Formalizing Anonymity-Delay Tradeoffs in Smart Grid Networks

Benjamin Lipton
Department of Computing Security
Rochester Institute of Technology
Rochester, NY 14623
email: benlipton@mail.rit.edu

Sumita Mishra, PhD
Department of Computing Security
Rochester Institute of Technology
Rochester, NY 14623
email: sumita.mishra@rit.edu

***Abstract*—Privacy concerns about smart metering have proven to be a significant roadblock to adoption of the smart grid. In the smart grid, electricity usage data are collected in near real time, exposing users to attacks that reveal their habits, location, and preferences. Other research in this area has demonstrated the possibility of using pseudonyms to decouple low-frequency billing data from high-frequency grid maintenance data. Such schemes must provide a way to authenticate the pseudonym as belonging to a real customer, to defend against the submission of false data. However, this authentication process exposes the system to timing attacks that can correlate the authentication of a particular customer with the appearance of a new pseudonym on the network. These timing attacks have been addressed informally in the literature with techniques such as adding a random delay to transmission. The anonymity thus provided is not quantified or guaranteed, and can be subverted with active attacks via the service provider. This paper introduces a method for defeating timing attacks based on the theory of mix networks. The use of a mix formalizes a method of randomizing the order in which authentication attempts are serviced. The mix design is described in detail and it is shown that this method is more resistant to attack and allows for an explicit tradeoff between anonymity and delay.**

Keywords-anonymity; data privacy; mix networks; smartgrids; smart grid privacy; smart metering; meter reading

I. Introduction

The established electrical grids of most countries are relatively static in design. Power is generated at a few centralized points and distributed outward to the consumer via a large network of cables. Companies that generate and distribute electricity receive feedback from the customer once a month via a meter reading, or in emergencies via customer support calls. Such an infrastructure can be inefficient and slow to respond to problems because of a lack of data.

Increasingly there has been a push to provide a more responsive infrastructure for power distribution, the Smart Grid. In the new environment, electricity usage data are monitored in near real time by Smart Meters installed in homes, and transmitted to electricity suppliers so they can make faster decisions about generation, repairs, and maintenance. With this increased data collection comes increased privacy concerns. Research like [1] and [2] has shown that sensitive data like the number of people in a home or the channel being watched on TV can be revealed by electricity usage data collected at a sufficiently high frequency. In the US, NIST has explicitly recognized privacy as an area of concern with the increase of data collection inherent in smart grid deployment, and has laid out recommendations for privacy protection in [3]. However, these recommendations are nontechnical and would not protect against an attacker that is not bound by organizational policy.

This paper acts as a complement to other work that has applied principles of privacy-enabled design to the problem of smart metering. We address limitations in the solution of [4], analyzing the effectiveness of passive and active timing attacks against the protocol presented there. Further, we present a framework for designing systems with better defensive properties using the formalism of cryptographic mixes, whose properties have already been studied in the literature.

The paper is organized as follows: Section II presents a selection of prior work in the areas of smart grid privacy, anonymity, and mix networks. Section III outlines our chosen environment and the behavior of the attacker against which we will try to defend. Section IV presents the criteria we will use to evaluate potential solutions to the problem, and Section V describes and evaluates the solutions themselves. Finally, Section VI presents our conclusions from the analysis and describes areas for future work.

II. Background

A variety of techniques for protecting consumer privacy in the smart grid have been proposed. Some researchers have chosen to preserve privacy by altering the data sent back to the utility. For example, [5] changes real-world electricity usage by using a battery to smooth out power draw, erasing the peaks that can be used to distinguish particular types of usage. Several authors [6], [7] have proposed group protocols to create data streams that are individually inaccurate but meaningful when aggregated. These techniques ensure that no individual user's data can be tracked, but also reduce the usefulness of the data provided to the electricity supplier because only aggregate results are available rather than per-household metrics.

In order to allow fine-grained per-household data to be provided without compromising the privacy of the residents, work such as [4] and [8] has proposed the use of a pseudonym for transmitting high-frequency metering data. This allows the data to be sent to the electricity supplier unmodified,

978-1-5090-5111-3/16/$31.00 ©2016 IEEE

but unassociated with any particular customer identity. The architecture of [4] will be discussed in more detail in Section III-A. Although the use of pseudonyms does provide an attractive option for privacy-enabled smart metering, it is not without issues even when implemented well; [9] shows that attacks are possible that use *only the consumption data* to reduce the privacy granted by a pseudonym.

