Authenticated Down-Sampling for Privacy-Preserving Energy Usage Data Sharing

Daisuke Mashima
Fujitsu Laboratories of America
1240 East Arques Avenue
Sunnyvale, CA, 94085
dmashima@us.fujitsu.com

***Abstract*—Thanks to the availability of fine-grained electricity usage data, electricity customers can benefit from a variety of services from utility companies as well as third party service providers, for energy efficiency, cost saving and incentives, social gaming, and so forth. However, at the same time, advancement of data analytics techniques may jeopardize customers' private, sensitive information. For instance, non-intrusive load monitoring (NILM) techniques could reveal customers' life cycle and style. In general, lowering granularity (sampling frequency) of electricity usage data readings can reduce the amount of information derived. In this direction, we propose a mechanism to allow electricity customers, in order to manage privacy risks, to flexibly down-sample their electricity usage data before sharing it with third-party service providers. On the other hand, our scheme does not invalidate a digital signature on the energy usage data (e.g., one made by utility companies) even after the down-sampling and thus retains verifiability of the data integrity. We then evaluate the overheads in terms of computation and communication and present possible integration into the Green Button information model with minor schema extension.**

I. Introduction

Smart electricity meters have capability to frequently monitor and report electricity usage information at each customer's premise. Such fine-grained, near-real-time information has enabled utility companies to provide efficient, stable electricity services. At the same time, various services and demand-side control technologies can be implemented on top of such data, for example demand response and dynamic pricing.

Collected energy usage data has also been made accessible to electricity customers, for example via Green Button Download My Data and Connect My Data services [1] provided by utility companies. Moreover, while typical smart meter data and Green Button data are usually recorded in 15-minute or longer intervals, customers can introduce additional metering devices that can provide them with much finer-grained information, around or over 1Hz sampling rate. For instance, neurio [2] sells a WiFi-enabled metering device that is installed in a distribution panel of a household and provides services for energy efficiency and home monitoring services. Besides, a number of utility companies allow customers to install Home Area Network (HAN) devices that can retrieve electricity usage data directly out of utility-installed smart meters. For example, PG&E authorizes the devices listed in [3] to be used with their smart meters. We envision the same trend in other countries. For instance, in Japan, such data flow is called "Route B" [4], and utilization of the data will rapidly grow towards the country's upcoming electricity market deregulation. Using such data and services provided based on it, customers can have deep insight and actionable information to not only improve energy efficiency but also to reduce electricity cost by taking advantage of utility's incentive and rebate programs.

Unfortunately, such a benefit does not come without negative side effect, namely privacy risks. Specifically, a number of emerging data analytics techniques, such as electricity load disaggregation or non-intrusive load monitoring [5], [6], [7], "behavioral privacy" [8] of customers might be compromised, through, for example, identification of what type of appliances are used and when. Combined with a number of side information, which can be obtained via a variety of sensors and/or IoT devices, the effectiveness of such privacy attacks would be further enhanced [9]. Moreover, appearance of open-source software for such analytics, such as NILMTK [10], though the toolkit was implemented for a benign purpose, might attract even non-skilled attackers to attempt the sophisticated analytics to extract sensitive information.

In this context, we propose a mechanism to enable authenticated down-sampling of electricity usage data, under which customers can flexibly adjust granularity of data. In general, lowering the resolution is expected to reduce the accuracy of privacy attacks mentioned above so is an effective way to mitigate privacy risks. At the same time, our scheme retains, even after the down-sampling by a customer, validity of data issuer's (e.g., utility companies' or metering devices') digital signature for ensuring data authenticity and integrity, which allows third parties to still utilize the data even for critical operations that require trustworthy data. Using the proposed mechanism, once electricity usage data is collected or downloaded, a customer can down-sample data before each data sharing solely based on her own privacy preference without relying on any other entities, which we believe establishes customer centricity regarding electricity usage data.

This paper is organized as follows. In Section II, we discuss related work. Section III discusses the proposed mechanism for authenticated down-sampling, followed by the evaluation of security and overhead in Section IV. We present integration into Green Button [1] in Section V and finally conclude the paper in Section VI.

