

Differential Privacy and Sublinear Time Are Incompatible Sometimes

Jeremiah Blocki  

Purdue University, West Lafayette, IN, USA

Hendrik Fichtenberger  

Google Research, Zürich, Switzerland

Elena Grigorescu  

University of Waterloo, ON, Canada

Tamalika Mukherjee  

Columbia University, New York, NY, USA

Abstract

Differential privacy and sublinear algorithms are both rapidly emerging algorithmic themes in times of big data analysis. Although recent works have shown the existence of differentially private sublinear algorithms for many problems including graph parameter estimation and clustering, little is known regarding hardness results on these algorithms. In this paper, we initiate the study of lower bounds for problems that aim for both differentially-private and sublinear-time algorithms. Our main result is the incompatibility of both the desiderata in the general case. In particular, we prove that a simple problem based on one-way marginals yields both a differentially-private algorithm, as well as a sublinear-time algorithm, but does not admit a “strictly” sublinear-time algorithm that is also differentially private.

2012 ACM Subject Classification Theory of computation → Theory of database privacy and security; Theory of computation → Streaming, sublinear and near linear time algorithms; Security and privacy → Privacy-preserving protocols

Keywords and phrases differential privacy, sublinear algorithms, sublinear-time algorithms, one-way marginals, lower bounds

Digital Object Identifier 10.4230/LIPIcs.ITCS.2025.19

Related Version *Full Version*: <https://arxiv.org/abs/2407.07262>

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

Elena Grigorescu: Funded in part by NSF CCF-1910659, NSF CCF-2205811, and NSF CCF-2228814.

Tamalika Mukherjee: Funded in part by NSF CCF-1910659, NSF CCF-2205811, and NSF CCF-2228814. Supported in part by Amazon through the Columbia Center of Artificial Intelligence Technology, a Google Cyber NYC Award, and DARPA under contract number W911NF2110371. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the United States Government or DARPA.

1 Introduction

While individuals have long demanded privacy-preserving analysis and processing of their data, their adoption and enforcement by governmental and private standards, policies, and jurisdictions are now accelerating. This urgency stems, in part, from the dramatic growth in the amount of data collected, aggregated, and analyzed per individual in recent years. The sheer volume of data also poses a computational challenge as resource demands scale with data size. Thus, it is expedient to develop privacy preserving algorithms for data-analysis whose resource requirements scale sub-linearly in the size of the input dataset. Two algorithmic



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

16th Innovations in Theoretical Computer Science Conference (ITCS 2025).

Editor: Raghu Meka; Article No. 19; pp. 19:1–19:18



Leibniz International Proceedings in Informatics

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

concepts that formalize these two objectives are *differential privacy* (DP) and *sublinear algorithms*. A randomized algorithm is differentially private if its output distribution does not change significantly when we slightly modify the input dataset to add/remove an individual's data i.e., a row in the dataset. Sublinear algorithms comprise classes of algorithms that have time or space complexity that is sublinear in their input size. Previous work in the intersection of both fields has promoted several classical sublinear algorithms to differentially private sublinear time algorithms for the same problems, e.g., sublinear-time clustering [3], graph parameter estimation [4] or sublinear-space heavy hitters in streaming [27, 6].

Intuitively, one might expect that privacy and sublinear time necessarily have a symbiotic relationship, i.e., if only a fraction of the data is processed, a significant amount of sensitive information may remain unread. Recent work [5] demonstrated that if a function $f: D \rightarrow \mathbb{R}$ has low global sensitivity (i.e., f is amenable to DP) and there exists a *sufficiently accurate* sublinear-time approximation algorithm for f , then there exists an accurate sublinear time DP approximation for f . This lead [5] to ask whether or not a similar transformation might apply for functions $f: D \rightarrow \mathbb{R}^d$ with multi-dimensional output. In this paper, we provide an example of a function $f: D \rightarrow \mathbb{R}^d$ with the following properties (1) there is an efficient sublinear time approximation algorithm for f , (2) there is a differentially private approximation algorithm for f running in time $O(|D|)$, and (3) any accurate differentially private approximation algorithm must run in time $\tilde{\Omega}(|D|)$. Thus, the intuition that privacy and sublinear time algorithms necessarily have a symbiotic relationship is incorrect.

Existing Model for DP Lower Bounds. Consider a database $D \in \{0, 1\}^{n \times d}$ with n rows where each row corresponds to an individual's record with d binary attributes. The model in this case is that the database D consists of a sample (according to a uniform or possibly adversarial distribution) from a larger population (or universe), and we are interested in answering queries on the sample (e.g., what fraction of individual records in D satisfy some property q ?). One of the interesting questions in DP lower bounds is the following: Suppose we fix a family of queries \mathcal{Q} and the dimensionality of database records d , then what is the sample complexity required to achieve DP and statistical accuracy for \mathcal{Q} ? Here the *sample complexity* is defined as the minimum number of records n such that there exists a (possibly computationally unbounded) algorithm that achieves both DP and accuracy.

A key problem that has been at the center of addressing this question is the *one-way marginal problem*. The one-way marginal problem takes a database $D \in \{0, 1\}^{n \times d}$ and releases the average values of all d columns. The best known private algorithm which has running time polynomial in the universe size is based on the multiplicative weights mechanism, and it achieves $(O(1), o(1/n))$ -DP for $n \in O(\sqrt{d} \log |\mathcal{Q}|)$ [17]. Any pure DP algorithm for this problem requires $n \geq \Omega(d)$ samples [18], while [8] showed that $n \geq \tilde{\Omega}(\sqrt{d})$ samples are necessary to solve this problem with approximate DP and within an additive error of $1/3$.

Existing Models for Sublinear-Time Algorithms. The works on sublinear-time algorithms utilize different input models, many of them tailored to the representation of the input, e.g., whether it is a function or a graph. These models typically define query oracles, i.e., mechanisms to access the input in a structured way. For example, the dense graph model [16] defines an adjacency query oracle that, given a pair of indices (i, j) , returns the entry $A(i, j)$ of the adjacency matrix A of the input graph. Query oracles enable an analysis of what parts, and more generally how much of the input was accessed by an algorithm. Since the fraction of input read is a lower bound for the time complexity of an algorithm, query access models are crucial to prove both sublinear-time upper and lower bounds.

Our Model. A challenge of proving lower bounds for sublinear time, differentially private algorithms lies in devising and applying a technique for analysis that combines the properties of both models. Lower bounds for differential privacy state a lower bound on the sample complexity that is required to guarantee privacy and non-trivial accuracy. These bounds do not state how much of the input an algorithm needs to *read* to guarantee privacy and accuracy, but only what input size *is required* to (potentially) enable such an algorithm. On the other hand, lower bounds for sublinear time algorithms state a bound on the time complexity as a function of the input size m . Note that time complexity is at least query complexity, and a lower bound on the latter immediately implies a lower bound on time complexity as well.

In our setting, we fix the number of records n , as well as the dimensionality d of the database, i.e., our problem size is $m = n \cdot d$. We define queries in our model to be *attribute queries*, i.e., querying the j -th attribute of a row i in the database D is denoted as $D(i, j)$. We emphasize that our use of the term *query* in our model is as in the sublinear-time algorithms model, and it is different from its use in conventional DP literature. Specifically, queries in DP literature refer to types of questions that the data analyst can make to the database to infer something about the population, whereas queries in the sublinear algorithms model refer to how the algorithm can access the input dataset D . In our work, we fix a problem of interest \mathcal{P} on database D (e.g., the one-way marginal problem), and we consider an algorithm that solves the problem \mathcal{P} on input D . Then we are interested in understanding the minimum number of (attribute) queries that an algorithm can make to solve the problem \mathcal{P} and satisfy both DP and accuracy, which we call the *query complexity*.

Result of [8] does not apply in our model. For the problem of one-way marginals, we know that $n \in \tilde{\Omega}(\sqrt{d})$ [8] records are required for any algorithm to achieve both DP and accuracy. For $m = \tilde{\Omega}(d^{3/2})$, there exists a DP algorithm that can solve this problem with $\tilde{O}(m)$ queries, i.e., the algorithm can query the entire dataset and add Gaussian noise. Using Hoeffding bounds one can analyze a simple non-private algorithm with accuracy $1/3$ that has query complexity $O(d \log d)$, which is sublinear in the problem size m . However, it is not clear whether $O(m)$ queries are necessary to achieve both DP and accuracy in our model. One might be tempted to directly apply the result of [8] to say that $\tilde{\Omega}(m)$ queries are necessary, but this does not work as the results of [8] focus on sample complexity. In particular, in our model it would be possible to distribute attribute queries across all rows (making $o(d)$ attribute queries in each row) so that every row is (partially) examined but the total number of queries is still $o(m)$. In particular, a sublinear time algorithm can substantially reduce ℓ_1 and ℓ_2 sensitivity by ensuring that the maximum number of queries in each row is $o(d)$.¹

The sample complexity lower bound of [8] uses fingerprinting codes to show that the output of an algorithm that is both DP and reasonably accurate for the one-way marginal problem can be used to reidentify an individual record, which contradicts the DP property. Intuitively, fingerprinting codes provide the guarantee that if an algorithm obtains an accurate answer after examining *at most* c rows in the database then it is possible to reidentify at

