Composition Theorems for Interactive Differential Privacy

Xin Lyu
Department of Electrical Engineering and Computer Science
University of California, Berkeley
Berkeley, CA, 94720
xinlyu@berkeley.edu

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.

Concurrently and independently of our work, Vadhan and Zhang [2022] proved an optimal concurrent composition theorem for f-DP [Dong et al., 2022], which implies our result for the approximate DP case.

1 Introduction

By now, differential privacy [Dwork et al., 2006b] 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. Roughly speaking, $\mathcal{M}$ 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 more complex and powerful differentially-

36th Conference on Neural Information Processing Systems (NeurIPS 2022).

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 (see, e.g., Dwork et al. [2006a, 2010], Kairouz et al. [2015], Murtagh and Vadhan [2018], Bassily et al. [2021]). However, most existing composition theorems only consider the scenario where an analyst runs several private algorithms *sequentially*. That is, the analyst will only move on to the next algorithm after finishing their computation with the previous one. In contrast, 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. Also, in practice, 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 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 $(x_1,y_1),(x_2,y_2),\dots,(x_t,y_t)\in(\mathcal{X}\times\mathcal{Y})^t$ 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 $(x_1,y_1),\dots,(x_t,y_t)$. 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. Suppose $\mathcal{M}_1,\mathcal{M}_2,\dots,\mathcal{M}_k$ are k systems. The concurrent composition of them is an interactive system $\mathrm{COMP}(\mathcal{M}_1\dots\mathcal{M}_k)$ 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 runs as follows. $\mathcal{A}$ first[1] 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

[1] We assume it is always the adversary who sends the first message. This is without loss of generality: we can let the first message sent from the adversary to each system be an "Initiliazation" query. Having received the initialization query, the system returns either a starting message or simply a "SUCCESS" symbol.

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)-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

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

Let $\mathcal{M}$ be an interactive mechanism. $\mathcal{M}$ is called (ε, δ)-approximate differentially private (or (ε, δ)-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)-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')-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 $\mathbf{IT}(\mathcal{A} : \mathcal{M}_1^1, \ldots, \mathcal{M}_k^1)$ 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 $\mathcal{M}_1^b, \ldots, \mathcal{M}_k^b$, then the concurrent composition $\mathrm{COMP}(\mathcal{M}_i)$ also enjoys the same (ε', δ')-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 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 is also interesting to 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)-DP 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 $(\varepsilon_i, \delta_i)$-DP. Then $\mathrm{COMP}(\mathcal{M}_1 \ldots \mathcal{M}_k)$ is (ε', δ')-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 for each mechanism is the same (ε, δ), their concurrent composition satisfies $O(\sqrt{k \log(1/\delta')}\varepsilon, \delta' + k\delta)$-DP for all $\delta' \in (0, 1)$.

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

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

Let $\mathcal{M}$ be a mechanism. $\mathcal{M}$ is called (α, ε)-Rényi differentially private (or (α, ε)-RDP for short), 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)-RDP, respectively, then the composition of $\mathcal{M}_1$ and $\mathcal{M}_2$ is $(\alpha, \varepsilon_1 + \varepsilon_2)$-RDP. Our next theorem generalizes this result to the concurrent composition setting.

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)-RDP. Then* $\mathrm{COMP}(\mathcal{M}_1 \ldots \mathcal{M}_k)$ *is $(\alpha, \sum_{i=1}^k \varepsilon_i)$-RDP.*

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 compose nicely under the concurrent composition. We state the corollary below, and prove it in Appendix A.3 for completeness.

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-zCDP (resp. (ρ_i, ω)-tCDP). Then* $\mathrm{COMP}(\mathcal{M}_1 \ldots \mathcal{M}_k)$ *is $(\sum_i \eta_i)$-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.

Concurrent and independent work. Concurrently to our work, a recent work by Vadhan and Zhang [2022] proves an optimal concurrent composition theorem for f-DP [Dong et al., 2022]. By the standard connection, their result implies the optimal concurrent composition theorem for approximate DP. However, our techniques are very different than theirs. Their result is stronger, as it is known that approxiamte-DP can be seen as a special case of f-DP [Dong et al., 2022]. However, our proof for approximate DP is more elementary: we do not need to work through f-DP as their proof does. Furthermore, our proof comes with several interesting technical ingredients that might be of independent interests. This includes a structural result for interactive mechanisms (Theorem 3), as well as a dual perspective to reason about Rényi divergences (Lemma 4).

3 Implications of Our Results

In this section, we discuss implications of our results, and demonstrate how they offer more than the sequential composition theorems.

Designing new algorithms. The optimal concurrent composition theorem makes it possible to design new differentially private algorithm that involves running several building blocks concurrently. As one motivating example, consider the Sparse Vector Technique (SVT). The standard SVT (as in Dwork and Roth [2014]) and its variants have been studied extensively in the literature. In particular, it was observed by Lyu et al. [2017], Zhu and Wang [2020] that one can add noise to the threshold only *once*, and then use the noisy threshold to answer $c > 1$ "meaningful" queries (namely, after reporting each meaningful query, the SVT algorithm does NOT refresh the noisy threshold). It was argued in [Lyu et al., 2017, Zhu and Wang, 2020] that this variant of SVT can offer a higher accuracy while consuming the same amount of privacy budget, both theoretically and empirically.

However, this variant of SVT has received relatively less attention in literature. One reason might be that it is unclear what happens if we compose this SVT with other mechanisms. In particular, the standard SVT refreshes its threshold after answering each "meaningful" query, which allows one to decompose the algorithm into c pieces of smaller SVT algorithms, and then compose with other mechanisms via the sequential computation. In contrast, the variants by Lyu et al. [2017], Zhu and Wang [2020] work by answering each "meaningful query" using the *same* noisy threshold, which do not seem to admit such a decomposition. This makes this variant less appealing: in most applications, people want to use SVT as a supporting subroutine for other algorithms. Therefore, it is crucial to understand the (concurrent) composition behavior of SVT with other mechanisms.

Now, with the new concurrent composition theorem, we can plug this variant of SVT in any algorithm, and argue the privacy guarantee of the whole computation by black-box applying Theorems 1 and 2 (depending on whether we are working with (ε, δ)-DP or RDP). To illustrate the idea, in Appendix B, we apply Theorem 1 to analyze a simple algorithm: private "Guess-and-Check" with the aforementioned variant of SVT as a subroutine. We hope our example can motivate people to design more powerful algorithms by concurrently composing simple building blocks.

Practical Implication. Besides the theoretical interests, our theorem has implications for practical deployments of interactive DP mechanisms. For one example, suppose there is a data center that holds the private information of individuals and offers data analysts access to the database (interactively and differentially-privately). Without knowing the concurrent composition theorem, it might be possible that some $k > 1$ analysts can collude by coordinating their (interactive) queries to the database and extracting much more sensitive information. Our result refutes the possibility of such an attack. In particular, suppose each data analyst has only an (ε, δ)-DP amount of privacy "quota". Then, even if they collude and spend their privacy budget in whatever way, their computation result is still $(O(\sqrt{k \log(1/\delta')}\varepsilon), k\delta + \delta')$-DP with respect to the private database.

4 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) randomized response mechanisms. This generalizes the idea developed by 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 an alternative 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 characterization of Rényi divergence might find itself useful in other applications.

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.

4.1 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 randomized 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 by Vadhan and Wang [2021] first. Let $\mathcal{M}^0, \mathcal{M}^1$ be the pair of systems by running the private mechanism on a pair of neighboring data sets. 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 $\frac{e^\varepsilon}{1+\varepsilon}\mathbf{IT}(\mathcal{A} : \mathcal{N}^b) + \frac{1}{1+e^\varepsilon}\mathbf{IT}(\mathcal{A} : \mathcal{N}^{1-b})$. 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 $\mathrm{RR}_\varepsilon^b$.

Now, suppose the adversary $\mathcal{A}$ is interacting with k mechanisms $\mathcal{M}_1^b, \dots, \mathcal{M}_k^b$ 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 $b_1', \dots, b_k'$. 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

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

Note that the left hand side of (2) can be simulated by a 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)-indistinguishable 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 A.1.

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 A.1.

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

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

Roughly, Lemma 1 says that there are two systems $\mathcal{E}^0, \mathcal{E}^1$ that capture the low-probability "bad behavior" of $\mathcal{M}^0, \mathcal{M}^1$. It is the primary technical contribution of this subsection. We prove Lemma 1 by explicitly constructing the two systems $\mathcal{E}^0, \mathcal{E}^1$. That is, we specify the probability density functions $\Pr[\mathcal{E}^b((x_j, y_j)_{j<i}, x_i) = y_i]$ for $\mathcal{E}^0, \mathcal{E}^1$ step by step, in the increasing order of $i = 1, 2, \dots, T$.

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 one can easily find a distribution Q' such that $P \equiv \delta Q + (1 - \delta)Q'$. The proof of Lemma 2 extends this simple idea.

Lemma 3 (Vadhan and Wang [2021])**.** *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 conclude the proof for Theorem 3 now. The "if" direction is obvious: 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).

4.2 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.

However, unlike some 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 elusive 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. Nonetheless, we manage to find a new approach to do so.