II. Related Work

[11] employs redactable digital signature scheme to achieve the goals similar to ours. By allowing a customer to hide some part of her energy usage data without losing verifiability of data authenticity, it accomplishes privacy protection by customers themselves and meaningful energy consumption data utilization. However, the redactable signature scheme only allows us to exercise hide-or-show control. Even though it meets minimal disclosure principle for privacy preservation [12], it could restrict the utilization of data. For instance, when the middle part of the daily energy consumption pattern is redacted, it would be difficult to apply effective data analytics techniques like time-series modeling. On the other hand, the mechanism proposed in this paper does not have such a limitation since the resulting data is still continuous, but with lower granularity.

In the same paper, a customer-centric energy data management system model is proposed. Its concept advocates each customer's ownership and control over her own energy usage data and is consistent with the perspective of major utility companies in the US [13] as well as EUs smart grid data protection and privacy policy [14]. The authenticated down-sampling scheme proposed in this paper also fits in this model.

In the related paper [15], the authors proposed another privacy-preserving mechanism that can work in the same context. Namely, their scheme allows customers to add bounded noise without losing third parties' ability to verify "loose" data authenticity. The mechanism proposed in this paper can be considered as an alternative privacy preservation mechanism because, in use cases that require exact data, their scheme using noise may not be suitable. On the other hand, the data generated our mechanism is exactly accurate but simply with lower granularity.

A number of privacy-preserving schemes applicable to smart meter data have been proposed [16], [17], [18], [19], [20], [21], [22], [23]. All of these focused on privacy against data collectors (i.e., utility companies) by means of cryptographic primitives and / or aggregation of data from multiple smart meters. On the other hand, this paper focuses on privacy protection against third-party service providers, by enabling electricity customers to control the amount of information to be disclosed. Thus, our scheme is orthogonal and complementary to these schemes. Moreover, our scheme is applied to individual customer's energy consumption data, without requiring coordination with other customers for the sake of aggregation, based on his / her own privacy preference.

To the best of our knowledge, the scheme proposed by Peer et al. [24] is the most closely related to our mechanism. The proposed scheme in that paper also allows customers to disclose different resolution of smart meter data stream. In their scheme, time-series data with different resolution are encrypted with different keys. Thus, by disclosing only a key corresponding to a certain resolution, the receiver of data can obtain the data with the resolution with the corresponding key. One key difference from our scheme is that theirs does not consider data authenticity. Another difference is that, the resolution that a customer can select upon data sharing must be one of the resolution options prepared by the data source. On the other hand, our scheme does not have that limitation.

III. Authenticated Down-Sampling

A. Motivation

The accuracy of non-intrusive load monitoring or non-intrusive appliance load monitoring techniques [5], [6], [7] highly depend on the granularity of the data. Intuitively, energy usage data collected at higher sampling rate captures more detailed characteristics, thereby allowing us to do accurate identification of events happening in a household by means of pattern matching etc. However, as demonstrated in Figure 1, in which the right-hand side figure shows 2-minute interval data from Household Energy Survey (HES) dataset [25] while the left-hand side one shows the down-sampled version (30-minute interval), many of the fluctuation and oscillation can be hidden after down-sampling. This intuition can be backed by the fact that most of the load disaggregation techniques utilize electricity usage data with over 1Hz, as summarized in [26]. Thus, enabling users to perform down-sampling based

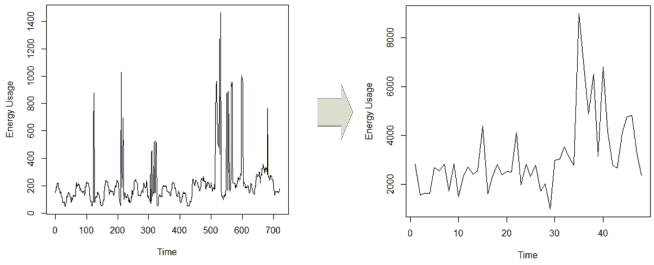

Figure 1. Effect of down-sampling of electricity usage data

on their privacy preference, nature of expected services, and trustworthiness of service providers is considered effective for controlling potential privacy risks.