For any system that attempts to preserve the anonymity of its users, an important question is how to quantify the anonymity provided. Several different metrics have been proposed for this purpose. The most basic is the "anonymity set," the number of different parties who could be the sender of a particular message. This metric is straightforward to understand, but limited in situations where some senders are more likely than others. If one sender is significantly more likely than all others, the system provides less anonymity, but the size of the anonymity set is the same. Several metrics have been proposed to take these relative probabilities into account. Among the most widely-used are the "effective size" of an anonymity probability distribution presented in [10] and the "degree of anonymity" presented in [11]. These very similar metrics both interpret sender anonymity as entropy in the probability distribution of senders for a message. As an alternative, the authors of [12] contrast these "global" metrics with a proposed "local" anonymity metric that focuses on the anonymity that each party achieves within the system rather than the quality of anonymity in the system as a whole.

This paper presents a framework for anonymity-preserving authentication protocols for the smart grid based on the idea of a cryptographic mix, as introduced by Chaum [13]. A mix is a system that accepts messages as input, transforms them so as to be unrecognizable, then releases them in a random order. This prevents an attacker with knowledge of the inputs and outputs from associating a particular input message with a particular output message. The anonymity and delay properties of various types of mixes have been studied [14], giving results that can be used to formalize the anonymity provided by the protocol. Later work such as [15] and [16] showed that the mixes previously analyzed are just examples of a larger, more general design space. By altering the probability distributions that govern the number of messages emitted each time it fires, a mix can be modified to have desirable properties of anonyimity and delay.

III. System Architecture and Threat Model

A. Architecture

The architecture considered will be similar to that introduced in [4]. Smart meters transmit two types of data to the utility: usage data for billing, which are sent at a low frequency (say every month), and data used by the grid operator for maintenance, which are sent at a high frequency (such as every few minutes). Each smart meter has two built-in identifiers, called the HFID (High-Frequency Identifier) and LFID (Low Frequency Identifier), that it uses to identify itself when sending these two types of data.

Since the low-frequency data are used for billing, the LFID must be associated with personally-identifiable customer information like a name and address. The HFID, on the other hand, must be validated as belonging to a real customer, but must not be associated with that customer's identity. This is achieved using an escrow service that operates autonomously from the utility. The pairs of LFID, HFID are provided to the escrow service by meter manufacturers, so that when a meter is brought online it can make a request using its LFID and receive a cryptographic certification from the escrow service that its HFID belongs to a valid meter. This certification is then used when trasmitting data to demonstrate its authenticity. This architecture is shown in Figure 1. At a high level, the protocol is as follows:

1) Customer transmits personal information and LFID to utility.
2) Utility grants authorization to LFID.
3) Smart meter transmits authorization, LFID, HFID to escrow service.
4) Escrow service sends authorization for HFID to utility, if it is correctly paired with LFID.
5) Smart meter uses HFID to transmit anonymized high-frequency data to utility.

The architecture described in [4] has been simplified by merging the roles of the utility, data concentrator, and distribution substation, as their separation grants no additional privacy.