An alternative characterization for Rényi DP. We introduce the following 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 $\mathbb{E}_{y \sim P}[h(y)] \leq e^{\frac{B(\alpha-1)}{\alpha}} \mathbb{E}_{y \sim Q}[h(y)^\beta]^{1/\beta}$.*

Note that if we let $\alpha \to \infty$, then Lemma 4 converges to a characterization of pure-DP. That is, $D_\infty(P \| Q) \leq B$ if and only if $\Pr[P = y] \leq e^B \Pr[Q = y]$ for every $y \in \mathcal{Y}$.

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 measures), we say that P is β-dominated by Q, denoted by $P \preceq_\beta Q$, if for every measurable function $f : \mathcal{Y} \to \mathbb{R}^{\geq 0}$, it holds that

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

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* $\operatorname{supp}(P) = \operatorname{supp}(Q) = \mathcal{Y}_1 \times \mathcal{Y}_2$. *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 the 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 A.2. 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 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 $(\mathcal{M}_1^0, \mathcal{M}_2^0)$ and $(\mathcal{M}_1^1, \mathcal{M}_2^1)$, 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 $D_\alpha(Q\|P)$ is symmetric. Be Lemma 4, it suffices to show that for every $h : \mathcal{Y} \times \mathcal{Y} \to \mathcal{Y} \to \mathbb{R}^{\ge 0}$ 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

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

[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 $\mathcal{A}$.

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_1 contains x_1). Formally, for every $b \in \{0,1\}$, $(x_2, y_2), \dots, (x_t, y_t)$ and x_{t+1}, define

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

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

$$\sum_{y_1} \left(P_1(y_1) \sum_{y_2} \left(P_2|_{y_1}(y_2) \sum_{y_3} P_3|_{y_{<3}}(y_3) \underline{h(y)} \right) \right) \tag{8}$$

$$\le \sum_{y_1} \left(P_1(y_1) \sum_{y_2} \left(P_2|_{y_1}(y_2) \underline{\left(\ell_1(y_1) \sum_{y_3} Q_3|_{y_{<3}}(y_3) h(y)^\beta \right)^{1/\beta}} \right) \right) \tag{9}$$

$$\le \sum_{y_1} \left(P_1(y_1) \left(e^{\varepsilon_2} \sum_{y_2} \left(Q_2|_{y_1}(y_2) \ell_1(y_1) \sum_{y_3} Q_3|_{y_{<3}}(y_3) h(y)^\beta \right) \right)^{1/\beta} \right) \tag{10}$$

$$= \sum_{y_1} \left(P_1(y_1) \ell_1(y_1)^{1/\beta} \underline{\left(e^{\varepsilon_2} \sum_{y_2} \left(Q_2|_{y_1}(y_2) \sum_{y_3} Q_3|_{y_{<3}}(y_3) h(y)^\beta \right) \right)^{1/\beta}} \right) \tag{11}$$

$$\le \left(e^{\varepsilon_1+\varepsilon_2} \sum_{y_1} \left(Q_1(y_1) \sum_{y_2} \left(Q_2|_{y_1}(y_2) \sum_{y_3} Q_3|_{y_{<3}}(y_3) h(y)^\beta \right) \right) \right)^{1/\beta} \tag{12}$$

$$= \left(e^{\varepsilon_1+\varepsilon_2} \sum_{y} Q(y) h(y)^\beta \right)^{1/\beta}. \tag{13}$$

Here, we used inequalities of the form $\sum_y P(y) \cdot h(y) \le \left(C \cdot \sum_y Q(y) \cdot h(y)^\beta \right)^{1/\beta}$ three times (they are (8) $\Rightarrow$ (9) $\Rightarrow$ (10) and (11) $\Rightarrow$ (12)). We use underlines to highlight the "h" part of each step in the deductions above.

(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 *both* y_1 and y_2. Now, note that $P_3|_{y_{<3}}$ (resp. $Q_3|_{y_{<3}}$) exactly describes one round of interaction between the 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. Having verified (8) $\Rightarrow$ (9), the steps (9) $\Rightarrow$ (10) and (11) $\Rightarrow$ (12) are straightforward.

Having justified (13) for every measure function h, we conclude that $D_\alpha(P\|Q) \le e^{\varepsilon_1+\varepsilon_2}$. A symmetric argument shows that $D_\alpha(Q\|P) \le e^{\varepsilon_1+\varepsilon_2}$. 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. 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. We refer to Appendix A.2 for the detail of the proof.

5 Conclusion and Future Directions

In this work, we consider the concurrent composition of interactive mechanisms. Regarding the general privacy guarantee under the concurrent composition, our result gives optimal composition theorems for several popular definitions of differential privacy, including (ε, δ)-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 explore more practical implications of the concurrent composition phenomena.

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. Our work developed several new tools and techniques to reason about interactive mechanisms. Can our technique be useful in studying fully adaptive compositions?

Acknowledgements

I am grateful to Jelani Nelson for advising this project and providing useful comments on an early draft of this paper. I would also like to thank Salil Vadhan and Wanrong Zhang for insightful discussions about their work.

X. Lyu was supported by ONR DORECG award N00014-17-1-2127.

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, 2006a. doi: 10.1007/11761679_29. URL https://doi.org/10.1007/11761679_29.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2006b. doi: 10.1007/11681878_14. URL https://doi.org/10.1007/11681878_14.

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.

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.

Min Lyu, Dong Su, and Ninghui Li. Understanding the sparse vector technique for differential privacy. *Proc. VLDB Endow.*, 10(6):637–648, 2017. doi: 10.14778/3055330.3055331. URL http://www.vldb.org/pvldb/vol10/p637-lyu.pdf.

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 Vadhan and Wanrong Zhang. Concurrent composition theorems for differential privacy. 2022. doi: 10.48550/ARXIV.2207.08335. URL https://arxiv.org/abs/2207.08335.

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.

Yuqing Zhu and Yu-Xiang Wang. Improving sparse vector technique with renyi differential privacy. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/e9bf14a419d77534105016f5ec122d62-Abstract.html.

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]

A Appendix: Missing Proofs

In this appendix, we show the formal proofs for all the lemmas and claims in the main paper.

A.1 Proofs for Approximate Differential Privacy

This subsection present omitted proofs in Section 4.1.

A.1.1 The key lemma

We start with the proof for Lemma 1.

Reminder of Lemma 1. *Suppose $\mathcal{M}^0, \mathcal{M}^1$ are (ε, δ)-indistinguishable. There are two systems $\mathcal{E}^0, \mathcal{E}^1$ 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*
$$\begin{aligned}&\Pr[\mathbf{IT}(\mathcal{A} : \mathcal{M}^b) \in S] - \delta \Pr[\mathbf{IT}(\mathcal{A} : \mathcal{E}^b) \in S]\\ \leq\ & e^{\varepsilon}\left(\Pr[\mathbf{IT}(\mathcal{A} : \mathcal{M}^{1-b}) \in S] - \delta \Pr[\mathbf{IT}(\mathcal{A} : \mathcal{E}^{1-b}) \in S]\right).\end{aligned}$$

Proof. Without loss of generality, we assume that the interaction between $\mathcal{M}^{0/1}$ and $\mathcal{A}$ consists of exactly $T \in \mathbb{N}$ rounds. For every $t \in [T]$, each $(x_i, y_i)_{i\in[t]}$ and $b \in \{0,1\}$, denote

$$M^b((y_i)_{i\in[t]}, (x_i)_{i\in[t]}) := \prod_{i=1}^{t} \Pr[\mathcal{M}^b((x_j, y_j)_{j<i}, x_i) = y_i]. \tag{14}$$

Intuitively, $M^b((y_i)_{i\in[t]}, (x_i)_{i\in[t]})$ is the probability of $\mathcal{M}^b$ responding $(y_1, \ldots, y_t)$, conditioning on that the query messages are fixed to $(x_1, \ldots, x_t)$. Note that knowing $M^b((y_i)_{i\in[T]}, (x_i)_{i\in[T]})$ for every $(x_i, y_i)_{i\in[T]}$ *uniquely* determines the system.

Let $\mathcal{A}$ be an arbitrary adversary. For each $(x_i, y_i)_{i\in[t-1]}$ and x_t, denote

$$A((x_i)_{i\in[t]}, (y_i)_{i\in[t-1]}) := \prod_{i=1}^{t} \Pr[\mathcal{A}((x_j, y_j)_{j<i}) = x_i]. \tag{15}$$

Note that $A((x_i)_{i\in[t]}, (y_i)_{i\in[t-1]})$ is the probability of $\mathcal{A}$ sending queries $(x_1, \ldots, x_t)$, conditioning on that the responses to the first $t-1$ queries are fixed to $(y_1, \ldots, y_{t-1})$.

In the following, when the size of a list $(\ell_i)_{i\in[L]}$ is clear from context, we may omit the subscript and simply write (ℓ_i) to denote the list. Now, having defined (14) and (15), we observe for each transcript $(x_i, y_i)_{i\in[T]}$ that

$$\Pr[\mathbf{IT}(\mathcal{A}, \mathcal{M}^b) = (x_i, y_i)_{i\in[T]}] = M^b((y_i), (x_i)) \cdot A((x_i), (y_i)). \tag{16}$$

In the following, we will construct two systems $\mathcal{E}^{0/1}$ such that, for every $(x_i, y_i)_{i\in[T]}$, it holds that

$$M^b((y_i), (x_i)) \geq \delta E^b((y_i), (x_i)) \tag{17}$$

and

$$M^b((y_i),(x_i)) - \delta E^b((y_i),(x_i)) \le e^{\varepsilon}\big(M^{1-b}((y_i),(x_i)) - \delta E^{1-b}((y_i),(x_i))\big). \tag{18}$$

If we have two systems $\mathcal{E}^{0/1}$ satisfying the above, then we can verify that they satisfy the lemma statement by combining (16), (17) and (18).

