Concurrent Composition of Differential Privacy

Salil Vadhan, Tianhao Wang

Harvard University

TL; DR: We initiate a study of the *concurrent* composition properties of *interactive* differentially private mechanisms, and derived the *optimal* composition bound for pure interactive DP mechanisms.

Background: DP under Composition

- Goal: analyze the privacy loss under the composition of multiple different DP mechanisms on the same dataset.
- Examples of existing DP composition theorems: Basic Composition, Advanced Composition, Optimal Composition, Moment Accountant, etc.

Motivation

- Existing composition theorems: assume that the underlying DP mechanisms are "one-shot" algorithms.
- We want to compose interactive mechanisms, e.g., Sparse Vector Technique (SVT).

Interactive DP under Composition

- There could be more than one composition operations for interactive mechanisms.
- **Sequential Composition**: all of the queries to the current mechanism must be completed before the session with another mechanism can be spawned.

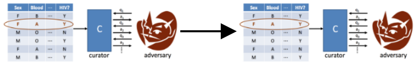

- **Concurrent Composition**: multiple interactions can be spawned and be executed simultaneously, queries to the mechanisms can be arbitrarily interleaved with each other.

- Unfortunately, none of the existing composition theorems for non-interactive DP can be directly applied to the setting of concurrent compositions.

References

Kairouz, P., Oh, S., & Viswanath, P. (2015, June). The composition theorem for differential privacy. In International conference on machine learning (pp. 1376-1385). PMLR.

Murtagh, J., & Vadhan, S. (2016, January). The complexity of computing the optimal composition of differential privacy. In Theory of Cryptography Conference (pp. 157-175). Springer, Berlin, Heidelberg.

Interactive Protocol

- Interactive protocol between two parties A and B
  - Viewing each party as a potentially randomized function.
  - (**private input, received messages, random coins**) => Next message to be sent out.

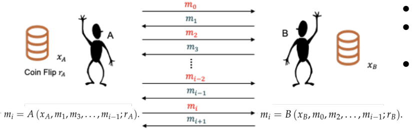

$m_i = A(x_A, m_1, m_3, \ldots, m_{i-1}; r_A)$. $m_i = B(x_B, m_0, m_2, \ldots, m_{i-1}; r_B)$.

View of a Party

$$\texttt{View}_A\langle A(x_A), B(x_B)\rangle = (r_A, x_A, m_1, m_3, \ldots)$$

Randomness Private Input received messages

- Party A: adversary.
- Party B is an interactive mechanism whose input is dataset x.
- Since we will only be interested in the adversary's view and the adversary does not have an input, we will drop the subscript and write A's view as $\texttt{View}\langle A, B(x)\rangle$

Formalizing Interactive Differential Privacy

The interactive differentially privacy as **a type of interactive protocol** between an adversary (without any computational limitations) and an interactive mechanism of special properties.

Definition 4 (Interactive Differential Privacy). *A randomized algorithm $\mathcal{M}$ is (ε, δ)-differentially private interactive mechanism if for every pair of adjacent datasets x, x', for every adversary algorithm $\mathcal{A}$, for every possible output set $T \subseteq$ Range $(\texttt{View}\langle\mathcal{A}, \mathcal{M}(\cdot)\rangle)$ we have*

$$\Pr[\texttt{View}\langle\mathcal{A}, \mathcal{M}(x)\rangle \in T] \le e^{\varepsilon} \Pr[\texttt{View}\langle\mathcal{A}, \mathcal{M}(x')\rangle \in T] + \delta$$

Group Privacy-like Bound for Concurrent Composition

Result: The concurrent composition of k (ϵ, δ) interactive DP mechanisms has a group privacy-like bound $(k\epsilon, ke^{k\epsilon}\delta)$.

