Protecting Consumer Privacy from Electric Load Monitoring

Stephen McLaughlin Patrick McDaniel
Systems and Internet Infrastructure Security Lab
Pennsylvania State Univesity
University Park, PA, USA
{smclaugh,mcdaniel}@cse.psu.edu

William Aiello
Networks, Systems, and Security Lab
University of British Columbia
Vancouver, B.C., Canada
aiello@cs.ubc.ca

ABSTRACT

The smart grid introduces concerns for the loss of consumer privacy; recently deployed smart meters retain and distribute highly accurate profiles of home energy use. These profiles can be mined by Non Intrusive Load Monitors (NILMs) to expose much of the human activity within the served site. This paper introduces a new class of algorithms and systems, called Non-Intrusive Load Leveling (NILL) to combat potential invasions of privacy. NILL uses an in-residence battery to mask variance in load on the grid, thus eliminating exposure of the appliance-driven information used to compromise consumer privacy. We use real residential energy use profiles to drive four simulated deployments of NILL. The simulations show that NILL exposes only 1.1 to 5.9 useful energy events per day hidden amongst hundreds or thousands of similar battery-suppressed events. Thus, the energy profiles exhibited by NILL are largely useless for current NILM algorithms. Surprisingly, such privacy gains can be achieved using battery systems whose storage capacity is far lower than the residence's aggregate load average. We conclude by discussing how the costs of NILL can be offset by energy savings under tiered energy schedules.

Categories and Subject Descriptors

J.m [**Computer Applications**]: Miscellaneous

General Terms

Security

Keywords

smart meter, privacy, load monitor

1. INTRODUCTION

Smart meters are being aggressively deployed in homes and businesses as part of a move to global smart grids [23]. This digitization of grid systems offers substantial benefits for society; increased efficiencies and information availability can enable cheaper and greener energy generation, less loss in energy storage and transmission, better fault isolation and recovery, and support for alternative energy sources, e.g., consumer generated wind and solar energy.

The move to digital grid control systems also introduces concerns about security [19, 24, 27]. The smart grid is a complex system of sensors, networks, and computing resources. Attacks against the smart-grid networks and computing elements can range from fraud, to denial of service, to privacy loss [27]. While some regulatory agencies have begun to explore security concerns, no comprehensive system has emerged to address these threats.

One area of particular concern is the loss of consumer privacy. Replacements for the antiquated in-home electromechanical meters, smart meters are embedded systems that use power and voltage sensors to collect and report *load profiles*. Load profiles are histories of energy usage collected at a configured granularity, e.g., seconds or minutes. While instrumental to managing energy use at the local and regional levels, such profiles are also sufficient to determine occupant behavior in residential settings [22, 21, 11]. Depicted here, this behavioral inference is made possible by a class of algorithms known as Non-Intrusive Load Monitoring (NILM):

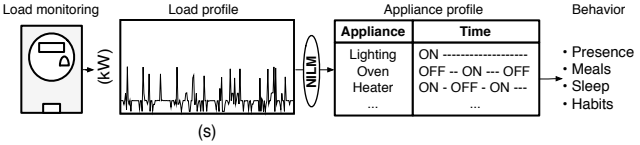

To simplify, NILM algorithms decompose load profiles into composite *appliance profiles* based on known or learned signatures. For example, traces of discrete changes in energy use can be mapped directly to ON/OFF events associated with identifiable appliances. The profile is a detailed description of appliance use and indirectly a surprisingly accurate model of human activity [22].

The concerns surrounding potential invasions of privacy via energy profiles appear to be more than hypothetical [15, 6]. Reuse of data by direct marketers, criminals, or law enforcement without prior approval or notification is often in conflict with privacy regulations, but may be occurring anyway [22]. Undercutting existing regulatory structures is a maze of often conflicting laws and court decisions relating to consumer privacy. For example, the 1939 Supreme Court *United State v. Miller* decision indicates that there is no reasonable expectation of privacy for information shared with third parties. State and regional agencies have built legal and regulatory structures to buttress privacy in the face of such decisions, but consumer rights remain, at best, murky.

The potential exposure of living conditions, occupancy, and family routines, through energy profiles warrants vigilance [15]. This prompts the goal of this work: we aim to protect consumer privacy

CCS'11, October 17–21, 2011, Chicago, Illinois, USA.

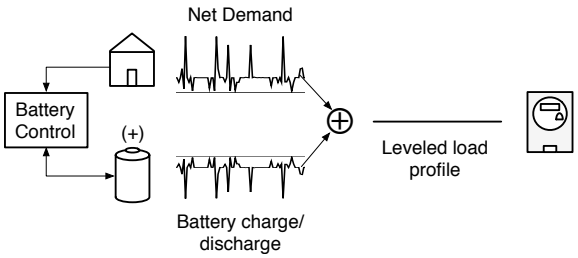

Figure 1: Idealized non-invasive load leveling (NILL).

in the face of profile-exposing smart meters while acknowledging two practical constraints of the current grid environment:

- *Energy usage must always be accurately reported.* Any modification of usage data would undermine grid management, and introduce inaccuracies in billing and grid controls.
- *The privacy solution must not require any modification to the meters, appliances, grid control systems, or provider operation*. The metering infrastructure, which is not under the control of the consumer, is assumed to be untrusted. Moreover, with millions of smart meters already installed and many more in deployment,[1] any solution requiring new grid systems will not be logistically and economically viable [27].

We address threats to electricity consumer privacy through *Non-Intrusive Load Leveling* (NILL), a novel technique to mask appliance features in a home's net load. Illustrated in Figure 1, NILL is conceptually simple; a consumer places a battery and control system between the smart meter and the circuit breaker of their residence. The load observed by the meter is smoothed by offsetting spikes and dips in usage by charging or discharging the battery. Hence, NILL removes the information content that reveals appliance usage. Because we make no assumptions about adversary motivations, NILL aims to smooth all appliance features in a house. However, because of the physical limits and structure of electrical systems, this is a more challenging task than one might initially surmise. Note that NILL is currently not designed to mask longer term energy usage such as day/night diurnal energy patterns, but only the instantaneous energy transitions that expose minute-to-minute human behavior exploited by NILM algorithms. However, we do explore the challenges and countermeasures posed by currently undeveloped NILM techniques that use more sophisticated learning and inference techniques in Appendix A.

NILL is an algorithm and control system that attempts to remove the fine-grained appliance signal represented by changes in the reported load. The control system directs the battery charges and discharges to obscure energy usage. This is conceptually similar to queue delay perturbation countermeasures that prevent networking timing analysis [5]. By de-correlating both the timing and amplitude of ON/OFF events in the load profile, we remove the signal that NILM algorithms use to identify behavior.

The idea of using a battery to provide "best effort" privacy protection is not new. For example, one short paper [16] has suggested the use of a power router to allow a battery to offset appliance loads, though the existing technology in this area limits the battery to handling one appliance at a time. Furthermore, none of the physical challenges of introducing a battery into a residential setting were evaluated. Ours is the first work to perform a rigorous physical simulation of such a system under substantial real-word data. The NILL approach presented here attempts to provide privacy for all appliances under all battery states. Our analysis also extends beyond previous results by examining NILL's effect on the basic unit of load monitoring, the feature pair, and by measuring the amount of privacy afforded over time due to changes in battery states.

The remainder of this paper identifies and evaluates a candidate NILL algorithm. A simulation of NILL is built on the widely-used SimPowerSystems [26] platform. We simulate four homes using energy profile data collected from real residential use. These experiments show that NILL exposed 1.1 and 5.9 identifiable appliance events per day. Such features reside amongst hundreds or thousands of battery-suppressed events, making reliable recovery of appliance profiles virtually impossible under current NILM. Further, we showed that such privacy can be achieved in the tested environments using only a moderately priced 50 amp-hour battery system–far smaller than the aggregate loads of protected residences.

[1]$4.3 billion dollars has been allocated by the U.S government for the smart grids [28], with similar programs in progress in Asia and the EU.

Table 1: Commodity smart meters [29].

