Fingerprinting codes and the price of approximate differential privacy

Citation

Bun, Mark, Jonathan Ullman, and Salil Vadhan. 2014. "Fingerprinting Codes and the Price of Approximate Differential Privacy." In Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC '14), New York, NY, May 31-June 3, 2014: 1-10

Published Version

10.1145/2591796.2591877

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:32186909>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Fingerprinting Codes and the Price of Approximate Differential Privacy*

Mark Bun[†]
Harvard University SEAS
mbun@seas.harvard.edu

Jonathan Ullman[‡]
Harvard University SEAS
jullman@seas.harvard.edu

Salil Vadhan[§]
Harvard University SEAS
salil@seas.harvard.edu

ABSTRACT

We show new lower bounds on the sample complexity of (ϵ, δ) -differentially private algorithms that accurately answer large sets of counting queries. A counting query on a database $D \in (\{0, 1\}^d)^n$ has the form “What fraction of the individual records in the database satisfy the property q ?” We show that in order to answer an arbitrary set $\mathcal{Q}$ of $\gg nd$ counting queries on D to within error $\pm\alpha$ it is necessary that

$$n \geq \tilde{\Omega} \left(\frac{\sqrt{d} \log |\mathcal{Q}|}{\alpha^2 \epsilon} \right).$$

This bound is optimal up to poly-logarithmic factors, as demonstrated by the Private Multiplicative Weights algorithm (Hardt and Rothblum, FOCS’10). It is also the first to show that the sample complexity required for (ϵ, δ) -differential privacy is asymptotically larger than what is required merely for accuracy, which is $O(\log |\mathcal{Q}|/\alpha^2)$. In addition, we show that our lower bound holds for the specific case of k -way marginal queries (where $|\mathcal{Q}| = 2^k \binom{d}{k}$) when α is a constant.

Our results rely on the existence of short *fingerprinting codes* (Boneh and Shaw, CRYPTO’95; Tardos, STOC’03), which we show are closely connected to the sample complexity of differentially private data release. We also give a new method for combining certain types of sample complexity lower bounds into stronger lower bounds.

Categories and Subject Descriptors

F.2 [Theory of Computation]: Analysis of Algorithms and Problem Complexity

*A full version of this paper is available at <http://arxiv.org/abs/1311.3158>

[†]Supported by an NDSEG Fellowship and NSF grant CNS-1237235.

[‡]Supported by NSF grant CNS-1237235.

[§]Supported by NSF grant CNS-1237235, a gift from Google, and a Simons Investigator Award.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

STOC ’14, May 31 - June 03 2014, New York, NY, USA
Copyright 2014 ACM 978-1-4503-2710-7/14/05 ...\$15.00.
<http://dx.doi.org/10.1145/2591796.2591877>

General Terms

Security, Theory

Keywords

differential privacy, fingerprinting codes

1. INTRODUCTION

Consider a database $D \in \mathcal{X}^n$, in which each of the n rows corresponds to an individual’s record, and each record consists of an element of some data universe $\mathcal{X}$ (e.g. $\mathcal{X} = \{0, 1\}^d$, corresponding to d binary attributes per record). The goal of privacy-preserving data analysis is to enable rich statistical analyses on such a database while protecting the privacy of the individuals. It is especially desirable to achieve (ϵ, δ) -*differential privacy* [14, 13], which (for suitable choices of ϵ and δ) guarantees that no individual’s data has a significant influence on the information released about the database. A natural way to measure the tradeoff between these two goals is via *sample complexity*—the minimum number of records n that is sufficient in order to achieve both differential privacy and statistical accuracy.

Some of the most basic statistics are *counting queries*, which are queries of the form “What fraction of individual records in D satisfy some property q ?” In particular, we would like to design an algorithm that takes as input a database D and, for some family of counting queries $\mathcal{Q}$, outputs an approximate answer to each of the queries in $\mathcal{Q}$ that is accurate to within, say, ± 0.01 . Suppose we are given a bound on the number of queries $|\mathcal{Q}|$ and the dimensionality of the database records d , but otherwise allow the family $\mathcal{Q}$ to be arbitrary. What is the sample complexity required to achieve (ϵ, δ) -differential privacy and statistical accuracy for $\mathcal{Q}$?

Of course, if we drop the requirement of privacy, then we could achieve perfect accuracy when D contains any number of records. However, in many interesting settings the database D consists of random samples from some larger population, and an analyst is actually interested in answering the queries on the population. Thus, even without a privacy constraint, D would need to contain enough records to ensure that for every query $q \in \mathcal{Q}$, the answer to q on D is close to the answer to q on the whole population, say within ± 0.01 . To achieve this form of *statistical accuracy*, it is well-known that it is necessary and sufficient for D to contain $\Theta(\log |\mathcal{Q}|)$ samples.¹ In this work we consider whether

¹For a specific family of queries $\mathcal{Q}$, the necessary and suffi-

there is an additional “price of differential privacy” if we require both statistical accuracy and (ϵ, δ) -differential privacy (for, say, $\epsilon = O(1)$, $\delta = o(1/n)$). This benchmark has often been used to evaluate the utility of differentially private algorithms, beginning with the seminal work of Dinur and Nissim [12].

Some of the earliest work in differential privacy [12, 18, 5, 14] gave an algorithm—the so-called *Laplace mechanism*—whose sample complexity is $\tilde{\Theta}(|\mathcal{Q}|^{1/2})$, and thus incurs a large price of differential privacy. Fortunately, a remarkable result of Blum, Ligett, and Roth [6] showed that the dependence on $|\mathcal{Q}|$ can be improved exponentially to $O(d \log |\mathcal{Q}|)$ where d is the dimensionality of the data. Their work was improved on in several important aspects [16, 19, 30, 25, 22, 24]. The current best upper bound on the sample complexity is $O(\sqrt{d} \log |\mathcal{Q}|)$, which is obtained via the private multiplicative weights mechanism of Hardt and Rothblum [25].

These results show that the price of privacy is small for datasets with few attributes, but may be large for high-dimensional datasets. For example, if we simply want to estimate the mean of each of the d attributes without a privacy guarantee, then $\Theta(\log d)$ samples are necessary and sufficient to get statistical accuracy. However, the best known (ϵ, δ) -differentially private algorithm requires $\Omega(\sqrt{d})$ samples—an exponential gap. In the special case of *pure* $(\epsilon, 0)$ -differential privacy, a lower bound of $\Omega(d \log |\mathcal{Q}|)$ is known ([23], using the techniques of [26]). However, for the general case of *approximate* (ϵ, δ) -differential privacy the best known lower bound is $\Omega(\log |\mathcal{Q}|)$ [12]. More generally, there are no known lower bounds that separate the sample complexity of (ϵ, δ) -differential privacy from the sample complexity required for statistical accuracy alone.

In this work we close this gap almost completely, and show that there is indeed a “price of approximate differential privacy” for high-dimensional datasets.

THEOREM 1.1 (INFORMAL). *Any algorithm that takes as input a database $D \in \{0, 1\}^d$, satisfies approximate differential privacy, and estimates the mean of each of the d attributes to within error $\pm 1/3$ requires $n \geq \tilde{\Omega}(\sqrt{d})$ samples.*

We establish this lower bound using a combinatorial object called a *fingerprinting code*, introduced by Boneh and Shaw [9] for the problem of watermarking copyrighted content. The use of “secure content distribution schemes” to prove lower bounds for differential privacy originates with the work of Dwork et al. [16], who used cryptographic “traitor-tracing schemes” to prove computational hardness results for differential privacy. Extending this connection, Ullman [33] used fingerprinting codes to construct a novel traitor-tracing scheme and obtain a strong computational hardness result for differential privacy.² Here we show that a *direct* use of fingerprinting codes yields information-theoretic lower bounds (namely on sample complexity).

We then give a *composition theorem* that allows us to combine our new lower bound of $\tilde{\Omega}(\sqrt{d})$ with (variants of) known lower bounds to obtain nearly-optimal sample complexity lower bounds for certain families of queries.

cient number of samples is proportional to the *VC-dimension* of $\mathcal{Q}$, which can be as large as $\log |\mathcal{Q}|$.

²In fact, one way to prove Theorem 1.1 is to replace the one-way functions in [33] with a random oracle, and thereby obtain an information-theoretically secure traitor-tracing scheme.

In addition to its dependence on d and $|\mathcal{Q}|$, we can consider how the sample complexity changes if we want to answer counting queries accurately to within $\pm \alpha$. As above, if we assume the database contains samples from a population, and require only that the answers to queries on the sampled database and the population are close, to within $\pm \alpha$, then $\Theta(\log |\mathcal{Q}|/\alpha^2)$ samples are necessary and sufficient for just statistical accuracy. When $|\mathcal{Q}|$ is large (relative to d and $1/\alpha$), the best sample complexity is again achieved by the private multiplicative weights algorithm, and is $O(\sqrt{d} \log |\mathcal{Q}|/\alpha^2)$. On the other hand, the best known lower bound is $\Omega(\max\{\log |\mathcal{Q}|/\alpha, 1/\alpha^2\})$, which follows from the techniques of [12]. Using our composition theorem, as well as our new lower bound, we are able to obtain a nearly-optimal sample complexity lower bound in terms of all these parameters. The result shows that the private multiplicative weights algorithm achieves nearly-optimal sample-complexity as a function of $|\mathcal{Q}|$, d , and α .

