The Limits of Two-Party Differential Privacy

Andrew McGregor*, Ilya Mironov†, Toniann Pitassi‡, Omer Reingold†, Kunal Talwar†, Salil Vadhan§
*Department of Computer Science University of Massachusetts, Amherst.
†Microsoft Research Silicon Valley.
‡Department of Computer Science University of Toronto.
§School of Engineering and Applied Sciences and Center for Research on Computation and Society
Harvard University.

***Abstract*—We study differential privacy in a distributed setting where two parties would like to perform analysis of their joint data while preserving privacy for both datasets. Our results imply almost tight lower bounds on the accuracy of such data analyses, both for specific natural functions (such as Hamming distance) and in general. Our bounds expose a sharp contrast between the two-party setting and the simpler client-server setting (where privacy guarantees are one-sided). In addition, those bounds demonstrate a dramatic gap between the accuracy that can be obtained by differentially private data analysis versus the accuracy obtainable when privacy is relaxed to a computational variant of differential privacy.**

The first proof technique we develop demonstrates a connection between differential privacy and deterministic extraction from Santha-Vazirani sources. A second connection we expose indicates that the ability to approximate a function by a low-error differentially-private protocol is strongly related to the ability to approximate it by a low communication protocol. (The connection goes in both directions.)

I. Introduction

A common architecture for database access is client-server, where the server manages the data and answers clients' queries according to its access policy. In such an architecture, there may be two very distinct privacy considerations. The first has to do with client's privacy and is highly motivated in cases where the server's knowledge of client's queries may be harmful, for instance, in patent litigation or market research. In such cases, without an expectation of privacy, clients may be discouraged from querying the database in the first place. Such concerns can be answered using various cryptographic solutions such as oblivious transfer [1], [2], single-server private-information retrieval (PIR) [3], and more generally, secure function evaluation (SFE) [4], which may be used to restore privacy for the clients.

The focus of this paper has to do with a complementary privacy concern: what kind of access should a server allow to the database while preserving the privacy of sensitive data that it may contain. In other words, the question we study is not *how* data analysis can be performed while preserving client's privacy (the cryptographic question) but rather *what* kind of data analysis preserves data privacy. While the answer to this question may be dependent on the nature of the data, a very powerful general-purpose notion is that of *differential privacy* [5], [6]. Informally, a randomized function of a database is *differentially private* if its output distribution is insensitive to the presence or absence of any particular record in the database. Therefore, if the analyses allowed on a database are guaranteed to preserve differential privacy, there is little incentive for an individual to conceal his or her information from the database (and in this respect the privacy of individual records is preserved).

Assume that a query to a database is a deterministic real-valued function. In such a case, differential privacy may be enforced by adding a small amount of noise, calibrated to the *sensitivity* of that function (defined as the largest change in its output that can be caused by adding or removing a record from its input). In the basic client-server setting, queries of constant sensitivity can be answered by adding Laplacian (symmetric exponential) noise with standard deviation inversely proportional to the privacy parameter [5], and indeed this mechanism can be shown to be optimal for counting queries as well as for a large class of clients' preferences [7].

Two-party differential privacy. In this paper we contrast the client-server setting with a setting where the database is distributed between two parties which would like to perform data analysis on their joint data. In this setting we would like to guarantee two-sided differential privacy, protecting the data of both parties. That is, each party's view of the protocol should be a differentially private function of the other party's input. Differential privacy for distributed databases was first considered in the seminal work on privacy-preserving distributed datamining by Dwork and Nissim [8]. More accurately, the definition of privacy in [8] is a precursor (and indeed a special case) of the now-standard definition of approximate differential privacy. Differential privacy in a highly distributed setting (which is less related to our work), was also considered in [9].

Although the distributed setting was considered earlier in the line of research on differential privacy, the state of knowledge in this setting was very minimal. While there were protocols given for specific functions (e.g., in [8], [10], [11]), there were no *general* results or lower bounds for computing functions

Andrew McGregor is supported by the NSF CAREER Award CCF-0953754. Work done in part while visiting Microsoft Research.

Omer Reingold was at the Weizmann Institute during some of this research, supported by US-Israel BSF grant 2006060.

Tonian Pitassi is supported by NSERC. Work done in part while visiting Microsoft Research.

Salil Vadhan is supported by NSF grant CNS-0831289 and US-Israel BSF grant 2006060. Work done in part while visiting Microsoft Research.

with two-sided differential privacy guarantees (in sharp contrast with the case of one-sided differential privacy). The goal of this paper is to start filling that gap.

The limitations of two-party differential privacy. Motivated by the work of Dwork and Nissim [8], we start our study with two related and very natural problems: the Hamming distance between two binary vectors (in how many locations they differ) and their scalar product.[1] We formulate the following prototypical problem for privacy-preserving two-party computations:

> Question 1. *What is the least additive error of any protocol for computing the Hamming distance between two binary vectors that is differentially private for both sides?*

Note that the Hamming distance is a function of sensitivity one (changing one bit can change the function by at most one). Therefore in the client-server setting this function could be approximated up to a constant additive error, while ensuring differential privacy (as discussed above). In this paper we show that the case of two-sided privacy is very different: *Any protocol for computing the Hamming distance of two n-bit vectors that is differentially private for both sides incurs additive error of $\tilde{\Omega}(\sqrt{n})$ and this is tight up to the a hidden log factor.*

A natural approach to approximating the Hamming distance by two parties is to use secure function evaluation in order to emulate a trusted third party, which has access to both parties' inputs, and operates as in the client-server setting (i.e., evaluates the Hamming distance and adds appropriate Laplacian noise). Similarly, every function with small sensitivity can be approximated well using secure-function evaluation. The "catch" (and the reason this does not contradict our aforementioned result on the Hamming distance) is that this approach only achieves a relaxed notion of *computational* differential privacy [11]. Loosely, this notion of differential privacy only holds against computationally-bounded adversaries. In other words, our result regarding the Hamming distance implies a separation between (information theoretic) differential privacy and computational differential privacy. It is natural to ask, if this separation can be made even stronger:

> Question 2. *What is the largest gap in accuracy between optimal differentially-private and computationally differentially-private protocols?*

Indeed, we show that the gap between accuracy can be as large as linear. We do so by exhibiting a function on two n-bit strings with constant sensitivity that cannot be privately approximated within error $o(n)$. Such a strong separation between (information-theoretic) differential privacy and computational differential privacy again stands in sharp contrast with the client-server setting where all of the known positive results have achieved information-theoretic differential privacy and there are not even candidates for a separation. In this respect, differential privacy in the two-party setting is closer to cryptography where most interesting tasks can only be obtained with computational rather than information theoretic security.

The techniques we develop to address the above questions rely on intriguing new connections: the first is a connection between differential privacy in the two-party setting and deterministic extractors for Santha-Vazirani sources. The second connection is with the communication complexity of two-party protocols. We further develop this latter connection and in particular demonstrate that the connection works in both directions. Loosely, and ignoring the relation between the various parameters, we show that a small-communication protocol for a function exists if and only if a low-error differentially private protocol exists. We now discuss our results in more detail and elaborate on the new connections we discover.

Hamming distance and deterministic extraction. We resolve the first question discussed above by establishing a connection between differentially private protocols and deterministic extractors for Santha-Vazirani sources.

