

A Privacy-Aware Architecture For Demand Response Systems

Stephen Wicker, Robert Thomas
School of ECE, Cornell University

Abstract

We explore the privacy issues implicated by the development of demand response systems. We begin by highlighting the invasive nature of fine-granularity power consumption data, showing that the data collected by Advanced Metering Infrastructure (AMI) reveals detailed information about behavior within the home. We then show how privacy-aware design principles lead to novel system architectures that realize the benefits of demand response without requiring that AMI data be centrally collected. The resulting systems avoid both harm to subscribers and the potential need to scrap AMI-based demand response efforts in the face of public outcry. We also show that Trusted Platform Modules can be used to develop privacy-sensitive metering infrastructure.

1. Introduction

Demand response systems balance daily power consumption patterns by showing consumers the cost of electricity at different times throughout the day. By reducing variation in load, demand response has the potential for an up to 20% reduction in peak load during summer months [1]. Demand response depends on fine granularity power consumption data to predict load, provide future pricing information, and to show the consumer the cost of his or her consumption. Such consumption data is provided through Advanced Metering Infrastructure (AMI) [2]. It has been shown that power consumption data creates a serious privacy concern in that it can be used to deduce details personal information regarding behavior within the home [3]. In this paper we address this problem by applying privacy-aware design practices to the development of demand response architectures.

Privacy-aware design is a design methodology that highlights privacy concerns and guides the practicing engineer or computer scientist in the creation of systems that minimize the privacy concerns of users and the public at large [4]. These design principles were derived from the Fair Information Practices developed by HEW in the 1970s [5]. At their core lie requirements that the collection of personally

identifying information be minimized. We propose that such collection should be a functional requirement of the system, and that when collected, data be used locally wherever possible. The latter results in a distributed processing requirement that drives the design of architectures for a wide variety of information processing systems.

We apply these privacy-aware design principles to the development of demand response systems, showing that a demand response architecture can be developed that meets all of the mission requirements for demand response without the centralized collection of privacy-sensitive power consumption data.

We begin by reviewing the potential for demand response systems, and then show that AMI data presents a serious privacy concern. We then present a series of privacy-aware design practices, highlighting the need to minimize data collection and to use distributed processing wherever possible. We then demonstrate, through an investigation of demand response system design, how privacy-aware design can be applied.

2. Demand response systems

Utilities are adopting microgrids and other systems that will provide cost savings in power generation, increase grid reliability and flexibility, and create new modes of consumer-utility interaction [6]. Demand response systems will play a key role in this effort. Generally speaking, demand response systems modify electricity consumption behavior by end-use customers in response to changes in the price of electricity over time [12]. The modifications, whether induced by presenting pricing information to the customer or through direct control of appliances by the utility, may alter the timing of demand, the level of instantaneous demand, or the total demand over a given period of time [14]. The overall goal is to balance electricity consumption over time, alleviating the utilities' (expensive) need to take generators on and offline.

Demand response systems require power consumption information at a level of granularity far finer than that required for monthly billing. The reason is simple – if consumption is to be modified in accord

with price over the course of the day, then consumption information must be available at the same level of granularity as the pricing information in order to properly bill the customer. The solution lies in Advanced Metering Infrastructure (AMI) – technology that can sample and record power consumption on a minute-by-minute basis, as opposed to the once-a-month meter readings of the past. AMI deployment has been underway for several years. The Federal Energy Regulatory Commission estimated that there were 7.95 million advanced meters installed nationwide in 2009 [6]. By 2009, twenty-six utilities in 19 states had announced or pursued advanced metering pilots or full-deployment programs.

The potential impact of demand response is immense. As seen in the figure below, depending on the extent of the distribution of AMI, the potential savings in energy in the United States during the peak summer period for electrical demand ranges from 4% - 20% of total load. The subsequent positive impact on the U.S. need for foreign oil and related resources would be difficult to overstate.

