

# Do Distributed Differentially-Private Protocols Require Oblivious Transfer?\*

Vipul Goyal<sup>1</sup>, Dakshita Khurana<sup>2</sup>, Ilya Mironov<sup>3</sup>, Omkant Pandey<sup>4</sup>, and Amit Sahai<sup>5</sup>

- 1 Microsoft Research India, Bangalore, India  
vipul@microsoft.com
- 2 UCLA and Center for Encrypted Functionalities, Los Angeles, USA  
dakshita@cs.ucla.edu
- 3 Google, USA  
mironov@google.com
- 4 Drexel University, Philadelphia, USA  
omkant@drexel.edu
- 5 UCLA and Center for Encrypted Functionalities, Los Angeles, USA  
sahai@cs.ucla.edu

---

## Abstract

We study the cryptographic complexity of two-party differentially-private protocols for a large natural class of boolean functionalities. Information theoretically, McGregor et al. [FOCS 2010] and Goyal et al. [Crypto 2013] demonstrated several functionalities for which the maximal possible accuracy in the distributed setting is significantly lower than that in the client-server setting. Goyal et al. [Crypto 2013] further showed that “highly accurate” protocols in the distributed setting for any non-trivial functionality in fact imply the existence of one-way functions. However, it has remained an open problem to characterize the exact cryptographic complexity of this class. In particular, we know that semi-honest oblivious transfer helps obtain optimally accurate distributed differential privacy. But we do not know whether the reverse is true.

We study the following question: *Does the existence of optimally accurate distributed differentially private protocols for any class of functionalities imply the existence of oblivious transfer (or equivalently secure multi-party computation)?* We resolve this question in the affirmative for the class of boolean functionalities that contain an XOR embedded on adjacent inputs. We give a reduction from oblivious transfer to:

- Any distributed optimally accurate  $\epsilon$ -differentially private protocol with  $\epsilon > 0$  computing a functionality with a boolean XOR embedded on adjacent inputs.
- Any distributed non-optimally accurate  $\epsilon$ -differentially private protocol with  $\epsilon > 0$ , for a constant range of non-optimal accuracies and constant range of values of  $\epsilon$ , computing a functionality with a boolean XOR embedded on adjacent inputs.

Enroute to proving these results, we demonstrate a connection between optimally-accurate two-party differentially-private protocols for functions with a boolean XOR embedded on adjacent inputs, and noisy channels, which were shown by Crépeau and Kilian [FOCS 1988] to be sufficient for oblivious transfer.

---

\* Research of the second and fifth authors supported in part from a DARPA/ARL SAFEWARE award, NSF Frontier Award 1413955, NSF grants 1228984, 1136174, 1118096, and 1065276, a Xerox Faculty Research Award, a Google Faculty Research Award, an equipment grant from Intel, and an Okawa Foundation Research Grant. This material is based upon work supported by the Defense Advanced Research Projects Agency through the ARL under Contract W911NF-15-C-0205. The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense, the National Science Foundation, or the U.S. Government. The work of the third author was done while at Microsoft Research. The full version of this paper is available on-line at <http://eprint.iacr.org/2015/1090>.



1998 ACM Subject Classification E.3 Data Encryption

Keywords and phrases Oblivious Transfer, Distributed Differential Privacy, Noisy Channels, Weak Noisy Channels

Digital Object Identifier 10.4230/LIPIcs.ICALP.2016.29

## 1 Introduction

*Differential privacy* [7, 8, 11, 14] is a very well-studied and popular privacy notion of recent years<sup>1</sup>. It provides powerful input privacy guarantees to participants of a statistical query database. Informally a randomized function computed on a database is said to be differentially private, if the output distribution of the function evaluated on the database, induced by the presence of a particular record, is statistically close to the output distribution induced when the record is absent. Note that this can be trivially achieved by computing an output that is independent of the entries in the database. Therefore, to be useful, a non-trivial differentially private algorithm must compute outputs that satisfy some meaningful notion of accuracy.

Consider a confidential dataset owned by a trusted server. The server must release the output of some statistic evaluated on the dataset, to an untrusted client. Even in this setting, where privacy is a concern only at the server's end, there is an evident tradeoff between privacy and accuracy. In fact, for any given privacy parameter  $\epsilon$ , there is a maximum possible accuracy (which we call the *optimal accuracy*) such that any algorithm with better than optimal accuracy will fail to remain differentially private. Such privacy-accuracy tradeoffs are reasonably well-understood in the client-server setting [7, 12, 15, 23]. There has also been a huge body of work in designing algorithms that achieve close to optimal accuracies for various functionalities and data mining tasks in the client-server setting.

The focus of this work is the distributed setting, where a database is jointly hosted by multiple mutually distrusting servers. This was first studied by Dwork et al. [10]. As an illustrative example, consider two hospitals which together wish to compute the correlation between the occurrence of smoking and lung cancer by taking into account their combined patient records. In this setting, we require the servers to engage in a protocol, at the end of which the privacy of each record of both the servers must be guaranteed without a significant loss in accuracy. Note that the privacy requirements must be met for both servers, *given their entire view of the protocol transcript*, not just the computed output; possibly necessitating an additional loss in accuracy (over and above the loss in the client-server setting).

The intuition that the distributed setting would necessitate a greater accuracy loss than the client-server setting has been proved to be correct in the information theoretic world for different classes of functions. Beimel, Nissim and Omri [1] showed accuracy limits for distributed differentially-private protocols for  $n$  parties each holding their own inputs. McGregor, Mironov, Pitassi, Reingold, Talwar and Vadhan [36] showed large accuracy gaps in the two-party setting for several natural functionalities with  $n$ -bit inputs. Goyal, Mironov, Pandey and Sahai [19] demonstrated a constant gap between the maximal achievable accuracies in the client-server and distributed settings for any non-trivial boolean functionality.

In the computational setting this gap vanishes, if a semi-honest protocol for oblivious transfer exists. In this case, both servers can use secure multi-party computation [18] to simulate the client-server differentially private function evaluation, thereby obtaining the

---

<sup>1</sup> See [9] for a survey of results.

optimally accurate output evaluated on the union of their databases. Although oblivious transfer is sufficient to close this gap, it is not clear whether this is a *necessary* assumption.

Indeed, there has been a separate line of work, starting with Haitner, Omri and Zorosim [21] demonstrating black-box separations between one-way functions and distributed differentially private algorithms with optimal accuracies, for two-party functionalities with long outputs. Khurana, Maji and Sahai [25] showed a black-box separation between public-key encryption and distributed differentially private algorithms with optimal accuracies for two-party boolean functionalities. These separations also extend to a range of non-optimal accuracies that are information theoretically impossible to achieve in the distributed setting. These results provide evidence that some “strong” cryptographic assumption is likely necessary for optimally accurate (or close to optimally accurate) distributed differentially private function evaluation. Despite this research, the following question has remained elusive:

*“Does there exist any class of functionalities whose distributed differentially private evaluation with optimal accuracy – necessitates the existence of oblivious transfer?”*

We prove that any protocol computing the boolean XOR functionality in a distributed differentially private manner with optimal accuracy and overwhelming probability of agreement (on the output) between both parties, implies the existence of oblivious transfer. Our result also directly lends itself to *any* boolean functionality that contains an embedded XOR on two adjacent inputs. Roughly, a function  $f$  is said to contain an embedded XOR if and only if the ideal functionality for  $f$  can be used to compute the boolean XOR functionality in the semi-honest setting. We give a formal definition of what it means for a function to contain an embedded XOR in the technical sections of the paper.

Interestingly, in the setting of secure computation, the ideal XOR functionality is known to be trivial. This is because the output of this functionality combined with the input of any individual party completely reveals the input of the other party. Thus, parties can simply send each other their inputs – and this corresponds to a secure evaluation of the XOR functionality. However, an optimally accurate distributed differentially private (noisy) protocol for XOR is not trivial, in fact we show that it implies oblivious transfer. Our proof makes use of the fact that an ideal (non-noisy) XOR is fully informative about the input of the other party.