Now we describe the construction. We start by defining for each $b \in \{0,1\}$, $t \le T$ and every partial history $(x_i, y_i)_{i\le t-1} \in (\mathcal{X}\times\mathcal{Y})^{t-1}$, $x_t \in \mathcal{X}$ a control function as

$$\mathsf{Lower}^b((x_i,y_i)_{i<t}, x_t) := \begin{cases} \sum_{y_t\in\mathcal{Y}} \max_{x_{t+1}}\{\mathsf{Lower}^b((x_i,y_i)_{i\le t}, x_{t+1})\} & t<T \\ \sum_{y_t\in\mathcal{Y}} \max\left\{M^b((y_i),(x_i)) - e^{\varepsilon}M^{1-b}((y_i),(x_i)), 0\right\} & t=T \end{cases}. \tag{19}$$

For every $t \le T-1$ and $(x_i, y_i)_{i\le t}$, we also define the following control function:

$$\mathsf{Upper}^b((x_i,y_i)_{i\le t}) := M^b((y_i)_{i\le t},(x_i)_{i\le t}) - e^{-\varepsilon}M^{1-b}((y_i)_{i\le t},(x_i)_{i\le t}). \tag{20}$$

We need the following two facts regarding the control functions.

Claim 1. *For each $b \in \{0,1\}$ and $x_1 \in \mathcal{X}$, it holds that $\mathsf{Lower}^b(\emptyset, x_1) \le \delta$.*

Proof. We construct an adversary $\mathcal{A}^*$ as follows. $\mathcal{A}^*$ is deterministic. It always sends x_1 as the first query. For every $1 \le t \le T-1$ and history $(x_i,y_i)_{i\in[t]}$, $\mathcal{A}^*$ computes the next query as

$$\mathcal{A}^*((x_i,y_i)_{i\in[t]}) = \arg\max_{x_{t+1}}\{\mathsf{Lower}^b((x_i,y_i)_{i\le t}, x_{t+1})\}.$$

Now, define

$$S^b := \{(x_i,y_i)_{i\in[T]} : \Pr[\mathbf{IT}(\mathcal{A}^*,\mathcal{M}^b) = (x_i,y_i)_{i\in[T]} > e^{\varepsilon}\Pr[\mathbf{IT}(\mathcal{A}^*,\mathcal{M}^{1-b}) = (x_i,y_i)_{i\in[T]}]\}.$$

Given that $\mathcal{M}^0$ and $\mathcal{M}^1$ are (ε,δ)-indistinguishable, we know that

$$\sum_{(x_i,y_i)\in S^b} \Pr[\mathbf{IT}(\mathcal{A}^*,\mathcal{M}^b) = (x_i,y_i)_{i\in[T]}] - e^{\varepsilon}\Pr[\mathbf{IT}(\mathcal{A}^*,\mathcal{M}^{1-b}) = (x_i,y_i)_{i\in[T]}] \le \delta.$$

On the other hand, by the definition of $\mathcal{A}^*$ and (19), it holds that

$$\mathsf{Lower}^b(\emptyset, x_1) = \sum_{(x_i,y_i)\in S^b} \Pr[\mathbf{IT}(\mathcal{A}^*,\mathcal{M}^b) = (x_i,y_i)] - e^{\varepsilon}\Pr[\mathbf{IT}(\mathcal{A}^*,\mathcal{M}^{1-b}) = (x_i,y_i)].$$

This can be verified by tracing how $\mathsf{Lower}^b(\emptyset, x_1)$ is determined from queries $(x_1,\ldots,x_T)$ (in the "max" operator), and noting that $\mathcal{A}^*$ follows exactly the same queries. Combining two equations above concludes the proof of Claim. □

Claim 2. *For every $b \in \{0,1\}$, $t \le T-1$ and $(x_i,y_i)_{i\le t-1} \in (\mathcal{X}\times\mathcal{Y})^{t-1}, x_t \in \mathcal{X}$, it holds that*

$$\mathsf{Lower}^b((x_i,y_i)_{i<t}, x_t) \le \mathsf{Upper}^b((x_i,y_i)_{i<t}) + e^{-\varepsilon}\mathsf{Lower}^{1-b}((x_i,y_i)_{i<t}, x_t).$$

Proof. We prove this claim by *downward* induction on t. For the case $t = T-1$, we have by definition that

$$\begin{aligned} & e^{-\varepsilon}\mathsf{Lower}^{1-b}((x_i,y_i)_{i<t}, x_t) \\ &= \sum_{y_t\in\mathcal{Y}} \max\left\{e^{-\varepsilon}M^{1-b}((y_i),(x_i)) - M^b((y_i),(x_i)), 0\right\} && (21) \\ &= \sum_{y_t\in\mathcal{Y}} e^{-\varepsilon}M^{1-b}((y_i),(x_i)) - M^b((y_i),(x_i)) + \\ &\quad \sum_{y_t\in\mathcal{Y}} \max\left\{M^b((y_i),(x_i)) - e^{-\varepsilon}M^{1-b}((y_i),(x_i)), 0\right\} && (22) \\ &\ge -\mathsf{Upper}^b((x_i,y_i)_{i<t}) + \mathsf{Lower}^b((x_i,y_i)_{i<t}, x_t). && (23) \end{aligned}$$

We justify the deductions briefly. (21) is by definition. (22) uses a simple trick that $\max\{a,0\} = a + \max\{-a,0\}$. The last step (23) is by definition again. In particular, we observe that for every $x_T \in \mathcal{X}$, it holds that

$$\sum_{y_T\in\mathcal{Y}} M^b((y_i),(x_i)) = \prod_{i=1}^{T-1} \Pr[\mathcal{M}^b((x_j,y_j)_{j<i},x_i) = y_i].$$

This proves the base case for $t = T-1$.

Assume the claim is true for $t+1 \le T-1$. We consider the case of t. We have

$$\begin{aligned}
\mathsf{Lower}^b((x_i,y_i)_{i<t},x_t) &= \sum_{y_t} \max_{x_{t+1}}\{\mathsf{Lower}^b((x_i,y_i)_{i\le t},x_{t+1})\} \\
&\le \sum_{y_t} \max_{x_{t+1}}\{\mathsf{Upper}^b((x_i,y_i)_{i\le t}) + e^{-\varepsilon}\mathsf{Lower}^{1-b}((x_i,y_i)_{i\le t},x_{t+1})\} \\
&\le \sum_{y_t} \mathsf{Upper}^b((x_i,y_i)_{i\le t}) + e^{-\varepsilon}\max_{x_{t+1}}\{\mathsf{Lower}^{1-b}((x_i,y_i)_{i\le t},x_{t+1})\} \\
&= \mathsf{Upper}^b((x_i,y_i)_{i<t}) + e^{-\varepsilon}\mathsf{Lower}^{1-b}((x_i,y_i)_{i<t},x_t).
\end{aligned}$$

The first inequality is due to the induction hypothesis. The second inequality is straightforward. This completes the proof for the claim. □

The construction. We are ready to describe the construction. In the following, we will assume $\varepsilon > 0$. Having shown the construction for every $\varepsilon > 0$, the case for $\varepsilon = 0$ can be argued by continuity. We will construct $\mathcal{E}^0, \mathcal{E}^1$ by specifying for every $t \in [T]$ and $(x_i,y_i)_{i\le t}$ the following:

$$E^b((y_i)_{i\le t},(x_i)_{i\le t}) := \prod_{i=1}^{t} \Pr[\mathcal{E}^b((x_j,y_j)_{j<i},x_i) = y_i].$$

Note that a valid $E^b(\cdot)$ uniquely defines a system $\mathcal{E}^b$. For brevity, we also define $E^0(\emptyset) = E^1(\emptyset) = 1$. Intuitively, we use $\emptyset$ to denote two "empty lists" (i.e., two lists $(y_i)_{i\le t},(x_i)_{i\le t}$ with $t = 0$).

We will construct $E^b((y_i)_{i\le t},(x_i)_{i\le t})$ for $t = 1,2,\ldots,T$ in order. Throughput the construction, we maintain the following property. For every $0 \le t \le T$, $(x_i,y_i)_{i\le t}$ and $b \in \{0,1\}$, we require

$$\delta\cdot E^b((y_i)_{i\le t},(x_i)_{i\le t}) \ge \begin{cases} \max_{x_{t+1}\in\mathcal{X}}\{\mathsf{Lower}^b((x_j,y_j)_{j\le t},x_{t+1})\} & t<T \\ \max\left\{M^b((y_i),(x_i)) - e^{\varepsilon}M^{1-b}((y_i),(x_i)),0\right\} & t=T \end{cases} \tag{24}$$

and

$$\delta\cdot E^b((y_i)_{i\le t},(x_i)_{i\le t}) \le \mathsf{Upper}^b((x_i,y_i)_{i\le t}) + e^{-\varepsilon}\delta\cdot E^{1-b}((y_i)_{i\le t},(x_i)_{i\le t}). \tag{25}$$

Meanwhile, for $E^b((y_i),(x_i))$ to describe a valid system, it is necessary and sufficient for it to be non-negative and satisfy the following equation for every $(x_i,y_i)_{i\le t} \in (\mathcal{X}\times\mathcal{Y})^t$ and x_{t+1}:

$$\sum_{y_{t+1}\in\mathcal{Y}} E^b((y_i)_{i\le t+1},(x_i)_{i\le t+1}) = E^b((y_i)_{i\le t},(x_i)_{i\le t}). \tag{26}$$

Next, we shall prove that we can construct a valid $E^{0/1}$ satisfying (24), (25) and (26). As we have said, we will construct E^b gradually in the increasing order of $t \in [T]$. For $t = 0$, we have set $E^0(\emptyset) = E^1(\emptyset) = 1$. (24) holds by Claim 1, and (25) holds trivially.

