Concentrated Differential Privacy: Simplifications, Extensions, and Lower Bounds*

Mark Bun[1]** and Thomas Steinke[2]***

[1] mbun@seas.harvard.edu
John A. Paulson School of Engineering and Applied Sciences, Harvard University
[2] Thomas.Steinke@ibm.com
IBM, Almaden Research Center

Abstract. "Concentrated differential privacy" was recently introduced by Dwork and Rothblum as a relaxation of differential privacy, which permits sharper analyses of many privacy-preserving computations. We present an alternative formulation of the concept of concentrated differential privacy in terms of the Rényi divergence between the distributions obtained by running an algorithm on neighboring inputs. With this reformulation in hand, we prove sharper quantitative results, establish lower bounds, and raise a few new questions. We also unify this approach with approximate differential privacy by giving an appropriate definition of "approximate concentrated differential privacy."

1 Introduction

Differential privacy [DMNS06] is a formal mathematical standard for protecting individual-level privacy in statistical data analysis. In its simplest form, (pure) differential privacy is parameterized by a real number $\varepsilon > 0$, which controls how much "privacy loss"[3] an individual can suffer when a computation (i.e., a statistical data analysis task) is performed involving his or her data.

One particular hallmark of differential privacy is that it degrades smoothly and predictably under the *composition* of multiple computations. In particular, if one performs k computational tasks that are each ε-differentially private and combines the results of those tasks, then the computation as a whole is $k\varepsilon$-differentially private. This property makes differential privacy amenable to the type of modular reasoning used in the design and analysis of algorithms: When a

* The full version of this work appears at https://arxiv.org/abs/1605.02065

** Supported by an NDSEG Fellowship and NSF grant CNS-1237235. Part of this work was done while the author was visiting Yale University.

*** Part of this work was done while the author was a Harvard University, supported by NSF grants CCF-1116616, CCF-1420938, and CNS-1237235.

[3] The privacy loss is a random variable which quantifies how much information is revealed about an individual by a computation involving their data; it depends on the outcome of the computation, the way the computation was performed, and the information that the individual wants to hide. We discuss it informally in this introduction and define it precisely in Definition 2 on page 3.

sophisticated algorithm is comprised of a sequence of differentially private steps, one can establish that the algorithm as a whole remains differentially private.

A widely-used relaxation of pure differential privacy is *approximate* or (ε, δ)-differential privacy [DKM+06], which essentially guarantees that the probability that any individual suffers privacy loss exceeding ε is bounded by δ. For sufficiently small δ, approximate (ε, δ)-differential privacy provides a comparable standard of privacy protection as pure ε-differential privacy, while often permitting substantially more useful analyses to be performed.

Unfortunately, there are situations where, unlike pure differential privacy, approximate differential privacy is not a very elegant abstraction for mathematical analysis, particularly the analysis of composition. The "advanced composition theorem" of Dwork, Rothblum, and Vadhan [DRV10] (subsequently improved by [KOV15,MV16]) shows that the composition of k tasks that are each (ε, δ)-differentially private is $(\approx\sqrt{k}\varepsilon, \approx k\delta)$-differentially private. However, these bounds can be unwieldy; computing the tightest possible privacy guarantee for the composition of k arbitrary mechanisms with differing $(\varepsilon_i, \delta_i)$-differential privacy guarantees is #P-hard [MV16]! Moreover, these bounds are not tight even for simple privacy-preserving computations. For instance, consider the mechanism that approximately answers k statistical queries on a given database by adding independent Gaussian noise to each answer. Even for this basic computation, the advanced composition theorem does not yield a tight analysis.[4]

Dwork and Rothblum [DR16] recently put forth a different relaxation of differential privacy called *concentrated differential privacy*. Roughly, a randomized mechanism satisfies concentrated differentially privacy if the privacy loss has small mean and is subgaussian. Concentrated differential privacy behaves in a qualitatively similar way as approximate (ε, δ)-differential privacy under composition. However, it permits sharper analyses of basic computational tasks, including a tight analysis of the aforementioned Gaussian mechanism.

Using the work of Dwork and Rothblum [DR16] as a starting point, we introduce an alternative formulation of the concept of concentrated differential privacy that we call "zero-concentrated differential privacy" (zCDP for short). To distinguish our definition from that of Dwork and Rothblum, we refer to their definition as "mean-concentrated differential privacy" (mCDP for short). Our definition uses the Rényi divergence between probability distributions as a different method of capturing the requirement that the privacy loss random variable is subgaussian.

[4] In particular, consider answering k statistical queries on a dataset of n individuals by adding noise drawn from $\mathcal{N}(0, (\sigma/n)^2)$ independently for each query. Each individual query satisfies $(O(\sqrt{\log(1/\delta)}/\sigma), \delta)$-differential privacy for any $\delta > 0$. Applying the advanced composition theorem shows that the composition of all k queries satisfies $(O(\sqrt{k}\log(1/\delta)/\sigma), (k+1)\delta)$-differential privacy for any $\delta > 0$. However, it is well-known that this bound can be improved to $(O(\sqrt{k\log(1/\delta)}/\sigma), \delta)$-differential privacy.

1.1 Our Reformulation: Zero-Concentrated Differential Privacy

As is typical in the literature, we model a dataset as a multiset or tuple of n elements (or "rows") in $\mathcal{X}^n$, for some "data universe" $\mathcal{X}$, where each element represents one individual's information. A (privacy-preserving) computation is a randomized algorithm $M : \mathcal{X}^n \to \mathcal{Y}$, where $\mathcal{Y}$ represents the space of all possible outcomes of the computation.

Definition 1 (Zero-Concentrated Differential Privacy (zCDP)) *A randomised mechanism $M : \mathcal{X}^n \to \mathcal{Y}$ is (ξ, ρ)-zero-concentrated differentially private (henceforth (ξ, ρ)-zCDP) if, for all $x, x' \in \mathcal{X}^n$ differing on a single entry and all $\alpha \in (1, \infty)$,*

$$\mathrm{D}_\alpha \left(M(x) \| M(x') \right) \leq \xi + \rho\alpha, \tag{1}$$

where $\mathrm{D}_\alpha \left(M(x) \| M(x') \right)$ is the α-Rényi divergence[5] between the distribution of $M(x)$ and the distribution of $M(x')$.

We define ρ-zCDP to be $(0, \rho)$-zCDP.[6]

Equivalently, we can replace (1) with

$$\mathbb{E}\left[e^{(\alpha-1)Z} \right] \leq e^{(\alpha-1)(\xi+\rho\alpha)}, \tag{2}$$

where $Z = \mathsf{PrivLoss}\left(M(x) \| M(x') \right)$ is the privacy loss random variable:

Definition 2 (Privacy Loss Random Variable) *Let Y and Y' be random variables on Ω. We define the* privacy loss random variable between Y and Y' *– denoted $Z = \mathsf{PrivLoss}\left(Y \| Y' \right)$ – as follows. Define a function $f : \Omega \to \mathbb{R}$ by $f(y) = \log(\mathbb{P}\left[Y = y \right] / \mathbb{P}\left[Y' = y \right])$. Then Z is distributed according to $f(Y)$.*

Intuitively, the value of the privacy loss $Z = \mathsf{PrivLoss}\left(M(x) \| M(x') \right)$ represents how well we can distinguish x from x' given only the output $M(x)$ or $M(x')$. If $Z > 0$, then the observed output of M is more likely to have occurred if the input was x than if x' was the input. Moreover, the larger Z is, the bigger this likelihood ratio is. Likewise, $Z < 0$ indicates that the output is more likely if x' is the input. If $Z = 0$, both x and x' "explain" the output of M equally well.

A mechanism $M : \mathcal{X}^n \to \mathcal{Y}$ is ε-differentially private if and only if $\mathbb{P}\left[Z > \varepsilon \right] = 0$, where $Z = \mathsf{PrivLoss}\left(M(x) \| M(x') \right)$ is the privacy loss of M on arbitrary inputs $x, x' \in \mathcal{X}^n$ differing in one entry. On the other hand, M being (ε, δ)-differentially

[5] Rényi divergence has a parameter $\alpha \in (1, \infty)$ which allows it to interpolate between KL-divergence ($\alpha \to 1$) and max-divergence ($\alpha \to \infty$). It should be thought of as a measure of dissimilarity between distributions. We define it formally in Section 2. Throughout, we assume that all logarithms are natural unless specified otherwise — that is, base $e \approx 2.718$.

[6] For clarity of exposition, we consider only ρ-zCDP in the introduction and give more general statements for (ξ, ρ)-zCDP later. We also believe that having a one-parameter definition is desirable.

private is equivalent, up to a small loss in parameters, to the requirement that $\mathbb{P}[Z > \varepsilon] \leq \delta$.

In contrast, zCDP entails a bound on the *moment generating function* of the privacy loss Z — that is, $\mathbb{E}\left[e^{(\alpha-1)Z}\right]$ as a function of $\alpha-1$. The bound (2) implies that Z is a *subgaussian* random variable with small mean. Intuitively, this means that Z resembles a Gaussian distribution with mean $\xi + \rho$ and variance 2ρ. In particular, we obtain strong tail bounds on Z. Namely (2) implies that

$$\mathbb{P}[Z > \lambda + \xi + \rho] \leq e^{-\lambda^2/4\rho}$$

for all $\lambda > 0$.[7]

Thus zCDP requires that the privacy loss random variable is concentrated around zero (hence the name). That is, Z is "small" with high probability, with larger deviations from zero becoming increasingly unlikely. Hence we are unlikely to be able to distinguish x from x' given the output of $M(x)$ or $M(x')$. Note that the randomness of the privacy loss random variable is taken only over the randomnesss of the mechanism M.

