Protecting against Physical Resource Monitoring

Gergely Acs, Claude Castelluccia, William Lecat
INRIA Rhone Alpes, Montbonnot, France
{gergely.acs, claude.castelluccia, william.lecat}@inrialpes.fr

ABSTRACT

This paper considers the problem of resource monitoring. We consider the scenario where an adversary is physically monitoring on the resource access, such as the electricity line or gas pipeline, of a user in order to learn private information about his victim. Recent works, in the context of smart metering, have shown that a motivated adversary can basically profile a user or a family solely from his electricity traces. However, these works only consider the case of a semi-honest-but-non-intrusive adversary that is only trying to learn information from the consumption reports sent by the user.

This paper, instead, considers the much more challenging case of a *intrusive semi-honest* adversary, i.e. a semi-honest adversary that is in addition physically monitoring the resource by modifying the distribution network. We aim at answering to the following question: is it possible to design a resource distribution scheme that prevents resource monitoring and provides strong protection? This paper proposes and analyzes several possible solutions. The proposed solutions provide different privacy bounds and performance results.

1. INTRODUCTION

Communication wiretapping or eavesdropping is often referring to telephone and Internet conversation monitoring by a third party. There exist different wiretapping or monitoring techniques, but most of them consist in inserting listening devices in the network. In this paper, we assume that the adversary is not interested in listening to the communications of a victim, but instead is interested in his resource, such as electricity or gas, consumption. We define this attack as "resource monitoring". It can typically be performed by inserting a meter on the pipe that is used to deliver the resource to the victim.

Communication wiretapping has been studied for several years. The adversary's objectives might be to listen to the content of the communication, i.e. emails, phone communication, or to have access to the communication patterns (traffic analysis). It was shown that by simply analyzing traffic, the adversary can gain a lot of information about the communications and the communicating parties [22].

In resource monitoring, the adversary is not interested in the content (he knows the nature of the resource he is monitoring), but rather in the consumption usage (for example the electricity usage). Several studies, in the context of smart metering [18, 14], have shown that a motivated adversary can profile a user or a family solely from his electricity traces. Extracting complex usage patterns of appliances from the raw consumption profile (e.g., using NALM [13] or simple off-the-shelf statistical tools [18]), one can infer detailed information about household activity (e.g, how many people are in home and what they are doing at a given time). This extracted information can be used to profile and monitor users for various purposes, creating serious privacy risks and concerns.

Resource monitoring is in fact quite similar to performing traffic analysis in the communication wiretapping case. However, existing prevention techniques, such as encryption and traditional traffic shaping, are not viable nor practical in the context of resource monitoring due to the nature of the considered resource. For instance, in communication systems, applying padding means sending extra traffic in order to hide the real traffic pattern. This dummy traffic is then filtered out at the receiver side. By contrast, in resource networks, applying dummy padding would consists of consuming extra resource (to hide the actual usage) which would have to be charged to the user. Most users would oppose to it.

This paper is the first work, to our knowledge, that considers the challenging problem of resource monitoring. It first describes formally the problem. It then proposes several possible "anti-monitoring" resource solutions that use the concept of smart (random) buffering and provide strong privacy guarantees under the differential privacy model. The security of the different solutions are formally analysed and their performance compared.

We acknowledge that the results presented in this paper are still preliminary and more work is needed.

However, we believe that with the advent of smart grid and smart metering of resource, the problem of resource monitoring is going to become a very important issue in the coming years and more research is needed on this topic.

2. RELATED WORK

Smart metering: Smart meters allow the utility provider to monitor, almost in real-time, consumption and possibly adjust generation and prices according to the demand. Several papers addressed the privacy problems of smart meters in the recent past [10, 18, 4, 5, 6, 20, 14, 11]. In [4, 5], the authors discuss the different security aspects of smart metering and the conflicting interests among stakeholders. The privacy of billing is considered in [20, 18]. Another line of works [3, 6, 11, 17] consider the problem of monitoring the aggregate consumption of multiple clients without leaking private information of any individual's consumption. All these works assumed a semi-honest-but-non-intrusive adversary who cannot install any extra hardware in the distribution network to collect more information about users and only uses the measurements provided by the users. In this paper, we consider the strongest semi-honest-but-intrusive adversary who can invade the distribution network and modify it to gather more information about users.

Differential privacy: The notion of differential privacy was first proposed in [9] (for a survey of recent results refer to [8]). Differential privacy says that releasing data using a differentially private algorithm will not increase the adversary's chance to infer any information about any users in the dataset. The main advantage of differential privacy over other privacy models is that it does not specify the prior knowledge of the adversary and provides rigorous privacy guarantee if each users' data is statistically independent [16].

Anti-wiretapping and traffic shaping: The usual techniques of defeating wiretapping (traffic analysis) in communication networks consists of traffic encryption and channel masking. Channel masking includes traffic shaping like padding, traffic delaying, and re-ordering of packet sequences [7, 15, 12, 19, 2]. Another research deals with telephone wiretap and dialed number recording [21]. Most of these works are concentrated on information systems where bitstreams can be easily transformed in order to hide their information content against a malicious eavesdropper. Moreover, these transformations (like padding) mostly causes minimal costs to benign users which is not the case for resource monitoring, where asking more resource than the real demand can cause significant extra costs to the clients.

3. MODEL

Suppose provider P sells some resource, such as electricity, to client C where the resource can be measured by an appropriate metric in units. The price of the resource is described by a pricing function: the price that the client pays for x units to the provider is denoted by $f(x)$, where $f : \mathbb{R} \to \mathbb{R}$.

We assume that the distribution process between P and C is periodic and happens in consecutive slots. In each slot, C consumes x units and pays $f(x)$ to P. We assume that P and C are fully cooperative, i.e., P always provides C with the requested amount and it always accepts the amount given by P.

C has an overall demand of $\sum_{j=1}^{N} X_i^j \stackrel{\mathsf{def}}{=} \mathbf{X}_i$ units in slot i where X_i^j denotes the jth sub-consumption (or sub-demand) in slot i. These sub-demands can correspond to the demands of other customers, if for example C resells the resource, or the consumption demand of C itself (e.g., if C watches TV then the electrical consumption of the TV during the watching period corresponds to a single sub-demand). The demand (or consumption) profile of C over n slots is defined by $(\mathbf{X}_1, \ldots, \mathbf{X}_n)$.

3.1 Adversary model

Several papers addressed the privacy problems of smart meters in the recent past [10, 18, 4, 5, 6, 20, 14, 11]. All of these papers, consider what is called in the literature, *semi-honest adversaries*. A semi-honest adversary faithfully follows all protocol specifications and does not misrepresent any information related to their inputs, e.g., size and content. However, during or after protocol execution, he passively attempts to infer additional information about the victim from the collected data.

This paper considers a stronger adversary model. We consider a *semi-honest intrusive* adversary, i.e. an adversary that faithfully follows all protocols, but that can, in addition, invade the distribution network to gather more information about clients. In other words, we are assuming that the adversary can monitor the electricity or gas consumption of the clients by installing meters on the power line or gas pipeline that is outside of the client's control (like outside from his household).

In general, the objective of the adversary is to infer detailed information about the victim's activity by monitoring her consumption. Our goal is to protect against a strong adversary that might know the total consumption $(\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_n)$ and all the sub-demands composing it, except from one i.e. X_i^j. We say that a scheme is secure if the adversary is unable to recover X_i^j whatever external information it has.

