

## CONCURRENT COMPOSITION FOR INTERACTIVE DIFFERENTIAL PRIVACY WITH ADAPTIVE PRIVACY-LOSS PARAMETERS

SAMUEL HANEY<sup>1</sup>, MICHAEL SHOEMATE<sup>2</sup>, GRACE TIAN<sup>2</sup>, SALIL VADHAN<sup>2</sup>, ANDREW VYRROS<sup>2</sup>,  
VICKI XU<sup>2</sup>, AND WANRONG ZHANG<sup>2</sup>

<sup>1</sup>Tumult Labs, Raleigh, USA  
*e-mail address*: sam.haney@tmlt.io

<sup>2</sup>Harvard University, Cambridge, USA  
*e-mail address*: shoematem@seas.harvard.edu, gracetian6@gmail.com, salil\_vadhan@harvard.edu, andrew\_vyrros@g.harvard.edu, vickixu@college.harvard.edu

---

**ABSTRACT.** In this paper, we study the concurrent composition of interactive mechanisms with adaptively chosen privacy-loss parameters. In this setting, the adversary can interleave queries to existing interactive mechanisms, as well as create new ones. We prove that every valid privacy filter and odometer for noninteractive mechanisms extends to the concurrent composition of interactive mechanisms if privacy loss is measured using  $(\epsilon, \delta)$ -DP,  $f$ -DP, or Rényi DP of fixed order. Our results offer strong theoretical foundations for enabling full adaptivity in composing differentially private interactive mechanisms, showing that concurrency does not affect the privacy guarantees. We also provide an implementation for users to deploy in practice.

### 1. INTRODUCTION

**1.1. Differential Privacy.** Differential privacy is a framework for protecting the privacy of individuals when analyzing data. It is a mathematical definition of privacy that ensures that the results of an analysis do not reveal too much information about any individual in the dataset. Because of its powerful worst-case guarantee, differential privacy has become a leading approach in privacy-preserving data analysis, where it is used to enable the analysis of sensitive data while preserving the privacy of individuals.

Differential privacy can be defined in terms of a general database space  $\mathcal{X}$  and a binary neighboring relation on  $\mathcal{X}$ . For example, if databases contain an ordered and known number  $n$  of real-valued entries, then  $\mathcal{X} = \mathbb{R}^n$ . The binary relation on  $\mathcal{X}$  specifies which datasets

---

*Key words and phrases:* Differential Privacy, Interactive Mechanisms, Concurrent Composition, Adaptivity.

\* This article was prepared for TPDP 2023. A preliminary version of this paper has appeared at the ACM Conference on Computer and Communications Security (CCS) 2023.

are *neighboring*, meaning that they differ on one individual’s data. For example, if  $x \in \mathbb{R}^n$ , then  $x, x' \in \mathcal{X}$  are neighboring if they differ on one coordinate. Differential privacy requires that the output distributions should be roughly the same on the two neighboring datasets.

**Definition 1.1** (Differential Privacy). A randomized mechanism  $\mathcal{M} : \mathcal{X} \rightarrow \mathcal{R}$  is  $(\epsilon, \delta)$ -*differentially private* if for every pair of neighboring datasets  $x, x' \in \mathcal{X}$ , and for every subset of possible outputs  $\mathcal{S} \subseteq \mathcal{R}$ ,

$$\Pr[\mathcal{M}(x) \in \mathcal{S}] \leq \exp(\epsilon) \cdot \Pr[\mathcal{M}(x') \in \mathcal{S}] + \delta.$$

By requiring that an analysis be robust to changes in neighboring datasets, differential privacy provides a guarantee that the privacy of individuals in the dataset is protected, regardless of what other data might be included or excluded.

In recent years, other forms of differential privacy have enjoyed use to address various shortcomings of  $(\epsilon, \delta)$ -DP (also known as approximate-DP) regarding composition. Some standard variants include Rényi DP (RDP) [Mironov \(2017\)](#),  $f$ -DP [Dong et al. \(2022\)](#) (the formal definition is given in Section 3), and zero-concentrated differential privacy (zCDP) [Dwork and Rothblum \(2016\)](#); [Bun and Steinke \(2016\)](#).

**Definition 1.2** (Rényi divergence [Rényi \(1961\)](#)). For two probability distributions  $P$  and  $Q$ , the *Rényi divergence* of order  $\alpha > 1$  is

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

**Definition 1.3** (Rényi DP [Mironov \(2017\)](#)). A randomized mechanism  $\mathcal{M}$  is  $(\alpha, \epsilon)$ -*Rényi differentially private* ( $\epsilon$ -RDP $_\alpha$ ) if for every two neighboring datasets  $x$  and  $x'$ ,

$$D_\alpha(\mathcal{M}(x)||\mathcal{M}(x')) \leq \epsilon$$

**1.2. Composition of Differentially Private Mechanisms.** A fundamental question in differential privacy concerns how privacy degrades under multiple mechanisms run on the same database.

**1.2.1. Composition Theorems for Noninteractive Mechanisms.** Definition 1.1 defines differential privacy for *noninteractive mechanisms*  $\mathcal{M}$ . Composition for such noninteractive mechanisms has been extensively studied in the literature. We will denote the (noninteractive, non-adaptive) *composition* of  $k$  noninteractive mechanisms  $\mathcal{M}_1, \dots, \mathcal{M}_k$  on a dataset  $x$  as  $\mathcal{M} := \text{Comp}(\mathcal{M}_1, \dots, \mathcal{M}_k)$ , where  $\mathcal{M}(x)$  returns  $\mathcal{M}_1(x), \dots, \mathcal{M}_k(x)$ , with each  $\mathcal{M}_j$  executed independently on its own random coins.

Basic composition of  $(\epsilon, \delta)$ -differential privacy [Dwork et al. \(2006\)](#) finds that the privacy-loss parameters scale at most linearly with the number of mechanisms composed. Advanced composition [Dwork et al. \(2010b\)](#) provides a tighter bound where the parameters scale sublinearly with the number of mechanisms, and optimal composition [Kairouz et al. \(2017\)](#); [Murtagh and Vadhan \(2016\)](#) provides an exact guarantee of this composition. DP variants such as Rényi DP,  $f$ -DP, and  $z$ -CDP can provide tighter composition bounds by capturing more information about the mechanisms  $\mathcal{M}_i$  being composed than just the two parameters  $\epsilon_i$  and  $\delta_i$ .

1.2.2. *Composition Theorems for Interactive Mechanisms.* While many differential privacy mechanisms are noninteractive, some mechanisms are expressly desired to be interactive, receiving and responding to adaptive queries from analysts. One example is the *adaptive* formulations of the aforementioned composition theorems Vadhan and Wang (2021); Lyu (2022); Vadhan and Zhang (2023). These allow an adversary to select noninteractive mechanisms adaptively based on the result received from  $\mathcal{M}_{i-1}$ . However, there are other, more sophisticated interactive mechanisms that answer queries in a correlated fashion based on secret state or randomness. (In adaptive composition Dwork et al. (2010b), each mechanism  $\mathcal{M}_i$  is run on the dataset using fresh randomness independently of how previous queries are answered.) Other examples include the Sparse Vector Technique Dwork et al. (2009, 2010a); Dwork and Roth (2014), and Private Multiplicative Weights Hardt and Rothblum (2010). Thus, interactive mechanisms have been used as the basic abstraction in the programming frameworks of the open-source software project OpenDP Gaboardi et al. (2020) as well as the Tumult Analytics platform Labs (2022) Berghel et al. (2022). An interactive mechanism  $\mathcal{M}$  is a party interacting with an analyst adversary in an interactive protocol, wherein each party has its *random coin*, which captures the randomness used by the mechanism.

**Definition 1.4** (Interactive Algorithms). An *interactive algorithm*, also known as a *randomized state machine*, consists of a randomized algorithm  $\mathcal{M} : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^* \times \{0, 1\}^*$  that takes the current state  $s \in \{0, 1\}^*$ , a query  $q \in \{0, 1\}^*$ , and returns a new state  $s' \in \{0, 1\}^*$  and an answer  $a \in \{0, 1\}^*$ , written  $(s', a) = \mathcal{M}(s, q)$ . We require that  $s' \neq \lambda$  and  $a \neq \lambda$ , where  $\lambda$  is the empty string. When we wish to make the randomness  $r$  of  $\mathcal{M}$  more explicit, we write  $(s', a) = \mathcal{M}(s, q; r)$ .

An interactive algorithm interacts with an analyst or adversary  $\mathcal{A}$  as follows.

**Definition 1.5** (Interaction between two mechanisms). For two interactive algorithms  $\mathcal{M}$  and  $\mathcal{A}$ , the *interaction* between  $\mathcal{M}$  and  $\mathcal{A}$  on an input  $x \in \{0, 1\}^*$  to  $\mathcal{M}$  is the following random process (denoted  $(\mathcal{A} \leftrightarrow \mathcal{M}(x))$ ):

- (1) Initialize  $s_1^{\mathcal{M}} \leftarrow x$ ,  $m_0 \leftarrow \lambda$ ,  $s_1^{\mathcal{A}} \leftarrow \lambda$ , where  $\lambda$  is the empty string.
- (2) Repeat the following for  $i = 1, 2, \dots$ 
  - (a) If  $i$  is odd, let  $(s_{i+1}^{\mathcal{M}}, m_i) \leftarrow \mathcal{M}(s_i^{\mathcal{M}}, m_{i-1})$  and  $s_{i+1}^{\mathcal{A}} \leftarrow s_i^{\mathcal{A}}$ .
  - (b) If  $i$  is even, let  $(s_{i+1}^{\mathcal{A}}, m_i) \leftarrow \mathcal{A}(s_i^{\mathcal{A}}, m_{i-1})$  and  $s_{i+1}^{\mathcal{M}} \leftarrow s_i^{\mathcal{M}}$ .
  - (c) if  $m_i = \text{halt}$ , then exit loop.

In the context of DP, think of the input  $x$  as a sensitive dataset that needs to be protected. In this context, an *interactive mechanism* is a randomized state machine whose input space upon initialization is the space of datasets  $\mathcal{X}$ , and takes in a database  $x \in \mathcal{X}$ . An *adversary* is a randomized state machine that takes in an empty string. The view of an adversary captures everything the adversary receives during the execution.

**Definition 1.6** (View of the adversary in an interactive mechanism). Let  $\mathcal{M}$  be an interactive mechanism and  $\mathcal{A}$  be an adversary interacting with the mechanism.  $\mathcal{A}$ 's *view* of  $(\mathcal{A}, \mathcal{M}(x))$  is the tuple  $\text{View}_{\mathcal{A}}(\mathcal{A} \leftrightarrow \mathcal{M}(x)) = (r_0, m_1, r_2, m_3, r_4, \dots)$  consisting of all the messages  $m_i$  received by  $\mathcal{A}$  from  $\mathcal{M}$  together with random coins  $r_i$  that  $\mathcal{A}$  tosses when computing  $m_i$  (i.e. for even  $i$ ,  $(s_{i+1}^{\mathcal{A}}, m_i) = \mathcal{A}(s_i^{\mathcal{A}}, m_{i-1}; r_i)$ ).

For shorthand, we will drop the subscript  $\mathcal{A}$  when referring to **View** for the rest of this paper, because future references to **View** in the rest of this paper always concern the adversary's view.

We can define interactive differential privacy based on measuring the same  $(\epsilon, \delta)$ -closeness, as in noninteractive differential privacy, between the views of the adversary on two neighboring datasets.

**Definition 1.7** ( $(\epsilon, \delta)$ -DP interactive mechanisms). An interactive mechanism  $\mathcal{M}$  is an  $(\epsilon, \delta)$ -*differentially private interactive mechanism*, or *I.M.* for short, if for every pair of neighboring datasets  $x, x' \in \mathcal{X}$ , every interactive adversary algorithm  $\mathcal{A}$ , and every subset of possible views  $\mathcal{S} \subseteq \text{Range}(\text{View})$  we have

$$\Pr[\text{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x)) \in \mathcal{S}] \leq e^\epsilon \Pr[\text{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x')) \in \mathcal{S}] + \delta.$$

To capture variants such as Rényi DP and  $f$ -DP, we will define a broader version of I.M.s based on a generalized notion of differential privacy in Section 2.

When we consider the composition of interactive mechanisms, it is straightforward to extend the composition theorems for noninteractive mechanisms to the *sequential composition* of interactive mechanisms. A subtler case is the *concurrent composition*, in which an adversary can interleave queries to multiple mechanisms concurrently, first studied by Vadhan and Wang [Vadhan and Wang \(2021\)](#). The queries can therefore depend on the answers received from other mechanisms. This is useful in many practical settings. For example, an analyst may run data analyses using multiple interactive analyses on the same dataset simultaneously and the queries in multiple analyses might be correlated. Concurrent composition, formally defined in Section 2.1, maintains  $k$  interactive mechanisms  $\mathcal{M}_1, \dots, \mathcal{M}_k$ . It is itself an interactive mechanism and the query it received from an adversary is of the form of  $(j, q)$ , meaning it issues a standard query  $q$  to  $\mathcal{M}_j$ . Concurrent composition theorems allow us to understand the privacy guarantee of the overall analysis given the privacy guarantees for each interactive analysis when they are executed independently.

Previous work provides concurrent composition theorems for several different types of differential privacy. Vadhan and Wang [Vadhan and Wang \(2021\)](#) shows that every composition theorem for noninteractive  $\epsilon$ -DP mechanisms extends to concurrent composition for interactive  $\epsilon$ -DP mechanisms. Lyu [Lyu \(2022\)](#) and Vadhan and Zhang [Vadhan and Zhang \(2023\)](#) generalize this result to  $(\epsilon, \delta)$ -DP when  $\delta > 0$ .

**Theorem 1.8** ([Lyu \(2022\)](#); [Vadhan and Zhang \(2023\)](#)). Suppose that for all noninteractive mechanisms  $\mathcal{M}_1, \dots, \mathcal{M}_k$  such that  $\mathcal{M}_i$  is  $(\epsilon_i, \delta_i)$ -DP for  $i = 1, \dots, k$ , their composition  $\text{Comp}(\mathcal{M}_1, \dots, \mathcal{M}_k)$  is  $(\epsilon, \delta)$ -DP. Then for all interactive mechanisms  $\mathcal{M}_1, \dots, \mathcal{M}_k$  with finite communication<sup>1</sup> such that  $\mathcal{M}_i$  is  $(\epsilon_i, \delta_i)$ -DP for  $i = 1, \dots, k$ , their concurrent composition  $\text{ConComp}(\mathcal{M}_1, \dots, \mathcal{M}_k)$ , as defined formally in Algorithm 2.4, is  $(\epsilon, \delta)$ -DP.

This composition theorem also extends to  $f$ -DP [Vadhan and Zhang \(2023\)](#). Moreover, Lyu [Lyu \(2022\)](#) shows that the privacy adds up under concurrent composition for any fixed order of  $\alpha > 1$  for Rényi DP (RDP).

**Theorem 1.9** ([Lyu \(2022\)](#)). For all  $\alpha > 1$ ,  $k \in \mathbb{N}$ ,  $\epsilon_1, \dots, \epsilon_k > 0$ , and all interactive mechanisms  $\mathcal{M}_1, \dots, \mathcal{M}_k$  such that  $\mathcal{M}_i$  is  $(\alpha, \epsilon_i)$ -RDP for  $i = 1, 2, \dots, k$ , their concurrent composition  $\text{ConComp}(\mathcal{M}_1, \dots, \mathcal{M}_k)$ , as defined formally in Algorithm 2.4, is  $(\alpha, \sum_{i=1}^k \epsilon_i)$ -RDP.

In all of the above cases, privacy-loss parameters are set upfront prior to data analysis and are fixed for all queries or computations performed on the data. However, in practice,

<sup>1</sup>Their proof relies on an induction argument on the number of messages exchanged, which requires an assumption of *finite communication*.

data analysts may not know in advance what they want to do with the data and may want to adaptively choose the privacy-loss parameters of subsequent mechanisms as they go along. This can lead to more efficient use of privacy resources. For example, we may wish to spend more privacy-loss budget on fine-grained analysis, and less privacy-loss budget on exploratory analysis. Therefore, in practice, allowing *full adaptivity* of mechanisms and privacy-loss parameters themselves and also adaptive choice of the length of composition as a function of intermediate analyses is important. This consideration motivates the study of privacy *filters* and *odometers*, which we discuss in the following section.

**1.3. Odometers and Filters.** Rogers, Roth, Ullman, and Vadhan [Rogers et al. \(2016\)](#) define two primitives for the adaptive composition of DP mechanisms, privacy *filters* and privacy *odometers*, to allow for the ability to track privacy loss during an interaction. Their definitions were given specifically for  $(\epsilon, \delta)$ -DP; here we follow Lécuyer [Lécuyer \(2021\)](#) and work with a more general formalism that applies to arbitrary privacy measures, including  $f$ -DP and Rényi-DP. In particular, our definition of privacy *odometers* is different from the odometers previously defined in [Rogers et al. \(2016\)](#).

A privacy *filter* is a mechanism that halts computation on a dataset once a preset privacy-loss budget is exceeded. It takes a global privacy-loss budget as an input and is equipped with a continuation rule that halts computation whenever the budget is exceeded. At each round, the continuation rule takes all privacy-loss parameters up to the current round, and outputs either `continue` or `halt`. If `continue`, the mechanism answers the query with the current privacy parameter. Once the mechanism outputs `halt`, no further computation is allowed. (An equivalent alternative is to refuse to answer the correct query, but allow the analyst to try again with new queries. However, the `halt` formulation is more convenient for presentation in this paper. For notational convenience in this paper, our algorithm pseudocode will have the filter go into a looping state that always returns the same state and a dummy message once `halt` is reached.) The filter guarantees that the interaction is differentially private according to the desired privacy-loss budget.

Existing work on privacy filters and odometers has only studied cases in which the adversary adaptively chooses the next mechanism to compose with, but the mechanisms themselves are non-interactive (so there is no notion of concurrent composition, since each mechanism returns only one answer).

**Definition 1.10** ( $\mathcal{F}$ -filtered composition of noninteractive  $(\epsilon, \delta)$ -DP mechanisms ( $\mathcal{F}$ -Filt( $NIM$ ))).

Let  $\mathcal{F}$  be a continuation rule that takes in a sequence of privacy-loss parameters  $(\epsilon_1, \delta_1), \dots$ , and a target privacy-loss budget  $(\epsilon, \delta)$ , and maps to a binary decision:  $\mathcal{F} : (\mathbb{R}_{\geq 0})^2 \times (\mathbb{R}_{\geq 0})^2 \rightarrow \{0, 1\}$ , where 1 means `continue` and 0 means `halt`, and  $\mathbb{R}_{\geq 0}^* = \cup_{k=0}^{\infty} \mathbb{R}^k$ . The  $\mathcal{F}$ -filtered composition of noninteractive mechanisms, denoted as  $\mathcal{F}$ -Filt( $NIM$ ), is an interactive mechanism.  $\mathcal{F}(\cdot; (\epsilon, \delta))$ -Filt( $NIM$ )( $s, m$ ) is defined as follows:

- (1) If  $m = \lambda$ , let  $s \leftarrow x$ , and set  $s' \leftarrow (s, [])$ , where  $[]$  is an empty list. Return  $(s', \lambda)$ , where  $\lambda$  is an empty string.
- (2) Parse  $s$  as  $s = (x, [(\epsilon_1, \delta_1), \dots, (\epsilon_k, \delta_k)])$
- (3) If  $\mathcal{M}_{k+1}$  is  $(\epsilon_{k+1}, \delta_{k+1})$ -DP and  $\mathcal{F}((\epsilon_1, \delta_1), \dots, (\epsilon_{k+1}, \delta_{k+1}); (\epsilon, \delta)) = 1$ :
  - (a) Let  $s' \leftarrow (x, [(\mathcal{M}_1, (\epsilon_1, \delta_1)), \dots, (\mathcal{M}_{k+1}, (\epsilon_{k+1}, \delta_{k+1}))])$
  - (b) Let  $m' \leftarrow \mathcal{M}_{k+1}(x)$
- (4) Else, let  $s' \leftarrow \text{Halt}$ , and let  $m'$  store the reason for halting
- (5) Return  $(s', m')$

Note that at each round  $k$  of computation, whether the mechanism  $\mathcal{M}_k$  is  $(\epsilon_k, \delta_k)$ -DP is something the implementation needs to be able to verify, e.g. by having verified privacy-loss parameters attached to every mechanism.

**Definition 1.11** (Valid  $(\epsilon, \delta)$ -DP filter for noninteractive mechanisms). We say a continuation rule  $\mathcal{F}$  is a valid  $(\epsilon, \delta)$ -DP NIM-filter for noninteractive mechanisms if for every pair of  $(\epsilon, \delta)$ ,  $\mathcal{F}(\cdot; (\epsilon, \delta))\text{-Filt}(\text{NIM})(\cdot)$  is an  $(\epsilon, \delta)$ -DP interactive mechanism.

We define privacy *odometer* as a mechanism that allows the analyst to keep track of the privacy loss at each step of computation.

**Definition 1.12** (Odometers). An odometer  $\mathcal{O} : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^* \times \{0, 1\}^*$  is an I.M. that supports a specially-designated *deterministic privacy-loss query*, denoted `privacy_loss`, that is deterministic and does not change state. Upon receiving a privacy-loss query,  $\mathcal{O}$  will output the upper bound of the privacy loss at that point in computation.