Epoch	Product(s)	Deployed
monthly	electromechanical meters	N/A
(no data)	Sensus iCon [34]	7.6 million
15 min	Elster REX 2 [8]	4.0 million
5 min	Echelon NES echelon	4.4 million
1 min	Itron Centron [14]	14.4 million
1 s	TED 5000 [37]	(no data)

2. BACKGROUND

2.1 Load Profiles

Conventional electromechanical watt-hour meters do not record instantaneous demand, only net energy consumed over time. Thus, they act as memoryless accumulators whose readouts are physically spinning dials. Energy use is measured by computing dial position changes since the last reading (typically by a human meter reader once a month). In contrast, smart meters generate *load profile*s, time series of electric demand, that are delivered to the provider at or near real time. The level of detail in load profiles is useful for load forecasting and fraud detection [23]. Common low cost meters measure epochs at 15 minutes, but more sophisticated models can generate profiles at a second or lower granularity. Table 1 summarizes the capabilities of several market-leading meters with different capabilities.[2]

The three-day load profile for a large 5-bedroom home is shown in Figure 2. A diurnal pattern is observable: peaks are felt in the morning, mid-day, and evening. The drop-out box shows an event occurring about 7pm on the 18th. A plasma television connected to a home theater system was turned on and then off about 5 seconds later using a master switch. The initial large spike represents the power-hungry television, followed by the theater receiver and speaker system powering on. The OFF event shows a symmetric decrease in power draw. NILM algorithms match these *sister* features (ON/OFF features of equal amplitude) against known appliance profiles to uncover in-residence behavior.

2.2 Non-Intrusive Load Monitoring

NILM algorithms extract *appliance profiles* from load profiles. It is considered "non-intrusive" because it does this at the electric meter without instrumenting individual appliances. An appli-

[2]Note that TED identified in the table is not a smart meter, but a in-home device used to monitor energy usage (see Section 4.1).

Basic Definitions

t	A time variable (t_0 is used for an initial time when needed.)
$d(t)$	The net demand from all appliances in the house over time
$u(t)$	The load measured by the smart meter (This includes battery charging)
$c(t)$	The battery's state of charge over time
$b(t)$	The battery's rate of charging over time
	$b(t) > 0$ The battery is charging
	$b(t) < 0$ The battery is discharging
H	The upper safe limit on the battery's state of charge
L	The lower safe limit on the battery's state of charge
K_{SS}	The target constant load value for u(t)

Relations

$u(t) = d(t) + b(t)$ (utility observable profile)

$c(t) = \int_{t_0}^{t} b(t)\,dt + c(t_0) = K_{SS}[t - t_0] - \int_{t_0}^{t} d(t)\,dt + c(t_0)$ (state of charge)

NILL Constraints

$u(t) = K_{SS}$ for some constant K_{SS} (leveled load)

$L < c(t) < H$ (safe state of charge)

Figure 3: Summary of the house and battery model used for NILL.

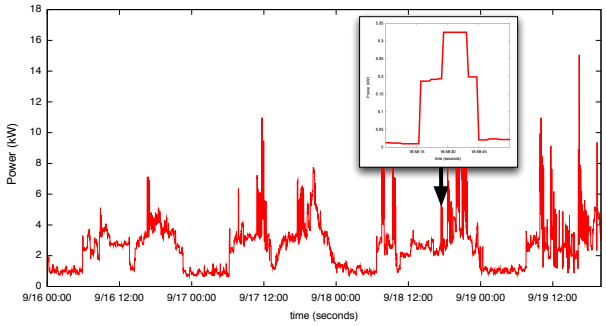

Figure 2: Three-day load profile for large home taken in September 2010. The drop-out box shows a higher resolution view of a television power ON and OFF.

ance profile consists of the types of appliances and the times during which each is operational during the day. Load profiling techniques classify devices by the changes in steady state load caused by their being turned ON and OFF [18]. The approach is to decompose the load profile into a composite of individual appliances features, i.e., representative pairs of ON/OFF events. For example, periodic spikes in energy use are visible in homes with electric furnaces during cold weather. NILM algorithms can extract the expected furnace load from the load profile to expose other, possibly smaller, appliance loads. These techniques use appliance models and information learned about a residence over time to reconstruct behavior from a single aggregate signature. Such techniques have been shown to be highly accurate in practice [25, 21, 33, 22]. We discuss other classes of NILM that are not relevant to residential smart meters in Section 5.3.

2.3 Energy Storage

NILL requires one of a particular class of *deep-cycle* batteries. Deep cycle batteries are designed to be able to operate adequately during long cycles of charging and discharging without significantly reducing their lifetime. Such batteries are frequently used in recreational vehicles such as RVs and boats. There are several types of deep cycle batteries that support highly variable load profiles (at short timescales) present in home energy consumption. The Absorbed Glass Mat (AGM) battery (which has a lead-acid chemistry) has several properties that make it ideal for home use; they work well at extreme temperatures, have low internal resistance, can be charged at high voltages, and are designed to prevent leakage. To avoid sulfation (inability to hold charge due to crystallization of the lead sulfate), deep cycle batteries should not be allowed to discharge below 20% of their total capacity, and staying above 50% is optimal. When a battery is to be charged beyond 90%, its charger should switch to a lower constant voltage than what was used for previous charging [13].

For our evaluation of NILL, we model a 50 Ah[3] lead-acid battery operating at a nominal voltage of 120V. This is achievable by connecting typical 50 Ah sealed DC batteries, which typically retail for approximately $100 [2], in series. One of the most common voltages for these types of batteries is 12V, requiring 10 such batteries (approx. $1,000) to achieve the necessary characteristics. We use a 60 ampere (A) maximum discharge current system as available in modern home solar setups [1].

3. Non-Intrusive Load Leveling

The goal of a NILL system is to level the load profile to a constant *target load*, thus removing appliance features. To achieve this, NILL relies on a battery to offset the power consumed by appliances. When an appliance turns ON, it will exert a load beyond the target load. Thus, NILL will discharge the battery to partially supply the load created by the appliance, maintaining the target load.[4] Similarly, if an appliance enters the OFF state, the load profile will decrease below the target load. These opportunities are used to charge the battery while restoring the target load. The NILL system presented here consists of two parts: a battery and a control system that regulates the battery's charge and discharge based on the present load and battery state. The controller attempts to maintain a steady state target load K_{SS}, but will go into one of two special states K_L or K_H if the battery needs to recover from a low or high state of charge. This section describes the NILL runtime control system and the calculation of the initial system parameters.

3.1 Run Time Control

In a perfect NILL implementation, there would be no runtime control as the battery would have sufficient capacity for maintaining the target load. For any reasonably sized battery, there will be times when the state of charge is insufficient to maintain the target load under a heavy load. We call this a *low recovery state* because the battery's SOC has become too low to maintain the target load.

[3]Ah stands for amp-hours, which is a measure of the battery charge capacity.

[4]The battery is only used to supply appliances in the house. It is never discharged back into the grid as is done in *net metering*.

Similarly, in times of light load, the battery will draw from the utility to maintain the constant load. If however, the load remains light, eventually the battery will reach its maximum SOC. We call this a *high recovery state*.

We use the model shown in Figure 3 in describing the control system and bootstrapping phase in the next section. The model captures both the actual load profile of the house $d(t)$, as well as the load under the influence of NILL as perceived by the electric meter $u(t)$. The essence of NILL is described by the equation, $u(t) = d(t) + b(t)$, where b is the battery's rate of charge over time. If $b(t) > 0$, the battery is charging, otherwise $b(t) < 0$ and the battery is discharging. Finally, $c(t)$ is used to represent the battery's state of charge (SOC).[5] The NILL controller must maintain the target load and respond to low and high recovery states. The controller sits next to the battery at the service-panel or electric meter. A sensor placed on the same line as the meter is used to monitor d, and one on the battery to monitor c. Using these two parameters, the controller selects b to maintain a constant u within the battery's operational constraints.

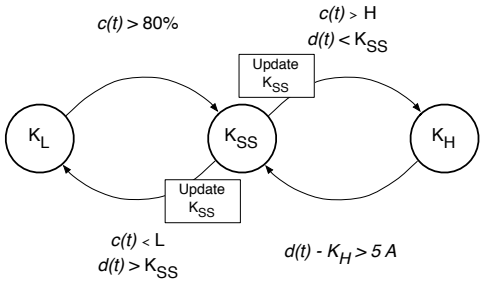

Figure 4: The transitions between the steady state target load and the high and low recovery loads. In state K_i, the battery output is calculated using $K_i - d$.

If the controller enters one of two recovery states, the target load will be altered to allow the battery to reach a safe state while still attempting to mask features present in the load profile. We denote the target load during steady state operation as K_{SS}, and the targets for the low and high recovery states as K_L and K_H respectively. Because entering a recovery state is a sign that the target load was either too high or too low to maintain, it is adjusted each time the controller transitions to a recovery state. Figure 4 illustrates the conditions under which the controller changes states and adjusts K_{SS}. While in state K_i, the battery is controlled using $b = K_i - d$.

