# Privacy Buckets: Upper and Lower Bounds for r-Fold Approximate Differential Privacy

Sebastian Meiser[1], Esfandiar Mohammadi[2*]
[1] University College London, United Kingdom, e-mail: s.meiser@ucl.ac.uk
[2] ETH Zurich, Switzerland, e-mail: mohammadi@inf.ethz.ch

May 8, 2018

**Abstract**

Many applications require robust guarantees against thousands and sometimes millions of observations, such as anonymous communication systems, privacy-enhancing database queries, or privacy-enhancing machine-learning methods. The notion of r-fold Approximate Differential Privacy (ADP) offers a well-established framework with a precise characterization of the degree of privacy after r observations of an attacker. However, existing bounds for r-fold ADP are loose and, if used for estimating the required degree of noise for an application, can lead to overcautious choices for perturbation randomness and thus to suboptimal accuracy.

We present a numerical (although widely applicable) method for capturing the privacy loss of differentially private mechanisms under composition, which we call privacy buckets. With privacy buckets we compute provable upper and lower bounds for ADP for a given number of observations. We compare our bounds with state-of-the-art bounds for r-fold ADP, including Kairouz, Oh, and Viswanath's composition theorem (KOV), concentrated dfferential privacy and the moment's accountant. We compare these bounds for the Laplace mechanism, the Gauss mechanism, for real-world timing leakage data and for the stochastic gradient descent and we significantly improve over their results (with the exception that the KOV bound seems tight for the Laplace mechanism). Moreover, our lower bounds almost meet our upper bounds, showing that no significantly tighter bounds are possible.

*The authors are in alphabetical order. Both authors equally contributed to this work.

# Contents



# 1 Introduction

Approximate differential privacy (ADP [6]) has been designed to quantify, with two parameters $(\varepsilon, \delta)$, the privacy leakage of systems that require a careful trade-off between the system's usefulness and the system's privacy leakage. Since its introduction, ADP has been successfully used to quantify the privacy leakage of privacy-enhancing mechanisms in various applications, including query-response of sensitive databases, training a deep neural networks [1] while hiding the training data, and even anonymous communication [24]. This privacy leakage, i.e., the $(\varepsilon, \delta)$ parameters, inevitably grows under continual observation; thus privacy eventually deteriorates (see Apple's case [23]). In many application scenarios, continual attacker-observation is unavoidable, e.g., an attacker may have thousands if not hundreds of thousands of observation points.

On one hand, precisely computing the $(\varepsilon, \delta)$ parameters after $r$ observations, called $r$-fold ADP bounds, is hard. On the other hand, imprecise bounds on $(\varepsilon, \delta)$ can lead to either a wrong perception of the privacy leakage (resulting, e.g., in unsatisfied customers) or to an over-cautious choice of system parameters (resulting in unnecessarily high costs). There is a rich body of work on approximating $r$-fold ADP-bounds, where $r$ is the number of observations, for privacy-enhancing mechanisms [9, 14, 1, 18, 8, 2]. Early work did not take the shape of the mechanism's output distribution into account [9, 14]. We show that such bounds are tight w.r.t. the Laplace mechanism but imprecise for many other mechanisms, e.g., the Gaussian mechanism, as these mechanism-oblivious bounds inherently assume a worst-case behavior under composition. Recent work [1, 18, 8, 2], in contrast, introduced mechanism-aware bounds that take the shape of output distribution of the mechanism into account and achieve significantly tighter bounds for some particular mechanism, such as the Gaussian mechanism. However, it is not clear how tight previous mechanism-aware bounds are and how much these bounds can be further improved.

### 1.1 Contribution

We introduce a numerical method—*privacy buckets*—for computing upper and lower $r$-fold $(\varepsilon, \delta)$-ADP bounds that take the mechanisms and their (fixed) noise parameters into account. To this end, we utilize a discretized version of the privacy loss random variable introduced by Dwork and Rothblum [8]. Our approach is sufficiently general enough to subsume the generic adaptive $r$-fold ADP bounds of prior work [14, 19]. We compare our upper bounds with state-of-the-art bounds on adaptive $r$-fold ADP and significantly improve over all of them. Moreover, our lower bounds almost meet the upper bounds, showing that no significantly tighter bounds are possible.

Our evaluations also illustrate the usefulness of privacy buckets in gaining insights about the composition behavior of various mechanisms. In particular, we find that for the right choice of scale parameter and standard-deviation, the Laplace mechanism and the Gauss mechanism converge to the same privacy leakage, i.e., their $(\varepsilon, \delta)$ parameters coincide from a sufficiently high number of observations $r$ onward. [1]

Our method is not only useful for deriving tight bounds for classical differential privacy mechanisms but can be applied to any privacy analysis resulting in differential privacy. We exemplify this statement by computing bounds for the anonymous communication system Vuvuzela [24], the stochastic gradient descent mechanism for deep learning [1] and for timing-leakage measurement histograms of a recently introduced browser extension for deniable communication [22].

## 2 Background and Related work

In this section, we review the notion of differential privacy, highlight an often implicit assumption in the analysis of differentially privacy mechanisms, generalize differential privacy to pairs of distributions, and position our work in the work from the literature.

**Differential privacy** Differential privacy (DP) [5] quantifies how closely related the outputs are of a mechanism on two similar inputs, from an information-theoretic perspective. We say that a mechanism $M$ is $\varepsilon$-DP, if for any two closely related inputs $D_1, D_2$, $\forall S \subseteq \mathcal{U}$. $\Pr[M(D_1) \in S] \leq e^{\varepsilon} \cdot \Pr[M(D_2) \in S]$. To extend the applicability of DP, approximate differential privacy [6] (ADP) has been introduced, which allows for distributions to exceed a limiting factor $\varepsilon$, as long as this deviation can be limited to a small value $\delta$ in the following way: $\forall S \subseteq \mathcal{U}$. $\Pr[M(D_1) \in S] \leq e^{\varepsilon} \cdot \Pr[M(D_2) \in S] + \delta$. This work focuses on ADP.

### 2.1 Worst case distributions for ADP

Classically, differential privacy argues about the output of a probabilistic mechanism $M$ that is run on similar inputs (e.g., neighboring databases). Since $M$ is probabilistic, the application of $M$ to any input $D$ can be seen as a random variable with outputs from a distribution $M(D)$. Differential privacy requires the outputs of $M$ on all pairs of neighboring databases w.r.t. an application specific sensitivity-metric, i.e., all pairs of

[1] For the expert reader, this observation indicates that the result from Dwork and Rothblum [8], that subgaussian privacy loss variables compose (at most as badly as) a Gaussian privacy loss variable, can be generalized.

distributions $M(D_1), M(D_2)$, where $D_1$ and $D_2$ are neighboring, to be closely related (quantified via the privacy parameters $\varepsilon$ and $\delta$).

Our approach works on individual pairs of distributions. While at first glance a focus on single pairs of distributions might seem to restrict the applicability to particular queries, our approach leads to far more general results. In this section, we explain that our approach can be used to analyze mechanisms in the presence of arbitrary adversarial queries. Many differential privacy proofs analyze inputs that are worst-case in the following sense: in terms of privacy these inputs are as bad as any other pair of inputs for a given sensitivity (see Definition 11). Applying such worst-case inputs $(x_0, x_1)$ to a mechanism $M$ leads to a pair of *worst-case distributions* $(M(x_0), M(x_1))$. Such worst-case distributions are typically considered in differential privacy proofs. Our approach computes adaptive $r$-fold ADP for a pair of distributions and can thus be used in the analysis of a mechanisms $M$ if it is applied to such worst-case inputs (see Corollary 1). To understand why worst-case distributions are an integral part of proofs for differential privacy, we now discuss, on an abstract level, how we typically prove that a mechanism is differentially private.

**Most differential privacy analyses implicitly use worst-case distributions** For illustration, let us consider a mechanism, where $M(D, q)$ (for a database $D$ and a query $q$) can be divided into a precise response $f(D, q)$ to a query and an independent noise distribution $N$. In the simplest case, mechanisms are of the form $M(D, q) = f(D, q) + N$, such as the Laplace mechanism, the Gauss mechanism, as well as any other distribution of noise $N$ added to some function $f(D, q)$, where $N$ does not depend on $f(D, q)$. In all these cases, differential privacy guarantees can be calculated based solely on the distribution of the noise and on the sensitivity $\Delta f$, where[2]

$$\Delta f = \max_{\substack{D_1, D_2 \text{ neighboring} \\ \text{query } q}} |f(D_1, q) - f(D_2, q)|.$$

To show that $M$ satisfies (approximate) differential privacy the proof then typically analyze the following two distributions: $N$ and $N + \Delta f$, implicitly assuming that $f(D_1, q) = 0$ and $f(D_2, q) = \Delta f$. Moreover, the proof then argues that for any value $\Delta' < \Delta f$ the distributions $N$ and $N + \Delta'$ also satisfy differential privacy. From this simplified analysis, it can then (implicitly) be derived that for all other values of $f(D_1, q)$ and $f(D_2, q)$ s.t., $|f(D_1, q) - f(D_2, q)| \leq \Delta f$, $f(D_1, q) + N$ and $f(D_2, q) + N$ also satisfy differential privacy, which concludes the analysis. In any such analysis, the worst-case distributions are the ones that are given $N$ and $N + \Delta f$ and those distributions can be used for our bucket analysis.

A prominent example of such an analysis is a recent work on a differentially private a mechanism for privacy-preserving stochastic gradient descent [1]. In this work, Abadi et al. first prove that a pair of a Gaussian ($\mu_0$) and a mixed Gaussian ($(1 - q)\mu_0 + q\mu_1$), where $\mu_i$ is a Gaussian with mean $i$, is worst case for their analysis, and they then estimate differential privacy for this pair of distributions.

**What if my differential privacy analysis doesn't implicitly use worst-case distributions?** If the mechanism does not consist of and cannot be reduced to independent noise being added to a numerical value, the above simplified description does not immediately apply. However, Murtagh and Vadhan [20, Lemma 3.2 & Lemma 3.7] show that for any $(\varepsilon, \delta)$-DP mechanism $M$ there is a worst case mechanism $M_{\varepsilon,\delta}$ operating on a single bit such that ADP guarantees for $M_{\varepsilon,\delta}$ translate directly to ADP guarantees for $M$, even under $r$-fold adaptive composition.[3]

In more detail, they show [20, Lemma 3.2] that there is a probabilistic translation $T$ that relates every differentially private mechanism $M$ on two neighboring inputs $D_0$ and $D_1$ to a generic pair of distributions $M_{\varepsilon,\delta}(0)$ and $M_{\varepsilon,\delta}(1)$. They then leverage the post-processing property of differential privacy to show that analyzing $M_{\varepsilon,\delta}(0)$ and $M_{\varepsilon,\delta}(1)$ is sufficient, even under composition.

[2]The notion of neighboring databases differs from application to application. For counting queries, e.g., two databases are typically called neighboring if they differ in at most one row.

[3]There are mechanisms in the literature for which the privacy parameters (e.g., the standard deviation $\sigma$ of a Gaussian noise distribution) can be adaptively chosen for every run and depending on previous adversarial observations. If privacy cannot be bounded per response or the number of responses relevant for computing privacy is unbounded, we may not find worst-case distributions [21].

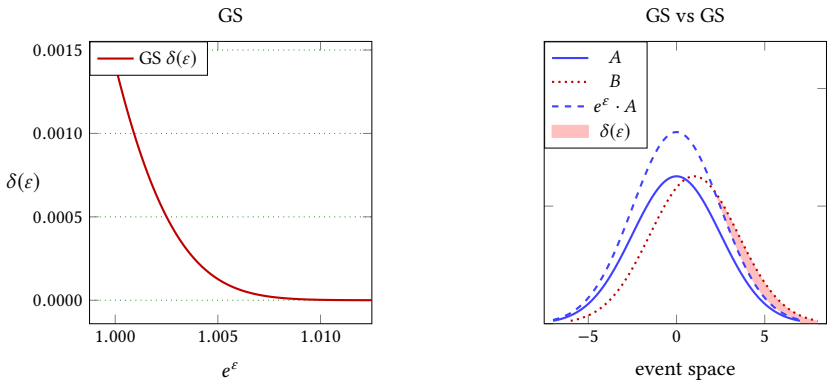


Figure 1: A graph depicting $\delta(\varepsilon)$ for the truncated Gauss mechanism (left) and a graphical depiction of how to compute $\delta(\varepsilon)$ for the truncated Gauss mechanism (right). Note that $e^{\varepsilon} \cdot A$ is not a probability distribution.

Thus, considering such worst-case distributions and their behavior under composition is sufficient for deriving bounds on the mechanism $M$. Moreover, even for different mechanisms and adaptively chosen neighboring inputs, one can compute sound bounds on differential privacy by only considering the distributions $M_{\varepsilon,\delta}(0)$ and $M_{\varepsilon,\delta}(1)$. This technique gives us a fall-back plan for computing bounds on the distributions of $M_{\varepsilon,\delta}(0)$ and $M_{\varepsilon,\delta}(1)$ if no more than $\varepsilon$ and $\delta$ of the mechanism $M$ is known. If more is known (ideally two exact output distributions) we can derive significantly tighter bounds.

**Worst-case distributions for the Laplace mechanism** As an example, let us consider counting queries $q$ with sensitivity 1 to which Laplace noise is added: the mechanism $M$ that gets a database $D$ as input is defined as $M(D) := q(D) + \mathrm{LP}_{\lambda,0}$, where $\mathrm{LP}_{\lambda,\mu}$ is the Laplace distribution with scale parameter $\lambda$ and mean $\mu$ (and $f(D,q) := q(D)$). In this example, it suffices to only consider $\mathrm{LP}_{\lambda,0}$ and $\mathrm{LP}_{\lambda,1}$, with means 0 and 1, instead of considering $M(D_0)$ and $M(D_1)$ for all possible combinations of neighboring databases $D_0$ and $D_1$. Let $D_0$ and $D_1$ be two such neighboring databases where the true answers to a query $q$ are $q(D_0) = x$ and $q(D_1) = x + 1$, respectively, for some numerical value $x$.[4] We can map any output $y$ drawn from $\mathrm{LP}_{\lambda,\mu}$ (for $\mu \in \{0,1\}$) to $y+x$ to obtain the correct adversarial view for the respective scenario $M(D_i) = q(D_i) + \mathrm{LP}_{\lambda,0}$.

## 2.2 Tight ADP on distributions

Approximate differential privacy is typically captured with two parameters $\varepsilon$ and $\delta$. In this work we show that considering not just one such pair of parameters, but a parameter space helps to derive tight adaptive $r$-fold ADP composition results. This parameter space of two distributions can be represented as a function $\delta(\varepsilon)$ such that $(\varepsilon, \delta(\varepsilon))$-ADP holds and $\delta(\varepsilon)$ is minimal (for $\varepsilon \geq 0$). To capture this minimality, we define (tight) differential privacy (generalized to pairs of distributions) and show how to precisely compute $\delta(\varepsilon)$.

**Definition 1** ((Tight) ADP)**.** *Two distributions $A$ and $B$ over the universe $\mathcal{U}$ are $(\varepsilon, \delta)$-ADP, if $\forall$ sets $S \subseteq \mathcal{U}$,*

$$P_A(S) \leq e^{\varepsilon} P_B(S) + \delta(\varepsilon) \text{ and}$$
$$P_B(S) \leq e^{\varepsilon} P_A(S) + \delta(\varepsilon),$$

*where $P_A(x)$ denotes the probability of the event $x$ in $A$ and $P_B(x)$ denotes the probability of the event $x$ in $B$ $A$ and $B$ are* tightly *$(\varepsilon, \delta(\varepsilon))$-ADP if they are $(\varepsilon, \delta(\varepsilon))$-ADP, and $\forall \delta' \leq \delta(\varepsilon)$ such that $A$ and $B$ are $(\varepsilon, \delta')$-ADP we have $\delta(\varepsilon) = \delta'$.*

[4]Since the differential privacy guarantee and analysis are symmetric, we can assume w.l.o.g. that $q(D_0) < q(D_1)$.

**A note on utility and sensitivity.** In the remainder of the paper, we consider pairs of distributions. In particular analyzing the ADP-parameter of mechanisms amounts to analyzing the respective worst-case for a given sensitivity. Hence, we can abstract away from the sensitivity of two inputs, and we can abstract away from the utility functions of a task.

**Computing a tight ADP bounds.** To see that and how we can compute $\delta(\varepsilon)$, first consider that any pair of distributions is $(\varepsilon, 1)$-ADP for arbitrary $\varepsilon \geq 0$. More precise bounds for distributions $A$ and $B$ can be captured by setting $\delta(\varepsilon)$ to the area between the probability distributions of $B$ and a scaled-up version of $A$: we multiply every point of the curve of $A$ with $e^{\varepsilon}$ (which is not a probability distribution anymore, because it sums up to $e^{\varepsilon}$ instead of to 1). We refer to Figure 1 for the $(\varepsilon, \delta(\varepsilon))$-graph of possible (tight) ADP-bounds for two truncated Gaussian distributions (left side) and for a graphical depiction of this intuition (right side). Any area where $B$ is larger than this scaled-up curve contains probability mass for events $x$ outside of the multiplicative bound, i.e., for which we have $P_A(x) \leq e^{\varepsilon} P_B(x)$. The difference between those terms is precisely what we need to characterize.

**Lemma 1.** *For every $\varepsilon$, two distributions $A$ and $B$ over a finite universe $\mathcal{U}$ are tightly $(\varepsilon, \delta)$-ADP with*

$$\delta = \max\left(\sum_{x \in U} \max\left(P_A(x) - e^{\varepsilon} P_B(x), 0\right), \sum_{x \in U} \max\left(P_B(x) - e^{\varepsilon} P_A(x), 0\right)\right)$$

*Proof.* Let $\varepsilon \geq 0$ and let $A$ and $B$ be two distributions over the universe $\mathcal{U}$. We show the equivalence by first showing that (1) for every set $S$, the calculation describes an upper bound and then that (2) there exists a set $S$ such that this bound is tight.

**(1)** We show that $\forall S \subseteq \mathcal{U}$,

$$\begin{aligned} & P_A(x \in S : x) - e^{\varepsilon} P_B(x \in S : x) \\ \leq & \sum_{x \in U} \max\left(P_A(x) - e^{\varepsilon} P_B(x), 0\right) \end{aligned}$$

The inverse direction then follows analogously.

$$\begin{aligned} & P_A(x \in S : x) - e^{\varepsilon} P_B(x \in S : x) \\ = & \sum_{x \in S} P_A(x) - e^{\varepsilon} P_B(x) \\ \leq & \sum_{x \in S} \max\left(P_A(x) - e^{\varepsilon} P_B(x), 0\right) \\ \leq & \sum_{x \in U} \max\left(P_A(x) - e^{\varepsilon} P_B(x), 0\right) \end{aligned}$$

**(2)** Let $S := \{x \in \mathcal{U} \text{ s.t. } \Pr[x \in A] \geq e^{\varepsilon} \Pr[x \in B]\}$. Then,

$$\begin{aligned} & P_A(x \in S : x) - e^{\varepsilon} P_B(x \in S : x) \\ = & \sum_{x \in S} P_A(x) - e^{\varepsilon} P_B(x) \\ = & \sum_{x \in U} \max\left(P_A(x) - e^{\varepsilon} P_B(x), 0\right). \end{aligned}$$

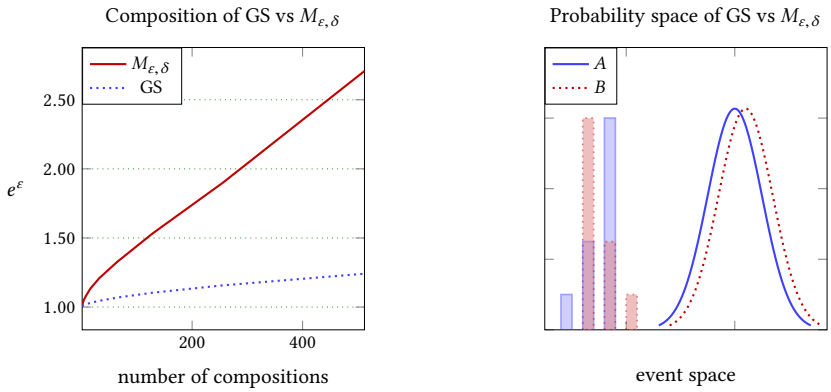


Figure 2: Comparison between the truncated Gauss mechanisms (blue) and a randomized response mechanism $M_{\varepsilon,\delta}$ (red). We show how $e^{\varepsilon}$ evolves in both cases for a fixed $\delta = 10^{-4}$ (left side). Nore that both graphs start at the same point, but quickly diverge. For ease of understanding we depict the probability distributions of interest for both mechanisms (right); here, randomized response directly consists of only 4 possible events. For both mechanisms, we portray the two distributions $A$ and $B$.

Analogously, for $S := \{x \in \mathcal{U} \text{ s.t. } \Pr[x \in B] \geq e^{\varepsilon}\Pr[x \in A]\}$,

$$\begin{aligned} & P_B(x \in S : x) - e^{\varepsilon} P_A(x \in S : x) \\ = & \sum_{x \in S} P_B(x) - e^{\varepsilon} P_A(x) \\ = & \sum_{x \in U} \max\left(P_B(x) - e^{\varepsilon} P_A(x), 0\right). \end{aligned}$$

Thus, for every pair of distributions $A$ and $B$ and for every $\varepsilon \geq 0$ the distributions are tightly $(\varepsilon, \delta)$-differentially private, where $\delta$ is calculated as described.

□

If only one pair $\varepsilon, \delta(\varepsilon)$ is considered, composition can only be based on very limited information about the distributions. In this case, for all we know, the distributions could actually have the shape of the randomized response distributions $M_{\varepsilon,\delta}(0)$ and $M_{\varepsilon,\delta}(1)$.[5] However, by considering more information we can derive much better composition bounds. We refer to Figure 2 both for the guarantees of those distributions (Gaussian versus $M_{\varepsilon,\delta}(0/1)$) under 512 compositions (left side) and for a graphical depiction of those distributions (right side).

## 2.3 Practical relevance of tight privacy bounds

To further highlight the importance of tight privacy bounds for actual mechanisms and protocols we briefly discuss as a case study the Vuvuzela [24] protocol, which is an anonymous communication system tailored towards messengers. The Vuvuzela paper argues that the only leakage of their strong anonymity mechanisms is the patterns of communication between entities. To limit this leakage, they apply noise to the patterns by sending a random number of dummy messages, where the number of messages follows a truncated Laplace distribution. Vuvuzela has two relevant protocol parts that can be analyzed separately: the dialing protocol (to establish contact) and the communication protocol (to transfer messages).

To quantify the improvements of a tight analysis, Figure 19 plots the growth of $\varepsilon$ and the growth of $\delta$, respectively, with an increasing number of observations $r$ for the conversation protocol, one of the two relevant parts of their system. The original paper [24] proposed to increase privacy with dummy messages that are distributed according to a Laplace distribution. We propose to use a Gaussian distribution with a

[5] Kairouz, Oh, and Viswanath [14] proved that if a mechanism satisfies $(\varepsilon, \delta)$-ADP, it cannot have more leakage than $M_{\varepsilon,\delta}$.

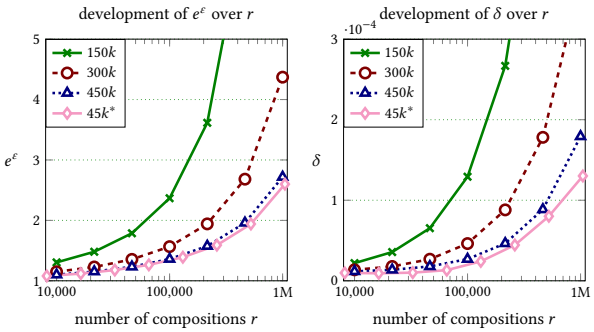


Figure 3: Vuvuzela conversion protocol: bounds on $\varepsilon$ and $\delta$ over $r$ (log-scale). We compare the original bounds for the originally recommended mechanisms with $150k$, $300k$, $450k$, analyzed with previous bounds [9] and our recommended mechanism with $45k$ messages overhead per round, analyzed using privacy buckets.

smaller mean that significantly reduces the noise-overhead. The original paper proposed three configurations with different noise overhead ($150k$, $300k$, and $450k$ noise messages per round) and privacy guarantees. We show that with Gaussian noise and a tighter analysis, we can achieve a higher degree of privacy than the previous highest-noise configuration with a lower overhead ($45k$ noise messages per round) than the lowest-noise configuration. Our proposal, achieves with a factor 10 reduction in the overhead than the highest-noise configuration the same degree of privacy.

Moreover, a tight analysis reveals several other potential improvements. For $r = 500,000$ the high configuration of Laplace noise (LP-*high*) provides a privacy of $\delta$ almost 4 orders of magnitude lower, and the corresponding Gaussian noise (with the same variance and mean) more than 6 orders of magnitude lower in comparison to their original guarantees. Furthermore, we can see that even a Gaussian noise with 1/3 of the average noise of even GS-new$_2$ meets the privacy requirements of $e^\varepsilon \leq 2$ and $\delta \leq 10^{-4}$ for $r = 500,000$ observations. We also analyze the dialing protocol, where similar improvements are possible: 5 times lower Gaussian noise suffice for matching their best guarantees and using Laplace noise incurs a privacy improvement of 3 orders of magnitude, whereas comparable Gaussian noise allows an improvement of 4 orders of magnitude.

We refer the interested reader to Section 8 for a more detailed description of our Vuvuzela analysis.

## 2.4 Composition of differential privacy

One of the main advantages of differential privacy is the fact that guarantees are still sound under composition, albeit with increasing values for $\varepsilon$ and $\delta$.

**Definition 2** ($k$-fold DP of a mechanism)**.** *A randomized algorithm $M$ with domain $\mathcal{D}$ and range $\mathcal{U}$ is $k$-fold $(\varepsilon, \delta)$-differentially private for sensitivity $s$ if for all $S \subseteq \mathcal{U}^k$ and for all $(x_1, \ldots, x_k), (y_1, \ldots, y_k) \in \mathcal{D}^k$ such that $\forall 1 \leq i \leq k.\ ||x_i - y_i||_1 \leq s$:*

$$\begin{aligned}&\Pr[(M(x_1), \ldots, M(x_k)) \in S] \\ \leq &e^\varepsilon \Pr[(M(y_1), \ldots, M(y_k)) \in S] + \delta\end{aligned}$$

Note that when we describe differential privacy in terms of distributions over the worst-case inputs, the composition of differential privacy is equivalent to considering differential privacy for product distributions. If $x_0, x_1$ are the worst-case inputs for a mechanism $M$, resulting in the distributions $M(x_0)$ and $M(x_1)$, then the k-fold composition is described in Definition 1 on the distributions $A = M(x_0)^k$ and $B = M(x_1)^k$. Similarly, a composition of two different mechanisms $M$ and $M'$ with worst-case inputs (in the sense of Section 2.1) $x_0, x_1$ and $x'_0, x'_1$ respectively, boils down to Definition 1 on the distributions $A = M(x_0) \times M'(x'_0)$ and $B = M(x_1) \times M'(x'_1)$.

The main composition results we compare our work with are: naive composition, slightly less naive composition and two composition result with improved bounds [9, 14]. We recall these results here.

**Lemma 2** (Naïve Composition)**.** *Let $(A_1, B_1)$ and $(A_2, B_2)$ be two pairs of distributions, such that $A_1$ and $B_1$ are $(\varepsilon_1, \delta_1)$-differentially private and $A_2$ and $B_2$ are $(\varepsilon_2, \delta_2)$-differentially private. Then $A_1 \times A_2$ and $B_1 \times B_2$ are $(\varepsilon_1 + \varepsilon_2, \delta_1 + \delta_2)$-differentially private.*

**Lemma 3** (Adaptive Composition)**.** *Let $(A_1, B_1)$ and $(A_2, B_2)$ be two pairs of distributions, such that $A_1$ and $B_1$ are $(\varepsilon_1, \delta_1)$-differentially private and $A_2$ and $B_2$ are $(\varepsilon_2, \delta_2)$-differentially private. Then $A_1 \times A_2$ and $B_1 \times B_2$ are $(\varepsilon_1 + \varepsilon_2, \delta_1 + (1 - \delta_1) \cdot \delta_2)$-differentially private.*

**Lemma 4** (Boosting and Differential Privacy (Advanced Composition) [9])**.** *Let $(A_1, B_1), \ldots, (A_k, B_k)$ be pairs of distributions, such that $A_i$ and $B_i$ are $(\varepsilon, \delta)$-differentially private for all $i \in \{1, \ldots, k\}$. Then $A_1 \times \ldots \times A_k$ and $B_1 \times \ldots \times B_k$ are $(\hat{\varepsilon}_{\hat{\delta}}, \hat{\delta})$-differentially private, where $\hat{\delta} = k \cdot \delta$ and $\hat{\varepsilon}_{\hat{\delta}} = O\left(k\varepsilon^2 + \varepsilon\sqrt{k \log\left(e + (\varepsilon\sqrt{k}/\hat{\delta})\right)}\right)$*

**Lemma 5** (Kairouz et al.'s Composition [14])**.** *For any $\varepsilon \geq 0$ and $\delta \in [0, 1]$, the class of $(\varepsilon, \delta)$-differentially private mechanisms satisfies*

$$(\varepsilon', \delta')\text{-differential privacy}$$

*under $k$-fold composition, for all $i \in \{0, \ldots, \lfloor k/2 \rfloor\}$ where $\varepsilon' = (k - 2i)\varepsilon$ and $\delta' = 1 - (1 - \delta)^k(1 - \delta_i)$*

$$\delta_i = \frac{\sum_{\ell=0}^{i-1} \binom{k}{\ell} \left(e^{(k-\ell)\varepsilon} - e^{(k-2i+\ell)\varepsilon}\right)}{(1 + e^{\varepsilon})^k}$$

These composition results allow for deriving differential-privacy guarantees under composition in a black-box manner, i.e., only depending on $\varepsilon$ and $\delta$. Consequently, these results are oblivious to how the underlying distributions actually compose and present, in a way, worst-case results under composition. Thus, we cannot expect that they come close to the tight differential privacy guarantee of the composed distributions. In the remainder of this paper we introduce, prove sound and discuss our main idea: approximating the distributions $A_1, A_2, B_1, B_2$ in a way that allows for a sound calculation of a differential-privacy guarantee that takes into account features of the distribution even under manifold composition. Moreover, we use the same technique to derive a lower bound for the guarantee, to bound the (unknown) tight differential privacy guarantee from both directions.

