Appendix: Proof of Theorem 6

We begin our proof by introducing a parallel composition theorem of f-DP recently introduced in Smith et al. (2021). This theorem gives a cumulative privacy guarantee on a sequence of analysis on a private dataset, in which each step is done using a new set of samples and is informed by previous steps only through accessing intermediate outputs.

Let $M_1 : \mathbb{Z}^{n_1} \to \mathcal{H}_1$ be the first algorithm and $M_2 : \mathbb{Z}^{n_2} \times \mathcal{H}_1 \to \mathcal{H}_2$ be the second algorithm. $\mathcal{H}_i$ denotes the image space of M_i. The joint algorithm $M : \mathbb{Z}^{n_1+n_2} \to \mathcal{H}_1 \times \mathcal{H}_2$ is defined as

$$M(S_1, S_2) = (h_1, M_2(S_2, h_1)),$$

where $S_1 \in \mathbb{Z}^{n_1}$ and $S_2 \in \mathbb{Z}^{n_2}$ are disjoint datasets and $h_1 = M_1(S_1)$. We can follow this recipe to define an l-fold composed algorithm with disjoint inputs iteratively and get the following privacy guarantee.

Lemma 10 *Let $M_i(\cdot; h_1, \ldots, h_{i-1}) : \mathbb{Z}^{n_i} \to \mathcal{H}_i$ be f_i-DP for all $h_1 \in \mathcal{H}_1, \ldots, h_{i-1} \in \mathcal{H}_{i-1}$. The l-fold composed algorithm with disjoint inputs $M : \mathbb{Z}^n \to \mathcal{H}_1 \times \cdots \times \mathcal{H}_l$ is f-DP, where $n = n_1 + \ldots + n_l$ and $f = \min\{f_1, \ldots, f_l\}^{**}$.*

Intuitively, lemma 10 builds on a combination of parallel composition property and post-processing processing property of differential privacy. Since different algorithm modules use disjoint sets of samples, the whole algorithm uses each single individual observation only once, thereby avoiding the privacy costs of sequential composition.

Our proof of Theorem 6 follows from lemma 10.

Proof Use $\tilde{\mathcal{B}}(\cdot) = \mathcal{B} \circ \mathcal{T}(\cdot; h_1, \ldots, h_k)$ to denote the combined procedure of the data transformation $\mathcal{T}$ followed by algorithm module $\mathcal{B}$. Because $\mathcal{B}$ is f_2-DP and $\mathcal{T}(\cdot; h_1, \cdots, h_k)$ is a 1-stable transformation, i.e. neighboring datasets are still neighbors after transformation for any $h_1 \in \mathcal{H}_1, \ldots, h_k \in \mathcal{H}_k$, $\tilde{\mathcal{B}}(\cdot; h_1, \ldots, h_k)$ is f_2-DP for any $h_1, \ldots, h_k$.

The full algorithm, $\mathcal{M} : \mathcal{S} \mapsto (\mathcal{A}_1(S_{1,1}), \ldots, \mathcal{A}_k(S_{1,k}), \tilde{\mathcal{B}}(S_2; \mathcal{A}_1(S_{1,1}), \mathcal{A}_2(S_{1,2}), \ldots, \mathcal{A}_k(S_{1,k})))$, is a $k+1$-composed algorithm with disjoint inputs. Each algorithm module has its corresponding DP guarantee. The conditions of lemma 10 check. Hence, $\mathcal{M}$ is f-DP with $f = \min\{f_{1,1}, \ldots, f_{1,k}, f_2\}^{**}$. ■