Privacy Concerns in Upcoming Residential and Commercial Demand-Response Systems

Mikhail A. Lisovich and Stephen B. Wicker
School of Electrical and Computer Engineering
Cornell University, Ithaca, New York 14853–3801
Email: mal86@cornell.edu, wicker@ece.cornell.edu

***Abstract*—We explore the privacy concerns arising from the collection of power consumption data in current and future demand-response systems. We claim that in a lax regulatory environment, the detailed household consumption data gathered by advanced metering projects can and will be repurposed by interested parties to reveal personally identifying information such as an individual's activities, preferences, and even beliefs. To develop this claim, we begin with an overview of demand-response technologies and their deployment trends, mentioning both the parties interested in the data and their motivations. We proceed to formalize the notion of privacy and list the types of personal information which can be estimated with current and upcoming monitoring technologies. To support our list, we conduct a small-scale monitoring experiment on a private residence. Our results show that personal information can be estimated with a high degree of accuracy, even with relatively unsophisticated hardware and algorithms. We discuss the implications of our results for future demand-response projects. Our paper concludes with guidelines for data-handling policies which ensure the protection of privacy.**

***Index Terms*—NG-SCADA, Protection, Privacy.**

I. Introduction

The next decades will see a transformation of our nation's power distribution systems. Next generation Supervisory Control and Data Acquisition (NG-SCADA) architectures will precipitate an exponential increase in both the data and control available to consumers and utilities. Utilities are increasingly adopting automated metering, advanced demand response architectures, microgirds, and other systems which will provide cost savings in power generation, increase grid reliability, and create new modes of consumer-utility interaction.

This transformation is already well underway. Recent years have seen several pilot microgrid projects [1], as well as increased deployment of Advanced Metering Infrastructure (AMI) systems by major utilities across the US. AMI systems in particular have been deployed on a large scale by entities such as California Public Utilities Commission [2]. According to a 2006 Federal Energy Regulatory Commission staff report [3], six percent of meters installed in the US are 'smart' meters supporting some advanced metering project, and the number continues to steadily increase.

Next generation SCADA projects will provide many advantages to both the utilities and the consumer. For the power companies, automated metering will reduce collection costs, while the ability to capture detailed usage information will allow for large-scale load research. The results of this research will allow utilities to improve planning and test the effects of various demand side management programs. For the consumer, the projects will result in more information, more control over power use, and the ability to actively participate in power generation. However, increased availability of data, along with emerging use cases, will inevitably create or exacerbate issues of privacy and security.

This paper is part of a larger effort by the TRUST[1] group to explore the confluence of sensor networking, power distribution, privacy, and security issues that will emerge from a substantial increase in power system monitoring at the consumer level. We collaborate with researchers from the Berkeley School of Law in choosing to focus on the privacy risks arising from the collection of power consumption data in current and future demand-response systems. Our joint claim (we will refer to it as the *main claim*) is that in the present regulatory and judicial environment, the detailed household consumption data gathered by advanced metering projects can and will be repurposed by interested parties to reveal personally identifying information about the programs' participants.

Although current projects implement measures to safeguard individuals' privacy and confidentiality, we believe that there exist strong motivations for entities involved in law enforcement, advertising, and criminal enterprises to collect and repurpose power consumption data.

Consumption data in the hands of these entities raises serious ethical concerns - without proper safeguards, these data may be used to commit fraud, initiate unsolicited and invasive advertising, and in the case of law enforcement, to conduct warrantless searches and monitoring that may infringe on individuals' Fourth Amendment rights.

These concerns couple with a disconcerting history of court cases relating to the issue. A recent case (Kyllo v. US [5]) has affirmed the primacy of privacy within the home, prohibiting in-home data-gathering for surveillance purposes without a warrant. However, other landmark cases have treated power consumption data as third-party business records, which have historically not been constitutionally protected against warrantless seizure [6]. Effectively, repurposing of data for law enforcement purposes is currently legal, provided that information is obtained from sources outside the home.

[1]TRUST is a multi-university collaboration focused on privacy & security. You can find out more about it at the TRUST website [4].

Due to these issues, there is a need for discussion between industry, academia, and lawmaking bodies on the privacy aspect of data collection. Both government and industry must be made aware of this erosion of consumer privacy, then urged to adopt attitudes and data-handling policies which will allow demand response technology to evolve while keeping the spirit of the Fourth Amendment intact.

Discussion and advocacy efforts are already underway. Our colleagues from Berkeley have put out an article [6] in the Stanford Technology Law Review chronicling the evolution of court opinion toward energy data privacy and calling for its constitutional protection. They have also collaborated with the California Public Utilities Commission (CPUC) to develop a set of draft guidelines [7] for a secure and privacy-preserving demand response infrastructure. In this paper, we contribute to the discussion by exploring both the legal and the technological/technical aspects of the main claim. On the legal side, we provide our own perspective on the definition of privacy and its relationship to monitoring technology, supporting our arguments with material from landmark court cases as well as new anecdotal evidence. On the technical side, we focus on the manner in which information can be collected and repurposed. Our contributions include highlighting the importance of certain algorithms for extrapolating activity, a formal way of evaluating privacy risks, and a proof of concept technical study.

The rest of this paper is concerned with systematically developing and substantiating certain aspects of our claim. In Section II we discuss the meaning of privacy and provide our own perspective on the legal landscape. In Section III, we familiarize the reader with the current state of advanced metering technology. We also describe Non-Intrusive Load Monitoring (NILM) systems and algorithms, singling them out as a fundamental tool for extrapolating activity. In section IV, we mention some of the parties interested in the data and their motivations for obtaining and repurposing it. In Section V, we aim to formalize these parties' impact on individual privacy by discussing a 'privacy metric' which encompasses the ways that privacy can be infringed. In Section VI, we prove that repurposing is feasible from a technical standpoint by conducting a small-scale monitoring experiment on a private residence. Our results show that personal information can be estimated with a high degree of accuracy, even with relatively unsophisticated hardware and algorithms. In Section VII, we comment on the algorithm's robustness and possible technological solutions to the privacy problem. In Section VIII, we discuss how our experimental methods can be extended to large scales. Finally, in Section IX we summarize data-handling guidelines suggested by our TRUST collaborators and discuss how our findings fit into the ongoing discussion.