We note that our terminology is slightly different from previous work by Rogers et al. Rogers et al. (2016), where they refer to the privacy-loss accumulator as the odometer, whereas our odometer is a mechanism that has a specially-designated privacy loss query. When posing this query, the mechanism outputs the current privacy loss up to that point.

We will stipulate that `privacy_loss` queries are encoded as binary strings in a fixed and recognizable way, such as using all strings that begin with a 1 to be `privacy_loss` queries, and all strings that begin with a 0 to be ordinary queries that can then be further parsed.

To measure the privacy loss at each round, we use a *truncated view* of an adversary  $\mathcal{A}$  interacting with an odometer  $\mathcal{O}$  as defined in Definition 1.13.

**Definition 1.13** (View between adversary  $\mathcal{A}$  with approximate-DP odometer  $\mathcal{O}$  truncated at  $(\epsilon, \delta)$ ). Given a privacy-loss parameter  $(\epsilon, \delta)$ , an adversary  $\mathcal{A}$ , and approximate-DP odometer  $\mathcal{O}$ ,  $\text{Trunc}_{(\epsilon, \delta)}(\text{View}(\mathcal{A} \leftrightarrow \mathcal{O}(x)))$  is constructed as follows:

- (1) Initialize  $s_1^{\mathcal{M}} \leftarrow x$ ,  $q_0 \leftarrow \lambda$ ,  $s_1^{\mathcal{A}} \leftarrow \lambda$ , where  $\lambda$  is the empty string.
- (2) Repeat the following for  $i = 1, 2, \dots$ 
  - (a) If  $i$  is odd, let  $(s_{i+1}^{\mathcal{M}}, a_i) \leftarrow \mathcal{O}(s_i^{\mathcal{M}}, q_{i-1})$ , and let  $(s_i^{\mathcal{M}}, (\epsilon_i, \delta_i)) \leftarrow \mathcal{O}(s_i^{\mathcal{M}}, \text{privacy\_loss})$ ,  $s_{i+1}^{\mathcal{A}} \leftarrow s_i^{\mathcal{A}}$ , with the assumption that `privacy_loss` queries do not change state.
  - (b) If  $\epsilon_i \geq \epsilon$  or  $\delta_i \geq \delta$ , then exit loop.
  - (c) If  $i$  is even, let  $(s_{i+1}^{\mathcal{A}}, q_i) \leftarrow \mathcal{A}(s_i^{\mathcal{A}}, a_{i-1}; r_i^{\mathcal{A}})$  and  $s_{i+1}^{\mathcal{M}} \leftarrow s_i^{\mathcal{M}}$ .
  - (d) If  $a_i = \text{halt}$ , then exit loop.
- (3) Return  $\text{Trunc}_{(\epsilon, \delta)}(\text{View}(\mathcal{A} \leftrightarrow \mathcal{O}(x))) = (r_0^{\mathcal{A}}, a_1, \dots, r_{i-2}^{\mathcal{A}}, a_{i-1})$ , where  $i$  is the index that the for loop was halted on.

Using truncated view for analyzing odometers is first introduced by Lécuyer Lécuyer (2021) for RDP. We extend their definition to a general notion of truncated view, which allows us to quantify privacy loss with all other DP variants, where we can simply replace  $(\epsilon, \delta)$  with other privacy-loss parameters for other DP variants. In contrast, the previous definition of odometer in Rogers et al. Rogers et al. (2016) only holds for  $(\epsilon, \delta)$ -DP.

Then the privacy guarantee of an odometer is defined as measuring the closeness between the distributions of the truncated views of an adversary.

**Definition 1.14** (Valid approximate-DP odometer for noninteractive mechanisms). We say  $\mathcal{O}$  is a valid approximate-DP *odometer* if for every pair of  $(\epsilon, \delta)$ , every adversary  $\mathcal{A}$ , and

every pair of adjacent datasets  $x, x'$ ,

$$\begin{aligned} & D(\text{Trunc}_{(\epsilon, \delta)}(\text{View}(\mathcal{A} \leftrightarrow \mathcal{O}(x))) || \\ & \text{Trunc}_{(\epsilon, \delta)}(\text{View}(\mathcal{A} \leftrightarrow \mathcal{O}(x')))) \preceq (\epsilon, \delta), \end{aligned}$$

where  $\preceq$  is the partial order defined by  $(\epsilon_1, \delta_1) \preceq (\epsilon_2, \delta_2)$  iff  $\epsilon_1 \leq \epsilon_2$  and  $\delta_1 \leq \delta_2$ .

A common example of an odometer is an I.M. equipped with a privacy-loss accumulator  $\mathcal{G}$ , which is a function that gives an upper bound on the accumulated privacy loss at each step.

**Definition 1.15** (Valid approximate-DP privacy-loss accumulator for noninteractive mechanisms). We say  $\mathcal{G}$  is a *valid approximate-DP NIM-privacy-loss accumulator* if  $\mathcal{G}\text{-Odom}(NIM)$  is a valid approximate-DP odometer.

We term these odometers built from privacy-loss accumulators  $\mathcal{G}$ -odometers. (However, a more general odometer need not be based on a privacy-loss accumulator.) Below, we provide a definition sketch for a  $\mathcal{G}$ -odometer, which is further defined in Section 2.3.

**Definition 1.16** (Definition sketch for  $\mathcal{G}$ -odometer for composition of noninteractive mechanisms ( $\mathcal{G}\text{-Odom}(NIM)$ )). Let  $\mathcal{G}$  be a privacy-loss accumulator that takes privacy-loss parameters  $(\epsilon_1, \delta_1), \dots, (\epsilon_k, \delta_k)$  in a sequence, and maps to a global privacy loss  $(\epsilon, \delta)$ , so  $\mathcal{G} : (\mathbb{R}_{\geq 0}^2)^k \rightarrow (\mathbb{R}_{\geq 0}^2)$ , for every  $k = 1, 2, \dots$ . A  $\mathcal{G}$ -odometer for composition of noninteractive mechanisms is denoted as  $\mathcal{G}\text{-Odom}(NIM)$ .  $\mathcal{G}\text{-Odom}(NIM)(s, m)$  is executed as follows:

- (1) If  $m = \lambda$ , let  $x \leftarrow s$ , set  $s' \leftarrow (x, [])$ , where  $[]$  is an empty list, and return  $(s', \lambda)$
- (2) Parse  $s$  as  $s = (x, [(\epsilon_1, \delta_1), \dots, (\epsilon_k, \delta_k)])$
- (3) If  $m$  is of the form  $(\mathcal{M}_{k+1}, (\epsilon_{k+1}, \delta_{k+1}))$  and  $\mathcal{M}_{k+1}$  is  $(\epsilon_{k+1}, \delta_{k+1})$ -DP:
  - (a) Let  $s' \leftarrow (x, [(\mathcal{M}_1, (\epsilon_1, \delta_1)), \dots, (\mathcal{M}_{k+1}, (\epsilon_{k+1}, \delta_{k+1}))])$
  - (b) Let  $m' \leftarrow \mathcal{M}_{k+1}(x)$
- (4) If  $m = \text{privacy\_loss}$ , then let  $s' \leftarrow s$  and  $m' \leftarrow \mathcal{G}((\epsilon_1, \delta_1), \dots, (\epsilon_k, \delta_k))$
- (5) Return  $(s', m')$ .

A number of previous works construct privacy filters and odometers for noninteractive mechanisms. The original odometer definition proposed by Rogers et al. [Rogers et al. \(2016\)](#) was defined specifically for  $(\epsilon, \delta)$ -DP and requires a simultaneous guarantee about the privacy loss being bounded at all points *in time* when composing with adaptive privacy-loss parameters. Specifically, the privacy loss is defined as  $\text{Loss}(v) = \log \left( \frac{\Pr[\text{View}(\mathcal{A} \leftrightarrow \mathcal{O}(x))=v]}{\Pr[\text{View}(\mathcal{A} \leftrightarrow \mathcal{O}(x'))=v]} \right)$ , and the odometer is defined as follows.

**Definition 1.17** (Previous definition of Odometers in [Rogers et al. \(2016\)](#)). We say  $\mathcal{G}$  is a valid approximate DP *odometer* if for every adversary  $\mathcal{A}$  in  $k$ -fold composition (as in, composition of  $k$  mechanisms), with adaptive privacy-loss parameters, and every pair of adjacent datasets  $x, x'$ , the following holds with probability at most  $\delta_g$  over  $v \leftarrow \text{View}(\mathcal{A} \leftrightarrow \mathcal{O}(x))$

$$|\text{Loss}(v)| > \mathcal{G}(\epsilon_1, \delta_1, \dots, \epsilon_k, \delta_k).$$

Note that this definition yields a form of approximate DP in its final guarantee, because there is always a  $\delta_g$  probability of failure.

Whitehouse, Ramdas, Rogers, and Wu [Whitehouse et al. \(2023\)](#) gave constructions of odometers and filters that quantitatively improve on the results presented in Rogers et al. [Rogers et al. \(2016\)](#), including when the composed mechanisms satisfy zCDP or Rényi

DP (but their odometer only guarantees the same,  $(\epsilon, \delta)$  property of Definition 1.17). For providing Rényi DP guarantees when composing Rényi DP mechanisms, Feldman and Zrnic [Feldman and Zrnic \(2022\)](#) constructed privacy filters. Definitions and constructions of privacy odometers for Rényi DP were given by Lécuyer [Lécuyer \(2021\)](#). Our general definition of odometers follows Lécuyer’s, which we find more natural and general, since it is applicable to all forms of DP. This definition can be thought of as being in analogy to the definition of a  $p$ -value in statistics. A  $p$ -value has an interpretation if we set a significance level  $\alpha$  before we observe our random events, and then reject the null hypothesis if the  $p$ -value is larger than  $\alpha$ . Similarly the privacy-loss reported by our odometer definition has an interpretation if we set a privacy-loss threshold  $(\epsilon, \delta)$  before the mechanism is executed and halt when the odometer’s privacy-loss query is not going to be smaller than  $(\epsilon, \delta)$ .

**1.4. Our Results on Concurrent Filters and Odometers.** In this paper, we consider privacy filters and odometers for composing *interactive* mechanisms. Previous work ([Vadhan and Wang \(2021\)](#); [Lyu \(2022\)](#); [Vadhan and Zhang \(2023\)](#)) has studied the concurrent composition of interactive mechanisms with pre-specified privacy-loss parameters, and in this paper we extend it to the filter and odometer settings where the privacy-loss parameters, and not just the mechanisms, can be chosen adaptively. We analogously define  $\mathcal{F}$ -filtered sequential composition of interactive mechanisms, or  $\mathcal{F}$ -*FiltSeq(IM)* for short, and  $\mathcal{F}$ -filtered concurrent composition of interactive mechanisms, or  $\mathcal{F}$ -*FiltCon(IM)* for short. The continuation rule  $\mathcal{F}$  here determines whether the adversary can create new interactive mechanisms with adaptively-chosen privacy-loss parameters. Similarly, we define  $\mathcal{G}$ -odometer for the sequential composition of interactive mechanisms, or  $\mathcal{G}$ -*OdomSeq(IM)* for short, and  $\mathcal{G}$ -odometer for the concurrent composition of interactive mechanisms, or  $\mathcal{G}$ -*OdomCon(IM)* for short. It should be noted that in both our filters and odometers, the “privacy-loss budget” is paid at the launch of each additional interactive mechanism. In particular, any `privacy_loss` queries asked to the odometer between mechanism launches will return the same value.

As composition theorems for noninteractive mechanisms typically can be easily extended to the sequential composition of interactive mechanisms, we focus on the case of concurrent composition, as defined formally in Definition 2.4. In this setting, the adversary can interleave queries to existing interactive mechanisms. In the filter and odometer case, the adversary can also create new interactive mechanisms with adaptively-chosen privacy-loss parameters (for new I.M. that are filters, with the added stipulation that the privacy loss with the new interactive mechanism added does not exceed the privacy budget). The complex dependence between the adversary’s interacting with the different mechanisms makes it non-trivial to extend filters and odometers to this case. Nevertheless, we prove that every valid privacy filter and odometer for noninteractive mechanisms extends to a privacy filter and odometer for interactive mechanisms for  $(\epsilon, \delta)$ -DP,  $f$ -DP, and RDP.

**Definition 1.18** (Valid  $(\epsilon, \delta)$ -DP filter for interactive mechanisms). We say a continuation rule  $\mathcal{F}$  is a *valid  $(\epsilon, \delta)$ -DP IM-filter* for interactive mechanisms if  $\mathcal{F}(\cdot; (\epsilon, \delta))$ -*FiltCon(IM)* is an  $(\epsilon, \delta)$ -DP interactive mechanism.

**Definition 1.19** (Valid  $(\epsilon, \delta)$ -DP privacy-loss accumulator for interactive mechanisms). We say  $\mathcal{G}$  is a *valid approx-DP IM-privacy-loss accumulator* if  $\mathcal{G}$ -*OdomCon(IM)* is a valid approx-DP odometer.

**Theorem 1.20** ( $(\epsilon, \delta)$ -DP filters and privacy-loss accumulators).

- (1)  $\mathcal{F}$  is a valid  $(\epsilon, \delta)$ -DP NIM-filter if and only if  $\mathcal{F}$  is a valid  $(\epsilon, \delta)$ -DP IM-filter.
- (2)  $\mathcal{G}$  is a valid  $(\epsilon, \delta)$ -DP NIM-privacy-loss accumulator if and only if  $\mathcal{G}$  is a valid  $(\epsilon, \delta)$ -DP IM-privacy-loss accumulator.

**Theorem 1.21** (*f*-DP filters and privacy-loss accumulators).

- (1)  $\mathcal{F}$  is a valid *f*-DP NIM-filter if and only if  $\mathcal{F}$  is a valid *f*-DP IM-filter.
- (2)  $\mathcal{G}$  is a valid *f*-DP NIM-privacy-loss accumulator if and only if  $\mathcal{G}$  is a valid *f*-DP IM-privacy-loss accumulator.

**Theorem 1.22** (RDP filters and privacy-loss accumulators).

- (1)  $\mathcal{F}(\epsilon_1, \epsilon_2, \dots; \epsilon) = \mathbb{I}(\sum_i \epsilon_i \leq \epsilon)$  is a valid  $(\alpha, \epsilon)$ -RDP IM-filter for every fixed order of  $\alpha > 1$ .
- (2)  $\mathcal{G}(\epsilon_1, \dots, \epsilon_i) = \sum_{i=1}^i \epsilon_i$  is a valid  $(\alpha, \epsilon)$ -RDP IM-privacy loss accumulator for every fixed order of  $\alpha > 1$ .

These theorems offer strong theoretical foundations for enabling *full adaptivity* in composing differentially private interactive mechanisms. Firstly, they allow data analysts to adaptively choose privacy loss parameters as well as the subsequent interactive analyses. By adjusting the strength of the privacy guarantee to reflect the actual needs of the analysis as they go along, this important feature allows us to optimize the trade-off between privacy and utility. Second, our results are particularly useful for DP libraries. For example, OpenDP<sup>2</sup> and Tumult Labs<sup>3</sup> use interactive mechanisms as a core abstraction for building differentially-private algorithms. Prior to our work, supporting adaptive selection of privacy-loss parameters has necessitated enforcing sequentiality on the use of interactive mechanisms. This restriction makes the libraries less user-friendly, as we cannot allow analysts or even DP programs to interact with multiple interactive mechanisms simultaneously. Therefore, there is a gap between the existing and desired functionality. It also increases complexity in the libraries, necessitating a control system to prevent interleaving queries. Given the results in our paper, we can now remove this interleaving restriction, while maintaining the *full adaptivity* feature. We provide further discussion and our implementation for privacy filters and odometers for interactive mechanisms in Section 5.

**1.5. Proof Strategy.** In this section, we explain our proof strategy, and we defer the detailed proof to Section 3 and Section 4.

To prove the above theorems, our strategy is to leverage interactive postprocessing from interactive mechanisms to interactive mechanisms. We note that previous work [Vadhan and Zhang \(2023\)](#) only considers interactive postprocessing from non-interactive mechanisms to interactive mechanisms. Here, we provide a general formulation called “person-in-the-middle”, which is formally defined in Definition 2.15.

A *person-in-the-middle (PIM) mechanism* is a randomized mechanism that acts as an interlocutor between two interacting mechanisms  $\mathcal{A}$  and  $\mathcal{M}$ . When  $\mathcal{A}$  attempts to send a message to  $\mathcal{M}$ , the PIM mechanism can undergo an interaction with  $\mathcal{A}$  until finally passing on the message at the end of this interaction to  $\mathcal{M}$ . Similarly, when  $\mathcal{M}$  attempts to send a message to  $\mathcal{A}$ , the PIM mechanism can interact with  $\mathcal{M}$  to modify the message before passing it to  $\mathcal{A}$ .

<sup>2</sup><https://opendp.org>

<sup>3</sup><https://www.tmlt.io>

This construction of the postprocessing by a PIM mechanism allows us the following theorem.

**Theorem 1.23** (Privacy preserved under postprocessing PIM mechanism, informally stated). If  $\mathcal{P}$  is an interactive postprocessing mechanism and  $\mathcal{M}$  is a differentially-private interactive mechanism with respect to any privacy measure (such as  $(\epsilon, \delta)$ -DP,  $f$ -DP, and Rényi-DP), then  $\mathcal{P} \circ \mathcal{M}$  is also differentially-private with the same privacy-loss parameters.

A useful corollary follows from Theorem 1.23.

**Corollary 1.24.** Suppose  $\mathcal{N}$  is an interactive mechanism over database space  $\mathcal{X}$  such that for every pair of neighboring datasets  $x, x'$  and every deterministic adversary  $\mathcal{A}$ , there exists an  $(\epsilon, \delta)$ -DP I.M.  $\mathcal{M}$  and an interactive postprocessing  $\mathcal{P}$  such that

$$\begin{aligned} \text{View}(\mathcal{A} \leftrightarrow \mathcal{N}(x)) &\equiv \text{View}(\mathcal{A} \leftrightarrow \mathcal{P}(\mathcal{M}(x))) \\ \text{View}(\mathcal{A} \leftrightarrow \mathcal{N}(x')) &\equiv \text{View}(\mathcal{A} \leftrightarrow \mathcal{P}(\mathcal{M}(x'))) \end{aligned}$$

Then  $\mathcal{N}$  is  $(\epsilon, \delta)$ -DP.

We will use this corollary to prove Theorem 1.21. Our proof also relies on the following reduction theorem Vadhan and Zhang (2023) showing that every interactive  $f$ -DP mechanism can be simulated by an interactive postprocessing of a noninteractive mechanism, for a fixed pair of neighboring datasets  $x, x'$ .

**Theorem 1.25** (Vadhan and Zhang (2023)). For every trade-off function  $f$ , every interactive  $f$ -DP mechanism  $\mathcal{M}$  with finite communication, and every pair of neighboring datasets  $x, x'$ , there exists a noninteractive  $f$ -DP mechanism  $\mathcal{N}$  and a randomized interactive postprocessing mechanism  $\mathcal{P}$  such that for every adversary  $\mathcal{A}$ , we have

$$\begin{aligned} \text{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x)) &\equiv \text{View}(\mathcal{A} \leftrightarrow \mathcal{P}(\mathcal{N}(x))) \\ \text{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x')) &\equiv \text{View}(\mathcal{A} \leftrightarrow \mathcal{P}(\mathcal{N}(x'))) \end{aligned}$$

Theorem 1.25 allows us to reduce the composition of interactive mechanisms  $\mathcal{M}$  to the composition of noninteractive mechanisms  $\mathcal{N}$ , as needed to prove Theorem 1.21. Fixing adjacent datasets  $x, x'$ , we can then define an interactive postprocessing  $\mathcal{P}$  such that for every deterministic adversary  $\mathcal{A}$ ,  $\text{View}(\mathcal{A} \leftrightarrow \mathcal{P} \circ \mathcal{F}\text{-Filt}(\mathcal{NIM})) \equiv \text{View}(\mathcal{A} \leftrightarrow \mathcal{F}\text{-FiltCon}(\mathcal{IM}))$  on any database  $x$ . Therefore we can use composition theorems for noninteractive  $f$ -DP mechanisms to construct concurrent  $f$ -DP filters. Theorem 1.20 follows directly from Theorem 1.21, as  $f$ -DP captures  $(\epsilon, \delta)$ -DP as a special case. In fact, the special case of Theorem 1.25 for  $(\epsilon, \delta)$ -DP was shown independently by Lyu Lyu (2022).

Our proof of Theorem 1.22 for RDP follows a different approach. Our proof relies on the concurrent composition theorem in Theorem 1.9 for two interactive RDP mechanisms, and our strategy is to apply induction on the number of mechanisms being composed.

We will make use of a *universal mechanism*  $\mathcal{U}_{\alpha, \epsilon}$  with parameters  $(\alpha, \epsilon)$  being a general mechanism that can simulate any valid  $(\alpha, \epsilon)$ -RDP mechanism (which is specified as the first query to the universal mechanism).