If customers' privacy would be the only goal, we could have simply allowed customers freely edit (or even forge!) their data before sharing. However, as discussed in [11], some third-party service providers, such as demand response aggregators, require reliable data for billing and accounting purposes. To enable them to do verification of data authenticity, the data needs to be digitally signed by a trusted entity, for instance a utility company that plays a role as an issuer and a custody of electricity usage data or a HAN device from a trusted vendor. Besides the use cases just discussed, the cryptographic verifiability will be very effective when the electricity usage data is autonomously verified and processed on emerging crypto-currency platforms, such as [27].

To simultaneously meet needs of privacy control and data verifiability, in this paper we propose a mechanism to allow customers to down-sample electricity usage data without invalidating the data issuer's digital signature.

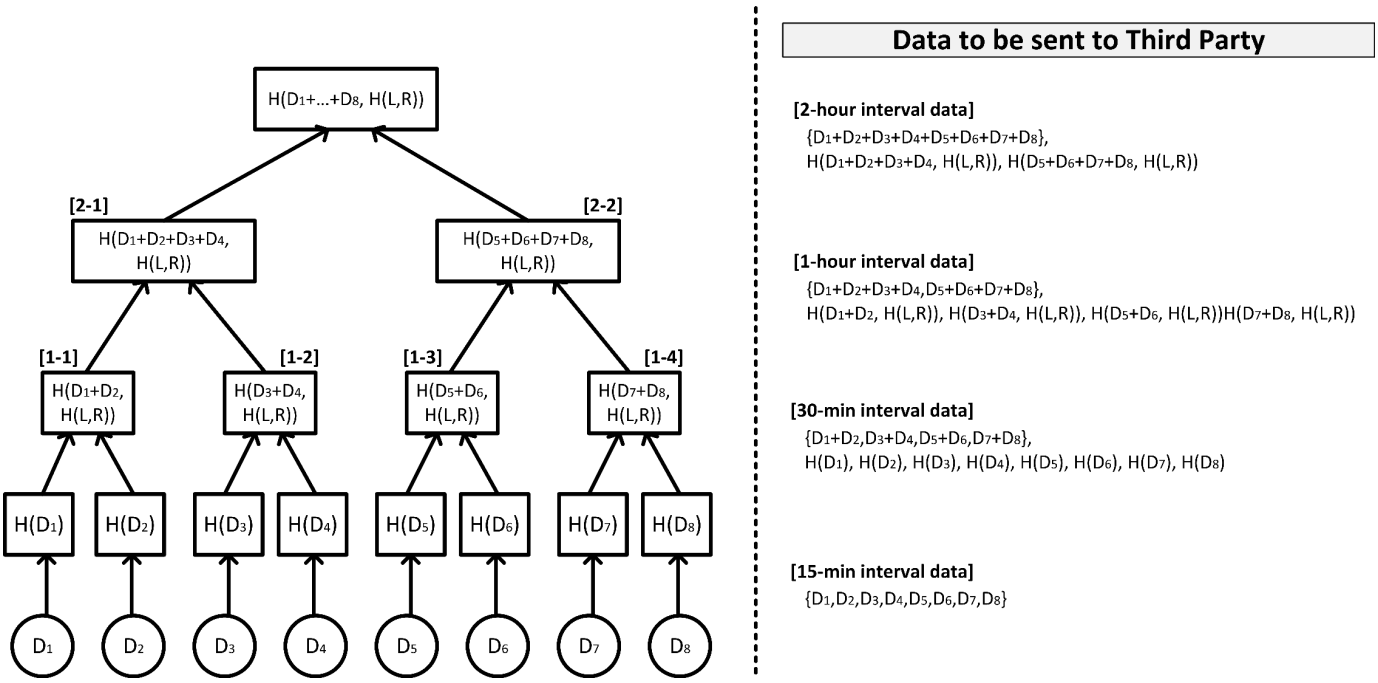

Figure 2. Construction of Authenticated Down-sampling Mechanism based on Merkle Hash Tree and examples of data to be sent to third parties at each down-sampling level

B. Construction

This section explains the construction of the authenticated down-sampling scheme we propose. The scheme is designed based on the redactable digital signature scheme [28] using Merkle Hash Tree [29], [30]. The illustration in the case of binary tree is shown in Figure 2. Although the discussion in this section focuses on a small-size example using binary-tree structure, extension to larger data and n-ary tree is straightforward.

