

Differentially Private Billing with Rebates

George Danezis¹, Markulf Kohlweiss¹, and Alfredo Rial²

¹ Microsoft Research, Cambridge, UK

{gdane,markulf}@microsoft.com

² ESAT-COSIC / IBBT, KU Leuven, Belgium

alfredo.rial@esat.kuleuven.be

Abstract. A number of established and novel business models are based on fine grained billing, including pay-per-view, mobile messaging, voice calls, pay-as-you-drive insurance, smart metering for utility provision, private computing clouds and hosted services. These models apply fine-grained tariffs dependent on time-of-use or place-of-use to readings to compute a bill.

We extend previously proposed billing protocols to strengthen their privacy in two key ways. First, we study the monetary amount a customer should add to their bill in order to provably hide their activities, within the differential privacy framework. Second, we propose a cryptographic protocol for oblivious billing that ensures any additional expenditure, aimed at protecting privacy, can be tracked and reclaimed in the future, thus minimising its cost. Our proposals can be used together or separately and are backed by provable guarantees of security.

1 Introduction

A number of business models are based on billing customers for fine grained use of a system or resource: mobile network providers charge per call length and type, pay-per-view TV providers charge for the actual requested content. Newer businesses rely heavily on fine grained recordings of activity for billing. Pay-as-you-drive automotive insurance bills drivers per mile depending on the type of road and time of travel. Electronic tolling and congestion charging schemes have been proposed on similar lines. Smart-metering for electricity and gas is being rolled out in the EU and the US in the next few years. Finally, private cloud provision as well as hosted on-line service provision might rely on fine-grained measurements of CPU time usage, memory allocation, disk storage, peak bandwidth, or even the demand and network congestion at the time of day.

The downside of fine-grained metering and billing is the potential threat to privacy [1, 2]. A common privacy-invasive architecture to support such billing consists of providers collecting all usage information in order to apply the appropriate tariffs. Privacy-friendly protocols have also been developed: it is possible to cryptographically combine certified readings with a tariff policy to produce a certified bill that leaks no additional information about the detailed readings [3–5]. Yet, even the final bill, which is for instance aggregated over a period of usage, may leak information or be used to leak specific readings.

This work makes two contributions to the field of privacy-friendly metering and billing. First, we discuss how to eliminate incidental, accidental or deliberate leakages

of information resulting from disclosing the final bill. We show that by adding some, in the long run small, amount of noise it is possible to offer strong privacy guarantees on what an adversary can infer from the final bill. This problem is similar to covert channel minimization [6], and we use techniques from differential privacy that could be more widely applicable. Second, we attempt to minimise the cost of privacy through a cryptographic oblivious billing mechanism. The true cost of service provision is tracked across billing periods, but not revealed to the service provider, which can only verify the deposited funds cover costs. This allows customers to determine the levels of privacy they require and even get a rebate for the additional funds they used to protect their privacy.

Throughout this work we motivate our protocols through the example of a leased private computation cloud. A service provider installs a cloud of 10000 CPUs in the premisses of a large government intelligence agency. In our example, billing is performed on the basis of the compute hours actually used at a fixed rate of \$0.12 per CPU instance / hour³. A more complex tariff scheme where each hour in the year is costed differently is also supported. The government agency needs to settle the bill each month, but is worried that the amount of computation on particular days is leaked to its adversaries. We will show how our protocols can be used to reduce any leakage below a desired level.

Discussion of the state-of-the-art. Deployed systems for fine grained billing usually employ procedural access control mechanisms to protect privacy: usage data is gathered, and often stored centrally for the purposes of billing. Access control allows only designated parties and processes to access the data, and encryption technology might be used to protect storage and communications. Despite those protections, the fact that personal information is under the control of a service provider raises privacy concerns. A pilot deployment of a pay-as-you-drive insurance scheme by Norwich Union failed, stating privacy concerns as a leading reason for low uptake.⁴

Two types of privacy preserving metering and billing have been proposed in the literature. First, a meter can be entrusted with applying a fine grained tariff to the usage data and only communicating to the service provider a final total fee. In this setting the meter has to be trusted by the users and the service providers both for privacy and correctness. This is usually achieved through trusted hardware and certification. In the automotive setting, where meters record positions of cars for tolling, spot checks have also been proposed to verify the correctness of the meter operation [3]. The second architecture requires meters to cryptographically certify readings and securely hand them over to a user device or service. Cryptographic operations can then be used to apply a tariff scheme, and output a bill along with the necessary cryptographic proofs that certify its correctness. Meters are simpler, and any device can be used to compute bills [5]. Both architectures achieve the same goal: the bill and other necessary information are

³ The value of a standard compute instance / hour on Amazon EC2 and Microsoft Azure in December 2010.

⁴ Insurer stops 'pay as you drive', BBC Radio 4's Money Box <http://news.bbc.co.uk/2/hi/programmes/moneybox/7453546.stm>

made available to the service provider, but further information on detailed readings is hidden from it and only available to the consumer.

In this work we are concerned with the remaining information leakage from privacy-preserving billing systems. The value revealed by the protocols, namely the value of the bill, could leak information or be used as a covert channel.

To illustrate the threat, consider a resource consumed in a number of i_{max} distinct time periods i , for $i \in [0, i_{max}]$. Some consumption takes place at each time period i denoted by $c_i \in [0, c_{max}]$, that should be billed at a tariff of p_i per unit. Thus the final bill for all periods should be $B = \sum_{i=0}^{i_{max}} c_i \cdot p_i$. Without making any assumptions on the consumption patterns, as they are out of the system designer's control, it is difficult to estimate what information may be leaking from the final value B . For example an adversary may know, through some side information, that the user consumed only in a single time period T . In such a case the exact value of c_T can be inferred straightforwardly by computing $c_T = B/p_T$. This example threat illustrates that a solution to this problem should make no assumptions about the consumption pattern, assume that arbitrary side-information is available to the adversary, and work for arbitrary (but known) tariff schemes.