## 2.5 Related work

**Mechanism-oblivious bounds for adaptive composition** Early composition bounds for adaptive $r$-fold ADP [9, 14, 19] only provide mechanism-oblivious bounds, i.e., these bounds are oblivious to the actual mechanisms. These results only rely on the initial values $(\varepsilon_0, \delta_0)$. Our work, in contrast, is mechanism-aware in the sense that it takes the shape of the distributions (/mechanisms) into account. Our results yield mechanism-aware $\delta$-tight bounds for adaptive composition and thereby lead to significantly tighter bounds.

**Mechanism-aware bounds for adaptive composition** Recent work [8, 2, 1, 18] partially take the shape of the mechanism into account by computing the Rényi divergence of the corresponding worst-case distributions, i.e., the moments of the distribution of ratios, to achieve tighter privacy bounds. Similarly, Abadi et al. [1] use the moments accountant based on Rényi divergence to find tighter bounds. These approaches indeed find tighter bounds in comparison to composition results and, in special cases, better than the best mechanism-oblivious composition theorem. As shown in our comparisons , however, our bounds are even tighter and—in contrast to all previous work—also include lower bounds and thereby a means to estimating their precision. Additionally, our work provides tight bounds for very low epsilon, even epsilon = 0, i.e., the statistical distance (also called total variation), which is used to formalize statistical indistinguishability.

**Adaptively chosen privacy parameters** As in previous work [9, 1, 14] our technique satisfies adaptive composition [9] in the following sense: sequences of mechanisms are composed where each query can be adaptively chosen by the attacker and depend on previously observed responses, but the noise distributions of each mechanism have to be independent of these previously observed responses to the attacker. This kind

of adaptive composition results does not hold for some mechanisms that achieve ADP under continual observation that use carefully correlated noise and/or only use noise when necessary [7, 10, 12, 13]. Nevertheless, the proofs of these adaptive mechanisms can still benefit from our results as they often over-approximate a subset of these correlated distributions with independent distributions

**Probabilistic differential privacy (PDP) vs ADP** It might appear preferable to only use $\delta$ such that it is only the probability of distinguishing events, in order to guarantee pure $\varepsilon$-DP with probability $(1-\delta)$ (which is also called PDP). However, if delta would only contain distinguishing events, both $\varepsilon$ and $\delta$ would grow linearly in the number of compositions. Thus, better $\varepsilon$-bounds can only be achieved by allowing some of the probability mass of the non-distinguishing events to be hidden within the $\delta$ parameter. While using PDP with distinguishing events has an intuitive interpretation, it is not closed under post-processing [17]. Hence, this work concentrates on ADP.

**Optimal mechanisms for a given utility function** Recent work [11, 15] made progress on finding optimal mechanisms for DP for a large class of utility functions. These results concentrate on single observations and do not characterize how these mechanism behave under $k$-fold composition.

**Dependencies** The work of Liu, Chakraborty and Mittal [16] discusses the importance of correctly measuring the sensitivity of databases for differential privacy. They show that in real-world examples entries can be correlated and thus cannot be independently exchanged as in DP's basic definition. Their approach, however, finally results in the same techniques as in previous work being used to achieve the same goal: noise applied to database queries results in differential privacy, although the sensitivity is calculated in a more complex manner. Our results can directly be applied in such a setting as well: given the (final) distributions that potentially consider dependent entries we calculate differential privacy guarantees for these distributions.

# 3 Privacy buckets of two distributions

## 3.1 Informal description of privacy buckets

Generic bounds for differential privacy under continual observation [9, 14] are stated independently of the shape of the underlying distributions, simply based on the ADP guarantees before the composition. This obliviousness is both strength and weakness: the exact shape of the distribution does not need to be characterized to apply these results, but they cannot devise tight bounds that are derivable from the shape of the distributions. We now introduce an alternative approach: we approximate the distributions with an explicit focus on their most important features for ADP, the privacy loss of atomic events.

Recall from Lemma 1 that for distributions $A$ and $B$ over the universe $\mathcal{U}$ we can calculate a value $\delta(\varepsilon)$ for every value $\varepsilon \geq 0$ so that $A$ and $B$ are tightly $(\varepsilon, \delta(\varepsilon))$-ADP:

$$\delta(\varepsilon) = \max\Bigg(\sum_{x\in\mathcal{U}} \max\left(P_A(x) - e^{\varepsilon}P_B(x), 0\right), \\ \sum_{x\in\mathcal{U}} \max\left(P_B(x) - e^{\varepsilon}P_A(x), 0\right)\Bigg),$$

For simplicity we consider $\delta(\varepsilon) = \sum_{x\in\mathcal{U}} \max\left(P_A(x) - e^{\varepsilon}P_B(x), 0\right)$ for now. Consequently, the contribution of each atomic event $x \in \mathcal{U}$ to $\delta(\varepsilon)$ is $\delta(x,\varepsilon) = \max(P_A(x) - e^{\varepsilon}P_B(x), 0)$ and their sum is $\sum_x \delta_x = \delta(\varepsilon)$. This is of course not surprising. Let us observe that if $P_B(x) = 0$, we have $\delta(x,\varepsilon) = P_A(x)$. We can combine all atomic events $x$ with $P_B(x) = 0$ into one non-atomic event $x_\infty$ of all such events.

For events $x$ with $P_B(x) > 0$, let $\mathcal{L}^{(x)}_{(A||B)} = \ln \frac{P_A(x)}{P_B(x)}$ be the logarithmic privacy loss of $x$ [8]. For ease of use, we do define the privacy loss without the logarithm as $e\mathcal{L}^{(x)}_{(A||B)} = e^{\mathcal{L}^{(x)}_{(A||B)}} = \frac{P_A(x)}{P_B(x)}$, which is simply the ratio between the two probabilities. Based on the privacy loss we can calculate the contribution $\delta_x$ of an

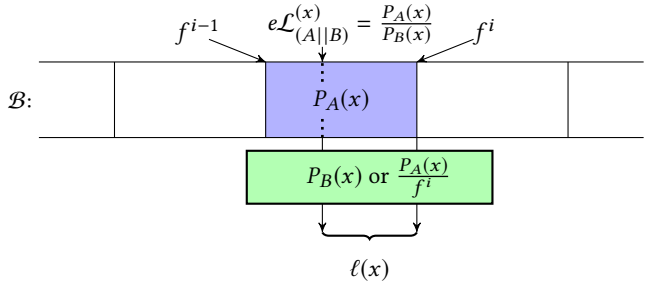


Figure 4: Depiction of how an element $x$ is placed into a bucket when $f^{i-1} < e\mathcal{L}^{(x)}_{(A||B)} \leq f^i$. Buckets store $f^i$ and $P_A(x)$ (accumulated over all elements in the bucket). We approximate $P_B(x)$ with $\frac{P_A(x)}{f^i}$, accepting an error of $\ell(x) = P_B(x) - \frac{P_A(x)}{f^i}$.

atomic event $x$ as

$$\begin{aligned}\delta(x,\varepsilon) &= \max\left(P_A(x) - e^\varepsilon P_B(x), 0\right) \\ &= \max\left(P_A(x) - e^\varepsilon \frac{P_A(x)}{e\mathcal{L}^{(x)}_{(A||B)}}, 0\right) \\ &= P_A(x) \cdot \max\left(\left(1 - \frac{e^\varepsilon}{e\mathcal{L}^{(x)}_{(A||B)}}\right), 0\right).\end{aligned}$$

**Combining the contributions of several events** For any two disjoint events $x, y$ with the same privacy loss $p = e\mathcal{L}^{(x)}_{(A||B)} = e\mathcal{L}^{(y)}_{(A||B)}$, their contribution can be combined without loss of information to

$$\begin{aligned}\delta(x \cup y, \varepsilon) &= \delta(x,\varepsilon) + \delta(y,\varepsilon) \\ &= (P_A(x) + P_A(y)) \cdot \max\left(\left(1 - \frac{e^\varepsilon}{p}\right), 0\right),\end{aligned}$$

requiring us to only remember the privacy loss $p$ and the sum of their probabilities $P_A(x) + P_A(y)$. In other words, we can combine all atomic events with the same ratio without losses. If we allow for a slight imprecision, we can soundly combine disjoint events $x$ and $y$ with approximately the same privacy loss by summing the probabilities $P_A(x) + P_A(y)$ and choosing $p = \max(e\mathcal{L}^{(x)}_{(A||B)}, e\mathcal{L}^{(y)}_{(A||B)})$ and yield $\delta(\varepsilon)(x \cup y) \geq \delta(x,\varepsilon) + \delta(y,\varepsilon)$.

**Constructing privacy buckets from atomic events** To render our approach feasible, we fix a finite set of privacy loss values $\{f^i | i \in \{-n, \ldots, n\}\}$ based on a factor $f$ that parametrizes the coarseness of the values and a limit $n \in \mathbb{N}$ that limits the number of values we consider. We then collect all atomic events $x$ with a similar privacy loss into one combined event, which we call a *bucket* as follows. Given a factor $f > 1$, the bucket $\mathcal{B}(i)$ summarizes all atomic events where $f^{i-1} < e\mathcal{L}^{(x)}_{(A||B)} \leq f^i$ (illustrated in Figure 4). The value of $\mathcal{B}(i)$ is the sum over the probabilities $P_A(x)$ of all those atomic events (according to distribution $A$). Here, $e\mathcal{L}^{(x)}_{(A||B)} \leq f^i$ guarantees soundness, whereas $f^{i-1} < e\mathcal{L}^{(x)}_{(A||B)}$ limits the imprecision of our approximation: For each $P_B(x)$ we introduce an error of $\ell(x) = P_B(x) - \frac{P_A(x)}{f^i} \leq P_A(x) \cdot \left(\frac{1}{f^{i-1}} - \frac{1}{f^i}\right)$. We define "corner buckets" that collect all atomic events with privacy loss outside of $[f^{-n}, f^n]$, where $\mathcal{B}(-n)$ contains all atomic events with a very small privacy loss and $\mathcal{B}(\infty)$ contains all atomic events with a very large (or even infinite) privacy loss. Using these $n+2$ buckets, we can now compute a bound for ADP as $\delta'_\varepsilon = \sum_{i \in \{-n,\ldots,n\}} \mathcal{B}(i) \cdot \max\left(\left(1 - \frac{e^\varepsilon}{f^i}\right), 0\right) + \mathcal{B}(\infty)$.

Buckets for given parameters $f$ and $n$.

| Bucket factor: | $f^{-n}$ | $f^{-n+1}$ | $\cdots$ | $f^{-2}$ | $f^{-1}$ | $f^0$ | $f^1$ | $f^2$ | $\cdots$ | $f^{n-1}$ | $f^n$ | $> f^n$ |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Index: | $-n$ | $-n+1$ | $\cdots$ | $-2$ | $-1$ | $0$ | $1$ | $2$ | $\cdots$ | $n-1$ | $n$ | $\infty$ |

Figure 5: Depiction of the buckets (separately) constructed for both $\mathcal{B}_A$ and $\mathcal{B}_B$. For $\mathcal{B}_A$ each bucket $\mathcal{B}_A(i)$ with $i \in \{-n+1, \ldots, n\}$ contains all elements $x \in \mathcal{U}$ with $f^{i-1}P_B(x) \leq P_A(x) \leq f^i P_B(x)$, the bucket $\mathcal{B}_A(-n)$ contains all elements with $P_A(x) \leq f^{-n}P_B(x)$ and the bucket $\mathcal{B}_A(\infty)$ contains all elements with $P_A(x) > f^n P_B(x)$.

This buckets representation does not only allow us to directly derive a bound $\delta$ on $\delta(\varepsilon)$. It also is particularly well suited for calculating ADP after composing several pairs of distributions. Note that we include "privacy loss" values that are smaller than 1. Those values by definition cannot influence $\delta(\varepsilon)$ directly, but they are crucial for computing tight bounds under composition.

**Composition** Consider a pair of distributions, say $(A_1, B_1)$ and $(A_2, B_2)$ over universes $\mathcal{U}_1, \mathcal{U}_2$, where $A_1, A_2$ are independent and $B_1, B_2$ are independent. We first create $\mathcal{B}_1$ from $(A_1, B_1)$ and $\mathcal{B}_2$ from $(A_2, B_2)$. For each event $(x, y) \in \mathcal{U}_1 \times \mathcal{U}_2$ where $x$ was placed in $\mathcal{B}_1(i)$ and $y$ was placed in $\mathcal{B}_2(j)$ we can now immediately derive an upper bound for the privacy loss: $\frac{P_{A_1}(x)}{P_{B_1}(x)} \cdot \frac{P_{A_2}(y)}{P_{B_2}(y)} \leq f^{i+j}$ (c.f. Figure 8).

As the $A_1$ and $A_2$ are independent and $B_1$ and $B_2$ are independent, we can generate a new set of buckets $\mathcal{B}'(i) = \sum_{j,k,j+k=i} \mathcal{B}_1(j) \cdot \mathcal{B}_2(k)$ and for these buckets $\mathcal{B}'$ we can directly compute $\delta'_\varepsilon$ s.t. for every choice of $\varepsilon \geq 0$, $(A_1 \times A_2, B_1 \times B_2)$ satisfy $(\varepsilon, \delta'_\varepsilon)$-ADP.

**Squaring** When composing privacy buckets, the bucket list naturally "broadens", i.e., the buckets that are farther away from the middle bucket (with factor $f^0$) gain higher values. When creating privacy buckets for a given number $n$, this effect leads to a trade-off between the granularity (i.e., the choice of the bucket factor $f$) and the expected number of compositions: the smaller the value of $f$, the more precise the privacy buckets model of the features of the distributions, but the fewer compositions before a significant amount of events reaches the corner buckets $\big(\mathcal{B}(-n)$ and $\mathcal{B}(\infty)\big)$, which again reduces the precision. To counter this effect, we introduce an additional operation which we call *squaring*: we square the factor $f$, thus halving the precision of the privacy buckets, and merge the privacy buckets into these new, more coarse-grained privacy buckets. Squaring allows us to start with much more fine-grained privacy buckets and reduce the granularity as we compose, which can significantly improve the overall precision of the approach. We choose to square $f$ instead of increasing it to an arbitrary $f'$ to ease the computation of the new privacy buckets: we simply combine buckets $2i-1$ and $2i$ with factors $f^{2i-1}$ and $f^{2i}$ into the new bucket $i$ with factor $(f^2)^i = f^{2i}$. We refer to Figure 9 for a graphical depiction of squaring.

## 3.2 A formal description of privacy buckets

We now formalize privacy buckets, our approximation of the pair of distributions based on the privacy loss of all atomic events, which is sufficient for calculating $(\varepsilon, \delta)$-ADP, the *privacy buckets*, and that comes with an efficient way for computing $r$-fold $(\varepsilon, \delta)$-ADP from a sequence of privacy buckets.

**The infinity symbol $\infty$** In this paper we will write $\infty$ to describe the corner case accumulated in the largest bucket $\mathcal{B}_\infty$ of our bucket lists. We consider $\infty$ to be a distinct symbol and in an abuse of notation, we use the following mathematical rules to interact with it:

- $\infty > i$ for all $i \in \mathbb{Z}$.
- $\infty + i = \infty$ for all $i \in \mathbb{Z}$.

The composition of privacy buckets is commutative but not associative. Moreover, when and how often the squaring is performed influences the resulting privacy buckets. Hence, we need to keep track of the order in which we applied composition and squaring. To this end, we define *composition trees*.

| | |
|---|---|
| $P_A(x)$ | the prob. that $x$ happens in $A$. |
| $\varepsilon, \delta$ | parameters for ADP. |
| $\mathcal{U}$ | universe of all atomic events. |
| $f$ | factor (close to 1) with $f > 1$. |
| $\infty$ | symbol for any ratio $> f^n$. |
| $n$ | index of the last bucket before $\infty$. |
| $N$ | bucket indexes $\{-n, \ldots, n\}$. |
| $N_\infty$ | bucket indexes with $\infty$, $N \cup \{\infty\}$. |
| $T$ | composition tree |
| $\mathbb{P}(A, B, f, n)$ | leaf node of A/B privacy buckets without error correction with indexes $N_\infty$ and ratios $\{\leq f^{-n}, \ldots, \leq f^n, > f^n\}$. |
| $T_1 \times T_2$ | node for composition of $T_1$ and $T_2$ |
| $\blacktriangledown T$ | node for squaring of $T$ |
| $A_T$ | $(A_1, \ldots, A_k)$ for a composition tree $T$ with leafs $\mathbb{P}(A_i, B_i, f, n)_{i=1}^k$ |
| $B_T$ | $(B_1, \ldots, B_k)$ for a composition tree $T$ with leafs $\mathbb{P}(A_i, B_i, f, n)_{i=1}^k$ |
| $\mathcal{B}_T(i)$ for $i \in N_\infty$ | privacy bucket of tree $T$ with index $i$. |
| $\mathcal{B}_T^{\circledast}(x)$ for $x \in \mathcal{U}$ | impact of the atomic event $x$ in tree $T$. |
| $\ell_T(i), \ell_T^{\circledast}(x)$ | "real" error correction term for index $i$ or atomic event $x$. |
| $\tilde{\ell}_T(i), \tilde{\ell}_T^{\circledast}(x)$ | bound on the maximum error, "virtual" error correction term. |
| $\iota_T(x)$ | index of $x$ w.r.t. composition tree $T$. |
| $j_\varepsilon$ | smallest $j$ such that $f^j \geq e^\varepsilon$. |
| $S_i$ | the set of atomic events that contribute to $\mathcal{B}(i)$. |

Figure 6: Notation for our privacy buckets.

**Definition 3** (Composition trees)**.** *For two tuples of distributions $(A_1, \ldots, A_\mathcal{W})$ and $(B_1, \ldots, B_\mathcal{W})$ of the same size $\mathcal{W}$, a* composition tree *is a tree with three kinds of nodes: leaves ($T = \mathbb{P}(A_i, B_i, f, n)$) that are labeled with a pair of distributions $A_i$ and $B_i$ a factor $f > 1$ and a $n \in \mathbb{N}$, squaring nodes ($T = \blacktriangledown T_1$) with exactly one child node, and composition nodes ($T = T_1 \times T_2$) with exactly two child nodes.*

*The* bucket factor *$f_T$ for a composition tree $T$ is $f_{\mathbb{P}(A,B,f,n)} := f$, $f_{\blacktriangledown T_1} := (f_{T_1})^2$, and $f_{T_1 \times T_2} := f_{T_1}$ if $f_{T_1} = f_{T_2}$ and undefined otherwise. The* number of buckets *$n_T$ of a composition tree $T$ is constant: $n_{\mathbb{P}(A,B,f,n)} := n$, $n_{\blacktriangledown T} = n_T$, and $n_{T_1 \times T_2} := n_{T_1}$ if $n_{T_1} = n_{T_2}$ and undefined otherwise.*

*For the distributions over which each composition tree is defined, we write $A_{\mathbb{P}(A,B,f,n)} = A, B_{\mathbb{P}(A,B,f,n)} = B$, $A_{T_1 \times T_2} = A_{T_1} \times A_{T_2}$, and $A_{\blacktriangledown T_1} = A_{T_1}$ and analogously $B_{T_1 \times T_2} = B_{T_1} \times B_{T_2}$, and $B_{\blacktriangledown T_1} = B_{T_1}$. We write $\mathcal{U}_T$ for the support of the distributions as $\mathcal{U}_T = [A_T] \cup [B_T]$.*

*We call a composition tree $T$* valid *if for the product distributions $A_T = \Pi_{k=1}^\mathcal{W} A_k$ all $A_j, A_s$ are pairwise independent ($j, s \in \{1, \ldots, \mathcal{W}\}$) and analogously for $B_T = \Pi_{k=1}^\mathcal{W} B_k$ all $B_j, B_s$ are pairwise independent ($j, s \in \{1, \ldots, \mathcal{W}\}$), $f_T$ and $n_T$ are defined, $f_T > 1$ and $n_T$ is an even natural number (i.e., there is a $q \in \mathbb{N}$ such that $n = 2q$). We sometimes write $f$ instead of $f_T$, $n$ instead of $n_T$, $A$ instead of $A_T$, and $B$ instead of $B_T$ if the composition tree is clear from the context.*

We now define the privacy buckets associated with a valid composition tree $T$, starting with the base case of leaf nodes.

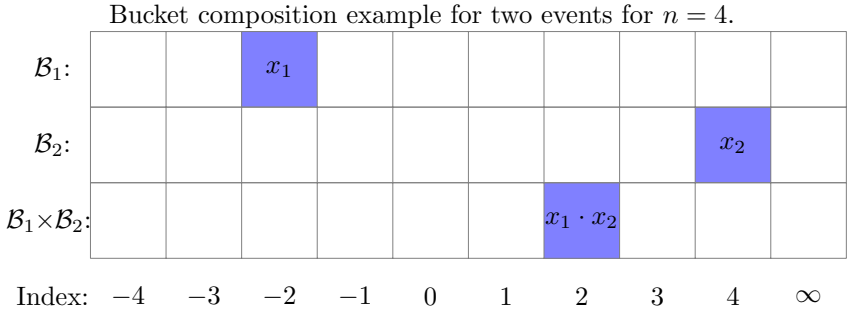


Figure 7: Depiction of how individual events $x_1$ with index $-2$ and $x_2$ with index 4 compose into their new bucket with index $-2 + 4 = 2$.

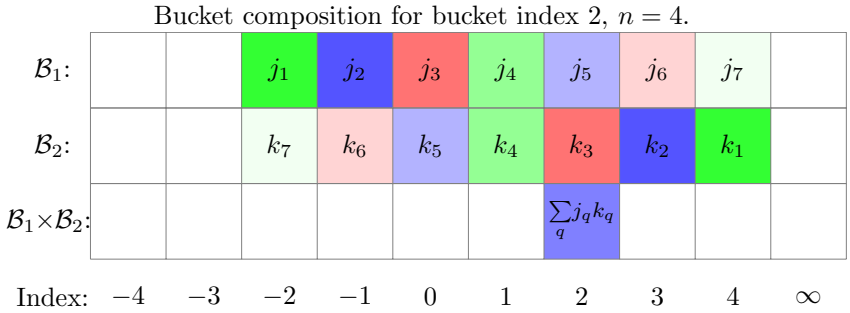


Figure 8: Depiction of the bucket composition for the (new) bucket with index $i = 2$. We calculate the value of bucket $i$ by summing over the product of all $\mathcal{B}_1(j_t) \cdot \mathcal{B}_2(k_t)$ for $t \in \{1, \ldots, 7\}$. Graphically, buckets with the same color are combined. Note that none of the buckets $\infty, -3$ and $-4$ are used for the composition, as for all $j \in \{-4, \ldots, 4\}$, $\infty + j \neq 2, -3 + j \neq 2$ and $-4 + j \neq 2$.

**Computing delta and evaluating a composition tree**

**Definition 4** (Privacy buckets of a composition tree)**.** *Let $T$ be a valid composition tree with $f := f_T$, $\mathcal{U} := \mathcal{U}_T$ and $n := n_T$. For $N_\infty := \{-n, -n+1, \ldots, n\} \cup \{\infty\}$, we define the $A_T/B_T$* privacy buckets *$\mathcal{B}_T : N_\infty \to [0, 1]$ recursively as follows.*

*If $T = \mathcal{D}(A, B, f, n)$, we define for $i \in N_\infty$,*

$$\mathcal{B}_{\mathcal{D}(A,B,f,n)}(i) = \textstyle\sum_{x \in S_i} P_A(x),$$

*where the sets $S_i$ are defined as follows:*

$$\begin{aligned} S_\infty &= \{x \in \mathcal{U}. P_A(x) > f^n P_B(x)\} \\ \forall i \in \{-n+1, \ldots, n\}\ S_i &= \left\{x \in \mathcal{U}.\ f^{i-1} P_B(x) < P_A(x) \leq f^i P_B(x)\right\} \\ S_{-n} &= \left\{x \in \mathcal{U}. P_A(x) \leq f^{-n} P_B(x)\right\}. \end{aligned}$$

*If $T = T_1 \times T_2$, we define*

$$\mathcal{B}_{T_1 \times T_2}(i) := \begin{cases} \sum_{j+k=i} \mathcal{B}_{T_1}(j) \cdot \mathcal{B}_{T_2}(k) & i \in \mathbb{N} \setminus \{-n\} \\ \sum_{j+k \leq -n} \mathcal{B}_{T_1}(j) \cdot \mathcal{B}_{T_2}(k) & i = -n \\ \sum_{j+k > n} \mathcal{B}_{T_1}(j) \cdot \mathcal{B}_{T_2}(k) & i = \infty \end{cases}$$

*If $T = \blacktriangledown T_1$,*

$$\mathcal{B}_{\blacktriangledown T_1}(i) := \begin{cases} \mathcal{B}_{T_1}(2i-1) + \mathcal{B}_{T_1}(2i) & i \in [-n/2 + 1, n/2] \\ \mathcal{B}_{T_1}(\infty) & i = \infty \\ 0 & \textit{otherwise} \end{cases}$$

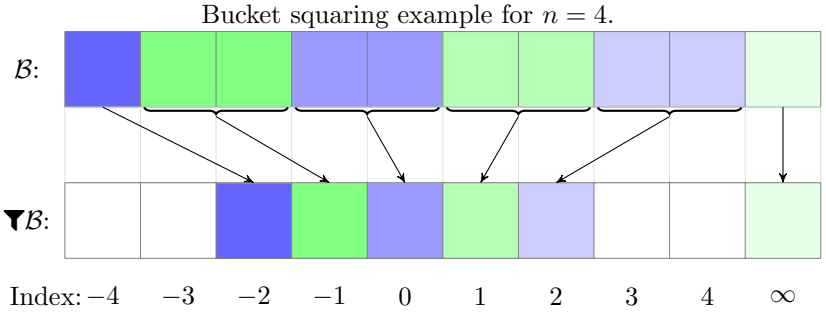


Figure 9: Depiction of the bucket squaring. Events from each bucket $\mathcal{B}(i)$ are moved into bucket $\mathcal{B}(\lceil i/2 \rceil)$, with the exception of $\mathcal{B}(\infty)$, which remains unchanged.

Note that since the sets $S_i$ for $i \in \{-n, \ldots, n\} \cup \{\infty\}$ describe a partitioning of $\mathcal{U}$, we have

$$\sum_{i \in \{-n,\ldots,n\} \cup \{\infty\}} \mathcal{B}(i) = 1.$$

We next define ADP directly on a privacy bucket list. For all atomic events $x$ in $S_i \neq S_\infty$, we know that $P_A(x) \leq f^i P_B(x)$. We perform a slight over-approximation by treating this inequality as an equality and then use $P_A(x) - P_A(x)/f^i$ as in Lemma 1. For $x \in S_\infty$, we add $P_A(x)$ to $\delta$, counting them as total privacy-breakdowns.

**Definition 5** (Delta)**.** *Let $T$ be a valid composition tree labeled with $f := f_T$ and $n := n_T$, then*

$$\delta_T(\varepsilon) = \mathcal{B}_T(\infty) + \sum_{i \in \{-n,\ldots,n\}} \max(0, \mathcal{B}_T(i) \cdot (1 - \tfrac{e^\varepsilon}{f^i}))$$

*We say that the privacy buckets with composition tree $T$ are $(\varepsilon, \delta)$-ADP, if $\delta_T(\varepsilon) \leq \delta$.*

### 3.3 Buckets per atomic event

For discussing our results and their soundness, we compare the differential privacy guarantees of privacy buckets with the real differential privacy guarantees (calculating which might not be feasible). To this end and for talking about individual atomic events, we assign an index to each such event. The index specifies the (one) bucket the respective event influences. For privacy buckets that have been created from distributions (and not composed), this index is simply the bucket the event was assigned to. After composition, the index depends on how the indexes of the respective buckets interacted: in the most simple case, if $x_1$ and $x_2$ are events with indexes $i$ and $j$, then the event $(x_1, x_2)$ will have the index $i + j$. However, the corner cases can modify the index, as the index can only be in the set $\{-n, \ldots, n, \infty\}$.

