

Gradual Release of Sensitive Data under Differential Privacy

Fragkiskos Koufogiannis,* Shuo Han[†] and George J. Pappas[‡]

Abstract. We introduce the problem of releasing private data under differential privacy when the privacy level is subject to change over time. Existing work assumes that privacy level is determined by the system designer as a fixed value before private data is released. For certain applications, however, users may wish to relax the privacy level for subsequent releases of the same data after either a re-evaluation of the privacy concerns or the need for better accuracy. Specifically, given a database containing private data, we assume that a response y_1 that preserves ϵ_1 -differential privacy has already been published. Then, the privacy level is relaxed to ϵ_2 , with $\epsilon_2 > \epsilon_1$, and we wish to publish a more accurate response y_2 while the joint response (y_1, y_2) preserves ϵ_2 -differential privacy. How much accuracy is lost in the scenario of gradually releasing two responses y_1 and y_2 compared to the scenario of releasing a single response that is ϵ_2 -differentially private? Our results consider the more general case with multiple privacy level relaxations and show that there exists a composite mechanism that achieves *no loss* in accuracy.

We consider the case in which the private data lies within \mathbb{R}^n with an adjacency relation induced by the ℓ_1 -norm, and we initially focus on mechanisms that approximate identity queries. We show that the same accuracy can be achieved in the case of gradual release through a mechanism whose outputs can be described by a *lazy Markov stochastic process*. This stochastic process has a closed form expression and can be efficiently sampled. Moreover, our results extend beyond identity queries to a more general family of privacy-preserving mechanisms. To this end, we demonstrate the applicability of our tool to multiple scenarios including Google’s project RAPPOR, trading of private data, and controlled transmission of private data in a social network. Finally, we derive similar results for the approximated differential privacy.

1 Introduction

Differential privacy is a framework that provides rigorous privacy guarantees for the release of sensitive data. The intrinsic trade-off between the privacy guarantees and accuracy of the privacy-preserving mechanism is controlled by the privacy level $\epsilon \in [0, \infty)$; smaller values of ϵ imply stronger privacy and less accuracy. Specifically, *end users*, who are interested in the output of the mechanism, demand acceptable accuracy

*Department of Electrical and Systems Engineering, University of Pennsylvania.
<mailto:fkouf@seas.upenn.edu>

[†]Department of Electrical and Systems Engineering, University of Pennsylvania.
<mailto:hanshuo@seas.upenn.edu>

[‡]Department of Electrical and Systems Engineering, University of Pennsylvania.
<mailto:pappasg@seas.upenn.edu>

of the privacy-preserving mechanism, whereas, *owners* of sensitive data are interested in strong enough privacy guarantees.

Existing work on differential privacy assumes that the privacy level is determined prior to release of any data and remains constant throughout the life of the privacy-preserving mechanism. However, for certain applications, the privacy level may need to be revised *after* data has been released, due to either users' need for improved accuracy or after owners' re-evaluation of the privacy concerns. One such application is trading of private data, where the owners re-evaluate their privacy concerns after monetary payments. Specifically, the end users initially access private data under ϵ_1 privacy guarantees and they later decide to “buy” more accurate data, relax privacy level to ϵ_2 , and enjoy better accuracy. Furthermore, the need for more accurate responses may dictate a change in the privacy level. In particular, a database containing sensitive data is persistent over time; e.g. a database of health records contains the same patients with the same health history over several years. Future uses of the database may require better accuracy, especially, after a threat is suspected (e.g. virus spread, security breach). These two example applications share the same core questions.

Is it possible to release a preliminary response with ϵ_1 -privacy guarantees and, later, release a more accurate and less private response with overall ϵ_2 -privacy guarantees? How would this scenario compare to publishing a single response under ϵ_2 -privacy guarantees? In fact, is the performance of the second response damaged by the preliminary one?

Composition theorems (McSherry and Talwar, 2007) provide a simple, but suboptimal, solution to gradually releasing sensitive data. Given an initial privacy level ϵ_1 , a noisy, privacy-preserving response y_1 is generated. Later, the privacy level is relaxed to a new value ϵ_2 , where $\epsilon_2 > \epsilon_1$, and a new response y_2 is published. For an overall privacy level of ϵ_2 , it is sufficient for the second response y_2 to be $(\epsilon_2 - \epsilon_1)$ -private, according to the composition theorem. Therefore, the accuracy of the second response deteriorates because of the initial release y_1 .

More sophisticated approaches can be defined such that the suboptimality of subsequent responses remains bounded. For instance, initially, we independently generate the responses $\{y_i\}_{i=-\infty}^{\infty}$, where y_i is 2^i -private. For an ϵ_1 -private response, we release the sequence $\{y_i\}_{i=-\infty}^{\lfloor \log_2 \epsilon_1 \rfloor - 1}$. According to composition theorems, this sequence is ϵ_1 -private and its accuracy is no worse than the accuracy of the last term $y_{\lfloor \log_2 \epsilon_1 \rfloor - 1}$, which is that of an $\frac{\epsilon_1}{2}$ -private mechanism. As soon as the privacy level is relaxed from ϵ_1 to ϵ_2 , more elements of the sequence are released. Such a setting partially handles the loss of accuracy in gradually relaxing the privacy level.

However, in this work, we derive an exact solution where there is *no* accuracy loss incurred. This mechanism employs correlation between successive responses, and, to the best of our knowledge, is the first mechanism that performs gradual release of sensitive data without loss of accuracy.

1.1 Our Results

This work introduces the problem of gradually releasing sensitive data. Our results focus on the case of vector-valued sensitive data $u \in \mathbb{R}^n$ with an ℓ_1 -norm adjacency relation. Our first result states that, for the one-dimensional ($n = 1$) identity query, there is an algorithm which relaxes privacy in two steps without sacrificing any accuracy. Although our technical treatment focuses on identity queries, our results are applicable to a broader family of queries; in particular, to a family of differentially private mechanisms that add Laplace noise. We also prove the *Markov property* for this algorithm and, thus, we can easily relax privacy in any number of steps; time complexity is linear in number of steps and memory complexity is constant. These two results provide a different perspective of differential privacy, and lead to the definition of a *lazy* Markov stochastic process indexed by the privacy level ϵ . Gradually releasing sensitive data is performed by sampling once from this stochastic process. We also extend the results to the high-dimensional case.

On a theoretical level, within differential privacy, our contributions explore the setting of a privacy level ϵ that *varies* over time. We focus on the mechanism that adds Laplace-distributed noise V_ϵ to the private data $u \in \mathbb{R}^n$:

$$Q_\epsilon u = u + \begin{bmatrix} v_\epsilon^{(1)} \\ \vdots \\ v_\epsilon^{(n)} \end{bmatrix}, \quad (1)$$

where ϵ is the privacy level and $\{v_\epsilon^{(i)}\}_{i=1}^n$ are independent and identically distributed samples from the stochastic process $\{V_\epsilon\}_{\epsilon>0}$ which has the following properties:

1. $\{V_\epsilon\}_{\epsilon>0}$ is Markov, i.e., for privacy levels $\epsilon_3 \geq \epsilon_2 \geq \epsilon_1 > 0$, the random variables V_{ϵ_1} and V_{ϵ_3} are independent conditioned on the value of V_{ϵ_2} . Formally: $V_{\epsilon_1} \perp V_{\epsilon_3} | V_{\epsilon_2}$.
2. V_ϵ is Laplace-distributed: $\mathbb{P}(V_\epsilon = x) = \frac{\epsilon}{2} e^{-\epsilon|x|}$.
3. $\{V_\epsilon\}_{\epsilon>0}$ is *lazy*, i.e., there is positive probability of not changing value:

$$\mathbb{P}(V_{\epsilon_1} = x | V_{\epsilon_2} = y) = \left(\frac{\epsilon_1}{\epsilon_2}\right)^2 \delta(x - y) + \left(1 - \left(\frac{\epsilon_1}{\epsilon_2}\right)^2\right) \frac{\epsilon_1}{2} e^{-\epsilon_1|x-y|}, \text{ where } \epsilon_2 \geq \epsilon_1 > 0, \quad (2)$$

where δ is Dirac's delta function.¹ For a fixed ϵ , mechanism (1) reduces to the Laplace mechanism. Equivalently, the conditional distribution $\mathbb{P}(V_{\epsilon_2} = y | V_{\epsilon_1} = x)$ can be derived from Distribution (2).

¹Dirac's delta $\delta(x)$ is a generalized function which is zero everywhere except for $x = 0$ and its integral is $\int_{-\infty}^{\infty} \delta(x) dx = 1$.

Mechanism (1) has the following properties and, thus, performs gradual release of private data:

- **Privacy:** For any set of privacy levels $\{\epsilon_i\}_{i=1}^m$, the mechanism that responds with $\{Q_{\epsilon_i}u\}_{i=1}^m$ is $(\max_{i=1}^m \epsilon_i)$ -private.
- **Accuracy:** For a fixed ϵ , the mechanism Q_ϵ is the optimal ϵ -private mechanism.

In practice, gradual release of private data is achieved by sampling the stochastic process $\{V_\epsilon\}_{\epsilon>0}$:

1. Draw a single sample $\{v_\epsilon\}_{\epsilon>0}$ from the stochastic process $\{V_\epsilon\}_{\epsilon>0}$.
2. Compute the signal $y_\epsilon = u + v_\epsilon$, $\epsilon > 0$.
3. For ϵ_1 -privacy guarantees, release the random variable y_{ϵ_1} .
4. Once privacy level is relaxed from ϵ_1 to ϵ_2 , where $\epsilon_2 \geq \epsilon_1$, release the random variable y_{ϵ_2} .
5. In order to relax privacy level in an arbitrarily many times, $\epsilon_1 \rightarrow \epsilon_2 \rightarrow \dots \rightarrow \epsilon_m$, repeat the last step.

More formally, our main result derives a composite mechanism that gradually releases private data by relaxing the privacy level in an arbitrary number of steps.

Theorem 1 (A. Gradual Privacy as a Composite Mechanism). *Let \mathbb{R}^n be the space of private data equipped with an ℓ_1 -norm adjacency relation. Consider m privacy levels $\{\epsilon_i\}_{i=1}^m$ such that $0 \leq \epsilon_1 \leq \dots \leq \epsilon_m$ which successively relax the privacy level. Then, there exists a composite mechanism Q of the form*

$$Qu := (u + V_1, \dots, u + V_m), \quad (3)$$

such that:

1. The restriction of the mechanism Q to the first j coordinates $(u + V_1, \dots, u + V_j)$ is ϵ_j -private, for any $j \in \{1, \dots, m\}$.
2. Each coordinate $j \in \{1, \dots, m\}$ of the mechanism $u + V_j$ achieves the optimal mean-squared error $\mathbb{E}\|V_j\|_2^2 = 2n/\epsilon_j^2$.

The mechanism that satisfies Theorem 1A has a closed-form expression which allows for computationally lightweight implementation. Furthermore, instead of considering m discrete privacy levels, Theorem 1B considers the *continuum* of privacy levels $\epsilon \in [0, \infty)$. To this end, our results are more succinctly stated in terms of a stochastic process $\{V_\epsilon\}_{\epsilon>0}$. A composite mechanism is recovered from the stochastic process by sampling the process $\{V_\epsilon\}_{\epsilon>0}$ at a finite set of privacy levels $\{\epsilon_i\}_{i=1}^m$.