We will use a trivial solution as a benchmark to evaluate our own proposals: the user could always pay an amount equivalent to the maximum possible consumption. In the example used so far, this would be: $\max B = c_{max} \cdot \sum_{i=0}^{i_{max}} p_i$. While this is an adequate solution from a privacy perspective, it nullifies the benefits of fine-grained billing as users end up paying a fixed premium irrespective of their consumption. Furthermore it is very wasteful, if the objective is to hide usage of the private cluster at the granularity of an hour or a day.

Outline. Our techniques provide guarantees of privacy depending on the level of protection required by the customers, as well as a cryptographic scheme to amortise the cost of such privacy provision. In Section 2 we study how much noise one needs to add to a bill to ensure specific consumption windows are protected. In Section 3 we propose a cryptographic rebate protocol that keeps a hidden track of the actual amounts due across multiple billing periods, allowing users to reclaim some of the extra payments made. The rebate protocols also support deposits, anonymous payments using e-cash, and negative bill noise, and prevent abuse by ensuring the funds paid cover the costs of consumption.

2 Differential Privacy for Billing

We start from the premise that customers can add some “noise” to their bill in order to hide their exact usage at specific times. Of course this billing noise represents real money, so they wish to minimise it for a given level of protection required. The first problem we tackle is to determine how much more a customer should pay to hide their pattern of activity for a particular time frame.

Differential privacy was developed as a framework for hiding personal records within databases [7]. A statistic extracted from a database is differentially private if it is nearly as likely as if it was extracted from a database with an arbitrary record removed. This definition encapsulates the intuition that a single individual's record does

not overwhelmingly affect the statistic in a way that information about the record might leak.

We have to modify this definition as well as its precise mathematical counterpart to make it applicable to the billing setting. We consider as our database the set of all readings from a meter. In the case of billing private cloud usage each record represents the number of CPUs used for each hour of the billing period. The customer then has to specify its privacy goal: for example they may wish to hide their activity at any arbitrary hour or any arbitrary day of computing. Then they should determine the quality of the protection provided, in terms of how much information the bill reveals about any particular period. Using those parameters we can calculate the additional amount to bill in order to achieve the desired privacy goals.

2.1 Privacy definitions

For simplicity we consider fixed size databases corresponding to a fixed term billing period. For our application this is sufficient, as we are primarily interested in the number of CPU instances used during each hour of the pricing period. For this reason the domain of all possible data sets is described as the Cartesian product: $\mathcal{D} = \{0, \dots, c_{max}\}^{i_{max}}$. For our private cloud scenario c_{max} is the number of instances in the private cloud, and i_{max} is the number of records per billing period. In our concrete example $c_{max} = 10000$ and i_{max} is the number of hours in a month or a year.

First we define the “distance” between two sets of readings, and repeat some key definitions and results from differential privacy [7], upon which we will be building.

Definition 1. *The record distance $RDist(D_1, D_2)$ between two data sets $D_1, D_2 \in \mathcal{D}$ corresponds to the number of elements (records) in which D_1 and D_2 differ.*

Definition 2. *A randomized function K gives ϵ -differential privacy if for all data sets $D_1, D_2 \in \mathcal{D}$ with $RDist(D_1, D_2) \leq 1$, and all $S \in \Sigma_{Image(K)}$,⁵*

$$Pr[K(D) \in S | D = D_1] \leq \exp(\epsilon) \times Pr[K(D) \in S | D = D_2] .$$

The probability is taken over the randomness of K .

Intuitively, mechanisms fulfilling this definition address concerns that an individual might have about filling in one record truthfully, rather than arbitrarily. Differential privacy guarantees that no output (and thus consequences of outputs) becomes significantly more or less likely. In our case the randomized function K will be the billing amount increased by some random value.

A further observation about hiding multiple records k from a database will also prove useful:

Definition 3. *A randomized function K gives (k, ϵ) -differential privacy if for all data sets $D_1, D_2 \in \mathcal{D}$ with $RDist(D_1, D_2) \leq k$, and all $S \in \Sigma_{Image(K)}$,*

$$Pr[K(D) \in S | D = D_1] \leq \exp(\epsilon \cdot k) \times Pr[K(D) \in S | D = D_2] .$$

The probability is taken over the randomness of K .

⁵ A σ -algebra over a set X is a set $\Sigma_X \subset 2^X$ such that $\emptyset \in \Sigma_X$; $S \in \Sigma_X \Rightarrow (X \setminus S) \in \Sigma_X$; and for any $(S_i)_{i \in \mathbb{N}}$, $S_i \in \Sigma_X$, $\bigcap S_i \in \Sigma_X$.

Lemma 1. *A ϵ -differentially private privacy mechanism K is also (k, ϵ) -differentially private.*

Lemma 1 follows from Definition 3, and shows that the same privacy mechanism K can obstruct inferences on multiple records. In such cases it provides a lower amount of privacy (i.e. $\epsilon' = \epsilon \cdot k$). Hence if a mechanism is to be used to protect multiple records suitable security margins should be provided.

Differentially private mechanisms. The classical differential privacy mechanism by Dwork [7] adds Laplacian noise to the outcome of a query, parametrised by the “sensitivity” of the function f .

Definition 4. *The sensitivity of a function $f : \mathcal{D} \rightarrow R^n$ is the maximum distance between output values for which the domain differs in at most one record:*

$$\Delta_f = \max_{\substack{D_1, D_2 \in \mathcal{D} \\ RDist(D_1, D_2) \leq 1}} \|f(D_1) - f(D_2)\|_1$$

For $n = 1$ the sensitivity of f is the maximum difference $|f(D_1) - f(D_2)|$ between pairs of databases D_1, D_2 that differ in only one element. It is shown in [7] that if $f : \mathcal{D} \rightarrow R$ is a function with sensitivity Δ_f , then $K(D) = Lap(f(D), \Delta_f/\epsilon)$ is differentially private.

Our adaptations of the differential privacy definitions. Instead of bounding the ratio between output probabilities of actual vs. arbitrary information for a single hourly record, we want to give customers the option of hiding an arbitrary period of time. For example we may want to hide specifics of daily (chunks of 24 records) or weekly (chunks of 168 records) consumption. We call the period length a user is concerned with the *privacy unit*. Furthermore we need to achieve this for statistics in discrete domains (not continuous function), that can only make the bills bigger, never smaller.

