Composition Theorems for Interactive Differential Privacy

Anonymous Author(s)
Affiliation
Address
email

Abstract

An interactive mechanism is an algorithm that stores a data set and answers adaptively chosen queries to it. The mechanism is called differentially private, if any adversary cannot distinguish whether a specific individual is in the data set by interacting with the mechanism. We study composition properties of differential privacy in concurrent compositions. In this setting, an adversary interacts with k interactive mechanisms in parallel and can interleave its queries to the mechanisms arbitrarily. Previously, Vadhan and Wang [2021] proved an optimal concurrent composition theorem for pure-differential privacy. We significantly generalize and extend their results. Namely, we prove optimal parallel composition properties for several major notions of differential privacy in the literature, including approximate DP, Rényi DP, and zero-concentrated DP. Our results demonstrate that the adversary gains no advantage by interleaving its queries to independently running mechanisms. Hence, interactivity is a feature that differential privacy grants us for free.

1 Introduction

By now, differential privacy [Dwork et al., 2016] has been widely accepted as a standard framework for protecting individual privacy when performing data analysis on data sets that may contain sensitive information of individuals (see, e.g., the surveys by Dwork and Roth [2014], Vadhan [2017]).

Let $\mathcal{M}$ be an algorithm that runs on a data set x and calculates some information about it. Informally speaking, the algorithm $\mathcal{M}$ is called differentially private, if the output distribution of $\mathcal{A}$ remains nearly identical when we arbitrarily modify a single entry in x.

One essential feature of differential privacy is its *composability*. Composition captures the scenario where a data analyst runs k differentially private algorithms sequentially, and releases the results afterward. Typically, a composition theorem has the following form: if each of the k algorithms satisfies differential privacy, then the analyst's output is still differentially private with moderately degraded privacy parameters.

Composition theorems are important for at least two reasons. First, we might want to perform computation tasks on the same data set multiple times and still have reasonable control over the privacy loss. In this case, composition theorems reveal how the privacy guarantee degrades over time. More importantly, composition theorems allow us to build complex and versatile differentially-private algorithms from simple primitives, and argue the privacy guarantee of the combined algorithm in a straightforward way.

There is a rich literature concerning the composition property of differential privacy (a partial list includes Dwork et al. [2006, 2010], Kairouz et al. [2015], Murtagh and Vadhan [2018], Bassily et al. [2021]). However, most existing composition theorems only consider the scenario where an

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

analyst runs several private algorithms *sequentially*. Namely, the analyst will only move on to the next algorithm after finishing the computation with the previous one. However, many fundamental primitives in differential privacy are interactive in nature, such as the sparse vector technique [Dwork et al., 2009, Roth and Roughgarden, 2010] and private multiplicative weight updates [Hardt and Rothblum, 2010]. Hence, the interactivity issue appears to be a significant limitation of current composition theorems. Namely, the data analyst may want to communicate with several interactive mechanisms *concurrently*, and interleave its queries to the mechanisms arbitrarily. A sequential composition theorem completely fails to capture this scenario. In practice, the deployments of DP algorithms often demand a better understanding of concurrent compositions of interactive mechanisms [Hay et al., 2020].

Recently, Vadhan and Wang [2021] initiated a study of concurrent compositions and proved an optimal concurrent composition theorem for pure differential privacy. In this work, we significantly advance this research direction by proving optimal concurrent composition theorems for several popular notions of differential privacy, including (ε, δ)-approximate DP, Rényi DP, zero-concentrated DP and truncated concentrated DP.

1.1 Setup

Before we continue, we set up necessary pieces of notation. We use $\mathcal{X}$ and $\mathcal{Y}$ to denote the domain of query messages and responses, respectively. We assume that both $\mathcal{X}$ and $\mathcal{Y}$ are finite sets. This assumption is for easing some mathematical manipulation and is not restrictive: all practical applications of differential privacy have finite input and output spaces anyway.

For a set S, denote by $\Delta(S)$ the set of all possible distributions supported over S. We define interactive systems below.

Definition 1 (Interactive system)**.** An interactive system is a (randomized) algorithm $\mathcal{M}\colon (\mathcal{X} \times \mathcal{Y})^* \times \mathcal{X} \to \Delta(\mathcal{Y})$. The input to $\mathcal{M}$ is an interaction history $(x_1, y_1), (x_2, y_2), \ldots, (x_t, y_t) \in (\mathcal{X} \times \mathcal{Y})^t$ together with a query x_{t+1}. The output of $\mathcal{M}$ is denoted by $y_{t+1} \sim \mathcal{M}((x_i, y_i)_{i \in [t]}, x_{t+1})$.

A technicality worth mentioning is that due to the internal memory and randomness of an interactive system $\mathcal{M}$, the response of $\mathcal{M}$ to the $(t+1)$-th query might be correlated with its responses to previous queries. Although the internal randomness of $\mathcal{M}$ is not explicitly stated as a parameter, Definition 1 captures this correlation by requiring that each query x_{t+1} to $\mathcal{M}$ is attached with the interaction history $(x_1, y_1), \ldots, (x_t, y_t)$. This history is sufficient for determining the conditional distribution of the response $\mathcal{M}((x_i, y_i)_{i \in [t]}, x_{t+1})$ without specifying the internal randomness and memory.

We make a distinction between mechanisms and systems. By "mechanism" we mean a differentially private algorithm $\mathcal{M}$ that holds a sensitive input d and answers queries about it. When applied to a concrete input d, $\mathcal{M}$ induces an interactive system, denoted by $\mathcal{M}^d$. According to the definition of differential privacy, studying the privacy of a mechanism boils down to studying the pair of systems $(\mathcal{M}^d, \mathcal{M}^{d'})$ induced by running $\mathcal{M}$ on every pair of neighboring inputs (d, d'). For brevity, we usually assume W.L.O.G. that the input only consists of a single bit $b \in \{0, 1\}$, and we compare the two systems $\mathcal{M}^0, \mathcal{M}^1$ induced by $\mathcal{M}$.