When composing two mechanisms (i.e.  $K = 2$  in Theorem 1.22), by the concurrent composition theorem in Theorem 1.9, an adversary can only start a second mechanism  $\mathcal{M}_2$  after  $\mathcal{M}_1$  if and only if its privacy-loss parameter  $\epsilon_2$  is at most  $\epsilon - \epsilon_1$ , where  $\epsilon$  is the target privacy-loss budget and  $\epsilon_1$  is the privacy-loss parameter for  $\mathcal{M}_1$ . We define the following

filter

$$\mathcal{F}_2(\epsilon_1, \dots, \epsilon_k; \epsilon) = \begin{cases} \mathbb{I}(\sum_{i=1}^k \epsilon_i \leq \epsilon) & \text{if } k \leq 2 \\ 0 & \text{otherwise.} \end{cases}$$

$\mathcal{F}_2$  is a valid  $(\alpha, \epsilon)$ -RDP IM-filter for composing two mechanisms. This is because if we fix a deterministic adversary  $\mathcal{A}$  as Lemma 2.3 allows us to do, then  $\mathcal{M}_1, \epsilon_1$  are predetermined adaptively, which means that the adversary interacting with  $\mathcal{F}_2\text{-FiltCon}(IM)$  is equivalent to the adversary interacting with  $\text{ConComp}(\mathcal{M}_1, \mathcal{U}_{\alpha, \epsilon - \epsilon_1})$ , where  $\mathcal{U}_{\alpha, \epsilon - \epsilon_1}$  is the universal  $(\alpha, \epsilon - \epsilon_1)$ -RDP mechanism, and the adversary's first query to  $\mathcal{U}_{\alpha, \epsilon - \epsilon_1}$  is an  $(\alpha, \epsilon_2)$ -RDP mechanism  $\mathcal{M}_2$ . This allows  $\mathcal{U}_{\alpha, \epsilon - \epsilon_1}$  to serve as a “placeholder” that can simulate an adaptively chosen  $\mathcal{M}_2$ . Upon verifying  $\epsilon_2 \leq \epsilon - \epsilon_1$ ,  $\mathcal{U}_2$  interacts with the adversary just as  $\mathcal{M}_2(x)$  does. We then apply an induction argument on the number of mechanisms being composed. For  $K > 2$ , we consider  $\mathcal{F}_K\text{-FiltCon}(IM)$ , which is a filter for composing up to  $K$  interactive mechanisms rather than two mechanisms. Assuming the privacy of  $\mathcal{F}_{K-1}\text{-FiltCon}(IM)$ , we argue that we can construct a postprocessing  $\mathcal{P}$  interacting with  $\mathcal{F}_2\text{-FiltCon}(IM)$  such that an adversary's interaction with  $\mathcal{P} \circ \mathcal{F}_2\text{-FiltCon}(IM)$  is equivalent to interacting with  $\mathcal{F}_K\text{-FiltCon}(IM)$ , allowing us to proceed by induction on  $K$ .

Crucially, the proof for Rényi DP uses the fact that Rényi DP of a fixed order  $\alpha$  is measured by a single real number  $\epsilon$ , and the optimal composition theorem is additive, i.e.  $\epsilon = \epsilon_1 + \dots + \epsilon_K$ . This is not satisfied in  $(\epsilon, \delta)$ -DP or  $f$ -DP, hence we use a different strategy in Theorem 1.21, as discussed above.

Feldman and Zrnic [Feldman and Zrnic \(2022\)](#) prove that  $\mathcal{F}(\epsilon_1, \epsilon_2, \dots; \epsilon) = \mathbb{I}(\sum_i \epsilon_i \leq \epsilon)$  is a valid RDP NIM-filter for every fixed order of  $\alpha > 1$ . They construct a supermartingale for the privacy loss, and then apply the optional stopping theorem to bound the overall privacy loss when the algorithm halts. The main idea of our proof strategy is to leverage interactivity, which allows us to proceed by induction, describing  $\mathcal{F}_k$  in terms of  $\mathcal{F}_{k-1}$ . This strategy can significantly simplify the proof in [Feldman and Zrnic \(2022\)](#) of the privacy filter for even noninteractive RDP mechanisms, by similarly doing induction that uses the sequential composition of two interactive RDP mechanisms.

We then can convert this privacy filter to a valid RDP IM-privacy loss accumulator based on the following lemma.

**Lemma 1.26.** A function  $\mathcal{G}$  is a valid RDP IM-privacy loss accumulator if and only if  $\mathcal{F}(\cdot; \epsilon) = \mathbb{I}(\mathcal{G}(\cdot) \leq \epsilon)$  is a valid RDP IM-filter for all  $\epsilon > 0$ .

We give a generalized version of Lemma 1.26 in Section 2, meaning that an analogue of Lemma 1.26 also holds for all other variants of DP.

## 2. CONCURRENT FILTER AND ODOMETER

**2.1. Preliminaries.** Vadhan and Zhang [Vadhan and Zhang \(2023\)](#) define the generalized notion of privacy, termed  $\mathcal{D}$  DP, based on the following generalized probability distance. The distance is defined on a partially ordered set, which allows us to compare privacy guarantees at different protection levels. It should enjoy the postprocessing and joint convexity properties.

**Definition 2.1** (Generalized probability distance [Vadhan and Zhang \(2023\)](#)). A *generalized probability distance measure* is a tuple  $(\mathcal{D}, \preceq, D)$  such that

- (1)  $(\mathcal{D}, \preceq)$  is a partially ordered set (poset).
- (2)  $D$  is a mapping that takes any two random variables  $X, X'$  over the same measurable space to an element  $D(X, X')$  of  $\mathcal{D}$ .
- (3) (Postprocessing.) The generalized distance mapping  $D$  is closed under postprocessing, meaning that for every measurable function  $g$ ,  $D(g(X), g(X')) \preceq D(X, X')$ .
- (4) (Joint Convexity.) Suppose we have a collection of random variables  $(X_i, X'_i)_{i \in \mathcal{I}}$  and a random variable  $I$  distributed on  $\mathcal{I}$ . If  $D(X_i, X'_i) \preceq d$  for all  $i \in \mathcal{I}$ , then  $D(X_I, X'_I) \preceq d$ .

With this definition, the difficulty of distinguishing two neighboring datasets is measured by the generalized distance between the distributions of an adversary's views.

**Definition 2.2** ( $\mathcal{D}$  DP Vadhan and Zhang (2023)). Let  $(\mathcal{D}, \preceq, D)$  be a generalized probability distance. For  $d \in \mathcal{D}$ , we call an interactive mechanism  $\mathcal{M}$   $d$ - $\mathcal{D}$  DP if for every interactive algorithm  $\mathcal{A}$  and every pair of neighboring datasets  $x, x'$ , we have

$$D(\text{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x)), \text{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x'))) \preceq d.$$

A convenient property of  $\mathcal{D}$  DP is that it suffices to consider deterministic adversaries, because  $\mathcal{D}$  DP satisfies joint convexity. Since the filters and odometers we define will be  $\mathcal{D}$  DP, we will only consider deterministic adversaries in this paper.

**Lemma 2.3** (Deterministic adversaries for  $d$ - $\mathcal{D}$  DP Vadhan and Wang (2021) Vadhan and Zhang (2023)). An interactive mechanism  $\mathcal{M}$  is  $d$ - $\mathcal{D}$  DP, if and only if for every pair of neighboring datasets  $x, x'$ , for every deterministic adversary  $\mathcal{A}$ ,

$$D(\text{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x)) \parallel \text{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x'))) \preceq d.$$

For the rest of this paper, we will notate  $\mathcal{D}^* = \bigcup_{k=0}^{\infty} \mathcal{D}^k$ .

Finally, we define concurrent composition formally. In this setting, the adversary can arbitrarily and adaptively interleave queries between several differentially-private mechanisms, meaning that the queries can be dependent.

**Definition 2.4** (Concurrent composition of interactive mechanisms Vadhan and Wang (2021)). Let  $\mathcal{M}_1, \dots, \mathcal{M}_k$  be interactive mechanisms taking private datasets  $x_1, \dots, x_k$  respectively. The concurrent composition of  $\mathcal{M}_1, \dots, \mathcal{M}_k$ , denoted  $\mathcal{M} = \text{ConComp}(\mathcal{M}_1, \dots, \mathcal{M}_k)$ , is defined as in Algorithm 1.

**2.2. Filters.** We can now define a generalization of the privacy filter for the composition of noninteractive mechanisms as introduced in Section 1.3. Recall that filters are based on continuation rules. In line with the fully-adaptive setting, the number of interactive mechanisms need not be specified beforehand.

**Definition 2.5** ( $\mathcal{F}$ -filtered composition of  $\mathcal{D}$  DP noninteractive mechanisms ( $\mathcal{F}$ -Filt( $NIM$ ))). For a continuation rule  $\mathcal{F} : \mathcal{D}^* \times \mathcal{D} \rightarrow \{0, 1\}$ , the  $\mathcal{F}$ -filtered composition of  $\mathcal{D}$  DP interactive mechanisms with privacy-loss budget  $d$  is the interactive mechanism given in Algorithm 2.

If there is a finite maximum  $K$  of noninteractive privacy-loss parameters  $d_1, d_2, \dots, d_K$  to take in, we denote the continuation rule  $\mathcal{F}_K(\cdot; d) : \mathcal{D}^* \times \mathcal{D} \rightarrow \{1, 0\}$ , where  $d$  is the target privacy budget.

**Definition 2.6** (Valid  $\mathcal{D}$  DP filter for noninteractive mechanisms). Let  $\mathcal{F}$  be a continuation rule  $\mathcal{F} : \mathcal{D}^* \times \mathcal{D} \rightarrow \{0, 1\}$ . We say  $\mathcal{F}$  is a valid  $\mathcal{D}$  DP  $NIM$ -filter for noninteractive mechanisms if for every  $d \in \mathcal{D}$ ,  $\mathcal{F}(\cdot; d)$ -Filt( $NIM$ ) is a  $d$ - $\mathcal{D}$  DP interactive mechanism.

**Algorithm 1:** Concurrent composition of interactive mechanisms

---

```

1 def  $\mathcal{M}(s, m)$ :
2   if  $m = \lambda$  then
3      $s \leftarrow (x, [(\mathcal{M}_1, s_1), \dots, (\mathcal{M}_k, s_k)])$  // initialize  $k$  mechanisms return
4      $(s, \lambda)$ 
5   Parse  $s = (x, [(\mathcal{M}_1, s_1), \dots, (\mathcal{M}_k, s_k)])$ 
6   Parse  $m$ 
7   if  $m$  is of the form  $m = (j, q)$  where  $j = 1, \dots, k$  and  $q$  is a query to  $\mathcal{M}_j$  then
8      $(s'_j, m') \leftarrow \mathcal{M}_j(s_j, q)$ 
9      $s' \leftarrow (x, [(\mathcal{M}_1, s_1), \dots, (\mathcal{M}_j, s'_j), \dots, (\mathcal{M}_k, s_k)])$ 
10    else if  $m$  cannot be parsed correctly then
11       $s' \leftarrow s, m' \leftarrow$  invalid query
12    return  $(s', m')$ 

```

---

**Algorithm 2:**  $\mathcal{F}$ -filtered composition of  $\mathcal{D}$  DP noninteractive mechanisms ( $\mathcal{F}$ -Filt(NIM)), for  $\mathcal{F}_K : \mathcal{D}^* \rightarrow \mathcal{D}$ , where  $K$  could be  $\infty$ .

---

```

1 def  $\mathcal{F}(\cdot; d)$ -Filt(NIM)( $s, m$ ):
2   if  $m = \lambda$  then
3      $x \leftarrow s, s' \leftarrow (x, [])$ , where  $[]$  is an empty list // initialize filter
4     return  $(s', \lambda)$ 
5   Parse  $s$  as  $s = (x, [(\mathcal{M}_1, d_1), (\mathcal{M}_2, d_2), \dots, (\mathcal{M}_k, d_k)])$ 
6   Parse  $m$ 
7   if  $m$  is of the form  $m = (\mathcal{M}', d')$  then
8     if  $k = K$  then
9        $s' \leftarrow$  Halt,  $m' \leftarrow$  at mechanism count limit
10    else if  $\mathcal{M}'$  is  $d'$ - $\mathcal{D}$  DP then
11       $\mathcal{M}_{k+1} \leftarrow \mathcal{M}'$ 
12       $d_{k+1} \leftarrow d'$ 
13    if  $\mathcal{F}(d_1, \dots, d_{k+1}) \neq 1$  then
14       $m' \leftarrow$  insufficient budget,  $s' \leftarrow$  Halt
15    else
16      Random sample  $r_{k+1} :=$  coin tosses for  $\mathcal{M}_{k+1}$ 
17       $s' \leftarrow (x, [d_1, \dots, d_{k+1}])$ 
18       $m' \leftarrow \mathcal{M}_{k+1}(x; r_{k+1})$ 
19    else if  $m$  cannot be parsed correctly then
20       $s' \leftarrow s, m' \leftarrow$  invalid query
21    return  $(s', m')$ 

```

---

We can now similarly define the privacy filter for concurrent composition of interactive mechanisms.

**Definition 2.7** ( $\mathcal{F}$ -filtered concurrent composition of  $\mathcal{D}$  DP interactive mechanisms). The  $\mathcal{F}$ -filtered concurrent composition of  $\mathcal{D}$  DP interactive mechanisms, denoted as  $\mathcal{F}\text{-FiltCon}(IM)$ , is an interactive mechanism as executed in Algorithm 3.

---

**Algorithm 3:**  $\mathcal{F}$ -filtered concurrent composition of  $\mathcal{D}$  DP interactive mechanisms ( $\mathcal{F}\text{-FiltCon}(IM)$ ), for  $\mathcal{F}_K : \mathcal{D}^* \rightarrow \mathcal{D}$ , where  $K$  could be  $\infty$ .

---

```

1 def  $\mathcal{F}(\cdot; d)\text{-FiltCon}(IM)(s, m)$ :
2   if  $m = \lambda$  then // initialize filter
3      $x \leftarrow s, s' \leftarrow (x, [])$ , where  $[]$  is an empty list
4     return  $(s', \lambda)$ 
5   Parse  $s$  as  $s = (x, [(\mathcal{M}_1, d_1, s_1), \dots, (\mathcal{M}_k, d_k, s_k)])$ 
6   Parse  $m$ 
7   if  $m$  is of the form  $m = (\mathcal{M}', d')$  then
8     if  $k = K$  then // already at mechanism limit
9        $s' \leftarrow \text{Halt}, m' \leftarrow$  at mechanism count limit
10    else if  $\mathcal{M}'$  is  $d'$ - $\mathcal{D}$  DP then
11       $\mathcal{M}_{k+1} \leftarrow \mathcal{M}'$ 
12       $d_{k+1} \leftarrow d'$ 
13      if  $\mathcal{F}(d_1, \dots, d_{k+1}) \neq 1$  then
14         $s' \leftarrow \text{Halt}, m' \leftarrow$  insufficient budget
15      else
16         $(s_{k+1}, m') \leftarrow \mathcal{M}_{k+1}(x, \lambda)$ 
17         $s' \leftarrow (x, [(\mathcal{M}_1, d_1, s_1), \dots, (\mathcal{M}_{k+1}, d_{k+1}, s_{k+1})])$ 
18    else if  $m$  is of the form  $m = (j, q)$  where  $j = 1, \dots, k$  and  $q$  is a query to  $\mathcal{M}_j$ 
19      then
20         $(s'_j, m') \leftarrow \mathcal{M}_j(s_j, q)$ 
21         $s' \leftarrow (x, d, [(\mathcal{M}_1, d_1, s_1), \dots, (\mathcal{M}_j, d_j, s'_j), \dots, (\mathcal{M}_k, d_k, s_k)])$ 
22    else if  $m$  cannot be parsed correctly then
23       $s' \leftarrow s, m' \leftarrow$  invalid query
24    return  $(s', m')$ 

```

---

A privacy filter is one in which Definition 2.7 holds for all distances  $d \in \mathcal{D}$ .

**Definition 2.8** (Valid  $\mathcal{D}$  DP filter for interactive mechanisms). Let  $\mathcal{F} : \mathcal{D}^* \times \mathcal{D} \rightarrow \{0, 1\}$  be a continuation rule.  $\mathcal{F}$  is a valid  $\mathcal{D}$  DP concurrent IM-filter if for every  $d \in \mathcal{D}$ ,  $\mathcal{F}(\cdot; d)\text{-FiltCon}(IM)$  is a  $d$ - $\mathcal{D}$  DP interactive mechanism.

**2.3. Odometer.** We can similarly generalize our definition for privacy odometers to  $\mathcal{D}$  DP.

**Definition 2.9** ( $\mathcal{G}$ -odometer for composition of noninteractive mechanisms ( $\mathcal{G}\text{-Odom}(NIM)$ )). A  $\mathcal{G}$ -odometer for the composition of noninteractive  $\mathcal{D}$  DP mechanisms, denoted as  $\mathcal{G}\text{-Odom}(NIM)$ , is executed as in Algorithm 4.

---

**Algorithm 4:**  $\mathcal{G}$ -odometer for composition of noninteractive mechanisms ( $\mathcal{G}$ - $\text{Odom}(NIM)$ ).

---

```

1 def  $\mathcal{G}$ - $\text{Odom}(NIM)(s, m)$ :
2   if  $m = \lambda$  then
3      $x \leftarrow s, s' \leftarrow (x, [])$ , where  $[]$  is an empty list
4     return  $(s', \lambda)$ 
5   Parse  $s$  as  $s = (x, [d_1, d_2, \dots, d_k])$ 
6   Parse  $m$ 
7   if  $m$  is of the form  $m = (\mathcal{M}', d')$  then
8     if  $\mathcal{M}'$  is  $d'$ - $\mathcal{D}$  DP then
9        $\mathcal{M}_{k+1} \leftarrow \mathcal{M}'$ 
10       $d_{k+1} \leftarrow d'$ 
11       $s' \leftarrow (x, [d_1, \dots, d_{k+1}])$ 
12      Random sample  $r_{k+1} :=$  coin tosses for  $\mathcal{M}_{k+1}$ 
13       $m' \leftarrow \mathcal{M}_{k+1}(x; r_{k+1})$ 
14     else
15        $s' \leftarrow s, m' \leftarrow$  Divergence cannot be parsed
16   else if  $m = \text{privacy\_loss}$  then
17      $s' \leftarrow s, m' \leftarrow \mathcal{G}(d_1, \dots, d_k)$ 
18   else if  $m$  cannot be parsed correctly then
19      $s' \leftarrow s, m' \leftarrow$  invalid query
20   return  $(s', m')$ 

```

---

Odometers come with a per-mechanism privacy guarantee, meaning that the divergence between the views after each new mechanism is started should be at most some set distance  $d \in \mathcal{D}$  apart. Because the odometer can receive a `privacy_loss` query at any point in computation, a notion that captures the during-computation guarantee is necessary. The generalized *truncated view* of the adversary up to the  $n$ -th round of interaction to an odometer equipped with continuation rule  $\mathcal{G}$  is defined as in Algorithm 5.

**Definition 2.10** (Truncated view for odometers Lécuyer (2021)). For a deterministic adversary  $\mathcal{A}$  and odometer  $\mathcal{O}$  on dataset  $x$ , the *truncated view*

$$\text{Trunc}_d(\text{View}(\mathcal{A} \leftrightarrow \mathcal{O}(x)))$$

is the return value of Algorithm 5, including the randomness and query answers.

Having the generalized truncated view to  $\mathcal{D}$  DP, we can now define a *valid*  $\mathcal{D}$  DP odometer, and a valid privacy-loss accumulator for an odometer of noninteractive mechanisms.

**Definition 2.11** (Valid  $\mathcal{D}$ -DP odometer).  $\mathcal{O}$  is a *valid*  $\mathcal{D}$  DP odometer if for every  $d \in \mathcal{D}$ , adversary  $\mathcal{A}$ , and every pair of adjacent datasets  $x, x'$ ,

$$D(\text{Trunc}_d(\text{View}(\mathcal{A} \leftrightarrow \mathcal{O}(x))) || \text{Trunc}_d(\text{View}(\mathcal{A} \leftrightarrow \mathcal{O}(x')))) \preceq d$$

**Definition 2.12** (Valid  $\mathcal{D}$ -DP privacy-loss accumulator for noninteractive mechanisms).  $\mathcal{G} : \mathcal{D}^* \rightarrow \mathcal{D}$  is a valid  $\mathcal{D}$  DP noninteractive-privacy-loss accumulator (NIM-privacy-loss-accumulator) if for every  $d \in \mathcal{D}$ ,  $\mathcal{G}(\cdot)$ - $\text{Filt}(NIM)$  is a valid  $d$ - $\mathcal{D}$  DP odometer.