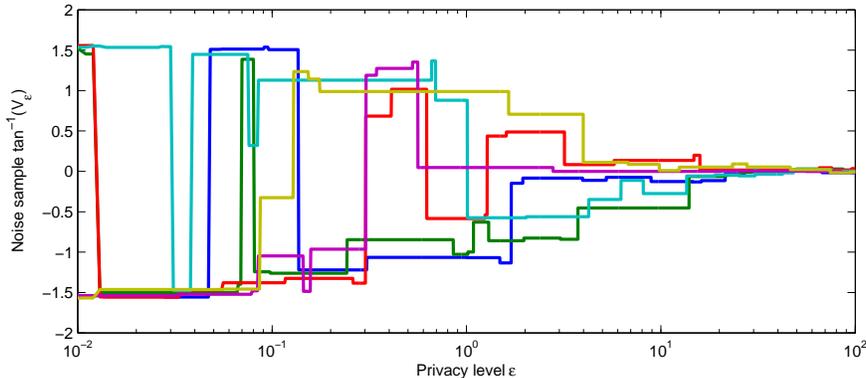


Figure 1: Gradual release of identity queries is achieved with the use of the stochastic process V_ϵ for $\epsilon \geq 0$, several samples of which are shown with different colors. In practice, a single sample of this process is drawn and, for a privacy level ϵ_0 and private data u , the noise version $u + V_{\epsilon_0}$ is published. For tight values of privacy ($\epsilon \rightarrow 0$), high values of noise ($|\tan^{-1} V_\epsilon| \rightarrow \frac{\pi}{2}$) are returned, whereas, almost zero samples ($V_\epsilon \rightarrow 0$) are returned for large privacy budgets ($\epsilon \rightarrow \infty$). The process V_ϵ is Markov; future samples depend only on the current value of the process which eases implementation. Furthermore, the process is lazy; the value of the process changes only a few times.

Theorem 1 (B. Gradual Privacy as a Stochastic Process). *Let \mathbb{R}^n be the space of privacy data equipped with the ℓ_1 -norm. Then, there exists a stochastic process $\{V_\epsilon\}_{\epsilon>0}$ that defines the family of mechanisms Q_ϵ parametrized by ϵ :*

$$Q_\epsilon u := u + V_\epsilon, \quad \epsilon \in (0, \infty), \quad (4)$$

such that:

- **Privacy:** For any $\epsilon > 0$, the mechanism that releases the signal $\{u + V_\sigma\}_{\sigma \in (0, \epsilon]}$ is ϵ -private.
- **Accuracy:** The mechanism Q_ϵ that releases the random variable $u + V_\epsilon$ is the optimal ϵ -private mechanism, i.e. the noise sample V_ϵ achieves the optimal mean-squared error $\mathbb{E}\|V_\epsilon\|_2^2 = 2n/\epsilon_j^2$.

From a more practical point of view, our results are applicable to cases beyond identity queries. Specifically, as shown later in Corollary 15 in Section 4.2, our results are directly applicable to a broad family of privacy-preserving mechanisms that are built upon the Laplace mechanism and, informally, have the following form. The sensitive data is initially pre-processed, then, the Laplace mechanism is invoked, and, finally, a post-processing step occurs. Under the assumption that the pre-processing step is invariant of the privacy level—the pre-processing function should be the same for any privacy level—gradual release of sensitive data is possible. We demonstrate the

applicability of our results on Google’s RAPPOR project (Erlingsson et al., 2014), which analyzes software features that individuals use while respecting their privacy. In particular, if a software feature is suspected to be malicious, privacy level can be gradually relaxed and a more accurate analysis can be performed. On another direction, our results broaden the spectrum of applications of differential privacy. To this end, we present an application to social networks where users have different privacy concerns against close friends, acquaintances, and strangers.

We conclude our paper with a discussion of possible further work. Although present work focuses on mechanisms that add Laplace-distributed noise, we conjecture that the feasibility of gradually releasing sensitive data is a more general property of differential privacy.

1.2 Previous Work

Differential privacy is an active field of research and a rich spectrum of differential private mechanisms has appeared in the literature. The exponential mechanism (McSherry and Talwar, 2007) is a powerful and generic tool for building differential private mechanisms. In particular, mechanisms that efficiently approximate linear (counting) queries have received a lot of attention (Hardt and Talwar, 2010; Li et al., 2010; Ullman, 2013). Besides counting queries, privacy-aware versions of more complex quantities have been introduced such as signal filtering (Le Ny and Pappas, 2014), optimization problems (Gupta et al., 2010; Han et al., 2014; Huang et al., 2015), machine learning (Shokri and Shmatikov, 2015), and allocation problems (Hsu et al., 2014). In addition to the theoretical work, differential privacy has been deployed in software tools (Reed and Pierce, 2010).

The aforementioned work assumes that the privacy level ϵ is a designer’s choice that is held fixed throughout the life of the privacy-aware mechanism. To the best of our knowledge, our work is the first approach that considers privacy-aware mechanisms with a varying privacy level ϵ . Gradually releasing private data resembles the setting of differential privacy under continuous observation, which was first studied in Dwork et al. (2010). In that setting of Dwork et al. (2010), the privacy level remains fixed while more sensitive data is being added to the database and more responses are released. In contrast, our setting assumes that both the sensitive data and the quantity of interest are fixed and the privacy level ϵ is varying.

Gradual release of sensitive data is closely related to optimality results. Work in Hardt and Talwar (2010) established optimality results in an asymptotic sense (with the size of the database). Instead, our work requires *exact* optimality results and, therefore, is presented within a tighter version of differential privacy that was explored in Chatzikokolakis et al. (2013), Koufogiannis et al. (2015), where exact optimality results exist. This tighter notion which is targeted for metric spaces and we call *Lipschitz privacy*, allows for the use of optimization techniques and calculus tools. Prior work on Lipschitz privacy includes the *exact* optimality of the Laplace mechanism is established under Lipschitz privacy (Koufogiannis et al., 2015; Wang et al., 2014).

On a more technical level, most prior work on differential privacy (Gupta et al., 2010; Han et al., 2014; Hsu et al., 2014) introduces differential private mechanisms that are built upon the Laplace mechanism and variations of it. Although building upon the Laplace mechanism limits the solution space, there is a good reason for doing so. Specifically, for non-trivial applications, the space of probability measures can be extremely rich and hard to deal with. Technically, our approach deviates from prior work by searching over the whole space of differential private mechanisms. Work in Xiao et al. (2011) is another example that proposes a non-Laplace-based private mechanism. Specifically, Xiao et al. (2011) considers a query with a single privacy level and suggests incrementally adding correlated noise when answering a query in order to reduce the resulting relative error. Although we also allow for correlated noise samples, our work substantially differs since we consider a different problem which includes multiple privacy levels.

2 Background Information

2.1 Differential Privacy

The framework of differential privacy (Dwork et al., 2006) dictates that, whenever sensitive data is accessed, a noisy response is returned. The statistics of the injected noise are deliberately designed to ensure two things. First, an adversary that observes the noisy response cannot *confidently infer* the original sensitive data. For a survey of results on differential privacy, we refer the reader to Dwork (2006) and Dwork and Roth (2013). The privacy level is parameterized by $\epsilon \in [0, \infty)$, where smaller values of ϵ imply stronger privacy guarantees. Second, the noisy response can still be used as a surrogate of the exact response *without severe performance degradation*. On the other hand, the accuracy of the noisy response is quantified by the mean-squared error from the exact response.

Work in Dwork et al. (2006) defined differential privacy, which provides strong privacy guarantees against a powerful adversary.

Definition 2 (Differential Privacy). *Let \mathcal{U} be a set of private data, $\mathcal{A} \subseteq \mathcal{U}^2$ be a symmetric binary relation (called adjacency relation) and \mathcal{Y} be a set of possible responses. For $\epsilon > 0$, the randomized mapping $Q : \mathcal{U} \rightarrow \Delta(\mathcal{Y})$ (called mechanism) is ϵ -differentially private if*

$$\mathbb{P}(Qu \in \mathcal{S}) \leq e^\epsilon \mathbb{P}(Qu' \in \mathcal{S}), \forall (u, u') \in \mathcal{A}, \forall \mathcal{S} \subseteq \mathcal{Y}. \quad (5)$$

Remark 1. *We assume the existence of a rich-enough σ -algebra $M \subseteq 2^{\mathcal{Y}}$ on the set of possible responses \mathcal{Y} . Then, $\Delta(\mathcal{Y})$ denotes the set of probability measures over (M, \mathcal{Y}) .*

Let² $y \sim Qu$ be a noisy response produced by the ϵ -differentially private mechanism Q . For brevity, we say that “output y preserves ϵ -privacy of the input u ”.

²We denote the output of the mechanism Q on the private data u with Qu instead of $Q(u)$ in order to simplify expressions.

The adjacency relation \mathcal{A} captures the aspects of the private data u that are deemed sensitive. Consider a scheme with n users, where each user i contributes her real-valued private data $u_i \in \mathbb{R}$, and a private database $u = [u_1, \dots, u_n] \in \mathbb{R}^n$ is composed. For $\alpha > 0$, an adjacency relation that captures the participation of a single individual to the aggregating scheme is defined as:

$$(u, u') \in \mathcal{A}_{\ell_0} \Leftrightarrow \exists j \text{ s.t. } u_i = u'_i, \forall i \neq j \text{ and } |u_j - u'_j| \leq \alpha. \quad (6)$$

Adjacency relation \mathcal{A}_{ℓ_0} can be relaxed to \mathcal{A}_{ℓ_1} , which is induced by the ℓ_1 -norm and is defined as:

$$(u, u') \in \mathcal{A}_{\ell_1} \Leftrightarrow \|u - u'\|_1 \leq \alpha, \quad (7)$$

where it holds that $\mathcal{A}_{\ell_0} \subseteq \mathcal{A}_{\ell_1}$.

Resilience to post-processing establishes that any post-processing on the output of an ϵ -differentially private mechanism cannot hurt the privacy guarantees.

Proposition 3 (Resilience to Post-Processing). *Let $Q : \mathcal{U} \rightarrow \Delta(\mathcal{Y})$ be an ϵ -differentially private mechanism and $g : \mathcal{Y} \rightarrow \mathcal{Z}$ be a possibly randomized function. Then, the mechanism $g \circ Q$ is also ϵ -differentially private.*

More complicated mechanisms can be defined from simple ones using the composition theorem.

Proposition 4 (Composition). *Let mechanisms $Q_1, Q_2 : \mathcal{U} \rightarrow \Delta(\mathcal{Y})$ respectively satisfy ϵ_1 and ϵ_2 -differential privacy. Then, the composite mechanism $Q : \mathcal{U} \rightarrow \Delta(\mathcal{Y}^2)$ defined by $Q = (Q_1, Q_2)$ is $(\epsilon_1 + \epsilon_2)$ -differentially private.*

Proposition 4 provides privacy guarantees whenever the *same* sensitive data is repeatedly used. Moreover, the resulting privacy level $\epsilon_1 + \epsilon_2$ given by Proposition 4 is an upper bound and can severely over-estimate the actual privacy level. The mechanism presented in this paper introduces correlation between mechanisms Q_1 and Q_2 , so that it provides much stronger privacy guarantees.