¹ Suppose for example, that $m = d^2$ so that there are d attribute columns and d rows and consider two sublinear time algorithms: Algorithm 1 examines the first \sqrt{d} rows and outputs the marginals for these samples. By contrast, Algorithm 2 uses rows $i\sqrt{d} + 1$ to $(i + 1)\sqrt{d}$ to compute the marginals columns $i\sqrt{d} + 1$ to $(i + 1)\sqrt{d}$ for each $i < \sqrt{d}$. Both algorithms examine the same number of cells in the database $d\sqrt{d}$, but the ℓ_1 (resp. ℓ_2) sensitivity of the algorithms are quite different. Algorithm 1 has ℓ_1 (resp. ℓ_2) sensitivity \sqrt{d} (resp. 1) while Algorithm 2 has ℓ_1 (resp. ℓ_2) sensitivity 1 (resp. $d^{-0.25}$).

least one of the corresponding users. However, if the attacker examines more than c rows, then we cannot prove that privacy is violated as fingerprinting codes no longer provide the guarantee that we can reidentify one of the corresponding users. In our model, an algorithm is allowed to make arbitrary attribute queries and is not restricted to querying all attributes corresponding to a fixed row, thus it is more difficult to prove a lower bound of this nature in our model. In particular, instead of sampling c rows an attacker could distribute the attribute queries across all rows (making $c = o(d)$ attribute queries in each row). The total number of cells examined is still cd , but the overall coalition has size $d \geq c$. Fingerprinting codes provide no guarantee of being able to trace a colluder since the overall coalition (number of rows in which some query was made) is larger than c . Thus, we cannot prove that privacy is violated.

Crucially, their construction relies on the algorithm being able to query the entire row (aka record) of the database and the fact that for a fixed coalition size c a fingerprinting code can trace an individual in any coalition of size $\leq c$ with high probability, as long as the individual *actively* colluded.

Our Contribution. We give the first separation between the query complexity of a non-private, sublinear-time algorithm and a DP sublinear-time algorithm (up to a log factor). We remind the reader that a lower bound on query complexity naturally implies a lower bound on time complexity as the time taken by an algorithm must be at least the number of queries made. Thus our theorem on query complexity also gives a lower bound for sublinear-time DP algorithms.

► **Theorem 1 (Informal Theorem).** *There exists a problem \mathcal{P} of size m such that*

1. \mathcal{P} can be solved privately with $O(m)$ query and time complexity.
2. \mathcal{P} can be solved non-privately with $O(m^{2/3} \log(m)) \in o(m)$ query and time complexity.
3. Any algorithm that solves \mathcal{P} with $(1/3, 1/3)$ -accuracy and $(O(1), o(1/m))$ -DP must have $\Omega(m/\log(m)) = \tilde{\Omega}(m)$ query and time complexity.

We note that [8] implies that any accurate, DP algorithm for \mathcal{P} requires $n \in \Omega(\sqrt{d})$, and we can in fact invoke the theorem for the *hardest* case and choose m so that $n = \Theta(\sqrt{d})$.² For full details on the definition of the problem and the formal version of the theorem, see Definition 21 and Theorem 26. Summarized in words, \mathcal{P} is solvable under differential privacy, and there exists a non-private sublinear-time algorithm to solve \mathcal{P} , but any DP algorithm must read (almost) the entire dataset, and thus have at least (nearly) linear running time. Our techniques build upon a rich literature of using fingerprinting codes in DP lower bounds. We note that the $\log(m)$ factor in our main result (Item 3) of Theorem 1 arises from the nearly-optimal Tardos fingerprinting code used in our lower bound construction. Thus it seems unlikely that this result can be improved unless one bypasses using fingerprinting codes entirely in the DP lower bound construction.

1.1 Technical Overview

Fingerprinting code (FPC) and DP. We start the construction of our lower bound with the privacy lower bounds based on fingerprinting codes [8]. For a set of n users and a parameter $c \leq n$, an (n, d, c) -FPC consists of two algorithms $(Gen, Trace)$. The algorithm Gen on input n outputs a codebook $C \in \{0, 1\}^{n \times d}$ where each row is a codeword of user $i \in [n]$ with code

² We call $n = \Theta(\sqrt{d})$ the hardest case because, when n becomes larger as a function of d , [1] show that subsampling can improve the privacy/accuracy trade-off of existing DP algorithms.

length $d = d(n, c)$. It guarantees that if at most c users collude to combine their codewords into a new codeword c' and the new code satisfies some (mild) *marking condition* – namely that if every colluder has the same bit b in the j -th bit of their codeword, then the j -th bit of c' must also be b – then the *Trace* algorithm of the fingerprinting code can identify at least one colluder with high probability.³ Bun et al. [8] used fingerprinting codes to show that the output of any accurate algorithm for the one-way marginal problem must satisfy the marking condition for sufficiently large d , and therefore, at least one row (i.e., individual) from the database is identifiable – making this algorithm not private. In more detail, given a fingerprinting code $(Gen, Trace)$, suppose a coalition of n users builds dataset $D \in \{0, 1\}^{n \times d}$ where each row corresponds to a codeword of length d from the codebook Gen . For $j \in [d]$, if every user has bit b in the j -th bit of their codeword then the one-way marginal answer for that column will be b . It is shown that any algorithm that has non-trivial accuracy for answering the one-way marginals on D can be used to obtain a codeword that satisfies the marking condition. Therefore, using *Trace* on such a codeword leads to identifying an individual in dataset D . Since an adversary is able to identify a user in D based on the answer given by the algorithm, this clearly violates DP.

Techniques of [8] do not directly apply. In our model, an algorithm only sees a subset of the entries in the entire database via attribute queries. Suppose a coalition of $c \leq n$ users belongs to a dataset $D \in \{0, 1\}^{n \times d}$ where each row corresponds to a codeword of length d . As a warm-up, let us first assume that an algorithm that solves the one-way marginal problem on input $D \in \{0, 1\}^{n \times d}$ always queries for entire rows and that an adversary can simulate a query oracle to the algorithm’s queries, i.e., respond with rows that exactly correspond to the set of c colluders (for more details see Section 5.1). To apply a fingerprinting code argument to such an algorithm, an adversary must identify a row from this subset of c rows by examining the output of the algorithm. However, since the accuracy guarantee of the algorithm applies (only) to the *whole* dataset, we cannot make the same argument as above to conclude that the marking condition holds for the *subsample* of rows. In other words, we need to ensure that the output of an accurate algorithm that only sees a subsample of rows can also satisfy the marking condition. The techniques of [8] do not ensure such a property.

Permute Rows and Pad and Permute Columns Fingerprinting codes (PR-PPC FPC).

In order to achieve the property described above, we need to ensure that any attempt of the algorithm to spoil the marking condition would contradict its accuracy guarantees. We achieve this property by padding $O(d)$ additional columns to the codebook C to obtain $C' \in \{0, 1\}^{n \times d'}$, where $d' = O(d)$, so that (codebook) columns whose output could be modified to violate the marking condition and (padded) columns whose modification would violate the accuracy guarantee are indistinguishable in the subsample with good probability. Padded columns have been used to define a variant of a fingerprinting code in previous work to achieve smooth DP lower bounds [22]. In our work, we not only need a variant of FPC with padded columns, but we also need to permute the rows of the codebook (see Section 4 for the construction). This is because we need to define a sampling procedure with certain properties for the adversary to obtain a dataset on c rows from a distribution over databases of n rows, and one way to do so is permuting the rows of the codebook and outputting the first c rows (e.g., see Theorem 15 and Theorem 27).

³ The idea of fingerprinting codes becomes colorful when imagining a publisher who distributes advance copies to press and wants to add watermarks that are robust, e.g., against pirated copies that result from averaging the copies of multiple colluders. To see that the marking assumption is a mild condition, consider that the codeword is hidden in the much larger content.

► **Remark 2.** We note that [8] used a similar padding technique to argue about obtaining error-robust codes from “weakly-robust” codes (see Lemma 6.4 in [8]). In particular, we could argue that the property that we need for our purpose is achieved by an error robust code. However, we choose to start with a weaker construction, as we do not inherently need the error robustness property.

Secret Sharing Encoding. Finally, we overcome the assumption that the algorithm queries for entire rows by applying a secret sharing scheme to the padded codebook (see Section 5.2). In particular, an adversary can encode each row $x_i \in \{0, 1\}^{d'}$ with respect to a random polynomial of degree $2d' - 1$ as a share of size $2d$. The shares are defined by the d codebook values and d random values from the field. For each query of the algorithm, the adversary answers with a share from the second half. Information theory implies that the algorithm can only recover the d codebook values after querying for all d random value shares. Thus, we obtain a derivate of the one-way marginal problem that requires the algorithm to query an entire row to reveal the padded code book row. While there exist a DP algorithm (Theorem 24) and a sublinear-time algorithm (Theorem 25) for this derived problem as well, we show that there exists no sublinear-time DP algorithm (up to a log factor) that can query for arbitrary entries in the database.