---

**Algorithm 5:** Truncated view of adversary interacting with odometer  $\mathcal{O}$  given input  $d$ .

---

```

1 def TruncView(View( $\mathcal{A} \leftrightarrow \mathcal{O}(x)$ )):
2   Initialize  $s_1^M \leftarrow x$ ,  $q_0 \leftarrow \lambda$ ,  $s_1^A \leftarrow \lambda$ 
3   for  $i = 1, 2, \dots$  do
4     if  $i$  is odd then
5        $(s_{i+1}^M, a_i) \leftarrow \mathcal{O}(s_i^M, q_{i-1})$ 
6        $(s_i^M, d_i) = \mathcal{O}(s_i^M, \text{privacy\_loss})$ 
7        $s_{i+1}^A \leftarrow s_i^A$ 
8       if  $d_i \not\leq d$  then
9          $\perp$  Exit loop
10      else if  $i$  is even then
11         $(s_{i+1}^A, q_i) \leftarrow \mathcal{A}(s_i^A, a_{i-1}; r_i^A)$ 
12         $s_{i+1}^M \leftarrow s_i^M$ 
13        if  $s_{i+1}^M = \text{Halt}$  or  $s_{i+1}^A = \text{Halt}$  then
14           $\perp$  Exit loop
15  return  $\text{Trunc}_d(\text{View}(\mathcal{A} \leftrightarrow \mathcal{O}(x))) = (r_0^A, a_1, r_2^A, a_3, \dots, r_{i-2}^A, a_{i-1})$ 

```

---

We now define the privacy odometer that allows us to keep track of the total privacy loss across multiple interactive mechanisms, where the adversary can interleave queries to existing interactive mechanisms, as well as create new interactive mechanisms with privacy-loss budgets chosen adaptively. Crucially, our constructions of odometers have a specially-designated privacy-loss query, which outputs a conservative upper bound on the privacy loss for the interaction up to that point. It should be noted that this odometer model for I.M.s does not define a per-query privacy loss as in the odometer for noninteractive mechanisms. Instead, the “privacy-loss budget” is paid at the launch of each additional interactive mechanism. Any `privacy_loss` queries asked between mechanism launches will return the same value.

**Definition 2.13** ( $\mathcal{G}$ -odometer for concurrent composition of interactive mechanisms ( $\mathcal{G}\text{-OdomCon}(IM)$ )). A  $\mathcal{G}$ -odometer for the *concurrent composition of interactive mechanisms*, denoted as  $\mathcal{G}\text{-OdomCon}(IM)$ , is defined as in Algorithm 6.

We are now ready to define the notion of a *valid*  $\mathcal{D}$  DP privacy-loss accumulator.

**Definition 2.14** (Valid  $\mathcal{D}$  DP privacy-loss accumulator for interactive mechanisms). We say  $\mathcal{G}$  is a valid  $\mathcal{D}$  DP *concurrent IM-privacy-loss accumulator* if for every  $d \in \mathcal{D}$ ,  $\mathcal{G}(\cdot; d)\text{-OdomCon}(IM)$  is a valid  $d\text{-}\mathcal{D}$  DP odometer.

**2.4. I.M.-to-I.M. Interactive Postprocessing.** Before we move on to defining concurrent filters and odometers, we will formulate interactive postprocessing in terms of state machines, which will become a helpful tool for us to prove important theorems related to concurrent filters and odometers for  $f$ -DP and Rényi-DP mechanisms. We will first give a formal definition of person-in-the-middle mechanisms.

---

**Algorithm 6:**  $\mathcal{G}$ -odometer for concurrent composition of interactive mechanisms ( $\mathcal{G}$ -OdomCon(IM)).

---

```

1 def  $\mathcal{G}$ -OdomCon(IM)(s, m):
2   if  $m = \lambda$  then
3      $x \leftarrow s, s' \leftarrow (x, [])$ , where  $[]$  is an empty list
4     return  $(s', \lambda)$ 
5   Parse  $s = (x, [(\mathcal{M}_1, d_1, s_1), (\mathcal{M}_2, d_2, s_2), \dots, (\mathcal{M}_k, d_k, s_k)])$ 
6   Parse  $m$ 
7   if  $m$  is of the form  $m = (\mathcal{M}', d')$  then
8     if  $\mathcal{M}'$  is  $d'$ - $\mathcal{D}$  DP then
9        $\mathcal{M}_{k+1} \leftarrow \mathcal{M}'$ 
10       $d_{k+1} \leftarrow d'$ 
11       $(s_{k+1}, m') \leftarrow \mathcal{M}_{k+1}(x, \lambda)$ 
12       $s' \leftarrow (x, [(\mathcal{M}_1, d_1, s_1), \dots, (\mathcal{M}_{k+1}, d_{k+1}, s_{k+1})])$ 
13    else
14       $s' \leftarrow s, m' \leftarrow$  ‘Divergence cannot be parsed’
15  else if  $m$  is of the form  $m = (j, q)$  where  $j = 1, \dots, k$  and  $q$  is a query to  $\mathcal{M}_j$ 
16    then
17       $(s'_j, m') \leftarrow \mathcal{M}_j(s_j, q)$ 
18       $s' \leftarrow (x, [(\mathcal{M}_1, d_1, s_1), \dots, (\mathcal{M}_j, d_j, s'_j), \dots, (\mathcal{M}_k, d_k, s_k)])$ 
19  else if  $m = \text{privacy\_loss}$  then
20     $s' \leftarrow s, m' \leftarrow \mathcal{G}(d_1, \dots, d_k)$ 
21  else if  $m$  cannot be parsed correctly then
22     $s' \leftarrow s, m' \leftarrow$  invalid query
23  return  $(s', m')$ 

```

---

**Definition 2.15** (Person-in-the-middle mechanism). An *interactive postprocessing mechanism* (or ‘PIM mechanism’ for ‘person-in-the-middle’) is a randomized mechanism  $\mathcal{P} : \{0, 1\}^* \times \{Q, A\} \times \{0, 1\}^* \rightarrow \{0, 1\}^* \times \{Q, A\} \times \{0, 1\}^*$ .  $\mathcal{P}$  takes its current state  $s \in \{0, 1\}^*$ , a value  $v \in \{Q, A\}$  to indicate whether it is receiving a query (from an analyst or adversary) or an answer (from the mechanism it is postprocessing), and a message  $m$  (which is either a query or an answer) and returns a new state  $s' \in \{0, 1\}^*$ , a value  $v' \in \{Q, A\}$  to indicate whether it is asking a query (of the mechanism it is postprocessing) or providing an answer (if it is answering a query) and message  $m'$  (which is either a query or an answer), denoted  $(s', v', m') = \mathcal{P}(s, v, m)$ .

With the PIM primitive, we can now more precisely define a postprocessed interactive mechanism.

**Definition 2.16** (Postprocessed interactive mechanism). Let  $\mathcal{M}$  be an interactive mechanism and  $\mathcal{P}$  be an interactive postprocessing. Then the postprocessing of  $\mathcal{M}$  by a PIM algorithm  $\mathcal{P}$ , denoted  $\mathcal{P} \circ \mathcal{M}$ , is the interactive mechanism defined in Algorithm 7.

---

**Algorithm 7:** Postprocessing of  $\mathcal{M}$  by PIM algorithm  $\mathcal{P}$ , denoted  $\mathcal{P} \circ \mathcal{M}$ .

---

```

1 def  $(\mathcal{P} \circ \mathcal{M})(s, m)$ :
2   if  $m = \lambda$  then
3      $s^{\mathcal{M}} \leftarrow s, s^{\mathcal{P}} \leftarrow \lambda, m \leftarrow \lambda, v \leftarrow Q$            // initialize  $\mathcal{M}$  and  $\mathcal{P}$ 
4   else
5     Parse  $s$  as  $s = (s^{\mathcal{M}}, s^{\mathcal{P}})$ 
6     Let  $(s^{\mathcal{P}}, v, m) \leftarrow \mathcal{P}(s^{\mathcal{P}}, Q, m)$ 
7   while  $v = Q$  do /* while  $\mathcal{P}$  is equipped to interact with  $\mathcal{M}$ , conduct
      interaction                                                                    */
8
9      $(s^{\mathcal{M}}, m) \leftarrow \mathcal{M}(s^{\mathcal{M}}, m)$ 
10     $(s^{\mathcal{P}}, v, m) = \mathcal{P}(s^{\mathcal{P}}, A, m)$ 
11  return  $((s^{\mathcal{M}}, s^{\mathcal{P}}), m)$ 

```

---

**Theorem 2.17** (Privacy holds under I.M.-to-I.M. postprocessing). If  $\mathcal{P} : \{0, 1\}^* \times \{Q, A\} \times \{0, 1\}^* \rightarrow \{0, 1\}^* \times \{Q, A\} \times \{0, 1\}^*$  is an interactive postprocessing algorithm and  $\mathcal{M}$  is a  $d$ - $\mathcal{D}$  DP interactive mechanism, then  $\mathcal{P} \circ \mathcal{M}$  is  $d$ - $\mathcal{D}$  DP.

*Proof.* Let  $\mathcal{A}$  be a deterministic adversary against  $\mathcal{P} \circ \mathcal{M}$ . By Definition 1.6,  $\mathbf{View}(\mathcal{A} \leftrightarrow \mathcal{P} \circ \mathcal{M}(x)) = (r_0, m_1, r_2, m_3, r_4, \dots)$ , where  $r_i$  are the random coins that  $\mathcal{A}$  tosses and  $m_i$  are the messages received by  $\mathcal{A}$ . Define an adversary  $\mathcal{A}'$  that maintains the states of  $\mathcal{A}$  and  $\mathcal{P}$  as submachines, constructed as in Algorithm 8.  $\mathcal{A}'$  is a normal adversary, so it will always get started by receiving an answer and end by making a query.

Informally, Algorithm 8 says that when  $\mathcal{A}'$  is to generate a query (determined by  $v = Q$ ), it will run  $\mathcal{A}$  with  $\mathcal{P}$  to generate a query. When  $\mathcal{A}'$  is to receive an answer (determined by  $v = A$ ), it will run  $\mathcal{P}$  to first check if  $\mathcal{P}$  wants to interactively process it with  $\mathcal{M}$  first. If so, the interaction between  $\mathcal{P}$  and  $\mathcal{M}$  would proceed until  $\mathcal{P}$  determines that the answer is in a state to be passed to  $\mathcal{A}$ , done by checking the output value  $v_{out}$  of  $\mathcal{P}$  given the state, the value  $A$ , and the message as inputs at that round.

Because  $\mathcal{M}$  is  $d$ - $\mathcal{D}$  DP, the interaction between  $\mathcal{A}'$  and  $\mathcal{M}$  is  $d$ - $\mathcal{D}$  DP. We define a noninteractive postprocessing function  $g$  that transforms  $\mathbf{View}(\mathcal{A}' \leftrightarrow \mathcal{M}(x)) = (r'_0, m'_1, r'_2, m'_3, \dots)$  on any dataset  $x$  as in Algorithm 9. Informally, the new view object generated by Algorithm 9 makes visible the submachine interaction between  $\mathcal{A}$  and  $\mathcal{P}$  in the operation of  $\mathcal{A}'$ , and snips the interaction between  $\mathcal{A}'$  and  $\mathcal{M}$  such that if there is an extended sequence between  $\mathcal{M}$  and submachine  $\mathcal{P}$  of  $\mathcal{A}$ , only the final answer from  $\mathcal{M}$  before  $\mathcal{P}$  and  $\mathcal{A}$  begin interacting will be included. Since noninteractive postprocessing preserves privacy properties by Definition 2.1 (3), we know that,

$$\begin{aligned} D(g(\mathbf{View}(\mathcal{A}' \leftrightarrow \mathcal{M})(x)) \| g(\mathbf{View}(\mathcal{A}' \leftrightarrow \mathcal{M})(x'))) \\ \leq D(\mathbf{View}(\mathcal{A}' \leftrightarrow \mathcal{M})(x) \| \mathbf{View}(\mathcal{A}' \leftrightarrow \mathcal{M})(x')) \leq d. \end{aligned}$$

The view of  $\mathcal{A}$  against  $\mathcal{P} \circ \mathcal{M}$  on any dataset  $x$  can be computed by postprocessing the view of  $\mathcal{A}'$  against  $\mathcal{M}$  on  $x$  using  $g$ ; in other words,  $\mathbf{View}(\mathcal{A} \leftrightarrow \mathcal{P} \circ \mathcal{M}) \equiv g(\mathbf{View}(\mathcal{A}' \leftrightarrow \mathcal{M}))$  on any dataset  $x$ . Therefore,  $D(\mathbf{View}(\mathcal{A} \leftrightarrow \mathcal{P} \circ \mathcal{M})(x) \| \mathbf{View}(\mathcal{A} \leftrightarrow \mathcal{P} \circ \mathcal{M})(x')) \leq d$ . This means  $\mathcal{P} \circ \mathcal{M}$  is also  $d$ - $\mathcal{D}$  DP.  $\square$

---

**Algorithm 8:** Adversary  $\mathcal{A}'$  that maintains the states of  $\mathcal{A}$  and  $\mathcal{P}$  as submachines.

---

```

1 def  $\mathcal{A}'(s, m)$ :
2   if  $s = \lambda$  then                                     // need to initialize  $\mathcal{A}'$ 
3      $(s^{\mathcal{A}}, m_0) \leftarrow \mathcal{A}(\lambda, \lambda)$       // initialize adversary submachine
4      $(s^{\mathcal{P}}, v, m') \leftarrow \mathcal{P}(\lambda, A, m)$     // initialize PIM submachine
5   else
6     Parse  $s = (s^{\mathcal{A}}, s^{\mathcal{P}})$ 
7      $v_{out} \leftarrow A$  // initialize  $v_{out}$  if  $v = Q$  then //  $\mathcal{P}$  will provide a query
8     to  $\mathcal{M}$ 
9     while  $v_{out} = A$  do
10      Let  $((s^{\mathcal{A}})', -) \leftarrow \mathcal{A}(s^{\mathcal{A}}, m)$ 
11      Let  $((s^{\mathcal{P}})', v_{out}, m') \leftarrow \mathcal{P}(s^{\mathcal{P}}, v, -)$ 
12       $s^{\mathcal{A}} \leftarrow (s^{\mathcal{A}})', s^{\mathcal{P}} \leftarrow (s^{\mathcal{P}})', m \leftarrow m'$ 
13       $s' \leftarrow ((s^{\mathcal{A}})', (s^{\mathcal{P}})')$ 
14       $v \leftarrow A$  // prime  $\mathcal{A}$  to receive an answer
15   else if  $v = A$  then
16      $(s^{\mathcal{P}}, v_{out}, m') \leftarrow \mathcal{P}(s^{\mathcal{P}}, v, m)$ 
17      $s' \leftarrow (s^{\mathcal{A}}, s^{\mathcal{P}})$ 
18     if  $v_{out} = A$  then //  $\mathcal{P}$  will provide an answer to  $\mathcal{A}$ 
19        $v \leftarrow Q$  // prime  $\mathcal{A}$  to ask a query
20   return  $(s', m')$ 

```

---

A useful corollary follows directly from Theorem 2.17.

**Corollary 2.18.** Suppose  $\mathcal{N} : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^* \times \{0, 1\}^*$  is an interactive mechanism such that for every pair of neighboring datasets  $x, x'$ , and every deterministic adversary  $\mathcal{A}$ , there exists an interactive mechanism  $\mathcal{M}$  that is  $d$ - $\mathcal{D}$  DP on dataset universe  $\{x, x'\}$  and a PIM mechanism  $\mathcal{P}$  such that

$$\mathbf{View}(\mathcal{A} \leftrightarrow \mathcal{N}(x)) \equiv \mathbf{View}(\mathcal{A} \leftrightarrow \mathcal{P} \circ \mathcal{M}(x))$$

$$\mathbf{View}(\mathcal{A} \leftrightarrow \mathcal{N}(x')) \equiv \mathbf{View}(\mathcal{A} \leftrightarrow \mathcal{P} \circ \mathcal{M}(x'))$$

*Proof.* By Theorem 2.17,  $\mathcal{P} \circ \mathcal{M}$  is  $d$ - $\mathcal{D}$  DP on dataset universe  $\{x, x'\}$ . Therefore,  $D(\mathbf{View}(\mathcal{A} \leftrightarrow \mathcal{N}(x)) \parallel \mathbf{View}(\mathcal{A} \leftrightarrow \mathcal{N}(x'))) = D(\mathbf{View}(\mathcal{A} \leftrightarrow \mathcal{P} \circ \mathcal{M}(x)) \parallel \mathbf{View}(\mathcal{A} \leftrightarrow \mathcal{P} \circ \mathcal{M}(x'))) \leq d$ . Therefore,  $\mathcal{N}$  is also  $d$ - $\mathcal{D}$  DP.  $\square$

A particularly convenient property of privacy-loss accumulators and privacy filters is that a valid privacy-loss accumulator can be converted into a valid privacy filter, and vice versa. This bijective property allows us to use a filter to build an odometer, and an odometer to build a filter of the form  $\mathbb{I}(\mathcal{G}(\cdot) \leq d)$ . Any nice properties of one would hold for the other.

**Lemma 2.19.** (1) A function  $\mathcal{G} : \mathcal{D}^* \rightarrow \mathcal{D}$  is a valid NIM-privacy-loss accumulator if and only if  $\mathcal{F} : \mathcal{D}^* \times \mathcal{D} \rightarrow \{0, 1\}$  constructed from  $\mathcal{G}(\cdot)$  such that  $\mathcal{F}(\cdot; d) = \mathbb{I}(\mathcal{G}(\cdot) \leq d)$  is a valid  $\mathcal{D}$ -DP NIM-filter.

---

**Algorithm 9:** Noninteractive postprocessing  $g$  that transforms  $\text{View}(\mathcal{A}' \leftrightarrow \mathcal{M}(x))$  to  $\text{View}(\mathcal{A} \leftrightarrow (\mathcal{P} \circ \mathcal{M})(x))$

---

```

1 def  $g(\text{View}(\mathcal{A}' \leftrightarrow \mathcal{M}(x)))$ :
2    $i \leftarrow 1$            //  $i$  indexes input  $j \leftarrow 0$            //  $j$  indexes output Parse
    $\text{View}(\mathcal{A}' \leftrightarrow \mathcal{M}(x))$  as  $(r'_0, m'_1, r'_2, m'_3, \dots)$ 
3    $v \leftarrow Q, v_{out} \leftarrow A$ 
4    $(s^A, m^A) \leftarrow \mathcal{A}(\lambda, \lambda)$ 
5    $(s^P, v, m^P) \leftarrow \mathcal{P}(\lambda, v, m'_1)$ 
6   while  $m'_i$  do
7     if  $v = Q$  then                                     // adding interaction  $\mathcal{A} \leftrightarrow \mathcal{P}$ 
8       while  $v_{out} \neq Q$  do                             // simulate interaction of  $\mathcal{A}$  and  $\mathcal{P}$ 
9          $(s^A, m^A) \leftarrow \mathcal{A}(s^A, m^A, r_j^A)$ 
10         $(s^P, v_{out}, m^P) \leftarrow \mathcal{P}(s^P, v, m^A, r_j^P)$  // using the randomness of  $\mathcal{P}$ 
           to get the new state  $m_{j+1} \leftarrow m^P$            // record message  $\mathcal{A}$ 
           receives from  $\mathcal{P}$   $j \leftarrow j + 2$ 
11       $v \leftarrow A$ 
12      else if  $v = A$  then                               // abbreviating interaction  $\mathcal{P} \leftrightarrow \mathcal{M}$ 
13         $(s^P, v_{out}, m) \leftarrow \mathcal{P}(s_i^P, v, m'_i, r_j^P)$ 
14        if  $v_{out} = Q$  then
15           $v \leftarrow A$ 
16        else if  $v_{out} = A$  then                         // record the history
17           $m_{j+1} \leftarrow m$ 
18           $v \leftarrow Q$ 
19           $j \leftarrow j + 2$ 
20       $i \leftarrow i + 2$ 
21  return  $(r_0^A, m_1, r_2^A, m_3, \dots, r_j^A, m_{j+1}, \dots)$ 

```

---

- (2) A function  $\mathcal{G} : \mathcal{D}^* \rightarrow \mathcal{D}$  is a valid concurrent IM-privacy-loss accumulator if and only if  $\mathcal{F} : \mathcal{D}^* \times \mathcal{D} \rightarrow \{0, 1\}$  constructed from  $\mathcal{G}(\cdot)$  such that  $\mathcal{F}(\cdot; d) = \mathbb{I}(\mathcal{G}(\cdot) \preceq d)$  is a valid  $\mathcal{D}$ -DP concurrent IM-filter.

The full proof of this theorem is enclosed in Appendix B.1. The approach is to first define an I.M.  $\mathcal{M}$  that such that for any adversary  $\mathcal{A}$  and dataset  $x$ ,  $\text{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x))$  is exactly the same as  $\text{Trunc}_d(\text{View}(\mathcal{A} \leftrightarrow \mathcal{G}\text{-OdomCon}(\text{IM})(x, \cdot)))$ . Then, using Theorem 2.17, we can show that  $\mathcal{M}$  is an *interactive post-processing* of  $\mathcal{F}(\cdot; d)\text{-FiltCon}(\text{IM})$ .

### 3. CONCURRENT FILTER & ODOMETER FOR $f$ -DP

$f$ -DP is based on a hypothesis testing interpretation of differential privacy. Consider a hypothesis testing problem that attempts to quantify the difficulty of measuring the output of a mechanism on two adjacent datasets:

$$H_0: \text{the dataset is } x \text{ versus } H_1: \text{the dataset is } x'$$

Let  $Y := \mathcal{M}(x)$  and  $Y' := \mathcal{M}(x')$  be the output distributions of mechanism  $\mathcal{M}$  on neighboring datasets  $x, x'$ . For a given rejection rule  $\phi$ , the type I error  $\alpha_\phi = \mathbb{E}[\phi(Y)]$  is the probability of rejecting  $H_0$  when  $H_0$  is true, while the type II error  $\beta_\phi = 1 - \mathbb{E}[\phi(Y')]$  is the probability of failing to reject  $H_0$  when  $H_1$  is true.

The trade-off function characterizes the optimal boundary between achievable type I and type II errors, which in our case is calculated by finding the minimal achievable type II error after fixing the type I error at any level.