Now let $t < T$. Also let $(y_i)_{i\le t} \in \mathcal{Y}^t, (x_i)_{i\le t} \in \mathcal{X}^t$ be two lists. Suppose we have constructed $E^{0/1}((y_i)_{i\le t},(x_i)_{i\le t})$ that satisfies (24) and (25). For every $x_{t+1} \in \mathcal{X}$ and $y_{t+1} \in \mathcal{Y}$, we construct $E^{0/1}((y_i)_{i\le t+1},(x_i)_{i\le t+1})$ in the following.

Fix $x_{t+1} \in \mathcal{X}$. We temporarily set

$$\widetilde{E}^b((y_i)_{i\le t+1},(x_i)_{i\le t+1}) = \frac{1}{\delta}\cdot\begin{cases} \max_{x_{t+2}\in\mathcal{X}}\{\mathsf{Lower}^b((x_j,y_j)_{j\le t+1},x_{t+2})\} & t+1<T \\ \max\left\{M^b((y_i),(x_i)) - e^{\varepsilon}M^{1-b}((y_i),(x_i)),0\right\} & t+1=T \end{cases}.$$

By Claim 2, we know that $\widetilde{E}^b$ satisfies (25). By the construction, $\widetilde{E}^b$ satisfies (24). However, $\widetilde{E}^b$ may fail to satisfy (26). Still, we have

$$\sum_{y_{t+1}} \widetilde{E}^b((y_i)_{i\le t+1},(x_i)_{i\le t+1}) \le \frac{1}{\delta}\mathsf{Lower}^b((x_i,y_i)_{i\le t},x_{t+1}) \le E^b((y_i)_{i\le t},(x_i)_{i\le t}).$$

In the following, we show that one can adjust $\widetilde{E}^b$ by increasing some $\widetilde{E}^b((y_i)_{i\le t+1},(x_i)_{i\le t+1})$ properly, so that the new $\widetilde{E}^b$ satisfies all of (24), (25) and (26).

To begin with, we define for each $b\in\{0,1\}$ the quantity

$$\mathrm{Gap}_b := E^b((y_i)_{i\le t},(x_i)_{i\le t}) - \sum_{y_{t+1}} \widetilde{E}^b((y_i)_{i\le t+1},(x_i)_{i\le t+1}). \tag{27}$$

Our goal is to decrease $\mathrm{Gap}_0,\mathrm{Gap}_1$ to zero by increasing $\widetilde{E}$. Since we only increase $\widetilde{E}^b((y_i)_{i\le t+1},(x_i)_{i\le t+1})$, (24) can never be compromised and we only need to consider (25). Consider y_{t+1} and $b\in\{0,1\}$. We say that $\widetilde{E}^b$ is *tight* at y_{t+1}, if

$$\delta\widetilde{E}^b((y_i)_{i\le t+1},(x_i)_{i\le t+1}) = \mathsf{Upper}^b((x_i,y_i)_{i\le t+1}) + e^{-\varepsilon}\delta\widetilde{E}^{1-b}((y_i)_{i\le t+1},(x_i)_{i\le t+1}).$$

Intuitively, $\widetilde{E}^b$ being tight at y_{t+1} means that we cannot increase $\widetilde{E}^b((y_i)_{i\le t+1},(x_i)_{i\le t+1})$ without increasing $\widetilde{E}^{1-b}((y_i)_{i\le t+1},(x_i)_{i\le t+1})$.

Here shows our adjustment strategy. We consider each $y_{t+1}\in\mathcal{Y}$ in an arbitrary but fixed order. For each y_{t+1}, we gradually increase $\widetilde{E}^{0/1}((y_i)_{i\le t+1},(x_i)_{i\le t+1})$ until one of the following events happens.

- Both $\widetilde{E}^0$ and $\widetilde{E}^1$ get tight at y_{t+1}.
- $\mathrm{Gap}_0 = 0$, and $\widetilde{E}^1$ is tight at y_{t+1}.
- $\mathrm{Gap}_1 = 0$, and $\widetilde{E}^0$ is tight at y_{t+1}.
- $\mathrm{Gap}_0 = \mathrm{Gap}_1 = 0$.

It is easy to see that if none of the above happens, we can keep increasing $\widetilde{E}^{0/1}$ at y_{t+1}[3]. This completes the description of the adjustment strategy.

Now, we claim that after the adjustment, we must have $\mathrm{Gap}_0 = \mathrm{Gap}_1 = 0$. Suppose it is not the case. For example, suppose $\mathrm{Gap}_0 \neq 0$. Then we know that $\widetilde{E}^0$ is tight at every y_{t+1}. This means that

$$\sum_{y_{t+1}} \delta\widetilde{E}^0((y_i)_{i\le t+1},(x_i)_{i\le t+1}) = \mathsf{Upper}^0((x_i,y_i)_{i\le t}) + e^{-\varepsilon}\sum_{y_{t+1}} \delta\widetilde{E}^1((y_i)_{i\le t+1},(x_i)_{i\le t+1}).$$

Recall that

$$\delta E^0((y_i)_{i\le t},(x_i)_{i\le t}) \le \mathsf{Upper}^0((x_i,y_i)_{i\le t}) + e^{-\varepsilon}\delta E^1((y_i)_{i\le t},(x_i)_{i\le t}).$$

Subtracting the inequality with the equality above, we deduce that $\mathrm{Gap}_0 \le e^{-\varepsilon}\mathrm{Gap}_1$. It implies that $\mathrm{Gap}_1 > 0$. And we can use a symmetric argument to show that $\mathrm{Gap}_1 \le e^{-\varepsilon}\mathrm{Gap}_0$. Since $e^{-\varepsilon} < 1$, the only solution to the system of inequalities is $\mathrm{Gap}_0 = \mathrm{Gap}_1 = 0$, a contradiction!

Having proven the claim, we know there is a way to adjust $\widetilde{E}^{0/1}$ so that they satisfy all of (24), (25), (26). We then set $E^{0/1}((y_i)_{i\le t+1},(x_i)_{i\le t+1})$ to be $\widetilde{E}^{0/1}((y_i)_{i\le t+1},(x_i)_{i\le t+1})$ and finish the construction for $(x_i,y_i)_{i\le t}$ and x_{t+1}.

We use the construction above for $t = 0,1,\dots,T-1$ in order to construct $E^{0/1}$. It remains to verify that $E^{0/1}$ satisfies the lemma statement. It suffices to verify for every $(x_i,y_i)_{i\le T}\in(\mathcal{X}\times\mathcal{Y})^T$ and $b\in\{0,1\}$ that

$$M^b((y_i),(x_i)) \ge \delta E^b((y_i),(x_i))$$

[3] To see this, note that for $b\in\{0,1\}$, we can keep increasing $\widetilde{E}^b$ until either (1) $\widetilde{E}^b$ gets tight at y_{t+1}, or (2) $\mathrm{Gap}_b = 0$. Therefore, if we cannot increase both $\widetilde{E}^0$ and $\widetilde{E}^1$, it must be one of the four cases above.

and

$$(M^b((y_i),(x_i)) - \delta E^b((y_i),(x_i))) \le e^{\varepsilon}(M^{1-b}((y_i),(x_i)) - \delta E^{1-b}((y_i),(x_i))).$$

In fact, since $e^{\varepsilon} > 1$, it suffices to verify the second inequality for $b \in \{0,1\}$. This can be verified by utilizing (25): note that $\mathsf{Upper}^b((x_i,y_i)_{i\le T}) = M^b((y_i),(x_i)) - e^{-\varepsilon}M^{1-b}((y_i),(x_i))$, and (25) tells us that

$$\delta E^b((y_i),(x_i)) \le M^b((y_i),(x_i)) - e^{-\varepsilon}M^{1-b}((y_i),(x_i)) + e^{-\varepsilon}\delta E^{1-b}((y_i),(x_i)).$$

Re-arranging proves the desired inequality. □

Remark 1. *Note that to verify the correctness of $E^{0/1}$, we only used the condition (25). It seems that (24) is useless in this proof. However, note that it is possible that $\mathsf{Upper}^b((y_i),(x_i))$ is negative for some $(x_i,y_i)_{i\le t}$, which makes it unclear whether (25) can always be satisfied by a positive valuation of E. This is why we need the other control function* Lower.

A.1.2 Wrap-up

Next, we quickly prove Lemma 2.

