Differential-Privacy Capacity

Wael Alghamdi*[1], Shahab Asoodeh[2], Flavio P. Calmon[1], Oliver Kosut[3] and Lalitha Sankar[3]
[1] School of Engineering and Applied Sciences, Harvard University (emails: alghamdi@g.harvard.edu, flavio@seas.harvard.edu)
[2] Department of Computing and Software, McMaster University (email: asoodehs@mcmaster.ca)
[3] School of Electrical, Computer, and Energy Engineering, Arizona State University (emails: {okosut, lsankar}@asu.edu)

2024 IEEE International Symposium on Information Theory (ISIT) | 979-8-3503-8284-6/24/$31.00 ©2024 IEEE | DOI: 10.1109/ISIT57864.2024.10619626

***Abstract*—We formulate a fundamental limit in differential privacy under growing composition. We introduce the universal composition curve: the best privacy guarantee under repeated composition of a given privacy mechanism given only the sensitivity of the query. We define *privacy capacity* as the slowest growth rate of this universal composition curve among all privacy mechanisms. We show that, in the limit of large compositions, privacy capacity "single-letterizes" as a minimax KL-divergence term. Our privacy capacity formula extends previous literature results that connect differential privacy and KL-divergence via concentration theorems.**

I. Introduction and Problem Formulation

Differential Privacy (DP) [1], [2] is the standard mathematical method for quantifying privacy in statistical and machine learning (ML) applications that rely on querying individual-level and privacy-sensitive datasets (e.g., [3]–[5]). DP mechanisms aim to balance privacy and utility by randomizing queries to a dataset instead of disclosing exact query values. Naturally, DP-ensuring methods have an inherent utility-privacy trade-off: the more randomness introduced to a query value, the more privacy is gained, but less utility is preserved. This trade-off becomes more difficult to navigate in the practical setting of *composition*, i.e., when multiple queries are computed over the same dataset (such as DP-SGD [6] in ML, and statistical disclosure control as deployed by the US Census [7]).

The competing challenges of ensuring privacy and maintaining utility under multiple queries to sensitive data have motivated a stream of recent work that aims to solve one of two problems: quantifying privacy under composition (commonly referred to as DP accountants), and designing DP mechanisms that optimize the utility-privacy trade-off [6], [8]–[17]. We aim to shed light on these two problems by uncovering a fundamental limit of DP under growing composition, namely, we prove that the optimal privacy-per-composition rate among all possible DP mechanisms is given by the following minimax KL-divergence term:

$$\begin{aligned} \underset{P_{Y|X}}{\text{minimize}} \quad & \sup_{\|x-x'\|\le s} D(P_{Y|X=x} \parallel P_{Y|X=x'}) \\ \text{subject to} \quad & \sup_{x\in\mathbb{R}^m} \mathbb{E}\left[\|Y-x\|^{\alpha} \mid X=x\right] \le C. \end{aligned} \tag{1}$$

Due to the resemblance of this fundamental limit to other information-theoretic capacity results (i.e., an operational definition that single-letterizes into an information-measure minimax optimization), we term this DP fundamental limit the *privacy capacity*.[1]

A. Primitives of Differential Privacy

The basic DP primitives include: a dataset $d \in \mathcal{D}$, a query function $q : \mathcal{D} \to \mathcal{X}$, and a randomized version Y of the query value $X = q(d)$. In the sequel, we will consider real vector-valued queries, i.e., $\mathcal{X} = \mathbb{R}^m$, and we will also take Y to be an $\mathbb{R}^m$-valued random vector. For example, d could be a table comprised of demographic information for the individuals living in a given county; $X = q(d)$ the median income of the individuals included in d; and Y the privacy mechanism of adding independent Gaussian noise. Thus, given any dataset $d \in \mathcal{D}$, there is an associated output distribution, and we may denote its underlying probability measure by $P_{Y|d}$ (which would be $\mathcal{N}(q(d), \sigma^2)$ for the Gaussian mechanism). In short, the randomized algorithm governed by $\{P_{Y|d}\}_{d\in\mathcal{D}}$ is decomposable into two parts: computing the query $X = q(d)$, then randomizing it as $(Y \mid X = x) \sim P_{Y|X=x}$, where we are now thinking of $P_{Y|X}$ as a Markov kernel on $\mathcal{X} = \mathcal{Y} = \mathbb{R}^m$. We will call the randomization $P_{Y|X}$ the *privacy mechanism*.

The most widely used variant of DP, known as approximate DP, quantifies the privacy afforded by the ensuing randomized algorithm as follows. We say that $\{P_{Y|d}\}_{d\in\mathcal{D}}$ satisfies (ε, δ)-DP, for some $\varepsilon \ge 0$ and $\delta \in [0, 1]$, if and only if the inequality

$$P_{Y|d}(A) \le e^{\varepsilon} P_{Y|d'}(A) + \delta \tag{2}$$

holds for all possible events A and for all pairs of *neighboring* datasets $d, d' \in \mathcal{D}$ (i.e., d and d' differ by one record, wich is commonly denoted by $d \simeq d'$). When ε and δ are both small, inequality (2) requires that the output distribution of the privacy mechanism does not change by much when exactly one individual is added to, removed from, or swapped in the dataset. In particular, this condition ensures the privacy of each individual, as an observer of the privatized version Y cannot do much better than randomly guessing whether a certain individual falls within the dataset. The definition in (2) is referred to as the *single-shot* setting, which we lift to the composition setting next.