**Definition 3.1** (Trade-off function [Dong et al. \(2022\)](#)). For any two probability distributions  $Y$  and  $Y'$  on the same space, the *trade-off function*  $T(Y, Y') : [0, 1] \rightarrow [0, 1]$  is defined as

$$T(Y, Y')(\alpha) = \inf\{\beta_\phi : \alpha_\phi < \alpha\},$$

where the infimum is taken over all measurable rejection rules  $\phi$ .

**Definition 3.2** ( $f$ -DP [Dong et al. \(2022\)](#)). Let  $f$  be a tradeoff function. A mechanism  $\mathcal{M} : \mathcal{X} \rightarrow \mathbb{R}$  is  $f$ -differentially private if for every pair of neighboring datasets  $x, x' \in \mathcal{X}$ , we have

$$T(\mathcal{M}(x), \mathcal{M}(x')) \geq f.$$

In our generalized DP framework in [Definition 2.1](#), the probability distance measure for  $f$ -DP is  $(S, \preceq, T)$ . The distance mapping is the trade-off function  $T$  in [Definition 3.2](#). The partially ordered set  $(S, \preceq)$  consists of the set  $S$  of all trade-off functions  $g : [0, 1] \rightarrow [0, 1]$  such that  $g$  is convex, continuous, non-increasing, and  $g(x) \leq 1 - x$  for  $x \in [0, 1]$ . The partial ordering is defined as  $f_1 \preceq f_2$  if  $f_1(\alpha) \geq f_2(\alpha)$  holds for all  $\alpha \in [0, 1]$ . A larger trade-off function means less privacy loss.

The recent result of [Vadhan and Zhang \(2023\)](#) shows that every interactive  $f$ -DP mechanism can be simulated by an interactive postprocessing of a noninteractive  $f$ -DP mechanism.

**Theorem 3.3** (Reduction of  $f$ -DP to noninteractive mechanism [Vadhan and Zhang \(2023\)](#)). For every  $f$  in the set  $\mathcal{S}$  of trade-off functions, every pair of adjacent datasets  $x, x'$ , and every interactive  $f$ -DP mechanism  $\mathcal{M}$  with finite communication complexity, there exists a pair of random variables  $Y, Y'$  with support  $[0, 1]$  and a randomized interactive mechanism  $\mathcal{P}$  such that  $D(Y, Y') \preceq f$ .

For every adversary  $\mathcal{A}$ , we have

$$\mathbf{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x)) \equiv \mathbf{View}(\mathcal{A} \leftrightarrow \mathcal{P}(Y))$$

$$\mathbf{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x')) \equiv \mathbf{View}(\mathcal{A} \leftrightarrow \mathcal{P}(Y')).$$

Essentially, the interactive postprocessing function in [Theorem 3.3](#) takes the output of the non-interactive mechanism as the input, and interacts with the adversary just as the interactive mechanism does. It will use the output of the non-interactive mechanism to simulate all the answers to the adversary. This result is useful because it implies that for any fixed pair of datasets  $x, x'$ , to analyze the concurrent composition of interactive mechanisms  $\mathcal{M}_i$ , it suffices to consider the composition of their noninteractive versions  $\mathcal{N}_i$ . As a result, composition theorems for noninteractive mechanisms extend to the concurrent composition of interactive  $f$ -DP mechanisms.

Since  $(\epsilon, \delta)$ -DP is a specific case of  $f$ -DP [Dong et al. \(2022\)](#), where  $f_{\epsilon, \delta} = \max\{0, 1 - \delta - \exp(\epsilon)\alpha, \exp(-\epsilon)(1 - \delta - \alpha)\}$ , a corollary of this result is that an interactive  $(\epsilon, \delta)$ -DP mechanism can be postprocessed to a noninteractive  $(\epsilon, \delta)$ -DP mechanism, a result shown independently by [Lyu \(2022\)](#).

**3.1. Concurrent Filter for  $f$ -DP.** Using Theorem 3.3, we derive a concurrent  $f$ -DP filter from an  $f$ -DP filter for noninteractive mechanisms.

**Theorem 3.4** (Concurrent  $f$ -DP Filter). Suppose that  $\mathcal{F} : \mathcal{D}^* \times \mathcal{D} \rightarrow \{1, 0\}$  is a valid  $f$ -DP continuation rule for noninteractive mechanisms. Then  $\mathcal{F}$  is a valid  $f$ -DP continuation rule for interactive mechanisms with finite communication.

*Proof.* We need to prove that for every  $f \in \mathcal{S}$ ,  $\mathcal{F}(\cdot; f)$ - $FiltCon(IM)$  is an  $f$ -DP I.M., which is equivalent to showing that for every pair of adjacent datasets  $x, x'$ ,  $\mathcal{F}(\cdot; f)$ - $FiltCon(IM)$  is  $f$ -DP on the dataset universe  $\{x, x'\}$ . Fixing  $x, x'$ , we will define an I.M.-to-I.M. postprocessing  $\mathcal{P}$  that takes  $\mathcal{F}$ - $Filt(NIM)$  to  $\mathcal{F}$ - $FiltCon(IM)$ . For simplicity, we will assume that  $\mathcal{A}$  is deterministic, as Lemma 2.3 allows us to do. By the postprocessing property of  $f$ -DP IMs given in Theorem 3.3, we know that an  $f_j$ -DP mechanism  $\mathcal{M}_i$  can be simulated by an interactive postprocessing  $\mathcal{P}_j$  of a noninteractive  $f_j$ -DP mechanism  $\mathcal{N}_j$ .

$\mathcal{P}$  is constructed as follows.  $\mathcal{P}$  depends on whether it is primed to receive a query  $v = Q$  from the adversary or an answer  $v = A$  from the mechanism, to pass to the other party. If the adversary asks to start a new interactive mechanism,  $\mathcal{P}$  will make note of the corresponding noninteractive mechanism and the noninteractive-to-interactive postprocessing and pass the request on  $\mathcal{F}$ - $FiltSeq(NIM)$ . Otherwise, if the adversary queries an existing interactive mechanism  $\mathcal{M}_j$ ,  $\mathcal{P}$  will use the corresponding interactive postprocessing  $\mathcal{P}_j$  to answer the query instead of passing the message forward to  $\mathcal{M}$ . The only time  $\mathcal{P}$  will interact with  $\mathcal{M}$  is to start a new mechanism and to pass back the confirmation that a new mechanism has begun to  $\mathcal{A}$ . The algorithmic pseudocode of  $\mathcal{P}$  is in Appendix A.1. in Algorithm 11.

Note that this algorithm only interacts with  $\mathcal{F}$ - $Filt(NIM)$  in starting a new mechanism. For the rest of the interactive queries, the algorithm directly uses the interactive postprocessing of the corresponding noninteractive mechanism in question to answer the queries. By assumption,  $\mathcal{F}$ - $Filt(NIM)$  is  $f$ -DP and  $\mathcal{F}$ - $FiltCon(IM) \equiv \mathcal{P} \circ \mathcal{F}$ - $Filt(NIM)$  on the dataset universe  $\{x, x'\}$ . Thus, by Corollary 2.18,  $\mathcal{F}$ - $FiltCon(IM)$  is  $f$ -DP.  $\square$

An analogue of Theorem 3.4 holds for  $(\epsilon, \delta)$ -DP.

**Corollary 3.5** (Concurrent  $(\epsilon, \delta)$  Filter). Every filter for noninteractive  $(\epsilon, \delta)$ -DP mechanisms is also a concurrent filter of interactive  $(\epsilon, \delta)$ -DP mechanisms.

Concurrent  $f$ -DP odometers can be defined similarly from  $f$ -DP odometers for noninteractive mechanisms.

**Theorem 3.6** (Concurrent  $f$ -DP odometer). Suppose that  $\mathcal{G} : \mathcal{D}^* \rightarrow \mathcal{D}$  is a valid  $f$ -DP privacy-loss accumulator for noninteractive mechanisms. Then  $\mathcal{G}$  is a valid  $f$ -DP privacy-loss accumulator for interactive mechanisms with finite communication.

*Proof.* By Lemma 2.19, if  $\mathcal{G}$  is a valid  $f$ -DP privacy-loss accumulator, then  $\mathcal{F}(\cdot; f) = I(\mathcal{G}(\cdot) \preceq f)$  is a valid  $f$ -DP filter for noninteractive mechanisms. By Theorem 3.4, if  $\mathcal{F}$  is an  $f$ -DP filter for noninteractive mechanisms, then  $\mathcal{F}$  will be an  $f$ -DP concurrent filter for interactive mechanisms. Applying Lemma 2.19 again,  $\mathcal{G}$  will be a valid  $f$ -DP privacy-loss accumulator for interactive mechanisms, and  $\mathcal{G}$ - $OdomCon(IM)$  is an  $f$ -DP interactive mechanism.  $\square$

Now we give an example of how to construct a valid  $(\epsilon, \delta)$ -DP privacy filter and odometer.

**Theorem 3.7** ( $(\epsilon, \delta)$ -DP privacy filter Whitehouse et al. (2023)). For every  $\delta' > 0$ ,

$$\begin{aligned} & \mathcal{F}((\epsilon_1, \delta_1), \dots, (\epsilon_k, \delta_k); (\epsilon, \delta)) \\ &= \mathbb{I} \left( \epsilon \leq \sqrt{2 \log \left( \frac{1}{\delta'} \right) \sum_{m \leq k} \epsilon_m^2 + \frac{1}{2} \sum_{m \leq k} \epsilon_m^2} \right) \cdot \mathbb{I} \left( \delta' + \sum_{m \leq k} \delta_k \leq \delta \right) \end{aligned}$$

is a valid approx-DP continuation rule for noninteractive mechanisms and interactive mechanisms.

**Theorem 3.8** ( $(\epsilon, \delta)$ -DP privacy odometer). Let  $\delta = \delta' + \delta''$  be a target approximation parameter such that  $\delta' > 0$ ,  $\delta'' \geq 0$ . Then

$$\begin{aligned} & \mathcal{G}((\epsilon_1, \delta_1), \dots, (\epsilon_k, \delta_k)) \\ &= \begin{cases} \left( \sqrt{2 \log \left( \frac{1}{\delta'} \right) \sum_{m \leq k} \epsilon_m^2 + \frac{1}{2} \sum_{m \leq k} \epsilon_m^2}, \delta \right) & \text{if } \delta' + \sum_{m \leq k} \delta_m \leq \delta \\ (\infty, \infty) & \text{otherwise} \end{cases} \end{aligned}$$

is a valid approx-DP privacy-loss accumulator for noninteractive and interactive mechanisms.

#### 4. CONCURRENT FILTER & ODOMETER FOR RÉNYI DIVERGENCE-BASED DP

Rényi DP is a relaxation of DP based on Rényi divergences. In this section, we show that privacy loss adds up for Rényi DP, in the setting of concurrent composition with adaptive privacy-loss parameters. This bound matches naturally with the previous concurrent composition result where all privacy-loss parameters are fixed upfront.

As a note, we inherit the limitation for when approximate-Rényi DP and zCDP are converted to approximate DP, in which there is some loss in the privacy parameters. This is an open question even for standard concurrent composition without adaptive privacy-loss parameters Lyu (2022); Vadhan and Zhang (2023), and we just inherit this limitation. The reason is that the current proof techniques for approximate DP and for Rényi DP are very different from each other, and it's not clear how to combine them. Hence, concurrent composition bounds for approximate RDP and zCDP are still an open question.

**Definition 4.1** (Rényi divergence Rényi (1961)). For two probability distributions  $P$  and  $Q$ , the *Rényi divergence* of order  $\alpha > 1$  is

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

**Definition 4.2** (Rényi DP Mironov (2017)). A randomized mechanism  $\mathcal{M}$  is  $(\alpha, \epsilon)$ -Rényi differentially private ( $\epsilon$ -RDP $_\alpha$ ) if for every two neighboring datasets  $x$  and  $x'$ ,

$$D_\alpha(\mathcal{M}(x)||\mathcal{M}(x')) \leq \epsilon$$

In our generalized DP framework in Definition 2.1, for Rényi DP of order  $\alpha$ , the partially ordered set  $\mathcal{D}$  is  $((\mathbb{R}^{\geq 0}) \cup \{\infty\}, \leq)$ . The distance mapping is  $\alpha$ -Rényi divergence for  $\alpha \in (1, \infty)$ . Rényi divergence is also closed under postprocessing due to the data-processing inequality, and it satisfies the joint convexity. We will notate RDP $_\alpha$  as the family of distance measures corresponding to RDP, where  $\mathcal{D} = (\alpha, \cdot)$ -RDP = RDP $_\alpha$ .

Our proof leverages the  $K = 2$  case of the concurrent composition theorem proved by Lyu Lyu (2022), where the privacy-loss parameters are specified upfront.

**Theorem 4.3** (Concurrent composition of Rényi-DP mechanisms Lyu (2022)). For all  $\alpha > 1$ ,  $k \in \mathbb{N}$ ,  $\epsilon_1, \dots, \epsilon_k > 0$ , and all interactive mechanisms  $\mathcal{M}_1, \dots, \mathcal{M}_k$  such that  $\mathcal{M}_i$  is  $\epsilon_i$ -RDP $_\alpha$ , the concurrent composition  $\text{ConComp}(\mathcal{M}_1, \dots, \mathcal{M}_k)$  is  $\sum_{i=1}^k \epsilon_i$ -RDP $_\alpha$ .

Another definition of differential privacy derived from Rényi divergence is zero-concentrated differential privacy.

**Definition 4.4** (Zero-concentrated differential privacy (zCDP) Bun and Steinke (2016)). An interactive mechanism  $\mathcal{M}$  is  $\rho$ -zero-concentrated differentially private ( $\rho$ -zCDP) if for every two neighboring datasets  $x$  and  $x'$ , every adversary  $\mathcal{A}$ , and all orders  $\alpha \in (1, \infty)$ ,

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

where  $D_\alpha$  is the Rényi divergence defined in Definition 1.2.

In other words,  $\mathcal{M}$  is  $\rho$ -zCDP iff  $\mathcal{M}$  is  $(\alpha, \rho\alpha)$ -RDP for every  $\alpha \in (1, \infty)$ .

Since zCDP bounds Rényi divergence for all orders  $\alpha$ , it is stronger than Rényi-DP of a fixed order.

In our generalized probability framework in Definition 2.1, the partially ordered set  $\mathcal{D}$  for zCDP is  $(\mathbb{R}^{\geq 0} \cup \{\infty\}, \leq)$ . The distance mapping is the worst-case Rényi divergence normalized by  $\alpha$ , over all  $\alpha \in (1, \infty)$ . In other words,  $D(\text{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x)) \parallel \text{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x'))) = \sup_{\alpha \in (1, \infty)} \left( \frac{D_\alpha(\text{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x)) \parallel \text{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x')))}{\alpha} \right)$ . zCDP is closed under postprocessing Bun and Steinke (2016); Lyu (2022) and satisfies the joint convexity property as well Bun and Steinke (2016).

Similar to Rényi-DP, the concurrent composition of zCDP mechanisms is additive.

**Theorem 4.5** (Concurrent composition of zCDP mechanisms Lyu (2022)). For  $k \in \mathbb{N}$ ,  $\rho_1, \dots, \rho_k > 0$ , and all interactive mechanisms  $\mathcal{M}_1, \dots, \mathcal{M}_k$  such that  $\mathcal{M}_i$  is  $\rho_i$ -zCDP, their concurrent composition  $\text{ConComp}(\mathcal{M}_1, \dots, \mathcal{M}_k)$  is  $\sum_{i=1}^k \rho_i$ -zCDP.

We will see that once we have a concurrent Rényi-DP filter, the concurrent odometer case follows directly. Concurrent filters and odometers for zCDP will follow directly from the Rényi filter result as well.

## 4.1. Concurrent Filter & Odometer for Rényi-DP.

### 4.1.1. Filter.

**Theorem 4.6** (Theorem 1.22 restated). An RDP filter with the continuation rule  $\mathcal{F}(\epsilon_1, \epsilon_2, \dots; \epsilon)$   $= \mathbb{I}(\sum_i \epsilon_i \leq \epsilon)$  is a valid RDP IM-filter for every fixed order of  $\alpha > 1$ .

Though generalized probability distance also covers RDP, the reduction in Theorem 3.3 holds only if the probability distance also satisfies the property of coupling Vadhan and Zhang (2023), which RDP does not. Therefore, we use a different strategy to prove 4.6. We prove this theorem by inducting on a finite number of mechanisms being maintained by the  $\mathcal{F}$ -filtered concurrent composition of interactive mechanisms  $\mathcal{F}\text{-FiltCon}(IM)$ , and then taking the limit of those mechanisms. To do so, we first consider that  $\mathcal{F}\text{-FiltCon}(IM)$  maintains up to  $K$  mechanisms.

**Lemma 4.7.** Let  $\alpha > 1$  and define

$$\mathcal{F}_K(\epsilon_1, \epsilon_2, \dots, \epsilon_k; \epsilon) = \begin{cases} \mathbb{I}(\sum_{i=1}^k \epsilon_i \leq \epsilon) & \text{if } k \leq K \\ 0 & \text{otherwise} \end{cases}$$

for every  $k \in \mathbb{N}$ . Then  $\mathcal{F}_K$  is a valid  $\text{RDP}_\alpha$  IM-filter.

*Proof.* We prove this by inducting on the maximum number  $K$  of mechanisms that  $\mathcal{F}_K\text{-FiltCon}(IM)$  maintains. Fix a pair of datasets  $x, x'$ . By Lemma 2.3, it suffices to consider a deterministic adversary to prove this theorem.

*Base case:*  $K = 2$ . Intuitively, with two mechanisms, the budget partition is determined before any mechanisms are queried, so we can reduce the problem to concurrent RDP composition. Suppose the adversary starts a  $\epsilon_1\text{-RDP}_\alpha$  mechanism  $\mathcal{M}_1$  through  $\mathcal{F}_2\text{-FiltCon}(IM)$ . By the constraints of  $\mathcal{F}_2\text{-FiltCon}(IM)$ , the second mechanism  $\mathcal{M}_2$ , that is able to start, must have privacy parameter  $\epsilon_2 \leq \epsilon - \epsilon_1$ . The adversary can adaptively interleave queries between  $\mathcal{M}_1$  and  $\mathcal{M}_2$ .

We will first define the notion of a *universal* mechanism to help us with the reduction. Informally, a “universal” mechanism can be any mechanism with the privacy loss bounded by the preset privacy loss budget. The universal mechanism will first verify if an  $\mathcal{M}$  attempted by an adversary to start satisfies the given privacy loss parameter guarantee as well as if the given privacy loss parameter does not exceed the preset privacy loss budget. Upon verifying this, it will interact with the adversary as  $\mathcal{M}$ . Given a preset privacy budget  $\epsilon$ , an  $\text{RDP}_\alpha$  universal mechanism  $\mathcal{U}_{\alpha, \epsilon}(s, m)$  would operate as follows:

---

**Algorithm 10:** Operation of  $\mathcal{U}_{\alpha, \epsilon}(s, m)$ .

---

```

1 def  $\mathcal{U}_{\alpha, \epsilon}(s, m)$ :
2   if  $m = \lambda$  then
3      $x \leftarrow s, s' \leftarrow (0, x)$            // initialize state of  $\mathcal{U}$ 
4     return  $(s, \lambda)$ 
5   if  $s$  is of form  $s = (0, x)$  and  $m = \mathcal{M}$ , where  $\mathcal{M}$  is an  $\text{RDP}_\alpha$  mechanism then
6     if  $\mathcal{M}$  is  $(\alpha, \epsilon_1)\text{-RDP}$  then
7        $(m', s_{\mathcal{M}}) \leftarrow \mathcal{M}(s, \lambda), s' \leftarrow (1, s_{\mathcal{M}})$ 
8     else
9        $m' \leftarrow \text{invalid query}, s' \leftarrow s$ 
10  else if  $s$  is of form  $s = (1, s_{\mathcal{M}})$  then
11     $(s'_{\mathcal{M}}, m') \leftarrow \mathcal{M}(s_{\mathcal{M}}, m)$ 
12     $s' \leftarrow (1, s'_{\mathcal{M}})$ 
13  return  $(s', m')$ 

```

---

The state of  $\mathcal{U}_{\alpha, \epsilon}(s, m)$  only needs the state of  $\mathcal{M}$ , the first mechanism it starts, and a bit to indicate where the first mechanism has started.

Given a deterministic adversary  $\mathcal{A}$  against  $\mathcal{F}_2\text{-FiltCon}(IM)$ , let  $(\mathcal{M}_1, \epsilon_1) = \mathcal{A}(\lambda, \lambda)$ , which is  $\mathcal{A}$ 's first query to  $\mathcal{F}_2\text{-FiltCon}(IM)$ . Assume that  $\mathcal{M}_1$  is  $\epsilon_1\text{-RDP}_\alpha$  (else the query will be rejected by  $\mathcal{F}_2\text{-FiltCon}(IM)$ , so we should use the first valid query of  $\mathcal{A}$ ).