THEOREM 1.2 (INFORMAL). *For every sufficiently small α and $s \geq d/\alpha^2$, there exists a family of queries $\mathcal{Q}$ of size s such that any algorithm that takes as input a database $D \in \{0, 1\}^d$, satisfies approximate differential privacy, and outputs an approximate answer to each query in $\mathcal{Q}$ to within $\pm \alpha$ requires $n \geq \tilde{\Omega}(\sqrt{d} \log |\mathcal{Q}|/\alpha^2)$.*

The previous theorem holds for a worst-case set of queries, but the sample complexity can be smaller for certain interesting families of queries. One family of queries that has received considerable attention is k -way marginals (see e.g. [1, 27, 21, 32, 10, 17]). A k -way marginal query on a database $D \in \{0, 1\}^d$ is specified by a set $S \subseteq [d]$, $|S| \leq k$, and a pattern $t \in \{0, 1\}^{|S|}$ and asks “What fraction of records in D has each attribute j in S set to t_j ?” The number of k -way marginal queries on $\{0, 1\}^d$ is about $2^k \binom{d}{k}$. For the special case of $k = 1$, the queries simply ask for the mean of each attribute, which was discussed above. We prove that our lower bound holds for the special case of k -way conjunction queries when α is a constant. The best previous sample complexity lower bound for constant α is $\Omega(\log |\mathcal{Q}|)$, which again follows from the techniques of [12].

THEOREM 1.3 (INFORMAL). *Any algorithm that takes a database $D \in \{0, 1\}^d$, satisfies approximate differential privacy, and outputs an approximate answer to each of the k -way marginal queries to within $\pm \alpha_0$, for a universal constant α_0 , requires $n \geq \tilde{\Omega}(k\sqrt{d})$.*

1.1 Our Techniques

We now describe the main technical ingredients used to prove these results. For concreteness, we will describe the main ideas for the case of k -way conjunction queries.

Fingerprinting Codes.

Fingerprinting codes, introduced by Boneh and Shaw [9], were originally designed to address the problem of watermarking copyrighted content. Roughly speaking, a (fully-collusion-resilient) fingerprinting code is a way of generating codewords for n users in such a way that any codeword can be uniquely traced back to a user. Each legitimate copy of a piece of digital content has such a codeword hidden in it, and thus any illegal copy can be traced back to the user who copied it. Moreover, even if an arbitrary subset of the users collude to produce a copy of the content, then under

a certain *marking assumption*, the codeword appearing in the copy can still be traced back to one of the users who contributed to it. The standard marking assumption is that if every colluder has the same bit b in the j -th bit of their codeword, then the j -th bit of the “combined” codeword in the copy they produce must be also b . We refer the reader to the original paper of Boneh and Shaw [9] for the motivation behind the marking assumption and an explanation of how fingerprinting codes can be used to watermark digital content.

We show that the existence of short fingerprinting codes implies sample complexity lower bounds for 1-way marginal queries. Recall that a 1-way marginal query q_j is specified by an integer $j \in [d]$ and asks simply “What fraction of records in D have a 1 in the j -th bit?” Suppose a coalition of users takes their codewords and builds a database $D \in (\{0, 1\}^d)^n$ where each record contains one of their codewords, and d is the length of the codewords. Consider the 1-way conjunction query $q_j(D)$. If every user in S has a bit b in the j -th bit of their codeword, then $q_j(D) = b$. Thus, if an algorithm answers 1-way conjunction queries on D with non-trivial accuracy, its output can be used to obtain a combined codeword that satisfies the marking assumption. By the tracing property of fingerprinting codes, we can use the combined codeword to identify one of the users in the database. However, if we can identify one of the users from the answers, then the algorithm cannot be differentially private.

This argument can be formalized to show that if there is a fingerprinting code for n users with codewords of length d , then the sample complexity of answering 1-way conjunctions must be at least n . The nearly-optimal construction of fingerprinting codes due to Tardos [31], gives fingerprinting codes with codewords of length $d = \tilde{O}(n^2)$, which implies a lower bound of $n \geq \tilde{\Omega}(\sqrt{d})$ on the sample complexity required to answer 1-way conjunction queries.

Composition of Sample Complexity Lower Bounds.

Given our lower bound of $\tilde{\Omega}(\sqrt{d})$ for 1-way conjunctions, and the known lower bound of $\Omega(k)$ for answering k -way conjunctions implicit in [12, 29], a natural approach is to somehow compose the two lower bounds to obtain a nearly-optimal lower bound of $\tilde{\Omega}(k\sqrt{d})$. Our composition technique uses the idea of the $\Omega(k)$ lower bound from [12, 29] to show that if we can answer k -way conjunction queries on a large database D with n rows, then we can obtain the answers to the 1-way conjunction queries on a subdatabase of roughly n/k rows. Our lower bound for 1-way marginals tell us that $n/k = \tilde{\Omega}(\sqrt{d})$, so we deduce $n = \tilde{\Omega}(k\sqrt{d})$.

Actually, this reduction only gives accurate answers to *most* of the 1-way marginals on the subdatabase, so we need an extension of our lower bound for 1-way marginals to differentially private algorithms that are allowed to answer a small fraction of the queries with arbitrarily large error. Proving a sample complexity lower bound for this problem requires a “robust” fingerprinting code whose tracing algorithm can trace codewords that have errors introduced into a small fraction of the bits. We show how to construct such a robust fingerprinting code of length $d = \tilde{O}(n^2)$, and thus obtain the desired lower bound. Fingerprinting codes satisfying a weaker notion of robustness were introduced by Boneh and Naor [8, 7].³

³In the fingerprinting codes of [8, 7] the adversary is allowed

1.2 Related Work

We have mostly focused on the sample complexity as a function of the number of queries, the number of attributes d , and the accuracy parameter α . There have been several works focused on the sample complexity as a function of the specific family $\mathcal{Q}$ of queries. For $(\epsilon, 0)$ -differential privacy, Hardt and Talwar [26] showed how to approximately characterize the sample complexity of a family $\mathcal{Q}$ when the accuracy parameter α is sufficiently small. Nikolov, Talwar, and Zhang [28] extended their results to give an approximate characterization for (ϵ, δ) -differential privacy and for the full range of accuracy parameters. Specifically, [28] give an (ϵ, δ) -differentially private algorithm that answers any family of queries $\mathcal{Q}$ on $\{0, 1\}^d$ with error α using a number of samples that is optimal up to a factor of $\text{poly}(d, \log |\mathcal{Q}|)$ that is independent of α . Thus, their algorithm has sample complexity that depends optimally on α . However, their characterization may be loose by a factor of $\text{poly}(d, \log |\mathcal{Q}|)$. In fact, when α is a constant, the lower bound on the sample complexity given by their characterization is always $O(1)$, whereas their algorithm requires $\text{poly}(d, \log |\mathcal{Q}|)$ samples to give non-trivially accurate answers. In contrast, our lower bounds are tight to within $\text{poly}(\log d, \log \log |\mathcal{Q}|, \log(1/\alpha))$ factors, and thus give meaningful lower bounds even when α is constant, but apply only to certain families of queries.

For the particular family of k -way conjunction queries, there have been attempts to prove optimal sample complexity lower bounds. In particular, when k is a constant, Kasiviswanathan et al. [27] give a lower bound of $\min\{|\mathcal{Q}|^{1/2}/\alpha, 1/\alpha^2\}$ on the sample complexity. Their result was improved by De [11], who proved that the same sample complexity lower bound applies even for algorithms that can introduce arbitrarily large error on a constant fraction of the queries. In the regime we consider where α is constant, these lower bounds are $O(1)$.

There have also been attempts to explicitly and precisely determine the sample complexity of even simpler query families than k -way conjunctions, such as point functions and interval functions [2, 3, 4]. These works show that these families can have sample complexity lower than $\tilde{O}(\sqrt{d} \log |\mathcal{Q}|/\alpha^2)$.

2. PRELIMINARIES

2.1 Differential Privacy

We define a *database* $D \in \mathcal{X}^n$ to be an ordered tuple of n rows $(x_1, \dots, x_n) \in \mathcal{X}$ chosen from a *data universe* $\mathcal{X}$. We say that two databases $D, D' \in \mathcal{X}^n$ are *adjacent* if they differ only by a single row, and we denote this by $D \sim D'$. In particular, we can replace the i th row of a database D with some fixed element of $\mathcal{X}$ to obtain another database $D_{-i} \sim D$.

DEFINITION 2.1 (DIFFERENTIAL PRIVACY [14]). *Let $\mathcal{A} : \mathcal{X}^n \rightarrow \mathcal{R}$ be a randomized algorithm (where n is a varying parameter). $\mathcal{A}$ is (ϵ, δ) -differentially private if for every two adjacent databases $D \sim D'$ and every subset $S \subseteq \mathcal{R}$,*

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

2.2 Counting Queries and Accuracy

to erase a large fraction of the coordinates of the combined codeword, and must reveal which coordinates are erased.

