

The Limits of Two-Party Differential Privacy

Andrew McGregor* Ilya Mironov[†] Toniann Pitassi[‡] Omer Reingold[†]
Kunal Talwar[†] Salil Vadhan[§]

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.)

*Department of Computer Science, University of Massachusetts, Amherst. Support by NSF CAREER Award CCF-0953754. Work done in part while visiting Microsoft Research.

[†]Microsoft Research Silicon Valley.

[‡]Department of Computer Science, University of Toronto. Research supported by NSERC. Work done in part while visiting Microsoft Research.

[§]School of Engineering and Applied Sciences and Center for Research on Computation and Society, Harvard University. Supported by NSF grant CNS-0831289. Work done in part while visiting Microsoft Research.

1 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 [Rab81, EGL85], single-server private-information retrieval (PIR) [KO97], and more generally, secure function evaluation (SFE) [Yao82], 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* [DMNS06, Dwo06]. 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 [DMNS06], and indeed this mechanism can be shown to be optimal for counting queries as well as for a large class of clients' preferences [GRS09].

Two-party differential privacy. In this paper, we contrast the client-server setting with a setting where the database is distributed between two parties who 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 [DN04]. More accurately, the definition of privacy in [DN04] 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 [BNO08].

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 [DN04, DKM⁺06, MPRV09]), there were no *general* results or lower bounds for computing functions 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 [DN04], 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.* This result also extends to the commonly used notion of approximate differential privacy (specifically, (ϵ, δ) -differential privacy for $\delta = o(1/n)$).

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 [MPRV09]. 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 for two-party protocols. This 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. (Indeed, subsequent to our work, Groce, Katz, and Yerukhimovich [GKY11] have shown that a wide class of computationally differentially private mechanisms in the client-server setting can be converted into ones that achieve information-theoretic differential privacy with essentially the same accuracy and efficiency.) 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.

It is natural to ask if the above separations can be made even stronger:

Question 2. *What is the largest gap in accuracy between two-party and client-server differentially private protocols?*

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 approximated within error $o(n)$ by a 2-party differentially private protocol. Unlike our result about Hamming distance, here our proof does not apply to approximate differential privacy. (In the conference version [MMP⁺10], we claimed that it did apply, but there was an error in the proof and a counterexample to the key lemma was found by De [De11].) Thus, this result does not provide a separation between information-theoretic and computational differential privacy, since the information-theoretic analogue of computational differential privacy is approximate differential privacy. Whether a stronger separation can be proved remains an intriguing open question:

(Open) Question 3. *What is the largest gap in accuracy between information-theoretic and computationally differentially private 2-party protocols?*

The techniques we develop to address the above questions rely on new connections with other topics in the theory of computation: the first is a connection between differential privacy in the two-party setting and

¹In [DN04], a central data-mining problem (detecting correlations between two binary attributes) was reduced to approximating the scalar product between two binary vectors.

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 speaking, 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 these connections.

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 [Vaz87] 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 [War65] 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 3.9 (Section 3). *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 naively) a lower bound on additive error better than $O(\sqrt{n})$ (see Section 3.3). 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. [BYJKS02]) and the *partition bound* (as defined by Jain and Klauck [JK10]) 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. [BBCR10] 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 4.8).

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 and client-server

differential privacy is linear: On the one hand, by property (1) the function can be approximated with differential privacy by a trusted third party (or server) with error proportional to $1/\epsilon$. On the other hand, every (information-theoretic) differentially private protocol has linear additive error. More precisely, the following theorem claims these properties of our construction:

Theorem 4.11 (Section 4.3). *There exists an absolute constant $\beta > 0$ such that for every n , there is an efficiently computable function $f: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \mathbb{R}$ and a distribution \mathcal{D} over its inputs, with the following properties*

- (a) *for every $\epsilon < \beta/3$, every ϵ -differentially private protocol P has expected additive error at least βn .*
- (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' .*

We note that the connection between differential privacy and the partition bound can be used to prove lower bounds on the error in computing specific functions for which lower bounds on the partition bound is known. For the specific example of Hamming distance (see [CR11]), the results obtained in this manner are incomparable with those obtained via deterministic extraction; we get a error bound of \sqrt{n} instead of $\Omega(\sqrt{n}/\log n)$ but it applies only for pure (not approximate) differential privacy.

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 $O(\epsilon rn)$ while keeping it differentially private. Compression is implemented using privacy-preserving consistent sampling [Man94, Hol09] and has a negligible probability of failure (which affects accuracy, not privacy). The formal theorem is stated as follows:

Theorem 4.12 (Section 4.5). *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 $O(r \cdot (\epsilon n + \log \log \frac{1}{\delta}))$ 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 4.14 (Section 4.6). *Let P be a deterministic protocol with communication complexity $\text{CC}(P)$ and the number of rounds r approximating a sensitivity-1 function $f: \Sigma^n \times \Sigma^n \rightarrow \mathbb{Z}$ 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 $\Delta + O(\text{CC}(P)r/\epsilon)$.*

The linear dependency on the communication complexity in the last theorem is unlikely to be improved due to the lower bound of Theorem 4.11.

2 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 [BNO08].

Definition 2.1 (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 2.2 (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 $\text{VIEW}_P^A(x, y)$ 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 , $\text{VIEW}_P^A(x, y)$ is a mechanism over the y 's. Let $\text{VIEW}_P^B(x, y)$ be similarly defined view of B whose input is y .

Definition 2.3 (Differential privacy for two-party protocols). *We say that a protocol P has ϵ -differential privacy if the mechanism $\text{VIEW}_P^A(x, y)$ is ϵ -differentially private for all values of x and same holds for $\text{VIEW}_P^B(x, y)$ 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 they may share private random bits 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 2.4 (Sensitivity). *For a real-valued function $f: \Sigma^n \rightarrow \mathbb{R}$ 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 3 and 4:

Definition 2.5 (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 [Yao79] introduced the following, by now classical, two-player communication game: Alice and Bob want to collaboratively compute a function $f: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$. 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 $\text{CC}(P)$, 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.

3 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 [SV86] and studied in the literature on randomness extractors.

We then generalize a result of Vazirani [Vaz87] 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 2.2) appears in Section A.

3.1 Unpredictable Sources from Differential Privacy

The model of random sources introduced by Santha and Vazirani [SV86] is one where each bit is somewhat unpredictable given the previous ones:

Definition 3.1 (α -unpredictable bit source²). *For $\alpha \in [0, 1]$, random variable $X = (X_1, \dots, X_n)$ taking values in $\{0, 1\}^n$ is an α -unpredictable bit source if for every $i \in [n]$, and every $x_1, \dots, x_{i-1} \in \{0, 1\}$, we have*

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

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 $n \rightarrow \infty$. 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 [RVW04], is to require that each bit is somewhat unpredictable given *all* of the other bits, even the future ones:

Definition 3.2 (Strongly α -unpredictable bit source). *For $\alpha \in [0, 1]$, a random variable $X = (X_1, \dots, X_n)$ taking values in $\{0, 1\}^n$ is a strongly α -unpredictable bit source if for every $i \in [n]$, and every $x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n \in \{0, 1\}^n$, we have*

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

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):

²In the terminology of Santha and Vazirani [SV86], this is an $\alpha/(1 + \alpha)$ semi-random source.

Lemma 3.3. *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 $\pi \in \text{Supp}(\Pi)$, 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 ,

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

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

3.2 Randomness Extraction and Lower Bounds for Inner Product

Vazirani [Vaz87] 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 [SV86]). 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 3.4. *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 [DO03].

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

Lemma 3.5. *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*

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

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

$$\|Z - U\|_{SD} = \frac{1}{2} \|p_Z - p_U\|_1 \leq \frac{\sqrt{m}}{2} \|p_Z - p_U\|_2 = \frac{1}{2} \sqrt{\sum_{k=0}^{m-1} |\hat{p}_Z(k) - \hat{p}_U(k)|^2}.$$

Plugging in the Fourier coefficients $\hat{p}_Z(0)$ and $\hat{p}_U(\cdot)$, the claim follows. \square

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.

Lemma 3.6. *Let X be any random variable taking values in $\{0, 1\}^n$, $\omega \in \mathbb{C}$ a primitive m th root of unity, and $t \in \mathbb{N}$. Then*

$$\sum_{y \in \mathbb{Z}_m^n} \left| \mathbb{E} \left[\omega^{\langle X, y \rangle} \right] \right|^{2t} \leq m^n \cdot \Pr \left[\sum_i (X(i) - X(i')) \equiv 0^n \pmod{m} \right],$$

where $X(1), \dots, X(t), X(1)', \dots, X(t)'$ are iid copies of X .

Proof. For every complex number u , we have $|u|^2 = u\bar{u}$, where \bar{u} is the complex conjugate of u . Thus

$$\begin{aligned} \sum_{y \in \mathbb{Z}_m^n} \left| \mathbb{E} \left[\omega^{\langle X, y \rangle} \right] \right|^{2t} &= \sum_{y \in \mathbb{Z}_m^n} \mathbb{E} \left[\omega^{\langle X, y \rangle} \right]^t \cdot \overline{\mathbb{E} \left[\omega^{\langle X, y \rangle} \right]^t} \\ &= \sum_{y \in \mathbb{Z}_m^n} \mathbb{E} \left[\omega^{\langle X, y \rangle} \right]^t \cdot \mathbb{E} \left[\omega^{-\langle X, y \rangle} \right]^t \\ &= \sum_{y \in \mathbb{Z}_m^n} \mathbb{E} \left[\omega^{\langle \sum_{i=1}^t (X(i) - X(i)'), y \rangle} \right], \\ &= \Pr \left[\sum_{i=1}^t (X(i) - X(i')) \equiv 0^n \pmod{m} \right] \cdot m^n, \end{aligned}$$

where the last equality uses the fact that for every nonzero vector $x \in \mathbb{Z}_m^n$, $\sum_{y \in \mathbb{Z}_m^n} \omega^{\langle x, y \rangle} = 0$. (Note that this fact is only true if ω is a primitive m th root of unity; indeed, if all the components of x were divisible by the order of ω , then the sum would equal m^n .) \square

Next we show that $2t$ independent unpredictable sources on $\{0, 1\}^n$ sum to zero with probability approaching $1/m^n$ at rate that vanishes exponentially in t .

