
Concurrent Composition of Differential Privacy

Salil Vadhan¹ Tianhao Wang¹

Abstract

We initiate a study of the composition properties of *interactive* differentially private mechanisms. An interactive differentially private mechanism is an algorithm that allows an analyst to adaptively ask queries about a sensitive dataset, with the property that an adversarial analyst’s view of the interaction is approximately the same regardless of whether or not any individual’s data is in the dataset. Previous studies of composition of differential privacy have focused on non-interactive algorithms, but interactive mechanisms are needed to capture many of the intended applications of differential privacy and a number of the important differentially private primitives. We focus on *concurrent composition*, where an adversary can arbitrarily interleave its queries to several differentially private mechanisms, which may be feasible when differentially private query systems are deployed in practice. We prove that when the interactive mechanisms being composed are *pure* differentially private, their concurrent composition achieves privacy parameters (with respect to pure or approximate differential privacy) that match the (optimal) composition theorem for noninteractive differential privacy. We also prove a composition theorem for interactive mechanisms that satisfy approximate differential privacy. That bound is weaker than even the basic (suboptimal) composition theorem for noninteractive differential privacy, and we leave closing the gap as a direction for future research, along with understanding concurrent composition for other variants of differential privacy.

1. Introduction

1.1. Composition of Differential Privacy

A crucial property of differential privacy is its behavior under composition. If we run multiple distinct differentially private algorithms on the same dataset, the resulting composed algorithm is also differentially private, with some degradation in the privacy parameters (ϵ, δ) . This property is especially important and useful since in practice we rarely want to release only a single statistic about a dataset. Releasing many statistics may require running multiple differentially private algorithms on the same database. Composition is also a very useful tool in algorithm design. In many cases, new differentially private algorithms are created by combining several simpler algorithms. The composition theorems help us analyze the privacy properties of algorithms designed in this way.

Formally, let $\mathcal{M}_0, \mathcal{M}_1, \dots, \mathcal{M}_{k-1}$ be differentially private mechanisms, we define the composition of these mechanisms by independently executing them. Specifically, we define $\mathcal{M} = \text{Comp}(\mathcal{M}_0, \mathcal{M}_1, \dots, \mathcal{M}_{k-1})$ as $\mathcal{M}(x) = (\mathcal{M}_0(x), \dots, \mathcal{M}_{k-1}(x))$ where each \mathcal{M}_i is run with independent coin tosses. For example, this is how we might obtain a mechanism answering a k -tuple of queries.

A handful of composition theorems already exist in the literature. The Basic Composition Theorem (Dwork et al., 2006) says that the privacy degrades at most linearly with the number of mechanisms executed. However, if we are willing to tolerate an increase in the δ term, we obtain the Advanced Composition Theorem (Dwork et al., 2010) where the privacy parameter ϵ only needs to degrade proportionally to the square root of number of mechanisms executed. Despite giving an asymptotically correct upper bound for the global privacy parameter, the Advanced Composition Theorem is not exact. Kairouz, Oh, and Viswanath (Kairouz et al., 2015) shows how to compute the *optimal* bound for composing k mechanisms where all of them are (ϵ, δ) -differentially private. Murtagh and Vadhan (Murtagh & Vadhan, 2016) further extends the optimal composition for the more general case where the privacy parameters may differ for each algorithm in the composition.

¹Harvard University. Correspondence to: Tianhao Wang <tianhaowang@fas.harvard.edu>.

1.2. Interactive Differential Privacy

The standard treatment of differential privacy refers to a *noninteractive* algorithm \mathcal{M} that takes a dataset x as input and produces a statistical release $\mathcal{M}(x)$. However, in many of the motivating applications of differential privacy, we don't want to perform all of our releases in one shot, but rather allow analysts to make *adaptive* queries to a dataset. Thus, we should view the mechanism \mathcal{M} as a party in a two-party protocol, interacting with a (possibly adversarial) analyst.