Definition 5. *The u -distance $Dist_u(D_1, D_2)$, e.g., $u \in \{\text{hourly, daily, weekly}\}$ between two data sets $D_1, D_2 \in \mathcal{D}$ corresponds to the number of u -units (collection of records) in which D_1 and D_2 differ.*

Our pricing scheme maps each $D \in \mathcal{D}$, $D = (c_1, \dots, c_{i_{max}})$ to a discrete price: $price(D) = \sum_{i=1}^{i_{max}} c_i \cdot p_I$, where i_{max} is the number of records per billing period, and p_I is the price per hour per instance. Rather than having continuous positive and negative noise as in the original Laplacian differential privacy mechanism, we want to only add discrete positive noise.

If we consider only privacy mechanisms with discrete outputs, we can simplify the differential privacy definition. For discrete distributions, $\Sigma_{Image(K)} = 2^{Image(K)}$, and $Pr[K(D) \in S] = \sum_{r \in S} Pr[K(D) = r]$. Definition 2 can thus be restated as the following equation: $\sum_{r \in S} Pr[K(D) = r | D = D_1] \leq \exp(\epsilon) \cdot \sum_{r \in S} Pr[K(D) = r | D = D_2]$. From this we derive an alternative definition for differential privacy for discrete distributions:

Definition 6. A randomized function K gives ϵ -differential u -privacy if for all data sets $D_1, D_2 \in \mathcal{D}$ with $\text{Dist}_u(D_1, D_2) \leq 1$, and all $r \in \text{Image}(K)$,

$$\Pr[K(D) = r | D = D_1] \leq \exp(\epsilon) \times \Pr[K(D) = r | D = D_2] .$$

The probability is taken over the randomness of K .

Lemma 2. Definition 2, Definition 3, and Lemma 1 apply to u -privacy:

1. For discrete privacy mechanisms Definition 2 and Definition 6 for $u = \text{hourly}$ are equivalent.
2. Let n_u be the number of records in a u -unit. If K is (n_u, ϵ) -differential hourly-private, then K is also $(n_u \cdot \epsilon)$ -differential u -private.

Dwork [8] notes that, because of the multiplicative nature of the definition, an output whose probability is zero on a given database must also have probability zero on any neighboring database, and therefore, by repeated application of the definition, on any other database.

Handling privacy mechanisms that result in distributions for which the support of $K(D_1)$ and $K(D_2)$ may differ requires extra care. Such a situation arises, e.g., when K adds only positive noise. If for instance $\text{price}(D_1) < \text{price}(D_2)$ to which K adds positive noise. Let ν_{\min} be the minimum amount of noise that is added, then the value $r = \text{price}(D_1) + \nu_{\min}$ is in the support of $K(D_1)$ but has 0 probability for $K(D_2)$. It follows that such a mechanism can never be differentially private.

To overcome this problem, we define partial differential privacy. A statistic offers partially differential u -privacy if it is differentially private for all outputs in the overlapping support of any two databases D_1 and D_2 with $\text{Dist}_u(D_1, D_2) \leq 1$. Furthermore we require the probability that the output of the statistic is not in the overlapping domains to be bound by a small probability δ . This means that the function is differentially private most of the time (or with probability at least $1 - \delta$).

Definition 7. A randomized function K gives δ -partially ϵ -differential u -privacy if the following two properties hold:

1. For all $D_1, D_2 \in \mathcal{D}$ with $\text{Dist}_u(D_1, D_2) \leq 1$, and all $r \in \text{Supp}(K(D_1)) \cap \text{Supp}(K(D_2))$,

$$\Pr[K(D_1) = r] \leq \exp(\epsilon) \times \Pr[K(D_2) = r] .$$

2. For all data sets $D_1, D_2 \in \mathcal{D}$ with $\text{Dist}_u(D_1, D_2) \leq 1$,

$$\Pr[r \leftarrow K(D_1) : r \notin \text{Supp}(K(D_2))] < \delta .$$

For both properties, the probability is taken over the randomness of K .

As for the traditional differential privacy definitions, longer periods of privacy can be guaranteed with lower security parameters:

Lemma 3. Let n_u be the number of records in a u -unit. If K is δ -partially (n_u, ϵ) -differential hourly-private, then K is also $(n_u \cdot \delta)$ -partially $(n_u \cdot \epsilon)$ -differential u -private

Proof. Consider the joint distribution of K for all D_1 and D_2 with $RDist(D_1, D_2) \leq n_u$. The probability of drawing a value r not in the domain of at least one of $K(D_i)$ is $\delta' \leq 1 - (1 - \delta)^{n_u} \leq n_u \cdot \delta$. This proves Property 2 for partial differential privacy. If r is in the domain, Property 1 is proved as in Lemma 2. \square

Given the above definition for privacy we propose a concrete mechanism to obscure readings. We simply add to the bill $f(D)$ for consumption D an amount of noise drawn from a Geometric distribution with parameter $p = \epsilon/\Delta_{f,u}$.⁶ The sensitivity $\Delta_{f,u}$ is the maximum difference of a bill between two databases D_1 and D_2 differing in at most 1 u-unit (e.g. an hour, a day, or a week). Similarly, ϵ is a security parameter expressing information leakage.

Theorem 1. *Let $f : \mathcal{D} \rightarrow \mathcal{R}$ be a function with sensitivity $\Delta_{f,u}$, then $K(D) = f(D) + Geo(\epsilon/\Delta_{f,u})$ is $(2 \cdot \epsilon)$ -partially ϵ -differentially u -private. (See [10] for the proof.)*

As also noted by [11], the application of a public function on the outputs of a differentially private statistic does not leak any additional information. We can modify the billing function to only charge up to the maximum possible consumption: $K'(D) = \min(f(D) + Geo(\epsilon/\Delta_{f,u}), \max_{D'} f(D'))$. Intuitively we use geometric noise, as this adds the maximal uncertainty for a given mean. The variant of the geometric distribution with support for negative and positive integers defined as $\Pr[k] = \frac{1}{2}(1-p)^{|k|}p$ is the discrete equivalent of the Laplace distribution, and would also provide differentially private guarantees. We limit ourselves to the proposed noise distribution to ensure users only add positive noise to their bills.

