Private Memoirs of a Smart Meter

Andrés Molina-Markham, Prashant Shenoy, Kevin Fu, Emmanuel Cecchet, and David Irwin
Department of Computer Science, University of Massachusetts Amherst
{amolina,shenoy,kevinfu,cecchet,irwin}@cs.umass.edu

Abstract

Household smart meters that measure power consumption in real-time at fine granularities are the foundation of a future smart electricity grid. However, the widespread deployment of smart meters has serious privacy implications since they inadvertently leak detailed information about household activities. In this paper, we show that even without *a priori* knowledge of household activities or prior training, it is possible to extract complex usage patterns from smart meter data using off-the-shelf statistical methods. Our analysis uses two months of data from three homes, which we instrumented to log aggregate household power consumption every second. With the data from our small-scale deployment, we demonstrate the potential for power consumption patterns to reveal a range of information, such as how many people are in the home, sleeping routines, eating routines, etc. We then sketch out the design of a privacy-enhancing smart meter architecture that allows an electric utility to achieve its net metering goals without compromising the privacy of its customers.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous; K.4.1 [**Computers and Society**]: Public Policy Issues—*Privacy*; K.6.2 [**Management of Computing and Information Systems**]: Installation Management—*Performance and usage measurement*

General Terms

Design, Standardization, Measurement, Security

Keywords

Smart Meters, Smart Grid, Privacy, Security

1 Introduction

Recently, there has been an increasing focus on "greening the home" using a combination of fine-grained power consumption monitoring, smart appliances, and renewable energy sources, e.g., rooftop solar panels. The trends have led to the design of smart electric grids that provide support for various technologies, including net metering, demand response, distributed generation, and microgrids [15]. An important component of a future smart grid is the installation of smart (or net) meters in homes that support both dynamic pricing and a two-way flow of electricity between homes (or microgrids) and the larger grid. As these meters become more sophisticated, they are able to measure household power consumption at ever finer time-scales. Initial deployments of the Advanced Metering Infrastructure (AMI) in Ontario, Canada support meter readings at 5 to 60 minute intervals [4]. The next generation of smart meters will reduce these time intervals to one minute or less. For instance, in July 2010, PECO, one of the largest providers of electricity and gas in the U.S., selected *Sensus* to provide an AMI with meters that support one minute intervals [24].

In this paper, we argue that the widespread deployment of smart meters has serious privacy implications since they inadvertently leak detailed information about household activities. The information leaks directly correlate with the time granularity that a meter measures power consumption. Unlike traditional dumb meters that record aggregate monthly usage for a utility, today's smart meters allow an utility, or a malicious party, to glean detailed information about household activity in real-time from fine-grained usage measurements. Further, research on nonintrusive load monitoring (NILM) has shown that it is possible to disambiguate individual appliance usage from an aggregate smart meter power trace by using prior knowledge of an appliance's power signature [17]. Such techniques reduce or eliminate the need for outlet- or appliance-level meters, since they are able to extract detailed usage information for individual appliances from an aggregate household power trace.

We show that even without detailed knowledge of appliance signatures *a priori* or prior training, it is possible to extract complex usage patterns from smart meter data using off-the-shelf statistical methods. Our methods are able to label specific types of activity in the home over time based on a number of characteristics, including the level of power consumption, its intermittency, and its duration. In a Facebook-world where users willingly share invasive details of their private lives with friends and strangers, the ability to extract this information may not appear to be an egregious violation of privacy. However, with our limited data, we argue that

it is possible to infer detailed information about household activity—questions such as how many people are in a home at a given time and whether a resident went out for dinner on a particular evening, for example. Entities that gather large amounts of data would potentially be able to predict even more detailed facts, such as residents' genders and ages.

