Smart Metering De-Pseudonymization

Marek Jawurek
SAP Research
Vincenz-Priessnitz Str.1
Karlsruhe, Germany
marek.jawurek@sap.com

Martin Johns
SAP Research
Vincenz-Priessnitz Str.1
Karlsruhe, Germany
martin.johns@sap.com

Konrad Rieck
Technische Universität Berlin
Franklinstrasse 28/29
Berlin, Germany
konrad.rieck@tu-berlin.de

ABSTRACT

Consumption traces collected by Smart Meters are highly privacy sensitive data. For this reason, current best practice is to store and process such data in pseudonymized form, separating identity information from the consumption traces. However, even the consumption traces alone may provide many valuable clues to an attacker, if combined with limited external indicators. Based on this observation, we identify two attack vectors using anomaly detection and behavior pattern matching that allow effective de-pseudonymization. Using a practical evaluation with real-life consumption traces of 53 households, we verify the feasibility of our techniques and show that the attacks are robust against common countermeasures, such as resolution reduction or frequent re-pseudonymization.

1. INTRODUCTION

The deployment of Smart Metering—the digital recording and processing of electricity consumption—is ever increasing. A Smart Meter is an electrical meter that records a fine-grained consumption trace of a household and sends it to the respective electricity supplier. These consumption traces, in contrast to traditional single annual consumption values, allow the realization of time-of-use tariffs and demand response schemes.

This flexibility, however, comes at a price. Every activity that takes place in the household and makes use of electrical appliances is reflected in the consumption trace. In consequence, Smart Metering has repeatedly been called a privacy invasion into households [7, 8] and a large body of previous work [5, 6, 11, 12, 14, 15, 20] has been concerned with inferring private information from energy consumption traces.

Based on the identified privacy implications, there is consensus that consumption data of Smart Metering needs to be adequately protected. Such protection entails the protection during storage by the supplier and during the use of the data by the supplier and 3rd party contractors. Pseudonymization of consumption traces is considered an effective defense against privacy attacks, as it allows for unlinking the identity of the household and its consumption trace. The consumer's identity can be stored independently from consumption traces, only linked by the pseudonym. In such a scenario, the privacy-invading methods developed in previous work can only be applied by the owner of both, the identity database and the consumption traces.

An attacker faces two problems, if he has only access to pseudonymized traces: First, deduction from pseudonymous consumption traces is error-prone as no identity information can be used as contextual data. Second and more important, all information inferred from consumption traces can not be attributed to a specific household due to the unlinkability introduced by pseudonymization. This makes consumption traces and its contained information unattractive for targeted abuse and apparently the consumers' privacy is protected.

In this paper, we develop two attack vectors targeting the privacy of pseudonymized consumption traces. The first attack allows to create a link between a household's identity and its consumption trace, and therefore enables an attacker to undo pseudonymization. If successful, this attack allows all existing deduction attacks to be applied again. The second method enables an attacker to track the origin of a consumption trace across re-pseudonymization or across different databases. For conducting these attacks in practice, we provide a data analysis framework that allows an attacker to apply either method to consumption databases.

The paper's main contributions are as follows:

1. An abstract definition of attack vectors on the unlinkability of pseudonymous Smart Metering consumption traces.

2. A machine learning framework for the analysis of consumption traces and subsequent execution of aforementioned attack vectors.

3. Experimental findings about the anomaly detection in consumption traces and the tracking of consumption traces across pseudonyms.

4. An evaluation of different mitigation techniques with respect to their effectiveness against those attacks.

The rest of this paper is structured as follows: In Section 2 we provide an overview of the terminology used in this paper. Section 3 describes the two attack vectors that we identified for de-pseudonymization. In Section 4, we present

ACSAC '11 Dec. 5-9, 2011, Orlando, Florida USA

a data analysis framework for conducting these attacks in practice. Section 5 shows our experimental results regarding the applicability of our attacks. Approaches for limiting linkability like reduced resolution, frequent re-pseudonymization and cryptography solutions are discussed in Section 6. Finally, in Section 7 we discuss related work before we conclude in Section 8.

2. TERMINOLOGY AND ASSUMPTIONS

This paper's understanding of identity, pseudonymity and other privacy related terms is based on terminology definitions provided by Pfitzmann and Hansen [19].

In the following we will use the term consumption trace for a recording of electric consumption in discrete time slots of equal length. The resolution of a consumption trace is the number of time slots per day. The term consumer stands for the household that makes use of the utility that is measured by a Smart Meter. The subject supplier stands exemplary for any subject (actual supplier or grid operator) that participates in Smart Metering and has legitimate interest in consumption traces, e.g. for billing electricity consumption. The Smart Meter records the consumption constantly and communicates it to the supplier. The Smart Meter, respectively the consumption trace, are attributed with a pseudonym for the consumer. This pseudonym is pre-arranged by the supplier and non-public.

The supplier stores consumption traces by pseudonym in a consumption trace database. The supplier accesses this database for billing or analysis in general. Contractors of the supplier also have access to this database but not to the identity database. Contractors provide analysis services to suppliers based on pseudonymized consumption traces. For the purpose of invoice creation the supplier owns another database, the identity database, that connects the consumer's identity with the pseudonym.

Optionally, a supplier may re-pseudonymize a consumption trace repeatedly, creating a 1 to n relationship between identities and pseudonyms. Consumers can switch suppliers which leads to the following situation: They have an old consumption trace, identity information and linking pseudonym in the old supplier's databases and also a current consumption trace and new pseudonym together with their identity information in their new supplier's databases. Contractors and suppliers may behave semi-honest, i.e, they stick to the protocols and respective laws but try to learn as much about consumers as possible. Our attacker model includes contractors, suppliers but also external malicious agents that might illegally obtain consumption traces.

3. ATTACK DESCRIPTION

The final goal of our attacker model is to create a link between the identity of consumers and their energy consumption. To this end, we present two different attacks that can be applied to achieve this goal. We name the attacks: *linking by behavior anomaly* and *linking by behavior pattern*.