*Corresponding author, and the remaining authors are in alphabetical order.

This work is supported in part by the NSF awards CAREER-1845852, CIF-1900750, CIF-2231707, CIF-2312667, CIF-1901243, CIF-2312666, SCH-2205080, and CIF-2007688, and by NSERC Canada.

[1] We note that our setup is mathematically more similar to the rate-distortion setting rather than, say, channel capacity. Nevertheless, we use the word "capacity" as it relates to the "highest privacy" for a given distortion constraint.

B. Universal Composition Curve

The relevant setting in practice is often *composition* of DP mechanisms, i.e., when we have multiple query functions[2] $\{q_j : \mathcal{D} \to \mathcal{X}\}_{j\in[k]}$. Applying these query functions on the same dataset $d \in \mathcal{D}$, one obtains k query values $X^k = (X_1, \cdots, X_k) = (q_1(d), \cdots, q_k(d))$. Then, a privatized version $Y^k = (Y_1, \cdots, Y_k)$ of X^k is computed by letting the Y_j be randomizations of the X_j independently of each other, i.e., the Y_j are mutually independent given d. The composed mechanism can also be seen as a randomized algorithm $\{P_{Y^k|d}\}_{d\in\mathcal{D}}$. When the privacy mechanisms $P_{Y_j|X_j} = P_{Y|X}$ are all the same (hence, $P_{Y^k|X^k=x^k} = \prod_{j\in[k]} P_{Y|X=x_j}$), we will call the full randomized algorithm $\{P_{Y^k|d}\}_{d\in\mathcal{D}}$ the *k-fold composition* of $P_{Y|X}$. Note that the composition satisfies (ε, δ)-DP when the analogous inequality to (2) holds, i.e., when

$$\sup_{d,d'\in\mathcal{D}\,:\,d\simeq d'} \sup_{A \text{ measurable}} P_{Y^k|d}(A) - e^{\varepsilon} P_{Y^k|d'}(A) \le \delta. \quad (3)$$

One additional primitive in DP of relevance to this work is the query *sensitivity*, which measures the maximal deviation of the query value for neighboring datasets. More precisely, we define the sensitivity in the single-shot setting by

$$\Delta(\mathcal{D}, \simeq, q) \coloneqq \sup_{d,d'\in\mathcal{D}\,:\,d\simeq d'} \|q(d) - q(d')\|, \quad (4)$$

where $\|\cdot\|$ denotes the Euclidean norm in $\mathbb{R}^m$. The sensitivity allows us to abstract away the triplet $(\mathcal{D}, \simeq, q)$ and quantify privacy in a universal, data-agnostic way. In other words, via considering the sensitivity, we may derive DP guarantees that hold universally over the class of triplets

$$\mathcal{T}(s) \coloneqq \{(\mathcal{D}, \simeq, q) \;:\; \mathcal{D} \text{ a set, } \simeq \text{ a binary relation on } \mathcal{D}, \\ q : \mathcal{D} \to \mathcal{X},\ \Delta(\mathcal{D}, \simeq, q) \le s\}.$$

In addition, sensitivity can be defined in the composition setting naturally by $\Delta(\mathcal{D}, \simeq, \{q_j\}_{j\in[k]}) \coloneqq \max_{j\in[k]} \Delta(\mathcal{D}, \simeq, q_j)$. Then, the class $\mathcal{T}(s)$ is naturally generalized for the composition setting into the class $\mathcal{T}^{\otimes k}(s)$ comprised of all possible triplets $(\mathcal{D}, \simeq, \{q_j\}_{j\in[k]})$ such that $(\mathcal{D}, \simeq, q_j) \in \mathcal{T}(s)$ for each $j \in [k]$.

Putting all the above DP primitives together, we arrive at the definition of the *universal composition curve*, which measures the best DP guarantees of the k-fold composition of a given privacy mechanism if one has access to only the sensitivity.

Definition 1 (Universal composition curve)**.** Consider any Markov kernel $P_{Y|X}$ on $\mathbb{R}^m$, and let $s > 0$ and $k \in \mathbb{N}$ be fixed. We define the *k-fold universal composition curve* of $P_{Y|X}$ to be the function $\delta^{\otimes k}_{P_{Y|X}} : \mathbb{R} \to [0, 1]$ defined at each $\varepsilon \in \mathbb{R}$ as the minimal number $\delta \in [0, 1]$ so that the k-fold composition of $P_{Y|X}$ satisfies (ε, δ)-DP for every possible triplet of bounded sensitivity $(\mathcal{D}, \simeq, \{q_j\}_{j\in[k]}) \in \mathcal{T}^{\otimes k}(s)$. In addition, we define the dual curve $\varepsilon^{\otimes k}_{P_{Y|X}}(\delta) \coloneqq \inf\{\varepsilon \ge 0 \;:\; \delta^{\otimes k}_{P_{Y|X}}(\varepsilon) \le \delta\}$.