Comparison to the Definition of Dwork and Rothblum For comparison, Dwork and Rothblum [DR16] define (μ, τ)-*concentrated differential privacy* for a randomized mechanism $M : \mathcal{X}^n \to \mathcal{Y}$ as the requirement that, if $Z = \mathsf{PrivLoss}(M(x)\|M(x'))$ is the privacy loss for $x, x' \in \mathcal{X}^n$ differing on one entry, then

$$\mathbb{E}[Z] \leq \mu \qquad \text{and} \qquad \mathbb{E}\left[e^{(\alpha-1)(Z-\mathbb{E}[Z])}\right] \leq e^{(\alpha-1)^2 \frac{1}{2}\tau^2}$$

for all $\alpha \in \mathbb{R}$. That is, they require both a bound on the mean of the privacy loss and that the privacy loss is tightly concentrated around its mean. To distinguish our definitions, we refer to their definition as *mean-concentrated differential privacy* (or mCDP).

Our definition, zCDP, is a *relaxation* of mCDP. In particular, a (μ, τ)-mCDP mechanism is also $(\mu - \tau^2/2, \tau^2/2)$-zCDP (which is tight for the Gaussian mechanism example), whereas the converse is not true. (However, a partial converse holds; see Lemma 24.)

1.2 Results

Relationship between zCDP and Differential Privacy Like Dwork and Rothblum's formulation of concentrated differential privacy, zCDP can be thought of as providing guarantees of (ε, δ)-differential privacy *for all* values of $\delta > 0$:

[7] We only discuss bounds on the upper tail of Z. We can obtain similar bounds on the lower tail of $Z = \mathsf{PrivLoss}(M(x)\|M(x'))$ by considering $Z' = \mathsf{PrivLoss}(M(x')\|M(x))$.

Proposition 3 *If M provides ρ-zCDP, then M is $(\rho+2\sqrt{\rho \log(1/\delta)}, \delta)$-differentially private for any $\delta > 0$.*

There is also a partial converse, which shows that, up to a loss in parameters, zCDP is equivalent to differential privacy with this $\forall \delta > 0$ quantification (see Lemma 22).

There is also a direct link from pure differential privacy to zCDP:

Proposition 4 *If M satisfies ε-differential privacy, then M satisfies $(\frac{1}{2}\varepsilon^2)$-zCDP.*

Dwork and Rothblum [DR16, Theorem 3.5] give a slightly weaker version of Proposition 4, which implies that ε-differential privacy yields $(\frac{1}{2}\varepsilon(e^\varepsilon - 1))$-zCDP; this improves on an earlier bound [DRV10] by the factor $\frac{1}{2}$.

Propositions 3 and 4 show that zCDP is an intermediate notion between pure differential privacy and approximate differential privacy. Indeed, many algorithms satisfying approximate differential privacy do in fact also satisfy zCDP.

Gaussian Mechanism Just as with mCDP, the prototypical example of a mechanism satisfying zCDP is the *Gaussian mechanism*, which answers a real-valued query on a database by perturbing the true answer with Gaussian noise.

Definition 5 (Sensitivity) *A function $q : \mathcal{X}^n \to \mathbb{R}$ has* sensitivity Δ *if for all $x, x' \in \mathcal{X}^n$ differing in a single entry, we have $|q(x) - q(x')| \leq \Delta$.*

Proposition 6 (Gaussian Mechanism) *Let $q : \mathcal{X}^n \to \mathbb{R}$ be a sensitivity-Δ query. Consider the mechanism $M : \mathcal{X}^n \to \mathbb{R}$ that on input x, releases a sample from $\mathcal{N}(q(x), \sigma^2)$. Then M satisfies $(\Delta^2/2\sigma^2)$-zCDP.*

We remark that either inequality defining zCDP — (1) or (2) — is exactly tight for the Gaussian mechanism for all values of α. Thus the definition of zCDP seems tailored to the Gaussian mechanism.

Basic Properties of zCDP Our definition of zCDP satisfies the key basic properties of differential privacy. Foremost, these properties include smooth degradation under composition, and invariance under postprocessing:

Lemma 7 (Composition) *Let $M : \mathcal{X}^n \to \mathcal{Y}$ and $M' : \mathcal{X}^n \to \mathcal{Z}$ be randomized algorithms. Suppose M satisfies ρ-zCDP and M' satisfies ρ'-zCDP. Define $M'' : \mathcal{X}^n \to \mathcal{Y} \times \mathcal{Z}$ by $M''(x) = (M(x), M'(x))$. Then M'' satisfies $(\rho + \rho')$-zCDP.*

Lemma 8 (Postprocessing) *Let $M : \mathcal{X}^n \to \mathcal{Y}$ and $f : \mathcal{Y} \to \mathcal{Z}$ be randomized algorithms. Suppose M satisfies ρ-zCDP. Define $M' : \mathcal{X}^n \to \mathcal{Z}$ by $M'(x) = f(M(x))$. Then M' satisfies ρ-zCDP.*

These properties follow immediately from corresponding properties of the Rényi divergence outlined in Lemma 15.

We remark that Dwork and Rothblum's definition of mCDP is not closed under postprocessing; we provide a counterexample in the full version of this work. (However, an arbitrary amount of postprocessing can worsen the guarantees of mCDP by at most constant factors.)

Group Privacy A mechanism M guarantees *group privacy* if no small group of individuals has a significant effect on the outcome of a computation (whereas the definition of zCDP only refers to individuals, which are groups of size 1). That is, group privacy for groups of size k guarantees that, if x and x' are inputs differing on k entries (rather than a single entry), then the outputs $M(x)$ and $M(x')$ are close.

Dwork and Rothblum [DR16, Theorem 4.1] gave nearly tight bounds on the group privacy guarantees of concentrated differential privacy, showing that a $(\mu = \tau^2/2, \tau)$-concentrated differentially private mechanism affords $(k^2\mu \cdot (1 + o(1)), k\tau \cdot (1 + o(1)))$-concentrated differential privacy for groups of size $k = o(1/\tau)$. We are able to show a group privacy guarantee for zCDP that is exactly tight and works for a wider range of parameters:

Proposition 9 *Let $M : \mathcal{X}^n \to \mathcal{Y}$ satisfy ρ-zCDP. Then M guarantees $(k^2\rho)$-zCDP for groups of size k — i.e. for every $x, x' \in \mathcal{X}^n$ differing in up to k entries and every $\alpha \in (1, \infty)$, we have*

$$\mathrm{D}_\alpha\left(M(x) \| M(x')\right) \leq (k^2\rho) \cdot \alpha.$$

In particular, this bound is achieved (simultaneously for all values α) by the Gaussian mechanism. Our proof is also simpler than that of Dwork and Rothblum; see Section 5.

Lower Bounds The strong group privacy guarantees of zCDP yield, as an unfortunate consequence, strong lower bounds as well. We show that, as with pure differential privacy, zCDP is susceptible to information-based lower bounds, as well as to so-called packing arguments [HT10,MMP+10,De12]:

Theorem 10 *Let $M : \mathcal{X}^n \to \mathcal{Y}$ satisfy ρ-zCDP. Let X be a random variable on $\mathcal{X}^n$. Then*

$$I\left(X; M(X)\right) \leq \rho \cdot n^2,$$

where $I(\cdot;\cdot)$ denotes the mutual information between the random variables (in nats, rather than bits). Furthermore, if the entries of X are independent, then $I(X; M(X)) \leq \rho \cdot n$.

Theorem 10 yields strong lower bounds for zCDP mechanisms, as we can construct distributions X such that $M(X)$ reveals a lot of information about X (i.e. $I(X; M(X))$ is large) for any accurate M.

In particular, we obtain a strong separation between approximate differential privacy and zCDP. For example, we can show that releasing an accurate approximate histogram (or, equivalently, accurately answering all point queries) on a data domain of size k requires an input with at least $n = \Theta(\sqrt{\log k})$ entries to satisfy zCDP. In contrast, under approximate differential privacy, n can be *independent* of the domain size k [BNS13]! In particular, our lower bounds show that "stability-based" techniques (such as those in the propose-test-release framework [DL09]) are not compatible with zCDP.

Our lower bound exploits the strong group privacy guarantee afforded by zCDP. Group privacy has been used to prove tight lower bounds for pure differential privacy [HT10,De12] and approximate differential privacy [SU15a]. These results highlight the fact that group privacy is often the limiting factor for private data analysis. For (ε, δ)-differential privacy, group privacy becomes vacuous for groups of size $k = \Theta(\log(1/\delta)/\varepsilon)$. Indeed, stability-based techniques exploit precisely this breakdown in group privacy.

As a result of this strong lower bound, we show that any mechanism for answering statistical queries that satisfies zCDP can be converted into a mechanism satisfying pure differential privacy with only a quadratic blowup in its sample complexity. More precisely, the following theorem illustrates a more general result we prove in Section 7.