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

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

Proof. We follow the notation in Section A.1.1. Namely, for each $(x_i,y_i)_{i\le t}$, define

$$M((y_i)_{i\le t},(x_i)_{i\le t}) = \prod_{i=1}^{t}\Pr[\mathcal{M}((x_j,y_j)_{j<i},x_i) = y_i].$$

Also define the same notation for E. Then we construct

$$N((y_i)_{i\le t},(x_i)_{i\le t}) = \frac{1}{1-\delta}\left(M((y_i)_{i\le t},(x_i)_{i\le t}) - \delta E((y_i)_{i\le t},(x_i)_{i\le t})\right).$$

Since $\mathcal{M} \ge \delta\mathcal{E}$, we know that $N((y_i),(x_i))$ is always non-negative. Moreover, N encodes a valid system because

$$\begin{aligned}
&\sum_{y_{t+1}\in\mathcal{Y}} N((y_i)_{i\le t+1},(x_i)_{i\le t+1})\\
&= \frac{1}{1-\delta}\sum_{y_{t+1}\in\mathcal{Y}} M((y_i)_{i\le t+1},(x_i)_{i\le t+1}) - \delta E((y_i)_{i\le t+1},(x_i)_{i\le t+1})\\
&= \frac{1}{1-\delta}\left(M((y_i)_{i\le t},(x_i)_{i\le t}) - \delta E((y_i)_{i\le t},(x_i)_{i\le t})\right)\\
&= N((y_i)_{i\le t},(x_i)_{i\le t}).
\end{aligned}$$

Finally, it is easy to verify $\mathbf{IT}(\mathcal{A}:\mathcal{M}) \equiv \delta\mathbf{IT}(\mathcal{A}:\mathcal{E}) + (1-\delta)\mathbf{IT}(\mathcal{A}:\mathcal{M})$. □

As we have shown in Section 4.1, combining Lemma 1, 2 and 3 together, we can prove Theorem 3 easily. Next, we show how Theorem 3 implies Theorem 1.

Proof of Theorem 1. Let $\mathcal{M}_1,\ldots,\mathcal{M}_k$ be k mechanisms, where $\mathcal{M}_i$ is (ε_i,δ_i)-approximate differentially private. We assume without loss of generality that all of $\mathcal{M}_i$'s hold a bit $b \in \{0,1\}$ as the sensitive data.

Let $\mathcal{A}$ be an arbitrary adversary interacting with $\mathrm{COMP}(\mathcal{M}_1,\ldots,\mathcal{M}_k)$. Next, we show how one can simulate $\mathbf{IT}(\mathcal{A},\mathrm{COMP}(\mathcal{M}_1^b,\ldots,\mathcal{M}_k^b))$ by running k (approximate) randomized response mechanisms. For each $\mathcal{M}_i^b$, construct an approximate randomized response mechanism $\mathrm{RR}^b_{\varepsilon_i,\delta_i}$. The output distribution of $\mathrm{RR}^b_{\varepsilon_i,\delta_i}$ is:

$$\mathrm{RR}^b_{\varepsilon_i,\delta_i} = \begin{cases} (b,\top) & \text{w.p. } \delta \\ (b,\bot) & \text{w.p. } (1-\delta)\frac{e^\varepsilon}{1+e^\varepsilon} \\ (1-b,\bot) & \text{w.p. } (1-\delta)\frac{1}{1+e^\varepsilon} \end{cases}.$$

We also prepare the decomposition of $\mathcal{M}_i^{0/1}$ with $\mathcal{N}_i^{0/1}, \mathcal{E}_i^{0/1}$ as promised by Theorem 3.

Now, we construct a simulator $\mathcal{S}$ as follows. For each $i \in [k]$, $\mathcal{S}$ runs $\mathrm{RR}^b_{\varepsilon_i,\delta_i}$ and gets a pair (b_i, σ_i). If $\sigma_i = \top$, then let $\mathcal{B}_i \leftarrow \mathcal{E}_i^{b_i}$. Otherwise, let $\mathcal{B}_i \leftarrow \mathcal{N}_i^{b_i}$. In this way, $\mathcal{S}$ gets a list of k systems $(\mathcal{B}_1, \dots, \mathcal{B}_k)$. The simulator then simulates the interaction between $\mathcal{A}$ and $\mathcal{B}_1, \dots, \mathcal{B}$, and outputs the interaction history. Let $\mathbf{Output}(\mathcal{S}, b)$ denote the output distribution of $\mathcal{S}$. We claim that

$$\mathbf{Output}(\mathcal{S}, b) \equiv \mathbf{IT}(\mathcal{A} : \mathcal{M}_1^b, \dots, \mathcal{M}_k^b). \tag{28}$$

To see this, for each $\mathcal{M}_i^b$, consider a two-party communication, where one party is $\mathcal{M}_i^b$, and the other party consists of $\mathcal{A}$ and $\mathcal{M}_j^b$ for $j \neq i$. The second party simulates all the interactions between $\mathcal{A}$ and $\mathcal{M}_j^b$, and only sends queries to $\mathcal{M}_i^b$ when $\mathcal{A}$ queries $\mathcal{M}_i^b$. From the second party's viewpoint, $\mathcal{M}_i^b$ looks identical to $\delta_i \mathcal{E}_i^b + (1-\delta_i)\frac{e^\varepsilon}{1+e^\varepsilon}\mathcal{N}_i^b + (1-\delta_i)\frac{1}{1+e^\varepsilon}\mathcal{N}_i^{1-b}$. Therefore,

$$\mathbf{IT}(\mathcal{A} : \mathcal{M}_1^b, \dots, \mathcal{M}_i^b, \dots, \mathcal{M}_k^b) \equiv \sum_{j=1}^{3} p_j \cdot \mathbf{IT}(\mathcal{A} : \mathcal{M}_1^b, \dots, \mathcal{M}_{i,j}^b, \dots, \mathcal{M}_k^b).$$

Here, we use $(p_1, p_2, p_3) = (\delta_i, (1-\delta_i)\frac{e^\varepsilon}{1+e^\varepsilon}, (1-\delta_i)\frac{1}{1+e^\varepsilon})$ and $(\mathcal{M}_{i,1}^b, \mathcal{M}_{i,2}^b, \mathcal{M}_{i,3}^b) = (\mathcal{E}_i^b, \mathcal{N}_i^b, \mathcal{N}_i^{1-b})$ for convenience. Applying this decomposition for every $i \in [k]$ proves (28).

Finally, note that $\mathbf{Output}(\mathcal{S}, b)$ is just a post-processing of the sequential composition of k (approximate) randomized response mechanisms. Hence, the optimal sequential composition theorem holds for $\mathbf{Output}(\mathcal{S}, b)$, which completes the proof. □

A.2 Proofs for Rényi Differential Privacy

In this section, we show omitted proofs for Theorem 2.

A.2.1 Preliminaries

We need some technical preparations first. Consider a measure space (X, μ). For two measurable functions f, g, define their inner product as

$$\langle f, g \rangle_\mu = \int f \cdot g d\mu.$$

For a real $\alpha \geq 1$, define the ℓ_α-norm of a function f as

$$\|f\|_{\mu,\alpha} := \left(\int f^\alpha d\mu \right)^{1/\alpha}.$$

Recall Hölder's inequality, which is essential for our proof.

Fact 1. *Suppose $\alpha, \beta \geq 1$ are Hólder conjugates of each other (i.e., $\frac{1}{\alpha} + \frac{1}{\beta} = 1$). Suppose f, g are two measurable functions. Then we have*

$$\langle f, g \rangle_\mu \leq \|f\|_{\mu,\alpha} \cdot \|g\|_{\mu,\beta}.$$

The inequality is sharp in the sense that for every measurable function f, we have

$$\|f\|_{\mu,\alpha} = \sup_{h : h \not\equiv 0} \frac{\langle f, h \rangle_\mu}{\|h\|_{\mu,\beta}}.$$

Recall our definition of dominance. For two measures P, Q on a space $\mathcal{Y}$, we say that P is β-dominated by Q, denoted by $P \preceq_\beta Q$, if for every measurable function $f : \mathcal{Y} \to \mathbb{R}^{\geq 0}$, it holds that $\|f\|_{P,1} \leq \|f\|_{Q,\beta}$.

A.2.2 Proof for lemmas

We are ready to show the proofs now. We start with Lemma 4.

Reminder of Lemma 4. *Suppose P, Q are two distributions supported on $\mathcal{Y}$. 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 $\mathbb{E}_{y\sim P}[h(y)] \leq e^{\frac{B(\alpha-1)}{\alpha}}\, \mathbb{E}_{y\sim Q}[h(y)^\beta]^{1/\beta}$.*

Proof. First, if there is $y \in \mathcal{Y}$ such that $0 = \Pr[Q = y] < \Pr[P = y]$, then we have $D_\alpha(P\|Q) = \infty$ and Condition 2 does not hold for any $B < \infty$. In the following, we assume $\mathrm{supp}(P) = \mathrm{supp}(Q) = \mathcal{Y}$. Note that in this case, we have $D_\alpha(P\|Q) < \infty$.

