

How to Make Your Approximation Algorithm Private: A Black-Box Differentially-Private Transformation for Tunable Approximation Algorithms of Functions with Low Sensitivity

Jeremiah Blocki ✉

Purdue University, West Lafayette, IN, USA

Elena Grigorescu ✉

Purdue University, West Lafayette, IN, USA

Tamalika Mukherjee ✉

Purdue University, West Lafayette, IN, USA

Samson Zhou ✉

University of California Berkeley, CA, USA

Rice University, Houston, TX, USA

Abstract

We develop a framework for efficiently transforming certain approximation algorithms into differentially-private variants, in a black-box manner. Specifically, our results focus on algorithms A that output an approximation to a function f of the form $(1 - \alpha)f(x) - \kappa \leq A(x) \leq (1 + \alpha)f(x) + \kappa$, where $\kappa \in \mathbb{R}_{\geq 0}$ denotes additive error and $\alpha \in [0, 1)$ denotes multiplicative error can be “tuned” to small-enough values while incurring only a polynomial blowup in the running time/space. We show that such algorithms can be made differentially private without sacrificing accuracy, as long as the function f has small “global sensitivity”. We achieve these results by applying the “smooth sensitivity” framework developed by Nissim, Raskhodnikova, and Smith (STOC 2007).

Our framework naturally applies to transform non-private FPRAS and FPTAS algorithms into ϵ -differentially private approximation algorithms where the former case requires an additional postprocessing step. We apply our framework in the context of sublinear-time and sublinear-space algorithms, while preserving the nature of the algorithm in meaningful ranges of the parameters. Our results include the first (to the best of our knowledge) ϵ -edge differentially-private sublinear-time algorithm for estimating the number of triangles, the number of connected components, and the weight of a minimum spanning tree of a graph whose accuracy holds with high probability. In the area of streaming algorithms, our results include ϵ -DP algorithms for estimating L_p -norms, distinct elements, and weighted minimum spanning tree for both insertion-only and turnstile streams. Our transformation also provides a private version of the smooth histogram framework, which is commonly used for converting streaming algorithms into sliding window variants, and achieves a multiplicative approximation to many problems, such as estimating L_p -norms, distinct elements, and the length of the longest increasing subsequence.

2012 ACM Subject Classification Security and privacy → Usability in security and privacy

Keywords and phrases Differential privacy, approximation algorithms

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2023.59

Category RANDOM

Related Version *Full Version*: <https://arxiv.org/abs/2210.03831>

Funding *Jeremiah Blocki*: Supported in part by NSF CCF-1910659, NSF CNS-1931443, and NSF CAREER award CNS-2047272.

Elena Grigorescu: Supported in part by NSF CCF-1910659, NSF CCF-1910411, and NSF CCF-2228814.



© Jeremiah Blocki, Elena Grigorescu, Tamalika Mukherjee, and Samson Zhou; licensed under Creative Commons License CC-BY 4.0

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2023).

Editors: Nicole Megow and Adam D. Smith; Article No. 59; pp. 59:1–59:24



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Tamalika Mukherjee: Supported in part by NSF CCF-1910659 and NSF CCF-2228814.

Samson Zhou: Work done in part while at Carnegie Mellon University and supported in part by a Simons Investigator Award and by NSF CCF-1815840.

1 Introduction

Approximation algorithms are often used to efficiently approximate a function $f : \mathcal{D} \rightarrow \mathbb{R}^+$ in settings where resource constraints prevent us from computing the function exactly. For example, problems such as Knapsack are NP-Hard and, unless $P = NP$, do not admit a polynomial time solution. However, the Knapsack problem admits a fully polynomial time approximation scheme (FPTAS) i.e., for any $\alpha > 0$ there is a deterministic algorithm, running in time $\text{poly}(n, 1/\alpha)$, which outputs a solution that is guaranteed to be nearly as good (up to multiplicative factor $1 \pm \alpha$) as the optimal solution. As a second example, even if the problem is computationally tractable it may still be the case that the input dataset $D \in \mathcal{D}$ is extremely large, making it infeasible to load the entire dataset into RAM, or impractical to execute a linear time algorithm. To remedy such shortcomings, models such as sublinear-space and sublinear-time algorithms have been proposed. For example, one may want to estimate frequencies of elements that appear in a stream of n elements up to a multiplicative $1 \pm \alpha$ factor, while using only $\text{poly}(\log n, \frac{1}{\alpha})$ memory cells. Or, one may want to estimate the number of connected components of a dense graph on n vertices up to (relative) additive error κ by only inspecting $\text{poly}(\log n, \frac{1}{\kappa})$ many edges.

In addition to time and space efficiency, user privacy is another important consideration in contexts where the input to our function f is sensitive user data. Differential privacy (DP) [20, 22] is a rigorous mathematical concept that gives provable guarantees on what it means for an algorithm to preserve the privacy of individual information in the input dataset. Informally, a randomized function computed on a dataset D is *differentially private* if the distribution of the function's output does not change significantly with the presence or absence of an individual data point. Thus, a natural goal is to develop efficient differentially private algorithms to approximate functions/queries of interest.

One general way to preserve differential privacy is to add noise scaled to the global sensitivity Δ_f of our function f , i.e., the maximum amount $|f(D) - f(D')|$ that the answer could change by modifying a single record in our dataset D to obtain a new dataset D' . This general approach yields efficient and accurate approximations for f as long as we have an efficient algorithm to compute f *exactly* and the global sensitivity of f is sufficiently small. However, in some resource-constrained settings, we may need to use an approximation algorithm \mathcal{A}_f instead of evaluating f exactly. Unfortunately, the mechanism that computes $\mathcal{A}_f(D)$ and then adds noise scaled to the global sensitivity Δ_f of our function f is not necessarily differentially private. In particular, even if we are guaranteed that $|\mathcal{A}_f(D) - f(D)| \leq \alpha f(D)$ we might still have $|\mathcal{A}_f(D) - \mathcal{A}_f(D')| \geq 2\alpha f(D) \gg \Delta_f$ for neighboring datasets D and D' , e.g., suppose $\mathcal{A}_f(D) = (1 + \alpha)f(D)$ and $\mathcal{A}_f(D') = (1 - \alpha)f(D')$. Thus, the global sensitivity of \mathcal{A}_f can be quite large and adding noise proportional to $\Delta_{\mathcal{A}_f}$ would prevent us from providing meaningful accuracy guarantees. This raises a natural question: Suppose that our function f admits an accurate (but not necessarily resource-efficient) differentially private approximation algorithm and that f also admits an efficient (but not necessarily private) approximation algorithm. Is it necessarily the case that there is also an equally efficient differentially private approximation algorithm?

Unfortunately, a result of [36, 18] suggests that the answer to the previous question may be no. Suppose our dataset D consists of n users x_1, \dots, x_n with n binary attributes i.e., $x_i \in \{0, 1\}^n$. Consider the function $f(D)$ that computes all of the one-way marginals i.e.,

$f(D) = \langle \frac{1}{n} \sum_{i=1}^n x_i[j] \rangle_{j=1}^n \in \mathbb{R}^n$. In particular, there is a non-private sublinear time algorithm that samples $\mathcal{O}(\log n)$ users and (with high probability) outputs a good approximation to *all* n one-way marginals. However, if we require that our algorithm satisfy pure, i.e., ϵ -differential privacy (resp. approximate, i.e., (ϵ, δ) -differential privacy) then we need to look at *at least* $\Omega(n/\epsilon)$ (resp. $\Omega(\sqrt{n \log(1/\delta)})$) samples [36, 18]. In light of this, we pose the following general questions:

What are sufficient conditions for an approximation algorithm to be made differentially private?

Can an approximation algorithm be made differentially private in an efficient black-box manner?

Over the years, many differentially private approximation algorithms have been developed for problems in optimization, machine learning, and distribution testing (see for e.g., [35, 1, 30, 34]), in a somewhat ad-hoc manner. Often, these results give a differentially private algorithm for that specific problem and do not easily generalize to give differentially private algorithms for a large class of problems. A general framework for developing differentially private approximation algorithms for a large class of problems is desirable as this would not only make DP approximation algorithms more easily accessible to non-DP experts, but more importantly, it would shed light on what kinds of algorithms are more amenable to differential privacy. Furthermore, a framework that uses the underlying approximation algorithm as a *black-box* is desirable as this avoids the need to (re)design, (re)analyze, and (re)implement the new differentially private versions of these approximation algorithms. We emphasize that this type of framework has been well-studied for *computing* functions privately by calibrating noise proportional to their global or smooth sensitivity [23, 41] (see Section A for more discussion).

Our work makes partial progress towards answering these general questions. In particular, we give an efficient black-box framework for converting a non-private approximation algorithm \mathcal{A}_f with tunable accuracy parameters into a differentially private approximation algorithm \mathcal{A}'_f with reasonable accuracy guarantees as long as the global sensitivity Δ_f of the function f being approximated is sufficiently low. For the case when \mathcal{A}_f is deterministic, we achieve a pure ϵ -differentially private approximation algorithm via a direct transformation, and when \mathcal{A}_f is randomized, i.e., has a small failure probability, we achieve ϵ -differential privacy by first applying a transformation that gives a (ϵ, δ) -differentially private algorithm and then apply a postprocessing step to achieve ϵ -differential privacy. For example, suppose that for any $\alpha > 0$ our algorithm \mathcal{A}_f , taking α and our dataset D as input, provides the guarantee that $|\mathcal{A}_f(D) - f(D)| \leq \alpha f(D)$ e.g., any FPTAS algorithm would satisfy our tunable accuracy requirement. In such a case, for any $\alpha > 0$ we can transform our non-private algorithm \mathcal{A}_f into a differentially private version with multiplicative error α and small additive error term which (necessarily) comes from the noise that we added. Intuitively, we exploit the fact that we can run \mathcal{A}_f with an even smaller accuracy parameter $\rho \ll \alpha$ which can be tuned to ensure that the smooth sensitivity of our algorithm is sufficiently small. Our same general framework still applies if we allow that the approximation algorithm \mathcal{A}_f has a small additive error term i.e., $|\mathcal{A}_f(D) - f(D)| \leq \alpha f(D) + \kappa$. If $\mathcal{A}_f(D)$ is only guaranteed to output a good approximation (i.e., $|\mathcal{A}_f(D) - f(D)| \leq \alpha f(D) + \kappa$) with probability $1 - \delta/2$ (e.g., an FPRAS algorithm would satisfy this requirement with additive error $\kappa = 0$) then our framework achieves ϵ -differential privacy by first obtaining an approximate (ϵ, δ) -differential privacy algorithm and then a postprocessing step. In cases where the approximation algorithm is not tunably

accurate our black-box framework does not necessarily apply¹. For example, the best known approximation algorithms for vertex cover achieve the guarantee $f(G) \leq \mathcal{A}_f(G) \leq 2f(G)$ i.e., because there is no way to control the smooth sensitivity of our approximation algorithm.

1.1 Our Contributions

We introduce a generic black-box framework for converting certain approximation algorithms for a function $f : \mathcal{D} \rightarrow \mathbb{R}^+$ into a differentially private approximation algorithm using smooth sensitivity [41]. We first introduce a definition for *tunable* approximation algorithms used throughout our paper, and then present our main results for the DP framework, and then give new differentially private algorithms for a variety of approximation algorithms obtained via this unifying framework. We refer to our full version [11] for missing details.

► **Definition 1** ((α, κ, δ) -approximation). *An algorithm \mathcal{A}_f is a (α, κ, δ) -approximation for f if for every $D \in \mathcal{D}$ with probability at least $1 - \delta$, we have that $(1 - \alpha)f(D) - \kappa \leq \mathcal{A}_f(D) \leq (1 + \alpha)f(D) + \kappa$.*

We may abuse notation and omit the failure probability δ parameter in the above definition, if it is clear from the context. Some algorithms \mathcal{A}_f may take the approximation parameters $\alpha, \kappa, \delta \geq 0$ as input².

► **Definition 2** (tunable approximation). *$\mathcal{A}_f(D, \alpha, \kappa, \delta)$ provides a tunable approximation of f if for every $\alpha, \kappa, \delta \geq 0$ the algorithm $\mathcal{A}_f(\cdot, \alpha, \kappa, \delta)$ obtained by hardcoding α, κ and δ is a (α, κ, δ) -approximation for f .*

When the parameters α, κ, δ are clear from the context, we may abuse notation and just write $\mathcal{A}_f(D)$. For a tunable approximation algorithm we will use $R(n, \alpha, \kappa, \delta)$ to denote the amount of a particular resource used by the algorithm. The resources we consider in this work include time, space and query complexity of the algorithm (depending on the model) which we denote by $T(\cdot, \cdot, \cdot, \cdot)$, $S(\cdot, \cdot, \cdot, \cdot)$, and $Q(\cdot, \cdot, \cdot, \cdot)$ respectively.