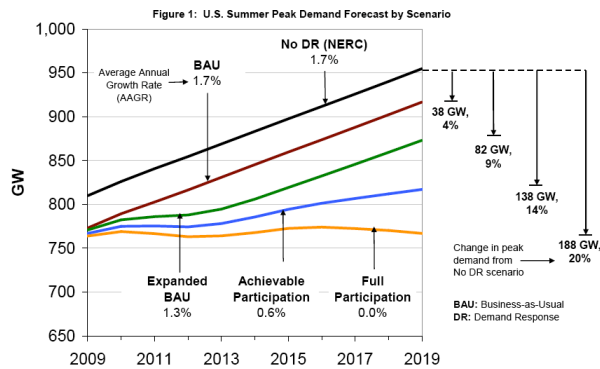


Figure 1. Assessment of the potential for demand response [1]

Looking more closely at Figure 1, one can see that the extent of the power savings is a function of AMI participation. An explanation of the various scenarios is provided below.

Table 1. Key Differences in Scenario Assumptions, [1]

Assumption	Business-as-Usual	Expanded BAU	Achievable Participation	Full Participation
AMI deployment	Partial Deployment	Partial deployment	Full deployment	Full deployment
Dynamic pricing participation (of eligible)	Today's level	Voluntary (opt-in); 5%	Default (opt-out); 60% to 75%	Universal (mandatory); 100%
Eligible customers offered enabling tech	None	None	95%	100%
Eligible customers accepting enabling tech	None	None	60%	100%
Basis for non-pricing participation rate	Today's level	"Best practices" estimate	"Best practices" estimate	"Best practices" estimate

In comparing the above table to Figure 1, note that the energy savings from the “opt-in” participation scenario is estimated at 9%, while that of the mandatory, universal approach is 20%. The reduction in peak consumption is thus more than doubled if regulators require that consumers have advanced metering installed at their homes. This will be an issue of national significance, for unless AMI is used properly, it poses a serious privacy threat.

3. AMI and the threat to privacy

It has been shown that the detailed power consumption data collected by advanced metering systems reveals information about in-home activities. Furthermore, such data can be combined with other readily available information to discover even more about occupant's activities [3, 15].

Reference [3] describes an experiment conducted in a standard student residence. A Brultech EML energy usage monitor was attached to the residence's breaker panel to collect real-time power consumption data. The data, obtained at intervals of 1 or 15 second(s) with a resolution of 1 Watt, was transferred to a non-intrusive load monitor (NILM) application running on a workstation. A behavior extraction algorithm was then run on the workstation in an attempt to predict behavior based solely on power consumption. Video data was used to establish a control for the experiment.

Some of the results from the experiment are reproduced below in Figure 2. Figure 2(a) below depicts aggregate power consumption data over the course of several days. Figure 2(b) then shows fine granularity data over the course of several hundred seconds, depicting an ability to isolate specific device switching events. Figure 2(c) and 2(d) then show how specific load events can be tied to individual behavior.

