

# Combined Heat and Privacy: Preventing Occupancy Detection from Smart Meters

Dong Chen, David Irwin, Prashant Shenoy, and Jeannie Albrecht<sup>†</sup>

University of Massachusetts Amherst

<sup>†</sup>Williams College

**Abstract**—Electric utilities are rapidly deploying smart meters that record and transmit electricity usage in real-time. As prior research shows, smart meter data indirectly leaks sensitive, and potentially valuable, information about a home’s activities. An important example of the sensitive information smart meters reveal is *occupancy*—whether or not someone is home and when. As prior work also shows, occupancy is surprisingly easy to detect, since it highly correlates with simple statistical metrics, such as power’s mean, variance, and range. Unfortunately, prior research that uses chemical energy storage, e.g., batteries, to prevent appliance power signature detection is prohibitively expensive when applied to occupancy detection. To address this problem, we propose preventing occupancy detection using the thermal energy storage of large elastic heating loads already present in many homes, such as electric water and space heaters. In essence, our approach, which we call Combined Heat and Privacy (CHPr), controls the power usage of these large loads to make it look like someone is always home. We design a CHPr-enabled water heater that regulates its energy usage to mask occupancy without violating its objective, e.g., to provide hot water on demand, and evaluate it in simulation and using a prototype. Our results show that a 50-gallon CHPr-enabled water heater decreases the Matthews Correlation Coefficient (a standard measure of a binary classifier’s performance) of a threshold-based occupancy detection attack in a representative home by 10x (from 0.44 to 0.045), effectively preventing occupancy detection at no extra cost.

## I. INTRODUCTION

The design of “smart” grids that optimize electricity generation and consumption to make it greener and more efficient has emerged as an important research area.<sup>1</sup> Smart grids are envisioned to leverage a variety of techniques to optimize their operation, including distributed generation from renewables, demand-side management, and variable time-of-use pricing, that require timely, fine-grained knowledge of electricity consumption at buildings throughout the grid. To support these optimizations, utilities are rapidly replacing existing electromechanical meters, which are read manually once a month, with smart meters that transmit a building’s electricity usage every few minutes. In 2011, an estimated 493 utilities in the U.S. had collectively installed more than 37 million smart meters [1].

Unfortunately, smart meters also indirectly leak private, and potentially valuable, information about a building’s occupants’ activities [2], [3], [4], [5], [6], [7]. To extract this information, third-party companies are now employing cloud-based, “big data” platforms to analyze smart meter data en masse [8], [9], [10]. While the purpose is, ostensibly, to provide consumers energy-efficiency recommendations, companies are mining the data for any profitable information. For example, detecting *power signatures*—sequences of changes in power unique to a device—for specific appliance brands could aid manufacturers

in guiding their marketing campaigns, e.g., identifying homes with GE versus Maytag appliances [8]. Many utilities are providing third-party companies access to troves of smart meter data. For instance, a recent report highlights one utility’s practice of requiring its customers to consent to sharing their data with third parties before permitting them to use an online web portal [11]. Such privacy violations have led to a small, but growing, backlash against smart meters [12].

An important example of simple and private information that smart meters leak is *occupancy*—whether or not someone is home and when. Tech-savvy criminals are already exploiting similar types of unintentional occupancy leaks, e.g., via publicly-visible online calendars and Facebook posts [13], to select victims for burglaries. In addition, occupancy may also indirectly reveal private information that is of interest to insurance companies, marketers, potential employers, or the government, e.g., in setting rates, directing ads, vetting an applicant’s background, or monitoring its citizens, respectively. Such information could include whether a home’s occupants: i) include a stay-at-home spouse, ii) keep regular working hours and daily routines, iii) frequently go on vacation, or iv) regularly eat out for dinner. As recent work demonstrates [14], [15], launching attacks that extract occupancy from smart meter data is surprisingly easy, since occupancy highly correlates with simple statistical metrics, such as power’s mean, variance, and range. Intuitively, users’ interaction with electrical devices, e.g., turning them on and off, lends itself to straightforward attacks that detect changes in these metrics and associates them with changes in occupancy. Figure 1 emphasizes the point by overlaying a home’s average power usage every minute with its occupancy—one is occupied and zero is unoccupied—between 6am and 11pm: power usage clearly increases and becomes more variable whenever people are home. The correlation between occupancy and power is not unique to this particular home—prior work [14], [15] has observed it in several homes.

Prior research proposes techniques to thwart privacy attacks on smart meter data [2], [3], [5], [7]. Broadly, these techniques use chemical energy storage, in the form of batteries, to power, or absorb, a fraction of a building’s total load, thereby changing the pattern of external grid power usage the smart meter records. By carefully controlling when batteries charge and discharge, the techniques prevent detecting appliance power signatures using sophisticated algorithms for Non-Intrusive Load Monitoring (NILM) [16], [17], [18]. However, these prior approaches do not change the statistical properties, e.g., high mean power, variance, and range, that imply occupancy, and are not designed to prevent occupancy detection. Thus, new techniques are necessary. To address the problem, we propose *Combined Heat and Privacy* (CHPr), which regulates thermal, rather than chemical, energy storage to make it look like someone is always home.

<sup>1</sup>Research supported by NSF grants CNS-1253063, CNS-1143655, CNS-0916577, CNS-0855128, CNS-0834243, and CNS-0845349.

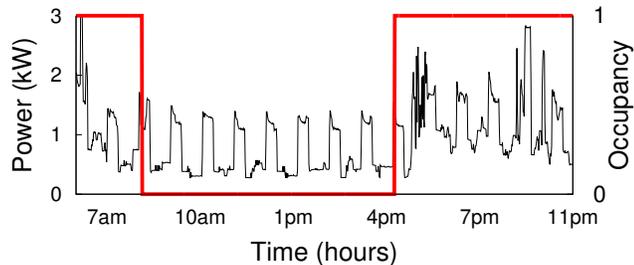


Fig. 1. When occupied, a home’s average power demand typically becomes larger and more variable due to occupants turning loads on and off.