As a concrete example any FPTAS algorithm \mathcal{A}_f for f would be a tunable approximation for f with running time $T(n, \alpha, \kappa, \delta) = \text{poly}(n, 1/\alpha)$ for any $\alpha > 0$ and any $\kappa, \delta \geq 0$ – an FPTAS has no additive error ($\kappa = 0$) and zero failure probability ($\delta = 0$). Similarly a FPRAS algorithm would be a tunable approximation with running time $T(n, \alpha, \kappa, \delta) = \text{poly}(n, \alpha, \log(1/\delta))$ for any $\alpha, \delta > 0$ and any $\kappa \geq 0$ – an FPRAS also has no additive error ($\kappa = 0$).

General Framework for Approximation Algorithms. Our main result gives a framework for converting any existing non-DP algorithm \mathcal{A}_f that provides an (α, κ, δ) -approximation of f into an ε -DP algorithm \mathcal{A}_f'' in the following manner: (1) Apply Algorithm 1 to obtain an (ε, δ) -DP algorithm \mathcal{A}_f' that achieves an $(\alpha', \kappa', \delta')$ -approximation (see Theorem 3), (2) Apply a postprocessing step on the output of \mathcal{A}_f' outlined in Theorem 5 to achieve an ε -DP algorithm \mathcal{A}_f'' with the same accuracy guarantees as \mathcal{A}_f' barring an additive error of $o(1)$. We emphasize that \mathcal{A}_f is a tunable approximation, in other words, \mathcal{A}_f takes the parameters (α, κ, δ) as input.

¹ One could still apply our black-box transformation. However, the accuracy guarantees would be degraded and we would only achieve (ε, δ) -differential privacy for sufficiently large values of $\varepsilon, \delta > 0$ which depend on the approximation error parameter α .

² We allow that $\alpha = \kappa = \delta = 0$ in which case \mathcal{A}_f can simply compute f exactly – whether or not this computation is efficient.

► **Theorem 3** ((ε, δ) -privacy). *Suppose that \mathcal{A}_f is a tunable approximation of $f : \mathcal{D} \rightarrow \mathbb{R}^+$. Then for all $\varepsilon > 0$, $\delta = \delta(n) > 0^3$, $\alpha \geq 0$ and $\kappa \geq 0$, there is an algorithm \mathcal{A}'_f such that*

- (1) (Privacy) \mathcal{A}'_f is (ε, δ') -differentially private where $\delta' = \delta(1 + \exp(\varepsilon/2))$.
- (2) (Accuracy) For all $D \in \mathcal{D}$, and $0 < \gamma$, with probability $1 - \delta - \exp(-\gamma)$,

$$(1 - \alpha')f(D) - \kappa' - \frac{2\Delta_f}{\varepsilon} \cdot \gamma \leq \mathcal{A}'_f(D) \leq (1 + \alpha')f(D) + \kappa' + \frac{2\Delta_f}{\varepsilon} \cdot \gamma$$

where $\alpha' = \frac{\alpha(\varepsilon+16\gamma)}{12\log(4/\delta)}$, $\kappa' = \kappa \left(\frac{2\gamma\alpha}{3\log(4/\delta)} + \frac{8\gamma}{\varepsilon} + 1 \right)$, and $\Delta_f := \max_{D, D' \in \mathcal{D}, D \sim D'} \|f(D) - f(D')\|_1$.

- (3) (Resource) \mathcal{A}'_f uses $R\left(n, \frac{\varepsilon\alpha}{\log(4/\delta)}, \kappa, \delta\right)$ resource, where $R(\cdot, \cdot, \cdot, \cdot)$ is the resource used by \mathcal{A}_f .

We illustrate the utility of Theorem 3 with specific parameters – if we have a non-private algorithm \mathcal{A}_f that guarantees an $(\alpha, 0, \delta)$ -approximation, then for constant ε , $\delta = \frac{1}{n^c}$ and $\gamma = c \log(n)$, we see that the DP algorithm \mathcal{A}_f achieves an $\left(\alpha(1 + o(1)), \mathcal{O}\left(\frac{\Delta_f \log(n)}{\varepsilon}\right), \frac{1}{n^c}\right)$ -approximation. We typically use these parameters for δ, γ in our applications for streaming and sublinear-time algorithms.

Our reduction in Theorem 3 is quite simple – we describe the associated Algorithm 1 below.

■ **Algorithm 1** (ε, δ) -differentially private framework \mathcal{A}'_f for tunable approximation algorithm \mathcal{A}_f .

Input: Input set D , accuracy parameters $\alpha \in (0, 1)$ and κ , DP parameter ε , DP failure probability $\delta \in (0, 1)$, approx. algorithm \mathcal{A}_f .

- 1: Let $x_A := \mathcal{A}_f(D, \rho, \tau, \delta/2)$, where $\rho := \left(\frac{\varepsilon\alpha}{12\log(4/\delta)}\right)$, and $\tau := \kappa$.
- 2: **return** $x_A + X$ where $X \sim \text{Lap}\left(\frac{2(4\rho x_A + 4\tau + \Delta_f)}{\varepsilon}\right)$

Note that in Algorithm 1, we leave our additive parameter κ as is when running \mathcal{A}_f , but we still choose to define $\tau := \kappa$. This is because depending on the problem, and the accuracy/efficiency guarantees desired, we can set τ to be a tuned version of κ (for e.g., we set $\tau := \kappa/\log(n)$ for the problem of estimating the number of connected components).

► **Remark 4.** We also note that, even if the failure probability $\delta > 0$ of \mathcal{A}_f is non-negligible, that we can always boost the success probability by running $\mathcal{A}_f(D)$ multiple times and computing the median over all outputs. Even if the error rate $0 < \delta < 1/2$ is a constant we can always reduce the failure probability to a lower target $0 < \delta' \ll \delta$ while increasing the running time by a multiplicative factor $\mathcal{O}(\log(1/\delta'))$. In particular, we can set δ' to be a negligible function of n such as $\delta' = n^{-\log n}$ whilst only incurring a $\mathcal{O}(\log^2 n)$ blowup in our running time.

We stress that we can only apply Theorem 3 to existing non-DP algorithms \mathcal{A}_f that give an approximation guarantee of the form $(1 - \alpha)f(D) - \kappa \leq \mathcal{A}_f \leq (1 + \alpha)f(D) + \kappa$. For example, we cannot apply Theorem 3 to obtain an (ε, δ) -DP algorithm for estimating the minimum vertex cover size in sublinear time. This is because the non-DP sublinear-time algorithm \mathcal{A}_{vc} has an approximation guarantee of the form $2VC(G) - \kappa n \leq \mathcal{A}_{vc} \leq 2VC(G) + \kappa n$. On the other hand, we can use our DP framework to obtain an (ε, δ) -DP algorithm for obtaining a

³ typically we set $\delta = \text{negl}(n)$ or $\delta = n^{-c}$ for some constant $c > 0$. In particular $\delta(n)$ may approach zero as $n \rightarrow \infty$.

$(0, \kappa n, \delta)$ -approximation of the maximum matching size in sublinear time (see full version for details [11]). Intriguingly, both the minimum vertex cover size and the maximum matching size algorithms use the same underlying strategy of estimating a greedy maximal matching in a local fashion, but since they return different estimators based on the objective and we can only use our framework as a black-box, we cannot apply our framework to the former while we can still apply it to the latter.

Finally, by applying a post-processing step described below, we show how to obtain an ε -DP algorithm from the (ε, δ) -DP algorithm obtained in Theorem 3. Importantly, the accuracy guarantee of the resulting ε -DP algorithm only differs by a small additive factor of $1/(KM)$, where $M = \max_D f(D)$ is the maximum possible output value, e.g., $M \leq n^3$ for triangle counting and $K > 0$. Moreover for negligible δ , the accuracy guarantee of the resulting pure DP algorithm still holds with high probability.

► **Theorem 5.** *Let $M = \max_D f(D)$ and let parameter $K > 0$. If $\mathcal{A}'_f(D)$ is (ε, δ) -DP algorithm with accuracy guarantee $(1 - \alpha)f(D) - \kappa \leq \mathcal{A}'_f(D) \leq (1 + \alpha)f(D) + \kappa$ holding with probability $1 - \eta$ then there exists an algorithm $\mathcal{A}''_f(D)$ which is ε -DP with accuracy guarantee $(1 - \alpha)f(D) - \kappa - \frac{1}{KM} \leq \mathcal{A}''_f(D) \leq (1 + \alpha)f(D) + \kappa + \frac{1}{KM}$ with probability at least $1 - \eta - p$ where $p = \frac{\delta K(M+1)}{e^\varepsilon - 1 + \delta K(M+1)}$.*

Our second result is an analogous framework for converting any existing deterministic non-DP approximation algorithm \mathcal{A}_f that provides an $(\alpha, \kappa, 0)$ -approximation of f into an ε -DP algorithm \mathcal{A}'_f .

► **Theorem 6 (ε -privacy).** *Suppose that \mathcal{A}_f is a deterministic tunable approximation of $f : \mathcal{D} \rightarrow \mathbb{R}^+$. Then for all $\varepsilon > 0$, $\alpha \geq 0$ and $\kappa \geq 0$, there is an algorithm \mathcal{A}'_f such that*

- (1) (Privacy) \mathcal{A}'_f is ε -differentially private.
- (2) (Accuracy) For all $D \in \mathcal{D}$, we have that with probability $\geq 9/10$,

$$(1 - \alpha')f(D) - \kappa' - \frac{7\Delta_f}{\varepsilon} \leq \mathcal{A}'_f(D) \leq (1 + \alpha')f(D) + \kappa' + \frac{7\Delta_f}{\varepsilon}$$

where $\alpha' := \alpha C_1(\varepsilon + C_2\gamma)$, $\kappa' := \kappa C_3(\alpha + \frac{C_4}{\varepsilon})$ for some constants $C_1, C_2, C_3, C_4 > 0$ and $\Delta_f := \max_{D, D' \in \mathcal{D}, D \sim D'} \|f(D) - f(D')\|_1$.

- (3) (Resource) \mathcal{A}'_f uses $R(n, \frac{\varepsilon\alpha}{36}, \kappa)$ resource, where $R(\cdot, \cdot, \cdot)$ is the resource used by \mathcal{A}_f .

DP Sublinear-time Results. We use Theorem 3 in conjunction with Theorem 5 in a black-box manner to obtain pure differentially-private sublinear time algorithms for several problems (see Table 1 for a summary).

In many models of sublinear-time computation the efficiency of the algorithm is measured in the number of queries made to the input, rather than the time complexity of the algorithm. It is often the case that the two are polynomially related, but there are instances in which the actual time complexity of the algorithm may be exponentially larger than the query complexity, in terms of the approximation factor. Nevertheless, in these instances too, the literature uses time and query complexity interchangeably. This is because the sublinear-time model assumes restricted or expensive access to the input, while further computation on local machines with the answers obtained from queries is considered to be cheap. We use query complexity for the sake of clarity.

We note that in the sublinear-time literature, the approximation parameters α, κ are usually considered to be a constant, but the analyses for most of these theorems hold for $\alpha = \alpha(n), \kappa = \kappa(n) \in (0, 1)$, where n is the input size.

Here we do not explicitly define the sublinear model (or the queries allowed) for each problem. For a graph G we denote the number of vertices as n , the number of edges as m , and the average degree of the graph as \bar{d} .

Typically, the accuracy guarantees of the non-DP results are presented with probability at least $2/3$ – in order to apply our framework, we apply the median trick (see Remark 4) to boost the probability of success to $1 - \delta$. For simplicity of comparing our results, for any constant $c > 0$, we set $\delta := 1/n^c$ in the sequel.

■ **Table 1** Summary of Sublinear-time DP graph algorithms obtained via our black-box DP transformation. According to our notation multiplicative error α means a multiplicative factor of $(1 \pm \alpha)$.