In the figure, the round nodes at the bottom contain actual energy usage measurements ordered chronologically. The nodes at one level above store cryptographic hash values of each measurement, just like the normal Merkle Hash Tree construction. Calculation of nodes at the next level, ones marked **[1-1]**, ..., and **[1-4]**, is a bit different. Here, $H(L, R)$ represents a hash value of concatenation of two (left and right) children nodes. Then, each node is computed as a hash value of the sum of the original measurements under it, for example D_1 and D_2 in the case of node **[1-1]**, and the hash value $H(L, R)$. Nodes **[2-1]** and **[2-2]** are also calculated in the same way. In this case, the sum of D_1, D_2, D_3 and D_4 and the sum of D_5, D_6, D_7 and D_8 are used respectively. Eventually, the root node is calculated using hash values of **[2-1]** and **[2-2]** as well as the sum of all measurements, which is then signed by a data issuer's private key (e.g., a utility company's or metering device's). After that, the resulting signature is passed to an electricity customer (i.e., data subject) along with the series of electricity usage measurements.

After receiving the time-series data with the issuer's digital signature, a customer can prepare down-sampled data as follows. Let us illustrate with examples in Figure 2. Assume that the raw measurements are sampled every 15 minutes, which is typical in the current smart meter network deployment. If the customer wants to share 30-minute granularity data in a verifiable way, she can do so as follows.

1) Compute down-sampled time series: $\{(D_1+D_2), (D_3+D_4), (D_5+D_6), (D_7+D_8)\}$
2) Attach hash values: $H(D_1), H(D_2), \ldots, H(D_8)$

Then the customer can send these with the digital signature to a third party. The third party receiving the data can then calculate the hash values of **[1-1]**, ..., and **[1-4]**, which further allows it to calculate **[2-1]** and **[2-2]** and eventually the root hash value. The root hash can be verified against the data issuer's digital signature by using the issuer's public key. Likewise, if the customer wants to share data of 1-hour granularity, she can share the following data with a third party:

1) Compute down-sampled time series: $\{(D_1+D_2+D_3+D_4), (D_5+D_6+D_7+D_8)\}$
2) Attach hash values: **[1-1]**, **[1-2]**, **[1-3]**, and **[1-4]**

Verification at the third party can be done in the similar way.

C. Flexibility for Privacy Preservation

This section illustrates, by showing further examples, how the proposed scheme meets flexible privacy preferences an electricity customer may have.

1) Combination of Different Down-Sampling Rate: Under the proposed scheme, down-sampling can be applied at

different levels even within the single set of singed time-series energy usage data. For instance, in the case of demand response (DR) services, data within time period corresponding to a DR event or peak hours may require higher sampling rate, while data during the other time periods may allow lower sampling rate. Let us consider a case where only the first half of data points in Figure 2 must have higher granularity while the latter can be down-sampled to 30-minute interval. An electricity customer that has the 15-minute interval energy consumption data signed by the utility can construct and share such a data set by including:

- $\{D_1, D_2, D_3, D_4, (D_5 + D_6), (D_7 + D_8)\}$
- $H(D_5), H(D_6), H(D_7), H(D_8)$

For the first half, the receiver of this data set can calculate $H(D_1)$, ..., and $H(D_4)$ first and then proceed to calculation of **[1-1]** and **[1-2]**. For the second half, the receiver can immediately calculate **[1-3]** and **[1-4]** by using the hash values provided. Then, eventually the same root hash value is calculated, and the receiver can verify the authenticity of data. Likewise, if a customer wants to send $\{D_1, D_2, (D_3 + D_4), (D_5 + D_6), D_7, D_8\}$, hash values to be sent together are $H(D_3), H(D_4), H(D_5)$, and $H(D_6)$

2) Combination with Redaction: Our construction is based on Merkle Hash Tree, which is also a popular building block to construct redactable signatures [28]. Thus, our scheme can also apply redaction, in addition to down-sampling, to completely hide unnecessary part of the time-series energy usage data without invalidating the digital signature. Suppose a case where, in Figure 2, the customer wants to disclose only data from D_3 to D_4 with 1-level down-sampling and wants to entirely hide the others. Such transformation is made possible by sending the following to the receiver.

