Private Interactive Communication
Across an Adversarial Channel

Ran Gelles, Amit Sahai, and Akshay Wadia

Department of Computer Science, University of California, Los Angeles, CA, USA
{gelles, sahai, awadia}@cs.ucla.edu

Abstract. Consider two parties Alice and Bob, who hold private inputs x and y, and wish to compute a function $f(x, y)$ *privately* in the information theoretic sense; that is, each party should learn nothing beyond $f(x, y)$. However, the communication channel available to them is *noisy*. This means that the channel can introduce errors in the transmission between the two parties. Moreover, the channel is *adversarial* in the sense that it knows the protocol that Alice and Bob are running, and maliciously introduces errors to disrupt the communication, subject to some bound on the total number of errors. A fundamental question in this setting is to design a protocol that remains private in the presence of large number of errors.

If Alice and Bob are only interested in computing $f(x, y)$ correctly, and not privately, then quite robust protocols are known that can tolerate a *constant* fraction of errors. However, none of these solutions is applicable in the setting of privacy, as they inherently leak information about the parties' inputs. This leads to the question whether we can simultaneously achieve privacy and error-resilience against a constant fraction of errors.

We show that privacy and error-resilience are contradictory goals. In particular, we show that for every constant $c > 0$, there exists a function f which is privately computable in the error-less setting, but for which no private and correct protocol is resilient against a c-fraction of errors.

Keywords: Interactive communication, coding, adversarial noise, private function evaluation, information-theoretic security.

1 Introduction

Consider two parties, A and B that wish to compute some function $f(x, y)$ of their respective private inputs x, y. The channel connecting the parties might be prone to error, and in order to compute f the parties run some error-resilient interactive protocol that is guaranteed to output $f(x, y)$ as long as the amount of error is small (e.g., the fraction of corrupt messages is below some threshold $c > 0$).

If efficiency is not concerned, party A can simply encode its input using a standard (Shannon) error-correcting code that corrects a fraction $c < 1/2$ of errors [Sha48]; party B computes the output $f(x, y)$ and sends it back (encoded with the same parameters) to A. This allows a global error rate of up to $1/4$. However, x might be very large, and more efficient protocols (communication-wise) can be found. In 1993, Schulman [Sch93, Sch96] showed a way to compile any interactive protocol π that assumes no error, into an error-resilient protocol Π that withstands a fraction of errors of up to $1/240$, whose overhead is linear (i.e., the communication of Π, is linear in the communication of π). This was followed by a flow of other works [RS94, BR11, GMS11, FGOS12, BK12, BN13] trying to improve the efficiency or error-resilience in various settings. Notably, the recent work of Braverman and Rao [BR11] showed how to compile a (noiseless) π into an error-resilient Π that correctly computes f as long as the error fraction is less than $1/4$. Their construction, similar to the construction of Schulman, has a linear overhead (such constructions are also called *constant rate*). The recent elegant work of Brakerski and Kalai [BK12] shows how to achieve *efficient* computation in the presence of adversarial noise of rates less than $1/32$.

All the above works consider only the correctness of the protocol in the presence of adversarial error. In this work we aim at achieving other properties as well, specifically, privacy. We ask the following question:

> *Can error-resilient protocols tolerating some constant fraction of errors be devised, if we also require the protocol to be* ***private****, that is, not to leak any information besides $f(x, y)$?*

Evidently, privacy is no longer guaranteed when compiling a private (error-free) protocol π for f using the methods of Schulman or Braverman and Rao. For instance, Schulman's scheme works by 'running' π until the users suspect their views are 'inconsistent' due to channel errors. Then, the users 'backtrack' to the last consistent point and continue from there. This "rewinding" is vital for the protocol's ability to recover from errors, yet it is fatal for its privacy.

We answer the above question in the negative, and show that privacy and error-resilience are in fact contradicting aspirations.

Theorem 1.1 (Separation of error-free and adversarial error for private protocols). *For any constant $c > 0$, there exists a function f such that a private protocol for f exists in the error-free setting, but no protocol is both private and correctly computes f over a noisy channel with at most a c-fraction of adversarial errors.*

In fact, our impossibility result is even stronger, as it rules out protocols even when the error-rate is below constant:

Theorem 1.2 (main, informal). *For any $d = O(2^n)$, there is a function f taking two inputs of size n, such that a private protocol for f exists in the error-free setting, but such that no private protocol correctly computes f over a noisy channel with at most $O(1/d)$-fraction of adversarial errors.*

With our result, it is very interesting to compare private computations in the *error-free* model, the *random-noise* model, and the *adversarial-error* model. As first shown by Kushilevitz [Kus89, Kus92] and Beaver [Bea91], in the error-free model, some functions are privately-computable and others are not. The set of functions that can be computed in a privately is fully characterized in [Bea91, Kus92], wherein an optimal protocol for any function in this set is also provided. In the random-noise model, a recent work by Ishai, Kushilevitz, Ostrovsky, Prabhakaran, Sahai, and Wullschleger [IKO+11] shows how to obtain *oblivious transfer* (OT) over the *binary symmetric channel*[1] (BSC) with constant communication. Since any function f can be privately computed assuming OT [GMW87] (in the semi-honest model), the result of [IKO+11] implies a constant rate protocol for *any* function f. In the adversarial-error model, on the other hand, we show that depending on the error rate, there are functions that can be privately computed in the error-free setting but not in the adversarial-error model. Thus, there is no hope to achieve constant-rate schemes for any function f, or even for any f which can be computed privately in the error-free setting.