Consider two uniformly distributed n-bit strings x and y which are the inputs of two parties that would like to approximate the Hamming distance. For any two-party protocol, conditioned on the transcript of the protocol, x and y are independent. Furthermore, if the protocol is differentially-private then each bit of x has some entropy even conditioned on *all other bits of* x (and similarly for y). In other words, conditioned on the transcript, x and y are two independent Santha-Vazirani sources. We then generalize a result of Vazirani [12] to argue that the inner product modulo $\lfloor\sqrt{n}\rfloor$ is a good (deterministic) extractor for such sources (i.e., it is distributed nearly uniformly over its range). This implies that no party is able to estimate the inner product (and consequently, the Hamming distance) of the inputs with accuracy $o(\sqrt{n}/\log n)$. This is almost tight, as standard randomized response [13] allows parties to approximate their Hamming distance with error $\Theta(\sqrt{n}/\epsilon)$ (both bounds assume that the privacy parameter ϵ is smaller than 1). More formally, the following theorem answers Question 1 from above:

Theorem 5 (Section III)**.** *Let $P(x,y)$ be a randomized protocol with ϵ-differential privacy for inputs $x,y \in \{0,1\}^n$, and let $\delta > 0$. Then, with probability at least $1-\delta$ over $x,y \leftarrow \{0,1\}^n$ and the coin tosses of P, party B's output differs from $\langle x,y\rangle$ by at least $\Delta = \Omega\left(\frac{\sqrt{n}}{\log n}\cdot\frac{\delta}{e^{\epsilon}}\right)$.*

Communication complexity and differential privacy. Towards answering the second question posed above, we note that the method based on deterministic-extraction from Santha-Vazirani sources is unlikely to yield (at least naïvely) a lower bound on additive error better than $O(\sqrt{n})$ (see Section III-C). We therefore develop a different approach based on a new connection between differentially-private protocols and communication complexity. We systematically explore these connections.

We first prove that the *information cost* (as defined by Bar-Yossef et al. [14]) and the *partition bound* (as defined by Jain

[1] In [8], a central data-mining problem (detecting correlations between two binary attributes) was reduced to approximating the scalar product between two binary vectors.

and Klauck [15]) of an ϵ-differentially-private protocol are both $O(\epsilon n)$. Loosely, information cost measures the amount of information that is shared between the transcript and the input of both parties. Therefore, the $O(\epsilon n)$ bound on the information cost in particular is quite natural, since differential privacy condition limits the amount of information learned on each individual bit of the inputs (and is thus only stronger). Motivated by applications in direct-sum theorems for communication complexity, Barak et al. [16] proved that a protocol over a product distribution can be compressed down to its information cost (up to a polylogarithmic factor in its original communication complexity). We can conclude that every ϵ-differentially-private protocol can be compressed to a small (roughly $O(\epsilon n)$) communication protocol (see Theorem 10).

Given the reduction from differential privacy to information cost, we construct a function with two properties: (1) the function has sensitivity 1 and range $\Theta(n)$; (2) approximating the function to within $o(n)$ by a 2-party protocol requires linear (in its input length) information cost. We construct such a function by taking an arbitrary boolean function with high information cost, embedding it in the space of codewords and extending its domain to all inputs in a manner consistent with the sensitivity condition. Such a function proves that the answer to Question 2 on the gap between two-party computational and information-theoretic differential privacy is linear: On the one hand, by property (1) the function can be approximated with differential privacy by a trusted third party (and thus in the computational setting too, by simulating the third party using SFE) with error proportional to $1/\epsilon$. On the other hand, every (information-theoretic) differentially-private protocol linear additive error. More precisely, the following theorem claims these properties of our construction:

Theorem 13 (Section IV-C)**.** *There exists an absolute constant $\beta > 0$ such that for every n, there is an efficiently computable function $f\colon \{0,1\}^n \times \{0,1\}^n \to \mathbb{R}$ and a distribution $\mathcal{D}$ over its inputs, with the following properties:*
(a) for every $\epsilon < \beta/3$, every ϵ-differentially private protocol P must incur additive error at least βn with probability at least $\frac{1}{10}$.
(b)for every ϵ, δ with $\epsilon + 1.2(\delta/\epsilon) < \beta/10$, every (ϵ, δ)-differentially private protocol P must incur additive error at least βn with probability at least $\frac{1}{10}$.
(c) f has sensitivity 1, i.e., $|f(x,y) - f(x',y')| \leq |(x,y) - (x',y')|_H$ for every x, y, x', y'.

We note that the connection between differential-privacy and communication complexity can be used to prove lower bounds on the error in computing specific functions for which lower bounds on the communication complexity is known. For the specific example of Hamming distance (see [17]), the results obtained in this manner are incomparable with those obtained via deterministic extraction.

The connection between differential privacy and communication complexity is quite strong and we explore it beyond our original motivation discussed above. In particular, for the application above we only cared that differentially-private protocols can be compressed into protocols that have low communication but are not necessarily differentially-private. Our next result demonstrates that every differentially-private protocol with r rounds can be compressed down to $O(\epsilon r n)$ *while keeping it differentially private*. Compression is implemented using privacy-preserving consistent sampling [18], [19] and has a negligible probability of failure (which affects accuracy, not privacy). The formal theorem is stated as follows:

Theorem 14 (Section IV-D)**.** *Let P be an ϵ-differentially private protocol with r rounds. Then, for every $\delta > 0$, there exists an $O(r\epsilon)$-differentially-private protocol P^* that has communication complexity $O(r(\epsilon n + \log\log \frac{1}{\epsilon\delta}))$ and except with probability $r\delta$, simulates P perfectly.*

In our final result we show that the connection between differential privacy and communication complexity goes the other way too: a deterministic protocol with r rounds and communication C can be transformed into an ϵ-differentially-private protocol with additive error $O(Cr/\epsilon)$:

Theorem 15 (Section IV-E)**.** *Let P be a deterministic protocol with communication complexity $\mathsf{CC}(P)$ and the number of rounds r approximating a sensitivity-1 function $f\colon \Sigma^n \times \Sigma^n \to \mathbb{Z}$ with error bounded by Δ. Then there exists an ϵ-differentially-private protocol with the same communication complexity and number of rounds that computes f with expected additive error $\Delta + O(\mathsf{CC}(P)r/\epsilon)$.*

The linear dependency on the communication complexity in the last theorem is unlikely to be improved due to the lower bound of Theorem 13.

Other Related Work Recently, Feigenbaum et al. [20] proposed a notion of approximate privacy for communication protocols. Their work is rather different from ours in that it only applies to deterministic protocols for evaluating functions exactly. Additionally, like in SFE, it is assumed that the function value itself is non-sensitive and can be revealed. Their notion of approximate privacy is based on the ratio of the number of inputs that are consistent with the final answer to the number of inputs consistent with the transcript (which includes the final answer). If this ratio is one, then the protocol is deemed *perfectly private* (see also, [21], [22]) in the sense that no more is information about the input is revealed beyond that revealed by the final answer.

II. Definitions

Let Σ be a finite alphabet and for strings $x, y \in \Sigma^n$, let $|x - y|_H = |\{i \in [n] : x_i \neq y_i\}|$ denote the Hamming distance between x and y. We recall the standard definition of differential privacy for mechanisms defined over strings from a finite alphabet Σ and generalize it to interactive protocols, following [9].

Definition 1 (Differential privacy)**.** *A mechanism M on Σ^n is a family of probability distributions $\{\mu_x : x \in \Sigma^n\}$ on $\mathcal{R}$. The mechanism is* ϵ-differentially private *if for every x and x'*

such that $|x - x'|_H = 1$ *and every measurable subset* $S \subset \mathcal{R}$ *we have*

$$\mu_x(S) \leq \exp(\epsilon)\mu_{x'}(S).$$

A common relaxation of ϵ-differential privacy is the following definition of δ-approximate ϵ-differential privacy, abbreviated as (ϵ, δ)-differential privacy:

Definition 2 (Approximate differential privacy)**.** *The mechanism* M *satisfies* δ-approximate ϵ-differential privacy *if for every* x *and* x' *such that* $|x - x'|_H = 1$ *and every measurable subset* $S \subset \mathcal{R}$ *we have*

$$\mu_x(S) \leq \exp(\epsilon)\mu_{x'}(S) + \delta.$$

The definition of differential privacy naturally extends to interactive protocols, by requiring that the *views* of all parties be differentially private in respect to other parties' inputs. The following definition assumes semi-honest parties, i.e., parties that are guaranteed to follow the protocol. Since the focus of this work is on establishing lower bounds on accuracy of differentially-private protocols, its results apply to models with weaker restrictions on adversarial parties as well.