- $\{(D_3 + D_4), (D_5 + D_6)\}$
- **[1-1]**, $H(D_3)$, $H(D_4)$, $H(D_5)$, $H(D_6)$, **[1-4]**

In this case, a receiver of the data can compute **[1-2]** and **[1-3]**, using $H(D_3)$, ..., and $H(D_6)$, and then **[2-1]** and **[2-2]** by using **[1-1]** and **[1-4]** provided. Eventually, the same root hash is calculated, and through the attached digital signature the receiver can be convinced of the integrity of down-sampled time-series data provided in the plain text.

IV. Evaluation

In this section, we first discuss security of the proposed scheme. After that, using the prototype implementation, we present basic benchmarks to evaluate the performance and overhead in practical situations.

A. Security Discussion

Let us discuss the security aspect of the authenticated down-sampling scheme. We have two major security goals to be met:

- Confidentiality of higher-granularity data when down-sampling is done, beyond what is inferred from the down-sampled value
- Integrity of down-sampled time-series data

Regarding the confidentiality goal, as long as secure hash function is used, it is in theory very difficult for attackers to identify a raw data from its hash value. Thus, attackers, with knowledge of data down-sampled by one level, can not recover the original data. Likewise, an attacker can not recover $(D_1 + D_2)$ or $(D_3 + D_4)$ from **[1-1]** and **[1-2]**. Thus, an attacker with data down-sampled by two level can not recover data down-sampled by one level.

However, we need to take the possibility of brute-force attacks into consideration. Since the data we are handling is only numbers within a certain range that can be easily narrowed down by educated guess. Moreover, a down-sampled value provided in plain text would also help attackers. In this way, an attacker could reveal the original values by exhaustively trying all possibilities. In nature, it is impossible to fully prevent this type of attacks, but we can employ a mechanism using per-customer key and initialization vector as implemented in [11] to make brute-force attacks difficult.

Next, to illustrate how the designed scheme can meet the second goal, let us consider the following two attack scenarios.

1) Tampering Data: Probably the most straightforward attacks to compromise our scheme would be tampering with some or all of the energy usage measurements or down-sampled values. However, modified values generate different hash values, either at leaf nodes or intermediate nodes, which will propagate to the root hash value. Thus, the verifier of the data can detect unauthorized modification of data by using the issuer's digital signature.

2) Reordering Data: Another type of attack would be to change the order of raw measurements or down-sampled values. For instance, in the example of Figure 2, when sending 30-minute granularity data, a malicious customer may send $\{(D_3 + D_4), (D_1 + D_2), (D_5 + D_6), (D_7 + D_8)\}$ instead of $\{(D_1 + D_2), (D_3 + D_4), (D_5 + D_6), (D_7 + D_8)\}$. In this case, when a verifier calculates the hash value corresponding to the node **[2-1]**, the resulting value would be $H((D_3 + D_4) + (D_1 + D_2), H(R, L))$. Even though the sum is the same as the original, H(R,L)≠H(L,R), which results in a different hash value at **[2-1]**. Since this difference will affect the root hash, the signature verification process will fail. Moreover, if $\{(D_1 + D_2), (D_5 + D_6), (D_3 + D_4), (D_7 + D_8)\}$ is sent to the verifier, it is clear that both **[2-1]** and **[2-2]** become different hash values, which again can be detected upon signature verification.

B. Performance and Overhead

To evaluate the computational overhead of the proposed scheme, we implemented the prototype module that calculates the root hash value in a binary tree for randomly generated time-series data. As a hash function, *HmacSHA256* was used with a hard-coded key. We employed a keyed hash function for the experiment following the proposed countermeasure against brute-force attacks discussed in [11]. The module is implemented using Java, and measurements were taken on a Windows 7 laptop equipped with Intel Core i7-3667U processor and 8GB RAM. Figure 3 summarizes the results. Each plot in the figure is the average of 20 measurements, and maximum and minimum calculation time observed are