One obvious, albeit wasteful, option for implementing CHPr is to simply consume power by dissipating heat without directing it to a useful purpose. Interestingly, a naïve CHPr strategy that masks occupancy by consuming (and wasting) energy to ensure demand is nearly always *flat*—close to the peak demand—is cost-competitive with comparable battery-based techniques due to the low price of electricity (12¢ per kWh on average [19]) and the high cost of batteries (~\$118 per kWh per year [20]). However, to prevent both wasting energy and increasing electricity costs, we propose integrating CHPr functionality into the large electric heating loads already found in many homes, such as water and space heaters. These loads effectively serve as thermal energy storage devices that CHPr can control to mask occupancy. In particular, we design a CHPr-enabled water heater with the goal of preventing occupancy detection without running out of hot water. Our approach combines multiple techniques to accomplish this goal: it i) uses partial demand flattening to eliminate a large majority of power variations, ii) injects artificial power signatures to obscure the relationship between occupancy and high, variable demand, and iii) adjusts its operation based on home activity patterns. CHPr is inspired by Combined Heat and Power (CHP) [21], which leverages the waste heat produced as a byproduct of generating electricity for water and space heating in buildings.

Our hypothesis is that a CHPr-enabled water heater is capable of regulating its power usage to prevent occupancy detection while still providing hot water on demand. In evaluating our hypothesis, this paper makes the following contributions. **Design Alternatives.** We describe the different design alternatives for preventing occupancy detection, including using both chemical and thermal energy storage, from smart meter data. We focus on our analysis on a simple threshold-based occupancy detection attack described in prior work [15]. **CHPr-enabled Water Heater.** We present the design of our CHPr-enabled water heater and its algorithm for regulating energy usage to prevent occupancy detection without running out of hot water. The approach combines partial demand flattening, artificial power signature injection, and an activity-based optimization to minimize its total energy requirements. **Implementation and Evaluation.** We experiment with our CHPr-enabled water heater in simulation and by deploying a proof-of-concept prototype in a real home. Our evaluation quantifies CHPr’s effectiveness using data from both the home and a real water heater. We show that our approach decreases the Matthews Correlation Coefficient—a standard measure of a binary classifier’s overall performance—of a threshold-based attack on the home’s smart meter data by a factor of 10 (from 0.44 to 0.045), effectively preventing occupancy detection.

## II. BACKGROUND

Our work assumes a building equipped with a smart meter that monitors aggregate electricity usage, and records the building’s average power  $P(t)$  over a sampling interval  $T$ , yielding a time-series of power values. Today’s newer utility-grade smart meters support sampling intervals from one to five minutes, while older meters support fifteen minutes to an hour. As a result, we focus on preventing occupancy detection from smart meter data with a one-minute sampling interval. Adapting our techniques to commercial power meters that offer higher resolution monitoring, e.g., 1Hz or greater, is future work. Given such a time-series  $P(t)$ , we represent *occupancy* as a binary function  $O(t)$ , over each sampling period  $t$ , where zero represents an unoccupied home and one represents a home with at least one person in it. Our work focuses on *masking* occupancy to prevent inferring  $O(t)$  from  $P(t)$ .

Since we are not aware of a general metric that applies to any possible occupancy detection attack, we evaluate CHPr using a threat model based on a specific and straightforward threshold-based attack, which signals occupancy if power’s mean, variance, or range exceeds some pre-defined threshold. In particular, we define an interval length  $T_{interval}$ , and then compute power’s mean, variance, and range over each interval. Anytime a metric exceeds a pre-defined power threshold, e.g., the nighttime average, we record a potential occupancy point, resulting in a series of points in time. We then cluster points to infer occupancy over time, such that if two points are within a time threshold, we consider the home occupied during the interval between those points. Our attack leverages the intuition from Figure 1 that occupancy correlates with periods of high, variable demand. A detailed description and evaluation of the attack is available in recent work [15], and is outside the scope of this paper. The attack above is simple, effective, and applies to the minute-level power data resolutions supported by smart meters. Using CHPr to prevent other types of attack vectors [14] and other sampling resolutions is future work.

### A. Prior Work

One way to prevent leaking any information, including occupancy, through smart meter data is to employ cryptographic techniques within the meter itself [6], [22], [23]. These techniques enable utilities to verify the correctness of various functions applied to smart meter data, e.g., a monthly bill using time-of-use rates, without requiring access to raw meter readings. However, this approach requires utilities to implement these protocols, including modifying the software of millions of already-installed smart meters [6], [22], [23]. As indicated in Section I, since smart meter data is valuable to utilities, they have little incentive to upgrade their infrastructure. Thus, an alternative approach, which does not require utility cooperation, is for consumers to obscure their smart meter data by actively altering their home’s grid power consumption. Prior techniques propose to alter grid power usage by controlling battery charging and discharging, called *Battery-based Load Hiding (BLH)* [2], [3], [5], [7].

BLH techniques focus primarily on preventing Non-Intrusive Load Monitoring (NILM) [16], [18], which analyzes changes in  $P(t)$  to compute a separate power time-series  $p_i(t)$  for each  $i = 1 \dots n$  appliances in a home. While

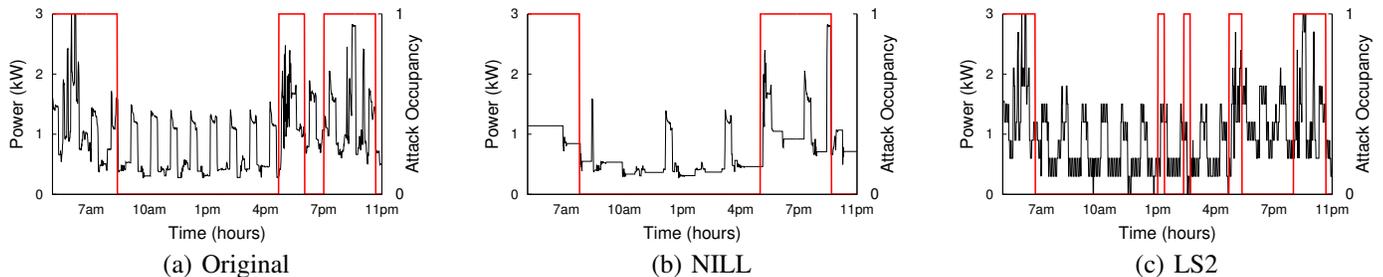


Fig. 2. A threshold-based attack is effective at detecting occupancy in smart meter data (a) when altered by BLH techniques, such as NILL (b) or LS2 (c).