Concurrent composition. We define concurrent composition of interactive systems, first defined and studied by Vadhan and Wang [2021]. Suppose $\mathcal{M}_1, \mathcal{M}_2, \ldots, \mathcal{M}_k$ are k systems. The concurrent composition of them is an interactive system $\mathrm{COMP}(\mathcal{M}_1 \ldots \mathcal{M}_k)$ with query domain $[k] \times \mathcal{X}$ and response domain $\mathcal{Y}$. An adversary is a (possibly randomized) query algorithm $\mathcal{A} : ([k] \times \mathcal{X} \times \mathcal{Y})^* \to \Delta([k] \times \mathcal{X})$. The interaction between $\mathcal{A}$ and $\mathrm{COMP}(\mathcal{M}_i)$ is a stochastic process that proceeds as follows. $\mathcal{A}$ first computes a pair $(i_1, x_1) \in [k] \times \mathcal{X}$, sends a query x_1 to $\mathcal{M}_{i_1}$ and gets the response y_1. In the t-th step, $\mathcal{A}$ calculates the next pair (i_t, x_t) based on the history, sends the t-th query x_t to $\mathcal{M}_{i_t}$ and receives y_t. There is no communication or interaction between the interactive systems. Each system $\mathcal{M}_i$ can only see its own interaction with $\mathcal{A}$. Let $\mathbf{IT}(\mathcal{A} : \mathcal{M}_1, \ldots, \mathcal{M}_k)$ denote the random variable recording the transcript of the interaction.

In the special case $k = 1$, there is only one system $\mathcal{M}$ and the adversary is interacting with it. We define approximate differential privacy for interactive mechanisms in this case.

Definition 2 (Indistinguishability and (ε, δ)-approximate DP)**.** Two interactive systems $\mathcal{M}^0, \mathcal{M}^1$ are called (ε, δ)-indistinguishable, if for every $b \in \{0, 1\}$, every adversary $\mathcal{A}$ and every collection of

transcripts $S \subseteq \{(x_i, y_i)_{i\in[T]}\}$, it holds that

$$\Pr[\mathbf{IT}(\mathcal{A} : \mathcal{M}^b) \in S] \leq e^{\varepsilon} \Pr[\mathbf{IT}(\mathcal{A} : \mathcal{M}^{1-b}) \in S] + \delta. \tag{1}$$

Let $\mathcal{M}$ be an interactive mechanism. $\mathcal{M}$ is called (α, ε)-approximate differentially private (or DP for short), if for every two neighboring data sets d and d', the systems $\mathcal{M}^d$ and $\mathcal{M}^{d'}$ are (ε, δ)-indistinguishable.

1.2 Differential Privacy in Concurrent Compositions

We study the privacy guarantee under concurrent compositions. Let $\mathcal{M}_1^b, \ldots, \mathcal{M}_k^b$ be k interactive mechanisms, each satisfying (ε, δ)-approximate DP. Consider their concurrent composition $\mathrm{COMP}(\mathcal{M}_1^b \ldots \mathcal{M}_k^b)$. We want to find out the smallest parameters ε', δ' such that $\mathrm{COMP}(\mathcal{M}_i^b)$ satisfies (ε', δ')-approximate DP. In the sequential composition, the adversary $\mathcal{A}$ interacts with $\mathcal{M}_i$'s in order and cannot interleave its queries. In this case, it is known by the advanced composition theorem [Dwork et al., 2010] that $\mathbf{IT}(\mathcal{A} : \mathcal{M}_1^0, \ldots, \mathcal{M}_k^0)$ and $\mathbf{IT}(\mathcal{A} : \mathcal{M}_1^0, \ldots, \mathcal{M}_k^0)$ are $(O(\sqrt{k \log(1/\delta')}\varepsilon), k\delta + \delta')$-indistinguishable.

However, in general, the adversary can interleave its queries arbitrarily, and the differential privacy guarantee warranted by $\mathrm{COMP}(\mathcal{M}_i)$ is less clear. Vadhan and Wang [2021] were the first to formally study this question. They showed that in the special case $\delta = 0$, an optimal composition holds for $\mathrm{COMP}(\mathcal{M}_i)$. That is, if we can prove an (ε', δ') upper bound on the privacy parameter for sequential compositions of $\mathcal{M}_1^b, \ldots, \mathcal{M}_i^b$, then the concurrent composition $\mathrm{COMP}(\mathcal{M}_i)$ also enjoys the same (ε', δ')-approximate DP.

Vadhan and Wang [2021] also considered the case $\delta > 0$ (i.e., approximate DP). However, for this case, they only showed an upper bound on ε', δ' that is even inferior to the basic composition in the sequential setting. It was asked as an open question in Vadhan and Wang [2021] whether the optimal composition theorem for approximate DP still holds in the concurrent composition.

Besides pure and approximate DP, there are also other notions of differential privacy that are extensively studied in the literature. A non-exhaustive list includes Rényi DP Mironov [2017], concentrated DP Dwork and Rothblum [2016], Bun and Steinke [2016], Bun et al. [2018], Gaussian DP and f-DP Dong et al. [2022] etc. Compared with the standard notion of (ε, δ)-approximate DP, these variants of DP either allow for a simplified analysis of private algorithms or give sharper bounds of privacy guarantee. In the sequential composition, the composition property of these variants has been well understood. It would also be interesting to also extend these composition theorems to the concurrent composition, thereby expanding the potential applicability of these DP notions.

2 Our Results

In this work, we give an affirmative answer to the open question mentioned above. Moreover, our result confirms that several major differential privacy definitions in the literature enjoy the same composition guarantee in the concurrent composition, just as they do in the sequential composition.

Approximate differential privacy. (ε, δ)-approximate differential privacy is arguably the most widely studied notion of differential privacy and is deemed the "standard" definition of DP. As our first main result, we show an optimal concurrent composition theorem for approximate DP.

Theorem 1. *Let $\mathcal{M}_1, \ldots, \mathcal{M}_k$ be k interactive mechanisms that run on the same data set. Suppose that each mechanism $\mathcal{M}_i$ satisfies $(\varepsilon_i, \delta_i)$-approximate DP. Then $\mathrm{COMP}(\mathcal{M}_1 \ldots \mathcal{M}_k)$ is (ε', δ')-approximate DP, where ε', δ' are given by the optimal (sequential) composition theorem [Kairouz et al., 2015, Murtagh and Vadhan, 2018].*

In particular, when the privacy parameter is the same (ε, δ) for each mechanism. Their concurrent composition satisfies $O(\sqrt{k \log(1/\delta')}\varepsilon, \delta' + k\delta)$-approximate DP for all $\delta' \in (0, 1)$.

Rényi differential privacy. Rényi differential privacy was defined by Mironov [2017]. We first recall its definition.