3.2 Privacy model

Our adversary model is included in the differential privacy model, which was first proposed in [9]. Differential privacy guarantees that the client's privacy should

not be threatened substantially more regardless what external knowledge the provider has.

Definition 1 ((ε, δ)-differential privacy) *An algorithm $\mathcal{A}$ is (ε, δ)-differential private, if for all data sets D_1 and D_2, where D_1 and D_2 differ in at most a single user, and for all subsets of possible answers $S \subseteq Range(\mathcal{A})$,*

$$P(\mathcal{A}(D_1) \in S) \leq e^{\varepsilon} \cdot P(\mathcal{A}(D_2) \in S) + \delta$$

Differential private algorithms produce indistinguishable outputs for similar inputs (more precisely, differing by a single entry), and thus, the modification of any single user's data in the dataset (including its removal or addition) changes the probability of any output only up to a multiplicative factor e^{ε} plus an additive constant δ. The parameters ε and δ allow us to control the level of privacy. Lower values of ε and δ imply stronger privacy, as they restrict further the influence of a user's data on the output.

The following theorem suggests a simple technique to differential privately release an aggregate by adding a random noise to the that, where the noise distribution is carefully calibrated to the global sensitivity of the aggregate. The global sensitivity of an aggregate is the maximum variation of its value when we change one of its data component.

Theorem 1 ([9]) *For all $h : \mathbb{D} \to \mathbb{R}^r$, the following mechanism $\mathcal{A}$ is ε-differential private (with δ=0): $\mathcal{A}(D) = h(D) + \mathcal{L}(S(h)/\varepsilon)$, where $\mathcal{L}(S(h)/\varepsilon)$ is an independently generated random variable following the Laplace distribution and $S(h)$ denotes the global sensitivity of h[1].*

3.3 Resource distribution model

Suppose a client has an overall demand of $\mathbf{X}_i$ units and consumes m_i units in slot i (note that $\mathbf{X}_i$ and m_i may be different). Let $\mathcal{D}_n$ be the random variable describing the difference $\sum_{i=1}^{n}(m_i - \mathbf{X}_i)$ after n slots. A distribution scheme

- has *satisfiability* φ over n slot, if $P(\mathcal{D}_i \geq 0) \geq \varphi$ for all $1 \leq i \leq n$,
- has *expected deficit* γ over n slot, i.e., $\mathbb{E}(f(\mathcal{D}_n)) = \gamma$,
- has *maximum deficit* θ with confidence μ over n slot, if $P(f(\mathcal{D}_i) \leq \theta) \geq \mu$ for all $1 \leq i \leq n$,
- is (ε, δ)-private over n slot, if it satisfies Definition 1 with parameters (ε, δ) over n slot.

[1]Formally, let $h : \mathbb{D} \to \mathbb{R}^r$, then the global sensitivity of h is $S(h) = \max ||h(D_1) - h(D_2)||_1$, where D_1 and D_2 differ in a single entry and $|| \cdot ||_1$ denotes the L_1 distance.

Intuitively, satisfiability measures the guarantee that the client can satisfy all its consumption demands over n slots, while expected deficit measures the average loss of the client (i.e., the expected value of the extra price that it pays beyond the price of its real demand). Similarly, the maximum deficit says that the maximum loss of the client in any slot over n will not exceed a certain bound with a given confidence. Observe that satisfiability, the expected and maximum deficit describe the utility of a distribution scheme, while (ε, δ) characterizes its privacy in the differential privacy model.

4. TOWARD SMART BUFFERING

In current smart metering systems, a client that needs $\mathbf{X}_i$ in slot i will directly get it from the resource provider. Although this approach has always 0 expected deficit, 0 maximum deficit and provides complete satisfiability over any n slots, it has no privacy guarantee; in fact, adding a new sub-demand differing from 0 will always result in $\delta = 1$ in Definition 1. For sake of simplicity, we assume in the sequel that the pricing function is defined as $f(x) = x$ for all x.

In this section, we present three new resource distribution schemes and analyze their privacy as well as their utility in the model described previously. These three schemes use buffers that sit between the client and provider and hide the consumed resource to the provider. A buffer is actually implemented as a battery that stores the resource.

The main idea of using buffers is to decorrelate the actual resource consumption from the resource obtained from the utility. When a customer needs energy, it gets part of it from its buffer and part of it from the provider. Similarly, the energy obtained from the provider is partly consumed by the user's devices and partly used to load the buffer.

The rest of this section presents three different buffer-based schemes. These schemes differ in how the buffer is managed (i.e. charged or discharged). The first proposed scheme has perfect privacy, but it is not resilient to profile dynamics and can result in low utility depending on the consumption profile. The last two approaches provide better privacy-utility trade-offs. The performance of these algorithms are compared in Section 5.

4.1 Constant rate buffering

A straightforward approach to improve privacy is that the client maintains a buffer and, in each slot, consumes the same amount, denoted by c. If $\mathbf{X}_i < c$, then the client puts the difference $c - \mathbf{X}_i$ into the buffer. If $\mathbf{X}_i > c$, then it takes $\mathbf{X}_i - c$ units out of her buffer to satisfy all sub-demands.

If c is large enough, then we obtain a perfectly private scheme where $\varepsilon = \delta = 0$ as the client will always

consume the same amount independently from its real demand. However, if $\sum_{i=1}^{n} \mathbf{X}_i > n \cdot c$, then the total consumed amount is not enough to satisfy all sub-demands which means that it has 0 satisfiability. The other problem is that the deficit highly depends on *all* sub-consumptions, and in general, on the client's consumption profile. Moreover, this profile must be known a priori in order to correctly calibrate c which makes it an impractical approach.

4.2 Smart buffering with symmetric Laplace noise

The intuition behind smart (random) buffering is as follows. The client perturbs its real consumption $\mathbf{X}_t$ and consumes $\hat{\mathbf{X}}_t = \mathbf{X}_t + \mathsf{s}$ units in each slot t, where s is a random value drawn from some distribution with mean 0. If $\mathsf{s} > 0$, then the client puts the extra s units into the buffer. If $\mathsf{s} < 0$ then it gets the missing κ units from the buffer. It is easy to see that this approach is more resilient to profile dynamics as the expected deficit is 0 independently from the demand profile of the client (i.e., $\mathbb{E}(\sum_{i=1}^{n} \kappa_t) = 0$ for all n where κ_t is the random variable describing s in slot t). Moreover, if the distribution of κ_t is chosen carefully, then one can minimize the maximum deficit as well as the information that is leaked by $\hat{\mathbf{X}}_t$: higher/lower noise variance results in stronger/weaker privacy and lower/higher maximum deficit with higher/lower satisfiability.

4.2.1 Operation

The algorithm works as follows in a slot t. First, the client draws a random sample s_t from an appropriate distribution and computes $\hat{\mathbf{X}}_t = \mathbf{X}_t + \mathsf{s}_t$.