Lemma 3.7. *Let $X(1), \dots, X(t), X(1)', \dots, X(t)'$ be independent, α -unpredictable bit sources and $m \in \mathbb{N}$. Then*

$$\begin{aligned} \Pr \left[\sum_i (X(i) - X(i)') \equiv 0^n \pmod{m} \right] &\leq \left[\frac{1}{m} + \left(1 - \frac{1 - \cos(2\pi/m)}{1 + (\alpha + \alpha^{-1})/2} \right)^t \right]^n \\ &= \left[\frac{1}{m} + \exp \left(-\Omega \left(\frac{\alpha t}{m^2} \right) \right) \right]^n. \end{aligned}$$

Proof. We observe that it suffices to prove the case $n = 1$, because

$$\begin{aligned} &\Pr \left[\sum_i (X(i) - X(i)') \equiv 0^n \pmod{m} \right] \\ &= \Pr \left[\sum_i (X(i)_1 - X(i)'_1) \equiv 0 \pmod{m} \right] \\ &\quad \cdot \Pr \left[\sum_i (X(i)_{2\dots n} - X(i)'_{2\dots n}) \equiv 0^{n-1} \pmod{m} \mid \sum_i (X(i)_1 - X(i)'_1) \equiv 0 \pmod{m} \right], \end{aligned}$$

and conditioned on the values of all the first bits of the sources (namely $X(i)_1$ and $X(i)'_1$), the remaining $n - 1$ bits of the sources (namely $X(i)_{2\dots n}$ and $X(i)'_{2\dots n}$) are independent α -unpredictable sources.

Assume $n = 1$ (so each $X(i) \in \{0, 1\}$), let $S = \sum_{i=1}^t (X(i) - X(i)') \in \mathbb{Z}$, and let Ω be a uniformly random m th root of unity. Then

$$\Pr[S \equiv 0 \pmod{m}] = \mathbb{E}[\Omega^S] \leq \frac{1}{m} + \max_{\omega \neq 1} |\mathbb{E}[\omega^S]| = \frac{1}{m} + \max_{\omega \neq 1} \prod_{i=1}^t |\mathbb{E}[\omega^{X(i)}]| \cdot |\mathbb{E}[\omega^{-X(i)}]|.$$

So now we only need to bound $|\mathbb{E}[\omega^W]|$ and $|\mathbb{E}[\omega^{-W}]|$ for an α -unpredictable 1-bit random variable W . Thinking of \mathbb{C} as \mathbb{R}^2 , $|\mathbb{E}[\omega^W]|$ is the length of a convex combination of two unit vectors, namely $u = \omega$ and $v = 1$, where the ratio of coefficients is at least α . This length is maximized when the two coefficient have ratio exactly α , namely when u has coefficient $p = 1/(1+\alpha)$ and v has coefficient $1-p = \alpha \cdot p = 1/(1+\alpha^{-1})$. In this case, the length is

$$\|pu + (1-p)v\| = \sqrt{1 - 2p \cdot (1-p) \cdot (1 - \langle u, v \rangle)} = \sqrt{1 - \frac{1}{1 + (\alpha + \alpha^{-1})/2} \cdot (1 - \cos \theta)}.$$

where $\omega = \exp(2\pi i\theta)$. Since ω is an m th root of unity other than 1, we have $\cos \theta \leq \cos(2\pi/m)$. The same analysis works to bound the length of $|\mathbb{E}[\omega^{-W}]|$, since ω^{-1} is also an m th root of unity.

Therefore,

$$\Pr[S \equiv 0 \pmod{m}] \leq \frac{1}{m} + \left(\sqrt{1 - \frac{1}{1 + (\alpha + \alpha^{-1})/2} \cdot (1 - \cos \theta)} \right)^{2t},$$

as desired. □

It follows that:

Lemma 3.8. *Let X be an α -unpredictable source on $\{0, 1\}^n$, $\omega \in \mathbb{C}$ a primitive m th root of unity, and $t \in \mathbb{N}$. Then*

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

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 3.4. 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

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

and let $L = \bigcup_\omega L_\omega$. By Lemma 3.5, 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 $(1/2)\sqrt{(m-1) \cdot (\delta/\sqrt{m})^2} \leq \delta/2$. Thus it suffices to prove that $\Pr[Y \in L] \leq \delta/2$, which in turn follows if $\Pr[Y \in L_\omega] \leq \delta/2m$ 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 3.8, we have

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

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.

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

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

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 3.9. *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 3.3, we know that for every $\pi \in \text{Supp}(\Pi)$, 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 $\beta = \log(1 + \alpha) \geq \alpha$. By Theorem 3.4, $(Y_\pi, \langle X_\pi, Y_\pi \rangle \bmod m)$ has statistical distance at most $\delta/2$ from (Y_π, U) , provided

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

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

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

for some universal constant c_1 , which in turn follows if

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

for a small universal constant $c_2 > 0$.

Consider the set $S = \{(\pi, y, z) : (f_B(\pi, y) - z) \bmod m \in \{m - \Delta, \dots, m - 1, 0, 1, \dots, \Delta\}\}$. 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:

$$\Pr[(\Pi, Y, \langle X, Y \rangle \bmod m) \in S] \leq \Pr[(\Pi, Y, U) \in S] + \delta/2 \leq 2\Delta/m + \delta/2 = \delta.$$

□

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.

3.3 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 3.10. *Let $f: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \mathbb{R}$ be a sensitivity-1 function. Then for every distribution μ such that for every input y , the conditional distribution $\mu(X | Y = y)$ is a product distribution $\prod_{i=1}^n \mu_i(X_i | Y = y)$, 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. [DP09]) say that every sensitivity-1 function on a product distribution is well concentrated around its expectation. Specifically, for every $h: \{0, 1\}^n \rightarrow \mathbb{R}$, and every product distribution ν on X ,

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

Applying this result to the function $f(X, y)$ and setting $g(y) = \mathbb{E}_{x \in \mu(X|Y=y)}[f(x, y)]$ 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 3.3) requires that the inputs come from a product distribution, we cannot get a lower bound better than $\Theta(\sqrt{n})$ from that approach.

4 Differential Privacy and Communication Complexity

In this section we relate differentially private protocols to communication complexity. We first show that techniques used to prove lower bounds for communication complexity can also be used to show lower bounds on differentially private protocols. There are two main approaches to proving communication complexity lower bounds. The information-theoretic approach lower bounds the communication of any protocol computing a function f by the *Information Cost* of computing f on some distribution (see Section 4.1 for a formal definition). The combinatorial approaches such as richness and the rectangle bound are generalized by the *Partition Bound* (see Section 4.2 for a formal definition). We show that both these approaches can be used to prove lower bounds for differentially private protocols. Specifically, we show that if an ϵ -differentially private protocol computes a function f , then both the information cost (Section 4.1) and the (logarithm of the) partition bound (Section 4.2) are $O(\epsilon n)$. Thus a linear lower bound on the communication complexity using either of the approaches yields a lower bound for differentially private protocols. We can extend the information cost bound to approximate differential privacy for distributions where all private input bits are drawn independently of the others. However, analogous results for general distributions, and a corresponding result for the partition bound do not hold (Section 4.4).

We then prove stronger separations between information-theoretic and computational differential privacy (Section 4.3). We also note that the message compression technique of Barak et al. [BBCR10], 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 4.5). Finally, we show that low-communication protocols can be converted into privacy-preserving ones with some loss of accuracy (Section 4.6).

4.1 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. [BYJKS02] (based on a earlier concept introduced by Chakrabarti et al. [CSWY01].)

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

Definition 4.1 (Mutual Information). *Given two random variables X and Y over the same probability space, their mutual information is defined as follows:*

$$I(X; Y) = H(X) - H(X | Y),$$

where H denotes Shannon entropy. The conditional mutual information is $I(X; Y | Z) = H(X | Z) - H(X | YZ)$.

Intuitively, $I(X; Y)$ 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 4.2 (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*

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

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.

Proposition 4.3. *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:*

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

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

Proof. Consider the mutual information between the random input $Z = (X_1, \dots, X_n, Y_1, \dots, Y_n)$ and the protocol's transcript Π . For every z, z' , differential privacy implies that

$$\exp(-2\epsilon n) \leq \frac{\Pr[\Pi(z) = \pi]}{\Pr[\Pi(z') = \pi]} \leq \exp(2\epsilon n).$$

so that

$$\exp(-2\epsilon n) \leq \frac{\Pr[\Pi(z) = \pi]}{\Pr[\Pi(Z') = \pi]} \leq \exp(2\epsilon n).$$

where Z' is an independent sample from μ .

$$\begin{aligned} I(\Pi(Z); Z) &= H(\Pi(Z)) - H(\Pi(Z)|Z) \\ &= \mathbb{E}_{(z, \pi) \leftarrow (Z, \Pi(Z))} \log \frac{\Pr[\Pi[Z] = \pi | Z = z]}{\Pr[\Pi(Z) = \pi]} \\ &\leq 2(\log_2 e)\epsilon n. \end{aligned}$$

If μ is the uniform distribution and $\Sigma = \{0, 1\}$ (i.e., each bit of Z is uniform and independent), then we can improve the bound as follows. By additivity of mutual information

$$I(Z; \Pi(Z)) = \sum_{i \in [2n]} I(Z_i; \Pi(Z) | Z_1 Z_2 \dots Z_{i-1})$$

Each term of the last expression can be written as:

$$I(Z_i; \Pi(Z) | Z_1 \dots Z_{i-1}) = H(Z_i | Z_1 \dots Z_{i-1}) - H(Z_i | \Pi(Z) Z_1 \dots Z_{i-1}).$$

Since each Z_i is independent and uniform in μ , the first term is 1. By the differential privacy property,

$$\frac{\Pr[\Pi[Z] = \pi | Z_1, \dots, Z_i = z_1, \dots, z_{i-1}, 0]}{\Pr[\Pi[Z] = \pi | Z_1, \dots, Z_i = z_1, \dots, z_{i-1}, 1]} \in (\exp(-\epsilon), \exp(\epsilon))$$

so that by Bayes' rule, we can conclude that the ratio

$$\frac{\Pr[Z_i = 0 | Z_1, \dots, Z_{i-1} = z_1, \dots, z_{i-1}, \Pi[Z] = \pi]}{\Pr[Z_i = 1 | Z_1, \dots, Z_{i-1} = z_1, \dots, z_{i-1}, \Pi[Z] = \pi]} \in (\exp(-\epsilon), \exp(\epsilon))$$

for all z_1, \dots, z_{i-1}, π . Thus the second term above is bounded by $H(\exp(\epsilon)/2)$. An easy calculation shows that the difference $(1 - H(\exp(\epsilon)/2))$ is bounded by $\epsilon^2/(2 \ln 2)$. The bound of $\log_2(e)\epsilon^2 n$ on the information cost of the protocol follows by summing over $2n$ terms. \square

In the conference version of this paper [MMP⁺10], we claimed an extension of the above proposition to (ϵ, δ) -differential privacy, but there was an error in the proof and De [De11] gave a counterexample. (See Section 4.4.) However, we can still show such a bound in the case when components of x and y are all independent.