To formalize the concept of interactive DP, we recall one of the standard formalizations of an *interactive protocol* between two parties A and B . We do this by viewing each party as a function, taking its private input x , all messages it has received (m_0, m_1, \dots) , and the party's random coins r , to the party's next message to be sent out. Based on that, we define the *view* of a party in an interactive protocol to capture everything the party "sees" during the execution:

Definition 1 (View of a party in an interactive protocol). *Let (A, B) be an interactive protocol. Let r_A and r_B be the random coins for A and B , respectively. A 's view of $(A(x_A; r_A), B(x_B; r_B))$ is the tuple $\text{View}_A(A(x_A; r_A), B(x_B; r_B)) = (r_A, x_A, m_0, m_1, \dots)$ consisting of all the messages received by A in the execution of the protocol together with the private input x_A and random coins r_A . If we drop the random coins r_A and/or r_B , $\text{View}_A(A(x_A), B(x_B))$ becomes a random variable. B 's view of $(A(x_A), B(x_B))$ is defined symmetrically.*

In our case, A is the adversary and B is the mechanism whose input is usually a database x . Since A does not have an input in our case, we will denote the interactive protocol as $(A, B(x))$ for the ease of notation. Since we will only be interested in A 's view and A does not have an input, we will drop the subscript and write A 's view as $\text{View}\langle A, B(x) \rangle$.

Now we are ready to define the interactive differential privacy as a type of interactive protocol between an adversary (without any computational limitations) and an interactive mechanism of special properties.

Definition 2 (Interactive Differential Privacy). *A randomized algorithm \mathcal{M} is an (ϵ, δ) -differentially private interactive mechanism if for every pair of adjacent datasets $x, x' \in \text{MultiSets}(\mathcal{X})$, for every adversary algorithm \mathcal{A} we have: we have*

$$\forall T \subseteq \text{Range}(\text{View}\langle \mathcal{A}, \mathcal{M}(\cdot) \rangle), \Pr[\text{View}\langle \mathcal{A}, \mathcal{M}(x) \rangle \in T] \leq e^\epsilon \Pr[\text{View}\langle \mathcal{A}, \mathcal{M}(x') \rangle \in T] + \delta \quad (1)$$

where the randomness is over the coin flips of both the algorithm \mathcal{M} and the adversary \mathcal{A} .

In addition to being the "right" modelling for many applications of differential privacy, interactive differential privacy also captures the full power of fundamental DP mechanisms such as the Sparse Vector Technique (Dwork et al., 2009; Roth & Roughgarden, 2010) and Private Multiplicative Weights (Hardt & Rothblum, 2010), which are in turn useful in the design of other DP algorithms (which can use these mechanisms as subroutines and issue adaptive queries to them). Interactive DP was also chosen as the basic abstraction in the programming framework for the new open-source software project OpenDP (Gaboardi et al., 2020), which was our motivation for this research.

Despite being such a natural and useful notion, interactive DP has not been systematically studied before. It has been implicitly studied in the context of distributed forms of DP, starting with (Beimel et al., 2008), where the sensitive dataset is split amongst several parties, who execute a multiparty protocol to estimate a joint function of their data, while each party ensures that their portion of the dataset has the protections of DP against the other parties. Indeed, in an m -party protocol, requiring DP against malicious coalitions of size $m - 1$ is equivalent to requiring that each party's strategy is an interactive DP mechanism in the sense of Definition 2. An extreme case of this is the *local model* of DP, where each party holds a single data item in \mathcal{X} representing data about themselves (Kasiviswanathan et al., 2011). There been extensive research about the power of interactivity in local DP; see (Chen et al., 2020) and the references therein. In contrast to these distributed models, in Definition 2 we are concerned with the *centralized DP* scenario where only one party (\mathcal{M}) holds sensitive data, and how an adversarial data analyst (\mathcal{A}) may exploit adaptive queries to extract information about the data subjects.

1.3. Our Contributions.

In this paper, we initiate a study of the composition of interactive DP mechanisms. Like in the context of cryptographic protocols, there are several different forms of composition we can consider. The simplest is *sequential composition*, where all of the queries to \mathcal{M}_{i-1} must be completed before any queries are issued to \mathcal{M}_i . It is straightforward to extend the proofs of the noninteractive DP composition theorems to handle sequential composition of interactive DP mechanisms.

Thus, we turn to *concurrent composition*, where an adversary can arbitrarily interleave its queries to the k mechanisms. Although the mechanisms use independent randomness, the adversary may create correlations between the executions

by coordinating its actions; in particular, its queries in one execution may also depend on messages it received in other executions. Concurrent composability is important for the deployment of interactive DP in practice, as one or more organizations may set up multiple DP query systems on datasets that refer to some of the same individuals, and we would not want the privacy of those individuals to be violated by an adversary that can concurrently access those systems. Concurrent composability may also be useful in the design of DP algorithms; for example, one might design a DP machine learning algorithm that uses adaptive and interleaved queries to two instantiations of an interactive DP mechanism like the Sparse Vector Technique (Dwork et al., 2009; Roth & Roughgarden, 2010).