BLH techniques have not been explicitly designed to prevent occupancy detection, we briefly describe two representative examples of BLH below to i) demonstrate that BLH techniques do not prevent occupancy detection simply as a side-effect of preventing NILM and ii) highlight the inherent difficulties in using batteries to mask occupancy. Our results show that BLH cannot defend against our simple threshold-based occupancy detection attack using practical battery capacities.

**Non-Intrusive Load Leveling** or NILL [5] removes changes in  $P(t)$  that reveal appliance power signatures by leveling, or flattening, the home’s external grid demand recorded by the smart meter. In essence, NILL charges batteries when actual demand is below a target demand, and then discharges batteries when it is above the target demand, to maintain meter readings as near to the target as possible. Ideally, demand is flat and always equal to the target demand, thereby revealing only the home’s average power usage and nothing else. Of course, only revealing the building’s average power would effectively prevent accurate occupancy detection over time. Unfortunately, for practical battery capacities, NILL diverges from this ideal. As we show below, once NILL fully discharges its battery, it can no longer alter the demand. Since battery depletion often occurs during the high demand periods that strongly correlate with occupancy, NILL does not prevent occupancy detection.

**Lazy Stepping** or LS [7] is an improvement to NILL that requires much less battery capacity to obscure appliance power signatures. The idea behind LS is that, rather than flatten a home’s demand, it controls battery charging and discharging to transform demand into a step function that removes the fine-grained changes in power useful in identifying appliances. However, as we show below, LS does not prevent occupancy detection: the periods of high demand that strongly correlate with occupancy remain clearly identifiable. LS’s technique highlights an important point: since occupancy correlates with periods of high demand, preventing occupancy detection is, in part, related to the amount of energy a building is able to shift across time. While LS is capable of judiciously using a much smaller battery than NILL to hinder NILM, preventing occupancy detection necessitates a larger capacity battery capable of masking periods of high demand, e.g., either by flattening them or injecting artificial high demand periods.

### B. Problems with Masking Occupancy using Batteries

Figure 2 empirically demonstrates the points above by using our simple attack to detect occupancy, even after demand has been altered by NILL and LS2.<sup>2</sup> The graphs overlay a

home’s average power usage every minute with the results of our occupancy detection attack for the same home and day as in Figure 1. In this case, we set the thresholds above equal to each metric’s value at night, which is similar to its value in an unoccupied home, with a clustering threshold of one hour. Figure 2(a) shows that, for the unaltered demand, with the exception of two brief periods, the attack’s predicted occupancy nearly exactly matches the ground truth from Figure 1.

Figure 2(b) then shows the results of the same attack on demand altered by NILL using a 6kWh battery, as in [5]. Despite the altered demand, the attack is still able to accurately detect occupancy. The NILL-altered demand demonstrates that, in practice, battery capacity limitations prevent ideal demand flattening. As expected, NILL does not prevent the high demand periods that correlate with occupancy, since it tends to deplete its battery during these periods, eliminating the option to later discharge its batteries to mask high demand. Of course, there exists some larger battery capacity, such that NILL would completely flatten demand at a home’s average, thereby preventing accurate occupancy detection. However, 6kWh of usable capacity<sup>3</sup> already imposes an excessively high cost—\$708 per year amortized over a battery’s lifetime based on recent cost estimates [20], which would increase an average U.S. home’s annual electricity bill by roughly 50% [24].

Likewise, Figure 2(c) shows the results of the attack on demand altered by the LS2 algorithm, which uses much less battery capacity—0.5kWh in this case, as in [7]—than NILL to hinder NILM. As the graph demonstrates, with 0.5kWh of battery capacity, LS2’s battery is simply too small to mask the periods of high demand by discharging its battery. Instead, LS2 discretizes demand to obscure the many small changes in power that NILM might leverage to identify appliances. As Figure 2(c) shows, due to the small capacity battery, demand altered by LS2 retains the general shape of the original demand profile including the periods of high, variable demand that indirectly reveal the home’s occupancy status.

Table I quantifies the effectiveness of both approaches by showing the percentage of time our attack yields true positives (detects occupancy and the home is occupied), true negatives (detects no occupancy and the home is not occupied), false positives (detects occupancy but the home is not occupied), and false negatives (detects no occupancy but the home is occupied). The accuracy is then the sum of the true positive and true negative percentages. The table also shows the Matthews Correlation Coefficient (MCC) [25], a standard measure of a

<sup>2</sup>LS2 is the best performing variant of LS [7].

<sup>3</sup>Cost estimates are based on a commercially-available sealed AGM/VRLA deep-cycle lead-acid battery designed for home solar panel installations.

External Demand	True Positives	True Negatives	False Positives	False Negatives	Accuracy	MCC
<b>Original (a)</b>	41.86%	46.57%	<b>1.27%</b>	10.29%	88.43%	<b>0.78</b>
<b>NILL (b)</b>	37.25%	47.84%	<b>0%</b>	14.90%	85.09%	<b>0.74</b>
<b>LS2 (c)</b>	23.53%	43.92%	<b>3.92%</b>	28.63%	67.45%	<b>0.42</b>

TABLE I. PERFORMANCE OF OUR THRESHOLD-BASED ATTACK ON A HOME’S ORIGINAL DEMAND AND AFTER BEING ALTERED BY NILL AND LS2.

binary classifier’s performance, where values are in the range  $-1.0$  to  $1.0$ , with  $1.0$  being perfect detection,  $0.0$  being random prediction, and  $-1.0$  indicating detection is always wrong. MCC values closer to  $0.0$ , or random prediction, are better for masking occupancy. The table shows that our simple threshold-based occupancy detection attack is effective on demand altered by NILL or LS2: it yields an MCC of  $0.74$  and  $0.42$  on the NILL-altered and LS2-altered demand, respectively, which is near the MCC ( $0.78$ ) of the attack on the original demand.