Interpretation of differential privacy in terms of inference. From the attackers perspective the goal of collecting statistics about the output of the privacy mechanism K is to infer something about the underlying database. For instance, the attacker might want to distinguish between two databases D_1 and D_2 , in the sense of semantic security.

Differential privacy does not guarantee anything about the probability ratio (likelihood ratio) between databases D_1 and D_2 with $Dist_u(D_1, D_2) \leq 1$ given an observed outcome of K ; it merely says that this ratio will differ only by a small factor from the ratio of the prior. Note that because D_1 and D_2 are interchangeable, the new ratio is also bounded from below.

Lemma 4. *Given an observed outcome of a differentially private K the probability ratio (likelihood ratio) between databases D_1 and D_2 with $Dist_u(D_1, D_2) \leq 1$ differs by less than a factor $\exp(\epsilon)$ from the ratio of the prior.*

$$\frac{\Pr[D = D_1 | K(D) = r]}{\Pr[D = D_2 | K(D) = r]} \leq \exp(\epsilon) \times \frac{\Pr[D = D_1]}{\Pr[D = D_2]}.$$

Proof. From Bayes theorem we can write:

$$\Pr[D = D_i | K(D) = r] = \frac{\Pr[K(D) = r | D = D_i] \times \Pr[D = D_i]}{\Pr[K(D) = r]}$$

⁶ Two-sided Geometric noise was also proposed in [9] as a differential privacy mechanism.

Privacy units	Security (ϵ)	Pay Monthly	Pay Yearly	Fixed Rate
Hourly	0.1	$\beta + \$144,000$	$\beta + \$12,000$	\$10,512,000
(units = 1)	0.01	$\beta + \$1,440,000$	$\beta + \$120,000$	\$10,512,000
Daily	0.1	$\beta + \$3,456,000$	$\beta + \$288,000$	\$10,512,000
(units = 24)	0.01	$(\$10,512,000)$	$\beta + \$2,880,000$	\$10,512,000
Weekly	0.1	$(\$10,512,000)$	$\beta + \$2,016,000$	\$10,512,000
(units = 168)	0.01	$(\$10,512,000)$	$(\$10,512,000)$	\$10,512,000

Table 1. Yearly average bill after the application of the privacy mechanism K' compared with the fixed rate privacy mechanism. Different values of the security parameter (ϵ), different privacy units (hourly, daily and weekly) as well as the options of paying monthly or yearly are presented. Amounts in parenthesis indicate that the expected cost is higher than paying for the maximum consumption.

whence, since K is differentially private, we can write:

$$\begin{aligned} \frac{Pr[D = D_1 | K(D) = r]}{Pr[D = D_2 | K(D) = r]} &= \frac{Pr[K(D) = r | D = D_1]}{Pr[K(D) = r | D = D_2]} \times \frac{Pr[D = D_1]}{Pr[D = D_2]} \\ &\leq \exp(\epsilon) \times \frac{Pr[D = D_1]}{Pr[D = D_2]}. \end{aligned} \quad \square$$

The cost of privacy. Obscuring bills by adding noise may lead to paying extra for a service. Customers have incentives to minimise their costs for a desired level of privacy protection. We provide a few illustrative examples of the average extra cost involved in settling a bill for different privacy units of an hour, a day or a week. In our usual example we consider a private cloud of 10K CPUs, billed as \$0.12 a CPU / hour. We denote as $\beta = f(D)$ the actual service cost associated with the use of the service for a year.

It is clear from Table 1 that providing a differentially private bill for more than a single hourly period is an expensive business. The proposed mechanism allows for lower overheads for yearly bills when customers wish to protect arbitrary hours or days in the year. When it comes to protecting arbitrary weeks this protection is only offered with a low security parameter ($\epsilon = 0.1$). Why is the cost so high? It is because the privacy guarantee offered is very strong: no matter what side information the adversary has, including the detailed readings for other periods, they should not be able to infer information about an arbitrary privacy unit. For example if the adversary knows the exact consumption for the other 364 days they should still not learn more than permitted about the last day. This is a very strong guarantee and as a result it comes at a high cost, when applied directly.

Table 1 also contains the cost of paying bills monthly, which incur a 12 fold overhead for the same level of protection. It is clear that there are advantages in paying in batches if in fact the desired property is to hide any fixed period of time within the billing period (an hour, a day, a week). We will see in the next section how we can do better than this: we can aggregate the true cost of service provision, and use cryp-

tographic methods to reclaim most of the additional cost of privacy in the long term without sacrificing any security.

Longer guarantees. Degradation of privacy in our framework is graceful, since some privacy guarantees are provided for periods longer than what is strictly defined by the chosen u-units. For example a user may choose a partially ϵ -differential function $K_{\epsilon,24}$ providing u-privacy for a day (i.e. 1 u-unit = 24 hourly periods) with $\epsilon = 0.01$. In our standing example this means he should add an extra amount to his bill drawn from a Geometric distribution with parameter \$2,880,000. What does that guarantee? Let's assume the adversary knows the exact consumption about all days except for one. Furthermore the adversary knows that the consumption on the target day could only have taken one out of two values with equal probability: this means that the ratio of priors $\frac{Pr[D=D_1]}{Pr[D=D_2]} = 1$. Then after receiving information about the bill the adversary would at best know that $0.99 \approx 1/(1 + \epsilon) \approx 1/e^\epsilon \leq \frac{Pr[D=D_1|K'(D)]}{Pr[D=D_2|K'(D)]} \leq e^\epsilon \approx 1 + \epsilon = 1.01$. This is a small amount of information.

Now let's consider an adversary that tries to infer something over a longer period, e.g., a week. The adversary knows all user consumption outside this target week, and furthermore knows that user consumption within the week could only have been one of two possibilities D'_1 or D'_2 with equal probability as before. Due to Lemma 3 we know that the $K'(D)$ scheme is also partially ϵ -differentially private for a longer u-unit of a week (1 weakly-unit = 7×24 hourly-units), with a new security parameter $\epsilon' = \epsilon \cdot 7$. This means that the new posterior ratio of probabilities over the two only possible outcomes is $0.93 \approx 1/(1 + \epsilon') \approx 1/e^{\epsilon'} \leq \frac{Pr[D=D_1|K'(D)]}{Pr[D=D_2|K'(D)]} \leq e^{\epsilon'} \approx 1 + \epsilon' = 1.07$. Despite the lower degree of privacy, some quantifiable protection is still available against longer-term profiling.