In a low recovery state, a new target load K_L is chosen to allow the battery to recharge while still hiding the majority of load events. Thus, K_L is set equal to the battery's maximum sustained charge amperage, masking all load events with amperages less than or equal to this maximum. To reduce the frequency of recovery states over time, the controller will adapt K_{SS} each time a low or high recovery state occurs. This is done using the exponential weighted moving average of the instantaneous demand since the last recovery state at t_r: $K_{SS} \leftarrow \alpha \frac{D}{t - t_r} + (1 - \alpha) K_{SS}$.[6] Once the battery reaches 80% SOC, the system returns to steady state. The effects of low recovery states in our experimental results are shown in Section 4.4.

In a high recovery state, the battery is at its maximum SOC, and the load is below K_{SS}. Once in this state, the only choices are idling the battery, which allows all events to appear in the load profile, or discharging the battery. Because NILL's goal is to cancel appliance level features, we choose the latter. The only question left is the choice of K_H. For this, the controller uses the most recent load samples to guess at a K_H that will be just below the current load (by approximately 1 to 5 amperes). If this guess is not successful after the first few seconds, a more conservative guess is made. With $K_H < d$, the battery can discharge minimally while producing a flat area in the load profile. If the load increases by 5 or more amps, the system returns to steady state. An example of a low recovery state in our experimental results is shown in section 4.3.

[5]Some literature uses Depth of Discharge (DOD), the complementary quantity to SOC.

[6]Note that this is not an EWMA over continuous time samples, but over discrete steady state periods.

3.2 Determining Initial Target Loads

A NILL system requires an initial value for K_{SS} to bootstrap normal operation. A good target load is one that can be sustained with little variation over time. Target load selection is also useful for battery sizing, i.e., if there is no feasible target load for a given battery capacity, a larger battery should be used. For the remainder of this section, we refer to the initial K_{SS} as simply K to distinguish from the steady state target load during run time operation. Two constraints must be satisfied for the target load K to be considered feasible. First, the utility observable profile should be leveled to K at all times. Second, the battery charge and discharge required to achieve $u(t) = K$ must not cause the battery to exceed its safe capacity limits L and H. Under these constraints, the equation for battery SOC can be rewritten in terms of d and K as shown in the RHS of the state of charge relation. This rewriting is used in the algorithm for finding a minimal target load.

Algorithm 1 $FindMinTargetLoad$

1: Given demand $d(t)$, start time t_0, end time t and initial charge $c(t_0)$
2: $D \leftarrow \int_{t_0}^{t} d(t)\, dt$
3: $d_{max} \leftarrow \max_{[t_0,t]} d(t)$
4: Binary search $K \in [0, d_{max}]$
5: Check that $L \leq K[t - t_0] - D + Hc(t_0) \leq H$ over $[t_0, t]$
6: Output minimal satisfactory K

We use Algorithm 1 to find the minimal target load for a given battery capacity. The minimal feasible K is chosen to put the least stress on the battery when coming on line. The input $d(t)$ is a sample of a load profile from the residence hosting the NILL installation. The output is the initial target K. In practice, we select $c(t_0) = 50\%$ SOC, $L = 20\%$ SOC, and $H = 90\%$ SOC to model the safe bounds on battery charge.

4. EVALUATION

4.1 Source Data Collection

The data used in the experiments was collected from devices installed in four homes in the Northeast United States. A TED 5000 [37] measuring transmitting unit (MTU) device was installed in each monitored location and collected real power (kW) and voltage data. TEDs passively measure power and voltage crossing the main circuit between the meter and the circuit panel. Energy readings are transmitted to a TED gateway over the house electrical circuits via power line communications. The gateway is connected via a wired network to a personal computer, which can access readings via HTTP. The TEDs were polled at half-hour intervals to collect per-second load profiles.

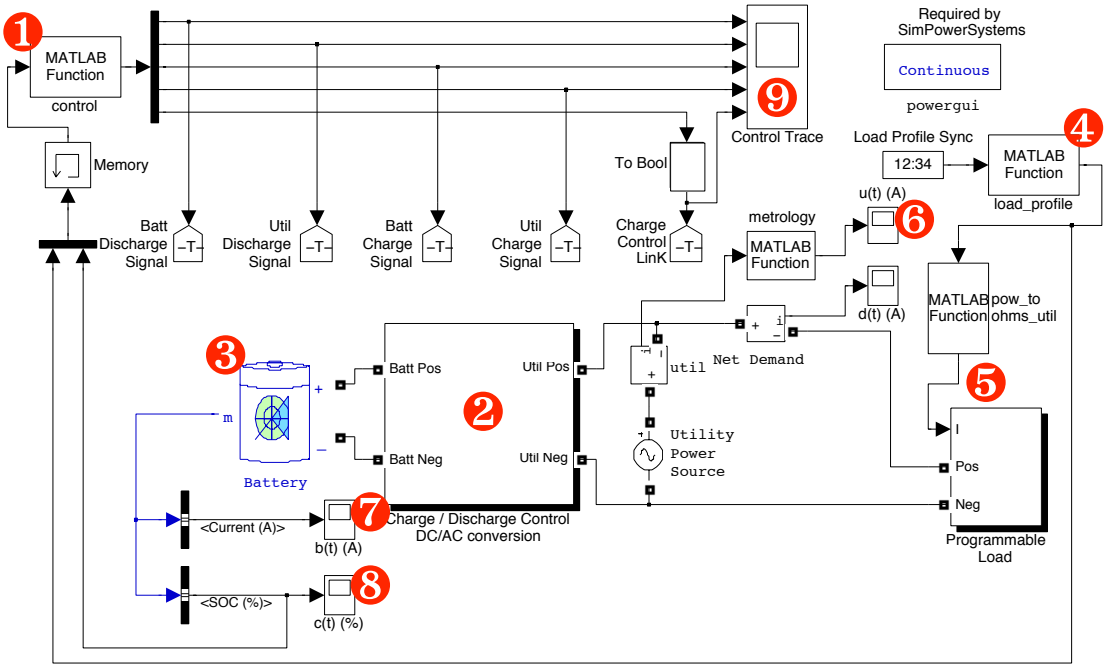

Figure 6: The NILL Simulator. Signals from the control algorithm (1) regulate the switching circuitry (2) that connects the battery (3) and load. The load profile data is inserted at (4) and exerted by a variable resistor (5). The sum of load and battery ($u(t)$) is measured by the meter (6), and rate and state of charge are measured at (7) and (8) respectively. A trace of the control signals (9) was used for debugging purposes.

Table 2: Experimental data sources - energy usage traces collected from residences in spring of 2010

Residence	H1	H2	A1	T1
Start	3/1 9am	4/17 12am	3/15 12am	4/18 11:15am
End	5/1 9am	5/16 12am	4/14 12am	5/17 11am
Length	61 days	30 days	30 days	29 days
Bedrooms	5	2	3	2
Residents	4	3	3	2
Init. K_{SS}	4.64 kW	4.08 kW	3.85 kW	8.20 kW

The experiments in the following section use load profiles collected in the spring of 2010 in the four residences, as described in Table 2. We refer to the data sets as H1, H2, A2, and T1 throughout (homes 1 and 2, apartment 1 and townhouse 1, respectively). The data collection process introduced a small number of sample outages in which no data was collected. This was due to brief power cycles of the TED or lost communication between the TED collector and usage sensor. We repair these gaps by placing repeated samples of the constant average of the surrounding 100 seconds. The H1 data contained two such gaps (1 hour, 3 minutes), H2 contained one gap (19 minutes), A1 contained no outages, and T1 had 2 gaps (11 minutes, 13 minutes). Given their relatively small size, these gaps have little influence on our experimental results.

Figure 5 illustrates energy use over different time scales for one residence, H1. The month profile highlights the relatively constant rate of use over time. Note that during the week of March 3rd, the usage drops off substantially. The occupants left for a spring break during this period, turned down the thermostat, and unplugged appliances throughout the house. The daily energy use exhibits similar diurnal patters as described in Section 2.1. The periodic usage spikes observed in the day-scale data was the result of the home's furnace turning ON and OFF blower motors to force heat in the

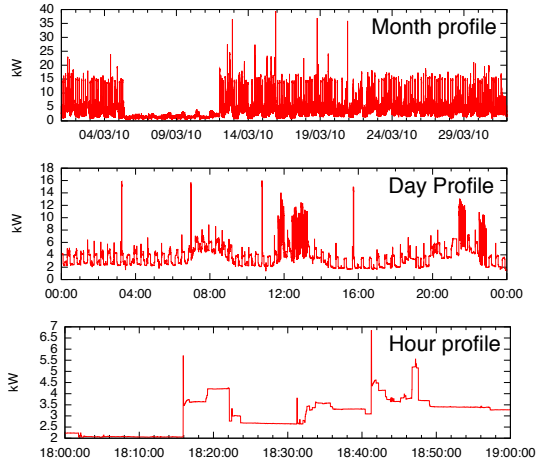

Figure 5: Residence monthly, daily, and hourly profiles.

gas-heated house. Finally, the hourly readings show the enormous sensitivity of these devices to internal use; while spikes caused by starting dryers and heat pump blowers are very clear, small changes caused by appliances such as lamps are also visible.