Theorem 11 *Let $n \in \mathbb{N}$ and $\alpha \geq 1/n$ be arbitrary. Set $\varepsilon = \alpha$ and $\rho = \alpha^2$. Let $q : \mathcal{X} \to [0,1]^k$ be an arbitrary family of statistical queries. Suppose $M : \mathcal{X}^n \to [0,1]^k$ satisfies ρ-zCDP and*

$$\mathop{\mathbb{E}}_{M}\left[\|M(x) - q(x)\|_\infty\right] \leq \alpha$$

for all $x \in \mathcal{X}^n$. Then there exists $M' : \mathcal{X}^{n'} \to [0,1]^k$ for $n' = 5n^2$ satisfying ε-differential privacy and

$$\mathop{\mathbb{E}}_{M'}\left[\|M'(x) - q(x)\|_\infty\right] \leq 10\alpha$$

for all $x \in \mathcal{X}^{n'}$.

For some classes of queries, this reduction is essentially tight. For example, for k one-way marginals, the Gaussian mechanism achieves sample complexity $n = \Theta(\sqrt{k})$ subject to zCDP, whereas the Laplace mechanism achieves sample complexity $n = \Theta(k)$ subject to pure differential privacy, which is known to be optimal. For more details, see Sections 6 and 7.

Approximate zCDP To circumvent these strong lower bounds for zCDP, we consider a relaxation of zCDP in the spirit of approximate differential privacy that permits a small probability δ of (catastrophic) failure:

Definition 12 (Approximate zCDP) *A randomized mechanism $M : \mathcal{X}^n \to \mathcal{Y}$ is δ-approximately (ξ, ρ)-zCDP if, for all $x, x' \in \mathcal{X}^n$ differing on a single*

entry, there exist events E (depending on $M(x)$) and E' (depending on $M(x')$) such that $\mathbb{P}[E] \geq 1-\delta$, $\mathbb{P}[E'] \geq 1-\delta$, and

$$\forall \alpha \in (1,\infty) \qquad \mathrm{D}_\alpha\left(M(x)|_E \| M(x')|_{E'}\right) \leq \xi + \rho \cdot \alpha \\ \wedge \qquad \mathrm{D}_\alpha\left(M(x')|_{E'} \| M(x)|_E\right) \leq \xi + \rho \cdot \alpha,$$

where $M(x)|_E$ denotes the distribution of $M(x)$ conditioned on the event E. We further define δ-approximate ρ-zCDP to be δ-approximate $(0,\rho)$-zCDP.

In particular, setting $\delta = 0$ gives the original definition of zCDP. However, this definition unifies zCDP with approximate differential privacy:

Proposition 13 *If M satisfies (ε,δ)-differential privacy, then M satisfies δ-approximate $\frac{1}{2}\varepsilon^2$-zCDP.*

Approximate zCDP retains most of the desirable properties of zCDP, but allows us to incorporate stability-based techniques and bypass the above lower bounds. This also presents a unified tool to analyse a composition of zCDP with approximate differential privacy; see Section 8.

Related Work Our work builds on the aforementioned prior work of Dwork and Rothblum [DR16].[8] We view our definition of concentrated differential privacy as being "morally equivalent" to their definition of concentrated differential privacy, in the sense that both definitions formalize the same concept.[9] (The formal relationship between the two definitions is discussed in Section 4.) However, the definition of zCDP generally seems easier to work with than mCDP. In particular, our formulation in terms of Rényi divergence simplifies many analyses.

Dwork and Rothblum prove several results about concentrated differential privacy that are similar to ours. Namely, they prove analogous properties of mCDP as we prove for zCDP. However, as noted, some of their bounds are weaker than ours; also, they do not explore lower bounds.

Several of the ideas underlying concentrated differential privacy are implicit in earlier works. In particular, the proof of the advanced composition theorem of Dwork, Rothblum, and Vadhan [DRV10] essentially uses the ideas of concentrated differential privacy.

We also remark that Tardos [Tar08] used Rényi divergence to prove lower bounds for cryptographic objects called *fingerprinting codes*. Fingerprinting codes turn out to be closely related to differential privacy [Ull13,BUV14,SU15b], and Tardos' lower bound can be (loosely) viewed as a kind of privacy-preserving algorithm.

[8] Although Dwork and Rothblum's work only appeared publicly in March 2016, they shared a preliminary draft of their paper with us before we commenced this work. As such, our ideas are heavily inspired by theirs.

[9] We use "concentrated differential privacy" (CDP) to refer to the underlying *concept* formalized by both definitions.

Further Work We believe that concentrated differential privacy is a useful tool for analysing private computations, as it provides both simpler and tighter bounds. We hope that CDP will be prove useful in both the theory and practice of differential privacy.

Furthermore, our lower bounds show that CDP can really be a much more stringent condition than approximate differential privacy. Thus CDP defines a "subclass" of all (ε, δ)-differentially private algorithms. This subclass includes most differentially private algorithms in the literature, but not all — the most notable exceptions being algorithms that use the propose-test-release approach [DL09] to exploit low local sensitivity.

This "CDP subclass" warrants further exploration. In particular, is there a "complete" mechanism for this class of algorithms, in the same sense that the exponential mechanism [MT07,BLR13] is complete for pure differential privacy? Can we obtain a simple characterization of the sample complexity needed to satisfy CDP? The ability to prove stronger and simpler lower bounds for CDP than for approximate DP may be useful for showing the limitations of certain algorithmic paradigms. For example, any differentially private algorithm that only uses the Laplace mechanism, the exponential mechanism, the Gaussian mechanism, and the "sparse vector" technique, along with composition and postprocessing will be subject to the lower bounds for CDP.

There is also room to examine how to interpret the zCDP privacy guarantee. In particular, we leave it as an open question to understand the extent to which ρ-zCDP provides a stronger privacy guarantee than the implied (ε, δ)-DP guarantees (cf. Proposition 3).

In general, much of the literature on differential privacy can be re-examined through the lens of CDP, which may yield new insights and results.

2 Rényi Divergence

Recall the definition of Rényi divergence:

Definition 14 (Rényi Divergence [Rén61, Equation (3.3)]) *Let P and Q be probability distributions on Ω. For $\alpha \in (1, \infty)$, we define the* Rényi divergence of order α between P and Q *as*

$$\begin{aligned} \mathrm{D}_\alpha(P\|Q) =& \frac{1}{\alpha - 1} \log \left(\int_\Omega P(x)^\alpha Q(x)^{1-\alpha} \mathrm{d}x \right) \\ =& \frac{1}{\alpha - 1} \log \left(\mathop{\mathbb{E}}_{x \sim P} \left[\left(\frac{P(x)}{Q(x)} \right)^{\alpha - 1} \right] \right), \end{aligned}$$

where $P(\cdot)$ and $Q(\cdot)$ are the probability mass/density functions of P and Q respectively or, more generally, $P(\cdot)/Q(\cdot)$ is the Radon-Nikodym derivative of P with respect to Q.

We also define the KL-divergence

$$\mathrm{D}_1(P\|Q) = \lim_{\alpha \to 1} \mathrm{D}_\alpha(P\|Q) = \int_\Omega P(x) \log \left(\frac{P(x)}{Q(x)} \right) \mathrm{d}x$$

and the max-divergence

$$\mathrm{D}_\infty\left(P\|Q\right) = \lim_{\alpha\to\infty}\mathrm{D}_\alpha\left(P\|Q\right) = \sup_{x\in\Omega}\log\left(\frac{P(x)}{Q(x)}\right).$$

Alternatively, Rényi divergence can be defined in terms of the privacy loss (Definition 2) between P and Q:

$$e^{(\alpha-1)\mathrm{D}_\alpha(P\|Q)} = \mathop{\mathbb{E}}_{Z\sim\mathsf{PrivLoss}(P\|Q)}\left[e^{(\alpha-1)Z}\right]$$

for all $\alpha\in(1,\infty)$. Moreover, $\mathrm{D}_1\left(P\|Q\right) = \mathop{\mathbb{E}}_{Z\sim\mathsf{PrivLoss}(P\|Q)}[Z]$.

We record several useful and well-known properties of Rényi divergence. We refer the reader to [vEH14] for proofs and discussion of these (and many other) properties.