- If $\mathsf{s}_t \geq 0$, then the client consumes $\hat{\mathbf{X}}_t$ units. Out of these, it serves its sub-demands using $\mathbf{X}_t$ units, and stores the rest s_t units into the buffer.
- If $\mathsf{s}_t < 0$ (i.e., $\mathsf{s}_t > \mathbf{X}_t$), then the client takes the following actions depending on the value of $\hat{\mathbf{X}}_t$:
  - if $\hat{\mathbf{X}}_t \geq 0$, it consumes $\hat{\mathbf{X}}_t$ units and gets s_t units from its buffer to serve its real demand $\mathbf{X}_t$.
  - if $\hat{\mathbf{X}}_t < 0$, it gets $|\mathsf{s}_t|$ units from the buffer, from which it uses $\mathbf{X}_t$ units to serve its demand and gives $|\mathbf{X}_t + \mathsf{s}_t|$ units back to the provider.

Note that the client may need to give back resource to the provider in the last step when $\hat{\mathbf{X}}_t$ is negative. In other words, the model defined in Section 3, which assumed that in each slot C consumes x units and pays $f(x)$ to P, must be extended. It is now assumed that in each slot, either C consumes x units and pays $f(x)$ to P or C gives y units back to P (where $x = \hat{\mathbf{X}}_t$ and $y = |\mathbf{X}_t + \mathsf{s}_t|$ in the current scheme).

Algorithm 1 SmartBufferLaplace

Require: M - buffer size (capacity), ν - initial buffer size, λ - scale parameter of the noise

1: $\mathsf{L}_1 := \nu$ // L_t is the buffer level at slot t
2: **for all** t **do**
3: $\quad \mathsf{s}_t \leftarrow \mathcal{L}(\lambda)$
4: $\quad$ **push** $\max(0, X_t + \mathsf{s}_t)$
5: $\quad$ **pop** $\min(0, X_t + \mathsf{s}_t)$
6: $\quad \mathsf{L}_{t+1} := \mathsf{L}_t + \tau$
7: **end for**

We believe that this scenario will be realistic in electrical distribution networks in the near future with the rise of Smart Grid technologies [1]. In Smart Grids, clients will also be able to produce electrical energy (e.g., by using solar panels or plug-in hybrid electric vehicles) and are envisioned to sell the excess to other clients or the provider itself.

4.2.2 Analysis

Privacy: Let κ_i be the random variable describing s_i in slot i. Following Theorem 1, if κ_i follows a symmetric Laplace distribution with scale parameter λ, where $\lambda = \max_j \sum_{i=1}^{n} X_i^j / \varepsilon$ (i.e., ε is calibrated to the maximum sub-consumption over n slots) and each κ_i is drawn independently, then the client obtains ε-differential privacy ($\delta = 0$) over n slots.

Utility: At first sight, the expected deficit $\mathbb{E}(\sum_{i=1}^{n} \kappa_i)$ is 0 independently from the client's demand profile, as all κ_i have a mean of 0. However, in that case, the satisfiability is 0.5 as in the very first slot when the buffer is empty, it has probability 0.5 that $\mathsf{s}_1 < 0$. Hence, before computing s_1, the client should ask for ν units to the provider in order to initialize the buffer. The value of ν should be calibrated to the maximum client demand (i.e., to $\max_i \mathbf{X}_i$). After all, the expected deficit becomes ν.

In order to compute the exact satisfiability and the maximum deficit, we need the following lemma.

Lemma 1 *Let κ_i are i.i.d random variables having Laplace distribution with pdf $f(x, \lambda) = \frac{1}{2\lambda} e^{-\frac{|x|}{\lambda}}$. Then, $P\left(\sum_{i=1}^{n} \kappa_i > c\right) \leq e^{-\frac{c^2}{8n\lambda^2}}$, if $0 < c < 2\sqrt{2}n\lambda$.*

The proof (as of all Theorems in this paper) is detailed in the Appendix. Using this lemma, we obtain the following results.

Theorem 2 *Let κ_i are i.i.d random variables having Laplace distribution with pdf $f(x, \lambda) = \frac{1}{2\lambda} e^{-\frac{|x|}{\lambda}}$. Then,* SmartBufferLaplace *has satisfiability $1 - e^{-\frac{\nu^2}{8n\lambda^2}}$ if $0 < \nu < 2\sqrt{2}n\lambda$, expected deficit ν, where ν is the initial buffer level, and maximum deficit c with confidence $1 - e^{-\frac{(c-\nu)^2}{8n\lambda^2}}$ if $0 < c - \nu < 2\sqrt{2}n\lambda$*

The proof is straightforward using Lemma 1 and the fact that $\sum_{i=1}^{n} \kappa_i$ is a symmetric random variable, which means that $P(\nu + \sum_{i=1}^{n} \kappa_i < 0) = P(\sum_{i=1}^{n} \kappa_i > \nu)$, if κ_i follows a Laplace distribution.

Observe that both satisfiability and maximum deficit depend on n, as the probability that $\sum_{i=1}^{n} \kappa_i$ exceeds a certain threshold becomes higher by increasing n. This means that the privacy as well as the utility deteriorates over time. Intuitively, the confidence of maximum deficit M is the probability that the buffer level remains below the buffer size (i.e, there is no *buffer overflow*), and satisfiability is the probability that the buffer level will not fall below 0 (i.e., there is no *buffer underflow*) over n slots.

Note that this random buffering approach is more resilient to profile dynamics than the constant rate buffering scheme. In particular, we only need to know the maximum sub-consumption (i.e., $\max_j \sum_{i=1}^{n} X_i^j$) to calibrate the Laplace noise to a given ε, and the maximum client demand (i.e., $\max_i \mathbf{X}_i$) to calibrate the initial buffer level ν to the desired satisfiability value. Both are easier to estimate than the whole profile which is needed for constant rate buffering.

4.3 Smart buffering with truncated geometric noise

This approach improves the performance of the previous scheme while keeping its resiliency against profile dynamics: it guarantees that the buffer will never run out (the satisfiability is 1), and it also never exceeds a certain threshold M (i.e., the maximum deficit is bounded almost surely). The idea is that instead of using noise with infinite domain we truncate the noise into a finite interval such that the buffer level will always be within $[0, M]$. In particular, if the buffer has a size of M and the buffer level is L_i in slot i, then the noise s_i is sampled from a distribution that can only take values from $[-\mathsf{L}_i, M - \mathsf{L}_i]$. This ensures that $0 \leq \mathsf{L}_i + \mathsf{s}_i \leq M$ for all i.

Before describing the operation, we define the truncated distribution that we use, which is the discrete approximation of the Laplace distribution. As we will later see, assuming discrete noise makes the analysis easier without losing generality: in most practical scenarios, continuous values are approximated using integers.

Definition 2 (Geometric Distribution) *Let $\alpha > 1$ and $\beta \in \mathbb{Z}$. The probability mass function of the symmetric geometric distribution centred at β is $\frac{1-\alpha}{1+\alpha} \cdot \alpha^{-|x-\beta|}$ where x takes integer values.*

The symmetric geometric distribution corresponds to the discrete version of the Laplace distribution, where $\alpha = e^{\frac{1}{\lambda}}$. It is easy to check that using symmetric geometric distribution in Theoreom 1 instead of the continuous Laplace, we will have a differentially private algorithm that outputs integer values.