4.2 Full System Simulation

The NILL system simulation answers the question of how effectively the battery and controller can remove appliance features from a metered load profile. To achieve a realistic battery model, we implemented the simulation in Simulink with the SimPowerSystems extension [26]. This was necessary, as using an oversimplified battery model, e.g. one with linear discharge characteristics, would lead to an inaccurate assessment of NILL's capabilities.

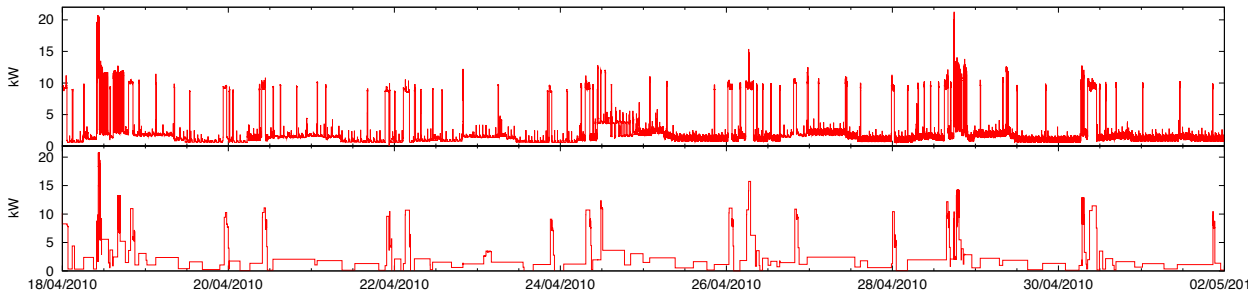

Figure 7: Top: The net load (d) consumed by T1 between Apr. 18 and Apr 30. Bottom: The same load under the simulated NILL system with a 50 Ah battery (u).

Our simulation is shown in Figure 6. At the heart of the system is a circuit with three elements: a net load exerted by appliances, a voltage source for the utility, and a battery. The appliance load is supplied by demand data collected from the TEDs. It is injected into the system as a real valued signal, and converted to an electrical load by a variable resistor in the programmable load. (Like several other residential NILM studies [10, 4, 22], we focus on real power consumption.) The programmable load is supplied by the utility and potentially the battery, which sit in parallel on the circuit. The part of the net load not supplied by the battery will be supplied by the utility voltage source according to Kirchhoff's circuit laws. Battery charging is controlled by the switching connecting the battery to the rest of the circuit. To avoid severe performance slowdowns in our simulation, we implement the battery switching using variable resistance and voltage sources in place of simulating non-linear power electronics. When the battery is discharging, a variable resistance is used to draw current from the battery, and a current source supplies the load proportionally. This process is reversed during charging.

With regards to timing in the circuit, the load profile data is inserted at the rate of one sample per second, the same rate at which it was recorded. The metrology unit on the utility meter decimates the signal to have 1/4 the resolution of the control unit. In reality, control units with MHz clocks can have orders of magnitude higher resolution than solid state metering, giving a full implementation a greater advantage. Forcing sampling at these rates however would have made the simulation run prohibitively slow. Finally, the memory block placed before the control serves to add a one-step delay between the propagation of the sensor signals to the simulator, and the propagation of the resulting control signals to the circuit.

The battery model used in the simulator was originally designed for SimPowerSystems in [38]. It maintains only a single state variable, the state of charge, as well as several other parameters to model voltage during the two phases of battery discharge: the exponential and nominal zones. The additional parameters, including internal resistance, are derived from the battery type, nominal voltage, and capacity. The model has been shown to accurately produce the original manufacturer voltage curves for lead acid batteries, as well as several other basic battery types.

The control unit receives two variables, the net demand and the battery's state of charge. Based on the control model presented in section 3.1, these are used to calculate the charge state and rate of charge for the target K_{SS}. The control also checks the conditions for entering a recovery state and adjusts K_{SS} using EWMA. The control system is implemented in just under 900 lines of Matlab code.

4.2.1 Simulation Procedure

We ran the simulator on each of the four load profile data sets. The runs for H2, A1, and T1 lasted for approximately 3 hours each, and H1 ran for 6 hours. The initial target load K_{SS} for each simulation was calculated by running Algorithm 1 over one week of sample data occurring prior to the load profiles used in the simulation. The initial K_{SS} for each residence is shown in Table 2. Though the initial values are not closely aligned with the longer term run-time steady states, they err on the side of caution due to the strict SOC constraint on line 5 of Algorithm 1. In each case, K_{SS} was either close to the immediate load at the beginning of the trace, or led to a high recovery state with no leaked features. Along with the metered power under NILL $u(t)$, the SOC over time $c(t)$ was included in the output of each simulation.

The results report below were collected using a simulated 50 Ah battery with a 60 A maximum discharge. The α value for the EWMA used to adjust K_{SS} was calibrated experimental to 0.5. Methods for choosing an optimal α and other parameters are covered in Section 5.2. The lower bound on charge was chosen to be 20% SOC. The initial state of the battery was set to 80% SOC.

4.3 Simulation Results

A two-week example of the load profile witnessed by the electric meter (u) compared with the total load (d) is shown in Figure 7. From this example, several things are apparent. First, there are extended periods of hours to days during which no appliance features are visible in u. Second, the majority of features that do occur in u have high amplitudes, and are among the largest present in d. Third, periods of light load in d are accompanied by periods of very light or nearly zero load in u. And finally, the emergent shape of the load profile remains the same for both d and u, and the area under the two curves is approximately the same. We now inspect each of these facets in detail.

Steady state loads comprise the majority of the NILL load profile. During these periods, the controller is able to maintain the target steady state load K_{SS} without depleting or fully charging the battery. Thus, the better the choice of K_{SS}, the longer the steady state can be maintained. Following several high amplitude loads at the beginning of the trace, the steady state converges to durations of half a day or more by Apr. 22$^{\text{nd}}$.

A number of large features appear in u outside of the steady state. These features tend to appear at the peaks of comparatively steady high-amplitude loads. Examples of such events occur on Apr. 18$^{\text{th}}$ and 28$^{\text{th}}$. These features, which are also some of the largest in d, either cause or occur during low recovery state. When the battery is recharging, the target load K_L is set to allow the maximum rate of charge, causing the large yet steady load below the visible features.

Table 3: NILL feature reduction

Residence	Non-NILL	NILL	Change
Total Features			
H1	1047099	61793	-94.10%
H2	286960	20713	-92.78%
A1	430214	24893	-94.21%
T1	384847	33413	-91.32%
Features per hour			
H1	358	21	-94.10%
H2	199	14	-92.79%
A1	289	16	-94.21%
T1	277	24	-91.32%
Sister feature pairs			
H1	340986	10552	-96.91%
H2	110994	4735	-95.73%
A1	176540	6030	-96.58%
T1	147982	8120	-94.51%

After sustained periods of steady state operation or light loads in d, a sharp decrease will occur in u as NILL enters a high recovery state. This happens when the battery has reached full charge while maintaining K_{SS}. In this state, the target load K_H is chosen to be just below some of the recently sampled demand values, so as to allow the battery to discharge at a low rate. As can be seen, in the majority of cases, the choice of K_H is correct on the first try. On a few occasions, most notably after mid day on Apr. 18th, the initial guess is too high, and K_H is lowered by 0.6 kW to allow the battery to discharge. In several instances (such as just prior to Apr. 26th) the NILL controller chose K_H to be slightly below zero, meaning that the battery is effectively discharging to the grid. This is only a concern if the meter can detect reverse power flows, in which case the NILL system can be configured to never set K_L below zero.