2.2 Lipschitz Privacy

Lipschitz privacy (Chatzikokolakis et al., 2013; Koufogiannis et al., 2015) is a slightly stronger version of differential privacy and is often used when the data is defined on metric spaces.

Definition 5 (Lipschitz Privacy). *Let (\mathcal{U}, d) be a metric space and \mathcal{Y} be the set of possible responses. For $\epsilon > 0$, the mechanism Q is ϵ -Lipschitz private if the following Lipschitz condition holds:*

$$|\ln \mathbb{P}(Qu \in \mathcal{S}) - \ln \mathbb{P}(Qu' \in \mathcal{S})| \leq \epsilon d(u, u'), \quad \forall u, u' \in \mathcal{U}, \forall \mathcal{S} \subseteq \mathcal{Y}. \quad (8)$$

Lipschitz privacy is closely related to the original definition of differential privacy, where the adjacency relation \mathcal{A} in differential privacy is defined through the metric d . In fact, any Lipschitz private mechanism is also differentially private.

Proposition 6. *For any $\alpha > 0$, an ϵ -Lipschitz private mechanism Q is $\alpha\epsilon$ -differentially private under the adjacency relation \mathcal{A} :*

$$(u, u') \in \mathcal{A} \Leftrightarrow d(u, u') \leq \alpha. \quad (9)$$

Adjacency relation \mathcal{A}_{ℓ_1} defined in (7) can be captured by the ℓ_1 -norm under the notion of Lipschitz privacy; the metric d is $d(u, u') = \|u - u'\|_1$.

Our results are stated within the Lipschitz privacy framework. Proposition 6 implies that our privacy results remain valid within the framework of differential privacy. For brevity, we call an ϵ -Lipschitz private mechanism as ϵ -private and imply that a differentially private mechanism can be derived.

Similar to differential privacy, Lipschitz privacy is preserved under post-processing (Proposition 3) and composition of mechanisms (Proposition 4). Compared to differential privacy, Lipschitz privacy is more convenient to work with when the data and adjacency relation are defined on a metric space, which allows for the use of calculus tools. Under mild assumptions, the Lipschitz condition (8) is equivalent to a derivative bound. In particular, for $\mathcal{U} = \mathbb{R}^n$ equipped with the metric induced by the norm $\|\cdot\|$, a mechanism Q is ϵ -Lipschitz private if

$$\|\nabla_u \ln \mathbb{P}(Qu = y)\|_* \leq \epsilon, \quad (10)$$

where $\|\cdot\|_*$ is the dual norm of $\|\cdot\|$. In practice, we check condition (10) to establish the privacy properties of mechanism Q .

Originally, differential privacy was stated in terms of databases, where two databases are considered adjacent if they differ by a single row (Dwork et al., 2006). Subsequent work (e.g. Geng and Viswanath (2012), Le Ny and Pappas (2014)) considers arrays of real-valued private data and the adjacency relation that we employ here. In practice, the two approaches are equivalent. For example, consider the private data $u \in \{0, 1\}^n$ which is an aggregation of a private bit across n users. This bit can indicate a private attribute of the user such as her health or her actual participation to the aggregation scheme. By setting $\alpha = 1$ in Theorem 6, we retrieve the original version of differential privacy.

2.3 Optimality of the Laplace Mechanism

Computing the optimal private mechanism for a fixed privacy level ϵ is considered an open problem. The Laplace mechanism is a special instance of the exponential mechanism (McSherry and Talwar, 2007) for real spaces (\mathbb{R}^n, ℓ_1) .

Definition 7 (Laplace Mechanism). *Let (\mathbb{R}^n, ℓ_1) be the space of private data. The Laplace mechanism is defined as:*

$$Qu = u + V, \text{ where } V \sim e^{-\epsilon\|V\|_1}. \quad (11)$$

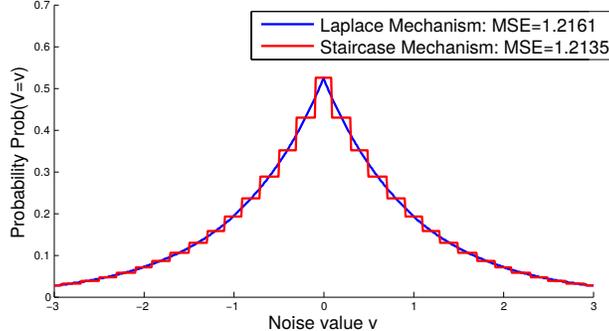


Figure 2: The staircase mechanism is the optimal ϵ -differential private mechanism, whereas the Laplace mechanism is the optimal ϵ -Lipschitz private mechanism. The two distributions are similar and there is only a small performance gap. Therefore, the Laplace distribution is often used in practice.

The Laplace mechanism can be shown to be ϵ -differentially private. In general, however, the Laplace mechanism is suboptimal in the sense of minimum mean-squared error. For the single-dimensional case, the staircase mechanism (Geng and Viswanath, 2012) is the optimal ϵ -differentially private mechanism; the mechanism which adds noise V whose distribution is shown in Figure 2. However, the Laplace mechanism is widely used and has several optimality results. Specifically, it is proven to be “universally” optimal—optimally approximating a single linear query, under any prior on private data—and, additionally, it is the optimal ϵ -Lipschitz private mechanism in the sense of both minimum entropy (Wang et al., 2014) and minimum mean-squared error (Koufogiannis et al., 2015), whereas the staircase mechanism fails to satisfy Lipschitz privacy due to its discontinuous probability density function.

Theorem 8 (Optimality of Laplace (Koufogiannis et al., 2015)). *Consider the ϵ -Lipschitz private (in (\mathbb{R}^n, ℓ_1)) mechanism $Q : \mathbb{R}^n \rightarrow \Delta(\mathbb{R}^n)$ of the form $Qu = u + V$, with $V \sim g(V) \in \Delta(\mathbb{R}^n)$. Then, the Laplace mechanism that adds noise with density $l_1^n(v) = (\frac{\epsilon}{2})^n e^{-\epsilon\|v\|_1}$ minimizes the mean-squared error. Namely, for any density g , we have:*

$$\mathbb{E}\|Qu - u\|_2^2 = \mathbb{E}_{V \sim g} \|V\|_2^2 \geq \mathbb{E}_{V \sim l_1^n} \|V\|_2^2 = \frac{2n}{\epsilon^2}. \quad (12)$$

The optimal private mechanism characterizes the privacy-performance trade-off and is required for gradually releasing sensitive data. Thus, optimality of the Laplace mechanism in Theorem 8 is a key ingredient in our results and renders the problem tractable.

3 Gradual Release of Private Data

We now formulate the problem of gradually releasing private data. Initially, we focus on a single privacy level relaxation from ϵ_1 to ϵ_2 and a single-dimensional space of

private data $\mathcal{U} = \mathbb{R}$. Subsections 3.2 and 3.3 present extensions to high-dimensional spaces and multiple rounds of privacy level relaxations, respectively.

Consider two privacy levels ϵ_1 and ϵ_2 with $\epsilon_2 \geq \epsilon_1 > 0$. We wish to design a composite mechanism $Q_{\epsilon_1 \rightarrow \epsilon_2} : \mathcal{U} \rightarrow \Delta(\mathcal{Y} \times \mathcal{Y})$ that performs gradual release of data. The first and second coordinates respectively refer to the initial ϵ_1 -private and the subsequent ϵ_2 -private responses. In practice, given privacy levels ϵ_1 and ϵ_2 and an input $u \in \mathcal{U}$, we sample (y_1, y_2) from the distribution $Q_{\epsilon_1 \rightarrow \epsilon_2} u$. Initially, only coordinate y_1 is published satisfying ϵ_1 -privacy guarantees. Once privacy level is relaxed to ϵ_2 , response y_2 is released as a more accurate response of the *same* query on the *same* private data.

An adversary that wishes to infer the private input u eventually has access to both responses y_1 and y_2 . Therefore, the pair (y_1, y_2) needs to satisfy ϵ_2 -privacy. On the other hand, an honest user wishes to maximize the accuracy of the response and, therefore, she is tempted to use an estimator $y_M = \theta(y_1, y_2)$ and infer a more accurate response y_M . In order to relieve honest users from any computational burden and without loss of generality, we wish that the best estimator to be the truncation:

$$\theta(y_1, y_2) = y_2. \quad (13)$$

The composition theorem (McSherry and Talwar, 2007) provides a trivial, yet highly conservative, approach. Specifically, compositional rules imply that, if y_1 satisfies ϵ_1 -privacy and (y_1, y_2) satisfies ϵ_2 -privacy, coordinate y_2 itself should be $(\epsilon_2 - \epsilon_1)$ -private. In the extreme case that $\epsilon_2 - \epsilon_1 = \delta \ll 1$, response y_2 alone is expected to be δ -private and, therefore, is highly corrupted by noise. This is unacceptable, since estimator (13) yields an even noisier response than the initial response y_1 . Even if honest users are expected to compute more complex estimators than the truncation one in (13), the approach dictated by composition theorem can still be unsatisfactory.

Specifically, consider the following two scenarios:

1. An ϵ_1 -private response y_1 is initially released. Once privacy level is relaxed from ϵ_1 to ϵ_2 , an supplementary response y_2 is released.
2. No response is initially released. Response y_2 is released as soon as the privacy level is relaxed to ϵ_2 .

Then, there is no guarantee that the best estimator $\theta(y_1, y_2)$ in Scenario 1 will match the accuracy of the response y_2 in Scenario 2. An accuracy gap between the two scenarios would severely impact differential privacy. Specifically, the system designer needs to be strategic when choosing a privacy level. Differently stated, a market of private data based on composition theorems would exhibit *friction*.

The key idea to overcome this friction is to introduce correlation between responses y_1 and y_2 . In this work, we focus on Euclidean spaces $\mathcal{U} = \mathbb{R}^n$ and mechanisms $Qu = u + V$ that approximate the identity query $q(u) = u$. Our main result states that a *frictionless* market of private data is feasible and Scenarios 1 and 2 are equivalent. This result has multi-fold implications:

- A system designer is not required to be strategic with the choice of the privacy level. Specifically, she can initially under-estimate the required privacy level with ϵ_1 and she can later fine-tune it to ϵ_2 without hurting the accuracy of the final response.
- A privacy data market can exist and private data can be traded “*by the pound*”. An ϵ_1 -private response y_1 can be initially purchased. Next, a supplementary payment can be made in return for a privacy level relaxation to ϵ_2 and a refined response y_2 . The accuracy of the refined response y_2 is, then, unaffected by the initial transaction and is controlled only by the final privacy level ϵ_2 .