We write $P(y), Q(y)$ as shorthands for $\Pr[P = y]$ and $\Pr[Q = y]$ for brevity. Now, note that $D_\alpha(P\|Q) \leq B$ is equivalent to $e^{D_\alpha(P\|Q)} \leq e^B$, which is further equivalent to

$$\mathop{\mathbb{E}}_{y\sim Q}\left[\frac{P(y)^\alpha}{Q(y)^\alpha}\right]^{1/\alpha} = \mathop{\mathbb{E}}_{y\sim P}\left[\frac{P(y)^{\alpha-1}}{Q(y)^{\alpha-1}}\right]^{1/\alpha} \leq e^{\frac{B(\alpha-1)}{\alpha}}.$$

Consider the measure space $M = (\mathcal{Y}, Q)$. By Holder's inequality, we have

$$\mathop{\mathbb{E}}_{y\sim Q}\left[\left(\frac{P(y)}{Q(y)}\right)^\alpha\right]^{1/\alpha} = \left\|\frac{P}{Q}\right\|_{Q,\alpha} = \sup_{h:h\not\equiv 0}\left\{\frac{\langle h, \frac{P}{Q}\rangle_Q}{\|h\|_{Q,\beta}}\right\}.$$

Moreover, since $\frac{P(y)}{Q(y)}$ is non-negative, it suffices to consider only non-negative h in the supremum above. Now we are ready to verify the equivalence.

- If Condition 1 holds, we have

  $$\sup_{h:h\not\equiv 0}\left\{\frac{\langle h, \frac{P}{Q}\rangle_Q}{\|h\|_{Q,\beta}}\right\} = \left\|\frac{P}{Q}\right\|_{Q,\alpha} \leq e^{\frac{B(\alpha-1)}{\alpha}}.$$

  Therefore, for every $h : \mathcal{Y} \to \mathbb{R}^{\geq 0}$, it holds that

  $$\mathop{\mathbb{E}}_{y\sim P}[h(y)] = \mathop{\mathbb{E}}_{y\sim Q}\left[h(y)\cdot\frac{P(y)}{Q(y)}\right] \leq \left\|\frac{P}{Q}\right\|_{Q,\alpha}\|h\|_{Q,\beta} \leq e^{\frac{B(\alpha-1)}{\alpha}}\mathop{\mathbb{E}}_{y\sim Q}[h(y)^\beta]^{1/\beta}.$$

- On the other hand, if Condition 2 holds, we have

  $$\left\|\frac{P}{Q}\right\|_{Q,\alpha} = \sup_{h:h\not\equiv 0}\left\{\frac{\langle h, \frac{P}{Q}\rangle_Q}{\|h\|_{Q,\beta}}\right\} \leq e^{\frac{B(\alpha-1)}{\alpha}}.$$

This completes the proof. □

The next lemma is Lemma 5.

Reminder of 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 $\beta \geq 1, B \geq 0$ be two reals. For each $y_1 \in \mathcal{Y}_1$, define

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

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

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

Proof. Suppose by contradiction that the conclusion of the lemma does not hold. That is, there is a function $g : \mathcal{Y}_1 \to \mathbb{R}^{\geq 0}$ such that

$$\|g\|_{P_1 \ell_1^{1/\beta},1} > \|g\|_{e^B Q_1,\beta}.$$

In the following, we show this contradicts with $P \preceq_\beta e^B Q$. First off, for each $y_1 \in \mathcal{Y}_1$, by the definition of $\ell_1(y_1)$, there is a function $f_{y_1} : \mathcal{Y}_2 \to \mathbb{R}^{\geq 0}$ such that

$$\|f_{y_1}\|_{P_2|_{P_1=y_1},1} = \int f_{y_1} dP_2|_{P_1=y_1} = \left(\int f_{y_1}^{\beta} d(\ell_1(y_1) Q_2|_{Q_1=y_1})\right)^{1/\beta} = \|f_{y_1}\|_{\ell_1(y_1)Q_2|_{Q_1=y_1},\beta}.$$

By scaling f_{y_1} properly, we can ensure that $\|f_{y_1}\|_{P_2|_{P_1=y_1},1} = \ell_1(y_1)^{1/\beta}$. Consequently, we have

$$\|f_{y_1}\|_{Q_2|_{Q_1=y_1},\beta} = \|f_{y_1}\|_{\ell_1(y_1)Q_2|_{Q_1=y_1},\beta} \cdot \ell_1(y_1)^{-1/\beta} = 1.$$

Define a new function $f : \mathcal{Y}_1 \times \mathcal{Y}_2 \to \mathbb{R}^{\geq 0}$ as $f(y_1, y_2) = g(y_1) \cdot f_{y_1}(y_2)$. Then, we have

$$\begin{aligned}
\|f\|_{P,1} &= \iint f(y_1, y_2) dP \\
&= \int \left(\int f_{y_1}(y_2) dP_2|_{P_1=y_1}\right) g(y_1) dP_1 \\
&= \int g(y_1) d(\ell_1^{1/\beta} P_1) \\
&> \left(\int g(y_1)^\beta d(e^B Q_1)\right)^{1/\beta} \\
&= \left(\int g(y_1)^\beta \|f_{y_1}\|^\beta_{Q_2|_{Q_1=y_1},\beta} d(e^B Q_1)\right)^{1/\beta} \\
&= \left(\int \left(g(y_1)^\beta \int f_{y_1}(y_2)^\beta d(Q_2|_{Q_1=y_1})\right) d(e^B Q_1)\right)^{1/\beta} \\
&= \left(\iint g(y_1)^\beta f_{y_1}(y_2)^\beta d(e^B Q)\right)^{1/\beta} \\
&= \left(\iint f^\beta d(e^B Q)\right)^{1/\beta} = \|f\|_{e^B Q,\beta}.
\end{aligned} \tag{29}$$

This contradicts to the assumption that $P \preceq e^B Q$. Therefore, we conclude that such function g does not exist and $P_1 \ell_1^{1/\beta} \preceq e^B Q_1$. □

A.2.3 Proof of the composition theorem

We prove the following theorem, which is equivalent to Theorem 2.

Theorem 4. *Let $\mathcal{M}_1, \mathcal{M}_2$ be two interactive mechanisms that run on the same data set. Suppose that $\mathcal{M}_1$,$\mathcal{M}_2$ are (α, ε_1), (α, ε_2)-Rényi DP, respectively. Then $\mathrm{COMP}(\mathcal{M}_1, \mathcal{M}_2)$ is $(\alpha, \varepsilon_1 + \varepsilon_2)$-Rényi DP.*

Theorem 4 implies Theorem 2 because we can interpret $\mathrm{COMP}(\mathcal{M}_1, \dots, \mathcal{M}_k)$ as $\mathrm{COMP}(\mathrm{COMP}(\mathcal{M}_1, \dots, \mathcal{M}_{k-1}), \mathcal{M}_k)$ and use Theorem 4 inductively. Now we prove Theorem 4.

Proof. Suppose without loss of generality that both mechanisms run on a single sensitive input bit $b \in \{0, 1\}$. Also suppose that there are $2T$ rounds of interactions. Starting with $\mathcal{M}_1$, the adversary communicates with two mechanisms alternately. This is without loss of generality: suppose the adversary $\mathcal{A}$ can decide the next query object based previous responses. Let $\mathbf{IT}(\mathcal{A} : \mathcal{M}_1, \mathcal{M}_2)$ be the transcript of the interaction between $\mathcal{A}$ and $\mathcal{M}_1, \mathcal{M}_2$. We reduce the interaction to a new protocol where the adversary speaks with two mechanism alternately. Let $\mathcal{A}'$ denote a modification of $\mathcal{A}$, defined as follows. $\mathcal{A}'$ simulates $\mathcal{A}$ while always alternating between two mechanisms. If the current mechanism is not the one that $\mathcal{A}$ wants to speak with, $\mathcal{A}'$ will send a special "SKIP" query, and the mechanism responds with an "ACK" message. After this round of interaction, $\mathcal{A}'$ will switch to interact with the other mechanism, which allows it to continue simulating $\mathcal{A}$. Let $\mathbf{IT}(\mathcal{A}' : \mathcal{M}_1, \mathcal{M}_2)$ denotes the transcript of the new interaction. Therefore, for $b \in \{0,1\}$, it is easy to establish a bijection between $\mathrm{supp}(\mathbf{IT}(\mathcal{A}' : \mathcal{M}_1^b, \mathcal{M}_2^b))$ and $\mathrm{supp}(\mathbf{IT}(\mathcal{A} : \mathcal{M}_1^b, \mathcal{M}_2^b))$. Moreover, the bijection mapping is independent of b[4]. Therefore, bounding the divergences between $\mathbf{IT}(\mathcal{A} : \mathcal{M}_1^b, \mathcal{M}_2^b)$, $b \in \{0,1\}$ is equivalent to bounding those between $\mathbf{IT}(\mathcal{A}' : \mathcal{M}_1^b, \mathcal{M}_2^b)$, $b \in \{0,1\}$.

Let $\mathcal{Y}, \mathcal{Z}$ denote the response domains of $\mathcal{M}_1, \mathcal{M}_2$ respectively. Also let $y_1, \dots, y_T$, $z_1, \dots, z_T$ denote the lists of responses returned by $\mathcal{M}_1$ and $\mathcal{M}_2$ respectively. We assume that each response y_i, z_j contains a copy of the corresponding query message (so that we can recover the whole interaction history just from the responses).

Now, fix $\mathcal{A}$ to be an arbitrary adversary. Let $P, Q \in \Delta((\mathcal{Y} \times \mathcal{Z})^T)$ denote the output distributions when $\mathcal{A}$ interacts with $(\mathcal{M}_1^0, \mathcal{M}_2^0)$ and $(\mathcal{M}_1^1, \mathcal{M}_2^1)$ 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 $D_\alpha(Q\|P)$ is symmetric.