Proposition 4.4. *Let μ be a product distribution over Σ^{2n} , i.e., $\mu(x, y) = \prod_i \mu_i(x_i) \mu'_i(y_i)$. If $P(x, y)$ has (ϵ, δ) -differential privacy, where $x, y \in \Sigma^n$ for a finite alphabet Σ , $\epsilon < 1$, and $\delta < \epsilon/4|\Sigma|^2$, then the information cost of P is bounded as follows:*

$$\text{ICost}_\mu(P) \leq \left(2\epsilon + \frac{4\delta|\Sigma|^2}{\epsilon} \cdot \log \frac{\epsilon}{2\delta|\Sigma|} \right) \cdot n.$$

Proof. Let μ be a product distribution over $Z = (X_1, \dots, X_n, Y_1, \dots, Y_n)$. For $i \in [2n]$, let Z_{-i} denote $(Z_1, \dots, Z_{i-1}, Z_{i+1}, \dots, Z_{2n})$. By the chain rule and Lemma B.1 in the appendix, we have

$$\begin{aligned} \text{I}(\Pi(Z); Z) &= \sum_i \text{I}(\Pi(Z); Z_i | Z_1 \dots Z_{i-1}) \\ &\leq \sum_i \text{I}(\Pi(Z); Z_i | Z_{-i}) \\ &= \sum_i \mathbb{E}_{(z_{-i} \leftarrow Z_{-i})} [\text{I}(\Pi(Z); Z_i | Z_{-i} = z_{-i})]. \end{aligned}$$

We will upper bound each term in the summation. Fix a set of values for z_{-i} . Now $\Pi(Z)$ is a mechanism dependent on Z_i alone; let $g(\cdot)$ denote $\Pi(Z)$ as a function of Z_i . Then by the definition of differential privacy, for any subset $A \subseteq \mathcal{R}$, and for any $a, b \in \Sigma$,

$$\Pr[g(a) \in A] \leq \exp(\epsilon) \cdot \Pr[g(b) \in A] + \delta.$$

Let B_{ab} be a maximal subset $A \subseteq \mathcal{R}$ such that $\Pr[g(a) \in A] > \exp(2\epsilon) \cdot \Pr[g(b) \in A]$. Thus every $A \subseteq B_{ab}^c$ satisfies $\Pr[g(a) \in A] \leq \exp(2\epsilon) \cdot \Pr[g(b) \in A]$. Moreover from the definition of B_{ab} , we have

$$\exp(2\epsilon) \cdot \Pr[g(b) \in B_{ab}] \leq \Pr[g(a) \in B_{ab}] \leq \exp(\epsilon) \cdot \Pr[g(b) \in B_{ab}] + \delta.$$

Rearranging we get

$$\Pr[g(a) \in B_{ab}] \leq \frac{\exp(\epsilon)}{\exp(\epsilon) - 1} \cdot \delta \leq 2\delta/\epsilon$$

for $\epsilon < 1$. Let $B_i = \cup_{a,b \in \Sigma} B_{ab}$ so that $\Pr[g(a) \in B_i] < 2\delta|\Sigma|^2/\epsilon$.

Now we write

$$\begin{aligned} \text{I}(\Pi(Z); Z_i | Z_{-i} = z_{-i}) &\leq H(\mathbf{1}_{\Pi(Z) \in B_i} | Z_{-i} = z_{-i}) + \text{I}(\Pi(Z); Z_i | Z_{-i} = z_{-i}, \mathbf{1}_{\Pi(Z) \in B_i}) \\ &\leq H\left(\frac{2\delta|\Sigma|^2}{\epsilon}\right) + \Pr[\Pi(Z) \in B_i | Z_{-i} = z_{-i}] \cdot \text{I}(\Pi(Z); Z_i | Z_{-i} = z_{-i}, \Pi(Z) \in B_i) \\ &\quad + \Pr[\Pi(Z) \notin B_i | Z_{-i} = z_{-i}] \cdot \text{I}(\Pi(Z); Z_i | Z_{-i} = z_{-i}, \Pi(Z) \notin B_i) \\ &\leq H\left(\frac{2\delta|\Sigma|^2}{\epsilon}\right) + \left(\frac{2\delta|\Sigma|^2}{\epsilon}\right) \cdot H(Z_i) + \text{I}(\Pi(Z); Z_i | Z_{-i} = z_{-i}, \Pi(Z) \notin B_i) \\ &\leq H\left(\frac{2\delta|\Sigma|^2}{\epsilon}\right) + \frac{2\delta|\Sigma|^2 \log |\Sigma|}{\epsilon} \\ &\quad + \mathbb{E}_{(a,\pi) \leftarrow (Z_i, g(Z_i)) | g(Z_i) \notin B_i} \left[\log \frac{\Pr[g(Z_i) = \pi | Z_i = a]}{\Pr[g(Z_i) = \pi]} \right]. \end{aligned}$$

By definition of B_i , $\Pr[g(a) = \pi] \leq \exp(2\epsilon) \cdot \Pr[g(b) = \pi]$ for every $a, b \in \Sigma$ and $\pi \notin B_i$, so that the last term is at most 2ϵ . Thus we conclude that when μ is a product distribution,

$$I(\Pi(Z); Z_i | Z_{-i}) \leq 2\epsilon + \frac{4\delta|\Sigma|^2}{\epsilon} \cdot \log \frac{\epsilon}{2\delta|\Sigma|}.$$

□

We can in fact extend the result allowing some limited dependencies.

Definition 4.5. Let X_1, \dots, X_{kn} be a sequence of random variables taking values in Σ and let $Z_i \stackrel{\text{def}}{=} (X_{(i-1)k+1}, \dots, X_{ik})$ denote the Σ^k -valued random variable denoting the i th block. We say the collection $\{X_1, \dots, X_{kn}\}$ is k -block independent if under some permutation of indices, the resulting random variables Z_1, \dots, Z_n are independent.

In other words, the random variables can be represented by a graphical model with n components of size k each. The following corollary extends this result to distributions that are mixtures of k -block independent distributions.

Corollary 4.6. Let μ be a distribution over Σ^{2n} such that there is a random variable W with the property that $\mu|W$ is a k -block independent distribution. If $P(x, y)$ has (ϵ, δ) -differential privacy, where $x, y \in \Sigma^n$ for a finite alphabet Σ , $\epsilon < 1$, and $\delta < \epsilon/4|\Sigma|^2$, then the information cost of P is bounded as follows:

$$I\text{Cost}_\mu(P) \leq \left(2k\epsilon + \frac{4\delta|\Sigma|^{2k}}{\epsilon} \cdot \log \frac{\epsilon}{2\delta|\Sigma|} \right) \cdot n + H(W).$$

Proof. We can handle the case that μ is k -block independent by observing that then $\Pi(Z)$ is $(k\epsilon, k\delta)$ -differentially private with respect to the Σ^k -valued blocks, which are fully independent. For the case that $\mu|W$ is k -block independent for a random variable W , we have $I(\Pi(X, Y); XY) \leq H(W) + I(\Pi(X, Y); XY|W)$. □

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

Theorem 4.7 (Barak et al. [BBCR10]). For every product distribution μ , for every protocol randomized P with output $\text{out}(P)$, and every $\gamma > 0$, there exists functions f_A, f_B , and protocol Q such that

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

where $I\text{Content}_\mu(P) = I(X; \Pi(X, Y) | Y) + I(Y; \Pi(X, Y) | X)$ which satisfies $I\text{Content}_\mu(P) = O(I\text{Cost}_\mu(P))$.

It follows that differentially private protocols can be compressed.

Theorem 4.8. Let P be an ϵ -differentially private protocol P with output $\text{out}(P)$ 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 $\|f_A(X, Q(X, Y)) - \text{out}(P)\|_{SD} < \gamma$, $\Pr[f_A(X, Q(X, Y)) \neq f_B(Y, Q(X, Y))] < \gamma$ and $\text{CC}(Q) \leq 3\epsilon\gamma^{-1}n \cdot \text{polylog}(\text{CC}(P)/\gamma)$.

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

Figure 1: Linear program for the Partition Bound for a function f .

4.2 Differential Privacy and the Partition Bound

Jain and Klauck [JK10] define the *partition bound* for a partial function $f: \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{Z}$. 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 $\text{prt}_{\gamma}(f)$ 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 $\log \text{prt}_{\gamma}(f)$. 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 [JK10] 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 $g(x, y) < t - \alpha$.

Theorem 4.9. *Suppose that an ϵ -differentially private protocol P computes a partial function $f: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \mathcal{Z}$ with error probability at most γ . Then $\log \text{prt}_{\gamma}(f) \leq 3\epsilon n$.*

Proof. Given P , we show how to construct a solution for the linear program defining $\text{prt}_{\gamma}(f)$ with objective function value at most $2^{3\epsilon n}$. Let the inputs (X, Y) to the protocol be chosen independently and uniformly at random. Let $\Pi(x, y)$ be the distribution of transcripts of the protocol on inputs (x, y) and let f_{π} be the output of the protocol on transcript π . Thus for every $(x, y) \in \text{Supp}(f)$, we have $\Pr_{\pi \leftarrow \Pi(x, y)}[f_{\pi} = f(x, y)] \geq 1 - \gamma$.

For a transcript π , let X_{π} be the distribution on x conditioned on the transcript being π and let Y_{π} be similarly defined. Since P is differentially private, for every $x, x' \in \{0, 1\}^n$,

$$\exp(-\epsilon n) \leq \frac{\Pr[\Pi(x, y) = \pi]}{\Pr[\Pi(x', y) = \pi]} \leq \exp(\epsilon n).$$

Thus by Bayes' rule, we conclude that

$$\exp(-\epsilon n) \leq \frac{\Pr[X_{\pi} = x]}{\Pr[X_{\pi} = x']} \leq \exp(\epsilon n).$$

In other words, the min entropy of the distribution X_π is at least $n - \epsilon'n$ where $\epsilon' = \epsilon \log_2 e$. Let $\mathbf{p} \in \mathfrak{R}^{\{0,1\}^n}$ be defined by $p_x = \Pr[X_\pi = x]$. Consider the greedy algorithm that starts with all $\alpha_{\pi,S}$ as zero, and increases $\alpha_{\pi,S}$ for $S = \text{Supp}(p - \sum_S \alpha_{\pi,S} \mathbf{1}_S)$ as long as each coordinate of $p - \sum_S \alpha_{\pi,S} \mathbf{1}_S$ stays non-negative. Repeating until all coordinates become zero, we get numbers $\alpha_{\pi,S} : S \subseteq \{0,1\}^n$ such that for all x ,

$$\Pr[X_\pi = x] = \sum_{S \subseteq \{0,1\}^n : x \in S} \alpha_{\pi,S},$$

and

$$\sum_{S \subseteq \{0,1\}^n} \alpha_{\pi,S} \leq 2^{-n+\epsilon'n}.$$

Doing a similar decomposition of Y_π , and taking the pairwise product, we can define numbers $\beta_{\pi,R}$ for $R \in \mathcal{R}$ such that for all (x, y) ,

$$\Pr[(X_\pi, Y_\pi) = (x, y)] = \sum_{R \in \mathcal{R} : (x,y) \in R} \beta_{\pi,R},$$

and

$$\sum_{R \in \mathcal{R}} \beta_{\pi,R} \leq 2^{-2n+2\epsilon'n}.$$

For a $z \in \mathcal{Z}$ and $R \in \mathcal{R}$, we now define $w_{z,R} = 2^{2n} \sum_{\pi: f_\pi=z} \Pr[\Pi(X, Y) = \pi] \cdot \beta_{\pi,R}$. We will show that this setting of variables satisfies all the constraints of the above linear program.