**Definition 6** (Index of an event according to a composition tree)**.** *For a valid composition tree $T$ with $A_T = \Pi_{k=1}^{\mathcal{W}} A_k$ and $B_T = \Pi_{k=1}^{\mathcal{W}} B_k$, and $\mathcal{U}_T = \Pi_{k=1}^{\mathcal{W}} \mathcal{U}_i$, $f := f_T$, and $n := n_T$, we define the set of indexes for atomic events $x = (x_1, \ldots, x_{\mathcal{W}}) \in \Pi_{k=1}^{\mathcal{W}} \mathcal{U}_k$ as follows.*

*First, we define for $T = \circledcirc(A_k, B_k, f, n)$ and consequently for atomic elements $x_k \in \mathcal{U}_k$ with $k \in \{1, \ldots, \mathcal{W}\}$, the index of $x_k$ as*

$$\iota_T(x_k) := \begin{cases} l & \text{if } l \in \{-n+1, \ldots, n\} \wedge \\ & f^{l-1} P_{B_k}(x_k) < P_{A_k}(x_k) \leq f^l P_{B_k}(x_k) \\ \infty & \text{if } P_{A_k}(x_k) > f^n P_{B_k}(x_k) \\ -n & \text{otherwise} \end{cases}$$

*For a pair of composition trees $T_1, T_2$ and for $T = T_1 \times T_2$ we define the index of $x = (x_1, x_2) \in A_{T_1} \times A_{T_2}$*

*as*

$$\iota_T(x) = \iota_{T_1 \times T_2}(x_1, x_2) := \begin{cases} -n & \text{if } \iota_{T_1}(x_1) + \iota_{T_2}(x_2) < -n \\ \infty & \text{if } \iota_{T_1}(x_1) + \iota_{T_2}(x_2) > n \\ \iota_{T_1}(x_1) + \iota_{T_2}(x_2) & \text{otherwise,} \end{cases}$$

*where we assume that* $\forall y, z \in \mathbb{Z},\ y + \infty = \infty > z$.
*For* $T = \blacktriangledown T_1$ *we define the index of* $x \in A_T$ *as*

$$\iota_T(x) = \iota_{\blacktriangledown T_1}(x) := \begin{cases} \lceil \iota_{T_1}(x)/2 \rceil & \text{if } \iota_{T_1}(x) \neq \infty \\ \infty & \text{otherwise,} \end{cases}$$

Recall that composition is not necessarily associative, i.e., there are composition trees $T_1, T_2, T_3$ and $x_1, x_2, x_3$ such that

$$\iota_{(T_1 \times T_2) \times T_3}(x_1, x_2, x_3) \neq \iota_{T_1 \times (T_2 \times T_3)}(x_1, x_2, x_3).$$

**Soundness of differential privacy guarantees for privacy buckets** We can now show the soundness of the bounds on ADP we calculate using privacy buckets. We will show that if privacy buckets are $(\varepsilon, \delta)$-ADP, then the distributions from which they were created (either directly or via composition) are also $(\varepsilon, \delta)$-ADP. Simply put, the guarantees we calculate are sound.

We begin by showing a helpful lemma that directly follows our main strategy: all atomic events $x$ that are assigned an index $\iota_{T(x)} = i \neq \infty$ (according to a composition tree $T$) satisfy $P_A(x) \leq f^i P_B(x)$.

**Lemma 6.** *Let* $T$ *be a valid composition tree For all* $x \in \mathcal{U}_T$ *with* $\iota_T(x) \neq \infty$ *and for* $f = f_T$*, we have* $P_{A_T}(x) \leq f^{\iota_T(x)} P_{B_T}(x)$*, i.e.,* $f^{\iota_T(x)} \leq e\mathcal{L}^{(x)}_{(A_T || B_T)}$.

*Proof.* We show the lemma by a structural induction over the composition tree $T$. Let $x \in \mathcal{U}_T$ with $\iota_T(x) \neq \infty$.

Case $T = \mathbb{P}(A, B, f, n)$.

If $\iota_{\mathbb{P}(A,B,f,n)}(x) = -n$, it follows that

$$P_A(x) \leq f^{-n} P_B(x) \tag{1}$$

$$P_A(x) \leq f^{\iota_{\mathbb{P}(A,B,f,n)}(x)} P_B(x). \tag{2}$$

Thus, for all $\iota_{\mathbb{P}(A,B,f,n)}(x) \neq \infty$ we get from Definition 6 and Equation (2) that for

$$P_{A_k}(x_k) \leq f^{\iota_{\mathbb{P}(A_k,B_k,f,n)}(x_k)} P_{B_k}(x_k). \tag{3}$$

For composition nodes (i.e., $T = T_1 \times T_2$), where $T_1$ and $T_2$ are valid composition trees, let $x = (x_1, x_2)$ with $x_1 \in \mathcal{U}_1$ and $x_2 \in \mathcal{U}_2$, $f := f_{T_1} = f_{T_2}$. We know from Definition 6 that $\iota_{T_1 \times T_2} \neq \infty \Rightarrow \iota_{T_1} \neq \infty \wedge \iota_{T_2} \neq \infty$. Moreover,

$$\begin{aligned} P_{A_{T_1 \times T_2}}(x) &= P_{A_{T_1}}(x_1) \cdot P_{A_{T_2}}(x_2) \\ &\overset{\text{IH}}{\leq} \left( f^{\iota_{T_1}(x_1)} P_{B_{T_1}}(x_1) \right) \cdot \left( f^{\iota_{T_2}(x_2)} P_{B_{T_2}}(x_2) \right) \\ &= f^{\iota_{T_1}(x_1) + \iota_{T_2}(x_2)} \left( P_{B_{T_1}}(x_1) P_{B_{T_1}}(x_2) \right) \\ &\leq f^{\iota_{T_1 \times T_2}(x)} P_{B_{T_1 \times T_2}}(x) \end{aligned}$$

Note that $f^{\iota_{T_1}(x_1) + \iota_{T_2}(x_2)} \leq f^{\iota_{T_1 \times T_2}(x)}$ holds by definition.

For squaring nodes (i.e., $T = \blacktriangledown T_1$) with $f := f_{T_1}$ (and consequently $f_{\blacktriangledown T_1} = f^2$), we know that $\iota_T(x) \neq \infty \Leftrightarrow \iota_{T_1}(x) \neq \infty$. By definition, we have

$$\begin{aligned} P_{A_{\blacktriangledown T_1}}(x) = P_{A_{T_1}}(x) &\overset{\text{IH}}{\leq} f^{\iota_{T_1}(x)} P_{B_{T_1}}(x) = f^{2\iota_{T_1}(x)/2} P_{B_{T_1}}(x) \\ &\leq (f^2)^{\lceil \iota_{T_1}(x)/2 \rceil} P_{B_{T_1}}(x) = (f_{\blacktriangledown T_1})^{\iota_{\blacktriangledown T_1}(x)} P_{B_{\blacktriangledown T_1}}(x) \end{aligned}$$

□

**BucketDelta**$(A, B, f, n, t, \varepsilon)$**:**
$T = \text{𝒟}(A, B, f, n)$
**for** $i$ **from** $0$ **to** $t$ **do**
  $T' = T \times T$
  **if** $\mathcal{B}_{T'}(\infty) > 2.2 \cdot \mathcal{B}_T(\infty)$ **then**
    $T = \blacktriangledown T$
  $T = T \times T$
**return** $\delta_T(\varepsilon)$

Figure 10: Depiction of how we create buckets – for simplicity without error correction terms and for the common special case where we compose the same distributions ($A_1 = A_2 = \ldots = A_r$ and $B_1 = B_2 = \ldots = B_r$). We use repeated squaring to compute $r$-fold DP for $r = 2^t$ compositions.

**Lemma 7** (Bucket values are sums over atomic events)**.** *Let $T$ be a valid composition tree. Then, for all $i \in \{-n_T, \ldots, n_T, \infty\}$,*

$$\mathcal{B}_T(i) = \sum\nolimits_{x \in \mathcal{U}_T \ s.t. \ \iota_T(x) = i} P_{A_T}(x).$$

*Proof.* We show the lemma via structural induction over $T$. Let $N := \{-n_T, \ldots, n_T\}$.

If $T = \text{𝒟}(A, B, f, n)$: Let $i \in N \cup \{\infty\}$. By Definitions 4 and 6 with $S_i$ as in Definition 4,

$$\mathcal{B}_{\text{𝒟}(A,B,f,n)}(i) = \textstyle\sum_{x \in S_i} P_A(x) = \sum_{x, \iota(x) = i} P_A(x).$$

Otherwise, assume the lemma holds for composition trees $T_1$ and $T_2$. If $T = T_1 \times T_2$, we have for $i \in N \setminus \{-n_T\}$,

$$\begin{aligned}
\mathcal{B}_{T_1 \times T_2}(i) &= \sum_{j,k \in N. j+k=i} \mathcal{B}_{T_1}(j) \cdot \mathcal{B}_{T_2}(k) \\
&\overset{\text{IH}}{=} \sum_{j,k \in N \ s.t. \ j+k=i} \left( \sum_{x_1 \in \mathcal{U}_{T_1} s.t. \ \iota_{T_1}(x_1) = j} \mathcal{B}_{T_1}(x_1) \right) \cdot \left( \sum_{x_2 \in \mathcal{U}_{T_2} s.t. \ \iota_{T_2}(x_2) = k} \mathcal{B}_{T_2}(x_2) \right) \\
&= \sum_{x=(x_1,x_2) \in \mathcal{U}_{T_1} \times \mathcal{U}_{T_2} \ s.t. \ \iota_{T_1}(x_1) + \iota_{T_2}(x_2) = i} \mathcal{B}_{T_1}(x_1) \cdot \mathcal{B}_{T_2}(x_2)
\end{aligned}$$

We know from Definition 6 that $\iota_T(x) = \iota_{T_1}(x_1) + \iota_{T_2}(x_2)$, since $\iota_T(x) \in \{-n+1, \ldots, n\}$.

$$\begin{aligned}
&= \sum_{x=(x_1,x_2) \in \mathcal{U}_{T_1} \times \mathcal{U}_{T_2} \ s.t. \ \iota_T(x) = i} \mathcal{B}_{T_1}(x_1) \cdot \mathcal{B}_{T_2}(x_2) \\
&\overset{\text{Definition 6}}{=} \sum_{x=(x_1,x_2) \in \mathcal{U} \ s.t. \ \iota_T(x) = i} \mathcal{B}_{T_1 \times T_2}(x).
\end{aligned}$$

For $i \in \{-n_T, \infty\}$ the proof follows analogously, where for $-n$ we have $j + k \leq -n$ and we know from Definition 6 that $\iota_T(x) = -n$ is equivalent to $\iota_{T_1}(x_1) + \iota_{T_2}(x_2) \leq -n$ and for $\infty$ we have $j + k > n$ and we know from Definition 6 that $\iota_T(x) = \infty$ is equivalent to $\iota_{T_1}(x_1) + \iota_{T_2}(x_2) \geq n$.

Figure 10 describes the algorithm of our bucketing that we suggest for practice.

If $T = \blacktriangledown T_1$, we have for $i \in \{-n_T, \ldots, -n_T/2 - 1, n_T/2 + 1, \ldots, n_T\}$, $\mathcal{B}_{\blacktriangledown T_1}(i) = 0 = \sum_{x \in \emptyset} P_{A_1}(x) = \sum_{x \in \mathcal{U}_1, \iota_{\blacktriangledown T_1} = i} P_{A_{T_1}}(x)$.

For $i = \infty$, we have $\mathcal{B}_{\blacktriangledown T_1}(\infty) = \mathcal{B}_{T_1}(\infty)$, so the statement follows from the IH. For $i \in \{-n_T/2 + 1, \ldots, n_T/2\}$ we have

$$\begin{aligned}
&\mathcal{B}_{\blacktriangledown T_1}(i) = \mathcal{B}_{T_1}(2i) + \mathcal{B}_{T_1}(2i-1) \\
&\overset{\text{IH}}{=} \textstyle\sum_{x \in \mathcal{U}_{T_1}, \iota_{T_1}(x) = 2i} P_{A_{T_1}}(x) + \sum_{x \in \mathcal{U}_{T_1}, \iota_{T_1}(x) = 2i-1} P_{A_{T_1}}(x) \\
&= \textstyle\sum_{x \in \mathcal{U}_{\blacktriangledown T_1} \iota_T(x) = i} P_{A_{\blacktriangledown T_1}}(x).
\end{aligned}$$

The statement for $\mathcal{B}_{\blacktriangledown T_1}(-n_T/2)$ follows analogously. □

We now state the first theorem of our paper: the buckets are sound.

**Theorem 1** (Buckets are sound)**.** *Let $X$ and $Y$ be two distributions and let $T_{X||Y}$ and $T_{Y||X}$ be valid composition trees with $A_{T_{X||Y}} = B_{T_{Y||X}} = X$ and $B_{T_{X||Y}} = A_{T_{Y||X}} = Y$.*

*Then for every $\varepsilon \geq 0$ and for any $\delta \geq \max\left(\delta_{T_{X||Y}}(\varepsilon), \delta_{T_{Y||X}}(\varepsilon)\right)$, $X$ and $Y$ are $(\varepsilon, \delta)$-ADP.*

The theorem follows quite trivially from the proof of Lemma 13 in the subsequent chapter. We still present a self-contained proof as it could be helpful in understanding the soundness of our privacy buckets.

*Proof.* We show that $\delta_{T_{X||Y}}(\varepsilon) \leq \delta$ implies $\delta \geq \sum_{x \in \mathcal{U}_{T_{X||Y}}} \max(P_X(x) - e^{\varepsilon} P_Y(x), 0)$ (one direction in Lemma 1); the proof for $T_{Y||X}$ and the other direction follows analogously. Let $n = n_{T_{X||Y}}$, $N = \{-n, \ldots, n\}$, $\mathcal{U} = \mathcal{U}_{T_{X||Y}}$ and $f = f_{T_{X||Y}}$. By definition,

$$\delta_{T_{X||Y}}(\varepsilon) = \sum_{i \in N} \left(\max\left(0, \mathcal{B}_{T_{X||Y}}(i) \cdot (1 - e^{\varepsilon}/f^i)\right)\right) + \mathcal{B}_{T_{X||Y}}(\infty).$$

We ignore $\mathcal{B}_{T_{X||Y}}(\infty)$ for now and apply Lemma 7 and get

$$\begin{aligned} &\sum_{i \in N} \left(\max\left(0, \sum_{x \in \mathcal{U}.\iota_{T_{X||Y}}(x) = i} P_X(x) \cdot (1 - \tfrac{e^{\varepsilon}}{f^i})\right)\right) \\ =&\sum_{i \in N.f^i > e^{\varepsilon}} \left(\sum_{x \in \mathcal{U}.\iota_{T_{X||Y}}(x) = i} P_X(x) \cdot (1 - \tfrac{e^{\varepsilon}}{f^i})\right) \end{aligned}$$

Using Lemma 6 we get

$$\sum_{x \in \mathcal{U}.\iota_T(x) = \in N.f^{\iota_{T_{X||Y}}(x)} > e^{\varepsilon}} \max\left(0, P_X(x) - e^{\varepsilon} P_Y(x)\right).$$

With $\mathcal{B}_{T_{X||Y}}(\infty)$ (where we also apply Lemma 7) we yield

$$\begin{aligned} &\sum_{x \in \mathcal{U}.\iota_{T_{X||Y}}(x) \in N.f^{\iota_{T_{X||Y}}(x)} > e^{\varepsilon}} \max\left(0, P_X(x) - e^{\varepsilon} P_Y(x)\right) \\ &+ \sum_{x \in \mathcal{U}.\iota_{T_{X||Y}}(x) = \infty} P_X(x) \\ \geq &\sum_{x \in \mathcal{U}} \max\left(0, P_X(x) - e^{\varepsilon} P_Y(x)\right). \end{aligned}$$

We repeat the calculation analogously for $T_{Y||X}$ and then we use Lemma 1 to see that $X$ and $Y$ are indeed $(\varepsilon, \delta)$-ADP.

□

# 4 Reducing and bounding the error

We have already presented a sound way of approximating a distribution pair by creating privacy buckets. Our calculations from the previous section lead to sound and, in many cases, better results than generic composition theorems from the literature. In this section we explore the precision of our results: we define error (correction) terms that help us to both find a lower bound on the differential privacy guarantee for the considered distributions even under manifold composition, and to find a tighter guarantee for differential privacy.

We distinguish between two types of error correction (EC) terms: the *real EC term* $\ell$ that captures the value we use to tighten our result in a sound way and the *virtual EC term* $\tilde{\ell}$ that captures the maximal influence an EC term can have. The virtual EC term accurately captures the difference between the probability an event $x$ appears to have in the alternative distribution (using the bucket factor) $\frac{P_A(x)}{f^{\iota}}$ and the probability that it actually has in the alternative $P_B(x)$. In some cases, however, we misplace an event such that it ends up in a bucket with an index that is too large: events $x$ that should not be considered for the

overall guarantee, i.e., that have $P_A(x) - e^{\varepsilon} P_B(x)) < 0$ can appear in a bucket with index $i$ s.t. $e^{\varepsilon} < f^i$. Thus, correctly calculating the EC term while possibly misplacing events can lead to wrong results.

There are two reasons for why events can be misplaced: First, when composing privacy buckets, events can be misplaced by one bucket. We take care of this by not including the EC terms of a certain number of buckets, depending on the number of compositions. Second, when events are put into the smallest bucket (with index $-n$), they can be arbitrarily "misplaced", particularly after a composition. To counter this effect, we introduce the real EC term, in which we do not include the error of the smallest bucket (with index $-n$).

## 4.1 Buckets with error correction terms

Our strategy is as follows. Assume two distributions $A$ and $B$: Whenever we add an event $x$ to a bucket $\mathcal{B}(i)$, we store the difference between the probability that the event occurs in $A$, adjusted by the bucket factor, and the probability that the same event occurs in $B$: $P_B(x) - \frac{P_A(x)}{f^i}$. Recall that the main purpose of the buckets is to keep track of the ratio between those two probabilities. We sum up all these *error correction terms* (or EC terms) per individual bucket $\mathcal{B}(i)$ and yield EC terms $\ell(i)$. We refer to Figure 4 (in Section 3.1) for a graphical intuition of our error correction.

As an example consider one bucket $\mathcal{B}(i)$, containing events $x \in S_i$ for a set $S_i$:

$$\begin{aligned}\frac{\mathcal{B}(i)}{f^i} - \ell(i) =& \frac{\sum_{x \in S_i} P_A(x)}{f^i} \\ & - \sum_{x \in S_i} \left( P_B(x) - \frac{P_A(x)}{f^i} \right) \\ =& \sum_{x \in S_i} P_B(x).\end{aligned}$$

Thus, only considering one additional value per bucket, we can precisely remember the probability that the events occurred in $B$ and we can then use this probability to calculate a more precise differential privacy guarantee. We omit the EC terms for the bucket $\mathcal{B}(\infty)$, as there is no bucket factor attached to it (so there is no value the error correction term could correct).

Although the error correction precisely captures the error per event $x$, we need to be careful which events we consider for calculating $\delta$. Consider the bucket $\mathcal{B}(j)$ with $f^{j-1} < e^{\varepsilon} \leq f^j$. If we were precise in our calculations, we would only consider *some* of the events from the bucket, namely the ones with $P_A(x) \leq e^{\varepsilon} P_B(x)$, but since we combined them all into one bucket, we cannot distinguish the individual events anymore. To retain a sound guarantee, we don't consider the EC term of this bucket when calculating $\delta$. Under composition this error slightly increases, as events can be "misplaced" by more than one bucket when we compose the buckets. Consequently, every composition increases the number of buckets for which we don't consider an EC term. Whenever events land in the bucket with index $-n$, an arbitrary "misplacement" can occur and our aforementioned strategy does not suffice. Thus, we distinguish between the *virtual EC term* $\tilde{\ell}$, which applies to index $-n$ and the *real EC term* $\ell$, where we always set $\ell(-n) = 0$. For our sound upper bound on $\delta$ we will use the real EC term $\ell$, and we will use the slightly too large $\tilde{\ell}$ to derive a lower bound on $\delta$.

For the composition, we want to calculate the error correction (EC) term for the combined events: given events $x_1$ and $x_2$ with (individual) error terms $P_{B_1}(x_1) - \frac{P_{A_1}(x)}{f^{\iota_1}}$ and $P_{B_2}(x_2) - \frac{P_{A_2}(x_2)}{f^{\iota_2}}$ we want (in the typical case, ignoring corner cases) to have an EC term for the pair of the form $P_{B_1 \times B_2}((x_1, x_2)) - \frac{P_{A_1 \times A_2}((x_1, x_2))}{f^{\iota_1 + \iota_2}}$. However, the buckets cannot keep track of the value for $P_{B_1 \times B_2}((x_1, x_2))$– recall that this is precisely why we have introduced the error terms. Fortunately, we can calculate the desired EC terms from the previous EC terms $\ell_{T_1}, \ell_{T_2}$, the bucket values $\mathcal{B}_{T_1}, \mathcal{B}_{T_2}$, and the bucket factor $f$ as

$$\ell_{T_1 \times T_2}(i) := \sum_{j+k=i} \frac{\mathcal{B}_{T_1}(j)}{f^j} \ell_{T_2}(k) + \frac{\mathcal{B}_{T_2}(k)}{f^k} \ell_{T_1}(j) + \ell_{T_1}(j) \ell_{T_2}(k).$$

Similarly, for the squaring, we quantify how the error terms change when we modify the buckets. Although each new bucket is composed of two previous buckets, the bucket factor actually only changes for one half of

the values: the evenly indexed buckets $\mathcal{B}_T(2i)$ with factor $f^{2i}$ are now moved into buckets $\mathcal{B}_{\blacktriangledown T}(i)$ with the same factor $(f^2)^i$ and thus their EC terms are still correct. The other half of buckets $\mathcal{B}_T(2i-1)$ with factor $f^{2i-1}$ are moved into the same buckets $\mathcal{B}_{\blacktriangledown T}(i)$ with factor $(f^2)^i$ and thus the EC terms need to be modified to capture this change in the bucket factor, based on the previous EC terms $\ell_T$ and bucket values $\mathcal{B}_T$:

$$\ell_{\blacktriangledown T}(i) := \ell_T(2i-1) + \mathcal{B}_T(2i-1)\left(\frac{1}{f^{2i-1}} - \frac{1}{f^{2i}}\right) + \ell_T(2i).$$

**Definition 7** (Privacy buckets with error correction terms)**.** *Let $T$ be a valid composition tree with $n := n_T$ and let $N = \{-n, \dots, n\}$. We define $A_T/B_T$* privacy buckets with EC terms *as as follows. $\mathcal{B}_T$, $f_T$, and $n_T$ are exactly defined as in Definition 4, and $\tilde{\ell}_T$, $\ell_T$, and $u_T$ are defined as follows*

$$\tilde{\ell}_{\mathscr{D}(A,B,f,n)}(i) := \begin{cases} \sum_{x \in \mathcal{U}, \iota(x)=i} P_B(x) - \frac{P_A(x)}{f^i} & i \in N \\ 0 & i = \infty \end{cases}$$

$$\ell_{\mathscr{D}(A,B,f,n)}(i) := \begin{cases} \tilde{\ell}_{\mathscr{D}(A,B,f,n)}(i) & i \in N \setminus \{-n\} \\ 0 & i \in \{-n, \infty\} \end{cases}$$

$$u_{\mathscr{D}(A,B,f,n)} := 1$$

*For composition we require that $f_{T_1} = f_{T_2}$ and we write $f = f_{T_1}$. To ease readability we define $V(j,k,x,y) = \frac{\mathcal{B}_{T_1}(j)}{f^j} y(k) + \frac{\mathcal{B}_{T_2}(k)}{f^k} x(j) + x(j)y(k)$ and based on $V$ we define the EC terms as*

$$\tilde{\ell}_{T_1 \times T_2}(i) := \begin{cases} \sum_{j+k=i} V(j,k,\tilde{\ell}_{T_1}, \tilde{\ell}_{T_2}) & i \in N \setminus \{-n\} \\ \sum_{j+k \le -n} V(j,k,\tilde{\ell}_{T_1}, \tilde{\ell}_{T_2}) & i = -n \\ 0 & i = \infty \end{cases}$$

$$\ell_{T_1 \times T_2}(i) := \begin{cases} \sum_{j+k=i} V(j,k,\ell_{T_1}, \ell_{T_2}) & i \in N \setminus \{-n\} \\ 0 & i \in \{-n, \infty\} \end{cases}$$

$$u_{T_1 \times T_2}(i) := u_{T_1} + u_{T_2}$$

*To ease the readability we define a function $W(i,x) := x(2i-1) + \mathcal{B}_1(2i-1)\left(\frac{1}{f_{T_1}^{2i-1}} - \frac{1}{f_{T_1}^{2i}}\right) + x(2i)$. We define the EC terms as*

$$\tilde{\ell}_{\blacktriangledown T_1}(i) := \begin{cases} W(i, \tilde{\ell}_{T_1}) & i \in [-n/2+1, n/2] \\ \tilde{\ell}_{T_1}(-n) & i = -n/2 \\ 0 & \textit{otherwise} \end{cases}$$

$$\ell_{\blacktriangledown T_1}(i) := \begin{cases} W(i, \ell_{T_1}) & i \in [-n/2+1, n/2] \\ 0 & \textit{otherwise} \end{cases}$$

$$u_{\blacktriangledown T_1}(i) := \lceil u_{T_1} \rceil + 1$$

## 4.2 Buckets and error correction terms per element

Before we can show the first helpful lemmas for the soundness of our error correction (EC) terms, we introduce the impact that each individual event $x$ has on the bucket terms that are influenced by $x$. We first simply define these terms per element separately and then continue by showing that each bucket value (and EC term) is simply the sum over the respective terms of all elements contributing to this bucket. This marks a significant step in the correctness (and tightness) of our results: Although we only consider a few values (one bucket value and one EC value per bucket) we still capture all individual events. The only exception to this precision then comes from misplaced events, which we will analyze subsequently. To distinguish terms per element from our previous (accumulated) terms, we mark terms considering only individual (atomic) events with a special symbol ⚛.

**Definition 8** (Privacy buckets with EC terms per element)**.** *Let $T$ be a valid composition tree with $n := n_T$ an $f := f_T$ and $N = \{-n, \dots, n\}$.*

*For $T = \mathbb{P}(A, B, f, n)$ with $\mathcal{U}_T =: \mathcal{U}$, we define for all $x \in \mathcal{U}$*

$$\mathcal{B}^{\mathbb{P}}_{\mathbb{P}(A,B,f,n)}(x) := P_A(x)$$

$$\tilde{\ell}^{\mathbb{P}}_{\mathbb{P}(A,B,f,n)}(x) := \begin{cases} P_B(x) - \frac{P_A(x)}{f^{\iota_{\mathbb{P}(A,B,f,n)}(x)}} & \iota_{\mathbb{P}(A,B,f,n)}(x) \in N, \\ 0 & \iota_{\mathbb{P}(A,B,f,n)}(x) = \infty, \end{cases}$$

$$\ell^{\mathbb{P}}_{\mathbb{P}(A,B,f,n)}(x) := \begin{cases} \tilde{\ell}^{\mathbb{P}}_{\mathbb{P}(A,B,f,n)}(x) & \iota_{\mathbb{P}(A,B,f,n)}(x) \in N \setminus \{-n\}, \\ 0 & \iota_{\mathbb{P}(A,B,f,n)}(x) \in \{-n, \infty\}. \end{cases}$$

*For $T = T_1 \times T_2$ with $\mathcal{U}_i = \mathcal{U}_{T_i}$ (for $i \in \{1, 2\}$), we define for all $x = (x_1, x_2)$ with $x_1 \in \mathcal{U}_1$ and $x_2 \in \mathcal{U}_2$*

$$\mathcal{B}^{\mathbb{P}}_{T_1 \times T_2}(x) := \mathcal{B}_{T_1}(x_1) \cdot \mathcal{B}_{T_2}(x_2)$$

*and we define the EC terms as*

$$\begin{aligned} &\text{if } \iota_{T_1 \times T_2}(x) \in \{-n, \dots, n\} \\ \tilde{\ell}^{\mathbb{P}}_{T_1 \times T_2}(x) &:= \left( \frac{\mathcal{B}^{\mathbb{P}}_{T_1}(x_1)}{f^{\iota_{T_1}(x_1)}} + \tilde{\ell}^{\mathbb{P}}_{T_1}(x_1) \right) \tilde{\ell}^{\mathbb{P}}_{T_2}(x_2) + \tilde{\ell}^{\mathbb{P}}_{T_1}(x_2) \left( \frac{\mathcal{B}^{\mathbb{P}}_{T_2}(x_2)}{f^{\iota_{T_2}(x_2)}} + \tilde{\ell}^{\mathbb{P}}_{T_2}(x_2) \right) - \tilde{\ell}^{\mathbb{P}}_{T_1}(x_1) \tilde{\ell}^{\mathbb{P}}_{T_2}(x_2) \\ &\text{if } \iota_{T_1 \times T_2}(x) \in \{\infty\} \\ &\tilde{\ell}^{\mathbb{P}}_{T_1 \times T_2}(x) := 0 \\ &\text{if } \iota_{T_1 \times T_2}(x) \in \{-n+1, \dots, n, \infty\} \\ \ell^{\mathbb{P}}_{T_1 \times T_2}(x) &:= \left( \frac{\mathcal{B}^{\mathbb{P}}_{T_1}(x_1)}{f^{\iota_{T_1}(x_1)}} + \ell^{\mathbb{P}}_{T_1}(x_1) \right) \ell^{\mathbb{P}}_{T_2}(x_2) + \ell^{\mathbb{P}}_{T_1}(x_2) \left( \frac{\mathcal{B}^{\mathbb{P}}_{T_2}(x_2)}{f^{\iota_{T_2}(x_2)}} + \ell^{\mathbb{P}}_{T_2}(x_2) \right) - \ell^{\mathbb{P}}_{T_1}(x_1) \ell^{\mathbb{P}}_{T_2}(x_2) \\ &\text{if } \iota_{T_1 \times T_2}(x) \in \{-n, \infty\} \\ &\ell^{\mathbb{P}}_{T_1 \times T_2}(x) := 0. \end{aligned}$$

*For a squaring node ($T = \blacktriangledown T_1$), we keep the bucket value as $\mathcal{B}^{\mathbb{P}}_{\blacktriangledown T_1}(x) := \mathcal{B}^{\mathbb{P}}_{T_1}(x_1)$ and we define the EC terms as follows for $f = f_{T_1}$:*

$$\begin{aligned} &\text{if } \iota_{T_1}(x) \in \{-n, \dots, n\} \\ \tilde{\ell}^{\mathbb{P}}_{\blacktriangledown T_1}(x) &:= \tilde{\ell}^{\mathbb{P}}_{T_1}(x) + \mathcal{B}^{\mathbb{P}}_{T_1}(x) \cdot \left( \frac{1}{f^{\iota_{T_1}(x)}} - \frac{1}{f^{2 \cdot \lceil \iota_{T_1}(x)/2 \rceil}} \right) \\ &\text{if } \iota_{T_1}(x) \in \{\infty\} \\ &\tilde{\ell}^{\mathbb{P}}_{\blacktriangledown T_1}(x) := 0 \\ \\ &\text{if } \iota_{T_1}(x) \in \{-n+1, \dots, n\} \\ \ell^{\mathbb{P}}_{\blacktriangledown T_1}(x) &:= \ell^{\mathbb{P}}_{T_1}(x) + \mathcal{B}^{\mathbb{P}}_{T_1}(x) \cdot \left( \frac{1}{f^{\iota_{T_1}(x)}} - \frac{1}{f^{2 \cdot \lceil \iota_{T_1}(x)/2 \rceil}} \right) \\ &\text{if } \iota_{T_1}(x) \in \{-n, \infty\} \\ &\ell^{\mathbb{P}}_{\blacktriangledown T_1}(x) := 0. \end{aligned}$$

We now show our first important lemma for the soundness of our buckets and EC terms: the terms we defined just previously indeed characterize the impact of each individual event on the overall bucket values and EC terms. These terms indeed are just the sum of the respective values per element for all elements of an index that equals the bucket index.