II. Privacy and the Law

In 1890 the Harvard Law Review published an article by Samuel Warren and Louis Brandeis entitled The Right to Privacy. This article, often proclaimed the most influential law review article ever written, identified a right to privacy in existing law while decrying the impact of novel technologies; in this case, the instantaneous photograph [8]:

> *Instantaneous photographs and newspaper enterprise have invaded the sacred precincts of private and domestic life; and numerous mechanical devices threaten to make good the prediction that what is whispered in the closet shall be proclaimed from the housetops.*

In what follows we will consider the proposed use and impact of yet another novel technology, power consumption sensors. We maintain that residential power monitoring, while a useful tool in the development of demand response systems, may constitute a significant invasion of the sacred precincts of private and domestic life in the current regulatory environment.

To understand the impact of monitoring, it is necessary to have a working definition for privacy. Privacy has often been characterized in terms of the right to be let alone; this frequently cited reference is from *Cooley on Torts*, (1880) [9], and is referenced, for example, in the Warren and Brandeis article. Subsequent commentators have added further dimensions and complexity to the definition by focusing on control over the revelation of personal information. Erving Goffmans *The Presentation of Self in Everyday Life* is a key reference that focuses on the definition of self in terms of selective disclosure [10]. Others further refine this definition by viewing privacy in terms of personal dignity and liberty. Building on the work of Bentham [11] and Foucault [12], scholars such as Julie Cohen have recognized that questions of informational privacy implicate self-definition, perceptions of public spaces, and the foundation of a liberal democratic state [13]. It follows from such work that when power consumption is monitored at a level at which personally identifying information is collected, issues of behavior modification and personal liberty are implicated.

It is possible and perhaps likely that the utilities will use collected information for purposes other than those stated. Oscar Gandy, Daniel Solove and other have noted the alacrity with which corporations will convert personal information that is collected in the course of business into a commodity and sell it, often without any acknowledgment of the rights of the parties whose personal information is being sold [14] [15]. State agencies are not immune from this tendency state departments of motor vehicles have sold information from driving records to commercial firms. Public outcry forced Congress to legislate against the practice [16].

Assuming that the collection of power consumption data implicates privacy, the question then arises as to whether federal law provides some protection. In Katz v United States (1967), the Supreme Court of the United States ruled that fourth amendment protections against unreasonable search and seizure covered conversations by individuals in a telephone booth [17]. In a concurring opinion, Justice Harlan formulated a two-part test in which it is determined that a search has occurred when the individual (1) has exhibited an actual (subjective) expectation of privacy, and (2) society is prepared to recognize that this expectation is (objectively) reasonable.

The question of an expectation of privacy became the deciding (and negating) element in subsequent decisions involving data held by third parties. In Smith v. Maryland, the U.S. Supreme Court held that the capturing and recording of

dialed telephone numbers did not constitute a search, as any user of a telephone must be aware that they are conveying dialed numbers to a third party the telephone company [18]. Similarly, in United States v Miller, the U. S. Supreme Court held that cancelled checks were the business records of the bank, and that the banks customers had no reasonable expectation of privacy in the date reflected in those records [19].

Protections against actual physical intrusion into the home are still in place, however, and it would seem that the use of data collecting technologies that are somehow equivalent to an intrusion are also forbidden without a warrant. In Kyllo v United States, the U.S. Supreme Court held in a 5-4 ruling that the thermal imaging of a suspects home from outside that home constituted a search [5]. Justice Scalia, writing for the majority compared the sensing to a physical intrusion:

> *Where, as here, the Government uses a device that is not in general public use, to explore details of the home that would previously have been unknowable without physical intrusion, the surveillance is a 'search' and is presumptively unreasonable without a warrant.*

The tenuous status of this understanding is reflected in Justice Stevens dissent. He seems to revisit the same third party, outside the home arguments seen in Smith, Miller, and the earlier Olmstead majority opinion that sanctioned wiretapping without warrants (with a notable dissent from Justice Brandeis). Stevens focused on the primitive nature of the technology in sanctioning its use without a warrant:

> *The Court has crafted a rule that purports to deal with direct observations of the inside of the home, but the case before us merely involves indirect deductions from "off-the-wall" surveillance, that is, observations of the exterior of the home. Those observations were made with a fairly primitive thermal imager that gathered data exposed on the outside of petitioner's home but did not invade any constitutionally protected interest in privacy.*

At least on its surface, it seems that power consumption data falls into the same constitutional bin as dialed numbers and checks information freely given to a third party. We do not agree with this position, but it stands as the basis for a law of information privacy which is the basis for the research described in this paper if one wishes to regard information as private, then one should not let it pass through the walls of ones home. In this paper we will explore means for quantifying and limiting the extent to which personally-identifying information passes across the boundaries of the home create a concern for individual privacy and autonomy.

III. Technology Overview

To familiarize the reader with the technical aspects of the issue, we begin with a brief overview of demand response technologies. We focus primarily on Advanced Metering Infrastructure (AMI) systems. While discussing AMI, we highlight the types of available raw data, as well as access points at which it can be collected by authorized and/or unauthorized parties.

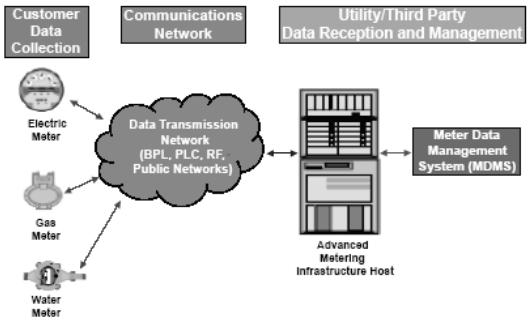

Figure taken directly from [20]

Fig. 1. AMI Building Blocks

We also describe Nonintrusive Load Monitoring (NILM) systems - a more powerful version of AMI capable of extrapolating appliance usage patterns[2]. Our interest in NILM algorithms stems from the fact that they are critical tools for extrapolating activity within the home.