[2]Strictly speaking, what we describe here is *non-adaptive* composition. In the *adaptive* counterpart, the j-th query function can take as input the previous $j-1$ mechanism outputs. Due to space limitation, and for clarity of presentation, we focus here on non-adaptive composition. Nevertheless, our privacy capacity results are exactly the same for adaptive composition too, which will be the subject of the fuller version of the present paper.

C. Privacy-Utility Trade-off

The discussion thus far has been on quantifying the privacy of a given mechanism, which we complement next by introducing a measure of its *utility*. We will measure the utility loss introduced by the privacy mechanism $P_{Y|X}$ via an input-worst-case and conditional-output-average of $\|Y - X\|^{\alpha}$:

$$L_{P_{Y|X}}(\alpha) \coloneqq \sup_{x\in\mathbb{R}^m} \mathbb{E}\left[\|Y - x\|^{\alpha} \mid X = x\right], \quad (5)$$

where $\alpha > 0$ is any prescribed constant. Thus, having $L_{P_{Y|X}}(\alpha) = 0$ corresponds to having $Y = X$ surely (so maximal utility is attained), and the larger $L_{P_{Y|X}}(\alpha)$ is, the less utility Y provides as a replacement for X.

In this paper, we formalize the following statement:

The best possible DP per composition (i.e., the lowest ratio $\frac{1}{k}\varepsilon^{\otimes k}_{P_{Y|X}}(\delta)$) is the minimax KL-divergence in (1)*.*

We call this fundamental limit on DP guarantees under composition the *privacy capacity*, formally defined in Definition 3. As we discuss in the next motivation section, the right normalizing factor for the composition curve is $\frac{1}{k}$.

D. Motivation: The Dichotomy of Linearly Growing DP Guarantees and Complete Randomness

To motivate the study of privacy capacity, we point to the lower bound on the composition curve shown in Theorem 3. In particular, this lower bound yields a dichotomy: for any mechanism $P_{Y|X}$, one of the following two scenarios holds,

1) $\varepsilon^{\otimes k}_{P_{Y|X}}(\delta_k) = 0$ for every $k \in \mathbb{N}$ and every $\delta_k \in [0, 1]$; or
2) $\varepsilon^{\otimes k}_{P_{Y|X}}(\delta_k) = \Omega(k)$ as $k \to \infty$ for $\limsup_{k\to\infty} \delta_k < 1$.

Furthermore, the first scenario is a degenerate case: it holds if and only if Y is independent of X, i.e., the DP mechanism $P_{Y|X}$ has no meaningful utility. Therefore, we naturally arrive at the following problem on the fundamental limits of DP.

Problem 1. *Among all mechanisms $P_{Y|X}$ enjoying the same non-trivial utility, what is the minimal achievable value of the privacy-per-composition rate $\frac{1}{k}\varepsilon^{\otimes k}_{P_{Y|X}}(\delta)$ as $k \to \infty$?*

In the spirit of information-theoretic analysis of DP, we term a rate such as above an *achievable privacy rate*, and we call the best achievable rate the *privacy capacity* (see Definition 3). The above dichotomy readily implies a lower bound on the privacy capacity (i.e., it shows a *converse* result); a matching *achievability* result is also proved in this paper.

E. Related Literature

The connection between DP under composition and the KL-divergence has been noted before [13], [17]–[21]. The common thread in those works is that a concentration theorem (namely, one of: the law of large numbers, central-limit theorem, or Berry-Esseen theorem) is applied to conclude that privacy of certain well-behaved mechanisms is governed in various forms by a KL-divergence term. Our present paper is differentiated from those works in the following respects.

First, for the concentration theorems to be applicable, the underlying mechanism was restricted to satisfy the relevant

concentration theorem's premises in [13], [17]–[21]. More precisely, with the exception of [13, Theorem 7], those results require finiteness of the first few moments of the so-called privacy-loss random variable (PLRV). One key differentiating factor in the present paper, thus, is that we go beyond the concentration theorems and derive the privacy capacity result in Theorem 1 universally for all possible mechanisms. In particular, our results remove the restriction to boundedness of the variance of the PLRV in our prior works [19]–[21], which aim to derive optimal DP mechanisms via solving the minimax KL-divergence problem (1).

Second, the universal composition curve (Definition 1) is different from the privacy curves considered in [13], [17], [18]. The curves considered in those works are roughly the curves obtained by the products of pairs of the so-called tightly-dominating pairs of measures. Such curves satisfy certain tightness in their guarantees: they give the optimal DP guarantees under composition for *some* mechanisms (namely, those who have realizable worst-case pairs of inputs). In contrast, the curve we introduce in Definition 1 provides a tight guarantee under composition for *every* mechanism; indeed, that is exactly how it is defined to begin with. We note that the universal composition curve has appeared in our prior work [19]–[21], but it was defined therein via its explicit formula directly (in terms of the hockey-stick divergence) rather than via its universal property as we do herein.