Such information is a foundation for building powerful analytic tools for predicting behavior that could potentially be misused by companies or even criminals. To mitigate these problems, we sketch out a design for a privacy-preserving smart meter architecture that enables an electric utility to achieve its net metering goals, while respecting the privacy of its consumers. The approach leverages the notion of Zero-Knowledge proofs and provides cryptographic guarantees for the integrity, authenticity, and correctness of payments, while allowing variable pricing without revealing the power measurements gathered during a billing period. In Section 2, we describe our infrastructure for gathering power traces in homes and outline simple data mining techniques to identify and label types of household activities. Residents of each home kept power journals for a few days during the sampling period to corroborate our measurements. Next, in Section 3, we describe a secure multi-party computation protocol that uses neighborhood gateways to preserve each home's privacy while enabling net metering.

2 Privacy Concerns with Smart Meters

Recent work by Quinn [23] provides an overview of the privacy implications of fine-grained power consumption monitoring. While Quinn does not present specific techniques or conduct a detailed data analysis, he posits that those with access to smart meter data will be able to infer answers to many questions about a household's personal, and potentially private, activity. While the answers to some of these questions may seem innocuous, e.g., when do people watch TV, others are quite disturbing, e.g., is there a newborn in the house. Table 1 highlights a few of these private questions, along with the power consumption pattern that may reveal their answer. The table also lists the monitoring granularity we believe a smart meter requires to accurately identify the necessary pattern. For instance, a relatively low level of power consumption and variation may indicate that no one is home, while power activity every few hours throughout every night may indicate regular nighttime feedings for a newborn. Even answers to seemingly innocuous questions may prove valuable to third-parties, e.g. for adjusting insurance rates, targeting advertising campaigns, resolving legal disputes, or conducting criminal investigations.

We demonstrate that revealing these complex usage patterns *is not difficult* by developing a simple approach that *opaquely labels different types of household activity*. The approach leverages simple off-the-shelf clustering and pattern recognition techniques on 2 months of power consumption data from 3 homes. To gather the data, we instrumented each home's main circuit breaker with a TED energy monitor [1] that logs household power consumption every second. Figure 1 graphically depicts the architecture. The TED monitor uses the home's power circuits to transmit power readings to a TED gateway that makes them available via a built-in web
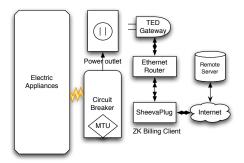

Figure 1. Our architecture using TED monitors/gateways and SheevaPlug computers.

browser. We then use an embedded SheevaPlug computer in each home to download the second-level data each hour from the TED gateway and transmit it to a central repository for analysis. Each entry in the TED data log consists of a power tuple (t, p) that includes a timestamp t and the average power consumption p in kilowatts over the previous second. The one-second logging granularity is smaller than that of existing smart meters [20], which allows us to identify many patterns that are not possible with current meters.

We distill our analysis into 4 steps: 1) pre-process power traces using an off-the-shelf clustering algorithm to identify and label similar types of power events, 2) tag each power event with one or more defining characteristics, 3) filter out automated appliances by observing their signatures during periods of low power activity, and 4) map opaque labels to real-life events using a small amount of externally gathered knowledge.

Label Power Events. We first pre-process each power trace using a density-based clustering algorithm (DBSCAN [7] as implemented by WEKA [11]) to group together power tuples into *power segments*. A *power segment* is simply a collection of tuples with a particular pattern of power consumption values that are adjacent in time. Power segments often have a constant power consumption over a given time period, although this is not required. In some cases, we identify events of the same shape, such as a steep ramp-up and then leveling off. The algorithm labels the power segments such that segments with a similar pattern receive the same label. In many of our figures, we distinguish these labels using different colors. We chose DBSCAN because of its simplicity and our prior familiarity with it, and have not compared it against similar, and potentially more sophisticated, algorithms, such as CLIQUE [2], MAFIA [9], DENCLUE [12]. However, we have found that even our simple approach is able to detect household activities with high accuracy.