Limitations. Our variant of differential privacy relies on only introducing positive noise. This is desirable as it guarantees that the bill at least covers the cost of service provision. At the same time this provides a one sided security property: a final bill can always be confused with a lower bill, but not always with a higher bill. For example there is a positive probability that a sensitive day passes with no consumption and then no noise is added to the bill. If an adversary knows all other consumptions in the year, they can infer that indeed no consumption took place on the unknown day. Our mechanism thus assumes that the baseline of no consumption is not as sensitive as high consumption.

While information leakage about low levels of consumption is possible, it is not very likely for high levels of security as characterised by the security parameter δ .

Summary. We have shown that adding noise to the bill can provide high levels of security parametrised by a parameter ϵ and a privacy unit. This security holds even against adversaries with knowledge of many readings. At the same time this comes with a high overhead. In the next section we show that the bulk of the cost of providing privacy can be recuperated in the long run. We achieve this by keeping hidden accounts of what is actually due for service provision, versus what has been paid. In the long run users can only add the necessary noise to keep their accounts positive, including negative noise – while ensuring that their funds cover their consumption.

3 Private Billing with Rebates

We have seen that one way of protecting privacy involves adding ‘noise’ to the bill to be paid for a certain period. Yet, the amount of noise can become significant particularly to achieve a high quality of privacy or privacy for longer periods within the billing time frame. For this reason we develop a complementary oblivious billing protocol that can be used to alleviate those concerns. Its key features include:

- The ability to maintain a hidden bill of actual consumption that can be used to reclaim any excess used for protecting privacy at a later time.
- A mechanism for proving that the amount paid to the utility provider exceeds the bill for actual consumption without revealing the actual bill.
- Support for an initial deposit to support later use of positive as well as negative noise for the bills.
- Compatibility with anonymous e-cash schemes allowing bills to be settled anonymously, as well as advanced privacy friendly payment mechanisms that allow users to hide the amounts actually paid to the utilities.

We discuss in detail and prove the correctness of the billing protocols, and the mechanisms to ensure payments exceed the amount consumed. The specifics of optional e-cash protocols that allow hidden payments are beyond the scope of this work, and we leave their detailed description to future work.

Our oblivious payment protocols can be used to reclaim in the long run an excess paid as a result of a differentially private billing mechanism as presented in the previous sections. With the deposit facility, adding negative noise is possible, as long as the overall balance of payments stays positive. The protocols can also be used to support the naive mechanism where a bill for maximal consumption is paid, and allow parties to later reclaim some of it back. Finally given anonymous e-cash they can be used to provide full oblivious payments without the need to add any noise to the bills, as they never need to be revealed (technically: $noise = -fee$). Which variant to use therefore depends on the infrastructure available and the degree of complexity parties are ready to accept.

3.1 The PSM protocol

We will be building upon PSM (Privacy-Preserving Smart Metering), a cryptographic protocol for privacy-friendly billing [5]. PSM mediates interactions among three parties: a meter M that outputs consumption data $cons$ and related information $other$; a service provider P that establishes a pricing policy \mathcal{T} and a user U that receives consumption readings from meter M and at each billing period pays a fee to provider P . The pricing policy \mathcal{T} is a public function that takes consumption data $cons$ together with other information $other$ (e.g., the time of consumption) and computes a price. The overall price $price(D) = \sum_{i=1}^{|D|} price_i$ is computed by adding the prices corresponding to the individual consumptions in a billing period. For our running private cloud example, $\mathcal{T}(cons, other) = cons \cdot 0.12$ and does not depend on $other$. As in the original protocols we assume a tamper resistant meter is used to provide accurate and appropriately cryptographically packaged readings. These can be processed by the user to prove

their bill in zero-knowledge to the provider. At this point users may also choose to add some noise to ensure differential privacy.

The security of PSM is shown in the simulation-based security paradigm [12, 13]. In the real world, the protocol $\text{PSM}(\mathcal{M}, \mathcal{P}, \mathcal{U})$ is run in an adversarial environment that may corrupt some of the protocol parties, indicated by $\tilde{\mathcal{M}}, \tilde{\mathcal{P}}, \tilde{\mathcal{U}}$. Corrupted parties just forward messages between the environment and honest protocol participants. In the ideal world, dummy protocol parties $D_{\mathcal{M}}, D_{\mathcal{P}}, D_{\mathcal{U}}$ run an ideal protocol $\text{Ideal}(\mathcal{F}_{\text{PSM}}, D_{\mathcal{M}}, D_{\mathcal{P}}, D_{\mathcal{U}})$ by just forwarding messages to an ideal functionality \mathcal{F}_{PSM} . Uncorrupted D_x , $x \in \{\mathcal{M}, \mathcal{P}, \mathcal{U}\}$ interact with the environment while corrupted dummy parties \tilde{D}_x interact with a simulator Sim .

We consider w.l.o.g. a corrupted provider $\tilde{\mathcal{P}}$ and say that a protocol is secure against $\tilde{\mathcal{P}}$, if there exists a simulator Sim such that no environment Env can tell whether it is interacting with $\text{PSM}(\mathcal{M}, \tilde{\mathcal{P}}, \mathcal{U})$ or with $\text{Sim} \parallel \text{Ideal}(\mathcal{F}_{\text{PSM}}, D_{\mathcal{M}}, \tilde{D}_{\mathcal{P}}, D_{\mathcal{U}})$. Conceptually Sim translates influence that Env has through $\tilde{\mathcal{P}}$ on the protocol into influence on \mathcal{F}_{PSM} through $\tilde{D}_{\mathcal{P}}$, and leakage that $\tilde{D}_{\mathcal{P}}$ receives from \mathcal{F}_{PSM} into leakage that Env could learn from $\tilde{\mathcal{P}}$. Similarly, PSM is proven secure against a corrupted user $\tilde{\mathcal{U}}$.