Additionally, we relate the ε parameter with the KL-divergence via an explicit formula (equations (9), (13), (14)). Such an explicit formula does not appear in [13], [18].

Finally, the growing-composition regime in [13] can be seen to be complementary to ours. As listed in the premises of [13, Theorem 7], the (ε, δ) parameters to be composed are assumed therein to satisfy $\max_{1\le i\le k} \varepsilon_i^{(k)} \to 0$ and $\max_{1\le i\le k} \delta_i^{(k)} \to 0$ as $k \to \infty$, which is a regime in which the composed mechanisms get closer and closer to being perfectly random. In contrast, we take the complementary approach: we require the amount of randomness of the privacy mechanism to be bounded, then study how fast the privacy guarantees grow.

Notation and Assumptions. The queries and privacy mechanisms are all assumed to be $\mathbb{R}^m$-valued, where m is arbitrary but fixed. We denote by $\|\cdot\|$ the ℓ^2 norm on $\mathbb{R}^m$, and sensitivity with respect to it is fixed to be some arbitrary constant $s > 0$ that will be dropped from notation. Expectation will be denoted by $\mathbb{E}_P[f] = \int_{\mathbb{R}^m} f(x) dP(x)$ and $\mathbb{E}_p[f] = \mathbb{E}_P[f]$ if P has probability density function (PDF) p. The set of all Markov kernels on $\mathbb{R}^m$ is denoted by $\mathscr{R}$, and the KL-divergence by $D(\cdot \parallel \cdot)$. If $D(P\|Q) < \infty$, the variance of information density is defined by $\mathrm{V}(P\|Q) \coloneqq \mathbb{E}_P[(\log dP/dQ - D(P\|Q))^2]$.

Although our privacy capacity result holds for every possible privacy mechanism, some of our intermediary results are derived under the following assumption of boundedness of the variance of information densities.

Assumption 1. The Markov kernel $P_{Y|X}$ on $\mathbb{R}^m$ is such that $\sup_{\|x-x'\|\le s} \mathrm{V}(P_{Y|X=x} \parallel P_{Y|X=x'}) < \infty$.

All omitted proofs can be found in the extended version [22].

II. Main Result: Privacy Capacity Is the Minimax KL-Divergence

We give the formal definition of privacy capacity first, then state the main result in Theorem 1.

A. Utility-Aware DP

Naturally, one should impose a cost constraint on $P_{Y|X}$ for the DP optimization problem to be nontrivial. Indeed, if we can choose any mechanism $P_{Y|X}$ without a utility constraint restricting how far Y can be from X, then nothing prevents us from choosing a Y that is independent of X. In this case, we obtain the best possible privacy guarantee $\varepsilon_{P_{Y|X}}^{\otimes k}(\delta) = 0$ for any $\delta \in [0,1]$ and any composition number k. This degenerate situation implies that we must restrict $P_{Y|X}$ to fall within a strict subset of the collection $\mathscr{R}$ of all Markov kernels on $\mathbb{R}^m$, namely, a subset of mechanisms with a known utility bound.

We parametrize the utility of a DP mechanism $P_{Y|X}$ using a cost function and a cost bound. Specifically, we measure the deviation of Y from X using an input-worst-case (i.e., supremum over $x \in \mathbb{R}^m$) and an output-average (i.e., expectation over $P_{Y|X=x}$) cost on the difference $Y - X$ between the input and output of the DP mechanism. Formally, we consider the following subsets of Markov kernels with controllable distortion. (See (5) for the definition of $L_{P_{Y|X}}(\alpha)$.)

Definition 2. We collect all mechanisms satisfying a cost bound $C \ge 0$ as measured by the loss-function exponent $\alpha > 0$ into a set denoted by $\mathscr{R}(\alpha, C) \subset \mathscr{R}$, i.e.,

$$\mathscr{R}(\alpha, C) \coloneqq \left\{ P_{Y|X} \in \mathscr{R} \ : \ L_{P_{Y|X}}(\alpha) \le C \right\}. \tag{6}$$

Example 1. Consider the squared ℓ^2-norm utility, i.e., $\alpha = 2$. One element of the set $\mathscr{R}(2, \sigma^2)$ (for fixed $\sigma > 0$) is the m-dimensional Gaussian mechanism with coordinate-wise noise variance σ^2/m (i.e., $\mathcal{N}(0, (\sigma^2/m) I_m)$). In fact, any additive mechanism $Y = X + Z$ for noise Z that is independent of X and satisfying $\mathbb{E}[\|Z\|^2] \le \sigma^2$ will fall within $\mathscr{R}(2, \sigma^2)$. Further, those additive mechanisms comprise only a strict subset of $\mathscr{R}(2, \sigma^2)$, e.g., the mechanism $Y = X + \min(1, \|X\|)Z$ will also belong to $\mathscr{R}(2, \sigma^2)$.