Once a link between a consumption trace and a specific household has been established all information contained in the trace can be attributed to this household. If, on the other hand, a consumption trace cannot be linked to one household with a significant higher probability than to another household, this means that the data and its contained information cannot be attributed to a single household. In this case, the contained information would not have a privacy impact on its origin. Thus, linkability is a sufficient condition for privacy loss in Smart Metering.

3.1 Linking by Behaviour Anomaly

Linking by behavior anomaly (LA) can be used by the attacker to link either an identity to a consumption trace or two consumption traces with each other (see Figures 1 and 2). This is accomplished by *identifying and correlating anomalies* in both data sources that occur at the same time.

We characterize an anomaly as a series of unusual events, where an event is some consumer behavior that is reflected in the energy consumption of the respective household. The rarity of an anomaly is based on different factors: Length of the series, the resolution of the time stamp (day, hours or minutes) and rarity of the singular events among the population of consumption traces under consideration. Length and resolution make up for singular events that happen very often among the population: Leaving home or coming home at slightly different times during the course of one week. On the other hand, there are events that neither require a series nor a high resolution: moving in/out, death/birth or holidays. A series of events for a low resolution can be observed if household inhabitants leave every weekend or stay at home always at specific work days. Here the rarity originates from the length of the series.

With respect to linking a consumption trace to an identity, LA can be used in both ways. Either identifying the household for a consumption trace or vice versa. It really depends on the final purpose of the identification which approach is taken, whether specific households should be targeted (like both examples in Section 3.3) or specific consumption profiles are of interest. The main requirement for this attack is to have two data sources that overlap for the time interval where an anomaly has been identified.

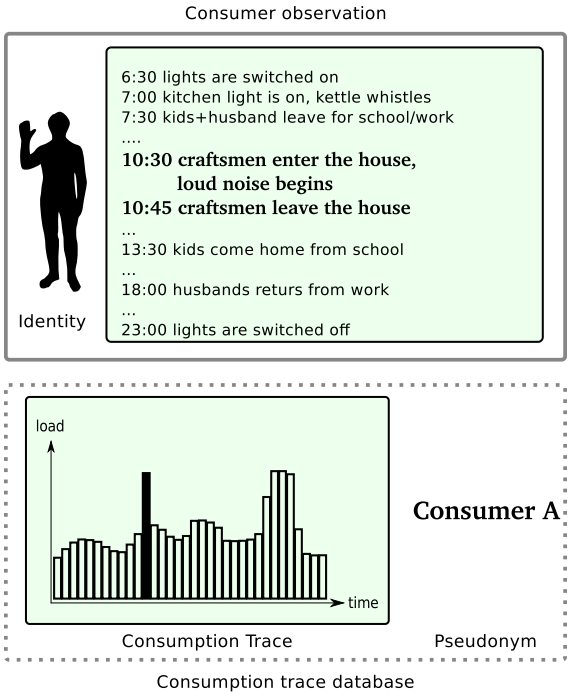

Figure 1: Behavior and consumption anomaly

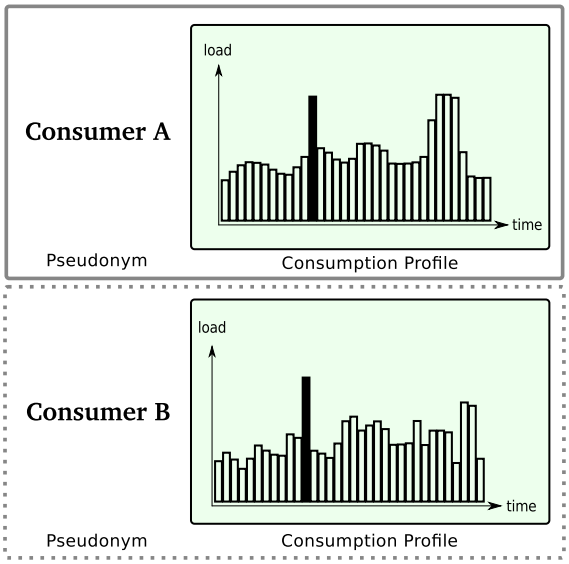

Figure 2: Anomalies in two consumption traces

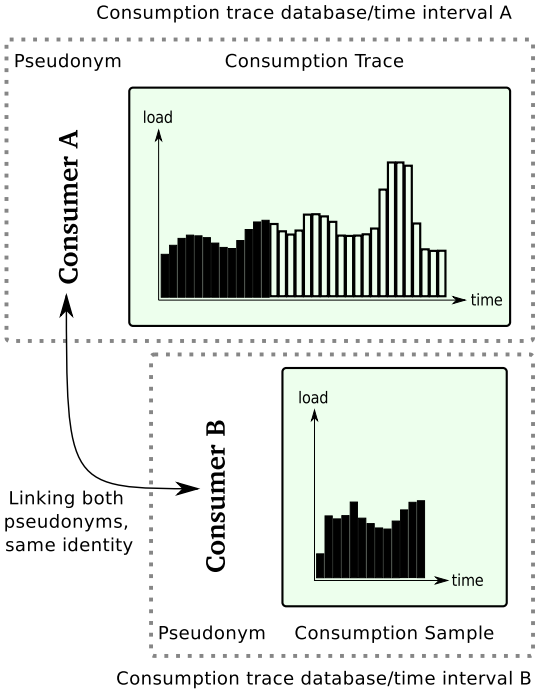

Figure 3: Identifying consumption traces with the same origin

3.2 Linking by Behavior Pattern

The goal of linking by behavior pattern (LP) is to link different pseudonyms of one consumer (see Figure 3). There are two reasons for one consumer to have different pseudonyms attached to his identity: Either one supplier re-pseudonymizes his consumption trace database or the consumer has consumption traces in different databases with different pseudonyms respectively. The latter happens when a consumer changes his supplier. The old supplier still possess the consumption trace with one pseudonym and the new supplier starts to collect a consumption trace under a new pseudonym. This is equivalent to tracking consumers across different databases.

