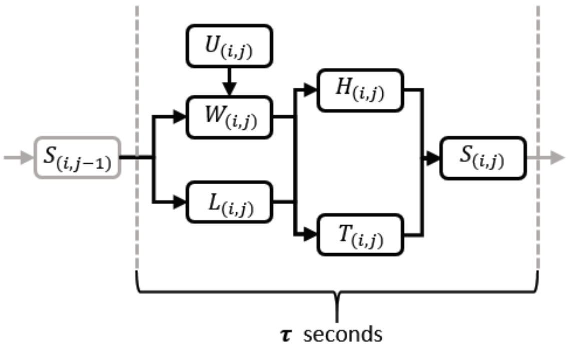


Figure 2: ThermalThrift predicts the state of the tank $S_{i,j}$ though a series of modeling constraints. The predicted hot water use $W_{i,j}$ is calculated from historical use $U_{i,j}$ and the previous tank temperature $S_{i,j-1}$. Thermal losses $L_{i,j}$ are also calculated from $S_{i,j-1}$. These losses $T_{i,j}$ are then used to predict the heating requirement $H_{i,j}$ and the final tank temperature $S_{i,j}$ for the $\tau$ second interval.

the predicted cost of a conventional setpoint heater. ThermalThrift then calculates the storage start time, $j_{TES}$, needed to reach $t_{TES}$ when peak hours begin and starts storing thermal energy at that time.

The details of ThermalThrift's consumer-facing and utility-facing variants are described below. While both approaches use different methods to select $j_{TES}$ and $t_{TES}$, they share common components for predicting heating costs. Hence, the following sections describe the common water heater tank model, derived from the physical properties of water heaters, as well as the novel components specific to the consumer-facing and utility-facing approaches to TES.

## 3.1 Water Heater Tank Model

To predict the cost of TES, ThermalThrift estimates the amount of heating that will be required in the future through a series of modeling constraints. To ensure predicted costs accurately reflect the usage in a household, ThermalThrift estimates the costs based on the historical hot water usage of that household. This usage is stored in a matrix $U$, where each row is a different day of historical data and each column is a time step over the course of that day. Hence, $U_{i,j}$ is the amount of how water drawn from the tank (litres) for use on day $i$ in time step $j$. The main task of prediction is to convert this historical data $U$ into an estimated heating requirement matrix $H$ based on a model of the tank and its operation in either the consumer-facing or utility-facing variant. Depending on the variant, the matrix $H$ can reflect a variety of control decisions (e.g. TES at a specific start time $j_{TES}$ or conventional heating) and the control decision that produces the optimal cost for that heating is selected. Each tank model can be individualized to a home though the tank parameters (size, insulation, surface area, heating power, ambient air temp, cold water temp, etc.), listed in Table 1. The water tank model, represented in equations 1-6, is derived from the physical properties of water heaters and apply to any electrically heated tank.

To determine the predicted heating $H$, ThermalThrift implements six matrices to model the operation of the tank: $S$, $H$, $W$, $U$, $L$, and $T$. The matrices $S$, $W$, $L$, and $T$ represent the predicted tank temperature in each interval, the predicted hot water drawn from the tank, the temperature change due to standby loss, and the tank temperature before heating respectively, with entries corresponding to the matrices $U$ and $H$. Hence, each matrix holds predicted values for all historical days (rows, $i$) and each time period (of $\tau$ seconds) during those days (columns, $j$). The tank's state (i.e. its temperature) at the end of each time period is represented by the matrix $S$. Each entry $S_{i,j}$ is calculated from the previous temperature $S_{i,j-1}$ by calculating the temperature in the tank after losses and usage, $T_{i,j}$, and adding the necessary heating $H_{i,j}$. ThermalThrift calculates $T_{i,j}$ from the temperature change due to standby loss $L_{i,j}$ and the influx of cold water $W_{i,j}$ (of temperature $t_c$) entering the tank due to hot water usage given the size of the tank $g_{tank}$. A flow diagram for the interactions of these matrices is shown in Figure 2. Formally:

$$T_{i,j} = \frac{(S_{i,j-1} - L_{i,j}) * (g_{tank} - W_{i,j}) + W_{i,j} * t_c}{g_{tank}}, \forall i,j \quad (1)$$

$$S_{i,j} = T_{i,j} + H_{i,j}, \forall i,j \quad (2)$$

where $S_{i,0}$ is initialized to the current measured temperature of the tank. The upper tank temperature is bounded below $t_{limit}$ (conservatively 93°C) to prevent boiling by:

$$S_{i,j} <= t_{limit}, \forall i,j \quad (3)$$

All components of equations 1 and 2 ($H$, $L$, and $W$) must be calculated to predict the tank state for the next interval.

The component $L$, the change in temperature due to standby loss, is calculated from the predicted standby loss rate over the time interval of length $\tau$. The loss rate is calculated from the tank surface area ($s_{tank}$), the insulation rating ($R$), the ambient air temperature ($t_a$), the specific heat of water ($c_w$), and the mass of the tank water ($m_{tank}$) as shown in:

$$L_{i,j} = \frac{s_{tank} * (S_{i,j-1} - t_a)}{R * c_w * m_{tank}} * \tau, \forall i,j \quad (4)$$

This equation shows that higher tank temperatures cause proportionally more standby loss, hence why tank setpoints are often kept in the lower thermostatic range (49-57°C). Additionally, this increase in loss at higher temperatures highlights the need for standby loss considerations in TES.

The component $W$ represents the amount of water needed from the predicted tank to fulfill the historical hot water use $U$. If no thermal energy is stored, the predicted usage $W_{i,j}$ equals the historical usage $U_{i,j}$, since any needed hot water is drawn directly from the tank. However, high TES temperatures storing thermal energy (above 60°C) are too hot for direct human use due to potential scalding. To provide comfortable temperatures, thermal storage water is mixed with cold water from the mains before being dispatched to the consumer. Hence, when thermal energy has been stored, $W$ is calculated from the mix of cold water (of temperature $t_c$) with predicted tank water (of temperature $S_{i,j-1}$) to produce a comfortable $t_{conv}$ temperature for the user. Formally:

$$W_{i,j} = \frac{t_{conv} * U_{i,j} - U_{i,j} * t_c}{S_{i,j-1} - t_c}, \forall i,j \text{ when } S_{i,j-1} > t_{conv} \quad (5)$$

While the $S$ equations 1, 2 and 3, the $L$ equation 4, and the $W$ equation 5 are common to both the consumer-facing and utility-facing variants, the heating component $H$ and overall control mechanisms differ in each variant. Hence, the component $H$ and the final optimization for each variant are described in the following sections. To simplify the equations in both variants, we define a parameter *maxheat* to be the maximum rate of change in temperature (°C per second) possible due to heating. *maxheat* is calculated from the water mass ($m_{tank}$) and heating power ($p_{tank}$) of the tank and the specific heat of water ($c_w$):

$$maxheat = \frac{p_{tank}}{m_{tank} * c_w} \tag{6}$$

In both variants, when ThermalThrift stores energy for TES it uses the maximum heating possible (i.e. *maxheat*) to do so, since slower heating provides more time for energy to be lost due to standby. How each variant selects the start time $j_{TES}$ or temperature $t_{TES}$ of thermal storage is describe below.