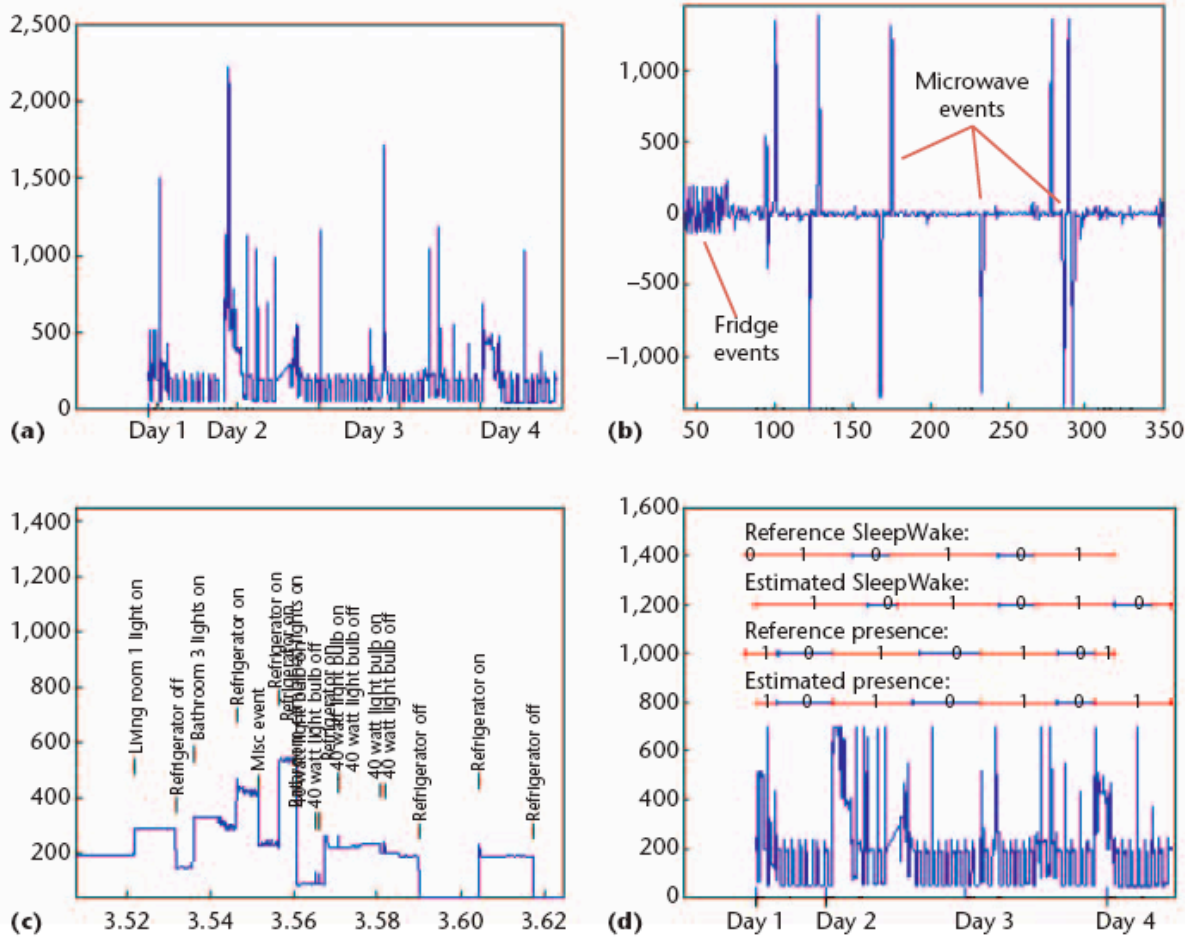


Figure 2. Behavior-extraction algorithm. (a) the aggregate power-consumption data, (b) the derived switch events, (c) several identified load events, and (d) a comparison between reference and estimated intervals. [3]

Table 2. Performance of Behavior Extraction Algorithm, [3]

	Sample Size	Ref. Events Detected	Percent Misdetects	Percent Interval Correctly Classified
Training Data				
Presence	8 Ref., 8 Est.	100%	0%	97.3%
Sleep Cycle	6 Ref., 6 Est.	100%	0%	93.4%
Microwave	8 Ref., 8 Est.	50%	78%	43%
Bathroom Lights	8 Ref., 8 Est.	72%	44%	52%
Passage Light	8 Ref., 82 Est.	38%	90%	57%
Living Room Lights	8 Ref., 8 Est.	55%	88%	58%
Experimental Data				
Presence	10 Ref., 10 Est.	80%	20%	97.4%
Sleep Cycle	12 Ref., 10 Est.	83%	0%	92.3%
Microwave	10 Ref., 58 Est.	80%	83%	99%
Bathroom Lights	60 Ref., 103 Est.	63%	42%	81%
Passage Light	8 Ref., 82 Est.	38%	90%	57%
Living Room Lights	19 Ref., 179 Est.	21%	89%	52%

4. Privacy-aware design

In this section we summarize a framework, developed in [4], for privacy-aware design. The framework consists of a set of principles that were derived from the Fair Information Practices proposed by the Department of Health, Education, and Welfare (HEW) in a 1973 study entitled *Records, Computers, and the Rights of Citizen* [5]. The framework is outlined below.