First note that for every (x, y) ,

$$\begin{aligned} 2^{-2n} &= \Pr[(X, Y) = (x, y)] \\ &= \sum_{\pi} \Pr[\Pi(X, Y) = \pi] \cdot \Pr[(X_\pi, Y_\pi) = (x, y)] \\ &= \sum_{\pi} \Pr[\Pi(X, Y) = \pi] \cdot \sum_{R \in \mathcal{R} : (x,y) \in R} \beta_{\pi,R} \\ &= \sum_{z \in \mathcal{Z}} \sum_{\pi: f_\pi=z} \Pr[\Pi(X, Y) = \pi] \cdot \sum_{R \in \mathcal{R} : (x,y) \in R} \beta_{\pi,R} \\ &= \sum_{R \in \mathcal{R} : (x,y) \in R} \sum_{z \in \mathcal{Z}} \sum_{\pi: f_\pi=z} \Pr[\Pi(X, Y) = \pi] \cdot \beta_{\pi,R} \\ &= 2^{-2n} \sum_{R \in \mathcal{R} : (x,y) \in R} \sum_{z \in \mathcal{Z}} w_{z,R}. \end{aligned}$$

Moreover, for every $(x, y) \in \text{Supp}(f)$, we similarly have

$$\begin{aligned} (1 - \gamma)2^{-2n} &\leq \Pr[(X, Y) = (x, y) \wedge f_\pi = f(x, y)] \\ &= \sum_{\pi: f_\pi=f(x,y)} \Pr[\Pi(X, Y) = \pi] \cdot \Pr[(X_\pi, Y_\pi) = (x, y)] \\ &= \sum_{\pi: f_\pi=f(x,y)} \Pr[\Pi(X, Y) = \pi] \cdot \sum_{R \in \mathcal{R} : (x,y) \in R} \beta_{\pi,R} \\ &= \sum_{R \in \mathcal{R} : (x,y) \in R} \sum_{\pi: f_\pi=f(x,y)} \Pr[\Pi(X, Y) = \pi] \cdot \beta_{\pi,R} \\ &= 2^{-2n} \sum_{R \in \mathcal{R} : (x,y) \in R} w_{f(x,y),R}. \end{aligned}$$

Finally,

$$\begin{aligned}
\sum_{z \in \mathcal{Z}} \sum_{R \in \mathcal{R}} w_{z,R} &= 2^{2n} \sum_{z \in \mathcal{Z}} \sum_{R \in \mathcal{R}} \sum_{\pi: f_\pi = z} \Pr[\Pi(X, Y) = \pi] \cdot \beta_{\pi,R} \\
&= 2^{2n} \sum_{\pi} \Pr[\Pi(X, Y) = \pi] \cdot \sum_{R \in \mathcal{R}} \beta_{\pi,R} \\
&\leq 2^{2n} \sum_{\pi} \Pr[\Pi(X, Y) = \pi] \cdot 2^{-2n+2\epsilon'n} \\
&= 2^{2\epsilon'n}.
\end{aligned}$$

The claim follows. \square

Chakrabarti and Regev [CR11] showed that the Gap Hamming problem — distinguishing inputs with Hamming distance at most $n/2 - c\sqrt{n}$ from those with Hamming distance at most $n/2 + c\sqrt{n}$ — has a smooth rectangle bound of $2^{\Omega(n)}$ (for some absolute constant c). By the results of Jain and Klauck [JK10], this implies that the partition bound for the Gap Hamming (partial) function is at least as large. Coupled with the above result, this implies that for some small constant ϵ , any ϵ -differentially private mechanism for Hamming distance must incur additive error $\Omega(\sqrt{n})$. This improves on the $\Omega(\sqrt{n}/\log n)$ lower bound that we showed using our deterministic extractor approach. However, this partition-bound approach does not apply to (ϵ, δ) -differential privacy. (See Section 4.4.)

4.3 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 4.10. *There exists an absolute constant $\beta > 0$ such that for every m , there is an efficiently computable function $f: \{0, 1\}^m \times \{0, 1\}^m \rightarrow \mathbb{R}$ 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 $g: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ and distribution \mathcal{D}_g over its inputs, we can transform it to a sensitivity-1 function $f_g: \{0, 1\}^m \times \{0, 1\}^m \rightarrow \mathbb{R}$, 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 $\text{ICost}_{\mathcal{D}_g, \gamma}(g)$. 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 $C \subseteq \{0, 1\}^m$, 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 $Enc: \{0, 1\}^n \rightarrow \{0, 1\}^m$ be an encoding algorithm for a linear-rate error-correcting code C with a decoding algorithm $Dec: \{0, 1\}^m \rightarrow \{0, 1\}^n$ that works up to decoding radius αm for some constant $\alpha > 0$. Such codes exist with $n = rm$ for some constant $r = r(\alpha) > 0$. Let $d(x, C)$ be the distance from x to the closest codeword in C . We then define

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

Note that when x and y are both codewords, $f_g(x, y)$ is exactly $\alpha m \cdot g(\text{Dec}(x), \text{Dec}(y))$. As we move away from C , $f_g(x, y)$ smoothly decays to 0. Moreover, since C has decoding radius αm , the function is well-defined and efficiently computable: if any of $\text{Dec}(x)$ or $\text{Dec}(y)$ fails to decode, it means that $d(x, C) + d(y, C) > \alpha m$ and the function is zero by definition.

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

We first argue that any communication protocol P_{f_g} approximating f_g to within error less than $\alpha m/2 = \alpha n/(2r)$ (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 $x' = \text{Enc}(x)$ and $y' = \text{Enc}(y)$, and Alice outputs 0 if her output $f_A(x', \Pi_{P_{f_g}})$ in protocol P_{f_g} is smaller than $\alpha m/2$, and 1 otherwise. Since $f_g(x', y')$ is equal to $\alpha m \cdot g(x, y)$, if P_{f_g} has error less than $\alpha m/2$ on (x', y') , then

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

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 $\cup_{x: g(x, \text{Dec}(y_1))=1} B(\text{Enc}(x), \alpha m - d(y_1, C))$. 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 αm , 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 [BYJKS02], we get the desired result. \square

Combining this result with Proposition 4.3, we conclude

Theorem 4.11. *There exists an absolute constant $\beta > 0$ such that for every n , there is an efficiently computable function $f: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \mathbb{R}$ and a distribution \mathcal{D} over its inputs, with the following properties*

- (a) for every $\epsilon < \beta/3$, every ϵ -differentially private protocol P has expected additive error at least βn .
- (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' .

4.4 Counterexample for Approximate Differential Privacy

As mentioned earlier, in the conference version of the paper [MMP⁺10], a theorem analogous to Theorem 4.11 was also claimed for (ϵ, δ) differential privacy, but there was an error in the proof and we do not know how to extend it. In fact, for inputs in the support of the distribution \mathcal{D} constructed in the proof of Theorem 4.11, there is an (ϵ, δ) differentially private protocol (with δ negligible in n) that allows the two parties to perfectly reconstruct each others' inputs and hence compute f exactly. This example is based on the work of De [De11].

Let C , f_g and \mathcal{D} be as in the proof of Theorem 4.11. We will show an approximately differentially private protocol that computes f_g exactly on $\text{Supp}(\mathcal{D})$ with high probability. Indeed observe that for any $i \in [n]$, the function $h_i: \{0, 1\}^m \rightarrow \mathbb{R}$ defined as $h_i(x) = \text{Dec}_i(x) \cdot (\alpha m - d(x, C)) \cdot \mathbf{1}(d(x, C) \leq \alpha m)$ has sensitivity 1, where $\text{Dec}_i(x)$ denoted the i th bit of the decoding of x . Thus Alice can communicate $h_i(x) + N(0, \sqrt{n \log(1/\delta)})/\epsilon$ for each i while preserving (ϵ, δ) -differential privacy [DKM⁺06]. For any $x \in C$, $h_i(x) \in \{0, \alpha m\}$ whereas the magnitude of the noise will be smaller than $\alpha m = \Theta(n)$ for all i with probability $1 - o(1)$, as long as $\sqrt{n \log(1/\delta)}/\epsilon = o(n/\log n)$. Such noisy measurements enable Bob to compute $\text{Dec}_i(x)$ for all i with high probability, and thus compute f_g . Thus, for example, there is an $(n^{-0.1}, 2^{-n^{0.1}})$ -differentially private protocol that computes f_g with high probability on \mathcal{D} .

This construction also shows that Proposition 4.3 does not extend to approximate differential privacy, as the information cost of the above protocol is close to n (since Bob learns Alice's input with high probability). Similarly, if we start with a function g with a partition bound of $2^{\Omega(n)}$ and consider the partial function f_g defined on $\text{Supp}(\mathcal{D})$ above, we can rule out an extension of Theorem 4.9 to approximate differential privacy.

4.5 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) [Man94, Hol09] 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.

Theorem 4.12. *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 $O(r \cdot (\epsilon n + \log \log \frac{1}{\epsilon\delta}))$ and except with probability $r\delta$, simulates P perfectly. In other words, there exist functions π_x, π_y such that $\Pr[\pi_x(\text{VIEW}_{P^*}^A(x, y)) = \text{VIEW}_P^A(x, y)] \geq 1 - r\delta$, and similarly for B .*

Proof. We show how to simulate one round of the protocol; the result for r rounds follows by composition. Suppose that Alice sends k bits in round t in protocol P . Thus given the messages sent in rounds $1, \dots, t-1$, there is a distribution μ_x on $\{0, 1\}^k$ for every $x \in \Sigma^n$. Let $\nu(z) = \max_{x \in \Sigma^n} \mu_x(z)$, the envelope of the distributions $\mu_x(z)$ for all x . Set $N = \exp(\epsilon n) \left(\frac{1}{\epsilon} + \log \frac{1+\epsilon}{\epsilon\delta} \right)$. The mechanism P^* simulates round t as follows:

1. With probability $\delta/(1+\epsilon)$ output fail and abort.
2. Interpret shared randomness as a sequence of N independent random values z_1, \dots, z_N from $\{0, 1\}^k$, where $z \in \{0, 1\}^k$ is picked with probability proportional to $\nu(z)$.
3. For each $i \in [N]$, Alice uses private randomness to pick $r_i \in [0, \nu(z_i)]$ uniformly at random.
4. Let $C = \{i: r_i \leq \mu_x(z_i)\}$. If C is nonempty, output a random element of C . Else output fail.

Bob, on receiving $i \neq \text{fail}$, interprets it as z_i using the shared randomness. In other words, the map π_y maps a message i to z_i . To prove the theorem, we need to establish three properties.