## 3.2 Consumer-Facing Cost Optimization

The consumer-facing ThermalThrift uses TES to minimize the cost of water heating for the homeowner. Though the goal of this variant is cost reduction, the TOU pricing monetarily rewards reductions in peak energy. Hence, minimizing costs for the consumer also tends to reduce peak energy consumption. However, if the cost of peak power is only slightly larger than off-peak, the increased standby loss due to higher temperatures may mean TES is predicted as not cost effective. In this situation, the consumer-facing ThermalThrift will neither store energy nor shift peak load, but will operate as a conventional water heater.

The consumer-facing approach implicitly selects the start time $j_{TES}$ of storage by evaluating the cost of TES at each time period ($\tau$ seconds) of its operation before peak hours. At each time step, ThermalThrift optimizes the average predicted cost of the heating component $H$ given the TOU price of electricity $\vec{p_{TOU}}$, subject to the tank modeling Equations 2, 3, 4, 5 and its specific constraints on the heating component $H$. The objective function can be formally stated:

$$\min_{c} avg\left(\left(\frac{H * m_{tank} * c_w}{3600}\right) * \vec{p_{TOU}}\right) \tag{7}$$

Each time ThermalThrift predicts costs, the values of component $H$ can represent three different control decisions: store thermal energy now ($j_{TES} = 1$), store it in the next time period ($j_{TES} = 2$), and do not store thermal energy. If $j_{TES} = 1$ is optimal, storage begins immediately. If either $j_{TES} = 2$ or no thermal storage is optimal, ThermalThrift maintains the conventional setpoint until the next time period, where it again evaluates the costs of TES. These three control decisions are represented in the prediction equations for $H$ using two control parameters: the immediate control $c = \{0,1\}$ and the delay control $d = \{0,1\}$. To store thermal energy now, $c$ must be 1. To store energy in the next time period, $c = 0$ and $d = 1$. To never store thermal energy, $c = 0$ and $d = 0$. Only the control decision $c$ is applied to the tank the optimization of the cost of $H$ is complete. While the delay decision $d$ is predicted for the current optimization, the next time step's optimization may make a different decision due to new usage information.

The modeling of $H$ is divided into four equations: heating predictions in the first time step, heating in the second time step, heating from the third time step to the start of peak hours, and heating during peak hours. Storage heating is performed in the first three equations. The fourth equation models the effect of storage on required heating during peak hours. For each time period $j$ in the equations, $H$ must hold either the heating amount required for storage or for maintaining the comfortable setpoint $t_{conv}$. Storage heating is denoted with the maximum possible change in temperature $maxheat * \tau$. Maintaining a comfortable setpoint is denoted with the temperature change $max(0, min(maxheat * \tau, t_{conv} - T_{i,j-1}))$, where the conventional temperature is maintained after temperature loss to the ability of the heating elements. Hence, for the first $H$ equation representing the first interval $j = 1$, either storage heating is performed or the conventional temperature is maintained according to:

$$\begin{aligned} H_{i,1} = {} & c * maxheat * \tau \\ & + (1-c) * max(0, min(maxheat * \tau, t_{conv} - T_{i,1})), \forall i \end{aligned} \tag{8}$$

In the second equation and prediction interval, heating could either continue storage heating from the previous interval ($c = 1$), have delayed the start of TES heating to this interval ($c = 0$, $d = 1$), or be maintaining a conventional temperature ($c = 0$, $d = 0$). Formally:

$$\begin{aligned} H_{i,2} = {} & c * maxheat * \tau + (1-c) * (d * maxheat * \tau \\ & + (1-d) * max(0, min(maxheat * \tau, t_{conv} - T_{i,2})), \\ & \forall i \end{aligned} \tag{9}$$

The third heating equation either continues increasing storage or maintains a conventional temperature until peak hours being. The beginning of peak hours is denoted by the interval $m$, when the price of electricity increases over the current price. Even when the current price is considered ”peak”, if the next price change increases the price then TES will be evaluated. This ensures that even 3-tiered peak systems (off-peak, peak, and super-peak) are evaluated for cost effective TES. Hence, until peak hours are reached at interval $m$, heating follows the equation:

$$\begin{aligned} H_{i,j} = {} & max(c,d) * maxheat * \tau + (1 - max(c,d)) * \\ & (max(0, min(maxheat * \tau, t_{conv} - T_{i,j})), \\ & \forall i, 2 < j < m \end{aligned} \tag{10}$$

Once peak hours being in interval $m$, thermal energy storage ceases since any electricity bought during these periods will not reduce peak consumption and will be more expensive. However, evaluation of heating costs extends past the start of peak hours $m$ to the time step where the price of electricity returns to or falls below the current price (i.e. the prediction horizon), denoted as interval $n$. This ensures that the predicted cost of TES includes the subsequent monetary savings during peak hours. Hence the fourth $H$ equation predicts the cost of only maintaining the conventional setpoint in the period between the start of peak hours $m$ and the end of peak hours $n$. Formally:

$$H_{i,j} = max(0, min(maxheat * \tau, t_{conv} - T_{i,j}),$$
$$\forall i, m \leq j < n \quad (11)$$

We end prediction at interval $n$, often far short of the end of a day, since there will be no monetary savings from the current TES decision after this point – electricity can be bought for immediate use at off-peak prices without the standby loss due to storage. Each time prediction takes place for the new time period, $n$ and $m$ are recalculated to reflect the new price-change times for that prediction. This allows pricing schemes with two peaks (e.g. one in the morning and one in the evening) to have both peaks evaluated for TES.

In the consumer facing approach, ThermalThrift stores thermal energy when costs are predicted to be lower than a conventional setpoint heater. As a byproduct of the TOU pricing, this also tends to shift peak load. For the utility-facing approach, described below, shifting peak load is the focus of prediction – provided TES is still cost effective.