**Summary.** BLH’s primary drawback when applied to the problem of preventing occupancy detection is that battery-based energy storage is expensive. Masking occupancy detection is, in part, related to the amount of total energy a building is able to shift across time. While NILL could effectively mask occupancy with a sufficiently large battery, e.g., by completely flattening a home’s demand at its average, the cost would be high, e.g., a yearly expense greater than  $50\%$  of a home’s annual electricity bill. While LS shows that judiciously controlling the charging and discharging for small capacity batteries is effective at hindering NILM, a small capacity battery simply cannot mask the periods of high, variable demand that indicate occupancy. In addition, any BLH technique, including both NILL and LS2, wastes a fraction of any energy it stores in its battery, due to energy conversion losses. These losses are at least  $20\%$  of the stored energy with existing battery and inverter technology [26]. The insights above lead to CHPr’s approach, which leverages the thermal energy storage inherent to large elastic heating loads, such as water heaters, to cheaply and efficiently mask occupancy. Since a  $4.5\text{kW}$  water heater that runs for a typical three hours per day consumes  $13.5\text{kWh}$  of energy [27], it is capable of shifting significantly more energy than the NILL or LS examples above. In addition, since CHPr only reschedules energy a water heater already consumes, it avoids energy conversion losses.

### III. USING THERMAL STORAGE: DESIGN ALTERNATIVES

We consider the design alternatives for using thermal energy storage to mask occupancy. Figure 3 highlights the differences between BLH and thermal energy storage. BLH flattens grid demand by controlling battery charging and discharging, such that, in the ideal (although not in practice for reasonable battery capacities), the smart meter always sees a steady, flat power consumption level (depicted by  $T$  in Figure 3(a)). Whenever the home’s demand rises above  $T$ , BLH discharges its battery to provide the home additional power, rather than drawing it from the grid. The approach thwarts occupancy detection attacks by “clipping” any power usage above  $T$ , exposing a constant power usage to the smart meter that effectively makes it look like *no one is ever home*.<sup>4</sup>

Thermal energy storage is also capable of flattening demand in a similar manner, although it cannot “clip” power usage in the same way as a battery, since it is incapable of

discharging general-purpose electricity, i.e., it cannot convert its heat back into electricity. Instead, thermal energy storage can only flatten demand by *raising* grid power usage, e.g., by converting electricity into heat, to its peak level (depicted by  $T'$  in Figure 3(b)). In this case, the thermal storage device controls its resistive heating elements to draw a variable amount of power (above the normal power draw) to ensure that the total power draw is always  $T'$ . Thus, thermal energy storage is able to thwart occupancy detection by “boosting” power usage such that the home always draws a steady power  $T'$  from the grid. The thermal device then stores the heat for later use.

Since the homes we monitor have a high peak-to-average power ratio, raising power usage to the peak value  $T'$  requires a substantial amount of energy, which in turn requires a large amount of thermal energy storage capacity to make use of the heat. To reduce the power necessary to mask occupancy, thermal energy storage can also leverage *artificial power signature injection*, which controls the thermal device to inject “noise” that resembles real electrical loads in the home (depicted in Figure 3(c)). By injecting fake signatures that resemble real loads during low-power periods when no one is home, the approach makes it appear that *someone is always home*, which also thwarts occupancy detection, but using less energy. As before, the thermal device stores its heat for later use. As we describe in the next section, CHPr leverages a hybrid approach (in Figure 3(d)) that combines artificial signature injection with partial demand flattening, such that it raises demand to an intermediate value  $T''$  (below the peak value  $T'$ ). Since partial demand flattening reveals peaks above  $T''$ , CHPr only injects artificial signatures larger than  $T''$ .

### IV. A CHPR-ENABLED WATER HEATER

A standard tank-based residential water heater includes a reserve tank with a cold-water inlet pipe at the bottom and a hot-water outlet pipe at the top, since heated water naturally rises to the top of the tank. Residential water heaters include tanks that range in size from  $30\text{--}100$  gallons (equivalent to  $113.6\text{--}378.5$  liters, respectively) with heating elements ranging from  $3500\text{W}$  to  $5500\text{W}$ . Importantly, a water heater’s average total energy usage (and its thermal energy capacity) is a significant fraction of an average home’s usage. For example, a standard  $50$  gallon (or  $189.3$  liter),  $4.5\text{kW}$  water heater that runs for three hours each day consumes  $13.5\text{kWh}$  [27], while an average U.S. home consumes only  $\sim 24\text{kWh}$  per day [24].

A typical water heater operates by always attempting to ensure that i) the tank is full and ii) the tank’s water temperature is equal to an adjustable target temperature that is typically set between  $120\text{F}$  and  $140\text{F}$  (or  $48.9\text{C}$  to  $60\text{C}$ ). Thus, when hot water is drawn from the tank, e.g., due to someone taking a shower, the water heater refills the tank with cold water, and then immediately begins heating it at maximum power until the tank’s water reaches the target temperature. The temperature of the intake water is usually in the range of  $50\text{F}\text{--}60\text{F}$  (or  $10\text{C}\text{--}15.6\text{C}$ ), but is dependent on the climate.

<sup>4</sup>An occupancy detector may still detect occupancy if  $T$  is sufficiently high.

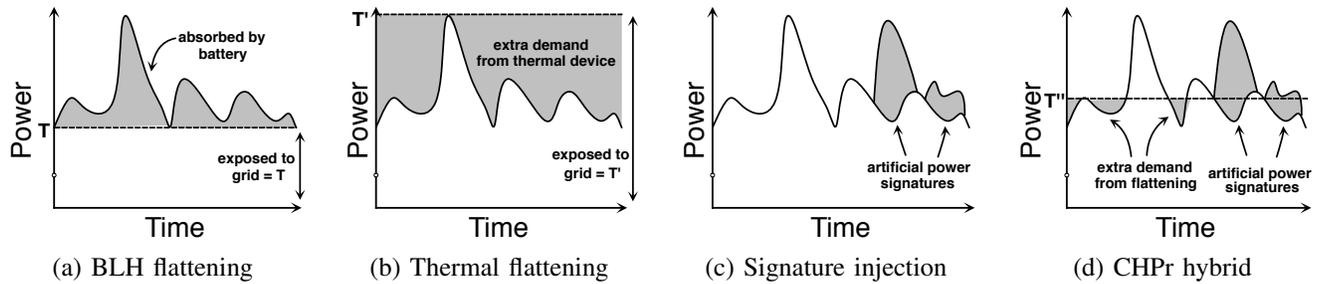


Fig. 3. Different options for masking occupancy, including i) demand flattening using both BLH (a) and thermal energy storage (b), ii) artificial power signature injection (c), and iii) CHPr’s hybrid approach (d) that combines demand flattening and artificial signature injection to minimize its energy requirements.

