Differentially Private Multi-party Computation

Peter Kairouz*, Sewoong Oh†, Pramod Viswanath*
*Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign
{kairouz2, pramodv}@illinois.edu
†Department of Industrial and Enterprise Systems Engineering
University of Illinois at Urbana-Champaign
{swoh}@illinois.edu

***Abstract*—We study the problem of multi-party computation under approximate (ε, δ) differential privacy. We assume an interactive setting with k parties, each possessing a private bit. Each party wants to compute a function defined on all the parties' bits. Differential privacy ensures that there remains uncertainty in any party's bit even when given the transcript of interactions and all the other parties' bits. This paper is a follow up to our work in [1], where we studied multi-party computation under $(\varepsilon, 0)$ differential privacy. We generalize the results in [1] and prove that a simple non-interactive randomized response mechanism is optimal. Our optimality result holds for all privacy levels (all values of ε and δ), heterogenous privacy levels across parties, all types of functions to be computed, all types of cost metrics, and both average and worst-case (over the inputs) measures of accuracy.**

I. Introduction

Multi-party computation (MPC) is a general framework where multiple parties exchange information over a broadcast channel towards the goal of computing a function over their inputs while keeping those inputs private [2], [3], [4], [5]. In this paper, we study the problem of multi-party computation under differential privacy [6], [7], [8], [9]. Each party possesses a single bit of information; the information bits are statistically independent. Each party is interested in computing a function, which could differ from party to party, and there could be a central observer (observing the entire transcript of the interactive communication protocol) interested in computing a separate function. The interactive communication is achieved via a broadcast channel that all parties and central observer can hear. It is useful to distinguish between two types of communication protocols: *interactive* and *non-interactive*. We say that a communication protocol is non-interactive if a message broadcasted by one party does not depend on the messages broadcasted by any other party. In contrast, interactive protocols allow the messages at any stage of the communication to depend on all the previous messages that were communicated over the broadcast channel.

Our contributions. Our main result is the exact optimality of a simple non-interactive protocol in terms of maximizing accuracy for any given privacy levels: each party randomizes (sufficiently) its own bit and broadcasts the noisy version. Each party and the central observer then separately compute their respective decision functions to maximize the appropriate notion of their accuracy measure. The optimality is general: it holds for all types of functions, heterogeneous privacy conditions on the parties, all types of cost metrics, and both average and worst-case (over the inputs) measures of accuracy. Finally, the optimality result is *simultaneous*, in terms of maximizing accuracy at each of the parties and the central observer. Each party only needs to know its own desired level of privacy, its own function to be computed, and its measure of accuracy. Optimal data release and optimal decision making are naturally separated.

Related work. Private MPC was first addressed in [7]. The study of accuracy-privacy tradeoffs in the MPC context was first initiated by [6], which studies a paradigm where differential privacy and secure function evaluation (SFE) coexist. Specific functions, such as the SUM function, were studied under this setting, but no exact optimality results were provided. In the context of two parties, privacy-accuracy tradeoffs have been studied in [8], [9] where a single function is computed by a "third-party" observing the transcript of an interactive protocol. [9] showed that every non-trivial privacy setting incurs loss on any non-trivial boolean function. Further, focusing on the specific scenario where each one of the two parties has a single bit of information, [9] characterized the exact accuracy-privacy tradeoff for AND and XOR functions; the corresponding optimal protocol turns out to be non-interactive. However, this result was derived under some assumptions: only two parties are involved, the central observer is the only entity that computes a function, the function has to be either XOR or AND, symmetric privacy conditions are used for both parties, and accuracy is measured only as worst-case over the four possible inputs. Further, their analysis technique does not generalize to the case when there are more than two parties.

The proof of our result critically relies on an operational interpretation of differential privacy which we present in Section III. Precisely, we show that a simple non-interactive randomized response protocol dominates all (ε, δ)-differentially private multi-party protocols. This powerful technique bypasses the previous results on the same setting, where weaker results were proved using more sophisticated proof techniques. Specifically, our work generalizes the results in [1], which only addressed $(\varepsilon, 0)$-differential privacy.