In this paper we study algorithms that answer *counting queries*. A counting query on $\mathcal{X}$ is defined by a predicate $q : \mathcal{X} \rightarrow \{0, 1\}$. Abusing notation, we define the evaluation of the query q on a database $D = (x_1, \dots, x_n) \in \mathcal{X}^n$ to be its average value over the rows,

$$q(D) = \frac{1}{n} \sum_{i=1}^n q(x_i).$$

DEFINITION 2.2 (ACCURACY FOR COUNTING QUERIES). Let $\mathcal{Q}$ be a set of counting queries on $\mathcal{X}$ and $\alpha, \beta \in [0, 1]$ be parameters. For a database $D \in \mathcal{X}^n$, a sequence of answers $a = (a_q)_{q \in \mathcal{Q}} \in \mathbb{R}^{|\mathcal{Q}|}$ is (α, β) -accurate for $\mathcal{Q}$ if $|q(D) - a_q| \leq \alpha$ for at least a $1 - \beta$ fraction of queries $q \in \mathcal{Q}$.

Let $\mathcal{A} : \mathcal{X}^n \rightarrow \mathbb{R}^{|\mathcal{Q}|}$ be a randomized algorithm. $\mathcal{A}$ is (α, β) -accurate for $\mathcal{Q}$ if for every $D \in \mathcal{X}^n$,

$$\Pr[\mathcal{A}(D) \text{ is } (\alpha, \beta)\text{-accurate for } \mathcal{Q}] \geq 2/3.$$

When $\beta = 0$ we may simply write that $\mathcal{A}$ is α -accurate for $\mathcal{Q}$.

In the definition of accuracy, we have assumed that $\mathcal{A}$ outputs a sequence of $|\mathcal{Q}|$ real-valued answers, with a_q representing the answer to q . Since we are not concerned with the running time of the algorithm, this assumption is without loss of generality.

An important example of a collection of counting queries is the set of *k-way marginals*. For all of our results it will be sufficient to consider only the set of *monotone k-way marginals*.

DEFINITION 2.3 (MONOTONE k-WAY MARGINALS). A (monotone) *k*-way marginal q_S over $\{0, 1\}^d$ is specified by a subset $S \subseteq [d]$ of size $|S| \leq k$. It takes the value $q_S(x) = 1$ if and only if $x_i = 1$ for every index $i \in S$. The collection of all (monotone) *k*-way marginals is denoted by $\mathcal{M}_{k,d}$.

2.3 Sample Complexity

In this work we prove lower bounds on the sample complexity required to simultaneously achieve differential privacy and accuracy.

DEFINITION 2.4 (SAMPLE COMPLEXITY). Let $\mathcal{Q}$ be a set of counting queries on $\mathcal{X}$ and let $\alpha, \beta > 0$ be parameters, and let ε, δ be functions of n . We say that $(\mathcal{Q}, \mathcal{X})$ has sample complexity n^* for (α, β) -accuracy and (ε, δ) -differential privacy if n^* is the least $n \in \mathbb{N}$ such that there exists an (ε, δ) -differentially private algorithm $\mathcal{A} : \mathcal{X}^n \rightarrow \mathbb{R}^{|\mathcal{Q}|}$ that is (α, β) -accurate for $\mathcal{Q}$.

We will focus on the case where $\varepsilon = O(1)$ and $\delta = o(1/n)$. This setting of the parameters is essentially the most-permissive for which (ε, δ) -differential privacy is still a meaningful privacy definition. However, pinning down the exact dependence on ε and δ is still of interest. Regarding ε , this can be done via the following standard lemma, which allows us to take $\varepsilon = 1$ without loss of generality.

LEMMA 2.5. For every set of counting queries $\mathcal{Q}$, universe $\mathcal{X}$, $\alpha, \beta \in [0, 1], \varepsilon \leq 1$. $(\mathcal{Q}, \mathcal{X})$ has sample complexity n^* for (α, β) -accuracy and $(1, o(1/n))$ -differential privacy if and only if it has sample complexity $\Theta(n^*/\varepsilon)$ for (α, β) -accuracy and $(\varepsilon, o(1/n))$ -differential privacy.

2.4 Re-identifiable Distributions

All of our eventual lower bounds will take the form a “re-identification” attack, in which we possess data from a large number of individuals, and identify one such individual who was included in the database. In this attack, we choose a distribution on databases and give an adversary 1) a database D drawn from that distribution and 2) either $\mathcal{A}(D)$ or $\mathcal{A}(D_{-i})$ for some row i , where $\mathcal{A}$ is an alleged sanitizer. The adversary’s goal is to identify a row of D that was given to the sanitizer. We say that the distribution is re-identifiable if there is an adversary who can identify such a row with sufficiently high confidence whenever $\mathcal{A}$ outputs accurate answers. If the adversary can do so, it means that there must be a pair of adjacent databases $D \sim D_{-i}$ such that the adversary can distinguish $\mathcal{A}(D)$ from $\mathcal{A}(D_{-i})$, which means $\mathcal{A}$ cannot be differentially private.

DEFINITION 2.6 (RE-IDENTIFIABLE DISTRIBUTION). For a data universe $\mathcal{X}$ and $n \in \mathbb{N}$, let $\mathcal{D}$ be a distribution on n -row databases $D \in \mathcal{X}^n$. Let $\mathcal{Q}$ be a family of counting queries on $\mathcal{X}$ and let $\gamma, \xi, \alpha, \beta \in [0, 1]$ be parameters. The distribution $\mathcal{D}$ is (γ, ξ) -re-identifiable from (α, β) -accurate answers to $\mathcal{Q}$ if there exists a (possibly randomized) adversary $\mathcal{B} : \mathcal{X}^n \times \mathbb{R}^{|\mathcal{Q}|} \rightarrow [n] \cup \{\perp\}$ such that for every randomized algorithm $\mathcal{A} : \mathcal{X}^n \rightarrow \mathbb{R}^{|\mathcal{Q}|}$, the following both hold:

1. $\Pr_{D \leftarrow \mathcal{D}} \left[\bigwedge_{\mathcal{A}(D) \text{ is } (\alpha, \beta)\text{-accurate for } \mathcal{Q}} (\mathcal{B}(D, \mathcal{A}(D)) = \perp) \right] \leq \gamma$.
2. For every $i \in [n]$, $\Pr_{D \leftarrow \mathcal{D}} [\mathcal{B}(D, \mathcal{A}(D_{-i})) = i] \leq \xi$.

Here the probability is taken over the choice of D and i as well as the coins of $\mathcal{A}$ and $\mathcal{B}$. We allow $\mathcal{D}$ and $\mathcal{B}$ to share a common state.

If $\mathcal{A}$ is an (α, β) -accurate algorithm, then its output $\mathcal{A}(D)$ will be (α, β) -accurate with probability at least $2/3$. Therefore, if $\gamma < 2/3$, we can conclude that $\Pr[\mathcal{B}(D, \mathcal{A}(D)) \in [n]] \geq 1 - \gamma - 1/3 = \Omega(1)$. In particular, there exists some $i^* \in [n]$ for which $\Pr[\mathcal{B}(D, \mathcal{A}(D)) = i^*] \geq \Omega(1/n)$. However, if $\xi = o(1/n)$, then $\Pr[\mathcal{B}(D, \mathcal{A}(D_{-i^*})) = i^*] \leq \xi = o(1/n)$. Thus, for this choice of γ and ξ we will obtain a contradiction to (ε, δ) -differential privacy for any $\varepsilon = O(1)$ and $\delta = o(1/n)$. We remark that this conclusion holds even if $\mathcal{D}$ and $\mathcal{B}$ share a common state.

3. LOWER BOUNDS VIA FINGERPRINTING CODES

In this section we prove that there exists a simple family of d queries that requires $n \geq \tilde{\Omega}(\sqrt{d})$ samples for both accuracy and privacy. Specifically, we prove that for the family of 1-way marginals on d bits, sample complexity $\tilde{\Omega}(\sqrt{d})$ is required to produce differentially private answers that are accurate even just to within $\pm 1/3$. In contrast, without a privacy guarantee, $\Theta(\log d)$ samples from the population are necessary and sufficient to ensure that the answers to these queries on the database and the population are approximately the same. The best previous lower bound for (ε, δ) -differential privacy is also $O(\log d)$, which follows from the techniques of [12, 29].

In Section 3.1 we give the relevant background on fingerprinting codes and in Section 3.2 we prove our lower bounds for 1-way marginals.

3.1 Fingerprinting Codes

Fingerprinting codes were introduced by Boneh and Shaw [9] to address the problem of watermarking digital content. A *fingerprinting code* is a pair of randomized algorithms $(Gen, Trace)$. The code generator Gen outputs a *codebook* $C \in \{0, 1\}^{n \times d}$. Each row c_i of C is the *codeword* of user i . For a subset of users $S \subseteq [n]$, we use $C_S \in \{0, 1\}^{|S| \times d}$ to denote the set of codewords of users in S .

The security property of fingerprinting codes asserts that any codeword can be “traced” to a user $i \in [n]$. Moreover, we require that the fingerprinting code is “fully-collusion-resilient”—even if any “coalition” of users $S \subseteq [n]$ gets together and “combines” their codewords in any way that respects certain constraints known as a *marking assumption*, then the combined codeword can be traced to a user $i \in S$. That is, there is a tracing algorithm $Trace$ that takes the codebook and combined codeword and outputs either a user $i \in [n]$ or $\perp$, and we require that if c' satisfies the constraints, then $Trace(C, c') \in S$ with high probability. Moreover, $Trace$ should accuse an innocent user, i.e. $Trace(C, c') \in [n] \setminus S$, with very low probability. Analogous to the definition of re-identifiable distributions (Definition 2.6), we allow Gen and $Trace$ to share a common state. When designing fingerprinting codes, one tries to make the marking assumption on the combined codeword as weak as possible.