Our findings is summarized as follows:

- We prove a composition theorem for interactive mechanisms that satisfy approximate differential privacy, where the bound for δ is weaker than the basic composition theorem for noninteractive differential privacy.
- We prove that when the interactive mechanisms being composed are *pure* differentially private, their concurrent composition achieves privacy parameters (with respect to pure or approximate differential privacy) that match the (optimal) composition theorem for noninteractive differential privacy.
- Based on computer experiments, we conjecture that the Optimal Composition Theorems can be extended to the concurrent composition of approximate DP mechanisms. We leave closing the gap as a direction for future research, along with understanding concurrent composition for other variants of differential privacy.

2. Basic Composition Theorem for Concurrent Composition

We use $\text{ConComp}(\mathcal{M}_0, \dots, \mathcal{M}_{k-1})$ to denote the concurrent composition of interactive mechanisms $\mathcal{M}_0, \dots, \mathcal{M}_{k-1}$. (See full version of paper for a formal definition.) Our first result is roughly an analogue of the Basic Composition Theorem.

Theorem 1. *If interactive mechanisms $\mathcal{M}_0, \dots, \mathcal{M}_{k-1}$ are each (ε, δ) -differentially private, then their concurrent composition $\text{ConComp}(\mathcal{M}_0, \dots, \mathcal{M}_{k-1})$ is $(k \cdot \varepsilon, \frac{e^{k\varepsilon} - 1}{e^\varepsilon - 1} \cdot \delta)$ -differentially private.*

More generally, if interactive mechanism \mathcal{M}_i is $(\varepsilon_i, \delta_i)$ -differentially private for $i = 0, \dots, k - 1$, then $\text{ConComp}(\mathcal{M}_0, \dots, \mathcal{M}_{k-1})$ is $(\varepsilon_g, \delta_g)$ -differentially private where $\varepsilon_g = \sum_{i=0}^{k-1} \varepsilon_i$, and $\delta_g = \sum_{i=0}^{k-1} e^{\sum_{j=0}^{i-1} \varepsilon_j} \cdot \delta_i$.

Just like in the Basic Composition Theorem for noninteractive DP, the privacy-loss parameters ε_i just sum up. However, the bound on δ_g is worse by a factor of at most e^{ε_g} . In the typical setting where we want to enforce a global privacy loss of $\varepsilon_g = O(1)$, this is only a constant-factor loss compared to the Basic Composition Theorem, but that constant can be important in practice. Note that expression for δ_g depends on the ordering of the k mechanisms $\mathcal{M}_0, \dots, \mathcal{M}_{k-1}$, so one can optimize it further by taking a permutation of the mechanisms that minimizes δ_g . The proof is by a standard hybrid argument, which is detailed in the full version of the paper. We note that the proof is very similar to the proof of the group privacy property of (noninteractive) differential privacy, where (ε, δ) -DP for datasets that differ on one record implies $(k \cdot \varepsilon, \frac{e^{k\varepsilon} - 1}{e^\varepsilon - 1} \cdot \delta)$ for datasets that differ on k records.

3. Advanced Composition Theorem for Pure Concurrent Composition

Next we show that the Advanced and Optimal Composition Theorems for noninteractive DP can be extend to interactive DP, provided that the mechanisms \mathcal{M}_i being composed satisfy pure DP (i.e. $\delta_i = 0$). Note that the final composed mechanism $\text{ConComp}(\mathcal{M}_0, \dots, \mathcal{M}_{k-1})$ can be approximate DP, by taking $\delta_g = \delta' > 0$, and thereby allowing for a privacy loss ε_g that grows linearly in \sqrt{k} rather than k .