[1] The BSC channel is parametrized by a probability $p \in [0, 1]$. For any input bit $b \in \{0, 1\}$, the output is $1 - b$ with probability p, and b with probability $1 - p$. That is, the channel flips the input bit with probability p.

Our Techniques. Our key insight is that resilience to errors implies that the protocol must be able to "backtrack" its course from an incorrect track reached due to channel errors, while, at least intuitively, privacy should prevent the protocol from taking any course other than the correct one and prevent "rewinding" the protocol and changing its intermediates inputs. Hence, showing that privacy forces the protocol to advance in a specific order, say, through specific ordered steps, would imply that error-correction can be performed only *within* a single step, but not *between steps*.

The line of work initiated by Kushilevitz [Kus89, Kus92] and Beaver [Bea91] helps us with this problem. In particular, the result of Maji, Prabhakaran and Rosulek [MPR09], shows that private protocols in the error-free model must advance in very specific ways. The progress of such a protocol can be split into steps where at each step only a single quantum of information is revealed (namely, a single decomposition of the domain, see formal definitions in Section 2). Moreover, revealing the information of step i before the information of step $j < i$ is revealed leads to violating privacy. The amount of steps depends only on f, and for any $1 \leq d \leq 2(2^n - 1)$ there exists a privately-computable function with exactly d steps [Kus92].

We revisit the works of [Bea91, Kus92, MPR09] and examine them in the adversarial error model. We show that a similar property exists in any private, error-resilient protocol. Next, we show that if a channel completely changes the messages of two consecutive steps, the privacy is compromised. This gives a lower bound of $\Omega(1/d)$ on the proportion of allowed error, and proves our main theorem.

More specifically, We identify the point in the protocol where A reveals that his input is in some set P rather than in Q (where $P \cup Q$ is some subset of his domain). The adversarial channel changes the messages of this "step" to lead B to believe that A's input is actually in Q. Then, at the next "step" it's B's turn to reveal whether his input is in some partition V, W of (a subset of) his domain. The channel keeps changing A's messages until the end of this step. However, this violates the privacy of B since his partition is done assuming A's input is in Q. Conditioned on the fact it is in P, the next step of B should have been decomposing his sub-domain into $\tilde{V}, \tilde{W} \neq V, W$. Hence, some information was leaked, and the protocol is not private.

Related Work. Concurrently and independently of our work, a very recent paper of Chung, Pass and Telang [CPT13] also examines impossibility results for coding schemes for secure interactive communication, however their work considers a model where a *party* can be adversarial together with the (rate-limited) adversarial channel, and the adversary can potentially also be computationally bounded. In contrast, our impossibility results are for a model where parties are *semi-honest*, and only the channel can be adversarial (and rate-limited). Thus, the impossibility results of [CPT13] and ours are largely incomparable; however, taken together, our paper and [CPT13] provide impossibility results for private interactive communication in both the cases of semi-honest parties and malicious parties.

Roadmap. We start with preliminaries and definitions in Section 2. Of special importance is Section 2.2, where we model the adversarial channel, and give definitions of privacy and correctness in the presence of such a channel. In Section 3 we study the structure of transcripts generated by a perfectly private and correct protocol in the presence of adversarial channel. The main result of this section is Theorem 3.11 which proves that a private protocol must advance in a very specific sequence of steps. This fact is used in Section 4 to construct a channel that introduces errors at well-chosen points in the protocol execution and violates privacy. Section 5 discusses extensions to the case of protocols which are not perfectly secure.

2 Interactive Computation With and Without Noise

2.1 Preliminaries: two-party computation in the noiseless model

We begin by recalling some definitions of two-party computation, assuming noiseless communication. Two parties, A and B hold inputs $x \in X$ and $y \in Y$ respectively. The parties wish to compute the value of $f : X \times Y \to Z$ for some finite set Z. The computation is done via an interactive probabilistic protocol $\pi = (\pi_A, \pi_B)$. The protocol works in rounds and defines for each party, at each round, the next message to be transmitted as a function of its input, randomness, and received messages.[2] The parties send massages $m \in \Sigma$ according to π in an alternating[3] way: for odd rounds, $m_i = \pi_A(x, R_A, m_2 m_4 \cdots m_{i-1})$ and for even rounds $m_i = \pi_B(y, R_B, m_1 m_3 \cdots m_{i-1})$. We denote the transcript of a specific instance $t = m_1 m_2 \cdots$ as the messages transmitted within that instance. We assume messages are delimited so it is possible to parse the transcript into specific messages.