Finally, it is interesting to observe the “philosophical” differences between differential privacy and secure computation:

- In (computationally) differentially-private protocols, “privacy comes first.” We would like to first ensure privacy of each individual input and then with this constraint, would like to compute an output which is as accurate as possible.
- In secure computation, “accuracy comes first.” We would like to release an accurate output to the function we are computing – and with this constraint, would like to ensure privacy of the inputs to the extent possible. Here, we require the transcript to leak no information about the inputs beyond what can be deduced from the output itself.

Nevertheless, as already mentioned, general secure computation immediately helps achieve the same (optimal) level of accuracy in distributed differentially-private protocols as the best achievable accuracy in the client-server setting. By relying completely on oblivious transfer for secure computation [26], our results show that the reverse is true as well (at least for the differentially private evaluation of any two-party functionality with an embedded XOR).

## 1.1 Our Contribution

Before elaborating upon our results, we briefly summarize what is known so far about accuracy gaps in the distributed differentially private computation of boolean functionalities.

Alice and Bob with inputs  $x$  and  $y$ , respectively, wish to compute  $f(x, y)$  in a differentially private manner in the distributed setting. An  $\epsilon$ -differentially private protocol for some functionality  $f$  ensures that the probability of Alice's views conditioned on  $y$  and  $y'$  are  $\lambda := e^\epsilon$  multiplicatively-close to each other, where  $y$  and  $y'$  represented as bit-strings differ only in one coordinate (i.e. they are adjacent inputs). A protocol between them is  $\alpha$ -accurate if for any  $x$  and  $y$ , the output of the protocol agrees with  $f(x, y)$ , with probability at least  $\alpha$ .

For boolean functionalities, the optimal accuracy (in the client-server model) is  $\alpha_\epsilon^* := \frac{\lambda}{\lambda+1}$ , where  $\lambda = e^\epsilon$ . Goyal et al. [19] showed that in the information theoretic setting,  $f = \text{AND}$  can only be computed  $\epsilon$ -differentially privately up to accuracy  $\alpha_\epsilon^{(\text{AND})} := \frac{\lambda(\lambda^2+\lambda+2)}{(\lambda+1)^3}$ . Similarly, for  $f = \text{XOR}$  the maximal achievable accuracy in the information theoretic setting is  $\alpha_\epsilon^{(\text{XOR})} := \frac{\lambda^2+1}{(\lambda+1)^2}$ . Note that  $\alpha_\epsilon^{(\text{XOR})} < \alpha_\epsilon^{(\text{AND})} < \alpha_\epsilon^*$ , for any finite  $\epsilon > 0$ .

We say that a function  $f$  contains an embedded XOR if there exist inputs  $x_0, x_1, y_0, y_1$  and outputs  $z_0, z_1$  such that  $f(x_a, y_b) = z_{\text{XOR}(a,b)}$  for all  $a, b \in \{0, 1\}$ . Similarly, we can define an embedded AND (equivalently, an embedded OR). By observing that any boolean function  $f$  which is sensitive to both parties' inputs either contains an embedded XOR or AND *on adjacent inputs* [3], the maximal achievable accuracy becomes

$$\alpha_\epsilon^{(f)} := \begin{cases} \alpha_\epsilon^{(\text{XOR})}, & \text{if } f \text{ contains an embedded XOR on adjacent inputs} \\ \alpha_\epsilon^{(\text{AND})}, & \text{otherwise.} \end{cases} \quad (1)$$

Given a semi-honest secure protocol for oblivious transfer, the optimal accuracy  $\alpha_\epsilon$  is achievable for any boolean  $f$ . With respect to the necessity of cryptographic assumptions, Goyal et al. [19] showed that achieving any accuracy between  $\alpha_\epsilon$  and  $\alpha_\epsilon^{(f)}$  for any function  $f$  in the distributed setting implies the existence of one-way functions. We strengthen their result to show that any two-party differentially private protocol that computes the XOR functionality in a differentially private manner with accuracy close to  $\alpha_\epsilon$  implies the existence of semi-honest secure oblivious transfer. Our result also extends to a weaker variant of differential privacy, namely *computational differential privacy* [37]. All our results hold for two-party functionalities where both parties obtain the same output with overwhelming probability. Our results can be summarized as follows (with  $k$  denoting the security parameter):

► **Informal Theorem 1.** *Semi-honest oblivious transfer reduces to any two-party  $\epsilon$  DP protocol for XOR with accuracy  $\rho (> 1/2)$  such that  $\rho \geq \alpha_\epsilon = \frac{e^\epsilon}{1+e^\epsilon}$  (the optimal accuracy).*

► **Informal Theorem 2.** *Semi-honest OT reduces to any two-party  $\epsilon_k$  computationally DP protocol for XOR with accuracy  $\rho_k \geq \alpha_{\epsilon_k} = \frac{e^{\epsilon_k}}{1+e^{\epsilon_k}}$  (the optimal accuracy).*

► **Informal Theorem 3.** *A  $(\rho_k, \frac{\lambda_k}{m_k} - 1, \frac{\lambda}{m_k} - 1)$  weak noisy channel [6, 43] reduces to any two-party  $\epsilon_k$  computationally DP protocol for XOR with (possibly non-optimal) accuracy  $\rho_k (> 1/2)$  where  $\lambda_k = e^{\epsilon_k}$  and  $m_k = \rho_k / (1 - \rho_k)$ .*

We prove the first two theorems via a reduction from (standard) noisy channels, which are known to imply semi-honest OT [5]. The first theorem is just a restriction of the second to the information-theoretic setting. The first two can also be viewed as special cases of the third. Furthermore, for a range of non-optimal accuracies we also show a reduction to weak noisy channels [6, 43]. Invoking known reductions of OT to weak binary symmetric channels, we obtain that for a small range of values of  $\epsilon_k$  and possibly non-optimal accuracies  $\rho_k$  such that  $\alpha_{\epsilon_k}^{(f)} \geq \rho_k \gg \alpha_{\epsilon_k}^*$ , there exist constants  $c_1, \left\{ c_2 < \frac{e^{c_1}}{1+e^{c_1}} \right\}$  such that for all  $\epsilon_k > c_1$  and  $\rho_k > c_2$ , any two-party  $\epsilon_k$ -private  $\rho_k$ -accurate computational DP protocol for XOR (or more generally any functionality with an embedded XOR on adjacent inputs) implies OT.

## 1.2 Related Work

The tradeoff between privacy and accuracy is quite central in designing differentially private algorithms. As mentioned before, in the client-server setting (where a single trusted server owns the entire database), the work of Dinur and Nissim [7] first showed limitations for a wide class of private algorithms. These limitations were further explored in [12, 15, 23].

The work of Dwork et al. [8, 11] proposed generic techniques for differentially private function evaluation, based on adding noise as a function of the sensitivity of database queries. The optimality of such techniques was studied in various settings in [13, 22, 41]. Variants of these techniques were shown to be optimal for certain classes of queries by [17, 20], and were shown to be non-optimal for other classes by Brenner and Nissim [2].

As mentioned before, there has also been a significant amount of work characterizing the accuracy of two-party differentially-private protocols. McGregor et al. [36] first showed that information theoretically, a large accuracy loss is inherent to the distributed differentially private computation of functionalities such as the inner product and hamming distance over  $n$ -bit inputs. This was followed by the work of Goyal et al. [19] who showed large gaps in the client-server and two-party accuracies for the differentially-private computation of boolean functionalities. Finally, the works of Haitner et al. [21] and Khurana et al. [25] showed that it is impossible to use one-way functions or even key-agreement in a black-box way to bridge any of these accuracy gaps. Our work subsumes these results for the case of XOR.