**Lemma 8** (All values are sums over atomic events)**.** *Let $T$ be a valid composition tree, labeled with $n \in \mathbb{N}$ . Then, the following statements hold for all $i \in \{-n, \dots, n, \infty\}$ and $x \in \mathcal{U}_T$:*

- $\mathcal{B}_T(i) = \sum_{x \in \mathcal{U} \text{ s.t. } \iota_T(x)=i} \mathcal{B}^{\circledast}_T(x)$
- $\tilde{\ell}_T(i) = \sum_{x \in \mathcal{U} \text{ s.t. } \iota_T(x)=i} \tilde{\ell}^{\circledast}_T(x)$
- $\ell_T(i) = \sum_{x \in \mathcal{U} \text{ s.t. } \iota_T(x)=i} \ell^{\circledast}_T(x)$

*Proof.* We show the lemma via structural induction over $T$.

**If $T = \circledast(A, B, f, n)$:** Let $i \in \{-n, \dots, n, \infty\}$ and $x \in \mathcal{U}_T$.

- By definition, $\mathcal{B}^{\circledast}_T(x) = P_A(x)$ (c.f., Definition 8). Thus, $\mathcal{B}^{\circledast}_T(i) = \sum_{x \text{ s.t. } \iota(x)=i} P_A(x) = \sum_{x \text{ s.t. } \iota(x)=i} \mathcal{B}^{\circledast}_T(x)$.
- If $i \in \{-n, \dots, n\}$, then $\tilde{\ell}_T(i) = \sum_{x \text{ s.t. } \iota(x)=i} P_B(x) - \frac{P_A(x)}{f^i} = \sum_{x \text{ s.t. } \iota(x)=i} \tilde{\ell}^{\circledast}_T(x)$. Otherwise $\tilde{\ell}_T(i) = 0 = \sum_{x \text{ s.t. } \iota(x)=i} 0 = \sum_{x \text{ s.t. } \iota(x)=i} \tilde{\ell}^{\circledast}_T(x)$.
- If $i \in \{-n+1, \dots, n\}$, then $\ell_T(i) = \sum_{x \text{ s.t. } \iota(x)=i} P_B(x) - \frac{P_A(x)}{f^i} = \sum_{x \text{ s.t. } \iota(x)=i} \ell^{\circledast}_T(x)$. Otherwise $\ell_T(i) = 0 = \sum_{x \text{ s.t. } \iota(x)=i} 0 = \sum_{x \text{ s.t. } \iota(x)=i} \ell^{\circledast}_T(x)$.

**If $T = T_1 \times T_2$:** We assume the lemma holds for the composition trees $T_1$ and $T_2$.
For $i \in \{-n+1, \dots, n\}$ and $x \in \mathcal{U}_T$ and $f := f_T$

$$\begin{aligned}
\mathcal{B}_{T_1 \times T_2}(i) &= \sum_{j,k \in \{-n,\dots,n\} \text{ s.t. } j+k=i} \mathcal{B}_{T_1}(j) \cdot \mathcal{B}_{T_2}(k) \\
&\overset{IV}{=} \sum_{j,k \in \{-n,\dots,n\} \text{ s.t. } j+k=i} \left( \sum_{x_1 \in \mathcal{U}_1 \text{s.t. } \iota_{T_1}(x_1)=j} \mathcal{B}^{\circledast}_{T_1}(x_1) \right) \cdot \left( \sum_{x_2 \in \mathcal{U}_2 \text{s.t. } \iota_{T_2}(x_2)=k} \mathcal{B}^{\circledast}_{T_2}(x_2) \right) \\
&= \sum_{x=(x_1,x_2) \in \mathcal{U}_1 \times \mathcal{U}_2 \text{ s.t. } \iota_{T_1}(x_1)+\iota_{T_2}(x_2)=i} \mathcal{B}^{\circledast}_{T_1}(x_1) \cdot \mathcal{B}^{\circledast}_{T_2}(x_2)
\end{aligned}$$

We know from Definition 6 that $\iota_T(x) = \iota_{T_1}(x_1) + \iota_{T_2}(x_2)$, since $\iota_T(x) \in \{-n+1, \dots, n\}$.

$$\begin{aligned}
&= \sum_{x=(x_1,x_2) \in \mathcal{U}_1 \times \mathcal{U}_2 \text{ s.t. } \iota_T(x)=i} \mathcal{B}^{\circledast}_{T_1}(x_1) \cdot \mathcal{B}^{\circledast}_{T_2}(x_2) \\
&= \sum_{x=(x_1,x_2) \in \mathcal{U} \text{ s.t. } \iota_T(x)=i} \mathcal{B}^{\circledast}_{T_1 \times T_2}(x).
\end{aligned}$$

For $i \in \{-n, \infty\}$ the proof follows analogously, where for $-n$ we have $j + k \leq -n$ and we know from Definition 6 that $\iota_T(x) = -n$ is equivalent to $\iota_{T_1}(x_1) + \iota_{T_2}(x_2) \leq -n$ and for $\infty$ we have $j + k > n$ and we know from Definition 6 that $\iota_T(x) = \infty$ is equivalent to $\iota_{T_1}(x_1) + \iota_{T_2}(x_2) \geq n$.

For the virtual error, we distinguish the following cases:

- $\iota_T(x) \in \{-n+1, \dots, n\}$. Then,

$$
\begin{aligned}
&\tilde{\ell}_{T_1 \times T_2}(i) \\
&= \sum_{(k,l) \in \{-n,\dots,n\}^2, k+l=i} \left( \frac{\mathcal{B}_{T_1}(k)}{f^k} + \tilde{\ell}_{T_1}(k) \right) \tilde{\ell}_{T_2}(l) + \tilde{\ell}_{T_1}(k) \left( \frac{\mathcal{B}_{T_2}(l)}{f^l} + \tilde{\ell}_{T_2}(l) \right) - \tilde{\ell}_{T_1}(k) \tilde{\ell}_{T_2}(l) \\
&= \sum_{(k,l) \in \{-n,\dots,n\}^2, k+l=i} \frac{\mathcal{B}_{T_1}(k)}{f^k} \tilde{\ell}_{T_2}(l) + \tilde{\ell}_{T_1}(k) \frac{\mathcal{B}_{T_2}(l)}{f^l} + \tilde{\ell}_{T_1}(k) \tilde{\ell}_{T_2}(l) \\
&= \sum_{(k,l) \in \{-n,\dots,n\}^2, k+l=i} \left( \frac{\sum_{x_1 \in \mathcal{U}_1 s.t.\ \iota_{T_1}(x_1)=k} \mathcal{B}^{\circledast}_{T_1}(x_1)}{f^k} \left( \sum_{x_2 \in \mathcal{U}_2 s.t.\ \iota_{T_2}(x_2)=l} \tilde{\ell}^{\circledast}_{T_2}(x_2) \right) \right. \\
&\quad + \left( \sum_{x_1 \in \mathcal{U}_1 s.t.\ \iota_{T_1}(x_1)=k} \tilde{\ell}^{\circledast}_{T_1}(x_1) \right) \frac{\sum_{x_2 \in \mathcal{U}_2 s.t.\ \iota_{T_2}(x_2)=l} \mathcal{B}^{\circledast}_{T_2}(x_2)}{f^l} \\
&\quad \left. + \left( \sum_{x_1 \in \mathcal{U}_1 s.t.\ \iota_{T_1}(x_1)=k} \tilde{\ell}^{\circledast}_{T_1}(x_1) \right) \left( \sum_{x_2 \in \mathcal{U}_2 s.t.\ \iota_{T_2}(x_2)=l} \tilde{\ell}^{\circledast}_{T_2}(x_2) \right) \right) \\
&= \sum_{(k,l) \in \{-n,\dots,n\}^2, k+l=i} \sum_{x_1 \in \mathcal{U}_1 s.t.\ \iota_{T_1}(x_1)=k} \sum_{x_2 \in \mathcal{U}_2 s.t.\ \iota_{T_2}(x_2)=l} \left( \frac{\mathcal{B}^{\circledast}_{T_1}(x_1)}{f^k} \tilde{\ell}^{\circledast}_{T_2}(x_2) \right. \\
&\quad \left. + \tilde{\ell}^{\circledast}_{T_1}(x_1) \frac{\mathcal{B}^{\circledast}_{T_2}(x_2)}{f^l} + \tilde{\ell}^{\circledast}_{T_1}(x_1) \tilde{\ell}^{\circledast}_{T_2}(x_2) \right) \\
&= \sum_{(x_1,x_2) \in \mathcal{U}_1 \times \mathcal{U}2 \ s.t.\ \iota_{T_1}(x_1) + \iota_{T_2}(x_2) = i} \left( \frac{\mathcal{B}^{\circledast}_{T_1}(x_1)}{f^{\iota_{T_1}(x_1)}} \tilde{\ell}^{\circledast}_{T_2}(x_2) + \tilde{\ell}^{\circledast}_{T_1}(x_1) \frac{\mathcal{B}^{\circledast}_{T_2}(x_2)}{f^{\iota_{T_2}(x_2)}} + \tilde{\ell}^{\circledast}_{T_1}(x_1) \tilde{\ell}^{\circledast}_{T_2}(x_2) \right)
\end{aligned}
$$

  We know from Definition 6 that $\iota_T(x) = \iota_{T_1}(x_1) + \iota_{T_2}(x_2)$, since $\iota_T(x) \in \{-n+1, \dots, n\}$.

$$
= \sum_{x \in \mathcal{U} \ s.t.\ \iota_T(x)=i} \tilde{\ell}^{\circledast}_{T_1 \times T_2}(x)
$$

- $\iota_T(x) = -n$. The proof of the case from above follows analogously with $k + l \leq -n$, since we know from Definition 6 that $\iota_T(x) = -n$ is equivalent to $\iota_{T_1}(x_1) + \iota_{T_2}(x_2) \leq -n$.
- $\iota_T(x) = \infty$.

$$
\begin{aligned}
&\tilde{\ell}_{T_1 \times T_2}(i) \\
= 0 = &\sum_{x \in \mathcal{U} \ s.t.\ \iota_T(x)=i} 0 \\
= &\sum_{x \in \mathcal{U} \ s.t.\ \iota_T(x)=i} \tilde{\ell}^{\circledast}_{T_1 \times T_2}(x).
\end{aligned}
$$

**If $T = \blacktriangledown T_1$:**

We assume the lemma holds for a composition tree $T_1$, we have for $i \in \{-n, \dots, -n/2 - 1, n/2 + 1, \dots, n\}$ and $x \in \mathcal{U}_T$

$$
\mathcal{B}_{\blacktriangledown T_1}(i) = 0 = \sum_{x \in \emptyset} \mathcal{B}^{\circledast}_{\blacktriangledown T_1}(x) = \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{\blacktriangledown T}=i} \mathcal{B}^{\circledast}_{\blacktriangledown T_1}(x)
$$

For $i = \infty$, we have

$$\mathcal{B}_{\blacktriangledown T_1}(\infty) = \mathcal{B}_{T_1}(\infty) \overset{\text{IH}}{=} \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{T_1}(x) = \infty} \mathcal{B}^{\circledast}_{T_1}(x) = \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{\blacktriangledown T_1}(x) = \infty} \mathcal{B}^{\circledast}_{\blacktriangledown T_1}(x).$$

For $i \in \{-n/2 + 1, \ldots, n/2\}$ we have

$$\begin{aligned}
\mathcal{B}_{\blacktriangledown T_1}(i) &= \mathcal{B}_{T_1}(2i) + \mathcal{B}_{T_1}(2i-1) \\
&\overset{\text{IH}}{=} \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{T_1}(x) = 2i} \mathcal{B}^{\circledast}_{T_1}(x) + \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{T_1}(x) = 2i-1} \mathcal{B}^{\circledast}_{T_1}(x) \\
&= \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{T_1}(x) = 2i} \mathcal{B}^{\circledast}_{\blacktriangledown T_1}(x) + \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{T_1}(x) = 2i-1} \mathcal{B}^{\circledast}_{\blacktriangledown T_1}(x) \\
&= \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{\blacktriangledown T}(x) = i} \mathcal{B}^{\circledast}_{\blacktriangledown T_1}(x).
\end{aligned}$$

For $i = -n/2$ we have

$$\begin{aligned}
\mathcal{B}_{\blacktriangledown T_1}(-n/2) &= \mathcal{B}_1(-n) \\
&\overset{\text{IH}}{=} \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{T_1}(x) = -n} \mathcal{B}_{T_1}(x) \\
&= \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{T_1}(x) = -n} \mathcal{B}^{\circledast}_{\blacktriangledown T_1}(x) \\
&= \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{\blacktriangledown T_1}(x) = -n/2} \mathcal{B}^{\circledast}_{\blacktriangledown T_1}(x).
\end{aligned}$$

We hence go forward to show the lemma for the EC terms.

For the EC terms and for $i \in \{-n, \ldots, -n/2 - 1, n/2 + 1, \ldots, n\}$

$$\tilde{\ell}_{\blacktriangledown T_1}(i) = 0 = \sum_{x \in \emptyset} \tilde{\ell}^{\circledast}_{\blacktriangledown T_1}(x) = \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{\blacktriangledown T_1} = i} \tilde{\ell}^{\circledast}_{\blacktriangledown T_1}(x)$$

For $i = \infty$, we have

$$\begin{aligned}
\tilde{\ell}_{\blacktriangledown T_1}(\infty) = 0 &= \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{\blacktriangledown T_1}(x) = \infty} 0 \\
&= \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{T_1}(x) = \infty} \tilde{\ell}^{\circledast}_{\blacktriangledown T_1}(x) = \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{\blacktriangledown T_1}(x) = \infty} \tilde{\ell}^{\circledast}_{\blacktriangledown T_1}(x).
\end{aligned}$$

For $i \in \{-n/2+1, \dots, n/2\}$, let $f := f_{T_1}$. Then we have

$$
\begin{aligned}
\tilde{\ell}_{\blacktriangledown T_1}(i) &= \tilde{\ell}_{T_1}(2i-1) + \mathcal{B}_{T_1}(2i-1)\left(\frac{1}{f^{2i-1}} - \frac{1}{f^{2i}}\right) + \tilde{\ell}_{T_1}(2i) \\
&\overset{\text{IH}}{=} \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{T_1}(x)=2i-1} \tilde{\ell}^{\circledast}_{T_1}(x) + \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{T_1}(x)=2i-1} \mathcal{B}^{\circledast}_{T_1}(x)\left(\frac{1}{f^{2i-1}} - \frac{1}{f^{2i}}\right) \\
&\quad + \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{T_1}(x)=2i} \tilde{\ell}^{\circledast}_{T_1}(x) \\
&= \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{T_1}(x)=2i-1} \tilde{\ell}^{\circledast}_{\blacktriangledown T_1}(x) - \mathcal{B}^{\circledast}_{T_1}(x) \cdot \left(\frac{1}{f^{\iota_{T_1}(x)}} - \frac{1}{f^{2\cdot\lceil \iota_{T_1}(x)/2 \rceil}}\right) \\
&\quad + \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{T_1}(x)=2i-1} \mathcal{B}^{\circledast}_{T_1}(x)\left(\frac{1}{f^{2i-1}} - \frac{1}{f^{2i}}\right) \\
&\quad + \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{T_1}(x)=2i} \tilde{\ell}^{\circledast}_{\blacktriangledown T_1}(x) - \mathcal{B}^{\circledast}_{T_1}(x) \cdot \left(\frac{1}{f^{\iota_{T_1}(x)}} - \frac{1}{f^{2\cdot\lceil \iota_{T_1}(x)/2 \rceil}}\right) \\
&= \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{T_1}(x)=2i-1} \tilde{\ell}^{\circledast}_{\blacktriangledown T_1}(x) - \mathcal{B}^{\circledast}_{T_1}(x) \cdot \left(\frac{1}{f^{2i-1}} - \frac{1}{f^{2i}}\right) + \mathcal{B}^{\circledast}_{T_1}(x)\left(\frac{1}{f^{2i-1}} - \frac{1}{f^{2i}}\right) \\
&\quad + \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{T_1}(x)=2i} \tilde{\ell}^{\circledast}_{\blacktriangledown T_1}(x) \\
&= \sum_{x \in \mathcal{U} \text{ s.t. } \iota_{\blacktriangledown T_1}(x)=i} \tilde{\ell}^{\circledast}_{\blacktriangledown T_1}(x)
\end{aligned}
$$

The proof for $\tilde{\ell}_{\blacktriangledown T_1}(i)$ in case $i = -n/2$ and the $\ell_{\blacktriangledown T_1}(i)$ follow analogously to the proof for $\tilde{\ell}_{\blacktriangledown T_1}(i)$ with the exception that the case $-n/2$ is analogous to the case $\infty$ instead to the cases $i \in \{-n+1, \dots, n\}$ for $\ell_{\blacktriangledown T_1}(i)$. □

With Lemma 8 we now have a powerful tool for proving a set of properties for our EC terms that will ultimately allow us to show the soundness of our results: We can relate every bucket value and every EC term to the underlying events and can thus analyze our properties per event.

## 4.3 Helpful properties of error correction terms

In this rather technical subsection we present and show a set of helpful properties of our EC terms that we require for our proof of soundness (and for our lower bound). We show that all error terms are positive (which means that not considering one of them can only increase the $\delta$ of our result), we show that our real EC term is always smaller than the virtual EC term, which we use for proving the soundness of the approximation (Lemma 13). Finally, we show that for every event $x$, the virtual EC term after an arbitrary amount of composition and squaring following the composition tree $T$ still precisely captures $P_B(x) - \frac{P_A(x)}{f^{\iota_T}}$.

**Lemma 9** (Positive real and virtual error correction terms)**.** *Let $T$ be a valid composition tree with $n := n_T$. Then for all $i \in \{-n, \dots, n, \infty\}$, both the real and virtual EC terms are positive, i.e., $\ell_T(i) \geq 0$ and $\tilde{\ell}_T(i) \geq 0$.*

*Proof.* We show the lemma via structural induction over $T$. For leaf nodes $T = \mathbb{P}(A, B, f, n)$, the real EC term of an initial bucketing is calculated as the sum of EC terms for each $x \in \mathcal{U}$, which are either $P_B(x) - \frac{P_A(x)}{f^{\iota_T(x)}}$ or 0. By definition we know that $P_A(x) \leq f^{\iota_T(x)} P_B(x)$, so all these values are positive. For composition $T_1 \times T_2$ with $f_{T_1} = f_{T_2} =: f$ we have either 0 or $V(j,k,x,y) = \frac{\mathcal{B}_{T_1}(j)}{f^j} y(k) + \frac{\mathcal{B}_{T_2}(k)}{f^k} x(j) + x(j)y(k)$, which is the sum and product of positive terms (the latter we know from the induction invariant). Analogously we notice

that for squaring $\blacktriangledown T_1$ with $f_{T_1} =: f$ we have either 0 or $\ell_{T_1}(2i-1) + \mathcal{B}_{T_1}(2i-1)\left(\frac{1}{f^{2i-1}} - \frac{1}{f^{2i}}\right) + \ell_{T_1}(2i)$, which again consists purely of positive terms (again via induction invariant).

More precisely, we distinguish the following cases:

**For $T = \circledcirc(A, B, f, n)$,** the real EC term of an initial bucketing is calculated as the sum of EC terms for each $x \in \mathcal{U}$, $\ell^{\circledast}_T(x) = P_B(x) - \frac{P_A(x)}{f^{\iota_T(x)}}$ if $\iota_T(x) \notin \{-n, \infty\}$ and 0 otherwise. For $\iota_T(x) \in \{-n, \dots, n\}$ by definition we have $P_A(x) \leq f^{\iota_T(x)} P_B(x)$. Thus, for all $i \in \{-n, \dots, n, \infty\}$ are positive, i.e., $\ell_T(i) \geq 0$.

**For $T = T_1 \times T_2$,** $\mathcal{B}_T$ with $f_{T_1} = f_{T_2} =: f$, by induction hypothesis, $\ell_{T_1}$ and $\ell_{T_2}$ are positive. We calculate the composed EC terms as either 0 (if $i \in \{-n, \infty\}$) or as

$$\ell_{T_1 \times T_2}(i) = \ell_{T_1 \times T_2}(i) = \sum_{j,k \ s.t. \ j+k=i} \left( \left( \frac{\mathcal{B}_{A_1}(j)}{f^j} \right) \ell_{T_2}(k) + \left( \frac{\mathcal{B}_{T_2}(k)}{f^k} \right) \ell_{T_1}(j) + \ell_{T_1}(j) \ell_{T_2}(k) \right),$$

which is positive as well since all the EC terms and all bucket terms are positive.

**For $T = \blacktriangledown T_1$,** We calculate, with $f := f_{T_1}$, the EC terms as either 0 (if $i \in \{-n, \dots, -n/2 - 1, n/2 + 1, \dots, n, \infty\}$) or as

$$\ell_{\blacktriangledown T_1}(i) = \blacktriangledown \ell_{T_1}(i) = \ell_{T_1}(2i-1) + \mathcal{B}_{T_1}(2i-1)\left(\frac{1}{f^{2i-1}} - \frac{1}{f^{2i}}\right) + \ell_{T_1}(2i),$$

which is positive as well since all the EC terms and all bucket terms are positive. [6] Analogously, we can show that the virtual EC terms $\tilde{\ell}$ are positive as well. □

We now show that the real EC term is smaller than the virtual EC term.

**Lemma 10** (The real error $\ell$ is smaller than the virtual error $\tilde{\ell}$)**.** *Let $T$ be a valid composition tree labeled $n \in \mathbb{N}$ with $\mathcal{U} := \mathcal{U}_T$. Then, the real error is always smaller than the virtual error: $\ell^{\circledast}_T(x) \leq \tilde{\ell}^{\circledast}_T(x)$.*

*Proof.* We show the lemma via structural induction over $T$.

**For $T = \circledcirc(A, B, f, n)$:** We know that $\tilde{\ell}^{\circledast}_{\circledcirc(A,B,f,n)}(x) \geq 0$. By definition, since $u_{\circledcirc(A,B,f,n)} = 1$, either $\ell^{\circledast}_{\circledcirc(A,B,f,n)}(x) = 0$ or $\ell^{\circledast}_{\circledcirc(A,B,f,n)}(x) = \tilde{\ell}^{\circledast}_{\circledcirc(A,B,f,n)}(x)$ holds. Thus, $\ell^{\circledast}_{\circledcirc(A,B,f,n)}(x) \leq \tilde{\ell}^{\circledast}_{\circledcirc(A,B,f,n)}(x)$.

**For $T = T_1 \times T_2$:** Let $\mathcal{U}_1$ be the universe of $T_1$ and $\mathcal{U}_2$ be the universe of $T_2$. By induction hypothesis, $\ell^{\circledast}_{T_1} \leq \tilde{\ell}^{\circledast}_{T_1}$ and $\ell^{\circledast}_{T_2} \leq \tilde{\ell}^{\circledast}_{T_2}$. Let $f = f_{T_1} = f_{T_2}$. For $\iota_T(x) = -n$, $\ell^{\circledast}_{T_1 \times T_2}(x) = 0$. By Lemma 9 we know that $0 \leq \tilde{\ell}^{\circledast}_{T_1 \times T_2}(x)$, hence $\ell^{\circledast}_{T_1 \times T_2}(x) = 0 \leq \tilde{\ell}^{\circledast}_{T_1 \times T_2}(x)$. For $\iota_{T_1 \times T_2}(x) \neq -n$, with $x_1 \in \mathcal{U}_1$ and $x_2 \in \mathcal{U}_2$ we have

$$\begin{aligned}
\ell^{\circledast}_{T_1 \times T_2}(x) &= \left( \frac{P_{A_1}(x_1)}{f^{\iota_{T_1}(x_1)}} + \ell^{\circledast}_{T_1}(x_1) \right) \ell^{\circledast}_{T_2}(x_2) + \left( \frac{P_{A_2}(x_2)}{f^{\iota_{T_2}(x_2)}} + \ell^{\circledast}_{T_2}(x_2) \right) \ell^{\circledast}_{T_1}(x_1) - \ell^{\circledast}_{T_1}(x_1) \ell^{\circledast}_{T_2}(x_2) \\
&= \left( \frac{P_{A_1}(x_1)}{f^{\iota_{T_1}(x_1)}} \right) \underbrace{\ell^{\circledast}_{T_2}(x_2)}_{\overset{\text{IH}}{\leq} \tilde{\ell}^{\circledast}_{T_2}(x_2)} + \left( \frac{P_{A_2}(x_2)}{f^{\iota_{T_2}(x_2)}} \right) \underbrace{\ell^{\circledast}_{T_1}(x_1)}_{\overset{\text{IH}}{\leq} \tilde{\ell}^{\circledast}_{T_1}(x_1)} + \underbrace{\ell^{\circledast}_{T_1}(x_1)}_{\overset{\text{IH}}{\leq} \tilde{\ell}^{\circledast}_{T_1}(x_1)} \underbrace{\ell^{\circledast}_{T_2}(x_2)}_{\overset{\text{IH}}{\leq} \tilde{\ell}^{\circledast}_{T_2}(x_2)} \\
&\overset{\text{IH}}{\leq} \left( \frac{P_{A_1}(x_1)}{f^{\iota_{T_1}(x_1)}} \right) \tilde{\ell}^{\circledast}_{T_2}(x_2) + \left( \frac{P_{A_2}(x_2)}{f^{\iota_{T_2}(x_2)}} \right) \tilde{\ell}^{\circledast}_{T_1}(x_1) + \tilde{\ell}^{\circledast}_{T_1}(x_1) \tilde{\ell}^{\circledast}_{T_2}(x_2) \\
&= \tilde{\ell}^{\circledast}_{T_1} \times \tilde{\ell}^{\circledast}_{T_2}(x) = \tilde{\ell}^{\circledast}_{T_1 \times T_2}(x)
\end{aligned}$$

[6] Note that in the case $-n/2$ there is only one term instead of two. This term, however, is still positive.

**For $T = \blacktriangledown T_1$:** This case directly holds by induction hypothesis, as the squaring operation is analogously defined for the real and the virtual error.

□

We now show our main lemma for the lower bound on $\delta$: the virtual EC term is precise for any event with an index other than $\infty$. We can directly use this lemma to get a lower bound for $\delta$ if we ignore the bucket with index $\infty$. Note that although the virtual error is precise on a per-event basis, events can still be misplaced and thus negatively contribute to $\delta$ if we use the virtual EC term. For our upper bound on $\delta$ we circumvent this problem by over-approximating misplaced events (using the real EC term) and by not using EC terms in some buckets with a bucket factor $f^i$ close to $e^\varepsilon$.

**Lemma 11** (Characterizing the virtual error after compositions and rescaling)**.** *Let $T$ be a valid composition tree with $A := A_T$, $B := B_T$, $\mathcal{U} := \mathcal{U}_T$, $n := n_T$, and $f := f_T$. Then, for $x \in \mathcal{U}$ with $\iota_T(x) \neq \infty$ we have*

$$\tilde{\ell}^{\circledast}_T(x) = P_B(x) - \frac{P_A(x)}{f^{\iota_T(x)}}$$

*Proof.* We show the lemma via structural induction over $T$. For $T = \circledcirc(A, B, f, n)$, the statement follows by construction:

$$\tilde{\ell}^{\circledast}_T(x), = P_B(x) - \frac{P_A(x)}{f^{\iota_T(x)}}$$

and $f_{\circledcirc(A,B,f,n)} = f$.