II. Problem Statement

Consider the setting where there are k parties, each with its own private binary data $x_i \in \{0,1\}$ generated independently. The independence assumption here is necessary because without it each party can learn something about others, which violates differential privacy, even without revealing any information. Differential privacy implicitly imposes independence in a multi-party setting. The goal of each party $i \in [k]$ is to compute an arbitrary function $f_i : \{0,1\}^k \to \mathcal{Y}$ of interest by interactively broadcasting messages. There might be a central observer who listens to all the messages being broadcasted, and wants to compute another arbitrary function $f_0 : \{0,1\} \to \mathcal{Y}$. The k parties are honest in the sense that once they agree on what protocol to follow, every party follows the rules. At the same time, they can be curious, and each party needs to ensure that other parties cannot learn its bit with sufficient confidence. This is done by imposing local differential privacy constraints. This setting is similar to the one studied in [10], [11] in the sense that there are multiple privacy barriers, each one separating an individual party from the rest of the world. However, the main difference is that we consider multi-party computation, where there are multiple functions to be computed, and each node might possess a different function to be computed.

Let $x = [x_1, \ldots, x_k] \in \{0,1\}^k$ denote the vector of k bits, and $x_{-i} = [x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_k] \in \{0,1\}^{k-1}$ is the vector of bits except for the i^{th} bit. The parties agree on an interactive protocol P to achieve the goal of multi-party computation. A 'transcript' τ is the output of P, and is it contains the the sequence of messages exchanged between the parties. Let the probability that a transcript τ is broadcasted (via a series of interactive communications) when the data is x be denoted by $P_{x,\tau} = \mathbb{P}(\tau \,|\, x)$ for $x \in \{0,1\}^k$ and for $\tau \in \mathcal{T}$. Then, a protocol can be represented as a matrix denoting the probability distribution over a set of transcripts $\mathcal{T}$ conditioned on x: $P = [P_{x,\tau}] \in [0,1]^{2^k \times |\mathcal{T}|}$.

In the end, each party makes a decision on what the value of function f_i is, based on its own bit x_i and the transcript τ that was broadcasted. A decision rule is a mapping from a transcript $\tau \in \mathcal{T}$ and private bit $x_i \in \{0,1\}$ to a decision $y \in \mathcal{Y}$ represented by a function $\hat{f}_i(\tau, x_i)$. We allow randomized decision rules, in which case $\hat{f}_i(\tau, x_i)$ can be a random variable. For the central observer, a decision rule is a function of just the transcript, denoted by a function $\hat{f}_0(\tau)$.

We consider two notions of accuracy: the average accuracy and the worst-case accuracy. For the i^{th} party, consider an accuracy measure $w_i : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$ (or equivalently a negative cost function) such that $w_i(f_i(x), \hat{f}_i(\tau, x_i))$ measures the accuracy when the function to be computed is $f_i(x)$ and the approximation is $\hat{f}_i(\tau, x_i)$. Then the average accuracy for this i^{th} party is defined as

$$\mathrm{ACC}_{\mathrm{ave}}(P, w_i, f_i, \hat{f}_i) \equiv \frac{1}{2^k} \sum_{x \in \{0,1\}^k} \mathbb{E}_{\hat{f}_i, P_{x,\tau}}[w_i(f_i(x), \hat{f}_i(\tau, x_i))] \,, \tag{1}$$

where the expectation is taken over the random transcript τ and any randomness in the decision function $\hat{f}_i$. For example, if the accuracy measure is an indicator such that $w_i(y, y') = \mathbb{I}_{(y=y')}$, then $\mathrm{ACC}_{\mathrm{ave}}$ measures the average probability of getting the correct function output. For a given protocol P, it takes $(2^k\,|\mathcal{T}|)$ operations to compute the optimal decision rule:

$$f^*_{i,\mathrm{ave}}(\tau, x_i) = \arg\max_{y \in \mathcal{Y}} \sum_{x_{-i} \in \{0,1\}^{k-1}} P_{x,\tau}\, w_i(f_i(x), y) \,, \tag{2}$$