Let $\mathcal{G}(\alpha, \beta)$ be a random variable having symmetric geometric distribution. Its truncated counterpart denoted by $\widetilde{\mathcal{G}}(\alpha, \beta, x_1, x_2)$ has a conditional probability distribution and it is defined as $P(\widetilde{\mathcal{G}}(\alpha, \beta, x_1, x_2) = k) = \frac{P(\mathcal{G}(\alpha,\beta)=k)}{\sum_{i=x_1}^{x_2} P(\mathcal{G}(\alpha,\beta)=i)}$ if $x_1 \leq k \leq x_2$ and 0 otherwise, where $[x_1, x_2]$ $(x_1, x_2 \in \mathbb{Z})$ is the truncation interval.

4.3.1 Operation

Suppose that the buffer size is M, where M is an even integer. The buffer level is initialized to $M/2$. Then, in each slot t, the client picks up a random sample τ from $\widetilde{\mathcal{G}}(\alpha, \mathsf{L}_t, 0, M)$, where L_t denotes the buffer level in slot t, and computes the noise as $\mathsf{s}_t = \tau - \mathsf{L}_t$. Therefore, $\mathsf{L}_t + \mathsf{s}_t$ will always be in $[0, M]$. Afterwards, it performs the same as SmartBufferLaplace described in Section 4.2.

Algorithm 2 SmartBufferGeometric

Require: M - buffer size (capacity), α - scale parameter of the noise

1: $\mathsf{L}_1 := M/2$
2: **for all** t **do**
3: $\quad \tau \leftarrow \widetilde{\mathcal{G}}(\alpha, \mathsf{L}_t, 0, M)$
4: $\quad \mathsf{s}_t := \tau - \mathsf{L}_t$
5: $\quad$ **push** $\max(0, X_t + \mathsf{s}_t)$
6: $\quad$ **pop** $\min(0, X_t + \mathsf{s}_t)$
7: $\quad \mathsf{L}_{t+1} := \tau$
8: **end for**

4.3.2 Analysis

Observe that the noise distribution may be different in each slot and the noise samples are not independent in consecutive slots; the distribution of the noise depends on the current buffer level which is shaped by the noise sample of the previous slot. In particular, the buffer level L_t can be described by a discrete time-homogeneous Markov chain. Indeed, the distribution of the buffer level in one slot only depends on the buffer level of the previous slot. The transition matrix of this Markov chain is $\mathbf{M}$ where $\mathbf{M}_{i,j} = P(\widetilde{\mathcal{G}}(\alpha, j, 0, M) = i)$ $(i, j \in [0, M])$[2] is the probability that the buffer level changes from j to i after a single slot.

Privacy: Although truncating the noise helps to ensure that the buffer level will always be within $[0, M]$, it also deteriorates privacy. To illustrate this, consider a demand $\mathbf{X}_1$ in the first slot and another demand $\mathbf{X}'_1$ which only differs in a single sub-demand. The added noise in both cases are drawn from the same distribution which

[2]To simplify the notation, we directly address an element of the transition matrix as well as the distribution vectors with the corresponding buffer level.

is $\widetilde{\mathcal{G}}(\alpha, M/2, 0, M)$. However, the possible set of outputs are different. Indeed, w.l.o.g, suppose that $\mathbf{X}_1 < \mathbf{X}'_1$. Then, $O = \mathbf{X}'_1 + M/2$ is a possible output with $\mathbf{X}'_1$ but not with $\mathbf{X}_1$, and hence, $P(\mathbf{X}_1 + \widetilde{\mathcal{G}}(\alpha, M/2, 0, M) = O) = 0$ and $P(\mathbf{X}'_1 + \widetilde{\mathcal{G}}(\alpha, M/2, 0, M) = O) > 0$. If O appears in the output, the two cases are distinguishable and we will have a privacy breach. The probability of such outputs is measured by δ in the (ε, δ)-differential privacy model.

Theorem 3 *Let π_i denote the distribution of the buffer levels in slot i and $\Delta_t = \sum_{i=1}^{t} \max_j X_i^j$.* SmartBufferGeometric *is (ε, δ)-private over n slots, where*

- $\varepsilon = \ln\left(\alpha^{\Delta_n} \cdot \prod_{i=1}^{n} \frac{\alpha^{\frac{M}{2}} - ch^{\alpha}(\frac{M}{2} - \Delta_{i-1})}{sh^{\alpha}(\frac{M}{2})}\right)$
- $\delta = 1 - \prod_{k=1}^{n} \left(1 - \delta^{(k)}\right)$

where $\delta^{(k)} = \sum_{i=0}^{\Delta_k} [\pi_k]_i$, and $sh^{\alpha}(x)$ and $ch^{\alpha}(x)$ are defined as $\frac{\alpha^x - \alpha^{-x}}{2}$ and $\frac{\alpha^x + \alpha^{-x}}{2}$, resp. ($\alpha > 1$).

Utility: As $0 \leq \mathsf{L}_i + \mathsf{s}_i \leq M$ holds for all i, it is straightforward to compute the utility.

Theorem 4 SmartBufferGeometric *has satisfiability of* 1*, expected deficit $\frac{M}{2}$, and maximum deficit M with confidence* 1

5. PERFORMANCE COMPARISON

Figure 1 shows how the probability that the buffer level remains below the buffer size (i.e., the confidence of maximum deficit M where M is the buffer size) changes depending on the buffer size using smart buffering with Laplace noise. The buffer is initialized to the half of the buffer size ($\nu = M/2$) which means that buffer overflow and underflow (i.e., the confidence of maximum deficit M and satisfiability) for the same M are identical. In addition, ε is set to 0.1 ($\delta = 0$), and the maximum value of any sub-demand in a single slot is 1 (i.e., $\Delta_t = t$). The buffer overflow and underflow becomes more likely for a larger number of time slots as the probability that the sum of random samples exceeds a fixed threshold increases by taking the sum of more samples. We can conclude that the buffer size needs to be 500 times larger (if $n > 20$) than the maximum sub-consumption in order to obtain reasonably low probability of buffer overflow/underflow (< 0.05).

Figure 2 shows how the privacy changes depending on the number of slots in smart buffering with truncated geometric noise. Figure 2(a) plots ε in function of the buffer size and the number of slots with $\alpha = 1.001$, while 2(c) depicts ε in function of the distribution parameter α and the buffer size. Apparently, the buffer size does not really affect ε: increasing the number of slots, ε still has acceptable value (i.e., < 0.1) with more than

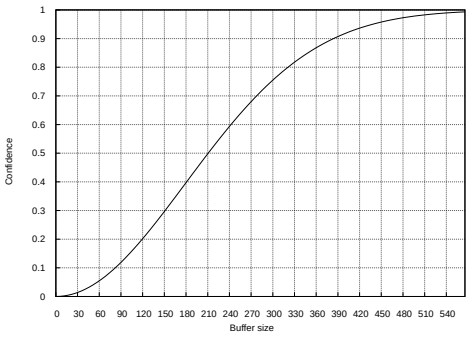

Figure 1: The confidence of maximum deficit M in smart buffering with Laplace noise ($\varepsilon = 0.1$) depending on M ($n = 20$). Note that $M/2 < 2\sqrt{2}n(1/\varepsilon)$.

20 slots. However, this is not the case for δ: Figure 2(b) shows that δ rapidly increases by increasing the number of slots. Namely, if the buffer size is 300 and $n > 20$, delta exceeds 0.5 (i.e., the probability that there is a privacy breach over 20 slots is 0.5). Intuitively, privacy breach occurs when the buffer level falls below a certain threshold (i.e., when such output appears that cannot appear after adding a new user). The probability of this increases by summing more noise samples. Also note that increasing the buffer size improves δ. However, the buffer size needs to be much larger than 300 in order to have reasonably low δ for $n > 20$ slots.