For $T = T_1 \times T_2$ with $A_i := A_{T_i}$, $B_i := B_{T_i}$, $\mathcal{U}_i := \mathcal{U}_{T_i}$, $f := f_T$, and set $A := A_1 \times A_2$ and $B := B_1 \times B_2$. By induction hypothesis, the statement holds for $\tilde{\ell}^{\circledast}_{T_1}$ and $\tilde{\ell}^{\circledast}_{T_2}$. By definition of the EC term composition we get for all $x_1 \in \mathcal{U}_1$ and $x_2 \in \mathcal{U}_2$

$$\begin{aligned}
\tilde{\ell}^{\circledast}_{T_1 \times T_2}(x) =& \left(\frac{P_{A_1}(x_1)}{f^{\iota_{T_1}(x_1)}} + \tilde{\ell}^{\circledast}_{T_1}(x_1)\right)\tilde{\ell}^{\circledast}_{T_2}(x_2) + \left(\frac{P_{A_2}(x_2)}{f^{\iota_{T_2}(x_2)}} + \tilde{\ell}^{\circledast}_{T_2}(x_2)\right)\tilde{\ell}^{\circledast}_{T_1}(x_1) - \tilde{\ell}^{\circledast}_{T_1}(x_1)\tilde{\ell}^{\circledast}_{T_2}(x_2) \\
=& \frac{P_{A_1}(x_1)}{f^{\iota_{T_1}(x_1)}} \cdot \tilde{\ell}^{\circledast}_{T_2}(x_2) + \frac{P_{A_2}(x_2)}{f^{\iota_{T_2}(x_2)}} \cdot \tilde{\ell}^{\circledast}_{T_1}(x_1) + \tilde{\ell}^{\circledast}_{T_1}(x_1)\tilde{\ell}^{\circledast}_{T_2}(x_2) \\
\overset{\text{IH}}{=}& \frac{P_{A_1}(x_1)}{f^{\iota_{T_1}(x_1)}} \cdot \left(P_{B_2}(x_2) - \frac{P_{A_2}(x_2)}{f^{\iota_{T_2}(x_2)}}\right) \\
&+ \frac{P_{A_2}(x_2)}{f^{\iota_{T_2}(x_2)}} \cdot \left(P_{B_1}(x_1) - \frac{P_{A_1}(x_1)}{f^{\iota_{T_1}(x_1)}}\right) \\
&+ \left(P_{B_1}(x_1) - \frac{P_{A_1}(x_1)}{f^{\iota_{T_1}(x_1)}}\right) \cdot \left(P_{B_2}(x_2) - \frac{P_{A_2}(x_2)}{f^{\iota_{T_2}(x_2)}}\right) \\
=& \frac{P_{A_1}(x_1)}{f^{\iota_{T_1}(x_1)}} \cdot P_{B_2}(x_2) - \frac{P_A(x)}{f^{\iota_T(x)}} \\
&+ \frac{P_{A_2}(x_2)}{f^{\iota_{T_2}(x_2)}} \cdot P_{B_1}(x_1) - \frac{P_A(x)}{f^{\iota_T(x)}} \\
&+ P_B(x) - \frac{P_{A_1}(x_1)}{f^{\iota_{T_1}(x_1)}} \cdot P_{B_2}(x_2) - \frac{P_{A_2}(x_2)}{f^{\iota_{T_2}(x_2)}} \cdot P_{B_1}(x_1) + \frac{P_A(x)}{f^{\iota_T(x)}} \\
=& P_B(x) - \frac{P_A(x)}{f^{\iota_T(x)}}
\end{aligned}$$

For $T = \blacktriangledown T_1$, where $T_1$ is a composition tree over the distributions $A/B$ over the universe $\mathcal{U}$, we know that for all $x \in \mathcal{U}$, $\iota_{\blacktriangledown T_1}(x) \in \{-n/2, \ldots, n/2\} \cup \{\infty\}$. Since the index $\infty$ is excluded in our lemma, we focus on the remaining values for the index. Note that the bucket factor in this case changes from $f_{T_1} =: f$ (of the child node) to $f_{\blacktriangledown T_1} = (f_{T_1})^2 = f^2$ (of the squaring node). By induction hypothesis, we have

$$\tilde{\ell}^{\circledast}_{T_1}(x) = P_B(x) - \frac{P_{A_1}(x)}{f^{\iota_T(x)}}$$

Consequently and since $\iota_{\blacktriangledown T_1}(x) \in \{-n/2, \ldots, n/2\}$ and $\mathcal{B}^{\text{❁}}_{T_1} = P_A(x)$, we get,

$$\begin{aligned}
\tilde{\ell}^{\text{❁}}_{\blacktriangledown T_1}(x) &= \tilde{\ell}^{\text{❁}}_{T_1}(x) + \mathcal{B}^{\text{❁}}_{T_1}(x) \cdot \left( \frac{1}{f^{\iota_{T_1}(x)}} - \frac{1}{f^{2 \cdot \lceil \iota_{T_1}(x)/2 \rceil}} \right) \\
&\overset{\text{IH}}{=} P_B(x) - \frac{P_A(x)}{f^{\iota_{T_1}(x)}} + P_A(x) \cdot \left( \frac{1}{f^{\iota_{T_1}(x)}} - \frac{1}{f^{2 \cdot \lceil \iota_{T_1}(x)/2 \rceil}} \right) \\
&= P_B(x) - \frac{P_A(x)}{f^{\iota_{T_1}(x)}} + P_A(x) \cdot \left( \frac{1}{f^{\iota_{T_1}(x)}} - \frac{1}{f^{2\iota_T(x)}} \right) \\
&= P_B(x) - \frac{P_A(x)}{(f^2)^{\iota_T(x)}}.
\end{aligned}$$

□

## 4.4 The approximated delta with error correction

Finally, we define how to calculate a sound upper bound on $\delta$ based on privacy buckets with EC terms. We note that when using the real EC term, events cannot harm the soundness by being misplaced as a result of parts of the event having been placed in the smallest bucket (with index $-n$). However, every composition can misplace events into the next larger bucket. This slight misplacement poses a problem for a small number of buckets with a bucket factor $f^i$ just slightly larger than $e^\varepsilon$, as they can now contain events that should have been placed in a lower bucket (with factor $f^{i^*} < e^\varepsilon$) and that now actually have a negative contribution to $\delta$: $P_A(x) - e^\varepsilon P_B(x) < 0$. All composition trees for privacy buckets carry a value $u = 1$ at each leaf that increases by 1 for every composition and that is halved by squaring. If $j_\varepsilon$ is the index of the bucket with the smallest bucket factor larger than $e^\varepsilon$, we don't consider the the EC term for buckets with index $i < j_\varepsilon + u$ and instead fall back to Definition 5 for those buckets. For the remaining buckets with $i \geq j_\varepsilon + u$, which typically is the vast majority of buckets, we make use of the real EC term to reduce the error.

**Definition 9** (Approximated delta with error correction)**.** *Let $T$ be a valid composition tree with $A := A_T$, $\mathcal{U} := \mathcal{U}_T$, $n := n_T$, and $f := f_T$.*

*We define $\delta_T(\varepsilon)$ with $j_\varepsilon \in \mathbb{N}$ such that $f^{j_\varepsilon - 1} < e^\varepsilon \leq f^{j_\varepsilon}$ as*

$$\begin{aligned}
\delta_T(\varepsilon) := &\sum_{i \in \{j_\varepsilon, \ldots, j_\varepsilon + u_T - 1\}} \mathcal{B}_T(i) - \frac{e^\varepsilon \mathcal{B}_T(i)}{f^i} \\
&+ \sum_{i \in \{j_\varepsilon + u_T, \ldots, n\}} \left( \mathcal{B}_T(i) - e^\varepsilon \left( \frac{\mathcal{B}_T(i)}{f^i} + \ell_T(i) \right) \right) + \mathcal{B}_T(\infty)
\end{aligned}$$

*Moreover, for all individual events $x \in \mathcal{U}$ we define*

$$\delta^{\text{❁}}_T(x, \varepsilon) := \begin{cases}
P_A(x) \cdot \left(1 - \frac{e^\varepsilon}{f^{\iota_T(x)}}\right) & 1.\ \textit{if } j_\varepsilon \leq \iota_T(x) \leq j_\varepsilon + u_T - 1 \\
P_A(x) - e^\varepsilon \left( \frac{P_A(x)}{f^{\iota_T(x)}} + \ell^{\text{❁}}_T(x) \right) & 2.\ \textit{if } j_\varepsilon + u_T \leq \iota_T(x) \leq n \\
P_A(x) & 3.\ \textit{if } \iota_T(x) = \infty \\
0 & 4.\ \textit{otherwise}
\end{cases}$$

*Let for any composition tree $T$, $\varepsilon \geq 0$ and $j_\varepsilon$ s.t., $f_T^{j_\varepsilon - 1} < e^\varepsilon \leq f_T^{j_\varepsilon}$,*

$$\delta^{\text{low}}_T := \textstyle\sum_{i \in \{j_\varepsilon, \ldots, n_T\}} \max \left( 0, \mathcal{B}_T(i) - e^\varepsilon \left( \frac{\mathcal{B}_T(i)}{(f_T)^i} + \tilde{\ell}_T(i) \right) \right)$$

Note that if $j > n$, we only consider elements in the bucket $B_\infty$.

Next we show that the real EC terms are bounded by the value of $u_T$: For every event $x$ the real EC term $\ell^{\text{❁}}_T(x)$ can never exceed a fraction of $\frac{1}{f^{\iota_T(x)-u}} - \frac{1}{f^{\iota_T(x)}}$ of the probability of the event. Intuitively, this means that the value of the real EC term can never be larger than what a *misplacement by u buckets* would result in.

**Lemma 12** (An upper bound for $\ell$)**.** *Let $T$ be a valid composition tree with $A := A_T$, $\mathcal{U} := \mathcal{U}_T$, $f := f_T$, and $u := u_T$. Let $\varepsilon \geq 0$ and with $j_\varepsilon \in \mathbb{N}$ such that $f^{j_\varepsilon - 1} < e^\varepsilon \leq f^{j_\varepsilon}$.*

*If $j_\varepsilon + u \leq \iota_T(x) \neq \infty$ ($x \in \mathcal{U}$), then the EC term never makes a negative contribution to the approximated delta with EC: $\ell^{\text{❁}}_T(x) \leq P_A(x)\left(\frac{1}{f^{\iota_T(x)-u}} - \frac{1}{f^{\iota_T(x)}}\right)$.*

*Proof.* We show the lemma via structural induction over $T$.

**Let $T = \text{𝔼}(A, B, f, n)$.** If $\iota_{\text{𝔼}(A,B,f,n)}(x) = -n$ then

$$\ell^{\text{❁}}_{\text{𝔼}(A,B,f,n)}(x) = 0 \leq P_A(x) \cdot \left(\frac{1}{f^{-n-1}} - \frac{1}{f^{-n}}\right).$$

Otherwise, if $\iota_{\text{𝔼}(A,B,f,n)}(x) > -n$, we know that by definition of $\iota_{\text{𝔼}(A,B,f,n)}(x)$ we have $f^{\iota_{\text{𝔼}(A,B,f,n)}(x)-1} P_B(x) \leq P_A(x)$

$$\begin{aligned}\ell^{\text{❁}}_{\text{𝔼}(A,B,f,n)}(x) =& P_B(x) - \frac{P_A(x)}{f^{\iota_{\text{𝔼}(A,B,f,n)}(x)}} \\ \leq& \frac{P_A(x)}{f^{\iota_{\text{𝔼}(A,B,f,n)}(x)-1}} - \frac{P_A(x)}{f^{\iota_{\text{𝔼}(A,B,f,n)}(x)}}.\end{aligned}$$

**Let $T = T_1 \times T_2$.** If $\iota_{T_1\times T_2}(x) = -n$, then $\ell^{\text{❁}}_{T_1\times T_2}(x) = 0 \leq \frac{P_A(x)}{f^{\iota_{T_1\times T_2}(x) - u_{T_1\times T_2}}} - \frac{P_A(x)}{f^{\iota_{T_1\times T_2}(x)}}$, since $u_{T_1\times T_2} \geq 0$. Otherwise, by induction hypothesis, the statement holds for $\ell_{T_1}$ $\ell_{T_2}$. For $x_1 \in \mathcal{U}_1$ and $x_2 \in \mathcal{U}_2$ we know that $\iota_T(x) = \iota_{T_1}(x_1) + \iota_{T_2}(x_2)$, since $\iota_T \neq \infty$. Moreover, we know that $P_A(x) = P_{A_1}(x_1) \cdot P_{A_2}(x_2)$ and $u := u_{T_1\times T_2} = u_{T_1} + u_{T_2}$ and we get

$$\begin{aligned}\ell^{\text{❁}}_{T_1\times T_2}(x) &= \left(\frac{P_{A_1}(x_1)}{f^{\iota_{T_1}(x_1)}} + \ell^{\text{❁}}_{T_1}(x_1)\right)\ell^{\text{❁}}_{T_2}(x_2) + \left(\frac{P_{A_2}(x_2)}{f^{\iota_{T_2}(x_2)}} + \ell^{\text{❁}}_{T_2}(x_2)\right)\ell^{\text{❁}}_{T_1}(x_1) - \ell^{\text{❁}}_{T_1}(x_1)\ell^{\text{❁}}_{T_2}(x_2) \\ &= \left(\frac{P_{A_1}(x_1)}{f^{\iota_{T_1}(x_1)}}\right)\ell^{\text{❁}}_{T_2}(x_2) + \left(\frac{P_{A_2}(x_2)}{f^{\iota_{T_2}(x_2)}}\right)\ell^{\text{❁}}_{T_1}(x_1) + \ell^{\text{❁}}_{T_1}(x_1)\ell^{\text{❁}}_{T_2}(x_2)\end{aligned}$$

$$
\begin{aligned}
&\overset{\text{IH}}{\leq} \left(\frac{P_{A_1}(x_1)}{f^{\iota_{T_1}(x_1)}}\right)\left(\frac{P_{A_2}(x_2)}{f^{\iota_{T_2}(x_2)-(u-u_{T_1})}} - \frac{P_{A_2}(x_2)}{f^{\iota_{T_2}(x_2)}}\right)\\
&\quad + \left(\frac{P_{A_2}(x_2)}{f^{\iota_{T_2}(x_2)}}\right)\left(\frac{P_{A_1}(x_1)}{f^{\iota_{T_1}(x_1)-u_{T_1}}} - \frac{P_{A_1}(x_1)}{f^{\iota_{T_1}(x_1)}}\right)\\
&\quad + \left(\frac{P_{A_1}(x_1)}{f^{\iota_{T_1}(x_1)-u_{T_1}}} - \frac{P_{A_1}(x_1)}{f^{\iota_{T_1}(x_1)}}\right)\left(\frac{P_{A_2}(x_2)}{f^{\iota_{T_2}(x_2)-(u-u_{T_1})}} - \frac{P_{A_2}(x_2)}{f^{\iota_{T_2}(x_2)}}\right)\\
&= \frac{P_A(x)}{f^{\iota_T(x)-(u-u_{T_1})}} - \frac{P_A(x)}{f^{\iota_T(x)}}\\
&\quad + \frac{P_A(x)}{f^{\iota_{T_2}(x_2)+\iota_{T_1}(x_1)-u_{T_1}}} - \frac{P_A(x)}{f^{\iota_{T_2}(x_2)+\iota_{T_1}(x_1)}}\\
&\quad + \frac{P_A(x)}{f^{\iota_T(x)-u_{T_1}-(u-u_{T_1})}} - \frac{P_A(x)}{f^{\iota_T(x)-(u-u_{T_1})}} - \frac{P_A(x)}{f^{\iota_T(x)-u_{T_1}}} + \frac{P_A(x)}{f^{\iota_T(x)}}\\
&= \frac{P_A(x)}{f^{\iota_T(x)-u_{T_2}}} - \frac{P_A(x)}{f^{\iota_T(x)}}\\
&\quad + \frac{P_A(x)}{f^{\iota_T(x)-u_{T_1}}} - \frac{P_A(x)}{f^{\iota_T(x)}}\\
&\quad + \frac{P_A(x)}{f^{\iota_T(x)-u}} - \frac{P_A(x)}{f^{\iota_T(x)-u_{T_2}}} - \frac{P_A(x)}{f^{\iota_T(x)-u_{T_1}}} + \frac{P_A(x)}{f^{\iota_T(x)}}\\
&= \frac{P_A(x)}{f^{\iota_T(x)-u}} - \frac{P_A(x)}{f^{\iota_T(x)}}
\end{aligned}
$$

**Let $T = \blacktriangledown T_1$.** For $f = f_{T_1}$ we have $f_{\blacktriangledown T_1} = f^2$ and $u = u_{\blacktriangledown T_1} = \lceil u_{T_1}/2 \rceil + 1$. We know that $\ell^{\text{❄}}_{\blacktriangledown T_1}(x) = \ell^{\text{❄}}_{T_1}(x) + \mathcal{B}^{\text{❄}}_{T_1}(x) \cdot \left(\frac{1}{f^{\iota_{T_1}(x)}} - \frac{1}{f^{2\cdot\lceil \iota_{T_1}(x)/2\rceil}}\right)$. Since we excluded $\iota_T(x) = \infty = \iota_{T_1}$ and $j_\varepsilon + u \leq \iota_T(x)$, we know that $\iota_T(x) \in \{0, \ldots, n/2\}$.

Thus,

$$
\begin{aligned}
\ell^{\text{❄}}_{\blacktriangledown T_1}(x) &= \ell^{\text{❄}}_{T_1}(x) + \mathcal{B}^{\text{❄}}_{T_1}(x) \cdot \left(\frac{1}{f^{\iota_{T_1}(x)}} - \frac{1}{f^{2\cdot\lceil \iota_{T_1}(x)/2\rceil}}\right)\\
&\overset{\text{IH}}{\leq} \frac{P_A(x)}{f^{\iota_{T_1}(x)-u_1}} - \frac{P_A(x)}{f^{\iota_{T_1}(x)}} + \mathcal{B}^{\text{❄}}_{T_1}(x) \cdot \left(\frac{1}{f^{\iota_{T_1}(x)}} - \frac{1}{f^{2\cdot\lceil \iota_{T_1}(x)/2\rceil}}\right)\\
&= \frac{P_A(x)}{f^{\iota_{T_1}(x)-u_1}} - \frac{P_A(x)}{f^{\iota_{T_1}(x)}} + \frac{P_A(x)}{f^{\iota_{T_1}(x)}} - \frac{P_A(x)}{f^{2\cdot\lceil \iota_{T_1}(x)/2\rceil}}\\
&= \frac{P_A(x)}{(f^2)^{\frac{\iota_{T_1}(x)-u_1}{2}}} - \frac{P_A(x)}{(f^2)^{\iota_T(x)}}\\
&\leq \frac{P_A(x)}{(f^2)^{\lceil \iota_{T_1}(x)/2\rceil - (\lceil u_1/2\rceil + 1)}} - \frac{P_A(x)}{(f^2)^{\iota_T(x)}}\\
&= \frac{P_A(x)}{(f^2)^{\iota_T(x)-u}} - \frac{P_A(x)}{(f^2)^{\iota_T(x)}}
\end{aligned}
$$

□

From Lemma 12 we can deduct that no event in a bucket with index $i \geq j_\varepsilon + u$ can have a negative impact on $\delta$. Since moreover for each event we consider an impact that is at least as large as the actual impact of the event (as in the precise calculation of $\delta$ from Lemma 1) we can show the soundness of our result:

**Lemma 13** (Soundness of the approximated delta with error correction)**.** *Let $T$ be a valid composition tree with $A := A_T$, $B := B_T$, and $\mathcal{U} := \mathcal{U}_T$. Then, for all $\varepsilon \geq 0$, the following statement holds:*

$$\delta_T(\varepsilon) \geq \sum_{x \in \mathcal{U}} \max\left(0, P_A(x) - e^{\varepsilon} P_B(x)\right)$$

*Proof.* Let $f = f_T$, $j_\varepsilon \in \mathbb{N}$, s.t. $f^{j_\varepsilon - 1} < e^{\varepsilon} \leq f^{j_\varepsilon}$.

We first show that $\delta_T(\varepsilon) = \sum_{x \in \mathcal{U}} \delta_T^{\text{❁}}(x, \varepsilon)$. Let $N^- = \{j_\varepsilon, \ldots, j_\varepsilon + u - 1\}$ and $N^+ = \{j_\varepsilon + u, \ldots, n\}$.

$$\begin{aligned}
\delta_T(\varepsilon) = & \sum_{i \in N^-} \mathcal{B}_T(i) \cdot \left(1 - \frac{e^{\varepsilon}}{f^i}\right) \\
& + \sum_{i \in N^+} \left(\mathcal{B}_T(i) - e^{\varepsilon}\left(\frac{\mathcal{B}_T(i)}{f^i} + \ell_T(i)\right)\right) + \mathcal{B}_T(\infty) \\
= & \sum_{i \in N^-} \left(\sum_{x \in \mathcal{U} \ s.t. \ \iota_T(x) = i} \mathcal{B}_T^{\text{❁}}(x)\right) \cdot \left(1 - \frac{e^{\varepsilon}}{f^i}\right) \\
& + \sum_{i \in N^+} \left(\left(\sum_{x \in \mathcal{U} \ s.t. \ \iota_T(x) = i} \mathcal{B}_T^{\text{❁}}(x)\right)\right. \\
& \left. - e^{\varepsilon}\left(\frac{\left(\sum_{x \in \mathcal{U} \ s.t. \ \iota_T(x) = i} \mathcal{B}_T^{\text{❁}}(x)\right)}{f^i} + \left(\sum_{x \in \mathcal{U} \ s.t. \ \iota_T(x) = i} \ell_T^{\text{❁}}(x)\right)\right)\right) \\
& + \sum_{x \in \mathcal{U} \ s.t. \ \iota_T(x) = \infty} \mathcal{B}_T^{\text{❁}}(x) \\
= & \sum_{x \in \mathcal{U}, \iota_T(x) \in N^-} \left(P_A(x) \cdot \left(1 - \frac{e^{\varepsilon}}{f^{\iota_T(x)}}\right)\right) \\
& + \sum_{x \in \mathcal{U}, \iota_T(x) \in N^-} P_A(x) - e^{\varepsilon}\left(\frac{P_A(x)}{f^{\iota_T(x)}} + \ell_T^{\text{❁}}(x)\right) \\
& + \sum_{x \in \mathcal{U} \ s.t. \ \iota_T(x) = \infty} P_A(x) \\
= & \sum_{x \in \mathcal{U}} \delta_T^{\text{❁}}(x, \varepsilon)
\end{aligned}$$

We next distinguish the the four cases of the definition of $\delta_T^{\text{❁}}(x, \varepsilon)$.

**Case 1.** This case occurs if $j_\varepsilon \leq \iota_T(x) \leq j_\varepsilon + u_T - 1$. By Lemma 6, we know the following

$$\begin{aligned}
& P_A(x) \leq f^{\iota_T(x)} P_B(x) \\
\Leftrightarrow \ & \frac{P_A(x)}{f^{\iota_T(x)}} \leq P_B(x) \\
\Leftrightarrow \ & P_A(x) - e^{\varepsilon}\frac{P_A(x)}{f^{\iota_T(x)}} \geq P_A(x) - e^{\varepsilon} P_B(x)
\end{aligned}$$

By definition of $\delta_T^{\text{❁}}(x, \varepsilon)$, we get

$$\delta_T^{\text{❁}}(x, \varepsilon) = P_A(x) - e^{\varepsilon}\frac{P_A(x)}{f^{\iota_T(x)}} \geq P_A(x) - e^{\varepsilon} P_B(x)$$

Moreover, as $\iota_T(x) >= j_\varepsilon$, we know that $e^\varepsilon \leq f^{\iota_T(x)}$. Hence, we also get

$$\delta_T^{\text{❁}}(x,\varepsilon) = \underbrace{P_A(x)}_{\geq 0} \cdot \left(1 - \frac{e^\varepsilon}{\underbrace{f^{\iota_T(x)}}_{\leq 1}}\right) \geq 0$$

**Case 2.** This case occurs if $\iota_T(x) \geq j_\varepsilon + u_T$, burt $\iota_T(x) \neq \infty$.

We first show that $\delta_T^{\text{❁}}(x,\varepsilon) \geq 0$. By Lemma 12 we know that $\ell_T^{\text{❁}}(x) \leq \frac{P_A(x)}{f^{\iota_T(x)-u_T}} - \frac{P_A(x)}{f^{\iota_T(x)}}$ holds; thus,

$$\begin{aligned}
&P_A(x) - e^\varepsilon \left(\frac{P_A(x)}{f^{\iota_T(x)}} + \ell_T^{\text{❁}}(x)\right)\\
\geq &P_A(x) - e^\varepsilon \left(\frac{P_A(x)}{f^{\iota_T(x)}} + \frac{P_A(x)}{f^{\iota_T(x)-u_T}} - \frac{P_A(x)}{f^{\iota_T(x)}}\right)\\
= &P_A(x) - e^\varepsilon \left(\frac{P_A(x)}{f^{\iota_T(x)-u_T}}\right)\\
\geq &P_A(x) - \frac{f^{j_\varepsilon}}{f^{\iota_T(x)-u_T}} P_A(x)\\
= &P_A(x) \cdot \left(1 - \frac{f^{j_\varepsilon}}{f^{\iota_T(x)-u_T}}\right)\\
\geq &0,
\end{aligned}$$

as by assumption $\iota_T(x) \geq j_\varepsilon + u_T$. We now show that $\delta_T^{\text{❁}}(x,\varepsilon) \geq P_A(x) - e^\varepsilon P_B(x)$.

Note that from Lemma 10 we know that $\ell_T^{\text{❁}}(x) \leq \tilde{\ell}_T^{\text{❁}}(x)$.

$$\frac{P_A(x)}{f^{\iota_T(x)}} + \underbrace{\ell_T^{\text{❁}}(x)}_{\leq \tilde{\ell}_T^{\text{❁}}(x)} \leq \frac{P_A(x)}{f^{\iota_T(x)}} + \tilde{\ell}_T^{\text{❁}}(x) \overset{Lemma\ 11}{=} \frac{P_A(x)}{f^{\iota_T(x)}} + P_B(x) - \frac{P_A(x)}{f^{\iota_T(x)}} = P_B(x)$$

Thus,

$$\delta_T^{\text{❁}}(x,\varepsilon) = P_A(x) - e^\varepsilon \left(\frac{P_A(x)}{f^{\iota_T(x)}} + \ell_T^{\text{❁}}(x)\right) \geq P_A(x) - e^\varepsilon P_B(x)$$

We combine these results and get

$$\begin{aligned}
\delta_T^{\text{❁}}(x,\varepsilon) =&P_A(x) - e^\varepsilon \left(\frac{P_A(x)}{f^{\iota_T(x)}} + \ell_T^{\text{❁}}(x)\right)\\
&\geq \max(0, P_A(x) - e^\varepsilon P_B(x))
\end{aligned}$$

**Case 3.** This case occurs if $\iota_T(x) = \infty$. By definition we have $\delta_T^{\text{❁}}(x,\varepsilon) = P_A(x) > \max(0, P_A(x) - e^\varepsilon P_B(x))$.

**Case 4.** This case occurs otherwise, i.e., if $\iota_T(x) < j_\varepsilon$. We calculate that

$$\begin{aligned}
&P_A(x) - e^\varepsilon P_B(x)\\
\overset{\text{since } e^\varepsilon \leq f^{j_\varepsilon}}{\leq} &P_A(x) - f^{j_\varepsilon} P_B(x)\\
\overset{\iota_T(x) < j_\varepsilon, P_B(x) \geq 0}{\leq} &P_A(x) - f^{\iota_T(x)} P_B(x) \overset{\text{Lemma 6}}{\leq} 0
\end{aligned}$$

and thus,

$$\delta_T^{\text{❁}}(x,\varepsilon) = 0 = \max(0, P_A(x) - e^\varepsilon P_B(x))$$

□

## 4.5 Main result

We now present our main technical theorem: Given any value for $\varepsilon \geq 0$ and a value $\delta(\varepsilon)$, s.t. the distributions are tightly $(\varepsilon, \delta(\varepsilon))$-differentially private, the value $\delta_T$ calculated as in Definition 9 presents a sound upper bound on $\delta(\varepsilon)$ from Lemma 1 and $\delta^{\text{low}}$ presents a lower bound on $\delta(\varepsilon)$.

**Definition 10** (Composition trees over distributions)**.** *Let $X$ and $Y$ be two distributions over the same universe $\mathcal{U}$. We call two composition trees $T_1$ and $T_2$* a pair of composition trees over the distributions $X$ and $Y$ *iff $A_{T_1} = B_{T_2} = X$ and $B_{T_1} = A_{T_2} = Y$.*

**Theorem 2** (Buckets with EC terms are sound)**.** *Let $A$ and $B$ be two distributions and let $T_1$ and $T_2$ be a pair of composition trees over $A$ and $B$ as in Definition 10. Then for every $\varepsilon \geq 0$ and with $\delta^{\text{up}}(\varepsilon) = \max\left(\delta_{T_1}(\varepsilon), \delta_{T_2}(\varepsilon)\right)$ and $\delta^{\text{low}}(\varepsilon) = \min\left(\delta^{\text{low}}_{T_1}(\varepsilon), \delta^{\text{low}}_{T_2}(\varepsilon)\right)$, the distributions $A$ and $B$ are $(\varepsilon, \delta^{\text{up}}(\varepsilon))$-ADP and*

$$\delta^{\text{low}}(\varepsilon) \leq \delta(\varepsilon) \leq \delta^{\text{up}}(\varepsilon),$$

*where $\delta(\varepsilon)$ is the tight $\delta$ as defined in Lemma 1.*

*Proof.* Lemma 1 shows that $A$ and $B$ are tightly $(\varepsilon, \delta(\varepsilon))$-differentially private for

$$\delta(\varepsilon) = \max\left(\sum_{x \in \mathcal{U}} \max\left(P_A(x) - e^{\varepsilon} P_B(x), 0\right), \sum_{x \in \mathcal{U}} \max\left(P_B(x) - e^{\varepsilon} P_A(x), 0\right)\right)$$

and Lemma 13 proves that $\delta(\varepsilon) \leq \delta_{T_{A||B}}(\varepsilon)$ holds true (for any composition tree $T$ and thus in particular for $T_{A||B}$).