More concretely, given privacy levels ϵ_1 and ϵ_2 with $\epsilon_2 > \epsilon_1$, we wish to design a composite mechanism $Q_{\epsilon_1 \rightarrow \epsilon_2} : \mathcal{U} \rightarrow \mathcal{Y} \times \mathcal{Y}$ with the following properties:

1. The restriction of $Q_{\epsilon_1 \rightarrow \epsilon_2}$ to the first coordinate should match the performance of the optimal ϵ_1 -private mechanism Q_{ϵ_1} . More restrictively, the first coordinate of the composite mechanism $Q_{\epsilon_1 \rightarrow \epsilon_2}$ should be distributed identically to the optimal ϵ_1 -private mechanism Q_{ϵ_1} :

$$\mathbb{P}(Q_{\epsilon_1 \rightarrow \epsilon_2} u \in \mathcal{S} \times \mathcal{Y}) = \mathbb{P}(Q_{\epsilon_1} u \in \mathcal{S}), \forall u \in \mathcal{U} \text{ and } \mathcal{S} \subseteq \mathcal{Y} \quad (14)$$

2. The restriction of $Q_{\epsilon_1 \rightarrow \epsilon_2}$ to the first coordinate should be ϵ_1 -private. This property is imposed by constraint 1.
3. The restriction of $Q_{\epsilon_1 \rightarrow \epsilon_2}$ to the second coordinate should match the performance of the optimal ϵ_2 -private mechanism Q_{ϵ_2} . Similarly to the first coordinate, the second coordinate of the composite mechanism $Q_{\epsilon_1 \rightarrow \epsilon_2}$ must be distributed identically to the optimal ϵ_2 -private mechanism Q_{ϵ_2} :

$$\mathbb{P}(Q_{\epsilon_1 \rightarrow \epsilon_2} u \in \mathcal{Y} \times \mathcal{S}) = \mathbb{P}(Q_{\epsilon_2} u \in \mathcal{S}), \forall u \in \mathcal{U} \text{ and } \mathcal{S} \subseteq \mathcal{Y} \quad (15)$$

4. Once both coordinates are published, ϵ_2 -privacy should be guaranteed. According to Lipschitz privacy, the requirement is stated as follows:

$$\mathbb{P}(Q_{\epsilon_1 \rightarrow \epsilon_2} u \in \mathcal{S}) \text{ is } \epsilon_2\text{-Lipschitz in } u, \text{ for all } \mathcal{S} \subseteq \mathcal{Y}^2. \quad (16)$$

Equations (14) and (15) require knowledge of the optimal ϵ -private mechanism. In general, computing the ϵ -private mechanism that maximizes a reasonable performance criterion is still an open problem. Theorem 8 establish the optimality of the Laplace mechanism as the optimal private approximation of the identity query.

3.1 Single-Dimensional Case

Initially, we consider the single-dimensional case where $\mathcal{U} = \mathbb{R}$ equipped with the absolute value. Theorem 8 establishes the optimal ϵ -private mechanism that is required by Equations (14) and (15):

$$Q_\epsilon u = u + V, \text{ where } V \sim e^{-\epsilon|V|}. \quad (17)$$

Mechanism (17) minimizes the mean-squared error from the identity query among all ϵ -private mechanisms that use additive and independent of the database noise:

$$\mathbb{E}_{V \sim e^{-\epsilon|V|}} (Q_\epsilon u - u)^2 \quad (18)$$

Theorem 9 establishes the existence of a composite mechanism that relaxes privacy from ϵ_1 to ϵ_2 *without any loss of performance*.

Theorem 9. *Consider privacy levels ϵ_1 and ϵ_2 with $\epsilon_2 \geq \epsilon_1 > 0$, and mechanisms of the form:*

$$Q_1 u := u + V_1 \text{ and } Q_2 u := u + V_2, \text{ with } (V_1, V_2) \sim g, \quad (19)$$

for some probability density $g \in \Delta(\mathbb{R}^2)$. Then, for $g = l_{\epsilon_1, \epsilon_2}$ with:

$$l_{\epsilon_1, \epsilon_2}(x, y) = \frac{\epsilon_1^2}{2\epsilon_2} e^{-\epsilon_2|y|} \delta(x - y) + \frac{\epsilon_1(\epsilon_2^2 - \epsilon_1^2)}{4\epsilon_2} e^{-\epsilon_1|x-y| - \epsilon_2|y|}, \quad (20)$$

where δ is the Dirac delta function, the following properties hold:

1. The mechanism Q_1 is ϵ_1 -private.
2. The mechanism Q_1 is optimal, i.e. Q_1 minimizes the mean-squared error $\mathbb{E}V_1^2$.
3. The mechanism (Q_1, Q_2) is ϵ_2 -private.
4. The mechanism Q_2 is optimal, i.e. Q_2 minimizes the mean-squared error $\mathbb{E}V_2^2$.

Proof. We straightforwardly verify that Distribution (20) has the aforementioned properties. The exact details are deferred to Appendix A. \square

The form of the probability distribution $l_{\epsilon_1, \epsilon_2}$ defined in (20) seems unintuitive. Informally, distribution $l_{\epsilon_1, \epsilon_2}$ can be derived as follows. Regarding the privacy constraints, it suffices (but is not necessary) that $l_{\epsilon_1, \epsilon_2}$ can be decomposed such that $V_1 - V_2$ is independent of V_2 , i.e.

$$l_{\epsilon_1, \epsilon_2}(V_1, V_2) = h(V_1 - V_2) l_{\epsilon_2}(V_2), \quad (21)$$

for some distribution h . In that case, response $Q_1 u$ can be viewed as a randomized post-processing of response $Q_2 u$. Next, distribution h can be computed by expressing the constraints that V_1 and V_2 need to be Laplace-distributed as integral equations and solving them. Unfortunately, there is no intuitive reason for the existence of the delta function (also called *atom* or the exact form of Distribution (20)).

Single Round of Privacy Relaxation

Theorem 9 achieves gradual release of sensitive data in two steps, first with ϵ_1 -privacy and, then, with ϵ_2 -privacy. In practice, Theorem 9 can be used as follows:

- Given the private value $u \in \mathbb{R}$, sample noise $V_1 \sim e^{-\epsilon_1|V_1|}$ and release response $y_1 = u + V_1$, which is optimal and respects ϵ_1 -privacy.
- Once privacy level is relaxed from ϵ_1 to ϵ_2 , sample noise V_2 from the conditioned on V_1 distribution:

$$\mathbb{P}(V_2 = y|V_1 = x) = \frac{\epsilon_1}{\epsilon_2} e^{-(\epsilon_2 - \epsilon_1)|x|} \delta(y - x) + \frac{\epsilon_2^2 - \epsilon_1^2}{2\epsilon_2} e^{-\epsilon_1|y-x| - \epsilon_2|y| + \epsilon_1|x|}, \quad (22)$$

and release response $y_2 = u + V_2$. Distribution (22) is derived from the joint distribution (20) and ensures both that (y_1, y_2) is ϵ_2 -private and that V_2 is optimally distributed.

Conditional distribution (22) is shown in Figure 3. Note that for $\epsilon_2 = \epsilon_1$, Distribution (22) is reduced to a delta function:

$$\mathbb{P}(V_2 = y|V_1 = x) = \delta(x - y). \quad (23)$$

In words, for $\epsilon_2 = \epsilon_1$ no privacy relaxation effectively happens and, thus, no updated response is practically released. Moreover, for $\epsilon_2 \rightarrow \infty$, a limiting argument shows that Distribution 22 is reduced to:

$$\mathbb{P}(V_2 = y|V_1 = x) = \delta(y). \quad (24)$$

Practically, letting $\epsilon_2 \rightarrow \infty$ cancel any privacy constraints and the exact value of private data u can be released $y_2 = u$. For general values of ϵ_1 and ϵ_2 , Pearson's correlation coefficient decreases for more aggressive privacy relations, $\rho_{V_1, V_2} = \frac{\epsilon_1}{\epsilon_2}$. Algorithm 1 provides a simple and efficient way to sample V_2 given V_1 . In fact, the first ϵ_1 -private response can be generated without knowing the next privacy level ϵ_2 in advance.

Single Round of Privacy Tightening

Tightening the privacy level is impossible, since it implies revoking already released data. Nonetheless, generating a more private version of the same data is still useful in cases such as *private data trading*. Specifically, one can consider a scenario where the data owner sells her private data at a privacy level ϵ_2 to a *private data trader*. Then, the trader may decide to re-sell the private data under a tighter privacy level ϵ_1 to an end-user. In that case, distribution (20) can be sampled in the opposite direction. Specifically, noise V_2 is initially sampled, $V_2 \sim e^{-\epsilon_2|V_2|}$, and the ϵ_2 -private response $y_2 = u + V_2$ is released. Next, private data u is traded to a *different agent* under the stronger ϵ_1 -privacy guarantees. Noise sample V_1 is drawn from distribution

$$\mathbb{P}(V_1 = x|V_2 = y) = \left(\frac{\epsilon_1}{\epsilon_2}\right)^2 \delta(x - y) + \left(1 - \left(\frac{\epsilon_1}{\epsilon_2}\right)^2\right) \frac{\epsilon_1}{2} e^{-\epsilon_1|x-y|}, \quad (25)$$

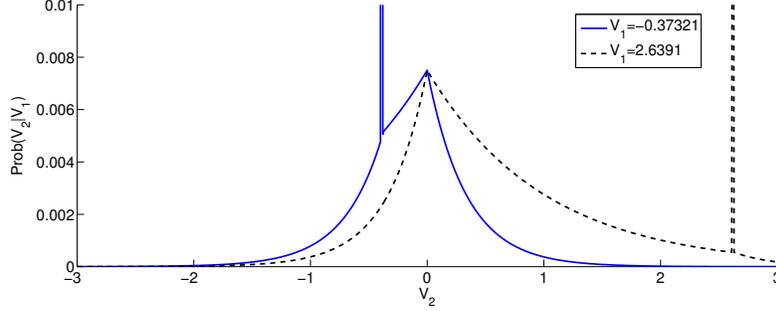


Figure 3: Gradual release of private data is performed in the following way. First, the ϵ_1 -private response $y_1 = u + V_1$ is released, where $V_1 \sim e^{-\epsilon_1|V_1|}$. Once privacy level is relaxed from $\epsilon_1 = 1$ to $\epsilon_2 = 2$, the supplementary response $y_2 = u + V_2$ is released, where the conditional distribution of V_2 given V_1 is shown above. The composite mechanism that releases (y_1, y_2) is ϵ_2 -private and V_2 is optimally distributed.

Algorithm 1 Sampling from Distribution (22) for the second noise sample $V_2 = y$ given the first noise sample $V_1 = x$ can be efficiently performed³.