Problem	Reference	Privacy	Mult. error	Add. error	Query Complexity
Number of Triangles	[24]	Non-Private	α	0	$\mathcal{O}\left(\left(\frac{n}{t^{1/3}} + \frac{m^{3/2}}{t}\right) \text{poly}(\log(n), \frac{1}{\alpha})\right)$
	This Work	ε -edge DP	α	$\mathcal{O}\left(\frac{n \log(n)}{\varepsilon}\right)$	$\mathcal{O}\left(\left(\frac{n}{t^{1/3}} + \frac{m^{3/2}}{t}\right) \text{poly}(\log(n), \frac{1}{\alpha\varepsilon})\right)$
Connected Components	[7]	Non-Private	0	κn	$\mathcal{O}\left(\frac{1}{\varepsilon^2} \log\left(\frac{1}{\kappa}\right) \log(n)\right)$
	This Work	ε -edge DP	0	$\mathcal{O}(\kappa n) + \mathcal{O}\left(\frac{\log(n)}{\varepsilon}\right)$	$\mathcal{O}\left(\frac{\log^2(n)}{\kappa^2} \log\left(\frac{\log(n)}{\kappa}\right)\right)$
Weighted MST	[19]	Non-Private	α	0	$\mathcal{O}\left(\bar{d} w \alpha^{-2} \log\left(\frac{\bar{d} w}{\alpha}\right) \log(n)\right)$
	This Work	ε -edge DP	α	$\mathcal{O}\left(\frac{\log(n)}{\varepsilon}\right)$	$\mathcal{O}\left(\bar{d} w \frac{\log^2(n)}{\alpha^2 \varepsilon^2} \log\left(\frac{\bar{d} w \log(n)}{\alpha \varepsilon}\right) \log(n)\right)$
Average Degree	[33]	Non-Private	α	0	$\mathcal{O}\left(\frac{n}{\sqrt{m}} \text{poly}\left(\frac{\log(n)}{\alpha}\right) \log(n)\right)$
	[10]	ε -edge DP	α	0	$\mathcal{O}\left(\sqrt{n} \text{poly}\left(\frac{\log(n)}{\alpha}\right) \text{poly}\left(\frac{1}{\varepsilon}\right) \log(n)\right)$ (analysis assumes $\bar{d} \geq 1$)
	This Work	ε -edge DP	α	0	$\mathcal{O}\left(\frac{n}{\sqrt{m}} \text{poly}\left(\frac{\log^2(n)}{\alpha \varepsilon}\right) \log(n)\right)$ for $\bar{d} = \Omega\left(\frac{\log(n)}{n\varepsilon}\right)$
Maximum Matching Size	[48]	Non-Private	0	κn	$\mathcal{O}\left(d^{\mathcal{O}(1/\kappa^2)} \log(n)\right)$
	This Work	ε -node DP	0	$\mathcal{O}\left(\frac{\kappa n}{\varepsilon}\right)$	$\mathcal{O}\left(d^{\mathcal{O}(1/\kappa^2)} \log(n)\right)$
Distance to Bipartiteness	[31]	Non-Private	0	κn^2	$\mathcal{O}\left((1/\kappa^3) \log(n)\right)$
	This Work	ε -edge DP	0	$\mathcal{O}(\kappa n^2) + \mathcal{O}\left(\frac{\log(n)}{\varepsilon}\right)$	$\mathcal{O}\left(\log^4(n)/\kappa^3\right)$

We give the first (to the best of our knowledge) ε -DP sublinear time algorithm for estimating the number of triangles, connected components, and the weight of a minimum spanning tree whose accuracy guarantees hold with high probability.

For estimating the average degree of a graph, in recent work, [10] gave a pure ε -DP algorithm that achieves an $(\alpha, 0)$ -approximation – a crucial observation is that their analysis only holds under the assumption that the average degree is at least one i.e., $\bar{d} \geq 1$ (see full version [11] for details). In this work, we remove the need for this assumption in the DP setting, by directly applying our black-box DP transformation to the original algorithm of [33] which works substantially better whenever we have $m = \omega(n)$ edges.

For estimating the maximum matching size in a graph, although [10] gave an ε -DP algorithm for estimating the maximum matching size that achieves a 2-multiplicative factor and κn additive factor, they left the task of finding an $(0, \kappa n)$ -approximation in the DP setting as an open problem. In this work, we partially resolve this problem by presenting an ε -DP algorithm that gives a $(0, \mathcal{O}\left(\frac{\kappa n}{\varepsilon}\right))$ -approximation of the maximum matching size. Crucially, our resulting analysis cannot guarantee that the added Laplace noise will be small with high probability, but only guarantees this will be the case with *constant probability*. This problem highlights a limitation of our black-box DP framework – if the non-DP algorithm that we want to apply our DP transformation on has a time/space/query complexity that has an exponential dependence on the approximation parameters then the resulting DP algorithm that achieves a similar approximation guarantee with high probability may be highly inefficient in terms of time/space/query complexity.

We also show how to apply our DP framework to an algorithm estimating the distance to bipartiteness in dense graphs [31, 3], which is accurate with probability $1 - o(1)$. The same reduction can be similarly applied to other natural properties that enjoy the feature that they admit distance-estimation algorithms with $\text{poly}(1/\kappa)$ query complexity, where κ is the additive (normalized) error. For example, in the fundamental results of [32] an efficient distance approximation algorithm for the maximum k -cut problem, and thus k -colorability is presented. [27], also based on results from [6], generalizes these properties to the notion of “semi-homogeneous partition properties” and show efficient distance estimation algorithms for properties such as Induced P_3 -freeness, induced P_4 -freeness, and chordality.⁴

DP Streaming Results. We also apply our framework given by Theorem 3 and Theorem 5 to obtain differentially-private streaming algorithms for many fundamental problems, i.e., see Table 2 and Table 3. We remark that while the accuracy guarantees of our resulting algorithms may be surpassed by recent works studying these problems on an individual basis, our applications are black-box reductions that avoid individual utility and privacy analysis of each non-private streaming algorithm, which can be heavily involved and quite non-trivial, e.g., [40, 9, 43, 17, 46, 13].

In the streaming model, elements of an underlying dataset arrive one-by-one and the goal is to compute or approximate some predetermined function on the dataset using space that is sublinear in the size of the dataset. Our reductions also have wide applications to various archetypes of data stream models, which we now discuss. In insertion-only streams, the updates of the stream increment the underlying dataset, such as adding edges to a graph, adding terms to a sequence, or increasing the coordinates of a frequency vector. In turnstile (or dynamic) streams, the updates of the stream can both increase and decrease (or insert and delete) elements of the underlying dataset. Finally, in the sliding window model, only the W most recent updates of the data stream define the underlying dataset. Both the turnstile streaming model and the sliding window model are generalizations of insertion-only streams, and our framework has implications in all three models.

We first show that our framework can be applied to existing non-private dynamic algorithms for weighted minimum spanning tree, L_p norm estimation for $p \geq 1$ (and also F_p moment estimation for $0 < p < 1$), and distinct elements estimation. Thus using our framework, we essentially get private dynamic algorithms for these problems for free (in terms of correctness, not optimality). Since the dynamic streaming model generalizes the insertion-only streaming model, we also obtain private streaming algorithms in the insertion-only model as well. We summarize these results in Table 2.

We then apply our framework in Theorem 3 to the sliding window model. To that end, we first recall that given a $(\alpha, 0)$ -approximation algorithm for the insertion-only streaming model, the smooth histogram framework [14] provides a transformation that obtains a $(\alpha, 0)$ -approximation algorithm in the sliding window model for a “smooth” function. Although there are problems that are known to not be smooth, e.g., [15, 12, 16, 25, 37], the smooth histogram framework does provide a $(\alpha, 0)$ -approximation to many important problems, such as counting, longest increasing subsequence, L_p norm estimation for $p \geq 1$ (and also F_p moment estimation for $0 < p < 1$), and distinct elements estimation. We remark that

⁴ In general, distance estimation is closely related to tolerant testing [42], and for dense graph properties it is known that if a property is testable with a number of queries of the form $f(\kappa)$, then they admit a distance estimator [28] with an exponential blowup in $\frac{1}{\kappa}$ in the query complexity. Hence, in its general form the query complexity of estimating the distance to “hereditary” graph properties is a tower of exponential of height $\text{poly}(1/\kappa)$ [4].

■ **Table 2** Summary of DP algorithms in the dynamic/turnstile model obtained via our black-box DP transformation. According to our notation multiplicative error α means a multiplicative factor of $(1 \pm \alpha)$.

Problem	Reference	Privacy	Mult. error	Add. error	Space Complexity
Weighted Minimum Spanning Tree	[2]	Non-Private	α	0	$\mathcal{O}\left(\frac{1}{\alpha} n \log^4 n\right)$
	This Work	ϵ -DP	α	$\mathcal{O}\left(\frac{M \log m}{\epsilon}\right)$	$\mathcal{O}\left(\frac{1}{\alpha \epsilon} n \log^5 n\right)$
L_p -norm, $p > 2$	[29]	Non-Private	α	0	$\mathcal{O}\left(\frac{1}{\alpha^2} n^{1-\frac{2}{p}} \log^2 n + \frac{1}{\alpha^{4/p}} n^{1-\frac{2}{p}} \log^{2/p} n \log^2 n\right)$
	This Work	ϵ -DP	α	$\mathcal{O}\left(\frac{\log m}{\epsilon}\right)$	$\mathcal{O}\left(\frac{1}{\alpha^{2+2/p}} n^{1-2/p}\right) \cdot \text{poly}\left(\log n, \log \frac{1}{\alpha \epsilon}\right)$
L_p -norm, $p = 2$	[5]	Non-Private	α	0	$\mathcal{O}\left(\frac{1}{\alpha^2} \log^2 n\right)$
	This Work	ϵ -DP	α	$\mathcal{O}\left(\frac{\log m}{\epsilon}\right)$	$\mathcal{O}\left(\frac{1}{\alpha^2} \log^4 n\right)$
L_p -norm, $p \in (0, 2)$	[38]	Non-Private	α	0	$\mathcal{O}\left(\frac{1}{\alpha^2} \log^2 n\right)$
	This Work	ϵ -DP	α	$\mathcal{O}\left(\frac{\log m}{\epsilon}\right)$	$\mathcal{O}\left(\frac{1}{\alpha^{2+2/p}} \log^4 n\right)$
L_p -norm, $p = 0$	[39]	Non-Private	α	0	$\mathcal{O}\left(\frac{1}{\alpha^2} \log^2 n \log \frac{1}{\alpha}\right)$
	This Work	ϵ -DP	α	$\mathcal{O}\left(\frac{\log m}{\epsilon}\right)$	$\mathcal{O}\left(\frac{1}{\alpha^{2+2/p}} \log^4 n\right)$

if we tried to apply the non-private smooth histogram framework to a DP insertion-only streaming algorithm, this might preserve privacy by post-processing, but may significantly increase the error in terms of accuracy. On the other hand, our framework avoids these issues and achieves private analogs of these algorithms in the sliding window model without compromising utility. We summarize our results for the sliding window model in Table 3. We note that in recent work, [26] give a generalized smooth histogram approach to convert a DP continual release streaming algorithm into a sliding window algorithm in the continual release setting. We focus on the one-shot streaming setting in our work.

■ **Table 3** Summary of DP algorithms in the sliding window model obtained via our black-box DP transformation. According to our notation multiplicative error α means a multiplicative factor of $(1 \pm \alpha)$.

Problem	Reference	Privacy	Mult. error	Add. error	Space Complexity
Longest Increasing Subsequence	[44]	Non-Private	α	0	$\mathcal{O}\left(\frac{k^2}{\alpha} \log^2 n\right)$
	This Work	ϵ -DP	α	$\mathcal{O}\left(\frac{\log m}{\epsilon}\right)$	$\mathcal{O}\left(\frac{k^2}{\alpha \epsilon} \log^4 n\right)$
Distinct Elements	[8]	Non-Private	α	0	$\mathcal{O}\left(\frac{1}{\alpha^3} \log^2 n\right)$
	This Work	ϵ -DP	α	$\mathcal{O}\left(\frac{\log m}{\epsilon}\right)$	$\mathcal{O}\left(\frac{1}{\alpha^3} \log^5 n\right)$
L_p -norm, $p = 2$	[47]	Non-Private	α	0	$\mathcal{O}\left(\frac{1}{\alpha^2} \log^3 n \log^3 \frac{1}{\alpha}\right)$
	This Work	ϵ -DP	α	$\mathcal{O}\left(\frac{\log m}{\epsilon}\right)$	$\mathcal{O}\left(\frac{1}{\alpha^{2+2/p}} \log^5 n \log^3 \frac{1}{\alpha \epsilon}\right)$
L_p -norm, $p \in (0, 2)$	[47]	Non-Private	α	0	$\mathcal{O}\left(\frac{1}{\alpha^2} \log^3 n (\log \log n)^2 \log^3 \frac{1}{\alpha}\right)$
	This Work	ϵ -DP	α	$\mathcal{O}\left(\frac{\log m}{\epsilon}\right)$	$\mathcal{O}\left(\frac{1}{\alpha^{2+2/p}} \log^5 n\right)$

1.2 Our Techniques

Given a tunable (α, κ, δ) -approximation algorithm \mathcal{A}_f for the function $f : \mathcal{D} \rightarrow \mathbb{R}^+$, our goal is to obtain a differentially private approximation algorithm that achieves a target $(\alpha', \kappa', \delta')$ -approximation of f where α', κ' are in terms of α, κ .