The former case, re-pseudonymization, means that a consumer's traces are stored under the pseudonym A in time interval t and stored under pseudonym B in time interval $t+1$ in the same database. This method could be applied by suppliers to prevent the de-pseudonymization of many years of consumption profiles under the same pseudonym in case one pseudonym is de-pseudonymized.

For this attack the attacker requires a database of pseudonymized consumption traces containing the pseudonym A. He tries to find a consumption trace with pseudonym B that has been created by the same consumer. Then the attacker can link the pseudonym A with pseudonym B.

In contrast to the LA attack, this attack can be applied even if the data sources do not overlap in time. This is because fundamental patterns in consumption are identified and subsequently looked for in the other data source. This means, that we can either consider consumption slices from two different consumption trace databases or two consumption traces from the same database but from different time intervals.

The feasibility of such an attack would also imply that everyone possessing current consumption traces and links to consumer identities can harvest all consumption traces that have been published (even in anonymized form) in the past or can be obtained for the past. On the other hand, legitimate holders of old consumption trace databases and the corresponding identities are able to "de-anonymize" data traces that are publicized in the future.

3.3 Exemplified Attacks

There is a multitude of attack scenarios that involve linking Smart Meter data and contained information to their origin household. We herein provide two examples:

Fake sick-day disclosure.

An employer could attempt to spy on his employees. In particular, he would like to know whether employees really spend their sick days at home. In order to do that, the employer obtains pseudonymized consumption traces for consumers in the geographic region of his employees' households. Using a correlation of vacation/company travel data with the consumption trace he links his employees' identities to the respective pseudonymized consumption traces. Then the employer can try to derive employee behavior on sick days from the employee's consumption trace.

Absence pattern deduction.

Another motivation for linking a household to its consumption trace is preparing a burglary [17]. Once a burglar has identified a worthwhile victim in the physical world, the burglar would like execute his plans undisturbed by the inhabitants. The burglar therefore performs an observation of the household in question and their weekend behavior. He simply finds out whether they stay at home or leave for the weekend over the course of several weeks. By correlating this information with a consumption trace database this household's trace can be identified in the database. Now, repeating long-term absence patterns of this household can be found, e.g. for a regular family meeting or a time share, and

subsequently the burglary can be scheduled for such a date. Traditionally, a burglar would need to observe a household for years to get these information. The linking of household identity and its consumption trace, however, allows him to tap into a wealth of information about a long time-frame of the household in question.

4. DATA ANALYSIS FRAMEWORK

So far we have studied the linking of consumption traces and consumer identities in an abstract manner. For conducting the presented attacks in practice, we now introduce a data analysis framework. This framework builds on concepts of machine learning and allows us to analyze consumption traces geometrically. To this end, the consumption trace of a consumer is mapped to a high-dimensional feature space, such that it can be analyzed by standard techniques of machine learning. In this geometric representation, the *linking by behavior anomaly* can be achieved using geometric anomaly detection, where unusual events are identified by a large distance from normal activity. Similarly, the *linking by behavior pattern* can be carried out using geometric classification, where the behavior of one consumer is separated from all other users in the feature space. Figure 4 illustrates this geometric interpretation of the two attacks.

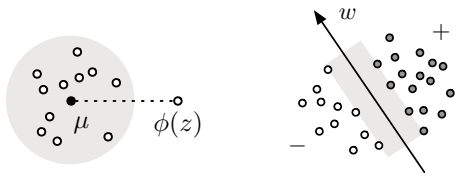

(a) Anomaly detection (b) Classification

Figure 4: Schematic depiction of data analysis: (a) distance to mean; (b) separating hyperplane.

Before presenting this data analysis framework, we introduce some basic notation. We refer to the known consumption trace of a consumer by $X = \{x_1, \ldots, x_n\}$, where X covers n days and each x_i corresponds to the measurements of one day. Depending on the resolution d of Smart Metering, each x_i is represented by a vector of d dimensions

$$x_i = \big(m_1, \ldots, m_d\big),$$

where m_j is the consumption measured at the j-th time slot of the Smart Metering resolution. Moreover, if we consider a grid of g bins associated with consumption values, we say that m_j falls into bin k, if m_j has the smallest difference to the consumption value associated with the k-th bin.

4.1 A Binary Feature Space

Mapping consumption traces to a vector space may seem trivial at a first glance, as the measurements of a day are already represented as a d-dimensional vector. However, for discriminating different patterns in this data, we require a more advanced representation that emphasizes the characteristics of each consumer and provides an expressive basis for application of machine learning.

Depending on the setup of electronic devices in a household, the consumption of a consumer changes between different states. Devices are switched on and off; thereby the consumption moves from one state to another. This discrete behavior is illustrated in Figure 5, which shows the consumption of one user over the period of one week on a fixed grid. Dark entries in the grid indicate frequent occurrences of a consumption value. It is notable that the consumption is neither a continuous nor a smooth function and several discrete states can be observed. For example, between 10:00 o'clock and 11:00 o'clock the consumption matches one of three possible states at roughly 300, 900 and 1500 W/h respectively.

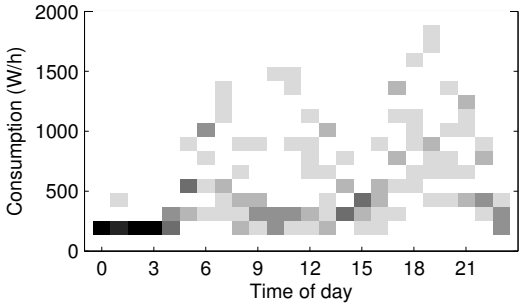

Figure 5: Grid representation of a consumption trace. Dark color indicates frequent occurrences.

Inspired by this observation, we construct a feature space that specifically reflects these states of consumption data. In particular, we employ a grid of g bins that spans the range of observable consumption values. Using this grid we define an indicator function $I_{j,k}$

$$I_{j,k}(x) = \begin{cases} 1 & \text{if } m_j \text{ falls into bin } k \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

that returns 1, if the j-th measurement of a day x falls into the k-th bin of our grid, and 0 otherwise.