Finally, Figure 3 compares smart buffering with constant rate buffering (CRB). We assumed that the client can compute the maximum sub-demand in each slot (which is 1 now) and also knows the maximum number of sub-demands, denoted by x, which can occur in a single slot. However, it cannot predict how many and what sub-demands will appear in each slot exactly.[3] Supposing this case, in CRB, the constant rate c needs to be $x \cdot \max_{i,j} X_i^j$ (i.e., the client prepares for the worst case and always consumes its maximum sum of sub-demand that it can have in a single slot).

Figure 3 plots the buffer size depending on x for $n = 10$ time slots. In smart buffering, the buffer size can be calculated from the desired privacy bound (ε and δ). While, in CRB, it is the product of n and $x \cdot \max_{i,j} X_i^j$, which is $10 \cdot x$ in the current example ($\max_{i,j} X_i^j = 1$). This is because, in the worst case, when there is no any demand over the n slots, all the resource asked to the provider has to be stored in the buffer.

As both smart buffering techniques are independent

[3]Clearly, it is a pessimistic assumption since the client usually has more a priori knowledge about its demand profile (e.g., the consumption of gas or electricity is lower at night than in daylight in most households) though this knowledge is application dependant. The modelling of such knowledge for specific applications (like gas or electricity distribution) belongs to future work.

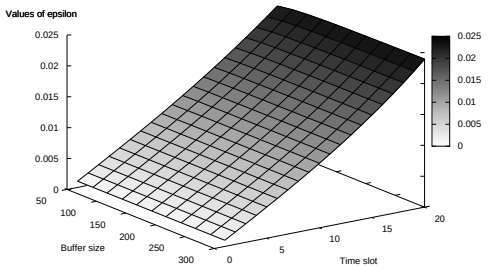

(a) ε depending on the buffer size and the number of slots ($\alpha = 1.001$)

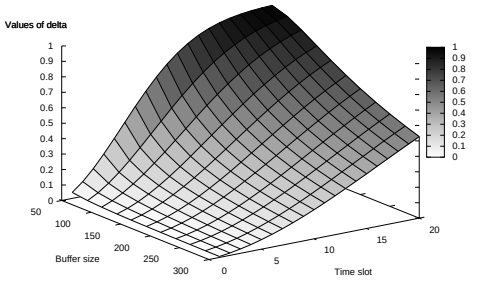

(b) δ depending on the buffer size and the number of slots ($\alpha = 1.001$)

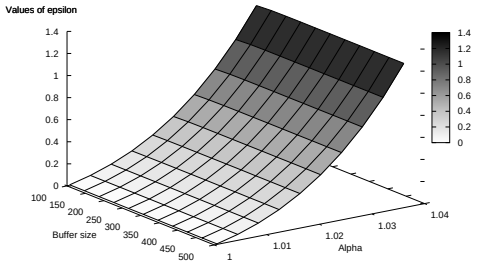

(c) ε depending on the buffer size and α

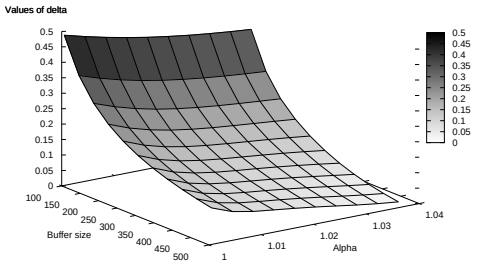

(d) δ depending on the buffer size and α

Figure 2: Smart buffering with truncated geometric noise ($\Delta = 1$).

from the number of sub-demands per slot, their buffer size is indepedent from x. SmartBufferLaplace with $\varepsilon = 0.06$ needs roughly the same size of buffer as SmartBufferGeometric with $\varepsilon = 0.06$ and $\delta = 0.09$. In other words, the guarantee that there is no buffer overflow/underflow comes at the expense of $\delta = 0.09$ (i.e., there is one privacy breach over 10 slots on average).

Figure 3 also shows that CRB is not scalable since the buffer needs to be expanded by having more sub-demands (and time slots). By contrast, the buffer size in smart buffering is independent from the number of sub-demands as the noise used to perturb the client's demand is only calibrated to the maximum sub-demand value and not to n and x. On the other hand, CRB can still have smaller buffer if the number of maximum sub-demands per slot is low: in Figure 3, CRB needs smaller buffer than smart buffering if $x < 50$, and smart buffering (with $\varepsilon = 0.06$ and $\delta = 0.09$) needs smaller buffer than CRB only if $x > 50$. On the other hand, recall that CRB always provides perfect privacy and there cannot be buffer overflow or underflow.

We can conclude that

- CRB is not scalable (the buffer size is a linear function of the number of sub-demands per slot) and can require large buffer if n and x is high. However, it provides perfect privacy and guarantees that there is no buffer overflow and underflow. Furthermore CRB is more practical (i.e., requires smaller buffer) than smart buffering if the number of sub-demands is low.

- SmartBufferLaplace is much more scalable (the buffer size is independent from the number of sub-demands per slot) and allows the client to adjust privacy to the probability of buffer overflow/underflow. However, the privacy deteriorates over time, and meaningful privacy ($\varepsilon < 0.1$) is achievable only for large buffers (500 times larger than the maximum sub-demand) with reasonably low probability of buffer overflow/underflow (< 0.05).

- SmartBufferGeometric is a trade-off solution between CRB and SmartBufferLaplace: Compared to SmartBufferLaplace, it guarantees that there is no buffer overflow and underflow at the expense of some privacy degradation (which is measured by δ) while retaining the scalability and flexibility of SmartBufferLaplace. However, the buffer size is also large (the same as for SmartBufferLaplace with less privacy).

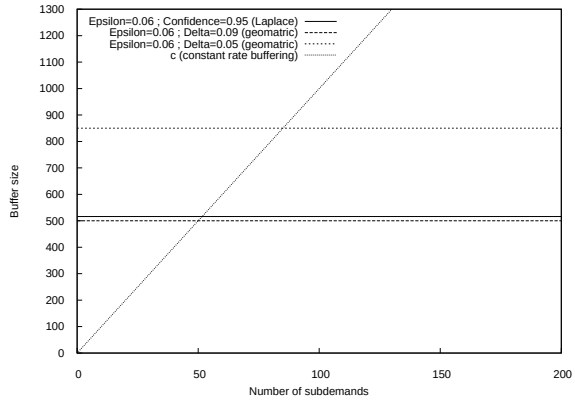

Figure 3: Comparison of different schemes ($n = 10$ slots).

6. CONCLUSION

In this paper, we made the first steps toward a privacy-aware resource distribution scheme and proposed three possible anti-monitoring approaches. The results show that there is no practical solution (yet): constant rate buffering can provide satisfactory performance if the client's profile originates from limited number of sub-demands (e.g., appliances) and/or the client's profile is known a priori. Although our new approaches, called smart buffering, is more scalable to the number of sub-demands and is resilient against profile dynamics, it provides reasonable privacy (or utility) only for larger buffer size.