The basic marking assumption is that each bit of the combined word c' must match the corresponding bit for some user in S . In order to prove sample-complexity lower bounds for (α, β) -accuracy with $\beta > 0$, we will need fingerprinting codes that are secure under an even weaker marking assumption. Specifically, that *most* bits of the combined word c' must match the corresponding bit for some user in S . Formally, for any $\beta \in [0, 1]$, we define

$$F_\beta(C_S) = \left\{ c' \in \{0, 1\}^d \mid \Pr_{j \leftarrow_R [d]} [\exists i \in S, c'_j = c_{ij}] \geq 1 - \beta \right\}$$

DEFINITION 3.1 (ROBUST FINGERPRINTING CODES). For any $n, d \in \mathbb{N}$, $\xi, \beta \in [0, 1]$, a pair of algorithms $(Gen, Trace)$ is an (n, d) -fingerprinting code with security ξ robust to a β fraction of errors if Gen outputs a codebook $C \in \{0, 1\}^{n \times d}$ and for every (possibly randomized) adversary $\mathcal{A}_{FP}$, and every coalition $S \subseteq [n]$, if we set $c' \leftarrow_R \mathcal{A}_{FP}(C_S)$, then

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

where the probability is taken over the coins of Gen , $Trace$, and $\mathcal{A}_{FP}$. The algorithms Gen and $Trace$ may share a common state.

Tardos [31] gave a construction of standard, non-robust fingerprinting codes for a nearly optimal number of users $\tilde{\Omega}(\sqrt{d/\log(1/\xi)})$. In the full version of this work, we show how to extend Tardos’ construction and analysis to yield error-robust fingerprinting codes with a nearly-optimal number of users that are tolerant to a constant fraction of errors.

THEOREM 3.2. For every $d \in \mathbb{N}$, and $\xi \in (0, 1]$, there exists an (n, d) -fingerprinting code with security ξ robust to a $1/75$ fraction of errors for $n = n(d, \xi) = \tilde{\Omega}(\sqrt{d/\log(1/\xi)})$.

Boneh and Naor [8] introduced a different notion of fingerprinting codes robust to adversarial “erasures”. In their

definition, the adversary is allowed to output a string in $\{0, 1, ?\}^d$, and in order to trace they require that the fraction of ? symbols is bounded away from 1 and that any non-? symbols respect the basic feasibility constraint. For this definition, constructions with nearly-optimal length $d = \tilde{O}(n^2)$, robust to a $1 - o(1)$ fraction of erasures are known [7]. In contrast, our codes are robust to adversarial “errors.” Robustness to a β fraction of errors can be seen to imply robustness to nearly a 2β fraction of erasures but the converse is false. Thus for corresponding levels of robustness our definition is strictly more stringent. Unfortunately we don’t currently know how to design a code tolerant to a $1/2 - o(1)$ fraction of errors, so our Theorem 3.2 does not subsume prior results on robust fingerprinting codes.

3.2 Lower Bounds for 1-Way Marginals

We are now ready to state and prove the main result of this section, namely that there is a distribution on databases $D \in (\{0, 1\}^d)^n$, for $n = \tilde{\Omega}(\sqrt{d})$, that is re-identifiable from accurate answers to 1-way marginals.

THEOREM 3.3. For every $n, d \in \mathbb{N}$, and $\xi \in [0, 1]$ if there exists an (n, d) -fingerprinting code with security ξ , robust to a β fraction of errors, then there exists a distribution on n -row databases $D \in (\{0, 1\}^d)^n$ that is (ξ, ξ) -re-identifiable from $(1/3, \beta)$ -accurate answers to $\mathcal{M}_{1,d}$.

In particular, if $\xi = o(1/n)$, then there is no algorithm $\mathcal{A} : (\{0, 1\}^d)^n \rightarrow \mathbb{R}^{|\mathcal{M}_{1,d}|}$ that is $(O(1), o(1/n))$ -differentially private and $(1/3, \beta)$ -accurate for $\mathcal{M}_{1,d}$.

By combining Theorem 3.3 with Theorem 3.2 we obtain a sample complexity lower bound for 1-way marginals, and thereby establish Theorem 1.1 in the introduction.

COROLLARY 3.4. For every $d \in \mathbb{N}$, the family of 1-way marginals on $\{0, 1\}^d$ has sample complexity at least $\tilde{\Omega}(\sqrt{d})$ for $(1/3, 1/75)$ -accuracy and $(O(1), o(1/n))$ -differential privacy.

PROOF OF THEOREM 3.3. Let $(Gen, Trace)$ be the promised fingerprinting code. We define the re-identifiable distribution $\mathcal{D}$ to simply be the output distribution of the code generator, Gen . And we define the privacy adversary $\mathcal{B}$ to take the answers $a = \mathcal{A}(D) \in [0, 1]^{|\mathcal{M}_{1,d}|}$, obtain $\bar{a} \in \{0, 1\}^{|\mathcal{M}_{1,d}|}$ by rounding each entry of a to $\{0, 1\}$, run the tracing algorithm $Trace$ on the rounded answers $\bar{a}$, and return its output. The shared state of $\mathcal{D}$ and $\mathcal{B}$ will be the shared state of Gen and $Trace$.

Now we will verify that $\mathcal{D}$ is (ξ, ξ) -re-identifiable. First, suppose that $\mathcal{A}(D)$ outputs answers $a = (a_{q_j})_{j \in [d]}$ that are $(1/3, \beta)$ -accurate for 1-way marginals. That is, there is a set $G \subseteq [d]$ such that $|G| \geq (1 - \beta)d$ and for every $j \in G$, the answer a_{q_j} estimates the fraction of rows having a 1 in column j to within $1/3$. Let $\bar{a}_{q_j}$ be a_{q_j} rounded to the nearest value in $\{0, 1\}$. Let j be a column in G . If column j has all 1’s, then $a_{q_j} \geq 2/3$, and $\bar{a}_{q_j} = 1$. Similarly, if column j has all 0’s, then $a_{q_j} \leq 1/3$, and $\bar{a}_{q_j} = 0$. Therefore, we have

$$a \text{ is } (1/3, \beta)\text{-accurate} \implies \bar{a} \in F_\beta(D). \quad (1)$$

By security of the fingerprinting code (Definition 3.1),

$$\Pr[\bar{a} \in F_\beta(D) \wedge Trace(D, \bar{a}) = \perp] \leq \xi. \quad (2)$$

Combining (1) and (2) implies that

$$\Pr[\mathcal{A}(D) \text{ is } (1/3, \beta)\text{-accurate} \wedge Trace(D, \bar{a}) = \perp] \leq \xi.$$

But the event $\text{Trace}(D, \bar{a}) = \perp$ is exactly the same as $\mathcal{B}(D, \mathcal{A}(D)) = \perp$, and thus we have established the first condition necessary for $\mathcal{D}$ to be (ξ, ξ) -re-identifiable.

The second condition for re-identifiability follows directly from the soundness of the fingerprinting code, which asserts that for every adversary $\mathcal{A}_{FP}$, in particular for $\mathcal{A}$, it holds that $\Pr[\text{Trace}(D, \mathcal{A}_{FP}(D_{-i})) = i] \leq \xi$. $\square$

Although our lower bound is stated for constant accuracy $\alpha = 1/3$, a simple argument, which appears in the full version of this work, shows that a lower bound of n^* for $(1/3, \beta)$ -accuracy implies a lower bound of $\Omega(n^*/\alpha)$ for (α, β) -accuracy. Thereby we obtain the following corollary, which is nearly optimal in both d and α .

COROLLARY 3.5. *For every $d \in \mathbb{N}$ and $0 < \alpha \leq 1/3$, the family of 1-way marginals on $\{0, 1\}^d$ has sample complexity at least $\tilde{\Omega}(\sqrt{d}/\alpha)$ for $(\alpha, 1/75)$ -accuracy and $(O(1), o(1/n))$ -differential privacy.*

REMARK 3.6. *As pointed out to us by Adam Smith, our connection can also be combined with standard algorithmic results in differential privacy (namely, the Gaussian mechanism) to give a simpler proof of Tardos' tight lower bound on the length of fingerprinting codes [31]. See the full version of our paper for details.*

4. A COMPOSITION THEOREM FOR SAMPLE COMPLEXITY

In this section we state and prove a composition theorem for sample complexity lower bounds. At a high-level the composition theorem starts with two pairs, $(\mathcal{Q}, \mathcal{X})$ and $(\mathcal{Q}', \mathcal{X}')$, for which we know sample-complexity lower bounds of n and n' respectively, and attempts to prove a sample-complexity lower bound of $n \cdot n'$ for a related family of queries on a related data universe.

Specifically, our sample-complexity lower bound will apply to the “product” of $\mathcal{Q}$ and $\mathcal{Q}'$, defined on $\mathcal{X} \times \mathcal{X}'$. We define the product $\mathcal{Q} \wedge \mathcal{Q}'$ to be

$$\mathcal{Q} \wedge \mathcal{Q}' = \{q \wedge q' : (x, x') \mapsto q(x) \wedge q'(x') \mid q \in \mathcal{Q}, q' \in \mathcal{Q}'\}.$$

Since q, q' are boolean-valued, their conjunction can also be written $q(x)q'(x')$.