1. *Low communication.* Observe that in round t of P^* , Alice sends $\lceil \log(N+1) \rceil$ bits to Bob.
2. *Accuracy of simulation.* For each $z \in \{0, 1\}^k$,

$$\Pr[(z_i = z) \wedge (i \in C)] = \Pr[z_i = z] \cdot \Pr[i \in C | z_i = z] = \frac{\nu(z)}{\nu^*} \times \frac{\mu_x(z)}{\nu(z)} = \frac{\mu_x(z)}{\nu^*},$$

where $\nu^* = \sum_z \max_{x \in \Sigma^n} \mu_x(z)$ is a normalization factor. It follows that if C is nonempty, $\{z_i: i \in C\}$ is a set of $|C|$ samples from μ_x . Hence whenever the mechanism does not fail, it respects the input distribution. By differential privacy of P , $\nu(z) \leq \exp(\epsilon n) \mu_x(z)$ for every $x \in \Sigma^n$. Thus $\Pr[i \in C] \geq \exp(-\epsilon n)$. This implies that $\Pr[|C| = 0] \leq (1 - \exp(-\epsilon n))^N \leq \epsilon\delta/(1+\epsilon)$. The probability of outputting fail is thus bounded as $\delta/(1+\epsilon) + \epsilon\delta/(1+\epsilon) = \delta$.

3. *Privacy.* To argue privacy, we first define X_j as a Bernoulli random variable with $\mathbb{E}[X_j] = \mu_x(z_j)/\nu(z_j)$ and observe that the size of set C is distributed as $\sum_j X_j$. The probability of outputting i is exactly

$$\Pr[P^* \text{ outputs } i \mid x, z_1, \dots, z_N] = \frac{\mu_x(z_i)}{\nu(z_i)} \mathbb{E}\left[\left(1 + \sum_{j \neq i} X_j\right)^{-1}\right].$$

Since $N \geq 1 + \exp(\epsilon n)/\epsilon$, it follows that $\mathbb{E}[\sum_{j \neq i} X_j] \geq 1/\epsilon$. Lemma C.1 in the Appendix then implies that

$$\frac{\mu_x(z_i)}{\nu(z_i) \left(1 + \sum_{j \neq i} \frac{\mu_x(z_j)}{\nu(z_j)}\right)} \leq \Pr [P^* \text{ outputs } i \mid x, z_1, \dots, z_N] \leq \exp(\epsilon) \frac{\mu_x(z_i)}{\nu(z_i) \left(1 + \sum_{j \neq i} \frac{\mu_x(z_j)}{\nu(z_j)}\right)}.$$

Since both the numerator and the denominator in the fraction change by at most $\exp(\pm\epsilon)$ when one moves from x to a neighboring database x' , the likelihood of outputting i changes by at most $\exp(\pm 4\epsilon)$. Thus, if the mechanism does not output fail, it is guaranteed to be 4ϵ -differentially private. Since the mechanism fails with probability between $\delta/(1 + \epsilon)$ and δ for all x , it follows that it has 5ϵ -differential privacy.

Repeating this argument for each of the r rounds leads to protocol with communication $O(\epsilon rn + r \log \log \frac{1}{\epsilon \delta})$ and that guarantees $5r\epsilon$ -differential privacy.

Finally, note that this protocol required the use of exponentially many bits of public randomness. By using the standard sampling-based reduction of [New95], we can reduce this to $O(\log n)$ bits of public randomness [New95], which can then be communicated by the protocol. Since the protocol was private for every choice of public random bits, this process does not affect the privacy guarantee. \square

4.6 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 4.3 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 $\Omega(C)$.

Our main tool in designing differentially private protocols is the exponential mechanism due McSherry and Talwar [MT07], whose definition and properties we recall:

Definition 4.13 (Exponential Mechanism). *A real-valued score function $q(x, r)$ is defined over the space of all possible inputs x and outputs r . For given x and privacy parameter ϵ the exponential mechanism denoted as $\mathcal{E}_q^\epsilon(x)$ outputs r with probability proportional to $\exp(-\epsilon q(x, r)/2)$.*

McSherry and Talwar prove that for a sensitivity-1 score function, the exponential mechanism satisfies ϵ -differential privacy. Moreover, if the number of possible outputs is $|\mathcal{R}|$, the loss in the value of the score function imposed by the mechanism, $q(x, \mathcal{E}_q^\epsilon(x)) - \min_r q(x, r)$, is dominated as a random variable by $2 \log |\mathcal{R}|/\epsilon + \text{Exp}(1/\epsilon)$, where $\text{Exp}(1/\epsilon)$ is the exponentially distributed random variable with parameter $1/\epsilon$. On expectation $\mathbb{E}[q(x, \mathcal{E}_q^\epsilon(x))] - \min_r q(x, r) < 4 \log |\mathcal{R}|/\epsilon$.

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. The score function $q(x, m)$, which specifies the exponential mechanism, is defined as the smallest number of bits one has to flip in the input x to make the protocol output m . More formally,

Theorem 4.14. *Let P be a deterministic protocol with communication complexity $\text{CC}(P)$ approximating a sensitivity-1 function $f: \Sigma^n \times \Sigma^n \rightarrow \mathbb{Z}$ 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 $\Delta + O(\text{CC}(P)r/\epsilon)$.*

Proof. Let π_i be the transcript up to and including the i th round of the protocol P , and let the protocol be specified as r functions $m_i(\cdot, \cdot)$, so that the first message of the protocol is $m_1(x, \pi_0)$, where π_0 is empty, the second message is $m_2(y, \pi_1)$, etc. The protocol approximates $f(x, y)$ with error Δ in the sense that $|f(x, y) - f_A(x, \pi_r)| \leq \Delta$ if r is even and $|f(x, y) - f_B(y, \pi_r)| \leq \Delta$ if r is odd for all $x, y \in \Sigma^n$.

We define a new differentially private protocol P^* by applying the exponential mechanism at each round to sample from the set of messages consistent with the transcript of the protocol so far. Assume wlog that i is odd, and let $X_i \subset \Sigma^n$ be Alice's set of inputs that are consistent with the transcript π_{i-1}^* under the original protocol P . In other words, if the j th message in π_{i-1}^* is μ_j , it holds that $\mu_j = m_j(x, \pi_{j-1}^*)$ for all $x \in X_i$ and odd $j < i$. If the length of i th message of P is k_i bits, let $M_i \subset \{0, 1\}^{k_i}$ be the set of all messages that the protocol P may output for the given transcript, i.e., $M_i = \{\mu \in \{0, 1\}^{k_i} : \exists x' \in X_i, \text{ s.t. } m_i(x', \pi_{i-1}^*) = \mu\}$. Define the score function $q: \Sigma^n \times M_i \rightarrow \mathbb{R}$ as

$$q_i(x, \mu) = \min_{m_i(x', \pi_{i-1}^*) = \mu, x' \in X_i} \|x - x'\|_1,$$

which counts the least number of bits one has to flip in x , such that P on that input would have produced the transcript π_{i-1}^* and output μ as the i th message.

Let the i th message of the new randomized protocol $P^*(x, y)$ be the output of the exponential mechanism $\mathcal{E}_{q_i}^{\epsilon/\lceil r/2 \rceil}(x)$. To compute the function, if the party receiving the last message of the protocol is Alice, she finds the closest $x' \in X_r$ to her input x and outputs $f_A(x', \pi_r^*)$, and similarly for Bob.

To prove the theorem we have to demonstrate the following three properties of the protocol P^* : (1) It is well-defined i.e., it always completes; (2) it is ϵ -differentially private; (3) its additive error is bounded as in the theorem statement. We address these points in turn.

1. P^* is well-defined. Since the i th round of P^* is the output of the exponential mechanism, the only possibility for the protocol's not completing is for M_i to be empty for some i . By construction it cannot happen, since for every feasible output μ there is an input x' , which is consistent with μ . As the sets X_i and M_i never become empty, the protocol never aborts.
2. P^* is ϵ -differentially private. P^* consists of sequential $\epsilon/\lceil r/2 \rceil$ -differentially-private applications of the exponential mechanism. The total privacy budget consumed by either party is thus at most ϵ .
3. P^* has bounded error. Let $\epsilon^* = \epsilon/\lceil r/2 \rceil$ and $K_i = \sum_{j=1}^{i-1} k_j$ —the length of π_{i-1}^* . We claim that for the closest to x element $x' \in X_i$, the distance between x and x' is dominated as a random variable by $K_i/\epsilon^* + \Gamma(i, 1/\epsilon^*)$. The proof is by induction on the round number i . For the first round X_1 includes all possible inputs, and the distance $\|x' - x\|$ is zero. For subsequent rounds, if x' is the closest to x element of X_i , the optimal (minimal) value of the score function $q_i(x, \mu)$ is $\|x' - x\|$. Since the exponential mechanism returns a message whose score exceeds the optimal value by less than $k_i/\epsilon^* + \text{Exp}(1/\epsilon^*)$, it means that there is a corresponding feasible input $x'' \in X_{i+1}$, at distance dominated by the random variable $K_{i-1}/\epsilon^* + \Gamma(i-1, 1/\epsilon^*) + k_i/\epsilon^* + \text{Exp}(1/\epsilon^*) = K_i/\epsilon^* + \Gamma(i, 1/\epsilon^*)$. Finally, when Alice approximates the value of the function by computing $f_A(x', \pi_r^*)$ for $x' \in X_r$ closest to her input x , the error of this approximation is dominated by

$$\begin{aligned} |f(x, y) - f_A(x', \pi_r^*)| &\leq |f(x, y) - f(x', y')| + |f(x', y') - f_A(x', \pi_r^*)| \\ &\leq |x - x'| + |y - y'| + \Delta = \Delta + 2K_r/\epsilon^* + \Gamma(r, 1/\epsilon^*), \end{aligned}$$

where y' is similarly defined value closest to Bob's input y and consistent with the protocol's transcript. In particular, it means that the expected approximation error is less than $\Delta + 2K_r/\epsilon^* + r/\epsilon^*$. Since $K_r > r$ and $\epsilon < 3r\epsilon^*$, the expected error of P^* is $\Delta + O(\text{CC}(P)r/\epsilon)$ as claimed. □

5 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 and 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 [CR11]. 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. [BBCR10], 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 two-party and client-server 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.

There are several immediate open questions left by our work. The most glaring is identifying the largest accuracy gap possible between two-party and client-server protocols that satisfy *approximate* differential privacy. We would also like to strengthen Theorems 4.12 and 4.14 to be independent of the number of rounds of communication, and extend Theorem 4.14 to randomized protocols.

In addition, there are connections between two-party differential privacy and *pan-privacy* [DNP⁺10]. 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 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 [DNP⁺10], 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, for a constant number of passes. (While not defined in [DNP⁺10], it is also natural to consider multi-pass pan-private algorithms.) Indeed, by a straightforward reduction, a k -pass pan-private algorithm for density estimation implies a $k\epsilon$ -differentially private two-party protocol for estimating the Hamming distance of two binary strings, with similar error. We sketch a proof of this observation in Appendix D. What further limitations for pan-privacy can be obtained?