We will construct an adversary  $\mathcal{A}'$  that interacts with  $\text{ConComp}(\mathcal{M}_1, \mathcal{U}_{\alpha, \epsilon - \epsilon_1})$ , where  $\mathcal{U}_{\alpha, \epsilon - \epsilon_1}$  is the universal  $(\epsilon - \epsilon_1)\text{-RDP}_\alpha$  mechanism, and the adversary's first query to  $\mathcal{U}_{\alpha, \epsilon - \epsilon_1}$

is an  $\epsilon_2$ -RDP $_\alpha$  mechanism  $\mathcal{M}_2$ . For each query  $\mathcal{A}$  makes to  $\mathcal{F}_2\text{-FiltCon}(IM)$  prior to starting  $\mathcal{M}_2$ ,  $\mathcal{A}'$  will make a query to  $\mathcal{M}_1$ . When  $\mathcal{A}$  starts  $\mathcal{M}_2$ ,  $\mathcal{A}'$  will make its first query to  $\mathcal{U}_{\alpha, \epsilon - \epsilon_1}$ . Upon verifying  $\epsilon_2 \leq \epsilon - \epsilon_1$ ,  $\mathcal{U}_{\alpha, \epsilon - \epsilon_1}$  interacts with the adversary just as  $\mathcal{M}_2$  does. Afterward,  $\mathcal{A}'$  can interleave queries at will between  $\mathcal{M}_1$  and  $\mathcal{M}_2$ , interacting with  $\text{ConComp}(\mathcal{M}_1, \mathcal{U}_{\alpha, \epsilon - \epsilon_1})$  just as  $\mathcal{A}$  does with  $\mathcal{F}_2\text{-FiltCon}(IM)$ .

In this case,  $\text{View}(\mathcal{A}' \leftrightarrow \text{ConComp}(\mathcal{M}_1, \mathcal{U}_{\alpha, \epsilon - \epsilon_1})(x)) \equiv \text{View}(\mathcal{A}' \leftrightarrow \mathcal{P} \circ \text{ConComp}(\mathcal{M}_1, \mathcal{M}_2)(x))$  for an interactive postprocessing  $\mathcal{P}$  that does not pass on the adversary's first query to start a second mechanism to  $\mathcal{U}_{\alpha, \epsilon - \epsilon_1}$ , but otherwise preserves all of the adversary's queries and the mechanism's answers. The same holds for the computation on  $x'$ . By Theorem 4.3,  $\text{ConComp}(\mathcal{M}_1, \mathcal{M}_2)$  is  $\epsilon$ -RDP $_\alpha$ , which means by Corollary 2.18 that  $\text{ConComp}(\mathcal{M}_1, \mathcal{U}_{\alpha, \epsilon - \epsilon_1})$  is as well. Therefore,  $\mathcal{F}_2\text{-FiltCon}(IM)$  is an  $\epsilon$ -RDP $_\alpha$  I.M. By Definition 2.8,  $\mathcal{F}_2(\cdot; \epsilon)$  is a valid RDP $_\alpha$  IM-filter.

*Induction on  $K$ .* By the inductive hypothesis, for every  $\epsilon \geq 0$ ,  $\mathcal{F}_{K-1}(\cdot; \epsilon)\text{-FiltCon}(IM)$  is  $\epsilon$ -RDP $_\alpha$ . We define the postprocessing  $\mathcal{P}$  interacting with  $\mathcal{F}_2(\cdot; \epsilon)\text{-FiltCon}(IM)$  so that for any  $K > 1$  and dataset  $x$ ,  $\text{View}(\mathcal{A} \leftrightarrow \mathcal{P} \circ \mathcal{F}_2(\cdot; \epsilon)\text{-FiltCon}(IM)(x)) \equiv \text{View}(\mathcal{A} \leftrightarrow \mathcal{F}_K\text{-FiltCon}(IM)(x))$ . The algorithmic pseudocode of  $\mathcal{P}$  is in Appendix A.2. in Algorithm 12.

When receiving the adversary's request to start  $\mathcal{M}_1$  and verifying the privacy loss is within budget,  $\mathcal{P}$  will start both  $\mathcal{M}_1$  and  $\mathcal{F}_{K-1}\text{-FiltCon}(IM)$  in  $\mathcal{F}_2\text{-FiltCon}(IM)$ .  $\mathcal{P}$  will directly pass on queries to  $\mathcal{M}_1$ . For queries  $q$  to  $\mathcal{M}_j$  where  $j > 1$  (in the form  $(j, q)$ ),  $\mathcal{P}$  will postprocess the message as  $(2, (j-1), q)$ , indicating that  $\mathcal{P}$  is passing  $q$  on to the  $(j-1)$ st mechanism of  $\mathcal{F}_{K-1}\text{-FiltCon}(IM)$ . (Concretely, suppose the adversary thinks it's interacting concurrently with  $\mathcal{F}_4\text{-FiltCon}(IM)$  maintaining  $\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3, \mathcal{M}_4$ . In this construction, the adversary will actually be interacting with  $\mathcal{F}_2\text{-FiltCon}(IM)$  maintaining  $\mathcal{M}_1$  and  $\mathcal{F}_3\text{-FiltCon}(IM)$ , the latter of which is comprised of  $\mathcal{M}_2, \mathcal{M}_3, \mathcal{M}_4$ , which are the first, second, and third mechanisms, respectively, of  $\mathcal{F}_3\text{-FiltCon}(IM)$ .)

Given a privacy-loss budget  $\epsilon$ , suppose the adversary starts a  $\epsilon_1$ -RDP $_\alpha$  mechanism  $\mathcal{M}_1$  with  $\epsilon_1 \leq \epsilon$ . As soon as  $\mathcal{M}_1$  gets created,  $\mathcal{P}$  will feed in  $\mathcal{M}_1$  to  $\mathcal{F}_2\text{-FiltCon}(IM)$ . Upon receiving confirmation that  $\mathcal{M}_1$  has properly started,  $\mathcal{P}$  will then feed in  $\mathcal{F}_{K-1}$  to  $\mathcal{F}_2\text{-FiltCon}(IM)$  with an  $\epsilon - \epsilon_1$  budget. We don't need to handle the parsing in the  $v = Q$  section because  $\mathcal{P}$  will always start two mechanisms at once. We can view  $\mathcal{F}_{K-1}\text{-FiltCon}(IM)$  as the first query to  $\mathcal{U}_2$ .

By construction, for any deterministic adversary  $\mathcal{A}$  and dataset  $x$ ,  $\text{View}(\mathcal{A} \leftrightarrow \mathcal{P} \circ \mathcal{F}_2\text{-FiltCon}(IM)(x))$  is identically distributed as  $\text{View}(\mathcal{A} \leftrightarrow \mathcal{F}_K\text{-FiltCon}(IM)(x))$ . By Theorem 2.17,  $\mathcal{P} \circ \mathcal{F}_2\text{-FiltCon}(IM)$  is an  $\epsilon$ -RDP $_\alpha$  I.M., which means  $\mathcal{F}_K\text{-FiltCon}(IM)$  is  $\epsilon$ -RDP $_\alpha$  I.M. By Definition 2.8,  $\mathcal{F}_K$  is a valid RDP $_\alpha$  IM-filter.  $\square$

To show that the filter works with an unbounded number of mechanisms, we examine the behavior as  $K \rightarrow \infty$ . For this to make sense, we observe that we are working with distributions on sequences that are well-behaved in that we can marginalize to finite subsequences. Suppose  $(\Omega^\infty, \Gamma^\infty)$  is the direct product of an infinite sequence of measurable spaces  $(\Omega_1, \Gamma_1), (\Omega_2, \Gamma_2), \dots$ , i.e.  $\Omega^\infty = \Omega_1 \times \Omega_2 \times \dots$  and  $\Gamma^\infty$  is the smallest  $\sigma$ -algebra containing all cylinder sets  $S_n(A) = \{x^\infty \in \Omega^\infty \mid x_1, x_2, \dots, x_n \in A\}$ ,  $A \in \Gamma^n$ , for  $n = 1, 2, \dots$  and  $\Gamma^n = \Gamma_1 \otimes \Gamma_2 \otimes \dots \otimes \Gamma_n$ . We call a sequence of probability distributions  $P^1, P^2, \dots$  *consistent* if  $P^n$  is a distribution on  $\Omega^n = \Omega_1 \times \Omega_2 \times \dots \times \Omega_n$  and  $P^{i+1}(A \times \Omega_{i+1}) = P^i(A)$  for  $A \in \Gamma^i$ ,  $i = 1, \dots, n$ . For any such sequence there exists a distribution  $P^\infty$  on  $(\Omega^\infty, \Gamma^\infty)$  such that its marginal distribution on  $\Omega^n$  is  $P^n$ , in the sense that  $P^\infty(S_n(A)) = P^n(A)$ ,

$A \in \Gamma^n$ , where  $\Omega^\infty = \Omega_1 \times \Omega_2 \times \dots$ ,  $S_n(A)$  is the *cylinder set*  $\{x^\infty \in \Omega^\infty \dots\}$ , and  $\Gamma^\infty$  is the smallest  $\sigma$ -algebra containing all the cylinder sets  $S_k(\mathcal{B})$  for  $k \in \mathbb{N}$  and  $\mathcal{B} \in \Gamma^k$ .

**Lemma 4.8** (Van Erven and Harremos (2014)). Let  $P^1, P^2, \dots$  and  $Q^1, Q^2, \dots$  be consistent sequences of probability distributions on  $(\Omega^1, \Gamma^1), (\Omega^2, \Gamma^2), \dots$ , where, for  $n = 1, 2, \dots, \infty$ ,  $(\Omega^n, \Gamma^n)$  is the direct product of the first  $n$  measurable spaces in the infinite sequence  $(\Omega_1, \Gamma_1), (\Omega_2, \Gamma_2), \dots$ . Then for any  $a \in (0, \infty]$

$$\lim_{n \rightarrow \infty} D_\alpha(P^n || Q^n) = D_\alpha(P^\infty || Q^\infty).$$

*Proof of Theorem 4.6.* Fix a divergence parameter  $\alpha \geq 1$ , and a pair of datasets  $x, x'$ . Consider  $\epsilon \geq 0$ , and initialize an adversary  $\mathcal{A}$ . Consider the views  $V = (r_1, a_1, r_2, a_2, \dots)$  and  $V' = (r'_1, a'_1, r'_2, a'_2, \dots)$  of  $\mathcal{A}$  interacting with  $\mathcal{F}(\cdot; \epsilon)$ -*FiltCon(IM)* on  $x$  and  $x'$ , respectively. Let  $q_i$  denote the adversary's  $i$ th query and  $a_i$  denote the answer to the adversary's  $i$ th query.

Define sequences of random variables  $V_1, V_2, V_3, \dots$  by  $V_i = (r_1, a_1, r_2, a_2, \dots, r_i, a_i)$ , and similarly for  $V'_1, V'_2, V'_3, \dots$ . By construction, both the distributions of  $V_1, \dots$  and  $V'_1, \dots$  are consistent.

By Lemma 4.7, for every  $k = 1, 2, \dots$  and  $\alpha > 1$ , we have  $D_\alpha(V_k || V'_k) \leq \epsilon$ . Then, invoking Lemma 4.8, we have

$$D_\alpha(X_\infty || Y_\infty) = \lim_{k \rightarrow \infty} D_\alpha(X_k || Y_k) \leq \epsilon,$$

where the last inequality is because the sequence of  $D_\alpha(X_k || Y_k)$  monotonically increases as  $k \rightarrow \infty$ , and is bounded above by  $\epsilon$ .  $\square$

Our approach can be used to provide another proof in Feldman and Zrnic (2022) for the bounds on privacy filters for sequentially-composed noninteractive DP mechanisms as well. They identify the sequence of losses incurred for each query – in this case the invocation of a noninteractive mechanism and the reception of its answers – as a supermartingale. They then apply the optional stopping theorem for martingales to bound the divergence between adjacent datasets and verify the validity of the filter.

The basis for our proof is that concurrent composition of two interactive RDP mechanisms implies a fully adaptive  $\mathcal{F}_2$ -*FiltCon(IM)*. Having interactivity allows us to proceed by induction in the proof of Lemma 4.7, describing  $\mathcal{F}_k$  in terms of  $\mathcal{F}_{k-1}$ . We could do the same for the sequential composition since induction preserves sequentiality. Then by Lemma 4.7,  $\mathcal{F}_k$ -*FiltSeq(IM)* is an  $\epsilon$ -RDP $_\alpha$  I.M. Thus, the sequential composition of two interactive Rényi DP mechanisms of order  $\alpha$  implies a fully adaptive filter for sequential composition of RDP mechanisms of order  $\alpha$ .

4.1.2. *Odometer.* Because of the bijective relationship between valid continuation rules and privacy-loss accumulators, we are able to define an additive relationship for concurrent odometers of Rényi mechanisms as well.

**Theorem 4.9** (Valid  $\epsilon$ -RDP $_\alpha$  odometer for concurrently-composed interactive mechanisms). Let  $\alpha > 1$ . Define  $\mathcal{G} : \mathbb{R}_{\geq 0}^* \rightarrow \mathbb{R}_{\geq 0}$  by  $\mathcal{G}(\epsilon_1, \dots, \epsilon_K) = \sum_{i=1}^K \epsilon_i$ .  $\mathcal{G}$  is a valid privacy-loss accumulator for the concurrent composition of a sequence of mechanisms  $\mathcal{M}_1, \mathcal{M}_2, \dots$ , where  $\mathcal{M}_i$  is an  $\epsilon_i$ -RDP $_\alpha$  interactive mechanism.

*Proof.* Consider the relation  $\mathcal{F}(\epsilon_1, \dots, \epsilon_K; \epsilon) : \mathcal{D}^* \times \mathcal{D} \rightarrow \{0, 1\}$  such that  $\mathcal{F}(\epsilon_1, \dots, \epsilon_K; \epsilon) = \mathbb{I}(\mathcal{G}(\epsilon_1, \dots, \epsilon_K \leq \epsilon))$ . By Theorem 4.6,  $\mathcal{F}$  is a valid RDP $_\alpha$  filter. By Lemma 2.19,  $\mathcal{G}$  is a valid privacy-loss accumulator for Rényi mechanisms of order  $\alpha$ .  $\square$

## 4.2. Concurrent Filter & Odometer for zCDP.

**Theorem 4.10** (Valid zCDP filter for concurrently-composed interactive mechanisms). Define the relation  $\mathcal{F} : \mathbb{R}_{\geq 0}^* \times \mathbb{R}_{\geq 0} \rightarrow \{0, 1\}$  by  $\mathcal{F}(\rho_1, \rho_2, \dots; \rho) = \mathbb{I}(\sum_i \rho_i \leq \rho)$ . Then  $\mathcal{F}$  is a valid  $\rho$ -zCDP filter for the concurrent composition of zCDP interactive mechanisms.

*Proof.* By Definition 4.4, a mechanism  $\mathcal{M}_i$  that is  $\rho_i$ -zCDP is also  $(\alpha, \rho_i\alpha)$ -RDP for all  $\alpha > 1$ . For a particular  $\alpha$ , by Theorem 4.6,  $\mathbb{I}(\sum_i \rho_i \alpha \leq \rho\alpha) = \mathbb{I}(\alpha \sum_i \rho_i \leq \rho\alpha) = \mathbb{I}(\sum_i \rho_i \leq \rho)$  is a valid RDP IM-filter, meaning that a filter  $\mathcal{F}^{(\alpha)}\text{-FiltCon}(IM)$  constructed from  $\mathcal{F}^{(\alpha)}(\rho_1\alpha, \rho_2\alpha, \dots; \rho\alpha) := \mathbb{I}(\sum_i \rho_i \alpha \leq \rho\alpha) = \mathcal{F}(\rho_1, \rho_2, \dots; \rho)$  is an  $(\alpha, \rho\alpha)$ -RDP I.M. Since Theorem 4.6 holds for all  $\alpha$ ,  $\mathcal{F}^{(\alpha)}\text{-FiltCon}(IM)$  is an  $(\alpha, \rho\alpha)$ -RDP I.M. for all  $\alpha > 1$ .

To show that  $\mathcal{F}(\rho_1, \rho_2, \dots; \rho) = \mathbb{I}(\sum_i \rho_i \leq \rho)$  is a valid  $\rho$ -zCDP IM-filter, we construct a postprocessing  $\mathcal{P}$  that transforms  $\mathcal{F}^{(\alpha)}(\cdot; \rho\alpha)\text{-FiltCon}(IM)$  to  $\mathcal{F}(\cdot; \rho)\text{-FiltCon}(IM)$ . Informally,  $\mathcal{P}$  converts all queries from the adversary that start  $\rho_i$ -zCDP mechanisms into  $(\alpha, \rho_i\alpha)$ -RDP mechanisms, and preserves all other messages. The algorithmic pseudocode of  $\mathcal{P}$  is in the appendix in Algorithm 13.

By Corollary 2.18, because  $\text{View}(\mathcal{A} \leftrightarrow \mathcal{F}\text{-FiltCon}(IM)) \equiv \text{View}(\mathcal{A} \leftrightarrow \mathcal{P} \circ \mathcal{F}^{(\alpha)}\text{-FiltCon}(IM))$ ,  $\mathcal{F}(\cdot; \rho)\text{-FiltCon}(IM)$  is a valid  $\rho$ -zCDP I.M., which means  $\mathcal{F}(\cdot; \rho)$  is a valid  $\rho$ -zCDP IM-filter.  $\square$

**Theorem 4.11** (Valid  $\rho$ -zCDP odometer for concurrently-composed interactive mechanisms). Define  $\mathcal{G} : \mathcal{D}^* \rightarrow \mathcal{D}$  to be  $\mathcal{G}(\rho_1, \rho_2, \dots) = \sum_{i=1} \rho_i$ . Then  $\mathcal{G}$  is a valid privacy-loss accumulator for the concurrent composition of a sequence of mechanisms  $\mathcal{M}_1, \mathcal{M}_2, \dots$ , where  $\mathcal{M}_i$  is an  $\rho_i$ -zCDP interactive mechanism.

*Proof.* Define a relation  $\mathcal{F}(\rho_1, \rho_2, \dots; \rho) : \mathcal{D}^* \times \mathcal{D} \rightarrow \mathcal{D}$  such that  $\mathcal{F}(\rho_1, \rho_2, \dots; \rho) = \mathbb{I}(\mathcal{G}(\rho_1, \rho_2, \dots) \leq \rho)$ . By Theorem 4.10,  $\mathcal{F}$  is a valid  $\rho$ -zCDP filter. By Lemma 2.19,  $\mathcal{G}$  is a valid zCDP privacy-loss accumulator.  $\square$

## 5. IMPLEMENTATION

In this section, we provide an overview of the implementation of differentially private mechanisms in the open-source software project OpenDP and the Tumult Analytics platform. We then discuss the implications of our results on the two platforms.

**5.1. Practical Application in the OpenDP Library.** The OpenDP Library [Gaboardi et al. \(2020\)](#) is a modular collection of statistical algorithms used to build differentially private computations. The OpenDP Library represents differentially private computations with measurements and odometers. Measurements and odometers fully characterize the privacy guarantees they give in terms of an input domain, input metric, and privacy measure. They both contain a function to make a differentially private release and a privacy map to reason about the privacy spend of a release. A privacy map is a function that takes in a bound on the distance between adjacent inputs/datasets ( $d_{in}$ ), and translates it to a

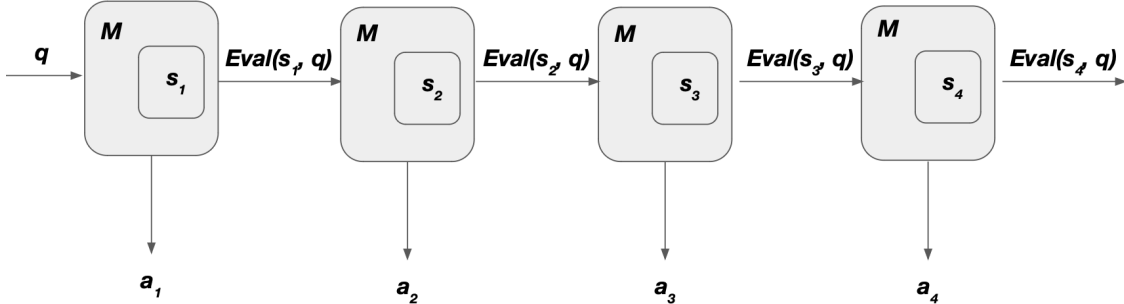


Figure 1: Diagram of a queryable.

bound on the distance between respective outputs ( $d_{out}$ , a privacy parameter). From the perspective of the library, an *interactive measurement* is simply a measurement for which the function emits a *queryable* (Figure 1), OpenDP terminology for an object that implements an interactive mechanism. A queryable is modeled as an opaque state machine, consisting of an internal state value and a transition function. When an analyst passes a query into the queryable, the state is updated, and an answer is returned.

All interactivity in the OpenDP Library uses queryables, and the privacy guarantee of an interactive measurement is defined exactly as in Definitions 1.7 and 2.2, quantitatively for all possible adversary’s strategies  $\mathcal{A}$ .

Interactive measurements are well-suited to capture interactive composition when the desired privacy parameters for each query are known up-front. For example, the internal state of a compositor queryable consists of the dataset and per-query privacy parameters. Each time a query  $q_i$  (which may itself be a measurement) is passed to the compositor queryable, a privacy parameter (like  $\epsilon_i$ ) is popped off of its internal state. Unlike measurements, odometers always emit a queryable from their function, and the privacy map is maintained *inside* the queryable. OpenDP’s implementation of odometers differs from the pseudocode in this paper, in that the input distance is not fixed up-front: the `privacy_loss` query also holds a  $d_{in}$ . Thus, queryables spawned by odometers can be queried with a bound on the distance between adjacent input datasets  $d_{in}$ , which returns the current corresponding privacy spend  $d_{out}$ .