for each $i \in [k]$. The computational cost of $(2^k\,|\mathcal{T}|)$ for computing the optimal decision rule is *unavoidable in general*, since that is the inherent complexity of the problem: describing the distribution of the transcript requires the same cost. We will show that the optimal protocol requires a set of transcripts of size $|\mathcal{T}| = 2^k$, and the computational complexity of the decision rule for a general function is 2^{2k}. However, for a fixed protocol, this decision rule needs to be computed only once before any message is transmitted. Further, it is also possible to find a closed form solution for the decision rule when f has a simple structure. One example is the XOR function where the optimal decision rule is as simple as evaluating the XOR of all the received bits, which requires $O(k)$ operations. When there are multiple maximizers y, we can choose either one of them arbitrarily, and it follows that there is no gain in randomizing the decision rule for average accuracy.

Similarly, the worst-case accuracy is defined as

$$\mathrm{ACC}_{\mathrm{wc}}(P, w_i, f_i, \hat{f}_i) \equiv \min_{x \in \{0,1\}^k} \mathbb{E}_{\hat{f}_i, P_{x,\tau}}[w_i(f_i(x), \hat{f}_i(\tau, x_i))] \,. \tag{3}$$

For worst-case accuracy, given a protocol P, the optimal decision rule of the i^{th} party with a bit x_i can be computed by solving the following convex program:

$$\begin{aligned} Q^{(x_i)} = \quad & \\ \arg\max_{Q \,\in\, \mathbb{R}^{|\mathcal{T}| \times |\mathcal{Y}|}} \quad & \min_{x_{-i} \in \{0,1\}^{k-1}} \sum_{\tau \in \mathcal{T}} \sum_{y \in \mathcal{Y}} P_{x,\tau}\, w_i(f_i(x), y) Q_{\tau,y} \\ \text{subject to} \quad & \sum_{y \in \mathcal{Y}} Q_{\tau,y} = 1 \,,\ \forall \tau \in \mathcal{T} \text{ and } Q \geq 0 \end{aligned} \tag{4}$$

The optimal (random) decision rule $f^*_{i,\mathrm{wc}}(\tau, x_i)$ is to output y given transcript τ according to $\mathbb{P}(y|\tau, x_i) = Q^{(x_i)}_{\tau,y}$. This can be formulated as a linear program with $|\mathcal{T}| \times |\mathcal{Y}|$ variables and $2^k + |\mathcal{T}|$ constraints. Again, it is possible to find a closed form solution for the decision rule when f has a simple structure: for the XOR function, the optimal decision rule is again evaluating the XOR of all the received bits requiring $O(k)$ operations.

For a central observer, the accuracy measures are defined similarly, and the optimal decision rule is now

$$f^*_{0,\text{ave}}(\tau) \quad = \quad \arg\max_{y\in\mathcal{Y}} \sum_{x\in\{0,1\}^k} P_{x,\tau}\, w_0(f_0(x), y)\ , \quad (5)$$

and for worst-case accuracy the optimal (random) decision rule $f^*_{0,\text{wc}}(\tau)$ is to output y given transcript τ according to $\mathbb{P}(y|\tau) = Q^{(0)}_{\tau,y}$.

$$Q^{(0)} = \quad (6)$$
$$\arg\max_{Q\,\in\,\mathbb{R}^{|\mathcal{T}|\times|\mathcal{Y}|}} \quad \min_{x\in\{0,1\}^k} \sum_{\tau\in\mathcal{T}}\sum_{y\in\mathcal{Y}} P_{x,\tau}\, w_0(f_0(x), y) Q_{\tau,y}$$
$$\text{subject to} \quad \sum_{y\in\mathcal{Y}} Q_{\tau,y} = 1\ ,\ \forall \tau\in\mathcal{T} \text{ and } Q \geq 0$$

where $w_0 : \mathcal{Y}\times\mathcal{Y} \to \mathbb{R}$ is the measure of accuracy for the central observer.