Water heaters generally employ a tight guardband of 15F (or 8.33C), such that if no hot water is drawn out, the water heater waits until the water is, for example, 105F (or 40.6C) before reheating it to the 120F (or 48.9C) target [28]. Since hot water rises, water heaters often employ two heating elements and thermostats, one at the top and bottom of the tank.

A CHPr-enabled water heater works by relaxing the operational requirements above and not always using the maximum power to immediately heat intake water. As an example, Figure 4 shows the power usage of a 50 gallon (or 189.3 liter), 4500W water heater over one day on the left y-axis. The short regular bursts of power are due to maintaining the water temperature within the 15F guardband, while the longer periods of power usage stem from heating the cold intake water that is replacing hot water drawn out of the tank. The right y-axis shows the amount of available hot water (at 120F), assuming ideal insulation where it takes  $2.93 \times 10^{-4}$  kWh to raise 1lb (or 0.45kg) of water by 1F (or 0.56C). We then compute the amount of 120F (or 48.9C) hot water by correlating the heater’s energy usage with a volume of heated water. Figure 4 indicates that, on this day, the tank never runs out of hot water. The figure also shows that the water heater could heat at a slower constant rate (indicated by the dotted red lines) using less than the maximum power without ever running out of hot water. Rather than heat at a slow constant rate, CHPr varies the heating element’s power usage to partially flatten demand and inject artificial signatures to mask occupancy, while using the same amount of energy over the period.

To determine how fast it must heat water to prevent running out, which dictates the energy it must consume over a given period, CHPr tracks the amount of remaining hot (120F/48.9C) water at the top of the tank and estimates the time until the next significant use of hot water. Our current implementation simply maintains an estimate of the average length  $t$  between usage periods greater than 25 gallons (or 94.66 liters), or roughly a single shower, and ensures that after a significant usage period all the water is heated within  $t$ . While more sophisticated methods for estimating  $t$  are possible, we did not explore them since our simple method proved effective. Given an energy budget and this time period estimate  $t$ , CHPr then determines how much to partially flatten demand and inject artificial signatures, as described below.

**Partial Demand Flattening.** Since a water heater does not use enough energy to completely flatten demand at the peak demand, CHPr employs a *flattening threshold*  $P_{flat}$  that only *partially* flattens demand to a target level less than the peak demand. To maintain  $P_{flat}$  at each  $t$  with current demand  $N(t)$ ,

CHPr must consume  $P_{flat} - N(t)$  whenever  $N(t) < P_{flat}$ . Since average demand is typically much lower than peak demand, a low flattening threshold is able to hide a large percentage of the changes in power without using much energy. Figure 5 illustrates this point by showing the energy required in our home deployment (as a percentage of the home’s total energy usage) to support various flattening thresholds, along with the amount of exposed readings above the flattening threshold. The figure shows that supporting flattening thresholds near 1kW do not require a significant amount of energy (<5%), but reduce the exposed readings by more than 25%.

**Artificial Power Signature Injection.** Partially flattening demand still exposes changes in power that occur above the threshold. To hide these changes, CHPr injects artificial power signatures. Importantly, CHPr does not simply inject demand randomly, since an attacker may be able to detect these random or atypical patterns in smart meter data. Instead, CHPr *replays realistic power signatures*. These power signatures are derived from the home’s aggregate data, by storing, in a database, sequences of the home’s power changes that occur above the flattening threshold. CHPr also takes additional steps to ensure artificial demand is difficult to discern from real demand. For example, the power signature database includes attributes for each signature, such as average power and duration. CHPr then divides power signatures into categories based on their attributes, e.g., small, medium, large and short, medium, and long, and computes the fraction of signatures in each category.

We use this fraction to weight each category’s random selection, such that the artificial demand matches the breakdown of real demand. In addition, to prevent attackers from detecting repeated signatures, CHPr introduces some randomness into the replayed signature by raising or lowering each point by a small random amount, e.g., 0-5% of usage. Finally, to reduce its energy requirements, CHPr only injects signatures when the home is unoccupied. Our premise is that injecting artificial power signatures should not be necessary when a home is occupied—there is no need to make the data look like someone is home when someone *actually* is home. When the home is unoccupied, CHPr randomly selects signatures from the database to inject and replay at an injection rate equal to the rate at which the home generates power signatures above the flattening threshold when occupied. Our prototype explicitly tracks home occupancy by monitoring occupants’ GPS coordinates in real time via a smartphone application.

**Activity-aware Optimization.** CHPr makes a home’s power profile look like someone is always home. However, even an occupied home’s expected power usage differs over time. For

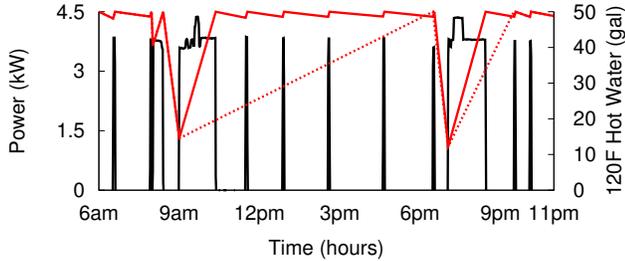


Fig. 4. A day’s power usage (black) for a 50 gallon (or 189.33 liter), 4.5kW water heater, and the remaining hot (120F/48.9C) water in its tank (red).

example, a home’s typical nighttime power usage is much lower than its daytime usage, even when occupied. Likewise, a home’s weekend power usage is often much greater than its weekday usage. Since these patterns are expected, there is no need to make low-power nighttime periods look like high-power daytime periods, or low-power weekdays look like high-power weekends. Instead, CHPr need only ensure that these time periods look the same with respect to each other, regardless of whether a home is occupied or unoccupied. Thus, CHPr adapts itself based on activity patterns by using a different flattening threshold, injection rate, and signature database at different times. CHPr sets different flattening thresholds and injection rates for days and nights, and weekends and weekdays. In addition, CHPr indexes its power signature database based on each signature’s real time-of-use. At any given time, CHPr randomly selects from past power signatures that also occurred near that time, e.g., within an hour, since typical power signatures in the morning, e.g., a coffee maker, are likely to be different from those in the evening, e.g., a TV. Indexing signatures by time is also important because an attacker could exploit usage patterns that appear unnatural.