Definition 3 (Rényi divergence and differential privacy)**.** Let P, Q be two distributions supported over $\mathcal{X}$. For each $\alpha > 1$, define the Rényi divergence of order α of P from Q as

$$D_\alpha(P\|Q) := \frac{1}{\alpha - 1} \log \left(\mathop{\mathbb{E}}_{x \sim P} \left[\left(\frac{P(x)}{Q(x)} \right)^{\alpha - 1} \right] \right).$$

Two interactive systems are called (α, ε)-Rényi close, if for every adversary $\mathcal{A}$ and every $b \in \{0, 1\}$, it holds that

$$D_\alpha(\mathbf{IT}(\mathcal{A} : \mathcal{M}^b) \| \mathbf{IT}(\mathcal{A} : \mathcal{M}^{1-b})) \leq \varepsilon.$$

Let $\mathcal{M}$ be a mechanism. $\mathcal{M}$ is called (α, ε)-Rényi differentially private, if for every two neighboring data sets d and d', the systems $\mathcal{M}^d$ and $\mathcal{M}^{d'}$ are (α, ε)-Rényi close.

A main advantage of Rényi DP is that it has a natural and simple composition. In the sequential setting, it is known that if two mechanisms $\mathcal{M}_1, \mathcal{M}_2$ are (α, ε_1) and (α, ε_2)-Rényi DP, respectively, then the composition of $\mathcal{M}_1$ and $\mathcal{M}_2$ is $(\alpha, \varepsilon_1 + \varepsilon_2)$-Rényi DP. Our next theorem generalizes this result to the concurrent composition.

Theorem 2. *Let $\mathcal{M}_1, \ldots, \mathcal{M}_k$ be k interactive mechanisms that run on the same data set. Suppose that each mechanism $\mathcal{M}_i$ is (α, ε_i)-Rényi DP. Then $\mathrm{COMP}(\mathcal{M}_1 \ldots \mathcal{M}_k)$ is $(\alpha, \sum_{i=1}^n \varepsilon_i)$-Rényi DP.*

One implication of Theorem 2 is that the zero-concentrated differential privacy by Bun and Steinke [2016] and the truncated concentrated differential privacy by Bun et al. [2018] also composes nicely under concurrent composition. Since this is a direct consequence of Theorem 2 we do not give the formal definitions here and refer interested readers to Bun and Steinke [2016], Bun et al. [2018] for more backgrounds. We only state the result as a corollary below.

Corollary 1. *Let $\mathcal{M}_1, \ldots, \mathcal{M}_k$ are k interactive mechanisms that run on the same data set. Suppose that each mechanism $\mathcal{M}_i$ is (η_i, ξ_i)-zCDP (resp. (ρ_i, ω)-tCDP). Then $\mathrm{COMP}(\mathcal{M}_1 \ldots \mathcal{M}_k)$ is $(\sum_i \eta_i, \sum_i \xi_i)$-zCDP (resp. $(\sum_i \rho_i, \omega)$-tCDP).*

Theorem 1, 2 and Corollary 1 provide compelling evidence that the adversary gains no advantage by interleaving its queries to independently running mechanisms. Consequently, interactivity can be viewed as a feature that differential privacy grants us for free.

3 Proof of Main Results

In this section, we show the proof of our results. We start with a very brief proof overview. We prove Theorem 1 by a reduction to the sequential composition of k (approximate) random response mechanisms. This generalizes the idea developed in Vadhan and Wang [2021]. To prove Theorem 2, we take a completely different approach, and our technique offers new tools to analyze Rényi DP. Namely, we propose a new characterization of Rényi divergence (Lemma 4), which allows for a fine-grained account of the privacy loss in the complex interaction involving multiple mechanisms. The new characterization of Rényi divergence might be of independent interest.

Notation. Let P, Q be two distributions supported over X. For a real $\eta > 0$, we say that $P \geq \eta Q$, if for every $S \subseteq X$, it holds that

$$\Pr_{x \sim P}[x \in S] \geq \eta \Pr_{x \sim Q}[x \in S].$$

Furthermore, we say $P \equiv Q$, if P and Q are identically distributed.

3.1 Approximate Differential Privacy

To prove Theorem 1, we follow the approach by Vadhan and Wang [2021], where they showed that one can simulate two $(\varepsilon, 0)$-indistinguishable interactive systems by post-processing a random response mechanism. This simulation enables them to reduce the concurrent composition to a sequential composition, and the optimal composition theorem follows. It was asked as an open question in Vadhan and Wang [2021] whether the same simulation can be carried out for approximate DP. We answer this question affirmatively.

Review of the Vadhan-Wang approach. It would be instructive to review the proof in Vadhan and Wang [2021] first. Let $\mathcal{M}^0, \mathcal{M}^1$ be the two systems that run on the real data set and a neighbouring

date set, respectively. The adversary $\mathcal{A}$ interacts with $\mathcal{M}^b$ for some $b \in \{0,1\}$ and wants to find out the value of b. An intuitive yet delicate fact due to Vadhan and Wang [2021] is that, if $\mathcal{M}^0$ and $\mathcal{M}^1$ are $(\varepsilon, 0)$-indistinguishable, then there exist two systems $\mathcal{N}^0, \mathcal{N}^1$ such that, for every adversary $\mathcal{A}$, the distribution of $\mathbf{IT}(\mathcal{A} : \mathcal{M}^b)$ is identical to $\frac{e^\varepsilon}{1+\varepsilon}\mathbf{IT}(\mathcal{N}^b : \mathcal{A}) + \frac{1}{1+e^\varepsilon}\mathbf{IT}(\mathcal{A} : \mathcal{N}^{1-b})$. This enables one to simulate the many-round interaction between $\mathcal{A}$ and $\mathcal{M}^b$ by running a one-round randomized response mechanism.