Warm-up: When \mathcal{A}_f is deterministic and only has multiplicative error. For simplicity, let us first consider an $(\alpha, 0, 0)$ -approximation algorithm \mathcal{A}_f , in other words, \mathcal{A}_f always outputs a value such that $(1 - \alpha)f(D) \leq \mathcal{A}_f(D) \leq (1 + \alpha)f(D)$. Since we want to make \mathcal{A}_f differentially private, intuitively, we need to add noise to the output of \mathcal{A}_f . The local sensitivity of \mathcal{A}_f at D (i.e., $LS_{\mathcal{A}_f}(D) = \max_{D' \sim D} |\mathcal{A}_f(D) - \mathcal{A}_f(D')|$) is upper bounded by $2\alpha f(D) + \Delta_f$. Since Δ_f is small and we can tune α to be arbitrarily small, it is tempting to think that we can just add noise proportional to $2\alpha f(D) + \Delta_f$. Unfortunately, scaling noise proportional to local sensitivity is not necessarily private. On the other hand we could ensure privacy by scaling noise proportional to the global sensitivity (i.e., $\max_{D \in \mathcal{D}} LS_{\mathcal{A}_f}(D) \leq \max_{D \in \mathcal{D}} 2\alpha f(D) + \Delta_f$)

but noise will likely be too large to obtain meaningful accuracy guarantees. We adopt the strategy of adding noise proportional to the smooth sensitivity [41] of \mathcal{A}_f instead. In particular, [41] observed that if we can find a “sufficiently smooth” function $S_f(D) \geq LS_{\mathcal{A}_f}(D)$ upper bounding the local sensitivity of \mathcal{A}_f then we can preserve privacy by computing $\mathcal{A}_f(D)$ and adding noise scaled according to $S_f(D)$.

We can show that the function $S_f(D) = 4\alpha\mathcal{A}_f(D) + \Delta_f$ is a β -smooth upper bound on the local sensitivity of \mathcal{A}_f for $\beta = 6\alpha$ where S_f is β -smooth if $S_f(D) \leq e^\beta S_f(D')$ for all pairs of neighboring datasets $D \sim D'$. To achieve privacy using the smooth sensitivity framework we need to ensure that β is sufficiently small relative to our privacy parameters ϵ and δ (if applicable). For example, we can achieve $(\epsilon, \delta(1 + \exp(\frac{\epsilon}{2})))$ -differential privacy by adding Laplace Noise scaled by $\frac{2S_f(D)}{\epsilon}$, but only if S_f is β -smooth for $\beta \leq \frac{\epsilon}{2\ln(2/\delta)}$. For pure differential privacy we require that $\beta < \frac{\epsilon}{2(\lambda+1)}$ where λ is a parameter of the noise distribution – smaller λ implies higher variance.

If we want to ensure that the output is accurate, we also need to ensure that the calibrated noise with $S_f(D)$ is small e.g., $o(f(D)) + \mathcal{O}(\Delta_f)$. Note that by definition since $S_f(D) = 4\alpha\mathcal{A}_f(D) + \Delta_f$, and we add noise proportional $S_f(D)$, we expect that the noise added may be $> \alpha f(D)$. Thus, in order to address this challenge, our basic strategy is to run the original (non-private) approximation algorithm \mathcal{A}_f with *tuned* error factors e.g., we decrease α by a multiplicative factor of $\frac{\epsilon}{\ln(n)}$, let $\rho := \frac{\epsilon\alpha}{\ln(n)}$. Since we are now running $\mathcal{A}_f(D, \rho, 0, 0)$, we have that the function $S_f(D) = 4\rho\mathcal{A}_f(D) + \Delta_f$ is a β -smooth upper bound on the local sensitivity of the algorithm $\mathcal{A}_f(\cdot, \rho, 0, 0)$. Assuming the global sensitivity Δ_f is small, we can now show that w.h.p. the noise sampled proportional to $S_f(D)$ is at most $\alpha f(D) + O(\Delta_f/\epsilon)$ thus resulting in an ϵ -differentially private algorithm with reasonable accuracy.

By tuning the parameter α we actually accomplish two useful properties (1. accuracy) we decrease both the local sensitivity and our smooth upper bound $S_f(D)$ which reduces the magnitude of the noise that we add, and (2. privacy) we achieve β -smoothness for increasingly small values of β so that the required condition $\beta \leq \frac{\epsilon}{2\ln(2/\delta)}$ (or $\beta < \frac{\epsilon}{2(\lambda+1)}$) can be satisfied if we want to scale noise according to $S_f(D)$.

Extending to deterministic \mathcal{A}_f with multiplicative and additive error. More generally, if we have an $(\alpha, \kappa, 0)$ -approximation algorithm \mathcal{A}_f then we can show that $S_f(D) = 4\alpha\mathcal{A}_f(D) + \Delta_f + 4\tau$ is a β -smooth upper bound on the local sensitivity of \mathcal{A}_f with $\beta = 6\alpha$ (see Lemma 14). In particular, note that the additive error term κ does not adversely impact smoothness. Thus, we can achieve pure differential privacy by tuning α such that $6\alpha < \beta < \frac{\epsilon}{2(\lambda+1)}$ and scaling our noise according to $S_f(D)$ (see Lemma 15). We can also obtain stronger accuracy guarantees by relaxing the requirement for pure DP and tuning α such that $6\alpha \leq \beta \leq \frac{\epsilon}{2\ln(2/\delta)}$ so that we can sample our noise from the Laplace distribution which has strong concentration guarantees.

When \mathcal{A}_f is randomized. The remaining challenge is to handle randomized approximation algorithms \mathcal{A}_f which are only guaranteed to output a good approximation with high probability i.e., with non-zero probability $\delta > 0$ the algorithm is allowed to output an arbitrarily bad approximation. In particular, let us consider an (α, κ, δ) -approximation algorithm \mathcal{A}_f . For any possible input D we are always guaranteed that with probability $\geq 1 - \delta$ the algorithm $\mathcal{A}_f(D)$ outputs a good approximation $(1 - \alpha)f(D) \leq \mathcal{A}_f(D) \leq (1 + \alpha)f(D)$. Unfortunately, the function $S_f(D) = 4\alpha\mathcal{A}_f(D) + \Delta_f + 4\kappa$ is no longer guaranteed to be a β -smooth upper bound on the local sensitivity of \mathcal{A}_f since \mathcal{A}_f may sometimes output a value outside the specified approximation bounds.

In order to address this challenge, we define a function $g_f(D)$ that matches $\mathcal{A}_f(D)$ with probability at least $1 - \delta/2$ and is *always* guaranteed to output a good approximation. We emphasize that $g_f(D)$ may not be efficiently computable, but it is well-defined and only used for the purpose of analysis. More specifically, we set $g_f(D) = \mathcal{A}_f(D)$ as long as $(1 - \alpha)f(D) - \kappa \leq \mathcal{A}_f(D) \leq (1 + \alpha)f(D) + \kappa$. If $\mathcal{A}_f(D) > (1 + \alpha)f(D)$, then we define $g_f(D) := (1 + \alpha)f(D)$, similarly, if $\mathcal{A}_f(D) < (1 - \alpha)f(D)$, then we define $g_f(D) := (1 - \alpha)f(D) + \kappa$. Observe that we are always guaranteed that $(1 - \alpha)f(D) - \kappa \leq g_f(D) \leq (1 + \alpha)f(D) + \kappa$. Thus, $S_f(D) = 4\alpha g_f(D) + \Delta_f + 4\tau$ is a $\beta = 6\alpha$ -smooth upper bound on the local sensitivity of g_f (see Lemma 8). As long as $6\alpha \leq \beta \leq \frac{\epsilon}{2 \ln(2/\delta)}$ we could preserve $(\epsilon, \delta(1 + \exp(\frac{\epsilon}{2})))$ -differential privacy by outputting $g_f(D)$ plus Laplace Noise scaled by $\frac{8\alpha g_f(D) + \Delta_f + 8\tau}{\epsilon}$ i.e., scaled according to our β -smooth upper bound on the local sensitivity of g_f . Unfortunately, the function g_f may not be efficiently computable. Thus, we substitute g_f for \mathcal{A}_f and instead output $\mathcal{A}_f(D)$ plus Laplace noise scaled according to $\frac{8\alpha \mathcal{A}_f(D) + \Delta_f + 8\tau}{\epsilon}$. While $4\alpha \mathcal{A}_f(D) + \Delta_f + 4\tau$ is not necessarily a β -smooth upper bound on the local sensitivity of \mathcal{A} , the key point is that the latter (efficiently computable) procedure is equivalent to the former (differentially private) procedure as long as $g_f(D) = \mathcal{A}_f(D)$ which happens as long as \mathcal{A}_f outputs a good approximation i.e., except with probability $\delta/2$. Thus, we can apply a hybrid argument to argue that the final efficiently computable algorithm is $(\epsilon, \frac{\delta}{2} + \delta(1 + \exp(\frac{\epsilon}{2})))$ -differential privacy (see Lemma 10). In order to ensure accuracy, we use the same strategy as before, i.e., we run $\mathcal{A}_f(D, \rho, \tau, \delta/2)$, where $\rho \leq \frac{\epsilon \alpha}{\log(1/\delta)}$. Sampling noise proportional to $S_f(D)$ (where $S_f(D)$ is now defined in terms of ρ), and absorbing the failure probability of algorithm \mathcal{A}_f into the DP failure probability term δ , results in an approximate differentially private algorithm. Finally applying the postprocessing step results in a pure differentially private algorithm. We refer to the full proofs (Section 2) for additional details.

Applications. We give some intuition on how we apply Theorem 3 to various applications by choosing appropriate parameters. Recall that with probability $1 - \delta - \exp(-\gamma)$, \mathcal{A}'_f outputs $(1 - \alpha')f(D) - \kappa' - \frac{2\Delta_f}{\epsilon} \cdot \gamma \leq \mathcal{A}'_f(D) \leq (1 + \alpha')f(D) + \kappa' + \frac{2\Delta_f}{\epsilon} \cdot \gamma$, where $\alpha' = \frac{\alpha(\epsilon + 16\gamma)}{12 \log(4/\delta)}$, and $\kappa' = \kappa \left(\frac{2\gamma\alpha}{3 \log(4/\delta)} + \frac{8\gamma}{\epsilon} + 1 \right)$ with a time/space/query complexity blow-up incurred by running the original algorithm \mathcal{A}_f with multiplicative accuracy parameter $\rho = \frac{\epsilon \alpha}{\log(4/\delta)}$. First, observe that if the original algorithm \mathcal{A}_f has time/space/query complexity with a dependence on $\text{poly}(\frac{1}{\alpha})$, then the resulting time/space/query complexities for \mathcal{A}'_f will still have a polynomial dependence, i.e., $\text{poly}(\frac{\log(4/\delta)}{\alpha})$ – this naturally leads to FPRAS or FPTAS applications, as well as other classes of approximation algorithms like sublinear time or space. On the otherhand, if the time/space/query complexity of \mathcal{A}_f has a non-polynomial dependence on $1/\alpha$, e.g., $\exp(\frac{1}{\alpha})$, then since δ is typically $\text{negl}(n)$ or $\frac{1}{n^c}$ for $c > 0$, the resulting DP algorithm \mathcal{A}'_f could have much worse time/space/query-guarantees with respect to n , e.g., in an extreme case if we set $\delta = 2^{-\text{poly}(n)}$ then $\rho = \Omega(\text{poly}(n)/\alpha)$ and we could incur a $\exp(\frac{\text{poly}(n)}{\alpha})$ multiplicative overhead in the running time. It is worth noting that one could optionally reduce the additive error term κ' for \mathcal{A}'_f by reducing the error term κ for \mathcal{A} .

We further emphasize this trade-off between obtaining small failure probability bounds and the accuracy or resource guarantees. Consider the following two examples – (Example 1) if we set the probability of failure, i.e., $\exp(-\gamma) = \delta = \frac{1}{n^c}$ for any $c > 0$, then the resulting approximation parameters are roughly $\alpha' = \alpha(1 + o(1))$, and $\kappa' = \kappa(\alpha + \frac{\log(n)}{\epsilon} + 1)^5$, and

⁵ In the applications we consider the original (non-private) approximation algorithm typically has only multiplicative or only additive error and not both. In particular, we typically either have $\alpha > 0$ and $\kappa = 0$ or $\kappa > 0$ and $\alpha = 0$, but not the case where $\alpha > 0$ and $\kappa > 0$. Considering the case when $\kappa \neq 0$, and $\alpha = 0$, we (roughly) have $\kappa' = \kappa(\frac{\log(n)}{\epsilon} + 1)$.

the additional error term depending on global sensitivity is roughly $\frac{\Delta_f \log(n)}{\varepsilon}$. We incur a time/space overhead by running \mathcal{A}_f with multiplicative accuracy parameter $\rho = \Omega(\varepsilon\alpha / \log n)$ instead of α . (Example 2) if we set the probability of failure, i.e., $\exp(-\gamma) = \delta = \frac{1}{n^{\log \log(n)}} - \text{negl}(n)$, then α' remains the same as before, and now $\kappa' = \kappa(\alpha + \frac{\log(n) \log \log(n)}{\varepsilon} + 1)$, but the additional error term depending on global sensitivity becomes $\frac{\Delta_f \log(n) \log \log(n)}{\varepsilon}$. In this latter case we incur time or space overhead by running \mathcal{A}_f with multiplicative accuracy parameter $\rho = \Omega\left(\frac{\varepsilon\alpha}{\log n \log \log n}\right)$ – to reduce the $\kappa' \log n \log \log n$ error term it could be useful to run \mathcal{A}_f with additive error parameter $\frac{\kappa}{\log n \log \log n}$ which may incur additional time or space overhead. Thus these two examples illustrate how, as we decrease the failure probability, the accuracy and resource (time/space in this case) guarantees become worse. See full version [11] for applications to sublinear time and streaming algorithms.