Having now covered how NILL reacts to loads, it should not be surprising that on the larger scale of days to weeks, the NILL graph looks similar to the original load. This is result of both the limited capacity batteries afforded by current technology, and typical weather and occupant behavioral patterns. While these larger trends reveal information about customers, most notably the likelihood of human presence in the house, they do not reveal the more fine-grained details such as how many occupants are at home or what their activities are. Thus, we now turn our analysis to quantifying NILL's effectiveness in removing individual appliance features from load profiles.

4.4 Countering NILM

NILM algorithms exploit the amplitude changes in load profiles. These "features" are indicative appliance ON/OFF events (and appliance-internal state changes that effect energy consumption, e.g., washing machine cycles [3]). Features exploited by NILM are represented by $(time, amplitude)$. Here, amplitude is not the absolute metric of energy, but the relative change in energy use from the last sample. For example, a system exhibiting a 5-second load profile: $\{t_0 : 0\text{W}, t_1 : 100\text{W}, t_2 : 200\text{W}, t_3 : 200\text{W}, t_4 : 200\text{W}, t_5 : 100\text{W}\}$ would yield features: $\{(t_1, +100\text{W}), (t_2, +100\text{W}), (t_3, 0\text{W}), (t_4, 0\text{W}), (t_5, -100\text{W})\}$. Samples with amplitude with no change are not features, and are ignored for analysis.

Table 3 shows the feature reduction for the simulated environments. The original TED profile data (Non-NILL) contains just over 1 million features on the 60 day profile of the large house H1, down to the 380 thousand features in the 29 day profile for the small townhouse T1. When the load profile is simulated in NILL-enabled residences (NILL), the number of features drops significantly (around 95% or more). This reduces the number of features from several hundred an hour to as few as 16 in the case of T1.

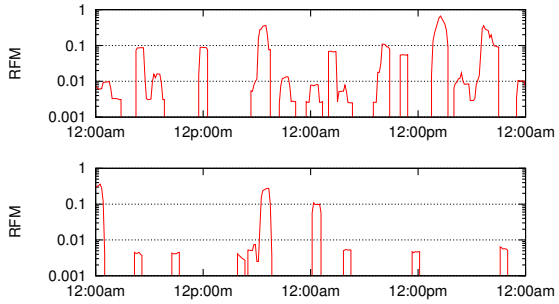

Figure 8: Relative feature mass over two day simulation for H1 (top) and A1 (bottom).

A key question asks if the residual features left after NILL are useful for the NILM analysis. Recall from Section 2 that NILM rely on pairs of symmetric ON/OFF "sister" features to infer appliance usage. We performed a sister matching algorithm that would match all events of greater than 10 W –about 1/6th the draw of a light bulb–as a pairing of positive charge increase and decrease (with a tolerance of 5 W). The matching algorithm linearly scans forward in the trace for the equal and opposite OFF event for each ON event. Table 3 shows that each simulated environment experienced 94% or greater reduction in identifiable sister features.

Raw feature reduction sketches the aggregate behavior of NILL, but only provides a crude estimate of privacy-preservation over time. A *feature mass* ($FM(w, D)$) is the number of non-zero features from over a discrete sample set $D = (d_0, d_1, \ldots, d_w)$, e.g., the number of features over a given time window of size w. In this work, the feature mass is the number of non-zero energy transitions occurring within a time-series interval sample D of length w:

$$FM(w, D) = \sum_{i=0}^{w} (d_i \neq 0)$$

The *relative feature mass* (RFM) for an interval w is the ratio of the feature mass for two sample sets (in this case NILL over non-NILL) over w, e.g.:

$$RFM(w, D) = \frac{FM(w, D_{NILL})}{FM(w, D_{NON\_NILL})}$$

RFM measures the relative number of features of the original and NILL data. Intuitively, as RFM approaches zero, there is very little signal relative to the original profile for a NILM algorithm to operate on. Figure 8 shows the RFM computed at ten minute intervals with a one hour sliding window, i.e., $w = 1\ hour$, over two days at A1 (11 days into the trace) and H1 (29 days into the trace).

The most visible characteristic in Figure 8 is the oscillation of RFM. Most of the trace shows the RFM at 1% or less, and at zero for some of the time (note that the Y axis is on a log scale). The reason for this is the nature of the battery system. When the battery is in the low recovery state, it is unable to suppress most appliance features, and they become visible in the load profile. (We investigate the impact of features occurring during low recovery state in the next section). RFM increases in this off state, peaking in the case of H1 at almost 0.5 RFM for a few minutes. T1 has fewer and less pronounced increases in RFM. Note that the prolonged RFM of zero represents perfect privacy.

To get a sense for the behavior of steady state NILL, consider the normal steady state. We arbitrarily define the system to be in a *NILL-effective* state when the RFM is less than 0.1–which repre-

Table 6: Residual Features

Residence	Features	Sisters	Residual Features (%)	Residual Sisters (%)
H1	1047099	340986	35969 (3.4%)	5526 (1.6%)
H2	286960	110994	13230 (4.6%)	3112 (2.8%)
A1	430214	176540	15556 (3.6%)	3648 (2.0%)
T1	384847	147982	30861 (8.0%)	7640 (5.1%)

Table 4: Feature reduction during NILL-effective ($RFM < 0.1$) state in simulated environments.

Residence	Non-NILL	NILL	Change
Total Features			
H1	879054	6808	-99.23%
H2	225088	1176	-99.48%
A1	354102	1508	-99.57%
T1	265400	1262	-99.52%
Features per hour			
H1	354	3	-99.09%
H2	186	1	-99.12%
A1	279	1	-99.32%
T1	260	2	-99.25%
Sister feature pairs			
H1	291718	561	-99.81%
H2	89558	118	-99.87%
A1	148446	128	-99.91%
T1	104012	161	-99.99%

Table 5: Simulated time in NILL-effective state

Residence	Non-NILL (sec)	NILL (sec)	% of sim.
H1	10527597	8928597	84.81%
H2	5183997	4345199	83.82%
A1	5356797	4570797	85.33%
T1	5009397	3255535	64.99%

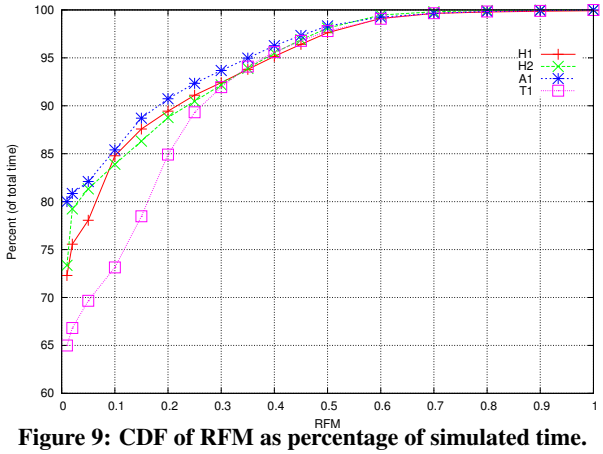

Figure 9: CDF of RFM as percentage of simulated time.

Table 7: Feature Exposure in the Low Recovery State

Res.	Time in Low State (%)	Residual Features (%)	Residual Sisters (%)
H1	891913 (8.5%)	18292 (50.8%)	3363 (60.8%)
H2	340685 (6.6%)	5248 (39.6%)	1414 (45.3%)
A1	435090 (8.1%)	8814 (56.6%)	2524 (69.1%)
T1	390279 (7.8%)	7456 (24.1%)	2172 (28.4%)

sents 10% residual features resulting from NILL. Table 4 shows the features present during intervals when the system is NILL-effective. In all experiments the features, features per hour, and sister features are reduced by over 99%. Table 5 shows that this state is the norm–where between 65% to over 85% of the simulated time was in this state in the simulated environment. One might (incorrectly) attribute the success of in these instances to periods when the load is quiescent. Quite the contrary, the majority of total features in the original data are present during these periods. 83% of the original features for H1 are observed during a NILL effective state, 78.5% of H2, 82% of A1, and 90% of T1.

Figure 9 shows RFM as a CDF–the Y-axis indicates the simulated time in which the system encounters a RFM (X-axis) or less. For all traces the feature mass is zero for 65% to over 80% of the time. Trace time rises steadily until around an RFM of 50%, which is around 2% of the trace for all simulated environments.

4.5 Residual Features

A residual feature is any feature that appears in the original profile that also appears in the NILL trace, e.g., an energy transition that appears at (or about) the same time with similar amplitude (within few watts). Because they precisely reflect appliance activities, these are the only features that are of value to current NILM algorithms. Such features slip through the NILL control system and are observable by the adversary. Table 6 shows the residual features occurring in the simulated environments.