In more detail, let $\mathrm{RR}_\varepsilon^b$ denote the standard randomized response mechanism, defined as follows. $\mathrm{RR}_\varepsilon^b$ only accepts one query. On the query, $\mathrm{RR}_\varepsilon^b$ ignores the query message, returns b with probability $\frac{e^\varepsilon}{1+e^\varepsilon}$, and returns $1-b$ otherwise. We modify $\mathcal{A}$ to a new adversary $\mathcal{A}'$: $\mathcal{A}'$ first sends a query to $\mathrm{RR}_\varepsilon^b$ and receives a bit b'. Then $\mathcal{A}'$ simulates the interaction between $\mathcal{A}$ and $\mathcal{N}^{b'}$, and outputs the transcript (i.e., $\mathbf{IT}(\mathcal{A} : \mathcal{N}^{b'})$). Let $\mathbf{Output}(\mathcal{A}' : \mathrm{RR}_\varepsilon^b)$ denote the output distribution of $\mathcal{A}'$ when interacting with $\mathrm{RR}_\varepsilon^b$. It is clear that

$$\mathbf{Output}(\mathcal{A}' : \mathrm{RR}_\varepsilon^b) \equiv \frac{e^\varepsilon}{1+\varepsilon}\mathbf{IT}(\mathcal{A} : \mathcal{N}^b) + \frac{1}{1+e^\varepsilon}\mathbf{IT}(\mathcal{A} : \mathcal{N}^{1-b}) \equiv \mathbf{IT}(\mathcal{A} : \mathcal{M}^b).$$

Therefore, $\mathcal{A}'$ simulates the interaction between $\mathcal{A}, \mathcal{M}^b$ faithfully, by a single query to RR_ε.

Now, suppose the adversary $\mathcal{A}$ is interacting with k mechanisms $\mathcal{M}_1^b, \ldots, \mathcal{M}_k^b$ in parallel. For each $i \in [k]$, assuming that $\mathcal{M}_i^0$ and $\mathcal{M}_i^1$ are $(\varepsilon_i, 0)$-indistinguishable, there is a decomposition of $\mathcal{M}_i^0, \mathcal{M}_i^1$ by some $\mathcal{N}_i^0$ and $\mathcal{N}_i^1$. Again, we modify $\mathcal{A}$ to a new mechanism $\mathcal{A}'$. $\mathcal{A}'$ first queries $\mathrm{RR}_{\varepsilon_i}^b, i \in [k]$ in order, and receives k bits $b_1', \ldots, b_k'$. Then $\mathcal{A}'$ simulates the interaction between $\mathcal{A}$ and $(\mathcal{N}_i^{b_i'})_{i\in[k]}$ and outputs the transcript. One can show that

$$\mathbf{Output}(\mathcal{A}' : \mathrm{RR}_{\varepsilon_1}^b, \ldots, \mathrm{RR}_{\varepsilon_k}^b) \equiv \mathbf{IT}(\mathcal{A} : \mathcal{M}_1^b, \ldots, \mathcal{M}_k^b). \tag{2}$$

Note that the left hand side of (2) can be simulated by the sequential composition of k randomized response mechanisms. Invoking the optimal composition theorem for sequential composition [Kairouz et al., 2015, Murtagh and Vadhan, 2018] concludes the proof.

Extension to approximate DP. Now, if $\mathcal{M}^0$ and $\mathcal{M}^1$ are (ε, δ)-distinguishable with $\delta > 0$, there might not be a nice decomposition of $\mathcal{M}^b$ into $\frac{e^\varepsilon}{1+e^\varepsilon}\mathcal{N}^b + \frac{1}{1+e^\varepsilon}\mathcal{N}^{1-b}$. Still, it is plausible to conjecture that there is a decomposition of $\mathcal{M}^0, \mathcal{M}^1$ with four systems $\mathcal{N}^0, \mathcal{N}^1, \mathcal{E}^0, \mathcal{E}^1$ such that for each $b \in \{0,1\}$,

$$\mathcal{M}^b = \delta\mathcal{E}^b + (1-\delta)\left(\frac{e^\varepsilon}{1+e^\varepsilon}\mathcal{N}^b + \frac{1}{1+e^\varepsilon}\mathcal{N}^{1-b}\right). \tag{3}$$

Our main technical result in this subsection proves the existence of such a decomposition.

Theorem 3. *Two systems $\mathcal{M}^0, \mathcal{M}^1$ are (ε, δ)-indistinguishable, if and only if there are four systems $\mathcal{N}^0, \mathcal{N}^1, \mathcal{E}^0, \mathcal{E}^1$ satisfying the following: for every adversary $\mathcal{A}$ and $b \in \{0,1\}$, it holds that*

$$\mathbf{IT}(\mathcal{A} : \mathcal{M}^b) \equiv \delta\mathbf{IT}(\mathcal{A} : \mathcal{E}^b) + (1-\delta)\left(\frac{e^\varepsilon}{1+\varepsilon}\mathbf{IT}(\mathcal{A} : \mathcal{N}^b) + \frac{1}{1+e^\varepsilon}\mathbf{IT}(\mathcal{A} : \mathcal{N}^{1-b})\right). \tag{4}$$

Theorem 3 implies Theorem 1 by a similar reduction to (approximate) random response. For completeness, we include a proof in Appendix.

We prove Theorem 3 by establishing a series of lemmas. In the following, we state these lemmas and explain their intuition. We defer the formal proof to Appendix.

Lemma 1. *Suppose $\mathcal{M}^0, \mathcal{M}^1$ are (ε, δ)-indistinguishable. There are two mechanism $\mathcal{E}^0, \mathcal{E}^1$ satisfying the following.*

- *For every adversary $\mathcal{A}$ and $b \in \{0,1\}$, it holds that $\mathbf{IT}(\mathcal{A} : \mathcal{M}^b) \geq \delta \cdot \mathbf{IT}(\mathcal{A} : \mathcal{E}^b)$.*
- *For every adversary $\mathcal{A}$, every set of transcripts $S \subseteq \{(x_i, y_i)_{i\in[T]}\}$ and $b \in \{0,1\}$, it holds that*

$$\begin{aligned}&\Pr[\mathbf{IT}(\mathcal{A} : \mathcal{M}^b) \in S] - \delta\Pr[\mathbf{IT}(\mathcal{A} : \mathcal{E}^b) \in S]\\ &\leq e^\varepsilon\left(\Pr[\mathbf{IT}(\mathcal{A} : \mathcal{M}^{1-b})] - \delta\Pr[\mathbf{IT}(\mathcal{A} : \mathcal{E}^{1-b})]\right).\end{aligned}$$

Roughly, Lemma 1 says that there are two systems $\mathcal{E}^0, \mathcal{E}^1$ that capture the low-probability "bad behavior" of $\mathcal{M}^0, \mathcal{M}^1$.