More concretely, let $\mathsf{VIEW}_P^A(x, y)$ be the joint probability distribution over x, the transcript of the protocol P, private randomness of the party A, where the probability space is private randomness of both parties. For each x, $\mathsf{VIEW}_P^A(x, y)$ is a mechanism over the y's. Let $\mathsf{VIEW}_P^B(x, y)$ be similarly defined view of B whose input is y.

Definition 3 (Differential privacy for two-party protocols)**.** *We say that a protocol* P *has* ϵ*-differential privacy if the mechanism* $\mathsf{VIEW}_P^A(x, y)$ *is* ϵ*-differentially private for all values of* x *and same holds for* $\mathsf{VIEW}_P^B(x, y)$ *and all values of* y*.*

Approximate differential privacy for interactive protocols is defined analogously. Without loss of generality, we assume that the parties do not share any public random bits since one party can choose those random bits and send them to the other without violating the privacy condition. Also, note that the above definition of privacy trivially maintains the privacy of x and y against a third party who only observes the transcript. In fact, this notion of privacy will be sufficient to imply many of the lower bounds we present.

The notion of (global) sensitivity of a function is useful in designing differentially-private protocol computing this function:

Definition 4 (Sensitivity)**.** *For a real-valued function* $f \colon \Sigma^n \to \mathbb{R}$ *define its* sensitivity *as the maximal difference in value on adjacent inputs, i.e.,* $\max_{|x-y|_H=1} |f(x) - f(y)|$.

The following definition plays a role in Sections III and IV:

Definition 5 (Statistical distance and δ-closeness)**.** *Given random variables* X *and* X' *taking values in* Ω*, we say that* X *and* X' *are* δ-close *if the* statistical distance *between their distributions is at most* δ*, i.e.,*

$$\|X - X'\|_{SD} := \frac{1}{2} \sum_{x \in \Omega} |\Pr[X = x] - \Pr[X' = x]| \leq \delta \ .$$

Communication Complexity. Yao [23] introduced the following, by now classical, two-player communication game: Alice and Bob want to collaboratively compute a function $f \colon \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$. Alice gets an n-bit string x and Bob gets another, called y. The players have unlimited computational power. They agree on a protocol beforehand, according to which they take turns in communicating with each other. At a player's turn, what that player communicates is a function of her input and what has been communicated so far. We call the sequences of messages, the *transcript* of the protocol and denote it by Π. The protocol also specifies a function $f_A(\cdot, \cdot)$ (resp. $f_B(\cdot, \cdot)$) that define the value computed by Alice (resp. Bob). Let P be a deterministic communication protocol. The cost of P, denoted $\mathsf{CC}(P)$, is the total number of bits that Alice and Bob communicate for the worst input. The deterministic complexity of f, denoted by $D(f)$, is the cost of the best deterministic protocol for f that outputs the correct answer for every input, i.e. $f_A(x, \Pi) = f_B(y, \Pi) = f(x, y)$. We also consider randomized communication protocols where the players may each flip private coins and we permit an arbitrarily small constant probability of failure, so that $\Pr[f_A(x, \Pi) = f(x, y)] \geq 1 - \gamma$ and similarly for B. For a randomized protocol, the cost of the protocol is defined as the maximum number of bits communicated over all inputs and coin flips.

III. Differential Privacy and Santha-Vazirani Sources

Differential privacy requires that a differentially-private protocol contains a limited amount of information about the parties' inputs. In particular, if the parties' inputs had a lot of entropy to begin with, then they still have a lot of entropy after we condition on the transcript of the protocol. In this section, we show that they retain much more structure than merely having high entropy. Specifically, if the parties' inputs were initially uniform and independent strings from $\{0,1\}^n$, then conditioned on any transcript of the protocol, the parties' inputs are *unpredictable bit sources* (also known as semi-random sources), as introduced by Santha and Vazirani [24] and studied in the literature on randomness extractors.

We then generalize a result of Vazirani [12] that shows that the inner product function has good randomness extraction properties on unpredictable bit sources, and use this to prove that no differentially-private two-party protocol can approximate the inner product (or the Hamming distance) to within additive error $o(\sqrt{n}/\log n)$. The extension of the result to protocols satisfying approximate differential privacy (Definition 2) appears in the full version.

A. Unpredictable Sources from Differential Privacy

The model of random sources introduced by Santha and Vazirani [24] is one where each bit is somewhat unpredictable

given the previous ones:

Definition 6 (α-unpredictable bit source[2]). *For $\alpha \in [0,1]$, random variable $X = (X_1, \ldots, X_n)$ taking values in $\{0,1\}^n$ is an* α-unpredictable bit source *if for every $i \in [n]$, and every $x_1, \ldots, x_{i-1} \in \{0,1\}$, we have*

$$\alpha \leq \frac{\Pr[X_i = 0 | X_1 = x_1, \ldots, X_{i-1} = x_{i-1}]}{\Pr[X_i = 1 | X_1 = x_1, \ldots, X_{i-1} = x_{i-1}]} \leq 1/\alpha.$$

Note that when $\alpha = 1$, the source must be the uniform distribution, and when $\alpha = 0$ the source is unconstrained. The larger α is, the more "randomness" the source is guaranteed to have. Commonly $\alpha \in (0,1)$ is thought of as being held constant as $n \to \infty$. Note also that under an α-unpredictable source, no string has probability mass greater than $1/(1+\alpha)^n$. Thus an α-unpredictable source always has min-entropy, defined as $\min_x \log(1/\Pr[X = x])$, at least βn, where $\beta = \log(1+\alpha) \geq \alpha$.

A more stringent requirement, previously studied in [25], is to require that each bit is somewhat unpredictable given *all* of the other bits, even the future ones:

Definition 7 (Strongly α-unpredictable bit source). *For $\alpha \in [0,1]$, a random variable $X = (X_1, \ldots, X_n)$ taking values in $\{0,1\}^n$ is a* strongly α-unpredictable bit source *if for every $i \in [n]$, and every $x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n \in \{0,1\}^n$, we have*

$$\alpha \leq \frac{\Pr[X_i = 0 | X_1 = x_1, \ldots, X_{i-1} = x_{i-1}, X_{i+1} = x_{i+1}, \ldots, X_n = x_n]}{\Pr[X_i = 1 | X_1 = x_1, \ldots, X_{i-1} = x_{i-1}, X_{i+1} = x_{i+1}, \ldots, X_n = x_n]} \leq 1/\alpha.$$

We now prove that conditioned on a differentially-private transcript, the parties' inputs not only have a lot of entropy, but in fact are strongly unpredictable sources (assuming they were initially uniform):

Lemma 1. *Let P be an ϵ-differentially private randomized protocol. Let X and Y be independent random variables uniformly distributed in $\{0,1\}^n$ and let random variable $\Pi(X,Y)$ denote the transcript of messages exchanged when protocol P is run on input (X,Y). Then for every $\pi \in \mathrm{Supp}(\Pi)$, the random variables corresponding to the inputs conditioned on transcript π, X_π and Y_π, are independent, strongly $e^{-\epsilon}$-unpredictable bit sources.*

Proof: The fact that independent inputs remain independent when conditioning on a transcript is a standard fact in communication complexity, which can be proved by induction on the number of rounds. (When we condition on the first message, the two inputs remain independent, and then what follows is a protocol with fewer rounds.)