One reason for the presence of residual features is the low recovery state, hereafter called simply low state for brevity, during which the battery is recharging. Table 7 describes NILL behavior in low battery states. Here, we see only 24-56% of the residual features are present in the low state (which occurs less than 10% of the time). Thus, while low state charging explains some of the residual features, many are present in the also steady (non-charging) state.

One can partition residual features into three classes. The first class of features contains very small changes, e.g., less than 60 W, about the load of a typical light bulb. Such loads, the most common in load profiles, are driven by numerous causes including small lighting fixtures, appliance state changes, and small electronics requiring standby power. Because they reside in the band on the smaller end of the power spectrum, it would not be difficult for a NILL system to mimic their arrivals and departures using small variations in the battery charge and discharge rate. We thus do not consider them for the remainder of this analysis.

A second class of features contains large transitions caused by often short-lived heavy loads, e.g., kitchen microwave. The aggregate draw of the large loads combined with ongoing draws exceeds the target load (K_L) plus the maximum output of the battery (60 A for the simulated battery we tested). There is little NILL can do about these loads because it simply cannot safely supply the necessary current. The last class are legitimate loads that are not managed by NILL because of low state. Table 8 provides a breakdown of the residual features present in the simulated environment. Interestingly, 70-85% of the features and 92-98% of the sister features are of the small variety.

Regardless of the reason for their presence, the features usable by NILM to correctly uncover are the large and medium residual sister features. To summarize:

Table 8: Classification of residual features

	Features			Sister Features		
Residence	**Small (%)**	**Medium (%)**	**Large (%)**	**Small (%)**	**Medium (%)**	**Large (%)**
H1	25334 (70.4%)	6257 (17.4%)	4378 (12.1%)	5167 (93.50%)	214 (3.87%)	145 (2.62%)
H2	11319 (85.5%)	1193 (9.0%)	718 (5.4%)	3079 (98.94%)	21 (0.67%)	12 (0.39%)
A1	13139 (84.4%)	1249 (8.0%)	1168 (7.5%)	3555 (97.45%)	45 (1.23%)	48 (1.32%)
T1	25823 (83.7%)	4717 (15.3%)	321 (1.0%)	7512 (98.32%)	118 (1.54%)	10 (0.13%)

Residence	**Sister features**	**Per Day**
H1	359 (0.11%)	5.9
H2	33 (0.03%)	1.1
A1	93 (0.05%)	3.1
T1	128 (0.09%)	4.4

Fewer than ten and as little as one true feature per day that are available to NILM algorithms in our simulated environment. This represents about 3 to 11 out of a thousand features exposed, a minuscule exposure of privacy.

One of the challenges a NILM algorithm might have in the presence of NILL is that it must identify which of the features out of the total trace are legitimate. For example, there are 473 true sister features hidden in 10,552 for H1. Unless there is a clear marker identifying a true feature in the timing/amplitude, then it would be virtually impossible to separate the real and synthetic transitions.

4.6 Entropy Analysis

In addition to comparing the number of features between the load seen by the battery and that seen by the utility, comparing the empirical entropy of two loads is also useful. As in the above discussion, distinguishing information about which appliance or type of appliance was just turned ON (or OFF) lies not in the d and u time series of aggregate load versus time but rather lies in the features (i.e., non-zero impulses) that. Define $\delta d(t) = d(t) - d(t-1)$ and $\delta u(t) = u(t) - u(t-1)$ as the time series of impulses generated by the appliances and seen by the utility, respectively. For each time series we define a probability mass function, $P_{\delta d}$ and $P_{\delta u}$, respectively, which takes a bin size b as a parameter. For concreteness define the ith bin as being the range $[ib, (i+1)b)$ in units of watts. $P_{\delta d}(i|b)$ is just the fraction of the values of δd that are in bin i. $P_{\delta u}$ is defined similarly. Given an empirical probability mass function $P(i)$ we calculate the empirical entropy using base 2 in the standard way:

$$H(P) = -\frac{1}{|I|} \sum_{i \in I} P(i) \log_2 P(i),$$

where I is the support of P, i.e., $I = \{i | P(i) > 0\}$.

Here, $H(P_{\delta d})$ can be interpreted as an upper bound on the number of bits of information a NILM algorithm could extract on average for each impulse sample in δd (if it had access to δd, as it would in the case that our NILL hardware and software were not deployed). Likewise, $H(P_{\delta u})$ can be interpreted as upper bound on the number of bits of information a NILM algorithm could extract on average for each impulse sample in δu, i.e., after the signal has been smoothed by our NILL algorithm. Table 9 reports the empirical entropies for our four data sets, computed using a bin size of 1 W. The bin size was chosen this small to allow for the potential of higher entropy due to higher precision but not any smaller in order to mitigate the introduction of undue noise.

As noted above, the empirical entropy is an upper bound on the information extractable by a NILM algorithm from a time series. Strictly speaking, showing a gap between an upper bound on the information extractable from d and an upper bound on the information extractable from u does not prove that the NILL algorithm decreases the information content of d. That is, the actual information content extractable from d may be below the upper bound on the information extractable from u. However, this would require that the information content of the two signals have significantly different characteristics. In the event that the sequences have similar characteristics, the ratio between the upper bound and the actual information extractable should be similar for each time series.

Note that if a time series is highly correlated, the empirical entropy will overestimate the information in the series. (For example, a series of 100 1's followed by 100 2's will have empirical entropy of 1 bit per time step although the series clearly does not contain 200 bits of information.) There are clearly examples of correlations in our time series. As mentioned above, when a home theatre system was turned on in H1 and H2, several features occurred in succession over a short period of time. In addition to such intra-appliance correlations, some home loads have quasi-periodic behavior. Examples of such loads include refrigerator compressors and home furnace systems.

To understand the extent of these time correlations we computed the autocorrelation function (Pearson variant) of δd and δu. The autocorrelation on input k of an n element times series s compares $s(i)$ with $s(i+k)$ for all $1 \leq i \leq n-k$. The autocorrelation is always at most 1 and at least -1, with values near 1 implying a high degree of correlation, values near -1 implying a high degree of anti-correlation, and values near 0 implying a high degree of independence. The autocorrelation of δd had absolute value less than 0.02 for all values of the lag k. Thus, in spite of the fact that there are some correlated impulses in the data set, the vast majority of impulses appear to be uncorrelated. The non-zero impulses of δu essentially constitute a subsequence of δd. (That is, if at time t, $\delta u(t)$ is non-zero then at the same time $\delta d(t)$ is non-zero as well. If fact, the size of the impulse is often the same in this case.) Since a subsequence of a sequence of independent events is itself independent, it follows that the impulses of δu are largely uncorrelated as well. Indeed, the autocorrelation of the δu time series was similarly small. This argues against the introduction of a systematic bias in empirical entropy of one series versus the other due to differences in time correlations between the two series.

In Table 9, the entropy is computed in two ways: one in which zero values of the respective time series are included when calculating a probability mass function for the time series, and one in which the zero values are excluded. Including the zero values has intuitive appeal as the fact that there are many more zeroes in the NILL time series does reflect that fact that the NILL time series conveys less information to a NILM algorithm. However, to ensure that the number of zeroes in the NILL time series were not biasing the results unduly, we also calculated the entropy of the various time series without the zeroes included. As the entropy represents the average amount of information per sample, and as there are more samples in the Non-NILL time series, we have in-

Table 9: Empirical Entropy

Residence	Non-NILL	NILL	Ratio	Non-NILL w/o zeroes	NILL w/o zeroes	Normalized Ratio
H1	0.633267	0.054265	0.085690	5.663782	7.519812	0.078346
H2	0.235150	0.029554	0.125681	4.516394	6.340144	0.101323
A1	0.363826	0.038459	0.1057071	4.468669	6.526919	0.039289
T1	0.360635	0.019137	0.0530647	5.114588	6.460732	0.084519

Table 10: NILL Cost Reduction per Month in Simulated Environments

Residence	O&R	Ont.	PG&E
H1	$8.94 (2.09%)	$11.11 (2.00%)	$18.67 (1.81%)
H2	$2.49 (5.16%)	$3.78 (4.27%)	$6.17 (4.28%)
A1	$3.41 (3.37%)	$4.96 (3.81%)	$10.22 (4.67%)
T1	$2.67 (2.53%)	$3.72 (2.97%)	$6.92 (2.62%)

cluded a ratio which normalizes for this fact. The numerator of the ratio represents the total information available in the NILL time series (i.e., the entropy times the number of terms in the NILL series without zeroes) and the denominator represents the total information available in the Non-NILL time series (i.e., the entropy times the number of terms in the Non-NILL series without zeroes).

