Achieving Differential Privacy in Smart Home Scenarios

Israel C. Vidal¹, Franck Rousseau², Javam C. Machado¹

¹Universidade Federal do Ceará (UFC), Fortaleza - CE - Brazil

²Université Grenoble Alpes, CNRS, Grenoble INP, LIG, Grenoble - France

{israel.vidal, javam.machado}@lsbd.ufc.br, Franck.Rousseau@imag.fr

Abstract. With the growth of the Internet of Things (IoT) and Smart Homes, there is an ever-growing amount of data coming from within people's houses. These data are intrinsically private and should be treated carefully, despite their high value for analysis. In this work, we propose a differentially private strategy to estimate frequencies of values in the context of Smart Home data.

1. Introduction

With the popularization of the Internet of Things (IoT) and the greater availability of various kinds of sensors in the market, there is an increasing amount of data being generated. We expect that by the year 2021, the amount of data generated by IoT devices, people and machines will reach the magnitude of zettabytes. These data can be beneficial for improving services, for example, by using Smart Meters data to gain a better understanding of the energy use in a city. However, careful attention to the privacy of these data is becoming more urgent. The lack of care with privacy can lead to severe problems, as shown in [Molina-Markham et al. 2010], where, through relatively simple statistical methods, one is able to identify crucial private information, such as if any household member watched the game on a given night or if the household members were late for work.

Although many works in the literature implement privacy through a trusted entity who has access to the raw data of a set of users, in real-world scenarios it is often not reasonable to depend on such an entity. More recent works focus on the local perspective of privacy, where the privacy process is done closer to the user without depending on a trusted third party entity. Besides that, IoT data often appears as streaming data, which brings some challenges due to its intrinsic characteristics, i.e., the data is potentially unbounded and happens in a non-predictable order.

In Differential Privacy (DP) [Dwork and Roth 2014], a mechanism \mathcal{M} is said to be differentially private if the probability of any output of \mathcal{M} does not vary significantly, by a threshold of ε , independent of the input. DP was initially proposed to work as an interactive model [Dwork et al. 2006], responding privately to statistical queries in a database. In this interactive scenario, a trusted entity that has access to the raw data is necessary. However, in more recent work, such as [Erlingsson et al. 2014], there has been significant interest in the local version of this model, called Local Differential Privacy (LDP), where a randomization process is done locally to ensure the definition of DP.

In this paper, we present a strategy that guarantees Local Differential Privacy for estimating frequencies of values in the context of Smart Homes. To evaluate our work, we have used real sensor data from Smart Meters [UK Power Networks 2015].

2. Related Works

The authors of [Molina-Markham et al. 2010] tackle the problem of privately charging energy consumption. In addition to proposing a statistical procedure capable of identifying house activities in fine-grained measurements, showing that there must be a meticulous privacy procedure to make use of smart meter data, they describe a protocol that allows smart meters to report a bill without revealing how the energy was used. The procedure uses cryptography and zero-knowledge proof to guarantee that the company will not have access to the information from which house comes a given data, even though the company will be able to charge for the energy used. A downside of this work is that they still have access to what they call *blinded data*, which consists of the data from all houses with the identification removed, and this is not enough for guaranteeing privacy.

The work [Ács and Castelluccia 2011] uses differential privacy to deal with the problem of using consumption data to learn privately about users. The approach is based on the Laplace Mechanism, which adds a noise sampled from a Laplace distribution to the result of a numerical query. The authors propose a Distributed Laplace Mechanism (DLM). The information they want to learn privately in this work is the consumption summation of N houses in a given time. To do this, each house adds a small noise n to its consumption c (which is not enough to guarantee DP), and encrypts the report r = c + n in a way that the server is not able to decrypt a report alone, but it is able to decrypt the summation of the reports. The server, then, has $S = \sum_{h=1}^{N} c_h + n_h = \sum_{h=1}^{N} c_h + N$, where N is a noise that follows the Laplace distribution as previously presented, i.e. that is enough to guarantee DP. This strategy is useful for learning the summation of consumption, but cannot be used to learn more information than that.