2 Related Work

Fingerprinting codes were first introduced in the context of DP lower bounds by [8]. Prior to their work, traitor-tracing schemes (which can be thought of as a cryptographic analogue of information-theoretic fingerprinting codes) were used by [12, 26] to obtain computational hardness results for DP. Subsequent works have refined and generalized the connection between DP and fingerprinting codes in many ways [23, 24, 14, 7, 21, 20, 9]. The fingerprinting code techniques of proving DP lower bounds have been used in many settings including principal component analysis [15], empirical risk minimization [2], mean estimation [8, 20], regression [9], gaussian covariance estimation [21]. Recently [22] use fingerprinting codes to give smooth lower bounds for many problems including the 1-cluster problem and k -means clustering.

In the streaming model, [10] give a separation between the space complexity of differentially private algorithms and non-private algorithms – under cryptographic assumptions they show that there exists a problem that requires exponentially more space to be solved efficiently by a DP algorithm vs a non-private algorithm. By contrast, our focus is on lower bounding the running time (query complexity) of a differentially private algorithm. Our bounds do not require any cryptographic assumptions. [19] give a lower bound in the continual release model, in particular they show that there exists a problem for which any DP continual release algorithm has error $\tilde{\Omega}(T^{1/3})$ times larger than the error of a DP algorithm in the static setting where T is the length of the stream.

3 Preliminaries

We define a database $D \in \mathcal{X}^n$ to be an ordered tuple of n rows $(x_1, \dots, x_n) \in \mathcal{X}$, where \mathcal{X} is the data universe. For our purposes, we typically take $\mathcal{X} = \{0, 1\}^d$. Databases D and D' are neighboring if they differ by a single row and we denote this by $D \sim D'$. In more detail, we can replace the i -th row of a database D with some fixed element of \mathcal{X} to obtain dataset $D_{-i} \sim D$. Importantly both D and D_{-i} are databases of the same size.

► **Definition 3** (Differential Privacy [11]). *Randomized algorithm $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{R}$ is (ε, δ) -differentially private if for every two neighboring databases $D \sim D'$ and every subset $S \subseteq \mathcal{R}$,*

$$\Pr[\mathcal{A}(D) \in S] \leq e^\varepsilon \Pr[\mathcal{A}(D') \in S] + \delta$$

► **Definition 4** (Accuracy). *Randomized algorithm $\mathcal{A} : \mathcal{X}^n \rightarrow \mathbb{R}^d$ is (α, p) -accurate for problem \mathcal{P} if for every $D \in \mathcal{X}^n$, with probability at least $1 - p$, the output of \mathcal{A} is a vector $a \in \{0, 1\}^d$ that satisfies $|a_{\mathcal{P}}(D) - a| \leq \alpha$ where $a_{\mathcal{P}}(D)$ denotes the exact solution of the problem \mathcal{P} on input D .*

The following definition of fingerprinting codes is “fully-collusion-resilient”. For any coalition of users S who collectively produce a string $y \in \{0, 1\}^d$ as output, as long as y satisfies the *marking condition* – for all positions $1 \leq j \leq d$, if the values x_{ij} for all users i in coalition S agree with some letter $s \in \{0, 1\}$, then $y_j = s$ – then the combined codeword y can be traced back to a user in the coalition. Formally, for a codebook $C \in \{0, 1\}^{n \times d}$, and a coalition $S \subseteq [n]$, we define the set of feasible codewords for C_S to be

$$F(C_S) = \{c' \in \{0, 1\}^d \mid \forall j \in [d], \exists i \in S, c'_j = c_{ij}\}$$

► **Definition 5** (Fingerprinting Codes [8]). *For any $n, d \in \mathbb{N}, \xi \in (0, 1]$, a pair of algorithms $(Gen, Trace)$ is an (n, d, c) -fingerprinting code with security ξ against a coalition of size c if Gen outputs a codebook $C \in \{0, 1\}^{n \times d}$ and secret state st and for every (possibly randomized) adversary \mathcal{A}_{FP} , and every coalition $S \subseteq [n]$ such that $|S| \leq c$, if we set $c' \leftarrow_R \mathcal{A}_{FP}(C_S)$ then*

1. $\Pr[c' \in F(C_S) \wedge Trace(c') = \perp] \leq \xi$
2. $\Pr[Trace(c') \in [n] \setminus S] \leq \xi$

where C_S contains the rows of C given by S , and the probability is taken over the coins of C , $Trace$, and \mathcal{A}_{FP} . The algorithms Gen and $Trace$ may share a common state denoted as st .

► **Remark 6.** Although the adversary \mathcal{A}_{FP} is defined as taking the coalition of users’ rows as input, we may abuse this notation and consider the entire codebook or a different input (related to the codebook) altogether. This does not change the security guarantees of the FPC against adversary \mathcal{A}_{FP} because the security guarantee holds as long as the output of \mathcal{A}_{FP} is a result of the users in the coalition S actively colluding.

► **Theorem 7** (Tardos Fingerprinting Code [25]). *For every $n \in \mathbb{N}$ and $4 \leq c \leq n$, there exists an (n, d, c) -fingerprinting code of length $d = O(c^2 \log(n/\xi))$ with security $\xi \in [0, 1]$ against coalitions of size c .*

► **Theorem 8** (Gaussian Mechanism, [13]). *Let $\varepsilon, \delta \in (0, 1)$ and $f : \mathbb{N}^d \rightarrow \mathbb{R}^d$. For $c > \sqrt{2 \ln(1.25/\delta)}/\varepsilon$, the Gaussian Mechanism with standard deviation parameter $\sigma \geq c\Delta_2 f$ is (ε, δ) -DP, where Δ_2 is the ℓ_2 -norm sensitivity of f .*

► **Lemma 9.** *For $n \geq \sqrt{200d \ln(20d) \ln(1.25/\delta)}/\varepsilon \in \tilde{\Omega}(\sqrt{d})$, given a dataset $D \in \{0, 1\}^{n \times d}$, there exists a $(1/10, 1/10)$ -accurate (ε, δ) -DP algorithm that solves the one-way marginals problem with $O(m)$ attribute queries, where $m = n \cdot d$.*

Proof. We note that the ℓ_2 -sensitivity of the one-way marginals problem on a database $\{0, 1\}^{n \times d}$ is \sqrt{d}/n . For $n \geq \sqrt{200d \ln(20d) \ln(1.25/\delta)}/\varepsilon \in \tilde{\Omega}(\sqrt{d})$, the Gaussian Mechanism is $(1/10, 1/10)$ -accurate with

$$\sigma = \frac{\sqrt{2 \ln(1.25/\delta)}}{\varepsilon} \cdot \frac{\varepsilon \sqrt{200d}}{\sqrt{2 \ln(1.25/\delta) \cdot d \ln(10d)}} = \frac{1}{\sqrt{200 \ln(10d)}},$$

$$\text{as } \Pr_{X \sim \mathcal{N}(0, \sigma^2)} \left[X \geq \frac{1}{10} \right] \leq 2e^{-\frac{1}{200\sigma^2}} \leq \frac{1}{20d}$$

by the Cramer-Chernoff inequality and a union bound over all d columns of the dataset. ◀

4 Permute Rows and Pad and Permute Columns Fingerprinting Codes (PR-PPC FPC)

In this section, we first introduce our pad and permute variant of the original fingerprinting codes where we Permute Rows and Pad and Permute Columns (PR-PPC (n, d, c, ℓ) -FPC). Given $(Gen, Trace)$ of an (n, d, c) -FPC, we construct Gen' and $Trace'$ in Algorithm 1 and Algorithm 2 to produce a PR-PPC (n, d, c, ℓ) -FPC. In more detail, Gen' samples codebook C and secret state st from Gen where $C \in \{0, 1\}^{n \times d}$. It then permutes the rows via a random permutation π_R , after which it pads 2ℓ columns and performs another random permutation π on the columns. Then, it releases the resulting codebook $C' \in \{0, 1\}^{n \times d'}$, where $d' = d + 2\ell$ and ℓ is the parameter which controls the number of columns padded to C . Note that the row permutation π_R is public while the column permutation π is part of the new secret state st' . The algorithm $Trace'$ receives an answer vector of dimension d' and uses its secret state $st' = (st, \pi)$ to feed the vector entries that correspond to the original first d columns via π^{-1} to $Trace$ and releases the output of $Trace$. We obtain the following result directly from the definition of Gen' and $Trace'$.

► **Corollary 10.** *Given an (n, d, c) -FPC, $\ell \geq 0$, and the corresponding PR-PPC (n, d, c, ℓ) -FPC, the properties of $Trace$ as stated by Definition 5 translate directly to $Trace'$.*

Algorithm 1 Gen' .