If we compute the indicator function $I_{j,k}$ for all d measurements of a day and all g bins of the grid, we obtain a mapping

$$\phi : X \mapsto \mathbb{R}^{d \cdot g}, \quad \phi(x) = \big(I_{j,k}(x)\big)_{\substack{1 \le j \le d \\ 1 \le k \le g}} \tag{2}$$

that maps the consumption of one day to a binary feature space with $d \cdot g$ dimensions. For example, if we have one measurement per hour, that is $d = 24$, and use a grid with $g = 100$, we embed the data in a feature space of 2,400 binary features, each corresponding to a particular consumption value at a particular hour. In contrast to a naive representation with d dimensions, this feature space describes states independently of absolute differences. That is, each state change induced by the consumer is treated equally, whether it involves switching on a small desk lamp or a powerful washing machine. For the experiments in Section 5, we consider $d = 24$ and make use of a grid with 100 bins of consumption values.

4.2 Anomaly Detection

By embedding the consumption traces into an expressive feature space, we are able to phrase our attacks in terms of geometric relationships between data points. For determining unusual activity in the data of a consumer we employ a standard technique for detecting geometric outliers. Given a consumption trace of n consecutive days, we first learn a

profile of normal activity by computing the mean μ of the data in the feature space as follows

$$\mu = \frac{1}{n} \sum_{i=1}^{n} \phi(x_i). \quad (3)$$

The profile μ captures the states shared by the majority of the consumption traces. As each vectors $\phi(x_i)$ contains only binary values, each dimensions of μ can be interpreted as the probability for observing a particular grid value at particular time of the day. Geometrically, the deviation of a day z from this profile can be determined by simply computing the distance

$$d(z) = ||\phi(z) - \mu||. \quad (4)$$

Note that $d(z)$ corresponds to the Euclidean distance in the vector space and can be efficiently computed with standard software libraries. This generic approach to anomaly detection is illustrated in Figure 4(a). If we notice a large distance $d(z)$ for a day z, some of the consumption states of this day differ from normal activity and z is likely to contain an anomalous event.

This technique for computing profiles can also be applied to compare different sets of consumption traces. For example, if we have two reference sets X_1 and X_2 from the same consumer, we can compute two mean values μ_1 and μ_2 and compare the distance to both. This setting allows us to study different classes of days during analysis, as shown in Section 5 for weekdays and weekends.

4.3 Classification

For *linking by behavior pattern*, we aim at inferring patterns from the consumption trace of a consumer. However, we are not interested in modelling the complete behavior of a consumer, but determining patterns that discriminate his behavior from others. Thus, we employ the technique of classification and learn a discrimination between users. A robust method for learning such a discrimination is a Support Vector Machine (SVM) [2, 16]. An SVM basically determines a hyperplane in the feature space that separates two classes with maximum margin. This geometric concept is illustrated in Figure 4(b). The hyperplane is constructed as a linear combination of the training data and separates the consumption trace of one consumer c from all others. Formally, this hyperplane is given by a direction vector

$$w_c = \sum_{i=1}^{n} y_i \alpha_i \phi(x_i) \quad (5)$$

and an offset term b_c, where $y_i \in \{-1, +1\}$ are training labels indicating whether day x_i corresponds to consumer c and α_i are the learned coefficients.

To account for multiple consumers, we make use of the *one-against-all approach* and learn a hyperplane for each consumer separating him from the rest of users. The discrimination function for each consumer is then given by

$$h_c(z) = \langle \phi(z), w_c \rangle + b_c. \quad (6)$$

The function $h_c(z)$ reflects the distance from day z to the hyperplane of consumer c. The more consumption states and patterns are shared with c, the higher $h_c(z)$ gets. Hence, if we want to link a day z to a consumer, we simply assign it to those consumer c with the largest value for $h_c(z)$.

Similarly to the anomaly detection, there exists several efficient libraries for computing SVMs. In our experiments, we make use of LibLinear [4]—a library capable of learning with millions of dimensions and data points.

5. EXPERIMENTAL EVALUATION

To study the impact of our attacks in practice, we conduct experiments using anonymized consumption traces for 53 households. The data stretches over 221 days and has a resolution of one value per hour. The goal of the experiments is to demonstrate the efficacy of the two attacks described in Section 3. However, due to privacy reasons we do not have identity data for the consumption traces and therefore cannot fully implement the LA attack. Yet, we can identify significant anomalies in the consumption trace that could greatly help in linking identities to pseudonyms.

5.1 Identification of Anomalies

In our first experiment, we apply the technique of anomaly detection to each of the 53 consumers. We are interested in identifying days that stand out of regular energy consumption and might provide a good basis for linking the consumption data with an external data source.

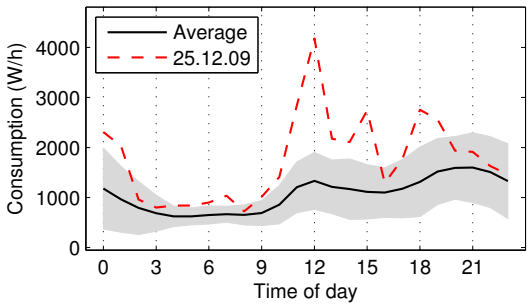

(a) Anomaly detected for consumer 21.

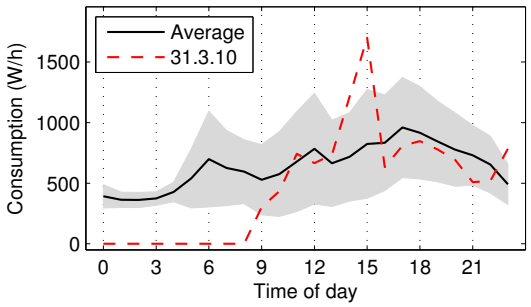

(b) Anomaly detected for consumer 49.

Figure 6: Exemplary anomalies for two consumers. The shaded area indicates the standard deviation.