Lemma 15 *Let P and Q be probability distributions and $\alpha\in[1,\infty]$.*

- Non-negativity: $\mathrm{D}_\alpha\left(P\|Q\right)\geq 0$ *with equality if and only if $P=Q$.*
- Composition: *Suppose P and Q are distributions on $\Omega\times\Theta$. Let P' and Q' denote the marginal distributions on Ω induced by P and Q respectively. For $x\in\Omega$, let P'_x and Q'_x denote the conditional distributions on Θ induced by P and Q respectively, where x specifies the first coordinate. Then*

  $$\mathrm{D}_\alpha\left(P'\|Q'\right)+\min_{x\in\Omega}\mathrm{D}_\alpha\left(P'_x\|Q'_x\right)\leq\mathrm{D}_\alpha\left(P\|Q\right)\leq\mathrm{D}_\alpha\left(P'\|Q'\right)+\max_{x\in\Omega}\mathrm{D}_\alpha\left(P'_x\|Q'_x\right).$$

  In particular if P and Q are product distributions, then the Rényi divergence between P and Q is just the sum of the Rényi divergences of the marginals.
- Quasi-Convexity: *Let P_0,P_1 and Q_0,Q_1 be distributions on Ω, and let $P=tP_0+(1-t)P_1$ and $Q=tQ_0+(1-t)Q_1$ for $t\in[0,1]$. Then* $\mathrm{D}_\alpha\left(P\|Q\right)\leq\max\{\mathrm{D}_\alpha\left(P_0\|Q_0\right),\mathrm{D}_\alpha\left(P_1\|Q_1\right)\}$. *Moreover, KL divergence is convex:*

  $$\mathrm{D}_1\left(P\|Q\right)\leq t\mathrm{D}_1\left(P_0\|Q_0\right)+(1-t)\mathrm{D}_1\left(P_1\|Q_1\right).$$

- Postprocessing: *Let P and Q be distributions on Ω and let $f:\Omega\to\Theta$ be a function. Let $f(P)$ and $f(Q)$ denote the distributions on Θ induced by applying f to P or Q respectively. Then* $\mathrm{D}_\alpha\left(f(P)\|f(Q)\right)\leq\mathrm{D}_\alpha\left(P\|Q\right)$. *Note that quasi-convexity allows us to extend this guarantee to the case where f is a randomized mapping.*
- Monotonicity: *For $1\leq\alpha\leq\alpha'\leq\infty$,* $\mathrm{D}_\alpha\left(P\|Q\right)\leq\mathrm{D}_{\alpha'}\left(P\|Q\right)$.

2.1 Gaussian Mechanism

The following lemma gives the Rényi divergence between two Gaussian distributions with the same variance.

Lemma 16 *Let $\mu,\nu,\sigma\in\mathbb{R}$ and $\alpha\in[1,\infty)$. Then*

$$\mathrm{D}_\alpha\left(\mathcal{N}(\mu,\sigma^2)\middle\|\mathcal{N}(\nu,\sigma^2)\right) = \frac{\alpha(\mu-\nu)^2}{2\sigma^2}$$

Consequently, the Gaussian mechanism, which answers a sensitivity-Δ query by adding noise drawn from $\mathcal{N}(0, \sigma^2)$, satisfies $\left(\frac{\Delta^2}{2\sigma^2}\right)$-zCDP (Proposition 6).

For the multivariate Gaussian mechanism, Lemma 16 generalises to the following.

Lemma 17 *Let $\mu, \nu \in \mathbb{R}^d$, $\sigma \in \mathbb{R}$, and $\alpha \in [1, \infty)$. Then*

$$\mathrm{D}_\alpha\left(\mathcal{N}(\mu, \sigma^2 I_d) \| \mathcal{N}(\nu, \sigma^2 I_d)\right) = \frac{\alpha \|\mu - \nu\|_2^2}{2\sigma^2}$$

Thus, if $M : \mathcal{X}^n \to \mathbb{R}^d$ is the mechanism that, on input x, releases a sample from $\mathcal{N}(q(x), \sigma^2 I_d)$ for some function $q : \mathcal{X}^n \to \mathbb{R}^d$, then M satisfies ρ-zCDP for

$$\rho = \frac{1}{2\sigma^2} \sup_{\substack{x, x' \in \mathcal{X}^n \\ \text{differing in one entry}}} \|q(x) - q(x')\|_2^2. \tag{3}$$

3 Relation to Differential Privacy

We now discuss the relationship between zCDP and the traditional definitions of pure and approximate differential privacy. There is a close relationship between the notions, but not an exact characterization.

Definition 18 (Differential Privacy (DP) [DMNS06,DKM+06]) *A randomized mechanism $M : \mathcal{X}^n \to \mathcal{Y}$ satisfies (ε, δ)-differential privacy if, for all $x, x' \in \mathcal{X}$ differing in a single entry, we have*

$$\mathbb{P}[M(x) \in S] \leq e^\varepsilon \mathbb{P}[M(x') \in S] + \delta$$

for all (measurable) $S \subset \mathcal{Y}$. Further define ε-differential privacy to be $(\varepsilon, 0)$-differential privacy.

3.1 Pure DP versus zCDP

We now show that ε-differential privacy implies $(\frac{1}{2}\varepsilon^2)$-zCDP (Proposition 4).

Proposition 19 *Let P and Q be probability distributions on Ω satisfying $\mathrm{D}_\infty(P\|Q) \leq \varepsilon$ and $\mathrm{D}_\infty(Q\|P) \leq \varepsilon$. Then $\mathrm{D}_\alpha(P\|Q) \leq \frac{1}{2}\varepsilon^2\alpha$ for all $\alpha > 1$.*

Remark 20 *In particular, Proposition 19 shows that the KL-divergence $\mathrm{D}_1(P\|Q) \leq \frac{1}{2}\varepsilon^2$. A bound on the KL-divergence between random variables in terms of their max-divergence is an important ingredient in the analysis of the advanced composition theorem [DRV10]. Our bound sharpens (up to lower order terms) and, in our opinion, simplifies the previous bound of $\mathrm{D}_1(P\|Q) \leq \frac{1}{2}\varepsilon(e^\varepsilon - 1)$ proved by Dwork and Rothblum [DR16].*

Proof (Proof of Proposition 19.). We may assume $\frac{1}{2}\varepsilon\alpha \le 1$, as otherwise $\frac{1}{2}\varepsilon^2\alpha > \varepsilon$, whence the result follows from monotonicity. We must show that

$$e^{(\alpha-1)\mathrm{D}_\alpha(P\|Q)} = \mathop{\mathbb{E}}_{x\sim Q}\left[\left(\frac{P(x)}{Q(x)}\right)^\alpha\right] \le e^{\frac{1}{2}\alpha(\alpha-1)\varepsilon^2}.$$

We know that $e^{-\varepsilon} \le \frac{P(x)}{Q(x)} \le e^{\varepsilon}$ for all x. Define a random function $A : \Omega \to \{e^{-\varepsilon}, e^{\varepsilon}\}$ by $\mathop{\mathbb{E}}_{A}[A(x)] = \frac{P(x)}{Q(x)}$ for all x. By Jensen's inequality,

$$\mathop{\mathbb{E}}_{x\sim Q}\left[\left(\frac{P(x)}{Q(x)}\right)^\alpha\right] = \mathop{\mathbb{E}}_{x\sim Q}\left[\left(\mathop{\mathbb{E}}_{A}[A(x)]\right)^\alpha\right] \le \mathop{\mathbb{E}}_{x\sim Q}\left[\mathop{\mathbb{E}}_{A}[A(x)^\alpha]\right] = \mathop{\mathbb{E}}_{A}[A^\alpha],$$

where A denotes $A(x)$ for a random $x \sim Q$. We also have $\mathop{\mathbb{E}}_{A}[A] = \mathop{\mathbb{E}}_{x\sim Q}\left[\frac{P(x)}{Q(x)}\right] = 1$. From this equation, we can conclude that

$$\mathop{\mathbb{P}}_{A}\left[A = e^{-\varepsilon}\right] = \frac{e^{\varepsilon}-1}{e^{\varepsilon}-e^{-\varepsilon}} \qquad \text{and} \qquad \mathop{\mathbb{P}}_{A}\left[A = e^{\varepsilon}\right] = \frac{1-e^{-\varepsilon}}{e^{\varepsilon}-e^{-\varepsilon}}.$$

Thus

$$\begin{aligned} e^{(\alpha-1)\mathrm{D}_\alpha(P\|Q)} &\le \mathop{\mathbb{E}}_{A}[A^\alpha] \\ &= \frac{e^{\varepsilon}-1}{e^{\varepsilon}-e^{-\varepsilon}} \cdot e^{-\alpha\varepsilon} + \frac{1-e^{-\varepsilon}}{e^{\varepsilon}-e^{-\varepsilon}} \cdot e^{\alpha\varepsilon} \\ &= \frac{\sinh(\alpha\varepsilon) - \sinh((\alpha-1)\varepsilon)}{\sinh(\varepsilon)}. \end{aligned}$$

The result now follows from the following inequality, which is proved in the full version of this work.

$$0 \le y < x \le 2 \implies \frac{\sinh(x)-\sinh(y)}{\sinh(x-y)} \le e^{\frac{1}{2}xy}.$$

3.2 Approximate DP versus zCDP

The statements in this section show that, up to some loss in parameters, zCDP is equivalent to a family of (ε,δ)-DP guarantees for all $\delta > 0$.