References

- [BBCR10] Boaz Barak, Mark Braverman, Xi Chen, and Anup Rao. How to compress interactive communication. In Leonard J. Schulman, editor, *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010*, pages 67–76. ACM, 2010.
- [BNO08] Amos Beimel, Kobbi Nissim, and Eran Omri. Distributed private data analysis: Simultaneously solving how and what. In David Wagner, editor, *Advances in Cryptology—CRYPTO 2008*, volume 5157 of *Lecture Notes in Computer Science*, pages 451–468. Springer, 2008.

- [BYJKS02] Ziv Bar-Yossef, T. S. Jayram, Ravi Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. In *43rd Symposium on Foundations of Computer Science, FOCS 2002*, pages 209–218. IEEE Computer Society, 2002.
- [CR11] Amit Chakrabarti and Oded Regev. An optimal lower bound on the communication complexity of gap-hamming-distance. In *Proceedings of the 43rd annual ACM symposium on Theory of computing, STOC '11*, pages 51–60, New York, NY, USA, 2011. ACM.
- [CS72] M. T. Chao and W. E. Strawderman. Negative moments of positive random variables. *Journal of the American Statistical Association*, 67(338):429–431, 1972.
- [CSWY01] Amit Chakrabarti, Yaoyun Shi, Anthony Wirth, and Andrew C. Yao. Informational complexity and the direct sum problem for simultaneous message complexity. In *Proc. 42nd Annual IEEE Symposium on Foundations of Computer Science*, pages 270–278, 2001.
- [De11] Anindya De. Lower bounds in differential privacy. Technical report, arXiv:1107.2183v1, July 2011.
- [DKM⁺06] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: privacy via distributed noise generation. In *Advances in Cryptology—EUROCRYPT 2006*, volume 4004 of *Lecture Notes in Computer Science*, pages 486–503. Springer, 2006.
- [DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2006.
- [DN04] Cynthia Dwork and Kobbi Nissim. Privacy-preserving datamining on vertically partitioned databases. In Matthew K. Franklin, editor, *Advances in Cryptology—CRYPTO 2004*, volume 3152 of *Lecture Notes in Computer Science*, pages 528–544. Springer, 2004.
- [DNP⁺10] Cynthia Dwork, Moni Naor, Toniann Pitassi, Guy Rothblum, and Sergey Yekhanin. Pan-private streaming algorithms. In *Proceedings of the First Symposium on Innovations in Computer Science (ICS 2010)*. Tsinghua University Press, Beijing, 2010.
- [DO03] Yevgeniy Dodis and Roberto Oliveira. On extracting private randomness over a public channel. In Sanjeev Arora, Klaus Jansen, José D. P. Rolim, and Amit Sahai, editors, *RANDOM-APPROX*, volume 2764 of *Lecture Notes in Computer Science*, pages 252–263. Springer, 2003.
- [DP09] Devdatt Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomised Algorithms*. Cambridge University Press, 2009.
- [Dwo06] Cynthia Dwork. Differential privacy. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Part II*, volume 4052 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2006.
- [EGL85] Shimon Even, Oded Goldreich, and Abraham Lempel. A randomized protocol for signing contracts. *Communications of the ACM*, 28(6):637–647, 1985.

- [GKY11] Adam Groce, Jonathan Katz, and Arkady Yerukhimovich. Limits of computational differential privacy in the client/server setting. In Yuval Ishai, editor, *TCC*, volume 6597 of *Lecture Notes in Computer Science*, pages 417–431. Springer, 2011.
- [GRS09] Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan. Universally utility-maximizing privacy mechanisms. In Michael Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009*, pages 351–360. ACM, 2009.
- [Hol09] Thomas Holenstein. Parallel repetition: Simplification and the no-signaling case. *Theory of Computing*, 5(1):141–172, 2009.
- [JK10] Rahul Jain and Hartmut Klauck. The partition bound for classical communication complexity and query complexity. In *Annual IEEE Conference on Computational Complexity*, 2010.
- [KO97] Eyal Kushilevitz and Rafail 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)*, pages 364–373. IEEE Computer Society, 1997.
- [Man94] Udi Manber. Finding similar files in a large file system. In *USENIX Winter*, pages 1–10, 1994.
- [MMNW10] Darakhshan Mir, S. Muthukrishnan, Aleksandar Nikolov, and Rebecca Wright. Pan-private algorithms: When memory does not help. In *Computer Science arXiv*, 2010.
- [MMP⁺10] Andrew McGregor, Ilya Mironov, Toniann Pitassi, Omer Reingold, Kunal Talwar, and Salil Vadhan. The limits of two-party differential privacy. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS ‘08)*. IEEE, 23–26 October 2010.
- [MPRV09] Ilya Mironov, Omkant Pandey, Omer Reingold, and Salil P. Vadhan. Computational differential privacy. In Shai Halevi, editor, *CRYPTO*, volume 5677 of *Lecture Notes in Computer Science*, pages 126–142. Springer, 2009.
- [MT07] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007)*, pages 94–103. IEEE Computer Society, 2007.
- [New95] Ilan Newman. Private vs. common random bits in communication complexity. *Information Processing Letters*, 39:67–71, 1995.
- [Rab81] Michael O. Rabin. How to exchange secrets with oblivious transfer. Technical Report TR-81, Aiken Computation Lab, Harvard University, May 1981.
- [RVW04] Omer Reingold, Salil Vadhan, and Avi Wigderson. A note on extracting randomness from Santha–Vazirani sources. Unpublished manuscript, 2004.
- [SV86] Miklós Sántha and Umesh V. Vazirani. Generating quasirandom sequences from semirandom sources. *Journal of Computer and System Sciences*, 33(1):75–87, 1986.
- [Vaz87] Umesh V. Vazirani. Strong communication complexity or generating quasirandom sequences from two communicating semirandom sources. *Combinatorica*, 7(4):375–392, 1987.

- [War65] Stanley L. Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, March 1965.
- [Yao79] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (preliminary report). In *STOC '79: Proceedings of the eleventh annual ACM symposium on Theory of computing*, pages 209–213, New York, NY, USA, 1979. ACM.
- [Yao82] Andrew Chi-Chih Yao. Protocols for secure computations (extended abstract). In *23rd Annual Symposium on Foundations of Computer Science (FOCS 1982)*, pages 160–164. IEEE, 1982.

A Approximate Differential Privacy and Unpredictable Sources

In this section, we generalize the results of the previous section to approximate differential privacy. We will do this by showing that approximate differentially private protocols give rise to the following approximate forms of unpredictable bit sources:

Definition A.1 (δ -approximate (strongly) α -unpredictable bit source). *For $\alpha \in [0, 1]$, a random variable $X = (X_1, \dots, X_n)$ taking values in $\{0, 1\}^n$ is a δ -approximate α -unpredictable bit source if with probability at least $1 - \delta$ over $i \leftarrow [n]$ and $(x_1, \dots, x_{i-1}) \leftarrow (X_1, \dots, X_{i-1})$, we have*

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

A random variable $X = (X_1, \dots, X_n)$ taking values in $\{0, 1\}^n$ is a δ -approximate strongly α -unpredictable bit source if with probability at least $1 - \delta$ over $i \leftarrow [n]$ and $(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n) \leftarrow (X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_n)$, we have

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

Here δ should be thought of as negligible in n , while α remains a constant. Now we relate the above two definitions (showing that the strong definition implies the other one, albeit with some loss in parameters), and prove that approximately unpredictable sources are statistically close to standard unpredictable sources. This will allow us to apply our extractor from the previous section to these sources as a black box.

Lemma A.2.

1. *If X is a δ -approximate strongly α -unpredictable bit source, then for every $\nu > 0$, X is a δ/ν -approximate α' -unpredictable bit source for*

$$\alpha' = \alpha \cdot \frac{1 - \nu}{1 + \alpha\nu} \geq \alpha \cdot \frac{1 - \nu}{1 + \nu}.$$

2. *If X is a δ -approximate α -unpredictable bit source, then X is $n\delta$ -close to some α -unpredictable bit source.*

Proof. 1. Let

$$B = \left\{ (i, x_1 \dots x_{i-1}, x_{i+1} \dots x_n) : \frac{\Pr[X_i = 0 | X_1 = x_1, \dots, X_{i-1} = x_{i-1}, X_{i+1} = x_{i+1}, \dots, X_n = x_n]}{\Pr[X_i = 1 | X_1 = x_1, \dots, X_{i-1} = x_{i-1}, X_{i+1} = x_{i+1}, \dots, X_n = x_n]} \notin [\alpha, 1/\alpha] \right\}.$$

We know that $\Pr[(I, X_1 \dots X_{I-1}, X_{I+1} \dots X_n) \in B] \leq \delta$, where I is a uniformly random element of $[n]$. Let

$$B' = \{(i, x_1 \dots x_{i-1}) : \Pr[(i, x_1 \dots x_{i-1}, X_{i+1} \dots X_n) \in B | X_1 = x_1, \dots, X_{i-1} = x_{i-1}] > \nu\}.$$

It follows by Markov's inequality that $\Pr[(I, X_1 \dots X_{I-1}) \in B'] \leq \delta/\nu$.

Now, for every $(i, x_1 \dots x_{i-1}) \notin B'$, we have

$$\Pr[X_i = 0 \mid X_1 = x_1, \dots, X_{i-1} = x_{i-1}] \leq \nu + \Pr[X_i = 1 \mid X_1 = x_1, \dots, X_{i-1} = x_{i-1}]/\alpha.$$

Writing $p = \Pr[X_i = 0 \mid X_1 = x_1, \dots, X_{i-1} = x_{i-1}]$ and $1 - p = \Pr[X_i = 1 \mid X_1 = x_1, \dots, X_{i-1} = x_{i-1}]$, we see that $p \leq (1 + \alpha\nu)/(1 + \alpha)$. Similarly, it can be shown that $1 - p \leq (1 + \alpha\nu)/(1 + \alpha)$. This implies that $p/(1 - p)$ and $(1 - p)/p$ are both at least

$$\alpha' = \frac{1 - \left(\frac{1 + \alpha\nu}{1 + \alpha}\right)}{\left(\frac{1 + \alpha\nu}{1 + \alpha}\right)} = \frac{\alpha \cdot (1 - \nu)}{1 + \alpha\nu}.$$

2. Call a prefix $x \in \{0, 1\}^{<n}$ *bad* if

$$\frac{\Pr[X_i = 0 \mid X_1 = x_1, \dots, X_{i-1} = x_{i-1}]}{\Pr[X_i = 1 \mid X_1 = x_1, \dots, X_{i-1} = x_{i-1}]} \notin [1/\alpha, \alpha].$$

Define X' as follows:

- (a) Sample $x \leftarrow X$.
- (b) If x has no bad prefix, output x .
- (c) Else let $x_1 \dots x_i$ be the shortest bad prefix of x . Choose $x'_{i+1}, \dots, x'_n \leftarrow \{0, 1\}$. Output $x_1 \dots x_i x'_{i+1} \dots x'_n$.