Figure 6(a) shows exceptionally high energy consumption throughout the day, probably in the course of Christmas preparations. Depending on the amount of time this consumption trace spans this could mean different things: If the trace spans several years (which our specific data source does not) this would indicate that this consumer has not has had such extensive Christmas preparations because he/she previously went away for Christmas or celebrates Christmas for the first time at home. As this particular trace only

spans approximately 7 months, this deviation just indicates that the 25th of December means something special to this household, which could in turn indicate that it is Christian. Depending on the context of this household this could mean incriminating information.

Figure 6(b) displays how a consumer apparently starts energy consumption shortly before 9 o'clock. This may indicate that the consumer has moved in and started the first electric devices in his new apartment. In the afternoon exceptionally high load can be observed, consistent with the use of machines by craftsmen. If one could correlate this data with data sources that hold information about moving households in this region this could lead to the identification of the inhabitants.

While we have shown only two strong anomalies from our data set, several others can be identified for the consumers. Provided external reference data, it is trivial to correlate these anomalies with unusual events and there is a realistic chance of unlinking pseudonyms.

As mentioned in Section 3.1, anomalies can also be identified using different profiles of consumption. In this second experiment, the consumption on workdays and weekends are analyzed to determine whether the household inhabitants stay or leave home. In particular, we compare the profile of weekends and workdays to identify workdays that match the consumption behavior of weekends.

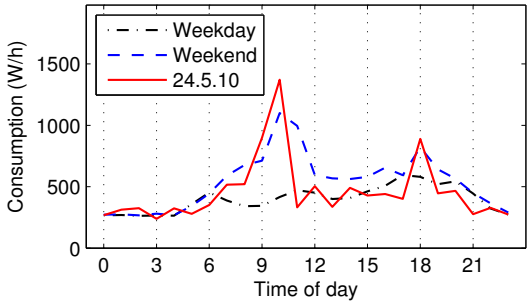

(a) Profiles for consumer 12 and "day-off".

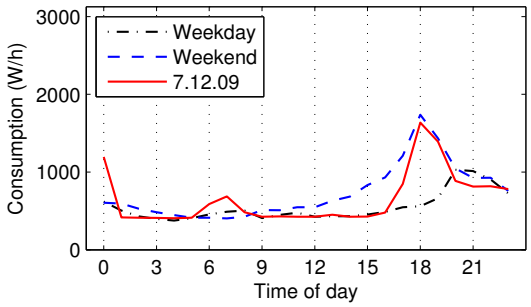

(b) Profiles for consumer 40 and "day-off".

Figure 7: Exemplary weekday and weekend profiles. The red line indicates a potential "day off".

Figure 7(a) displays how usual workday and weekend profile look like for a consumer. For the 24th of May however, we can identify a day that matches a weekend day judging by its energy consumption but is a Monday (workday). Figure 7(b) shows the same for another consumer for the 7th of December 2009 which is a Thursday.

It is notable that for the given resolution of Smart Metering accurate profiles can be learned for the consumption data. While we have shown examples for comparing weekends and weekdays, several other scenarios exist that would allow to further structure and analyze the electricity consumption of a household. Together with the anomaly detection, these analysis steps clearly allow to narrow down the search for a particular identity and help to unlink its pseudonym.

5.2 Linking by Behavior Pattern

In this experiment we conduct the LP attack by using the classification technique from Section 4.3 to determine whether consumption traces have the same origin household. We use one time interval of our consumption trace database for training our machine learning framework and subsequently provided it with test data from a different, non-overlapping time interval. The algorithm implemented by the framework then tries to link consumption traces that behave similarly. We measure its accuracy by the relative frequency of correct linking decisions.

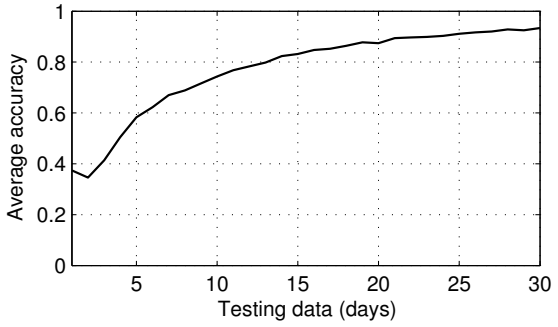

(a) Accuracy for varying size of test data.

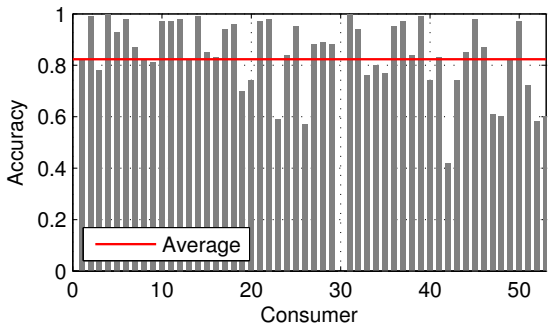

(b) Per-consumer accuracy for 14 days of test data.

Figure 8: Classification accuracy for varying sizes of test data (a) and individual users (b).

Figure 8(a) displays linking accuracies for training data of 60 days in dependency of test data size, where we assume fixed pseudonyms during the given time spans. The graph steadily climbs to over 90% accuracy for 30 days of test data. Figure 8(b) represents a breakdown of the linking accuracy for testing data of 14 days. For several pseudonyms an almost perfect unlinking is possible and on average an accuracy of 83% is attained, corresponding to 5 correct identifications out of 6 consumers. Note that pseudonym 30 undergoes significant perturbations over the course of our data

which leads to repeated mis-classification and subsequently zero accuracy.

Figure 9 displays the accuracy of our approach depending on the sizes of the training set and the test data in days. One can see, that the size of the test data has a slightly stronger impact on the accuracy then the size of the training data. Overall, the accuracy reaches approximately 83% if the training and test data is larger than 28 days. As a result, our attack is even effective if a re-pseudonymization is conducted every month.

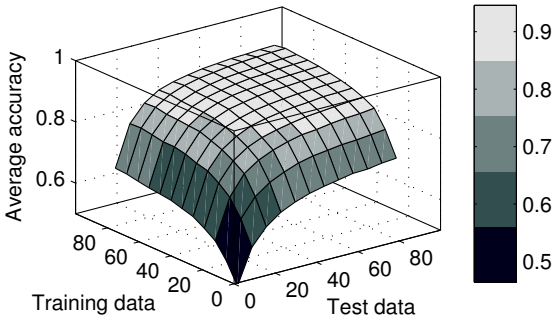

Figure 9: Impact of training and testing data size on classification accuracy.