Lemma 21 *Let $M : \mathcal{X}^n \to \mathcal{Y}$ satisfy (ξ,ρ)-zCDP. Then M satisfies (ε,δ)-DP for all $\delta > 0$ and*

$$\varepsilon = \xi + \rho + \sqrt{4\rho\log(1/\delta)}.$$

Thus to achieve a given (ε,δ)-DP guarantee it suffices to satisfy (ξ,ρ)-zCDP with

$$\rho = \left(\sqrt{\varepsilon-\xi+\log(1/\delta)} - \sqrt{\log(1/\delta)}\right)^2 \approx \frac{(\varepsilon-\xi)^2}{4\log(1/\delta)}.$$

Proof. Let $x, x' \in \mathcal{X}^n$ be neighbouring. Define $f(y) = \log(\mathbb{P}[M(x) = y] / \mathbb{P}[M(x') = y])$. Let $Y \sim M(x)$ and $Z = f(Y)$. That is, $Z = \mathsf{PrivLoss}(M(x) \| M(x'))$ is the privacy loss random variable. Fix $\alpha \in (1, \infty)$ to be chosen later. Then

$$\mathbb{E}\left[e^{(\alpha-1)Z}\right] = \mathop{\mathbb{E}}_{Y \sim M(x)}\left[\left(\frac{\mathbb{P}[M(x) = Y]}{\mathbb{P}[M(x') = Y]}\right)^{\alpha-1}\right] = e^{(\alpha-1)\mathrm{D}_\alpha\left(M(x)\middle\|M(x')\right)} \leq e^{(\alpha-1)(\xi+\rho\alpha)}.$$

By Markov's inequality

$$\mathbb{P}[Z > \varepsilon] = \mathbb{P}\left[e^{(\alpha-1)Z} > e^{(\alpha-1)\varepsilon}\right] \leq \frac{\mathbb{E}\left[e^{(\alpha-1)Z}\right]}{e^{(\alpha-1)\varepsilon}} \leq e^{(\alpha-1)(\xi+\rho\alpha-\varepsilon)}.$$

Choosing $\alpha = (\varepsilon - \xi + \rho)/2\rho > 1$ gives

$$\mathbb{P}[Z > \varepsilon] \leq e^{-(\varepsilon-\xi-\rho)^2/4\rho} \leq \delta.$$

This implies that for any measurable $S \subset \mathcal{Y}$,

$$\mathbb{P}[M(x) \in S] \leq e^{\varepsilon}\mathbb{P}[M(x') \in S] + \delta.$$

Lemma 21 is not tight, and we give a quantitative refinement in Lemma 38 (setting $\delta = 0$ there). There, we also show a partial converse to Lemma 21:

Lemma 22 *Let $M : \mathcal{X}^n \to \mathcal{Y}$ satisfy (ε, δ)-DP for all $\delta > 0$ and*

$$\varepsilon = \hat{\xi} + \sqrt{\hat{\rho}\log(1/\delta)} \tag{4}$$

for some constants $\hat{\xi}, \hat{\rho} \in [0, 1]$. Then M is $\left(\hat{\xi} - \frac{1}{4}\hat{\rho} + 5\sqrt[4]{\hat{\rho}}, \frac{1}{4}\hat{\rho}\right)$-zCDP.

Thus zCDP and DP are equivalent up to a (potentially substantial) loss in parameters and the quantification over all δ.

4 Zero- versus Mean-Concentrated Differential Privacy

We begin by recalling the definition of mean-concentrated differential privacy:

Definition 23 (Mean-Concentrated Differential Privacy [DR16]) *A randomized mechanism $M : \mathcal{X}^n \to \mathcal{Y}$ satisfies (μ, τ)-mCDP if, for all $x, x' \in \mathcal{X}^n$ differing in one entry, and letting $Z = \mathsf{PrivLoss}(M(x) \| M(x'))$, we have*

$$\mathbb{E}[Z] \leq \mu \quad \text{and} \quad \mathbb{E}\left[e^{\lambda\left(Z - \mathbb{E}[Z]\right)}\right] \leq e^{\lambda^2 \cdot \tau^2/2}$$

for all $\lambda \in \mathbb{R}$.

In contrast (ξ, ρ)-zCDP requires that, for all $\alpha \in (1, \infty)$, $\mathbb{E}\left[e^{(\alpha-1)Z}\right] \leq e^{(\alpha-1)(\xi+\rho\alpha)}$, where $Z \sim \mathsf{PrivLoss}\left(M(x) \| M(x')\right)$ is the privacy loss random variable. In the full version of this work, we show that these definitions are equivalent up to a (potentially significant) loss in parameters.

Lemma 24 *If $M : \mathcal{X}^n \to \mathcal{Y}$ satisfies (μ, τ)-mCDP, then M satisfies $(\mu - \tau^2/2, \tau^2/2)$-zCDP. Conversely, if $M : \mathcal{X}^n \to \mathcal{Y}$ satisfies (ξ, ρ)-zCDP, then M satisfies $(\xi + \rho, O(\sqrt{\xi + 2\rho}))$-mCDP.*

Thus we can convert (μ, τ)-mCDP into $(\mu - \tau^2/2, \tau^2/2)$-zCDP and then back to $(\mu, O(\sqrt{\mu + \tau^2/2}))$-mCDP. This may result in a large loss in parameters, which is why, for example, pure DP can be characterised in terms of zCDP, but not in terms of mCDP.

We view zCDP as a relaxation of mCDP; mCDP requires the privacy loss to be "tightly" concentrated about its mean and that the mean is close to the origin. The triangle inequality then implies that the privacy loss is "weakly" concentrated about the origin. (The difference between "tightly" and "weakly" accounts for the use of the triangle inequality.) On the other hand, zCDP direcly requires that the privacy loss is weakly concentrated about the origin. That is, zCDP gives a subgaussian bound on the privacy loss that is centered at zero, whereas mCDP gives a subgaussian bound that is centered at the mean and separately bounds the mean.

There may be some advantage to the stronger requirement of mCDP, either in terms of what kind of privacy guarantee it affords, or how it can be used as an analytic tool. However, it seems that for most applications, we only need what zCDP provides.

5 Group Privacy

In this section we show that zCDP provides privacy protections to small groups of individuals.

Definition 25 (zCDP for Groups) *We say that a mechanism $M : \mathcal{X}^n \to \mathcal{Y}$ provides (ξ, ρ)-zCDP for groups of size k if, for every $x, x' \in \mathcal{X}^n$ differing in at most k entries, we have*

$$\forall \alpha \in (1, \infty) \qquad \mathrm{D}_\alpha\left(M(x) \| M(x')\right) \leq \xi + \rho \cdot \alpha.$$

The usual definition of zCDP only applies to groups of size 1. Here we show that it implies bounds for all group sizes. We begin with a technical lemma.

Lemma 26 (Triangle-like Inequality for Rényi Divergence) *Let P, Q, and R be probability distributions. Then*

$$\mathrm{D}_\alpha\left(P \| Q\right) \leq \frac{k\alpha}{k\alpha - 1} \mathrm{D}_{\frac{k\alpha-1}{k-1}}\left(P \| R\right) + \mathrm{D}_{k\alpha}\left(R \| Q\right) \tag{5}$$

for all $k, \alpha \in (1, \infty)$.

Proof. Let $p = \frac{k\alpha-1}{\alpha(k-1)}$ and $q = \frac{k\alpha-1}{\alpha-1}$. Then $\frac{1}{p} + \frac{1}{q} = \frac{\alpha(k-1)+(\alpha-1)}{k\alpha-1} = 1$. By Hölder's inequality,

$$\begin{aligned}
e^{(\alpha-1)\mathrm{D}_\alpha(P\|Q)} &= \int_\Omega P(x)^\alpha Q(x)^{1-\alpha}\mathrm{d}x \\
&= \int_\Omega P(x)^\alpha R(x)^{-\alpha} \cdot R(x)^{\alpha-1} Q(x)^{1-\alpha} \cdot R(x)\mathrm{d}x \\
&= \mathop{\mathbb{E}}_{x\sim R}\left[\left(\frac{P(x)}{R(x)}\right)^\alpha \cdot \left(\frac{R(x)}{Q(x)}\right)^{\alpha-1}\right] \\
&\le \mathop{\mathbb{E}}_{x\sim R}\left[\left(\frac{P(x)}{R(x)}\right)^{p\alpha}\right]^{1/p} \cdot \mathop{\mathbb{E}}_{x\sim R}\left[\left(\frac{R(x)}{Q(x)}\right)^{q(\alpha-1)}\right]^{1/q} \\
&= e^{(p\alpha-1)\mathrm{D}_{p\alpha}(P\|R)/p} \cdot e^{q(\alpha-1)\mathrm{D}_{q(\alpha-1)+1}(R\|Q)/q}.
\end{aligned}$$

Taking logarithms and rearranging gives

$$\mathrm{D}_\alpha(P\|Q) \le \frac{p\alpha-1}{p(\alpha-1)}\mathrm{D}_{p\alpha}(P\|R) + \mathrm{D}_{q(\alpha-1)+1}(R\|Q).$$

Now $p\alpha = \frac{k\alpha-1}{k-1}$, $q(\alpha-1)+1 = k\alpha$, and $\frac{p\alpha-1}{p(\alpha-1)} = \frac{k\alpha}{k\alpha-1}$.

Proposition 27 *If $M : \mathcal{X}^n \to \mathcal{Y}$ satisfies (ξ,ρ)-zCDP, then M gives $(\xi \cdot k\sum_{i=1}^k \frac{1}{i}, \rho\cdot k^2)$-zCDP for groups of size k.*

In particular, (ξ,ρ)-zCDP implies $(\xi\cdot O(k\log k), \rho\cdot k^2)$-zCDP for groups of size k. The Gaussian mechanism shows that $k^2\rho$ is the optimal dependence on ρ. However, $O(k\log k)\xi$ is not the optimal dependence on ξ: $(\xi,0)$-zCDP implies $(k\xi,0)$-zCDP for groups of size k.