To prove Lemma 1, we first argue that for each adversary $\mathcal{A}$, there are systems $\mathcal{E}^0_{\mathcal{A}}, \mathcal{E}^1_{\mathcal{A}}$ (depending on $\mathcal{A}$) that capture the bad behavior in the interaction between $\mathcal{M}^{0/1}$ and $\mathcal{A}$. To construct $\mathcal{E}^0, \mathcal{E}^1$ that are independent of $\mathcal{A}$. We use duality of linear programming. That is, we formulate a zero-sum finite game, where the attacker chooses an adversary strategy $\mathcal{A}$, and the defender chooses two systems $\mathcal{E}^0, \mathcal{E}^1$. The outcome of the game measures how well $\mathcal{E}^0, \mathcal{E}^1$ capture the "bad behavior" in the interaction between $\mathcal{A}$ and $\mathcal{M}^{0/1}$. For this game, we have shown that for every strategy of the attacker, the defender always has a good response. We can then use the minimax theorem (i.e., strong duality) to show there are two systems $\mathcal{E}^0, \mathcal{E}^1$ that perform well on *every* adversary $\mathcal{A}$.

Lemma 2. *Suppose $\mathcal{M}, \mathcal{E}$ are two systems such that for every adversary $\mathcal{A}$, it holds that $\mathbf{IT}(\mathcal{A} : \mathcal{M}) \geq \delta \mathbf{IT}(\mathcal{A} : \mathcal{E})$. Then there is a system $\mathcal{N}$ such that for every adversary $\mathcal{A}$, it holds that*

$$\mathbf{IT}(\mathcal{A} : \mathcal{M}) \equiv \delta \mathbf{IT}(\mathcal{A} : \mathcal{E}) + (1 - \delta)\mathbf{IT}(\mathcal{A} : \mathcal{N}).$$

For intuition, suppose P, Q are two distributions such that $P \geq \delta Q$, then we can easily find a distribution Q' such that $P \equiv \delta Q + (1 - \delta)Q'$. The proof of Lemma 2 generalizes this simple idea.

The following lemma is proven by Vadhan and Wang [2021].

Lemma 3. *Suppose $\mathcal{N}^0, \mathcal{N}^1$ are $(\varepsilon, 0)$-indistinguishable. Then there are two systems $\mathcal{N}^{0'}, \mathcal{N}^{1'}$ such that for every adversary $\mathcal{A}$, it holds that*

$$\mathbf{IT}(\mathcal{A} : \mathcal{N}^b) \equiv \frac{e^\varepsilon}{1 + e^\varepsilon}\mathbf{IT}(\mathcal{A} : \mathcal{N}^{b'}) + \frac{1}{1 + e^\varepsilon}\mathbf{IT}(\mathcal{A} : \mathcal{N}^{(1-b)'}).$$

Wrap-up. We can sketch the proof for Theorem 3 now. First, the existence of a decomposition satisfying (4) implies that $\mathcal{M}^0, \mathcal{M}^1$ are (ε, δ)-indistinguishable. For the other direction, we start by constructing $\mathcal{E}^0, \mathcal{E}^1$ using Lemma 1. Then we construct $\mathcal{N}^0, \mathcal{N}^1$ by Lemma 2. Lemma 1 and 2 together ensure that $\mathcal{N}^0$ and $\mathcal{N}^1$ are $(\varepsilon, 0)$-indistinguishable, which enables us to invoke Lemma 3 and decompose $\mathcal{N}^0, \mathcal{N}^1$ into $\mathcal{N}^{0'}, \mathcal{N}^{1'}$. $(\mathcal{E}^0, \mathcal{E}^1, \mathcal{N}^{0'}, \mathcal{N}^{1'})$ forms the final decomposition. It is straightforward to verify that they satisfy (4).

3.2 Rényi Differential Privacy

Our result for Rényi differential privacy (Theorem 2) takes a completely different approach.

An intuition. Let $\mathcal{M}_1$ be an (α, ε)-Rényi DP mechanism. Intuitively, (α, ε)-Rényi DP means that $\mathcal{M}_1$ has ε unit of privacy budget and can distribute it to T queries. Viewing the privacy budget as a form of "deposit", we hope to argue that two or more independently running mechanisms spend their deposit independently, and an adversary cannot trigger any mechanism to spend more privacy budget than it holds by interacting with other mechanisms. This is rather intuitive: one cannot withdraw more money from a bank than they have deposited, even if they are using services from other banks.

However, unlike the intuitive and easy-to-measure resources such as time and energy, the notion of privacy loss looks somewhat illusive. Even worse, we need to reason about this weird resource in a stochastic process consisting of interactions with multiple systems. It was not clear how one can quantify the privacy loss in such an interactive and complex process. Our proof for Theorem 2 presents a new approach to do so.

A new characterization for Rényi DP. We introduce a new characterization of Rényi divergence based on Hölder's inequality and duality. That is, we prove

Lemma 4 (An alternative characterization of Rényi divergence)**.** *Suppose P, Q are two distributions supported over $\mathcal{Y}$. For every $\alpha > 1$ and $B \geq 0$, let $\beta = \frac{\alpha}{\alpha - 1}$ be the Hölder conjugate of α. The following statements are equivalent.*

- $D_\alpha(P \| Q) \leq B$.
- *For every function $h : \mathcal{Y} \to \mathbb{R}^{\geq 0}$, it holds that $\mathbb{E}_{y \sim P}[h(y)] \leq e^{\frac{B(\alpha-1)}{\alpha}} \mathbb{E}_{y \sim Q}[h(y)^\beta]^{1/\beta}$.*

We do not know if Lemma 4 has been known/used in the statistics and/or math community. As far as we are aware, we are the first to use Lemma 4 in the context of differential privacy. Note that if we let $\alpha \to \infty$, then Lemma 4 converges to a characterization of pure-DP. That is, $D_\infty(P\|Q) \le B$ if and only if $\Pr[P = y] \le e^B \Pr[Q = y]$ for every $y \in \mathcal{Y}$[1].

Lemma 4 provides a convenient tool to reason about the privacy loss in an interactive environment consisting of multiple rounds. Intuitively, this is because Condition 2 in the statement above is more amenable to a "hybrid argument". However, to quantify the privacy loss during an interaction, we still need to find a way to track the privacy loss.