Tag Power Events. We append to each power segment a few distinguishing attributes. The primary attributes are each power segment's duration and its power step, i.e, the power increase or decrease at the beginning of the segment. We also label power segments with a particular shape if the power level was not constant. In this case, we identify and label non-constant shapes manually, although it is possible to automate the process. The result is a 6-tuple that includes the segment's label, start time, average power, duration, beginning power step, and shape label. We are able

Question	Pattern	Granularity
Were you home during your sick leave?	Yes: Power activities during the day No: Low power usage during the day	Hour/Minute
Did you get a good night's sleep?	Yes: No power events overnight for at least 6 hours No: Random power events overnight	Hour/Minute
Did you watch the game last night?	Yes: Appliance activity matching TV program No: No power event in accordance with game showtime	Minute/Second
Did you leave late for work?	Yes: Last power event time later than Google maps estimated travel time No: Last power event time leaves enough time for commute	Minute
Did you leave your child home alone?	Yes: Single person activity pattern No: Simultaneous power events in distinct areas of the house	Minute/Second
Do you eat hot or cold breakfast?	Hot: Burst of power events in the morning (microwave/coffee machine/toaster) Cold: No power event matching hot breakfast appliances	Second

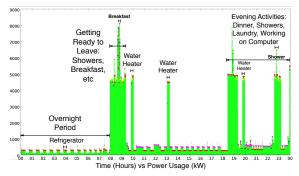Table 1. Private questions and answers that fine-grained power consumption data reveals.

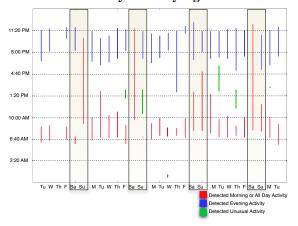Figure 2. Example day-long second-level power trace with labels from the day's activity log.

Figure 3. Identification of human presence with high probability for each day of the month.
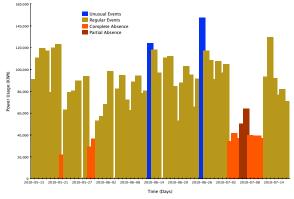

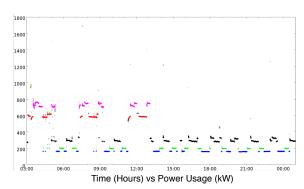Figure 4. Low power periods correspond to little human activity over our two-month trace for one home.

Figure 5. Power signatures for a dehumidifier and an air re-circulator. Note that the dehumidifier shuts off after it fills up.

to automatically process these 6-tuples to answer different types of queries on the data. For example, we identify repetitive usage patterns by filtering for power segments with the same duration, beginning power step, and shape. Figure 2 shows power segments (appended with labels from our activity logs) in a typical day for one of the homes. In this figure, a high variation in color corresponds to human activity, e.g., periods between 8:00 AM - 9:30 AM and 6:30 PM - midnight. Using the intuitive observation that relatively high power consumption and variation indicates human activity, Figure 3 reveals when people were in one of the homes over the course of a month with weekends highlighted.

Filter Automated Appliances. Figure 2 also demonstrates that while nearly all human-triggered power events correspond to the beginning of a power segment, there are many

segments that do not correspond to any human interaction. To obtain only power segments associated with human activity, we filter out the power signatures of automated appliances, such as refrigerators, heating or air conditioning. We leverage the intuitive observation that periods of low power activity correlate well with periods of little human activity to isolate signatures. We illustrate the point in Figure 4, which identifies periods of low activity in a home over our 60-day trace. Likewise, periods of high activity correlate with more people being inside the home, i.e., for a get-together or party. Figure 5 shows power signatures for appliances during an absence from the home. In this case, the signatures correspond to a dehumidifier that runs for 2 hours every 4 hours, and an air re-circulator that runs for 20 minutes every hour.
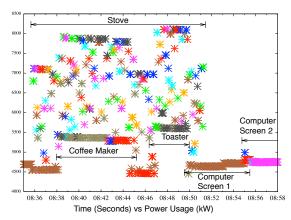
Figure 6. Power segments from eating breakfast. The clustering algorithm automatically generates the color scheme. The labels are from our activity logs.