1. Provide Full Disclosure of Data Collection

- 1.1. Description Requirement
- 1.2. Enforceability Requirement
- 1.3. Irrevocability Requirement
- 1.4. Intelligibility Requirement

2. **Require Consent to Data Collection**
 - 2.1. Acknowledgement Requirement
 - 2.2. Opt-In Requirement
3. **Minimize Collection of Personal Data**
 - 3.1. Functional Requirement for Collection
 - 3.2. Distributed Processing Requirement
4. **Minimize Identification of Data with Individuals**
 - 4.1. Non-Attribution Requirement
 - 4.2. Separate Storage Requirement
5. **Minimize and Secure Data Retention**
 - 5.1. Functional Requirement for Retention
 - 5.2. Security Requirement
 - 5.3. Non-Reusability Requirement

The principles are briefly summarized below.

5.1. Provide full disclosure of data collection

Disclosure has long been recognized as a critical element of public data collection, and was a prominent component of the Fair Information Practices. It has been proposed that the disclosure of data collection in the specific context of an information network should take the form of a public statement that personal data will be collected, and a full characterization of the type of data to be collected.

The first element of disclosure is the **description requirement**. An adequate description includes the type of data to be collected; this should be very specific, including details such as resolution or granularity. An adequate disclosure of resolution would specify the granularity – how often are power consumption data samples being taken?

There must also be a clear indication as to how long the data will be retained, and the means by which it will be retained. A focus on the means for retention opens the possibility for the advancement of privacy-aware storage technologies (for example, technologies that limit or prevent data reuse, or technologies that allow a subscriber to retain control over his or her data). Informed subscribers may prefer one storage technology to another, motivating operating companies through market forces to adopt the preferred technologies. Finally, the description should also include the use to which the collected data will be put.

The effectiveness of a disclosure requirement is strongly dependent on an **enforcement requirement**. The threat of punishment must be of sufficient magnitude and certainty that a collecting entity will be motivated to provide a clear disclosure and to comply with it.

It is also important that a given data set always be treated according to the privacy policy under which it was collected. An **irrevocability requirement** should be applied to collected data so that the customer will have some certainty as to how collected data will be treated for the duration of its retention.

The extent to which a subscriber feels secure in his or her communications will lie in part on that subscriber's understanding of the data collection disclosure. It follows that there must be an **intelligibility requirement** for data collection disclosures.

5.2. Require consent to data collection

The term “consent” is loaded with legal implications. For the purposes of privacy protection, consent is the flipside of disclosure – it establishes the disclosure as a contract. A requirement for consent also serves a pedagogical purpose – it alerts the user to the presence of data collection, and may heighten the awareness of the presence of a potential privacy issue.

It has been proposed that any subscriber/user of a given communication technology must acknowledge the data collection disclosure before they can use the technology. The **acknowledgement requirement** can take the same form as license agreements for software updates. The user must click an appropriate button on a screen before proceeding to use the technology. Such acknowledgements are also found in car GPS units.

The technology that underlies a given service may change over time. A residential consumer may be associated with a given power utility for a long period of time, during which power consumption monitoring technology has changed dramatically. If data collection practices change, the user should be notified. Furthermore, user consent to such alterations should take the form of an **opt-in requirement**, as opposed to one of opting out. The former clearly increases the extent to which the consumer understands and acknowledges data collection [7].

5.3. Minimize collection of personal data

“Personal data” is data that identifies or is correlated with the behavior, thoughts, and/or preferences of an individual. It has been shown that residential power consumption data can be correlated with the behavior of individuals within a house; the finer the resolution, the more detailed the disclosure [3].

The first requirement under this heading is probably the most important of all of the design requirements – it is that such collection be *necessary*. Specifically, there must be a **functional requirement for collection**: the collection of personal data must be tied to the functionality of the technology.

The **distributed processing requirement** calls for data to be used as close as possible to the point at which it is collected. There are two rationales for this requirement. First, it prevents the creation of a single database that will be a target for hackers, law enforcement, or others who wish to exploit the data. Second, it reduces the ability of the service provider to market the data to third parties, or to re-use the data for purposes other than that for which it was originally collected. We will explore this requirement in more detail when we consider privacy-aware demand response systems.