Table 9 reports the empirical entropies for our four data sets. When including the zeroes of the time series, the ratio of the NILL to Non-NILL entropies varies from 0.53 on the low side for data set T1 to 0.126 on the high side for data set H2. When the zeroes are included, the normalized ratio varies from 0.039 on the low side for data set A1 to 0.101 on the high side for data set H2.

5. DISCUSSION

5.1 Energy Efficiency and Consumer Cost

Because the cost of energy generation increases substantially with mid-day demand, the Time of Use (TOU) pricing scheme is being introduced to shave demand during peak hours [17]. TOU creates a cost schedule that is tied to the demand curve observed by the provider. Energy costs rise with expected demand. For example, Orange and Rockland power in New York charges 1.280¢/kWh for off-peak use (9:00pm-10:00am year round), and 7.17¢/kWh for peak use (10:00am-9:00pm) October through May. During June through September, peak charges increase to 7.17¢/kWh for shoulder peak use (10:00am-noon and 7:00pm-9:00pm) and 19.899¢/kWh for high peak use (noon-7:00pm) [30]. TOU is touted as a way to control costs by creating incentives for customers to use energy during non-peak hours.

NILL may positively or negatively impact consumer cost under a TOU schedule. Consider a simplified model of consumer cost in a NILL and non-NILL household. Assume that the total kWh usage for the home is U, the percentage of use at peak is U_p, and the provider costs for peak and off-peak are C_p and C_o per kWh, respectively. The total monthly bill for the home is:

$$T = (U * U_p) * C_p + (U * (1 - U_p)) * C_o$$

According to the U.S. Energy Information Administration, the average home uses 920 kWh per month [39]. Using this figure, we can bound on the cost or savings observable by an average residence under Orange and Rockland TOU cost schedule. In the extreme cases, the costs of using NILL can increase by $24.84 (where the home uses no energy during peak hours) or decrease by as much as $29.35 (where the home no energy during off-peak hours). Based on usage statistics, the load of the average home will use slightly more energy during peak hours ($U_p > \%50$), with more pronounced peak loads in the summer months (due to air conditioning).

Table 10 shows the energy costs in the simulated environments under three time-of-use pricing schemes including Orange and Rockland, Ontario Power Generation [9] (Ontario, Canada), and Pacific Gas and Electric [31] (California). The simulated environments observed cost savings ranging from 1.8% to 5.2%. This reflects the smoothing of load by the battery, where some amount of the peak load is shifted into the off-peak hours. Note that the price of energy is significantly higher in California than the other simulated environments, and thus the costs are about double that of the upper-mid Atlantic and Ontario regions. The costs savings of $2 to $10 may partially offset the costs of the NILL system.

Also contributing to the long-term cost of a NILL scheme is the cost of battery maintenance and replacement as a function of battery lifetime. Under a NILL scheme that exercises a lead-acid battery to between 50 to 80% DOD, the battery can be expected to last between 500 to 1,000 cycles [7], which equates to one or two years. One way to extend battery lifetime is to only use NILL when occupants are at home. Additional measures to reduce cycle frequency are the primary goal of future work for this project.

5.2 Optimizing NILL Parameters

The NILL system simulated in this paper relies on a number of important parameters for its operation. While some guidance has been given for parameter choice, such as the calculation of K_{SS}, others were chosen using expert knowledge. In practice, set rules should be used to choose K_H, K_L, and the EWMA parameter α.

The K_L value chosen in a low recovery state allows the battery to recharge quickly to just below a maximum SOC (80% in our experiments). For many lead acid batteries, optimal lifetime and performance can be achieved by recharging to between 50% and 80% SOC at a high rate, and then throttling back to charge to between 90% to 100% [13]. K_H was chosen using the previous few data points from the net load. While this results in a low discharge from the battery, it is not guaranteed to be minimal. As lead-acid batteries tend to perform best at a high SOC, K_H should be chosen to maximize efficient performance. The EWMA was used to adapt K_{SS} over time as the net load caused recovery states. The better the choice of K_{SS}, the longer a steady state may be maintained. While we achieved good results with $\alpha = 0.5$, this heuristic is not guaranteed to work under arbitrary conditions. Instead, a full implementation should consider techniques such as minimizing Mean Squared Error (MSE).

5.3 Power Signatures and Disaggregation

Two NILM techniques under research that may pose a challenge to NILL are power signature analysis and load disaggregation. Power signature based NILMs collect high-resolution data to detect unique appliance signatures. One example of a device power signature is the turn on transient in demand when starting a motor [35, 41]. Similarly, power signatures can be found in loads with a small ca-

pacitance such as fluorescent lighting and electronic devices [40]. This technique can be highly accurate, but requires equally accurate measurements with a sampling rate of at least 600 samples per second and a resolution of tenths of amps. In the most extreme case a signature-based NILM can characterize the line noise caused by different appliances during transitional phases using MHz sampling rates [32]. As modern residential smart meters simply do not support the sampling rate or computation needed for detecting load signatures, we do not consider them further here.

Load disaggregation techniques attempt to find the set of appliance loads contributing to a net load by solving a discrete knapsack problem. Though this technique was originally considered to be computationally impractical [12], it has recently been demonstrated as feasible for residential loads [20]. In this scenario, the NILL battery represents an unknown quantity in the knapsack problem. At one-watt resolution, there are thousands of possible values for the battery quantity, many of which will have valid solutions to the knapsack fitting. Thus, we believe that load disaggregation based NILMs are not a significant threat to the privacy afforded by NILL.

6. CONCLUSIONS

We have introduced a new class of privacy-preserving algorithms called non-intrusive load leveling (NILL). NILL blinds privacy-exposing NILM algorithms by removing the majority of useful energy use transitions sensed by recently introduced smart meters. Simulations of NILL over real usage profiles in four homes showed that between 1.1 and 5.9 meaningful events were exposed to NILM algorithms per day. Such features reside amongst hundreds or thousands of false events, making recovery of appliance profiles virtually impossible under current algorithms. Future efforts will expand on the analysis to explore tradeoffs between different battery systems and the integration with existing in-home alternative energy generation technology.

7. REFERENCES

[1] Blue Sky Solar Boost 6024HDL. http://www.ecodirect.com/Blue-Sky-Solar-Boost-6024HL-60Amp-12-24-Volt-p/blue-sky-solar-boost-6024hdl.htm.

[2] Rechargeable Sealed Lead Acid Battery. http://www.lowcostbatteries.com/product_p/12v-50ah.htm.

[3] BARANSKI, M., AND VOSS, J. Detecting Patterns of Appliances from Total Load Data Using a Dynamic Programming Approach. *IEEE International Conference on Data Mining* (2004).

[4] BARANSKI, M., AND VOSS, J. Genetic Algorithm for Pattern Detection in NIALM Systems. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics* (2004).

[5] BRUMLEY, D., AND BONEH, D. Remote Timing Attacks are Practical. In *Proceedings of the 12th USENIX Security Symposium* (2003).

[6] CARSON, A. Smart Grids and the Future of Privacy. http://www.greentechmedia.com/articles/read/smart-grids-and-the-future-of-privacy/, 2010.

[7] ELECTROPEDIA. Battery Life (and Death). http://www.mpoweruk.com/life.htm.

[8] ELSTER. REX2-EA meter. http://www.elstermetering.com, 2010.

[9] GREENTERRAFIRMA.COM. Time Of Use Electrical Pricing. http://greenterrafirma.com/time-of-use-electrical-pricing.html, 2010.

[10] GUEDRI, M. E., D'URSO, G., LAJAUNIE, C., AND FLEURY, G. Time-Frequency Characterisation for Electric Load Monitoring. In *Proceedings of the 17th European Signal Processing Conference (EUSIPCO)* (2009).

[11] HART, G. W. Residential Energy Monitoring and Computerized Surveillance via Utility Power Flows. *IEEE Technology and Society Magazine* (1989).

[12] HART, G. W. Nonintrusive Appliance Load Monitoring. *Proceedings of the IEEE* (2004).

[13] HUANG, B., HSU, P., WU, M., AND HO, P. System Dynamic Model and Charging Control of Lead-Acid Battery for Stand-Alone Solar PV System. *Solar Energy 84* (2010).

[14] ITRON. CENTRON Meter. http://www.itron.com, 2010.

[15] JOHN, J. S. Smart Grid Data: Too Much For Privacy, Not Enough For Innovation? http://gigaom.com, 2010.