We now begin to describe how we can prove a sample complexity lower bound for $\mathcal{Q} \wedge \mathcal{Q}'$. First, we describe a certain product operation on databases. Let $D \in \mathcal{X}^n$, $D = (x_1, \dots, x_n)$, be a database. Let $D'_1, \dots, D'_n \in (\mathcal{X}')^{n'}$ where $D'_i = (x'_{i1}, \dots, x'_{in'})$ be n databases. We define the product database $D^* = D \times (D'_1, \dots, D'_n) \in (\mathcal{X} \times \mathcal{X}')^{n \cdot n'}$ as follows: For every $i = 1, \dots, n, j = 1, \dots, n'$, let the (i, j) -th row of D^* be $x^*_{(i,j)} = (x_i, x'_{ij})$. Note that we index the rows of D^* by (i, j) . We will sometimes refer to $D'_1, \dots, D'_n$ as the “subdatabases” of D^* .

The key property of these databases is that we can use a query $q \wedge q' \in \mathcal{Q} \wedge \mathcal{Q}'$ to compute a “subset-sum” of the vector $s_{q'} = (q'(D'_1), \dots, q'(D'_n))$ consisting of the answers to q' on each of the n subdatabases. That is, for every $q \in \mathcal{Q}$ and $q' \in \mathcal{Q}'$,

$$(q \wedge q')(D^*) = \frac{1}{nn'} \sum_{i=1}^n \sum_{j=1}^{n'} (q \wedge q')(x^*_{(i,j)}) = \frac{1}{n} \sum_{i=1}^n q(x_i) q'(D'_i). \quad (3)$$

Thus, every approximate answer $a_{q \wedge q'}$ to a query $q \wedge q'$ places a subset-sum constraint on the vector $s_{q'}$. (Namely, $a_{q \wedge q'} \approx \frac{1}{n} \sum_{i=1}^n q(x_i) q'(D'_i)$) If the database D and family $\mathcal{Q}$ are chosen appropriately, and the answers are sufficiently accurate, then we will be able to reconstruct a good approximation to $s_{q'}$. Indeed, this sort of “reconstruction attack” is the core of many lower bounds for differential privacy, starting with the work of Dinur and Nissim [12]. The setting they consider is essentially the special case of what we have just described where $D'_1, \dots, D'_n$ are each just a single bit ($\mathcal{X}' = \{0, 1\}$, and $\mathcal{Q}'$ contains only the identity query). In Section 5 we will discuss choices of D and $\mathcal{Q}$ that allow for this reconstruction.

We now state the formal notion of reconstruction attack that we want D and $\mathcal{Q}$ to satisfy.

DEFINITION 4.1 (RECONSTRUCTION ATTACKS). *Let $\mathcal{Q}$ be a family of counting queries over a data universe $\mathcal{X}$. Let $n \in \mathbb{N}$ and $\alpha', \alpha, \beta \in [0, 1]$ be parameters. Let $D = (x_1, \dots, x_n) \in \mathcal{X}^n$ be a database. Suppose there is an adversary $\mathcal{B}_D : \mathbb{R}^{|\mathcal{Q}|} \rightarrow [0, 1]^n$ with the following property: For every vector $s \in [0, 1]^n$ and every sequence $a = (a_q)_{q \in \mathcal{Q}} \in \mathbb{R}^{|\mathcal{Q}|}$ such that*

$$\left| a_q - \frac{1}{n} \sum_{i=1}^n q(x_i) s_i \right| < \alpha$$

for at least a $1 - \beta$ fraction of queries $q \in \mathcal{Q}$, $\mathcal{B}_D(a)$ outputs a vector $t \in [0, 1]^n$ such that

$$\frac{1}{n} \sum_{i=1}^n |t_i - s_i| \leq \alpha'.$$

Then we say that $D \in \mathcal{X}^n$ admits an α' -reconstruction attack from (α, β) -accurate answers to $\mathcal{Q}$.

A reconstruction attack itself implies a sample-complexity lower bound, as in [12]. However, we show how to obtain stronger sample complexity lower bounds from the reconstruction attack by applying it to a product database D^* to obtain accurate answers to queries on its subdatabases. For each query $q' \in \mathcal{Q}'$, we run the adversary promised by the reconstruction attack on the approximate answers given to queries of the form $(q \wedge q') \in \mathcal{Q} \wedge \mathcal{Q}'$. As discussed above, answers to these queries will approximate subset sums of the vector $s_{q'} = (q'(D'_1), \dots, q'(D'_n))$. When the reconstruction attack is given these approximate answers, it returns a vector $t_{q'} = (t_{q',1}, \dots, t_{q',n})$ such that $t_{q',i} \approx s_{q',i} = q'(D'_i)$ on average over i . Running the reconstruction attack for every query q' gives us a collection $t = (t_{q',i})_{q' \in \mathcal{Q}', i \in [n]}$ where $t_{q',i} \approx q'(D'_i)$ on average over both q' and i . By an application of Markov's inequality, for most of the subdatabases D'_i , we have that $t_{q',i} \approx q'(D'_i)$ on average over the choice of $q' \in \mathcal{Q}'$. For each i such that this guarantee holds, another application of Markov's inequality shows that for most queries $q' \in \mathcal{Q}'$ we have $t_{q',i} \approx q'(D'_i)$, which is our definition of (α, β) -accuracy (later enabling us to apply a re-identification adversary for $\mathcal{Q}'$).

The algorithm we have described for obtaining accurate answers on the subdatabases is formalized in Figure 1.

We are now in a position to state the main lemma that enables our composition technique. The lemma says that if we are given accurate answers to $\mathcal{Q} \wedge \mathcal{Q}'$ on D^* and the database $D \in \mathcal{X}^n$ admits a reconstruction attack from ac-

Let $a = (a_{q \wedge q'})_{q \in \mathcal{Q}, q' \in \mathcal{Q}'}$ be an answer vector.
 Let $\mathcal{B}_D : \mathbb{R}^{|\mathcal{Q}|} \rightarrow [0, 1]^n$ be a reconstruction attack.
 For each $q' \in \mathcal{Q}'$
 Let $(t_{q',1}, \dots, t_{q',n}) = \mathcal{B}_D((a_{q \wedge q'})_{q \in \mathcal{Q}})$
 Output $(t_{q',i})_{q' \in \mathcal{Q}', i \in [n]}$.

Figure 1: The reconstruction $\mathcal{R}_D^*(a)$.

curate answers to $\mathcal{Q}$, then we can obtain accurate answers to $\mathcal{Q}'$ on the most of the subdatabases $D'_1, \dots, D'_n \in (\mathcal{X}')^{n'}$.

LEMMA 4.2. *Let $D \in \mathcal{X}^n$ and $D'_1, \dots, D'_n \in (\mathcal{X}')^{n'}$ be databases and $D^* \in (\mathcal{X} \times \mathcal{X}')^{n \cdot n'}$ be as above. Let $a = (a_{q \wedge q'})_{q \in \mathcal{Q}, q' \in \mathcal{Q}'} \in \mathbb{R}^{|\mathcal{Q} \wedge \mathcal{Q}'|}$. Let $\alpha', \alpha, \beta \in [0, 1]$ be parameters. Suppose that for some parameter $c > 1$, the database D admits an α' -reconstruction attack from $(\alpha, c\beta)$ -accurate answers to $\mathcal{Q}$. Then if $(t_{q',i})_{q' \in \mathcal{Q}', i \in [n]} = \mathcal{R}_D^*(a)$ (Figure 1),*

a is (α, β) -accurate for $\mathcal{Q} \wedge \mathcal{Q}'$ on $D^* \implies$

$$\Pr_{i \leftarrow_R [n]} \left[\begin{array}{c} (t_{q',i})_{q' \in \mathcal{Q}'} \text{ is} \\ (6c\alpha', 2/c)\text{-accurate for } \mathcal{Q}' \text{ on } D_i \end{array} \right] \geq 5/6.$$

We defer the proof to the full version of this work. The proof closely follows what we sketched above, but requires some additional bookkeeping to handle the case where a is only accurate for most queries. In this case the reconstruction attack may fail completely for certain queries $q' \in \mathcal{Q}'$ and we need to account for this additional source of error.

We now explain how the main lemma allows us to prove a composition theorem for sample complexity lower bounds. We start with a query family $\mathcal{Q}$ on a database $D \in \mathcal{X}^n$ that admits a reconstruction attack, and a distribution $\mathcal{D}'$ over databases in $(\mathcal{X}')^{n'}$ that is re-identifiable from answers to a family $\mathcal{Q}'$. We show how to combine these objects to form a re-identifiable distribution $\mathcal{D}^*$ for queries $\mathcal{Q} \wedge \mathcal{Q}'$ over $(\mathcal{X} \times \mathcal{X}')^{n \cdot n'}$, yielding a sample complexity lower bound of $n \cdot n'$.

A sample from $\mathcal{D}^*$ consists of $D^* = D \times (D'_1, \dots, D'_n)$ where each subdatabase D'_i is an independent sample from $\mathcal{D}'$. The main lemma above shows that if there is an algorithm $\mathcal{A}$ that is accurate for $\mathcal{Q} \wedge \mathcal{Q}'$ on D^* , then an adversary can reconstruct accurate answers to $\mathcal{Q}'$ on most of the subdatabases $D'_1, \dots, D'_n$. Since these subdatabases are drawn from a re-identifiable distribution, the adversary can re-identify a member of one of the subdatabases D'_i . Since the identified member of D'_i is also a member of D^* , we will have a re-identification attack against D^* as well.

We are now ready to formalize our composition theorem.