Finally, any data that is collected in bulk for later testing purposes should be aggregated and/or anonymized.

5.4. Minimize identification of data with individuals

There is a distinction to be drawn between collecting information about equipment and collecting information about the equipment's users. The first sub-requirement under this heading is the **Non-Attribution Requirement**, which calls for anonymizing the data collected about equipment wherever possible.

Given that many networked services are billed to individuals, there must be some connection between the usage of the service and personally identifying information, such as a name and address. We propose a **Separate Storage Requirement** such that billing and “functional” records are stored in separate places. The separation can be enforced through policies of mutually exclusive permissions, such as the Chinese Wall Security Policy. The Chinese Wall Security policy establishes “conflict of interest classes,” then puts in place mandatory, legally enforceable controls by which any given individual is allowed to have access to at most one data set belonging to each class [11]. It would thus be both difficult and illegal for any person to have access to both billing and functional records.

5.5. Minimize and secure data retention

Data retention must be directly related to the functionality of the technology. It is not sufficient that data is useful in some other context, or may be useful

at some future date. A **Functional Requirement for Retention** has been proposed: the storage of the data must be directly connected to the functionality of the technology. If data must be stored, then it must be protected. The **Basic Security Requirement** requires that data be stored in such a manner that inadvertent disclosure is difficult or impossible. This is a longstanding, general concern in many industries, so we will not dwell on it here except to note that a requirement that consumers be notified when data is lost or stolen has been shown to reduce the frequency of such events.

Finally, but perhaps most importantly, there is a **Non-Reusability Requirement** that calls for data to be stored in such a manner that its use in an undisclosed manner be difficult or impossible.

6. Privacy-aware demand response

When we view demand response systems through the lens of privacy-aware design a privacy-preserving solution is apparent. The goal of demand response systems is to modify consumption behavior, whether through inducement or direct control, by exploiting fine-grained pricing information. The behavior of interest – consumption – is highly distributed. With the distributed processing requirement in mind, it becomes clear that it is not the power consumption data that needs to be collected, but it is instead the pricing data that needs to be distributed. Fine-grained consumption information need never leave the immediate neighborhood, thus alleviating most privacy concerns.

Such a privacy-aware approach is not, however, what some utilities have in mind. In the following excerpt from the 2006 FERC “Assessment of Demand Response and Advanced Metering,” AMI is *defined* as a system that provides for centralized collection. There seems to be no allowance for architectural options that are more sensitive to the privacy needs of the consumer.

For purposes of this report, Commission staff defined “advanced metering” as follows: “Advanced metering is a metering system that records customer consumption [and possibly other parameters] hourly or more frequently and that provides for daily or more frequent transmittal of measurements over a communication network to a central collection point.”

Assessment of Demand Response and Advanced Metering,” Federal Energy Regulatory Commission Staff Report, Docket No. AD06-2-000, Washington DC, August 2006, p. vi.

The above definition has since been quoted by utilities; see, for example [9], pg. 14. It has also been represented graphically in AMI literature distributed by FERC, as seen below [2]. Note that reference is made to the potential for third party data reception and management. This arguably increases the potential for unregulated use of the acquired data, including commodification and subsequent re-use by marketers and others.

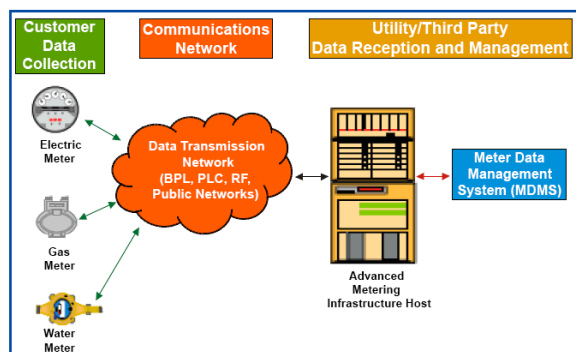


Figure 3. AMI building blocks [EPRI2007]

The need to secure AMI data has certainly been noted¹, but that is not the point. If it is not necessary to the mission of the system that the data be collected, then it should not be collected. The potential harm to consumers has been noted elsewhere (see [1] and [4]). It should also be noted that the utility and the long-term future of the demand response program are also at risk. Consumers may become alarmed at the privacy risk, motivating legislation calling for an expensive retooling of the system. Judicial action may also put the program at risk. Whether from public outcry or judicial action, systems that forsake privacy-awareness may find themselves shut down.