Figure 7. An example of the same power segments from Figure 6, but at a 30 second logging granularity.

Map Events to Real Life. After collecting and analyzing a sufficient amount of data, it is possible to identify patterns of recurring clusters according to their characteristics. Powerful data mining techniques could be applied to the obtained power segments. For example, the grouped power segments shown in Figure 5 could be filtered out automatically by entering them in a clustering algorithm, this time in *supervised* mode. Alternatively, tagged power segments could also be classified and matched to future occurrences. Further, pattern matching can be improved and past instances can be re-analyzed when new appliances are disambiguated. To illustrate this, Figure 6 shows in detail the disambiguation of power segments (identified by different colors). In this case, clustering distinguishes opaque *events* but not specific appliances or activities. An entity that had access to large amounts of data could then classify these events based on prior knowledge. In our case, we substituted knowledge from activity journals. Each home manually kept detailed power activity journals for at least 3 days over the 60 day period to provide some knowledge of activities in the home. These journals were as accurate as possible, and recorded rough timestamps for turning on and off every light switch and appliance throughout the day. Using the data from our activity journals, we map the opaque power segments to specific types of real-life events. The segments in Figure 6 that were identified by the clustering algorithm have been marked with arrows corresponding to activities logged by the individuals living in the home. The clustering algorithm finds power segments for the stove, coffee maker, toaster and two computer screens, which is enough to answer the question in Table 1 about whether a person had a hot or cold breakfast that morning. Note that the algorithm is able to delimit these segments despite the simultaneous operation of other appliances. To demonstrate the importance of the logging granularity, Figure 7 shows the same trace as Figure 6, but with a 30-second logging granularity. In this case, the pattern reveals little about the usage of each separate component.

Summary. Our analysis demonstrates how easy it is to identify private information from smart meters. We use simple
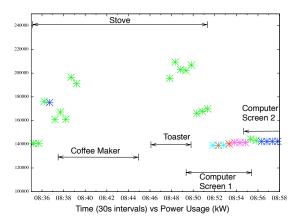
and well-known techniques to identify complex patterns of household activity. Note that utilities with access to thousands of homes will have even more data to leverage to build models to identify particular user behaviors. Further, utilities that have access to detailed power signatures for particular brands and models of household appliances will be able to significantly increase the detail and accuracy of our analysis.

3 Towards a Privacy-Enhancing Architecture

From the perspective of electric utilities, smart meters should meet the following goals: (i) enable critical peak billing and support dynamic pricing, (ii) support tamper and energy theft alarms, (iii) support power failure and restoration notifications, and (iv) support demand response for home automation. [1] In addition to these goals, we add the goal of safeguarding the privacy of the consumer. More precisely, the goals of the utility should be met while avoiding unnecessary information leakage, or, more explicitly, without revealing *when* and *how* energy is being used by a particular household. In this section, we outline a new smart meter architecture that reconciles the security goals required by electricity providers and the privacy goals of the consumer.

3.1 Adversarial Model

Our privacy-enhancing smart meter architecture consists of three components (see Figure 8):

Household Smart Meters gather fine-grained power readings (tuples) that consist of a pseudo-random tuple id or tag, a timestamp, and the corresponding power usage in kW, and relay blinded readings to neighborhood gateways. Our model assumes that consumers have an incentive to tamper with electricity meters to avoid paying bills.