Require: Privacy levels ϵ_1 and ϵ_2 , such that $\epsilon_2 > \epsilon_1 > 0$, and noise sample x .

```

function RELAXPRIVACY( $x, \epsilon_1, \epsilon_2$ )
  switch
    case with probability  $\frac{\epsilon_1}{\epsilon_2} e^{-(\epsilon_2 - \epsilon_1)|x|}$ :
      return  $y = x$ .
    case with probability  $\frac{\epsilon_2 - \epsilon_1}{2\epsilon_2}$ :
      draw  $z \sim \begin{cases} e^{(\epsilon_1 + \epsilon_2)z}, & \text{for } z \leq 0 \\ 0, & \text{otherwise.} \end{cases}$ 
      return  $y = \text{sgn}(x)z$ .
    case with probability  $\frac{\epsilon_1 + \epsilon_2}{2\epsilon_2} (1 - e^{-(\epsilon_2 - \epsilon_1)|x|})$ :
      draw  $z \sim \begin{cases} e^{-(\epsilon_2 - \epsilon_1)z}, & \text{for } 0 \leq z \leq |x| \\ 0, & \text{otherwise.} \end{cases}$ 
       $y = \text{sgn}(x)z$ .
    case with probability  $\frac{\epsilon_2 - \epsilon_1}{2\epsilon_2} e^{-(\epsilon_2 - \epsilon_1)|x|}$ :
      draw  $z \sim \begin{cases} e^{-(\epsilon_1 + \epsilon_2)z}, & \text{for } z \geq |x| \\ 0, & \text{otherwise.} \end{cases}$ 
      return  $y = \text{sgn}(x)z$ .
  end switch
end function

```

and the ϵ_1 -private response $y_1 = u + V_1$ is released. Remarkably, response y_1 can be generated conditioning only on y_2 :

$$y_2 = y_1 + V_{2 \rightarrow 1}, \quad (26)$$

where $V_{2 \rightarrow 1} = V_1 - V_2$ is independent of V_2 , $V_{2 \rightarrow 1} \perp V_2$. In words, tightening privacy under the Laplace mechanism does not require access to the original data u and can be performed by an agent other than the private data owner. Theorem 3 suggests that the randomized post-processing $y_1 = y_2 + V_{2 \rightarrow 1}$ of the ϵ_2 -private response y_2 is at least ϵ_2 -private. For $V_{2 \rightarrow 1}$ given by Distribution (25), this tightening of privacy level is precisely quantified, i.e. $\epsilon_2 \rightarrow \epsilon_1$. Recall that our results are *tight*; no excessive accuracy is sacrificed in the process.

3.2 High-Dimensional Case

Theorem 9 can be generalized for the case that the space of private data is Euclidean \mathbb{R}^n equipped with the ℓ_1 -norm. Theorem 8 establishes that the Laplace mechanism:

$$Q_\epsilon u = u + V, \text{ where } V \sim e^{-\epsilon \|V\|_1}. \quad (27)$$

minimizes the mean-squared error from the identity query among all ϵ -private mechanisms that use additive noise $V \in \mathbb{R}^n$:

$$\mathbb{E}_{V \sim e^{-\epsilon \|V\|_1}} \|Q_\epsilon u - u\|_2^2. \quad (28)$$

Theorem 8 shows that each coordinate of V is independently sampled. This observation implies that Theorem 9 can be applied to n dimensions independently.

Theorem 10. *Consider privacy levels ϵ_1, ϵ_2 with $\epsilon_2 > \epsilon_1 > 0$. Let Q_1 be an ϵ_1 -private mechanism and Q_2 an ϵ_2 -private mechanism of the form:*

$$Q_1 u := u + V_1 \text{ and } Q_2 u := u + V_2 \text{ with } (V_1, V_2) \sim g \in \Delta(\mathbb{R}^{2n}), \quad (29)$$

where $u \in \mathbb{R}_{\ell_1}^n$. Then, gradual release of sensitive data u from ϵ_1 to ϵ_2 is achieved by the probability distribution $l_{\epsilon_1, \epsilon_2}^n$:

$$l_{\epsilon_1, \epsilon_2}^n(V_1, V_2) = \prod_{i=1}^n l_{\epsilon_1, \epsilon_2}(V_1^{(i)}, V_2^{(i)}), \quad (30)$$

where $V_i = [V_i^{(1)}, \dots, V_i^{(n)}]$, $i = 1, 2$. Namely:

- Mechanism Q_1 is ϵ_1 -private and optimal.
- Mechanism Q_2 is the optimal ϵ_2 -private mechanism.
- Mechanism (Q_1, Q_2) is ϵ_2 -private.

Proof. See Appendix A. □

³Distribution (22) is split in four regions; the atom and three exponentially decaying ones. Note that the absolute value $|x|$ is used in comparisons which requires the multiplication by $\text{sgn}(x)$ in the end.

3.3 Multiple Privacy Relaxations

Theorems 9 and 10 perform privacy relaxation from ϵ_1 to ϵ_2 . However, the privacy level is possibly updated multiple times. Theorem 11 handles the case where the privacy level is successively relaxed from ϵ_1 to ϵ_2 , to ϵ_3 , until ϵ_m . Specifically, Theorem 11 enables the use of Theorem 9 multiple times while relaxing privacy level from ϵ_i to ϵ_{i+1} for $i \in \{1, \dots, m-1\}$. We call this statement the *Markov property* of the Laplace mechanism.

Theorem 11. *Consider m privacy levels $\{\epsilon_i\}_{i=1}^m$ with $0 < \epsilon_1 < \dots < \epsilon_m$ and mechanisms Q_i of the form:*

$$Q_i u = u + V_i, \text{ with } (V_1, \dots, V_m) \sim g \in \Delta(\mathbb{R}^m). \quad (31)$$

Consider the distribution $g = l_{\epsilon_1, \dots, \epsilon_m}$, with:

$$l_{\epsilon_1, \dots, \epsilon_m}(v_1, \dots, v_m) = l_{\epsilon_1}(v_1) \prod_{i=1}^{m-1} \frac{l_{\epsilon_i, \epsilon_{i+1}}(v_i, v_{i+1})}{l_{\epsilon_i}(v_i)}, \quad (32)$$

where $l_\epsilon(v) = \frac{\epsilon}{2} e^{-\epsilon|v|}$. Then, distribution $l_{\epsilon_1, \dots, \epsilon_m}$ has the following properties:

1. *Each prefix mechanism (Q_1, \dots, Q_i) is ϵ_i -private, for $i \in \{1, \dots, m\}$.*
2. *Each mechanism Q_i is the optimal ϵ_i -private mechanism, i.e. it minimizes the mean-squared error $\mathbb{E}V_i^2$.*

Proof. See Appendix A. □

Additionally, performing multiple rounds of privacy relaxations can be performed in the context of Theorem 10 is possible. In that case, Theorem 11 is independently applied to each component.

Multiple Rounds of Privacy Relaxation

Theorem 11 states that it is possible to repeatedly use Theorem 9 to perform multiple privacy level relaxations. An intuitive proof of Theorem 11 can be constructed by considering Scenarios 1 and 2 introduced in the beginning of Section 3. Specifically, Theorem 9 constructs a *coupling* such that Scenario 1 replicates Scenario 2. Therefore, once the first round of privacy relaxation $\epsilon_1 \rightarrow \epsilon_2$ occurs, the two scenarios are indistinguishable. The second round of privacy relaxation $\epsilon_2 \rightarrow \epsilon_3$ is performed by starting at the first step of Scenario 1.

In practice, Theorem 11 allows for an efficient implementation of an arbitrary number or privacy relaxation rounds $\epsilon_1 \rightarrow \epsilon_2 \rightarrow \dots \rightarrow \epsilon_m$. In particular, only the most recent privacy level ϵ_i and noise sample V_i need to be stored in memory. Sampling for V_{i+1} depends only on current privacy level ϵ_i , current noise sample V_i and next privacy level ϵ_{i+1} . Past privacy levels $\{\epsilon_j\}_{j < i}$, past noise samples $\{V_j\}_{j < i}$, and future privacy levels $\{\epsilon_j\}_{j > i+1}$ are not needed.

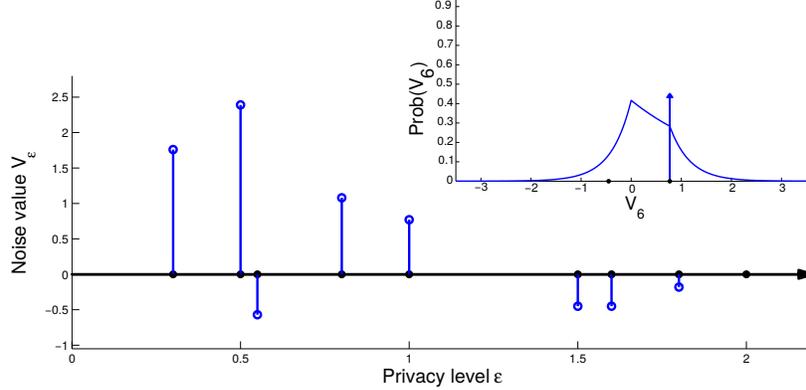


Figure 4: Privacy level can be repeatedly relaxed. For each round of relaxation $\epsilon_i \rightarrow \epsilon_{i+1}$, the distribution of the next noise sample V_{i+1} depends only on the last noise sample V_i . Past noise samples $\{V_j\}_{j<i}$ can be discarded from memory, thus, there is no complexity incurred from repeatedly relaxing privacy level.

3.4 A Private Stochastic Process

Theorems 9 and 11 offer a novel dimension to the Laplace mechanism. Specifically, these results establish a real-valued stochastic process $\{V_\epsilon : \epsilon > 0\}$. Sampling from the process $\{V_\epsilon\}_{\epsilon>0}$ performs gradual release of sensitive data for the continuum of privacy levels $(0, \infty)$. Consider the mechanisms Q_ϵ that respond with $Q_\epsilon u = y_\epsilon = u + V_\epsilon$. Then:

- V_ϵ is optimally distributed, i.e. Laplace-distributed with parameter $\frac{1}{\epsilon}$
- Any “pre-fix” response $\{y_\sigma\}_{\sigma \in (0, \epsilon]}$ is ϵ -private.

Samples of the process $\{V_\epsilon\}_{\epsilon>0}$ are plotted in Figure 1. This process features properties that allow efficient sampling:

- It is *Markov*, $V_s \perp V_t | V_q$, for $s < q < t$. Thus, a sample of the process V_ϵ over an interval $[\epsilon_1, \epsilon_2]$ can be extended to $[\epsilon_1, \epsilon_3]$, for $\epsilon_3 > \epsilon_2$.
- It is *lazy*, i.e. $V_{\epsilon+\delta} = V_\epsilon$ with high probability, for $\delta \ll 1$. Therefore, A sample of the process $\{V_\epsilon\}_{\epsilon_1 \leq \epsilon \leq \epsilon_2}$ can be efficiently stored; only a finite (random) number m of points $(\epsilon_i, V_{\epsilon_i})_{i=1}^m$ where jumps occur need to be stored for *exact* re-construction of the process.

4 Extensions

4.1 Approximate Privacy

For completeness we mention that gradual release of private data can be performed within approximate differential privacy and the Gaussian mechanism. Approximate differential privacy (Dwork et al., 2006) is another variant of differential privacy. The privacy level is governed by two constants ϵ and δ . In the special case that $\delta = 0$, we retrieve the original notion of differential privacy.

Definition 12 (Approximate Differential Privacy). *Let \mathcal{U} be a set of private data, $\mathcal{A} \subseteq \mathcal{U}^2$ be a symmetric binary relation (called adjacency relation) and \mathcal{Y} be a set of possible responses. For $\epsilon > 0$, the randomized mapping $Q : \mathcal{U} \rightarrow \Delta(\mathcal{Y})$ (called mechanism) is (ϵ, δ) -differentially private if*

$$\mathbb{P}(Qu \in \mathcal{S}) \leq e^\epsilon \mathbb{P}(Qu' \in \mathcal{S}) + \delta, \quad \forall (u, u') \in \mathcal{A}, \quad \forall \mathcal{S} \subseteq \mathcal{Y}. \quad (33)$$

For real-valued responses $\mathcal{Y} = \mathbb{R}^n$, the Gaussian mechanism (Dwork and Roth, 2013) is used to provide privatized responses under approximate differential privacy. Contrary to the Laplace mechanism, there are no exact optimality guarantees for the Gaussian mechanism.