6.1. Distributed collection strategies

A privacy-aware demand response architecture must account for several different data flows. For each of them, a privacy analysis should be performed and a privacy-aware design adopted as necessary. First, in systems that seek to alter consumer behavior, pricing data must be presented to the consumer so that he/she has a basis upon which to make consumption decisions. This does not present a privacy concern, as the utility can simply broadcast the pricing to the

residential meter and/or to an application on the consumer's home computer.

Second, in direct control systems, the utility has to send signals to appliances to control their electricity consumption over the course of the day. Though this may create a significant security issue, it may not provide substantial information about consumer behavior and preferences within the home.

The third flow is more problematic. Consumer-specific consumption data must be provided to the utility for billing purposes. There is an issue here, as one cannot stream consumption data to the utility without creating the aforementioned privacy issue. One also cannot stream real-time cost data, as it would be trivial to convert this information back into consumption data. The solution lies in accumulating price-weighted consumption data at the residence and then sending the aggregate cost to the utility on a weekly or monthly basis. This implies a level of security at the meter that requires a Trusted Platform Module or the equivalent.

Finally, the utility needs consumption data, temporally precise, but aggregated at the level of the consumer, in order to predict demand and maintain a price model. Typically, aggregated real power consumption data at the substation level is sufficient to predict the need for new transmission and distribution lines and generation necessary to service the predicted demand. A neighborhood aggregator can be used to combine and anonymize data so that the desired temporal granularity is provided without generating information about individual behavior. Aggregator contractual obligations to the utility provides it with information sufficient to determine how much of the predicted demand can be mitigated through pricing mechanisms. In any case the utility's need for consumption data should not be at the level of the individual consumer. Anonymization can be performed by summing the power consumption data for a sufficient number of customers so that a single customer's data cannot be isolated.

The above approaches result in the architecture depicted in Figure 4.

¹ See, for example, "AMI System Security Requirements" at http://www.oe.energy.gov/DocumentsandMedia/14-AMI_System_Security_Requirements.pdf.

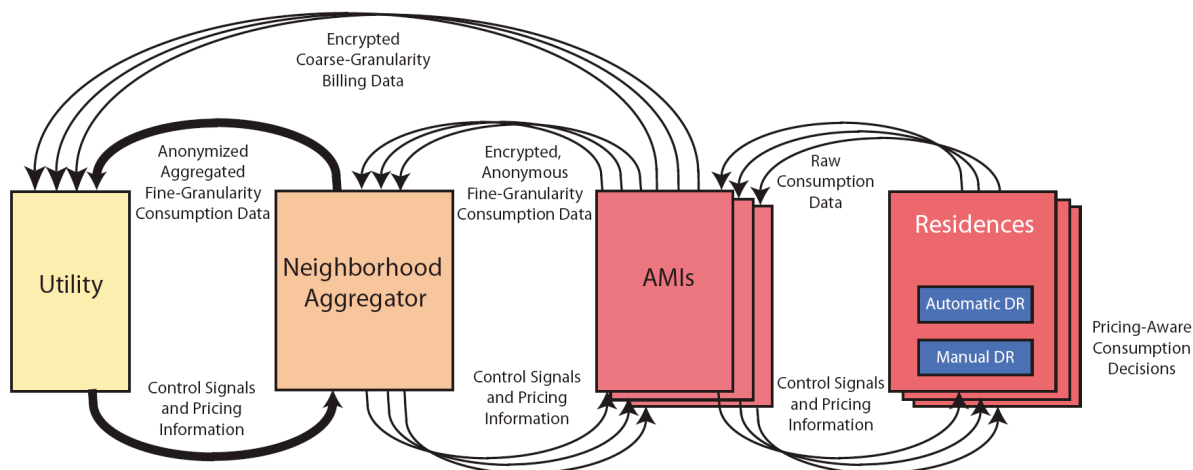


Figure 4. Privacy-aware demand response architecture

6.2. Tamper-proof meters

We propose the development of tamper-proof, privacy-sensitive metering infrastructure that is based on the use of the Trusted Platform Module. The Trusted Platform Module (TPM) was developed as a set of standards by the Trusted Computing Group (TCG) [8]. The TPM performs a wide variety of functions, including the secure generation, storage and use of cryptographic keys. These keys are used for several standardized purposes, including remote attestation, binding, signing, and sealing.