B. Definition of Privacy Capacity

We define the *privacy capacity* as the best possible per-composition ε value—as the number of compositions grows without bound—among all mechanisms with $\delta \to 0$ and common utility.

Definition 3 (Privacy Capacity)**.** For a loss-function exponent $\alpha > 0$ and utility bound $C \ge 0$, we say that $\varepsilon \ge 0$ is an *achievable privacy rate for* (α, C) if there is a mechanism $P_{Y|X} \in \mathscr{R}(\alpha, C)$ whose privacy budget under composition is upper bounded by

$$\frac{1}{k} \varepsilon_{P_{Y|X}}^{\otimes k}(\delta_k) \le \varepsilon \tag{7}$$

for all $k \in \mathbb{N}$ and some $\delta_k \to 0$ (as $k \to \infty$). The *privacy capacity* is defined as the infimal achievable privacy rate

$$\varepsilon^\star(\alpha, C) \coloneqq \inf \{ \varepsilon : \varepsilon \text{ achievable privacy rate for } (\alpha, C) \}.$$

Remark 1. We make an analogy with the rate-distortion function as follows:

- The cost function $\|\cdot\|^\alpha$ and cost bound C play the role of distortion function and distortion bound, respectively;
- DP parameter δ plays the role of decoding error probability;
- number of compositions k plays the role of the blocklength;
- the DP parameter ε after composition plays the role of the (logarithm of) the size of the reconstruction codebook;
- the privacy rate and privacy capacity play the roles of coding rate and rate-distortion function, respectively; and
- as we prove in this paper, the minimax KL-divergence from (1) plays the role of the minimal mutual information.

The largest difference between the rate-distortion function and our privacy capacity is that rate-distortion assumes a known source distribution, whereas we make no assumption on the distribution of the mechanism input X beyond the sensitivity.

C. Statement of the Main Theorem

Viewing the definition of privacy capacity in Definition 3 as an *operational* one, we show that the corresponding *information* definition is given by the minimax KL divergence in (1). We denote this KL-divergence term using the following notation.

Definition 4 (Minimax KL-divergence)**.** For any loss-function exponent $\alpha > 0$ and utility bound $C \geq 0$, we denote the optimal value in (1) by

$$\mathrm{KL}^\star(\alpha, C) := \inf_{P_{Y|X} \in \mathscr{R}(\alpha, C)} \sup_{\|x - x'\| \leq s} D\left(P_{Y|X=x} \| P_{Y|X=x'}\right). \tag{8}$$

Our privacy capacity result shows that the two definitions $\varepsilon^\star(\alpha, C)$ and $\mathrm{KL}^\star(\alpha, C)$ coincide.

Theorem 1. *The privacy capacity is equal to the minimax KL-divergence in (1), i.e., for every $\alpha, C > 0$ we have*

$$\varepsilon^\star(\alpha, C) = \mathrm{KL}^\star(\alpha, C). \tag{9}$$

Furthermore, additive, continuous, spherically-symmetric mechanisms can achieve the privacy capacity: for any rate $\varepsilon > \varepsilon^\star(\alpha, C)$, there is spherically-symmetric PDF p over $\mathbb{R}^m$ such that the privacy rate ε is achievable by the additive DP mechanism $Y = X + Z$ with continuous noise $Z \sim p$ that is independent of X and satisfying the utility $P_{Y|X} \in \mathscr{R}(\alpha, C)$.

One powerful aspect of this privacy capacity result is its "single-letterization": whereas the initial definition of privacy capacity $\varepsilon^\star(\alpha, C)$ inherently depends on composition (hence on product measures), Theorem 1 shows that privacy capacity is in fact equally written as the minimax KL-divergence $\mathrm{KL}^\star(\alpha, C)$, a term that pertains to a single use of the DP mechanism and does not depend on composition.

III. Proof of the Main Theorem

We present the proof steps in this section and relegate the details of the intermediary results to the extended version [22] due to space limitation.

A. Preliminary Results

We will use the following shorthands:

$$\mathrm{KL}^{\max}_{P_{Y|X}} := \sup_{x, x' \in \mathbb{R}^m : \|x - x'\| \leq s} D(P_{Y|X=x} \| P_{Y|X=x'}), \tag{10}$$

$$\mathrm{V}^{\max}_{P_{Y|X}} := \sup_{x, x' \in \mathbb{R}^m : \|x - x'\| \leq s} \mathrm{V}(P_{Y|X=x} \| P_{Y|X=x'}). \tag{11}$$

We will also use the same notations with the subscript $P_{Y|X}$ replaced by the PDF p if the mechanism is given by $Y = X + Z$ for $Z \sim p$ independent of X.

We borrow and refine some of the results in [20], which will be required for the proof of our main theorem. First, the following limit is shown in [20, Theorem 1].