Measure theory setup. Before we continue, it would be more convenient to switch to a measure-theoretic language. Consider two measures P, Q on a space $\mathcal{Y}$ (P and Q are not necessarily probability measure), we say that P is β-dominated by Q, denoted by $P \preceq_\beta Q$, if for every measurable function $f : \mathcal{Y} \to \mathbb{R}^{\ge 0}$, it holds that

$$\|f\|_{P,1} := \int f(y)dP(y) \ge \left(\int f(y)^\beta dQ(y)\right)^{1/\beta} =: \|f\|_{Q,\beta}.$$

When $\mathcal{Y}$ is a finite set, the integral coincides with an equivalent summation. i.e.,

$$\int f(y)dP(y) = \sum_y P(y)f(y).$$

We will use integral and summation interchangeably.

In this notation, Lemma 4 can be equivalently stated as $D_\alpha(P\|Q) \le B$ if and only if P is β-dominated by $e^B Q$ for $\beta = \frac{\alpha}{\alpha-1}$.

The following lemma is essential for us.

Lemma 5. *Let $\mathcal{Y}_1 \times \mathcal{Y}_2$ be a space. Consider two distributions P, Q on $\mathcal{Y}_1 \times \mathcal{Y}_2$. Assume* $\mathrm{supp}(P) = \mathrm{supp}(Q) = \mathcal{Y}_1 \times \mathcal{Y}_2$. *Let P_1, P_2 be the margin of P on $\mathcal{Y}_1, \mathcal{Y}_2$. For each $y_1 \in \mathcal{Y}_1$, denote by $P_2|_{P_1=y_1}$ the marginal distribution of y_2 conditioning on y_1. Also define the same notation for Q.*

Let $\beta \ge 1, B \ge 0$ be two reals. Let $\alpha = \frac{\beta}{\beta-1}$. For each $y_1 \in \mathcal{Y}_1$, define

$$\ell_1(y_1) = \inf_K \{K : P_2|_{P_1=y_1} \preceq_\beta K \cdot Q_2|_{Q_1=y_1}\} = \exp(D_\alpha(P_2|_{P_1=y_1} \| Q_2|_{Q_1=y_1})).$$

Suppose $P \preceq_\beta e^B Q$. Consider the measure spaces $(\mathcal{Y}_1, P_1(y_1) \cdot \ell_1(y_2)^{1/\beta})$ and $(\mathcal{Y}_1, Q_2)$. We have

$$P_1 \ell_1^{1/\beta} \preceq_\beta e^B Q_1.$$

Intuitively, the function $\ell(y_1)$ serves as the role of "privacy budget monitor". To see this, fix an adversary $\mathcal{A}$ and think of (y_1, y_2) as the responses of the system to the adversary[2]. After observing y_1, the adversary wants to distinguish between two conditional distributions $P_2|_{P_1=y_1}$ and $Q_2|_{Q_1=y_1}$. At this moment, $\ell(y_1)$ shows up as an upper bound of "extra information" that the adversary can extract by utilizing their second query. Alternatively, $\ell(y_1)$ quantifies the amount of the remaining privacy budget the mechanism has after outputting y_1. On average, the function $\ell_1(y_1)$ provides a fine-grained control of privacy loss in the sense that $P_1 \ell_1^{1/\beta} \preceq_\beta e^B Q_1$.

Proof for a 3-round toy example. We are ready to describe the proof for Theorem 2. To illustrate the idea, we prove a toy case here and defer the full proof to Appendix. The proof for the toy case includes all the important ideas. Extending it to a full proof is straightforward.

We describe the toy scenario now. Suppose there are two mechanisms $\mathcal{M}_1, \mathcal{M}_2$ that run on a sensitive input bit $b \in \{0, 1\}$. The interaction consists of 3 rounds. The adversary $\mathcal{A}$ communicates with $\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_1$ in order, and outputs the response (y_1, y_2, y_3). For brevity, we assume that

[1] In fact, this is what inspired us to devise the characterization.

[2] Although the query made by $\mathcal{A}$ is not explicitly recorded, the pair (y_1, y_2) can capture this information by requiring that each response y_i must be attached with the query message x_i. This does not leak any additional information because x_i is solely chosen by $\mathcal{A}$.

each response y_i contains a copy of the query message x_i, so that we recover the whole transcript $((x_1, y_1), (x_2, y_2), (x_3, y_3))$ only from the responses.

Let $P, Q \in \Delta(\mathcal{Y} \times \mathcal{Y} \times \mathcal{Y})$ be the output distribution when $\mathcal{A}$ interacts with $\mathcal{M}_1^b$ and $\mathcal{M}_2^b$. respectively. Suppose $\mathcal{M}_1, \mathcal{M}_2$ are (α, ε_1), (α, ε_2)-Rényi DP respectively. Our goal is to prove that

$$\max(D_\alpha(P\|Q), D_\alpha(Q\|P)) \le \varepsilon_1 + \varepsilon_2.$$

We bound $D_\alpha(P\|Q)$ below. The bound for $D_\alpha(P\|Q)$ is symmetric. Be Lemma 4, it suffices to show that for every $h : \mathcal{Y} \times \mathcal{Y} \to \mathcal{Y} \to \mathbb{R}$, it holds that

$$\sum_{y=(y_1,y_2,y_3)} P(y)h(y) \le \left(e^{\varepsilon_1+\varepsilon_2} \sum_{y=(y_1,y_2,y_3)} Q(y)h(y)^\beta \right)^{1/\beta} \tag{5}$$

where $\beta = \frac{\alpha}{\alpha-1}$ is the Hölder conjugate of α. Let P_1, P_2, P_3 be the projection of P onto the three rounds, and let $P_i|_{y_{<i}}$ denote the distribution of y_i conditioning on $y_1, \ldots, y_{i-1}$. Also define the same notation for Q. Then we write

$$\sum_{y=(y_1,y_2,y_3)} P(y)h(y) = \sum_{y_1} \left(P_1(y_1) \sum_{y_2} \left(P_2|_{y_1}(y_2) \sum_{y_3} P_3|_{y_{<3}}(y_3)h(y) \right) \right) \tag{6}$$

For every $y_1 \in \mathcal{Y}$, let $\mathcal{M}_1^0|_{y_1}$ (resp. $\mathcal{M}_1^1|_{y_1}$) denote the interactive system $\mathcal{M}_1^0$ (resp. $\mathcal{M}_1^1$) *conditioning on* that it has answered y_1 to the first query. Recall we have assumed that y_i contains the query message x_1. Formally, for every $(x_2, y_2), \ldots, (x_t, y_t)$ and x_{t+1} and $b \in \{0, 1\}$ define

$$\mathcal{M}_1^b|_{y_1}((x_j, y_j)_{2\le j\le t}, x_{t+1}) \equiv \mathcal{M}_1^b((x_j, y_j)_{1\le j\le t}, x_{t+1}).$$