We stress that our work is still preliminary and we acknowledge that more research is needed in this new area. It still remains an open question whether more efficient schemes exist. Furthermore, this work does not consider the financial cost of the schemes. Resources, such as electricity, gas, are getting more and more expensive. Furthermore, their price tend to vary with time according to the consumption need/peek. These dynamic pricing might have an impact on the cost of the proposed schemes. Designing resource anti-monitoring solutions that are secure and that optimize resource usage remains a challenging and exciting problem.

Acknowledgements

The work presented in this paper was supported in part by the European Commission within the STREP WSAN4CIP project. The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official policies or endorsement of the WSAN4CIP project or the European Commission.

7. REFERENCES

[1] U.S. Department of Energy prepared by Litos Strategic Communication. The Smart Grid: An introduction. `http://www.oe.energy.gov/DocumentsandMedia/DOE_SG_Book_Single_Pages.pdf`.

[2] A. Acquisti, R. Dingledine, and P. Syverson. On the economics of anonymity. In *Financial Cryptography*, 2003.

[3] G. Acs and C. Castelluccia. I have a DREAM! (DiffeRentially privatE smArt Metering. In *Proceedings of Information Hiding Conference*, 2011.

[4] R. Anderson and S. Fuloria. On the security economics of electricity metering. In *Proceedings of the WEIS*, June 2010.

[5] R. Anderson and S. Fuloria. Who controls the off switch? In *Proceedings of the IEEE SmartGridComm*, June 2010.

[6] J.-M. Bohli, C. Sorge, and O. Ugus. A Privacy Model for Smart Metering. In *Proceedings of IEEE ICC*, 2010.

[7] D. Chaum. Untraceable electronic mail, return addresses and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, February 1981.

[8] C. Dwork. Differential Privacy: A Survey of Results. In *In Proceedings of Theory and Applications of Models of Computation (TAMC)*, 2008.

[9] C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating Noise to Sensitivity in Private Data Analysis. In *Proceedings of the 3rd IACR TCC*, 2006.

[10] C. Efthymiou and G. Kalogridis. Smart Grid Privacy via Anonymization of Smart Metering Data. In *Proceedings of IEEE SmartGridComm*, October 2010.

[11] F. D. Garcia and B. Jacobs. Privacy-friendly Energy-metering via Homomorphic Encryption. In *Proceedings of the STM*, 2010.

[12] D. M. Goldschlag, M. G. Reed, and P. F. Syverson. Hiding routing information. In *Proceedings of Information Hiding*, pages 137–150, 1996.

[13] G. Hart. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12):1870–1891, December 1992.

[14] G. Kalogridis, C. Efthymiou, S. Denic, T. A. Lewis, and R. Cepeda. Privacy for Smart Meters: Towards Undetectable Appliance Load Signatures. In *Proceedings of IEEE SmartGridComm*, October 2010.

[15] D. Kesdogan, J. Egner, and R. Büschkes. Stop-and-go-mixes providing probabilistic anonymity in an open system. In *Information Hiding*, pages 83–98, 1998.

[16] D. Kifer and A. Machanavajjhala. No Free Lunch

in Data Privacy. In *to appear in SIGMOD 2011*, 2011.
[17] K. Kursawe, G. Danezis, and M. Kohlweiss. Privacy-friendly Aggregation for the Smart-grid. In *Proceedings of PETS*, 2011.
[18] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin. Private memoirs of a smart meter. In *Proceedings of ACM Buildsys*, 2010.
[19] M. G. Reed, P. F. Syverson, and D. M. Goldschlag. Anonymous connections and onion routing. *IEEE Journal on Selected Areas in Communications*, 16(4):482–494, May 1998.
[20] A. Rial and G. Danezis. Privacy-Preserving Smart Metering. In *Technical Report, MSR-TR-2010-150*. Microsoft Research, 2010.
[21] M. Sherr, E. Cronin, S. Clark, and M. Blaze. Signaling Vulnerabilities in Wiretapping Systems. *IEEE Security & Privacy Magazine*, 3(6):13–25, November 2005.
[22] C. Wright, F. M. L. Ballard, and G. Masson. Language Identification of Encrypted VoIP Traffic: Alejandra y Roberto or Alice and Bob? In *Proceedings of the 16th USENIX Security Symposium*, 2007.

APPENDIX

A. PROOF OF LEMMA 1

PROOF. We derive a Chernoff bound. The moment generating function of κ_i is $\mathbf{E}(e^{t\cdot\kappa_i}) = \frac{1}{1-t^2\lambda^2}$, if $|t| < 1/\lambda$. In addition, $\frac{1}{1-x} \leq 1+2x \leq e^{2x}$ for all $0 < x < \frac{1}{2}$. Hence, $\mathbf{E}(e^{t\cdot\kappa_i}) \leq e^{2t^2\lambda^2}$, if $0 < t < \frac{1}{\sqrt{2}\lambda}$. Then,

$$\begin{aligned} P\left(\sum_{i=1}^{n}\kappa_i > c\right) &= P\left(e^{t\cdot\sum_{i=1}^{n}\kappa_i} > e^{t\cdot c}\right) \\ &\leq \inf_{0<t<\frac{1}{\sqrt{2}\lambda}} e^{-t\cdot c}\prod_{i=1}^{n}\mathbf{E}(e^{t\cdot\kappa_i}) \\ &\leq \inf_{0<t<\frac{1}{\sqrt{2}\lambda}} e^{-t\cdot c+2t^2n\lambda^2} \\ &= e^{-\frac{c^2}{8n\lambda^2}} \end{aligned}$$

where the last equality holds if $c < 2\sqrt{2}n\lambda$.

B. PROOF THEOREM 3

Lemma 2 (Truncated geometric distribution) *Let $\widetilde{\mathcal{G}}(\alpha,\beta,-M/2,M/2)$ denote a random variable having geometric distribution truncated to $\left[-\frac{M}{2},\frac{M}{2}\right]$. Then,*

$$P(\widetilde{\mathcal{G}}(\alpha,\beta,-M/2,M/2)=k) = \frac{\alpha^{-|k-\beta|}}{1+\frac{2}{\alpha^{\frac{M}{2}}}\cdot\frac{\alpha^{\frac{M}{2}}-ch^{\alpha}(\beta)}{\alpha-1}}$$

where $ch^{\alpha}(x)$ is defined as $\frac{\alpha^x+\alpha^{-x}}{2}$ and $\alpha > 1$.

PROOF OF LEMMA 2. We will compute the density function of $\widetilde{\mathcal{G}}(\alpha,\beta)$ in interval $[-M/2, M/2]$:

$$P(\widetilde{\mathcal{G}}(\alpha,\beta)=k) = \frac{P(\mathcal{G}(\alpha,\beta)=k)}{\sum_{i=-M/2}^{M/2}P(\mathcal{G}(\alpha,\beta)=i)} =$$

$$= \frac{\frac{1-\alpha}{1+\alpha}\alpha^{-|k-\beta|}}{\frac{1-\alpha}{1+\alpha}\sum_{i=\frac{-M}{2}}^{\frac{M}{2}}\alpha^{-|i-\beta|}}$$