**Tuning CHPr.** CHPr sets the flattening threshold  $P_{flat}$  for each period based on the excess energy available after estimating the energy required to inject artificial signatures (based on the rate of signatures observed when the home is occupied). Of course, CHPr could run out energy if its i) estimated energy budget over a time period  $t$  is inaccurate or ii) occupants leave for extended periods, such that the water heaters does not have enough thermal capacity to partially flatten demand and inject artificial signatures over the period. As with BLH, whenever CHPr runs out of energy it has no choice but to expose the home’s raw usage to the smart meter. We evaluate the frequency and impact of running out of energy in Section VI.

## V. CHPR IMPLEMENTATION

We implement both a CHPr simulator and proof-of-concept prototype. The simulator, written in R, takes as input a home’s aggregate power trace and its water heater power trace, and reschedules the water heater’s power consumption based on the approach outlined in the previous section. In addition to the simulator, we also deploy a proof-of-concept CHPr prototype in a real 3-bedroom, 2-bath house. Rather than implement a full-fledged water heater, the prototype’s purpose is only to demonstrate the ability to modulate a home’s power usage to mask occupancy based on CHPr’s approach. To do this, we employ 18 Insteon LampLinc programmable dimmer switches [29], which enable a computer to remotely set their

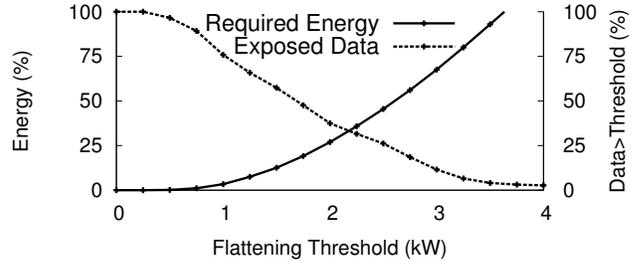


Fig. 5. Partial demand flattening hides many changes in power (below the  $P_{flat}$  threshold) without requiring a significant fraction of the home’s energy.

dim level via the Insteon powerline networking protocol. We use LampLincs because they are widely available and have open-source software support. Controlling the LampLincs’ dim level enables us to control their power usage in the same way as a water heater’s heating element.

The prototype uses an eGauge power meter [30] in the home’s electrical panel to query the real-time power readings for the whole home every second via a web-based API. In addition, each of the home’s two adult occupants run a real-time geolocation application on their cell phone, which our software queries in real time to determine the home’s ground truth occupancy (based on the occupants’ GPS coordinates). We have collected GPS data for roughly one year, and power data for over three years. We implement CHPr’s algorithm from the previous section in Python and run it on an embedded Linux-based DreamPlug server, which connects to an Insteon Powerline Modem via USB to programmatically control the power usage of the LampLincs. The system is able to use simple Insteon command-line tools for Linux to turn each LampLinc on and off and adjust its dim level between 0% and 100% in increments of 1%. Since a light bulb’s power usage scales linearly with its dim level, the system is capable of controlling power usage at 3W granularities (1% of one 300W LampLinc). In total, the 18 LampLincs enable the system to control  $18 \times 300W = 5400W$  of power, which is sufficient for replaying even sizable loads in the home. The software stores the set of artificial power signatures, indexed by time period, that are available for replay in a sqlite3 database, and then queries the database to select signatures for replay. We seed the power signature database with real power signatures based on the past six months of usage, indexed by their time-of-use to support our activity-aware optimization. The size of the implementation is less than 1500 lines of code.

## VI. EXPERIMENTAL EVALUATION

We evaluate CHPr’s effectiveness against our threshold-based occupancy detection attack in the home above [15], both in simulation (using data collected from the home) and using our prototype deployed in the home. While the home’s occupancy rate appears high, based on our own data collection at three other homes and national statistics [19], [24], we believe the home’s power usage and occupancy pattern are representative of a large class of homes. For instance, consider that even if all occupants are away for a standard 40-hour work week (8 hours per day), and home otherwise, the resulting occupancy rate is still 76.2% (128 out of 168 hours).

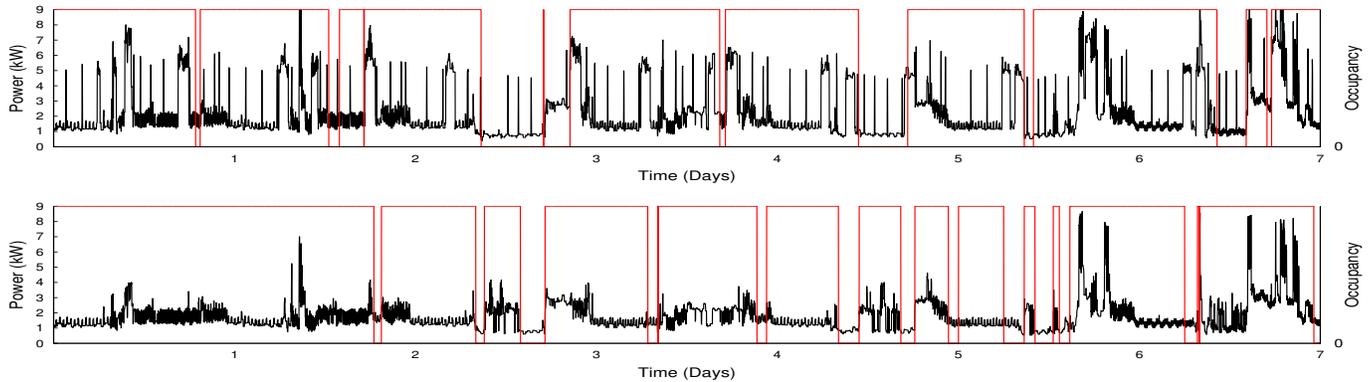


Fig. 6. A home’s original week-long power usage and ground truth occupancy (top), as well as its power usage when using a CHPr-enabled water heater and detected occupancy when using the threshold-based occupancy detection attack outlined in Section II (bottom) .