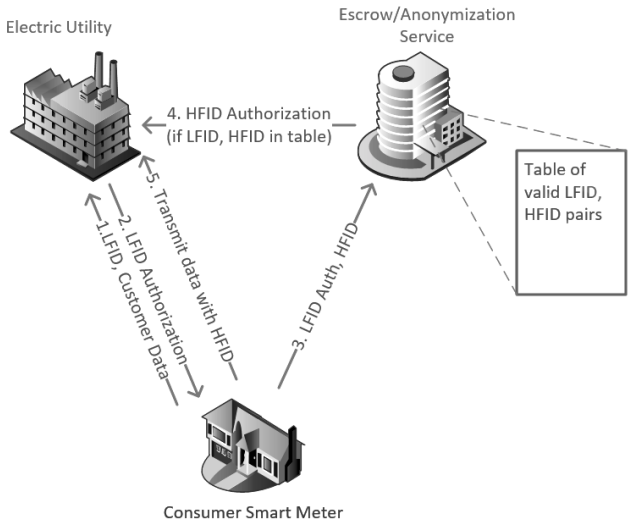

Figure 1. Simplified diagram of authorization flow

B. Threat Model

It is important to note that steps 2 and 4 are subject to a timing attack by the utility. If all steps proceed as quickly as possible, the utility will receive the HFID authorization shortly after it provides the LFID authorization to the HFID. Since it already knows the identity of the customer with a particular LFID, it would then be able to associate the HFID with that customer as well. It is this type of association attack

that we aim to prevent by randomizing the order that HFID authorizations reach the utility. We will provide a formal analysis of the technique proposed by [4], making smart meters add a random delay to step 3. We will also propose a design that moves the randomization into the escrow service in step 4. When the escrow service plays an active part in maintaining customers' anonymity by randomizing the output order, we will refer to it as the *anonymizer*.

An attacker operating within the utility's systems has several strategies available to it to help in its effort to associate HFID with LFID:

- Passive timing analysis
- Actively delaying LFID authorizations
- Actively creating false LFID authorizations, which will be processed like real ones by the escrow service but can be recognized by the attacker when the response is released

Although the attacker can use active techniques to reduce a user's anonymity, these attacks have some cost as well. Delaying of authorization is costly because an unauthorized meter can not transmit data, and the data are valuable to the utility for billing and grid maintenance purposes. Inserting of false authorizations is even more expensive because it can only be effective if the submitted LFID, HFID pairs are recognized by the escrow service. This could even mean the attacker would need to purchase actual smart meters and connect them to the network in order to mount this attack.

In addition to expense, these active attacks could only be carried out by a sophisticated adversary with the ability to control the systems used for meter authorization, and modify their operation for an extended period of time. This would be difficult even for an internal actor, and likely beyond the reach of an attacker external to the utility. A passive attack is much more likely because it only requires access to already-collected data. Such data could be stolen by an employee of the utility, or an external attacker, without requiring such advanced capabilities.

C. Temporary IDs

One disadvantage to the architecture described above is that the identifiers come pre-programmed into the smart meter, rendering the meter useless in case an attacker identifies the owner of a particular HFID. A preferable architecture would be one in which the HFID can be periodically re-set.

The techniques described in the following sections are also effective in a system where the HFID is not assigned in advance, but instead generated by the escrow service. This way, a compromised meter can be re-used simply by reauthorizing the LFID and generating a new HFID, and even if there is no compromise HFIDs can be rotated periodically to increase the size of the anonymity set.

The details of authentication in such a system will be left for future work. For now it is sufficient to note that, although the role of the escrow service has changed, the need for randomization of the output order is still present, otherwise the attacker could easily correlate between a just-authorized LFID and a new HFID that begins sending right after it.

IV. Evaluation Criteria

A. Anonymity

Anonymity is the goal that represents the needs of the customer. For their privacy, customers would prefer that their data be indistiguishable from the data of other customers, and the more confusion about the source of a data stream the better.

As discussed in section II, a number of metrics have been proposed to quantify this confusion. In this paper we will use the source hiding property of [12] to evaluate the anonymity provided by the analyzed systems. Although the information-theoretic metrics ([10], [11]) are widely used, the local aspect of the source hiding metric better reflects the individual consumer's desire for anonymity. In certain cases, a system could score well on the information-theoretic metrics, but still expose one sender as having a much larger probability than all others, allowing the attacker to guess the association with a high level of accuracy. A system that is source-hiding avoids this type of "guessing attack" by placing a limit on the probability of the senders, and therefore on the accuracy of the guess.

More formally, let Ψ be the set of information about the system visible to the attacker, such as the number of messages entering and leaving the mix at each round since the attacker began recording. The attacker attempts to analyze this information to make an accurate guess about the sender of a message. Model his analysis technique as a function $P(\Psi, s, \beta)$ that outputs the probability that identity s is the sender of message β, based on the available data Ψ. The attacker applies this function to all the possible senders s for a given message β, to determine which sender is most likely to be the real sender of the message. A system is source-hiding with parameter Θ if

$$\forall\ \Psi, \beta, s : P(\Psi, s, \beta) < \Theta$$

That is, no matter what has happened in the system (including active manipulations by the attacker), the attacker can never assign a probability greater than Θ to any possible sender of a given message. So, no matter which sender he selects as his guess, the probability of the guess being correct will be less than Θ.