We compute the denominator as follows:

$$\sum_{i=\frac{-M}{2}}^{\frac{M}{2}}\alpha^{-|i-\beta|} = \sum_{i=\frac{-M}{2}}^{\beta-1}\alpha^{-|i-\beta|} + 1 + \sum_{i=\beta+1}^{\frac{M}{2}}\alpha^{-|i-\beta|} =$$

$$= \sum_{i=1}^{\frac{M}{2}+\beta}\alpha^{-i} + 1 + \sum_{i=1}^{\frac{M}{2}-\beta}\alpha^{-i} = 1 + \frac{1}{\alpha}\frac{1-\frac{1}{\alpha^{\frac{M}{2}+\beta}}}{1-\frac{1}{\alpha}} + \frac{1}{\alpha}\frac{1-\frac{1}{\alpha^{\frac{M}{2}-\beta}}}{1-\frac{1}{\alpha}}$$

$$= 1 + \frac{1}{\alpha^{\frac{M}{2}}(\alpha-1)}\left(2\alpha^{\frac{M}{2}} - 2\left(\frac{\alpha^{\beta}+\alpha^{-\beta}}{2}\right)\right)$$

Here, $ch^{\alpha}(x)$ is $\frac{\alpha^x+\alpha^{-x}}{2}$ by definition. Thus, we obtain:

$$P(\widetilde{\mathcal{G}}(\alpha,\beta)=k) = \frac{\alpha^{-|k-\beta|}}{1+\frac{2}{\alpha^{\frac{M}{2}}}\frac{\alpha^{\frac{M}{2}}-ch^{\alpha}(\beta)}{\alpha-1}}$$

Lemma 3 *Let $k, \ell > 0$. Then, the function $f : \mathbb{R} \to \mathbb{R}$ defined below is increasing on the interval $\left[-\frac{M}{2}+\ell, \frac{M}{2}\right]$.*

$$f(x) = \frac{k - ch^{\alpha}(x-\ell)}{k - ch^{\alpha}(x)}$$

PROOF OF LEMMA 3.

$$\frac{df(x)}{dx} = \frac{\ln(\alpha)\cdot(2k\cdot sh^{\alpha}(\ell)ch^{\alpha}(2x-\ell) - sh^{\alpha}(\ell))}{(k-ch^{\alpha}(x))^2}$$

$$= \frac{\ln(\alpha)\cdot sh^{\alpha}(\ell)\cdot(2k\cdot ch^{\alpha}(2x-\ell)-1)}{(k-ch^{\alpha}(x))^2} > 0$$

PROOF OF THEOREM 3. Let $\mathbb{D}_1 = (\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_n)$ and $\mathbb{D}_2 = (\mathbf{X}'_1, \mathbf{X}'_2, \ldots, \mathbf{X}'_n)$ be two datasets that differ in only one sub-demand over n slots. Without loss of generality, we assume that $\mathbb{D}_2$ has one more sub-demand than $\mathbb{D}_1$ (along n slots). We partition the possible outputs into two subsets $S = S_1 \cup S_2$, where S_1 contains the outputs that can appear in both $\mathcal{A}(\mathbb{D}_1)$ and $\mathcal{A}(\mathbb{D}_2)$, while the outputs in S_2 can appear only in $\mathcal{A}(\mathbb{D}_2)$, where $\mathcal{A}$ denotes algorithm SMARTBUFFERGEOMETRIC. In slot i, the buffer levels are denoted by L_i and L'_i, and the picked random values are s_i and s'_i used by $\mathcal{A}$ with $\mathbb{D}_1$ and $\mathbb{D}_2$, resp.

We have

$$\frac{P(\mathcal{A}(\mathbb{D}_1)\in S)}{P(\mathcal{A}(\mathbb{D}_2)\in S)} \le \frac{P(\mathcal{A}(\mathbb{D}_1)\in S_1)}{P(\mathcal{A}(\mathbb{D}_2)\in S_1)} + \frac{P(\mathcal{A}(\mathbb{D}_1)\in S_2)}{P(\mathcal{A}(\mathbb{D}_2)\in S)}$$

We first show that $P(\mathcal{A}(\mathbb{D}_1)\in S_1) \le e^{\varepsilon}\cdot P(\mathcal{A}(\mathbb{D}_2)\in S_1)$, where $\varepsilon = \ln\left(\alpha^{\Delta_n}\cdot\prod_{i=1}^{n}\frac{\alpha^{\frac{M}{2}}-ch^{\alpha}(\frac{M}{2}-\Delta_{i-1})}{sh^{\alpha}(\frac{M}{2})}\right)$. Let $(O_1,\ldots,O_n)\in S_1$ be a possible output of $\mathcal{A}(\mathbb{D}_1)$ and $\mathcal{A}(\mathbb{D}_2)$. In addition, u_i denotes the additional sub-demand in $\mathbb{D}_2$ in slot i, where $0\le u_i\le \max_j X_i^j$ for all i.

$$P(\mathbf{X}_1+\mathsf{s}_1=O_1,\ldots,\mathbf{X}_n+\mathsf{s}_n=O_n) =$$

$$= P(\mathsf{s}_1=O_1-\mathbf{X}_1,\ldots,\mathsf{s}_n=O_n-\mathbf{X}_n) =$$

$$= P(\mathsf{L}_1=O_1-\mathbf{X}_1,\ldots,\mathsf{L}_n=\sum_{k=1}^{n}(O_k-\mathbf{X}_k)) =$$

$$= P(\mathsf{L}_1=a_1)\cdot P(\mathsf{L}_2=a_2|\mathsf{L}_1=a_1)\cdot\ldots$$
$$\ldots\cdot P(\mathsf{L}_n=a_n|\mathsf{L}_{n-1}=a_{n-1})$$

where $a_i=\sum_{k=1}^{i}(O_k-\mathbf{X}_k)$. Similarly,

$$P(\mathbf{X}'_1+\mathsf{s}'_1=O_1,\ldots,\mathbf{X}'_n+\mathsf{s}'_n=O_n) =$$

$$= P(\mathsf{L}'_1=a'_1)\cdot P(\mathsf{L}'_2=a'_2|\mathsf{L}'_2=a'_1)\cdot\ldots$$
$$\ldots\cdot P(\mathsf{L}'_n=a'_n|\mathsf{L}'_{n-1}=a'_{n-1})$$

where $a'_i=a_i-\sum_{j=1}^{i}u_j$. Thus,

$$\frac{P(\mathsf{L}_i=a_i|\mathsf{L}_{i-1}=a_{i-1})}{P(\mathsf{L}'_i=a'_i|\mathsf{L}'_{i-1}=a'_{i-1})} \le \frac{P(\widetilde{\mathcal{G}}(\alpha,a_{i-1},0,M)=a_i)}{P(\widetilde{\mathcal{G}}(\alpha,a'_{i-1},0,M)=a'_i)} =$$

$$= \frac{P(\widetilde{\mathcal{G}}(\alpha,a_{i-1}-M/2,-M/2,M/2)=a_i-M/2)}{P(\widetilde{\mathcal{G}}(\alpha,a'_{i-1}-M/2,-M/2,M/2)=a'_i-M/2)} \le$$

(based on Lemma 2)

$$\le \alpha^{u_i}\frac{1+\frac{2}{\alpha^{\frac{M}{2}}}\frac{\alpha^{\frac{M}{2}}-ch^{\alpha}(a_{i-1}-\sum_{k=1}^{i-1}u_i-M/2)}{\alpha-1}}{1+\frac{2}{\alpha^{\frac{M}{2}}}\frac{\alpha^{\frac{M}{2}}-ch^{\alpha}(a_{i-1}-M/2)}{\alpha-1}} \le$$

$$\le \alpha^{u_i}\cdot\frac{\alpha^{\frac{M}{2}}-ch^{\alpha}(a_{i-1}-\sum_{k=1}^{i-1}u_i-M/2)}{\alpha^{\frac{M}{2}}-ch^{\alpha}(a_{i-1}-M/2)} \le$$