We do this by extending the main proof technique of (Kairouz et al., 2015; Murtagh & Vadhan, 2016) to interactive DP mechanisms. Specifically, we reduce the analysis of interactive $(\varepsilon, 0)$ -DP mechanisms to that of analyzing the following simple “randomized response” mechanism:

Definition 3 ((Kairouz et al., 2015)). *For $\varepsilon > 0, \delta \in [0, 1]$, define a randomized noninteractive algorithm $\text{RR}_{(\varepsilon, \delta)} : \{0, 1\} \rightarrow \{0, 1, \text{'Iam0'}, \text{'Iam1'}\}$ as follows:*

$$\begin{array}{ll}
 \Pr [\text{RR}_{(\varepsilon, \delta)}(0) = \text{'Iam0'}] = \delta & \Pr [\text{RR}_{(\varepsilon, \delta)}(1) = \text{'Iam0'}] = 0 \\
 \Pr [\text{RR}_{(\varepsilon, \delta)}(0) = 0] = (1 - \delta) \cdot \frac{e^\varepsilon}{1 + e^\varepsilon} & \Pr [\text{RR}_{(\varepsilon, \delta)}(1) = 0] = (1 - \delta) \cdot \frac{1}{1 + e^\varepsilon} \\
 \Pr [\text{RR}_{(\varepsilon, \delta)}(0) = 1] = (1 - \delta) \cdot \frac{1}{1 + e^\varepsilon} & \Pr [\text{RR}_{(\varepsilon, \delta)}(1) = 1] = (1 - \delta) \cdot \frac{e^\varepsilon}{1 + e^\varepsilon} \\
 \Pr [\text{RR}_{(\varepsilon, \delta)}(0) = \text{'Iam1'}] = 0 & \Pr [\text{RR}_{(\varepsilon, \delta)}(1) = \text{'Iam1'}] = \delta
 \end{array}$$

Note that $\text{RR}_{(\varepsilon, \delta)}$ is a noninteractive (ε, δ) -DP mechanism. For simplicity, when $\delta = 0$, we denote the mechanism as RR_ε . We show that every interactive $(\varepsilon, 0)$ -DP mechanism can be, in some sense, simulated from RR_ε :

Theorem 2. *Suppose that \mathcal{M} is an interactive $(\varepsilon, 0)$ -differentially private mechanism. Then for every pair of adjacent datasets x_0, x_1 there exists an interactive mechanism T s.t. for every adversary \mathcal{A} and every $b \in \{0, 1\}$ we have*

$$\text{View}(\mathcal{A}, \mathcal{M}(x_b)) \equiv \text{View}(\mathcal{A}, T(\text{RR}_\varepsilon(b)))$$

Here T is an interactive mechanism that depends on \mathcal{M} as well as a fixed pair of adjacent datasets x_0 and x_1 . It receives a single bit as an output of $\text{RR}_\varepsilon(b)$, and then interacts with the adversary \mathcal{A} just like \mathcal{M} would. Kairouz, Oh, and Viswanath (Kairouz et al., 2015) proved Theorem 2 result for the case that \mathcal{M} and T are noninteractive. The interactive case is more involved because we need a single T that works for all adversary strategies \mathcal{A} . (If we allow T to depend on the adversary strategy \mathcal{A} , then the result would readily follow from that of (Kairouz et al., 2015), but this would not suffice for our application to concurrent composition.)

Given the Theorem 2, to analyze $\text{ConComp}(\mathcal{M}_0(x_b), \dots, \mathcal{M}_{k-1}(x_b))$ on $b = 0$ vs. $b = 1$, it suffices to analyze $\text{ConComp}(T_0(\text{RR}_{\varepsilon_0}(b)), \dots, T_{k-1}(\text{RR}_{\varepsilon_{k-1}}(b)))$. An adversary's view interacting with the latter concurrent composition can be simulated entirely from the output of $\text{Comp}(\text{RR}_{\varepsilon_0}(b), \dots, \text{RR}_{\varepsilon_{k-1}}(b))$, which is the composition of entirely noninteractive mechanisms. Thus, we conclude:

Corollary 3. *The Advanced and Optimal Composition Theorems extend to the concurrent composition of $(\varepsilon_i, \delta_i)$ -interactive DP mechanisms \mathcal{M}_i provided that $\delta_0 = \delta_1 = \dots = \delta_{k-1} = 0$.*

4. Future Work

In the full version of the paper, we present empirical evidence for our conjecture that the Optimal Composition Theorems can be extended to the concurrent composition of approximate DP mechanisms. Specifically, we experimentally evaluate the conjecture for 3-message interactive mechanisms with 1-bit messages, and we test whether instantiations of this 2-round interactive mechanism that are (ε, δ) -DP can be simulated as the interactive post-processing of randomized response $\text{RR}_{(\varepsilon, \delta)}$. In all of our trials, we find a feasible interactive post-processing algorithm. The implementation details of the experiment can be found in the full version of the paper. Based on the above findings, we conjecture that the concurrent composition of interactive DP mechanisms may still have the same bound as the composition for noninteractive DP mechanisms. Besides, we might be able to prove it through a similar construction of interactive post-processing mechanisms as we did in Theorem 2. This means that every interactive DP mechanisms can be reduced to noninteractive randomized response. We leave the resolution of these conjectures for future work.