## 3.3 Utility-Facing Peak Optimization

The utility-facing ThermalThrift takes advantage of TES to reduce peak energy consumption at no more cost to the consumer than their daily, non-TES average. The goal of peak load reduction, and not cost minimization, allows ThermalThrift to store more thermal energy than the consumer-facing approach and cover more usage during peak hours – at the cost of higher standby losses and overall costs. To prevent costs from exceeding a conventional heater, the utility-facing approach incorporates a constraint to keep predicted TES costs below the predicted cost of a non-TES water heater to ensure storage is still cost effective.

The utility-facing variant optimizes for the TES temperature, $t_{TES}$, for each peak period during a day. The TES temperature designates the amount of thermal energy to have in storage when peak hours being. Unlike the consumer-facing approach, the prediction evaluation for this control decision is performed only once per peak period – rather than on an interval by interval basis. This allows the utility-facing variant to select the TES temperature that optimizes load shifting, while ensuring costs do not exceed the predicted conventional setpoint heater costs for that entire off-peak and peak period. Prediction is performed either when the day starts, or at the start of the off-peak hours before a peak period. Once the TES temperature is chosen, ThermalThrift calculates the storage start time $j_{TES}$ necessary to reach the temperature and beings energy storage at that time.

To predict what $t_{TES}$ temperature will save the most peak energy while remaining cost effect, ThermalThrift optimizes over a sweep of possible TES temperatures. The TES temperature that is predicted to shift the most energy on average, when evaluated with predictions from $k$ days of historical data, is selected using this objective function:

$$\min_{t_{TES}} \frac{\sum_{i=1}^{k} \sum_{j=m}^{n} H_{i,j} * m_{tank} * c_w}{k} \quad (12)$$

To ensure any chosen $t_{TES}$ is cost effective, the average predicted cost of heating the tank must be below the average predicted cost of operating a conventional setpoint heater. For brevity, the predicted cost of operating a normal water heater on each historical day is denoted $cost_{conv}$. Hence, the main constraint that ensures the selected TES temperature is cost effective is:

$$avg((\frac{H * m_{tank} * c_w}{3600}) * \vec{p_{TOU}}) \leq avg(cost_{conv}) \quad (13)$$

To ensure comfort for the user, $t_{TES}$ must be at or above the comfortable water temperature (e.g. 49 °C):

$$t_{TES} \geq t_{conv} \quad (14)$$

Additionally, the $t_{TES}$ temperature must be reached just before peak hours begin in interval $m$. Formally:

$$S_{i,m} = t_{TES}, \forall i \quad (15)$$

The heating component $H$ can take on values representing each of the $t_{TES}$ temperatures in the temperature sweep subject to the constraints in Equations 13, 14, and 15. The modeling of $H$ is divided into four equations: predicted conventional heating up to the start of storage, heating for storage, ensuring storage exactly reaches $t_{TES}$, and conventional heating during peak hours. Because each historical day of data might require slightly different start times to reach $t_{TES}$, due to hot water usage during energy storage, the storage start time of each predicted day is denoted by the vector $\vec{s}$. Before this start time for each day the tank is predicted as heating only to maintain a conventional temperature in the first $H$ equation:

$$H_{i,j} = max(0, min(maxheat * \tau, t_{conv} - T_{i,j})),$$
$$\forall i, 1 \leq j < \vec{s_i} \quad (16)$$

The next two $H$ equations model the heating between this $\vec{s_i}$ interval for each predicted day and the start of the peak period $m$. First, the tank is heated with *maxheat* until the $m$-1 interval just before peak hours. Then, if less than *maxheat* is required in the $m$-1 interval, only enough heating is used to reach the TES temperature. Formally:

$$H_{i,j} = maxheat * \tau, \forall i, \vec{s_i} \leq j < m - 1 \quad (17)$$
$$H_{i,m-1} = t_{TES} - T_{i,m-1}, \forall i \quad (18)$$

The final $H$ equation for the utility-facing approach models conventional setpoint heating during peak hours. As in the consumer-facing approach, this prediction only lasts to the end of the peak period, $n$, since energy bought after this period will have the same cost as the energy being stored. Formally:

$$H_{i,j} = max(0, min(maxheat * \tau, t_{conv} - T_{i,j})),$$
$$\forall i, m < j < n \quad (19)$$

Given these constraints on the prediction model, the utility-facing variant chooses the optimal TES temperature for load shifting while not increasing costs. If costs cause $t_{TES} = t_{conv}$, then ThermalThrift acts as a conventional water heater and performs no thermal storage. As a byproduct of TOU pricing, reducing peak load also tends to reduce consumer costs and allows ThermalThrift to shift load cost effectively.

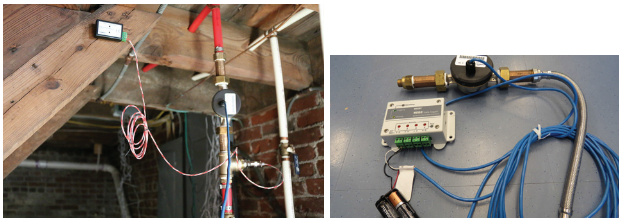

Figure 3: An inline water sensor (bottom) with attached data logger collects second granularity data on water drawn from the hot water tank (top) in the test home.

# 4 Experimental Setup

To assess the potential of ThermalThrift to cost effectively shift peak load given dynamic usage and standby loss in TES, we collected data on real world hot water usage in a test home and TOU pricing from various utility companies across the US. To ensure we had behavioral variation in our water usage data, we collected data from 6 different pairs of participants. We used this historical data and TOU pricing along with our approach to control a water heater modeled by the energy simulation software TRNSYS [20].

## 4.1 Usage Data

Hot water use traces were collected *in-situ* from 6 pairs of participants (G1-6) living in an instrumented test home (5 pairs for 2 weeks, the 6th for 3). The test home had all the hot water fixtures common to a residential home (kitchen sink, 2 bathroom sinks, dishwasher, washing machine, and shower). To ensure the data had the personal usage patterns of the participants, participants were encouraged to live in and use the home as they normally would use their own. The study had IRB approval and each participant received a $100 incentive. Due to a sensor malfunction the third participant pair, study group G3, has only 8 days of collected data. In total, 78 days of usage data is used in this work.

Usage data was collected externally to the water heater as hot water left the tank. A Seametrics SEA Series Turbine Flow Meter, shown in Figure 3, recorded the flow of hot water as it was used by the participants. Flow data was collected using a Hobo Data logger at a sample rate of 1Hz and stored locally. Due to the length of the in-situ studies, and the storage limits of other sensors in the home, data stored on the logger was extracted manually at two week intervals.