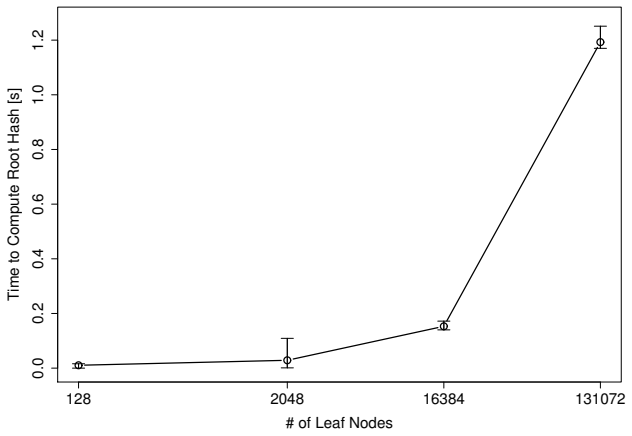

Figure 3. Time to calculate a root hash value for different number of leaf nodes

also shown as error bars. Regarding the size of time-series data, we considered 2^7 measurements, which can cover 15-minute interval readings for a day, 2^{11} that covers 1-minute interval measurements for a day, 2^{14}, which roughly matches the number of measurements with 5-second interval sampling rate for a day, and 2^{17} to cover 1-second interval sampling for a day.

As can be seen, for data with sampling interval that is typical to smart-metered data, the processing time is negligible, so the scheme is scalable enough to be used by utility companies with a large number of customers. In the case of high sampling rate, it takes over 1 second. However, as of today (and in the near future), such a data is only available on a HAN device in customer's premise, which directly obtains data from a smart meter. Because it is supposed to handle data of only a single customer, we believe this processing time is still acceptable.

Next, let us discuss the communication overhead when a customer is sending data to a third party. If we use Hmac-SHA256 as a hash function, the length of each hash value is 32 Bytes. If the number of measurements included is 2^7, and down-sampled by one level, we need to include 2^7 hash values, which add up to 4 KBytes. In case there are 2^{17} leaf nodes and they are down-sampled by one level, the size will be around 4 MBytes. Increase by this level is still acceptable when the data is sent over the Internet. Also, when the degree of down-sampling is higher, the overhead caused by the attached hash values becomes smaller.

V. Integration into Green Button

Recent years, Green Button [1] has been attracting attention as a standardized format of electricity usage data exchange. A number of major utility companies in the US have started services based on it to empower customers to access and utilize their own data. Green Button data has been already employed by a number of services for electricity customers, including data analytics for energy cost saving, social gaming and so forth. Thus, integration of our scheme into Green Button data model will be valuable to facilitate the broad adoption of the technology and expand the applicability. In this section we briefly discuss an enhanced Green Button data model that can accommodate the proposed mechanism.

One possible design can be illustrated in Figure 4. The essential enhancement is the modification of *IntervalReading* elements. At the high level, each *IntervalReading* element can be considered as an intermediate node in a hash tree. Thus, as described in Figure 2, it contains the sum of energy consumption measurements under it in the tree as well as hash values of its immediate children nodes, which can be stored in *value* and newly-added *ChildHash* element. In addition, when down-sampling is done, the contents of *timePeriod* element must be updated accordingly. As described in [11], the specification of a hash function etc. and a key for keyed hash calculation can be stored under a *HashInformation* element, and utility's digital signature can be stored under *SignatureInformation*. As can be seen, only a minor modification would be required.

VI. Conclusions and Future Work

The utilization of fine-grained electricity usage data collected by smart meters and / or home area network devices are increasing while the privacy risks associated with it has been argued. In this paper, we presented a mechanism to allow electricity customers to flexibly down-sample their own data, for the sake of privacy preservation, before sharing it with third-party service providers. Moreover, such transformation can be done without invalidating data issuer's digital signature, thereby retaining verifiability of data integrity by any third parties. Computation and communication overheads and implementation on top of Green Button information model [1] are also discussed to demonstrate practicality of its deployment.

An important future work is quantitative evaluation of data utility and privacy trade-offs. To conduct such an evaluation, we need to define a way to quantify how much privacy an electricity customer can attain by applying a certain level of down-sampling to assist the customers' informed decision. Implementation of user-friendly tool for visualizing such a trade-off would be also desired. For instance, if accuracy of typical non-intrusive load monitoring attacks is the primary concern of electricity customers, we can utilize an open-source tool like NILMTK [10] to mount simulated attacks on down-sampled data and visualize results. However, privacy in practice incorporates a variety of different aspects, so development of more generic, comprehensive quantification schemes and tools is an very interesting future research direction.