B. Delay

Delay is the goal that represents the needs of the utility. In order to operate the grid and bill customers, the utility requires data. If the system holds onto smart meter registrations for too long, preventing the meters from entering service, this prevents the utility from getting the data it needs. The results of this (for example, decreased energy efficiency or slower response to incidents) are potentially costly but difficult to quantify.

V. Results: Performance of Sample Anonymizer Implementations

A. Random delay

In [4], meters introduce a random delay of a few days or weeks between receiving LFID authorization and submitting it to the escrow service for HFID authorization. Let us define $D(t)$ as the cumulative distribution function of the probability distribution of this delay. That is, $D(t) = P(delay \leq t)$. If the support of this distribution has an upper bound ($D(t \geq T_{max}) = 1$) we will call it bounded; otherwise it is unbounded.

1) Passive attack resistance: With a passive attacker, the anonymity provided by this scheme depends on the flow of new authentiation requests into the system. Thus the system is not source-hiding, because an unlucky sequence of request times combined with an unlucky selection of random delays can result in the identity of a meter being completely revealed. That is, for any sequence of delay times selected from the distribution, there exists a sequence of arrival times such that one of the requests arrives and is released entirely during a period when all previous delays have expired and no newer requests have arrived.

2) Active attack resistance: When the attacker can delay messages, the situation gets even worse, because he can now actively push the system towards the situation described in the previous section. However, his likelihood of success depends on the distribution.

If the delay distribution is bounded, this attack is highly effective. The attacker delays all requests for T_{max} seconds, then submits a single request, and then delays all requests for another period of T_{max} seconds. By the time the single request occurs, all of the previously-authorized meters will already have finished their delays and begun transmitting. Then, at some time during the next T_{max} seconds, the requesting meter will begin transmitting as well. Since it is the only meter that could have begun during that period, its identity is not hidden from the attacker. Using the terminology in [14], this is an *exact*, *certain* attack.

An unbounded delay distribution performs better; the attack is now *uncertain* because the attacker may have to delay messages for an arbitrarily long time to wait for authorizations to go through. On the other hand, the unbounded distribution also means the system has no maximum delay. A request could be delayed for an arbitrarily long time, even when no attack is taking place.

Note that the active attacker in this case only needs the ability to delay messages. He does not need to forge any messages, nor would he gain any benefit from doing so, as all delays are determined independently by the meters.

B. Mix-based implementation

In this section we discuss a design for implementing an anonymizing escrow service using a mix. In the context of the anonymizer, the incoming messages are authentication requests from the smart meter, authorized by the utility. The outgoing messages are HFID authorizations from the escrow service. Upon release of the authorization, the associated smart meter will begin transmitting data to the utility under its HFID.

We focus on a mix whose firing algorithm has two parameters: T, the period and n_{min}, the threshold. As the anonymizer authorizes or generates HFIDs for meters, they are entered into a pool. Every T seconds, the mix checks N_p, the number of entries in the pool. If $N_p \geq n_{min}$, the mix fires as follows. An integer N_s is chosen from $[0, N_p]$ with uniform distribution. This determines the number of HFIDs that will be released this round. Those HFIDs are chosen uniformly from all those in the pool.

1) Passive attack resistance: No HFID will be released unless at least n_{min} possibilities are in the pool. This strictly limits the effectiveness of a passive attacker. Even if the attacker is completely sure that a given customer is in the pool (for instance, that customer's LFID was authorized since the last release cycle), any HFID that is released from the pool must be one of at least n_{min} entries. The maximum probability that the attacker could assign to any customer as the owner of that HFID is therefore $\frac{1}{n_{min}}$. Thus the system is source-hiding with parameter $\frac{1}{n_{min}}$.