The usage data collected from the participants varied widely in the amount of water, type of fixtures, and time of use. For example, bathroom sink usage ranged from 62 to 359 uses across groups and the dishwasher ranged from 0 to 12 uses. Figure 4 shows the average hourly usage for each group. Most of the participants' usage follows expected peak patterns, peaking in the morning and evening. However, personal variation is visible in group G4's usage peak in the early afternoon.

## 4.2 TOU Pricing

Pricing schedules for the evaluation were taken from real world TOU pricing schemes from utilities across the United States. The 12 TOU schemes we selected to represent a diverse set of possible pricing schedules and from different geographic locations. Two power utilities (National Grid and

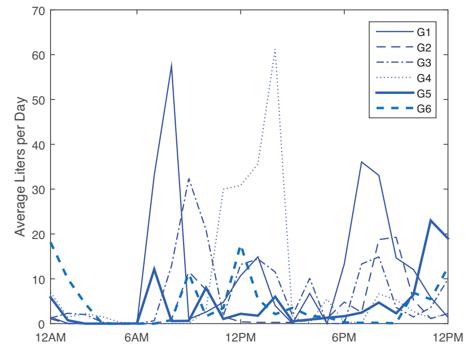


Figure 4: Participants' hot water use generally peaks in the morning and evening, though group G4 used the majority of their hot water around noon.

Texas Utilities) have year round schedules, where pricing does not change according to the season. The 5 other utilities have seasonal pricing, where prices and peak power tiers change. Many of the utilities use a 2 tier peak pricing system, with off-peak and on-peak pricing. Others use a 3 tier system with off-peak, mid-peak, and on-peak pricing (sometimes called off-peak, on-peak, and super-peak). The pricing schedules we chose have a variety of each of these properties. For simplicity, the pricing models are hereby referred to as P1-12 in descending order of average conventional water heater peak load used by the study groups (G1-6). In general, the lower the pricing number (P1), the more hot water averaged across the study groups is used during peak hours. The prices themselves can be seen for the summer in Figure 5 and winter in Figure 6. The peak price increase ranges from $0.02 per kWh (P5) to $0.18 per kWh (P2). The prices and times were obtained from each utility's website.

## 4.3 Baseline and Optimal Algorithms

We compare against two baselines. First, we model a *conventional* water heater that maintains a constant temp of 49°C. The 49°C temperature was chosen as our baseline because it is traditionally recommended as the most energy efficient setpoint for conventional residential water heaters in the U.S. Any lower, and the bacteria Legionella would be capable of growth in the tank. Any higher, and standby loss would increase, thereby increasing the amount of energy used to maintain the temperature. In many European countries and Canada, 60 °C is the regulated setpoint temperature with a required mixer for lower temperature delivery, since the bacteria Legionella dies after 30 minutes at this temperature. Despite this, we chose 49°C as our baseline because of the prevalence of hot water tanks in the U.S. and other countries, it's lower consumption of peak energy due to the lower setpoint, and because, with high temperature TES storage, ThermalThrift could be modified to also consider holding a temp of 60°C or higher for 30 minutes to kill the bacteria – allowing a 49°C setpoint temperature to reduce costs while still providing sanitary tank water. Both the consumer-facing and utility-facing approaches are compared to this baseline. If the consumer-facing approach chooses never perform TES, it duplicates the baseline results. If the

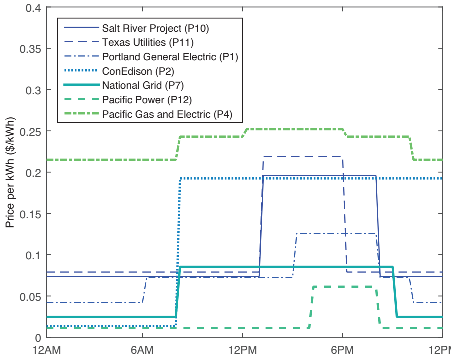


Figure 5: Many summer TOU pricing schedules have peak hours that span the entire daytime period from 8am-9pm.

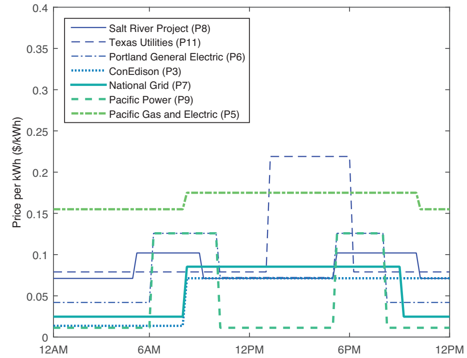


Figure 6: Some winter TOU pricing schedules have two peaks: one in the morning and one in the evening.

utility-facing approach chooses a 49°C TES temperature, it also duplicates the baseline results.

Second, we compare against a *usage agnostic* version of ThermalThrift. This baseline represents approaches that statically choose high TES temperatures for off peak hours. However, the agnostic approach only heats to this temperature just before peak hours to minimize easily avoidable off-peak standby loss. Hence, the usage agnostic approach is a stronger baseline than current static approaches, but still does not account for individual household usage. The highest temperature of 93°C is chosen as the pre-peak temperature to represent the main motivation for this type of TES: load shifting. Like ThermalThrift and the conventional baseline, it maintains a minimum comfortable temperature of 49°C.

The results for both approaches are evaluated against their ThermalThrift *optimal* values. We define optimal by using the day being evaluated as the only historical day in $U$ when an optimization is performed. Essentially, the optimal evaluation has perfect prediction for future hot water usage as it is modeling from oracle data.

### 4.4 Water Heater Simulation

We used the TRNSYS energy simulation software developed by the University of Wisconsin Madison to evaluate ThermalThrift's control on a water heater [20, 31]. We used the Type4a water heater (non-stratified with no hot water inlet) provided with the TRNSYS version 17 library for our simulation. The water heater is modeled after the water heater present in the test home. It uses a 151 litre tank, insulation of $R = 2.6$, $4.5kW$ heating elements, and a conventional or comfortable temperature of $t_{conv} = 49$°C. As the tank was situated in the basement, an ambient air temperature of 10°C is used and cold water entering the tank is set to 10°C. Overall, the tank represents a typical, well insulated residential water heater used in a 1-3 person home.

For the TRNSYS simulation, TES controls are applied to the TRNSYS water heater at each time step according to the two ThermalThrift variants. The electrical use of the TRNSYS water heater and the cost of that usage is recorded. Then, ThermalThrift's control is recalculated by the ThermalThrift variants and applied for the next time step. For our evaluation we chose to use a 2 minute interval since ThermalThrift models the tank as a discrete system. We found no appreciable difference in results at smaller intervals.