IoT data often appear as data streams. Works that deals with the problem of guaranteeing privacy in the context of streaming data deal with additional complexity because streaming data is potentially unbounded and continuously generated at rapid rates. The work [Leal et al. 2018] proposes a strategy to estimate the sensitivity and also presents a microaggregation algorithm that is capable of enhancing the utility for publishing differentially private data using the Laplace Mechanism in the context of streaming data. This work depends on a trusted third party entity to achieve its privacy, which may not be acceptable in the context of IoT data and smart homes.

The work [Cao and Yoshikawa 2015] uses differential privacy to publish statistics about streaming of trajectories. The objective is to publish, for a defined set of possible locations, how many people are in each location at a given time. The authors use the concept of a l-trajectory, i.e., a trajectory of size l. They show that it is possible to guarantee that a l-trajectory is DP. The concept of l-trajectory proposed is relevant and gives us an insight that to solve the privacy problem in streaming scenarios it may be useful to simplify the problem in order to be able to achieve a solution. On the other hand, this work, besides depending on a third party trusted entity, need to know beforehand the set of locations, which may not be reasonable in real-world scenarios.

3. Proposal

Our proposal aims to provide LDP for users, in the context of Smart Homes, that agree to provide their data to entities, so those entities can learn privately from data produced inside peoples' houses and, for example, provide better services. From now on, we will consider that this entity is the Service Provider (SP), but in real-world scenarios, it would be possible that they were two separated entities. We do not tackle the problem of charging for consumption, as described in [Molina-Markham et al. 2010].

It is necessary to collect Smart Home data privately because those data are intrinsically sensitive and the lack of proper care in their management could lead to harmful inferences, e.g., the energy producer entity could learn, using fine-grained energy measurements, when a given house is empty.

The solution was thought to work over an edgeOS, i.e., a specialized operating system that runs in an edge gateway, from now on called edgeBox, and manages smart things. In this paper, we have omitted the full architecture of an edgeOS, but [Shi et al. 2016] can be checked for more details. For our proposal, what is important about an edgeOS is that there is a data abstraction layer in it that gathers data produced by things inside a house. Our solution works between the data abstraction layer and all external communication to guarantee that all data that goes outside the house is private. A possible exception for this is that for the SP to be able to charge for the consumption, it may need to have access to coarse-grained measurements. As shown in section 2, there are possible strategies to charge privately.

Figure 1 illustrates the scenario where the SP has access to the total consumption in order to charge for it (ideally in a private way) and uses the Privacy Gateway proposed in this paper to learn from the data generated by things inside houses. The process executed by the Privacy Gateway will be detailed next.

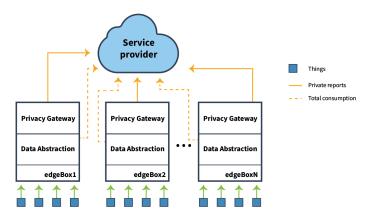


Figure 1. Communication between Service Provider and the N selected houses

The solution works as follows: the SP contacts a number N of people (the greater, the better) in different houses, and offers them something, e.g., a discount on the energy bill, in exchange for a defined privacy budget ε to be consumed in a number k of reports. Notice that a smaller budget means a more private output. The SP sends, then, a set H of parameters to each participant p. Every participant must use these parameters in H, so the SP can later decode the reports and get utility from them. The parameters are: the number of bins, the ranges of each bin, the number k of reports and the time δ after which each report will be sent. Given that the value to be sent is likely to be a real number in the context of IoT data, we will transform this value into a discrete form, in order to be able to make it private so that we can later get information from it. To do this, we will use a histogram representation of the values.