[16] KALOGRIDIS, G., EFTHYMIOU, C., DENIC, S., LEWIS, T., AND CEPEDA, R. Privacy for Smart Meters: Towards Undetectable Appliance Load Signatures. In *First IEEE International Conference on Smart Grid Communications (SmartGridComm)* (2010).

[17] KING, C. S. The Economics of Real-Time and Time-of-Use Pricing For Residential Consumers. Tech. rep., American Energy Institute, 2001.

[18] LAUGHMAN, C., LEE, K., COX, R., SHAW, S., LEEB, S., NORFORD, L., AND ARMSTRONG, P. Power signature analysis. *IEEE Power and Energy Magazine 1* (2003).

[19] LEMAY, M., GROSS, G., GUNTER, C. A., AND GARG, S. Unified Architecture for Large-Scale Attested Metering. In *Hawaii International Conference on System Sciences* (2007).

[20] LEMAY, M., HAAS, J. J., AND GUNTER, C. A. Collaborative Recommender Systems for Building Automation. In *Hawaii International Conference on System Sciences* (2009).

[21] LEO, A. The measure of power. *MIT Technology Review* (2001).

[22] LISOVICH, M. A., MULLIGAN, D. K., AND WICKER, S. B. Inferring Personal Information from Demand-Response Systems. *IEEE Security and Privacy 8* (2010).

[23] LITOS STRATEGIC COMMUNICATION. The Smart Grid: An Introduction. http://www.oe.energy.gov/SmartGridIntroduction.htm, 2008.

[24] LIU, Y., NING, P., AND REITER, M. K. False Data Injection Attacks against State Estimation in Electric Power Grids. In *Proceedings of the 16th ACM Conference on Computer and Communications Security* (2009).

[25] MARCEAU, M. L., AND ZMEUREANU, R. Nonintrusive Load Disaggregation Computer Program to Estimate the Energy Consumption of Major End Uses in Residential Buildings. *Energy Conversion and Management 41*, 13 (2000).

[26] MATHWORKS. Simpowersystems 5.3. http://www.mathworks.com/products/simpower/, 2010.

[27] MCDANIEL, P., AND MCLAUGHLIN, S. Security and Privacy Challenges in the Smart Grid. *IEEE Security & Privacy Magazine* (2009).

[28] MERITT, R. Stimulus: DoE Readies $4.3 Billion for Smart Grid. *EE Times* (2009).

[29] METERPEDIA.COM. Meter Manufactuer Data, 2010.

[30] ORANGE, AND ROCKLAND. Time of Use Rate. http://www.oru.com/programsandservices/incentivesandrebates/timeofuse.html, 2010.

[31] PACFIC GAS AND ELECTRIC COMPANY. Electric Schedule E-6, Residential Time-Of-Use Service. http://greenterrafirma.com/time-of-use-electrical-pricing.html, 2010.

[32] PATEL, S. N., ROBERTSON, T., KIENTZ, J. A., REYNOLDS, M. S., AND ABOWD, G. D. At the Flick of a Switch: Detecting and Classifying Unique Electrical Events

on the Residential Power Line. In *Proceedings of Ubicomp* (2007).

[33] PRUDENZI, A. A Neuron Nets Based Procedure for Identifying Domestic Appliances Pattern-of-Use From Energy Recordings at Meter Panel. In *Power Engineering Society Winter Meeting* (2002).

[34] SENSUS. iCon APX C&I Meter. http://na.sensus.com/, 2010.

[35] SULTANEM, F. Using Appliance Signatures for Monitoring Residential Loads at Meter Panel Level. *IEEE Transactions on Power Delivery 6*, 4 (1991).

[36] TANG, H., AND MCMILLIN, B. M. Security Property Violation in CPS through Timing. *International Conference on Distributed Computing Systems Workshops 0* (2008).

[37] THE ENERGY DETECTIVE. The TED 5000 Overview. http://www.theenergydetective.com/, 2010.

[38] TREMBLAY, O., DESSAINT, L. A., AND DEKKICHE, A. A Generic Battery Model for the Dynamic Simulation of Hybrid Electric Vehicles. In *2007 IEEE Vehicle Power and Propulsion Conference* (September 2007), IEEE, pp. 284–289.

[39] US ENERGY INFORMATION ADMINISTRATION. Frequently Asked Questions – Electricity. http://tonto.eia.doe.gov/ask/electricity_faqs.asp#electricity_use_home, 2010.

[40] W. K. LEE AND G. S. K FUNG AND H. Y. LAM AND F. H. Y. CHAN AND MARK LUCENTE. Exploration on Load Signatures. In *International Conference on Electrical Engineering* (2004).

[41] YANG, H.-T., CHANG, H.-H., AND LIN, C.-L. Design a Neural Network for Features Selection in Non-intrusive Monitoring of Industrial Electrical Loads. In *11th International Conference on Computer Supported Cooperative Work in Design* (2007).

APPENDIX

A. FUTURE NILM ALGORITHMS

The analysis discussed in the previous section only considered currently available NILM algorithms. Future approaches will apply more powerful machine learning and data mining techniques. Consider the threat models of as-yet-unknown NILM algorithms that may use features more profitably or exploit knowledge of NILL to expose energy usage. In particular, the NILM algorithms considered thus far have been oblivious to the potential presence of NILL. This raises the question, could a NILM aware of the operations of a NILL algorithm recover more information about the actual demand d from the observed demand u than a basic feature-based NILM? While the answer to this question is almost certainly yes, it is currently unclear how this inference could be done, or how it could be countered.

A.1 Feature Correlation

In the evaluation, we saw that there are features in u that do not overlap with those visible in d, i.e., the non residual features. To a basic steady state NILM, these features reveal little to nothing about the appliance that caused the feature. An adversarial NILM, however, may be able to learn a mapping from some of the features present only in u to those in d. This may be done under three separate circumstances. First, the NILL system may have been installed before the NILM. When this is the case, there is no non-NILL data available for training, forcing the NILM to first cluster features from the NILL load profile. If a correlation is found between the NILL clustering and clusters for known appliance features, a mapping may potentially be found between the two sets.

In the second scenario, NILL is installed some time after NILM, revealing both NILL and non-NILL load profiles over mutually exclusive time span–thus allowing a differential analysis of the pre- and post-NILL data. Similar to the previous case, the NILM must find the clustering of the NILL data that most closely matches that of the non-NILL data. Any assumption about the significance of such a correlation, however, can be further leveraged NILL algorithms to cause false inference by the adversary. In the final scenario, which is both the least likely yet most advantageous for the adversary, there may be temporally overlapping NILL and non-NILL samples. If the non-identical features in these two sample sets are highly correlated, a linear mapping from NILL to non-NILL would give a fair approximation of the mapping between features. Once again, a NILM's reliance on any such mapping could be leveraged by NILL to cause false positive appliance classifications by adding random noise or phantom features.

A.2 Inference from NILL Internal State

It has been previously shown that the high level behavior of cyber-physical systems can be determined based on their low level signaling [36]. Similarly, adversarial NILMs may attempt to infer properties of the NILL system's internal state based on features in the NILL load profile. For example, a sudden spike in feature mass is likely indicative of a low recovery state. A low recovery state in turn reveals the state of the battery, i.e. that the state of charge $c \leq L$ and rate of charge $b(t) > 0$. Once the state is known at a point in time, it is possible to determine the state at any future time by observing transitions in the target load. The question that remains is whether or not an adversarial NILM can infer more about d from knowing the internal system state and u, than from knowing u alone.

To infer anything about d beyond the features already overlapping with u, the NILM must know something about b, because $u - b = d$. But can anything be implied about b from knowing NILL's internal state? The previous example demonstrated that in the low recovery scenario it is known that $b > 0$ and likely close to the battery's maximum safe charge rate. However, as was shown experimentally in Section 4.4, there is already a high RFM during a low recovery state, making the overlap between d and u fairly substantial. Whether the remainder of features in u reveal information about d during a recovery scenario is left to future analysis.

If the NILL is in steady state, then b may be either positive or negative depending completely on $K_{SS} - d$. Thus it seems highly unlikely that a NILM can make inferences based on knowing that the NILL is in steady state. Finally, in a high recovery scenario, K_H is chosen to be slightly below the most recent several d values, allowing the battery to discharge at a low rate. In this case, it is known that $b < 0$, and that $K_H - d$ is fairly small. While this is an indication that d is a light load for the duration of the high recovery state, it reveals only the general size of appliances, but not specific information as to which appliances or whether they are automatic or manually operated.