As seen in the following excerpt, binding is the encryption of a message using a public key. Public key cryptography uses a pair of keys, one public and one private, to facilitate information security without the need for secure key transfer. Note that the TPM stores the private key as a “nonmigratable” key – the private key cannot be transferred to another device. The TPM assures the security of such keys by maintaining secure locations that cannot be tampered with or accessed.

4.2.6.1 Binding

Binding is the traditional operation of encrypting a message using a public key. That is, the sender uses the public key of the intended recipient to encrypt the message. The message is only recoverable by decryption using the recipient’s private key. When the private key is managed by the TPM as a nonmigratable key only the TPM that created the key may use it. Hence, a message encrypted with the public key, “bound” to a particular instance of a TPM. It is possible to create migratable private keys that are transferable between multiple TPM devices. As

such, binding has no special significance beyond encryption.

[TCG Specification Architecture Overview, Revision 1.4, pg. 15]

We propose that the metering infrastructure use this capability to encrypt data transferred to the local aggregator, thus ensuring that any interception of this data will not lead to the leakage of personal information.

Signing is the generation of a digital signature. Digital signatures are often used to enforce non-repudiation; the focus is more on insuring that the party who signs the message is who they say they are, as opposed to preventing others from reading the message.

Signing also in the traditional sense, associates the integrity of a message with the key used to generate the signature. The TPM tags some managed keys as signing only keys, meaning these keys are only used to compute a hash of the signed data and encrypt the hash. Hence, they cannot be misconstrued as encryption keys.

[TCG Specification Architecture Overview, Revision 1.4, pg. 15]

We propose that signing be used to ensure that data received by the local aggregator and the utility was indeed transmitted by the metering infrastructure. This will prevent a variety of attacks (e.g. wormhole attacks and DNS attacks) based on the insertion of false data or false messaging into the network.

Remote attestation is a mechanism for verifying, often through an unforgeable hash algorithm, the state of the hardware and software of a computing device (see, for example, [13]). A platform can attest to its description of platform characteristics that affect the integrity (trustworthiness) of a platform. All forms of attestation require reliable evidence of the attesting entity.

Attestation can be understood along several dimensions, attestation by the TPM, attestation to the platform, attestation of the platform and authentication of the platform.

[TCG Specification Architecture Overview, Revision 1.4, pg. 5]

Given that the AMI is accessible by residents and potentially by third parties, remote attestation will be an important tool for insuring that the hardware and software of the meter has not been altered.

Finally we note that the TPM (or an equivalent device) can be programmed to keep accumulating electricity purchase information in a cryptographically secure vault. Fine granularity consumption data is thus unavailable to anyone wishing to transfer it to another platform. The privacy of the consumer will thus be assured.

6.3. Managing electric automobiles

Demand flattening needs to take place at the level of the individual residence. This can, as we have shown, be implemented at the residence without the need to collect data at a central facility. To improve the performance of such a system, both a “manual” and an “automatic” demand response element can be implemented.

The manual piece will consist of a demand response system in which pricing data is delivered to the home and presented to the homeowner. He or she would then be motivated to make choices that will reduce cost and presumably flatten demand. The automatic element would monitor these choices and control the charging of large home consumption elements such as an electric car.

As cars would typically be plugged in when the consumer returns home, there remains the potential for every automatic charging element making the decision to charge at the same time, resulting in an aggregate demand curve that is not flat. The community aggregator can randomize the charging using the power consumption data stream described above. The aggregator can cause the initiation of all local charging to be staggered across the normal sleep hours.