Each p will report, after some time δ , its private value v, as defined at the beginning of the process. The method of privately reporting consists of encoding the value v in a bit-array B of length equal to the number of bins, where only the bit corresponding to the range that contains v is set to 1, and the other positions are set to 0. This procedure is called Unary Encoding (UE) [Wang et al. 2017].

After encoding v into B, the next step is to bitwise perturb B in a differentially private manner. Let B' be the differentially private version of B. The process of obtaining B' consists in keeping the bit value of B with a probability p and changing it with a probability q = 1-p. In order for this process to achieve differential privacy it is necessary that $\frac{P[B|v1]}{P[B|v2]} \le e^{\varepsilon_i}$. As shown in [Wang et al. 2017], $p = \frac{e^{\frac{\varepsilon_i}{2}}}{e^{\frac{\varepsilon_i}{2}}+1}$ and $q = \frac{1}{e^{\frac{\varepsilon_i}{2}}+1}$ are sufficient for this property to happen.

Notice that, as we want to send a number k of reports privately, we cannot consume all the privacy budget ε in a single report. Therefore, for each single private report, we will use $\varepsilon_i = \frac{\varepsilon}{k}$. This strategy will guarantee that, if continuously reported, any window of k consecutive reports is ε -differentially private. The guarantee comes from the sequential composition property of differential privacy [McSherry 2009].

With B' adequately generated, the Privacy Gateway can finally send the differentially private version of v to the SP, which will then gather it with the reports from the other N houses to obtain information from it. The process of obtaining information from the differentially private reports consists in constructing the histogram Hist. Each bin i of Hist is obtained with the summation of the i^{th} element from all reports, $Hist[i] = \sum_{j=1}^{N} B_j[i]$. It is important to remember that each Hist[i] contains not only the reports that were truly reported for the i^{th} bin, but also some noise added by the differentially private mechanism.

Therefore, the next step is to get an unbiased estimation for Hist, which we will call Unb_Hist . To calculate Unb_Hist , we need to get rid of the noise added for each bin, which can be done in the following way: $Unb_Hist[i] = \frac{Hist[i]-Nq}{p-q}$, where p is the probability of keeping the bit value and q is the probability of inverting a bit used in the process of creating the private reports. N is the number of reports. For the sake of space, we let the reader refers to [Wang et al. 2017] for the proof that this yields an unbiased estimation. Notice that if we have an unpopular bin, i.e., the number of reports (Hist[i]) is small, our unbiased estimation can be negative. When this happens, the considered value for $Unb_Hist[i]$ will be zero.

As the SP have selected N different participants to send k reports spaced by a δ time, the process of calculating $Unb_Hist[i]$ could be performed in one of two ways: (i) using each set of size N separately, which keeps the notion of time and (ii) using the union of all data with size k * N. This strategy may yield more accurate frequencies for the generated histogram since there is a more significant number of reports, but the notion of time is lost. In our evaluation, we have opted to use strategy (i), since we believe the time attribute is essential in the context of IoT data and Smart Cities.

4. Evaluation

For the experimental evaluation, we have used real sensor data that consists of energy consumption readings from 5,567 London households generated between 2011 and 2014

as part of the Low Carbon London project. There are 167 million rows. We have used the attribute "KWH/hh (per half hour)."

In order to better evaluate our proposal, we have sampled rows from the data set to simulate a fixed number N of houses. The number of samples k to be sent from each house was fixed in 10, thus, the privacy budget ε_i used for a single report is equal to $\frac{\varepsilon}{k} = \frac{\varepsilon}{10}$. We have tested the following values for ε : 1, 2, 3, 4, 10. Therefore, the values used for ε_i were: 0.1, 0.2, 0.3, 0.4, 1. Notice that the choice of using a single value for k does not have a significant impact on the results since it has a direct influence on ε , which we have tested for different values.