III. Differentially Private Multi-party Computation

Privacy is measured by approximate differential privacy [12], [13]. Since we allow for heterogeneous privacy constraints across parties, we use $(\varepsilon_i, \delta_i)$ to denote the desired privacy level of the i^{th} party. We say that a protocol P is $(\varepsilon_i, \delta_i)$-differentially private for the i^{th} party if for $x_i, x_i' \in \{0,1\}$, $x_{-i} \in \{0,1\}^{k-1}$, and $S \subseteq \mathcal{T}$, we have that

$$\mathbb{P}(\tau \in S | x_i, x_{-i}) \quad \leq \quad e^{\varepsilon_i}\, \mathbb{P}(\tau \in S | x_i', x_{-i}) + \delta_i\ . \quad (7)$$

A mechanism P is differentially private if it is $(\varepsilon_i, \delta_i)$-differentially private for all $i \in [k]$. Differential privacy ensures that no adversary can infer the private data x_i with high enough confidence, no matter what auxiliary information or computational power she might have.

Consider the following simple protocol known as the *randomized response*, which is a term first coined by [14] and commonly used in many private communications including the multi-party setting [8]. We will show in Section IV that this is the optimal protocol that simultaneously maximizes the accuracy for all the parties. Each party broadcasts a randomized version of its bit denoted by $\tilde{x}_i$ such that

$$\tilde{x}_i = \begin{cases} 0 & \text{if } x_i = 0 \text{ with probability } \delta_i\ , \\ 1 & \text{if } x_i = 0 \text{ with probability } \dfrac{(1-\delta_i)e^{\varepsilon_i}}{1+e^{\varepsilon_i}}\ , \\ 2 & \text{if } x_i = 0 \text{ with probability } \dfrac{(1-\delta_i)}{1+e^{\varepsilon_i}}\ , \\ 3 & \text{if } x_i = 0 \text{ with probability } 0\ , \end{cases}$$

$$\tilde{x}_i = \begin{cases} 0 & \text{if } x_i = 1 \text{ with probability } 0\ , \\ 1 & \text{if } x_i = 1 \text{ with probability } \dfrac{(1-\delta_i)}{1+e^{\varepsilon_i}}\ , \\ 2 & \text{if } x_i = 1 \text{ with probability } \dfrac{(1-\delta_i)e^{\varepsilon_i}}{1+e^{\varepsilon_i}}\ , \\ 3 & \text{if } x_i = 1 \text{ with probability } \delta_i\ . \end{cases} \quad (8)$$

The proof of optimality of this randomized response depends on an operational definition of differential privacy which we now present.

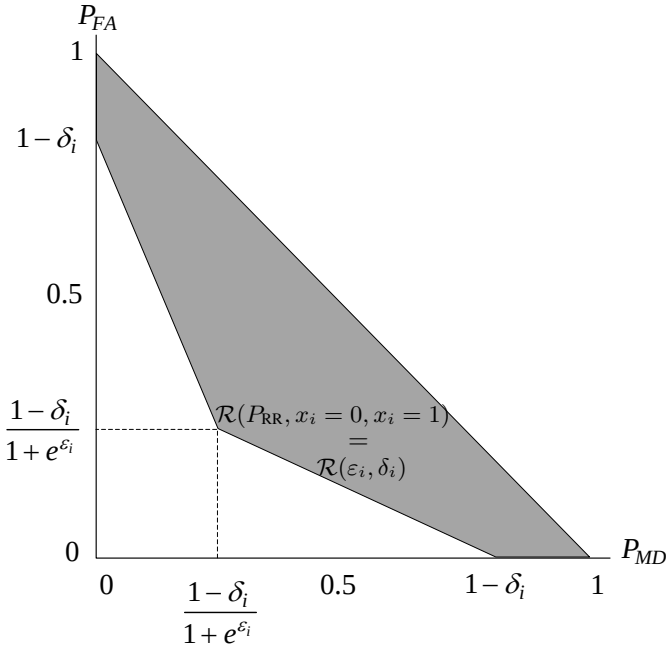

Fig. 1: Error region dictated by $(\varepsilon_i, \delta_i)$-differential privacy