**Preventing Occupancy Detection.** Figure 6 uses our simulator (with input data from the home above over a representative week in the summer) to demonstrate CHPr’s ability to mask occupancy. The top graph shows both the home’s power usage, including a standard 50 gallon (or 189.3 liter) water heater, as well as its ground truth occupancy using the occupants’ GPS coordinates. The brief spikes in electricity usage throughout the week are due to heating water. The lower graph then shows the power usage after rescheduling the water heater’s power consumption using CHPr, as well as the detected occupancy of this modified power trace using our threshold-based attack. A good example of CHPr’s capabilities occurs between days four and five when the home is unoccupied for an extended period. Using the original demand, the low power usage clearly indicates the occupants are away, while the CHPr-modified demand makes the power usage appear similar to an occupied home. While there are a few instances where the water heater runs out of energy, i.e., fully heats all of its tank’s water, that cause it to expose a low power usage that may reveal an unoccupied home, e.g., between days two and three, the data exposes much less occupancy information overall. In addition, there are no instances where our (simulated) reserve tank runs out of hot water due to CHPr’s operation.

We also quantify the performance of the occupancy detection attack on both the original demand and the CHPr-modified demand in terms of the Matthews Correlation Coefficient (MCC) [25]. Recall from Section II, that the MCC is a standard measure of a binary classifier’s performance, where values are in the range  $-1.0$  to  $1.0$ , with  $1.0$  being perfect detection,  $0.0$  being random prediction, and  $-1.0$  indicating detection is always wrong. MCC values closer to  $0.0$ , or random prediction, are better for masking occupancy. In this case, our results show that the MCC of the attack on the CHPr-modified data is only  $0.045$ , which is nearly the same as random prediction, i.e., an MCC of  $0.0$ , and is a factor of 10 less than the MCC of the attack on the original data, which is  $0.44$ .

**Result:** *By lowering the MCC to 0.045, our CHPr-enabled 50 gallon (or 189.3 liter) water heater effectively prevents occupancy detection from the threshold-based attack in our simulated home without exhausting the hot water supply. Unlike battery-based techniques, CHPr requires no additional power usage and does not increase power costs.*

**Optimizations.** Figure 7(a) shows CHPr’s energy requirements

(as a percentage of the home’s total demand) from employing our various optimizations in our simulated home. The graph demonstrates the benefit of i) the activity-aware optimization that determines different flattening thresholds and injection rates at different times, in this case day versus night and weekday versus weekend, and ii) the occupancy-aware optimization that only injects artificial power signatures when occupants are away. Since nights and weekdays exhibit lower power usage than days and weekends, CHPr requires much less energy during these periods to mask occupancy. The graph also indicates that only injecting artificial signatures when the home is unoccupied also results in a significant energy reduction. Ultimately, the result shows that combining these optimizations versus using none of them reduces CHPr’s energy requirements by nearly a factor of two, enabling it to make efficient use of the water heater’s limited thermal energy storage capacity.

**Result:** *CHPr’s occupancy- and activity-aware optimizations are important in reducing, in this case by over a factor of two, the energy required to prevent occupancy detection from our threshold-based attack in our simulated home.*

**Prototype Demonstration.** While the results above use our simulator, Figure 7(b) demonstrates the performance of our CHPr prototype for an eight hour period for both 1-minute and 5-minute average power data. In this case, the unmodified demand is the home’s demand without CHPr’s contribution, while the CHPr-modified demand is the external usage seen by the smart meter, which includes using the prototype to reschedule the power usage of an emulated water heater. We can extract these separate power values because eGauge records power for each of the home’s 26 individual circuits. The experiment shows how our prototype modifies a home’s demand in real time, including both flattening it and accurately replaying artificial power signatures, to mask the usage trends in the real data. Ultimately, our prototype demonstrates that CHPr’s approach permits a straightforward implementation using widely-used, off-the-shelf components. Controlling the power usage of the resistive heating elements found in water heaters uses the same functionality as a basic dimmer switch, which rapidly cuts power for some fraction of each alternating current cycle, e.g., 50 or 60Hz, to precisely adjust power usage.

**Result:** *CHPr functionality is simple to implement and deploy, requiring only the mechanism of a basic dimmer switch and the ability to programmatically adjust its dim level in real time.*

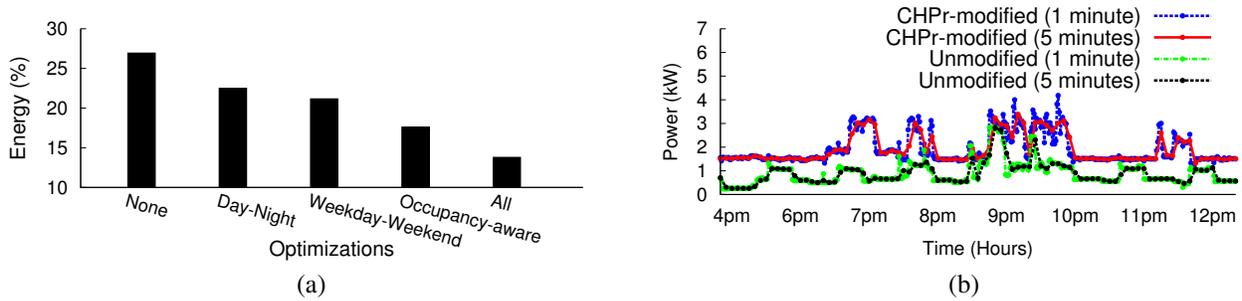


Fig. 7. CHPr’s optimizations significantly reduce its energy requirements (a), while (b) shows a demonstration of masking occupancy with our CHPr prototype.