Besides, we plan to explore further application scenarios, including one using crypto-currency mechanisms. For example, in services like SolarCoin [31], amount of energy generated by PV needs to be reported, and in return, to encourage renewable generation, coins are granted based on the reported amount. In such a scenario, verifiability of data integrity is crucial for fairness and soundness of the service, but down-sampled data may suffice.

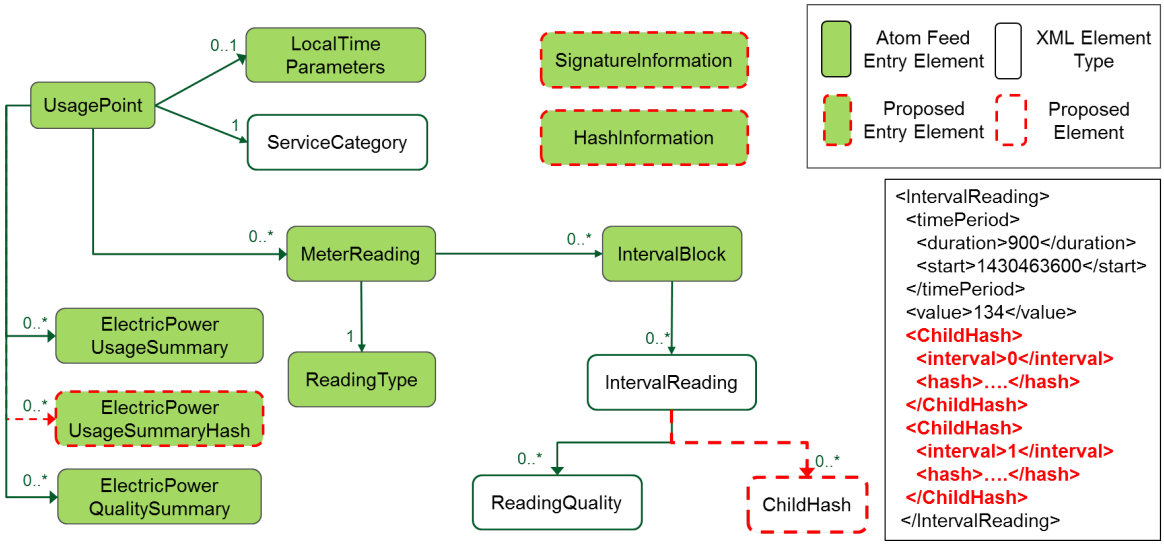

Figure 4. Extended Green Button information model

Finally, although in this work we only focused on electricity usage data, we believe the scheme itself is applicable to other types of time-series data in smart grid systems as well as in other domains. Considering application in other areas is also part of our future work.

REFERENCES

[1] "Green Button," http://www.greenbuttondata.org/.

[2] "neurio," http://www.neurio.io/.

[3] "Validated han devices," http://www.pge.com/en/myhome/saveenergymoney/rebates/han/validateddevice/index.page.

[4] I. Kitagawa and S. Sekiguchi, "Technologies supporting smart meter networks," *Fujitsu Science Technology Journal*, vol. 49, no. 3, pp. 307–312, 2013.

[5] D. Bergman, D. Jin, J. Juen, N. Tanaka, C. Gunter, and A. Wright, "Non-intrusive load-shed verification," *Pervasive Computing, IEEE*, vol. 10, no. 1, pp. 49 –57, jan.-march 2011.

[6] G. Hart, "Nonintrusive appliance load monitoring," *Proceedings of the IEEE*, vol. 80, no. 12, pp. 1870 –1891, dec 1992.

[7] M. Zeifman and K. Roth, "Nonintrusive appliance load monitoring: Review and outlook," *IEEE Transactions on Consumer Electronics*, vol. 57, no. 1, pp. 76–84, 2011.

[8] SGIP Smart Grid Cybersecurity Committee, "Guidelines for smart grid cybersecurity: Vol. 2, privacy and the smart grid," http://csrc.nist.gov/publications/drafts/nistir-7628-r1/draft_nistir_7628_r1_vol2.pdf.

[9] D. E. Phillips, R. Tan, M.-M. Moazzami, G. Xing, J. Chen, and D. K. Yau, "Supero: A sensor system for unsupervised residential power usage monitoring," in *Pervasive Computing and Communications (PerCom), 2013 IEEE International Conference on*. IEEE, 2013, pp. 66–75.