2) Active attack resistance: An active attacker is more effective at damaging the anonymity of the system, but the mix adds significant additional cost to this process. We will consider an active attack consisting of the following steps (other attacks are possible and can be analyzed similarly):

1) Attempt to remove all real messages from the pool: delay real messages and insert false messages until probability that all real messages have been cleared exceeds P_{target}.
2) Submit 1 real message into the mix.
3) Delay real messages and insert false messages until a message is released that doesn't match any false message.

If step 1 succeeds, the released message definitely corresponds to the one submitted in step 2. Thus the probability of revealing a message with this technique is at least P_{target}. With greater expenditure of time and false messages, the attacker can make P_{target} arbitrarily close to 1. So, the system is not source-hiding against such an attacker. It should be noted, however, that this is an extraordinarily powerful attacker, who is able to delay and modify communications out of the utility for what could be an arbitrary length of time.

3) Delay: There is no maximum to the delay that may be applied to an element in the pool. If the pool contains fewer than n_{min} entries, or the random selection of N_s comes up 0, there may be a round with no entries released at all, and even if some are released a particular one could stay in the pool for an arbitrary length of time. However, it is unlikely to do so, and we can compute the probability of it being released each round as follows.

Each round, the value of N_s is chosen uniformly from the number of items in the pool, so each value in $[0, N_p(t)]$ is equally likely, with probability

$$P_{N_s}(t) = \frac{1}{N_p(t) + 1}$$

The probability that the single real entry in the pool is selected in each round is

$$\begin{aligned} P_{sel}(t) &= \sum_{N_s=0}^{N_p(t)} \left(P_{N_s}(t) \frac{N_s}{N_p} \right) \\ &= P_{N_s}(t) \frac{1}{N_p(t)} \sum_{N_s=0}^{N_p(t)} N_s \\ &= \frac{1}{(N_p(t))(N_p(t)+1)} \frac{(N_p(t))(N_p(t)+1)}{2} \\ &= \frac{1}{2} \end{aligned}$$

This probability remains constant regardless of the input to the mix (except of course if the pool size drops below n_{min}, in which case $P_{sel} = 0$). Since the probability that a particular authorization will be released each round is $\frac{1}{2}$, the expected number of rounds until that authorization is released is $\frac{1}{(\frac{1}{2})} = 2$. So, although the delay of a message is theoretically unbounded, most messages will be delayed only a few rounds.

VI. Conclusions

The privacy concerns of smart metering are substantial, and existing solutions that use pseudonyms do not do enough to quantify and mitigate the impact of timing-based attacks. This paper analyzes one suggested solution, that of adding random delays during authentication, and proposes a means for implementing new solutions using a mix. The following summarizes the results of comparing random delays to the specific mix-based solution presented here, though other mixes would have different properties:

- Passive attack resistance: Random delays are not source-hiding, while the mix is source-hiding with parameter $\frac{1}{n_{min}}$.
- Active attack resistance: Neither system is source-hiding, but the attack against the mix requires an attacker that can insert messages, while delaying messages is sufficient to attack the random delays.
- Delay (no attack): For random delays the expected and maximum delay depends on the distribution. If the distribution is unbounded there is no maximum delay. Similarly, with the mix there is no maximum delay, but the expected delay is 2 rounds.

Thus, the mix implementation improves the resistance to passive attacks, though an active attacker with unlimited resources can still overcome the defenses provided by this system. There is reason to believe that an attacker with limited time to mount an active attack on each meter would face limits on the effectiveness of his attack, since uncertainty about the number of messages in the mix would lead to uncertainty about whether the HFID returned by the attack matches the sender being attacked. However, a full analysis of this new threat model will be left for future work.

Beyond the performance of this specific mix, this paper provides a design that is adaptable to different situations by using a mix with appropriate properties for the desired scenario. This would be achieved by altering the distribution from which N_s is sampled. A mix could be made to favor greater anonymity by skewing in the direction of small N_s values, or lower delay by tending to make N_s larger. Exploring this design space to find mixes with desirable properties is one avenue of future work. Other future work will include detailing a protocol that allows HFIDs to be regenerated periodically and investigating the impact of reauthorization on anonymity.