Car batteries are in the 4kW to 8 kW range, depending on the model of the car. The batteries will not be discharged to more than 80% capacity. That means that a demand of 6kW is likely. On a 120V 20A circuit, charging would take between 2 and 3 hours. On a 240V 30A circuit, charging would take about an hour. It follows that sufficient staggering can take place throughout the sleeping hours to balance the load placed on these grid by charging vehicles.

7. Conclusions

A framework for privacy-aware design practices was developed as a roadmap for embedding privacy-awareness into information networks. The framework was then applied to the problems of the collection of power consumption data in demand response systems. In closing we wish to emphasize that privacy-aware design is still in its infancy. There are many interesting technical problems to be solved as the design toolbox for privacy-aware information networks is developed.

8. References

- [1] “A National Assessment of Demand Response Potential,” Staff Report, Federal Energy Regulatory Commission, June 2009. Available at <http://www.ferc.gov/legal/staff-reports/06-09-demand-response.pdf>
- [2] “Advanced Metering Infrastructure,” Engineering Power Research Institute, Available at <http://www.ferc.gov/eventcalendar/Files/20070423091846-EPRI%20-%20Advanced%20Metering.pdf>
- [3] Mikhail Lisovich, Deirdre Mulligan, and Stephen B. Wicker, “Inferring Personal Information from Demand-Response Systems,” *IEEE Security and Privacy Magazine*; January/February 2010.
- [4] Stephen B Wicker and Dawn E. Schrader, “Privacy-Aware Design Principles for Information Networks,” *Proceedings of the IEEE*, submitted 2010, to appear.
- [5] *Records, Computers, and the Rights of Citizen*, HEW, 1973.
- [6] “2009 Assessment of Demand Response and Advanced Metering,” staff report, Sept. 2009; <http://www.ferc.gov/legal/staff-reports/sep-09-demand-response.pdf>.
- [7] Jeff Sovern, “Opting In, Opting Out, Or No Options At All: The Fight For Control Of Personal Information,” *Washington Law Review*, 74, 1033, 1999.
- [8] TPM Main, Part 1 Design Principles, Specification Version 1.2, Level 2 Revision 103, Trust Computing Group, 9 July 2007.

- [9] "Empowering Consumers Through a Modern Electric Grid" Illinois Smart Grid Initiative Final Report 2009, available at <http://www.cnt.org/repository/ISGI.FinalReport.pdf>.
- [10] P.A. Subrahmanyam, David Wagner, Deirdre Mulligan, Erin Jones, Umesh Shankar & Jack Lerner, Network Security Architecture For Demand Response/Sensor Networks 14 (2005, Rev. 2006), <http://www.Ucop.Edu/Ciee/Dret/d/>(Report For The Network Security Architecture For Demand Response/Sensor Networks Project, Ciee Award No. Dr-04-03a, B, Wa No. Dr-005, Under Cec/Cii Prime Contract No. 300-01-043, Conducted By Cyberknowledge and the University Of California At Berkeley)
- [11] David F.C. Brewer and Michael J. Nash, "The Chinese Wall Security Policy," *Proceedings of the 1989 IEEE Symposium on Security and Privacy*, 1989, pp.206 – 214.
- [12] M.H. Albadi and E.F. El-Saadany, "A Summary Of Demand Response In Electricity Markets," *Electric Power Systems Research* Volume 78, 2008, pp. 1989–1996.
- [13] LeMay, M. and Gunter, C. A. 2009. Cumulative attestation kernels for embedded systems. In *Proceedings of the 14th European Conference on Research in Computer Security* (Saint-Malo, France, September 21 - 23, 2009). M. Backes and P. Ning, Eds. Lecture Notes In Computer Science. Springer-Verlag, Berlin, Heidelberg, 655-670.
- [14] International Energy Agency, *The Power to Choose—Demand Response in Liberalized Electricity Markets*, OECD, Paris, 2003.
- [15] George W. Hart, "Residential Energy Monitoring and Computerized Surveillance via Utility Power Flows," *IEEE Technology and Society Magazine*, June 1989.