6. MITIGATION TECHNIQUES

We finally investigate in this section three mitigation alternatives and their ability to contain and mitigate the aforementioned attacks.

6.1 Lower Resolution

A simple yet promising approach to mitigating our attacks is to lower the resolution of Smart Metering. The idea behind is that consumption traces are blurred and therefore anomalies and patterns are covered. In order to evaluate this mitigation technique we scale the consumption traces down in several steps and executed our experiments on the down-scaled data.

Figure 10 represents our anomaly detection for resolutions of 6, 3 or 1 value per day for the anomaly that was identified in Figure 6(a). One can see that the anomaly remains visible, even if the Smart Meter only records one value per day. The reason for this is that the anomaly spanned a larger part of the day and hence had a high impact on the total energy consumption of that day. Figure 11 shows an anomaly that behaves differently. While for resolutions of 6 or 3 values per day the anomaly still can be identified, it is not recognizable in the 1 value per day.

Regarding the LP attack the reduction of Smart Metering resolution has a bigger effect. Figure 12 shows the linking accuracy in dependency of the test data size for different resolutions. In contrast to Figure 9 the linking accuracy drops significantly with a reduction of the resolution. While the accuracy still reaches almost 70% for 8 measurements per day it drops to approx. 4% for one measurement a day.

These results show that a reduction of the Smart Metering resolution has mixed effects on our attacks. For anomaly detection (and subsequent anomaly linking) the attacker will probably be still quite successful if he manages to find

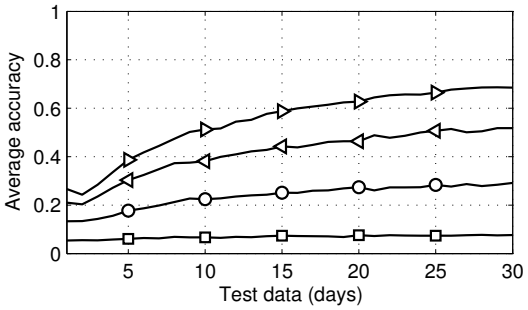

Figure 12: Classification accuracy with different resolution. The accuracy is given for resolutions of 8 (▷), 6 (◁), 3 (○) and 1 (□) value per day.

anomalies that have a big impact on the total energy consumption. However, short bursts of energy consumption will not be recognizable at low resolutions. Moreover, our results show that the LP attack can be successfully mitigated. The linking accuracy drops significantly with every reduction of the resolution.

Whether lower resolutions are a viable mitigation also depend on the supplier's requirements: The resolution of real-time tariffs and time-of-use prices is limited by the chosen Smart Metering resolutions.

6.2 Frequent Re-pseudonymization

Re-pseudonymization could be considered another mitigation technique. The holder of the identity and consumption databases introduces new pseudonyms for identities every now and then. Subsequently arriving consumption trace items of a household will be recorded under the new pseudonym. This leads to the effect that a holder of only the consumption trace database has only short intervals of data per pseudonym. Which means, that the training set for the LP attack is limited by the re-pseudonymization time frame.

If we assume that the attacker tries to track origins of consumption traces across re-pseudonymization he would try to match two data sets of different pseudonyms to determine whether they belong to the same origin. If we assume that the re-pseudonymization time frame is constant then the size of the training set and test set are the same. The potential effect on the LP attack can be seen in Figure 9.

Even if the re-pseudonymization time frame has only 20 days we can link two pseudonyms of the same identity with 80% accuracy. Thus, for the method to be effective one would need to re-pseudonymize in very short intervals. There are two major drawbacks of re-pseudonymization: First, analysis over frequently re-pseudonymized consumption traces can only span intervals of the re-pseudonymization time frame. Analysis by contractors that requires long-term consumption data is not possible because they would need information about the pseudonymization for that. Second, frequent re-pseudonymization incurs an overhead of storing the linking between the different pseudonyms and the identity.

6.3 Privacy-preserving Techniques

Another mitigation technique is the prevention of transmitting and storing consumption traces in the first place. The approaches described in [9] or [21] reduce the amount

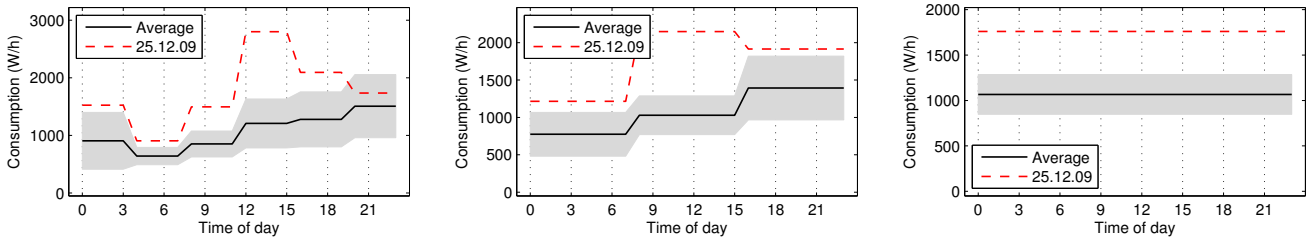

(a) 6 values per day. (b) 3 values per day. (c) 1 value per day.

Figure 10: Anomaly detection with different resolution. The detected anomaly from Figure 10(a) is shown.