References

[1] U. Greveler, P. Glsektterz, B. Justusy, and D. Loehr, "Multimedia content identification through smart meter power usage profiles," in *Proceedings of the International Conference on Information and Knowledge Engineering (IKE)*. The Steering Committee of The World Congress in Computer Science, Computer Engineering and Applied Computing (WorldComp), 2012, p. 1.

[2] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin, "Private Memoirs of a Smart Meter," in *Proceedings of the 2Nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, ser. BuildSys '10. New York, NY, USA: ACM, 2010, pp. 61–66.

[3] Smart Grid Interoperability Panel Cyber Security Working Group, "Introduction to NISTIR 7628 guidelines for smart grid cyber security," *NIST Special Publication*, no. 800-53 Revision 3, 2010.

[4] C. Efthymiou and G. Kalogridis, "Smart grid privacy via anonymization of smart metering data," in *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, Oct 2010, pp. 238–243.

[5] G. Kalogridis, C. Efthymiou, S. Z. Denic, T. A. Lewis, and R. Cepeda, "Privacy for Smart Meters: Towards Undetectable Appliance Load Signatures," in *2010 First IEEE International Conference on Smart Grid Communications (SmartGridComm)*, Oct. 2010, pp. 232–237.

[6] K. Kursawe, G. Danezis, and M. Kohlweiss, "Privacy-friendly Aggregation for the Smart-grid," in *Proceedings of the 11th International Conference on Privacy Enhancing Technologies*, ser. PETS'11. Berlin, Heidelberg: Springer-Verlag, 2011, pp. 175–191.

[7] S. Li, K. Choi, and K. Chae, "An enhanced measurement transmission scheme for privacy protection in smart grid," in *2013 International Conference on Information Networking (ICOIN)*, Jan. 2013, pp. 18–23.

[8] C. Rottondi, G. Mauri, and G. Verticale, "A data pseudonymization protocol for smart grids," in *Online Conference on Green Communications (GreenCom), 2012 IEEE*. IEEE, 2012, pp. 68–73.

[9] M. Jawurek, M. Johns, and K. Rieck, "Smart Metering De-pseudonymization," in *Proceedings of the 27th Annual Computer Security Applications Conference*, ser. ACSAC '11. New York, NY, USA: ACM, 2011, pp. 227–236.

[10] A. Serjantov and G. Danezis, "Towards an Information Theoretic Metric for Anonymity," in *Privacy Enhancing Technologies*, ser. Lecture Notes in Computer Science, R. Dingledine and P. Syverson, Eds. Springer Berlin Heidelberg, Apr. 2002, no. 2482, pp. 41–53.

[11] C. Díaz, S. Seys, J. Claessens, and B. Preneel, "Towards Measuring Anonymity," in *Privacy Enhancing Technologies*, ser. Lecture Notes in Computer Science, R. Dingledine and P. Syverson, Eds. Springer Berlin Heidelberg, Apr. 2002, no. 2482, pp. 54–68.

[12] G. Tóth, Z. Hornák, and F. Vajda, "Measuring Anonymity Revisited," in *Proceedings of the Ninth Nordic Workshop on Secure IT Systems*, Espoo, Finland, 2004, pp. 85–90.

[13] D. L. Chaum, "Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms," *Commun. ACM*, vol. 24, no. 2, pp. 84–90, Feb. 1981.

[14] A. Serjantov, R. Dingledine, and P. Syverson, "From a trickle to a flood: Active attacks on several mix types," in *Information Hiding*, ser. Lecture Notes in Computer Science, F. A. P. Petitcolas, Ed. Springer Berlin Heidelberg, 2003, vol. 2578, pp. 36–52.

[15] C. Díaz and A. Serjantov, "Generalising Mixes," in *Privacy Enhancing Technologies*, ser. Lecture Notes in Computer Science, R. Dingledine, Ed. Springer Berlin Heidelberg, Mar. 2003, no. 2760, pp. 18–31.

[16] A. Serjantov, "A Fresh Look at the Generalised Mix Framework," in *Privacy Enhancing Technologies*, ser. Lecture Notes in Computer Science, N. Borisov and P. Golle, Eds. Springer Berlin Heidelberg, Jun. 2007, no. 4776, pp. 17–29.