Another interesting question for future work is analyzing concurrent composition for variants of differential privacy, such as Concentrated DP (Dwork & Rothblum, 2016; Bun & Steinke, 2016; Bun et al., 2018), Rényi DP (Mironov, 2017), and Gaussian DP (Dong et al., 2019). Some of these notions require bounds on Rényi divergences, e.g. that

$$D_\alpha(\text{View}(\mathcal{A}, \mathcal{M}(x)) || \text{View}(\mathcal{A}, \mathcal{M}(x'))) \leq \rho,$$

for adjacent datasets x, x' and certain pairs (α, ρ) . Here sequential composition can be argued using a chain rule for Rényi divergence:

$$D_\alpha((Y, Z) || (Y', Z')) \leq D_\alpha(Y || Y') + \sup_y D_\alpha(Z|_{Y=y} || Z'|_{Y'=y}). \quad (2)$$

Taking Y to be the view of the analyst interacting with $\mathcal{M}_0(x)$, Z to be the view of the analyst in a subsequent interaction with $\mathcal{M}_1(x)$, and Y' and Z' to be analogously defined with respect to an adjacent dataset x' , we obtain the usual composition bound of $\rho_0 + \rho_1$ on the overall Rényi divergence of order α , where ρ_0 and ρ_1 are the privacy-loss parameters of the individual mechanisms. However, this argument fails for concurrent DP, since we can no longer assert the privacy properties of \mathcal{M}_1 conditioned on any possible value y of the adversary's view of the interaction with \mathcal{M}_0 . Unfortunately, the Chain Rule (2) does not hold if we replace the supremum with an expectation, so a new proof strategy is needed (if the composition theorem remains true).

Acknowledgements

We sincerely thank Boaz Barak, Cynthia Dwork, Marco Gaboardi, Michael Hay, Weiwei Pan, and Andy Vyrros for helpful comments and discussions.

References

- Beimel, A., Nissim, K., and Omri, E. Distributed private data analysis: Simultaneously solving how and what. In *Annual International Cryptology Conference*, pp. 451–468. Springer, 2008.
- Bun, M. and Steinke, T. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pp. 635–658. Springer, 2016.
- Bun, M., Dwork, C., Rothblum, G. N., and Steinke, T. Composable and versatile privacy via truncated cdp. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 74–86, 2018.
- Chen, L., Ghazi, B., Kumar, R., and Manurangsi, P. On distributed differential privacy and counting distinct elements. *arXiv preprint arXiv:2009.09604*, 2020.
- Dong, J., Roth, A., and Su, W. J. Gaussian differential privacy. *arXiv preprint arXiv:1905.02383*, 2019.
- Dwork, C. and Rothblum, G. N. Concentrated differential privacy. *arXiv preprint arXiv:1603.01887*, 2016.
- Dwork, C., McSherry, F., Nissim, K., and Smith, A. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pp. 265–284. Springer, 2006.
- Dwork, C., Naor, M., Reingold, O., Rothblum, G. N., and Vadhan, S. On the complexity of differentially private data release: efficient algorithms and hardness results. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pp. 381–390, 2009.
- Dwork, C., Rothblum, G. N., and Vadhan, S. Boosting and differential privacy. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pp. 51–60. IEEE, 2010.
- Gaboardi, M., Hay, M., and Vadhan, S. A programming framework for opendp. 2020.
- Hardt, M. and Rothblum, G. N. A multiplicative weights mechanism for privacy-preserving data analysis. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pp. 61–70. IEEE, 2010.
- Kairouz, P., Oh, S., and Viswanath, P. The composition theorem for differential privacy. In *International conference on machine learning*, pp. 1376–1385. PMLR, 2015.
- Kasiviswanathan, S. P., Lee, H. K., Nissim, K., Raskhodnikova, S., and Smith, A. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.
- Mironov, I. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pp. 263–275. IEEE, 2017.
- Murtagh, J. and Vadhan, S. The complexity of computing the optimal composition of differential privacy. In *Theory of Computing*, pp. 157–175. Theory of Computing, 2016.
- Roth, A. and Roughgarden, T. Interactive privacy via the median mechanism. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pp. 765–774, 2010.