Next, we show that $\delta^{\text{low}} \leq \delta(\varepsilon)$. We show the computation for $\delta^{\text{low}}_{A||B}$, the computation for $\delta^{\text{low}}_{B||A}$ follows analogously:

$$\begin{aligned}
\delta^{\text{low}}_{A||B} &= \sum_{i \in \{j_\varepsilon, \ldots, n\}} \max\left(0, \mathcal{B}_{T_{A||B}}(i) - e^{\varepsilon}\left(\frac{\mathcal{B}_{T_{A||B}}(i)}{f^i} + \tilde{\ell}_{T_{A||B}}(i)\right)\right) \\
&\overset{\text{Lemma 8}}{=} \sum_{i \in \{j_\varepsilon, \ldots, n\}} \max\left(0, \sum_{x \in \mathcal{U}, \iota_T(x) = i} P_A(x) - e^{\varepsilon}\left(\frac{P_A(x)}{f^i} + \tilde{\ell}^{\circledast}_{T_{A||B}}(x)\right)\right) \\
&\overset{\text{Lemma 11}}{=} \sum_{i \in \{-n, \ldots, n, \infty\}} \max\left(0, \sum_{x \in \mathcal{U}, \iota_T(x) = i} P_A(x) - e^{\varepsilon}\left(\frac{P_A(x)}{f^i} + \left(P_B(x) - \frac{P_A(x)}{f^{\iota_T(x)}}\right)\right)\right) \\
&= \sum_{i \in \{j_\varepsilon, \ldots, n\}} \max\left(0, \sum_{x \in \mathcal{U}, \iota_T(x) = i} P_A(x) - e^{\varepsilon} P_B(x)\right) \\
&\leq \sum_{i \in \{j_\varepsilon, \ldots, n\}} \sum_{x \in \mathcal{U}, \iota_T(x) = i} \max\left(0, P_A(x) - e^{\varepsilon} P_B(x)\right) \\
&\leq \sum_{x \in \mathcal{U}} \max\left(0, P_A(x) - e^{\varepsilon} P_B(x)\right)
\end{aligned}$$

Hence, we conclude that

$$\delta^{\text{low}}_{A||B} \leq \sum_{x \in \mathcal{U}} \max\left(0, P_A(x) - e^{\varepsilon} P_B(x)\right) \leq \delta(\varepsilon).$$

the computation for $\delta^{\text{low}}_{B||A}$ follows analogously, ending with $\delta^{\text{low}}_{B||A} \leq \sum_{x \in \mathcal{U}} \max\left(0, P_B(x) - e^{\varepsilon} P_A(x)\right) \leq \delta(\varepsilon)$. □

Thus, the bounds calculated present a sound over-approximation of the real differential privacy values. As discussed in Section 2, distributions can be used to calculate differential privacy in a variety of applications. To this end, we require the existence of *worst-case* inputs that are independent of the random coins used by the mechanism in the previous rounds.

**Definition 11** (Worst-case inputs)**.** *Inputs $x_0, x_1$ are worst-case inputs for a given sensitivity $s$ and a mechanism $M$ if $\Pr[M(x_0) \in S] \leq e^{\varepsilon} \Pr[M(x_1) \in S] + \delta$, implies $M$ is $(\varepsilon, \delta)$-ADP for all inputs with sensitivity $s$.*

These worst-case inputs commonly exist when differential privacy is applied (see Section 2.1) and they enable us to directly derive the relevant distributions.

As a corollary to Theorem 2, we see that we can compute upper and lower bounds for a sequence of privacy-enhancing mechanisms, where each of the mechanisms may be different from the others.

**Corollary 1.** *Let $M_1, \ldots, M_r$ be privacy-enhancing mechanisms for which there exist worst-case inputs $(x_{0,1}, x_{1,1}), \ldots (x_{0,r}, x_{1,r})$. Let $M_i(b)$ be the output distribution of the mechanism $M_i$ on input $x_{b,i}$. Let $T_1$ and $T_2$ be a pair of composition trees over $\prod_{i=1}^{r} M_i(0)$ and $\prod_{i=1}^{r} M_i(1)$ as in Definition 10. Then for every $\varepsilon \geq 0$ and with $\delta^{\text{up}}(\varepsilon) = \max\left(\delta_{T_1}(\varepsilon), \delta_{T_2}(\varepsilon)\right)$ and $\delta^{\text{low}}(\varepsilon) = \min\left(\delta_{T_1}^{\text{low}}(\varepsilon), \delta_{T_2}^{\text{low}}(\varepsilon)\right)$, the distributions $A$ and $B$ are $(\varepsilon, \delta^{\text{up}}(\varepsilon))$-ADP and*

$$\delta^{\text{low}}(\varepsilon) \leq \delta(\varepsilon) \leq \delta^{\text{up}}(\varepsilon),$$

*where $\delta(\varepsilon)$ is the tight $\delta$ as defined in Lemma 1.*

*Proof.* Consider the reduction that replaces all inputs of the attacker with sensitivity $s$ with the worst case inputs $((x_{0,1}, x_{1,1}), \ldots, (x_{0,r}, x_{1,r}))$ for sensitivity $s$. Theorem 2 implies the result for the product distributions $\prod_{i=1}^{r} M_i(x_{0,i})$ and $\prod_{i=1}^{r} M_i(x_{1,i})$. By the definition of worst-case inputs, we know that the result holds for any other sequence of inputs $((x'_{0,1}, x'_{1,1}), \ldots, (x'_{0,r}, x'_{1,r}))$. □

**Heterogenous adaptive $r$-fold composition** Bounds for (heterogenous) adaptive $r$-fold composition classically only restrict mechanisms to the class of all $(\varepsilon, \delta)$-ADP mechanisms. Thus, by only choosing worst-case mechanisms $M_{\varepsilon,\delta}$ (see Section 2) for each step, we get a for heterogeneous adaptive $r$-fold composition as in [14] .

When the class of mechanisms is restricted further, e.g., the structure of the mechanisms is partially known, we suggest to derive (and prove sound) tighter worst-case mechanisms for which we can then give significantly better results.

## 4.6 Implementation

We implemented $\delta_T(\varepsilon)$ (c.f. Theorem 2) in Python using the NumPy [3] and the SciPy [4] libraries in 655 LoC. The most time consuming part in the computation is the composition. We phrased the composition as a series of inner products and use the NumPy library, which has an efficient implementation of inner products. However, a series of inner products can be massively parallelized. While we added a simple form of parallelization (62 LoC), we expect that a massive parallelization via GPUs should improve the overall efficiency by several orders of magnitudes.

On an early 2015 MacBook Pro 13-inch with a 2.9 GHz Intel Core i5 (2 cores, with hyperthreading 4) and 16 GB of RAM, one composition operation took for $100,000$ buckets around 115 seconds, i.e., a little bit less than 2 minutes, with our unoptimized prototypical implementation. Our simple parallelization was configured to use 4 processes. Hence, with repeated squaring, computing $2^{18} = 262,144$ compositions took around 35 minutes. This benchmarking was done using the `example_gauss.py` script from the provided implementation (with $100,000$ instead of $10,000$ buckets and 4 instead of 1 parallel threads). All computations in Sections 5, 7 and 8 use $100,000$ buckets.

Given a bucket factor as well as a number of buckets $2n + 2$, our implementation constructs privacy buckets from any given histogram / distribution with a limed number of events. For Laplacian noise and Gaussian noise we have implemented special constructors that create privacy buckets for those functions in a more-or less precise fashion.

Given any privacy buckets and a number of rounds $r$, our implementation then calculates both upper bounds (with error correction) and lower bounds using repeated squaring: we compose the bucket distribution with itself in each round, thus calculating $2^r$ compositions in $r$ composition steps, with each composition step being quadratic in the number of buckets $n$.[7] Our implementation adaptively decides whether or not to perform "squaring", i.e., to rebase the factor depending on whether the bucket with index $\infty$ would otherwise grow too much. Empirically, we found that an increase of weight of the $\infty$ bucket by more than a factor of 2.2 is a good indicator that squaring should be performed. Additionally, we include a parameter that disables squaring as long as the $\mathcal{B}(\infty)$ is below this parameter, which is important for cases where $\mathcal{B}(\infty)$ is initially zero or very small. Finally, we compute an $(\varepsilon, \delta)$-graph by calculating $\delta$ as in Definition 9 for every $\varepsilon = f^i$ with $i \in \{0, \ldots, n\}$.

We refer to `https://ratiobuckets.rocks` for our implementation of privacy buckets.

# 5 Comparison to Kairouz et al.'s composition theorem

Kairouz et al. proved a composition theorem [14] that significantly improves on the standard and advanced composition theorem. This composition theorem [14] provides a composition result where each $\varepsilon, \delta$ pair after $r$-fold composition is solely derived from one $\varepsilon, \delta$ pair of the original pair of distributions. Hence, this composition result does take the entire shape of the distribution into account. In other words, the resulting epsilon and delta bounds are not necessarily tight in the sense of Definition 1.

Recall that we show that our privacy buckets approach provides an upper and a lower bound and that the distance between these two bounds can be made arbitrarily small by increasing the granularity of the buckets. The privacy buckets can be seen as an approximation of the two $\varepsilon, \delta$ graphs[8] of the original pair of distributions $A$ and $B$. As a consequence, our results show that the two $\varepsilon, \delta$ graphs of $A$ and $B$ capture all features that are relevant for computing the two $\varepsilon, \delta$ graphs after $k$-fold composition (i.e., of $A^k$ and $B^k$).

We show in this section that Kairouz et al.'s composition theorem seems to be tight for the Laplace mechanisms but not for all mechanisms, such as the Gauss mechanism or the measured timing-leakage of the CoverUp system [22]. While our approach does not provide significantly tighter bounds for Laplace mechanism, our privacy buckets significantly improve the privacy bounds on other mechanisms, such as Gauss mechanism and CoverUp-data. We first describe how we compute these mechanisms and then how we compute the composition theorem. Subsequently, we compare the tightness of the bounds from our privacy buckets approach to the bounds from Kairouz et al.'s composition theorem in these three scenarios. In the three case studies of this section we consider one-dimensional data, e.g., in responses to statistical queries over sensitive databases or leakage due to suspicious timing delays. However, our approach and our implementation can also deal with higher-dimensional data.

## 5.1 Embedding the Laplace mechanism

We analyze the Laplace mechanism, the classical mechanism to achieve DP, by comparing two distributions of Laplace noise with means 0 and 1 respectively. This case corresponds to many applications of the Laplace mechanism for DP, such as counting queries for databases with sensitivity 1. We choose in our case study a Laplace distribution with mean $\mu = 0$ and scale factor $\gamma = 200$, denoted as $\mathrm{LP}(\mu, \gamma)$. As a result, an attacker either makes observations from $\mathrm{LP}(0, 200)$ or from $\mathrm{LP}(1, 200)$ (as the sensitivity is 1). We consider truncated Laplace distributions, since that corresponds closer to real-world applications. If not mentioned otherwise, we truncate at $\mu - 2500$ and $\mu + 2500$.

We want to give strong evidence that both Kairouz et al.'s composition theorem and our privacy buckets are tight for the bounds of the Laplace mechanism. As a consequence, we carefully embed the Laplace mechanism in a way that has a small discretization error. The bucket method introduced in Definition 7 iterates over all atomic events in the support of the distributions. For modeling the Laplace distribution, or rather, two Laplace distributions $A$ and $B$, we consider the quotients of the probability mass functions and integrate distribution $A$ over the range of events that fall into each bucket: for $\mathcal{B}(i)$ we integrate over

[7] Using the Fast Fourier Transformation (FFT) each composition step can be reduced to $O(n \log n)$ by representing it as a convolution.

[8] There are two $\varepsilon, \delta$ graphs since the DP definition is asymmetric.

all events $x$ such that $f^i < p_A(x)/p_B(x) \leq f^{i+1}$. This technique can also be applied to other distributions with an infinitely large support, where all areas where $B$ has a probability of zero naturally contribute to the bucket $\mathcal{B}_\infty$.

Recall the probability density function for the Laplace distribution with mean $\mu$ and scale parameter $\gamma$ as $\text{Laplace}(x) := \frac{1}{2\gamma} e^{\frac{-|x-\mu|}{\gamma}}$. For differential privacy we often compare two such distributions with the same scale parameter $\gamma$ and different medians $\mu_1$ and $\mu_2$, where the means are the real values to which we add Laplace noise with scale parameter $\gamma$. We know that without composition, we get $(\varepsilon, 0)$-ADP with $\varepsilon = \frac{1}{\gamma}$. Consequently, we can describe the quotient $f$ at each point $x$ as We calculate the quotient $f(x) = \frac{\text{Laplace}_{\mu_1}(x)}{\text{Laplace}_{\mu_2}(x)}$ depending on the relation between the values for $x, \mu_1$ and $\mu_2$:

- $x \leq \min(\mu_1, \mu_2)$: $f(x) = e^{-(\mu_1 - x)\varepsilon} / e^{-(\mu_2 - x)\varepsilon} = e^{(-\mu_1 + x - x + \mu_2)\varepsilon} = e^{(\mu_2 - \mu_1)\varepsilon}$
- $\mu_1 \geq x \geq \mu_2$: $f(x) = e^{-(\mu_1 - x)\varepsilon} / e^{-(x - \mu_2)\varepsilon} = e^{(-\mu_1 + x + x - \mu_2)\varepsilon} = e^{(-\mu_1 - \mu_2 + 2x)\varepsilon}$
- $\mu_1 \leq x \leq \mu_2$: $f(x) = e^{-(x - \mu_1)\varepsilon} / e^{-(\mu_2 - x)\varepsilon} = e^{(-x + \mu_1 + \mu_2 - x)\varepsilon} = e^{(\mu_1 + \mu_2 - 2x)\varepsilon}$
- $x \geq \max(\mu_1, \mu_2)$: $f(x) = e^{-(x - \mu_1)\varepsilon} / e^{-(x - \mu_2)\varepsilon} = e^{(\mu_1 - x + x - \mu_2)\varepsilon} = e^{(\mu_1 - \mu_2)\varepsilon}$

It turns out that for a pair of Laplace distributions the quotient in the region $\min(\mu_1, \mu_2) \leq x \leq \max(\mu_1, \mu_2)$ is either monotonically increasing or monotonically decreasing. For any $x$ smaller than $\min(\mu_1, \mu_2)$, the quotient is stable at $e^{-\varepsilon}$ and for any $x$ larger than $\max(\mu_1, \mu_2)$ the quotient is stable at $e^{\varepsilon}$. Recall that our buckets capture a *range of quotients*: bucket $i$ captures all x such that $f^i < p_A(E)/p_B(E) \leq f^{i+1}$. As a result, each bucket $i$ contains contiguous points and defines an interval on the $x - axis$. For each interval we define the *bucket borders*, i.e., for the bucket with index $i$, we call the value $x$ with $f(x) = f^{i-1}$ the *left bucket border* $\text{lbb}(i)$ and the value $x$ with $f(x) = f^i$ the *right bucket border* $\text{rbb}(i)$.

For $\mu_1 > \mu_2$, the right bucket border $\text{rbb}(i)$ is the $x$ such that

$$
\begin{aligned}
e^{(2x - \mu_1 - \mu_2)\varepsilon} &= f^i = e^{(i\varepsilon/\text{gr})} =: e^j \\
\Leftrightarrow (2x - \mu_1 - \mu_2)\varepsilon &= j \\
\Leftrightarrow (2x - \mu_1 - \mu_2) &= j/\varepsilon \\
\Leftrightarrow 2x &= \mu_1 + \mu_2 + j/\varepsilon \\
\Leftrightarrow x &= (\mu_1 + \mu_2 + j/\varepsilon)/2 \\
\Leftrightarrow x &= (\mu_1 + \mu_2 + \frac{(i\varepsilon/\text{gr})}{\varepsilon})/2 \\
\Leftrightarrow x &= (\mu_1 + \mu_2 + i/\text{gr})/2 \\
\implies \text{rbb}(i) &= 1/2(\mu_1 + \mu_2 + i/\text{gr}) \\
\implies \text{rbb}(i-1) &= 1/2(\mu_1 + \mu_2 + i/\text{gr} - 1/\text{gr}) \\
&= \text{rbb}(i) - 1/(2\text{gr}) \\
&= \text{lbb}(i)
\end{aligned}
$$

For $\mu_1 < \mu_2$, the right bucket border $\text{rbb}(i)$ is the $x$ such that

$$\begin{aligned}
e^{(-2x+\mu_1+\mu_2)\varepsilon} =& f^i = e^{(i\varepsilon/\text{gr})} =: e^j \\
\Leftrightarrow (-2x+\mu_1+\mu_2)\varepsilon =& j \\
\Leftrightarrow (-2x+\mu_1+\mu_2) =& j/\varepsilon \\
\Leftrightarrow 2x =& \mu_1+\mu_2-j/\varepsilon \\
\Leftrightarrow x =& (\mu_1+\mu_2-j/\varepsilon)/2 \\
\Leftrightarrow x =& (\mu_1+\mu_2-\frac{(i\varepsilon/\text{gr})}{\varepsilon})/2 \\
\Leftrightarrow x =& (\mu_1+\mu_2-i/\text{gr})/2 \\
\implies \text{rbb}(i) =& 1/2(\mu_1+\mu_2-i/\text{gr}) \\
\implies \text{rbb}(i-1) =& 1/2(\mu_1+\mu_2-i/\text{gr}+1/\text{gr}) \\
=& \text{rbb}(i)+1/(2\text{gr}) \\
=& \text{lbb}(i)
\end{aligned}$$

As a result, the bucket $i$ has the value $\int_{\text{lbb}(i)}^{\text{rbb}(i)} \text{Laplace}(\mu_1, 1/\epsilon)$.

We compute the error correction term as $\ell(i) := \int_{\text{lbb}(i)}^{\text{rbb}(i)} \left(B(x) - \frac{A(x)}{f^i}\right)$ and we can directly compute the virtual error from this term.

For the buckets with index $\pm i$ s.t. $f^i = e^\varepsilon$ we integrate over the respective remaining areas $\mathcal{B}(-i) = \int_{-\infty}^{\text{rbb}(-i)} \text{Laplace}(\mu_1, 1/\epsilon)$ and to $\mathcal{B}(i)$ we add $\int_{\text{rbb}(i)}^{\infty} \text{Laplace}(\mu_1, 1/\epsilon)$. As we chose $f$ to fit $e^\varepsilon$ the events in these regions exactly have the respective quotient of the bucket and we don't have errors for these integrals. Consequently, the error terms for bucket $\mathcal{B}(-i)$ are zero and the error terms for bucket $\mathcal{B}(i)$ are composed of the error terms for the values x with $\text{lbb}(i) < x < \text{rbb}(i)$.

**Truncated Laplace distributions.** The truncation of each of either of these functions, causes the quotient of a region to be either 0 or to have 0 in the denominator, which we treat as infinity. The regions are captured by the outer buckets with indexes $-n$ and $\infty$ respectively.

## 5.2 Embedding the Gauss mechanism

The truncated Gauss mechanism is also an often-used mechanism in privacy-preserving applications. It works similar to the Laplace mechanism insofar as it convolves the input (e.g., a query response) with a Gaussian distribution. In this work, we use a mean $\mu = 0$ and a standard deviation $\sigma = 200\sqrt{2}$ (to achieve the same variance as $\text{LP}(0, 200)$), denoted as $\text{GS}(\mu, \sigma^2)$, and we truncate these distributions at $\mu - 2500$ and $\mu + 2500$, if not mentioned otherwise. For the truncated Gauss mechanism, we do not use a precise embedding but rather produce a histogram for each of the two distributions, using SciPy's `scipy.stat.norm` function. Then, we use the normal interface of our bucketing implementation that parses a pair of histograms and produces a bucketlist vector, a real error vector, and a virtual error vector. We accept that this implementation may produce discretization artifacts that, however, should be both small w.r.t. the values concerned and should not lead to a significantly different shape of the distributions under composition.

## 5.3 Embedding CoverUp's data

Classical anonymous communication networks (ACN) have the goal of hiding the IP address of the sender and the recipient of a communication. Such ACNs do however not hide the participation time, i.e., whether, when, and for how long a party uses an ACN. This participation time can be used for long-term attacks (e.g., intersection attacks) and can raise suspicion national state-level adversaries. Sommer et al. [22] propose a system, called CoverUp, that has the goal of hiding this participation time leakage. CoverUp assumes a collaborating popular web service with a significant amount of regular visitors. This webpage would be incorporated into the usage of an ACN and trigger all its visitors to produce cover traffic. This web page

would serve an iFrame that loads content from a trusted server, which in turn would serve a piece of JavaScript code that executes a dummy client for the ACN on the visitors browser. ACN users would act as a normal visitor, receive the JS code, but additionally have a dedicated CoverUp browser extension installed. The browser extension would enable a communication channel to an external application by replacing the dummy messages from the dummy client with actual messages from an external application and by forwarding all messages from the network to the external application. For CoverUp to properly hide the participation time ACN users (called *voluntary* participants) and normal website visitors (called *involuntary* participants) have to be indistinguishable. While both execute the same piece of JS code, the voluntary participants perform additional computations. As a consequence, the response time of the voluntary participants differs by a few milliseconds from the response time of the involuntary participants. CoverUp remedies this timing leakage by adding random delays in the JS code, i.e., for voluntary and involuntary participants.

The CoverUp paper presents an analysis of this timing leakage (after adding the noise) and aims for a high degree of privacy after more than 250k observations. The CoverUp authors experimentally measured the timing delays of voluntary and involuntary participants in the lab and produced histograms of these timing delays. These histograms are used as a model for the timing delays of voluntary and involuntary participants to assess the timing leakage of CoverUp. We apply our algorithm to these histograms of timing delays, to illustrate that and how well our approach works on measured data. We use data from the CoverUp project, which is openly available online.[9]

In this comparison, we only consider those measured delays on a Linux system that are observable after the webpage has been loaded, called the "periodic" measurements in the CoverUp paper.

## 5.4 Computing Kairouz et al.'s composition theorem

We directly implement the bounds from Kairouz et al.'s theorem. We do not use any statements specific to Gauss or Laplace, as those are simplified and provide worse bounds.

**Theorem 3** ([14]). *For any $\varepsilon \geq 0$ and $\delta \in [0,1]$, the class of $(\varepsilon, \delta)$-ADP mechanisms satisfies $(\varepsilon', \delta')$-ADP under $r$-fold composition, for all $i \in \{0, \ldots, \lfloor r/2 \rfloor\}$ where $\varepsilon' = (r-2i)\varepsilon$ and $\delta' = 1 - (1-\delta)^r(1-\delta_i)$*

$$\delta_i = \frac{\sum_{\ell=0}^{i-1} \binom{r}{\ell} \left(e^{(r-\ell)\varepsilon} - e^{(r-2i+\ell)\varepsilon}\right)}{(1+e^{\varepsilon})^r}$$

We compute for a given number $r$ of compositions the epsilon-delta graph by looking up for a fine-grid of $\varepsilon$ values the corresponding $\delta$ value of the original pair of distributions and then computing and storing all $(\varepsilon', \delta')$ pairs according to the theorem above, i.e., for all $i \in \{0, \ldots, \lfloor r/2 \rfloor\}$. From these stored $(\varepsilon', \delta')$ pairs, we remove all pairs for which we have stored lower $(\varepsilon'', \delta'')$ pairs, i.e., pairs such that $\varepsilon'' \leq \varepsilon'$ and $\delta'' \leq \delta'$. We output the remaining list of $(\varepsilon', \delta')$ pairs, which form a monotonically decreasing $(\varepsilon, \delta)$-graph. Due to our direct implementation of $\delta_i$, we can only evaluate the composition theorem up to $r = 512$ before the intermediate computation results (in particular, the $e^{O(k)}$-terms) become too large.

In our computation, the granularity of the grid of $\varepsilon$ values of the original pair of distributions naturally leads to an imprecision. We use a fine grid of

$$e^{\varepsilon} \in \{(1 + 10^{-14})^{1.1^j} \mid j \in \{0, \ldots, n\}\},$$

where we choose $n$ as a point where the $(\varepsilon, \delta)$ after $r$-fold composition becomes stationary. While we concede that it might be possible to obtain a slightly lower bound from the composition theorem, we are confident that, due to this fine grid, the resulting graphs for Kairouz et al.'s composition theorem that we compute are representative.

## 5.5 Comparing evaluations

We are finally in a position to evaluate how our privacy buckets compare against Kairouz et al.'s composition theorem. Figure 11 and Figure 12 show our results. We see that our upper and lower bounds coincide, i.e., that our results are tight. Also, Kairouz et al.'s composition theorem is tight with respect to a pair of Laplace

[9] Available under `http://coverup.tech`.

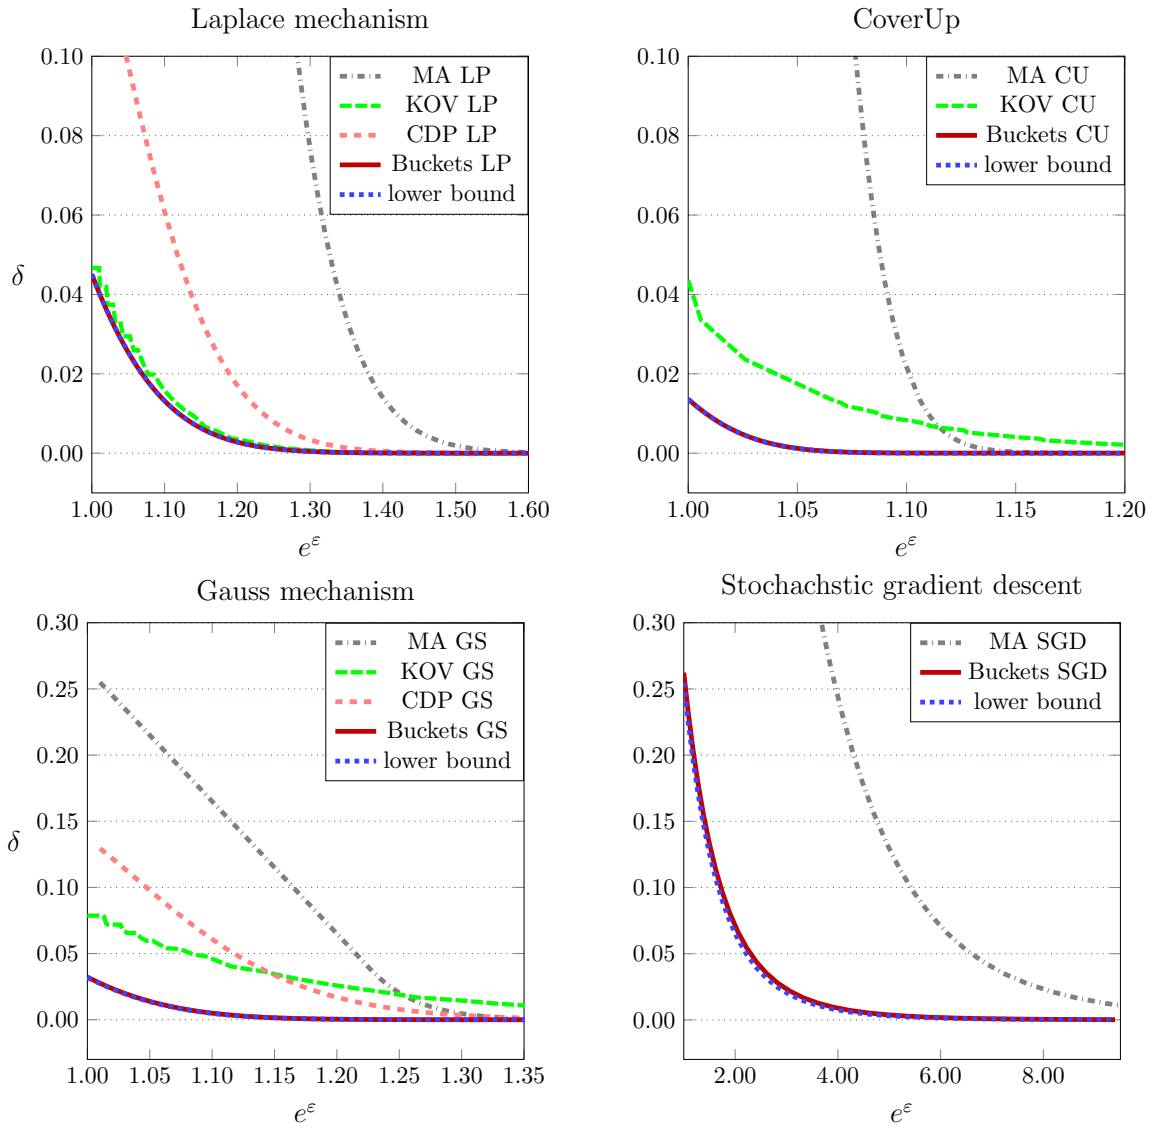


Figure 11: A group of $(\varepsilon, \delta)$-graphs for the Laplace mechanism, the Gauss mechanism, the CoverUp mechanism and the stochastic gradient descent mechanism. The first three are for $r = 512$, the latter for $r = 2^{16}$.

distributions (i.e., the Laplace mechanism). We see that for the Gauss mechanism that composition theorem is already after 512 compositions not very tight. For the CoverUp-data our privacy buckets are tight, while there is a large gap to the bounds from Kairouz et al's composition theorem. We used in the computation of all these graphs $100,000$ buckets.