To see that X_π is a strongly unpredictable bit source, we observe that by Bayes' Rule and the uniformity of X,

$$\begin{aligned}
&\frac{\Pr[X_i = 0 | X_1 = x_1, \ldots, X_{i-1} = x_{i-1}, X_{i+1} = x_{i+1}, \ldots, X_n = x_n, \Pi = \pi]}{\Pr[X_i = 1 | X_1 = x_1, \ldots, X_{i-1} = x_{i-1}, X_{i+1} = x_{i+1}, \ldots, X_n = x_n, \Pi = \pi]} \\
&\quad = \frac{\Pr[\Pi = \pi | X_1 = x_1, \ldots, X_{i-1} = x_{i-1}, X_i = 0, X_{i+1} = x_{i+1}, \ldots, X_n = x_n]}{\Pr[\Pi = \pi | X_1 = x_1, \ldots, X_{i-1} = x_{i-1}, X_i = 1, X_{i+1} = x_{i+1}, \ldots, X_n = x_n]} \\
&\quad = \frac{\Pr[\Pi(x_1 \cdots x_{i-1} 0 x_{i+1} \cdots x_n, Y) = \pi]}{\Pr[\Pi(x_1 \cdots x_{i-1} 1 x_{i+1} \cdots x_n, Y) = \pi]}.
\end{aligned}$$

By ϵ-differential privacy, the latter ratio is between $e^{-\epsilon}$ and e^{ϵ}. ∎

[2]In the terminology of Santha and Vazirani [24], this is an $\alpha/(1+\alpha)$ semi-random source.

B. Randomness Extraction and Lower Bounds for Inner Product

Vazirani [12] showed that the inner product function modulo 2 extracts an almost-uniform bit from any two independent unpredictable sources (in sharp contrast to the fact that from one unpredictable source, no function can extract a bit that is more than α-unpredictable [24]). We generalize this to show that the inner product function modulo m extracts an almost-uniform element of $\mathbb{Z}_m$, provided that the length n of the sources is at least roughly m^2. We then combine this with the results of the previous section to show that every two-party differentially-private protocol for approximating the inner product function must incur an error of roughly $m \approx \sqrt{n}$. Indeed, if a significantly better approximation could be computed given the transcript (and one party's input), then the inner product would be concentrated in an interval of size significantly smaller than m, contradicting the fact that it reduces to an almost-uniform element of $\mathbb{Z}_m$.

Our extractor is the following:

Theorem 2. *There is a universal constant c such that the following holds. Let X be an α-unpredictable bit source on $\{0,1\}^n$, let Y be a source on $\{0,1\}^n$ with min-entropy at least βn (independent from X), and let $Z = \langle X, Y \rangle \bmod m$ for some $m \in \mathbb{N}$. Then for every $\delta \in [0,1]$, the random variable (Y,Z) is δ-close to (Y,U) where U is uniform on $\mathbb{Z}_m$ and independent of Y, provided that*

$$n \geq c \cdot \frac{m^2}{\alpha\beta} \cdot \log\left(\frac{m}{\beta}\right) \cdot \log\left(\frac{m}{\delta}\right).$$

Notice that for constant α, β, and δ, we can take m as large as $\Omega(\sqrt{n}/\log n)$ and satisfy the condition of the theorem. Note also that the output Z is guaranteed to be close to uniform even given the source Y. Two-source extractors with this property have been studied in several papers, starting with [26].

The first step is to reduce proving near-uniformity of the extractor's output distribution Z to bounding the magnitude of its Fourier coefficients $\mathrm{E}[\omega^Z]$:

Lemma 3. *Let Z be a random variable taking values in $\mathbb{Z}_m$. Then the statistical distance between Z and the uniform distribution on $\mathbb{Z}_m$ is at most*

$$\frac{1}{2}\sqrt{\sum_{\omega \neq 1} \left|\mathrm{E}\left[\omega^Z\right]\right|^2},$$

where the sum is over all complex m'th roots of unity ω other than 1.

Proof: Let U be a uniformly distributed random variable in $\mathbb{Z}_m$. Let $p_Z(\cdot)$ and $p_U(\cdot)$ denote the probability mass

function of Z and U respectively. We have

$$\begin{aligned}\|Z-U\|_{SD} &= \frac{1}{2}\|p_Z - p_U\|_1 \le \frac{\sqrt{m}}{2}\|p_Z - p_U\|_2 \\ &= \frac{1}{2}\sqrt{\sum_{k=0}^{m-1} |\hat{p_Z}(k) - \hat{p_U}(k)|^2}\ .\end{aligned}$$

Plugging in the Fourier coefficients $\hat{p_Z}(0)$ and $\hat{p_U}(\cdot)$, the claim follows. ∎

Next, instead of estimating the Fourier coefficients of the output $Z = \langle X, Y\rangle \bmod m$ when both sources X and Y are random, we fix $Y = y$ and argue that there are not many y's for which the Fourier coefficients are large. To get a good bound on the number of y's, we estimate the $2t$'th moment of the Fourier coefficients. The following lemma is proved in the full version of the paper.

Lemma 4. *Let X be any random variable taking values in $\{0,1\}^n$, $\omega \in \mathbb{C}$ a* primitive *m'th root of unity, and $t \in N$. Then*

$$\sum_{y\in\mathbb{Z}_m^n} \left|\mathrm{E}\left[\omega^{\langle X,y\rangle}\right]\right|^{2t} \le \left[1 + m\exp\left(-\Omega\left(\frac{\alpha t}{m^2}\right)\right)\right]^n.$$

We will apply this taking t a bit larger than m^2/α, so that the $\exp(-\Omega(\alpha t/m^2))$ term is small. We now put the above pieces together to obtain our extractor:

Proof of Theorem 2: Let X be an α-unpredictable bit source on $\{0,1\}^n$, Y a βn-source on $\{0,1\}^n$. For every complex m'th root of unity $\omega \neq 1$, let

$$L_\omega = \left\{y \in \{0,1\}^n : \left|\mathrm{E}\left[\omega^{\langle X,y\rangle}\right]\right| > \frac{\delta}{\sqrt{m}}\right\},$$

and let $L = \bigcup_\omega L_\omega$. By Lemma 3, it holds that for every $y \notin L$, the statistical distance between $Z_y = \langle X, y\rangle \bmod m$ and the uniform distribution on $\mathbb{Z}_m$ is at most $(1/2)\sqrt{(m-1)\cdot(\delta/\sqrt{m})^2} \le \delta/2$. Thus it suffices to prove that $\Pr[Y \in L] \le \delta/2$, which in turn follows if $\Pr[Y \in L_\omega] \le \delta/2m$ for each $\omega \neq 1$.

Every m'th root of unity $\omega \neq 1$ is a primitive ℓ'th root of unity for some $\ell | m$. By Lemma 4, we have

$$\begin{aligned}|L_\omega| &\le \frac{\sum_{y\in\mathbb{Z}_\ell^n}\left|\mathrm{E}\left[\omega^{\langle X,y\rangle}\right]\right|^{2t}}{(\delta/\sqrt{m})^{2t}} \\ &\le \frac{\left[1+\ell\cdot\exp(-\Omega(\alpha t/\ell^2))\right]^n}{(\delta^2/m)^t} \\ &\le \frac{\left[1+m\cdot\exp(-\Omega(\alpha t/m^2))\right]^n}{(\delta^2/m)^t} \\ &\le \frac{2^{\beta n/2}}{(\delta^2/m)^t}.\end{aligned}$$

for $t = \lceil c_0\cdot(m^2/\alpha)\cdot\log(m/\beta)\rceil$ for a sufficiently large universal constant c_0.