Listing 1 Functionality \mathcal{F}_{PBR}

\mathcal{F}_{PBR} is parameterized by deposit relation R and a policy set Y and interacts with dummy parties $D_{\mathcal{M}}, D_{\mathcal{P}}$ and $D_{\mathcal{U}}$. Initially $T = \emptyset, d = 0, \text{account} = 0$.

```

On (Policy,  $\mathcal{T}$ ) from  $D_{\mathcal{P}}$  where  $\mathcal{T} \in Y$ 
  - store  $\mathcal{T}$ ; send (Policy,  $\mathcal{T}$ ) to  $D_{\mathcal{U}}$ 
On (Consume,  $\text{cons}, \text{other}$ ) from  $D_{\mathcal{M}}$ 
  - increment counter  $d$ ; add  $(d, \text{cons}, \text{other})$  to  $T$ ; send (Consume,  $\text{cons}, \text{other}$ ) to  $D_{\mathcal{U}}$ 
On (Deposit,  $(\text{inc}, \text{wit}), \text{instance}$ ) from  $D_{\mathcal{U}}$  where  $\text{balance} + \text{inc} \geq 0$ 
  - if  $((\text{inc}, \text{wit}), \text{instance}) \in R$ , let  $\text{balance} += \text{inc}$ , send (Deposit,  $\text{instance}$ ) to  $D_{\mathcal{P}}$ 
On (Payment,  $\text{from}, \text{until}, \text{noise}$ ) from  $D_{\mathcal{U}}$  where
   $0 \leq \text{from} \leq \text{until} \leq d$  and  $\text{balance} + \text{noise} \geq 0$ 
  - for  $i = \text{from}$  to  $\text{until}$ , calculate  $\text{price}_i = \mathcal{T}(\text{cons}_i, \text{other}_i)$ 
  - let  $\text{fee} = \sum_{i=\text{from}}^{\text{until}} \text{price}_i + \text{noise}$  and  $\text{balance} += \text{noise}$ 
  - send (Pay,  $\text{from}, \text{until}, \text{fee}$ ) to  $D_{\mathcal{P}}$ 

```

3.2 Rebate Ideal Functionality

We propose a new ideal functionality \mathcal{F}_{PBR} (see Listing 1) that extends the functionality \mathcal{F}_{PSM} . The functionality keeps track of the user's consumptions in a set T containing tuples $(i, \text{cons}, \text{other})$. During a payment, the policy \mathcal{T} is applied to all $(\text{cons}, \text{other})$ in the interval $[\text{from}, \text{until}]$ to compute the price $\text{price}_i = \mathcal{T}(\text{cons}_i, \text{other}_i)$ per consumption. The overall fee that the user has to pay is computed as $\text{fee} = \sum_{i=\text{from}}^{\text{until}} \text{price}_i + \text{noise}$. The value noise is added to the fee to improve the user's privacy. The ideal functionality also maintains a balance that corresponds to the sum of all the noise added to payments. Note that the user can get rebates by using negative noise, but that the balance is never allowed to be negative.

The ideal functionality also allows to increase the balance through a deposit. The user has to provide input $((\text{inc}, \text{wit}), \text{instance}) \in R$. The parameterization by relation R allows to support both standard deposit mechanisms that reveal the deposited amount

inc as well as advanced deposit mechanisms that hide this value from the provider. In the simple mechanism the user reveals how much he wants to deposit: $wit = \epsilon$ and R corresponds to simple equality, i.e. $R = \{(inc, \epsilon), inc) | inc \in \mathbb{Z}\}$.

To obtain a more advanced privacy-friendly deposit mechanism, the witness could correspond to a one-show anonymous credential $cred$. The relation requires that $cred$ is a one-show credential with an increment value inc and serial number s issued under public key pk_B , i.e. $R = \{((inc, cred), (s, pk_B)) | \text{Verify}(pk_B, cred, (inc, s)) = \text{accept}\}$. The real protocol cryptographically enforces this using a zero-knowledge proof of signature possession [14].⁷ To obtain such a one-show credential without revealing the value of inc to the provider, additional infrastructure is needed. In particular such a mechanism seems to require some form of anonymous payment, either physical cash or anonymous e-cash. Given such a payment mechanism, the provider's bank, after receiving an anonymous payment of value inc and depositing this amount on the provider's bank account, could blindly issue the signature $\text{Sign}(pk_B, (inc, s))$ using a partially-blind issuing protocol [14]. The issue protocol guarantees that the bank does not learn s , and thus even if the provider and his bank collude they cannot link the issuing of $cred$ to its use.

Listing 2 Protocol PBR(M, P, U)

Parties M, P, U are parameterized by R and Y and interact over secure channels. All participants have registered public keys generated by Mkeygen, Pkeygen, Ukeygen with a key registration authority \mathcal{F}_{REG} and keep their private keys secret. P also registers commitment parameters par_c .

- On (Policy, \mathcal{T}) from Env
 - P runs $\mathcal{T}_s \leftarrow \text{SignPolicy}(sk_P, \mathcal{T})$ and sends \mathcal{T}_s to U
 - Upon receiving \mathcal{T}_s , U extracts \mathcal{T} ; if $\mathcal{T} \notin Y$, he aborts
 - if $\text{VerifyPolicy}(pk_P, \mathcal{T}_s) = 1$, U stores \mathcal{T}_s , and sends (Policy, \mathcal{T}) to Env
- On (Consume, $cons$, $other$) from Env
 - M increments d_M , runs $SC \leftarrow \text{SignConsumption}(sk_M, par_c, cons, other, d_M)$ and sends (SC) to U
 - Upon receiving (SC), U runs $b \leftarrow \text{VerifyConsumption}(pk_M, par_c, SC, d_U + 1)$
 - if $b = 1$, U increments d_U , adds SC to T_U , parses SC as $(d_M, cons, open_{cons}, c_{cons}, other, open_{other}, c_{other}, sc)$, and sends (Consume, $cons$, $other$) to Env
- On (Deposit, (inc, wit) , $instance$) from Env where $balance + inc \geq 0$ and $((inc, wit), instance) \in R$
 - U runs $(aux', D) \leftarrow \text{Deposit}(par_c, (inc, wit), instance, aux, R)$
 - U sets $balance += inc$ and $aux = aux'$ and sends $(D, instance)$ to P
 - Upon receiving $(D, instance)$, P runs $(c'_{balance}, b) \leftarrow \text{VerifyDeposit}(par_c, D, c_{balance}, instance, R)$
 - if $b = 1$, he sets $c_{balance} = c'_{balance}$ and sends (Deposit, $instance$) to Env
- On (Payment, $from$, $until$, $noise$) from Env where $0 \leq from \leq until \leq d_U$ and $balance + noise \geq 0$
 - U runs $(aux', Q) \leftarrow \text{Pay}(sk_U, par_c, \mathcal{T}_s, T_U[from : until], noise, aux)$