2 General Transformation for Approximation Algorithms

In this section, we formally define our black-box differentially private transformation for (randomized) approximation algorithms. Given a tunable approximation (see Definition 2) algorithm of f , call it \mathcal{A}_f , that outputs an (α, κ, δ) -approximation, our framework for randomized algorithms involves two steps – (1) Apply Algorithm 1 to \mathcal{A}_f to obtain an (ε, δ) -DP algorithm \mathcal{A}'_f with accuracy guarantees outlined in Theorem 3 (2) Apply postprocessing step to the output of \mathcal{A}'_f to obtain an ε -DP algorithm (see Theorem 5).

We first prove Theorem 3 that provides theoretical guarantees for algorithm \mathcal{A}'_f (Algorithm 1). This is our main contribution as the postprocessing step to obtain pure DP applies a folklore result.

Observe that even for the case when the original algorithm \mathcal{A}_f gives an $(\alpha, 0, \delta)$ -approximation of f (i.e., $\kappa = 0$), the resulting DP algorithm \mathcal{A}'_f will still have an additive error, this additive error is inherent due to the requirement of adding Laplace noise to preserve DP. We emphasize that the Laplace noise added to the output of algorithm \mathcal{A}_f depends on the global sensitivity of the function f , therefore, we can only get meaningful DP approximation algorithms using this transformation for functions with low global sensitivity.

Proof of Theorem 3. \mathcal{A}'_f is defined in Algorithm 1 – it first runs $\mathcal{A}_f(D, \rho, \kappa, \delta/2)$ where $\rho := \frac{\varepsilon\alpha}{12 \log(4/\delta)}$ and then adds Laplace Noise. Thus, the resource used by \mathcal{A}'_f is $R(n, \rho, \kappa, \delta)$. The privacy guarantee follows from Lemma 10, and the accuracy guarantee follows from Lemma 12. ◀

► **Remark 7.** When \mathcal{A}_f is a PRAS, by definition, the output of \mathcal{A}_f is an $(\alpha, 0, \delta)$ -approximation of f running in time $T(n, \alpha, 0, \delta) = \text{poly}(n, 1/\alpha, \log(1/\delta))$. Applying Theorem 3 with negligible $\delta = n^{-\log n}$ and $\gamma = \log^2 n$ for any $\alpha' > 0$ we obtain a private $(\alpha', O(\frac{\Delta_f}{\varepsilon \log^2 n}), 2n^{-\log n})$ -approximation with polynomial running time $\text{poly}(n, 1/\varepsilon, 1/\alpha')$.

► **Lemma 8.** Let $0 < \rho < 1/2$. Suppose that \mathcal{A}_f outputs a (ρ, τ, δ) -approximation of a function $f : \mathcal{D} \rightarrow \mathbb{R}^+$ with global sensitivity Δ_f . Let $\mathcal{A}_{f,R}$ denote a deterministic run of \mathcal{A} using a fixed set of random coins R . Define function $g_{f,R}$ by

$$g_{f,R}(D) = \begin{cases} \mathcal{A}_{f,R}(D) & \text{if } (1-\rho)f(D) - \tau \leq \mathcal{A}_{f,R}(D) \leq (1+\rho)f(D) + \tau \\ (1-\rho)f(D) - \tau & \text{if } \mathcal{A}_{f,R}(D) < (1-\rho)f(D) - \tau \\ (1+\rho)f(D) + \tau & \text{if } \mathcal{A}_{f,R}(D) > (1+\rho)f(D) + \tau \end{cases}$$

Then the function $S_f(D) = 4\rho g_{f,R}(D) + 4\tau + \Delta_f$ is a β -smooth upper bound for $g_{f,R}$ where $\beta \geq 6\rho$.

Proof. Fix an arbitrary set of random coin tosses R . We frequently use the fact that $(1 - \rho)f(D) - \tau \leq g_{f,R}(D) \leq (1 + \rho)f(D) + \tau$. We also note that since $0 \leq \rho < 1/2$, we have that $\frac{1}{2}f(D) - \tau \leq g_{f,R}(D) \leq 2f(D) + \tau$.

First, we show that Condition 1 of Definition 24 holds. Without loss of generality, we assume $f(D) \geq f(D')$, where D' is any neighboring input. Then:

$$\begin{aligned}
LS_{g_{f,R}}(D) &= \max_{D': D \sim D'} \|g_{f,R}(D) - g_{f,R}(D')\| \\
&\leq \|(1 + \rho)f(D) + \tau - (1 - \rho)f(D') + \tau\| \\
&\leq \rho\|f(D) + f(D')\| + 2\tau + \Delta_f \\
&\leq 2\rho f(D) + 2\tau + \Delta_f \\
&\leq 4\rho g_{f,R}(D) + 2\rho\tau + 2\tau + \Delta_f && \text{as } \frac{1}{2}f(D) - \tau \leq g_{f,R}(D) \\
&\leq 4\rho g_{f,R}(D) + 4\tau + \Delta_f \\
&= S_f(D)
\end{aligned}$$

Next, we show that Condition 2 of Definition 24 holds below. We have:

$$\begin{aligned}
S_f(D) &= 4\rho g_{f,R}(D) + 4\tau + \Delta_f \\
&\leq 4\rho(1 + \rho)f(D) + 4\rho\tau + 4\tau + \Delta_f && \text{by def of } g_{f,R} \\
&\leq 4\rho(1 + \rho)(\Delta_f + f(D')) + 4\rho\tau + 4\tau + \Delta_f && \text{since } D \sim D' \\
&\leq 4\rho(1 + \rho)f(D') + (4\rho(1 + \rho) + 1)\Delta_f + 4\rho\tau + 4\tau \\
&\leq 4\rho \frac{(1 + \rho)}{1 - \rho} (g_{f,R}(D') + \tau) + (1 + 6\rho)\Delta_f + 4\rho\tau + 4\tau && \text{by def of } g_{f,R} \\
&\leq 4\rho(1 + \rho)(1 + 2\rho)(g_{f,R}(D') + \tau) + (1 + 6\rho)\Delta_f + 4\rho\tau + 4\tau \\
&\leq 4\rho(1 + \rho)(1 + 2\rho)g_{f,R}(D') + 12\rho\tau + (1 + 6\rho)\Delta_f + 4\rho\tau + 4\tau \\
&\leq 4\rho(1 + \rho)(1 + 2\rho)g_{f,R}(D') + (1 + 6\rho)(\Delta_f + 4\tau) \\
&\leq 4\rho(1 + 4\rho)g_{f,R}(D') + (1 + 6\rho)(\Delta_f + 4\tau) \\
&\leq (1 + 6\rho)(4\rho g_{f,R}(D') + \Delta_f + 4\tau) \\
&\leq e^{6\rho} \cdot (4\rho g_{f,R}(D') + \Delta_f + 4\tau) = e^\beta S_f(D'),
\end{aligned}$$

where $\beta \geq 6\rho$. ◀

► **Remark 9.** As a special case if we have a $(0, \kappa, 0)$ -approximation algorithm \mathcal{A}_f (i.e., no multiplicative error, zero failure probability) then applying Lemma 8 yields the smooth upper bound $S_f(D) = 4\tau + \Delta_f$. We observe that this smooth upper bound is independent of D and, therefore, S_f is just an upper bound on the global sensitivity of g_f . Furthermore, in this special case we are guaranteed that $\mathcal{A}_f(D) = g_f(D)$ with probability 1. Thus, in this special case, we can achieve pure ε -DP by computing $\mathcal{A}_f(D)$ and adding Laplace noise proportional to S_f .

If we have a $(0, \kappa, \delta)$ -approximation algorithm for $\delta \neq 0$ then we still have $S_f(D) = 4\tau + \Delta_f$ which means that $S_f(D)$ is an upper bound on the global sensitivity of g_f . However, computing $\mathcal{A}_f(D)$ and adding Laplace noise proportional to S_f does not necessarily yield a pure DP algorithm since we may have $\Delta_f(D) \neq \mathcal{A}_f(D)$ with non-zero probability δ . If we have $(\alpha, \kappa, 0)$ -approximation algorithm \mathcal{A}_f , and $\alpha \neq 0$, since $S_f(D) = 4\rho\mathcal{A}_f(D) + 4\tau + \Delta_f$ still depends on the input D . However, we can still achieve DP using Theorem 6.

59:14 How to Make Your Approximation Algorithm Private

Applying Lemma 8 to the theorem calibrating noise to smooth bounds on the smooth sensitivity [41] we show that the Algorithm 1 preserves privacy below.

► **Lemma 10** (Privacy). *Algorithm 1 is (ε, δ') -differentially private where $\delta' = \delta(1 + \exp(\varepsilon/2))$.*

Proof. Consider a modification of Algorithm 1, call it Algorithm 1' where instead of computing $\mathcal{A}_f(D, \rho, \tau, \delta/2)$ we instead sample the random coins R that \mathcal{A}_f would have used and replace the value $\mathcal{A}_f(D, \rho, \tau, \delta/2; R)$ (which we denote as $\mathcal{A}_{f,R}(D)$ in the sequel) with $g_{f,R}(D)$. The function $g_{f,R}$ may not be efficiently computable, but we only use Algorithm 1' for the purpose of analysis. We first observe that by Lemma 8, for any $\beta \geq 6\rho$ the function $S_f(D) = 4\rho g_{f,R}(D) + 4\tau + \Delta_f$ is a β -smooth upper bound on the sensitivity of $g_{f,R}$. Thus, by Theorem 25, it is sufficient to set $\rho := \frac{\varepsilon\alpha}{12 \log(4/\delta)}$ for $6\rho \leq \beta \leq \frac{\varepsilon}{2 \ln(\frac{4}{\delta})}$ and add noise proportional to $\text{Lap}\left(\frac{2S_f(D)}{\varepsilon}\right) = \text{Lap}\left(\frac{2(4\rho g_{f,R}(D) + 4\tau + \Delta_f)}{\varepsilon}\right)$ to preserve $(\varepsilon, \delta/2)$ -privacy of Algorithm 1'.

Since $g_{f,R}(D)$ is $\mathcal{A}_{f,R}(D)$ except with probability $\delta/2$, Algorithm 1 is identical to Algorithm 1' except with probability $\delta/2$. Thus, this shows that Algorithm 1 is (ε, δ) -private. ◀

► **Fact 11.** *If $Y \sim \text{Lap}(b)$, then $\Pr[|Y| \geq \ell \cdot b] = \exp(-\ell)$.*

► **Lemma 12** (Accuracy). *For all $\gamma > 0$, with probability $1 - \exp(-\gamma) - \delta$,*

$$\begin{aligned} (1 - \rho(1 + \frac{16\gamma}{\varepsilon}))f(D) - \tau(\frac{8\gamma\rho}{\varepsilon} + \frac{8\gamma}{\varepsilon} + 1) - \frac{2\Delta_f\gamma}{\varepsilon} &\leq \mathcal{A}'(D) \leq (1 + \rho(1 + \frac{16\gamma}{\varepsilon}))f(D) \\ &+ \tau(\frac{8\gamma\rho}{\varepsilon} + \frac{8\gamma}{\varepsilon} + 1) + \frac{2\Delta_f\gamma}{\varepsilon}. \end{aligned}$$

Proof. First, using Fact 11, for any $\gamma > 0$, we have that,

$$\Pr\left[|X| \geq \frac{2(4\rho\mathcal{A}(D) + 4\tau + \Delta_f)}{\varepsilon} \cdot \gamma\right] = \exp(-\gamma)$$

\mathcal{A}_f is a $(\rho, \tau, \delta/2)$ -approximation of f so for any $D \in \mathcal{D}$, we have that $\mathcal{A}_f(D) \leq (1 + \rho)f(D) + \tau$ with probability $1 - \delta/2$. Since $0 < \rho < 1/2$ we have $(1 + \rho)f(D) + \tau \leq 2f(D) + \tau$. Therefore, by a union bound,

$$\begin{aligned} \Pr\left[(\mathcal{A}_f(D) > (1 + \rho)f(D) + \tau) \vee (\mathcal{A}_f(D) < (1 - \rho)f(D) - \tau) \vee (|X| \geq \frac{2(4\rho\mathcal{A}_f(D) + 4\tau + \Delta_f)}{\varepsilon} \cdot \gamma)\right] &\leq \delta/2 + \exp(-\gamma) \end{aligned}$$