Theorem 2 ([20, Theorem 1])**.** *For any Markov kernel $P_{Y|X}$ on $\mathbb{R}^m$ that satisfies Assumption 1 and any $\delta \in (0, 1/2)$,*

$$\lim_{k \to \infty} \frac{\varepsilon^{\otimes k}_{P_{Y|X}}(\delta)}{k} = \mathrm{KL}^{\max}_{P_{Y|X}}. \tag{12}$$

We refine the above limit in the following two theorems (see the extended version [22] for the proofs). First, we show that the lower-bound part holds unconditionally.

Theorem 3. *For any Markov kernel $P_{Y|X}$ on $\mathbb{R}^m$, the k-fold dual composition curve satisfies the asymptotic lower bound*

$$\liminf_{k \to \infty} \frac{\varepsilon^{\otimes k}_{P_{Y|X}}(\delta_k)}{k} \geq \mathrm{KL}^{\max}_{P_{Y|X}}. \tag{13}$$

whenever $\limsup_{k \to \infty} \delta_k < 1$.

Second, we need an effective version of the upper bound (while keeping Assumption 1 this time).

Theorem 4. *Fix a Markov kernel $P_{Y|X}$ on $\mathbb{R}^m$ satisfying Assumption 1. Then, we have the finite-composition bound*

$$\varepsilon^{\otimes k}_{P_{Y|X}}(\delta) \leq k \cdot \mathrm{KL}^{\max}_{P_{Y|X}} + \sqrt{\frac{1}{\delta} \cdot k \cdot \mathrm{V}^{\max}_{P_{Y|X}}} \tag{14}$$

for every $k \in \mathbb{N}$ and $\delta \in (0, 1]$.

The other result we borrow from [20] is the fact that the minimax KL-divergence optimization (1) can be restricted to additive, continuous, spherically-symmetric mechanisms.

Theorem 5 ([20, Theorems 2–3])**.** *For any fixed constants $\alpha, C > 0$, there is an additive, continuous, spherically-symmetric mechanism solving the optimization problem (1).*

B. Proof of the Privacy Capacity Formula: Converse

We prove that the privacy capacity satisfies a fundamental lower bound, which is given by the solution to the KL-divergence optimization (1). Using information-theoretic terminology, we have the following *converse* result.

Theorem 6. *For any loss-function exponent $\alpha > 0$ and utility bound $C \geq 0$, the privacy capacity is lower bounded by the minimax KL-divergence in (1):*

$$\varepsilon^\star(\alpha, C) \geq \mathrm{KL}^\star(\alpha, C). \tag{15}$$

Proof. This converse result follows readily from the lower bound on the composition curve we prove in Theorem 3. Suppose $\varepsilon \geq 0$ is an achievable privacy rate for (α, C). Let $P_{Y|X} \in \mathscr{R}(\alpha, C)$ and $\delta_k \to 0$ be such that $\varepsilon^{\otimes k}_{P_{Y|X}}(\delta_k) \leq k \cdot \varepsilon$ for all k. In particular, $\limsup_{k\to\infty} \delta_k < 1$. Hence, by Theorem 3,

$$\varepsilon \geq \liminf_{k\to\infty} \frac{\varepsilon^{\otimes k}_{P_{Y|X}}(\delta_k)}{k} \geq \mathrm{KL}^{\max}_{P_{Y|X}} \geq \mathrm{KL}^{\star}(\alpha, C) \tag{16}$$

As this is true for all achievable rates ε, we conclude that $\varepsilon^{\star}(\alpha, C) \geq \mathrm{KL}^{\star}(\alpha, C)$, and the proof is complete. □

The derivation of the *achievability* result is more technically involved, because the corresponding upper bound in Theorem 4 does not hold for all possible mechanisms. Hence, new proof techniques are needed, which we illustrate next.

C. Proof of the Privacy Capacity Formula: Achievability

We complement the converse result of Theorem 6 with the following achievability result.

Theorem 7. *For any fixed constants $\alpha, C > 0$, all privacy rates above the minimax KL-divergence in (1) are achievable, i.e.,*

$$\varepsilon^{\star}(\alpha, C) \leq \mathrm{KL}^{\star}(\alpha, C). \tag{17}$$

Furthermore, for any rate $\varepsilon > \varepsilon^{\star}(\alpha, C)$, there is a continuous and spherically-symmetric noise Z (independent of X) such that the additive mechanism $Y = X + Z$ falls within the utility set $\mathscr{R}(\alpha, C)$ and achieves the privacy rate ε.

Proof. We first apply Theorem 5 to extract a spherically-symmetric PDF p over $\mathbb{R}^m$ that achieves the minimax KL-divergence: $\mathrm{KL}^{\star}(\alpha, C) = \mathrm{KL}^{\max}_p$. Fix any $C' > C$.