Thus, by the union bound.

$$\Pr[Y \in L_\omega] \le 2^{-\beta n}\cdot|L_\omega| \le \frac{2^{-\beta n/2}}{(\delta^2/m)^t} \le \frac{\delta}{2m},$$

provided that $n \ge (2/\beta)\cdot(t\cdot\log(m/\delta^2) + \log(2m/\delta))$, which holds by hypothesis. ∎

We now combine the fact that the inner product modulo m is good extractor for unpredictable sources with the connections between differentially-private protocols and unpredictable sources to show that no differentially-private protocol can estimate inner product to within error $o(\sqrt{n}/\log n)$:

Theorem 5. *Let P be a randomized protocol with ϵ-differential privacy and let $\delta > 0$. Then with probability at least $1-\delta$ over the inputs $x, y \leftarrow \{0,1\}^n$ and the coin tosses of P, party B's output differs from $\langle x, y\rangle$ by at least*

$$\Delta = \Omega\left(\frac{\sqrt{n}}{\log n}\cdot\frac{\delta}{e^\epsilon}\right).$$

Proof: Let X and Y be uniform and independent in $\{0,1\}^n$ and Π be the communication transcript. Party B's output is a function $f_B(Y,\Pi)$. Let $m = 4\Delta/\delta$.

By Lemma 1, we know that for every $\pi \in \mathrm{Supp}(\Pi)$, X_π and Y_π are independent α-unpredictable sources for $\alpha = e^{-\epsilon}$. This implies that Y_π has min-entropy at least βn for $\beta = \log(1+\alpha) \ge \alpha$. By Theorem 2, $(Y_\pi, \langle X_\pi, Y_\pi\rangle \bmod m)$ has statistical distance at most $\delta/2$ from (Y_π, U), provided

$$n \ge c_0\cdot\frac{m^2}{\alpha\beta}\cdot\log\left(\frac{m}{\beta}\right)\cdot\log\left(\frac{m}{\delta}\right),$$

for a universal constant c_0. Using the fact that $m = 4\Delta/\delta$ and $\beta \ge \alpha$, this follows if:

$$n \ge c_1\cdot\left[\frac{\Delta\cdot e^\epsilon}{\delta}\cdot\log\left(\frac{\Delta\cdot e^\epsilon}{\delta}\right)\right]^2,$$

for some universal constant c_1, which in turn follows if

$$\frac{\Delta\cdot e^\epsilon}{\delta} \le c_2\cdot\frac{\sqrt{n}}{\log n},$$

for a small universal constant $c_2 > 0$.

Consider the set $S = \{(\pi, y, z) : (f_B(\pi, y) - z) \bmod m \in \{m-\Delta, \ldots, m-1, 0, 1, \ldots, \Delta\}\}$. Notice that in every execution where B's output $f_B(\pi, y)$ differs from $\langle x, y\rangle$ by at most Δ, we have $(\pi, y, \langle x, y\rangle \bmod m) \in S$. We can bound the probability of this occurring by using the fact that $(\Pi, Y, \langle X, Y\rangle \bmod m)$ has statistical distance at most $\delta/2$ from (Π, Y, U). Specifically, we have:

$$\begin{aligned}\Pr[(\Pi, Y, \langle X, Y\rangle \bmod m) \in S] &\le \Pr[(\Pi, Y, U) \in S] + \delta/2 \\ &\le 2\Delta/m + \delta/2 = \delta\ .\end{aligned}$$

∎

This theorem implies a similar result for the Hamming distance, because the inner product between two bitstrings $x, y \in \{0,1\}^n$ can be expressed as $\langle x, y\rangle = |x|_H + |y|_H - |x-y|_H$. Thus, a differentially-private protocol for estimating the Hamming distance $|x-y|_H$ can be turned into one for the inner product by having the parties send differentially-private estimates of the Hamming weights of their inputs.

C. Limitation of the Extractor Technique

The deterministic extractor approach above depends crucially on the fact that the x_i's are independent, or nearly independent of each other. We observe that standard measure concentration techniques imply that such a technique cannot go beyond $\sqrt{n}$ for any function with sensitivity 1.

Theorem 6. *Let $f\colon \{0,1\}^n \times \{0,1\}^n \to \mathbb{R}$ be a sensitivity-1 function. Then for every distribution μ such that for every input y, the conditional distribution $\mu(X \mid Y = y)$ is a product distribution $\Pi_{i=1}^n \mu_i(X_i \mid Y = y)$, there is function $g(y)$ such that $\Pr_{(x,y)\sim\mu}[|g(y) - f(x,y)| > t] \leq 2\exp(-t^2/2n)$.*

Proof: Standard martingale concentration results (see e.g. [27]) say that every sensitivity-1 function on a product distribution is well concentrated around its expectation. Specifically, for every $h\colon \{0,1\}^n \to \mathbb{R}$, and every product distribution ν on X,

$$\Pr[|h(x) - \mathrm{E}_{x\sim\nu}[h(x)]| > t] \leq 2\exp(-t^2/2n).$$

Applying this result to the function $f(X,y)$ and setting $g(y) = \mathrm{E}_{x\in\mu(X|Y=y)}[f(x,y)]$ yields the result. ∎

In other words, $f(x,y)$ can be computed by Bob up to an expected additive error of $O(\sqrt{n})$ without any communication, provided that Alice's input comes from a product distribution (conditioned on Bob's). Since the connection to unpredictable bit sources (Lemma 1) requires that the inputs come from a product distribution, we cannot get a lower bound better than $\Theta(\sqrt{n})$ from that approach.

IV. Differential Privacy and Communication Complexity

In this section we characterize differentially-private protocols in terms of their communication complexity. Specifically, we present bounds on the information cost and the partition bound in terms of the privacy parameter (Section IV-A and IV-B). We then prove stronger separations between information-theoretic and computational differential privacy (Section IV-C). We also note that the message compression technique of Barak et al. [16], implies that all differentially-private protocols are compressible.

Furthermore, we show that if there exists a differentially-private protocol with a constant number of rounds, it can be compressed *while keeping it differentially private* (Section IV-D). Finally, we show that low-communication protocols can be converted into privacy-preserving ones with some loss of accuracy (Section IV-E).

A. Differential Privacy and Information Cost

As a first tool of proving feasibility of differentially-private protocol with certain accuracy, we establish a connection between differential privacy and the concept of *information cost* as defined by Bar-Yossef et al. [14] (based on a earlier concept introduced by Chakrabarti et al. [28].)

The definition of information cost is based on the following standard definitions of mutual information and conditional mutual information:

Definition 8 (Mutual Information)**.** *Given two random variables X and Y over the same probability space, their* mutual information *is defined as follows:*

$$\mathsf{I}(X;Y) = H(X) - H(X \mid Y),$$

where H is the entropy. The conditional mutual information *is $\mathsf{I}(X;Y \mid Z) = H(X \mid Z) - H(X \mid YZ)$.*

Intuitively, $\mathsf{I}(X;Y)$ captures the amount of information shared by two variables. For example, if the variables are identical, their mutual information equals their entropy; if they are independent, it is zero. Mutual information motivates the definition of *information cost* for protocols, which corresponds to the amount of information that is learnt about the players' inputs from the messages communicated.

Definition 9 (Information Cost)**.** *Given a distribution μ over inputs X and Y to the two parties of protocol P, we define* information cost *of P for distribution μ as*

$$\mathsf{ICost}_\mu(P) = \mathsf{I}(XY;\Pi(X,Y)),$$

where $\Pi(X,Y)$ is the random transcript of the protocol on input (X,Y).