Thus, with probability $1 - \exp(-\gamma) - \delta/2$, we have that

$$(1 - \rho)f(D) - \tau \leq \mathcal{A}_f(D) \leq (1 + \rho)f(D) + \tau \leq 2f(D) + \tau \tag{1}$$

and

$$|X| < \frac{2(4\rho\mathcal{A}_f(D) + 4\tau + \Delta_f)}{\varepsilon} \cdot \gamma \tag{2}$$

By plugging in Eq. 1 into Eq. 2, we have that with probability $1 - \exp(-\gamma) - \delta/2$,

$$|X| < \frac{2(4\rho(2f(D) + \tau) + 4\tau + \Delta_f)}{\varepsilon} \cdot \gamma$$

Overall, this means that with probability $1 - \exp(-\gamma) - \delta/2$,

$$\begin{aligned} (1 - \rho)f(D) - \tau - \frac{2(4\rho(2f(D) + \tau) + 4\tau + \Delta_f)}{\varepsilon} \cdot \gamma &\leq \mathcal{A}'_f(D) \\ &\leq (1 + \rho)f(D) + \tau + \frac{2(4\rho(2f(D) + \tau) + 4\tau + \Delta_f)}{\varepsilon} \cdot \gamma \end{aligned}$$

Grouping the like terms together gives the theorem statement. ◀

► **Theorem 5.** *Let $M = \max_D f(D)$ and let parameter $K > 0$. If $\mathcal{A}'_f(D)$ is (ε, δ) -DP algorithm with accuracy guarantee $(1 - \alpha)f(D) - \kappa \leq \mathcal{A}_f(D) \leq (1 + \alpha)f(D) + \kappa$ holding with probability $1 - \eta$ then there exists an algorithm $\mathcal{A}''_f(D)$ which is ε -DP with accuracy guarantee $(1 - \alpha)f(D) - \kappa - \frac{1}{KM} \leq \mathcal{A}_f(D) \leq (1 + \alpha)f(D) + \kappa + \frac{1}{KM}$ with probability at least $1 - \eta - p$ where $p = \frac{\delta K(M+1)}{e^\varepsilon - 1 + \delta K(M+1)}$.*

The proof of Theorem 5 is in Appendix D. At a high level our idea is to define $\mathcal{A}'_f(D) = \frac{\lceil K\mathcal{A}_f(D) \rceil}{KM}$. The rounding step introduces a small additive error term $\leq \frac{1}{KM}$ and ensures that $\mathcal{A}'_f(D)$ now has bounded range $|\mathcal{R}| \leq (M + 1)K$. Since \mathcal{A}'_f has bounded range we can apply a folklore result (see Theorem 19) to transform this (ε, δ) -DP algorithm into an ε -DP algorithm.

2.1 Achieving Pure DP for Approximation Algorithms with Zero Failure Probability

In this section we show how one can achieve pure differential privacy ($\delta = 0$) when we have a tunable $(\alpha, \kappa, 0)$ -approximation algorithm. The basic framework is the same except that we use the Cauchy distribution instead of Laplace when applying the Smooth Sensitivity framework – see Theorem 25. Since we assume $\delta = 0$ in this section we will sometimes simplify notation and write $T(n, \alpha, \kappa)$ (resp. $S(n, \alpha, \kappa)$) instead of $T(n, \alpha, \kappa, 0)$ (resp. $S(n, \alpha, \kappa, 0)$). We move the proofs in this section to Appendix C.

■ **Algorithm 2** ε -differentially private framework for tunable deterministic approximation algorithms.

Input: Input set D , accuracy parameter $\alpha \in (0, 1)$, differential privacy parameter ε , approx. algorithm \mathcal{A}_f .

- 1: Let $x_A := \mathcal{A}_f(D, \rho, \tau, 0)$ where $\rho := \frac{\varepsilon\alpha}{36}$ and $\tau := \kappa$.
- 2: **return** $x_A + X$ where $X \sim \mathcal{C}\left(\frac{6(4\rho x_A + \Delta_f)}{\varepsilon}\right)$

► **Remark 13.** When \mathcal{A}_f is a PTAS, by definition, the output of \mathcal{A}_f is an $(\alpha, 0, 0)$ -approximation of f running in time $T(n, \alpha, 0) = \text{poly}(n, 1/\alpha)$. Applying Theorem 6 with for any $\alpha > 0$ we obtain a private $(\alpha, O(\frac{\Delta_f}{\varepsilon}), 9/10)$ -approximation with polynomial running time $\text{poly}(n, 1/\varepsilon, 1/\alpha)$.

► **Lemma 14.** *Suppose that \mathcal{A}_f outputs a $(\rho, \tau, 0)$ -approximation where $0 < \rho < 1/2$ of a function $f : \mathcal{D} \rightarrow \mathbb{R}^+$ with global sensitivity Δ_f . Then the function $S_f(D) = 4\rho\mathcal{A}_f(D) + 4\tau + \Delta_f$ is a β -smooth upper bound for \mathcal{A}_f where $\beta \geq 6\rho$.*

The proof remains the same as in Lemma 8. Applying Lemma 14 to the theorem calibrating noise to smooth bounds on the smooth sensitivity [41] we show that the Algorithm 2 preserves privacy below.

► **Lemma 15.** *Algorithm 2 is ε -differentially private.*

► **Lemma 16.** *For all $\gamma > 6.5$, with probability at least $9/10$,*

$$\begin{aligned} & \left(1 - \rho \left(1 + \frac{48\gamma}{\varepsilon}\right)\right) f(D) - \frac{24(\rho + 1)\gamma\tau}{\varepsilon} - \frac{6\Delta_f}{\varepsilon} \cdot \gamma \leq \mathcal{A}'_f(D) \\ & \leq \left(1 + \rho \left(1 + \frac{48\gamma}{\varepsilon}\right)\right) f(D) + \frac{24(\rho + 1)\gamma\tau}{\varepsilon} + \frac{6\Delta_f}{\varepsilon} \cdot \gamma \end{aligned}$$

► Remark 17. For simplicity, we have chosen to sample from the standard cauchy distribution $\lambda = 2$, more generally, if we sample noise with density $h(z) \propto \frac{1}{1+|z|^\lambda}$, where $\lambda = c$, then with probability $1 - \delta$, $\gamma = \frac{1}{\delta^{1/c}}$ in Lemma 16.

Application to the Knapsack Problem. As a fun example we consider the knapsack problem. The knapsack problem is well known to be NP-Hard, but also admits an FPTAS. To define an instance of the knapsack problem we have a maximum weight capacity W for the knapsack and n items each with a value $v_{max} \geq v_i \geq 0$ and a weight $w_i \geq 0$. The goal is to find a subset $S \subseteq [n]$ of items to put in the knapsack maximizing the total value $v(S) = \sum_{i \in S} v_i$ subject to the constraint that the total weight $w(S) = \sum_{i \in S} w_i$ does not exceed our capacity i.e., $w(S) \leq W$.

For the purpose of this illustration let's fix the capacity W and weights w_1, \dots, w_n and let $f(v_1, \dots, v_n)$ denote the value of the optimal knapsack solution given values v_1, \dots, v_n . Let's say that two knapsack instances $(W, v_1, \dots, v_n, w_1, \dots, w_n)$ and $(W, v'_1, \dots, v'_n, w_1, \dots, w_n)$ are neighbors if $\sum_i \|v_i - v'_i\| \leq 1$. Thus, we are viewing the exact value of each item as sensitive and the goal of differential privacy is to prevent an attacker from inferring these sensitive values exactly. Observe that the global sensitivity of f is upper bounded by $\Delta_f \leq \max_{v \sim v'} \max_{S \subseteq [n]} |v(S) - v'(S)| \leq 1^6$.

Since there is an FPTAS algorithm for Knapsack we can find a non-private approximation algorithm $\mathcal{A}_f(\vec{v}, \alpha, \kappa = 0)$ running in time $T(n, \alpha) = \text{poly}(n, 1/\alpha)$. If we apply Theorem 6 then for any target α' our ϵ -DP algorithm \mathcal{A}'_f runs in time $\text{poly}(n, 1/\epsilon, 1/\alpha)$ and solves Knapsack with additive error $\mathcal{O}(1/\epsilon)$ and multiplicative error α' with probability at least $9/10$. If we don't require pure DP then we can also apply Theorem 3 then for any target α' our algorithm \mathcal{A}'_f runs in time $\text{poly}(n, 1/\epsilon, 1/\alpha, \log(1/\delta))$ and solves Knapsack with probability at least $1 - \delta - \exp(-\gamma)$ with additive error at most $\mathcal{O}(\gamma/\epsilon)$ and multiplicative error α' .

3 Conclusion and Open Questions

In this work, we introduce a general framework for transforming a non-private approximation algorithm into a differentially private approximation algorithm. We show specific applications of our framework for sublinear time and sublinear space algorithms. Although our framework applies to a large variety of problems and settings, it does incur a small penalty in both runtime and space for achieving differential privacy. A natural question is whether these losses are necessary for a general black-box framework and what are sufficient conditions for achieving a black-box reduction.

It also seems possible that our framework could provide a method for achieving differentially private algorithms when the important resource is not runtime, number of queries, or space. For example, in distributed algorithms, it is often desired to achieve sublinear communication while in learning/testing, it is often desired to achieve sublinear query complexity. We believe that exploring the limits and capabilities of our framework in those settings would be a natural future direction of work.

⁶ We could also define neighboring knapsack instances such that we can completely replace the value of any item i.e., v and v' are neighbors if there exists some index $i \in [n]$ such that $v_i \neq v'_i$ and $v_j = v'_j$ for all $j \neq i$. However, in this case we can we would have large global sensitivity $\Delta_f = v_{max}$. Thus, we won't be able to design an accurate differentially private approximation even if we are willing to solve the NP-Hard knapsack problem exactly.

References

- 1 Jayadev Acharya, Ziteng Sun, and Huanyu Zhang. Differentially private testing of identity and closeness of discrete distributions. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 6879–6891, 2018. URL: <https://proceedings.neurips.cc/paper/2018/hash/7de32147a4f1055bed9e4faf3485a84d-Abstract.html>.
- 2 Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Analyzing graph structure via linear measurements. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 459–467, 2012.
- 3 Noga Alon, Wenceslas Fernandez de la Vega, Ravi Kannan, and Marek Karpinski. Random sampling and approximation of max-csps. *J. Comput. Syst. Sci.*, 67(2):212–243, 2003. doi: 10.1016/S0022-0000(03)00008-4.
- 4 Noga Alon, Eldar Fischer, Ilan Newman, and Asaf Shapira. A combinatorial characterization of the testable graph properties: It’s all about regularity. *SIAM J. Comput.*, 39(1):143–167, 2009.
- 5 Noga Alon, Yossi Matias, and Mario Szegedy. The space complexity of approximating the frequency moments. *J. Comput. Syst. Sci.*, 58(1):137–147, 1999.
- 6 Gunnar Andersson and Lars Enggbretsen. Property testers for dense constraint satisfaction programs on finite domains. *Random Struct. Algorithms*, 21(1):14–32, 2002.
- 7 Petra Berenbrink, Bruce Krayenhoff, and Frederik Mallmann-Trenn. Estimating the number of connected components in sublinear time. *Inf. Process. Lett.*, 114(11):639–642, 2014. doi: 10.1016/j.ipl.2014.05.008.
- 8 Jaroslaw Blasiok. Optimal streaming and tracking distinct elements with high probability. *ACM Trans. Algorithms*, 16(1):3:1–3:28, 2020.
- 9 Jeremiah Blocki, Avrim Blum, Anupam Datta, and Or Sheffet. The johnson-lindenstrauss transform itself preserves differential privacy. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS*, pages 410–419, 2012.
- 10 Jeremiah Blocki, Elena Grigorescu, and Tamalika Mukherjee. Privately estimating graph parameters in sublinear time. In *49th International Colloquium on Automata, Languages, and Programming, ICALP*, pages 26:1–26:19, 2022.
- 11 Jeremiah Blocki, Elena Grigorescu, Tamalika Mukherjee, and Samson Zhou. How to make your approximation algorithm private: A black-box differentially-private transformation for tunable approximation algorithms of functions with low sensitivity, 2022. [arXiv:2210.03831](https://arxiv.org/abs/2210.03831).
- 12 Vladimir Braverman, Petros Drineas, Cameron Musco, Christopher Musco, Jalaj Upadhyay, David P. Woodruff, and Samson Zhou. Near optimal linear algebra in the online and sliding window models. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pages 517–528, 2020.
- 13 Vladimir Braverman, Joel Manning, Zhiwei Steven Wu, and Samson Zhou. Private data stream analysis for universal symmetric norm estimation. In *3rd Annual Symposium on Foundations of Responsible Computing FORC*, 2022.
- 14 Vladimir Braverman and Rafail Ostrovsky. Effective computations on sliding windows. *SIAM J. Comput.*, 39(6):2113–2131, 2010.
- 15 Vladimir Braverman, Rafail Ostrovsky, and Carlo Zaniolo. Optimal sampling from sliding windows. *J. Comput. Syst. Sci.*, 78(1):260–272, 2012.
- 16 Vladimir Braverman, Viska Wei, and Samson Zhou. Symmetric norm estimation and regression on sliding windows. In *Computing and Combinatorics - 27th International Conference, COCOON, Proceedings*, pages 528–539, 2021.
- 17 Zhiqi Bu, Sivakanth Gopi, Janardhan Kulkarni, Yin Tat Lee, Judy Hanwen Shen, and Uthaiapon Tantipongpipat. Fast and memory efficient differentially private-sgd via JL projections. *CoRR*, abs/2102.03013, 2021. [arXiv:2102.03013](https://arxiv.org/abs/2102.03013).