⁷ A zero-knowledge proof of knowledge is a two-party protocol between a prover and a verifier. The prover convinces the verifier, who knows only a public proof $instance$, that he knows a secret input (called *witness*) that allows him to prove that the public and the secret value together fulfill some relational statement $(witness, instance) \in R$ without disclosing the secret input to the verifier.

- U sets $aux = aux'$ and $balance += noise$; U sends $(Q, from, until)$ to P
- Upon receiving $(Q, from, until)$, P runs $(fee, c'_{balance}, b) \leftarrow \text{VerifyPayment}(pk_M, pk_U, pk_P, par_c, Q, c_{balance}, from, until)$
- if $b=1$, he sets $c_{balance} = c'_{balance}$ and sends $(Pay, from, until, fee)$ to Env

3.3 Rebate Protocol

We propose a new protocol for privacy-preserving billing with rebates (PBR) (see Listing 2) that extends PSM with a mechanism for adding noise, keeping a hidden balance, and making deposits. Like PSM, our protocol operates in the \mathcal{F}_{REG} hybrid-model [12] where parties register their public keys at a trusted registration entity. As in the original scheme the user receives signed policies from the utility provider P and signed readings from the meter M. The payment transaction only reveals the overall fee, which now can be subject to additional noise.

We extend this protocol with a novel oblivious rebate system that allows the user to get rebates (in the amount of his noise) in future payments. The rebate is implemented using a homomorphic update c_{noise} to a balance commitment $c_{balance}$ that commits the user to his balance towards the provider but keeps the *balance* itself secret. Our protocol supports an optional Deposit mechanism that allows the user to add or withdraw funds from the rebate *balance*. Value *aux* contains the opening for a commitment $c_{balance}$ to *balance*. Through the use of zero-knowledge proofs the provider is guaranteed that the value committed to in $c_{balance}$ is updated correctly and never becomes negative.

The protocol parties P, U, and M interact with each other using algorithms Pkeygen, Ukeygen, Mkeygen (for key generation); SignPolicy, SignConsumption, Deposit, and Pay (for generation of input); and VerifyPolicy, VerifyConsumption, VerifyDeposit, and VerifyPayment (for verification of input). The functionality of the meter as well as SignPolicy, SignConsumption, VerifyPolicy, and VerifyConsumption are unchanged from the original scheme.⁸ We describe the new Deposit and VerifyDeposit algorithms and the changes to Pay and VerifyPayment:

Listing 3 Algorithms

- Deposit($par_c, (inc, wit), instance, aux, R$). Parse *aux* as $(balance, open_{balance}, c_{balance})$. Compute commitment $(c_{inc}, open_{inc}) = \text{Commit}(par_c, inc)$ and a non-interactive proof π_{inc} .⁹

$$\pi_{inc} \leftarrow \text{NIPK}\{ (inc, open_{inc}, wit, balance, open_{balance}) : \\ (c_{balance}, open_{balance}) = \text{Commit}(par_c, balance) \wedge \\ (c_{inc}, open_{inc}) = \text{Commit}(par_c, inc) \wedge \\ ((inc, wit), instance) \in R \wedge balance + inc \geq 0 \} .$$

Let $D = (\pi_{inc}, c_{inc})$ and $aux' = (balance + inc, open_{balance} + open_{inc}, c_{balance} \otimes c_{inc})$. Output (aux', D) .
- VerifyDeposit($par_c, D, c_{balance}, instance, R$). Parse D as $(\pi_{balance}, c_{inc})$. Verify π_{inc} . If verification succeeds, set $b = 1$ and $c'_{balance} = c_{balance} \otimes c_{inc}$, otherwise set $b = 0$. Output $(c'_{balance}, b)$.

⁸ For the sake of brevity we omit the Reveal mechanism of PSM. It would add little new and could be implemented in a straight forward manner using trapdoor commitments.

⁹ If R corresponds to equality, the protocol can be simplified to avoid computing c_{inc} .

- $\text{Pay}(sk_U, par_c, \mathcal{Y}_s, T, noise, aux)$. Parse aux as $(balance, open_{balance}, c_{balance})$. Compute commitment $(c_{noise}, open_{noise}) = \text{Commit}(par_c, noise)$ and a non-interactive proof π_{noise} :

$$\begin{aligned} \pi_{noise} \leftarrow \text{NIPK}\{ & (noise, open_{noise}, balance, open_{balance}) : \\ & (c_{balance}, open_{balance}) = \text{Commit}(par_c, balance) \wedge \\ & (c_{inc}, open_{inc}) = \text{Commit}(par_c, inc) \wedge balance + noise \geq 0\}. \end{aligned}$$

Let $aux' = (balance + noise, open_{balance} + open_{noise}, c_{balance} \otimes c_{noise})$.

The rest of the algorithm follows [5]: For each $(i, cons, open_{cons}, c_{cons}, other, open_{other}, c_{other}, sc) \in T$ where $from \leq i \leq until$, calculate $price_i = \mathcal{Y}(cons, other)$, commitment $(c_{price_i}, open_{price_i}) = \text{Commit}(par_c, price_i)$, and a proof π_i that $price_i$ was computed correctly according to the policy and the commitments c_{cons}, c_{other} . The proof π_i depends on the policy \mathcal{Y} and can use auxiliary values in \mathcal{Y}_s , see [5] on how to implement different pricing policies.