The state of odometer queryables in OpenDP consist of the dataset and a vector of privacy maps (one from each query). When the odometer queryable is passed a privacy-loss query, the queryable passes  $d_{in}$  from the query into each of the child privacy maps stored in its state, which emit the  $d_i$  used in Algorithms 4 and 6. The privacy loss returned is the composition of these privacy loss parameters.

A filter is an instance of an interactive measurement. Any odometer can be converted into a filter IM by fixing an upper bound on the privacy parameters, as in Lemma 2.19. That is, any queryable spawned by the filter IM will refuse to answer any queries that would exceed the configured privacy parameters.

The OpenDP Library uses “wrapping” to handle situations where a queryable must influence the behavior of queryables it spawns in order to uphold its privacy guarantees. For example, a sequential composition queryable must wrap any queryables it spawns in a

logic that first asks the sequential compositor queryable for permission to execute (which allows the sequential compositor to maintain sequentiality), next executes the query, and then recursively wraps any queryable in the answer with the same logic.

Concurrent composition improves upon sequential composition in that concurrent composition does not need to influence the behavior of child queryables. Thus, concurrent compositors, odometers, and filters benefit from a simple implementation, in that they do not need to implement complicated sequentiality constraints via “wrapping”. More specifically, implementing *any* composition primitive for interactive mechanisms is more complex than for noninteractive mechanisms, as the composition primitive needs to provide the analyst/adversary access to many interactive mechanisms (in such a way that the only access to the sensitive data is through queries, and in particular ensuring that the private internal state of the mechanisms cannot be inspected directly), in contrast to simply sending the analyst the results of noninteractive mechanisms. In the case of sequential composition, the compositor implementation had to be even more complex in order to enforce sequentiality; our work removes the need for this additional complexity.

Concurrent composition also allows the analyst to work in a more exploratory fashion. By allowing non-sequential access to mechanisms, the analyst is free to sequence their queries arbitrarily. Analysts may also return to previous (un-exhausted) mechanisms without incurring additional privacy penalties for burn-in. This will also allow analysts to interact with multiple mechanisms simultaneously in settings where computational concurrency may make sequencing ambiguous. Our filter and odometer theorems also allow analysts to choose privacy parameters adaptively, and partition their privacy budget across different mechanisms, providing added flexibility.

We provide our implementation of concurrent composition based on the OpenDP Library on Github at <https://github.com/concurrent-composition/concurrent-composition>.

**5.2. Practical Application in the Tumult Framework.** Tumult Analytics Labs (2022) Berghel et al. (2022) is an open-source framework for releasing aggregate information from sensitive datasets with differential privacy. Analytics queries are evaluated using differentially private mechanisms in the context of a Tumult Analytics *session*, which is essentially a privacy filter. Analytics is written on top of Tumult Core, a framework for constructing differentially private mechanisms with automated privacy guarantees. Interactive mechanisms are the basic abstraction for interactivity in Tumult Core (an interactive mechanism is called a *queryable* in Tumult Core). Tumult Core currently restricts interleaving queries to the composition of queryables to prevent concurrent composition. For example, suppose we start with a privacy filter queryable, and our first query spawns a new queryable using some portion of the budget. Tumult Core requires that the user finish interacting with this spawned queryable before asking a new query to the parent privacy filter, otherwise we may (adaptively) spawn a second queryable and concurrently ask queries of both spawned queryables. Note that while certain concurrent composition results are known from prior work, Tumult Core currently does not allow concurrent composition under any circumstances to reduce code complexity and maintain UX consistency. The results in this paper allow for removing this restriction.

Although spawning queryables within queryables may seem niche, one common application in Tumult Analytics is with parallel composition. Parallel composition in differential privacy says that if we partition a dataset and run DP queries on each subset in the partition, the overall privacy loss is the maximum privacy loss across the mechanisms. In Tumult

Analytics, the mechanisms the user runs on each subset of the partitioned data can be interactive. Under the hood, this is implemented as follows: the top level privacy filter spawns a new interactive mechanism, consisting of multiple privacy filters (one for each subset in the partition). To prevent concurrent composition where it is not allowed, we require that the user finish interacting with this parallel composition interactive mechanism before they can ask any more queries of the top-level privacy filter, or spawn a new parallel composition mechanism from the top-level privacy filter.

The results in this paper allow for removing the restrictions on concurrent composition, and thereby improve the existing system in two ways. First, it reduces code complexity and improves auditability. The privacy guarantee of an interactive mechanism is contingent on the promise that there will be no concurrent composition in scenarios where it is not allowed. This property is more challenging to verify than other properties of the system that contribute to the privacy guarantee required, which tend to be localized and easy to verify in isolation.

Second, removing restrictions on concurrent composition improves the UX of Tumult Analytics. As mentioned previously, users of Analytics must finish interacting with one queryable before they start interacting with a subsequent queryable, and this multi-queryable scenario happens frequently when using parallel composition. This has the potential to cause issues:

- It is possible a user could have a reason for wanting to interleave queries to two or more queryables.
- Even if interleaving queries isn't required, the user may structure their code such that queries to multiple queryables are interleaved. This would cause the user to see an unfamiliar error, and require them to restructure their code.

Both of these issues would be solved if restrictions on concurrent composition for privacy filters was removed, which the results of our paper can enable.

#### ACKNOWLEDGEMENT

S.V. is supported by a gift from Apple, NSF grant BCS-2218803, a grant from the Sloan Foundation and a Simons Investigator Award. V.X. is supported in part by NSF grant BCS-2218803. W.Z. is supported in part by a Computing Innovation Fellowship from the Computing Research Association (CRA) and the Computing Community Consortium (CCC), and the BU Census grant.

We thank the JPC referees for their detailed reviews and valuable feedback. We are especially grateful to Roodabeh Safavi Hemami for identifying an error in Algorithm 7 in a previous version of this work.

#### REFERENCES

- S. Berghel, P. Bohannon, D. Desfontaines, C. Estes, S. Haney, L. Hartman, M. Hay, A. Machanavajjhala, T. Magerlein, G. Miklau, A. Pai, W. Sexton, and R. Shrestha. Tumult Analytics: a robust, easy-to-use, scalable, and expressive framework for differential privacy. *arXiv preprint arXiv:2212.04133*, Dec. 2022. <https://doi.org/10.48550/arXiv.2212.04133>.

- M. Bun and T. Steinke. Concentrated differential privacy: Simplifications, extensions, and lower bounds. In *Theory of Cryptography Conference*, pages 635–658. Springer, 2016. [https://doi.org/10.1007/978-3-662-53641-4\\_24](https://doi.org/10.1007/978-3-662-53641-4_24).
- J. Dong, A. Roth, and W. J. Su. Gaussian differential privacy. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(1):3–37, 02 2022. ISSN 1369-7412. <https://doi.org/10.1111/rssb.12454>.
- C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4):211–407, 2014. <https://doi.org/10.1561/04000000042>.
- C. Dwork and G. N. Rothblum. Concentrated differential privacy. *arXiv preprint arXiv:1603.01887*, 2016. <https://doi.org/10.48550/arXiv.1603.01887>.
- C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor. Our data, ourselves: Privacy via distributed noise generation. In *Annual international conference on the theory and applications of cryptographic techniques*, pages 486–503. Springer, 2006. [https://doi.org/10.1007/11761679\\_29](https://doi.org/10.1007/11761679_29).
- C. Dwork, M. Naor, O. Reingold, G. N. Rothblum, and S. P. Vadhan. On the complexity of differentially private data release: efficient algorithms and hardness results. In *Proceedings of the 41st ACM Symposium on Theory of Computing*, STOC '09, pages 381–390, 2009. <https://doi.org/10.1145/1536414.1536467>.
- C. Dwork, M. Naor, T. Pitassi, and G. N. Rothblum. Differential privacy under continual observation. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*, STOC '10, pages 715–724, 2010a. <https://doi.org/10.1145/1806689.1806787>.
- C. Dwork, G. N. Rothblum, and S. Vadhan. Boosting and differential privacy. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 51–60. IEEE, 2010b. <https://doi.org/10.1109/FOCS.2010.12>.
- V. Feldman and T. Zrnic. Individual privacy accounting via a renyi filter. *arXiv preprint arXiv:2008.11193*, 2022. <https://doi.org/10.48550/arXiv.2008.11193>.
- M. Gaboardi, M. Hay, and S. Vadhan. A programming framework for opendp. *Manuscript*, 2020. [https://scholar.harvard.edu/files/salil/files/opendp\\_programming\\_framework\\_10.pdf](https://scholar.harvard.edu/files/salil/files/opendp_programming_framework_10.pdf).
- M. Hardt and G. N. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *2010 IEEE 51st annual symposium on foundations of computer science*, pages 61–70. IEEE, 2010. <https://doi.org/10.1109/FOCS.2010.85>.
- P. Kairouz, S. Oh, and P. Viswanath. The composition theorem for differential privacy. *IEEE Transactions on Information Theory*, 63(6):4037–4049, 2017. <https://doi.org/10.1109/TIT.2017.2685505>.
- T. Labs. Tumult Analytics, Dec. 2022. <https://tmlt.dev>.
- M. Lécuyer. Practical privacy filters and odometers with renyi differential privacy and applications to differentially private deep learning. *arXiv preprint arXiv:2103.01379*, 2021. <https://doi.org/10.48550/arXiv.2103.01379>.
- X. Lyu. Composition theorems for interactive differential privacy. In *Thirty-sixth Conference on Neural Information Processing Systems*, 2022. <https://doi.org/10.48550/arXiv.2207.09397>.
- I. Mironov. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 263–275. IEEE, 2017. <https://doi.org/10.1109/CSF.2017.11>.

- J. Murtagh and S. Vadhan. The complexity of computing the optimal composition of differential privacy. In *Theory of Cryptography Conference*, pages 157–175. Springer, 2016. [https://doi.org/10.1007/978-3-662-49096-9\\_7](https://doi.org/10.1007/978-3-662-49096-9_7).
- A. Rényi. On measures of entropy and information. In *Proceedings of the fourth Berkeley symposium on mathematical statistics and probability*, volume 1. Berkeley, California, USA, 1961. [https://digitalassets.lib.berkeley.edu/math/ucb/text/math\\_s4\\_v1\\_article-27.pdf](https://digitalassets.lib.berkeley.edu/math/ucb/text/math_s4_v1_article-27.pdf).
- R. Rogers, A. Roth, J. Ullman, and S. Vadhan. Privacy odometers and filters: pay-as-you-go composition. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS’16*, page 1929–1937, Red Hook, NY, USA, 2016. Curran Associates Inc. ISBN 9781510838819. <https://dl.acm.org/doi/10.5555/3157096.3157312>.
- S. Vadhan and T. Wang. Concurrent composition of differential privacy. In *Theory of Cryptography Conference*, pages 582–604. Springer, 2021. [https://doi.org/10.1007/978-3-030-90453-1\\_20](https://doi.org/10.1007/978-3-030-90453-1_20).
- S. Vadhan and W. Zhang. Concurrent composition theorems for differential privacy. In *55th Annual ACM Symposium on Theory of Computing*, 2023. <https://doi.org/10.1145/3564246.3585241>.
- T. Van Erven and P. Harremoës. Rényi divergence and kullback-leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014. <https://doi.org/10.1109/TIT.2014.2320500>.
- J. Whitehouse, A. Ramdas, R. Rogers, and Z. S. Wu. Fully-adaptive composition in differential privacy. ICML’23. JMLR.org, 2023. <https://dl.acm.org/doi/10.5555/3618408.3619949>.

## APPENDIX A. ALGORITHMS

**A.1. Algorithm for proof of Theorem 3.4.** In this section, we enclose the IM-to-IM postprocessing algorithm  $\mathcal{P}$  for the proof of Theorem 3.4, as defined in Algorithm 11.  $\mathcal{P}$  will interact with  $\mathcal{F}\text{-Filt}(NIM)$  only to initialize and return  $\mathcal{F}\text{-Filt}(NIM)$ 's response for a new noninteractive  $f$ -DP mechanism, if an adversary asks to start an interactive  $f$ -DP mechanism. Otherwise, for a query to mechanism  $\mathcal{M}_j$ ,  $\mathcal{P}$  will use the NIM-to-IM postprocessing of  $\mathcal{M}_j$ ,  $\mathcal{P}_j$ , to answer the query. This allows any adversary  $\mathcal{A}$  to interact with  $\mathcal{F}\text{-Filt}(NIM)$  as if it were interacting with  $\mathcal{F}\text{-FiltCon}(IM)$ .

---

**Algorithm 11:** Interactive postprocessing  $\mathcal{P}$  that transforms  $f$ -DP  $\mathcal{F}\text{-Filt}(NIM)$  into  $\mathcal{F}\text{-FiltCon}(IM)$ .

---

```

1  def  $\mathcal{P}(s, v, m)$ :
2      if  $s = \lambda$  then // initialize  $\mathcal{P}$ 
3           $s \leftarrow (\square)$ , where  $\square$  is an empty list
4      if  $v = Q$  then //  $\mathcal{P}$  accepting queries
5          Parse  $s$  as  $([(\mathcal{P}_1, s^{\mathcal{P}_1}, f_1), \dots, (\mathcal{P}_k, s^{\mathcal{P}_k}, f_k)])$ 
6          if  $m$  is of the form  $m = (\mathcal{M}', f')$  then //  $\mathcal{A}$  starts a new mechanism
7              Let  $\mathcal{P}_{k+1}$  be the interactive postprocessing for noninteractive mechanism
               $\mathcal{N}_{k+1}$  so that  $\mathcal{P}_{k+1} \circ \mathcal{N}_{k+1} \equiv \mathcal{M}_{k+1}$  on dataset universe  $\{x, x'\}$ 
8               $s' \leftarrow ((\mathcal{P}_1, s^{\mathcal{P}_1}, f_1), \dots, (\mathcal{P}_{k+1}, \text{pending}, f_{k+1}))$ 
9               $m' \leftarrow (\mathcal{N}_{k+1}, f_{k+1})$ 
10             else if  $m$  is of form  $m = (j, q)$  where  $j = 1, \dots, k$  and  $q$  is a query to  $\mathcal{M}_j$ 
11                 then //  $\mathcal{A}$  queries an old mechanism
12                      $(s', m') \leftarrow \mathcal{P}_j(s, q)$ 
13                      $s' \leftarrow ((\mathcal{P}_1, s^{\mathcal{P}_1}, f_1), \dots, (\mathcal{P}_j, (s^{\mathcal{P}_j})', f_j), \dots, (\mathcal{P}_k, s^{\mathcal{P}_k}, f_k))$ 
14             else if  $v = A$  then //  $\mathcal{P}$  accepting answers
15                 if  $s$  is of form  $((\mathcal{P}_1, s^{\mathcal{P}_1}, f_1), \dots, (\mathcal{P}_k, \text{pending}, f_k))$  then
16                     if  $m = \text{invalid query}$  then // did not successfully start a new
17                         mechanism
18                          $s' \leftarrow ((\mathcal{P}_1, s^{\mathcal{P}_1}, f_1), \dots, (\mathcal{P}_{k-1}, s^{\mathcal{P}_{k-1}}, f_{k-1}))$ 
19                          $m' \leftarrow m$ 
20                     else // noninteractive output of new mechanism
21                          $(s^{\mathcal{P}_k}, m') \leftarrow \mathcal{P}(\lambda, A, m)$  // set up new interactive
22                         postprocessor's state
23                          $s' \leftarrow ((\mathcal{P}_1, s^{\mathcal{P}_1}, f_1), \dots, (\mathcal{P}_k, s^{\mathcal{P}_k}, f_k))$ 
24             return  $(s', A, m')$ 

```

---

**A.2. Algorithm for proof of Lemma 4.7.** In this section, we enclose the algorithm for I.M.-to-I.M. postprocessing algorithm  $\mathcal{P}$  for the proof of Lemma 4.7, as defined in Algorithm 12. This algorithm assumes inductively that  $\mathcal{F}_{K-1}\text{-FiltCon}(IM)$  is a valid  $\text{RDP}_\alpha$  filter. For

any adversary, when receiving the adversary's request to start  $\mathcal{M}_1$  and verifying the privacy loss is within budget,  $\mathcal{P}$  will start both  $\mathcal{M}_1$  and  $\mathcal{F}_{K-1}\text{-FiltCon}(IM)$  in  $\mathcal{F}_2\text{-FiltCon}(IM)$ .  $\mathcal{P}$  will directly pass on queries to  $\mathcal{M}_1$ . For queries  $q$  to  $\mathcal{M}_j$  where  $j > 1$  (in the form  $(j, q)$ ),  $\mathcal{P}$  will postprocess the message as  $(2, (j-1), q)$ , indicating that  $\mathcal{P}$  is passing  $q$  on to the  $(j-1)$ th mechanism of  $\mathcal{F}_{K-1}\text{-FiltCon}(IM)$ .

---

**Algorithm 12:** Interactive postprocessing  $\mathcal{P}$  that transforms  $\mathcal{F}_2(\cdot; \epsilon)\text{-FiltCon}(IM)$  into  $\mathcal{F}_K(\cdot; \epsilon)\text{-FiltCon}(IM)$ .

---

```

1 def  $\mathcal{P}(s, v, m)$ :
2    $k \leftarrow$  number of mechanisms (started or attempting to start)
3   if  $s = \lambda$  then // initialize  $\mathcal{P}$ 
4      $s \leftarrow (k, \square)$ , where  $\square$  is an empty list,  $v \leftarrow Q$ 
5   if  $v = Q$  then //  $\mathcal{P}$  accepting queries
6     if  $m$  is of the form  $m = (\mathcal{M}', (\alpha, \epsilon'))$  then //  $\mathcal{A}$  starts a new mechanism
7       if  $k = 0$  then // currently no mechanisms
8          $\mathcal{M}_1 \leftarrow \mathcal{M}', \epsilon_1 \leftarrow \epsilon'$ 
9          $b_1 \leftarrow \text{no}$  //  $b_1$  is the boolean keeping track of whether  $\mathcal{M}_1$  has
           successfully started  $k \leftarrow k + 1, s' \leftarrow (k, [(\mathcal{M}_1, (\alpha, \epsilon_1), b_1)])$ ,  $m' \leftarrow m$ 
10        else // both  $\mathcal{M}_1$  and  $\mathcal{F}_{K-1}$  started
11           $s' \leftarrow s, m' \leftarrow m$  // pass on query; the filter will handle any
            rejections (any request will be over budget)
12        else if  $m$  is of the form  $m = (j, q)$  where  $j = 1, \dots, K$  and  $q$  is a query to  $\mathcal{M}_j$  then
          //  $\mathcal{A}$  queries an old mechanism
13          if  $j = 1$  then // querying  $\mathcal{M}_1$ 
14             $m' \leftarrow (2, (j-1), q), s' \leftarrow s, m' \leftarrow m$ 
15           $v \leftarrow A$  // we will be talking to the mechanism next
16          return  $(s', Q, m')$ 
17        else if  $v = A$  then //  $\mathcal{P}$  accepting answers
18          if  $k = 1, b_1 = \text{no}$ , and  $m$  is not an error message then // started  $\mathcal{M}_1$ 
19             $b_1 \leftarrow \text{yes}, b_2 \leftarrow \text{no}, m' \leftarrow (\mathcal{F}_{K-1}\text{-FiltCon}(IM), (\alpha, \epsilon - \epsilon_1))$ 
20             $k \leftarrow k + 1, m_{\mathcal{M}} \leftarrow m$ 
21             $s' \leftarrow (k, [(\mathcal{M}_1, (\alpha, \epsilon_1), b_1), (\mathcal{F}_{K-1}\text{-FiltCon}(IM), (\alpha, \epsilon - \epsilon_1), b_2, m_{\mathcal{M}})])$  /* start
               $\mathcal{F}_{K-1}\text{-FiltCon}(IM)$  immediately, store  $m$  for return after  $\mathcal{F}_{K-1}$ 
              starts */
22            return  $(s', Q, m')$ 
23          else if  $k = 1$  and  $b_2 = \text{no}$  then // just started  $\mathcal{F}_{K-1}$ 
24            Parse  $s$  as  $s = (2, [(\mathcal{M}_1, (\alpha, \epsilon_1), b_1), (\mathcal{F}_{K-1}\text{-FiltCon}(IM), (\alpha, \epsilon - \epsilon_1), b_2, m_{\mathcal{M}})])$ 
25             $b_2 \leftarrow \text{yes}$ 
26             $m' \leftarrow m_{\mathcal{M}}$  // pass on the answer from  $\mathcal{M}_1$ 
27             $v \leftarrow Q$  // talking to adversary next
28            return  $(s', A, m')$ 
29          else
30             $s' \leftarrow s, v \leftarrow Q, m' \leftarrow m$ 
31            return  $(s', A, m')$ 

```

---

**A.3. Algorithm for proof of Theorem 4.10.** In this section, we enclose the algorithm for I.M.-to-I.M. postprocessing algorithm  $\mathcal{P}$  for the proof of Theorem 4.10, as defined in Algorithm 13. Essentially,  $\mathcal{P}$  only postprocesses queries that start new  $\rho$ -zCDP mechanisms into queries that start new  $(\alpha, \rho\alpha)$ -RDP mechanisms. All other interactions are preserved.

---

**Algorithm 13:** Interactive postprocessing  $\mathcal{P}$  that transforms  $\mathcal{F}^{(\alpha)}(\cdot; \epsilon)$ - $FiltCon(IM)$  to  $\mathcal{F}(\cdot; \rho)$ - $FiltCon(IM)$ .