Require: Number of users $n \in \mathbb{N}$, number of padded 0/1 columns ℓ

- 1: Sample codebook $(C, st) \leftarrow Gen(n)$ such that $C \in \{0, 1\}^{n \times d}$.
- 2: Sample random permutation $\pi : [d'] \rightarrow [d']$ where $d' := d + 2\ell$. For an $n \times d'$ matrix A , define $n \times d'$ matrix $\pi(A)$ such that $\pi(j)$ -th column of $\pi(A)$ equals to the j -th column of A for every $j \in [d']$.
- 3: Sample random permutation $\pi_R : [n] \rightarrow [n]$. For an n -row matrix A , define $\pi_R(A)$ such that $\pi_R(i)$ -th row of $\pi_R(A)$ equals to the i -th row of A for every $i \in [n]$.
- 4: $C^{\pi_R} \leftarrow$ Permute rows of C via random permutation π_R .
- 5: $C_{pad} \leftarrow$ Pad ℓ columns of all 1's and ℓ columns of all 0's to matrix C^{π_R} .
- 6: $C' \leftarrow$ Permute the columns of C_{pad} according to random permutation π .
- 7: Output C' along with the new secret state $st' := (st, \pi)$ and permutation π_R .

Algorithm 2 $Trace'$.

Require: Answer vector $\mathbf{a} = (a^1, \dots, a^{d'}) \in \{0, 1\}^{d'}$, secret state $st' = (\pi, st)$

- 1: Output $Trace(\mathbf{a}_{og}, st)$ where $\mathbf{a}_{og} = (a^{\pi(1)}, \dots, a^{\pi(d)}) \in \{0, 1\}^d$.

We define the *feasible sample property* of an FPC below. Informally, it states that if we have an algorithm that takes a sample of rows from a codebook as input, and the algorithm's output is a feasible codeword for the entire codebook, then the same output should be a feasible codeword for the sample.

► **Definition 11 (Feasible Sample Property).** *Let $C \in \{0, 1\}^{n \times d}$ be a codebook of an (n, d, c) -FPC, $S \subseteq [n]$ be a coalition and $C_S \subseteq C$ be the matrix consisting of the corresponding rows indexed by S . Given an algorithm \mathcal{A} that takes as input C_S and outputs a vector $\mathbf{o} \in \{0, 1\}^d$, the feasible sample property states that if $\mathbf{o} \in F(C)$, then $\mathbf{o} \in F(C_S)$.*

► **Lemma 12.** *PR-PPC (n, d, c, ℓ) -FPC satisfies the feasible sample property with probability at least $1 - \frac{d}{\ell}$.*

Proof. Given $(Gen', Trace')$ of PR-PPC (n, d, c, ℓ) -FPC which produces codebook $C' \in \{0, 1\}^{n \times d'}$ and sampling algorithm \mathcal{A} which takes as input $C'_S \subseteq C'$ and outputs a vector $\mathbf{o} \in \{0, 1\}^{d'}$ where $d' = d + 2\ell$, we define the event BAD_S as $\mathbf{o} \in F(C')$ but $\mathbf{o} \notin F(C'_S)$.

We denote the indices of columns of C' in which all the bits are 1 as $C'_{|1}$ and the indices of columns in which all the bits are 0 as $C'_{|0}$. Similarly, we define $C'_{S|1}$ and $C'_{S|0}$ for the columns that are all 1 and all 0 in C'_S , respectively. Note that $C'_{|1} \subseteq C'_{S|1}$ and $C'_{|0} \subseteq C'_{S|0}$. Since by definition, algorithm \mathcal{A} only has access to the set of rows in C'_S , in order for the output \mathbf{o} to satisfy $\mathbf{o} \in F(C')$ but $\mathbf{o} \notin F(C'_S)$, an adversary that aims for BAD_S must flip a bit of the resulting codeword that originates from $\{C'_{S|1} \cup C'_{S|0}\} \setminus \{C'_{|1} \cup C'_{|0}\}$. In other words, the event BAD_S occurs only if the adversary identifies a column from C' that contains at least one 0 and one 1, but reduces to an all-1 or all-0 column in C'_S .

More formally, the adversary can pick the bit $b \in \{0, 1\}$ resulting in a column from $C'_{S|b}$ to flip. The probability that the adversary correctly identifies a column from $C'_{S|b} \setminus C'_{|b}$ is at most $\frac{|C'_{S|b} \setminus C'_{|b}|}{|C'_{S|b}|}$. Observe that $|C'_{S|b}| \geq |C'_{|b}| \geq \ell$ due to the ℓ padded all- b columns, and therefore $|C'_{S|b} \setminus C'_{|b}| \leq |C'_{S|b}| - |C'_{|b}| \leq (d + \ell) - \ell = d$. Thus, the probability that event BAD_S occurs is at most $\max_{b \in \{0, 1\}} \frac{|C'_{S|b} \setminus C'_{|b}|}{|C'_{S|b}|} \leq \frac{d}{\ell}$. \blacktriangleleft

5 Lower Bound

We present our lower bound for sublinear-time DP algorithms in this section. In Section 5.1 we discuss a warm-up problem, where the algorithm can only make row queries to release the one-way marginals of the dataset. We present our main result and the lower bound construction for algorithms that can make arbitrary attribute queries in Section 5.2. In the sequel, our problem space has size $m = n \cdot d = \Omega(d\sqrt{d})$ and our results will be in terms of dimension d .

5.1 Warm Up: Using a Random Oracle

In this section, we first present a warm-up problem which we call the Random Oracle Problem (\mathcal{P}_{RO}). This is an extension of the one-way marginals problem in the following manner – for an input dataset $D = (x_1, \dots, x_n)$, and access to a random oracle H , the \mathcal{P}_{RO} problem takes as input an encoded dataset $D_H = (z_1, \dots, z_n)$ in which $z_i = H(i) \oplus x_i$, and outputs the one-way marginals of the underlying dataset D (see Definition 13 for a formal definition). The main intuition for introducing such a problem is that we want to force an algorithm that solves this problem to query an entire row. Recall that in our (final) model an algorithm is allowed to make arbitrary attribute queries – however, given D_H , in order to approximate or exactly compute the one-way marginals of D , an algorithm needs to query $H(i)$ for $i \in [n]$, as otherwise, by the properties of the random oracle and one-time pad (OTP), the value of x_i is information-theoretically hidden.

► Definition 13 (Random Oracle Problem \mathcal{P}_{RO}). *Given a random oracle $H : [n] \rightarrow \{0, 1\}^d$, and a dataset $D = (x_1, \dots, x_n)$ where $x_i \in \{0, 1\}^d$, define dataset $D_H := (z_1, \dots, z_n)$ where $z_i = H(i) \oplus x_i$. For simplicity of notation, we refer to the operation for obtaining D_H from D as $H(D)$. The problem \mathcal{P}_{RO} on input D_H releases the one-way marginals of D .*

We use $\mathcal{P}_{RO}(D)$ to denote that \mathcal{P}_{RO} releases the one-way marginals of the underlying dataset D .

19:10 Differential Privacy and Sublinear Time Are Incompatible Sometimes

Query Model. On input $D_H \in (\{0, 1\}^d)^n$, an algorithm can query the random oracle H through *row queries*, i.e., given a row index $i \in [n]$ of D_H , the answer given is $H(i) \in \{0, 1\}^d$. We note that our final result in this subsection will still be presented in the form of attribute queries as 1 row query translates to d attribute queries.

Observe that there exists an (ε, δ) -DP algorithm for $\mathcal{P}_{RO}(D)$ that on input D_H , queries the entire dataset via row queries to the random oracle H , i.e., it makes $dn = O(d\sqrt{d})$ queries. After obtaining the rows to the underlying dataset D it releases the one-way marginals using the Gaussian Mechanism (see Lemma 9). We also note that there exists a sublinear non-DP algorithm for $\mathcal{P}_{RO}(D)$ which makes $O(d \log d)$ queries, which is a simple corollary of Hoeffding bounds. Our goal in this section is to prove the lower bound below. Recall that $n \in \Omega(\sqrt{d})$, so the problem size is $\Omega(d\sqrt{d})$.

► **Theorem 14** (Lower Bound for \mathcal{P}_{RO}). *Any algorithm that solves the problem \mathcal{P}_{RO} with s attribute query complexity, $(1/3, 1/3)$ -accuracy and $(O(1), o(1/s))$ -DP must have $s = \Omega(d\sqrt{d}/\log(d))$.*

We present a high level overview of the proof of Theorem 14 here. We first show that given an (n, d, c) -FPC, there exists a distribution on c rows from which an adversary \mathcal{B} can sample and create an n -row input instance for an algorithm \mathcal{A} that accurately solves \mathcal{P}_{RO} (see Theorem 15). Next we argue that the rounded output of \mathcal{A} , denoted as \mathbf{a} , is a feasible codeword for the sample of c rows as long as \mathcal{A} is accurate in a non-trivial manner and \mathbf{a} is feasible for the entire dataset (see Lemma 17). The adversary \mathcal{B} can then use the output from \mathcal{A} to (potentially) reidentify an individual from the coalition of size c . Next we relate these claims back to DP through Lemma 18, which states that if there exists a distribution \mathcal{C} on $c \leq n$ row databases according to Theorem 15, then there is no (ε, δ) -DP algorithm \mathcal{A} that is $(1/3, 1/3)$ -accurate for \mathcal{P}_{RO} with $\varepsilon = O(1)$ and $\delta = o(1/c)$. Finally, invoking the Tardos construction for fingerprinting codes in Theorem 7 gives us our lower bound.