A. Advanced Metering

In a typical Advanced Metering setup, the customer is equipped with solid state electronic meters that collect time-based data at daily, hourly or sub-hourly intervals. The types of available devices differ from project to project, but may include electricity, gas, and water meters. These meters have the ability to transmit the collected data through commonly available fixed networks such as Broadband over Power Line (BPL), Power Line Communications (PLC), and public networks (e.g., landline, cellular, paging). The meter data are received by the AMI host system and sent to the Meter Data Management System (MDMS) that manages data storage and analysis, shaping the information into a form useful for the utility [20]. The typical building blocks of an AMI system are shown in Figure 1.

The data is required to be reasonably complete and accurate. In [22], the specifications are that more than 98% of all meter data make it to the intermediate node, and that the readings have a precision of at least 10 Watt-hours (0.01 kWh).

As mentioned in the introduction, AMI systems have already been deployed on large scales. The reader is referred to [3] for detailed statistics on deployment and system capabilities.

B. Non-Intrusive Load Monitoring

A NILM system collects data like its AMI counterpart, but goes a step further by processing the data to determine the operating schedules of individual electrical loads. This is typically done by disaggregating the collected data stream into individual load signatures and matching each signature with reference signatures stored in a database. For private

[2]For a more complete overview of AMI and NILM, we refer the reader to [20] and [21], respectively.

residences, these loads are usually appliances such as the refrigerator, air conditioner, or water heater.

Several NILM systems of varying capabilities exist, including a commercially available system which can distinguish between major appliances [23], a system based on genetic algorithm which does not need training data [24], and various experimental high-capability systems developed at MIT which identify complex loads and even pinpoint malfunctioning appliances [21].

These systems are used for a wide variety of purposes, including load research, evaluating impact of rate structure changes, implementing incentive programs for particular appliance usage patterns, and handling of high-bill complaints [23]. However, they are important to us because appliance usage information can easily be used to extract user behavior and demographic information.

Current NILM systems require data with a second/sub-second resolution. Because of this, processing is usually done locally, at the electricity meter. However, there are no technical constraints preventing NILM algorithms from running remotely, and useful results may be obtained even with data from an AMI system (i.e. performance degrades gracefully, although hourly readings yield predictably worse results). Therefore, when considering how power consumption data can be repurposed and the kinds of information that can be extracted from it, one should consider a NILM algorithm as an essential building block. We will develop this thought in Section V.

IV. Players, Use Cases and Motivations

Utilities typically have policies which provide a certain amount of protection for utility records and personal information. For example the California Energy Commission requires the consumer's written consent for the release of personal data related to billing, credit, and power usage [25]. Utility records may be released in certain circumstances if customer not identified, though exceptions are made for law enforcement.

Given these policies, there exist agencies, organizations and individuals who have natural motives to use power consumption data for purposes other than load research and demand response. These interested parties fall into two categories, those likely to obtain some/all of the information in the current regulatory environment, and those likely to seek it though illegal means. In the former case, the utility may engage in partnership in a for-profit venture or be required to cooperate by the federal government. In the latter case, the expected proliferation of access points may facilitate unauthorized access. We proceed to list and describe some of these entities, citing precedent where appropriate.

A. Law Enforcement Agencies

By far the most important entities to consider are law enforcement agencies, both on the federal and state levels. These agencies' motivations might range from counter-terrorism surveillance to anti-drug operations and routine law enforcement. They are aided by loopholes contained in current jurisprudence[3], which allow easy access to public utility records and provide legal precedent for their use in prosecuting criminal cases.

There exist pilot programs in several cities where police routinely use public utility records to seek out drug producers. KXAN Austin recently reported that the Austin Police Department has an agreement that allows it to access Austin Energy power usage records without a search warrant [26]. Investigators have used their access to screen consumers for possible drug production, relying on the fact the heat lamps and watering systems used to grow marijuana indoors can vastly increase an average energy bill.

Police and utility representatives claim that such techniques comply with all state and federal investigative laws. While this claim is disputed and the Austin incident is an exceptional case (since many utilities require a subpoena for releasing records), the case sets a precedent for increasingly sophisticated future use of consumption data for law enforcement purposes. Future uses might involve real-time and ex-post-facto tracking for any range of felonies and misdemeanors.

B. Employers

One parameter that can easily be estimated from power usage data is presence - whether or not person(s) are present within a residence (see Sections III, IV). An employer concerned with productivity or false sick-day claims might use presence information to monitor its employees. A 2006 article in The Denver Post [27] details the use of GPS technology embedded in phones to track employees during the work day. In the article, the director of the Electronic Privacy Information Center expresses concern that the technology may be used for off-work tracking, emphasizing the fact that no clear-cut privacy legislation exists to protect workers from potential abuse.

C. Marketing Partners

Behavoir and appliance usage information may potentially be used for directed advertisements. For example, some NILM systems are powerful enough identify specific appliance brands, and may even identify malfunctioning appliances [21]. A marketing company partnering with a utility may use this data to send customers targeted advertisements for repair/upgrade, or more generally derive demographic data for broader advertising claims. While not as invasive as the above use cases, targeted advertising of this sort may meet with consumers' disapproval and must be considered.

D. Criminals

In their article [6], our Trust colleagues give an excellent scenario for criminal abuse of power consumption data: criminals could tap into a Meter Data Management System or simply monitor the unencrypted traffic between it and the individual meters. From the information, they could process the data to obtain occupancy patterns of houses in the entire

[3]The reader is referred to both Section II of our paper and the Stanford Law article[6] for a more in-depth discussion.

neighborhood. Knowledge of these patterns would facilitate burglary or some other property crime.

V. Formalizing Privacy

The previous section showed by way of examples that the evolution of monitoring technology creates real threats to individual privacy. However, it is not apparent just how these threats can be quantified, especially as a function of available data. There is a need for a 'privacy metric', which takes associates the degree of data availability (accuracy of readings, time resolution, types of readings, etc) with potential privacy risks, providing a robust and reliable indicator of overall privacy.