Next, define

$$\ell_1(y_1) := \exp\left(\sup_{A:\text{adversary}} \left\{ D_\alpha\big(\mathbf{IT}(A : \mathcal{M}_1^0|_{y_1}) \| \mathbf{IT}(A : \mathcal{M}_1^1|_{y_1})\big) \right\} \right). \tag{7}$$

From Lemma 5, one can show that $P_1 \ell_1^{1/\beta} \preceq e^B Q_1$. Turning back to (5), we then have

$$\begin{aligned}
&\sum_{y_1} \left(P_1(y_1) \sum_{y_2} \left(P_2|_{y_1}(y_2) \sum_{y_3} P_3|_{y_{<3}}(y_3)h(y) \right) \right) && (8)\\
&\le \sum_{y_1} \left(P_1(y_1) \sum_{y_2} \left(P_2|_{y_1}(y_2) \left(\ell_1(y_1) \sum_{y_3} Q_3|_{y_{<3}}(y_3)h(y)^\beta \right)^{1/\beta} \right) \right) && (9)\\
&\le \sum_{y_1} \left(P_1(y_1) \left(e^{\varepsilon_2} \sum_{y_2} \left(Q_2|_{y_1}(y_2)\ell_1(y_1) \sum_{y_3} Q_3|_{y_{<3}}(y_3)h(y)^\beta \right) \right)^{1/\beta} \right) && (10)\\
&= \sum_{y_1} \left(P_1(y_1)\ell_1(y_1)^{1/\beta} \left(e^{\varepsilon_2} \sum_{y_2} \left(Q_2|_{y_1}(y_2) \sum_{y_3} Q_3|_{y_{<3}}(y_3)h(y)^\beta \right) \right)^{1/\beta} \right) && (11)\\
&\le \left(e^{\varepsilon_1+\varepsilon_2} \sum_{y_1} \left(Q_1(y_1) \sum_{y_2} \left(Q_2|_{y_1}(y_2) \sum_{y_3} Q_3|_{y_{<3}}(y_3)h(y)^\beta \right) \right) \right)^{1/\beta} && (12)\\
&= \left(e^{\varepsilon_1+\varepsilon_2} \sum_{y} Q(y)h(y)^\beta \right)^{1/\beta}. && (13)
\end{aligned}$$

We justify the deductions briefly. (8) $\Rightarrow$ (9) is the most critical step. To see this, observe that knowing y_2 does not change the view of the first mechanism, because the second query is sent to the independently running mechanism $\mathcal{M}_2^b$. Therefore, $\mathcal{M}_1^b|_{y_1}$ remains the same after conditioning on y_1 and y_2. Now, note that $P_3|_{y_{<3}}$ (resp. $Q_3|_{y_{<3}}$) exactly describes one round of interaction between an adversary and $\mathcal{M}_1^0|_{y_1}$ (resp. $\mathcal{M}_1^1|_{y_1}$). Consequently, the information leaked by y_3 must be subject

to the bound (7) and the inequality holds. Next, (9) $\Rightarrow$ (10) can be justified by identifying the term followed in $P_2|_{y_1}(y_2)$ in (9) as an "h" function and using the definition of β-dominance (recall that $\mathcal{M}_2^b$ satisfies (α, ε_2)-DP, which implies that $P_2|_{y_1} \preceq_\beta e^{\varepsilon_2} Q_2|_{y_1}$). (11) $\Rightarrow$ (12) is similar.

Having justified (13) for every measure function h, we conclude that $D_\alpha(P\|Q) \leq e^{\varepsilon_1+\varepsilon_2}$. A symmetric argument shows that $D_\alpha(Q\|P) \leq e^{\varepsilon_1+\varepsilon_2}$. This completes the proof for the toy example.

Proof sketch for the general case. The proof for the general case extends the idea above with some minor twists. First off, by induction, we only need to prove the composition theorem for the case with two mechanisms and many rounds. An issue worth noting is that $\mathcal{A}$ can choose the next query object based on previous responses. However, we can suppose without loss of generality that $\mathcal{A}$ always communicates with mechanisms alternately, by adding a vanilla query x^* to the query space. If the current mechanism is not the one $\mathcal{A}$ wishes to speak with, $\mathcal{A}$ just sends the vanilla query x^*. The mechanism then returns a fixed response, which does not leak any information.

Now, suppose $\mathcal{A}$ communicates with each mechanism for T rounds, with the total number of rounds being $2T$. Let $(y_1, \ldots, y_T), (z_1, \ldots, Z_T)$ denote the responses from two mechanisms, respectively. The order in which $\mathcal{A}$ sees the response are $y_1, z_1, \ldots, y_T, z_T$. Let P (resp. Q) denote the distribution of the responses when interacting with $(\mathcal{M}_1^0, \mathcal{M}_2^0)$ (resp. $(\mathcal{M}_1^1, \mathcal{M}_2^1)$). It suffices to show for every measurable function h that

$$\sum_{y,z} P(yz)h(yz) \leq \left(e^{\varepsilon_1+\varepsilon_2} \sum_{y,z} Q(yz)h(yz)^\beta \right)^{1/\beta}. \tag{14}$$

To show (14), one can define privacy loss function for each round $i \in [T]$ and do a similar argument as we one we did in (8) $\Rightarrow$ (13).

4 Conclusion and Future Directions

In this work, we consider the concurrent compositions of interactive mechanisms. Regarding the general privacy guarantee under the concurrent composition, our result gives an optimal answer for several popular definitions of differential privacy, including (ε, δ)-approximate DP and Rényi DP. Our work is purely theoretical, and we do not see any negative societal impacts it may cause.

For future directions, we ask whether one can use our composition theorems to design new differentially-private algorithms that may involve running several differentially-private mechanisms in parallel. It is also interesting to simplify/unify privacy proofs for existing algorithms using the new composition theorems. Another outstanding question is to prove or refute the optimal composition theorem for f-DP [Dong et al., 2022]. f-DP is argued to be the most "informative" notion of differential privacy. A concurrent composition theorem (or a refutation) will help us better understand the role of interactivity in differential privacy.