The difficulty is that we may not apply Theorem 4 on p directly, since it does not satisfy Assumption 1. Nevertheless, we approximate p via the convolution $q = p * \psi_\sigma$ where $\psi_\sigma \propto \exp(-\|z/\sigma\|^{1/2})$ and $\sigma > 0$ is small enough. Note that q is spherically symmetric.[3] This convolution satisfies several key properties. First, denoting $c_\alpha(z) = \|z\|^\alpha$, we prove in Proposition 8 in the extended version [22] that there is small enough $\sigma = \sigma(\alpha, C', p)$ so that

$$\mathbb{E}_q[c_\alpha] \leq C'. \tag{18}$$

In addition, by the data-processing inequality,

$$D(q \parallel T_a q) \leq D(p \parallel T_a p) \qquad \text{for all } a \in \mathbb{R}^m, \tag{19}$$

where we denote the shift by $(T_a f)(x) = f(x - a)$. Therefore,

$$\mathrm{KL}^{\star}(\alpha, C') \leq \mathrm{KL}^{\max}_q \leq \mathrm{KL}^{\max}_p = \mathrm{KL}^{\star}(\alpha, C). \tag{20}$$

Importantly, by Proposition 9 in the extended version [22],

$$\mathrm{V}^{\max}_q < \infty, \tag{21}$$

which follows from the key property that $\left|\log \frac{q(x)}{q(y)}\right| \leq \frac{\|x-y\|^\gamma}{\sigma^\gamma}$.

Now, we may proceed by applying Theorem 4 on q. Let $\delta_k = \min(1, \mathrm{V}^{\max}_q/\sqrt{k})$, so $\delta_k \to 0$ as $k \to \infty$. Applying Theorem 4 on the mechanism $Q_{Y|X}$ where $Y = X + Z$ for $Z \sim q$ independent of X, we get the upper bound

$$\varepsilon^{\otimes k}_{Q_{Y|X}}(\delta_k) \leq k \cdot \mathrm{KL}^{\max}_q + \sqrt{\frac{1}{\delta_k} \cdot k \cdot \mathrm{V}^{\max}_q} \tag{22}$$

$$\leq k \cdot \left(\mathrm{KL}^{\star}(\alpha, C) + \frac{1}{k^{1/4}}\right), \tag{23}$$

where the second inequality follows for all large k. Hence, any $\varepsilon > \mathrm{KL}^{\star}(\alpha, C)$ is an achievable privacy rate for (α, C'). In sum, we have shown that for any $0 < C < C'$,

$$\varepsilon^{\star}(\alpha, C') \leq \mathrm{KL}^{\star}(\alpha, C). \tag{24}$$

Note that this is not yet the desired inequality, since there is a mismatch in the cost bounds C and C'.

The last ingredient is the continuity of the function $C \mapsto \mathrm{KL}^{\star}(\alpha, C)$ over $(0, \infty)$ (Proposition 10 in the extended version [22]). Using this continuity, we infer from (24) that

$$\varepsilon^{\star}(\alpha, C') \leq \lim_{C \nearrow C'} \mathrm{KL}^{\star}(\alpha, C) = \mathrm{KL}^{\star}(\alpha, C'). \tag{25}$$

As $0 < C < C'$ were arbitrary, the proof of (17) is complete.

Finally, we note that the constructed $Q_{Y|X}$ can be made to satisfy the last statement in the theorem. Starting from $C' > 0$ and $\varepsilon > \varepsilon^{\star}(\alpha, C')$, pick small enough $\eta > 0$ so that $\varepsilon(\alpha, C') \leq \varepsilon^{\star}(\alpha, C' - \eta) + \eta < \varepsilon$; this is possible since inequality (17) combined with Theorem 6 imply $\varepsilon^{\star}(\alpha, \cdot) = \mathrm{KL}^{\star}(\alpha, \cdot)$, which is continuous by [22, Proposition 10]. Then, applying the above construction of p and q for $C = C' - \eta$ yields that q achieves the privacy rate ε in view of $\mathrm{KL}^{\star}(\alpha, C) = \varepsilon^{\star}(\alpha, C)$. This completes the proof of the theorem. □

IV. Concluding Remarks

We prove a new fundamental limit for universal approximate differential privacy guarantees under composition. First, we define the universal composition curve via its universal property as the best privacy guarantee under k-fold composition if one has access to only the sensitivity of the query. Then, we introduce the concept of privacy capacity via its operational definition as the best possible ε DP parameter per composition among all mechanisms satisfying a prescribed utility constraint and for which the other DP parameter δ converges to 0. Our main result is deriving the corresponding information definition of privacy capacity, showing that it is equal to a minimax KL-divergence term. Our result extends previous results by showing the explicit relationship between DP and the KL-divergence without restriction on the underlying privacy mechanisms.

We note that our proofs can be extended to the case of adaptive composition, where the same privacy capacity formula holds. In addition, our results also hold for utility functions beyond $\|\cdot\|^\alpha$. Due to space limitations, we relegate this discussion to the extended version [22].

It would be interesting to extend our results to the composition of distinct mechanisms. Another future line of work is analyzing the finite-composition regime, e.g., what is the fastest rate of decay of δ_k subject to ε_k achieving the privacy capacity?

[3] Indeed, $q(Mz) = \int_{\mathbb{R}^m} p(v)\psi_\sigma(Mz - v)dv = \int_{\mathbb{R}^m} p(M^T v)\psi_\sigma(z - M^T v)dv = \int_{\mathbb{R}^m} p(v)\psi_\sigma(z - v)dv = q(z)$ for all orthogonal matrices M.