Theorem 13 (Gaussian Mechanism). *For a query $q : \mathcal{U} \rightarrow \mathbb{R}^n$, let $\Delta = \inf_{(u, u') \in \mathcal{A}} \|q(u) - q(u')\|_2$, and consider the mechanism that adds Gaussian-distributed noise:*

$$Qu = q(u) + \begin{bmatrix} V^{(1)} \\ \vdots \\ V^{(n)} \end{bmatrix}, \quad (34)$$

where $V^{(i)}$ are i.i.d. zero-mean Gaussian noise with variance $\sigma(\epsilon, \delta) := \frac{\Delta}{\epsilon} \sqrt{2 \ln \left(\frac{1.25}{\delta} \right)}$. Then, mechanism Q satisfies (ϵ, δ) -approximate differential privacy.

Gradually release of sensitive data under the Gaussian mechanism is feasible. The following theorem is the adaptation of Theorem 1 for the case of approximated differential privacy.

Theorem 14. *Let \mathcal{U} be the space of privacy data, $q : \mathcal{U} \rightarrow \mathbb{R}^n$ be a query, and Δ be the ℓ_2 -sensitivity for an adjacency relation \mathcal{A} . Then, consider the privacy levels $\{(\epsilon_t, \delta_t)\}_{t \in [0, \infty)}$ such that $\sigma(\epsilon_t, \delta_t)$ is a decreasing function of t . Then, the family of mechanisms Q_t :*

$$Q_t u := q(u) + V_t, \quad \text{where } V_t = \begin{bmatrix} B_{\sigma(\epsilon_t, \delta_t)}^{(1)} \\ \vdots \\ B_{\sigma(\epsilon_t, \delta_t)}^{(n)} \end{bmatrix} \quad \text{and } t \in (0, \infty), \quad (35)$$

where $V_t = [V_t^{(1)}, \dots, V_t^{(n)}]$ and $V_t^{(i)} \stackrel{d}{=} B_{\sigma(\epsilon_t, \delta_t)}$ are independent samples of the Brownian motion B :

- **Privacy:** For any $t > 0$, the mechanism that releases the signal $\{q(u) + V_\tau\}_{\tau \in (0,t]}$ is (ϵ_t, δ_t) -private.
- **Accuracy:** The mechanism Q_t that releases the random variable $q(u) + V_t$ is the Gaussian mechanism with $\sigma(\epsilon_t, \delta_t)$.

Proof. See Appendix A. □

4.2 Non-identity Queries: Crowdsourcing Statistics with RAPPOR

Theorems 9 and 11 perform gradual release of private data by releasing responses that approximate the identity query $q(u) = u$. In practice, however, the end-user is interested in more expressive queries q such as the mean value $\frac{1}{n} \sum_{i=1}^n u_i$ of a collection of private data u_1, \dots, u_n and solutions to optimization problems (Han et al., 2014). Our results are directly applicable to a family of queries which are approximated by private mechanisms built around the Laplace mechanism. Specifically, consider mechanisms based on the Laplace mechanism and have the form shown in Figure 5. The database of private data is initially pre-processed and, then, additive Laplace-distributed noise is used. The result is post-processed in order to maximize the accuracy of the response. Informally stated:

Corollary 15. *Let (\mathcal{U}, d) be a metric space of sensitive data, \mathcal{Y} be a set of responses, and $\epsilon > 0$ be a privacy level. Let*

- $F : \mathcal{U} \rightarrow \Delta(\mathbb{R}^n)$ be a pre-processing step with sensitivity β . Function F is assumed to be invariant of ϵ , i.e. it does not change for different privacy levels,
- $\mathcal{L}_\epsilon : \mathbb{R}^n \rightarrow \Delta(\mathbb{R}^n)$ be the Laplace mechanism with parameter ϵ :

$$\mathcal{L}_\epsilon u = u + V, \text{ where } V \sim e^{-\frac{\epsilon}{\beta} \|V\|_1}, \quad (36)$$

- $G_\epsilon : \mathbb{R}^n \rightarrow \Delta(\mathcal{Y})$ be a post-processing step.

Consider the ϵ -private mechanism

$$G \circ \mathcal{L} \circ F : \mathcal{U} \rightarrow \Delta(\mathcal{Y}). \quad (37)$$

Then, there exists a composite mechanism that performs gradual release of private data $u \in \mathcal{U}$.

The term “gradual release of private data” should be understood in the *scenario-replicating* sense; using Corollary 15 to relax privacy from ϵ_1 to ϵ_2 , the resulting performance is the same as if the initial privacy level ϵ_1 had never occurred and the system was set at privacy level ϵ_2 from the beginning.

Examples of existing such privacy-aware mechanisms can be found in e.g. smart grids (Koufogiannis et al., 2014) and user’s reports (Erlingsson et al., 2014). On the other

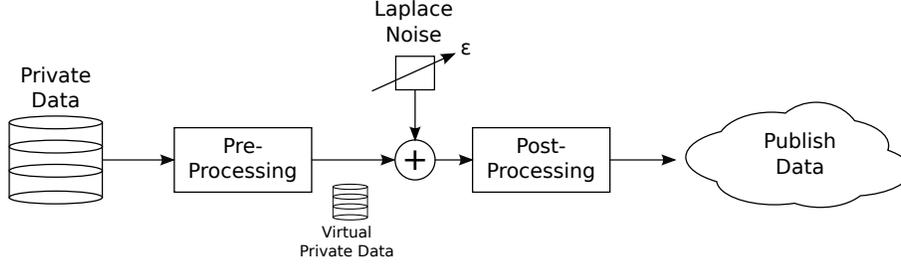


Figure 5: User 1 wants to share his sensitive data, such as his date of birth, in the a social network. Although, user 1 has no privacy concerns when sharing this information with his close friends 2 and 3, he has gradually increasing privacy issues for other members of the network. Specifically, a group A of distant users should not be able to collude and extract more information than what it is intended.

hand, our results do not address the gradual release of private data through mechanisms that do not fulfill this assumption, such as privately solving optimization problems with stochastic gradient descent (Han et al., 2014).

In particular, Google’s RAPPOR (Erlingsson et al., 2014) is a mechanism that collects private data from multiple users for “crowdsourcing statistics” and can be expressed in terms of the Laplace mechanism. RAPPOR collects personal information from users such as the software features they use and the URLs they visited, and provides statistics of this information over a population of users. Algorithmically, a Bloom filter B of size k —which is a bank of k hashing function in parallel— is applied to each user’s private data u :

$$B : \mathcal{U} \rightarrow \{0, 1\}^k, \quad y = [y_1, \dots, y_k] = B(u), \quad (38)$$

where \mathcal{U} is the space of private data, in particular, the set of all strings. Next, each bit y_i is perturbed with probability f and the result is memoised:

$$f : \{0, 1\}^k \rightarrow \{0, 1\}^k, \quad z = [z_1, \dots, z_k] = f(y), \quad \text{where } z_i = \begin{cases} 0, & \text{w.p. } \frac{1}{2}\alpha, \\ 1, & \text{w.p. } \frac{1}{2}\alpha, \\ y_i, & \text{w.p. } 1 - c_1, \end{cases} \quad (39)$$

where “w.p.” stands for “with probability” and $\alpha \in [0, 1]$ is a parameter. Finally, RAPPOR applies another perturbation each time a report is communicated to the server. This perturbation is equivalent to the map (39) but differently parametrized:

$$g : \{0, 1\}^k \rightarrow \{0, 1\}^k, \quad w = [w_1, \dots, w_k] = f(z), \quad \text{where } \mathbb{P}(w_i = 1) = \begin{cases} \beta, & \text{if } z_i = 1, \\ \gamma, & \text{if } z_i = 0, \end{cases} \quad (40)$$

where $\beta, \gamma \in [0, 1]$ are parameters. RAPPOR’s differential privacy guarantees relax (increased ϵ) for small values of α and γ , and large values of β .

An important limitation of RAPPOR is that parameters α , β , and γ are forever fixed. However, there are reasons that require the ability to update these values in a way that the privacy is relaxed and the accuracy is increased:

- Due to the non-trivial algorithm of decoding the reports, a tight accuracy-analysis is not possible. Instead, the accuracy of the system is evaluated once the system is bootstrapped.⁴ Our results make it possible to initialize the parameters with tight values $\alpha \rightarrow 1$, $\beta \rightarrow .5$, $\gamma \rightarrow .5$, and subsequently relax the parameters until a desired accuracy is achieved.
- Once a process or URL is suspected as malicious, the server would be interested in relaxing the privacy level and performing more accurate analysis of the potential threat. Once such a threat is identified, our result allows users to gradually relax their privacy parameters and the server can more confidently explore the potential threat.

In order to apply Theorems 9 and 11 to RAPPOR, we express the randomized maps (39) and (40) using the Laplace mechanism. Specifically, consider the functions \bar{f} and \bar{g} that add Laplace noise and project the result to $\{0, 1\}$:

$$\bar{f}(\psi) = \left[\psi + V_f > \frac{1}{2} \right], \quad \text{where } V_f \sim \text{Lap} \left(\frac{1}{-2 \ln \alpha} \right), \quad (41)$$

$$\bar{g}(\zeta) = \left[\zeta + V_g > \frac{\ln(2\gamma)}{\ln(4\gamma(1-\beta))} \right], \quad \text{where } V_g \sim \text{Lap} \left(\frac{1}{-\ln(4\beta(1-\gamma))} \right), \quad (42)$$

where $\psi, \zeta \in \{0, 1\}$, $\text{Lap}(b)$ is the Laplace distribution with parameter b , and the bracket symbol $[\cdot] \in \{0, 1\}$ is 1 if, and only if, the inside expression is true. Note that functions \bar{f} and \bar{g} have the structure of Figure 5. Moreover, it can be shown that \bar{f} and \bar{g} applied component-wise to y and z are reformulations of the maps f and g . Therefore, privacy level relaxation is achieved by sampling noises V_f and V_g as suggested by our results. Note that, due to the Markov property, past privacy levels need not be stored in memory; only the currently applicable privacy level is retained.

4.3 Privacy in Social Networks

The context of social networks provides another setting where gradually releasing private data is critical. Consider a social network as a graph $G = (V, E)$, where V is the set of users and E the set of friendships between them. Each user owns a set of sensitive data that can include the date of birth, the number of friends and the city he currently resides in. In the realm of social media, user’s privacy concerns scale with the distance to other users of the network. Specifically, an individual is willing to share his private data with his close friends without any privacy guarantees, is skeptical about sharing this information with friends of his friends, and is alarmed to release his sensitive data

⁴Even in that case, estimating the actual accuracy can be challenging since it should be performed in a differential private way.