There has also been a bulk of work on the complexity of two-party finite functionalities in the information theoretic setting [26, 3, 32, 27, 28, 29, 31, 33]. Chor and Kushilevitz [4] established that all Boolean functions either reduce to SFE or can be trivially simulated. In the computationally bounded setting, Maji et al. [34] give a complete characterization of deterministic two-party finite functionalities while a series of works [35, 30, 24] give an information-theoretic characterization of (randomized, fixed-role) two-party functionalities.

Note that all constant communication protocols for Boolean functionalities can be viewed as two-party ideal finite functionalities, and therefore characterized according to [35, 30, 24]. Yet, our characterization extends to any polynomial-round *protocols* for differentially private computation with optimal accuracy, of certain classes of Boolean functionalities. This requires extra techniques to account for the entire transcript of protocol execution, which may leak information over and above the output of the ideal functionality.

## 1.3 Technical Overview

We consider the simple setting of distributed differentially private evaluation of boolean functions. Alice and Bob, with inputs  $x$  and  $y$  respectively, execute a protocol to compute a Boolean function  $f(x, y)$ . The protocol must preserve privacy (according to the differential privacy guarantee) of the input of each party. From [19], we know that any non-trivial Boolean function must embed an AND or an XOR minor on adjacent inputs. In this work, we focus on the XOR functionality; and our proof directly extends to any functionality with an embedded XOR on adjacent inputs. We also consider protocols with perfect agreement, that is, where Alice and Bob always get the same output at the end of the protocol (which is equivalent to saying that the output is part of the transcript). However, our proof also extends to protocols where parties agree on the output with overwhelming probability.<sup>2</sup>

---

<sup>2</sup> Note that if we relax the requirement of output agreement, then there is a simple information theoretically secure protocol achieving optimal accuracy.

Our main idea will be to use any protocol that implements the XOR functionality to construct an ideal noisy channel. An ideal noisy channel with flip probability  $p < 1/2$  is a functionality that takes input a bit  $X$  from the sender, samples an independent bernoulli random variable (the ‘error’)  $E$ , where  $E \sim \text{Ber}(p)$ , computes  $\tilde{X} = (X \oplus E)$  and outputs it to the receiver.

Consider an optimally accurate differentially private evaluation of the boolean XOR functionality, where both parties agree on the output with overwhelming probability (and can then publish the output). In this case, the output of the differentially private functionality can be interpreted as a “noisy” version of the correct output. In the optimally accurate setting, the probability that the output is correct is exactly  $\alpha_\epsilon = \frac{e^\epsilon}{1+e^\epsilon}$ . In other words, let  $Z$  denote the output of the protocol, then for all inputs  $X, Y$ ; the output  $Z = (X \oplus Y) \oplus E$ , where  $E$  is a bernoulli random variable  $E \sim \text{Ber}(\frac{1}{1+e^\epsilon})$ .

Our protocol to realize a noisy channel is simple: the sender (Alice) and receiver (Bob) sample independent random (private) input bits  $X \xleftarrow{\$} \{0, 1\}$  and  $Y \xleftarrow{\$} \{0, 1\}$ . They invoke the differentially private protocol for XOR with inputs  $(X, Y)$  and obtain output  $Z$ , where  $Z = (X \oplus Y) \oplus E$ , and  $E$  is the error as defined above. The sender outputs  $X$  and the receiver outputs  $Z \oplus Y (= X \oplus E)$ . It is easy to see that this protocol *correctly* implements a noisy channel with noise  $E \sim \text{Ber}(\frac{1}{1+e^\epsilon})$ . However observe that the underlying differentially private protocol for XOR may not be an ideal secure computation protocol for the noisy XOR functionality. In particular, the protocol transcript may leak extra information over the official output. Thus, it remains to prove that the above protocol implements a secure noisy channel – while only relying on the security guarantee of the differential privacy condition.

In the computational setting, the ideal noisy channel functionality can be realized by a protocol with the following security property [5]: roughly, no efficient distinguisher on the sender’s end, or on the receiver’s end respectively, should be able to distinguish the cases when the error  $E$  was 0 from when  $E$  was 1. More formally, let  $\mathcal{D}_R$  denote a distinguisher that obtains the entire view of the receiver, and  $\mathcal{D}_S$  denote a distinguisher that obtains the entire view of the sender at the end of the protocol. Then, for any non-uniform PPT distinguisher  $\mathcal{D}_R$ , the following security guarantee is required to hold:  $\Pr[\mathcal{D}_R = 1 | E = 0] - \Pr[\mathcal{D}_R = 1 | E = 1] = \text{negl}(k)$  over the randomness of the protocol. Symmetrically, at the sender’s end, for any non-uniform PPT distinguisher  $\mathcal{D}_S$ ,  $\Pr[\mathcal{D}_S = 1 | E = 0] - \Pr[\mathcal{D}_S = 1 | E = 1] = \text{negl}(k)$  over the randomness of the protocol. Here  $\text{negl}(\cdot)$  denotes some function that is asymptotically smaller than the inverse of any polynomial function, and  $k$  denotes the security parameter.

Now, the challenge is to prove that the protocol outlined above satisfies these security properties – that is, no efficient distinguisher on the sender side, or on the receiver side can distinguish the case when  $E = 0$  from when  $E = 1$ . Here, we use the following properties of the optimally accurate differentially private XOR functionality.

- Because of optimal accuracy, the protocol output is correct with probability exactly  $\frac{e^\epsilon}{1+e^\epsilon}$ .
- The (ideal, non-noisy) XOR functionality is fully informative: its output along with any of the parties’ inputs, can be used to correctly compute the input of the other party.

Since the protocol is optimally accurate, the protocol output is correct – that is,  $Z = X \oplus Y$  with probability exactly  $\alpha_\epsilon = \frac{e^\epsilon}{1+e^\epsilon}$ . Moreover, by the fully-informative property of XOR, the correct output, together with the input of any party can be used to correctly compute the other party’s input. In other words, for all  $X, Y$ , the noisy output  $Z$  of the differentially private protocol, together with the input  $Y$ , helps compute a guess for the other party’s input that is correct with probability at least  $\alpha_\epsilon$  ( $Z \oplus Y$  equals  $X$  with probability  $\alpha_\epsilon$ ).

Note that if a party could guess the other party’s input with probability any better than  $\alpha_\epsilon$ , this would directly violate differential privacy. Therefore, the output already allows

computing the best possible guess (upto differential privacy limits) for the other party's input. Informally, this means that any extra information about the error (say, leaked from the transcript) could be used to obtain a better guess of the other party's input and directly violate differential privacy. To prove security of our noisy channel, we must formalize these arguments. This is done in Section 3 and forms the core of our proof of security.

## 2 Preliminaries

**Notation.** Let  $\pi := \langle A, B \rangle$  be a two-party protocol. Let  $\text{view}_\pi^P(x_A, x_B)$  be the random variable which, in a random execution of  $\pi$  with inputs  $(x_A, x_B)$  consists of  $(x_P, R_P, \text{trans})$ , where  $R_P$  is the randomness used by party  $P$  and  $\text{trans}$  is the sequence of messages exchanged between the parties in the sampled execution. Let  $\text{out}_P$  be the function applied by party  $P$  on  $\text{view}_\pi^P(x_A, x_B)$  to obtain the output for  $P$ ,  $\text{out}_P(\text{view}_\pi^P(x_A, x_B))$ . We say that the protocol is *symmetric* if both parties receive the same output, i.e., for every  $x, y$ :  $\text{out}_A(\text{view}_\pi^A(x_A, x_B)) = \text{out}_B(\text{view}_\pi^B(x_A, x_B))$ . This is called the *official* output of the protocol, denoted by  $\text{out}_\pi(x_A, x_B)$ . For the rest of this paper, we consider only symmetric protocols, however we note that our results can be easily extended to protocols in which both parties agree on the output with overwhelming probability.

In the computational setting, we consider a family of protocols  $\{\pi_k\}_{k \in \mathbb{N}}$ , where  $k$  is the security parameter. Then, the view of party  $P \in \{A, B\}$  is denoted by  $\text{view}_\pi^P(k, x_A, x_B)$ .