## 5 Results

We evaluate cost effectiveness of water heater TES with ThermalThrift by replaying the 2-3 week traces from the participant datasets. Each day is evaluated in order, with the temperature at the end of one day starting the next to ensure any standby loss across multiple days is accounted for. All days in a dataset, except the current day being evaluated, are used as historical data in the matrix $U$ in equation 5 for learning. Each dataset has a different number of days, and hence each dataset uses a different number of historical days in the evaluation. The minimum number of days the approaches learn on is 7 in dataset G3. The maximum number is 19 days for G6. ThermalThrift's results using this historical data for both the consumer-facing and utility-facing variants are called *learned*, with the title of each graph or section specifying the variant. Three evaluation metrics are used throughout: cost, peak load, and total load. *Load* is the amount of kWh drawn by the water heater to heat water. Load can be either *peak* energy, when the power is drawn during a pricing schedule defined peak period, or the *total* energy drawn over the course of the day. *Cost* is the cost of this energy over the course of the entire day. We present results as an average of all participant groups for each of the TOU pricing schedules. For all individual households, ThermalThrift never increased their average daily costs over a conventional setpoint heater.

### 5.1 Consumer-Facing Cost Minimization

The consumer-facing approach of ThermalThrift shows that water heater TES can be used to save both cost and peak load for all pricing schedules when performing load shifting, as shown in Figure 7. Overall, the consumer-facing approach taken by ThermalThrift uses TES to reduce 47% of the peak electric load and 25% of the cost to consumers over a conventional water heater. For the usage agnostic baseline, cost reductions are often comparable to the learned and optimal ThermalThrift, except for three pricing schedules (P4,5,8) where it costs consumers more than a conventional tank to attempt storage without using a learning method to manage standby loss with respect to usage. Additionally, the

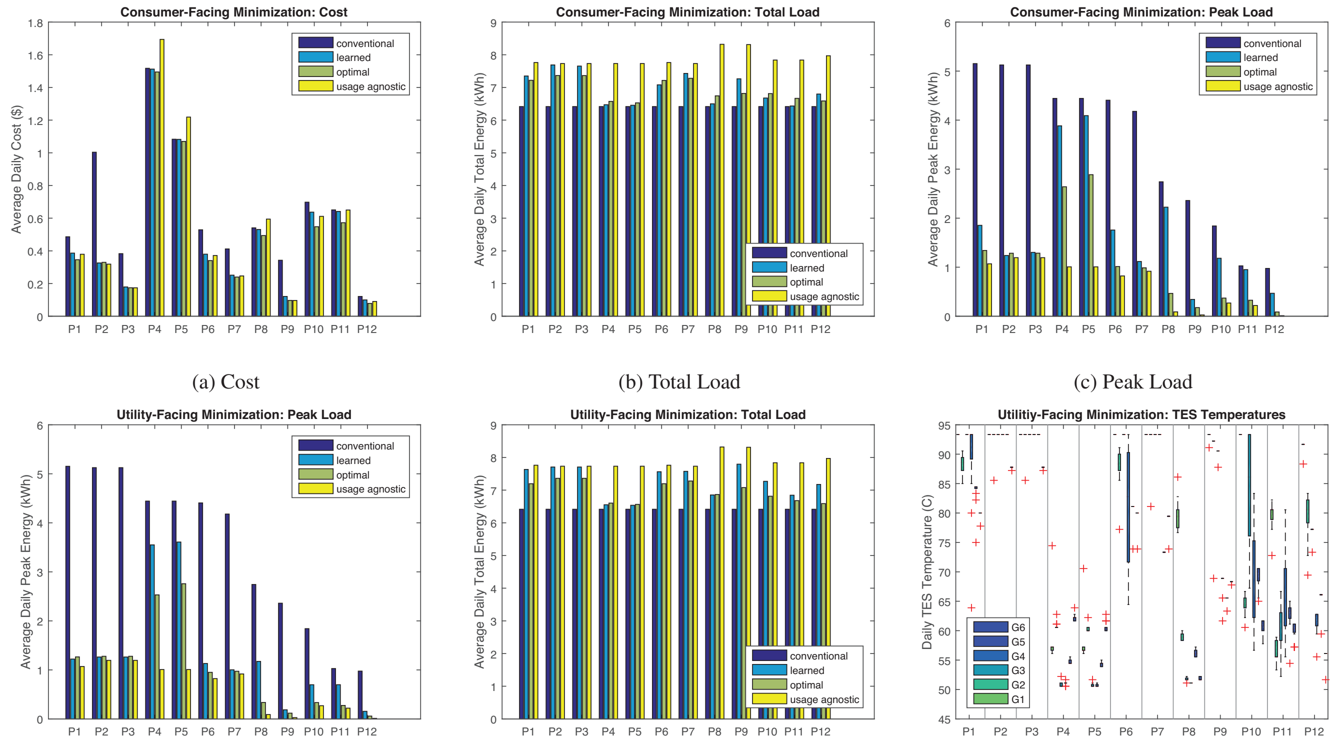


(a) Cost

(b) Total Load

(c) Peak Load

(d) Peak Load

(e) Total Load

(f) TES Temperature

Figure 7: (a-b)Consumer-facing ThermalThrift saves both money and peak energy for every pricing schedule (P1-12). Schedules that have low cost savings tend also to have lower peak energy savings (e.g. P4-5).(d) Utility-facing ThermalThrift saves more peak energy than the consumer-facing variant. (f) The chosen TES temperatures for these savings vary widely across pricing schedules, participant groups, and individual days. The most common TES temperature is the highest possible, 93°C. (b,e) Both variants use more total energy.

usage agnostic approach uses more total energy than ThermalThrift's 8% increase, as unused storage is wasted through standby loss as seen in Figure 7(b). While the usage agnostic approach does save more peak energy, it is unlikely to be adopted by consumers due to these costs and therefore fails to leverage the potential TES of water heaters.

While ThermalThrift reduces costs for all pricing schedules and shows that dynamically managing standby loss is costs effective for load shifting, the total savings vary with TOU pricing schedules. The lowest saving pricing schedules can be categorized into two groups: low relative increase in peak pricing and high relative increase in peak pricing during variable use hours. Pricing models P4, P5, and P8 fit into the first category; they have high peak usage, but low differences between peak and off-peak prices ($0.037, $0.02, and $0.0309). For these pricing schedules the standby loss cost of incorrectly storing too much TES is too large compared to the possible savings, so ThermalThrift performs little storage. Even in the optimal case, where ThermalThrift knows how much hot water will be used during peak hours, the optimal cost improves little over the learned cost due to the low price differential.