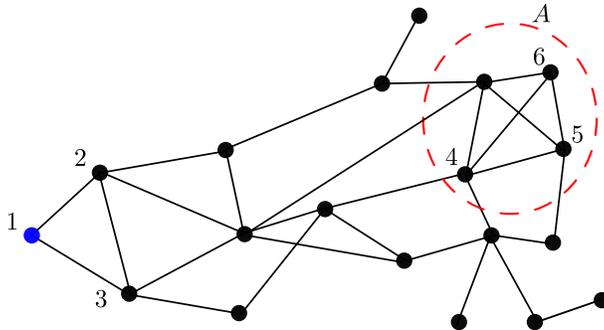


Figure 6: User 1 wants to share his sensitive data, such as his date of birth, in the a social network. Although, user 1 has no privacy concerns when sharing this information with his close friends 2 and 3, he has gradually increasing privacy issues for other members of the network. Specifically, a group A of distant users should not be able to collude and extract more information than what it is intended.

over the entire social network. Therefore, an individual i chooses a different privacy level ϵ_j for each user $j \in V$, for example, a decreasing function of the distance between users i and j :

$$\epsilon_j = \frac{1}{d(i, j)}, \quad (43)$$

where d is any distance function that quantifies how far away user i resides from user j , e.g. the length of the shortest path between nodes i and j . Then, user i could generate an ϵ_j -private response y_j independently for each member j of the network. However, more private information than desired is leaked. Specifically, consider the part of the social network shown in Figure 6, where user $i = 1$ wishes to share his sensitive data u , such as her date of birth. Then, consider a group $A \subseteq V$ of users residing far away from user 1 such that the privacy budget ϵ_j allocated by user i to each member j of the group A is small:

$$d(1, j) \gg 1 \Rightarrow \epsilon_j \ll 1.$$

In the case that members of the large group A decide to collude, they can infer more information about the sensitive data u . Specifically, if a *large* group A averages the received responses $\{y_j : j \in A\}$, the exact value of sensitive data u is recovered. Indeed, composition theorem implies that only $\left(\sum_{j \in A} \epsilon_j\right)$ -privacy of sensitive data u is guaranteed. For a large group A , this privacy level becomes very loose.

Our approach mitigates this issue. We assume that noisy versions of the private data are correlated and we design a mechanism that retains strong privacy guarantees. For real-valued sensitive data u , user 1 samples $\{v_\epsilon : \epsilon > 0\}$ from the stochastic process $\{V_\epsilon : \epsilon > 0\}$, and responds to user j with $y_j = u + v_{\epsilon_j}$, as shown in Figure 7. In

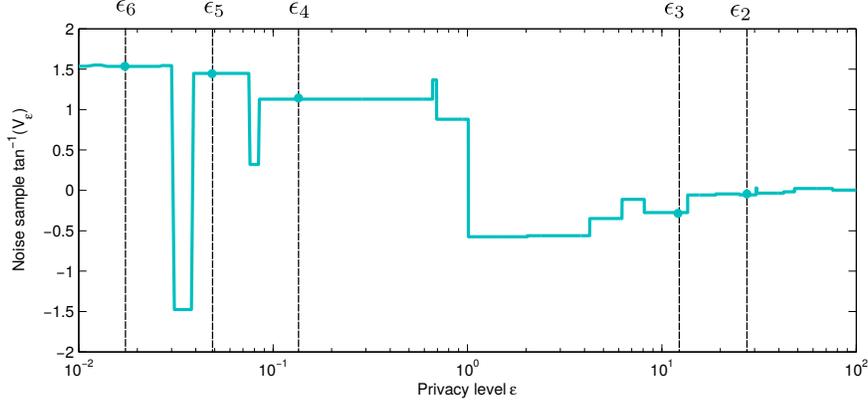


Figure 7: User 1 draws a single sample from the stochastic process $\{V_\epsilon\}_{\epsilon>0}$ and responds to user i with $y_i = u + V_{\epsilon_i}$, where ϵ_i is the privacy level against user i . Eventually, having access to more responses $\{y_i\}_{i \in A}$ does not reveal more information about private data u than the best response $\max_{i \in A} \epsilon_i$.

the case that a large group A of users colludes, they are unable to extract much more information. Specifically, such a collusion renders individual’s sensitive information at most $(\max_{j \in A} \epsilon_j)$ -differential private. This privacy budget is significantly tighter than the one derived in the naive application of differential privacy and corresponds to the best information that a member of the group A has. After all, if a close friend leaks sensitive information, it is impossible to revoke it.

5 Discussion

Finally, we speculate that gradually releasing private data can be extended to any query and is, therefore, a deeper property of differential privacy. Such a property is a key ingredient for the existence of a frictionless market of private data. In such a market, owners of private data can gradually agree to a rational choice of privacy level. Moreover, buying the exact private data is expected to be extremely costly. Instead, people may choose to buy private data in “chunks”, in the sense of increasing privacy budgets. Ideally, future work would explore whether gradually releasing sensitive data *without loss in accuracy* is feasible for a broader family of privacy-preserving mechanisms beyond mechanisms that approximate identity queries. This work was focused on mechanisms which are defined on real space or sensitive data $\mathcal{U} = \mathbb{R}^n$ under an ℓ_1 -norm adjacency relation, and approximate the identity query.

Acknowledgement

The authors would like to thank Aaron Roth for providing useful feedback and suggesting the application of our results to Google’s RAPPOR project and the anonymous reviewer for suggesting the suboptimal approach based on the composition theorems. This work was supported in part by the TerraSwarm Research Center, one of six centers supported by the STARnet phase of the Focus Center Research Program (FCRP) a Semiconductor Research Corporation program sponsored by MARCO and DARPA and in part by NSF CNS-1505799 and the Intel-NSF Partnership for Cyber-Physical Systems Security and Privacy.

References

- Chatzikokolakis, K., Andrés, M. E., Bordenabe, N. E., and Palamidessi, C. (2013). “Broadening the scope of Differential Privacy using metrics.” In *Privacy Enhancing Technologies*, 82–102. Springer.
- Dwork, C. (2006). “Differential privacy.” In *Automata, languages and programming*.
- Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006). “Calibrating noise to sensitivity in private data analysis.” In *Theory of cryptography*, 265–284. Springer.
- Dwork, C., Naor, M., Pitassi, T., and Rothblum, G. N. (2010). “Differential privacy under continual observation.” In *Proceedings of the 42nd ACM symposium on Theory of computing*, 715–724. ACM.
- Dwork, C. and Roth, A. (2013). “The algorithmic foundations of differential privacy.” *Theoretical Computer Science*, 9(3-4): 211–407.
- Erlingsson, Ú., Pihur, V., and Korolova, A. (2014). “RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response.” In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, 1054–1067. ACM.
- Geng, Q. and Viswanath, P. (2012). “The optimal mechanism in differential privacy.” *arXiv preprint arXiv:1212.1186*.
- Gupta, A., Ligett, K., McSherry, F., Roth, A., and Talwar, K. (2010). “Differentially private combinatorial optimization.” In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, 1106–1125. Society for Industrial and Applied Mathematics.
- Han, S., Topcu, U., and Pappas, G. J. (2014). “Differentially Private Distributed Constrained Optimization.” *arXiv preprint arXiv:1411.4105*.
- Hardt, M. and Talwar, K. (2010). “On the geometry of differential privacy.” In *Proceedings of the 42nd ACM symposium on Theory of computing*, 705–714. ACM.

- Hsu, J., Huang, Z., Roth, A., Roughgarden, T., and Wu, Z. S. (2014). “Private matchings and allocations.” In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, 21–30. ACM.
- Huang, Z., Mitra, S., and Vaidya, N. (2015). “Differentially private distributed optimization.” In *Proceedings of the 2015 International Conference on Distributed Computing and Networking*, 4. ACM.
- Koufogiannis, F., Han, S., and Pappas, G. (2015). “Optimality of the Laplace Mechanism in Differential Privacy.” *arXiv preprint arXiv:1504.00065*.
- Koufogiannis, F., Han, S., and Pappas, G. J. (2014). “Computation of privacy-preserving prices in smart grids.” In *IEEE Conference on Decision and Control*.
- Le Ny, J. and Pappas, G. J. (2014). “Differentially private filtering.” *Automatic Control, IEEE Transactions on*, 59(2): 341–354.
- Li, C., Hay, M., Rastogi, V., Miklau, G., and McGregor, A. (2010). “Optimizing linear counting queries under differential privacy.” In *Proceedings of the twenty-ninth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, 123–134. ACM.
- McSherry, F. and Talwar, K. (2007). “Mechanism design via differential privacy.” In *IEEE Symposium on Foundations of Computer Science*.
- Reed, J. and Pierce, B. C. (2010). “Distance makes the types grow stronger: a calculus for differential privacy.” In *ACM Sigplan Notices*.
- Shokri, R. and Shmatikov, V. (2015). “Privacy-Preserving Deep Learning.” In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, 1310–1321. ACM.
- Ullman, J. (2013). “Answering $n^{2+o(1)}$ counting queries with differential privacy is hard.” In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, 361–370. ACM.
- Wang, Y., Huang, Z., Mitra, S., and Dullerud, G. E. (2014). “Entropy-minimizing Mechanism for Differential Privacy of Discrete-time Linear Feedback Systems.” In *IEEE Conference on Decision and Control*.
- Xiao, X., Bender, G., Hay, M., and Gehrke, J. (2011). “iReduct: Differential privacy with reduced relative errors.” In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*, 229–240. ACM.

Appendix A: Full Proofs

We now prove Theorem 9 which performs a single round of privacy level relaxation. The proof consists of validating that Distribution (20) has Laplace-distributed marginals and satisfies the privacy constraints.

Proof. Consider the mechanism $Q = (Q_1, Q_2)$ induced by the noise density (20). We prove that this mechanism satisfies all the desired properties:

1. The first coordinate is Laplace-distributed with parameter $\frac{1}{\epsilon_1}$. For $x \geq 0$, we get:

$$\begin{aligned}
\mathbb{P}(V_1 = x) &= \int_{\mathbb{R}} l_{\epsilon_1, \epsilon_2}(x, y) dy = \frac{\epsilon_1^2}{2\epsilon_2} e^{-\epsilon_2 x} + \frac{\epsilon_1(\epsilon_2^2 - \epsilon_1^2)}{4\epsilon_2} \int_{\mathbb{R}} e^{-\epsilon_1|x-y| - \epsilon_2|y|} dy \\
&= \frac{\epsilon_1^2}{2\epsilon_2} e^{-\epsilon_2 x} + \frac{\epsilon_1(\epsilon_2^2 - \epsilon_1^2)}{4\epsilon_2} \left(\int_{-\infty}^0 e^{-\epsilon_1 x + (\epsilon_1 + \epsilon_2)y} dy + \int_0^x e^{-\epsilon_1 x - (\epsilon_2 - \epsilon_1)y} dy \right. \\
&\quad \left. + \int_x^{\infty} e^{\epsilon_1 x - (\epsilon_1 + \epsilon_2)y} dy \right) \\
&= \frac{\epsilon_1^2}{2\epsilon_2} e^{-\epsilon_2 x} + \frac{\epsilon_1(\epsilon_2 - \epsilon_1)}{4\epsilon_2} e^{-\epsilon_1 x} e^{(\epsilon_1 + \epsilon_2)y} \Big|_{-\infty}^0 - \frac{\epsilon_1(\epsilon_1 + \epsilon_2)}{4\epsilon_2} e^{-\epsilon_1 x} e^{-(\epsilon_2 - \epsilon_1)y} \Big|_0^x \\
&\quad - \frac{\epsilon_1(\epsilon_2 - \epsilon_1)}{4\epsilon_2} e^{-(\epsilon_1 + \epsilon_2)y} \Big|_x^{\infty} \\
&= \frac{\epsilon_1}{2} e^{-\epsilon_1 x}
\end{aligned} \tag{A-1}$$

The case $x \leq 0$ follows from the symmetry $(x, y) \rightarrow (-x, -y)$. Therefore, the first coordinate is ϵ_1 -private and achieves optimal performance.