Given a broadcasted transcript τ and x_{-i} (all private bits except for x_i), construct a binary hypothesis test on whether $x_i = 0$ or $x_i = 1$. A binary hypothesis test is completely characterized by a, possibly randomized, decision rule $\hat{x}_i : (\tau, x_{-i}) \to \{0,1\}$. The two types of error associated with $\hat{x}_i$ are: (1) false alarm: $\hat{x}_i = 1$ when $x_i = 0$, and (2) miss detection: $\hat{x}_i = 0$ when $x_i = 1$. The probability of false alarm is given by $P_{\text{FA}} = \mathbb{P}(\hat{x}_i = 1 | x_i = 0)$ while the probability of miss detection is given by $P_{\text{MD}} = \mathbb{P}(\hat{x}_i = 0 | x_i = 1)$. For a fixed privacy protocol P, the convex hull of all pairs $(P_{\text{MD}}, P_{\text{FA}})$ for all decision rules $\hat{x}_i$ defines a two-dimensional error region where P_{MD} is plotted against P_{FA}. For example, the randomized response mechanism P_{RR} given in (8) has an error region $\mathcal{R}(P_{\text{RR}}, x_i = 0, x_i = 1)$ shown in Figure 1.

The differential privacy constraints in Equation (7) impose the following conditions on the error regions of all $(\varepsilon_i, \delta_i)$-differentially private protocols

$$P_{\text{FA}} + e^{\varepsilon_i} P_{\text{MD}} \quad \geq \quad 1 - \delta_i,$$
$$e^{\varepsilon_i} P_{\text{FA}} + P_{\text{MD}} \quad \geq \quad 1 - \delta_i,$$

for any decision rule $\hat{x}_i$ and any $i \in [k]$. The above two conditions define an error region $\mathcal{R}(\varepsilon_i, \delta_i)$ shown in Figure 1. Interestingly, the next theorem shows that the converse result is also true.

Lemma 1: A mechanism P is differentially private if and only if $\mathcal{R}(P, x_i = 0, x_i = 1) \subseteq \mathcal{R}(\varepsilon_i, \delta_i)$ for all $i \in [k]$.

The proof of the above lemma can be found in [15] (see Corollary 2.3 on page 4). Notice that it is no coincidence

that $\mathcal{R}(P_{\text{RR}}, x_i = 0, x_i = 1) = \mathcal{R}(\varepsilon_i, \delta_i)$ (see Figure 1). This property will be essential in proving the optimality of the randomized response.

Lemma 1 allows us to benefit from the data processing inequality (DPI) and its converse, which follows from a celebrated result by [16]. These inequalities, while simple by themselves, lead to surprisingly strong technical results. Indeed, there is a long line of such a tradition in the information theory literature (see Chapter 17 of [17]).

Recall that τ contains the sequence of messages broadcasted by all k parties. Let $\tau(i)$ represent the messages broadcasted by the i^{th} party and observe that $\tau = \{\tau(1), \cdots, \tau(k)\}$. Consider two privatization protocols, P_1 and P_2, and let τ_1 and τ_2 denote the output transcripts under protocols P_1 and P_2, respectively. We say that P_1 dominates P_2 if there exists a sequence of stochastic transformations $\{W_1, \cdots, W_k\}$ such that for all $i \in [k]$, given x_{-i}, τ_2 can be simulated by applying W_i to $\tau_1(i)$ and x_{-i}. In other words, given x_{-i}, $W_i(\tau_1(i), x_{-i})$ has the same distribution as τ_2 .

Lemma 2: A multi-party privacy protocol P_1 dominates a protocol P_2 if and only if $\mathcal{R}(P_2, x_i = 0, x_i = 1) \subseteq \mathcal{R}(P_1, x_i = 0, x_i = 1)$ for all $i \in [k]$.

The proof of the above lemma can be found in [16]. Lemma 2 will be critical in proving the optimality of the randomized response.

Corollary 3.1: Any differentially private protocol P is dominated by the randomized response P_{RR} given in Equation (8). Therefore, there exists a sequence of stochastic transformations $\{W_1, \cdots, W_k\}$ such that $W_i(\tilde{x}_i, x_{-i})$ has the same distribution as τ for all $i \in [k]$.

Corollary 3.1 follows from Lemma 1, Lemma 2, and the fact that$\mathcal{R}(\varepsilon_i, \delta_i) = \mathcal{R}(P_{\text{RR}}, x_i = 0, x_i = 1)$ for all $i \in [k]$.

IV. Main Result

We show, perhaps surprisingly, that the simple randomized response presented in (8) is the unique optimal protocol in a very general sense.