Proof Idea: Suppose interactive mechanisms $M_0, \ldots, M_{\{k-1\}}$ are each (ϵ_i, δ_i)-differentially private. View A and $M_0, \ldots, M_{\{i-1\}}, M_{\{i+1\}}, \ldots, M_{\{k-1\}}$ as a combined adversary A^*, we can show that

$$\Pr[\texttt{View}\langle\mathcal{A}, \text{ConComp}(\mathcal{M}_0(x'), \ldots, \mathcal{M}_{i-1}(x'), \mathcal{M}_i(x), \ldots, \mathcal{M}_{k-1}(x))\rangle \in S]$$

$$\le e^{\epsilon_i} \Pr[\texttt{View}\langle\mathcal{A}, \text{ConComp}(\mathcal{M}_0(x'), \ldots, \mathcal{M}_{i-1}(x'), \mathcal{M}_i(x'), \ldots, \mathcal{M}_{k-1}(x))\rangle \in S] + \delta_i$$

which can be used for constructing a hybrid argument.

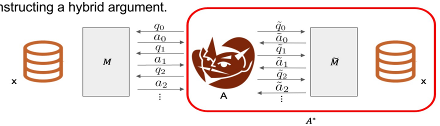

Optimal Concurrent Composition Bound for Pure DP

Result: **Every** interactive $(\varepsilon, 0)$ –DP mechanisms can be simulated by the post-processing of **randomized response** $RR_{(\varepsilon,0)}$ (a non-interactive mechanism).

=> Optimal (approx. DP, Renyi DP, f-DP, etc) bounds for concurrent composition of interactive pure DP mechanisms = optimal bounds for composition of **non-interactive** pure DP mechanisms.

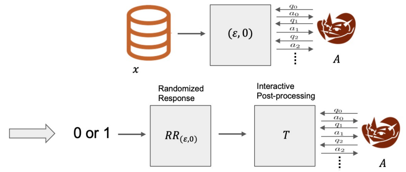

- (ε, δ)-DP version of Randomized Response:

  $\text{RR}_{(\varepsilon,\delta)} : \{0,1\} \to \{0, 1, \texttt{'I am 0'}, \texttt{'I am 1'}\}$

  $\Pr[\text{RR}_{(\varepsilon,\delta)}(0) = \texttt{'I am 0'}] = \delta$ $\quad \Pr[\text{RR}_{(\varepsilon,\delta)}(1) = \texttt{'I am 0'}] = 0$

  $\Pr[\text{RR}_{(\varepsilon,\delta)}(0) = 0] = (1-\delta)\cdot\frac{e^\varepsilon}{1+e^\varepsilon}$ $\quad \Pr[\text{RR}_{(\varepsilon,\delta)}(1) = 0] = (1-\delta)\cdot\frac{1}{1+e^\varepsilon}$

  $\Pr[\text{RR}_{(\varepsilon,\delta)}(0) = 1] = (1-\delta)\cdot\frac{1}{1+e^\varepsilon}$ $\quad \Pr[\text{RR}_{(\varepsilon,\delta)}(1) = 1] = (1-\delta)\cdot\frac{e^\varepsilon}{1+e^\varepsilon}$

  $\Pr[\text{RR}_{(\varepsilon,\delta)}(0) = \texttt{'I am 1'}] = 0$ $\quad \Pr[\text{RR}_{(\varepsilon,\delta)}(1) = \texttt{'I am 1'}] = \delta$

- Post-processing preserves differential privacy
  => To analyze the concurrent composition of arbitrary pure interactive DP mechanisms, it suffices to analyze the composition of randomized responses of the same parameters (analogue to the proof strategy in [KOV15] and [MV16]).
- Therefore, the optimal bound for concurrent composition of pure interactive DP is the same as the optimal bound for composing non-interactive counterpart.

Future Work

- We empirically test whether the Optimal Composition Theorems can be extended to the concurrent composition of approximate DP for 3-message interactive mechanisms with 1-bit message. In all our trials, we find a feasible interactive post-processing algorithm.
- We therefore conjecture that the concurrent composition of interactive DP mechanisms may still have the same bound as the composition for non-interactive DP.