## VII. CONCLUSION

This paper presents CHPr (Combined Heat and Privacy), which prevents occupancy detection using the thermal energy storage inherent to the large elastic heating loads already present in many homes, in particular electric water heaters. As we show in Section II, CHPr leverages thermal energy storage to mask occupancy because using chemical energy storage, in the form of batteries, requires a level of energy storage capacity that is prohibitively expensive. CHPr’s algorithm combines partial demand flattening, artificial power signature injection, and activity- and occupancy-aware optimizations to reduce its energy requirements. Importantly, CHPr does not waste any energy and does not increase electricity costs: it simply reschedules the energy a water heater already consumes to mask occupancy, while ensuring the reserve tank does not run out of hot water. Our evaluation shows that CHPr is effective at masking occupancy by regulating the power usage of a standard 50 gallon (or 189.3 liter) CHPr-enabled water heater, decreasing the MCC of a threshold-based occupancy detection attack in a representative home by 10x (from 0.44 to 0.045). Of course, more advanced attack vectors are possible. While we believe CHPr’s approach is general and could be adapted to combat more advanced attacks (both known [14] and unknown), this is the subject of future work. Our current approach only applies to homes that already have electric water heaters (38% of U.S. homes based on recent estimates [31]); generalizing it to use other types of background loads is also the subject of future work.

**Acknowledgements.** We would like to thank our shepherd Silvia Santini for providing us valuable feedback, which significantly improved the quality of this paper.

## REFERENCES

- [1] “U.S. Energy Information Administration, Frequently Asked Questions, How Many Smart Meters are Installed in the U.S. and who has them?” <http://www.eia.gov/tools/faqs/faq.cfm?id=108&t=3>, 2011.
- [2] M. Backes and S. Melsner, “Differentially Private Smart Metering with Battery Recharging,” *IACR Cryptology*, no. 183, April 2012.
- [3] G. Kalogridis, C. Efthymiou, S. Denic, T. Lewis, and R. Cepeda, “Privacy for Smart Meters: Towards Undetectable Appliance Load Signatures,” in *SmartGridComm*, October 2010.
- [4] P. McDaniel and S. McLaughlin, “Security and Privacy Challenges in the Smart Grid,” *IEEE Security and Privacy Magazine*, vol. 33, November 2009.
- [5] S. McLaughlin, P. McDaniel, and W. Aiello, “Protecting Consumer Privacy from Electric Load Monitoring,” in *CCS*, October 2011.
- [6] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin, “Private Memoirs of a Smart Meter,” in *BuildSys*, November 2010.
- [7] W. Yang, N. Li, Y. Qi, W. Qardaji, S. McLaughlin, and P. McDaniel, “Minimizing Private Data Disclosures in the Smart Grid,” in *CCS*, October 2012.
- [8] “Bidgely,” <http://bidgely.com>, May 2013.
- [9] “Chai Energy,” <http://www.mychai.co/>, May 2013.
- [10] “PlottWatt,” <https://plotwatt.com/>, May 2013.
- [11] V. Chadwick, C. Butt, and H. Cook, “Smart Meter Data Shared Far and Wide,” *The Age*, Tech. Rep., September 23rd 2012.
- [12] “Stop Smart Meters!” <http://stopsmartmeters.org/>, May 2013.
- [13] A. Bloxham, “The Telegraph, Most Burglars using Facebook and Twitter to Target Victims, Survey Suggests,” <http://www.telegraph.co.uk/technology/news/8789538/Most-burglars-using-Facebook-and-Twitter-to-target-victims-survey-suggests.html>, September 26th 2011.
- [14] W. Kleiminger, C. Beckel, T. Staake, and S. Santini, “Occupancy Detection from Electricity Consumption Data,” in *BuildSys*, 2013.
- [15] D. Chen, S. Barker, A. Subbaswamy, D. Irwin, and P. Shenoy, “Non-Intrusive Occupancy Monitoring using Smart Meters,” in *BuildSys*, November 2013.
- [16] K. Armel, A. Gupta, G. Shrimali, and A. Albert, “Is Disaggregation the Holy Grail of Energy Efficiency? the Case of Electricity,” *Energy Policy*, vol. 52, no. 1, January 2013.
- [17] G. Hart, “Nonintrusive Appliance Load Monitoring,” *IEEE*, vol. 80, no. 12, December 1992.
- [18] M. Zeifman and K. Roth, “Nonintrusive Appliance Load Monitoring: Review and Outlook,” *IEEE Transactions on Consumer Electronics*, vol. 57, no. 1, February 2011.
- [19] “Electric Power Monthly with Data for February 2013,” U.S. Energy Information Administration, Tech. Rep., April 2013.
- [20] A. Mishra, D. Irwin, P. Shenoy, J. Kurose, and T. Zhu, “SmartCharge: Cutting the Electricity Bill in Smart Homes with Energy Storage,” in *e-Energy*, May 2012.
- [21] “Combined Heat and Power: A Clean Energy Solution,” U.S. Department of Energy, Tech. Rep., August 2012.
- [22] A. Rial and G. Danezis, “Privacy-Preserving Smart Metering,” in *WPES*, October 2011.
- [23] A. Molina-Markham, G. Danezis, K. Fu, P. Shenoy, and D. Irwin, “Designing Privacy-preserving Smart Meters with Low-cost Microcontrollers,” in *FC*, February 2012.
- [24] D. Cauchon, “USAToday, Household Electricity Bills Skyrocket,” <http://www.usatoday.com/money/industries/energy/story/2011-12-13/electric-bills/51840042/1>, December 13th 2011.
- [25] B. Matthews, “Comparison of the Predicted and Observed Secondary Structure of T4 Phage Lysozyme,” *Biochimica et Biophysica Acta.*, vol. 405, no. 2, October 1975.
- [26] S. Schoenung, “Energy Storage Systems Cost Update: A Study for the DOE Energy Storage Systems Program,” Sandia National Laboratories, Tech. Rep., April 2011.
- [27] “WaterHeaterTimer.org,” <http://waterheatertimer.org/How-much-does-it-cost-to-run-water-heater.html>, May 2013.
- [28] “Water Heaters are not Water Delivery Temperature Control Devices,” [http://www.cashacme.com/legionella\\_related\\_info\\_art1.php](http://www.cashacme.com/legionella_related_info_art1.php).
- [29] “SwitchLinc Dimmer INSTEON Remote Control Dimmer Switch Owners Manual (rev 5.0+),” <http://www.smarthome.com/manuals/2476d.pdf>, July 24th 2012.
- [30] “eGauge Energy Monitoring Solutions,” <http://www.egauge.net/>, 2012.
- [31] “EnergyGuide: Smart Energy Choices,” <https://www.energyguide.com/library/EnergyLibraryTopic.asp?bid=austin&prd=10&TID=17240&SubjectID=8374>, May 2013.