By the definition of differential privacy, none of the input bits in a differentially-private protocol are fully revealed to the other party. This implies the following natural bound on the information cost of a differentially-private protocol. We defer the proof to the full version.

Proposition 7. *If $P(x,y)$ has ϵ-differential privacy, where $x,y \in \Sigma^n$ for a finite alphabet Σ, then for every distribution μ on $\Sigma^n \times \Sigma^n$, the information cost of P is bounded as follows:*

$$\mathsf{ICost}_\mu(P) \leq 3\epsilon n.$$

If $\Sigma = \{0,1\}$ and μ is the uniform distribution, then the bound can be improved to $\mathsf{ICost}_\mu(P) \leq 1.5\epsilon^2 n$.

The following proposition, also proved in the full version, extends this result to (ϵ,δ)-differential privacy.

Proposition 8. *If $P(x,y)$ has (ϵ,δ)-differential privacy, where $x,y \in \Sigma^n$ for a finite alphabet Σ and $\epsilon < 1$, then for every distribution μ on $\Sigma^n \times \Sigma^n$, the information cost of P is bounded as follows:*

$$\mathsf{ICost}_\mu(P) \leq (10\epsilon + 6\delta|\Sigma| \cdot (\log|\Sigma|)/\epsilon)n.$$

Compressing Differentially-Private Protocols The information cost of protocol is closely related to the communication complexity since $\mathsf{I}(XY;\Pi(X,Y)) \leq H(\Pi(X,Y)) \leq |\Pi(X,Y)|$ for every distribution on X and Y. Barak et al. recently proved a bound in the other direction.

Theorem 9 (Barak et al. [16])**.** *For every product distribution μ, for every protocol randomized P with output $\mathrm{out}(P)$, and every $\gamma > 0$, there exists functions f_A, f_B, and protocol Q*

$$
\begin{array}{lrll}
\text{Min. } \sum_{z\in\mathcal{Z}}\sum_{R\in\mathcal{R}} w_{z,R} & & & \\
\text{subject to} & & & \\
\quad \sum_{R:(x,y)\in R} w_{f(x,y),R} & \geq & 1-\gamma & \forall (x,y)\in \mathrm{Supp}(f) \\
\quad \sum_{R:(x,y)\in R}\sum_{z\in\mathcal{Z}} w_{z,R} & = & 1 & \forall (x,y)\in \mathcal{X}\times\mathcal{Y} \\
\quad w_{z,R} & \geq & 0 & \forall z\in\mathcal{Z}, R\in\mathcal{R}
\end{array}
$$

Fig. 1. Linear program for the Partition Bound for a function f.

such that

$$
\begin{aligned}
\|f_A(X,Q(X,Y)) - \mathrm{out}(P)\|_{SD} &< \gamma\ , \\
\Pr[f_A(X,Q(X,Y)) \neq f_B(Y,Q(X,Y))] &< \gamma\ , \textit{ and} \\
\mathsf{IContent}_\mu(P)\gamma^{-1}\mathrm{polylog}(\mathsf{CC}(P)/\gamma) &\geq \mathsf{CC}(Q)\ ,
\end{aligned}
$$

where $\mathsf{IContent}_\mu(P) = \mathsf{I}(X;\Pi(X,Y)\mid Y) + \mathsf{I}(Y;\Pi(X,Y)\mid X)$ *which satisfies* $\mathsf{IContent}_\mu(P) = O(\mathsf{ICost}_\mu(P))$.

It follows that differentially private protocols can be compressed.

Theorem 10. *Let P be an ϵ-differentially private protocol P with output $\mathrm{out}(P)$ where the input (X,Y) is distributed according to an arbitrary product distribution μ. Then for every $\gamma > 0$, there exists functions f_A, f_B, and a protocol Q such that $\|f_A(X,Q(X,Y)) - \mathrm{out}(P)\|_{SD} < \gamma$, $\Pr[f_A(X,Q(X,Y)) \neq f_B(Y,Q(X,Y))] < \gamma$ and $\mathsf{CC}(Q) \leq 3\epsilon\gamma^{-1}n\cdot\mathrm{polylog}(\mathsf{CC}(P)/\gamma)$.*

B. Differential Privacy and the Partition Bound

Jain and Klauck [15] define the *partition bound* for a partial function $f\colon \mathcal{X}\times\mathcal{Y}\to\mathcal{Z}$. This bound is given by the linear program in Figure 1. Here $\mathcal{R} = 2^{\mathcal{X}}\times 2^{\mathcal{Y}}$ is the set of all rectangles in $\mathcal{X}\times\mathcal{Y}$. Denoting by $prt_\gamma(f)$ the optimum of this linear program, Jain and Klauck show that every randomized γ-error public coin protocol computing f has communication complexity at least $\log prt_\gamma(f)$. Moreover, they showed that this lower bound dominates most other lower bounding techniques in randomized communication complexity such as (smooth) rectangle bound and (smooth) discrepancy bound (see [15] for precise definitions of these bounds).

In this subsection, we show that for any differentially private protocol computing a partial function f, the value of the partition bound is small. Thus a proof that f has large communication complexity using the partition bound also shows that f has no ϵ-differentially private protocol for some ϵ. Since the definition of the partition bound assumes that the transcript determines the output of the protocol (this is without loss of generality in communication protocols, but not necessarily so in private communication protocols), we assume that this is the case for the private protocol. A similar result can be proved without this assumption for an appropriately modified linear program.

We also note that considering *partial* functions allows us to also capture protocols that compute approximations (as is typically the case for differentially private protocols). For example, a differentially private protocol that computes function g to within additive error α whp yields, for any threshold t, a differentially private protocol that computes the partial function f whp, where $f(x,y) = 1$ when $g(x,y) > t+\alpha$ and $f(x,y) = 0$ when $g(x,y) < t-\alpha$.

The proof of the following result is deferred to the full version.

Theorem 11. *Suppose that an ϵ-differentially private protocol P computes a partial function $f\colon \{0,1\}^n\times\{0,1\}^n\to\mathcal{Z}$ with error probability at most γ. Then $\log prt_\gamma(f) \leq 3\epsilon n$.*

C. A Stronger Separation

In this section, we show that for worst case error, the gap between computational and information-theoretic differential privacy is essentially as large as possible. We first argue that there are low sensitivity functions such that any protocol approximating the function to a additive linear error must incur linear information cost.

Theorem 12. *There exists an absolute constant $\beta > 0$ such that for every m, there is an efficiently computable function $f\colon \{0,1\}^m\times\{0,1\}^m\to\mathbb{R}$ and distribution $\mathcal{D}$ over $\{0,1\}^m\times\{0,1\}^m$ with the following properties:*
(a) every protocol that outputs a βm additive approximation to f with probability at least $\frac{9}{10}$ over inputs from $\mathcal{D}$ must have information cost at least βm.
(b) f has sensitivity 1, i.e., $|f(x,y) - f(x',y')| \leq |(x,y) - (x',y')|_H$ for every x,y,x',y'.

Proof: We show that given a predicate function $g\colon \{0,1\}^n\times\{0,1\}^n\to\{0,1\}$ and distribution $\mathcal{D}_g$ over its inputs, we can transform it to a a sensitivity-1 function $f_g : \{0,1\}^m\times\{0,1\}^m\to\mathbb{R}$, and a distribution $\mathcal{D}$ over its inputs, for $\frac{m}{n}$ constant. This transformation has the property that every protocol approximating f_g within error cm (for some constant $c > 0$ to be determined) with probability $(1-\gamma)$ over $\mathcal{D}$ has information cost at least $\mathsf{ICost}_{\mathcal{D}_g,\gamma}(g)$. Plugging in a function g and distribution $\mathcal{D}_g$ with large information cost would then imply the result.