Neighborhood Gateways are computer appliances placed between the smart meters and the remote utility servers. In our architecture, power readings are relayed to utility companies by these gateways, without disclosing which home reported which power tuple, thus concealing their origin. These blinded power tuples consist of a timestamp and the corresponding power usage in kW. This information allows

---

[1]Another potential application includes intrusion detection. During our data collection, one of the authors experienced a car break-in at his apartment. Power traces show precisely when the lights of his carport, activated by a motion sensor, were turned on.
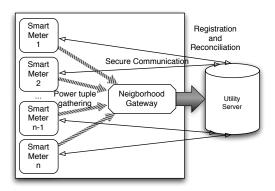
Figure 8. Privacy-enhanced smart metering using neighborhood gateways.

utility companies to meet their goals while protecting the privacy of individual households. Neighborhood gateways may also act as control and storage points in a micro-grid. Communication with both the smart meters and the utility servers is assumed to be over a secure channel that provides authenticity, confidentiality and integrity. For simplicity, in this model neighborhood gateways are assumed trusted. This, however, does not need to be the case. For example, it would be possible to use mixnets or other ways to communicate power tuples from the smart meters to the gateways, see for example [6] .

Remote Utility Servers collect the blinded power tuples from the neighborhood gateways corresponding to all the smart meters in the neighborhood, allowing them to meet their demand response goals. In order to negotiate billing, these servers periodically communicate directly with each smart meter via a Zero-Knowledge protocol, which will be described below. We assume that the utility company's server is an *honest but curious adversary*. That is, the server may attempt to obtain private information about the consumer, but it will follow the protocol correctly and will not provide false information.

3.2 Using Zero-Knowledge Protocols

While fine-grained measurements are necessary to support variable pricing and demand response, we showed in the previous section that a fine-grained power trace is vulnerable to the inferring of private information via powerful data mining techniques. To address the issue, we draw upon Zero-Knowledge protocols as a means to provide privacy to consumers while continuing to employ fine-grained measurements to meet smart metering goals. Zero-Knowledge (ZK) protocols [10] are challenge-response protocols that allow the *prover* to demonstrate the knowledge of a *secret* to the *verifier*, without revealing any partial information that would help the *verifier* infer the secret, other than the fact that the *prover* knows the secret. These protocols typically rely on interactive verifications in which the *verifier* presents a series of challenges to the *prover* that can easily be responded to when the *prover* knows the *secret*, but are extremely difficult to respond to reliably without knowledge of the *secret*; as the number of consecutive challenges increases, the probability of answering these challenges without knowing the secret decreases exponentially.

3.3 Achieving Privacy and Security Goals

Privacy In our case, the smart meter is the *prover*, the utility company's server is the *verifier* and the power trace is the *secret*. The protocol allows a smart meter to report its bill, computed from fine-grained measurements, without revealing *how* or *when* electricity was used, while guaranteeing to the provider that customers do not under-report their usage. Since fine-grained measurements are never sent to the server, privacy is enhanced.

Security ZK protocols are resilient to tampering. That is, not only is it extremely difficult for the *prover* to falsely convince the *verifier* that he knows the *secret*, but the *prover* would also be unable to respond satisfactorily to the *verifier*'s challenges using a *secret* that had been tampered with (in this case, an inaccurate power trace). Additionally, a small number of measurements will be made available to the utility company, which will help ensure that the sensing capabilities of the smart meter have not been compromised.

Further, as mentioned previously, it is important to emphasize that the protocol provides aggregate neighborhood-level usage information to the utility company. Such aggregate information is useful for predicting future demand in an area and facilitates long-term capacity planning for electricity generation.

3.4 A Zero-Knowledge Billing Protocol

The billing protocol that we propose consists of three phases for each billing cycle: *registration*, *tuple gathering*, and *reconciliation*. In practice, however, the registration of tags for the next billing cycle and the reconciliation for the current cycle can happen essentially simultaneously. As a consequence, the direct communication between smart meters and utility servers occurs only once per billing cycle. On the other hand, blinded power tuples are constantly being gathered and relayed by the neighborhood gateways.