The varying values for ε and N give us a better understanding of how these two variables impact the utility, as can be seen in Figure 2. To measure the utility we have evaluated the histogram intersection, given by $\frac{\sum_{i=1}^{nb} min(Ori_Hist[i],Unb_Hist[i])}{\sum_{i=1}^{nb} Unb_Hist[i]}$, where nb is the number of bins, Unb_Hist is the histogram generated by our strategy and Ori_Hist is the original histogram. The histogram intersection measures how similar does our proposal generate the histogram compared with the histogram of the original data. The number of bins used in the experiments was 100. The maximum and minimum values in the data set are 0.0 and 10.76, respectively, and each bin used has equal width.

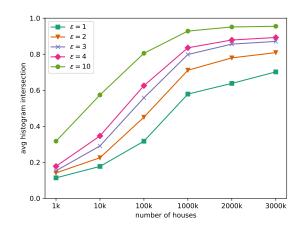


Figure 2. Average histogram intersection by number of houses. Varying ε .

Figure 2 shows, for each value of ε , the average histogram intersection of 10 different execution of our proposal which simulates the process of each house reporting k = 10 times and, after each set of reports (from N houses), the Service Provider calculates the unbiased histogram. It is possible to observe that for a small number of houses, the solution outputs a low histogram intersection, but as N grows, we get more accurate results. The reason why this happens is that when there are few reports, we cannot cancel out the noise added by the differential privacy mechanism. Remember that, the higher the ε , the weaker the privacy guarantee.

5. Conclusion

In this paper, we have proposed a practical solution for estimating the frequency of values issued from houses' IoT devices in a differentially private manner. It allows an energy

provider to collect house consumption for analytics and still provides privacy for individuals living in the house. Data utility depends on the number of houses and the available privacy budget. For one million houses, our preliminary results have reached around 80% of data utility with a privacy budget of 3 to 4, which is very reasonable. For the next steps, it might be valuable to adapt the strategy to work for consecutive windows. Notice that this is not a trivial problem and could demand sophisticated adaptations in order to be solved.

Acknowledgments

This research was supported by FUNCAP and LSBD/UFC.

References

- Ács, G. and Castelluccia, C. (2011). I have a dream!(differentially private smart metering). In *International Workshop on Information Hiding*, pages 118–132. Springer.
- Cao, Y. and Yoshikawa, M. (2015). Differentially private real-time data release over infinite trajectory streams. In 2015 16th IEEE International Conference on Mobile Data Management, volume 2, pages 68–73. IEEE.
- Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer.
- Dwork, C. and Roth, A. (2014). The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3–4):211–407.
- Erlingsson, Ú., Pihur, V., and Korolova, A. (2014). Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, pages 1054–1067. ACM.
- Leal, B. C., Vidal, I. C., Brito, F. T., Nobre, J. S., and Machado, J. C. (2018). δ-doca: Achieving privacy in data streams. In *Data Privacy Management, Cryptocurrencies* and Blockchain Technology, pages 279–295. Springer.
- McSherry, F. D. (2009). Privacy integrated queries: an extensible platform for privacypreserving data analysis. In *Proceedings of the 2009 ACM SIGMOD International Conference on Management of data*, pages 19–30. ACM.
- Molina-Markham, A., Shenoy, P., Fu, K., Cecchet, E., and Irwin, D. (2010). Private memoirs of a smart meter. In *Proceedings of the 2nd ACM workshop on embedded sensing systems for energy-efficiency in building*, pages 61–66. ACM.
- Shi, W., Cao, J., Zhang, Q., Li, Y., and Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5):637–646.
- UK Power Networks (2015). SmartMeter Energy Consumption Data in London Households. https://data.london.gov.uk/dataset/ smartmeter-energy-use-data-in-london-households. Accessed: 2019-06-28.
- Wang, T., Blocki, J., Li, N., and Jha, S. (2017). Locally differentially private protocols for frequency estimation. In 26th {USENIX} Security Symposium ({USENIX} Security 17), pages 729–745.