In this section, we briefly show how to approach the construction of such a metric. Although the actual construction is the subject of future work, the insights we gain while thinking about it can be applied to our 'proof of concept' demonstration.

To construct a privacy metric, we need to better understand the nature of the information which can be extracted from available sensor data. Thus, we will start by suggesting a formal framework for extrapolating activity.

A. Extrapolating Activity

Extrapolating activity may be thought of in two stages - during the first 'intermediate' stage, NILM in combination with data from other sensors is used to extract appliance usage, track an individual's position, and match particular individuals to particular observed events. During the second stage, the intermediate data is combined with contextual data (such as the number/age/sex of individuals in the residence, tax and income records, models of typical human behavior). Together, these data are used to identify activities, behaviors, preferences, beliefs, and so on. The two stages are not cleanly separated - raw data may be used directly to estimate a parameter of interest, and determination of some intermediate parameters may rely on contextual information. However, many parameters in the second stage rely on the same intermediate data (e.g. sleeping habits and eating habits may both be extrapolated from tracking data.)

Note that the nature of the sensors will necessarily lead to 'sample impoverishment' - the data collected will almost certainly be insufficient for accurate tracking and event assignment. For example, if several individuals arrive at the house at once, one can't assign the event 'living room light turns on' to a particular individual with any degree of certainty. Also, a person moving through a residence without triggering any appliances or temperature/humidity sensors is invisible to the system. This limitation has to be taken into account when defining second-stage parameters .

There is a clear upper limit for first stage - at most, the gathered information will reveal everything that's happening in the house, yielding precise information about all movements, activities, and even the condition of appliances (though it may not be possible to achieve this limit with current or future in-home sensing systems). However, it's more difficult to define an absolute performance metric for the second stage - the number of specific preferences and beliefs that can be estimated is virtually limitless. In order to develop a comprehensive privacy metric, one needs to carefully define a list of 'important' parameters, basing importance both on how fundamental a parameter is (how many other parameters may be derived from it) and on home/business owners' expectations of privacy. Expectations of privacy, in turn, are partially based on previous abuse incidents (such as the one in Section IV-A). The list of second stage parameters may be hierarchical, with more specific parameters being used to evaluate more general ones. Once an appropriate list is defined and 'importance' values assigned, it is possible to determine the sufficiency of available data based on requirements of current and future NILM, tracking, and other relevant algorithms.

The list of important second-stage parameters form the evaluation criteria. Algorithms for estimating the parameters, along with the corresponding data requirements, provide a method for evaluating the sufficiency of the available data. Together, these provide a metric for how much information may potentially be disclosed by a particular monitoring system. Developing a comprehensive privacy metric is the subject future work for the TRUST Center.

VI. Experiment

Although it is known that first-stage parameters such as appliance usage may be accurately estimated (see performance chart in [23]), to our knowledge no one has tried to extrapolate activity from power consumption data. In this paper we want to prove that activity extrapolation is feasible, thus lending credibility to our thesis and providing an experimental precedent which others can cite in future efforts. To do this, we conduct a small-scale monitoring experiment on a private residence.

A. Experimental Setup

We conducted our experiment in a typical student residence (Figure 2a). For data gathering, we used the Brultech EML energy usage monitor. Figure 2b shows the data gathering setup. The energy monitor was attached to the residence's breaker panel and sent real-time power usage information to a workstation responsible for data collection. The station recorded power usage at intervals of 1 or 15 second(s) and with a resolution of 1 Watt. The same workstation then ran the NILM and behavior extraction algorithms . To evaluate the system's performance, we placed a network of cameras around the residence. We elected to use the Axis 206 network camera (position shown in Figure 2a), which we connected to a workstation using an Ethernet switch. The workstation ran the AXIS Camera Station software and recorded motion events for later processing. The camera control setup is shown in Figure 2b.

B. Experimental Protocols

The experiment was run semi-continuously over a period of two weeks. This time frame allowed us to obtain repeat data for pattern matching while accounting for time constraints. Power and camera data collection software was shut down

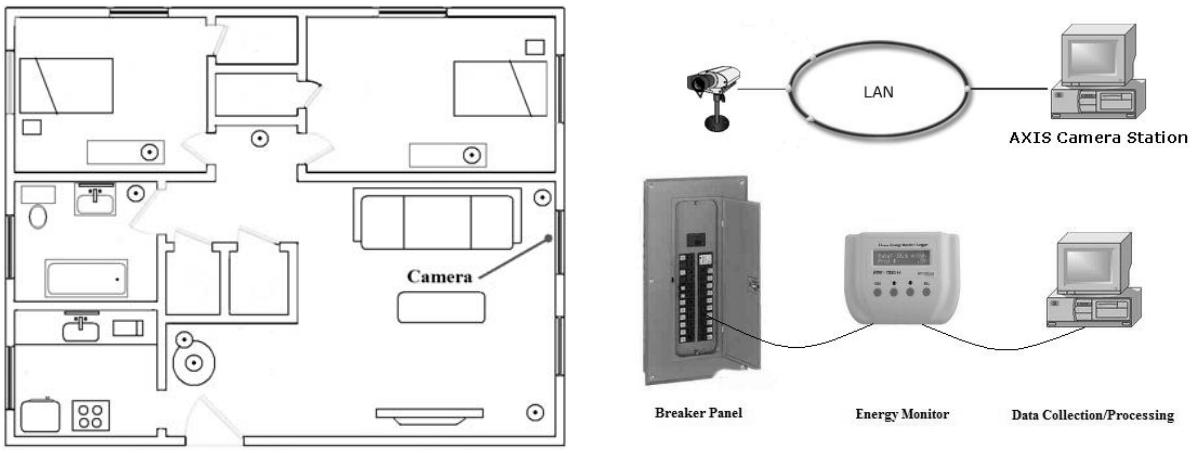

(a) (b)

Fig. 2. Experimental Setup: (a) shows the floorplan of the residence; (b) shows the camera and electrical data gathering setups

on a semi-daily basis for archiving, maintenance, and manual video data processing.