► **Theorem 15.** *For every $n, d \in \mathbb{N}$, $\xi \in [0, 1]$ and $c \leq n$, if there exists an (n, d, c) -fingerprinting code with security ξ , then there exists a distribution on c -row databases \mathcal{C}_S , a row permutation $\pi_R : [n] \rightarrow [n]$, and an adversary \mathcal{B} for every randomized algorithm \mathcal{A} with row query complexity c and $(1/3, 1/3)$ -accuracy for \mathcal{P}_{RO} such that*

1. $\Pr_{C'_S \leftarrow \mathcal{C}_S}[\mathcal{B}^{\mathcal{A}}(C'_S) = \perp] \leq \xi$
2. For every $i \in [c]$, $\Pr_{C'_S \leftarrow \mathcal{C}_S}[\mathcal{B}^{\mathcal{A}}(C'_{S-i}) = \pi_R^{-1}(i)] \leq \xi$.

The probabilities are taken over the random coins of \mathcal{B} and the choice of C'_S .

Let $(Gen, Trace)$ be the promised (n, d, c) -fingerprinting code in the theorem statement. We first construct a PR-PPC (n, d, c, ℓ) -FPC with $\ell := 100d$ (see Section 4 for details).

The distribution \mathcal{C}_S on c -row databases is implicitly defined through the sampling process below

1. Let $C' \leftarrow Gen'(n, 100d)$ (see Algorithm 1) where $C' \in \{0, 1\}^{n \times d'}$ and $d' = d + 100d = 101d$. Note that Gen' also outputs π_R which is a public permutation on rows.
2. Let $C'_S = (x_1, \dots, x_c) \in \{0, 1\}^{c \times d'}$ be the first c rows of $C' \in \{0, 1\}^{n \times d'}$
3. Output C'_S

Next we define the privacy adversary \mathcal{B} .

Adversary \mathcal{B} Algorithm. Adversary \mathcal{B} receives C'_S as input and does the following:

1. Create a database $D = (r_1, \dots, r_n) \in \{0, 1\}^{n \times d'}$ where each row $r_i \in \{0, 1\}^{d'}$ consists of 0/1 entries sampled independently and uniformly at random.

2. Given oracle access to randomized algorithm \mathcal{A} which solves \mathcal{P}_{RO} on input D , \mathcal{B} simulates the answer to the distinct i_j -th row query (where $j \in [c]$) made by \mathcal{A} to random oracle H as follows:
 - a. Return $H(i_j) := r_{i_j} \oplus x_j$.
3. Let \mathbf{a} be the output of $\mathcal{A}(D)$ where $\mathbf{a} \in [0, 1]^{d'}$. Round each entry of \mathbf{a} to $\{0, 1\}$, call this new vector $\bar{\mathbf{a}} \in \{0, 1\}^{d'}$.
4. Output $Trace'(\bar{\mathbf{a}})$

Analysis. We focus on proving that Property 1 and Property 2 of the theorem statement are indeed satisfied by adversary \mathcal{B} .

Recall the notation in Definition 13 where $\mathcal{P}_{RO}(C')$ means that \mathcal{P}_{RO} releases the one-way marginals of the underlying dataset C' . We first show that \mathcal{A} solving $\mathcal{P}_{RO}(H(D))$ is perfectly indistinguishable from \mathcal{A} solving $\mathcal{P}_{RO}(C')$ in Lemma 16. This is necessary as $Trace'$ can only identify an individual in the coalition of size c with respect to the codebook C' produced by Gen' .

► **Lemma 16.** *\mathcal{A} solving $\mathcal{P}_{RO}(H(D))$ is perfectly indistinguishable from \mathcal{A} solving $\mathcal{P}_{RO}(C')$.*

Proof. We define the following experiments.

Real World.

1. Given $C' = (x_1, \dots, x_n)$ where $(C', st') \leftarrow Gen'(\ell)$ with $\ell = 100d$, let $C'_S = (x_1, \dots, x_c)$.
2. Create a database $D = (r_1, \dots, r_n)$ where $r_i \in \{0, 1\}^{d'}$ are random entries.
3. Let \mathbf{a} be the output of $\mathcal{A}(D)$ where $\mathbf{a} \in [0, 1]^{d'}$. Simulate H as follows:
 - a. Let i_1, \dots, i_c be distinct queries made to H . For $j \in [c]$, fix $H(i_j) := r_{i_j} \oplus x_j$

Ideal World.

1. Given codebook $C' = (x_1, \dots, x_n)$ where $(C', st') \leftarrow Gen'(\ell)$ with $\ell = 100d$, let $H(C') = (z_1, \dots, z_n)$ (see Definition 13 for $H(\cdot)$ notation).
2. Let $\mathbf{a} \leftarrow \mathcal{A}(H(C'))$ where $\mathbf{a} \in [0, 1]^{d'}$.
 - a. Let i_1, \dots, i_c be distinct arbitrary queries made to H . For $j \in [c]$, H returns the following answer $H(i_j) := z_{i_j} \oplus x_j$

In the **Real World**, \mathcal{A} is provided $D = (r_1, \dots, r_n)$ as input (where D is generated in the same manner as by adversary \mathcal{B}), while the **Ideal World** is one in which \mathcal{A} takes $H(C')$ as input. We show that \mathcal{A} learns the same information in the **Real World** and the **Ideal World**, i.e., these views are perfectly indistinguishable. Observe that the only difference from the viewpoint of \mathcal{A} between the **Real World** and the **Ideal World** is that H is simulated in the former via indices fixed by C'_S whereas H is queried on arbitrary indices in the latter. Since the rows of C' have already been permuted (recall Algorithm 1), by nature of the random oracle H , these two instances are perfectly indistinguishable. ◀

Recall that the security condition of the fingerprinting code (see Definition 5) only holds if $\bar{\mathbf{a}}$ is a feasible codeword for the coalition of rows, i.e., C'_S in our case. The following lemma states that if \mathcal{A} is accurate for $\mathcal{P}_{RO}(C')$, then the rounded output of \mathcal{A} is indeed a feasible codeword for both C' and C'_S .

► **Lemma 17.** *Suppose \mathcal{A} is $(1/3, 1/3)$ -accurate for $\mathcal{P}_{RO}(C')$. Then the rounded output $\bar{\mathbf{a}}$ from algorithm \mathcal{A} is a feasible codeword for both C' and C'_S with probability at least $1 - \frac{1}{3} - \frac{1}{100}$.*

In other words, with probability at least $1 - \frac{1}{3} - \frac{1}{100}$, $\bar{\mathbf{a}} \in F(C')$ and $\bar{\mathbf{a}} \in F(C'_S)$.

19:12 Differential Privacy and Sublinear Time Are Incompatible Sometimes

Proof. Assuming that \mathcal{A} is $(1/3, 1/3)$ -accurate for $\mathcal{P}_{RO}(C')$, we first show that $\bar{\mathbf{a}}$ is a feasible codeword for C' with probability at least $2/3$. By the accuracy guarantee of \mathcal{A} , we know that for any column i_j $|\mathbf{a}_{i_j} - a_{i_j}| \leq 1/3$ where a_{i_j} is the actual 1-way marginal for column i_j with probability at least $2/3$. Thus for any column i_j of all 1's in C' , $\mathbf{a}_{i_j} \geq 2/3$ which means $\bar{\mathbf{a}}_{i_j} = 1$, thus satisfying the marking condition. A similar argument holds for the case when a column is all 0's.

Next, using the fact that we use a PR-PPC $(n, d, c, 100d)$ -FPC and that $\bar{\mathbf{a}} \in F(C')$ with probability at least $2/3$, we can invoke Lemma 12 which states that the feasible sample property is satisfied by our PR-PPC FPC construction. Note that in our case, the sampling algorithm described in Definition 11 is \mathcal{A} together with the postprocessing step of rounding the output of \mathcal{A} . Also, even though \mathcal{A} takes the entire dataset as input, it effectively only has access to the rows of the underlying sample via queries to \mathcal{B} and thus satisfies the properties required in Definition 11. Lemma 12 states that with probability $\leq \frac{1}{100}$, $\bar{\mathbf{a}}$ is not a feasible codeword for C'_S . By a union bound we have that $1 - \frac{1}{3} - \frac{1}{100}$, $\bar{\mathbf{a}}$ must be a feasible codeword for C'_S . ◀

Proof of Theorem 15. From the above Lemma 17, we have that \mathcal{A} is $(1/3, 1/3)$ -accurate for $\mathcal{P}_{RO}(C')$ implies that $\bar{\mathbf{a}}$ is a feasible codeword for C'_S . By the security of the fingerprinting code, Corollary 10 and Lemma 16, we have that $\Pr[\bar{\mathbf{a}} \in F(C'_S) \wedge \text{Trace}'(\bar{\mathbf{a}}) = \perp] \leq \xi$. Since \mathcal{B} releases the output of $\text{Trace}'(\bar{\mathbf{a}})$, the event $\mathcal{B}^{\mathcal{A}}(C'_S) = \perp$ is identical to $\text{Trace}'(\bar{\mathbf{a}}) = \perp$. Thus Property 1 of the theorem statement which states that the probability that \mathcal{B} outputs \perp is bounded by ξ follows. Property 2 follows directly from the soundness property of the fingerprinting code. ◀

► **Lemma 18.** *Suppose there exists a distribution on $c \leq n$ row databases \mathcal{C}_S according to Theorem 15. Then there is no (ε, δ) -DP algorithm \mathcal{A} with query complexity c that is $(1/3, 1/3)$ -accurate for \mathcal{P}_{RO} with $\varepsilon = O(1)$ and $\delta = o(1/c)$.*

Proof. Suppose C'_S is sampled from the distribution on c -row databases \mathcal{C}_S and \mathcal{B} is the adversary from Theorem 15. From the lemma statement we know that \mathcal{A} is $(1/3, 1/3)$ -accurate, thus using Lemma 17 and Theorem 15, we have that $\Pr[\pi_R(\mathcal{B}^{\mathcal{A}}(C'_S)) \in [c]] \geq 1 - \frac{1}{3} - \frac{1}{100} - \xi \geq \Omega(1)$. By an averaging argument, this means that there exists some $i^* \in [c]$ for which $\Pr[\pi_R(\mathcal{B}^{\mathcal{A}}(C'_S)) = i^*] \geq \Omega(1/c)$. However, if $\xi = o(1/c)$ by Property 2 in Theorem 15 we have that $\Pr[\pi_R(\mathcal{B}^{\mathcal{A}}(C'_{S_{i^*}})) = i^*] \leq \xi = o(1/c)$.