Figure 13 compares for fixed epsilon values the evolution of the delta bounds from Kairouz et al.'s composition theorem and from our approach. This comparison again uses the Laplace mechanism, the Gauss mechanism and the CoverUp data.

The tightness of Kairouz et al.'s[14] bounds for the Laplace mechanism suggests that there is no noise distribution with the same, or smaller, initial $\varepsilon$ and $\delta$ values that has a worse composition behavior than the Laplace mechanism.

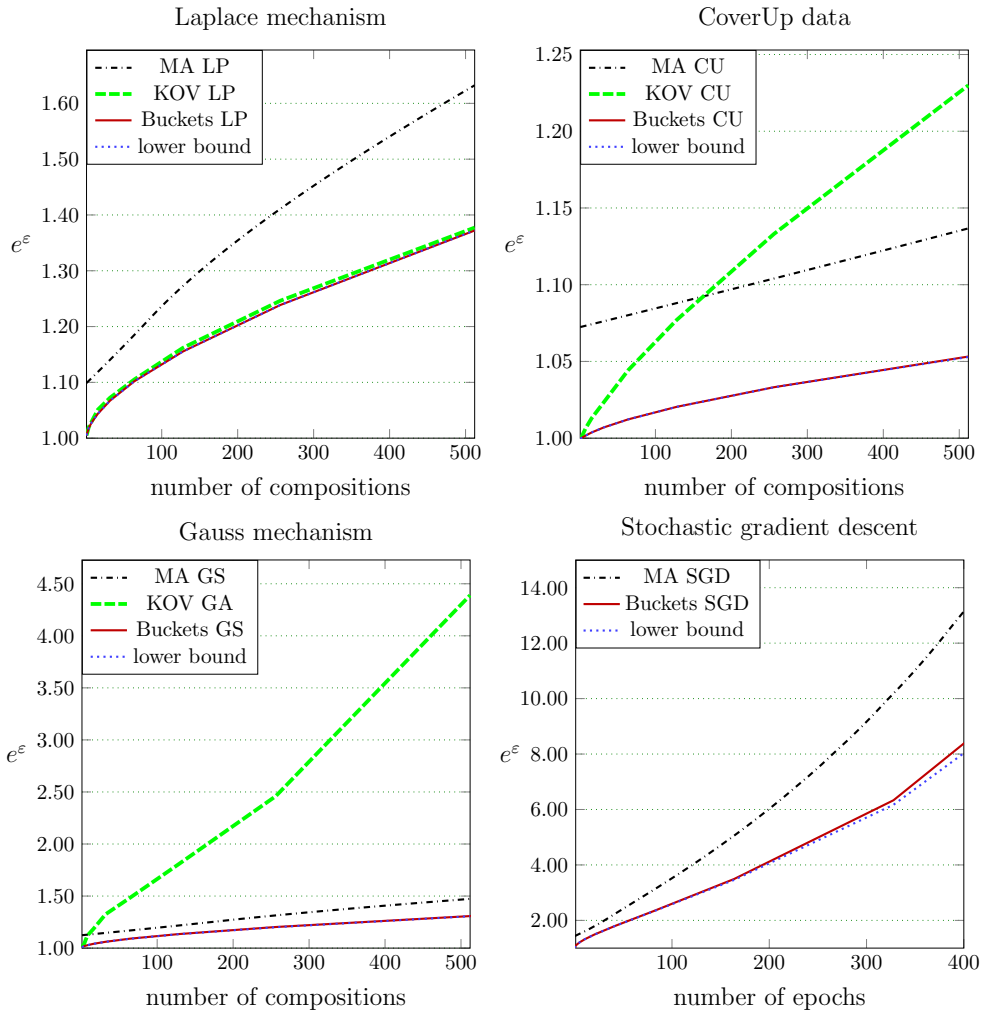


Figure 12: Comparison with the bounds from Rényi Privacy and CDP for the Gauss mechanism, i.e., the distributions $\mathrm{GS}(0, 2 * 200^2)$ and $\mathrm{GS}(1, 2 * 200^2)$. Left: $(\varepsilon, \delta)$-graph for 512 compositions. Right: growth of $e^\varepsilon$ over the number of compositions for $\delta \leq 10^{-5}$.

# 6 Comparison to bounds based on Rényi Divergence

General purpose bounds [9, 14] were solely based on one $(\varepsilon, \delta)$ pair from the pair of distributions – say $A, B$ – in question. In contrast, the recently introduced notions of Concentrated Differential Privacy (CDP) [8, 2], Rényi Differential Privacy (RDP) [18], and Moments Accountant [1] introduced mechanism-aware bounds for differential privacy by using the so-called Rényi Divergence, which can be expressed[10] as the higher log-moments of our privacy buckets distributions (for $\alpha > 0$):

$$D_{\alpha+1} = \frac{1}{\alpha} \log \sum_i \mathcal{B}(i) \cdot f^{\alpha i}$$

A pair of distributions $A, B$ satisfies $(\xi, \rho)$-Concentrated DP if the Rényi Divergence is bounded by an affine linear function: $D_\alpha \leq \xi + \rho\alpha$ (for all $\alpha \geq 0$). Rényi Differential Privacy directly characterizes the privacy by the Rényi Divergences: $(\alpha, D_\alpha)_\alpha$. Rényi Differential Privacy can be translated to $(\varepsilon, \delta)$-ADP as follows: whenever $(\alpha, D_\alpha)_\alpha$, then also $(\varepsilon, \alpha D_\alpha - \alpha\varepsilon)$-ADP holds. The Moments Accountant uses the same

[10] This theoretical characterization assumes a sufficiently high number of buckets such that no approximations need to be made.

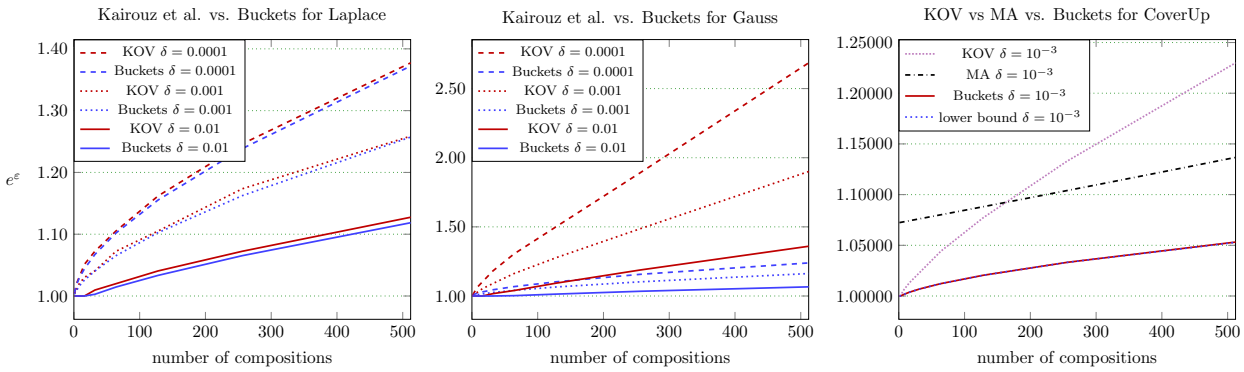


(a) Laplace (b) Gauss (c) CoverUp

Figure 13: Growth of $e^\varepsilon$ over the number of compositions (y-axis) for fixed $\delta$ values (different line-styles) for a growing number of compositions.

characterization and proposes $(\varepsilon, \min_\alpha(\alpha D_\alpha - \alpha\varepsilon))$ as ADP bounds. This section compares these mechanism-aware bounds with our privacy buckets bounds and illustrates that we achieve significantly lower bounds. We evaluate the Gauss mechanism, the randomized response, and a differentially privacy stochastic gradient descent mechanism (DP-SGD), which is useful for deep learning [1]. In these three evaluations, we illustrate two advantages of our approach: $(\varepsilon, \delta)$-graphs for a fixed number of compositions highlight that we achieve significantly reduced $\delta$-bounds for very small $e^\varepsilon$, plots about $e^\varepsilon$-bounds for a fixed bound on $\delta$ but a growing number of compositions illustrate that we achieve significantly reduced $e^\varepsilon$-bounds.

The Gauss mechanism uses the same variance and distance of the means as before: we compare $\mathrm{GS}(0, \sqrt{2}\cdot 200)$ with $\mathrm{GS}(1, \sqrt{2}\cdot 200)$. Our privacy buckets result in the lowest bounds and that these bounds are tight, as the lower bounds are very close to the upper bounds. While Rényi Privacy actually uses more information to characterize the leakage than CDP (as it maintains separate bounds for each moment), CDP [2] proves tighter ADP bounds, i.e., the translation from CDP to ADP is tighter. As the moments of the Gauss Mechanism and the randomized response can be bounded by an affine function, CDP yields tighter bound due to the tighter translation. We apply an optimization to MA and CDP: instead of computing the respective $\delta$ bounds for a given $\varepsilon$ (written as $\delta(\varepsilon)$) from the RDP and CDP papers, the graph plots the minimum of $\delta(\varepsilon)$ and $\min_\varepsilon(e^{\varepsilon'} - e^\varepsilon) + \delta(\varepsilon)$, which is a generic bound on $\delta$ for a given $\varepsilon$ that can be derived from and $(\varepsilon, \delta)$-ADP graph.

The randomized response mechanism $\mathrm{RR}_{p,f}$ is parametric in a bias $p$ and is defined for a binary predicate $f : X \to \{0, 1\}$. For an input database $D$, if $f(D)$ the mechanism $\mathrm{RR}_{p,f}(D)$ outputs 1 with probability $p$ and 0 with probability $1 - p$. If $f(D)$ is not true, $\mathrm{RR}_{p,f}(D)$ outputs 0 with probability $p$ and 1 with probability $1 - p$. It can be easily shown that the pair $(A, B)$ of distributions $A(0) = p, A(1) = 1 - p$ and $B(0) = 1 - p$ and $B(1) = p$ is a worst case distribution, which corresponds to the case where two databases $D_0, D_1$ are compared with $f(D_0) = 0$ and $f(D_1) = 1$.[11] In particular, since Differential Privacy considers worst-case inputs, we do not need to explicitly specify the predicate $f$ as long as there are at least two databases $D_0$ and $D_1$ for which $f(D_0) \neq f(D_1)$, i.e., as long as $f$ is not constantly true or false on all databases. Figure 14 shows that privacy buckets clearly lead to tighter ADP-bounds, followed by the CDP and then the RDP bounds. As for the Gauss mechanism, CDP yields tighter bounds than RDP since the translation to ADP is tighter.

Finally, we evaluate the DP-SGD mechanism and illustrate that this method leads to significantly tighter bounds. Abadi et al. prove that it suffices to consider the leakage of a Gaussian mixture model $(1 - q)\mathrm{GS}(0, \sigma^2) + q\mathrm{GS}(1, \sigma^2)$ versus a Gaussian distribution $\mathrm{GS}(0, \sigma^2)$, for some $q$ that depends on the training parameters. We construct privacy buckets for these two distributions and compute the composition with 100,000 buckets. To improve comparability, we use the same parameters as in the Moments Accountant

[11] For arbitrary $D'_0, D'_1$, the reduction checks whether $f(D'_0) = f(D'_1)$. If so, the reduction draws a $p$-biased coin and outputs the result. If $f(D'_0) = 1$ and $f(D'_1) = 0$, the reduction flips the output of the game with the worst-case distributions, and if $f(D'_0) = 1$ and $f(D'_1) = 0$ the reduction simply forwards the result of the game with the worst-case distributions.

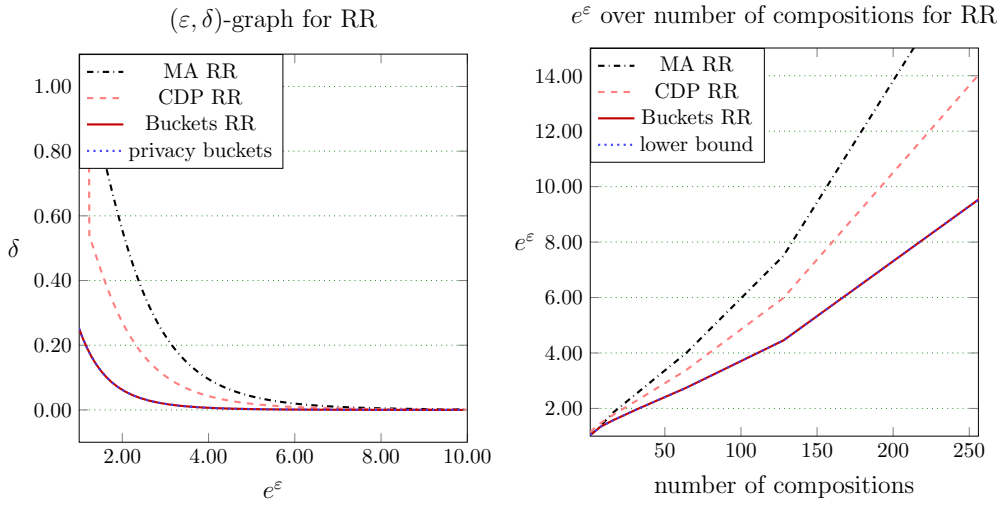


Figure 14: Comparison of bounds for randomized response (RR) mechanism with $p = 0.51$. Left: $(\varepsilon, \delta)$-graph for 512 compositions. Right: growth of $e^{\varepsilon}$ over the number of compositions for $\delta \leq 10^{-4}$.

paper, $\sigma = 4, q = 0.01, \delta = 10^{-5}$; however, we plot on the y-axis $e^{\varepsilon}$ and not $\varepsilon$. While the Moments Accountant-paper discusses the application to Deep Learning, we solely concentrate on the improvement of the privacy bound and refer to the applicability to Deep Learning to the Moments Accountant paper [1]. We see that also in this case for which the Moments Accountant/RDP is very well suited, the privacy buckets lead to significantly tighter bounds.[12] We stress that CDP is not applicable in this application, since $D_{\alpha}$ cannot be bounded by a affine linear function .

# 7 Comparison of the Gaussian and the Laplace mechanism

As we have seen in Section 5.5, Kairouz et al.'s composition theorem is fairly tight for the Laplace mechanism but not for the Gauss mechanism. Figure 15 (upper two graphs) compares a truncated Laplace and a truncated Gauss mechanism and find that for the same variance the Gauss mechanism provides a significantly higher degree of privacy.[13] For a fixed variance of $80,000$, a sensitivity of 1 ($mu_1 = 0$ and $\mu_2 = 2$), and a truncation at $-2500$ and 2500 for $\mu_1$ (and $-2499$ and 2501 for $\mu_2$), the upper left graph in Figure 15 depicts how, for different but fixed epsilon values, the delta increases over the course of 512 evaluations. The graph clearly shows that in the course of 512 compositions, the reduced leakage of the Gauss mechanism becomes visible. The lower left graph in Figure 15 shows the full epsilon-delta graphs of a Gaussian and a Laplace mechanism after 512 compositions, where the two mechanisms use noise that has the same variance $(80,000)$. In particular, the delta-value where the $(\varepsilon, \delta)$ graph levels out is 4 orders of magnitude lower for Gaussian noise than it is for Laplace noise, since the Gaussian distribution falls much steeper than Laplace distribution. This difference of the Gaussian and the Laplace mechanisms becomes even more pronounced in our analysis and improvement of the Vuvuzela protocol in Section 8. The analysis of Vuvuzela also illustrates that the steepness of the Gaussian distribution enables a much tighter truncation, i.e., the distribution can be truncated much earlier than a Laplace distribution without sacrificing privacy. This tighter truncation, in turn, leads to a smaller range of noise that is required to achieve the same privacy goals as with Laplace noise.

Additionally, we found evidence that the epsilon-delta graph of the Laplace mechanism converges toward the epsilon-delta graph of a Gauss mechanism with half the variance of the Laplace mechanism. For the same sensitivity, and truncations as above, the two right graphs in Figure 15 illustrate that after 512 compositions these two graphs converge toward each other. The upper right graph in Figure 15 depicts how, for different but fixed epsilon values, the delta increases over the course of 512 evaluations. The graph clearly shows

[12] We used the publicly available implementation of the Moments Accountant for our computations.
[13] All computations have been conducted with $100,000$ buckets.

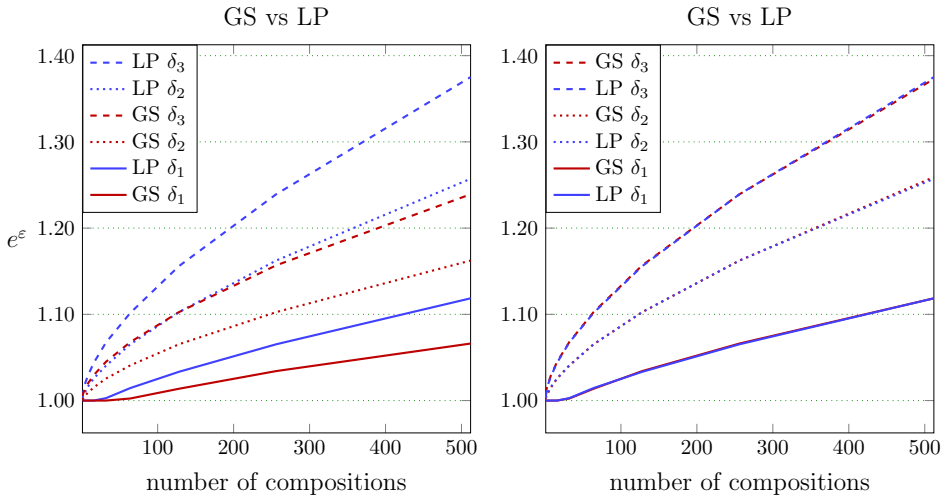


(a) Growth of $e^\varepsilon$ over the number of compositions (y-axis) for fixed $\delta$ values: $\delta_1 = 0.01, \delta_2 = 0.001, \delta_3 = 0.0001$ (different line-styles) for a growing number of compositions. The legend is in the same order as the graphs.

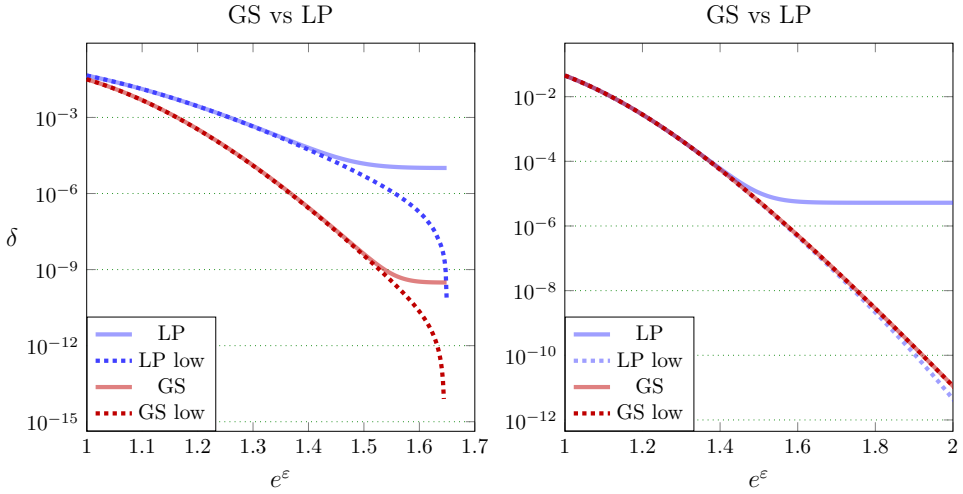


(b) The $\varepsilon, \delta$ graphs (upper and lower bounds) after $k = 512$ compositions applied to a Gaussian and a Laplace mechanism with $\delta$ on the y-axis and $e^\varepsilon$ on the x-axis.

Figure 15: Truncated Gauss mechanisms (red) vs. truncated Laplace mechanism (blue) both with sensitivity = 1. For both mechanism truncation is at $\mu_i - 2500$ and $\mu_i + 2500$ ($\mu_1 = 0$ and $\mu_2 = 1$). At twice the variance the Laplace mechanism converges towards the Gauss mechanism, so much that the blue lines almost completely cover the red lines.

how in the course of 512 compositions, the delta values of the Laplace mechanism converge toward the delta values of the Gauss mechanism. The lower right graph in Figure 15 shows the full epsilon-delta graphs of a Gaussian and a Laplace mechanism after 512 compositions, where the Laplace mechanism has twice the variance $(80,000)$ of the Gauss mechanism $(40,000)$. This figure shows how close the two epsilon-delta graphs are and that they almost only differ due to their different y-values at the point where they have been truncated. This difference, however, is crucial. As explained above, it is caused by the steepness of the Gaussian distribution and enables a much tighter truncation, which in turn can lead to significantly less noise overhead, as we illustrate in our analysis of Vuvuzela.

Dwork and Rothblum [8] presented a related result. They characterized the composition behavior of a mechanisms (i.e., a pair of distributions $A, B$) for which the privacy loss distribution $e\mathcal{L}_{(A||B)}$ is Subgaussian, i.e., the moment-generating function (MGF) is smaller than the MGF of a Gaussian. They show that such for

mechanisms the privacy loss distribution after $r$-fold composition can be bounded by a Gaussian. Moreover, showed that the privacy loss distribution, i.e., the privacy buckets distribution, of the Gauss mechanism is a Gaussian distribution. Complementarily, our findings show that the privacy loss distribution of a Laplace mechanism converges towards a Gaussian, already after 512 compositions. We leave it for future work to investigate the connection between the Laplace distribution and a Gaussian distribution with half the variance.

# 8 Application to Vuvuzela

In this section, we show how aiming for tight bounds in a privacy analysis can significantly improve the bandwidth overhead of a protocol. As a case study, we use the Vuvuzela [24] protocol, which is an anonymous communication system tailored towards messengers. Vuvuzela uses Laplace noise to achieve strong privacy properties. Using the insights from Section 7, we not only estimate tighter bounds for the Laplace noise but also propose to change the shape of the noise distribution to Gaussian noise. With our bucketing approach, we show that already 5 to 10 times less noise[14] suffices to achieve the same strong privacy properties. [15]

We refer to the original Vuvuzela paper for a full presentation and restrict our presentation to the bare bones that are needed to understand the noise messages that Vuvuzela uses to achieve strong privacy properties.

We stress that our work contributes to improving the epsilon-delta bounds and thus to improve a given privacy analysis. This work is not meant to help in finding a suitable attacker model, a suitable definition or accurate usage profiles. Hence, we stick to Vuvuzela's privacy analysis, as it was presented in the original paper.

## 8.1 Protocol overview

Vuvuzela clients communicate by deposing their encrypted messages in virtual locations in the one of the mixes (the locations are called *dead drops*). For agreeing on such a dead drops, Vuvuzela deploys a dialing protocol where the dialer sends the ID of a dead drop to dedicated invitation dead drops. This ID is encrypted with the peer's public key with an encryption schemes that is designed to hide the recipient's identity. On the dialer's side directly the conversation protocol is started where the client regularly retrieves the chat messages from and deposits chat messages to the dead drop from the invitation. If the recipient receives and accepts the invitation, the recipient also starts the conversation protocol.

**Privacy analysis** Vuvuzela assumes a global network-level attacker that is additionally able to compromise some mixes. To achieve strong resistance against compromised servers, each path in Vuvuzela traverses all nodes. To counter traffic correlation attacks, Vuvuzela clients produce dummy traffic at a constant rate. The Vuvuzela paper argues that the only remaining source of leakage is the patterns of registering invitations and patterns of access requests to these dead drops: single requests to dead drops, corresponding to dummy messages or messages before the peer accepted the conversation, and pairs of requests to the same dead drop, corresponding to an active conversation.

**Privacy-enhancing measures** Vuvuzela reduces the information that an attacker can learn by triggering each mix to produce cover stories for potentially communicating parties. For the dialing protocol, the mixes produce cover stories ($i$) by sending dummy invitation registrations and invitation requests to the dedicated invitation dead drops. The number of these dummy registrations and dummy requests is in each round drawn from the truncated Laplace distribution $\lceil \max(0, \text{Laplace}(\gamma_d, \mu_d)) \rceil$ for some system parameters $\gamma_d$ and $\mu_d$. For the conversation protocol, the mixes produce cover stories ($ii$) for idle parties, by sending pairs of dummy access requests to uniform-randomly chosen dead drops, and ($iii$) for (bi-directionally) communicating parties, by sending (single) dummy access requests to uniform-randomly chosen dead drops. The

[14] The more observations are estimated, the higher the error of the advanced composition result, which is used in the original analysis from the Vuvuzela paper; hence, in those cases the tightness of our bounds leads to a more significant improvement.

[15] We acknowledge that for the analysis of the Laplace noise previous results [14] would already yield tight results, but for the Gaussian noise our approach yields much tighter results (see Section 7).

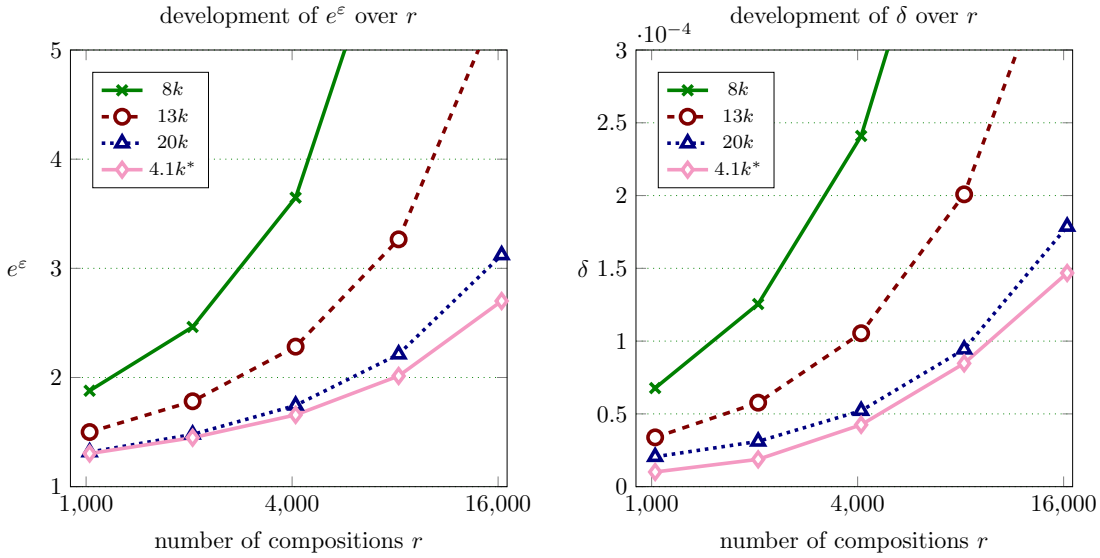


Figure 16: The privacy bounds for Vuvuzela's dialing protocol. The left graph shows the $e^{\varepsilon}$-values on the y-axis and the number of observations $r$ on the x-axis (i.e., $r$-fold composition) in log-scale and the right graph shows the corresponding $\delta$-values on the y-axis. The solid green ($\mu = 8k, \gamma = 500$), the dashed red ($\mu = 13k, \gamma = 770k$), and the dotted blue line ($\mu = 20k, \gamma = 1130$) are from the original Vuvuzela paper, and the solid magenta line (Gaussian noise, $\mu = 4.1k^*, \sigma = 320$) is computed with this work's technique.

number of (single) dummy access requests $(ii)$ is in each round drawn from the truncated Laplace distribution $\lceil \max(0, \text{Laplace}(\gamma_c, \mu_c)) \rceil$ for system parameters $\gamma_c$ and $\mu_c$, and the number of pairs of dummy access requests $(iii)$ is in each round drawn from the truncated Laplace distribution $\lceil \max(0, \text{Laplace}(\mu_c/2, \gamma_c/2)) \rceil$. The system parameters $\mu_d, \mu_c, \gamma_d, \gamma_c$ determine how much noise-overhead the protocol produces and how much privacy it will offer.

**Privacy-impact of the dummy requests** The goal of the these dummy requests and invitations is to produce a cover stories for dialing parties $(i)$, for idle parties $(ii)$, and for conversing $(iii)$. The Vuvuzela paper separately conducts a privacy analysis for the dialing protocol ($(i)$) and the conversation protocol ($(ii)$ and $(iii)$ combined). For the dialing protocol, the paper concludes that it suffices to bound the $r$-fold $(\epsilon, \delta)$ differential privacy of $\max(0, \text{Laplace}(\mu_d, \gamma_d))$ and $\max(0, \text{Laplace}(\mu_d + 2, \gamma_d))$, i.e., the $(\epsilon, \delta)$ differential privacy of the product distributions $\max(0, \text{Laplace}(\mu_d, \gamma_d))^r$ and $\max(0, \text{Laplace}(\mu_d + 2, \gamma_d))^r$. The parameter $r$ indicates the number of rounds at which that the attacker conducts an observation. For the conversation protocol, the paper concludes that it suffices to estimate the $r$-fold $(\epsilon, \delta)$ differential privacy of $\max(0, \text{Laplace}(\mu_c, \gamma_c)) + \max(0, \text{Laplace}(\mu_c/2, \gamma_c/2))$ and $\max(0, \text{Laplace}(\mu_c + 2, \gamma_c)) + \max(0, \text{Laplace}(\mu_c/2 + 1, \gamma_c/2))$. The Vuvuzela paper uses the advanced composition theorem for differential privacy [9] to bound $\epsilon$ and $\delta$. The paper analyzes for the conversation protocol three system parameters: $\mu = 150k, \gamma = 7.5k$, $\mu = 300k, \gamma = 13.8k$, and $\mu = 450k, \gamma = 20k$. We show that the resulting bounds can be significantly improved and we indicate all new bounds with a "$*$" sign in the respective figures.