### 2.1 Noisy Channels

Informally, a noisy channel takes as input a bit  $b$  and outputs bit  $b' = b \oplus e$  where *error bit*  $e \sim \text{Ber}(1 - \rho)$  is sampled independently, and  $\oplus$  is the bitwise exclusive-or operation. Roughly, the security requirement is that the error  $e$  remains “semantically secure” for both the sender and the receiver. Somewhat counterintuitively, we consider the flip probability of a  $\rho$ -noisy channel to be  $(1 - \rho)$ . This is done deliberately to match DP protocols.

A  $(\rho, \alpha, \beta)$ -weak binary symmetric channel [6, 43] is a noisy channel where the error is no longer “semantically secure”. In particular, a malicious sender or receiver obtains partial leakage on the error added by the channel, and  $(\alpha, \beta)$  denote sender and receiver leakage respectively. We defer the formal definitions to the full version.

Any protocol implementing a noisy channel is sufficient to implement the semi-honest OT functionality. A reduction between these primitives was first given by Crépeau and Kilian [5]. Furthermore, [43] showed that any protocol implementing the weak binary symmetric channel for a certain range of parameters of  $(\rho, \alpha, \beta)$ , is sufficient to implement OT. Although these reductions are information-theoretic, they also carry over to the computational setting. We use the following corollary from [43]:

► **Corollary 1.** *Let  $\rho, \alpha, \beta$  be constants, and let  $\bar{\epsilon} = \frac{(1-\rho)^2}{(1-\rho)^2 + \rho^2}$ . If at least one of the conditions  $2\alpha + \beta + \bar{\epsilon} \leq 0.12$ , or  $\beta + \bar{\epsilon} < \frac{(1-\alpha)^4}{44}$ , or  $88\alpha + 44\bar{\epsilon} < (1 - \beta)^2$ , or  $196\alpha + 98\beta + \frac{49}{2} < (1 - 2\bar{\epsilon})^2$  holds, then there exists a protocol that uses a  $(\rho, \alpha, \beta)$ -passive weak BSC and efficiently implements OT secure in the semi-honest model.*

### 2.2 Differential Privacy

We give the formal definition of a weak notion of computational differential privacy. Assuming dense sets, this is a strictly weaker definition than other (simulation-based) definitions of CDP [37]. Therefore, our reductions automatically extend to other definitions.

► **Definition 2** ( $\epsilon$ -Indistinguishable-Computational Differential Privacy, [37]). We say that an ensemble  $\{M_k\}_{k \in \mathbb{N}}$  of randomized functions  $M_k : \{0, 1\}^n \mapsto \mathcal{R}_k$  with finite range  $\mathcal{R}_k$ , provides  $\epsilon_k$ -IND-CDP if there exists a negligible function  $\text{negl} : \mathbb{N} \mapsto \mathbb{R}$  such that for every non-uniform PPT algorithm (“distinguisher”)  $D$ , every polynomial  $p(\cdot)$ , every sufficiently large  $k \in \mathbb{N}$ , every  $(x, x') \in \{0, 1\}^n \times \{0, 1\}^n$  satisfying  $|x - x'|_h = 1$ , and every advice string  $z_k$  of size at most  $p(k)$  it holds that:  $\Pr[D_k(M_k(x)) = 1] \leq e^{\epsilon_k} \times \Pr[D_k(M_k(x')) = 1] + \text{negl}(k)$ , where we write  $D_k(y)$  for  $D(1^k, z_k, y)$  and the probability is over the mechanism  $M_k$  and  $D_k$ .

► **Definition 3** (Differential Privacy over a subset of transcripts). We say that an ensemble  $\{M_k\}_{k \in \mathbb{N}}$  of randomized functions  $M_k : \{0, 1\}^n \mapsto \mathcal{R}_k$  with finite range  $\mathcal{R}_k$ , provides  $\epsilon_k$ -IND-CDP over some subset of executions  $\mathbb{S}$  if there exists a negligible function  $\text{negl} : \mathbb{N} \mapsto \mathbb{R}$  such that for every non-uniform PPT algorithm (“distinguisher”)  $D$ , every polynomial  $p(\cdot)$ , every sufficiently large  $k \in \mathbb{N}$ , every adjacent pair  $(x, x') \in \{0, 1\}^n \times \{0, 1\}^n$ , and every advice string  $z_k$  of size  $\leq p(k)$   $\Pr[D_k(M_k(x)) = 1 \wedge (M_k(x) \in \mathbb{S}_k)] \leq e^{\epsilon_k} \times \Pr[D_k(M_k(x')) = 1 \wedge (M_k(x') \in \mathbb{S}_k)] + \text{negl}(k)$  where we write  $D_k(y)$  for  $D(1^k, z_k, y)$  and the probability is taken over the randomness of mechanism  $M_k$  and distinguisher  $D_k$ .

► **Definition 4** (Two-Party Differential Privacy). Let  $\pi : \langle A, B \rangle$  be a protocol with inputs of  $A$  and  $B$  in  $\{0, 1\}^n$ . Then  $\pi$  provides  $\epsilon$ -DP if: (1) for every  $x \in \{0, 1\}^n$  the mechanism represented by the function  $\text{view}_\pi^A(x, \cdot)$  over the inputs  $y \in \{0, 1\}^n$  is  $\epsilon$ -DP, and (2) for every  $y \in \{0, 1\}^n$  the mechanism represented by  $\text{view}_\pi^B(\cdot, y)$  over the inputs  $x \in \{0, 1\}^n$  is  $\epsilon$ -DP.

In the two-party computational setting,  $\epsilon_k$ -IND-CDP is defined analogously. Formally, let  $\{\pi_k := \langle A, B \rangle(1^k)\}_{k \in \mathbb{N}}$  be an ensemble of interactive functions where the inputs of  $A$  and  $B$  are in  $\{0, 1\}^n$ . We say that  $\{\pi_k\}_{k \in \mathbb{N}}$  provides  $\epsilon_k$ -IND-CDP if: (1) for every  $x \in \{0, 1\}^n$   $\{\text{view}_\pi^A(k, x, \cdot)\}_k$  provides  $\epsilon_k$ -IND-CDP over the inputs  $y \in \{0, 1\}^n$ , and (2) for every  $y \in \{0, 1\}^n$   $\{\text{view}_\pi^B(k, \cdot, y)\}_k$  provides  $\epsilon_k$ -IND-CDP over the inputs  $x \in \{0, 1\}^n$ .

► **Definition 5** (Accuracy in Differential Privacy [19]). The *accuracy* of a randomized Boolean mechanism  $M : \{0, 1\}^n \mapsto \{0, 1\}$  with respect to a Boolean function  $f : \{0, 1\}^n \mapsto \{0, 1\}$  is defined as:  $\text{Acc}_f(M) = \min_x \{\Pr[M(x) = f(x)]\}$ , where the probability is taken over the randomness of  $M$ .

The accuracy of a *symmetric two-party protocol*  $\pi := \langle A, B \rangle$  w.r.t.  $f : \{0, 1\}^n \times \{0, 1\}^n \mapsto \{0, 1\}$  is the accuracy of the (Boolean) mechanism  $\text{out}_\pi : \{0, 1\}^n \times \{0, 1\}^n \mapsto \{0, 1\}$ ; where  $\text{out}_\pi$  returns the official output. Accuracy in the computational setting is defined analogously.