2. The second coordinate is Laplace-distributed with parameter $\frac{1}{\epsilon_2}$. We have:

$$\begin{aligned}
\mathbb{P}(V_2 = y) &= \int_{\mathbb{R}} l_{\epsilon_1, \epsilon_2}(x, y) dx = \frac{\epsilon_1^2}{2\epsilon_2} e^{-\epsilon_2|y|} + \frac{\epsilon_1(\epsilon_2^2 - \epsilon_1^2)}{4\epsilon_2} e^{-\epsilon_2|y|} \int_{\mathbb{R}} e^{-\epsilon_1|x-y|} dx \\
&= \frac{\epsilon_1^2}{2\epsilon_2} e^{-\epsilon_2|y|} + \frac{\epsilon_1(\epsilon_2^2 - \epsilon_1^2)}{4\epsilon_2} e^{-\epsilon_2|y|} \int_{\mathbb{R}} e^{-\epsilon_1|x|} dx \\
&= \frac{\epsilon_1^2}{2\epsilon_2} e^{-\epsilon_2|y|} + \frac{\epsilon_2^2 - \epsilon_1^2}{2\epsilon_2} e^{-\epsilon_2|y|} \\
&= \frac{\epsilon_2}{2} e^{-\epsilon_2|y|}
\end{aligned} \tag{A-2}$$

Thus, the second coordinate achieves optimal performance.

3. Lastly, we need to prove that the composite mechanism is ϵ_2 -private. We handle the delta part separately by defining $D = \{x : (x, x) \in S\}$ for a measurable

$S \subseteq \mathbb{R}^2$. The probability of landing in set S is:

$$\mathbb{P}(Qu \in S) = \frac{\epsilon_1^2}{2\epsilon_2} \int_D e^{-\epsilon_2|x-u|} dx + \frac{\epsilon_1(\epsilon_2^2 - \epsilon_1^2)}{4\epsilon_2} \iint_S e^{-\epsilon_1|(x-u)-(y-u)| - \epsilon_2|y-u|} dx dy \quad (\text{A-3})$$

We take the derivative and use Fubini's theorem to exchange the derivative with the integral:

$$\begin{aligned} \frac{d}{du} \mathbb{P}(Qu \in S) &= \frac{\epsilon_1^2}{2\epsilon_2} \int_D \epsilon_2 \text{sgn}(x-u) e^{-\epsilon_2|x-u|} dx \\ &\quad + \frac{\epsilon_1(\epsilon_2^2 - \epsilon_1^2)}{4\epsilon_2} \iint_S \epsilon_2 \text{sgn}(y-u) e^{-\epsilon_1|x-y| - \epsilon_2|y-u|} dx dy \Rightarrow \\ \left| \frac{d}{du} \mathbb{P}(Qu \in S) \right| &\leq \frac{\epsilon_1^2}{2\epsilon_2} \int_D \epsilon_2 e^{-\epsilon_2|x-u|} dx + \frac{\epsilon_1(\epsilon_2^2 - \epsilon_1^2)}{4\epsilon_2} \iint_S \epsilon_2 e^{-\epsilon_1|(x-u)-(y-u)| - \epsilon_2|y-u|} dx dy \Rightarrow \\ \left| \frac{d}{du} \mathbb{P}(Qu \in S) \right| &\leq \epsilon_2 \mathbb{P}(Qu \in S) \Rightarrow \left| \frac{d}{du} \ln \mathbb{P}(Qu \in S) \right| \leq \epsilon_2 \end{aligned} \quad (\text{A-4})$$

This completes the proof. \square

Theorem 10 performs gradual release of high-dimensional private data by component-wisely invokes Theorem 9—which handles the gradual release of scalar data—and we confirm that distribution $l_{\epsilon_1, \epsilon_2}^n$ possesses the desired properties.

Proof. Let $[x^{(1)}, \dots, x^{(n)}]$ denote the coordinates of a vector $x \in \mathbb{R}^n$. The desired probability distribution is defined by independently sampling each coordinate using Theorem 9. Let:

$$l_{\epsilon_1, \epsilon_2}^n(x, y) = g(x, y) = \prod_{i=1}^n l_{\epsilon_1, \epsilon_2}(x^{(i)}, y^{(i)}), \quad (\text{A-5})$$

The probability distribution satisfies the required marginal distributions:

$$\int_{\mathbb{R}^n} g(x, y) d^n y = \left(\frac{\epsilon_1}{2}\right)^n e^{-\epsilon_1 \|x\|_1} \quad \text{and} \quad \int_{\mathbb{R}^n} g(x, y) d^n x = \left(\frac{\epsilon_2}{2}\right)^n e^{-\epsilon_2 \|y\|_1}$$

Moreover, it satisfies ϵ_2 -privacy constraints:

$$\begin{aligned} &\|\nabla_u \ln \mathbb{P}(Q_1 u = z_1 \text{ and } Q_2 u = z_2)\|_\infty \\ &= \|\nabla_u \ln l_{\epsilon_1, \epsilon_2}^n(z_1 - u, z_2 - u)\|_\infty \\ &= \max_{1 \leq i \leq n} \left| \frac{\partial}{\partial u_i} \ln l_{\epsilon_1, \epsilon_2}^n(z_1 - u, z_2 - u) \right| \\ &= \max_{1 \leq i \leq n} \left| \frac{\partial}{\partial u_i} \ln l_{\epsilon_1, \epsilon_2}(z_1^{(i)} - u^{(i)}, z_2^{(i)} - u^{(i)}) \right| \\ &\leq \max_{1 \leq i \leq n} \epsilon_2 = \epsilon_2, \end{aligned}$$

where in the last line we used the fact that $l_{\epsilon_1, \epsilon_2}$ is ϵ_2 -private. This completes the proof. \square

Theorem 11 provides the Markov property which allows for multiple privacy level relaxations. This is performed by repeatedly summoning the single-round of privacy relaxation result stated in Theorem 9.

Proof. The proof uses induction on m . The case $m = 2$ is handled by Theorem 9. For brevity, we prove the statement for $m = 3$. Let $f(x, y) = l_{\epsilon_1, \epsilon_2}(x, y)$ and $g(y, z) = l_{\epsilon_2, \epsilon_3}(y, z)$. Consider the joint probability $l_{\epsilon_1, \epsilon_2, \epsilon_3}$:

$$h(x, y, z) = l_{\epsilon_1, \epsilon_2, \epsilon_3}(x, y, z) = \frac{f(x, y)g(y, z)}{l_{\epsilon_2}(y)}, \quad (\text{A-6})$$

where $l_{\epsilon_2}(y) = \frac{\epsilon_2}{2} e^{-\epsilon_2|y|}$. Measure h possesses all the properties that perform gradual release of private data:

- All marginal distributions of measure h are Laplace with parameters $\frac{1}{\epsilon_1}$, $\frac{1}{\epsilon_2}$, and $\frac{1}{\epsilon_3}$, respectively:

$$\int_{\mathbb{R}} \int_{\mathbb{R}} h(x, y, z) dy dz = l_{\epsilon_1}(x), \quad \int_{\mathbb{R}} \int_{\mathbb{R}} h(x, y, z) dx dz = l_{\epsilon_2}(y), \quad \text{and} \quad \int_{\mathbb{R}} \int_{\mathbb{R}} h(x, y, z) dx dy = l_{\epsilon_3}(z).$$

- Mechanism Q_1 is ϵ_1 -private since V_1 is Laplace-distributed with parameter $\frac{1}{\epsilon_1}$.
- Mechanism (Q_1, Q_2) is ϵ_2 -private. Margining out V_3 shows that $(V_1, V_2) \sim l_{\epsilon_1, \epsilon_2}$, which guarantees ϵ_2 -privacy according to Theorem 9.
- Mechanism (Q_1, Q_2, Q_3) is ϵ_3 -private. It holds that:

$$\begin{aligned} & \left| \frac{\partial}{\partial u} \mathbb{P}(Q_1 u = \psi_1, Q_2 u = \psi_2, \text{ and } Q_3 u = \psi_3) \right| \\ &= \left| \frac{\partial}{\partial u} h(\psi_1 - u, \psi_2 - u, \psi_3 - u) \right| \\ &= \left| \frac{\partial h(x, y, z)}{\partial x} + \frac{\partial h(x, y, z)}{\partial y} + \frac{\partial h(x, y, z)}{\partial z} \right| \Bigg|_{\substack{x=\psi_1-u, \\ y=\psi_2-u, \\ z=\psi_3-u}} \end{aligned} \quad (\text{A-7})$$

Algebraic manipulation of the last expression establishes the result:

$$\begin{aligned} \left| \frac{\partial h}{\partial x} + \frac{\partial h}{\partial y} + \frac{\partial h}{\partial z} \right| &= \left| \frac{f_x g}{l_{\epsilon_2}} + \frac{f_y g}{l_{\epsilon_2}} + \frac{f g_y}{l_{\epsilon_2}} - l_{\epsilon_2} \frac{f g}{l_{\epsilon_2}^2} + \frac{f g_z}{l_{\epsilon_2}} \right| \\ &= \left| -\text{sgn}(y) \epsilon_2 \frac{f g}{l_{\epsilon_2}} - \text{sgn}(z) \epsilon_3 \frac{f g}{l_{\epsilon_2}} + \text{sgn}(y) \epsilon_2 \frac{f g}{l_{\epsilon_2}} \right| \\ &= \left| -\text{sgn}(z) \epsilon_3 \frac{f g}{l_{\epsilon_2}} \right| \\ &= \epsilon_3 h, \end{aligned} \quad (\text{A-8})$$

where we used the properties $l'_{\epsilon_2} = -\text{sgn}(y)\epsilon_2 l_{\epsilon_2}$, $f_x + f_y = -\text{sgn}(y)\epsilon_2 f$, and $g_y + g_z = -\text{sgn}(z)\epsilon_3 g$, where the last two identities were derived in the proof of Theorem 9. \square

Lastly, we provide a proof of Theorem 14 which allows relaxing the privacy parameters within the framework of approximate differential privacy.

Proof. For a given t , the added noise V_t is distributed according to the Gaussian mechanism for parameters (ϵ_t, δ_t) . In order to prove the privacy property, we re-write the released signal as follows:

$$\{Q_\tau u\}_{\tau=-\infty}^t = Q_t u + \{V_\tau - V_t\}_{\tau=-\infty}^t. \quad (\text{A-9})$$

Mechanism (A-9) can, then, be viewed as the composition of the (ϵ_t, δ_t) -private mechanism with a randomized post-processing. Indeed, the post-processing is independent of the mechanism $Q_t u$ since:

$$B_{\sigma(\epsilon_\tau, \delta_\tau)} - B_{\sigma(\epsilon_t, \delta_t)} \perp B_{\sigma(\epsilon_t, \delta_t)}, \quad \forall \tau \leq t, \quad (\text{A-10})$$

where we used the monotonicity of $\sigma(\epsilon_t, \delta_t)$. \square