[10] N. Batra, J. Kelly, O. Parson, H. Dutta, W. Knottenbelt, A. Rogers, A. Singh, and M. Srivastava, "NILMTK: An Open Source Toolkit for Non-intrusive Load Monitoring," in *Fifth International Conference on Future Energy Systems (ACM e-Energy)*, Cambridge, UK, 2014.

[11] G. Lahoti, D. Mashima, and W. Chen, "Customer-centric energy usage data management and sharing in smart grid systems," in *Proc. of ACM Smart Energy Grid Security Workshop*, 2013.

[12] "The Laws of Identity," http://www.identityblog.com/stories/2004/12/09/thelaws.html, 2004.

[13] Pacific Gas and Electric Company, "Pg&e privacy policy," http://www.pge.com/en/about/company/privacy/customer/index.page.

[14] European Parliament, "Report on the local and regional consequences of the development of smart grids," http://www.europarl.europa.eu/sides/getDoc.do?pubRef=-//EP//TEXT+REPORT+A7-2014-0019+0+DOC+XML+V0//EN.

[15] D. Mashima and A. Roy, "Privacy preserving disclosure of authenticated energy usage data," in *Proc. of the 5th IEEE International Conference on Smart Grid Communications (SmartGridComm 2014)*. IEEE, 2014.

[16] K. Kursawe, G. Danezis, and M. Kohlweiss, "Privacy-friendly aggregation for the smart-grid," in *Privacy Enhancing Technologies*. Springer, 2011, pp. 175–191.

[17] E. Shi, T.-H. H. Chan, E. G. Rieffel, R. Chow, and D. Song, "Privacy-preserving aggregation of time-series data." in *NDSS*, vol. 2, 2011, p. 4.

[18] M. Jawurek, M. Johns, and F. Kerschbaum, "Plug-in privacy for smart metering billing," in *PETS*, 2011, pp. 192–210.

[19] A. Molina-Markham, G. Danezis, K. Fu, P. J. Shenoy, and D. E. Irwin, "Designing privacy-preserving smart meters with low-cost microcontrollers," in *Financial Cryptography*, 2012, pp. 239–253.

[20] G. Barthe, G. Danezis, B. Grégoire, C. Kunz, and S. Z. Béguelin, "Verified computational differential privacy with applications to smart metering," in *CSF*, 2013, pp. 287–301.

[21] G. Danezis, C. Fournet, M. Kohlweiss, and S. Z. Béguelin, "Smart meter aggregation via secret-sharing," in *ACM Smart Energy Grid Security Workshop*, 2013, pp. 75–80.

[22] A. Rial and G. Danezis, "Privacy-preserving smart metering," in *Proceedings of the 10th annual ACM workshop on Privacy in the electronic society*. ACM, 2011, pp. 49–60.

[23] R. Lu, X. Liang, X. Li, X. Lin, and X. S. Shen, "Eppa: An efficient and privacy-preserving aggregation scheme for secure smart grid communications," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 23, no. 9, pp. 1621–1631, 2012.

[24] C. D. Peer, D. Engel, and S. B. Wicker, "Hierarchical key management for multi-resolution load data representation," in *Proc. of the 5th IEEE International Conference on Smart Grid Communications (SmartGridComm 2014)*. IEEE, 2014.

[25] J.-P. Zimmermann, M. Evans, J. Griggs, N. King, L. Harding, P. Roberts, and C. Evans, "Household electricity survey: A study of domestic electrical product usage," *Intertek Testing & Certification Ltd*, 2012.

[26] K. Carrie Armel, A. Gupta, G. Shrimali, and A. Albert, "Is disaggregation the holy grail of energy efficiency? the case of electricity," *Energy Policy*, vol. 52, pp. 213–234, 2013.

[27] "Ethereum," https://www.ethereum.org/.

[28] R. Johnson, D. Molnar, D. X. Song, and D. Wagner, "Homomorphic signature schemes," in *CT-RSA*, 2002, pp. 244–262.

[29] R. C. Merkle, "Protocols for public key cryptosystems," in *IEEE Symposium on Security and Privacy*, 1980, pp. 122–134.

[30] R. Merkle, "A certified digital signature," in *CRYPTO*, 1989, pp. 218–238.

[31] "SolarCoin," http://solarcoin.org.