Proof. We show this by induction on k. The statement is clearly true for groups of size 1. We now assume the statement holds for groups of size $k-1$ and will verify it for groups of size k.

Let $x, x' \in \mathcal{X}^n$ differ in k entries. Let $\hat{x} \in \mathcal{X}^n$ be such that x and $\hat{x}$ differ in $k-1$ entries and x' and $\hat{x}$ differ in one entry.

Then, by the induction hypothesis,

$$\mathrm{D}_\alpha(M(x)\|M(\hat{x})) \le \xi\cdot(k-1)\sum_{i=1}^{k-1}\frac{1}{i} + \rho\cdot(k-1)^2\cdot\alpha$$

and, by zCDP,

$$\mathrm{D}_\alpha(M(\hat{x})\|M(x')) \le \xi + \rho\cdot\alpha$$

for all $\alpha \in (1, \infty)$. By (5), for any $\alpha \in (1, \infty)$,

$$\begin{aligned}
&\mathrm{D}_\alpha\left(M(x) \| M(x')\right) \\
&\leq \frac{k\alpha}{k\alpha - 1} \mathrm{D}_{\frac{k\alpha-1}{k-1}}\left(M(x) \| M(\hat{x})\right) + \mathrm{D}_{k\alpha}\left(M(\hat{x}) \| M(x')\right) \\
&\leq \frac{k\alpha}{k\alpha - 1}\left(\xi \cdot (k-1) \sum_{i=1}^{k-1} \frac{1}{i} + \rho \cdot (k-1)^2 \cdot \frac{k\alpha - 1}{k-1}\right) + \xi + \rho \cdot k\alpha \\
&= \xi \cdot \left(1 + \frac{k\alpha}{k\alpha - 1}(k-1) \sum_{i=1}^{k-1} \frac{1}{i}\right) + \rho \cdot \left(\frac{k\alpha}{k\alpha - 1}(k-1)^2 \frac{k\alpha - 1}{k-1} + k\alpha\right) \\
&\leq \xi \cdot k \sum_{i=1}^{k} \frac{1}{i} + \rho \cdot k^2 \cdot \alpha,
\end{aligned}$$

where the last inequality follows from the fact that $\frac{k\alpha}{k\alpha-1}$ is a decreasing function of α for $\alpha > 1$.

6 Lower Bounds

In this section we develop tools to prove lower bounds for zCDP. We will use group privacy to bound the mutual information between the input and the output of a mechanism satisfying zCDP. Thus, if we are able to construct a distribution on inputs such that any accurate mechanism must reveal a high amount of information about its input, we obtain a lower bound showing that no accurate mechanism satisfying zCDP can be accurate for this data distribution.

We begin with the simplest form of our mutual information bound, which is an analogue of the bound of [MMP+10] for pure differential privacy:

Proposition 28 *Let $M : \mathcal{X}^n \to \mathcal{Y}$ satisfy (ξ, ρ)-zCDP. Let X be a random variable in $\mathcal{X}^n$. Then*

$$I(X; M(X)) \leq \xi \cdot n(1 + \log n) + \rho \cdot n^2,$$

where I denotes mutual information (measured in nats, rather than bits).

Proof. By Proposition 27, M is $(\xi \cdot n \sum_{i=1}^{n} \frac{1}{i}, \rho \cdot n^2)$-zCDP for groups of size n. Thus

$$\mathrm{D}_1\left(M(x) \| M(x')\right) \leq \xi \cdot n \sum_{i=1}^{n} \frac{1}{i} + \rho \cdot n^2 \leq \xi \cdot n(1 + \log n) + \rho \cdot n^2$$

for all $x, x' \in \mathcal{X}^n$. Since KL-divergence is convex,

$$\begin{aligned}
I(X; M(X)) &= \mathop{\mathbb{E}}_{x \leftarrow X}\left[\mathrm{D}_1\left(M(x) \| M(X)\right)\right] \\
&\leq \mathop{\mathbb{E}}_{x \leftarrow X}\left[\mathop{\mathbb{E}}_{x' \leftarrow X}\left[\mathrm{D}_1\left(M(x) \| M(x')\right)\right]\right] \\
&\leq \mathop{\mathbb{E}}_{x \leftarrow X}\left[\mathop{\mathbb{E}}_{x' \leftarrow X}\left[\xi \cdot n(1 + \log n) + \rho \cdot n^2\right]\right] \\
&= \xi \cdot n(1 + \log n) + \rho \cdot n^2.
\end{aligned}$$

The reason this lower bound works is the strong group privacy guarantee — even for groups of size n, we obtain nontrivial privacy guarantees. While this is good for privacy it is bad for usefulness, as it implies that even information that is "global" (rather than specific to a individual or a small group) is protected. These lower bounds reinforce the connection between group privacy and lower bounds [HT10,De12,SU15a].

In contrast, (ε, δ)-DP is not susceptible to such a lower bound because it gives a vacuous privacy guarantee for groups of size $k = O(\log(1/\delta)/\varepsilon)$. This helps explain the power of the propose-test-release paradigm.

Furthermore, we obtain even stronger mutual information bounds when the entries of the distribution are independent:

Lemma 29 *Let $M : \mathcal{X}^m \to \mathcal{Y}$ satisfy (ξ, ρ)-zCDP. Let X be a random variable in $\mathcal{X}^m$ with independent entries. Then*

$$I\left(X; M(X)\right) \leq (\xi + \rho) \cdot m,$$

where I denotes mutual information (measured in nats, rather than bits).

Proof. First, by the chain rule for mutual information,

$$I(X; M(X)) = \sum_{i \in [m]} I(X_i; M(X) | X_{1 \cdots i-1}),$$

where

$$\begin{aligned} I(X_i; M(X) | X_{1 \cdots i-1}) &= \mathop{\mathbb{E}}_{x \leftarrow X_{1 \cdots i-1}} \left[I(X_i | X_{1 \cdots i-1} = x; M(X) | X_{1 \cdots i-1} = x) \right] \\ &= \mathop{\mathbb{E}}_{x \leftarrow X_{1 \cdots i-1}} \left[I(X_i; M(x, X_{i \cdots m})) \right], \end{aligned}$$

by independence of the X_is.

We can define mutual information in terms of KL-divergence:

$$\begin{aligned} I(X_i; M(x, X_{i \cdots m})) &= \mathop{\mathbb{E}}_{y \leftarrow X_i} \left[\mathrm{D}_1 \left(M(x, X_{i \cdots m}) | X_i = y \| M(x, X_{i \cdots m}) \right) \right] \\ &= \mathop{\mathbb{E}}_{y \leftarrow X_i} \left[\mathrm{D}_1 \left(M(x, y, X_{i+1 \cdots m}) \| M(x, X_{i \cdots m}) \right) \right]. \end{aligned}$$

By zCDP, we know that for all $x \in \mathcal{X}^{i-1}$, $y, y' \in \mathcal{X}$, and $z \in \mathcal{X}^{m-i}$, we have

$$\mathrm{D}_1 \left(M(x, y, z) \| M(x, y', z) \right) \leq \xi + \rho.$$

Thus, by the convexity of KL-divergence,

$$\mathrm{D}_1 \left(M(x, y, X_{i+1 \cdots m}) \| M(x, X_{i \cdots m}) \right) \leq \xi + \rho$$

for all x and y. The result follows.

More generally, we can combine dependent and independent rows as follows.

Theorem 30 *Let $M : \mathcal{X}^n \to \mathcal{Y}$ satisfy (ξ, ρ)-zCDP. Take $n = m \cdot \ell$. Let $X^1, \cdots, X^m$ be independent random variables on $\mathcal{X}^\ell$. Denote $X = (X^1, \cdots, X^m) \in \mathcal{X}^n$. Then*

$$I\left(X; M(X)\right) \leq m \cdot \left(\xi \cdot \ell(1 + \log \ell) + \rho \cdot \ell^2\right),$$

where I denotes the mutual information (measured in nats, rather than bits).

6.1 Example Applications of the Lower Bound

We informally discuss a few applications of our information-based lower bounds to some simple and well-studied problems in differential privacy.

One-Way Marginals Consider $M : \mathcal{X}^n \to \mathcal{Y}$ where $\mathcal{X} = \{0,1\}^d$ and $\mathcal{Y} = [0,1]^d$. The goal of M is to estimate the attribute means, or one-way marginals, of its input database x:

$$M(x) \approx \overline{x} = \frac{1}{n} \sum_{i \in [n]} x_i.$$

It is known that this is possible subject to ε-DP if and only if $n = \Theta(d/\varepsilon)$ [HT10,SU15a]. This is possible subject to (ε, δ)-DP if and only if $n = \tilde{\Theta}(\sqrt{d \log(1/\delta)}/\varepsilon)$, assuming $\delta \ll 1/n$ [BUV14,SU15a].

We now analyze what can be accomplished with zCDP. Adding independent noise drawn from $\mathcal{N}(0, d/2n^2\rho)$ to each of the d coordinates of $\overline{x}$ satisfies ρ-zCDP. This gives accurate answers as long as $n \gg \sqrt{d/\rho}$.