- For every Boolean mechanism  $M : \{0, 1\}^n \mapsto \{0, 1\}$  and Boolean function  $f : \{0, 1\}^n \mapsto \{0, 1\}$ , if  $M$  is  $\epsilon$ -DP then:  $\text{Acc}_f(M) \leq \frac{\lambda}{1+\lambda}$  where  $\lambda = e^\epsilon$ .<sup>3</sup> We call the bound  $\rho = \frac{\lambda}{1+\lambda}$ , the *optimal* accuracy, achieved by setting  $M(x) = f(x) \oplus e$  such that  $\Pr[e = 0] = \frac{\lambda}{1+\lambda}$ .
- If  $M$  satisfies  $\epsilon$ -IND-CDP, then  $\text{Acc}_f(M) \leq \frac{\lambda}{1+\lambda} + \text{negl}(k)$  for a negligible function  $\text{negl}(\cdot)$ .
- If a symmetric protocol ensemble  $\{\pi_k\}_{k \in \mathbb{N}}$  provides  $\epsilon$ -IND-CDP for a constant  $\epsilon > 0$ , then the accuracy of this ensemble w.r.t. the XOR function is at most  $\frac{\lambda + \text{negl}(k)}{1+\lambda} = \rho + \text{negl}'(k)$  for constant  $\epsilon$ . The accuracy  $\rho$  can be achieved using secure two-party computation [37].

### 3 Noisy Channels Reduce to Optimal Two-Party IND-CDP

► **Theorem 6.** *If there exists a two-party  $\epsilon_k$ -IND-CDP protocol with accuracy  $\rho_k (> 1/2)$  such that  $\rho_k \geq \frac{e^{\epsilon_k}}{1+e^{\epsilon_k}}$  with respect to the exclusive-or function for a constant  $\epsilon_k > 0$ , then there exists a protocol implementing the  $\rho_k$ -noisy-channel functionality.*

<sup>3</sup> Informally, if this is not the case, there exists a distinguisher such that the ratio between the probability that it guesses the input correctly versus incorrectly is greater than  $e^\epsilon$ , thereby violating  $\epsilon$ -DP.



**Proof.** Let  $\{\pi_k\}_k$  where  $\pi_k = \langle A, B \rangle(1^k)$  be an ensemble of  $\epsilon_k$ -IND-CDP protocols for computing the XOR function with accuracy  $\rho_k \geq \frac{\lambda}{1+\lambda}$  where  $\lambda = e^{\epsilon_k}$ , and  $\epsilon_k > 0$  is a constant. Note that since the protocol is  $\epsilon_k$ -IND-CDP and  $\epsilon_k > 0$ , we have that  $\rho_k \leq \frac{\lambda}{1+\lambda} + \text{negl}(k)$  for some negligible function  $\text{negl}(k)$ . For the rest of the proof, we denote  $\epsilon_k$  by  $\epsilon$ , and  $\rho_k$  by  $\rho$ .

The following protocol ensemble  $\{\pi_k := \langle S, R \rangle(1^k)\}_k$  implements a  $\rho$ -noisy-channel:  $S$  receives bit  $x$  as input, and  $R$  has no input.  $R$  samples a random bit  $y$  and the parties execute the  $\epsilon$ -IND-CDP protocol  $\langle A(x), B(y) \rangle(1^k)$  and obtain the (same) bit  $z$  as official output of this protocol.  $R$  outputs  $\tilde{x} = z \oplus y$  and  $S$  outputs  $\perp$ . The correctness of this protocol follows directly from the accuracy of the  $\epsilon$ -IND-CDP protocol. We now show that it satisfies sender-security.

**Sender security.** Assume to the contrary, that the protocol does not satisfy sender-security. That is, there exists a non-uniform PPT distinguisher  $\mathcal{D}_R$ , a fixed polynomial  $q(\cdot)$ , and infinitely many values  $k$  for which (there exists a polynomial-sized advice string  $z_k$  such that)  $\text{Adv}_\pi^R(k) \geq 1/q(k)$ . Fix one such  $k$  from now on and let:

$$p_k = \Pr[\mathcal{D}_R(z_k, \text{view}_{\pi_k}^R) = 1 | E = 0] - \Pr[\mathcal{D}_R(\text{view}_{\pi_k}^R) = 1 | E = 1]. \quad (2)$$

Note that  $\text{Adv}_\pi^R(k) = |p_k|$ . Without loss of generality, let  $p_k > 0$  for this  $k$ , and therefore by assumption  $p_k \geq 1/q(k)$ . We abuse notation and write  $\mathcal{D}_R = 1$  to denote the event that  $\mathcal{D}_R(1^k, z_k, \text{view}_{\pi_k}^R) = 1$ .<sup>4</sup> Since  $p_k \neq 0$  we must have that  $0 < \Pr[\mathcal{D}_R = 1] < 1$ .

Let  $E$  be the random variable denoting the *error bit* for the  $\epsilon$ -IND-CDP protocol. That is, for the  $\epsilon$ -IND-CDP protocol,  $E = \tilde{x} \oplus x$ . Since we are in the computational setting, the accuracy of the protocol may be different for each input, denoted by:  $\rho_{00}, \rho_{01}, \rho_{10}, \rho_{11}$ . However, they must all be within a negligible distance from each other and therefore lie within the interval  $[\rho - \text{negl}(k), \rho + \text{negl}(k)]$ , where  $\rho$  denotes  $\rho_{00}$ . Since a correct output is equivalent to  $E = 0$ , and each input is selected with equal probability,  $\Pr[E = 0]$  (which is equivalent to “average” accuracy) also lies in the same interval. We show that if  $p_k$  is noticeable then differential privacy is violated on the set of transcripts where  $\mathcal{D}_R$  outputs 1.

► **Claim 7.**  $\Pr[E = 0 \wedge \mathcal{D}_R = 1] > e^\epsilon \times \Pr[E = 1 \wedge \mathcal{D}_R = 1] + \frac{p_k}{2}$ .

**Proof.** Let  $\Pr[E = 0] = \rho^*$ , and  $\mu(k)$  be a negligible function so that  $\rho^* = \frac{\lambda}{1+\lambda} + \mu(k) > 1/2$ . Also,  $\Pr[E = 0] / \Pr[E = 1]$  is equal to  $\rho^* / (1 - \rho^*) = \lambda + \mu'(k)$  for some negligible function  $\mu'$ . Now, since  $\Pr[\mathcal{D}_R = 1] \neq 0$ , we can write (using Bayes’ rule):

$$\begin{aligned} \Pr[E = 0 \wedge \mathcal{D}_R = 1] &= \Pr[\mathcal{D}_R = 1 | E = 0] \times \Pr[E = 0] \\ &= (p_k + \Pr[\mathcal{D}_R = 1 | E = 1]) \times \Pr[E = 0] && \text{(By equation 2)} \\ &= \left( p_k + \frac{\Pr[E = 1 | \mathcal{D}_R = 1] \times \Pr[\mathcal{D}_R = 1]}{\Pr[E = 1]} \right) \times \Pr[E = 0] && \text{(Bayes' rule)} \\ &= p_k \cdot \Pr[E = 0] + \Pr[E = 1 \wedge \mathcal{D}_R = 1] \times \frac{\Pr[E = 0]}{\Pr[E = 1]} \end{aligned}$$

Note that:  $p_k \cdot \Pr[E = 0] = p_k \rho^* > p_k/2$ , and  $\frac{\Pr[E = 0]}{\Pr[E = 1]} = \frac{\rho^*}{1 - \rho^*} = \lambda + \mu'(k) > \lambda$ . Therefore,  $\Pr[E = 0 \wedge \mathcal{D}_R = 1] > \frac{p_k}{2} + \lambda \cdot \Pr[E = 1 \wedge \mathcal{D}_R = 1]$ . ◀

<sup>4</sup> Note that the input of the sender in sampling view  $\text{view}_{\pi_k}^R$  is uniformly chosen by definition of sender-security; and further, since  $k$  has been fixed, letting  $\mathcal{D}_R := \mathcal{D}_R(1^k, z_k, \text{view}_{\pi_k}^R)$  is unambiguous and well defined. Note that now,  $p_k = \Pr[\mathcal{D}_R = 1 | e = 0] - \Pr[\mathcal{D}_R = 1 | e = 1]$ .

- Obtain inputs  $M_k(x), S_k, D'_k$
- If  $S_k(M_k(x)) = 1, D_k(M_k(x)) = D'_k(M_k(x))$
- If  $S_k(M_k(x)) \neq 1, D_k(M_k(x)) = 0$

■ **Figure 1** Algorithm for  $\epsilon$ -IND-CDP Distinguisher  $D_k$ .