In other words, the probability of $\mathcal{B}^{\mathcal{A}}$ outputting a fixed output i^* on neighboring input databases C'_S and $C'_{S_{i^*}}$ is different, which violates (ε, δ) -DP for any $\varepsilon = O(1)$ and $\delta = o(1/c)$. We note here that since \mathcal{A} can make at most c row queries, the DP guarantee for \mathcal{A} must hold for any neighboring sample of c rows. Since \mathcal{B} does some postprocessing of the output from \mathcal{A} , and we have shown that \mathcal{B} cannot be (ε, δ) -DP, this implies that \mathcal{A} cannot be (ε, δ) -DP for any $\varepsilon = O(1)$ and $\delta = o(1/c)$. ◀

5.2 Using a Secret Sharing Encoding

In this section, we remove the requirement of an algorithm querying an entire row that we enforced in the previous section. We first define the security requirement of a general encoding scheme that is sufficient to construct our DP lower bound in Definition 19. We then show that the Shamir encoding as defined in Definition 20 satisfies the security requirement (see Theorem 23). We define a problem based on this secret sharing encoding called \mathcal{P}_{SS} (see Definition 21) that uses the encoding to release the one-way marginals of an underlying dataset. Finally, we show that this problem cannot have a sublinear time DP algorithm with reasonable accuracy (see Theorem 26).

► **Definition 19** (Security Game). Let $\text{Exp}(\text{Enc}_d, \mathcal{A}, q, d, x)$ denote the following experiment: (1) the challenger computes $y_0 \leftarrow \text{Enc}_d(x)$ and $y_1 \leftarrow \text{Enc}_d(0^d)$, picks a random bit b and outputs $y = y_b$. (2) $\mathcal{A}^y(d, q, x)$ is given oracle access to y and may make up to q queries to the string y . (3) The game ends when the attacker \mathcal{A} outputs a guess b' . (4) The output of the experiment is $\text{Exp}(\text{Enc}_d, \mathcal{A}, q, d, x) = 1$ if $b' = b$ and the attacker made at most q queries to y ; otherwise the output of the experiment is $\text{Exp}(\mathcal{A}, q, d, x) = 0$. We say that the scheme Enc_d is $(q(d), d, \gamma(q, d))$ -secure if for all $x \in \{0, 1\}^d$ and all attackers \mathcal{A} making at most q queries we have

$$\Pr[\text{Exp}(\text{Enc}_d, \mathcal{A}, q, d, x) = 1] \leq \frac{1}{2} + \gamma(q, d)$$

► **Definition 20** (Shamir Encoding). Given a row $x_i \in \{0, 1\}^d$ where $i \in [n]$ and a field \mathbb{F} s.t. $|\mathbb{F}| > 4d$ let $SS_d(x_i)$ be the following encoding (1) pick random field elements $\alpha_1^{(i)}, \dots, \alpha_d^{(i)}, \alpha_{d+1}^{(i)}, \dots, \alpha_{3d}^{(i)}$ (distinct) and $z_{d+1}^{(i)}, \dots, z_{2d}^{(i)}$ and define the polynomial $p_i(\cdot)$ of degree $2d - 1$ s.t. $p_i(\alpha_j^{(i)}) = x_j$ and $p_i(\alpha_{d+j}^{(i)}) = z_{d+j}^{(i)}$ for $j \leq d$. (2) publish $SS_d(x_i) = \left(\alpha_1^{(i)}, \dots, \alpha_d^{(i)}, \{(\alpha_j^{(i)}, p_i(\alpha_j^{(i)}))\}_{j=d+1}^{3d} \right)$ as share of x_i .

► **Definition 21** (Secret Sharing Problem \mathcal{P}_{SS}). Let dataset $D := (x_1, \dots, x_n) \in \{0, 1\}^{n \times d}$. Given $D_S := (SS_d(x_1), \dots, SS_d(x_n))$, the goal of the secret-sharing problem \mathcal{P}_{SS} is to release all the one-way marginals of dataset D .

We use $\mathcal{P}_{SS,d}(D)$ to denote that \mathcal{P}_{SS} releases the one-way marginals of the underlying dataset D with dimension d .

Query Model. On input D_S , an algorithm solving the \mathcal{P}_{SS} problem can make attribute queries to obtain the underlying dataset D and release its one-way marginals. For a row $i \in [n]$, the i_j -th attribute query returns the pair of field elements $(\alpha_{j+d}^{(i)}, p(\alpha_{j+d}^{(i)}))$ of share $SS_d(x_i)$ for $1 \leq j \leq 2d$. We note that the prefix of $SS_d(x_i)$ given by $\alpha_1^{(i)}, \dots, \alpha_d^{(i)}$ is published separately after an attribute query for the row i has been queried. In other words, the prefix does not count towards the query complexity of the algorithm.

► **Remark 22.** We remark that one can also define a different query model in which the prefix is released to the adversary whenever the i -th row is queried and our results still hold.

For completeness, we first show that the Shamir encoding SS_d defined in Definition 20 is $(q(d), d, 0)$ -secure (as defined in Definition 19) where $q(d) = d$.

► **Theorem 23.** The scheme SS_d is $(d, d, 0)$ -secure.

Proof. Let $x \in \{0, 1\}^d$ and field \mathbb{F} s.t. $|\mathbb{F}| > 4d$. Recall the secret sharing scheme $SS_d(x) = (\alpha_1, \dots, \alpha_d, \{(\alpha_j, p(\alpha_j))\}_{j=d+1}^{3d})$ defined in Definition 20. We describe two experiments below where the Real World experiment simulates the view of the adversary and the Ideal World experiment just randomly outputs field elements. We will show that these two experiments are perfectly indistinguishable, and the security claim follows.

Real World(x).

1. Query $SS_d(x)$ for the first $q(d) = d$ pairs of coordinates and let the answers be the prefix $\alpha_1, \dots, \alpha_d$ and $\{(\alpha_{j+d}, z_{j+d})\}_{j \in [d]}$.
2. Output $\alpha_1, \dots, \alpha_d$ and $\{(\alpha_{j+d}, z_{j+d})\}_{j \in [d]}$

Ideal World(x).

1. Uniformly sample $\alpha'_1, \dots, \alpha'_d, \alpha'_{d+1}, \dots, \alpha'_{2d}, r_{d+1}, \dots, r_{2d}$ from \mathbb{F} .
2. Output $\alpha'_1, \dots, \alpha'_d$ and $\{(\alpha'_{j+d}, r_{j+d})\}_{j \in [d]}$

Since by construction, the first d pairs of coordinates returned by SS_d and the prefix of size d correspond to $3d$ random field elements, the view of the **Real World** is therefore just the uniform distribution on $3d$ field elements and thus is identical to that of the view of the **Ideal World**. ◀

We present our main lower bound result in Theorem 26. Before we proceed, we first demonstrate the existence of a DP linear-time algorithm and non-DP sublinear-time algorithm for \mathcal{P}_{SS} below.