For a lower bound, consider sampling $X_1 \in \{0,1\}^d$ uniformly at random. Set $X_i = X_1$ for all $i \in [n]$. By Proposition 28,

$$I(X; M(X)) \leq n^2 \rho$$

for any ρ-zCDP $M : (\{0,1\}^d)^n \to [0,1]^d$. However, if M is accurate, we can recover (most of) X_1 from $M(X)$, whence $I(X; M(X)) \geq \Omega(d)$.This yields a lower bound of $n \geq \Omega(\sqrt{d/\rho})$, which is tight up to constant factors.

Histograms (a.k.a. Point Queries) Consider $M : \mathcal{X}^n \to \mathcal{Y}$, where $\mathcal{X} = [T]$ and $\mathcal{Y} = \mathbb{R}^T$. The goal of M is to estimate the histogram of its input:

$$M(x)_t \approx h_t(x) = |\{i \in [n] : x_i = t\}|$$

For ε-DP it is possible to do this if and only if $n = \Theta(\log(T)/\varepsilon)$); the optimal algorithm is to independently sample

$$M(x)_t \sim h_t(x) + \mathsf{Laplace}(2/\varepsilon).$$

However, for (ε, δ)-DP, it is possible to attain sample complexity $n = O(\log(1/\delta)/\varepsilon)$ [BNS16, Theorem 3.13]. Interestingly, for zCDP we can show that $n = \Theta(\sqrt{\log(T)/\rho})$ is sufficient and necessary:

Sampling

$$M(x)_t \sim h_t(x) + \mathcal{N}(0, 1/\rho)$$

independently for $t \in [T]$ satisfies ρ-zCDP. Moreover,

$$\mathbb{P}\left[\max_{t\in[T]} |M(x)_t - h_t(x)| \geq \lambda\right] \leq T \cdot \mathbb{P}\left[|\mathcal{N}(0, 1/\rho)| > \lambda\right] \leq T \cdot e^{-\lambda^2\rho/2}.$$

In particular $\mathbb{P}\left[\max_{t\in[T]} |M(x)_t - h_t(x)| \geq \sqrt{\log(T/\beta)/\rho}\right] \leq \beta$ for all $\beta > 0$. Thus this algorithm is accurate if $n \gg \sqrt{\log(T)/\rho}$.

On the other hand, if we sample $X_1 \in [T]$ uniformly at random and set $X_i = X_1$ for all $i \in [n]$, then $I(X; M(X)) \geq \Omega(\log T)$ for any accurate M, as we can recover X_1 from $M(X)$ if M is accurate. Proposition 28 thus implies that $n \geq \Omega(\sqrt{\log(T)/\rho})$ is necessary to obtain accuracy. This gives a strong separation between approximate DP and zCDP.

Lower Bounds with Accuracy The above examples can be easily discussed in terms of a more formal and quantitative definition of accuracy. For instance, in the full version of this work, we revisit the histogram example:

Proposition 31 *If $M : [T]^n \to \mathbb{R}^T$ satisfies ρ-zCDP and*

$$\forall x \in [T]^n \qquad \mathop{\mathbb{E}}_{M}\left[\max_{t\in[T]} \left|M(x)_t - h_t(x)\right|\right] \leq \alpha n,$$

then $n \geq \Omega(\sqrt{\log(\alpha^2 T)/\rho\alpha^2})$.

We remark that our lower bounds for zCDP can be converted to lower bounds for mCDP using Lemma 24.

7 Obtaining Pure DP Mechanisms from zCDP

We now establish limits on what more can be achieved with zCDP over pure differential privacy. In particular, we prove that any mechanism satisfying zCDP can be converted into a mechanism satisfying pure DP with at most a quadratic blowup in sample complexity. Formally, we show the following theorem.

Theorem 32 *Fix $n \in \mathbb{N}$, $n' \in \mathbb{N}$, $k \in \mathbb{N}$ $\alpha > 0$, and $\varepsilon > 0$. Let $q : \mathcal{X} \to \mathbb{R}^k$ and let $\|\cdot\|$ be a norm on $\mathbb{R}^k$. Assume $\max_{x\in\mathcal{X}} \|q(x)\| \leq 1$. Suppose there exists a (ξ, ρ)-zCDP mechanism $M : \mathcal{X}^n \to \mathbb{R}^k$ such that for all $x \in \mathcal{X}^n$,*

$$\mathop{\mathbb{E}}_{M}\left[\|M(x) - q(x)\|\right] \leq \alpha.$$

Assume $\xi \leq \alpha^2$, $\rho \leq \alpha^2$, and

$$n' \geq \frac{4}{\varepsilon\alpha}\left(\rho \cdot n^2 + \xi \cdot n \cdot (1 + \log n) + 1\right).$$

Then there exists a $(\varepsilon, 0)$-differentially private $M' : \mathcal{X}^{n'} \to \mathbb{R}^k$ satisfying

$$\mathop{\mathbb{E}}_{M'}[\|M'(x) - q(x)\|] \leq 10\alpha$$

and

$$\mathop{\mathbb{P}}_{M'}\left[\|M'(x) - q(x)\| > 10\alpha + \frac{4}{\varepsilon n'}\log\left(\frac{1}{\beta}\right)\right] \leq \beta$$

for all $x \in \mathcal{X}^{n'}$ and $\beta > 0$.

Before discussing the proof of Theorem 32, we make some remarks about its statement:

- Unfortunately, the theorem only works for families of statistical queries $q : \mathcal{X} \to \mathbb{R}^k$. However, it works equally well for $\|\cdot\|_\infty$ and $\|\cdot\|_1$ error bounds.
- If $\xi = 0$, we have $n' = O(n^2\rho/\varepsilon\alpha)$. So, if ρ, ε, and α are all constants, we have $n' = O(n^2)$. This justifies our informal statement that we can convert any mechanism satisfying zCDP into one satisfying pure DP with a quadratic blowup in sample complexity.
- The requirement that $\xi, \rho \leq \alpha^2$ is only used to show that

$$\max_{x \in \mathcal{X}^{n'}} \min_{\hat{x} \in \mathcal{X}^n} \|q(x) - q(\hat{x})\| \leq 2\alpha. \tag{6}$$

  However, in many situations (6) holds even when $\xi, \rho \gg \alpha^2$. For example, if $n \geq O(\log(k)/\alpha^2)$ or even $n \geq O(VC(q)/\alpha^2)$ then (6) is automatically satisfied. The technical condition (6) is needed to relate the part of the proof with inputs of size n to the part with inputs of size n'.

The proof of Theorem 32 is not constructive. Rather than directly constructing a mechanism satisfying pure DP from any mechanism satisfying zCDP, we show the contrapositive statement: any lower bound for pure DP can be converted into a lower bound for zCDP. Pure DP is characterized by so-called packing lower bounds and the exponential mechanism.

In the full version of this work, we use a greedy argument to show that for any output space and any desired accuracy, there is a set T that is simultaneously a "packing" and a "net:"

Lemma 33 *Let $(\mathcal{Y}, d)$ be a metric space. Fix $\alpha > 0$. Then there exists a countable $T \subset \mathcal{Y}$ such that both of the following hold.*

- *(Net:) Either T is infinite or for all $y' \in \mathcal{Y}$ there exists $y \in T$ with $d(y, y') \leq \alpha$.*
- *(Packing:) For all $y, y' \in T$, if $y \neq y'$, then $d(y, y') > \alpha$.*

It is well-known that a net yields a pure DP algorithm:

Lemma 34 (Exponential Mechanism [MT07,BLR13]) *Let $\ell : \mathcal{X}^n \times T \to \mathbb{R}$ satisfy $|\ell(x, y) - \ell(x', y)| \leq \Delta$ for all $x, x' \in \mathcal{X}^n$ differing in one entry and all*

$y \in T$. Then, for all $\varepsilon > 0$, there exists an ε-differentially private $M : \mathcal{X}^n \to T$ such that

$$\mathop{\mathbb{P}}_{M}\left[\ell(x, M(x)) \leq \min_{y \in T} \ell(x, y) + \frac{2\Delta}{\varepsilon} \log\left(\frac{|T|}{\beta}\right)\right] \geq 1 - \beta$$

and

$$\mathop{\mathbb{E}}_{M}\left[\ell(x, M(x))\right] \leq \min_{y \in T} \ell(x, y) + \frac{2\Delta}{\varepsilon} \log |T|$$

for all $x \in \mathcal{X}^n$ and $\beta > 0$.

On the other hand, in the full version of this work we use Proposition 28 to show that a packing yields a lower bound for zCDP:

Lemma 35 *Let $(\mathcal{Y}, d)$ be a metric space and $q : \mathcal{X}^n \to \mathcal{Y}$ a function. Let $M : \mathcal{X}^n \to \mathcal{Y}$ be a (ξ, ρ)-zCDP mechanism satisfying*

$$\mathop{\mathbb{P}}_{M}\left[d(M(x), q(x)) > \alpha/2\right] \leq \beta$$

for all $x \in \mathcal{X}^n$. Let $T \subset \mathcal{Y}$ be such that $d(y, y') > \alpha$, for all $y, y' \in T$ with $y \neq y'$. Assume that for all $y \in T$ there exists $x \in \mathcal{X}^n$ with $q(x) = y$. Then

$$(1 - \beta) \log |T| - \log 2 \leq \xi \cdot n(1 + \log n) + \rho \cdot n^2.$$

In particular, if $\xi = 0$, we have

$$n \geq \sqrt{\frac{(1 - \beta) \log |T| - \log 2}{\rho}} = \Omega(\sqrt{\log |T|/\rho}).$$

In the full version of this work, we combine these lemmas to prove Theorem 32.