Theorem 4.1: Let the optimal decision rule be defined as in (2) for the average accuracy and (5) for the worst-case accuracy. Then, for any privacy levels $(\varepsilon_i, \delta_i)$, any function $f_i : \{0,1\}^k \to \mathcal{Y}$, and any accuracy measure $w_i : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$ for $i \in [k]$, together with the optimal decision rule, the randomized response achieves the maximum accuracy for the i^{th} party among all differentially private interactive and non-interactive protocols. For the central observer, the randomized response with the optimal decision rule defined in (5) and (7) achieves the maximum accuracy among all $\{(\varepsilon_i, \delta_i)\}$-differentially private interactive protocols and all decision rules for any arbitrary function f_0 and any measure of accuracy w_0.

This is a strong optimality result. Every party and the central observer can simultaneously achieve the optimal accuracy, using a universal randomized response. Each party only needs to know its own desired level of privacy, its own function to be computed, and its measure of accuracy. Optimal data release and optimal decision making are naturally separated. It is not immediate at all that such a simple non-interactive randomized response mechanism would achieve the maximum accuracy. The proof critically harnesses the data processing inequalities and is provided in Appendix A.

V. Conclusion

In this paper, we studied the problem of differentially private multi-party computation. We showed that a simple non-interactive randomized response is optimal for all privacy levels (all values of ε and δ), heterogenous privacy levels across parties, all types of functions to be computed, all types of cost metrics, and both average and worst-case (over the inputs) measures of accuracy. Though our results are general, they only handle settings where each party possesses a single bit. In the more general scenario where parities can have multiple bits, interaction might be critical to achieving the optimal privacy-utility tradeoffs.

Appendix
Proof of Theorem 4.1

To prove Theorem 4.1, it is sufficient to prove Theorem A.1 stating that any other protocol can be simulated from the randomized response outputs. Let $\{x_i\}_{i\in[k]}$ and $\tau_{\text{RR}} = \{\tilde{x}_i\}_{i\in[k]}$ denote the k private bits and transcript under the randomized response P_{RR} (Equation (8)), respectively. We will prove that any differentially private multi-party protocol can be simulated from τ_{RR}. This proves the desired theorem, since the optimal protocol and the optimal decision rules can be simulated by each node (and the central observer) upon receiving the randomized responses. Hence, proving that randomized response is sufficient to achieve optimal performance (on any metric).

Theorem A.1: For any protocol P that generates a random transcript τ, there exists a stochastic transformation T such that the joint distribution of the bits and the transcript can be simulated from the randomized outputs:

$$(x_1, \ldots, x_k, \tau) \quad \stackrel{D}{=} \quad (x_1, \ldots, x_k, T(\tilde{x}_1, \ldots, \tilde{x}_k)) \,, \quad (9)$$

where $\stackrel{D}{=}$ denotes equality in distribution, and $\tilde{x}_i$ is a randomized response of x_i.

To prove the above theorem, our strategy is to apply an induction argument over a class of stochastic transformations $\{T_1, T_2, \cdots, T_k\}$, where T_ℓ operates on $\tilde{x}_1^\ell = (\tilde{x}_1, \ldots, \tilde{x}_\ell)$ and

$x_{\ell+1}^k = (x_{\ell+1}, \ldots, x_k)$. We will prove the following series of equations:

$$\begin{aligned}(x_1, \ldots, x_k, \tau) &\stackrel{D}{=} (x_1, \ldots, x_k, T_1(\tilde{x}_1, x_2^k)) &(10)\\ &\stackrel{D}{=} (x_1, \ldots, x_k, T_2(\tilde{x}_1^2, x_3^k)) &(11)\\ &\vdots\\ &\stackrel{D}{=} (x_1, \ldots, x_k, T_k(\tilde{x}_1^k)) \,, &(12)\end{aligned}$$