1. Registration: In this phase of the protocol, smart meters would *cryptographically commit* to a set of N pseudo-random tags $\{r_i\}$, and a set of *keys* $\{k_1, ..., k_m\}$, where N is the number of power tuples necessary to compute the electricity consumption for a billing period, and m is the number of rounds in the protocol. A higher number of rounds provides more resilience against customer cheating, but in practice a small number, e.g. $m = 10$, provides adequate guarantees. The challenge-response protocol would rely on the fact that the smart meter knows these parameters and that the power tuples generated during consumption are associated to these committed tags.

2. Tuple Gathering: Smart meters create tuples $[r_i, t_i, \rho_i]$, where r_i is a pseudorandom tag from the set of *committed* tags previously produced in the registration phase; t_i denotes the timestamp for the power tuple; and ρ_i denotes the power usage reported by the smart meter in kW at time t_i. The neighborhood gateways do not reveal which random tuples r_i correspond to which smart meter.

3. Reconciliation: During reconciliation, the client computes the bill using variable pricing as $E = \sum_{i=1}^{N} \text{Cost}(t_i, \rho_i)$ using the tuples $[r_i, t_i, \rho_i]$, and submits E to the server. To prove E is correctly computed, the client and server engage in m verification rounds. The smart meter passes each verification round if the cost E was computed using tuples with

tags as well as the keys committed in the registration phase.

Additionally, the billing protocol will be aided by sporadic *random spot checks* that reveal the client identity attached to an insignificant amount of tuples. This spot checking is designed to prevent clients from manipulating tuples. In our case, these *random spot checks* are implemented by allowing the neighborhood gateways to unblind a negligible amount of tuples. Such limited unblinding is useful in gathering limited data to design targeted incentives, e.g., for demand response. This feature will complement the capability of utilities to obtain neighborhood-level aggregates for computing seasonal usage in future demand estimation.

Typically Zero-Knowledge protocols are computationally expensive. However, they have been successfully applied to provide location privacy to vehicular services [22, 3]. While our goals are different from providing location privacy, we use similar insights to provide an efficient implementation for smart meters. Our initial results show that, at minute-level granularity, an entire day's usage measurement can be verified in less than 5 minutes. The same level of efficiency can be achieved for the computation of other aggregate functions that may be needed in smart meters, such as sample means, sample variances, maxima, linear regressions, and sample correlations [19], since our protocol uses *homomorphic encryption*, which ensures that the sum of two *ciphertexts* decrypts to the sum of their corresponding *plaintexts*. This property, in essence, enables computation on encrypted data.

4   Related Work

Prior work exists on distilling information about appliance usage from power traces. However, prior approaches assume knowledge of the appliances in a home or take measurements from appliances in order to use *supervised learning* techniques to disambiguate events. For example, Patel et al. [21] use individual traces from USB oscilloscopes to help filter and disambiguate power traces of particular appliances. Jiang et al. [13] propose a solution to a similar problem that relies on a wireless sensor network to monitor building energy usage. Similarly, Lam et al. [16] proposes to classify appliances based on individual fine-grained load signatures. More accurate monitoring for utilities has many benefits, as prior work discusses, e.g., [18], [5], [14] but they do not propose techniques to preserve privacy of the subjects. One alternative approach to protect privacy is through adding noise to load signatures using rechargeable batteries [8].

5   Conclusion

We highlight the issue of privacy and smart meters. To illustrate the privacy leaks smart meters allow, we develop a simple approach to label usage patterns using off-the-shelf statistical techniques. We then sketch the design of a privacy-enhancing architecture that enables utilities to satisfy their net metering goals. As part of ongoing research, we are implementing our protocol into a prototype smart meter using a TED energy monitor [1] and an embedded ARM-based Linux board that runs our ZK protocol. We are generalizing the billing protocol to include the computation of other useful information that allows utilities to improve their planning capabilities. Finally, we

are formalizing our leakage model to ensure that reporting aggregate information provides adequate privacy guarantees.

References

[1] www.theenergydetective.com.

[2] R. Agrawal, J. Gehrke, and D. Gunopulos. Automatic Subspace Clustering of High Dimensional Data. *Data Mining and Knowledge Discovery*, (11), 2005.