8 Approximate zCDP

Recall our definition of approximate zCDP:

Definition 36 (Approximate zCDP) *A randomised mechanism $M : \mathcal{X}^n \to \mathcal{Y}$ is δ-approximately (ξ, ρ)-zCDP if, for all $x, x' \in \mathcal{X}^n$ differing on a single entry, there exist events $E = E(M(x))$ and $E' = E'(M(x'))$ such that, for all $\alpha \in (1, \infty)$,*

$$\mathrm{D}_\alpha\left(M(x)|_E \| M(x')|_{E'}\right) \leq \xi + \rho \cdot \alpha \qquad \text{and} \qquad \mathrm{D}_\alpha\left(M(x')|_{E'} \| M(x)|_E\right) \leq \xi + \rho \cdot \alpha$$

and $\mathop{\mathbb{P}}_{M(x)}[E] \geq 1 - \delta$ and $\mathop{\mathbb{P}}_{M(x')}[E'] \geq 1 - \delta$.

Clearly 0-approximate zCDP is simply zCDP. Hence we have a generalization of zCDP. As we will show later in this section, δ-approximate $(\varepsilon, 0)$-zCDP is equivalent to (ε, δ)-DP. Thus we have also generalized approximate DP. Hence, this definition unifies both relaxations of pure DP.

Approximate zCDP is a three-parameter definition which allows us to capture many different aspects of differential privacy. However, three parameters is quite overwhelming. We believe that use of the one-parameter ρ-zCDP (or the two-parameter δ-approximate ρ-zCDP if necessary) is sufficient for most purposes.

It is easy to verify that the definition of approximate zCDP satisfies the usual composition and post-processing properties. However, the strong group privacy guarantees of Section 5 no longer apply to approximate zCDP and, hence, the strong lower bounds of Section 6 also no longer hold. Circumventing these lower bounds is part of the motivation for considering approximate zCDP.

In the full version of this work, we use techniques developed in [KOV15,MV16] to show that approximate DP can be converted to approximate zCDP.

Lemma 37 *If $M : \mathcal{X}^n \to \mathcal{Y}$ satisfies (ε, δ)-DP, then M satisfies δ-approximate $(\varepsilon, 0)$-zCDP, which, in turn, implies δ-approximate $(0, \frac{1}{2}\varepsilon^2)$-zCDP.*

Conversely, approximate zCDP also implies approximate DP. The following result sharpens Lemma 21.

Lemma 38 *Suppose $M : \mathcal{X}^n \to \mathcal{Y}$ satisfies δ-approximate (ξ, ρ)-zCDP. If $\rho = 0$, then M satisfies (ξ, δ)-DP. In general, M satisfies $(\varepsilon, \delta + (1-\delta)\delta')$-DP for all $\varepsilon \geq \xi + \rho$, where*

$$\delta' = e^{-(\varepsilon-\xi-\rho)^2/4\rho} \cdot \min \left\{ \begin{array}{l} 1 \\ \sqrt{\pi \cdot \rho} \\ \frac{1}{1+(\varepsilon-\xi-\rho)/2\rho} \\ \frac{2}{1+\frac{\varepsilon-\xi-\rho}{2\rho}+\sqrt{\left(1+\frac{\varepsilon-\xi-\rho}{2\rho}\right)^2+\frac{4}{\pi\rho}}} \end{array} \right. .$$

A result is that we can give a sharper version of the so-called advanced composition theorem [DRV10]. Note that the following results are subsumed by the bounds of Kairouz, Oh, and Viswanath [KOV15] and Murtagh and Vadhan [MV16]. However, these bounds may be extended to analyse the composition of mechanisms satisfying CDP with mechanisms satisfying approximate DP. We believe that such a "unified" analysis of composition will be useful.

Applying Lemma 37 and Lemma 38 yields the following result.

Corollary 39 *Let $M_1, \cdots, M_k : \mathcal{X}^n \to \mathcal{Y}$ and let $M : \mathcal{X}^n \to \mathcal{Y}^k$ be their composition. Suppose each M_i satisfies $(\varepsilon_i, \delta_i)$-DP. Then M satisfies*

$$\left(\frac{1}{2}\|\varepsilon\|_2^2 + \sqrt{2}\lambda\|\varepsilon\|_2, \sqrt{\frac{\pi}{2}} \cdot \|\varepsilon\|_2 \cdot e^{-\lambda^2} + \|\delta\|_1 \right) \text{-DP}$$

for all $\lambda \geq 0$. Alternatively M satisfies

$$\left(\frac{1}{2}\|\varepsilon\|_2^2 + \sqrt{2\log(\sqrt{\pi/2} \cdot \|\varepsilon\|_2/\delta') \cdot \|\varepsilon\|_2}, \delta' + \|\delta\|_1 \right) \text{-DP}$$

for all $\delta' \geq 0$.

In comparison to the composition theorem of [DRV10], we save modestly by a constant factor in the first term and, in most cases $\sqrt{\pi/2}\|\varepsilon\|_2 < 1$, whence the logarithmic term is an improvement over the usual advanced composition theorem.

Acknowledgements We thank Cynthia Dwork and Guy Rothblum for sharing a preliminary draft of their work with us. We also thank Ilya Mironov, Kobbi Nissim, Adam Smith, Salil Vadhan, and the Harvard Differential Privacy Research Group for helpful discussions and suggestions.

References

BLR13. Avrim Blum, Katrina Ligett, and Aaron Roth. A learning theory approach to noninteractive database privacy. *J. ACM*, 60(2):12, 2013.

BNS13. Amos Beimel, Kobbi Nissim, and Uri Stemmer. Private learning and sanitization: Pure vs. approximate differential privacy. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 16th International Workshop, APPROX 2013, and 17th International Workshop, RANDOM 2013, Berkeley, CA, USA, August 21-23, 2013. Proceedings*, pages 363–378, 2013.

BNS16. Mark Bun, Kobbi Nissim, and Uri Stemmer. Simultaneous private learning of multiple concepts. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, ITCS '16, pages 369–380, New York, NY, USA, 2016. ACM.

BUV14. Mark Bun, Jonathan Ullman, and Salil P. Vadhan. Fingerprinting codes and the price of approximate differential privacy. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 1–10, 2014.

De12. Anindya De. Lower bounds in differential privacy. In *Proceedings of the 9th International Conference on Theory of Cryptography*, TCC'12, pages 321–338, Berlin, Heidelberg, 2012. Springer-Verlag.

DKM+06. Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology - EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28 - June 1, 2006, Proceedings*, pages 486–503, 2006.

DL09. Cynthia Dwork and Jing Lei. Differential privacy and robust statistics. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 371–380, 2009.

DMNS06. Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, pages 265–284, 2006.

DR16. Cynthia Dwork and Guy Rothblum. Concentrated differential privacy. *CoRR*, abs/1603.01887, 2016.

DRV10. Cynthia Dwork, Guy N. Rothblum, and Salil P. Vadhan. Boosting and differential privacy. In *IEEE Symposium on Foundations of Computer Science (FOCS '10)*, pages 51–60. IEEE, 23–26 October 2010.

HT10. Moritz Hardt and Kunal Talwar. On the geometry of differential privacy. In *Proceedings of the Forty-second ACM Symposium on Theory of Computing*, STOC '10, pages 705–714, New York, NY, USA, 2010. ACM.

KOV15. Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, pages 1376–1385, 2015.

MMP+10. Andrew McGregor, Ilya Mironov, Toniann Pitassi, Omer Reingold, Kunal Talwar, and Salil P. Vadhan. The limits of two-party differential privacy. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 81–90, 2010.

MT07. F. McSherry and K. Talwar. Mechanism design via differential privacy. In *Foundations of Computer Science, 2007. FOCS '07. 48th Annual IEEE Symposium on*, pages 94–103, Oct 2007.

MV16. Jack Murtagh and Salil P. Vadhan. The complexity of computing the optimal composition of differential privacy. In *Theory of Cryptography - 13th International Conference, TCC 2016-A, Tel Aviv, Israel, January 10-13, 2016, Proceedings, Part I*, pages 157–175, 2016.

Rén61. Alfréd Rényi. On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, pages 547–561, Berkeley, Calif., 1961. University of California Press.

SU15a. Thomas Steinke and Jonathan Ullman. Between pure and approximate differential privacy. *CoRR*, abs/1501.06095, 2015.

SU15b. Thomas Steinke and Jonathan Ullman. Interactive fingerprinting codes and the hardness of preventing false discovery. In *COLT*, 2015. http://arxiv.org/abs/1410.1228.

Tar08. Gábor Tardos. Optimal probabilistic fingerprint codes. *J. ACM*, 55(2), 2008.

Ull13. Jonathan Ullman. Answering n {2+ o (1)} counting queries with differential privacy is hard. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 361–370. ACM, 2013.

vEH14. T. van Erven and P. Harremos. Rényi divergence and Kullback-Leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, July 2014.