For a distribution D, we always use $D(x)$ to denote $\Pr[D = x]$. Write $y = (y_1, \dots, y_T)$ where y_i denotes the i-th response. Also write $z = (z_1, \dots, z_T)$ and denote $yz := (y_1, z_1, \dots, y_T, z_T)$. By Lemma 4, it suffices to show that for every $h\colon (\mathcal{Y} \times \mathcal{Z})^T \to \mathbb{R}^{\ge 0}$, it holds that

$$\sum_{y \in \mathcal{Y}^T, z \in \mathcal{Z}^T} P(yz)h(yz) \le \left(e^{\varepsilon_1 + \varepsilon_2} \sum_{y \in \mathcal{Y}^T, z \in \mathcal{Z}^T} Q(yz)h(yz)^\beta \right)^{1/\beta} \tag{30}$$

where $\beta = \frac{\alpha}{\alpha - 1}$ is the Hölder conjugate of α.

For each $i \in [T]$, let P_i^y, P_i^z be the projection of P onto y_i, z_i. For each $i \in [T]$, let $y_{\le i}, z_{\le i}$ denote the first i responses from y and z. Denote $(yz)_{\le i} = (y_1, z_1, \dots, y_i, z_i)$. Then, let $P_i^y|_{yz_{<i}}$ denote the distribution of y_i conditioning on $(yz)_{<i}$, and $P_i^z|_{yz_{<i}, y_i}$ denote the distribution of z_i conditioning on $(yz)_{<i}$ and y_i. Also define the same notation for Q. Then we write

$$\sum_{y \in \mathcal{Y}^T, z \in \mathcal{Z}^T} P(yz)h(yz) = \sum_{(yz)_{\le T-1}} P((yz)_{\le T-1}) \sum_{y_T, z_T} P_T^y|_{yz_{<T}}(y_T) P_T^z|_{yz_{<T}, y_T}(z_T) h(yz) \tag{31}$$

For every $t < T$ and every $y_{\le t}$, let $\mathcal{M}_1^0|_{y_{\le t}}$ (resp. $\mathcal{M}_1^1|_{y_{\le t}}$) denote the interactive system $\mathcal{M}_1^0$ (resp. $\mathcal{M}_1^1$) conditioning on that it has answered $y_1, \dots, y_t$ to the first t queries. Formally, for every $(x_{t+1}, y_{t+1}), \dots, (x_{t'}, y_{t'})$ and $x_{t'+1}$, define

$$\mathcal{M}_1^b|_{y_{\le t}}((x_i, y_i)_{t < i \le t'}, x_{t'+1}) := \mathcal{M}_1^b((x_i, y_i)_{1 \le i \le t'}, x_{t'+1}).$$

We also define the same notation for the second mechanism $\mathcal{M}_2$. Next, define

$$\ell_t(y_{\le t}) := \exp\left(\sup_{A:\text{adversary}} \left\{ D_\alpha\big(\mathbf{IT}(A : \mathcal{M}_1^0|_{y_{\le t}}) \| \mathbf{IT}(A : \mathcal{M}_1^1|_{y_{\le t}})\big) \right\} \right) \tag{32}$$

and

$$r_t(z_{\le t}) := \exp\left(\sup_{A:\text{adversary}} \left\{ D_\alpha\big(\mathbf{IT}(A : \mathcal{M}_2^0|_{z_{\le t}}) \| \mathbf{IT}(A : \mathcal{M}_2^1|_{z_{\le t}})\big) \right\} \right). \tag{33}$$

By the assumed Rényi DP guarantee, we have that $\ell_0(\emptyset) \le e^{\varepsilon_1}$ and $r_0(\emptyset) \le e^{\varepsilon_2}$. We claim the following.

Claim 3. *For each $t \le T - 1$ and $y_{<t}, z_{<t}$, consider two measures $P_t^y|_{yz_{<t}}(y)\ell_t(y_{<t} \circ y)$ and $Q_t^y|_{yz_{<t}}(y)$ on the space $\mathcal{Y}$ (here $\circ$ denotes concatenation). It holds that*

$$P_t^y|_{yz_{<t}}(y)\ell_t(y_{<t} \circ y)^{1/\beta} \preceq_\beta \ell_{t-1}(y_{<t}) Q_t^y|_{yz_{<t}}(y).$$

A symmetric conclusion holds for P^z and z. Namely

$$P_t^z|_{yz_{<t}, y_t}(z) r_t(z_{<t} \circ z)^{1/\beta} \preceq_\beta r_{t-1}(z_{<t}) Q_t^z|_{yz_{<t}, y_t}(z).$$

[4] This is to say, suppose $\mathcal{M}_1^0, \mathcal{M}_2^0, \mathcal{M}_1^1, \mathcal{M}_2^1$ are four systems, then the bijection between $\mathrm{supp}(\mathbf{IT}(\mathcal{A} : \mathcal{M}_1^b, \mathcal{M}_2^b))$ and $\mathrm{supp}(\mathbf{IT}(\mathcal{A}' : \mathcal{M}_1^b, \mathcal{M}_2^b))$ would be the same for $b \in \{0,1\}$.

Proof. Construct an adversary $\mathcal{A}'$ interacting with $\mathcal{M}_1^b|_{y_{<t}}$ as follows. $\mathcal{A}'$ starts $\mathcal{A}$ with the conditioning that $\mathcal{A}$ has gone through the interaction history $yz_{<t}$. Then $\mathcal{A}'$ simulates one step of $\mathcal{A}$ and sends a query to $\mathcal{M}_1^b|_{y_{<t}}$. Upon receiving the response y, $\mathcal{A}'$ observes y_t and switches to run the optimal adversary against $\mathcal{M}_1^b|_{y_{\le t}}$ provided by (32). By definition, $\mathcal{M}_1^b|_{y_{<t}}$ is $(\alpha, \log(\ell_{t-1}(y_{<t})))$-Rényi DP. Applying Lemma 5 on $\mathbf{IT}(\mathcal{A}', \mathcal{M}_1^b|_{y_{<t}})$ completes the proof. The proof for P_t^z is similar. □

Turning back to (31), we first deduce that

$$\begin{aligned}
&\sum_{(yz)_{\le T-1}} P((yz)_{\le T-1}) \sum_{y_T, z_T} P_T^y|_{yz_{<T}}(y_T) P_T^z|_{yz_{<T}, y_T}(z_T) h(yz) \\
\le &\sum_{(yz)_{\le T-1}} P((yz)_{\le T-1}) \sum_{y_T} P_T^y|_{yz_{<T}}(y_T) \left(r_{T-1}(z_{<T}) \sum_{z_T} Q_T^z|_{yz_{<T}, y_T}(z_T) h(yz)^\beta \right)^{1/\beta} \\
\le &\sum_{(yz)_{\le T-1}} P((yz)_{\le T-1}) \left(r_{T-1}(z_{<T}) \ell_{T-1}(y_{<T}) \sum_{y_T, z_T} Q_T^y|_{yz_{<T}}(y_T)\, Q_T^z|_{yz_{<T}, y_T}(z_T) h(yz)^\beta \right)^{1/\beta} .
\end{aligned} \tag{34}$$

So far we haven't utilized Claim 3 yet. Denote

$$H(yz_{\le T-1}) := \left(\ell_{T-1}(y_{<T}) \sum_{y_T, z_T} Q_T^y|_{yz_{<T}}(y_T)\, Q_T^z|_{yz_{<T}, y_T}(z_T) h(yz)^\beta \right)^{1/\beta} .$$

Applying Claim 3 on (34) for P_{T-1}^z yields that

$$\begin{aligned}
&\sum_{(yz)_{\le T-2}, y_{T-1}} P((yz)_{\le T-2}, y_{T-1}) \sum_{z_{T-1}} P_{T-1}^z|_{yz_{\le T-2}, y_{T-1}}(z_{T-1}) r_{T-1}(z_{<T})^{1/\beta} H \\
\le &\sum_{(yz)_{\le T-2}, y_{T-1}} P((yz)_{\le T-2}, y_{T-1}) \left(r_{T-2}(z_{\le T-2}) \sum_{z_{T-1}} Q_{T-1}^z|_{yz_{\le T-2}, y_{T-1}}(z_{T-1}) H^\beta \right)^{1/\beta} .
\end{aligned} \tag{35}$$

We proceed to apply Claim 3 on (35) for $P_{T-1}^y, P_{T-2}^z, P_{T-2}^y \ldots, P_1^z, P_1^y$ in order. We can get

$$\begin{aligned}
&\sum_{(yz)_{\le T-1}} P((yz)_{\le T-1}) \sum_{y_T, z_T} P_T^y|_{yz_{<T}}(y_T) P_T^z|_{yz_{<T}, y_T}(z_T) h(yz) \\
\le &\left(\ell_0(\emptyset) r_0(\emptyset) \sum_{yz} Q(yz) h(yz)^\beta \right)^{1/\beta} .
\end{aligned} \tag{36}$$

This shows that $P \preceq e^{\varepsilon_1 + \varepsilon_2} Q$, which consequently implies that $D_\alpha(P \| Q) \le \varepsilon_1 + \varepsilon_2$. Similarly, we can bound $D_\alpha(Q \| P) \le \varepsilon_1 + \varepsilon_2$. Combining two bounds together completes the proof. □

A.3 Proof for Concentrated DP

In this section, we prove Corollary 1. We recall the definition of zero-concentrated DP and truncated concentrated DP.