Electrical data was collected from the house breaker panel and stored as a text file with a time resolution of 15 seconds and power resolution of 1 Watt. The data was kept in its raw form for the duration of the experiment and analyzed after its conclusion.

Camera data was collected by the Axis Camera Station software and stored in mpeg format at a resolution of 320x240 at 4 fps. At regular intervals, video data was manually analyzed and processed into activity logs. Upon completion of the logs, the original video data was deleted. Activity logs had the following format:

Date/Time Subject Activity

The subject could be any of the house's three residents or a guest. While residents were identified by name, guests were identified only as *Guest_x*. Possible activities included turning any of the household appliances on or off (ex: *kitchen_lamp_1_on*), entering or leaving the residence, sleeping, preparing meals, taking a bath, or having a party. Note that because the cameras were not put in individual rooms, the resulting activity logs were not fully complete. However, this arrangement respected the residents' privacy and lead to more natural behavior, while the collected data were sufficient to estimate parameters of interest (see Section VI-D for the parameters).

The experiment's participants interacted with the system simply by going about their daily routines. No specific action was required of them, other than notifying house guests about the experiment.

C. Privacy Protections

The experiment involved potentially serious intrusions into the participants' private lives. Therefore, when designing the experiment we took steps to maximize the participants' comfort, minimize potential for embarrassment, and protect their confidentiality.

First, each participant was given a consent form explaining the experiment, detailing their rights, urging them to ask questions, and highlighting the completely voluntary nature of their participation. Participants were free to withdraw from the experiment at any time without penalty. They were also given contact information which they could use to reach us if they had any questions or concerns.

Secondly, video logs were processed by one of the household's residents, which eased the participants' anxiety at being videotaped.

Thirdly, all electrical and video data was kept secure and confidential. Collected data was stored in a password-protected folder, able to be accessed only by individuals directly involved in the project. Also, all publicly available results were stripped of any potentially identifying information.

Finally, the experiment was specifically designed to comply with Cornell Human Subjects Testing guidelines. It has been reviewed and approved by the Cornell Institutional Review Board. The approval request form, consent form, and Experimental Setup & Protocol documents are available from the authors upon request.

D. Parameters to be Estimated

We chose several parameters which were both revealing and possible to estimate using our data gathering equipment and processing algorithms. They are:

- Presence/Absence - whether or not someone is present at the house
- Number of Individuals - if presence is detected, we estimate the number of individuals present.
- Appliance Use - microwave, stove, water heater, TV, misc appliances etc.
- Sleep/wake cycle - when, on average, each individual wakes up and falls asleep.
- Miscellaneous Events - Breakfast, Dinner, Shower, Party, etc.

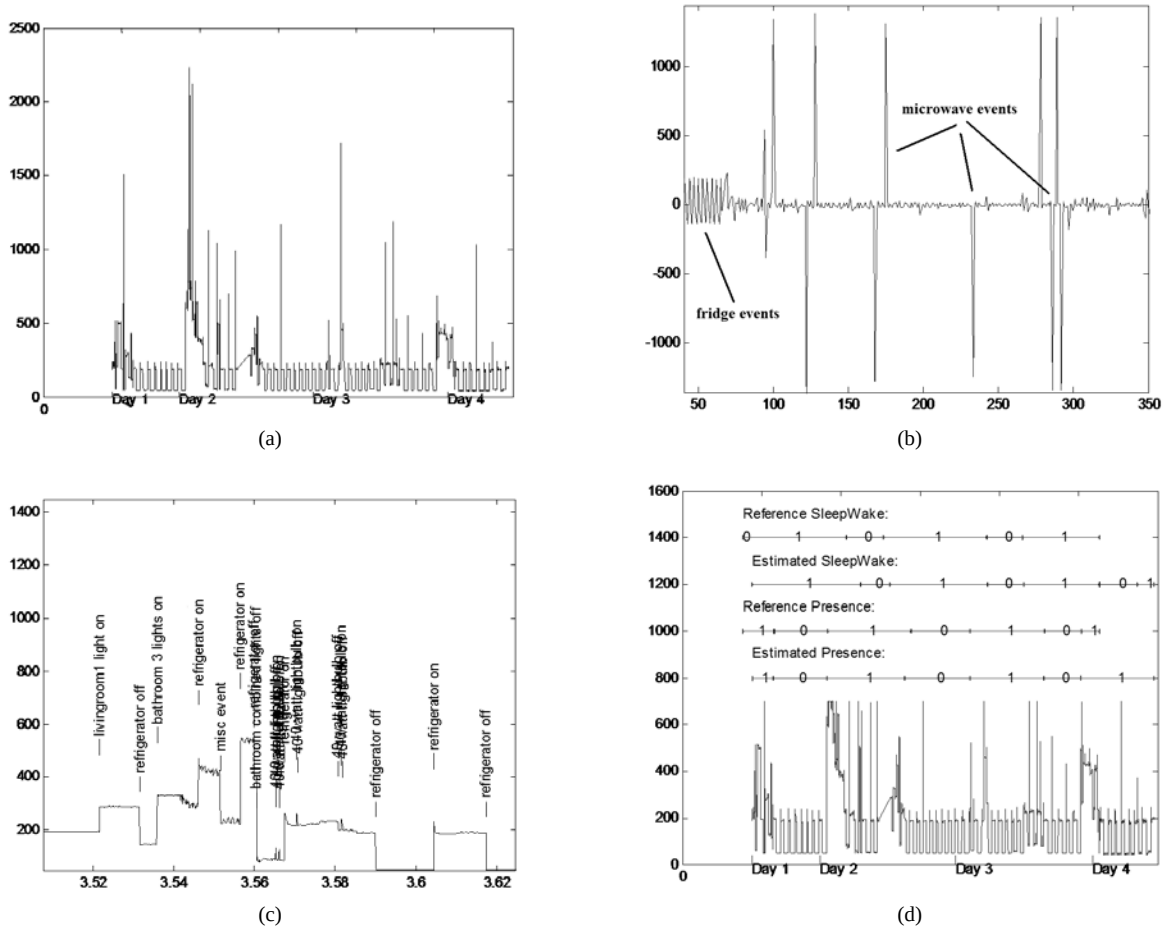

(a) (b) (c) (d)

Fig. 3. Behavoir Extraction Algorithm: (a) shows the aggregate power consumption data; (b) shows the derived switch events; (c) presents several identified load events; (d) compares reference and estimated intervals.