We also note that there is a recent interest in *fully adaptive* compositions of differential privacy, which studies how the data analyst can manage the privacy budget and monitor the privacy loss themselves. In particular, the notion of privacy odometers and filters were proposed to capture these demands Rogers et al. [2016], Feldman and Zrnic [2021], Whitehouse et al. [2022], Lécuyer [2021]. This question necessitates a better understanding of information leakage in an interactive environment. This paper developed several new tools and techniques to reason about interactive mechanisms. Can our technique be useful in studying fully adaptive compositions?

References

Raef Bassily, Kobbi Nissim, Adam D. Smith, Thomas Steinke, Uri Stemmer, and Jonathan R. Ullman. Algorithmic stability for adaptive data analysis. *SIAM J. Comput.*, 50(3), 2021. doi: 10.1137/16M1103646. URL `https://doi.org/10.1137/16M1103646`.

Mark Bun and Thomas Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In Martin Hirt and Adam D. Smith, editors, *Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part I*, volume 9985 of *Lecture Notes in Computer Science*, pages 635–658, 2016. doi: 10.1007/978-3-662-53641-4_24. URL `https://doi.org/10.1007/978-3-662-53641-4_24`.

Mark Bun, Cynthia Dwork, Guy N. Rothblum, and Thomas Steinke. Composable and versatile privacy via truncated CDP. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 74–86. ACM, 2018. doi: 10.1145/3188745.3188946. URL https://doi.org/10.1145/3188745.3188946.

Jinshuo Dong, Aaron Roth, and Weijie J. Su. Gaussian differential privacy. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 84, 2022.

Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4):211–407, 2014. doi: 10.1561/0400000042. URL https://doi.org/10.1561/0400000042.

Cynthia Dwork and Guy N. Rothblum. Concentrated differential privacy. *CoRR*, abs/1603.01887, 2016. URL http://arxiv.org/abs/1603.01887.

Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In Serge Vaudenay, editor, *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*, volume 4004 of *Lecture Notes in Computer Science*, pages 486–503. Springer, 2006. doi: 10.1007/11761679_29. URL https://doi.org/10.1007/11761679_29.

Cynthia Dwork, Moni Naor, Omer Reingold, Guy N. Rothblum, and Salil P. Vadhan. On the complexity of differentially private data release: efficient algorithms and hardness results. In Michael Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 381–390. ACM, 2009. doi: 10.1145/1536414.1536467. URL https://doi.org/10.1145/1536414.1536467.

Cynthia Dwork, Guy N. Rothblum, and Salil P. Vadhan. Boosting and differential privacy. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 51–60. IEEE Computer Society, 2010. doi: 10.1109/FOCS.2010.12. URL https://doi.org/10.1109/FOCS.2010.12.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. *J. Priv. Confidentiality*, 7(3):17–51, 2016. doi: 10.29012/jpc.v7i3.405. URL https://doi.org/10.29012/jpc.v7i3.405.

Vitaly Feldman and Tijana Zrnic. Individual privacy accounting via a rényi filter. In Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 28080–28091, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/ec7f346604f518906d35ef0492709f78-Abstract.html.

Moritz Hardt and Guy N. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 61–70. IEEE Computer Society, 2010. doi: 10.1109/FOCS.2010.85. URL https://doi.org/10.1109/FOCS.2010.85.

Michael Hay, Marco Gaboardi, and Salil Vadhan. A programming framework for opendp, 2020. URL https://projects.iq.harvard.edu/files/opendp/files/opendp_programming_framework_11may2020_1_01.pdf.

Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. In Francis R. Bach and David M. Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 1376–1385. JMLR.org, 2015. URL http://proceedings.mlr.press/v37/kairouz15.html.

Mathias Lécuyer. Practical privacy filters and odometers with rényi differential privacy and applications to differentially private deep learning. *CoRR*, abs/2103.01379, 2021. URL https://arxiv.org/abs/2103.01379.

Ilya Mironov. Rényi differential privacy. In *30th IEEE Computer Security Foundations Symposium, CSF 2017, Santa Barbara, CA, USA, August 21-25, 2017*, pages 263–275. IEEE Computer Society, 2017. doi: 10.1109/CSF.2017.11. URL https://doi.org/10.1109/CSF.2017.11.

Jack Murtagh and Salil P. Vadhan. The complexity of computing the optimal composition of differential privacy. *Theory Comput.*, 14(1):1–35, 2018. doi: 10.4086/toc.2018.v014a008. URL https://doi.org/10.4086/toc.2018.v014a008.

Ryan M. Rogers, Salil P. Vadhan, Aaron Roth, and Jonathan R. Ullman. Privacy odometers and filters: Pay-as-you-go composition. In Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 1921–1929, 2016. URL https://proceedings.neurips.cc/paper/2016/hash/58c54802a9fb9526cd0923353a34a7ae-Abstract.html.

Aaron Roth and Tim Roughgarden. Interactive privacy via the median mechanism. In Leonard J. Schulman, editor, *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 765–774. ACM, 2010. doi: 10.1145/1806689.1806794. URL https://doi.org/10.1145/1806689.1806794.

Salil P. Vadhan. The complexity of differential privacy. In Yehuda Lindell, editor, *Tutorials on the Foundations of Cryptography*, pages 347–450. Springer International Publishing, 2017. doi: 10.1007/978-3-319-57048-8_7. URL https://doi.org/10.1007/978-3-319-57048-8_7.

Salil P. Vadhan and Tianhao Wang. Concurrent composition of differential privacy. In Kobbi Nissim and Brent Waters, editors, *Theory of Cryptography - 19th International Conference, TCC 2021, Raleigh, NC, USA, November 8-11, 2021, Proceedings, Part II*, volume 13043 of *Lecture Notes in Computer Science*, pages 582–604. Springer, 2021. doi: 10.1007/978-3-030-90453-1_20. URL https://doi.org/10.1007/978-3-030-90453-1_20.

Justin Whitehouse, Aaditya Ramdas, Ryan Rogers, and Zhiwei Steven Wu. Fully adaptive composition in differential privacy. *CoRR*, abs/2203.05481, 2022. doi: 10.48550/arXiv.2203.05481. URL https://doi.org/10.48550/arXiv.2203.05481.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] We included a discussion for future directions.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes] The full proofs will be available in supplementary material.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]