► **Theorem 24.** *There exists a (ε, δ) -DP algorithm that solves the problem \mathcal{P}_{SS} with $O(d\sqrt{d})$ attribute query complexity and $(1/10, 1/10)$ -accuracy.*

Proof. On input D_S , the algorithm queries the entire dataset via attribute queries, i.e., it makes $dn = O(d\sqrt{d})$ queries. Given $SS(x_i) = (\alpha_1^{(i)}, \dots, \alpha_d^{(i)}, \{(\alpha_j^{(i)}, p_i(\alpha_j^{(i)}))\}_{j=d+1}^{3d})$ for a row $i \in [n]$, the algorithm first recovers the polynomial p_i of degree $2d - 1$ by doing Lagrange Interpolation over the $2d$ points given by $\{(\alpha_j^{(i)}, p_i(\alpha_j^{(i)}))\}_{j=d+1}^{3d}$. Then the original row x_i is obtained by evaluating $(p_i(\alpha_1^{(i)}), \dots, p_i(\alpha_d^{(i)}))$. Once the original rows x_1, \dots, x_n are recovered in this manner, the algorithm can release the one-way marginals by adding Gaussian noise as detailed in Lemma 9. ◀

► **Theorem 25.** *There exists a sublinear-time algorithm that solves the problem \mathcal{P}_{SS} with s attribute query complexity $O(d \log d)$ and $(1/10, 1/10)$ -accuracy.*

Proof. The algorithm makes $O(d \log d)$ attribute queries and performs the same decoding procedure as outlined in the proof of Theorem 24 to obtain the underlying $\log(d)$ rows and computes the one-way marginals on this subset of rows. The accuracy of this algorithm is a simple corollary of Hoeffding bounds. Recall that $n \in \Omega(\sqrt{d})$, so the problem size is $\Omega(d\sqrt{d})$. ◀

► **Theorem 26 (Main Theorem).** *Any algorithm that solves the problem \mathcal{P}_{SS} with s attribute query complexity, $(1/3, 1/3)$ -accuracy and $(O(1), o(1/n))$ -DP must have $s = \Omega(d\sqrt{d}/\log(d))$.*

In order to prove Theorem 26, we follow a similar strategy as presented in the warm-up Section 5.1. Given an (n, d, c) -FPC, we first show how to construct a c -row distribution and an adversary \mathcal{B} that can identify a user in the coalition of size c in Theorem 27.

► **Theorem 27.** *For every $n, d \in \mathbb{N}$, $\xi \in [0, 1]$ and $c \leq n$, if there exists an (n, d, c) -fingerprinting code with security ξ , then there exists a distribution on c -row databases \mathcal{C}_S , a row permutation $\pi_R : [n] \rightarrow [n]$ and an adversary \mathcal{B} for every randomized algorithm \mathcal{A} with attribute query complexity cd' and $(1/3, 1/3)$ -accuracy for \mathcal{P}_{SS} such that*

1. $\Pr_{C'_S \leftarrow \mathcal{C}_S}[\mathcal{B}^{\mathcal{A}}(C'_S) = \perp] \leq \xi$
2. For every $i \in [c]$, $\Pr_{C'_S \leftarrow \mathcal{C}_S}[\mathcal{B}^{\mathcal{A}}(C'_{S-i}) = \pi_R^{-1}(i)] \leq \xi$.

where $d' = 101d$ and the probability is over the random coins of \mathcal{B} and the choice of C'_S .

Let $(Gen, Trace)$ be the promised (n, d, c) -fingerprinting code in the theorem statement. We first construct a PR-PPC (n, d, c, ℓ) -FPC with $\ell := 100d$ (see Section 4 for details).

The distribution \mathcal{C}_S on c -row databases is implicitly defined through the sampling process below:

1. Let $C' \leftarrow \text{Gen}'(n, d')$ (see Algorithm 1) where $C' \in \{0, 1\}^{n \times d'}$ and $d' = 101d$.
2. Let $C'_S = (x_1, \dots, x_c) \in \{0, 1\}^{c \times d'}$ be the first c rows of $C' \in \{0, 1\}^{n \times d'}$.
3. Output C'_S .

Next we define the privacy adversary \mathcal{B} .

Adversary \mathcal{B} Algorithm. Let \mathbb{F} be a finite field of order q' where $q' > 4d'$. Adversary \mathcal{B} receives $C'_S = (x_1, \dots, x_c)$ as input and feeds the algorithm \mathcal{A} an input instance C'_B of $\mathcal{P}_{SS, d'}(C')$ by simulating answers to attribute queries made by \mathcal{A} as described in Step 2a below. \mathcal{B} then uses the rounded answer returned by \mathcal{A} (Step 2b) to obtain an individual in the coalition by invoking Trace' in Step 2c.

1. Initialize $q_i = 0$ for each row $i \in [n]$ and initialize a counter $t = 0$.
2. Simulate the oracle algorithm \mathcal{A} with query access to an $(n \times 2d')$ database C'_B :
 - a. When \mathcal{A} makes a fresh query (i, j) , update $q_i = q_i + 1$ and
 - If $q_i \leq d'$, then set $b = q_i + d'$. Respond with a random pair of field elements $(\alpha_b^{(i)}, z_b^i)$. Record this tuple.
 - If $q_i = d' + 1$, then
 - i. Increment t by one.
 - ii. Define the entire polynomial p_i randomly, subject to the constraints that it is consistent with row x_t and the previous responses sent for row i : $p_i(\alpha_j^{(i)}) = x_{t,j}$ for $j \leq d'$ and $p_i(\alpha_b^{(i)}) = z_b^i$ for $j > d'$.
 - iii. Send \mathcal{A} the response $(\alpha_{j+d'}^{(i)}, p_i(\alpha_{j+d'}^{(i)}))$.
 - If $q_i > d' + 1$, then the polynomial p_i is already defined. Send the response $(\alpha_{j+d'}^{(i)}, p_i(\alpha_{j+d'}^{(i)}))$.
 - b. When \mathcal{A} outputs a vector $\mathbf{a} \in [0, 1]^{d'}$, round its entries to 0, 1 and call it $\bar{\mathbf{a}} \in \{0, 1\}^{d'}$.
 - c. Return $\text{Trace}'(\bar{\mathbf{a}})$.

We emphasize that although algorithm \mathcal{A} can make attribute queries to more than c rows, the adversary \mathcal{B} never defines a secret sharing polynomial for more than $t \leq c$ rows of the input C'_S .

► **Lemma 28.** *Suppose \mathcal{A} is $(1/3, 1/3)$ -accurate for $\mathcal{P}_{SS, d'}(C')$. Then the rounded output $\bar{\mathbf{a}}$ from algorithm \mathcal{A} is a feasible codeword for both C' and C'_S with probability at least $1 - \frac{1}{3} - \frac{1}{100}$. In other words, with probability at least $1 - \frac{1}{3} - \frac{1}{100}$, $\bar{\mathbf{a}} \in F(C')$ and $\bar{\mathbf{a}} \in F(C'_S)$.*

Proof. Assuming that \mathcal{A} is $(1/3, 1/3)$ -accurate for $\mathcal{P}_{SS, d'}(C')$, we first show that $\bar{\mathbf{a}}$ is a feasible codeword for C' with probability at least $2/3$. By the accuracy guarantee of \mathcal{A} , we know that for any column i_j with probability at least $2/3$, $|\mathbf{a}_{i_j} - a_{i_j}| \leq 1/3$ where a_{i_j} is the actual one-way marginal for column i_j . Thus for any column i_j of all 1's in C' , $\mathbf{a}_{i_j} \geq 2/3$ which means $\bar{\mathbf{a}}_{i_j} = 1$, thus satisfying the marking condition. A similar argument holds for the case when a column is all 0's.

Next, using the fact that we use a PR-PPC $(n, d, c, 100d)$ -FPC and that $\bar{\mathbf{a}} \in F(C')$ with probability at least $2/3$, we can invoke Lemma 12 which states that the feasible sample property is satisfied by our PR-PPC FPC construction. Note that in our case, the sampling algorithm described in Definition 11 is \mathcal{A} together with the postprocessing step of rounding the output of \mathcal{A} . Also, even though \mathcal{A} takes the entire dataset as input, it effectively only has access to the rows of the underlying sample via queries to \mathcal{B} and thus satisfies the properties required in Definition 11. In particular, recall that the adversary maintains the invariant $t \leq c$. Lemma 12 states that with probability $\leq \frac{1}{100}$, $\bar{\mathbf{a}}$ is not a feasible codeword for C'_S . By a union bound we have that $1 - \frac{1}{3} - \frac{1}{100}$, $\bar{\mathbf{a}}$ must be a feasible codeword for C'_S . ◀

Proof of Theorem 27. From the above Lemma 28, we have that \mathcal{A} is $(1/3, 1/3)$ -accurate for $\mathcal{P}_{SS, d'}(C')$ implies that $\bar{\mathbf{a}}$ is a feasible codeword for C'_S . By the security of the underlying (n, d, c) -fingerprinting code and the corresponding security guarantee of the PR-PPC $(n, d, c, 100d)$ -FPC given by Corollary 10, we have that $\Pr[\bar{\mathbf{a}} \in F(C'_S) \wedge \text{Trace}'(\bar{\mathbf{a}}) = \perp] \leq \xi$. Since \mathcal{B} releases the output of $\text{Trace}'(\bar{\mathbf{a}})$, the event $\mathcal{B}^{\mathcal{A}}(C'_S) = \perp$ is identical to $\text{Trace}'(\bar{\mathbf{a}}) = \perp$. Thus Property 1 of the theorem statement which states that the probability that \mathcal{B} outputs \perp is bounded by ξ follows. Property 2 follows directly from the soundness property of the fingerprinting code. \blacktriangleleft

► **Corollary 29.** *\mathcal{A} must make at least $c \cdot d'$ attribute queries to $C_{\mathcal{B}}$ to obtain c rows of C'_S where $d' = 101d$.*

Proof. Recall that Theorem 23 states that $SS_{d'}$ is $(d', d', 0)$ -secure where $d' = 101d$. Thus, in order to obtain each row of C'_S , \mathcal{A} must make at least d' cell queries. The statement follows from the fact that \mathcal{A} queries for c rows in total. \blacktriangleleft