REFERENCES

[1] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Proc. Theory of Cryptography (TCC)*, Berlin, Heidelberg, 2006, pp. 265–284.

[2] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, "Our data, ourselves: Privacy via distributed noise generation," in *EUROCRYPT*, S. Vaudenay, Ed., 2006, pp. 486–503.

[3] Ú. Erlingsson, V. Pihur, and A. Korolova, "Rappor: Randomized aggregatable privacy-preserving ordinal response," in *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*. ACM, 2014, pp. 1054–1067.

[4] Differential privacy team Apple, "Learning with privacy at scale," 2017.

[5] D. Kifer, S. Messing, A. Roth, A. Thakurta, and D. Zhang, "Guidelines for implementing and auditing differentially private systems," *ArXiv*, vol. abs/2002.04049, 2020.

[6] 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*, 2016, pp. 308–318.

[7] J. Abowd, R. Ashmead, R. Cumings-Menon, S. Garfinkel, M. Heineck, C. Heiss, R. Johns, D. Kifer, P. Leclerc, A. Machanavajjhala, B. Moran, W. Sexton, M. Spence, and P. Zhuravlev, "The 2020 Census Disclosure Avoidance System TopDown Algorithm," *Harvard Data Science Review*, no. Special Issue 2, jun 24 2022, https://hdsr.mitpress.mit.edu/pub/7evz361i.

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

[9] J. Murtagh and S. Vadhan, "The complexity of computing the optimal composition of differential privacy," in *Proc. Int. Conf. Theory of Cryptography*, 2016, pp. 157–175.

[10] S. Asoodeh, J. Liao, F. P. Calmon, O. Kosut, and L. Sankar, "Three variants of differential privacy: Lossless conversion and applications," *IEEE Journal on Selected Areas in Information Theory*, vol. 2, no. 1, pp. 208–222, 2021.

[11] P. Kairouz, S. Oh, and P. Viswanath, "The composition theorem for differential privacy," in *Proceedings of the 32nd International Conference on Machine Learning*, F. Bach and D. Blei, Eds., vol. 37, 2015, pp. 1376–1385.

[12] S. Meiser and E. Mohammadi, "Tight on budget? tight bounds for r-fold approximate differential privacy," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '18, 2018, p. 247–264.

[13] J. Dong, A. Roth, and W. J. Su, "Gaussian Differential Privacy," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 84, no. 1, pp. 3–37, 02 2022. [Online]. Available: https://doi.org/10.1111/rssb.12454

[14] A. Koskela, J. Jälkö, and A. Honkela, "Computing tight differential privacy guarantees using fft," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020, pp. 2560–2569.

[15] S. Gopi, Y. T. Lee, and L. Wutschitz, "Numerical composition of differential privacy," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

[16] A. Koskela, J. Jälkö, L. Prediger, and A. Honkela, "Tight differential privacy for discrete-valued mechanisms and for the subsampled gaussian mechanism using fft," in *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, A. Banerjee and K. Fukumizu, Eds., vol. 130. PMLR, 13–15 Apr 2021, pp. 3358–3366. [Online]. Available: https://proceedings.mlr.press/v130/koskela21a.html

[17] W. Alghamdi, J. F. Gomez, S. Asoodeh, F. Calmon, O. Kosut, and L. Sankar, "The saddle-point method in differential privacy," in *Proceedings of the 40th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, Eds., vol. 202. PMLR, 23–29 Jul 2023, pp. 508–528. [Online]. Available: https://proceedings.mlr.press/v202/alghamdi23a.html

[18] D. M. Sommer, S. Meiser, and E. Mohammadi, "Privacy loss classes: The central limit theorem in differential privacy," *Proceedings on Privacy Enhancing Technologies*, no. 2, pp. 245–269, 2019.

[19] W. Alghamdi, S. Asoodeh, F. P. Calmon, O. Kosut, L. Sankar, and F. Wei, "Cactus mechanisms: Optimal differential privacy mechanisms in the large-composition regime," in *2022 IEEE International Symposium on Information Theory (ISIT)*, 2022, pp. 1838–1843.

[20] W. Alghamdi, S. Asoodeh, F. P. Calmon, J. Felipe Gomez, O. Kosut, and L. Sankar, "Optimal multidimensional differentially private mechanisms in the large-composition regime," in *2023 IEEE International Symposium on Information Theory (ISIT)*, 2023, pp. 2195–2200.

[21] ——, "Schrödinger mechanisms: Optimal differential privacy mechanisms for small sensitivity," in *2023 IEEE International Symposium on Information Theory (ISIT)*, 2023, pp. 2201–2206.

[22] W. Alghamdi, S. Asoodeh, F. P. Calmon, O. Kosut, and L. Sankar, "Differential-privacy capacity," 2024. [Online]. Available: https://github.com/WaelAlghamdi/Privacy-Capacity