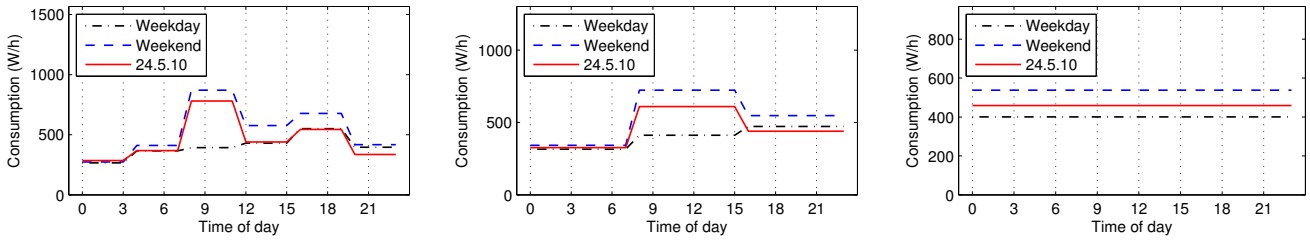

(a) 6 values per day. (b) 3 values per day. (c) 1 value per day.

Figure 11: Weekday and weekend profiles with different resolution. The profiles from Figure 7(a)) are shown.

of data that is effectively intelligible by the suppler to one value per reporting interval. This value represents the price for the energy consumption of this day. It allows application of high-resolution time-of-use or real-time prices and can be used to just report one value per year if desired. This way it combines the advantages of a very low resolution (one value per year) and the application of demand response tariffs.

It is noteworthy, that the single value per reporting interval does contain a weighting of the energy consumption by the applicable tariff. Therefore it contains more information than a single consumption value of similar low-resolution strategy. However, its applicability is constrained to billing and cannot be used for forecasting since the tariff skews the actual consumption.

7. RELATED WORK

Attacks on Smart Metering privacy.

The works [5, 6, 11, 12, 14, 15, 20] investigate in the broad area of consumption trace analysis and behavior analysis form energy consumption. The authors of [14], for instance, investigate the effectiveness of their developed behavior deduction using a NALM approach. Over the course of two weeks they conducted an experiment that collected electrical data and video surveillance of the inhabitants. Their developed system inferred behavior events and load events from the electrical data and evaluated the performance of behavior analysis with control data extracted from the video surveillance. Finally, they construct a sample disclosure metric that categorizes their behavior deductions into categories like presence, sleep schedule and others and rates the disclosure in those categories according to the ability of their behavior-extraction system. [11] on the other hand focuses on detecting and characterizing different appliances according to load signatures.

Mitigation by data prevention.

The following approaches are good at ensuring consumer privacy because no household specific consumption traces are stored at the supplier's and therefore the LA/LP attacks are not feasible:

In [1] a model for measuring privacy in Smart Metering is developed and subsequently two different solutions to ensure privacy are presented: A Trusted Third Party-based approach, where individual consumption profiles are aggregated at the third party and only sums are communicated to the supplier. The other approach attempts to mask consumption profiles by adding randomness to the actual profile with an expectation of the random distribution of zero.

In [18] another twofold approach is presented: The first solution employs a sophisticated Trusted Platform Module (TPM) in the Smart Meter to obtain signed tariff data from the supplier and calculate a trustworthy bill. The second solution makes use of the electrical grid infrastructure as a third party to anonymize up-to-date consumption values sent out constantly by Smart Meters.

In [13] the authors first perform an informal threat analysis of Smart Metering and provide a sketch for an attested Smart Meter architecture. Using virtualization, mandatory network access control and trusted computing techniques this architecture enables multiple applications to use the Smart Meter hardware and to work in a privacy-preserving and integer manner. The article identifies applications for billing the consumer very closely to the data origin (in the household) and applications that provide the consumer with a consumer portal. They achieve privacy-preserving Smart Metering billing by remote attestation of the billing software in the TPM of the Smart Meter.

A cryptographic approach to Smart Meter privacy has been presented in [9]. A privacy component homomorphically calculates the price locally in the household and only

reports the final price and cryptographic proofs to the supplier. With the help of those proofs the correct calculation with the correct tariff can be verified. In contrast to [21] this work focuses on how it can be built into existing Smart Meter reporting protocols.

Another cryptographic approach very similar to [9] is described in [21]. It focuses on realizing a variety of different tariff types with a cryptographic solution. Both approaches have been jointly evaluated in Section 6.3 with respect to their effectiveness to the developed attacks.

Mitigation by anonymization.

In [3] the authors propose a system to separate the data flow from the Smart Meter into two flows: One high- and one low-frequency data flow. The low frequency flow is attributed with the household's identity and can therefore be used for billing. The high-frequency data flow, that tells more about the habits of household inhabitants, is transmitted anonymously by the Smart Meter. An escrow service, potentially provided by the manufacturer of the Smart Meter, authenticates the anonymous high frequency flow towards the utility so that trust can be placed in its authenticity. The escrow service can disclose the identity of the high-frequency data flow in case of abuse.

This approach anonymizes the high-frequency data flow but cannot mitigate our attack vectors. Our assumptions are that attackers have access to anonymous (from their point of view) consumption traces but still manage to create linkability using correlation with secondary data sources.

Mitigation by hiding.

In [10] the authors propose to use a 'Load Signature Moderator' (LSM) and batteries to mask consumption events that represent 'privacy threats'. The LSM either detects or is notified by the appliances of approaching consumption events and could apply different algorithms like hiding, smoothing or obfuscation to hide those events from the Smart Meter. The batteries serve as energy buffer and enable the smoothing of actual loads. Then the authors measure the achieved privacy protection using three metrics: relative entropy cluster classification based similarity and regression analysis with different battery capacities and their best-effort moderation algorithm.

Our attack vectors would indeed be mitigated by the proposed approach. However, as the authors mention the proposed approach could conflict with cost-saving consumption strategies. On another note, it is questionable whether consumer see that their privacy value offsets the costs of such a system. If combined with cost-saving strategies, this might however be a desirable solution for consumers.

8. CONCLUSIONS

We have presented two attack vectors on the unlinkability of pseudonymized Smart Metering consumption traces. The first attack *linking by behavior anomaly* attempts to link a household identity to a pseudonymous consumption trace by anomaly correlation. The second attack *linking by behavior pattern* attempts to trace the origin of a consumption trace across different pseudonyms (due to re-pseudonymization or due to storage in different databases) by using patterns in their electricity consumption. To demonstrate the impact of the two attack vectors, we have presented a data analysis framework which allows us to conduct the attacks in practice and which we apply to perform experiments on real consumption traces.