Computing $fee = noise + \sum_{i=from}^{until} price_i$ and $open_{fee} = open_{noise} + \sum_{i=from}^{until} open_{price_i}$ gives an opening to a commitment to fee . Let $Q = (fee, open_{fee}, c_{noise}, \pi_{balance}, \{sc_i, i, c_{cons_i}, c_{other_i}, c_{price_i}, \pi_i\}_{i=1}^N)$.¹⁰ Output (aux', Q) .

- $\text{VerifyPayment}(pk_M, pk_U, pk_P, par_c, Q, c_{balance}, from, until)$. Parse Q as $(fee, open_{fee}, c_{noise}, \pi_{balance}, \{sc_i, d_i, c_{cons_i}, c_{other_i}, c_{price_i}, \pi_i\}_{i=1}^N)$. Verify π_{noise} . If verification fails, set $b = 0$. Otherwise set $c'_{balance} = c_{balance} \otimes c_{noise}$ and $b = 1$.

The rest of the algorithm follows [5]: For $i = from$ to $until$, run $\text{Mverify}(pk_M, sc_i, \langle i, c_{cons_i}, c_{other_i} \rangle)$ and verify π_i . Set $b = 0$ if any of the signatures or the proofs is not correct. Add the commitments to the prices $c'_{fee} = c_{noise} \otimes (\otimes_{i=1}^N c_{price_i})$ and execute $\text{Open}(par_c, c'_{fee}, fee, open_{fee})$. If the output is **reject** set $b = 0$. Output $(fee, c'_{balance}, b)$.

Theorem 2. Given the security of its building blocks, PBR is secure against a corrupted provider \tilde{P} and a corrupted user \tilde{U} . (See [10] for the proof.)

Using PBR for differential privacy. Even an ideal cryptographic billing mechanism as described by the PSM or PBR ideal functionalities cannot protect the user's privacy against an adversary/environment that already knows enough about the user's behavior – possibly including all consumption or additional random noise – to infer privacy sensitive information from the final fee alone. For our privacy analysis we assume that the environment Env is divided into a part Env_U that is controlled by the user, and a part $\text{Env}_{\tilde{P}}$ that is controlled by the adversary and that may have some influence on and knowledge about the user's behavior. In the original PSM protocol all the final fee is only the result of the individual consumptions of Env_U for which the provider may make inferences or gain side information. The PBR protocol gives Env_U the possibility to obscure the fee with random noise, which is easier to conceal from $\text{Env}_{\tilde{P}}$.

4 Conclusions

Our PBR protocol allows the user to add random noise to the final bill, to hide usage patterns that could otherwise be deduced from the fee. The rebate protocol supports

¹⁰ In [5], the user keys sk_U and pk_U are used to create and verify a signature on the payment message. This intuitively facilitates non-repudiation and non-exculpability properties, but is not modeled by the ideal functionality. This carries over to our adaptations.

deposits, anonymous payments using e-cash, and negative bill noise, while ensuring that the funds paid always cover the cost of consumption. The use of noise, however, comes at a cost, as it is money that the user has to pay upfront as a deposit and cannot invest elsewhere. Consequently, we adapt the differential privacy framework to study how much noise is needed to protect specific consumption windows at different security levels. The differential privacy framework protects users against worse case outcomes – we leave as an open problem crafting more economical noise regimes to protect privacy by making further assumption about the users’ typical behavior.

Acknowledgment. We would like to thank Claudia Diaz, Carmela Troncoso, Cedric Fournet, and Jorn Lapon for discussions that were most helpful in the preparation of this work. This work was supported in part by GOA TENSE (GOA/11/007), and by the IAP Programme P6/26 BCRYPT (Belgian Science Policy). Alfredo Rial is a Research Foundation - Flanders (FWO) doctoral researcher.

References

1. Anderson, R., Fuloria, S.: On the security economics of electricity metering. In: The Ninth Workshop on the Economics of Information Security. (2010)
2. Molina-Markham, A., Shenoy, P., Fu, K., Cecchet, E., Irwin, D.: Private memoirs of a smart meter. In: 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (BuildSys 2010), Zurich, Switzerland (November 2010)
3. Balasch, J., Rial, A., Troncoso, C., Preneel, B., Verbaauwhede, I., Geuens, C.: Pretp: Privacy-preserving electronic toll pricing. In: USENIX Security Symposium, USENIX Association (2010) 63–78
4. Bohli, J.M., Sorge, C., Ugus, O.: A privacy model for smart metering. In: Communications Workshops (ICC), 2010 IEEE International Conference on. (may 2010) 1–5
5. Rial, A., Danezis, G.: Privacy-preserving smart metering. Technical Report MSR-TR-2010-150, Microsoft Research (November 2010)
6. Lipner, S.B.: A comment on the confinement problem. In: Proceedings of the fifth ACM symposium on Operating systems principles. SOSP ’75, New York, NY, USA, ACM (1975) 192–196
7. Dwork, C.: Differential privacy. In Bugliesi, M., Preneel, B., Sassone, V., Wegener, I., eds.: ICALP (2). Volume 4052 of Lecture Notes in Computer Science., Springer (2006) 1–12
8. Dwork, C.: Differential privacy in new settings. In Charikar, M., ed.: SODA, SIAM (2010) 174–183
9. Ghosh, A., Roughgarden, T., Sundararajan, M.: Universally utility-maximizing privacy mechanisms. In: Proceedings of the 41st annual ACM symposium on Theory of computing. STOC ’09, New York, NY, USA, ACM (2009) 351–360
10. Danezis, G., Kohlweiss, M., Rial, A.: Differentially private billing with rebates. Cryptology ePrint Archive, Report 2011/134 (2011) <http://eprint.iacr.org/>.
11. Barak, B., Chaudhuri, K., Dwork, C., Kale, S., McSherry, F., Talwar, K.: Privacy, accuracy, and consistency too: a holistic solution to contingency table release. In Libkin, L., ed.: PODS, ACM (2007) 273–282
12. Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. In: FOCS. (2001) 136–145
13. Kusters, R.: Simulation-based security with inexhaustible interactive turing machines. In: Computer Security Foundations Workshop, 2006. 19th IEEE, IEEE (2006) 12–320
14. Camenisch, J., Lysyanskaya, A.: A signature scheme with efficient protocols. In: In SCN 2002, volume 2576 of LNCS, Springer (2002) 268–289