We say that a subset  $\mathbb{S}$  of transcripts is PPT-checkable, if there exists a probabilistic poly-time “checking” algorithm for deciding membership of a transcript in  $\mathbb{S}$ .

► **Claim 8.** *If  $\Pr[E = 0 \wedge \mathcal{D}_R = 1] > e^\epsilon \times \Pr[E = 1 \wedge \mathcal{D}_R = 1] + \frac{p_k}{2}$  is such that  $\frac{p_k}{2}$  is non-negligible over uniformly chosen sender input, then the protocol ensemble  $\{\pi_k\}_k$  does not preserve  $\epsilon$ -IND-CDP on the PPT-checkable subset of transcripts satisfying  $\mathcal{D}_R = 1$ .*

**Proof.** From  $\Pr[E = 0 \wedge \mathcal{D}_R = 1] > e^\epsilon \Pr[E = 1 \wedge \mathcal{D}_R = 1] + \frac{p_k}{2}$  it follows that  $\Pr[\tilde{x} = x \wedge \mathcal{D}_R = 1] > e^\epsilon \Pr[\tilde{x} \neq x \wedge \mathcal{D}_R = 1] + \frac{p_k}{2}$ , over the randomness of  $x$  where  $\tilde{x}$  denotes the output of the receiver. Since  $x$  is *uniformly chosen* in  $\{0, 1\}$ ,

$$\begin{aligned} & \Pr[\tilde{x} = 1 \wedge \mathcal{D}_R = 1 | x = 1] + \Pr[\tilde{x} = 0 \wedge \mathcal{D}_R = 1 | x = 0] \\ & > e^\epsilon (\Pr[\tilde{x} = 0 \wedge \mathcal{D}_R = 1 | x = 1]) + e^\epsilon (\Pr[\tilde{x} = 1 \wedge \mathcal{D}_R = 1 | x = 0]) + \frac{p_k}{2} \end{aligned}$$

Now, it is easy to observe that *either* of the following statements are true.

1.  $\Pr[\tilde{x} = 1 \wedge \mathcal{D}_R = 1 | x = 1] > e^\epsilon \times \Pr[\tilde{x} = 1 \wedge \mathcal{D}_R = 1 | x = 0] + \frac{p_k}{4}$  OR,
2.  $\Pr[\tilde{x} = 0 \wedge \mathcal{D}_R = 1 | x = 0] > e^\epsilon \times \Pr[\tilde{x} = 0 \wedge \mathcal{D}_R = 1 | x = 1] + \frac{p_k}{4}$

In either case, it is possible to claim the existence of a distinguisher. If  $p_k$  is noticeable and statement 1 holds, then there exists a distinguisher  $D_k^{1'}$  with output equal to receiver output  $\tilde{x}$ , which violates IND-CDP over the PPT checkable subset corresponding to  $\mathcal{D}_R = 1$ . On the other hand, if  $p_k$  is noticeable and statement 2 is true, then there exists a distinguisher  $D_k^{2'}$  with output equal to  $1 - \tilde{x}$ , which violates IND-CDP over the PPT checkable subset corresponding to  $\mathcal{D}_R = 1$ .

It follows from this claim that if  $p_k$  is noticeable, then the protocol ensemble  $\{\pi_k\}_k$  does not preserve  $\epsilon$ -IND-CDP on the PPT-checkable subset of transcripts on which  $\mathcal{D}_R = 1$ , because there exists distinguisher  $D'_k \in \{D_k^{1'}, D_k^{2'}\}$  and a corresponding pair of inputs  $(x^*, x^{*'}) \in (\{0, 1\} \times \{0, 1\})$  such that  $\Pr[D'_k = 1 \wedge \mathcal{D}_R = 1 | x = x^*] > e^\epsilon \times \Pr[D'_k = 0 \wedge \mathcal{D}_R = 1 | x = x^{*'}] + \frac{p_k}{4}$ . In other words, there exists a non-uniform distinguisher that violates  $\epsilon$ -IND-CDP on this subset. This proves the claim. ◀

► **Claim 9.** *A two-party protocol ensemble that provides  $\epsilon$ -IND-CDP over all executions also provides  $\epsilon$ -IND-CDP over any PPT-checkable subset of executions.*

**Proof.** Assume to the contrary that there exists a two-party  $\epsilon$ -IND-CDP protocol for which there is a non-uniform PPT distinguisher  $D'_k$  that violates  $\epsilon$ -IND-CDP over some PPT-checkable subset of executions (denoted by  $\mathbb{S}_k$ ). Let  $S_k$  denote the code of a PPT-checking algorithm that returns 1 if some execution  $M_k(x) \in \mathbb{S}_k$ , and 0 otherwise.

Then, we construct a non-uniform PPT distinguisher  $D_k$  (Figure 1) that accepts  $S_k, D'_k$  as advice  $z_k$ , and violates  $\epsilon$ -IND-CDP for the protocol.

We know that for some polynomial  $p(\cdot)$ , some sufficiently large  $k \in \mathbb{N}$ , some  $(x^*, x^{*'}) \in \{0, 1\} \times \{0, 1\}$ , some advice string  $z'_k$  of size at most  $p(k)$  and all functions  $\text{negl} : \mathbb{N} \mapsto \mathbb{R}$ , it holds that:  $\Pr[D'_k(M_k(x^*)) = 1 \wedge (M_k(x^*) \in \mathbb{S}_k)] > e^{\epsilon_k} \Pr[D'_k(M_k(x^{*'})) = 1 \wedge (M_k(x^{*'}) \in \mathbb{S}_k)] + \text{negl}(k)$

$\mathbb{S}_k$ )] +  $\text{negl}(k)$ , where the probability is taken over the randomness of mechanism  $M_k$  and distinguisher  $D'_k$ , and  $D'_k(y)$  represents  $D'(1^k, z_k, y)$ . Then, by a simple manipulation:

$$\begin{aligned} \Pr[D_k(M_k(x^*)) = 1] &= \Pr[D_k(M_k(x^*)) = 1 \wedge (M_k(x^*) \in \mathbb{S}_k)] \\ &> e^\epsilon \Pr[D'_k(M_k(x^{*'})) = 1 \wedge (M_k(x^{*'}) \in \mathbb{S}_k)] + \text{negl}(k) \\ &= e^\epsilon (\Pr[D_k(M_k(x^{*'})) = 1 \wedge (M_k(x^{*'}) \in \mathbb{S}_k)] + \Pr[D_k(M_k(x^{*'})) = 1 \wedge (M_k(x^{*'}) \notin \mathbb{S}_k)]) \\ &+ \text{negl}(k), = e^\epsilon \Pr[D_k(M_k(x^{*'})) = 1] + \text{negl}(k). \end{aligned}$$

Therefore, we have a non-uniform PPT distinguisher  $D_k$  such that for some polynomial  $p(\cdot)$ , some sufficiently large  $k \in \mathbb{N}$ , for the same  $(x^*, x^{*'}) \in (\{0, 1\} \times \{0, 1\})$ , some advice string  $z_k$  of size at most  $p(k)$  and all functions  $\text{negl} : \mathbb{N} \mapsto \mathbb{R}$  it holds that  $\Pr[D'_k(M_k(x^*)) = 1] > e^{\epsilon k} \times \Pr[D'_k(M_k(x^{*'})) = 1] + \text{negl}(k)$ . This completes the proof of this claim.  $\blacktriangleleft$

From these claims, it follows that if  $p_k$  is noticeable, then  $\{\pi_k\}_k$  does not preserve  $\epsilon$ -IND-CDP. This is a contradiction, therefore  $p_k = \text{negl}(k)$ , and the noisy channel is sender-secure.