We first prove Equation (10). To do so, we show an equivalent version of this equation, which is $(x_1, \tau) \stackrel{D}{=} (x_1, T(\tilde{x}_1, x_2^k))$ for all fixed values of x_2^k. Equation (10) follows by applying Bayes rule to this equation. First, note that for all fixed x_2^k,

$$\mathcal{R}\big(P, x_1 = 0, x_1 = 1\big) \quad \subseteq \quad \mathcal{R}(\varepsilon_1, \delta_1) \,, \tag{13}$$

by the fact that τ is $(\varepsilon_1, \delta_1)$-differentially private and Lemma 1. Next, notice that by construction, the randomized response achieves this outer bound, i.e.

$$\mathcal{R}\big(P_{\text{RR}}, x_1 = 0, x_1 = 1\big) \quad = \quad \mathcal{R}(\varepsilon_1, \delta_1) \,, \tag{14}$$

for all values of x_2^k which holds only under the current assumption that x_1^k are independent. Hence from the reverse data processing inequality in Corollary 3.1, it follows that for each instance of x_2^k, there exists a stochastic transformation such that τ is simulated from $\tilde{x}_1$, i.e. $(x_1, \tau) \stackrel{D}{=} (x_1, T(\tilde{x}_1, x_2^k))$. This proves the desired Equation (10).

We now prove an inductive step that allows us to recursively show Equations (11) and (12). We want to prove that there always exists a stochastic transformation $T_{\ell+1}$ such that

$$(x_1^k, T_\ell(\tilde{x}_1^\ell, x_{\ell+1}^k)) \quad \stackrel{D}{=} \quad (x_1^k, T_{\ell+1}(\tilde{x}_1^{\ell+1}, x_{\ell+2}^k)) \,, \tag{15}$$

for any stochastic transformation T_ℓ satisfying $(\varepsilon_{\ell+1}, \delta_{\ell+1})$-differential privacy. Again, we prove that $(x_{\ell+1}, T_\ell(\tilde{x}_1^\ell, x_{\ell+1}^k)) \stackrel{D}{=} (x_{\ell+1}, T_{\ell+1}(\tilde{x}_1^{\ell+1}, x_{\ell+2}^k))$ for all values of $(x_1^\ell, \tilde{x}_1^\ell, x_{\ell+1}^k)$. Then, Equation (15) follows from Bayes rule. First note that from the assumption that $T_\ell(\tilde{x}_1^\ell, x_{\ell+1}^k)$ is $(\varepsilon_{\ell+1}, \delta_{\ell+1})$-differentially private with respect to $x_{\ell+1}$, we know that for any fixed values of $(x_1^\ell, \tilde{x}_1^\ell, x_{\ell+2}^k)$, binary hypothesis testing on $x_{\ell+1}$ based on the observation $T_\ell(\tilde{x}_1^\ell, x_{\ell+1}^k)$ must obey the differential privacy constraint:

$$\begin{aligned}&\mathbb{P}(T_\ell(\tilde{x}_1^\ell, x_{\ell+1}^k) \in S | x_{\ell+1}, x_1^\ell, \tilde{x}_1^\ell, x_{\ell+2}^k) \leq\\ &e^{\varepsilon_{\ell+1}} \mathbb{P}(T_\ell(\tilde{x}_1^\ell, x_{\ell+1}^k) \in S | \overline{x_{\ell+1}}, x_1^\ell, \tilde{x}_1^\ell, x_{\ell+2}^k) + \delta_{\ell+1} \,,\end{aligned} \tag{16}$$

and since $T_\ell(\tilde{x}_1^\ell, x_{\ell+1}^k)$ is conditionally independent of x_1^ℓ given $\tilde{x}_1^\ell$, we get

$$\begin{aligned}&\mathbb{P}(T_\ell(\tilde{x}_1^\ell, x_{\ell+1}^k) \in S | x_{\ell+1}, \tilde{x}_1^\ell, x_{\ell+2}^k) \leq\\ &e^{\varepsilon_{\ell+1}} \mathbb{P}(T_\ell(\tilde{x}_1^\ell, x_{\ell+1}^k) \in S | \overline{x_{\ell+1}}, \tilde{x}_1^\ell, x_{\ell+2}^k) + \delta_{\ell+1} \,.\end{aligned} \tag{17}$$