In the second category, pricing schedules P10 and P11 have high price differences between peak and off-peak hours ($0.1219 and $0.14), meaning TES standby loss costs could easily be covered by peak savings. However, most usage within the participant groups occurs outside of these hours, causing the learning component of ThermalThrift to predict low amounts of usage during peak hours and store little thermal energy. However, the occasional use of high flow appliances (shower, dishwasher, washing machine) during these hours is costly and consumes a large amount of peak load. Since ThermalThrift had not learned these events from historical data, it did not store energy in preparation for them. This, and the higher savings exhibited by the optimal ThermalThrift for both cost and peak load, highlight the importance of prediction for water heater TES. The current approach uses a simple algorithm to the average cost across all historical days for prediction. The approach could be more complex, weighting predicted costs based on day of the week or similarity to the current day's usage, to achieve additional savings in P10-11. However, overall, the simple averaging approach already reduces 47% of the water heating peak load. Additionally, the algorithm's simplicity allows ThermalThrift to learn usage and operate effectively with only a short learning period (i.e. a few days).

## 5.2 Utility-Facing Peak Optimization

The utility-facing results in Figure 7(d) show an even greater reduction in peak energy over the consumer-facing variant. On average, ThermalThrift reduces 62% of peak

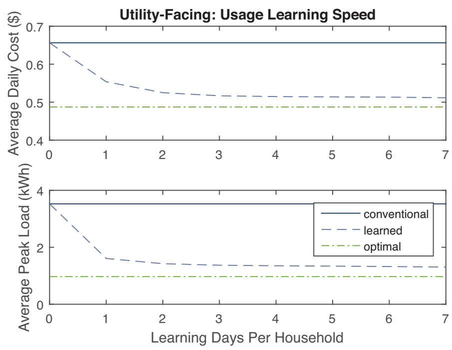


Figure 8: ThermalThrift is able to quickly learn usage patterns and successfully use predictions to manage standby loss. After only 7 days of training data for a given household, ThermalThrift is able to achieve within 5% and 36% of the optimal cost and peak load reduction.

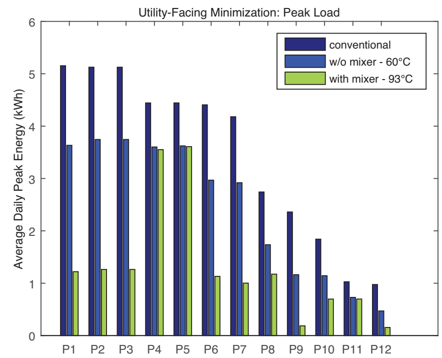


Figure 9: With a human tolerable TES limit of 60°C requiring no mixing valve, ThermalThrift shifts 24% of the peak load. For some pricing schemes no more shifting is possible (e.g. P4-5), but on average a 93°C TES limit shifts far more of the peak load (62%).

load while maintaining costs below that of a conventional, non-TES water heater. Since it maintains the conventional cost, ThermalThrift cannot reach the peak shifting achieved by the aggressively peak shifting usage agnostic approach as seen in Figure 7(d). This is most visible for schedules P4, 5, and 8 where cost constraints limit the learned ThermalThrift's peak savings. However, by maintaining conventional costs the utility-facing ThermalThrift is more likely to be accepted by consumers, facilitating peak load shifting for utilities using water heater TES. Additionally, due to the design of TOU pricing to monetarily incentivize load shifting, the utility-facing approach does still save consumers money as seen in Figure 8.

The chosen TES temperatures for ThermalThrift's savings varied significantly across pricing schedules and participant groups, as shown in Figure 7(e). Some schedules (P2,3,7) had the majority of their daily TES temperatures set to 93°C, indicating that TES is highly cost effective for these participant group/pricing schedule combinations. Three low peak load shifting schedules (P4, 5, 8) had TES temperatures clustered around lower temperatures and show the same saving issues as the consumer-facing approach: relative pricing between peak and off-peak is too low to risk TES. Many of the schedules (P1, 6, 9, 10, 11, 12) had TES temperatures that greatly varied between individual participant groups and across days due to the differences in usage patterns during peak hours, indicating that learning these usage patterns for individual households is necessary to manage standby loss for cost effective TES and load shifting. Overall, ThermalThrift shifted water heating peak load for every pricing scheme we evaluated and reduced 15% more peak load on average than the consumer-facing approach with only a 13% increase in total energy consumption.

## 5.3 Learning Period Analysis

The success of a learning water heater depends in part on how quickly it can learn. Therefore, we also evaluate ThermalThrift with limited learning days in Figure 8. The figure shows the average cost and peak load across all participant groups and TOU schedules for all combinations of learning days in 8 days of data using cross-fold validation. This indicates that after only 7 days of training data for a household, ThermalThrift achieves within 5% and 36% of the optimal cost and peak load reduction for the utility-facing approach, with similar results for the consumer-facing approach. Additionally, ThermalThrift halves the difference between the conventional and optimal baselines with only one day of learning data. This indicates that learning usage patterns can quickly make TES cost effective for real consumers.

## 5.4 TES Potential and the Mixing Valve

The addition of a mixing valve on tanks allows the tank to be heated much higher than human tolerable temperatures and therefore provides a larger opportunity for TES. However, ThermalThrift can still provide TES with only human tolerable temperatures (60°C and below), provided a low comfort temperature such as 49°C is used. Some pricing schedules P4, P5, P8, and P11 have many chosen TES temperatures at or below 60°C even when higher temperatures are available, as shown in Figure 7(f). Thus, for these pricing schemes, the highest potential load shifting is often available cost effectively without exceeding human tolerance limits. Figure 9 shows peak energy savings in a TES limited 60°C and 93°C ThermalThrift. It shows that ThermalThrift could save peak load over a conventional tank even without mixing capabilities. Additionally, a lower temperature limit consumes less energy for TES in the pre-peak period than the 93°C limit, reducing total energy consumption. However, pricing schemes that do benefit from the higher TES temperature limit can shift more than twice as much peak load due to the larger storage capacity. ThermalThrift with a 93°C limit can shift 62% of peak load on average across all pricing schemes, while a 60°C limit shifts only 24%.