We embed a large multiple of g in a low-sensitivity function f_g. We do so by first defining f_g on the set of well-separated points $C\subseteq\{0,1\}^n$, where C is the set of codewords of a code with linear distance. Low sensitivity is then ensured by interpolating the value of f_g appropriately.

Let $Enc\colon \{0,1\}^n\to\{0,1\}^m$ be an encoding algorithm for a linear-rate error-correcting code C with a decoding algorithm $Dec\colon \{0,1\}^m\to\{0,1\}^n$ that works up to decoding radius δm for some constant δ. Such codes exist with $n = rm$ for some constant $r = r(\delta) > 0$. Let $d(x,C)$ be the distance from x to the closest codeword in C. We then define $f_g(x,y) =$

$$
\begin{cases}
g(Dec(x),Dec(y))\cdot(\delta m - d((x,y),C)) & \text{if } d((x,y),C)\leq\delta m, \\
0 & \text{otherwise,}
\end{cases}
$$

where we have used $d((x,y),C)$ as shorthand for $d(x,C)+d(y,C)$. Note that when x and y are both codewords, $f_g(x,y)$ is exactly $\delta m\cdot g(Dec(x),Dec(y))$. As we move away from C, $f_g(x,y)$ smoothly decays to 0. Moreover, since C has decoding radius δm, the function is well-defined and efficiently computable: if any of $Dec(x)$ or $Dec(y)$ fails to decode, it

means that $d(x,C)+d(y,C)>\delta m$ and the function is zero by definition.

The distribution $\mathcal{D}$ is concentrated on the codewords with $p_{\mathcal{D}}(Enc(x),Enc(y))=p_{\mathcal{D}_g}(x,y)$.

We first argue that any communication protocol P_{f_g} approximating f_g to within error less than $\delta m/2=\delta n/(2r)$ (with probability $(1-\gamma)$ over $\mathcal{D}$) yields a γ-error communication protocol P_g for g on distribution $\mathcal{D}_g$, with the same communication complexity. This is done in the natural way: in P_g, Alice and Bob on input (x,y) simply run the protocol P_{f_g} on inputs $x'=Enc(x)$ and $y'=Enc(y)$, and Alice outputs 0 if her output $f_A(x',\Pi_{P_{f_g}})$ in protocol P_{f_g} is smaller than $\delta m/2$, and 1 otherwise. Since $f_g(x',y')$ is equal to $\delta m\cdot g(x,y)$, if P_{f_g} has error less than $\delta m/2$ on (x',y'), then

$$\left|f_A(x',\Pi_{P_{f_g}})-f_g(x',y')\right|<\frac{\delta m}{2},$$

in which case Alice's output of P_g on (x,y) is exactly $g(x,y)$. A similar claim holds for Bob. From the definition of $\mathcal{D}$ it follows that the failure probability of P_g is the same as that of P_{f_g}.

Next we bound the sensitivity of f_g. Let (x_1,y_1) and (x_2,y_2) be neighboring inputs and assume without loss of generality that $y_1=y_2$. The main observation is that $f_g(\cdot,y_1)$ is zero except for small neighborhoods around certain codewords, i.e., except for $\cup_{x:g(x,Dec(y_1))=1}B(Enc(x),\delta m-d(y_1,C))$. It is easily seen to have sensitivity 1 within each ball, since $d(x,C)$ has sensitivity 1. Since the decoding radius of C is δm, these balls are disjoint. As f_g is zero on the boundary of these balls, the sensitivity is at most 1 everywhere.

Finally, plugging in any function g which has $\Omega(n)$ information cost, e.g., the inner product function [14], we get the desired result. ∎

Combining this result with Proposition 7, we conclude

Theorem 13. *There exists an absolute constant $\beta>0$ such that for every n, there is an efficiently computable function $f\colon\{0,1\}^n\times\{0,1\}^n\to\mathbb{R}$ and a distribution $\mathcal{D}$ over its inputs, with the following properties:*
(a) for every $\epsilon<\beta/3$, every ϵ-differentially private protocol P must incur additive error at least βn with probability at least $\frac{1}{10}$.
(b) for every ϵ,δ with $\epsilon+1.2(\delta/\epsilon)<\beta/10$, every (ϵ,δ)-differentially private protocol P must incur additive error at least βn with probability at least $\frac{1}{10}$.
(c) f has sensitivity 1, i.e., $|f(x,y)-f(x',y')|\le|(x,y)-(x',y')|_H$ for every x,y,x',y'.

D. Private Message Compression

In this section, we argue that for protocols with a constant number of rounds, compression can be done while maintaining differential privacy. The basic idea is for Alice (resp. Bob) to use consistent sampling (dart throwing) [18], [19] from the distribution μ_x (resp. μ_y) to pick a message to be sent. Instead of sending the message itself which may be arbitrarily long, Alice and Bob use shared randomness to pick the darts, so that it suffices to send the index of the dart picked. We argue that this can be done privately with small communication. We defer the proof to the full version.

Theorem 14. *Let P be an ϵ-differentially private protocol with r rounds. Then for every $\delta>0$, there exists an $O(r\epsilon)$-differentially-private protocol P^* that has communication complexity $O(r(\epsilon n+\log\log\frac{1}{\epsilon\delta}))$ and except with probability $r\delta$, simulates P perfectly. In other words, there exist functions π_x, π_y such that $\Pr[\pi_x(\mathsf{VIEW}^A_{P^*}(x,y))=\mathsf{VIEW}^A_P(x,y)]\ge 1-r\delta$, and similarly for B.*

E. From Low Communication to Privacy

The previous sections show that, loosely speaking, differential privacy implies low communication complexity. In this section we demonstrate the converse: if there exists a protocol for computing a sensitivity-1 function, the function can be approximated in a differentially-private manner with error proportional to the communication and round complexity of the original protocol. The lower bound proven in Section IV-C suggests that the linear dependency on the communication complexity is best possible, at least without further restrictions on the functionality, as there are sensitivity-1 functions that can be computed exactly using communication C but cannot be approximated by any differentially-private protocol with error better than $\Omega(C)$.

Given a deterministic protocol for computing the sensitivity-1 function $f(x,y)$ we construct an ϵ-differentially-private protocol by sampling messages of the new protocol using the exponential mechanism [29]. The following result is proven in the full version of the paper.

Theorem 15. *Let P be a deterministic protocol with communication complexity $\mathsf{CC}(P)$ approximating a sensitivity-1 function $f\colon\Sigma^n\times\Sigma^n\to\mathbb{Z}$ with error bounded by Δ. Then there exists an ϵ-differentially-private protocol with the same communication complexity and the number of rounds which computes f with expected additive error $\Delta+O(\mathsf{CC}(P)r/\epsilon)$.*

V. Conclusions and Open Problems

We have investigated the limitations of two-party differential privacy and exposed interesting connections to deterministic extractors for Santha-Vazirani sources, as well as to communication complexity. In our first result we prove a lower bound on accuracy of approximating the Hamming distance between two vectors—the classical problem in two-party computations—with two-sided guarantees of differential privacy. The lower bound on the additive error, which is tight up to a logarithmic factor, is proportional to $\tilde{\Omega}(\sqrt{n})$ and matches the recently obtained bound on accuracy of sublinear communication protocols [17]. The connection between differential privacy and communication complexity seems to be a genuine phenomenon, exemplified by the following results:

— We present bounds on the information cost and the partition bound in terms of the privacy parameter. The information cost bound, in combination with the message compression technique of Barak et al. [16], implies that all differentially-private protocols are compressible.

Furthermore, using existing bounds on the information cost of specific communication problems allows us to construct a function that exhibits the largest possible gap between accuracy of optimal differentially-private and computationally differentially-private protocols.