Definition 4 (zero-concentrated differential privacy, Bun and Steinke [2016])**.** Let $\rho > 0$ be a real and $\mathcal{M}$ be a mechanism. $\mathcal{M}$ is called ρ-zero-concentrated DP (or ρ-zCDP for short), if for every $\alpha \in (1, +\infty)$, $\mathcal{M}$ is $(\alpha, \alpha \cdot \rho)$-RDP.

Definition 5 (truncated concentrated differential privacy Bun et al. [2018])**.** Let $\rho > 0, \omega > 1$ be two reals, and $\mathcal{M}$ be a mechanism. $\mathcal{M}$ is called (ρ, ω)-truncated DP (or (ρ, ω)-tCDP), if for every $\alpha \in (1, \omega)$, $\mathcal{M}$ is $(\alpha, \alpha \cdot \rho)$-RDP.

We are ready to prove Corollary 1 below.

Proof. We first prove for zCDP. Suppose $\mathcal{M}_1, \dots, \mathcal{M}_k$ are k interactive mechanisms, where for each $i \in [k]$, $\mathcal{M}_i$ is ρ_i-zCDP. By definition, we know that $\mathcal{M}_i$ is $(\alpha, \alpha\rho_i)$-RDP for every $\alpha > 1$. By Theorem 2, we know that $\mathrm{COMP}(\mathcal{M}_1, \dots, \mathcal{M}_k)$ satisfies $(\alpha, \alpha(\sum_i \rho_i))$-RDP. Since this argument holds for every $\alpha > 1$, we conclude that $\mathrm{COMP}(\mathcal{M}_1, \dots, \mathcal{M}_k)$ satisfies $(\sum_i \rho_i)$-zCDP.

The proof for tCDP is similar. Fix $\omega > 1$. Again let $\mathcal{M}_1, \dots, \mathcal{M}_k$ are k interactive mechanisms, where for each $i \in [k]$, $\mathcal{M}_i$ is (ρ_i, ω)-tCDP. Then, for every $\alpha \in (1, \omega)$, we know that $\mathcal{M}_i$ is $(\alpha, \alpha\rho_i)$-RDP by definition. Theorem 2 then shows that $\mathrm{COMP}(\mathcal{M}_1, \dots, \mathcal{M}_k)$ satisfies $(\alpha, \alpha(\sum_i \rho_i))$-RDP. Since the argument holds for every $\alpha \in (1, \omega)$, we conclude that $\mathrm{COMP}(\mathcal{M}_1, \dots, \mathcal{M}_k)$ satisfies $(\sum_i \rho_i, \omega)$-tCDP. □

B A Motivating Example of Concurrent Composition

To demonstrate the power of concurrent composition, in this section, we use Theorem 1 to analyze a simple private "Guess-and-Check" algorithm. We remark that this is a rather preliminary application: the weaker concurrent composition theorem by Vadhan and Wang [2021] is sufficient to do the job. However, the main purpose of this section is to highlight the importance of concurrent composition, and hopefully inspire researchers to design more sophisticated algorithms.

Setup. Now we describe the problem. The private algorithm holds a sensitive data set X. The user keeps issuing queries to the algorithm, where each query consists of a 1-Lipschitz function f_i and a guess $\tau_i \in \mathbb{R}$ for the value of $f_i(X)$. The algorithm's job is to verify if $f_i(X) \approx \tau_i$. If it is the case, the algorithm reports "PASS" and continues to the next query. Otherwise, the algorithm reports "WRONG" and a value v_i that is approximately equal to $f_i(X)$ (i.e., the algorithm not only declares the invalidity of the user's guess, but also provides a correct estimation for $f_i(X)$).

We consider the following algorithm.

Algorithm 1: The Private Guess-and-Check

```
Input: Private dataset $X$. Error tolerance parameter $E > 0$. Privacy-related parameters
       $c \geq 1, \varepsilon \in (0, 1)$.
1  Program:
2      $\rho \leftarrow \mathrm{Lap}\left(\frac{1}{\varepsilon}\right)$      // Note that this noise has standard deviation $\approx \frac{1}{\varepsilon}$
3      for $i = 1, 2, \dots,$ do
4          Receive the next query $(f_i, \tau_i)$
5          $\gamma_i \leftarrow \mathrm{Lap}(c/\varepsilon)$
6          if $|f_i(X) - \tau_i| + \gamma_i \geq E + \rho$ then
7              $v_i \leftarrow f_i(X) + \mathrm{Lap}(c/\varepsilon)$
8              Report $(\mathrm{WRONG}, v_i)$
9              $t \leftarrow t + 1$
10             if $t = c$ then
11                 HALT the algorithm.
12         else
13             Return PASS
```

Discussions. Algorithm 1 is parameterized by an error tolerance parameter $E > 0$ and two privacy parameters $c \geq 1, \varepsilon \in (0, 1)$. Roughly speaking, it can process queries until identifying at least c queries whose guesses deviate from the true value by at least (roughly) E. It works by (concurrently) composing a variant of the sparse vector technique by Lyu et al. [2017] with the standard Laplace noise-adding mechanism.

The main advantage of the Lyu et al. [2017] SVT is that it only adds noise to the threshold once (Line 2 of algorithm 1), using a *much smaller* noise, which makes the SVT algorithm more accurate. Since the utility guarantee of the algorithm is not the focus of this work, we omit more discussions here and refer interested readers to [Lyu et al., 2017, Zhu and Wang, 2020] for more detail.

We consider the privacy guarantee of Algorithm 1. In fact, without the concurrent composition framework, it is not clear whether or not Algorithm 1 is really private! If we replace Line 7 of the algorithm by $v_i \leftarrow 0$, then the algorithm is indeed $(3\varepsilon, 0)$-private, because it is just a faithful implementation of the Lyu et al. [2017] SVT. However, in Algorithm 1, the algorithm reports a correct estimation v_i for each inaccurate guess, which implies that the future query to the algorithm may depend on v_i, and thus on the private data set X. In this case, the original analysis from [Lyu et al., 2017] does not hold anymore.

Analyzing the privacy. While it is not hard to prove the privacy property of Algorithm 1 by examining the proof of Lyu et al. [2017] carefully and applying some modifications, here we show that Algorithm 1 admits a fairly straightforward privacy proof under the concurrent composition framework, using the privacy theorem by Lyu et al. [2017] as a black box. We do the analysis now. First, we have the following lemma from [Lyu et al., 2017].

Lemma 6 (Theorem 2 in Lyu et al. [2017])**.** *Consider replacing Line 7 of Algorithm 1 with $v_i \leftarrow 0$. The resulting algorithm is $(3\varepsilon, 0)$-DP.*

The following fact is well known.

Lemma 7 (Laplace mechanism)**.** *Consider the following algorithm: given a list of c adaptively chosen, 1-Lipschitz queries $(g_1, \cdots, g_c)$, answer each query with $g_i(X) + \mathrm{Lap}(c/\varepsilon)$. The algorithm is $(\varepsilon, 0)$-DP.*

Combining Lemmas 6 and 7 under the concurrent composition framework directly yields the following result.

Theorem 5. *Algorithm 1 is $(4\varepsilon, 0)$-DP.*

Proof. Consider simulating Algorithm 1 by concurrently composing two algorithms A_1, A_2. A_1 is just a modification of Algorithm 1 where we replace Line 7 in Algorithm 1 with $v_i \leftarrow 0$. By Lemma 6, A_1 is $(3\varepsilon, 0)$-DP. A_2 accepts at most c 1-Lipschitz query. For each query g_i, A_2 responds with $g_i(X) + \mathrm{Lap}(c/\varepsilon)$. By Lemma 7, A_2 is $(\varepsilon, 0)$-DP. By Theorem 1, $\mathrm{COMP}(A_1, A_2)$ is $(4\varepsilon, 0)$-DP.

We now describe how to simulate Algorithm 1 with $\mathrm{COMP}(A_1, A_2)$. For each query (f_i, τ_i) to Algorithm 1, we first feed it into A_1 and observe the outcome. We pass this query if the outcome is PASS. Otherwise, the outcome must be $(\mathrm{WRONG}, 0)$. We then query A_2 with f_i to get an estimation $f_i + \mathrm{Lap}(c/\varepsilon)$, and think of this estimation as the "v_i" returned by Algorithm 1. In this way, it is easy to see that we faithfully simulate Algorithm 1 by interacting with $\mathrm{COMP}(A_1, A_2)$. Since $\mathrm{COMP}(A_1, A_2)$ is $(4\varepsilon, 0)$-DP, Algorithm 1 must be $(4\varepsilon, 0)$-DP also. This completes the proof. □

Remark 2. *Finally, we remark that a similar private "Guess-and-Check" algorithm was also proposed and analyzed by Zhu and Wang [2020], where the authors also considered using a version of SVT* without *refreshing the threshold after answering each "meaningful" query. Therefore, their algorithm is also subject to the concurrent composition issue, which seems to be overlooked in the original analysis of Zhu and Wang [2020]. Since they were working with Rényi DP, our Theorem 2 provides a remedy to this issue easily.*