## 5.5 TES Power Peak

In managing standby loss by implementing storage just before peak hours, there is the potential for a slightly earlier water heating aggregate peak in a neighborhood. However, aggregate energy peaks across the grid consist of more than just water heating energy and include other peak energy use such as HVAC and lighting. A new peak in water heating use

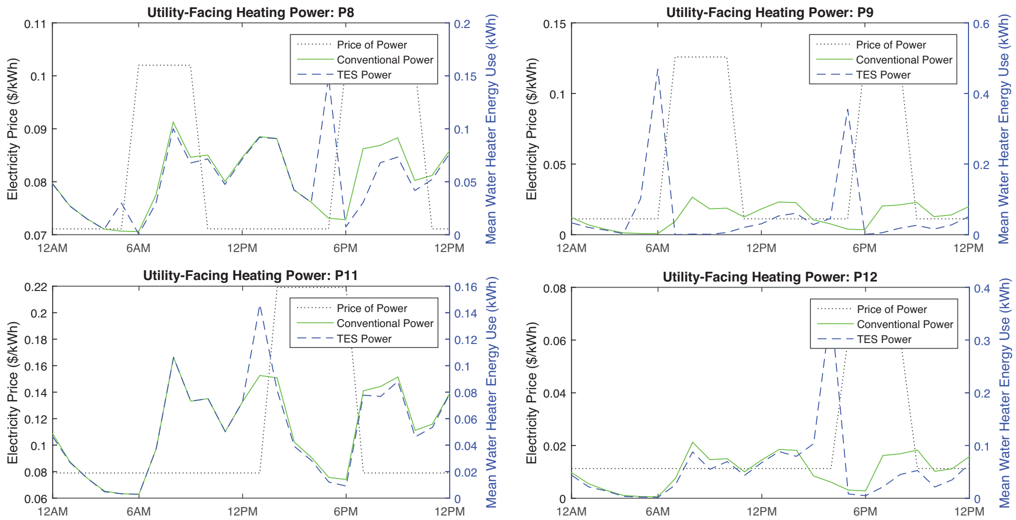


Figure 10: Heating TES just before peak hours to minimize standby loss can create a ”new peak” in water heating just before peak hours. However, the magnitude and time of this peak depends greatly on the pricing scheme. In widespread use, such prices could be adjusted to suppress this peak or move it to times of excess power generation (e.g. renewables).

may simply help fill the valley still present due to other peak energy use and can be used to perform peak leveling. Controlling this leveling may depend on the pricing scheme used. Figure 10 shows four pricing schemes and the shift in heating energy caused by the shifting focused utility-facing variant. In each case, shifted energy use creates a water heating peak just before peak hours. The amplitude and time of this spike depends greatly on the pricing scheme itself. While P8 and P9 have peaks that fill in the valleys of the aggregate water heating use of the study groups, scheme P11 heightens a peak in the middle of the day. Adjustments to the price and time of peak hours, in conjunction with information about the historical data and operation of the learning algorithm, could be used on a neighborhood scale to shift water heating to ideal times. Additionally, neighborhood schemes, such as one similar to that presented in [27] to account for the ”payback peak” when off water heaters are turned back on after peak hours, could constrain TES charging in a neighborhood to reduce or level the new peak. Hence, while such a peak is a concern when the learning algorithm is widely used, it can be mitigated or shifted on a neighborhood scale.

## 6 Limitations and Future Work

ThermalThrift's results indicate that learning hot water usage patterns in each home can produce cost effective TES. This approach can reduce costs for the consumer and reduce peak energy for the utility. Future work must explore whether these results extend across a wider variety of households and buildings, including homes with larger families and a wider variety of household water appliances, or other building types such as laundromats, restaurants, and hotels. Intuitively, buildings with more hot water usage in the beginning of peak hours will experience more savings with ThermalThrift – due to standby losses usage patterns during the ends of peak periods will fare less favorably.

With TES, ThermalThrift maintains a higher temperature during peak hours for a longer period of time than conventional water heaters. Most water activities will be unaffected by more available hot water, such as sink, dishwasher, or washing machine uses, but usage such as showers could potentially be prolonged due to the increased hot water – thereby increasing hot water use during peak hours. Future work must analyzing this potential effect, and how ThermalThrift might adapt to these changes in water use.

Future work may also look at the effect of changing temperatures in a tank over time. For water heaters, higher temperatures can increase sediment buildup, covering and/or wearing out the heating elements, shortening the lifespan. However, the effect of limited and dynamic higher temperatures is not well studied (ThermalThrift only uses high temperatures just before and during peak hours) and ThermalThrift shows savings even with water heater recommended temperatures as the upper TES bound (e.g. 60°C). The effect of these temperature changes on lifetime is an avenue of future work. In terms of temperature, future work may also look at adding a Legionella killing temperature constraint to TES operation if 49°C is maintained as the user comfort minimum. Most Canadian homes have hot water thermostats at 60 °C to kill Legionella at the costs of larger standby losses. However, 60 °C need only be maintained for 30 minutes to kill Legionella and such temperatures could be deliberately achieved periodically with TES.

In addition to thermal load shifting, ThermalThrift could also be used for cost effective storage in other demand-response services including storage of renewable energy and performing frequency regulation. Several works, evaluate water heating for ancillary services to aggregately store highly variable renewable energy, but note that heating demand does not necessarily match excess power [9, 14] and Pourmousavi et. al. use load shifting to match with wind generation but do not account for consumer costs [24]. Combining these approaches with ThermalThrift, which accounts for consumer costs and requires excess energy for TES charging, may provide a cost effective way to direct renewable energy to a consumer who will use it to decrease their own costs – something that will vary significantly based on individual usage and TOU pricing. Along with changing TOU pricing, this may also be used to control the potential for a slightly earlier aggregate peak in a neighborhood from ThermalThrift. Evaluating what response ThermalThrift has on changing peak periods, prices, and renewable energy is a direction for future work.

## 7 Conclusion

In this paper we demonstrate that learning hot water usage patterns for each household can help manage standby loss and ensures hot water TES is cost effective for consumers. Our findings show that the consumer-facing approach, designed to reduce costs, reduces peak electrical load 47% and reduces costs 25% over a conventional water heater. Our utility-facing approach, designed to reduce peak load without increasing consumer costs, reduces peak load 62% over a conventional water heater. Additionally, the standby loss caused by TES only increased total energy use 8% and 13% for the consumer and utility variants respectively. These savings show the large potential impact of TES for water heaters and indicate that homeowners and utilities may be sufficiently incentivized by ThermalThrift's savings to promote the technology for cost and peak load reduction.

## 8 References


[1] P. R. Babu and S. D. V. Prasanna. Intelligent water heater demand controller for consumer satisfaction and energy conservation. *i-Manager's Journal on Circuits & Systems*, 2(4):21, 2014.

[2] A. Belov, N. Meratnia, B. J. van der Zwaag, and P. Havinga. An efficient water flow control approach for water heaters in direct load control. 2014.