*Receiver security.* The output  $z$  of the  $\epsilon$ -IND-CDP-protocol, obtained by both parties, is symmetric with respect to the input of each party. Moreover, since the inputs of both parties are chosen uniformly at random, the security of the receiver follows in a manner similar to sender security. This completes the proof of the theorem.  $\blacktriangleleft$

Combining this with Crépeau-Kilian's reduction [5] of OT to noisy channels, we obtain:

**Corollary 10.** *If there exists a two-party  $\epsilon_k$ -IND-CDP protocol with accuracy  $\rho_k$  such that  $\rho_k \geq \frac{e^{\epsilon k}}{1+e^{\epsilon k}}$  with respect to the exclusive-or function for a constant  $\epsilon_k > 0$ , then there exists an ensemble of protocols implementing the semi-honest oblivious-transfer functionality in the computational setting.*

## 4 Noisy Channels Reduce to Non-Optimal Two-Party IND-CDP

**Theorem 11.** *If there exists a two-party  $\epsilon_k$ -IND-CDP protocol with non-optimal accuracy  $\rho_k^1 \leq \frac{e^{\epsilon k}}{1+e^{\epsilon k}}$  with respect to the exclusive-or function for a constant  $\epsilon_k > 0$ , then there exists a protocol implementing the  $(\rho_k^1, \frac{\lambda}{m} - 1, \frac{\lambda}{m} - 1)$ -passive weak binary symmetric channel functionality where  $\rho_k^1 > 1/2$ ,  $\lambda = e^{\epsilon k}$  and  $m = \frac{\rho_k^1}{1-\rho_k^1}$ .*

The proof of this theorem follows in a similar manner as Theorem 6, and can be found in the full version of our paper. While our reduction to weak noisy channels holds for all parameters  $\epsilon > 0$  and accuracies  $\rho_k^1$ , the range of parameters for which such channels give OT is small. The following corollary follows from Theorem 11, and Corollary 1.

**Corollary 12.** *If there exists a two-party  $\epsilon_k$ -IND-CDP protocol with non-optimal accuracy  $\rho_k^1 \leq \frac{e^{\epsilon k}}{1+e^{\epsilon k}}$  with respect to the exclusive-or function for a constant  $\epsilon > 0$ , then there exist constants  $c_1, \{c_2 < \frac{e^{c_1}}{1+e^{c_1}}\}$ , such that for all  $\epsilon_k > c_1$  and  $\rho_k^1 > c_2$ , there is a protocol implementing the semi-honest oblivious transfer functionality.*

## 5 Conclusion and Open Problems

### 5.1 Extension to Functionalities with an Embedded XOR

Recall that we say that a function  $f$  contains an embedded XOR on *adjacent* inputs if there exist adjacent inputs  $x_0, x_1, y_0, y_1$  and outputs  $z_0, z_1$  such that  $f(x_1, y_b) = z_{\text{XOR}(a,b)}$  for all

$a, b \in \{0, 1\}$ . It is easy to observe that any finite functionality  $f$  with an embedded XOR, which can be computed with optimal accuracy restricted to its embedded XOR on adjacent inputs, can be used to obtain a differentially private optimally accurate XOR functionality over boolean inputs. Accuracy of XOR follows from the accuracy of the original functionality  $f$ , and privacy of XOR follows because differential privacy is a worst-case guarantee which must be maintained even when restricted to a single bit of the adjacent inputs. The resulting differentially private optimally accurate XOR protocol can then be used to obtain a secure noisy channel and therefore, perform oblivious transfer.

## 5.2 Open Problems

**Characterizing All Functionalities.** It remains an intriguing open problem to obtain a complete characterization of functionalities whose differentially private evaluation with optimal accuracy in a distributed setting, is cryptographically complete. It is interesting to obtain a complete characterization even for boolean functionalities, since the differentially private evaluation of any non-trivial functionality with optimal accuracy (such as the inner product and hamming distance functionalities considered by McGregor et al. [36]) implies the differentially private evaluation of a non-trivial boolean functionality with optimal accuracy.

Consider, the case of boolean AND. This functionality is interesting, because any non-trivial boolean functionality must contain embedded AND or XOR on adjacent inputs [3]. Therefore, for instance, showing that any (possibly polynomial round) protocol that gives a differentially private protocol for the boolean AND functionality with optimal accuracy, is cryptographically complete – would imply the completeness of an optimally accurate distributed differentially private protocol for any non-trivial boolean functionality. However, unlike XOR, the AND functionality is not completely informative about the other party’s input. In case the input of a party is 0, even a non-noisy output of the ideal AND functionality conveys absolutely no information about the input of the other party. In case the input is 1, the output allows to exactly compute the other party’s input. Therefore, if a party has input 0, the differentially-private output would be completely useless for this party, while there could be additional leakage from the transcript (allowed by differential privacy) that we do not know how to use. Such functionalities appear to have interesting connections to weak oblivious transfer, from which it is not completely known how to obtain oblivious transfer.

**Characterizing non-optimal accuracies.** From the work of McGregor et al. [36] and Goyal et al. [19] in the information theoretic setting, it is clear that for any privacy parameter  $\epsilon$ , there is a constant gap in the maximal achievable accuracies of any  $\epsilon$  differentially private protocol in the client-server and distributed settings. Goyal et al. [19] additionally showed that any hope of bridging this gap would imply the existence of one-way functions. The black box separation results of [21, 25] also hold for differentially private protocols with any accuracy in this range. Yet, it is unclear whether all protocols with accuracies in this range must imply the existence of oblivious transfer.

Our techniques for non-optimal accuracies give rise to weak noisy channels and weak versions of oblivious transfer, which for a constant range of parameters, do imply full-fledged oblivious transfer. Yet, there is a large gap between the upper and lower bounds for weak oblivious transfer amplification, and since our reductions go via noisy channels – this gap lends itself to our setting. We believe that this provides additional motivation to revive (and continue) research on the characterization of weak noisy channels.

---

**References**

---

- 1 Amos Beimel, Kobbi Nissim, and Eran Omri. Distributed private data analysis: Simultaneously solving how and what. In Wagner [42], pages 451–468. doi:10.1007/978-3-540-85174-5\_25.
- 2 Hai Brenner and Kobbi Nissim. Impossibility of differentially private universally optimal mechanisms. In focs2010 [16], pages 71–80. doi:10.1109/FOCS.2010.13.
- 3 Benny Chor and Eyal Kushilevitz. A zero-one law for boolean privacy (extended abstract). In David S. Johnson, editor, *STOC*, pages 62–72. ACM, 1989. doi:10.1145/73007.73013.
- 4 Benny Chor and Eyal Kushilevitz. A zero-one law for Boolean privacy. *SIAM J. Discrete Math.*, 4(1):36–47, 1991. doi:10.1137/0404004.
- 5 Claude Crépeau and Joe Kilian. Weakening security assumptions and oblivious transfer (abstract). In Shafi Goldwasser, editor, *CRYPTO*, volume 403 of *Lecture Notes in Computer Science*, pages 2–7. Springer, 1988. doi:10.1007/0-387-34799-2\_1.
- 6 Ivan Damgård, Joe Kilian, and Louis Salvail. On the (im)possibility of basing oblivious transfer and bit commitment on weakened security assumptions. In Jacques Stern, editor, *Advances in Cryptology – EUROCRYPT’99, International Conference on the Theory and Application of Cryptographic Techniques, Prague, Czech Republic, May 2-6, 1999, Proceedings*, volume 1592 of *Lecture Notes in Computer Science*, pages 56–73. Springer, 1999. doi:10.1007/3-540-48910-X\_5.
- 7 Irit Dinur and Kobbi Nissim. Revealing information while preserving privacy. In Frank Neven, Catriel Beeri, and Tova Milo, editors, *PODS*, pages 202–210. ACM, 2003. doi:10.1145/773153.773173.
- 8 Cynthia Dwork. Differential privacy. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *ICALP (2)*, volume 4052 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2006. doi:10.1007/11787006\_1.
- 9 Cynthia Dwork. A firm foundation for private data analysis. *Commun. ACM*, 54(1):86–95, 2011. doi:10.1145/1866739.1866758.
- 10 Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In Serge Vaudenay, editor, *EUROCRYPT*, volume 4004 of *Lecture Notes in Computer Science*, pages 486–503. Springer, 2006. doi:10.1007/11761679\_29.
- 11 Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *TCC*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2006. doi:10.1007/11681878\_14.
- 12 Cynthia Dwork, Frank McSherry, and Kunal Talwar. The price of privacy and the limits of LP decoding. In David S. Johnson and Uriel Feige, editors, *STOC*, pages 85–94. ACM, 2007. doi:10.1145/1250790.1250804.
- 13 Cynthia Dwork, Moni Naor, Omer Reingold, Guy N. Rothblum, and Salil P. Vadhan. On the complexity of differentially private data release: efficient algorithms and hardness results. In Mitzenmacher [38], pages 381–390. doi:10.1145/1536414.1536467.
- 14 Cynthia Dwork and Kobbi Nissim. Privacy-preserving datamining on vertically partitioned databases. In Matthew K. Franklin, editor, *CRYPTO*, volume 3152 of *Lecture Notes in Computer Science*, pages 528–544. Springer, 2004. doi:10.1007/978-3-540-28628-8\_32.
- 15 Cynthia Dwork and Sergey Yekhanin. New efficient attacks on statistical disclosure control mechanisms. In Wagner [42], pages 469–480. doi:10.1007/978-3-540-85174-5\_26.
- 16 *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*. IEEE Computer Society, 2010.