(based on Lemma 3)

$$\le \alpha^{u_i}\cdot\frac{\alpha^{\frac{M}{2}}-ch^{\alpha}(M-\sum_{k=1}^{i-1}u_i-\frac{M}{2})}{\alpha^{\frac{M}{2}}-ch^{\alpha}(M-\frac{M}{2})} \le$$

$$\le \alpha^{\max_j X_i^j}\cdot\frac{\alpha^{\frac{M}{2}}-ch^{\alpha}(\frac{M}{2}-\Delta_{i-1})}{\alpha^{\frac{M}{2}}-ch^{\alpha}(\frac{M}{2})}$$

Thus,

$$\frac{P(\mathcal{A}(\mathbb{D}_1)\in S_1)}{P(\mathcal{A}(\mathbb{D}_2)\in S_1)} = \frac{P(\mathbf{X}_1+\mathsf{s}_1=O_1,\ldots,\mathbf{X}_n+\mathsf{s}_n=O_n)}{P(\mathbf{X}'_1+\mathsf{s}'_1=O_1,\ldots,\mathbf{X}'_n+\mathsf{s}'_n=O_n)} \le$$

$$\le \alpha^{\Delta_n}\prod_{i=1}^{n}\frac{\alpha^{\frac{M}{2}}-ch^{\alpha}(\frac{M}{2}-\Delta_{i-1})}{\alpha^{\frac{M}{2}}-ch^{\alpha}(\frac{M}{2})}$$

Now consider the ratio $\frac{P(\mathcal{A}(\mathbb{D}_1)\in S_2)}{P(\mathcal{A}(\mathbb{D}_2)\in S)}$. Assuming that $\Delta_i=\sum_{t=1}^{i}\max_j X_t^j$, $\mathcal{A}(\mathbb{D}_1)$ has output from S_2, if $\mathsf{L}_i\in[0,\Delta_i]$ (if it happens, there is a privacy breach). The probability of this is denoted by $\delta^{(i)}$ in slot i and it is bounded by the probability that the buffer level is within $[0,\Delta_i]$ in slot i. Recall that $\mathbf{M}^k_{i,j}$ is the probability that the buffer level changes from j to i after k slots. If the initial distribution of the buffer level is π_0, which is 0 for all coordinates of $[0,M]$ except for coordinate $M/2$ which is 1, then the distribution of the buffer level after k slots can be computed as $\pi_k=\mathbf{M}^k\pi_0$. Thus, $\delta^{(k)}=\sum_{i=0}^{\Delta_k}[\pi_k]_i$, and we obtain:

$$\frac{P(\mathcal{A}(\mathbb{D}_2)\in S_1)}{P(\mathcal{A}(\mathbb{D}_2)\in S)} \le \frac{1-\prod_{k=1}^{n}(1-\delta^{(k)})}{P(\mathcal{A}(\mathbb{D}_2)\in S)}$$

Due to the symmetry property of π_k (see Theorem 4), it is easy to show that $\frac{P(\mathcal{A}(\mathbb{D}_2)\in S_2)}{P(\mathcal{A}(\mathbb{D}_1)\in S)} \le 1+\frac{1-\prod_{k=1}^{n}(1-\delta^{(k)})}{P(\mathcal{A}(\mathbb{D}_1)\in S)}$ which completes the proof.

Note that if the total sub-demand over n slots is more than the buffer size then $\delta^{(n)}=1$ and the privacy breach is certain. The probability of a privacy breach (and so the δ) can be decreased by expanding the buffer.

C. PROOF THEOREM 4

PROOF. The proofs of satisfiability and maximum deficit are trivial. We will show that the expected value of the buffer level with transition matrix $\mathbf{M}$ and initial distribution π_0, where all coordinates are zero except the $(M/2)$th which is 1, is $M/2$. This implies that the expected deficit is $M/2$.

Note that the distribution of the buffer level after k slots can be computed as $\pi_k=\mathbf{M}^k\pi_0$. Since π_0 is a symmetrical distribution (i.e. $\pi_{M/2+i}=\pi_{M/2-i}$ where $i\in[0,M/2]$), using induction, it is sufficient to show that if π is a symmetrical distribution then $\pi'=\mathbf{M}\cdot\pi$ is also symmetrical. Let π be such a symmetrical distribution. By definition,

$$\pi'_{M/2+i}=\sum_{j=0}^{M}\mathbf{M}_{M/2+i,j}\cdot\pi_j$$

$$\pi'_{M/2-i}=\sum_{j=0}^{M}\mathbf{M}_{M/2-i,j}\cdot\pi_j$$

Recall that $\mathbf{M}_{i,j}=P(\widetilde{\mathcal{G}}(\alpha,j,0,M)=i)=P(\widetilde{\mathcal{G}}(\alpha,j-M/2,-M/2,M/2)=i-M/2)$. We have:

$$\pi'_{M/2-i}=\sum_{j=0}^{M}\frac{\alpha^{-|M/2-i-M/2-j+M/2|}\pi_j}{1+\frac{2}{\alpha^{\frac{M}{2}}}\frac{\alpha^{\frac{M}{2}}-ch^{\alpha}(j-M/2)}{\alpha-1}} =$$

$$= \sum_{j=0}^{M/2} \frac{\alpha^{-|-i-j|}\pi_{M/2+j}}{1+\frac{2}{\alpha^{\frac{M}{2}}}\frac{\alpha^{\frac{M}{2}}-ch^{\alpha}(j)}{\alpha-1}} + \sum_{j=0}^{M/2} \frac{\alpha^{-|-i+j|}\pi_{M/2-j}}{1+\frac{2}{\alpha^{\frac{M}{2}}}\frac{\alpha^{\frac{M}{2}}-ch^{\alpha}(-j)}{\alpha-1}}$$

Similarly,

$$\pi'_{M/2+i} = \sum_{j=0}^{M} \frac{\alpha^{-|M/2+i-M/2-j+M/2|}\pi_j}{1+\frac{2}{\alpha^{\frac{M}{2}}}\frac{\alpha^{\frac{M}{2}}-ch^{\alpha}(j-M/2)}{\alpha-1}} =$$

$$= \sum_{j=0}^{M/2} \frac{\alpha^{-|i-j|}\pi_{M/2+j}}{1+\frac{2}{\alpha^{\frac{M}{2}}}\frac{\alpha^{\frac{M}{2}}-ch^{\alpha}(j)}{\alpha-1}} + \sum_{j=0}^{M/2} \frac{\alpha^{-|i+j|}\pi_{M/2-j}}{1+\frac{2}{\alpha^{\frac{M}{2}}}\frac{\alpha^{\frac{M}{2}}-ch^{\alpha}(-j)}{\alpha-1}}$$

Since $\alpha^{-|i-j|} = \alpha^{-|j-i|}$, $\alpha^{-|i+j|} = \alpha^{-|-j-i|}$, $ch^{\alpha}(-j) = ch^{\alpha}(j)$, and $\pi_{M/2+i} = \pi_{M/2-i}$, we obtain

$$\pi'_{M/2+i} = \pi'_{M/2-i}$$