- 18 Mark Bun, Jonathan R. Ullman, and Salil P. Vadhan. Fingerprinting codes and the price of approximate differential privacy. *SIAM J. Comput.*, 47(5):1888–1938, 2018. doi:10.1137/15M1033587.
- 19 Bernard Chazelle, Ronitt Rubinfeld, and Luca Trevisan. Approximating the minimum spanning tree weight in sublinear time. *SIAM J. Comput.*, 34(6):1370–1379, 2005. doi:10.1137/S0097539702403244.
- 20 Cynthia Dwork. Differential privacy. In *Automata, Languages and Programming, 33rd International Colloquium, ICALP, Proceedings, Part II*, pages 1–12, 2006.
- 21 Cynthia Dwork and Jing Lei. Differential privacy and robust statistics. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC*, pages 371–380. ACM, 2009.
- 22 Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography, Third Theory of Cryptography Conference, TCC, Proceedings*, pages 265–284, 2006.
- 23 Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. *J. Priv. Confidentiality*, 7(3):17–51, 2016.
- 24 Talya Eden, Amit Levi, Dana Ron, and C. Seshadhri. Approximately counting triangles in sublinear time. *SIAM J. Comput.*, 46(5):1603–1646, 2017. doi:10.1137/15M1054389.
- 25 Alessandro Epasto, Mohammad Mahdian, Vahab S. Mirrokni, and Peilin Zhong. Improved sliding window algorithms for clustering and coverage via bucketing-based sketches. In *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 3005–3042, 2022.
- 26 Alessandro Epasto, Jieming Mao, Andres Muñoz Medina, Vahab Mirrokni, Sergei Vassilvitskii, and Peilin Zhong. Differentially private continual releases of streaming frequency moment estimations. In *14th Innovations in Theoretical Computer Science Conference, ITCS 2023, January 10-13, 2023, MIT, Cambridge, Massachusetts, USA*, volume 251 of *LIPICs*, pages 48:1–48:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.
- 27 Nimrod Fiat and Dana Ron. On efficient distance approximation for graph properties. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 1618–1637. SIAM, 2021.
- 28 Eldar Fischer and Ilan Newman. Testing versus estimation of graph properties. *SIAM J. Comput.*, 37(2):482–501, 2007.
- 29 Sumit Ganguly and David P. Woodruff. High probability frequency moment sketches. In *45th International Colloquium on Automata, Languages, and Programming, ICALP*, pages 58:1–58:15, 2018.
- 30 Badih Ghazi, Ravi Kumar, Pasin Manurangsi, and Thao Nguyen. Robust and private learning of halfspaces. In Arindam Banerjee and Kenji Fukumizu, editors, *The 24th International Conference on Artificial Intelligence and Statistics, AISTATS 2021, April 13-15, 2021, Virtual Event*, volume 130 of *Proceedings of Machine Learning Research*, pages 1603–1611. PMLR, 2021. URL: <http://proceedings.mlr.press/v130/ghazi21a.html>.
- 31 Arijit Ghosh, Gopinath Mishra, Rahul Raychaudhury, and Sayantan Sen. Tolerant bipartiteness testing in dense graphs. In *49th International Colloquium on Automata, Languages, and Programming, ICALP 2022, July 4-8, 2022, Paris, France*, volume 229 of *LIPICs*, pages 69:1–69:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
- 32 Oded Goldreich, Shafi Goldwasser, and Dana Ron. Property testing and its connection to learning and approximation. *J. ACM*, 45(4):653–750, 1998.
- 33 Oded Goldreich and Dana Ron. On estimating the average degree of a graph. *Electron. Colloquium Comput. Complex.*, 2004.
- 34 Maoguo Gong, Yu Xie, Ke Pan, Kaiyuan Feng, and Alex Kai Qin. A survey on differentially private machine learning [review article]. *IEEE Comput. Intell. Mag.*, 15(2):49–64, 2020. doi:10.1109/MCI.2020.2976185.

- 35 Anupam Gupta, Katrina Ligett, Frank McSherry, Aaron Roth, and Kunal Talwar. Differentially private combinatorial optimization. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 1106–1125, 2010.
- 36 Moritz Hardt and Kunal Talwar. On the geometry of differential privacy. In Leonard J. Schulman, editor, *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 705–714. ACM, 2010. doi:10.1145/1806689.1806786.
- 37 Rajesh Jayaram, David P. Woodruff, and Samson Zhou. Truly perfect samplers for data streams and sliding windows. In *PODS: International Conference on Management of Data*, pages 29–40, 2022.
- 38 Daniel M. Kane, Jelani Nelson, Ely Porat, and David P. Woodruff. Fast moment estimation in data streams in optimal space. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC*, pages 745–754, 2011.
- 39 Daniel M. Kane, Jelani Nelson, and David P. Woodruff. An optimal algorithm for the distinct elements problem. In *Proceedings of the Twenty-Ninth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS*, pages 41–52, 2010.
- 40 Darakhshan J. Mir, S. Muthukrishnan, Aleksandar Nikolov, and Rebecca N. Wright. Pan-private algorithms via statistics on sketches. In *Proceedings of the 30th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS*, pages 37–48, 2011.
- 41 Kobbi Nissim, Sofya Raskhodnikova, and Adam D. Smith. Smooth sensitivity and sampling in private data analysis. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pages 75–84, 2007.
- 42 Michal Parnas, Dana Ron, and Ronitt Rubinfeld. Tolerant property testing and distance approximation. *J. Comput. Syst. Sci.*, 72(6):1012–1042, 2006.
- 43 Adam D. Smith, Shuang Song, and Abhradeep Thakurta. The flajolet-martin sketch itself preserves differential privacy: Private counting with minimal space. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems, NeurIPS*, 2020.
- 44 Xiaoming Sun and David P. Woodruff. The communication and streaming complexity of computing the longest common and increasing subsequences. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 336–345, 2007.
- 45 Jakub Tetek. Additive noise mechanisms for making randomized approximation algorithms differentially private. *CoRR*, abs/2211.03695, 2022. doi:10.48550/arXiv.2211.03695.
- 46 Lun Wang, Iosif Pinelis, and Dawn Song. Differentially private fractional frequency moments estimation with polylogarithmic space. In *The Tenth International Conference on Learning Representations, ICLR*, 2022.
- 47 David P. Woodruff and Samson Zhou. Tight bounds for adversarially robust streams and sliding windows via difference estimators. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pages 1183–1196. IEEE, 2021.
- 48 Yuichi Yoshida, Masaki Yamamoto, and Hiro Ito. Improved constant-time approximation algorithms for maximum matchings and other optimization problems. *SIAM J. Comput.*, 41(4):1074–1093, 2012. doi:10.1137/110828691.

A Related Work

One of the first differentially private frameworks for computing general functions was introduced by [23] which released functions with additive noise, where the noise is calibrated according to the *global sensitivity* of the function f . This framework was generalized by [41], to handle functions which might have a high global sensitivity but are usually less sensitive in practice. The framework allows the release of functions with instance-specific noise, where the noise that is added is not just determined by f but by the input dataset as well. The noise magnitude calibrated is according to the *smooth sensitivity* of f on the input dataset

which is a smooth upper bound on the *local sensitivity* of f on an input dataset. The smooth sensitivity of a function may be hard to compute, therefore in the same work, [41] give a generic method called the *sample and aggregate* method that bypasses the explicit computation of the smooth sensitivity of the function and works even when the function is given as a black-box. [21] suggested a framework called *Propose-Test-Release* to release statistical estimators with additive noise where the noise is calibrated according to the *local sensitivity* of the estimator. Note that adding noise proportional to the local sensitivity of a function with respect to an input set usually does not preserve privacy, but their approach first proposes a bound on the local sensitivity and privately tests whether this bound holds for the specific input set, and then releases the noisy response to the query.

In the context of developing differentially private frameworks for approximation algorithms, [10] formally introduced the notion of *coupled global sensitivity* of a randomized algorithm, which gives an analogous framework as that of the global sensitivity framework [23], but for randomized approximation algorithms instead of deterministic functions. In this framework, one can run a non-private randomized approximation algorithm $\mathcal{A}_f(D)$ on the dataset, and privacy is obtained by adding noise proportional to the coupled global sensitivity of \mathcal{A}_f . More formally, the coupled global sensitivity measures the worst-case L_1 -sensitivity of the outputs of a randomized algorithm \mathcal{A}_f on neighboring inputs over a minimum coupling of the internal coin tosses of \mathcal{A}_f .

In independent work, Tetek [45] also explores the problem of transforming randomized approximation algorithms into (pure) differentially private approximation algorithms. In contrast to our results Tetek's transformation [45] assumes that the error of the original approximation algorithm either has small subexponential diameter or bounded mean error – assumptions that would not apply generically to every (tunable) approximation algorithm. Assuming subexponential error their work shows that it is possible to achieve ϵ -DP by adding Laplace Noise yielding accuracy guarantees that hold with high probability. However, the assumption of the error being subexponential is quite strong and does not often hold for many randomized approximation algorithms. While assuming bounded mean error is a weaker assumption on the error of the non-private randomized algorithm, however the DP noise is sampled from the Pareto distribution, which has polynomial tail bounds. This leads to accuracy guarantees which only hold with constant probability. Note that applying the median trick commonly used to amplify success probability in the non-private literature adversely affects the privacy budget and is thus not desirable. In contrast, our transformation applies generically to any (tunable) accurate approximation algorithm and we achieve accuracy guarantees that hold with high probability for the same problems studied in their paper. Finally, we correct an outdated claim⁷ from the comparison to our work detailed in [45] that says that we only achieve approximate privacy. We can achieve *pure* DP algorithms by applying a postprocessing step to the output of our transformation as outlined in Theorem 5.

B Preliminaries

We use the notation $\tilde{O}(f(n))$ to mean $f(n) \cdot \text{polylog}(f(n))$. We define datasets D and D' as *neighboring*, denoted as $D \sim D'$, if removing or adding one point in D results in D' ; alternatively, if changing one data point in D results in D' .

⁷ A prior version of the paper achieved pure DP, but that transformation (Theorem 1.5) only applied to deterministic tunable approximation algorithms

► **Definition 18** (Differential privacy). [22] An algorithm \mathcal{A} is (ϵ, δ) -DP if for every pair of neighboring datasets $D \sim D'$, and for all sets \mathcal{S} of possible outputs, we have that $\Pr[\mathcal{A}(D) \in \mathcal{S}] \leq e^\epsilon \Pr[\mathcal{A}(D') \in \mathcal{S}] + \delta$. When $\delta = 0$ we simply say that the algorithm is ϵ -DP.

Given an (ϵ, δ) -DP algorithm, one can obtain an ϵ -DP algorithm under certain conditions outlined below. We include the proof for completeness in Appendix D.

► **Theorem 19** (Approximate DP to Pure DP). Let $\mathcal{A} : \mathcal{D} \rightarrow \mathcal{R}$. If \mathcal{A} is an (ϵ, δ) -DP algorithm such that $\delta \leq \frac{(e^\epsilon - 1)p}{|\mathcal{R}|(1-p)}$ then there is an algorithm \mathcal{A}' such that \mathcal{A}' is ϵ -DP defined in the following manner.

$$\mathcal{A}'(D) = \begin{cases} \mathcal{A}(D) & \text{with probability } 1 - p \\ \text{random}(\mathcal{R}) & \text{with probability } p \end{cases}$$

where \mathcal{R} is the range of \mathcal{A}_f .

We define the distributions we will use to sample additive noise from below.

► **Definition 20** (Laplace distribution). We say a random variable X is drawn from a Laplace distribution with mean μ and scale $b > 0$ if the probability density function of X at x is $\frac{1}{2b} \exp\left(-\frac{|x-\mu|}{b}\right)$. We use the notation $X \sim \text{Lap}(b)$ to denote that X is drawn from the Laplace distribution with scale b and mean $\mu = 0$.

► **Definition 21** (Cauchy distribution). We say a random variable X is drawn from a Cauchy distribution with location parameter x_0 and scale $b > 0$ if the probability density function of X at x is $\frac{1}{\pi b} \left(\frac{b^2}{(x-x_0)^2 + b^2}\right)$. We use the notation $X \sim \text{C}(b)$ to denote that X is drawn from the Cauchy distribution with scale b and location parameter $x_0 = 0$.