We apply our method to estimate tighter $\varepsilon$ and $\delta$ bounds for Vuvuzela, and to reduce the recommended noise. Recall that we observed in Section 7 that Gaussian noise for the same variance behaves better under composition than Laplacian noise. This section studies how much our tighter bounds enable us to reduces the noise in the case that Gaussian noise is used or that Laplace noise is used, and this section studies how much the originally recommended amount of noise improves the degree of privacy, in case Gaussian noise is used or Laplace noise is used. We stress that while in the case of Vuvuzela there is no utility function that we have to preserve other than to minimize the bandwidth overhead, our approach is also suited for applications where a utility function has to be preserved. In those cases, we would probably reduce the variance to an appropriate level and then compute tight bounds.

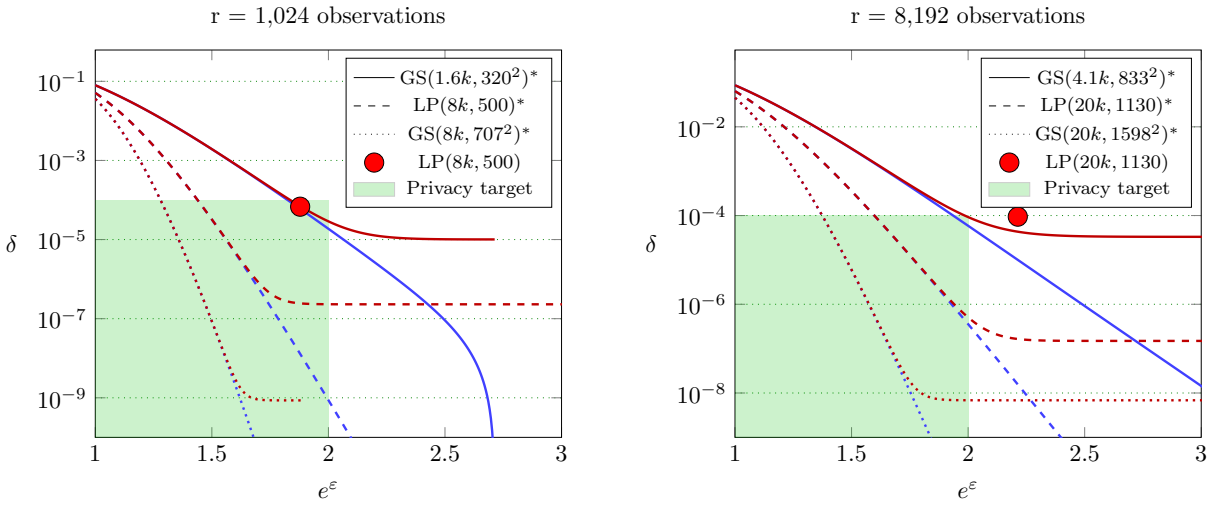


(a) After $r = 1,024$ observations with Gaussian noise with $\mu = 1.6k$ and $\sigma = 320$ (solid), Laplace noise $\mu = 8k, \gamma = 500$ (dashed), and Gaussian noise with $\mu = 8k$ and $\sigma = 707$ (dotted), and the red dot represents the $\varepsilon, \delta$ combination for $\mu = 8k, \gamma = 500$ from the original Vuvuzela paper.

(b) After $r = 8,192$ observations with Gaussian noise with $\mu = 4.1k$ and $\sigma = 833$ (solid), Laplace noise $\mu = 20k, \gamma = 1130$ (dashed), and Gaussian noise with $\mu = 20k$ and $\sigma = 1598$ (dotted), and the red dot represents the $\varepsilon, \delta$ combination for $\mu = 20k, \gamma = 1130$ from the original Vuvuzela paper.

Figure 17: The $(\varepsilon, \delta)$ graphs (y-axis and x axis, respectively, y-axis in $\log_{10}$-scale) from our method in comparison with the bound from the original Vuvuzela paper (for the dialing protocol). The figure depicts upper (red) and a lower bounds (blue) and Vuvuzela's privacy target (green, $\delta \leq 10^{-4}, e^{\varepsilon} \leq 2$).

## 8.2 Tighter privacy analysis for the dialing protocol

For the dialing protocol, we show that with Gaussian noise the noise rate can be reduced by a factor of almost 5 while still meeting the privacy requirements, and for the conversation protocol the noise rate can be reduced by a factor of 10 while still meeting the privacy requirements. With Laplace noise the noise rate can be reduced by a factor of 2 and for the conversation protocol by a factor of 4. We refer to Figures 21 and 22, placed in the appendix. As the conversation protocol produces more observations (i.e., more compositions) and the looseness of the bounds that the original Vuvuzela paper used amplifies more heavily for a high the number of observations, the tightness of our bounds is more pronounced for the conversation protocol.

For comparability, we depict in Figure 16 the original graphs from the Vuvuzela analysis, which show the epsilon graph and the delta graph with increasing $r$, respectively, for the dialing protocol and estimated with the advanced composition result. We extend those Figures with the lowest, magenta graphs (marked with a $*$) that show the performance of our proposed Gaussian noise that uses nearly 5 times less noise and is computed with our bucketing approach.[16] As our method computes not only one $\varepsilon, \delta$ pair for each number of observations $r$ but an entire $\varepsilon, \delta$ graph, we chose representative $\epsilon$ values that are close to (and even below) the epsilon values for the highest noise configuration $\text{LP}(20k, 1130)$ from the original Vuvuzela paper. The figure shows that our bounds with the reduced noise and with using Gaussian noise $\text{GS}(4.1k, 833^2)$ are below the previous bounds for the highest noise configuration $\text{LP}(20k, 1130)$, proving that a noise reduction of nearly a factor of 5 still yields for the dialing protocol to achieve the privacy requirements of $e^{\varepsilon} \leq 2$ and $\delta \leq 10^{-4}$.

Next, we illustrate that our method computes bounds that are several orders of magnitude better than Vuvuzela's original bounds. For $r = 8,192$ observations, Figure 17b illustrates that using the highest noise configuration with Laplace noise $\text{LP}(20k, 1130)$ results in a privacy bound that is almost 3 orders of magnitude lower, in terms of the delta, and with Gaussian noise $\text{GS}(20k, 1598^2)$ more than 4 orders of magnitude. The figure depicts the $\varepsilon, \delta$ graphs computed by our approach for the highest noise configuration $\text{LP}(20k, 1130)$, for the corresponding Gaussian noise $\text{GS}(20k, 1598^2)$, for the configuration that we propose

[16] All computations have been conducted with 100,000 buckets.

| name | $\mu$ | $\gamma$ for LP | $\sigma^2$ (used for GS) |
|---|---|---|---|
| new$_1$ | $15k$ | — | $2.5k^2$ |
| new$_2$ | $45k$ | — | $7.5k^2$ |
| low | $150k$ | $7.3k$ | $2 \cdot 7.3k^2$ |
| medium | $300k$ | $13.8k$ | — |
| high | $450k$ | $20k$ | $2 \cdot 20k^2$ |

Figure 18: Parameters of our Vuvuzela analysis.

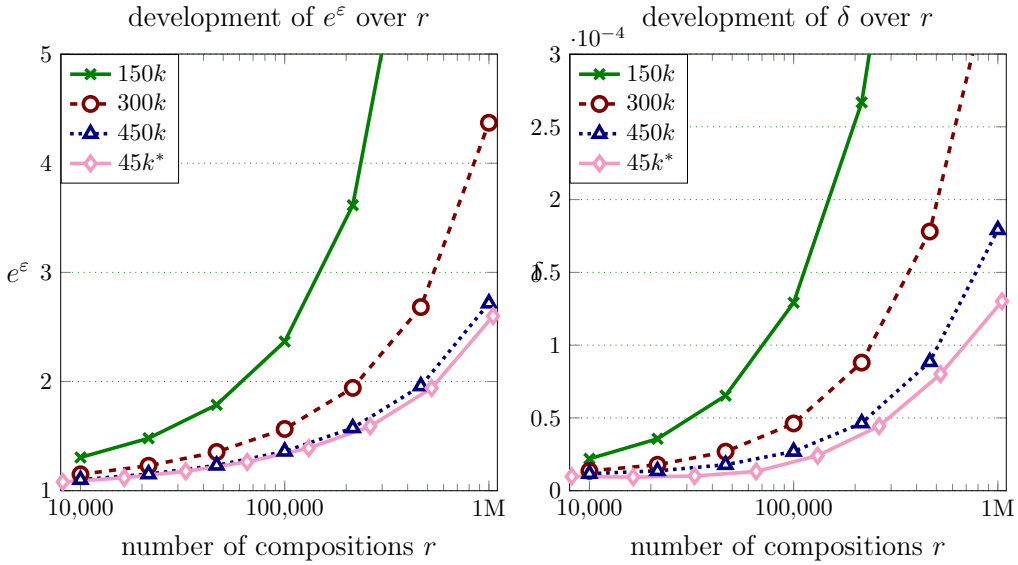


Figure 19: Vuvuzela conversion protocol: bounds on $\varepsilon$ and $\delta$ over $r$ (log-scale). We compare the original bounds for the originally recommended mechanisms with $150k$, $300k$, $450k$, analyzed with previous bounds [9] and our recommended mechanism with $45k$ messages overhead per round, analyzed using privacy buckets.

$\mathrm{GS}(4.1k, 833^2)$), and compares it against Vuvuzela's previous bounds $\mathrm{LP}(20k, 1130)$. We additionally depict the respective lower bounds, which show that our bounds are quite tight in the sense that there is not much room for improvement. Moreover, due to the more comprehensive view that a full $\varepsilon, \delta$ graph provides, we can see that the the highest noise configuration with Gaussian noise $\mathrm{GS}(20k, 1598^2)$ even achieves the privacy requirements ($\delta \leq 10^{-4}$) for less than $e^\varepsilon = 1.5$ after $8,192$ observations.[17]

We would like to stress that the lower bounds show that our result is tight up to $\delta \geq 10^{-4}$ for $\mathrm{GS}(4.1k, 833^2)$, $\delta \geq 10^{-6}$ for $\mathrm{LP}(20k, 1130)$, and $\mathrm{GS}(20k, 1598^2)$ for $\delta \geq 10^{-8}$. This tightness is solely a scalability issue and ultimately only depends on the number (and hence granularity) of the buckets. A more optimized implementation (e.g., based on GPUs) would be able to significantly increase the number of buckets, thus achieving even tighter upper and lower bounds.

For completeness, we also show in Figure 17a the $\varepsilon, \delta$ graphs for the dialing protocol for low $r$: $r = 1024$ and the recommended parameters $\mu = 8k, \gamma = 500$. Here, we can see that our bound is 2 orders of magnitude lower than Vuvuzela's previous bounds for the noise level. The figure also shows that reducing the noise by a factor of 5, i.e., $\mathrm{GS}(1.6k, 320)$, still achieves the privacy requirements ($e^\varepsilon \leq 2$ and $\delta \leq 10^{-4}$).

As a comparison, using Laplace noise only enables a noise reduction of a factor of 2, as shown in Figure 22 in the appendix. Interestingly, the reduced Laplace noise achieves the same privacy bounds as the reduced Gaussian noise if the Laplace noise has twice the variance as the Gaussian noise (i.e., $\gamma = \sigma$) but a 2.5 times wider range, as indicated in Section 7. This shows what a significant effect the steepness of the Gaussian noise can have in practice.

[17] Recall that the variance of $\mathrm{GS}(\mu, (\sqrt{2}x)^2) = 2x^2$ equals the variance of $\mathrm{LP}(\mu, x) = 2x^2$.

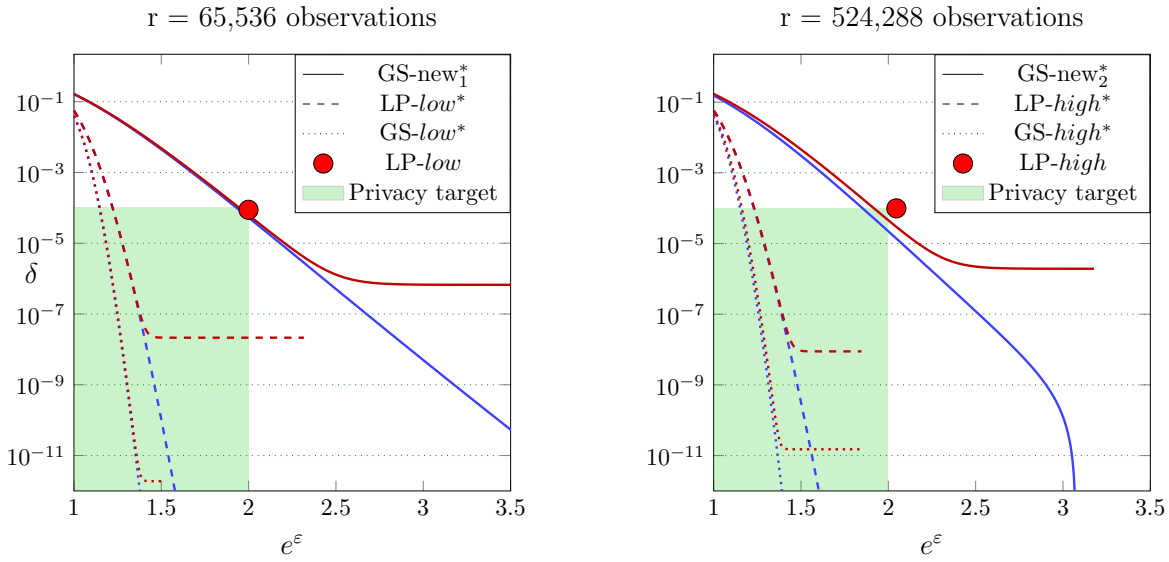


(a) After $r = 65,536$ observations with Gaussian noise with $\mu = 15k$ and $\sigma = 2.5k$ (solid), Laplace noise $\mu = 150k, \gamma = 7.3k$ (dashed), and Gaussian noise with $\mu = 150k$ and $\sigma = 10.3k$ (dotted), and the red dot represents the $\varepsilon, \delta$ combination for $\mu = 150k, \gamma = 7.3k$ from the original Vuvuzela paper.

(b) After $r = 524,288$ observations with Gaussian noise with $\mu = 45k$ and $\sigma = 7.5k$ (solid), Laplace noise $\mu = 450k, \gamma = 20k$ (dashed), and Gaussian noise with $\mu = 450k$ and $\sigma = 28.2k$ (dotted), and the red dot represents the $\varepsilon, \delta$ combination for $\mu = 450k, \gamma = 20k$ from the original Vuvuzela paper.

Figure 20: The $(\varepsilon, \delta)$ graphs (y-axis and x axis, respectively, y-axis in $\log_{10}$-scale) from our method in comparison with the bound from the original Vuvuzela paper (for the conversation protocol). The figure depicts upper (red) and a lower bounds (blue) and Vuvuzela's privacy target (green area, $\delta \leq 10^{-4}, e^{\varepsilon} \leq 2$).

### 8.3 Tighter privacy analysis for the conversation protocol

Figure 19 depicts the epsilon graph and the delta graph with increasing $r$, respectively, for the conversation protocol. We compare Gaussian noise GS-new$_2$ with the previous bounds for the recommended noise configurations. We see that although GS-new$_2$ adds significantly less noise, the bounds outperform the ones from the original analysis.

For $r = 524,288$ observations, Figure 20b shows that using LP-high results in bounds for $\delta$ that are almost 4 orders of magnitude lower, and for the corresponding Gaussian noise GS-high more than 6 orders of magnitude in comparison to their original result. Also, Figure 20b shows the corresponding lower bounds. We can see that our bounds are tight for reasonably small values of $\varepsilon$. Furthermore, we can see that even GS-new$_1$ meets the privacy requirements of $e^{\varepsilon} \leq 2$ and $\delta \leq 10^{-4}$ for $r = 524,288$ observations.

For completeness, we also show in Figure 20a the $\varepsilon, \delta$ graphs for the conversation protocol for $r = 65,536$. Here, we can also see the tightness of our bound for reasonably small $\varepsilon$. We can see that GS-low is more than 7 orders of magnitude lower than Vuvuzela's previous bounds for the same noise level. Moreover, we can see that even GS-new$_2$ meets the privacy requirements of $e^{\varepsilon} \leq 2$ and $\delta \leq 10^{-4}$ for $r = 65,536$ observations.

As a comparison, using Laplace noise only enables a noise reduction of a factor of 4, as shown in Figure 21 in the appendix. Also here, we can observe that the Laplace noise has twice the variance of the Gaussian noise and has a 2.5 times wider range, illustrating the advantages of Gaussian noise in practice.

## 9 Conclusion and future work

In this paper we have presented *privacy buckets*, a sound numerical approach for computing upper and lower bounds for differential privacy after r-fold composition. Our approach is based on concrete distributions, but can be applied in a variety of cases, which can include adaptive composition, evolving sequences of

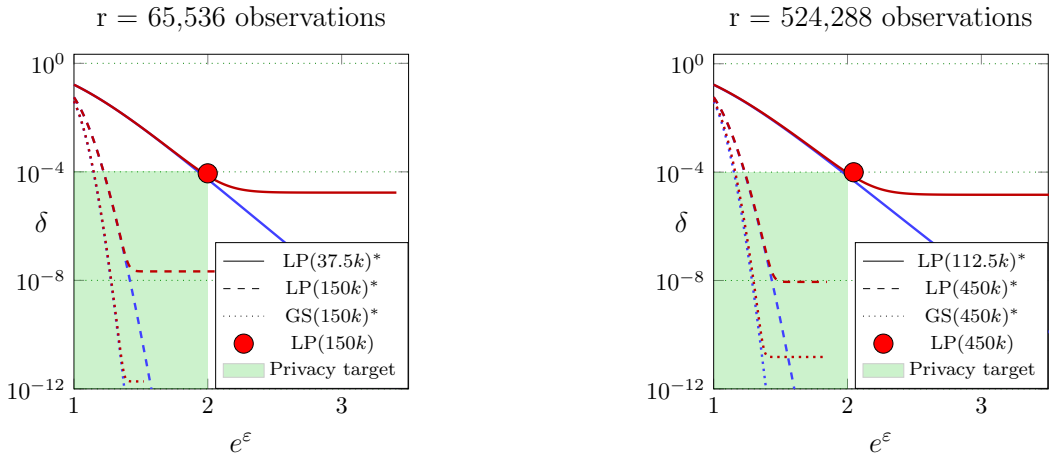


(a) After $r = 65,536$ observations with Laplace noise with $\mu = 37.5k$ and $\sigma = 2.3k$ (solid), Laplace noise $\mu = 150k, \gamma = 7.3k$ (dashed), and Gaussian noise with $\mu = 150k$ and $\sigma = 10.3k$ (dotted), and the red dot represents the $\varepsilon, \delta$ combination for $\mu = 150k, \gamma = 7.3k$ from the original Vuvuzela paper.

(b) After $r = 524,288$ observations with Laplace noise with $\mu = 112.5k$ and $\sigma = 6.9k$ (solid), Laplace noise $\mu = 450k, \gamma = 20k$ (dashed), and Gaussian noise with $\mu = 450k$ and $\sigma = 28.2k$ (dotted), and the red dot represents the $\varepsilon, \delta$ combination for $\mu = 450k, \gamma = 20k$ from the original Vuvuzela paper.

Figure 21: The $(\varepsilon, \delta)$ graphs (y-axis and x axis, respectively, y-axis in $\log_{10}$-scale) from our method in comparison with the bound from the original Vuvuzela paper (for the conversation protocol). The figure depicts upper (red) and a lower bounds (blue) and Vuvuzela's privacy target (green area, $\delta \leq 10^{-4}, e^\varepsilon \leq 2$).

distributions and static distributions. All compositions, as well as our reshaping operation of *squaring* the bucket factor have been shown sound and (empirically) tight in many cases.

We compared our approach to the Kairouz, Oh and Viswanath's (KOV) composition theorem, as well as to the Moment's Accountant (MA) bounds and bounds derived via Concentrated Differential Privacy (CDP). We found that the KOV theorem provides reasonably tight bounds for the Laplace mechanism but not for other distributions, such as the Gauss mechanism or for a pair of histograms of timing-leakage measurements from the CoverUp system. Our bounds significantly improve over MA bounds and CDP bounds, which is particularly relevant for smaller values of $e^\varepsilon$. We also observed that Gauss mechanism behaves much better under a high number of compositions than a Laplace mechanism with the same variance, and we found evidence that the $(\varepsilon, \delta)$-graph of a Laplace mechanism converges to the $(\varepsilon, \delta)$-graph of a Gauss mechanism with half the variance. By repeating the analysis of the anonymity network Vuvuzela we show that tighter bounds can have a significant impact on actual protocols. Moreover our analysis can help devise better protocols, e.g., to exchange the Laplace noise with Gaussian noise for which even better results can be achieved.

We encourage the application of our privacy buckets to other ADP mechanisms, such as to the optimal ADP mechanisms [11, 15] (e.g., comparing their composition behavior to the Gauss mechanism) and to measure the impact of our bounds on precision and recall of privacy-preserving ML methods [1], as well as to improve more existing privacy analyses. We consider exploring the relationship between ADP of the Gauss mechanism and ADP of the Laplace mechanism, as well as analyses probing the development of ADP provided by other noise distributions under composition interesting future work.

## 10 Acknowledgement

This work has been partially supported by the European Commission through H2020-DS-2014-653497 PANORAMIX, the EPSRC Grant EP/M013-286/1, and the Zurich Information Security Center (ZISC).

# References


[1] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang. Deep Learning with Differential Privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 308–318. ACM, 2016.

[2] M. Bun and T. Steinke. Concentrated Differential Privacy: Simplifications, Extensions, and Lower Bounds. In *Theory of Cryptography (TCC)*, pages 635–658. Springer, 2016.

[3] N. developers. Numpy.org: Scientific Computing with Python. Accessed in August 2017, available at `http://www.numpy.org`.

[4] S. developers. SciPy.org: Scientific Computing Tools for Python. Accessed in August 2017, available at `https://www.scipy.org`.

[5] C. Dwork. Differential Privacy. In *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part II*, pages 1–12. Springer, 2006.

[6] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor. Our Data, Ourselves: Privacy Via Distributed Noise Generation. In *Advances in Cryptology - EUROCRYPT 2006*, pages 486–503. Springer, 2006.

[7] C. Dwork, M. Naor, T. Pitassi, and G. N. Rothblum. Differential Privacy Under Continual Observation. In *Proceedings of the 42th Annual ACM Symposium on Theory of Computing (STOC)*, pages 715–724. ACM, 2010.

[8] C. Dwork and G. N. Rothblum. Concentrated Differential Privacy. *CoRR*, abs/1603.01887, 2016.

[9] C. Dwork, G. N. Rothblum, and S. Vadhan. Boosting and differential privacy. In *2010 51st Annual IEEE Symposium on Foundations of Computer Science*, pages 51–60. IEEE, 2010.

[10] C. Dwork, K. Talwar, A. Thakurta, and L. Zhang. Analyze Gauss: Optimal Bounds for Privacy-preserving Principal Component Analysis. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*, pages 11–20. ACM, 2014.

[11] Q. Geng and P. Viswanath. The optimal mechanism in differential privacy. In *2014 IEEE International Symposium on Information Theory (ISIT)*, pages 2371–2375. IEEE, 2014.

[12] M. Hardt and G. N. Rothblum. A Multiplicative Weights Mechanism for Privacy-Preserving Data Analysis. In *2010 51st Annual IEEE Symposium on Foundations of Computer Science*, pages 61–70. Springer, 2010.

[13] T.-H. Hubert Chan, E. Shi, and D. Song. Private and Continual Release of Statistics. In *Automata, Languages and Programming. ICALP 2010*, pages 405–417. Springer, 2010.

[14] P. Kairouz, S. Oh, and P. Viswanath. The composition theorem for differential privacy. *IEEE Transactions on Information Theory*, 63(6):4037–4049, 2017.

[15] K. Kalantari, L. Sankar, and A. D. Sarwate. Optimal differential privacy mechanisms under Hamming distortion for structured source classes. In *2016 IEEE International Symposium on Information Theory (ISIT)*, pages 2069–2073. IEEE, 2016.

[16] C. Liu, S. Chakraborty, and P. Mittal. Dependence Makes You Vulnberable: Differential Privacy Under Dependent Tuples. In *NDSS*, 2016.

[17] S. Meiser. Approximate and probabilistic differential privacy definitions, 2018.

[18] I. Mironov. Renyi Differential Privacy. In *Proceedings of the 30th IEEE Computer Security Foundations Symposium (CSF)*, pages 263–275. IEEE, 2017.

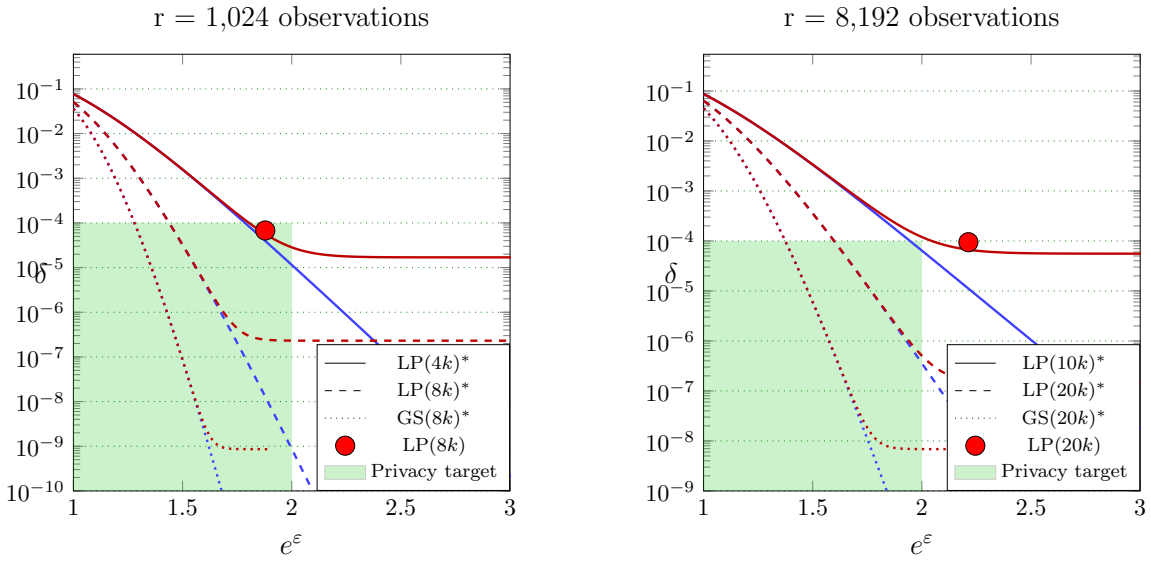


(a) After $r = 1,024$ observations with Laplace noise with $\mu = 4k$ and $\sigma = 330$ (solid), Laplace noise $\mu = 8k, \gamma = 500$ (dashed), and Gaussian noise with $\mu = 8k$ and $\sigma = 707$ (dotted), and the red dot represents the $\varepsilon, \delta$ combination for $\mu = 8k, \gamma = 500$ from the original Vuvuzela paper.

(b) After $r = 8,192$ observations with Laplace noise with $\mu = 10k$ and $\sigma = 827$ (solid), Laplace noise $\mu = 20k, \gamma = 1130$ (dashed), and Gaussian noise with $\mu = 20k$ and $\sigma = 1598$ (dotted), and the red dot represents the $\varepsilon, \delta$ combination for $\mu = 20k, \gamma = 1130$ from the original Vuvuzela paper.

Figure 22: The $(\varepsilon, \delta)$ graphs (y-axis and x axis, respectively, y-axis in $\log_{10}$-scale) from our method in comparison with the bound from the original Vuvuzela paper (for the dialing protocol). The figure depicts upper (red) and a lower bounds (blue) and Vuvuzela's privacy target (green, $\delta \leq 10^{-4}, e^{\varepsilon} \leq 2$).


[19] J. Murtagh and S. Vadhan. The Complexity of Computing the Optimal Composition of Differential Privacy. In *Proceedings of the 13th International Conference on Theory of Cryptography (TCC)*, pages 157–175. Springer, 2016.

[20] J. Murtagh and S. P. Vadhan. The Complexity of Computing the Optimal Composition of Differential Privacy. *CoRR*, abs/1507.03113, 2015.

[21] R. M. Rogers, A. Roth, J. Ullman, and S. Vadhan. Privacy odometers and filters: Pay-as-you-go composition. In *Advances in Neural Information Processing Systems 29*, pages 1921–1929. Curran Associates, Inc., 2016.

[22] D. Sommer, A. Dhar, L. Malitsa, E. Mohammadi, D. Ronzani, and S. Capkun. Anonymous Communication for Messengers via "Forced" Participation. Technical report, available under `https://eprint.iacr.org/2017/191`, 2017.

[23] J. Tang, A. Korolova, X. Bai, X. Wang, and X. Wang. Privacy Loss in Apple's Implementation of Differential Privacy on MacOS 10.12. *ArXiv e-prints*, 2017.

[24] J. van den Hooff, D. Lazar, M. Zaharia, and N. Zeldovich. Vuvuzela: Scalable Private Messaging Resistant to Traffic Analysis. In *Proceedings of the 25th ACM Symposium on Operating Systems Principles (SOSP 2015)*, pages 137–152. ACM, 2015.