THEOREM 4.3. *Let $\mathcal{Q}$ be a family of counting queries on $\mathcal{X}$, and let $\mathcal{Q}'$ be a family of counting queries on $\mathcal{X}'$. Let $\gamma, \xi, \alpha', \alpha, \beta \in [0, 1]$ be parameters. Assume that for some parameters $c > 1$, $\gamma, \xi, \alpha', \alpha, \beta \in [0, 1]$, the following both hold:*

1. *There exists a database $D \in \mathcal{X}^n$ that admits an α' -reconstruction attack from $(\alpha, c\beta)$ -accurate answers to $\mathcal{Q}$.*
2. *There is a distribution $\mathcal{D}'$ on databases $D \in (\mathcal{X}')^{n'}$ that is (γ, ξ) -re-identifiable from $(6c\alpha', 2/c)$ -accurate answers to $\mathcal{Q}'$.*

Then there is a distribution on databases $D^* \in (\mathcal{X} \times \mathcal{X}')^{n \cdot n'}$ that is $(\gamma + 1/6, \xi)$ -re-identifiable from (α, β) -accurate answers to $\mathcal{Q}$.

Let $D = (x_1, \dots, x_n) \in \mathcal{X}^n$ be a database that admits reconstruction.
 Let $\mathcal{D}'$ on $(\mathcal{X}')^{n'}$ be a re-identifiable distribution.
 For $i = 1, \dots, n$, choose $D'_i \leftarrow_R \mathcal{D}'$ (independently)
 Output $D^* = D \times (D'_1, \dots, D'_n) \in (\mathcal{X} \times \mathcal{X}')^{n \cdot n'}$

Figure 2: The new distribution $\mathcal{D}^*$.

Let $D^* = D \times (D'_1, \dots, D'_n)$.
 Run $\mathcal{R}_D^*(\mathcal{A}(D^*))$ (Figure 1) to reconstruct a set of approximate answers $(t_{q',i})_{q' \in \mathcal{Q}', i \in [n]}$.
 Choose a random $i \leftarrow_R [n]$.
 Output $\mathcal{B}'(D'_i, (t_{q',i})_{q' \in \mathcal{Q}'})$.

Figure 3: The privacy adversary $\mathcal{B}^*(D^*, \mathcal{A}(D^*))$.

PROOF SKETCH. Let $D = (x_1, \dots, x_n) \in \mathcal{X}^n$ be the database that admits a reconstruction attack (Definition 4.1). Let $\mathcal{D}'$ be the promised re-identifiable distribution on databases $D \in (\mathcal{X}')^{n'}$ and $\mathcal{B}' : (\mathcal{X}')^{n'} \times \mathbb{R}^{|\mathcal{Q}'|} \rightarrow [n'] \cup \{\perp\}$ be the promised adversary (Definition 2.6).

In Figure 2, we define a distribution $\mathcal{D}^*$ on databases $D^* \in (\mathcal{X} \times \mathcal{X}')^{n \cdot n'}$. In Figure 3, we define an adversary $\mathcal{B}^* : (\mathcal{X} \times \mathcal{X}')^{n \cdot n'} \times \mathbb{R}^{|\mathcal{Q} \wedge \mathcal{Q}'|}$ for a re-identification attack. The shared state of $\mathcal{D}^*$ and $\mathcal{B}^*$ will be the shared state of $\mathcal{D}'$ and $\mathcal{B}'$. The next two claims show that $\mathcal{D}^*$ satisfies the two properties necessary to be a $(\gamma + 1/6, \xi)$ -re-identifiable distribution (Definition 2.6).

CLAIM 4.4.

$$\Pr_{\substack{D^* \leftarrow_R \mathcal{D}^* \\ \text{coins}(\mathcal{A}), \text{coins}(\mathcal{B}^*)}} \left[\begin{array}{c} (\mathcal{B}^*(D^*, \mathcal{A}(D^*))) = \perp \\ \wedge (\mathcal{A}(D^*) \text{ is } (\alpha, \beta)\text{-accurate for } \mathcal{Q} \wedge \mathcal{Q}') \end{array} \right] \leq \gamma + 1/6.$$

PROOF OF CLAIM 4.4. Assume that $\mathcal{A}(D^*)$ is (α, β) -accurate for $\mathcal{Q} \wedge \mathcal{Q}'$. By the construction of $\mathcal{B}^*$ and by Lemma 4.2 it suffices to prove that

$$\Pr_{\substack{D^* \leftarrow_R \mathcal{D}^* \\ i \leftarrow_R [n]}} \left[\begin{array}{c} (\mathcal{B}'(D'_i, (t_{q',i})_{q' \in \mathcal{Q}'})) = \perp \\ \wedge ((t_{q',i}) \text{ is } (6c\alpha', 2/c)\text{-accurate for } \mathcal{Q}') \end{array} \right] \leq \gamma \quad (4)$$

We prove this inequality by giving a reduction to the re-identifiability of $\mathcal{D}'$. Consider the following sanitizer $\mathcal{A}'$: On input $D' \leftarrow_R \mathcal{D}'$, $\mathcal{A}'$ first chooses a random index $i^* \leftarrow_R [n]$. Next, it samples $D'_1, \dots, D'_{i^*-1}, D'_{i^*+1}, \dots, D'_n \leftarrow_R \mathcal{D}'$ independently, and sets $D'_{i^*} = D'$. Finally, it runs $\mathcal{A}$ on $D^* = D \times (D'_1, \dots, D'_n)$ and then runs the reconstruction attack $\mathcal{R}^*$ to recover answers $(t_{q',i})_{q' \in \mathcal{Q}', i \in [n]}$ and outputs $(t_{q',i^*})_{q' \in \mathcal{Q}'}$. The following random variables are identically distributed:

1. $(t_{q',i})_{q' \in \mathcal{Q}'}$, where $(t_{q',i})_{q' \in \mathcal{Q}', i \in [n]}$ is the output of $\mathcal{R}_D^*(\mathcal{A}(D^*))$ on $D^* \leftarrow_R \mathcal{D}^*$, and $i \leftarrow_R [n]$.

2. $\mathcal{A}'(D')$ where $D' \leftarrow_{\mathcal{R}} \mathcal{D}'$.

Thus the left-hand side of (4) is equal to

$$\Pr_{D' \leftarrow_{\mathcal{R}} \mathcal{D}'} \left[\bigwedge (\mathcal{A}'(D') \text{ is } (6c\alpha', 2/c)\text{-accurate for } \mathcal{Q}') = \perp \right] \leq \gamma$$

which follows because $\mathcal{D}'$ is a (γ, ξ) -re-identifiable from $(6c\alpha', 2/c)$ -accurate answers to $\mathcal{Q}'$. Thus we have established (4), completing the proof of the claim. $\square$

The next claim follows directly from the definition of $\mathcal{B}^*$ and the fact that $\mathcal{D}'$ is (γ, ξ) -re-identifiable.

CLAIM 4.5. *For every $(i, j) \in [n] \times [n']$,*

$$\Pr_{D \leftarrow_{\mathcal{R}} \mathcal{D}^*} [\mathcal{B}^*(D, \mathcal{A}(D_{-(i,j)})) = (i, j)] \leq \xi.$$

Combining Claims 4.4 and 4.5 suffices to prove that $\mathcal{D}^*$ is $(\gamma + 1/6, \xi)$ -re-identifiable from (α, β) -accurate answers to $\mathcal{Q} \wedge \mathcal{Q}'$, completing the proof of the theorem.

5. APPLICATIONS OF THE COMPOSITION THEOREM

In this section we show how to use our composition theorem (Section 4) to combine our new lower bounds for 1-way marginal queries from Section 3 with (variants of) known lower bounds from the literature to obtain our main results. In Section 5.1 we prove a lower bound for k -way marginal queries when α is a constant that establishes Theorem 1.3 from the introduction. Then in Section 5.2 we show a stronger lower bound for arbitrary counting queries when α is a varying parameter, thereby proving Theorem 1.2 in the introduction.

5.1 Lower Bounds for Answering k -Way Marginals with Constant Accuracy

In this section, we carry out the composition of sample complexity lower bounds for k -way conjunctions as described in the introduction. Recall that we obtain our new $\tilde{\Omega}(k\sqrt{d})$ lower bound by combining the re-identification based $\tilde{\Omega}(\sqrt{d})$ lower bound for 1-way marginals (Section 3.2) with a known $\Omega(k)$ lower bound based on a reconstruction attack. The lower bound of $\Omega(k)$ for k -way marginals is a special case of a lower bound of $\Omega(VC(\mathcal{Q}))$ due to [29] and based on [12], where $VC(\mathcal{Q})$ is the *Vapnik-Chervonenkis (VC) dimension* of $\mathcal{Q}$. To apply our composition theorem, we need to formulate this reconstruction attack in the language of Definition 4.1. In particular, we observe that the same proof generalizes to allow us to reconstruct fractional vectors $s \in [0, 1]^n$, instead of just Boolean vectors as in [12, 29]. We emphasize that we apply the same argument as in those works.

DEFINITION 5.1 (VC DIMENSION). *Let $\mathcal{Q}$ be a collection of counting queries over a data universe $\mathcal{X}$. We say a set $\{x_1, \dots, x_k\} \subseteq \mathcal{X}$ is shattered by $\mathcal{Q}$ if for every string $v \in \{0, 1\}^k$, there exists a query $q \in \mathcal{Q}$ such that $(q(x_1), \dots, q(x_k)) = (v_1, \dots, v_k)$. The VC-Dimension of $\mathcal{Q}$ denoted $VC(\mathcal{Q})$ is the cardinality of the largest subset of $\mathcal{X}$ that is shattered by $\mathcal{Q}$.*

The following fact is well-known.