We formally define the concept of global sensitivity which is a worst-case notion of sensitivity for deterministic functions below.

► **Definition 22** (Global sensitivity). The global sensitivity of a function $f : \mathcal{D} \rightarrow \mathbb{R}^d$ is defined by

$$\Delta_f = \max_{D, D' \in \mathcal{D}, D \sim D'} \|f(D) - f(D')\|_1.$$

We define the notion of local sensitivity for a fixed input, which can be much smaller than the global sensitivity, but in general, adding noise calibrated according to the local sensitivity does not preserve DP.

► **Definition 23** (Local sensitivity). For $f : \mathcal{D} \rightarrow \mathbb{R}$ and $D \in \mathcal{D}$, the local sensitivity of f at D is defined as

$$LS_f(D) = \max_{D' : D \sim D'} \|f(D) - f(D')\|_1.$$

Note: if $f : \mathcal{D} \times \mathcal{R} \rightarrow \mathbb{R}$ is a randomized function which, in addition to a dataset $D \in \mathcal{D}$ takes random coins $r \in \mathcal{R}$ as input we simply define $LS_f(D) = \max_{r \in \mathcal{R}} LS_{f_r}$ where $f_r(D) \doteq f(D; r)$.

In order to add instance-specific noise, we define the notions of β -smooth upper bound which is a smooth upper bound on the local sensitivity.

59:22 How to Make Your Approximation Algorithm Private

► **Definition 24** (Smooth upper bound on local sensitivity). For $\beta > 0$, a function $S : \mathcal{D} \rightarrow \mathbb{R}$ is a β -smooth upper bound on the local sensitivity of $f : \mathcal{D} \rightarrow \mathbb{R}$ if

(1) For all $D \in \mathcal{D}$, we have $S(D) \geq LS_f(D)$.

(2) For all $D, D' \in \mathcal{D}$ with $\|D - D'\|_1 = 1$, we have $S(D) \leq e^\beta \cdot S(D')$.

Finally, although one cannot add noise calibrated with local sensitivity, one can add noise proportional to a β -smooth upper bound on the local sensitivity as follows.

► **Theorem 25** (Corollary 2.4 in [41]). Let $f : \mathcal{D} \rightarrow \mathbb{R}$ and $S : \mathcal{D} \rightarrow \mathbb{R}$ be a β -smooth upper bound on the local sensitivity of f .

(1) If $\beta \leq \frac{\varepsilon}{2(\lambda+1)}$ and $\lambda > 1$, the algorithm $D \rightarrow f(D) + \frac{2(\lambda+1)S(D)}{\varepsilon} \cdot \eta$, where η is sampled from the distribution with density $h(z) \propto \frac{1}{1+|z|^\lambda}$, is ε -differentially private.

(2) If $\beta \leq \frac{\varepsilon}{2\ln(2/\delta)}$ and $\delta \in (0, 1)$, then the algorithm $D \rightarrow f(D) + \frac{2S(D)}{\varepsilon} \cdot \eta$ where $\eta \sim \text{Lap}(1)$ is (ε, δ') -differentially private for $\delta' = \delta (1 + \exp(\frac{\varepsilon}{2}))^8$.

C Proofs of Section 2.1

Proof of Theorem 6.

Proof. \mathcal{A}'_f is defined in Algorithm 2 – we first run $\mathcal{A}_f(D, \rho, \kappa)$ where $\rho := \frac{\varepsilon\alpha}{36}$ and then we add noise proportional to the standard Cauchy distribution. Thus, the resource used will be $R(n, \rho, \kappa)$.

The privacy guarantee follows from Lemma 15, and the accuracy guarantee follows from Lemma 16. ◀

Proof of Lemma 15

Proof. We first observe that by Lemma 14, $S_f(D) = 4\mathcal{A}_f(D) + 4\tau + \Delta_f$ is a β -smooth upper bound for \mathcal{A}_f . Recall that $\rho := \frac{\varepsilon\alpha}{36}$, thus we can apply Theorem 25 (with $\lambda = 2$) where $6\rho \leq \beta \leq \frac{\varepsilon}{6}$ and conclude that it is sufficient to add noise proportional to $C\left(\frac{2(2+1)S_f(x)}{\varepsilon}\right) = C\left(\frac{6(4\rho\mathcal{A}_f(D) + 4\tau + \Delta_f)}{\varepsilon}\right)$ to preserve ε -privacy. ◀

► **Fact 26.** If $Y \sim C(x; 0, b)$, then $\Pr[|Y| \geq \ell b] = 1 - \frac{2 \tan^{-1}(\ell)}{\pi}$.

Proof of Lemma 16.

Proof. First, we invoke Fact 26 below,

$$\Pr\left[|X| \geq \frac{6(4\rho\mathcal{A}_f(D) + 4\tau + \Delta_f)}{\varepsilon} \cdot \gamma\right] = 1 - \frac{2 \tan^{-1}(\gamma)}{\pi} \leq \frac{1}{10}$$

where the final inequality comes from using the fact that $\gamma > 6.5$. In other words, with probability $\geq 9/10$,

$$|X| \leq \frac{6(4\rho\mathcal{A}_f(D) + 4\tau + \Delta_f)}{\varepsilon} \cdot \gamma \tag{3}$$

\mathcal{A}_f is a $(\rho, \tau, 0)$ -approximation of f so for any $D \in \mathcal{D}$, we have that $\mathcal{A}_f(D) \leq (1 + \rho)f(D) + \tau$. Since $0 < \rho < 1/2$ we have $(1 + \rho)f(D) + \tau \leq 2f(D) + \tau$.

⁸ These bounds differ slightly from those listed in the original paper (Corollary 2.4 in [41]). We confirmed with the authors in private communication that δ should be multiplied by $(1 + \exp(\varepsilon/2))$.

By plugging in the relation $\mathcal{A}_f(D) \leq 2f(D) + \tau$ into Eq. 3, we have that with probability at least $9/10$,

$$|X| \leq \frac{6(8\rho f(D) + 4\rho\tau + 4\tau + \Delta_f)}{\varepsilon} \cdot \gamma$$

Thus with probability at least $9/10$,

$$\begin{aligned} (1 - \rho)f(D) - \tau - \frac{6(8\rho f(D) + 4\rho\tau + 4\tau + \Delta_f)}{\varepsilon} \cdot \gamma &\leq \mathcal{A}'_f(D) \\ &\leq (1 + \rho)f(D) + \tau + \frac{6(8\rho f(D) + 4\rho\tau + 4\tau + \Delta_f)}{\varepsilon} \cdot \gamma \end{aligned}$$

Rearranging the like terms together in the above expression completes the proof. \blacktriangleleft

D Proof of Approximate DP to Pure DP transformation

► **Theorem 5.** *Let $M = \max_D f(D)$ and let parameter $K > 0$. If $\mathcal{A}'_f(D)$ is (ε, δ) -DP algorithm with accuracy guarantee $(1 - \alpha)f(D) - \kappa \leq \mathcal{A}_f(D) \leq (1 + \alpha)f(D) + \kappa$ holding with probability $1 - \eta$ then there exists an algorithm $\mathcal{A}''_f(D)$ which is ε -DP with accuracy guarantee $(1 - \alpha)f(D) - \kappa - \frac{1}{KM} \leq \mathcal{A}_f(D) \leq (1 + \alpha)f(D) + \kappa + \frac{1}{KM}$ with probability at least $1 - \eta - p$ where $p = \frac{\delta K(M+1)}{e^\varepsilon - 1 + \delta K(M+1)}$.*

Proof. Note that WLOG we can assume that $\mathcal{A}_f(D)$ outputs a value between 0 and M since we can always truncate the output to this range – this operation preserves privacy by postprocessing and does not adversely affect accuracy. For some $K > 0$, define algorithm $\mathcal{A}''_f(D)$ as outputting $\frac{\lceil K\mathcal{A}_f(D) \rceil}{KM}$. Observe that \mathcal{A}''_f is (ε, δ) -DP by postprocessing and the accuracy guarantee of \mathcal{A}''_f is almost identical to that of \mathcal{A}_f since by definition $|\mathcal{A}''_f(D) - \mathcal{A}_f(D)| < \frac{1}{KM}$. By post-processing we can ensure that the range \mathcal{R} of $\mathcal{A}''_f(D)$ is small $|\mathcal{R}| = (M+1)K$ since $\mathcal{R} = \{\frac{i}{KM} : 0 \leq i \leq KM\}$. Thus, we can pick p such that $\delta \leq \frac{(e^\varepsilon - 1)p}{|\mathcal{R}|(1-p)}$ and apply a folklore theorem (see Theorem 19) to transform our (ε, δ) -DP algorithm $\mathcal{A}''_f(D)$ to an ε -DP algorithm $\mathcal{A}'_f(D)$ in the following manner:

$$\mathcal{A}'_f(D) = \begin{cases} \mathcal{A}''_f(D) & \text{with probability } 1 - p \\ \text{random}(\mathcal{R}) & \text{with probability } p \end{cases}$$

By combining the accuracy guarantees of \mathcal{A}_f and \mathcal{A}''_f we see that with probability $1 - \eta - p$, we have that $(1 - \alpha)f(D) - \kappa - \frac{1}{KM} \leq \mathcal{A}_f(D) \leq (1 + \alpha)f(D) + \kappa + \frac{1}{KM}$ where $p = \frac{\delta K(M+1)}{e^\varepsilon - 1 + \delta K(M+1)}$ as claimed. \blacktriangleleft

► **Theorem 19** (Approximate DP to Pure DP). *Let $\mathcal{A} : \mathcal{D} \rightarrow \mathcal{R}$. If \mathcal{A} is an (ε, δ) -DP algorithm such that $\delta \leq \frac{(e^\varepsilon - 1)p}{|\mathcal{R}|(1-p)}$ then there is an algorithm \mathcal{A}' such that \mathcal{A}' is ε -DP defined in the following manner.*

$$\mathcal{A}'(D) = \begin{cases} \mathcal{A}(D) & \text{with probability } 1 - p \\ \text{random}(\mathcal{R}) & \text{with probability } p \end{cases}$$

where \mathcal{R} is the range of \mathcal{A}_f .

Proof. Recall that we define \mathcal{A}' as follows:

$$\mathcal{A}'(D) = \begin{cases} \mathcal{A}(D) & \text{with probability } 1 - p \\ \text{random}(\mathcal{R}) & \text{with probability } p \end{cases}$$

59:24 How to Make Your Approximation Algorithm Private

Let $D, D' \in \mathcal{D}$ be neighboring databases and fix output $y \in \mathcal{R}$. We first give a general claim regarding the probability of $\mathcal{A}'(D) = y$ in terms of the $\Pr[\mathcal{A}(D) = y]$.

▷ **Claim 27.** For $D \in \mathcal{D}$,

$$\Pr[\mathcal{A}'(D) = y] = \Pr[\mathcal{A}(D) = y] (1 - p) + \frac{p}{|\mathcal{R}|}$$

Now we need to show that $\Pr[\mathcal{A}'(D) = y] \leq e^\epsilon \Pr[\mathcal{A}'(D') = y]$.

$$\begin{aligned} & \Pr[\mathcal{A}'(D) = y] \\ &= \Pr[\mathcal{A}(D) = y] (1 - p) + \frac{p}{|\mathcal{R}|} \\ &\leq (1 - p) (e^\epsilon \Pr[\mathcal{A}(D') = y] + \delta) + \frac{p}{|\mathcal{R}|} \\ &\leq e^\epsilon \Pr[\mathcal{A}(D') = y] (1 - p) + \delta (1 - p) + \frac{p}{|\mathcal{R}|} \end{aligned} \tag{4}$$

$$= e^\epsilon (\Pr[\mathcal{A}'(D') = y] - \frac{p}{|\mathcal{R}|}) + \delta (1 - p) + \frac{p}{|\mathcal{R}|} \tag{5}$$

$$\begin{aligned} &\leq e^\epsilon \Pr[\mathcal{A}'(D') = y] + \delta (1 - p) + \frac{p}{|\mathcal{R}|} (1 - e^\epsilon) \\ &\leq e^\epsilon \Pr[\mathcal{A}'(D') = y] \end{aligned} \tag{6}$$

The transition 4 to 5 follows from the observation that $\Pr[\mathcal{A}'(D') = y] = (1 - p) \Pr[\mathcal{A}(D') = y] + \frac{p}{|\mathcal{R}|}$ and therefore, $(1 - p) \Pr[\mathcal{A}(D') = y] = \Pr[\mathcal{A}'(D') = y] - \frac{p}{|\mathcal{R}|}$. The last equation 6 follows because $\delta \leq \frac{(e^\epsilon - 1)p}{|\mathcal{R}|(1 - p)}$ and thus

$$\delta (1 - p) + \frac{p}{|\mathcal{R}|} (1 - e^\epsilon) \leq 0 . \quad \blacktriangleleft$$