Our experiments indicate that the task of finding relevant anomalies in consumption traces for the *linking by behavior anomaly* is feasible and allows deduction of household behavior. Regarding the *linking by behavior pattern* attack our experiments suggest that tracking the consumption trace across different pseudonyms is also feasible and can be executed quite accurately. Finally, we analyzed different mitigation techniques like lower resolution, frequent re-pseudonymization or data prevention with respect to their ability to thwart our attacks: A lower Smart Metering resolution has a mixed/good effect on the LA/LP attacks respectively, frequent re-pseudonymization requires very frequent (more often than 20 days) changes of pseudonyms to have noticeable (less than 80% accuracy) effects on linkability. Data prevention, by privacy-preserving cryptographic approaches that calculate the price in the household, has the best effect on both attacks and is also the most flexible with respect to high resolution time-of-use and real-time tariffs.

We have shown that alleged unlinkability introduced by pseudonymity of consumption traces is not sufficient for consumer privacy. Using the right secondary data sources attackers can link pseudonymized consumption traces back to consumers or track consumers across different databases of consumption traces. To prevent a failure of Smart Metering due to consumer distrust solutions must be found, that allow legitimate calculations on consumption traces without endangering consumer privacy.

9. FUTURE WORK

As we lack control data for our *linking by behavior anomaly* attack we could not fully evaluate its practical impact. Future work in this area is thus the investigation of linkability with adequate secondary data sources for willing consumers.

Another interesting question is whether persons can be tracked across different residencies. New apartments/houses induce some change in consumption patterns but to some extend personal habits and preferences might still be encoded in the consumption trace. A research question could be whether there is a component in the energy consumption pattern that remains the same, i.e. that identifies the inhabitants even across residencies?

10. ACKNOWLEDGMENTS

Martin Johns' work in this paper was partly funded by the German Federal Ministry of Economics and Technology (BMWi) as part of the e-IKT project with reference number 01ME09012.

Marek Jawurek's work in this paper was partly funded by the German Federal Ministry of Economics and Technology (BMWi) as part of the MEREGIOmobil project with reference number 01ME09007.

References

[1] J.-M. Bohli, O. Ugus, and C. Sorge. A privacy model for smart metering. In *Proceedings of the First IEEE International Workshop on Smart Grid Communications (in conjunction with IEEE ICC 2010)*, 2010.

[2] N. Cristianini and J. Shawe-Taylor. *An Introduction to*

Support Vector Machines. Cambridge University Press, Cambridge, UK, 2000.

[3] C. Efthymiou and G. Kalogridis. Smart grid privacy via anonymization of smart metering data. *2010 First IEEE International Conference on Smart Grid Communications*, pages 238–243, 2010.

[4] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9: 1871–1874, 2008.

[5] G. Hart. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12):1870 –1891, dec 1992. ISSN 0018-9219.

[6] G. W. Hart. Residential energy monitoring and computerized surveillance via utility power flows. *IEEE Technology and Society Magazine*, June 1989.

[7] W. Heck. Smart energy meter will not be compulsory. NRC Handelsblad (online), April 2009. `http://www.nrc.nl/international/article2207260.ece/Smart_energy_meter_will_not_be_compulsory`.

[8] A. Jamieson. Smart meters could be 'spy in the home'. Telegraph (UK) (online), October 2009. `http://www.telegraph.co.uk/finance/newsbysector/energy/6292809/Smart-meters-could-be-spy-in-the-home.html`.

[9] M. Jawurek, M. Johns, and F. Kerschbaum. Plug-in privacy for smart metering billing. *CoRR*, abs/1012.2248, 2010.

[10] 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*, pages 232–237. IEEE, 2010.

[11] H. Lam, G. Fung, and W. Lee. A novel method to construct taxonomy electrical appliances based on load signaturesof. *Consumer Electronics, IEEE Transactions on*, 53(2):653 –660, may 2007. ISSN 0098-3063.

[12] C. Laughman, K. Lee, R. Cox, S. Shaw, S. Leeb, L. Norford, and P. Armstrong. Power signature analysis. *Power and Energy Magazine, IEEE*, 1(2):56 – 63, mar-apr 2003. ISSN 1540-7977. doi: 10.1109/MPAE.2003.1192027.

[13] M. Lemay, G. Gross, C. A. Gunter, and S. Garg. Unified architecture for large-scale attested metering. In *in Hawaii International Conference on System Sciences. Big Island*. ACM, 2007.

[14] M. A. Lisovich, D. K. Mulligan, and S. B. Wicker. Inferring personal information from demand-response systems. *IEEE Security & Privacy*, 8(1):11–20, January 2010.

[15] 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*, BuildSys '10, pages 61–66, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0458-0.

[16] K.-R. Müller, S. Mika, G. Rätsch, K. Tsuda, and B. Schölkopf. An introduction to kernel-based learning algorithms. *IEEE Neural Networks*, 12(2):181–201, May 2001.

[17] D. Mulligan and J. Lerner. Taking the long view on the fourth amendment: Stored records and the sanctity of the home. Talk (online), January 2007. URL `http://www.truststc.org/pubs/318.html`.

[18] R. Petrlic. A privacy-preserving concept for smart grids. In *Sicherheit in vernetzten Systemen: 18. DFN Workshop*, pages B1–B14. Books on Demand GmbH, 2010.

[19] A. Pfitzmann and M. Hansen. A terminology for talking about privacy by data minimization: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management. online, Aug. 2010. URL `http://dud.inf.tu-dresden.de/literatur/Anon_Terminology_v0.34.pdf`.

[20] A. Prudenzi. A neuron nets based procedure for identifying domestic appliances pattern-of-use from energy recordings at meter panel. In *Power Engineering Society Winter Meeting, 2002. IEEE*, volume 2, pages 941 – 946 vol.2, 2002.

[21] A. Rial and G. Danezis. Privacy-preserving smart metering. Technical report, Microsoft Research, November 2010. URL `http://research.microsoft.com/apps/pubs/?id=141726`.