[3] J. Balasch, A. Rial, and C. Troncoso. PrETP: Privacy-Preserving Electronic Toll Pricing. *USENIX Security*, 2010.

[4] A. Cavoukian. Privacy by Design... Take the Challenge. *Information and Privacy Commissioner of Ontario, Canada*, 2009.

[5] G. Cohn, S. Gupta, J. Froehlich, E. Larson, and S. Patel. GasSense: Appliance-level, Single-point Sensing of Gas Activity in the Home. *Pervasive*, 2010.

[6] C. Efthymiou, G. Kalogridis. Smart Grid Privacy via Anonymization of Smart Metering Data. *IEEE SmartGridComm*, 2010.

[7] M. Ester, H. Kriegel, J. Sander, and X. Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. *SIGKDD*, 1996.

[8] G. Kalogridis, C. Efthymiou, S. Denic,T. Lewis, R. Cepeda. Privacy for Smart Meters: Towards Undetectable Appliance Load Signatures. *IEEE SmartGridComm*, 2010.

[9] S. Goil, H. Nagesh, and A. Choudhary. MAFIA: Efficient and Scalable Subspace Clustering for Very Large Data Sets. *Technical Report CPDC-TR-9906-010, Northwestern University*, 1999.

[10] S. Goldwasser, S. Micali, and C. Rackoff. The Knowledge Complexity of Interactive Proof-systems. *SIAM Journal of Computing*, 18, 1989.

[11] M. Hall, E. Frank, G. Holmes, and B. Pfahringer. The WEKA Data Mining Software: An Update. *SIGKDD*, 2009.

[12] A. Hinneburg and H. Gabriel. Denclue2.0: Fast Clustering Based on Kernel Density Estimation. *Advances in Intelligent Data Analysis*, 2007.

[13] X. Jiang, S. Dawson-Haggerty, and P. Dutta. Design and Implementation of a High-fidelity AC Metering Network. *IPSN*, 2009.

[14] Y. Kim, T. Schmid, M. B. Srivastava, and Y. Wang. Challenges in resource monitoring for residential spaces. *BuildSys*, pages 1–6, 2009

[15] K. LaCommare and C. Marnay. Microgrids and Heterogeneous Power Quality and Reliability. *International Journal of Distributed Energy Resources*, 2007.

[16] H. Lam, G. Fung, and W. Lee. A Novel Method to Construct a Taxonomy of Electrical Appliances Based on Load Signatures. *IEEE Transactions on Consumer Electronics, IEEE Transactions on*, 53(2), 2007.

[17] C. Laughman, D. Lee, R. Cox, S. Shaw, S. Leeb, L. Norford, and P. Armstrong. Advanced Non-intrusive Monitoring of Electric Loads. *IEEE Power and Energy Magazine*, pages 56–63, 2003

[18] F. Mattern, T. Staake, and M. Weiss. ICT for green: How Computers Can Help Us to Conserve Energy. *Conference on Energy-Efficient Computing and Networking*, 2010.

[19] A. Molina, M. Salajegheh, and K. Fu. HICCUPS: Health Information Collaborative Collection Using Privacy and Security. *SPIMACS*, 2009.

[20] M. of Energy and O. C. Infrastructure. Functional Specification for an Advanced Metering Infrastructure Version 2. 2007.

[21] S. Patel, T. Robertson, J. Kientz, and M. Reynolds. At the Flick of a Switch: Detecting and Classifying Unique Electrical Events on the Residential Power Line. *UbiComp*, 2007.

[22] R. Popa, H. Balakrishnan, and A. Blumberg. VPriv: Protecting Privacy in Location-based Vehicular Services. *USENIX Security*, 2009.

[23] E. Quinn. Smart Metering and Privacy: Existing Law and Competing Policies. *A Report for the Colorado Public Utilities Commision*, 2009.

[24] Sensus. FlexNet AMI System.