This implies that for each value of $(\tilde{x}_1^\ell, x_{\ell+2}^k)$,

$$\mathcal{R}\big(T_\ell, x_{\ell+1} = 0, x_{\ell+1} = 1\big) \quad \subseteq \quad \mathcal{R}(\varepsilon_{\ell+1}, \delta_{\ell+1}) \,.$$

Next, notice that by construction, the randomized response achieves this outer bound, i.e.

$$\mathcal{R}\big(P_{\text{RR}}, x_{\ell+1} = 0, x_{\ell+1} = 1\big) \quad = \quad \mathcal{R}(\varepsilon_{\ell+1}, \delta_{\ell+1}) \,, \tag{18}$$

for all values of $(\tilde{x}_1^\ell, x_{\ell+2}^k)$ which holds only under the current assumption that x_1^k are independent. Hence from the reverse data processing inequality in Corollary 3.1, it follows that for each instance of $(\tilde{x}_1^\ell, x_{\ell+2}^k)$, there exists a stochastic transformation such that T_ℓ is simulated from $\tilde{x}_{\ell+1}$, i.e. $(x_{\ell+1}, T_\ell(\tilde{x}_1^\ell, x_{\ell+1}^k)) \stackrel{D}{=} (x_{\ell+1}, T_{\ell+1}(\tilde{x}_{\ell+1}, \tilde{x}_1^\ell, x_{\ell+2}^k))$. This proves the desired induction step in Equation (15). Consequently, by induction Equation (12) holds, and this proves Theorem A.1.

REFERENCES

[1] P. Kairouz, S. Oh, and P. Viswanath, "Secure multi-party differential privacy," in *Advances in Neural Information Processing Systems*, 2015.

[2] A. C. Yao, "Protocols for secure computations," in *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. IEEE, 1982, pp. 160–164.

[3] M. Ben-Or, S. Goldwasser, and A. Wigderson, "Completeness theorems for non-cryptographic fault-tolerant distributed computation," in *Proceedings of the twentieth annual ACM symposium on Theory of computing*. ACM, 1988, pp. 1–10.

[4] O. Goldreich, S. Micali, and A. Wigderson, "How to play any mental game," in *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing*, ser. STOC '87. New York, NY, USA: ACM, 1987, pp. 218–229. [Online]. Available: http://doi.acm.org/10.1145/28395.28420

[5] D. Chaum, C. Crépeau, and I. Damgard, "Multiparty unconditionally secure protocols," in *Proceedings of the twentieth annual ACM symposium on Theory of computing*. ACM, 1988, pp. 11–19.

[6] A. Beimel, K. Nissim, and E. Omri, "Distributed private data analysis: Simultaneously solving how and what," in *Advances in Cryptology–CRYPTO 2008*. Springer, 2008, pp. 451–468.

[7] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, "Our data, ourselves: Privacy via distributed noise generation," in *Advances in Cryptology-EUROCRYPT 2006*. Springer, 2006, pp. 486–503.

[8] A. McGregor, I. Mironov, T. Pitassi, O. Reingold, K. Talwar, and S. Vadhan, "The limits of two-party differential privacy," in *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*. IEEE, 2010, pp. 81–90.

[9] V. Goyal, I. Mironov, O. Pandey, and A. Sahai, "Accuracy-privacy tradeoffs for two-party differentially private protocols," in *Advances in Cryptology–CRYPTO 2013*. Springer, 2013, pp. 298–315.

[10] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, "Local privacy and statistical minimax rates," in *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*. IEEE, 2013, pp. 429–438.

[11] P. Kairouz, S. Oh, and P. Viswanath, "Extremal mechanisms for local differential privacy," in *Advances in Neural Information Processing Systems 27*, 2014, pp. 2879–2887.

[12] C. Dwork, "Differential privacy," in *Automata, languages and programming*. Springer, 2006, pp. 1–12.

[13] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of Cryptography*. Springer, 2006, pp. 265–284.

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

[15] P. Kairouz, S. Oh, and P. Viswanath, "The composition theorem for differential privacy," *arXiv preprint arXiv:1311.0776*, 2013.

[16] D. Blackwell, "Equivalent comparisons of experiments," *The Annals of Mathematical Statistics*, vol. 24, no. 2, pp. 265–272, 1953.

[17] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 2012.