FACT 5.2. *The set of k -way conjunctions $\mathcal{M}_{k,d}$ over any data universe $\{0, 1\}^d$ with $d \geq k$ has VC-dimension $VC(\mathcal{M}_{k,d}) \geq k$.*

LEMMA 5.3 (VARIANT OF [12, 29]). *Let $\mathcal{Q}$ be a collection of counting queries over a data universe $\mathcal{X}$ and let $n = VC(\mathcal{Q})$. Then there is a database $D \in \mathcal{X}^n$ which admits a 4α -reconstruction attack from $(\alpha, 0)$ -accurate answers to $\mathcal{Q}$.*

The proof is a simple variant of the arguments in [12, 29], and appears in the full version of this work.

We can now prove our sample-complexity lower bound for k -way marginals, thereby establishing Theorem 1.3 in the introduction.

THEOREM 5.4. *There exists a universal constant $\alpha_0 > 0$ such that for every $k, d \in \mathbb{N}$, $k \leq d$, there is an*

$$n = n(k, d) = \tilde{\Omega}(k\sqrt{d})$$

such that there exists a distribution on n -row databases $D \in (\{0, 1\}^d)^n$ that is $(1/3, o(1/n))$ -re-identifiable from $(\alpha_0, 0)$ -accurate answers to the k -way marginals $\mathcal{M}_{k,d}$.

PROOF. The previous results will imply the existence of two privacy attacks, and we obtain the result by applying the composition theorem (Theorem 4.3) to them.

1. By combining Theorem 3.3 and Theorem 3.2, there exists a distribution on databases $D' \in (\{0, 1\}^{d/2})^{n'}$ that is $(\gamma = 1/6, \xi = o(1/n'k))$ -re-identifiable from $(6c\alpha' = 1/3, 2/c = 1/75)$ accurate answers to the 1-way marginals $\mathcal{M}_{1,d/2}$ for $n' = \tilde{\Omega}(\sqrt{d}/\log(dk))$,
2. By Lemma 5.3 and Fact 5.2, there exists a database $D \in (\{0, 1\}^{d/2})^{k-1}$ that admits a $(\alpha' = 4\alpha_0)$ -reconstruction attack from $(\alpha = \alpha_0, \beta = 0)$ -accurate answers to the $(k-1)$ -way marginals $\mathcal{M}_{k-1,d/2}$ for any α_0 .

By applying Theorem 4.3 (with parameter $c = 150$) to these two distributions, we obtain a new distribution on $(\{0, 1\}^d)^{n'(k-1)}$ that is $(1/3, o(n'k))$ -re-identifiable from $(\alpha_0, 0)$ -accurate answers to $\mathcal{M}_{k-1,d/2} \wedge \mathcal{M}_{1,d/2}$ on $\{0, 1\}^{d/2} \times \{0, 1\}^{d/2}$. Note that this family of queries is a subset of $\mathcal{M}_{k,d}$ on $\{0, 1\}^d$, but $(\alpha_0, 0)$ -accuracy for $\mathcal{M}_{k,d}$, (accuracy for all queries in $\mathcal{M}_{k,d}$) implies $(\alpha_0, 0)$ -accuracy for any subset of $\mathcal{M}_{k,d}$. $\square$

Moreover, just as with the 1-way marginals (cf. Corollary 3.5), we can show the sample complexity increases at least linearly with $1/\alpha$ for vanishing α .

COROLLARY 5.5. *For every $d \in \mathbb{N}$ and $0 < \alpha \leq \alpha_0$, the k -way marginals $\mathcal{M}_{k,d}$ have sample complexity at least $\tilde{\Omega}(k\sqrt{d}/\alpha)$ for $(\alpha, 0)$ -accuracy and $(O(1), o(1/n))$ -differential privacy.*

5.2 Lower Bounds for Arbitrary Queries

Using our composition theorem, we can also prove a nearly-optimal sample complexity lower bound as a function of the $|\mathcal{Q}|, d$, and α and establish Theorem 1.2 in the introduction. The result will follow from three lower bounds: the $\tilde{\Omega}(\sqrt{d})$ lower bound for 1-way marginals and the $\Omega(VC(\mathcal{Q}))$ bound that we have already discussed, and a lower bound of $\Omega(1/\alpha^2)$ that is a simple variant of the seminal reconstruction attack of Dinur and Nissim [12], and related attacks such as [15, 20]. Roughly, the results of [12] can be interpreted in our framework as showing that there is an

$\Omega(1/\alpha^2)$ -row database that admits a $1/100$ -reconstruction attack from $(\alpha, 0)$ -accurate answers to some family of queries $\mathcal{Q}$, but only when the vector to be reconstructed is Boolean. That is, the attack reconstructs a bit vector accurately provided that every query in $\mathcal{Q}$ is answered correctly. Dwork et al. [15, 20] generalized this attack to only require (α, β) -accuracy for some constant $\beta > 0$, and we will make use of this extension (although we do not require computational efficiency, which was a focus of those works). Finally, we need an extension to the case of fractional vectors $s \in [0, 1]^n$, instead of Boolean vectors $s \in \{0, 1\}^n$.

The extension is fairly simple and the proof follows the same outline of the original reconstruction attack from [12]. We are given accurate answers to queries in $\mathcal{Q}$, which we interpret as approximate “subset-sums” of the vector $s \in [0, 1]^n$ that we wish to reconstruct. The reconstruction attack will output any vector t from a discretization $\{0, 1/m, \dots, (m-1)/m, 1\}^n$ of the unit interval that is “consistent” with these subset-sums. The main lemma we need is an “elimination lemma” that says that if $\|t - s\|_1$ is sufficiently large, then for a random subset $T \subseteq [n]$,

$$\frac{1}{n} \left| \sum_{i \in T} t_i - s_i \right| > 3\alpha$$

with suitable large constant probability. For $m = 1$ this lemma can be established via combinatorial arguments, whereas for the $m > 1$ case we establish it via the Berry-Esséen Theorem. The lemma is used to argue that for every t that is sufficiently far from s , a large fraction of the subset-sum queries will witness the fact that t is far from s , and ensure that t is not chosen as the output.

First we state the lemma that we just described, and then we will verify that it indeed leads to a reconstruction attack. The proof appears in the full version of this work.

LEMMA 5.6. *Let $\kappa > 0$ be a constant, let $\alpha > 0$ be a parameter with $\alpha \leq \kappa^2/240$, and let $n = 1/576\kappa^2\alpha^2$. Then for every $r \in [-1, 1]^n$ such that $\frac{1}{n} \sum_{i=1}^n |r_i| > \kappa$, and a randomly chosen $q \subseteq [n]$,*

$$\Pr_{q \subseteq [n]} \left[\left| \frac{1}{n} \sum_{i \in q} r_i \right| > 3\alpha \right] \geq \frac{3}{5}.$$

THEOREM 5.7. *Let $\alpha' \in (0, 1]$ be a constant, let $\alpha > 0$ be a parameter with $\alpha \leq (\alpha')^2/960$, and let $n = 1/144(\alpha')^2\alpha^2$. For any data universe $\mathcal{X} = \{x_1, \dots, x_n\}$ of size n , there is a set of counting queries $\mathcal{Q}$ of size at most $O(n \log(1/\alpha))$ such that the database $D = (x_1, \dots, x_n)$ admits a α' -reconstruction attack from $(\alpha, 1/3)$ -accurate answers to $\mathcal{Q}$.*

PROOF SKETCH. First we will give a reconstruction algorithm $\mathcal{B}$ for an arbitrary family of queries. We will then show that for a random set of queries $\mathcal{Q}$ of the appropriate size, the reconstruction attack succeeds for every $s \in [0, 1]^n$ with non-zero probability, which implies that there exists a set of queries satisfying the conclusion of the theorem. We will use the shorthand

$$\langle q, s \rangle = \frac{1}{n} \sum_{i=1}^n q(x_i) s_i$$

for vectors $s \in [0, 1]^n$.

In order to show that the reconstruction attack $\mathcal{B}$ from Figure 5.2 succeeds, we must show that $\frac{1}{n} \sum_{i=1}^n |t_i - s_i| \leq \alpha'$.

Input: Queries $\mathcal{Q}$, and $(a_q)_{q \in \mathcal{Q}}$ that are $(\alpha, 1/3)$ -accurate for s .
Let $m = \lceil \frac{1}{\alpha} \rceil$
Find any $t \in \{0, 1/m, \dots, (m-1)/m, 1\}^n$ such that

$$\Pr_{q \leftarrow \mathcal{R}(\mathcal{Q})} [|\langle q, t \rangle - a_q| < 2\alpha] > \frac{5}{6}.$$

Output: t .

Figure 4: The reconstruction adversary $\mathcal{B}$.