— Any deterministic protocol can be converted into a differentially-private one with accuracy proportional to its communication complexity and the number of rounds.

This work raises many new questions. There are connections between two-party differential privacy and *pan-privacy* [30]. A pan-private algorithm requires not only that its output be differentially private, but also that the internal state be differentially private as well. In other words, the algorithm must be privacy-preserving both inside and out. Such algorithms can be viewed as one-pass streaming algorithms, where the internal state is privacy-preserving at each point in time. (For streaming purposes the size of the internal state should also be kept small.) In [30], many important and natural statistics, such as density estimation, were shown to be computable pan-privately and with reasonable accuracy.

Our lower bound on the two-party complexity of the Hamming distance function implies a lower bound on *multi-pass* pan-private algorithms for density estimation, as well as for other natural statistics. (While not defined in [30], it is also natural to consider multi-pass pan-private algorithms.) Indeed, by a straightforward reduction, a pan-private algorithm for density estimation implies a two-party protocol for estimating the Hamming distance of two binary strings, with similar error. What further limitations for pan-privacy can be obtained?

Finally, we would like to strengthen Theorems 14 and 15 to be independent of the number of rounds of communication, and extend Theorem 15 to randomized protocols.

REFERENCES

[1] M. O. Rabin, "How to exchange secretes with oblivious transfer," Aiken Computation Lab, Harvard University, Technical Report TR-81, May 1981.

[2] S. Even, O. Goldreich, and A. Lempel, "A randomized protocol for signing contracts." *Communications of the ACM*, vol. 28, no. 6, pp. 637–647, 1985.

[3] E. Kushilevitz and R. Ostrovsky, "Replication is NOT needed: SINGLE database, computationally-private information retrieval," in *Proceedings of the 38th Annual Symposium on Foundations of Computer Science (FOCS 1997)*. IEEE Computer Society, 1997, pp. 364–373.

[4] A. C.-C. Yao, "Protocols for secure computations (extended abstract)," in *23rd Annual Symposium on Foundations of Computer Science (FOCS 1982)*. IEEE, 1982, pp. 160–164.

[5] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006*, ser. Lecture Notes in Computer Science, S. Halevi and T. Rabin, Eds., vol. 3876. Springer, 2006, pp. 265–284.

[6] C. Dwork, "Differential privacy," in *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Part II*, ser. Lecture Notes in Computer Science, M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, Eds., vol. 4052. Springer, 2006, pp. 1–12.

[7] A. Ghosh, T. Roughgarden, and M. Sundararajan, "Universally utility-maximizing privacy mechanisms," in *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009*, M. Mitzenmacher, Ed. ACM, 2009, pp. 351–360.

[8] C. Dwork and K. Nissim, "Privacy-preserving datamining on vertically partitioned databases," in *Advances in Cryptology—CRYPTO 2004*, ser. Lecture Notes in Computer Science, M. K. Franklin, Ed., vol. 3152. Springer, 2004, pp. 528–544.

[9] A. Beimel, K. Nissim, and E. Omri, "Distributed private data analysis: Simultaneously solving how and what," in *Advances in Cryptology—CRYPTO 2008*, ser. Lecture Notes in Computer Science, D. Wagner, Ed., vol. 5157. Springer, 2008, pp. 451–468.

[10] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, "Our data, ourselves: privacy via distributed noise generation," in *Advances in Cryptology—EUROCRYPT 2006*, ser. Lecture Notes in Computer Science, vol. 4004. Springer, 2006, pp. 486–503.

[11] I. Mironov, O. Pandey, O. Reingold, and S. P. Vadhan, "Computational differential privacy," in *CRYPTO*, ser. Lecture Notes in Computer Science, S. Halevi, Ed., vol. 5677. Springer, 2009, pp. 126–142.

[12] U. V. Vazirani, "Strong communication complexity or generating quasirandom sequences from two communicating semirandom sources," *Combinatorica*, vol. 7, no. 4, pp. 375–392, 1987. [Online]. Available: http://dx.doi.org/10.1007/BF02579325

[13] S. L. Warner, "Randomized response: A survey technique for eliminating evasive answer bias," *Journal of the American Statistical Association*, vol. 60, no. 309, pp. 63–69, Mar. 1965.

[14] Z. Bar-Yossef, T. S. Jayram, R. Kumar, and D. Sivakumar, "An information statistics approach to data stream and communication complexity," in *43rd Symposium on Foundations of Computer Science, FOCS 2002*. IEEE Computer Society, 2002, pp. 209–218.

[15] R. Jain and H. Klauck, "The partition bound for classical communication complexity and query complexity," in *Annual IEEE Conference on Computational Complexity*, 2010.

[16] B. Barak, M. Braverman, X. Chen, and A. Rao, "How to compress interactive communication," in *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010*, L. J. Schulman, Ed. ACM, 2010, pp. 67–76.

[17] A. Chakrabarti and O. Regev, "Tight lower bound for the Gap Hamming problem," 2010, unpublished manuscript.

[18] U. Manber, "Finding similar files in a large file system," in *USENIX Winter*, 1994, pp. 1–10.

[19] T. Holenstein, "Parallel repetition: Simplification and the no-signaling case," *Theory of Computing*, vol. 5, no. 1, pp. 141–172, 2009. [Online]. Available: http://www.theoryofcomputing.org/articles/v005a008

[20] J. Feigenbaum, A. D. Jaggard, and M. Schapira, "Approximate privacy: foundations and quantification (extended abstract)," in *ACM Conference on Electronic Commerce*, 2010, pp. 167–178.

[21] B. Chor and E. Kushilevitz, "A zero-one law for boolean privacy," *SIAM J. Discrete Math.*, vol. 4, no. 1, pp. 36–47, 1991.

[22] E. Kushilevitz, "Privacy and communication complexity," *SIAM J. Discrete Math.*, vol. 5, no. 2, pp. 273–284, 1992.

[23] A. C.-C. Yao, "Some complexity questions related to distributive computing (preliminary report)," in *STOC '79: Proceedings of the eleventh annual ACM symposium on Theory of computing*. New York, NY, USA: ACM, 1979, pp. 209–213.

[24] M. Sántha and U. V. Vazirani, "Generating quasirandom sequences from semirandom sources," *Journal of Computer and System Sciences*, vol. 33, no. 1, pp. 75–87, 1986. [Online]. Available: http://dx.doi.org/10.1016/0022-0000(86)90044-9

[25] O. Reingold, S. Vadhan, and A. Wigderson, "A note on extracting randomness from Santha–Vazirani sources," 2004, unpublished manuscript.

[26] Y. Dodis and R. Oliveira, "On extracting private randomness over a public channel," in *RANDOM-APPROX*, ser. Lecture Notes in Computer Science, S. Arora, K. Jansen, J. D. P. Rolim, and A. Sahai, Eds., vol. 2764. Springer, 2003, pp. 252–263.

[27] D. Dubhashi and A. Panconesi, *Concentration of Measure for the Analysis of Randomised Algorithms*. Cambridge University Press, 2009.

[28] A. Chakrabarti, Y. Shi, A. Wirth, and A. C. Yao, "Informational complexity and the direct sum problem for simultaneous message complexity," in *Proc. 42nd Annual IEEE Symposium on Foundations of Computer Science*, 2001, pp. 270–278.

[29] F. McSherry and K. Talwar, "Mechanism design via differential privacy," in *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007)*. IEEE Computer Society, 2007, pp. 94–103.

[30] C. Dwork, M. Naor, T. Pitassi, G. Rothblum, and S. Yekhanin, "Pan-private streaming algorithms," in *Proceedings of the First Symposium on Innovations in Computer Science (ICS 2010)*. Tsinghua University Press, Beijing, 2010.