- 17 Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan. Universally utility-maximizing privacy mechanisms. In Mitzenmacher [38], pages 351–360. doi:10.1145/1536414.1536464.
- 18 Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or a completeness theorem for protocols with honest majority. In Alfred V. Aho, editor, *STOC*, pages 218–229. ACM, 1987. doi:10.1145/28395.28420.
- 19 Vipul Goyal, Ilya Mironov, Omkant Pandey, and Amit Sahai. Accuracy-privacy tradeoffs for two-party differentially private protocols. In Ran Canetti and Juan A. Garay, editors, *CRYPTO (1)*, volume 8042 of *Lecture Notes in Computer Science*, pages 298–315. Springer, 2013. doi:10.1007/978-3-642-40041-4\_17.
- 20 Mangesh Gupte and Mukund Sundararajan. Universally optimal privacy mechanisms for minimax agents. In Jan Paredaens and Dirk Van Gucht, editors, *PODS*, pages 135–146. ACM, 2010. doi:10.1145/1807085.1807105.
- 21 Iftach Haitner, Eran Omri, and Hila Zarosim. Limits on the usefulness of random oracles. *Theory of Cryptography Conference (TCC, to appear)*, 2013.
- 22 Moritz Hardt and Kunal Talwar. On the geometry of differential privacy. In Schulman [40], pages 705–714. doi:10.1145/1806689.1806786.
- 23 Shiva Prasad Kasiviswanathan, Mark Rudelson, Adam Smith, and Jonathan Ullman. The price of privately releasing contingency tables and the spectra of random matrices with correlated rows. In Schulman [40], pages 775–784. doi:10.1145/1806689.1806795.
- 24 Dakshita Khurana, Daniel Kraschewski, Hemanta K. Maji, Manoj Prabhakaran, and Amit Sahai. All complete functionalities are reversible, 2015.
- 25 Dakshita Khurana, Hemanta K. Maji, and Amit Sahai. Black-box separations for differentially private protocols. In Palash Sarkar and Tetsu Iwata, editors, *Advances in Cryptology – ASIACRYPT 2014 – 20th International Conference on the Theory and Application of Cryptology and Information Security, Kaoshiung, Taiwan, R.O.C., December 7-11, 2014, Proceedings, Part II*, volume 8874 of *Lecture Notes in Computer Science*, pages 386–405. Springer, 2014. doi:10.1007/978-3-662-45608-8\_21.
- 26 Joe Kilian. Founding cryptography on oblivious transfer. In Janos Simon, editor, *STOC*, pages 20–31. ACM, 1988. doi:10.1145/62212.62215.
- 27 Joe Kilian. A general completeness theorem for two-party games. In *STOC*, pages 553–560. ACM, 1991. doi:10.1145/103418.103475.
- 28 Joe Kilian. More general completeness theorems for secure two-party computation. In F. Frances Yao and Eugene M. Luks, editors, *STOC*, pages 316–324. ACM, 2000. doi:10.1145/335305.335342.
- 29 Joe Kilian, Eyal Kushilevitz, Silvio Micali, and Rafail Ostrovsky. Reducibility and completeness in private computations. *SIAM J. Comput.*, 29(4):1189–1208, 2000. doi:10.1137/S0097539797321742.
- 30 Daniel Kraschewski, Hemanta K. Maji, Manoj Prabhakaran, and Amit Sahai. A full characterization of completeness for two-party randomized function evaluation. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology – EUROCRYPT 2014 – 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, volume 8441 of *Lecture Notes in Computer Science*, pages 659–676. Springer, 2014. doi:10.1007/978-3-642-55220-5\_36.
- 31 Robin Künzler, Jörn Müller-Quade, and Dominik Raub. Secure computability of functions in the it setting with dishonest majority and applications to long-term security. In Reingold [39], pages 238–255. doi:10.1007/978-3-642-00457-5\_15.
- 32 Eyal Kushilevitz. Privacy and communication complexity. In *FOCS*, pages 416–421. IEEE, 1989. doi:10.1109/SFCS.1989.63512.

- 33 Hemanta K. Maji, Manoj Prabhakaran, and Mike Rosulek. Complexity of multi-party computation problems: The case of 2-party symmetric secure function evaluation. In Reingold [39], pages 256–273. doi:10.1007/978-3-642-00457-5\_16.
- 34 Hemanta K. Maji, Manoj Prabhakaran, and Mike Rosulek. A zero-one law for cryptographic complexity with respect to computational security. In Tal Rabin, editor, *CRYPTO*, volume 6223 of *Lecture Notes in Computer Science*, pages 595–612. Springer, 2010. doi:10.1007/978-3-642-14623-7\_32.
- 35 Hemanta K. Maji, Manoj Prabhakaran, and Mike Rosulek. A unified characterization of completeness and triviality for secure function evaluation. In Steven D. Galbraith and Mridul Nandi, editors, *INDOCRYPT*, volume 7668 of *Lecture Notes in Computer Science*, pages 40–59. Springer, 2012. doi:10.1007/978-3-642-34931-7\_4.
- 36 Andrew McGregor, Ilya Mironov, Toniann Pitassi, Omer Reingold, Kunal Talwar, and Salil P. Vadhan. The limits of two-party differential privacy. In *focs2010* [16], pages 81–90. doi:10.1109/FOCS.2010.14.
- 37 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. doi:10.1007/978-3-642-03356-8\_8.
- 38 Michael Mitzenmacher, editor. *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 – June 2, 2009*. ACM, 2009.
- 39 Omer Reingold, editor. *Theory of Cryptography, 6th Theory of Cryptography Conference, TCC 2009, San Francisco, CA, USA, March 15-17, 2009. Proceedings*, volume 5444 of *Lecture Notes in Computer Science*. Springer, 2009. doi:10.1007/978-3-642-00457-5.
- 40 Leonard J. Schulman, editor. *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*. ACM, 2010.
- 41 Jonathan Ullman and Salil P. Vadhan. PCPs and the hardness of generating private synthetic data. In Yuval Ishai, editor, *TCC*, volume 6597 of *Lecture Notes in Computer Science*, pages 400–416. Springer, 2011. doi:10.1007/978-3-642-19571-6\_24.
- 42 David Wagner, editor. *Advances in Cryptology – CRYPTO 2008, 28th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2008. Proceedings*, volume 5157 of *Lecture Notes in Computer Science*. Springer, 2008.
- 43 Jürg Wullschlegler. Oblivious transfer from weak noisy channels. In *TCC*, pages 332–349, 2009.