[3] BSRIA. Trends in the world traditional and renewable heating markets. http://www.iea.org/media/workshops/2014/buildingwebinars/webinar4/4bsriakrystynapresentationforieafinal.pdf, 2014.

[4] K. Elgazzar, H. Li, and L. Chang. A centralized fuzzy controller for aggregated control of domestic water heaters. In *Electrical and Computer Engineering, 2009. CCECE'09. Canadian Conference on*, pages 1141–1146. IEEE, 2009.

[5] V. L. Erickson, M. Á. Carreira-Perpiñán, and A. E. Cerpa. Observe: Occupancy-based system for efficient reduction of hvac energy. In *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*, pages 258–269. IEEE, 2011.

[6] Eurostat. Number of private households. http://appsso.eurostat.ec.europa.eu/nui/show.do?dataset=lfst_hhnhtych&lang=en, 2015.

[7] A. Evans, V. Strezov, and T. J. Evans. Assessment of utility energy storage options for increased renewable energy penetration. *Renewable and Sustainable Energy Reviews*, 16(6):4141–4147, 2012.

[8] M. Fiorentini, P. Cooper, and Z. Ma. Development and optimization of an innovative hvac system with integrated pvt and pcm thermal storage for a net-zero energy retrofitted house. *Energy and Buildings*, 94:21–32, 2015.

[9] L. Gelažanskas and K. A. Gamage. Distributed energy storage using residential hot water heaters. *Energies*, 9(3):127, 2016.

[10] A. A. Harshal Upadhye, Ron Domitrovic. Evaluating peak load shifting abilities and regulation service potential of a grid connected residential water heater. *ACEEE Summer Study on Energy Efficiency in Buildings*, 2012.

[11] P. Kepplinger, G. Huber, and J. Petrasch. Autonomous optimal control for demand side management with resistive domestic hot water heaters using linear optimization. *Energy and Buildings*, 100:50–55, 2015.

[12] S. Kuravi, J. Trahan, D. Y. Goswami, M. M. Rahman, and E. K. Stefanakos. Thermal energy storage technologies and systems for concentrating solar power plants. *Progress in Energy and Combustion Science*, 39(4):285–319, 2013.

[13] M. Lacroix. Electric water heater designs for load shifting and control of bacterial contamination. *Energy conversion and management*, 40(12):1313–1340, 1999.

[14] K. Lajoie, D. A. Halamay, and T. K. Brekken. Residential water heaters as a grid-scale energy storage solution using model predictive control. In *Technologies for Sustainability (SusTech), 2013 1st IEEE Conference on*, pages 62–69. IEEE, 2013.

[15] J. Lu, T. Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, E. Field, and K. Whitehouse. The smart thermostat: using occupancy sensors to save energy in homes. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, pages 211–224. ACM, 2010.

[16] Y. Ma, F. Borrelli, B. Hencey, B. Coffey, S. Bengea, and P. Haves. Model predictive control for the operation of building cooling systems. *Control Systems Technology, IEEE Transactions on*, 20(3):796–803, 2012.

[17] P.-D. Moroşan, R. Bourdais, D. Dumur, and J. Buisson. Building temperature regulation using a distributed model predictive control. *Energy and Buildings*, 42(9):1445–1452, 2010.

[18] N. Nassif, C. Hall, and D. Freelnad. Optimization of ice thermal storage system design for hvac systems. *Cell*, 336:740–1601, 2013.

[19] NewsChannel 10. Newschannel 10 investigates: Xcel energy saver switch boxes. http://www.newschannel10.com/story/19185280/newschannel-10-investigates-xcel-energy-savor-switch-boxes, 2012.

[20] U. of Wisconsin-Madison. Solar Energy Laboratory and S. A. Klein. *TRNSYS, a transient system simulation program*. Solar Energy Labor-ataory, University of Wisconsin–Madison, 1979.

[21] F. Oldewurtel, A. Parisio, C. N. Jones, M. Morari, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and K. Wirth. Energy efficient building climate control using stochastic model predictive control and weather predictions. In *American control conference (ACC), 2010*, pages 5100–5105. IEEE, 2010.

[22] R. Parameshwaran and S. Kalaiselvam. Energy conservative air conditioning system using silver nano-based pcm thermal storage for modern buildings. *Energy and Buildings*, 69:202–212, 2014.

[23] M. Perlman and B. Mills. Development of residential hot water use patterns. *ASHRAE transactions*, 91(2A):657–679, 1985.

[24] S. A. Pourmousavi, S. N. Patrick, and M. H. Nehrir. Real-time demand response through aggregate electric water heaters for load shifting and balancing wind generation. *Smart Grid, IEEE Transactions on*, 5(2):769–778, 2014.

[25] M. Sebzali, B. Ameer, and H. Hussain. Comparison of energy performance and economics of chilled water thermal storage and conventional air-conditioning systems. *Energy and Buildings*, 69:237–250, 2014.

[26] A. Sepulveda, L. Paull, W. G. Morsi, H. Li, C. Diduch, and L. Chang. A novel demand side management program using water heaters and particle swarm optimization. In *Electric Power and Energy Conference (EPEC), 2010 IEEE*, pages 1–5. IEEE, 2010.

[27] M. Shaad, R. Errouissi, C. Diduch, M. Kaye, and L. Chang. Aggregate load forecast with payback model of the electric water heaters for a direct load control program. In *Electrical Power and Energy Conference (EPEC), 2014 IEEE*, pages 214–219. IEEE, 2014.

[28] J. Širokỳ, F. Oldewurtel, J. Cigler, and S. Prívara. Experimental analysis of model predictive control for an energy efficient building heating system. *Applied Energy*, 88(9):3079–3087, 2011.

[29] N. Soares, J. Costa, A. Gaspar, and P. Santos. Review of passive pcm latent heat thermal energy storage systems towards buildings energy efficiency. *Energy and buildings*, 59:82–103, 2013.

[30] Y. Sun, S. Wang, F. Xiao, and D. Gao. Peak load shifting control using different cold thermal energy storage facilities in commercial buildings: a review. *Energy Conversion and Management*, 71:101–114, 2013.

[31] University of Wisconsin–Madison. Solar Energy Laboratory and Klein, Sanford A. A transient systems simulation program. http://sel.me.wisc.edu/trnsys/, 2013.

[32] Wikipedia. "List of energy storage projects". Wikipedia, The Free Encyclopedia, Accessed online: 24-June-2016.

[33] D. Zhou, C.-Y. Zhao, and Y. Tian. Review on thermal energy storage with phase change materials (pcms) in building applications. *Applied energy*, 92:593–605, 2012.