More formally, we begin by combining all data into a single timeline. For each parameter, we partition this timeline into segments, with each segment assigned some value. For most parameters, the value is binary, indicating whether a person present or absent, asleep or awake, etc. The sole exception is the 'Number of Individuals' parameter, which is assigned a set value from the partition $\{0, 1, 2, 3, > 3\}$. For a specific parameter, the i^{th} 'on' interval is defined by T_i^{on} and T_i^{off}. An example partition for the Presence/Absence event is shown below:

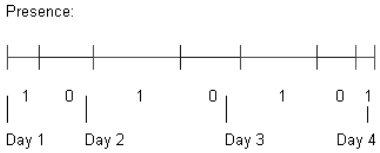

E. Performance Metrics and Evaluation

Once energy use data is gathered and processed with behavior extraction algorithms, we wish to compare the results against reference results obtained from camera data. To do this, we employ two classes of metrics. The first class is event-based and consists of the Failure-to-Detect/ Misdetection percentages for each parameter. These percentages are computed by using the following procedure:

1) Define the cutoff threshold T_{thresh}, choosing it based on experimentation with training data
2) For each parameter, examine the sequence of turn-on/turn-off events on both the reference and estimated intervals.
3) If a camera event occurs but a corresponding electrical event does not occur within T_{thresh} seconds, declare a *Failure to Detect.*
4) If an electrical event occurs but a corresponding camera event does not occur within T_{thresh} seconds, declare a *Misdetection.*

The second class of metrics takes a broader perspective by computing the percentage of the reference interval which is correctly classified. This may in some cases be a better indicator of long-term performance, since the algorithm may miss several short-duration events while classifying the vast majority of the interval correctly.

Together these metrics help one get a well-rounded picture of the algorithm's performance, providing both detail and global perspective.

TABLE I
ALGORITHM PERFORMANCE

	Ref. Events Detected	% Misdetects	% Interval Correctly Classified
Training Data			
Presence	100%	0%	97.3%
Sleep Cycle	100%	0%	93.4%
Shower/Bathroom	35%	38.4%	62%
Microwave Use	33%	76.9%	-
Misc. Events	Unreliable	-	-
Experimental Data			
Presence	80%	20%	97.4%
Sleep Cycle	83%	0%	92.3%
Shower/Bathroom	59%	21.1%	81.8%
Microwave Use	20%	68.7%	-
Misc. Events	Unreliable	-	-

F. Behavior Extraction Algorithms

Our behavior extraction system is implemented in MatLab and consists of two major components: a NILM algorithm and a suite of functions which estimate the high-level parameters mentioned in the previous section. The NILM algorithm we implemented is based on an early MIT prototype [28]. It analyzes the electrical data (Fig. 3a) gathered by the load monitor, performing edge detection, cluster matching, and anomaly resolution.

During edge detection, the algorithm computes a difference series $\Delta(t) = P(t) - P(t-1)$ from the electrical data $P(t)$. Adjacent $\Delta(t)$'s of the same sign and greater than a certain threshold are merged into *switch events* (Fig. 3b).

During cluster matching, switch events are matched against a database of load signatures and classified as either 'on' or 'off' events. A load signature may be a switch event of a certain magnitude (a 40-watt light bulb has a step turn-on signature of $\Delta(t) = 40$ Watts) or a series of such events (a refrigerator has a turn-on signature of $\Delta(t) = 1100$ W, $\Delta(t+1) = -960$ W). Unclassified events are either discarded as noise or labeled with a catchall 'misc. event' classifier. A sample of classified events is shown in Fig. 3c.

During anomaly resolution, the algorithm tries to classify the miscellaneous events as a combination of different turn-on/turn-off events. This allows for classification of events that occur close to each other.

Once the load events are classified, behavior extraction routines use them to determine presence schedules, sleeping cycles, shower & bahtroom use, mealtimes, and other activities. We briefly describe the most important routines:

- Presence - Because the refrigerator is the only load in the residence with automated turn-on/turn-off events, we assume that *any* non-refrigerator event indicates presence. On the other hand, absence is defined by low power usage and lack of events. An extended interval with low power usage during which no events occur implies that all subjects have left the residence.
- Sleep Cycle - Intervals of inactivity which occur between late evening and early morning are likely to imply that all people are sleeping (as opposed to absent). Therefore, all such absence intervals are reclassified as sleep intervals.
- Number of People Present - Estimated based on both frequency of events and the number of lights/appliances simultaneously in use.
- Bath/Lunch/Dinner - Derived from both load events (bathroom lights, microwave, stove fan) and timing information.
- Other appliances - Derived from load events.

The final major component of our system is the analysis suite. Reference data derived from camera logs is automatically processed into reference intervals, which are then compared against estimated intervals using metrics described in Section VI-E. Sample output, showing reference and estimated intervals for both presence and sleep cycles, is shown in Fig. 3d.

G. Results

Our algorithms were run on two sets of data: a smaller three day 'training' set and a larger seven day 'experimental' set. While we actively modified the algorithms to increase performance on the training set, we kept it completely unchanged on the experimental set. The results are shown in Table I. For each estimated quantity, the table's second column gives the percentage of successfully detected reference events, the third column gives the misdetection percentage, and the fourth column states the percentage of reference interval correctly classified.

One important appliance left out of Table I is the refrigerator, which autonomously cycles between high and low states. Unfortunately, we did not directly observe these state transitions directly (this would have required a separate energy monitor exclusively for the refrigerator). However, we can comment on the algorithm's performance by manually examining the electrical data readout. For the training data set, 101 of approximately 104 refrigerator events (more than 97%) were correctly classified. Sucess rate was similarly high for the experimental data set.