It can be verified that X' is an α -unpredictable bit source. X' is $n\delta$ -close to X because for each $i \in \{0, \dots, n - 1\}$, a random sample $x \leftarrow X$ has a bad prefix of length i with probability at most δ . \square

As a first step towards relating approximate differentially private protocols to approximate unpredictable sources, we show that approximate differentially private mechanisms behave like differentially private ones with high probability, in the following sense:

Lemma A.3. *Let $M: \{0, 1\}^n \rightarrow \mathcal{R}$ be a δ -approximate, ϵ -differentially private mechanism. Then for every $\gamma > 0$, and every $x, y \in \{0, 1\}^n$ such that $|x - y|_H = 1$, if we generate $m \leftarrow M(x)$, then we have*

$$e^{-(\epsilon + \gamma)} \leq \frac{\Pr[M(y) = m]}{\Pr[M(x) = m]} \leq e^{\epsilon + \gamma}$$

with probability at least $1 - \delta'$, for

$$\delta' = \delta \cdot \frac{1 + e^{-\epsilon - \gamma}}{1 - e^{-\gamma}}.$$

Note that for small values of $\epsilon, \gamma > 0$, δ' is roughly $2\delta/\gamma$.

Proof. Let $S = \{m: \Pr[M(x) = m] < e^{-\epsilon - \gamma} \cdot \Pr[M(y) = m]\}$. Then

$$\begin{aligned} \Pr[M(x) \in S] &< e^{-\epsilon - \gamma} \cdot \Pr[M(y) \in S] \\ &\leq e^{-\epsilon - \gamma} \cdot (e^\epsilon \cdot \Pr[M(x) \in S] + \delta) \\ &\leq e^{-\gamma} \cdot \Pr[M(x) \in S] + e^{-\epsilon - \gamma} \cdot \delta, \end{aligned}$$

and thus

$$\Pr[M(x) \in S] \leq \delta \cdot \frac{e^{-\epsilon - \gamma}}{1 - e^{-\gamma}}.$$

Similarly, for $T = \{m: \Pr[M(x) = m] > e^{\epsilon+\gamma} \cdot \Pr[M(y) = m]\}$, we have

$$\begin{aligned} \Pr[M(x) \in T] &\leq e^\epsilon \cdot \Pr[M(y) \in T] + \delta \\ &< e^\epsilon \cdot e^{-\epsilon-\gamma} \cdot \Pr[M(x) \in T] + \delta \\ &= e^{-\gamma} \cdot \Pr[M(x) \in T] + \delta, \end{aligned}$$

so

$$\Pr[M(x) \in T] < \delta \cdot \frac{1}{1 - e^{-\gamma}},$$

and the probability that $M(x)$ ends up either in S or T is bounded as

$$\Pr[M(x) \in S \cup T] < \delta \cdot \frac{1 + e^{-\epsilon-\gamma}}{1 - e^{-\gamma}}.$$

□

Now we show that approximate differentially private protocols give rise to approximate unpredictable sources:

Lemma A.4. *Let P be a randomized protocol with δ -approximate ϵ -differential privacy. Let X and Y be independent random variables uniformly distributed in $\{0, 1\}^n$, and let random variable Π denote the transcript of P when run on input (X, Y) . Then for every $\gamma > 0$, there are numbers $\{\delta_\pi\}_{\pi \in \text{Supp}(\Pi)}$ such that*

1. *For every $\pi \in \text{Supp}(\Pi)$, the random variables X_π and Y_π are independent δ_π -approximate strongly $e^{-(\epsilon+\gamma)}$ -unpredictable bit sources.*
2. $\mathbb{E}[\delta_\Pi] \leq \delta'$, for

$$\delta' = 2\delta \cdot \frac{1 + e^{-\epsilon-\gamma}}{1 - e^{-\gamma}}.$$

Proof. For $i \in [n]$ and $(x, \pi) \in \text{Supp}(X, T)$, define

$$\begin{aligned} \rho(i, x, t) &= \frac{\Pr[X_i = 0 \mid X_1 = x_1, \dots, X_{i-1} = x_{i-1}, X_{i+1} = x_{i+1}, \dots, X_n = x_n, \Pi = \pi]}{\Pr[X_i = 1 \mid X_1 = x_1, \dots, X_{i-1} = x_{i-1}, X_{i+1} = x_{i+1}, \dots, X_n = x_n, \Pi = \pi]} \\ &= \frac{\Pr[\Pi(x_1 \cdots x_{i-1} 0 x_{i+1} \cdots x_n, Y) = \pi]}{\Pr[\Pi(x_1 \cdots x_{i-1} 1 x_{i+1} \cdots x_n, Y) = \pi]}, \end{aligned}$$

where the equality is by the same calculation in the proof of Lemma 3.3. Let $B = \{(i, x, \pi): \rho(i, x, \pi) \notin [e^{-\epsilon-\gamma}, e^{\epsilon+\gamma}]\}$. By Lemma A.3, we have $\Pr[(i, X, \Pi) \in B] \leq \delta'/2$.

Similarly, we can switch the role of X and Y , and define $\rho'(i, y, \pi)$ for $(y, \pi) \in \text{Supp}(Y, \Pi)$ and a corresponding set B' . Again we have $\Pr[(i, Y, \Pi) \in B'] \leq \delta'/2$.

Now, for $\pi \in \text{Supp}(\Pi)$, if we define $\delta_\pi = \max\{\Pr[(I, X_\pi) \in B], \Pr[(I, Y_\pi) \in B']\}$, then X_π and Y_π are both δ_π -approximate strongly $(\epsilon + \gamma)$ -unpredictable bit sources. We have

$$\begin{aligned} \mathbb{E}[\delta_\Pi] &\leq \mathbb{E}_{\pi \leftarrow \Pi}[\Pr[(I, X_\pi) \in B] + \Pr[(I, Y_\pi) \in B']] \\ &= \Pr[(I, X, \Pi) \in B] + \Pr[(I, Y, \Pi) \in B'] \\ &\leq \delta'/2 + \delta'/2, \end{aligned}$$

where the last inequality is by approximate differential privacy and Lemma A.3. (Indeed, it holds even for each fixed value of i .) □

Finally, we again use the randomness extraction properties of the inner product function to show that differentially private protocols must incur an error of nearly \sqrt{n} .

Theorem A.5. *Let P be a randomized protocol with δ -approximate ϵ -differential privacy. Then for every $\delta' \geq 48n\delta$, 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*

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

Notice that the error bound here is the same as in Theorem 3.9, up to the hidden constant and the constraint $\delta' \geq 48n\delta$. The latter constraint is very reasonable as we typically think of δ as negligible in n , i.e., $\delta = n^{-\omega(1)}$.

Proof. Let X and Y be uniform and independent in $\{0, 1\}^n$ and Π the random transcript on input (X, Y) . Party B 's output is a function $f_B(Y, \Pi)$. Let $m = 6\Delta/\delta$, $\gamma = 1$ and $\nu = 1/2$.

By Lemmas A.2 and A.4, we know that there are numbers $\{\delta_\pi\}_{\pi \in \text{Supp}(\Pi)}$ such that

1. for every $\pi \in \text{Supp}(\Pi)$, the random variables X_π and Y_π are each $n\delta_\pi/\nu$ -close to an α -unpredictable bit-source for

$$\alpha = e^{-\epsilon-\gamma} \cdot \frac{1-\nu}{1+\nu} = \Omega(e^\epsilon),$$

and

2. $E[\delta_\Pi] \leq \delta''$, for

$$\delta'' = 2\delta \cdot \frac{1 + e^{-\epsilon-\gamma}}{1 - e^{-\gamma}} \leq 8\delta.$$

Setting $\beta = \log(1 + \alpha)$, item 1 and Theorem 3.4 imply that $(Y_\pi, \langle X_\pi, Y_\pi \rangle \bmod m)$ has statistical distance at most $n\delta_t/\nu + \delta'/3$ from (Y_π, U) where, as in the proof of Theorem 3.9, this follows provided:

$$\begin{aligned} \Delta &\geq c_2 \cdot \frac{\sqrt{n}}{\log n} \cdot \frac{\delta'/3}{e^{\epsilon+\gamma}} \\ &= c_3 \cdot \frac{\sqrt{n}}{\log n} \cdot \frac{\delta'}{e^\epsilon}, \end{aligned}$$

as guaranteed by the hypothesis of the theorem.

Now we observe that the statistical distance between $(\Pi, Y, \langle X, Y \rangle \bmod m)$ and (Π, Y, U) is at most $E[n\delta_\Pi/\nu + \delta'/3] \leq n \cdot (8\delta) \cdot 2 + \delta'/3 = 2\delta'/3$, so the probability of an error at most Δ is now at most $2\delta'/3 + 2\Delta/m = \delta'$. \square

B Missing Lemma from Section 4.1

The following lemma is used in the proof of proposition 4.4.

Lemma B.1. *Let X, Y, Z be random variables such that X and Y are independent. Then $I(X; Z) \leq I(X; Z|Y)$*

Proof.

$$\begin{aligned} I(X; Z) &= H(X) - H(X|Z) \\ &= H(X|Y) - H(X|Z) \\ &\leq H(X|Y) - H(X|ZY) \\ &= I(X; Z|Y). \end{aligned}$$

Here in the second step we have used the fact that X and Y are independent, and in the third step that conditioning decreases entropy. \square

C Expectation of the inverse

In this section, we prove that

Lemma C.1. *Let $\{X_i\}_{i=1}^N$ be a sequence of independent Bernoulli random variables with $E[X_i] = \mu_i$ and let $\mu \stackrel{\text{def}}{=} \sum_i \mu_i$. Then*

$$\frac{1}{1 + \mu} \leq E \left[\left(1 + \sum_{i=1}^N X_i \right)^{-1} \right] \leq \frac{1}{\mu}.$$

Proof. The first inequality follows by applying Jensen's inequality to the function $\phi(X) = 1/(1+X)$, which is convex for $X \geq 0$.

To prove the second inequality, we use a result from [CS72] who gave a formula for negative moments of random variables. We reproduce the proof of the case we use for completeness. Observe that for every $t, x > 0$,

$$\frac{t^x}{x} = \int_0^t u^{x-1} du.$$

Setting $t = 1$ and taking expectations over random x , we get

$$E\left[\frac{1}{X}\right] = \int_0^1 E[u^{X-1}] du.$$

For $X = 1 + \sum_i X_i$, we upper bound

$$\begin{aligned} E[u^{X-1}] &= \prod_i E[u^{X_i}] \\ &= \prod_i (1 - (1-u)\mu_i) \\ &\leq \prod_i \exp(-(1-u)\mu_i) \\ &= \exp(-(1-u)\mu). \end{aligned}$$

We conclude that

$$E\left[\frac{1}{X}\right] \leq \exp(-\mu) \int_0^1 \exp(\mu u) du = \frac{1 - \exp(-\mu)}{\mu} < \frac{1}{\mu}.$$

\square

D Pan-Privacy and Differentially Private Communication

In this section, we explore the connection between differentially private communication protocols and pan-private stream algorithms [DNP⁺10]. Pan-private stream algorithms are stream algorithms where at every point in time, the internal state of the algorithm is differentially private. A well-known connection between communication complexity and stream algorithms establishes space lower bounds for stream algorithms from two-party communication complexity lower bounds. Similarly, in this section, we will use the same

connection to establish negative results (on the error) for pan-private stream algorithms from our two-party differentially private communication complexity lower bounds. Other upper and lower bounds for pan-private algorithms have been obtained in [DNP⁺10] and more recently in [MMNW10].