---

```

1 def  $\mathcal{P}(s, v, m)$ :
2   if  $s = \lambda$  then // initialize  $\mathcal{P}$ 
3      $s \leftarrow (\square)$ , where  $\square$  is an empty list
4      $v \leftarrow Q$ 
5   if  $v = Q$  then //  $\mathcal{P}$  accepting queries
6     if  $m$  is of the form  $m = (\mathcal{M}', \rho')$  then //  $\mathcal{A}$  starts a new mechanism
7        $m' \leftarrow (\mathcal{M}', (\alpha, \rho'\alpha))$ 
8     else if  $m$  is of the form  $m = (j, q)$  where  $j = 1, 2, \dots$  and  $q$  is a query to
9        $\mathcal{M}_j$  then // ferry query to mechanism
10       $m' \leftarrow m$ 
11    else if  $m$  cannot be parsed correctly then
12       $m' \leftarrow$  invalid query
13     $v \leftarrow A$  // talk to the mechanism next
14    return  $(s, Q, m')$ 
15  else if  $v = A$  then //  $\mathcal{P}$  accepting answers
16     $v \leftarrow Q$  // talk to adversary next
17    return  $(s, A, m)$  // ferry answer back

```

---

## APPENDIX B. PROOFS

**B.1. Proof of Lemma 2.19.** In this section, we enclose the full proof of Lemma 2.19.

- Lemma B.1** (Lemma 2.19 restated). (1) A function  $\mathcal{G} : \mathcal{D}^* \rightarrow \mathcal{D}$  is a valid NIM-privacy-loss accumulator if and only if  $\mathcal{F} : \mathcal{D}^* \times \mathcal{D} \rightarrow \{0, 1\}$  constructed from  $\mathcal{G}(\cdot)$  such that  $\mathcal{F}(\cdot; d) = \mathbb{I}(\mathcal{G}(\cdot) \preceq d)$  is a valid  $\mathcal{D}$ -DP NIM-filter.
- (2) A function  $\mathcal{G} : \mathcal{D}^* \rightarrow \mathcal{D}$  is a valid concurrent IM-privacy-loss accumulator if and only if  $\mathcal{F} : \mathcal{D}^* \times \mathcal{D} \rightarrow \{0, 1\}$  constructed from  $\mathcal{G}(\cdot)$  such that  $\mathcal{F}(\cdot; d) = \mathbb{I}(\mathcal{G}(\cdot) \preceq d)$  is a valid  $\mathcal{D}$ -DP concurrent IM-filter.

*Proof.* We will prove the bijection relation for IM, as the proof for NIM holds similarly. Fix a pair of datasets  $x, x'$ . We will prove Lemma 2.19 for the interactive case, because the non-interactive case follows similarly.

Fix a pair of datasets  $x, x'$ .

*If.* Suppose that  $\mathcal{F}(\cdot; d) = \mathbb{I}(\mathcal{G}(\cdot) \preceq d)$  is a valid  $\mathcal{D}$  DP concurrent IM-filter. Then by Definition 2.8,  $\mathcal{F}(\cdot; d)$ - $FiltCon(IM)$  is a  $d$ - $\mathcal{D}$  DP interactive mechanism for every  $d \in \mathcal{D}$ . Thus, for an adversary  $\mathcal{A}$  interacting with  $\mathcal{G}$ - $OdomCon(IM)$ , we can construct an interactive

mechanism  $\mathcal{M}$  that will behave exactly like  $\mathcal{G}\text{-OdomCon}(IM)$  as in Algorithm 6, but on every round of operation also checks if  $\mathcal{G}(\cdot) \not\leq d$ , in which case  $\mathcal{M}$  will maintain its previous state and output `halt` message. It is easy to see that  $\text{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x))$  is exactly the same as  $\text{Trunc}_d(\text{View}(\mathcal{A} \leftrightarrow \mathcal{G}\text{-OdomCon}(IM)(x)))$  for any database  $x$ .

We now prove that  $\mathcal{M}$  is an interactive postprocessing of  $\mathcal{F}(\cdot; d)\text{-FiltCon}(IM)$ . Define a postprocessing  $\mathcal{P}$  as in Algorithm 14. This mechanism allows  $\mathcal{M}$  to behave the same as  $\mathcal{F}(\cdot; d)\text{-FiltCon}(IM)$ , just adding additional support for the deterministic `privacy_loss` queries, which do not change the state of the odometer and therefore do not affect divergence.

Therefore,

$$D(\text{Trunc}_d(\text{View}(\mathcal{A} \leftrightarrow \mathcal{G}\text{-OdomCon}(IM)(x)))) \quad (\text{B.1})$$

$$\begin{aligned} & \text{Trunc}_d(\text{View}(\mathcal{A} \leftrightarrow \mathcal{G}\text{-OdomCon}(IM)(x'))) \\ &= D(\text{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x)) \parallel \text{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x'))) \\ &= D(\text{View}(\mathcal{A} \leftrightarrow \mathcal{P} \circ \mathcal{F}\text{-FiltCon}(IM)(x)) \parallel \end{aligned} \quad (\text{B.2})$$

$$\begin{aligned} & \text{View}(\mathcal{A} \leftrightarrow \mathcal{P} \circ \mathcal{F}\text{-FiltCon}(IM)(x'))) \\ &\leq D(\text{View}(\mathcal{A} \leftrightarrow \mathcal{F}\text{-FiltCon}(IM)(x)) \parallel \text{View}(\mathcal{A} \leftrightarrow \mathcal{F}\text{-FiltCon}(IM)(x'))) \end{aligned} \quad (\text{B.3})$$

$$\leq d, \quad (\text{B.4})$$

where (B.3) follows from Theorem 2.17. By Definition 2.14,  $\mathcal{G} : \mathcal{D}^* \rightarrow \mathcal{D}$  is a valid  $\mathcal{D}$  DP privacy-loss accumulator.

*Only if.* Suppose that  $\mathcal{G} : \mathcal{D}^* \rightarrow \mathcal{D}$  is a valid  $\mathcal{D}$  DP privacy-loss accumulator. For every  $d \in \mathcal{D}$ , because the continuation rule is  $\mathcal{F}(\cdot; d) = \mathbb{I}(\mathcal{G}(\cdot) \leq d)$ ,  $\mathcal{F}(\cdot; d)\text{-FiltCon}(IM)$  will halt when  $\mathcal{G}(\cdot) \not\leq d$ , which is the same condition for outputting `View` truncated at  $d$  as in Algorithm 5. We can use the same  $\mathcal{M}$  as in the previous direction that is constructed such that, for an adversary  $\mathcal{A}$  and database  $x$ ,  $\text{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x))$  is exactly the same as  $\text{Trunc}_d(\text{View}(\mathcal{A} \leftrightarrow \mathcal{G}\text{-OdomCon}(IM)(x)))$ .

We now define an interactive postprocessing  $\mathcal{P}$  that transforms  $\mathcal{M}$  into  $\mathcal{F}(\cdot; d)\text{-FiltCon}(IM)$  in Algorithm 15, which removes support for `privacy_loss` queries since  $\mathcal{F}(\cdot; d)\text{-FiltCon}(IM)$  would not support that. Other interaction is essentially preserved, since  $\mathcal{M}$  and  $\mathcal{F}(\cdot; d)\text{-FiltCon}(IM)$  will both halt when  $\mathcal{G}(\cdot) \not\leq d$ .

Again, because `privacy_loss` queries do not affect divergence, we have

$$\begin{aligned} & D(\text{View}(\mathcal{A} \leftrightarrow \mathcal{F}\text{-FiltCon}(IM)(x)) \parallel \text{View}(\mathcal{A} \leftrightarrow \mathcal{F}\text{-FiltCon}(IM)(x'))) \\ &= D(\text{View}(\mathcal{A} \leftrightarrow \mathcal{P} \circ \mathcal{M}(x)) \parallel \text{View}(\mathcal{A} \leftrightarrow \mathcal{P} \circ \mathcal{M}(x'))) \\ &\leq D(\text{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x)) \parallel \text{View}(\mathcal{A} \leftrightarrow \mathcal{M}(x'))) \end{aligned} \quad (\text{B.5})$$

$$= D(\text{Trunc}_d(\text{View}(\mathcal{A} \leftrightarrow \mathcal{G}\text{-OdomCon}(IM)(x)) \parallel \text{Trunc}_d(\text{View}(\mathcal{A} \leftrightarrow \mathcal{G}\text{-OdomCon}(IM)(x')))) \quad (\text{B.6})$$

$$\begin{aligned} & \text{Trunc}_d(\text{View}(\mathcal{A} \leftrightarrow \mathcal{G}\text{-OdomCon}(IM)(x'))) \\ &\leq d \end{aligned} \quad (\text{B.7})$$

where (B.5) follows from Theorem 2.17. Thus,  $\mathcal{F}(\cdot; d)\text{-FiltCon}(IM)$  is a  $d\text{-}\mathcal{D}$  DP interactive mechanism. By Definition 2.8, since this is true for arbitrary  $d$ ,  $\mathcal{F}$  is a valid  $\mathcal{D}$  DP filter.  $\square$

**B.2. Proof of Example 3.7 (( $\epsilon, \delta$ )-DP filter).** In this section, we demonstrate how the ( $\epsilon, \delta$ )-DP filter and odometer for noninteractive mechanisms defined in Whitehouse et al.

---

**Algorithm 14:** Interactive postprocessing  $\mathcal{P}$  that transforms  $\mathcal{F}(\cdot; d)$ -*FiltCon*( $IM$ ) into  $\mathcal{M}$ .

---

```

1 def  $\mathcal{P}(s, v, m)$ :
2   if  $m = \lambda$  then                                     // initialize  $\mathcal{P}$ 
3      $s' \leftarrow ([\ ])$ , where  $[\ ]$  is an empty list
4   if  $v = Q$  then                                       // accepting queries from adversary
5     if  $m$  is of form  $m = (\mathcal{M}', d')$  then
6       Parse  $s$  as  $([\mathcal{M}_1, \dots, \mathcal{M}_k])$ 
7        $s' \leftarrow ([\mathcal{M}_1, \dots, \mathcal{M}_k, (\mathcal{M}_{k+1}, \text{pending})])$ 
8     else if  $m$  is of form  $m = (j, q)$  where  $j = 1, \dots, k$  and  $q$  is a query to  $\mathcal{M}_j$ 
9       then
10         $s' \leftarrow s$ 
11      else if  $m = \text{privacy\_loss}$  then
12         $m' \leftarrow \mathcal{G}(\cdot)$ 
13        return  $(s', A, m')$ 
14       $m' \leftarrow m$ 
15       $v' \leftarrow Q, v \leftarrow A$                        // pass query on to mechanism
16    else if  $v = A$  then                                       // accepting answers from filter
17      if  $s$  is of form  $s = ([\mathcal{M}_1, \dots, (\mathcal{M}_k, \text{pending})])$  then
18        if  $m$  is not an error message then // successfully started new
19          mechanism
20           $s' \leftarrow ([\mathcal{M}_1, \dots, \mathcal{M}_k])$ 
21        else
22           $s' \leftarrow ([\mathcal{M}_1, \dots, \mathcal{M}_{k-1}])$ 
23      else
24         $s' \leftarrow s$ 
25       $m' \leftarrow m, v' \leftarrow A, v \leftarrow Q$        // pass query on to adversary
26    return  $(s', v', m')$ 

```

---

(2023) correspond to the filter and odometer we construct in Algorithm 2 and Algorithm 4, respectively.

**Example B.2** (Example 3.7 restated). For every  $\delta' > 0$ ,

$$\begin{aligned}
& \mathcal{F}((\epsilon_1, \delta_1), \dots, (\epsilon_k, \delta_k); (\epsilon, \delta)) \\
&= \mathbb{I} \left( \epsilon \leq \sqrt{2 \log \left( \frac{1}{\delta'} \right) \sum_{m \leq k} \epsilon_m^2 + \frac{1}{2} \sum_{m \leq k} \epsilon_m^2} \right) \cdot \mathbb{I} \left( \delta' + \sum_{m \leq k} \delta_m \leq \delta \right)
\end{aligned}$$

is a valid  $(\epsilon, \delta)$ -DP privacy filter.

*Proof.* This filter is translated from the following theorem from Whitehouse et al. (2023), who define an  $(\epsilon, \delta)$ -DP privacy filter in their notation as follows:

---

**Algorithm 15:** Interactive postprocessing  $\mathcal{P}$  that transforms  $\mathcal{M}$  into  $\mathcal{F}(\cdot; d)$ - $\text{FiltCon}(\text{IM})$ .

---

```

1 def  $\mathcal{P}(s, v, m)$ :
2   if  $s = \lambda$  then                                     // initialize  $\mathcal{P}$ 
3      $s \leftarrow ([ ])$ , where  $[ ]$  is an empty list
4   if  $v = Q$  then                                       // accepting queries from adversary
5     if  $m$  is of form  $m = (\mathcal{M}', d')$  then
6       Parse  $s$  as  $([\mathcal{M}_1, \dots, \mathcal{M}_k])$ 
7        $s' \leftarrow ([\mathcal{M}_1, \dots, \mathcal{M}_k, (\mathcal{M}_{k+1}, \text{pending})])$ 
8     else if  $m$  is of form  $m = (j, q)$  where  $j = 1, \dots, k$  and  $q$  is a query to  $\mathcal{M}_j$ 
9       then
10       $s' \leftarrow s$ 
11     else if  $m$  is of form  $m = \text{privacy\_loss}$  then      // remove support for
12       privacy loss queries
13        $m' \leftarrow \text{invalid query}$ 
14       return  $(s', A, m')$ 
15      $m' \leftarrow m$ 
16      $v' \leftarrow Q, v \leftarrow A$                        // pass query on to mechanism
17   else if  $v = A$  then                                   // accepting answers from filter
18     if  $s$  is of form  $s = ([\mathcal{M}_1, \dots, (\mathcal{M}_k, \text{pending})])$  then
19       if  $m$  is not an error message then              // successfully started new
20       mechanism
21        $s' \leftarrow ([\mathcal{M}_1, \dots, \mathcal{M}_k])$ 
22     else
23        $s' \leftarrow ([\mathcal{M}_1, \dots, \mathcal{M}_{k-1}])$ 
24   else
25      $s' \leftarrow s$ 
26    $m' \leftarrow m, v' \leftarrow A, v \leftarrow Q$        // pass query on to adversary
27   return  $(s', v', m')$ 

```

---

**Theorem B.3** ( $(\epsilon, \delta)$ -DP privacy filter Whitehouse et al. (2023)). Suppose  $\mathcal{M}_1, \mathcal{M}_2, \dots$  is a sequence of mechanisms such that, for any  $k \geq 1$ ,  $\mathcal{M}_k$  is  $(\epsilon_k, \delta_k)$ -DP conditioned on  $\mathcal{M}_1, \dots, \mathcal{M}_{k-1}$ . ( $\mathcal{M}_1, \mathcal{M}_2, \dots$  are random variables that take on values that are the possible outputs of the respective mechanism on a dataset.) Let  $\epsilon > 0$  and  $\delta = \delta' + \delta''$  be a target approximation parameter such that  $\delta' > 0$ ,  $\delta'' \geq 0$ . Define  $N := N((\epsilon_1, \delta_1), (\epsilon_2, \delta_2), \dots) := \inf\{k \in \mathbb{N} : \epsilon < \sqrt{2 \log(\frac{1}{\delta'})} \sum_{m \leq k+1} \epsilon_m^2 + \frac{1}{2} \sum_{m \leq k+1} \epsilon_m^2 \text{ or } \delta'' < \sum_{m \leq k+1} \delta_m\}$ . Then the composition of these mechanisms  $\text{Comp}(\mathcal{M}_1, \dots, \mathcal{M}_{N(\cdot)})$  is  $(\epsilon, \delta)$ -DP.

$\mathcal{M}_k$  being  $(\epsilon_k, \delta_k)$ -DP conditioned on the previous mechanisms means the privacy-loss parameters are chosen adaptively. This is analogous to our setting for filters as defined in Algorithm 2, in which the adversary can send in a query that decomposes into a mechanism

and distance, a signal to the filter to start a new mechanism. Since the query could depend on the history of interaction, the new mechanism's parameters are adaptively chosen.

We define  $\mathcal{F} : \mathcal{D}^* \rightarrow \{1, 0\}$  as

$$\begin{aligned} & \mathcal{F}((\epsilon_1, \delta_1), \dots, (\epsilon_k, \delta_k); (\epsilon, \delta)) \\ &= \mathbb{I} \left( \sqrt{2 \log \left( \frac{1}{\delta'} \right) \sum_{m \leq k} \epsilon_m^2} + \frac{1}{2} \sum_{m \leq k} \epsilon_m^2 \leq \epsilon \right) \cdot \mathbb{I} \left( \delta' + \sum_{m \leq k} \delta_m \leq \delta \right), \end{aligned}$$

meaning that a mechanism equipped with  $\mathcal{F}$  will continue until for some  $k$ , either:

$$\sqrt{2 \log \left( \frac{1}{\delta'} \right) \sum_{m \leq k+1} \epsilon_m^2} + \frac{1}{2} \sum_{m \leq k+1} \epsilon_m^2 > \epsilon \quad \text{or} \quad \sum_{m \leq k+1} \delta_m > \delta''.$$

This is equivalent to [Whitehouse et al. \(2023\)](#)'s notion of  $N(\cdot)$ . Both the condition on  $\epsilon$  and the condition on  $\delta''$  are increasing in the number of mechanisms being composed. The difference is that in [Whitehouse et al. \(2023\)](#) setting, the number of mechanisms being composed  $N(\cdot)$  is the smallest number for adding another mechanism will cause  $\sqrt{2 \log \left( \frac{1}{\delta'} \right) \sum_{m \leq N(\cdot)+1} \epsilon_m^2} + \frac{1}{2} \sum_{m \leq N(\cdot)+1} \epsilon_m^2$  to exceed  $\epsilon$ , or in other words,  $N(\cdot)$  is the largest number for which  $\sqrt{2 \log \left( \frac{1}{\delta'} \right) \sum_{m \leq N(\cdot)} \epsilon_m^2} + \frac{1}{2} \sum_{m \leq N(\cdot)} \epsilon_m^2 \leq \epsilon$ . Since we define our filter in terms of a bound on the privacy parameter instead of a bound on the number of mechanisms, this is equivalent to checking that for  $k$  mechanisms being composed,  $\sqrt{2 \log \left( \frac{1}{\delta'} \right) \sum_{m \leq k} \epsilon_m^2} + \frac{1}{2} \sum_{m \leq k} \epsilon_m^2 \leq \epsilon$ . The same holds for the  $\sum_{m \leq k} \delta_m \leq \delta''$  term. Since  $\delta'$  is already set, we fold the  $\sum_{m \leq k} \delta_m \leq \delta''$  requirement into the  $\delta' + \delta'' \leq \delta$  requirement to get the condition  $\delta' + \sum_{m \leq k} \delta_m \leq \delta$ . Taking the lower bound of the  $k$  that satisfies either  $\sqrt{2 \log \left( \frac{1}{\delta'} \right) \sum_{m \leq k} \epsilon_m^2} + \frac{1}{2} \sum_{m \leq k} \epsilon_m^2 \leq \epsilon$  or  $\delta' + \sum_{m \leq k} \delta_m \leq \delta$  is equivalent to finding the  $k$  that satisfies both conditions.

With fully adaptive composition of noninteractive mechanisms, the mechanism will only halt if starting a new mechanism would exceed its privacy-loss budget. Since  $N(\cdot)$  mechanisms composed are  $(\epsilon, \delta)$ -DP in [Whitehouse et al. \(2023\)](#)'s setting, we know that a mechanism equipped with  $\mathcal{F}$  that has not yet halted will be  $(\epsilon, \delta)$ -DP as well.

Then by [Definition 2.6](#),  $\mathcal{F}\text{-Filt}(NIM)$  is an  $(\epsilon, \delta)$ -DP interactive mechanism, meaning that  $\mathcal{F}$  is a valid  $(\epsilon, \delta)$ -DP for the sequence of  $(\epsilon_i, \delta_i)$ -DP noninteractive mechanisms as defined above.  $\square$

To define a concurrent privacy odometer, we make use of [Lemma 1.26](#).

**Theorem B.4** ( $(\epsilon, \delta)$ -DP privacy odometer). Let  $\delta = \delta' + \delta''$  be a target approximation parameter such that  $\delta' > 0$ ,  $\delta'' \geq 0$ . Then

$$\begin{aligned} & \mathcal{G}((\epsilon_1, \delta_1), \dots, (\epsilon_k, \delta_k)) \\ &= \begin{cases} \left( \sqrt{2 \log \left( \frac{1}{\delta'} \right) \sum_{m \leq k} \epsilon_m^2} + \frac{1}{2} \sum_{m \leq k} \epsilon_m^2, \delta \right) & \delta' + \sum_{m \leq k} \delta_m \leq \delta \\ (\infty, \infty) & \text{otherwise} \end{cases} \end{aligned}$$

is a valid  $(\epsilon, \delta)$ -DP privacy-loss accumulator for noninteractive mechanisms, where  $\epsilon = \sqrt{2 \log \left( \frac{1}{\delta'} \right) \sum_{m \leq k} \epsilon_m^2} + \frac{1}{2} \sum_{m \leq k} \epsilon_m^2$ .