► **Lemma 30.** *Suppose there exists a distribution on $c \leq n$ row databases according to Theorem 27. Then there is no (ε, δ) -DP algorithm \mathcal{A} with row query complexity c that is $(1/3, 1/3)$ -accurate for \mathcal{P}_{SS} with $\varepsilon = O(1)$ and $\delta = o(1/c)$.*

Proof. Suppose C'_S is sampled from the distribution on c -row databases \mathcal{C}_S and \mathcal{B} is the adversary from Theorem 27. From the lemma statement we know that \mathcal{A} is $(1/3, 1/3)$ -accurate, thus using Lemma 28 and Theorem 27, we have that $\Pr[\pi_R(\mathcal{B}^{\mathcal{A}}(C'_S)) \in [c]] \geq 1 - \frac{1}{3} - \frac{1}{100} - \xi \geq \Omega(1)$. By an averaging argument, this means that there exists some $i^* \in [c]$ for which $\Pr[\pi_R(\mathcal{B}^{\mathcal{A}}(C'_S)) = i^*] \geq \Omega(1/c)$. However, if $\xi = o(1/c)$ by Property 2 in Theorem 27 we have that $\Pr[\pi_R(\mathcal{B}^{\mathcal{A}}(C'_{S_{-i^*}})) = i^*] \leq \xi = o(1/c)$.

In other words, the probability of $\mathcal{B}^{\mathcal{A}}$ outputting a fixed output i^* on neighboring input databases C'_S and $C'_{S_{-i^*}}$ is different which violates (ε, δ) -DP for any $\varepsilon = O(1)$ and $\delta = o(1/c)$. We note here that since \mathcal{A} can make at most c row queries, the DP guarantee for \mathcal{A} must hold for any neighboring sample of c rows. Since \mathcal{B} does some postprocessing of the output from \mathcal{A} , and we have shown that \mathcal{B} cannot be (ε, δ) -DP, this implies that \mathcal{A} cannot be (ε, δ) -DP for any $\varepsilon = O(1)$ and $\delta = o(1/c)$. \blacktriangleleft

Proof of Theorem 26 . Recall that Lemma 30 states that if there exists a distribution \mathcal{C}_S on $c \leq n$ row databases, then there is no (ε, δ) -DP algorithm \mathcal{A} that is $(1/3, 1/3)$ -accurate for \mathcal{P}_{SS} with $\varepsilon = O(1)$ and $\delta = o(1/c)$. From Theorem 27, such a distribution can be constructed from an (n, d, c) -fingerprinting code. Finally, invoking the Tardos construction for fingerprinting codes in Theorem 7, we get that the row query complexity must be $c = \Omega(\sqrt{d/\log(d)})$. Using Corollary 29, we know that the cell query complexity must be at least $c \cdot d' \geq \Omega(d\sqrt{d/\log(d)})$ where $d' = 101d$. \blacktriangleleft

References

- 1 Borja Balle, Gilles Barthe, and Marco Gaboardi. Privacy amplification by subsampling: Tight analyses via couplings and divergences. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018*, pages 6280–6290, 2018. URL: <https://proceedings.neurips.cc/paper/2018/hash/3b5020bb891119b9f5130f1fea9bd773-Abstract.html>.
- 2 Raef Bassily, Adam D. Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014*, pages 464–473. IEEE Computer Society, 2014. doi: 10.1109/FOCS.2014.56.

- 3 Jeremiah Blocki, Elena Grigorescu, and Tamalika Mukherjee. Differentially-private sublinear-time clustering. In *2021 IEEE International Symposium on Information Theory (ISIT)*, pages 332–337. IEEE, 2021. doi:10.1109/ISIT45174.2021.9518014.
- 4 Jeremiah Blocki, Elena Grigorescu, and Tamalika Mukherjee. Privately estimating graph parameters in sublinear time. In *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.ICALP.2022.26.
- 5 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. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2023*, volume 275 of *LIPIcs*, pages 59:1–59:24. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.APPROX/RANDOM.2023.59.
- 6 Jeremiah Blocki, Seunghoon Lee, Tamalika Mukherjee, and Samson Zhou. Differentially private L_2 -heavy hitters in the sliding window model. In *The Eleventh International Conference on Learning Representations, ICLR 2023*. OpenReview.net, 2023.
- 7 Mark Bun, Thomas Steinke, and Jonathan R. Ullman. Make up your mind: The price of online queries in differential privacy. *J. Priv. Confidentiality*, 9(1), 2019. doi:10.29012/JPC.655.
- 8 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.
- 9 T Tony Cai, Yichen Wang, and Linjun Zhang. The cost of privacy: Optimal rates of convergence for parameter estimation with differential privacy. *The Annals of Statistics*, 49(5):2825–2850, 2021.
- 10 Itai Dinur, Uri Stemmer, David P. Woodruff, and Samson Zhou. On differential privacy and adaptive data analysis with bounded space. In *Advances in Cryptology – EUROCRYPT 2023 – 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques*, volume 14006 of *Lecture Notes in Computer Science*, pages 35–65. Springer, 2023. doi:10.1007/978-3-031-30620-4_2.
- 11 Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography, Third Theory of Cryptography Conference, TCC, Proceedings*, pages 265–284, 2006. doi:10.1007/11681878_14.
- 12 Cynthia Dwork, Moni Naor, Omer Reingold, Guy N. Rothblum, and Salil P. Vadhan. On the complexity of differentially private data release: efficient algorithms and hardness results. In Michael Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009*, pages 381–390. ACM, 2009. doi:10.1145/1536414.1536467.
- 13 Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014. doi:10.1561/04000000042.
- 14 Cynthia Dwork, Adam D. Smith, Thomas Steinke, Jonathan R. Ullman, and Salil P. Vadhan. Robust traceability from trace amounts. In Venkatesan Guruswami, editor, *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015*, pages 650–669. IEEE Computer Society, 2015. doi:10.1109/FOCS.2015.46.
- 15 Cynthia Dwork, Kunal Talwar, Abhradeep Thakurta, and Li Zhang. Analyze gauss: optimal bounds for privacy-preserving principal component analysis. In David B. Shmoys, editor, *Symposium on Theory of Computing, STOC 2014*, pages 11–20. ACM, 2014. doi:10.1145/2591796.2591883.
- 16 Oded Goldreich, Shari Goldwasser, and Dana Ron. Property testing and its connection to learning and approximation. *Journal of the ACM (JACM)*, 45(4):653–750, 1998. doi:10.1145/285055.285060.

- 17 Moritz Hardt and Guy N. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010*, pages 61–70. IEEE Computer Society, 2010. doi:10.1109/FOCS.2010.85.
- 18 Moritz Hardt and Kunal Talwar. On the geometry of differential privacy. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010*, pages 705–714. ACM, 2010. doi:10.1145/1806689.1806786.
- 19 Palak Jain, Sofya Raskhodnikova, Satchit Sivakumar, and Adam Smith. The price of differential privacy under continual observation. In *International Conference on Machine Learning*, pages 14654–14678. PMLR, 2023. URL: <https://proceedings.mlr.press/v202/jain23b.html>.
- 20 Gautam Kamath, Jerry Li, Vikrant Singhal, and Jonathan R. Ullman. Privately learning high-dimensional distributions. In *Conference on Learning Theory, COLT 2019*, volume 99 of *Proceedings of Machine Learning Research*, pages 1853–1902. PMLR, 2019. URL: <http://proceedings.mlr.press/v99/kamath19a.html>.
- 21 Gautam Kamath, Argyris Mouzakis, and Vikrant Singhal. New lower bounds for private estimation and a generalized fingerprinting lemma. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022*, 2022.
- 22 Naty Peter, Eliad Tsfadia, and Jonathan R. Ullman. Smooth lower bounds for differentially private algorithms via padding-and-permuting fingerprinting codes. In *The Thirty Seventh Annual Conference on Learning Theory*, volume 247 of *Proceedings of Machine Learning Research*, pages 4207–4239. PMLR, 2024. URL: <https://proceedings.mlr.press/v247/peter24a.html>.
- 23 Thomas Steinke and Jonathan R. Ullman. Between pure and approximate differential privacy. *J. Priv. Confidentiality*, 7(2), 2016. doi:10.29012/JPC.V7I2.648.
- 24 Thomas Steinke and Jonathan R. Ullman. Tight lower bounds for differentially private selection. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017*, pages 552–563. IEEE Computer Society, 2017. doi:10.1109/FOCS.2017.57.
- 25 Gábor Tardos. Optimal probabilistic fingerprint codes. *J. ACM*, 55(2):10:1–10:24, 2008. doi:10.1145/1346330.1346335.
- 26 Jonathan R. Ullman. Answering $n^{2+o(1)}$ counting queries with differential privacy is hard. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13*, pages 361–370. ACM, 2013. doi:10.1145/2488608.2488653.
- 27 Jalaj Upadhyay. Sublinear space private algorithms under the sliding window model. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 6363–6372. PMLR, 09–15 June 2019. URL: <http://proceedings.mlr.press/v97/upadhyay19a.html>.