For the notion of pan-privacy, we consider an algorithm \mathcal{A} that processes a stream of data items $S = \langle s_1, s_2, \dots, s_m \rangle \in \Sigma^*$.

Definition D.1. *We say streams $S, S' \in \Sigma^*$ are Σ -adjacent if they differ only in the presence or absence of any number of occurrences of a single element $\sigma \in \Sigma$.*

As the algorithm processes each symbol, it updates its internal state according to some probabilistic mapping. The idea behind pan-privacy is that an adversary should not be able to distinguish between two Σ -adjacent streams by inspecting the internal state of the algorithm. The following definition makes this precise in the context of algorithms that may take multiple passes over the data stream.³

Definition D.2 (Multi-Pass Pan-Privacy). *Let \mathcal{A} be an algorithm that takes p passes over a finite length stream S , maintaining an internal state from a state space \mathcal{M} during its computation, and outputs a value $\text{out}(S) \in \mathcal{O}$. (One pass is defined to be a left-to-right scan over the input.) Let B be an adversary that can make intrusions during the computation of \mathcal{A} (getting to see the internal state of \mathcal{A} at time steps chosen by B). The time steps of the intrusions can be chosen adaptively (chosen based on what B has seen so far) or non-adaptively (fixed prior to the computation), and announced (\mathcal{A} is notified of the intrusions and can modify its state based on them) or unannounced (the computation of \mathcal{A} is oblivious to the intrusions). Let $\text{VIEW}_{\mathcal{A}}^B(S)$ denote the view of adversary B when making intrusions to the computation of \mathcal{A} on input stream S , and also seeing the final output of \mathcal{A} . (So if B makes t intrusions, then $\text{VIEW}_{\mathcal{A}}^B(S) \in \mathcal{M}^t \times \mathcal{O}$.)*

We say that \mathcal{A} is ϵ -differentially private against B if for every two data streams $S, S' \in \Sigma^$ that are Σ -adjacent, and every set T of possible views of B , we have*

$$\Pr[\text{VIEW}_{\mathcal{A}}^B(S) \in T] \leq e^\epsilon \cdot \Pr[\text{VIEW}_{\mathcal{A}}^B(S') \in T].$$

((ϵ, δ) -differential privacy is defined analogously, allowing an additional additive δ term.)

By exploiting the well-known connection between communication protocols and streaming algorithms, we deduce the following theorem from our lower bounds for differentially private two-party protocols:

Theorem D.3. *Any p -pass estimator for the density $F_0(S) = |\{\sigma : \sigma \in S\}|$, that is ϵ -pan-private for all adversaries that make $2p - 1$ non-adaptive announced intrusions, has error $\Omega(\delta e^{-2\epsilon} \sqrt{n/\log n})$ with probability $1 - \delta$.*

Proof. Let \mathcal{A} be a p -pass, ϵ -pan-private density estimator with error E . Suppose two players Alice and Bob, with inputs $x, y \in \{0, 1\}^n$ respectively, want to estimate the Hamming distance $|x - y|_H$. They may use the stream algorithm \mathcal{A} to construct a communication protocol as follows. Alice generates the stream $S_x = \langle i : x_i = 1 \rangle$ and Bob generates the stream $S_y = \langle i : y_i = 1 \rangle$. Consider the stream $S_x S_y$ and note that

$$F_0(S_x S_y) = |x|_H/2 + |y|_H/2 + |x - y|_H/2 .$$

Hence, an additive E approximation of F_0 yields an additive $2E + O(1/\epsilon)$ approximation of $|x - y|_H$ if $|x|_H + |y|_H$ is known up to additive error $O(1/\epsilon)$. To get an approximation of $F_0(S_x S_y)$, Alice runs \mathcal{A} on S_x and transmits the resulting internal state to Bob who instantiates \mathcal{A} with this state and continues running

³The corresponding definition for single-pass algorithms in [DNP⁺10] adds a further condition for pan-privacy. There, the authors suppose that the algorithm may be designed to output a sequence of values as the stream is processed. In that case, it is required that this sequence also does not compromise privacy. For the purposes of our lower bounds we may ignore the output sequence other than the final output.

the algorithm on S_y . If $p > 1$, he then sends the new memory state back to Alice and they continue in this manner until all p passes of the algorithm have been emulated. Note that because \mathcal{A} is ϵ -pan-private, the resulting transcript is ϵ -DP. The transcript can be augmented with a $O(1/\epsilon)$ approximation of $|x|_H$ and still be 2ϵ -DP. However, Theorem 3.9 states that any protocol whose transcript is a 2ϵ -DP mechanism for the player's inputs has error $\Omega(\delta e^{-2\epsilon} \sqrt{n}/\log n)$ with probability $1 - \delta$. Hence we deduce $E = \Omega(\delta e^{-2\epsilon} \sqrt{n}/\log n)$ with probability $1 - \delta$. \square

Note that the special case of $p = 1$ in the above theorem gives an accuracy lower-bound for only a single announced intrusion. Also, by using Theorem A.5 instead of Theorem 3.9, we also get a lower bound for approximate differential privacy.

We now present a result for multi-pass pan-private stream algorithms for the case of *two* announced intrusions. The result follows from the observation that the communication protocols that arise by simulating stream algorithms have a particular ‘‘Markovian’’ property which we now define.

Definition D.4 (Markovian Communication Protocols). *Consider a communication protocol P between two parties A and B in which the players take turns to send messages back and forth. Suppose there are t messages and let M_i denote the i th message sent in this protocol: if i is odd then M_i is sent by A and if i is even then M_i is sent by B . We say the protocol is Markovian if for all $i \geq 2$ and messages m_1, \dots, m_i*

$$\Pr[M_i = m_i | M_1 = m_1, M_2 = m_2, \dots, M_{i-1} = m_{i-1}] = \Pr[M_i = m_i | M_{i-1} = m_{i-1}] .$$

The following lemma establishes that if the transcript of a Markovian protocol compromises privacy, then there exists two successive messages in the protocol that together also compromise privacy.

Lemma D.5. *For a t -message Markovian communication protocol P , consider the transcript Π of the protocol as a mechanism for the input of the players $Z = (X, Y)$. If Π is not ϵ -DP then there exists two successive messages in the protocol M_i and M_{i+1} such that (M_i, M_{i+1}) is not $\epsilon/2t$ -DP.*

Proof. If Π is not ϵ -DP, there exists a sequence of t messages, m_1, \dots, m_t , such that for two adjacent inputs $z = (x, y)$ and $z' = (x', y')$,

$$\left| \ln \frac{\Pr_z[\Pi = (m_1, \dots, m_t)]}{\Pr_{z'}[\Pi = (m_1, \dots, m_t)]} \right| > \epsilon .$$

By the Markov property of the protocol we may write:

$$\frac{\Pr_z[\Pi = (m_1, \dots, m_t)]}{\Pr_{z'}[\Pi = (m_1, \dots, m_t)]} = \frac{\Pr_z[M_1 = m_1] \cdot \prod_{i \in [t-1]} \Pr_z[M_{i+1} = m_{i+1} | M_i = m_i]}{\Pr_{z'}[M_1 = m_1] \cdot \prod_{i \in [t-1]} \Pr_{z'}[M_{i+1} = m_{i+1} | M_i = m_i]}$$

and hence

$$\epsilon < \left| \ln \frac{\Pr_z[M_1 = m_1]}{\Pr_{z'}[M_1 = m_1]} + \sum_{i \in [t-1]} \ln \frac{\Pr_z[M_{i+1} = m_{i+1} | M_i = m_i]}{\Pr_{z'}[M_{i+1} = m_{i+1} | M_i = m_i]} \right| .$$

Consequently, either $\left| \ln \frac{\Pr_z[M_1 = m_1]}{\Pr_{z'}[M_1 = m_1]} \right| > \epsilon/t$ or there exists $i \in [t-1]$ such that $\left| \ln \frac{\Pr_z[M_{i+1} = m_{i+1} | M_i = m_i]}{\Pr_{z'}[M_{i+1} = m_{i+1} | M_i = m_i]} \right| > \epsilon/t$. Assume $\left| \ln \frac{\Pr_z[M_i = m_i]}{\Pr_{z'}[M_i = m_i]} \right| \leq \epsilon/(2t)$ for all $i \in [t]$, otherwise we are done. Therefore,

$$\left| \ln \frac{\Pr_z[M_i = m_i, M_{i+1} = m_{i+1}]}{\Pr_{z'}[M_i = m_i, M_{i+1} = m_{i+1}]} \right| = \left| \ln \frac{\Pr_z[M_i = m_i]}{\Pr_{z'}[M_i = m_i]} + \ln \frac{\Pr_z[M_{i+1} = m_{i+1} | M_i = m_i]}{\Pr_{z'}[M_{i+1} = m_{i+1} | M_i = m_i]} \right| > \frac{\epsilon}{t} - \frac{\epsilon}{2t} = \frac{\epsilon}{2t} ,$$

as required. \square

Note that above lemma is not true for communication protocols that are not Markovian. Consider an arbitrary t message protocol P that compromises privacy. This can be transformed into a $t + k$ message protocol Q such that observing any k messages does not compromise the privacy at all: in message $i \in [k]$ of Q the players exchange random strings R_{i1}, \dots, R_{it} and the $(k + j)$ th message sent is $M_j \oplus R_{1j} \oplus R_{2j} \oplus \dots \oplus R_{kj}$.

Theorem D.6. *Any p -pass density estimator that is ϵ -pan-private for all adversaries that make two announced intrusions, has error $\Omega(\delta e^{-4(2p-1)\epsilon} \sqrt{n}/\log n)$ with probability $1 - \delta$.*

Proof. The proof follows along similar lines to Theorem D.3 except this time \mathcal{A} is a p -pass stream algorithm that maintains ϵ -pan-privacy under two announced intrusions. Hence, the emulation of \mathcal{A} gives a $2p - 1$ message communication protocol such that observing any two of the messages maintains 2ϵ -DP. Hence, by Lemma D.5, this implies the entire transcript maintains $4(2p - 1)\epsilon$ -DP. By appealing to Theorem 3.9, we conclude that the error is $\Omega(\delta e^{-4(2p-1)\epsilon} \sqrt{n}/\log n)$ with probability $1 - \delta$. \square