Generally, the algorithm performed quite well in determining presence and sleep cycles. In both cases, over 90% of the total interval length was correctly classified, for both training and experimental data. We beleive this is due to our sucess in identifying the refrigerator load, the small number of autonomous appliances in the residence, and the consequent simplicity of presence / sleep-wake heuristics.

Unfortunately, due to time constraints we were not able to implement a routine to determine the number of individuals

present. Shower and bathroom use was detected with moderate success, at least for the experimental data. There was less reliability in classifying other appliances. (Say why this is so)

We note the possible sources of error. They include the limited capabilities of our data gathering system (which can detect only real power, and only at 15 sec intervals), as well as unoptimized decision heuristics. They also include possible errors in the camera logs, since our camera was not in a position to observe all loads directly and turn-on/turn-off events were sometimes missed during manual processing.

To provide perspective on the effectiveness of our algorithm, we compared it to a reference NILM algorithm [29] developed by M. Baranski and J. Voss. The results were reassuring - our algorithm, which was customized for a small residence and preloaded with a load signature library, performed just as well (and in some cases better) than the reference algorithm. Both algorithms identified microwave, refrigerator, and bathroom light events consistently. However, we were able to identify more specific appliances, while a hardwired bias gave us an edge in performance when classifying refrigerator events. Our algorithm identified 97% refrigerator events, while the reference algorithm identified 65%.

VII. Algorithm Robustness and Privacy Solutions

We would like to comment on the algorithm's robustness, and by extension on the informational content contained within the data. We do this by measuring the effect of increased data granularity on the estimation of presence intervals.

Although what follows is a comment rather than a complete analysis, the parameter's tolerance to data scarcity gives an upper bound on the dataset's informational content and provides sufficient ground for us to discuss the relationship between data granularity and privacy solutions.

We believe that privacy protection ultimately lies in policy (see Section IX). However, it's worthwhile to examine technological solutions. Privacy can be preserved through technological means by decreasing the data's information content through signal processing. Such processing may form a useful part of a policy solution - interested parties may be given lower resolution data (resolution depending on its intended use) as a way of ensuring their compliance with stated privacy policies. Additionally, consumers may choose to control the amount of information content leaving their home (in this case, the signal processing is performed in-residence by the meters), exchanging quality of service for privacy protection.

There are several ways to increase the granularity of data. The original dataset can be passed through a filter, downsampled, or corrupted by noise. In particular, a lowpass filter may be applied to remove events of high frequency, masking events which rapidly trigger between'on' and 'off' states . No matter what is done to the high-resolution data, it is important to retain weekly/monthly electricity usage numbers, since the data's analysts will want true averages and totals for billing and research purposes.

We test the algorithm's robustness by performing downsampling with interpolation (averaging over intervals of r datapoints). The results are given in Fig. VII. Our algorithm's

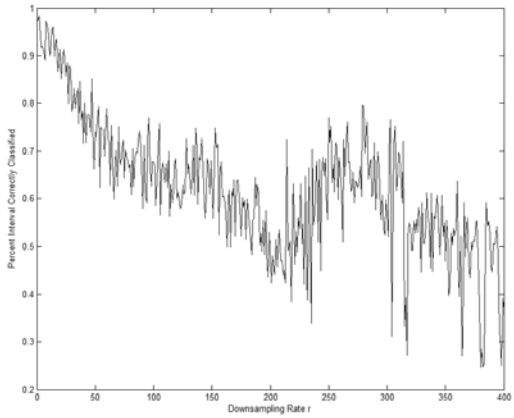

Fig. 4. Algorithm Performance During Downsampling

performance degrades quite gracefully - meaningful estimates are obtained through values as high as $r = 80$ (averaging over 20 minute intervals). This is due to the fact that refrigerator cycles (an integral part of determining presence, see Section VI-F) last for approximately 21 minutes, and are identified as such for lower values of r. In general, it seems that extensive averaging is needed to mask specific events, as the intervals between 'on' and 'off' events is often measured in dozens of minutes.

We note that presence was experimentally proved to be the most robust parameter, as the algorithm's performance for other parameters (especially detecting specific appliances) dropped much quicker. We also note that our algorithm is far from the best when handling low-resolution data, since other decision heuristics (involving power consumption levels as well as switch events) can be implemented to improve performance.

VIII. Discussion

Our experiment shows that presence events and sleep cycles can be estimated with high confidence, at least for a household with few appliances and relatively infrequent switch on/off events.

However, we believe there is potential for vast performance improvements. First, we note that the residence did not have an electric stove or a water boiler - two readily identifiable loads whose 'on' intervals directly correspond to mealtimes, laundry, and showers. Second, we've used only electrical data - a behavior extraction algorithm can combine data streams from electric, water, gas, humidity, and any other available sensors. Third our data resolution (15 seconds in most cases) was relatively low and our behavior extraction algorithms were relatively unsophisticated, as our aim was showcasing feasibility and not optimizing performance. NILM and behavior extraction systems of the near future will surely surpass our effort in performance, enabling person-to-event assignments and perhaps even limited tracking.

On the other hand, we believe that useful data can be extracted by less potent technology. Hourly power averages such as the ones produced by California's AMI system may

also be used to determine presence and sleep cycles (although to a coarser degree). Major appliances a large steady state power consumption (e.g. heat lamps) can also be identified

Note that future concerns are not limited to the performance of these systems the level of on an individual household. Because the algorithms are fully automated, analysis may be done on a extremely large scales, involving hundreds or thousands of residences. Easy access to such personal and demographic information will inevitably generate a market for it.

IX. Guidelines

A report to the California Energy Commission [7], written in part by our Berkeley colleagues, makes several recommendations for power-data handling. They recommend:

1) Multiple tiers of control and oversight, both by the utilities themselves and the state/federal government.
2) Explicit guidelines regulating access to data for customer service, load research, and other functions
3) Strong user control over information leaving the residence.
4) Protocols which do most of the data processing at stations located inside the residence, as well hard prohibitions against relaying certain types of data

One of the authors' main points is that data mining of hourly usage data by utilities be carefully monitored and regulated. The authors advise that utilities should become subject to more stringent rules on the release and re-use of personal data as data mining practices develop and new information in which consumers have a reasonable expectation of privacy is exposed.