Let $s \in [0, 1]^n$, and let $s' \in \{0, 1/m, \dots, (m-1)/m, 1\}^n$ be the vector obtained by rounding each entry of s to the nearest $1/m$. It is enough to show that the reconstruction attack outputs a vector close to s' . The vector s' itself satisfies $|\langle q, s' \rangle - a_q| \leq 2\alpha$ for any subset-sum query q , so the reconstruction attack always finds some vector t . To show that the reconstruction is successful, fix any $t \in \{0, 1/m, \dots, (m-1)/m, 1\}^n$ such that $\frac{1}{n} \sum_{i=1}^n |t_i - s'_i| > \frac{\alpha'}{2}$. If we write $r = s' - t \in \{-1, \dots, -1/m, 0, 1/m, \dots, 1\}^n$, then $\frac{1}{n} \sum_{i=1}^n |r_i| > \frac{\alpha'}{2}$ and $\langle q, r \rangle = \langle q, t \rangle - \langle q, s' \rangle$. In order to show that no t that is far from s' can be output by $\mathcal{B}$, we will show that for any $r \in \{-1, \dots, -1/m, 0, 1/m, \dots, 1\}^n$ with $\frac{1}{n} \sum_{i=1}^n |r_i| > \frac{\alpha'}{2}$,

$$\Pr_{q \leftarrow \mathcal{R}(\mathcal{Q})} [|\langle q, r \rangle| > 3\alpha] \geq \frac{1}{2}.$$

To prove this, we first observe by Lemma 5.6 (setting $\kappa = \frac{1}{2}\alpha'$) that for a randomly chosen query q defined on $\mathcal{X}$,

$$\Pr_q [|\langle q, r \rangle| > 3\alpha] \geq \frac{3}{5}.$$

The lemma applies because $\langle q, r \rangle = \frac{1}{n} \sum_{i=1}^n q(x_i) r_i$ is a random subset-sum of the entries of r . Applying a Chernoff bound and a union bound, we find that there exists a family of queries $\mathcal{Q}$ of size $O(n \log m)$ such that for every s, t such that $\frac{1}{n} \sum_{i=1}^n |t_i - s_i| > \alpha'$,

$$\Pr_{q \leftarrow \mathcal{R}(\mathcal{Q})} [|\langle q, s \rangle - \langle q, t \rangle| > 3\alpha] \geq \frac{1}{2}.$$

By $(\alpha, 1/3)$ -accuracy and a triangle inequality, we can conclude

$$\Pr_{q \leftarrow \mathcal{R}(\mathcal{Q})} [|a_q - \langle q, t \rangle| > 2\alpha] \geq \frac{1}{2} - \frac{1}{3} \geq \frac{1}{6},$$

which implies that t cannot be the output of $\mathcal{B}$. This completes the proof. $\square$

5.2.1 Putting Together the Lower Bound

Now we show how to combine the various attacks to prove Theorem 1.2 in the introduction. We obtain our lower bound by applying two rounds of composition. In the first round, we compose the reconstruction attack described above with the re-identifiable distribution for 1-way marginals. We then take the resulting re-identifiable distribution and apply a second round of composition using the reconstruction attack for query families of high VC-dimension, just as in the proof of Theorem 1.3. The formal proof is deferred to the full version of this work.

THEOREM 5.8. *For all $d \in \mathbb{N}$, all sufficiently small (i.e. bounded by an absolute constant) $\alpha > 2^{-d/6}$, and all $h \leq$*

$2^{d/3}$, there is a family of queries $\mathcal{Q}$ of size $O(hd \log(1/\alpha)/\alpha^2)$ and an

$$n = n(h, d, \alpha) = \tilde{\Omega} \left(\frac{\sqrt{d} \log h}{\alpha^2} \right)$$

such that there exists a distribution on n -row databases $D \in (\{0, 1\}^d)^n$ that is $(1/2, o(1/n))$ -re-identifiable from $(\alpha, 0)$ -accurate answers to $\mathcal{Q}$.

Acknowledgements

We thank Kobbi Nissim for drawing our attention to the question of sample complexity and for helpful discussions and the anonymous reviewers for helpful comments.

6. REFERENCES

- [1] BARAK, B., CHAUDHURI, K., DWORK, C., KALE, S., MCSHERRY, F., AND TALWAR, K. Privacy, accuracy, and consistency too: a holistic solution to contingency table release. In *PODS* (2007), pp. 273–282.
- [2] BEIMEL, A., KASIVISWANATHAN, S. P., AND NISSIM, K. Bounds on the sample complexity for private learning and private data release. In *TCC* (2010), pp. 437–454.
- [3] BEIMEL, A., NISSIM, K., AND STEMMER, U. Characterizing the sample complexity of private learners. In *ITCS* (2013), pp. 97–110.
- [4] BEIMEL, A., NISSIM, K., AND STEMMER, U. Private learning and sanitization: Pure vs. approximate differential privacy. In *APPROX-RANDOM* (2013), pp. 363–378.
- [5] BLUM, A., DWORK, C., MCSHERRY, F., AND NISSIM, K. Practical privacy: the SuLQ framework. In *PODS* (2005), pp. 128–138.
- [6] BLUM, A., LIGETT, K., AND ROTH, A. A learning theory approach to non-interactive database privacy. In *STOC* (2008), pp. 609–618.
- [7] BONEH, D., KAIYAS, A., AND MONTGOMERY, H. W. Robust fingerprinting codes: a near optimal construction. In *Digital Rights Management Workshop* (2010), pp. 3–12.
- [8] BONEH, D., AND NAOR, M. Traitor tracing with constant size ciphertext. In *ACM Conference on Computer and Communications Security* (2008), pp. 501–510.
- [9] BONEH, D., AND SHAW, J. Collusion-secure fingerprinting for digital data. *IEEE Transactions on Information Theory* 44, 5 (1998), 1897–1905.
- [10] CHANDRASEKARAN, K., THALER, J., ULLMAN, J., AND WAN, A. Faster private release of marginals on small databases. *ITCS 2014 (to appear)* (2014).
- [11] DE, A. Lower bounds in differential privacy. In *TCC* (2012), pp. 321–338.
- [12] DINUR, I., AND NISSIM, K. Revealing information while preserving privacy. In *PODS* (2003), pp. 202–210.
- [13] DWORK, C., KENTHAPADI, K., MCSHERRY, F., MIRONOV, I., AND NAOR, M. Our data, ourselves: Privacy via distributed noise generation. In *EUROCRYPT* (2006), pp. 486–503.
- [14] DWORK, C., MCSHERRY, F., NISSIM, K., AND SMITH, A. Calibrating noise to sensitivity in private data analysis. In *TCC* (2006), pp. 265–284.
- [15] DWORK, C., MCSHERRY, F., AND TALWAR, K. The price of privacy and the limits of lp decoding. In *STOC* (2007), pp. 85–94.
- [16] DWORK, C., NAOR, M., REINGOLD, O., ROTHBLUM, G. N., AND VADHAN, S. P. On the complexity of differentially private data release: efficient algorithms and hardness results. In *STOC* (2009), pp. 381–390.
- [17] DWORK, C., NIKOLOV, A., AND TALWAR, K. Efficient algorithms for privately releasing marginals via convex programming. *Manuscript* (2013).
- [18] DWORK, C., AND NISSIM, K. Privacy-preserving datamining on vertically partitioned databases. In *CRYPTO* (2004), pp. 528–544.
- [19] DWORK, C., ROTHBLUM, G. N., AND VADHAN, S. P. Boosting and differential privacy. In *FOCS* (2010), pp. 51–60.
- [20] DWORK, C., AND YEKHANIN, S. New efficient attacks on statistical disclosure control mechanisms. In *CRYPTO* (2008), pp. 469–480.
- [21] GUPTA, A., HARDT, M., ROTH, A., AND ULLMAN, J. Privately releasing conjunctions and the statistical query barrier. In *STOC* (2011), pp. 803–812.
- [22] GUPTA, A., ROTH, A., AND ULLMAN, J. Iterative constructions and private data release. In *TCC* (2012), pp. 339–356.
- [23] HARDT, M. *A Study in Privacy and Fairness in Sensitive Data Analysis*. PhD thesis, Princeton University, 2011.
- [24] HARDT, M., LIGETT, K., AND MCSHERRY, F. A simple and practical algorithm for differentially private data release. In *NIPS* (2012), pp. 2348–2356.
- [25] HARDT, M., AND ROTHBLUM, G. N. A multiplicative weights mechanism for privacy-preserving data analysis. In *FOCS* (2010), pp. 61–70.
- [26] HARDT, M., AND TALWAR, K. On the geometry of differential privacy. In *STOC* (2010), pp. 705–714.
- [27] KASIVISWANATHAN, S. P., RUDELSON, M., SMITH, A., AND ULLMAN, J. The price of privately releasing contingency tables and the spectra of random matrices with correlated rows. In *STOC* (2010), pp. 775–784.
- [28] NIKOLOV, A., TALWAR, K., AND ZHANG, L. The geometry of differential privacy: the sparse and approximate cases. In *STOC* (2013), pp. 351–360.
- [29] ROTH, A. Differential privacy and the fat-shattering dimension of linear queries. In *APPROX-RANDOM* (2010), pp. 683–695.
- [30] ROTH, A., AND ROUGHGARDEN, T. Interactive privacy via the median mechanism. In *STOC* (2010), pp. 765–774.
- [31] TARDOS, G. Optimal probabilistic fingerprint codes. *J. ACM* 55, 2 (2008).
- [32] THALER, J., ULLMAN, J., AND VADHAN, S. P. Faster algorithms for privately releasing marginals. In *ICALP (I)* (2012), pp. 810–821.
- [33] ULLMAN, J. Answering $n^{2+o(1)}$ counting queries with differential privacy is hard. In *STOC* (2013), pp. 361–370.