In effect, our paper fleshes out the details of this recommendation. Our discussion of interested entities and motivations shows that repurposing of consumption data creates very real privacy concerns for the consumer, and by extension highlights the reasonable expectations of privacy that he or she should develop. Our technical discussion and proof of concept demonstration shows what data mining may be capable of, illustrating the extent to which consumer privacy can be violated. Finally, our privacy metric framework, in combination with the technical discussions, allows one to more precisely define the permitted and prohibited uses of data mining.

We hope that this paper helps those campaigning for privacy and sways those responsible for creating NG-SCADA technologies toward making decisions which both respect and safeguard consumer privacy.

Acknowledgments

The authors would like to sincerely thank Devashree Trivedi, who provided a helpful presence and equally helpful input during every stage of the project, and who single-handedly ran data gathering during the experimental stage. The authors would also like to thank Dr. Michael Baranski, whose papers helped us during development and who provided comparison results using his own NILM algorithm. Finally, the authors would like to thank Judith Cardell, Deirdre Mulligan, Jack Lerner and others who were very helpful throughout the project's duration.

References

[1] "CERTS Microgrid Test Bed," Website. [Online]. Available: http://certs.aeptechlab.com/
[2] C. P. U. Commission, "Proceedings on Demand Response and Advanced Metering." [Online]. Available: http://www.cpuc.ca.gov/PUC/hottopics/1Energy/R0206001.htm
[3] "Assessment of Demand Response and Advanced Metering," Staff Report, aug 2006. [Online]. Available: http://ferc.gov/legal/staff-reports/demand-response.pdf
[4] "TRUST Website," 2007. [Online]. Available: http://www.truststc.org/
[5] "Kyllo v. United States," 533 U.S. 27.
[6] J. I. Lerner and D. K. Mulligan, "Taking the 'Long View' on the Fourth Amendment: Stored Records and the Sanctity of the Home," To be published in the Stanford Technology Law Review, 2008. [Online]. Available: http://certs.aeptechlab.com/
[7] P. Subrahmanyam, D. Wagner, D. K. Mulligan, J. Lerner, U. Shankar, and E. Jones, "Network security architecture for demand response/sensor networks," California Energy Commission, Public Interest Research Group, Tech. Rep., Jan. 2008.
[8] S. Warren and L. Brandeis, "The Right to Privacy," *Harvard Law Review*, vol. 4, 1890.
[9] T. Cooley, *A Treatise on the Law of Torts, Or The Wrongs Which Arise Independent of Contract*, 2nd ed., 1880.
[10] E. Goffman, *The Presentation of Self in Everyday Life*, 1st ed. Anchor, 1959.
[11] J. Bentham, *The Panopticon Writings*, M. Bozovic, Ed. London: Verso, 1993.
[12] M. Foucault, *Discipline And Punish: The Birth Of The Prison*, A. Sheridan, Ed. Pantheon Books, 1977.
[13] J. E. Cohen, "Examined Lives: Informational Privacy and the Subject as Object," *Stanford Law Review*, vol. 52, 2000.
[14] O. Gandy, *The Panoptic Sort: A Political Economy of Personal Information*. Westview Press, 1993.
[15] D. Solove, *The Digital Person: Technology And Privacy In The Information Age*. NYU Press, 2006.
[16] "Drivers Privacy Protection Act," 8 U.S.C. 2721 et. seq. (Public Law 103-322).
[17] "Katz v. United States," 389 U.S. 347, 1967.
[18] "Smith v. Maryland," 442 U.S. 735, 1979.
[19] "United States v. Miller," 425 U.S. 435, 1976.
[20] E. P. R. Institute, "Advanced Metering Infrastructure (AMI)," feb 2007. [Online]. Available: http://www.ferc.gov/EventCalendar/Files/20070423091846-EPRI\%20-\%20Advanced\%20Metering.pdf
[21] C. Laughman, K. Lee, R. Cox, S. Shaw, S. Leeb, N. Les, and P. Armstrong, "Power Signature Analysis," *IEEE Power and Energy Magazine*, vol. 1, no. 2, pp. 1540–7977, Mar. 2003.
[22] O. M. of Energy, "Functional Specification For an Advanced Metering Infrastructure," jul 2006. [Online]. Available: http://www.energy.gov.on.ca/english/pdf/electricity/smartmeters/Functional_Specification_for_Advanced_Metering_Infrastructure.pdf
[23] "Single Point End–Use Energy Disaggregation (SPEED) Marketing Brochure," 2001. [Online]. Available: http://www.enetics.com/downloads/SPEED\%20Brochure.pdf
[24] M. Baranski and V. Jurgen, "Genetic Algorithm for Pattern Detection in NIALM Systems," in *Proc. IEEE International Conference on Systems, Man and Cybernetics*, vol. 4, oct 2004, pp. 3462–3468.
[25] Mulligan, Deirdre K. and Lerner, Jack I. and Jones, Erin and King, Jen and Sislin, Catlin and Wilson, Bethelwel and Hall, Joseph, "Privacy and the Law in Demand Response Energy Systems," Samuelson Law, Technology and Public Policy Clinic, 2006. [Online]. Available: http://www.truststc.org/pubs/36/Jones_PrivacyAndLawInDemandResponse..pdf
[26] K. A. News, "High Utility Bills May Lead Police To Your Door," nov 2007. [Online]. Available: http://www.kxan.com/global/story.asp?s=7322955
[27] T. McGhee, "GPS Technology Tracks Employees," The Denver Post, dec 2006. [Online]. Available: http://certs.aeptechlab.com/
[28] S. Drenker and A. Kader, "Nonintrusive Monitoring of Electric Loads," in *Proc. IEEE Computer Applications in Power*, vol. 12, dec 1999, pp. 47–51.
[29] M. Baranski and V. Jurgen, "Detecting Patterns of Appliances from Total Load Data Using a Dynamic Programming Approach," in *Proc. IEEE Fourth International Conference on Data Mining, (ICDM '04)*, Nov. 2004, pp. 327–330.