At the end of an execution of π, the parties compute their outputs as a function of their respective inputs, randomness, and views. These functions will be denoted by $\mathsf{out}^{\pi}_{A}(\cdot,\cdot,\cdot)$ and $\mathsf{out}^{\pi}_{B}(\cdot,\cdot,\cdot)$ respectively. That is, for inputs $(x, y) \in X \times Y$, transcript t, and randomness R_A, R_B, party A's output will be $\mathsf{out}^{\pi}_{A}(x, R_A, t)$, and B's output is $\mathsf{out}^{\pi}_{B}(y, R_B, t)$. We will usually omit the randomness, and implicitly treat π (along with the output function) as a single randomized protocol.

Correctness and Privacy. We recall the standard definitions for correctness and privacy, assuming the computation is over an error-free channel.

Definition 2.1. *A protocol is said to be (perfectly) correct if*

$$\forall x, y : \Pr_{t \leftarrow \pi(x,y)} [\mathsf{out}^{\pi}_{A}(x, t) = \mathsf{out}^{\pi}_{B}(y, t) = f(x, y)] = 1.$$

As for privacy, we follow the notions developed by Chor, Kushilevitz and Beaver [Kus89, CK91, Bea91, Kus92].

Definition 2.2 (Privacy for Noiseless Protocols). *A protocol is private (with respect to party A) if, for any $x, x' \in X$ and $y \in Y$ such that $f(x, y) = f(x', y)$, the protocol generates the same distribution of transcripts for inputs (x, y) and (x', y). Formally, for any $t \in \Sigma^*$, $\Pr[t \mid x, y] = \Pr[t \mid x', y]$. Similarly, privacy of party B is given if for any $x \in X$ and $y, y' \in Y$ such that $f(x, y) = f(x, y')$, for any $t \in \Sigma^*$, $\Pr[t \mid x, y] = \Pr[t \mid x, y']$.*

Kushilevtiz [Kus92] and Beaver [Bea91] gave a full characterization for functions that can be computed privately (in the noiseless model). We now recall this characterization.

We say that f is *column-partitionable* if there exists a non-trivial partition of Y into disjoint $P, Q \subset Y$ such that for any $x \in X$, for any $y \in P, y' \in Q$, it holds that $f(x, y) \neq f(x, y')$. We refer to such P, Q as *valid column-partitions.* We let $f_{X \times P}$ and $f_{X \times Q}$ be f restricted to the domains $X \times P$ and $X \times Q$. Similarly, we say f is *row-partitionable*, if we can partition X into Q and P such that for any $y \in Y$, and any $x \in P$, $x' \in Q$, it holds that $f(x, y) \neq f(x', y)$. Domains that were achieved by a sequence of valid partitionings are called *restricted domains.* We say that some subdomain $X' \times Y'$ is a restricted domain of depth d if it takes d recursive decompositions to obtain $X' \times Y'$. Furthermore, we can generalize the notion of f being row-partitionable, if X can

[2] Note that previously sent messages are function of the above and need not be given explicitly to the protocol.

[3] This assumption can be avoided with changing our results by a factor of at most 2.

be partitioned into $m > 1$ disjoint sets $X_1, X_2, \ldots, X_m$ and for any $u \in X_i$ and any $v \in X_j$ where $j \neq i$, it holds that for every $y \in Y$, $f(u,y) \neq f(v,y)$. A partition of X into m^* disjoint sets is called *maximal* if any other partition of X has $m \leq m^*$. We will mostly consider *binary decomposable* functions—functions which are maximally partitionable into at most $m = 2$ partitions.

We say that f is *partitionable* if: (1) f is constant; or, (2) f is either row-partitionable or column-partitionable, and each of the restricted functions is also partitionable.

The following lemma which appears in [Bea91, Kus92], identifies the set of privately-computable functions exactly as the set of *partitionable* functions.

Lemma 2.3 ([Bea91, Kus92]). *A function f has a correct protocol π (in the noiseless model) which is private with respect to both parties if and only if f is partitionable.*

2.2 The noisy model: two-party computation over a noisy channel

Let us now augment the above model into a noisy version. First, let us define the notion of a channel. We assume an arbitrary, causal channel (possibly with memory). The channel's instantiation at the ℓ-th round $\mathsf{Ch}_\ell(m_\ell \mid (m_1, m'_1)(m_2, m'_2) \cdots (m_{\ell-1} m'_{\ell-1}))$ is a distribution on Σ characterized by the (input,output) channel-messages up to that round. Hence, the probability that the channel takes $\boldsymbol{m} = m_1 \cdots m_n$ and outputs $\boldsymbol{m'} = m'_1 \cdots m'_n$ is given by

$$\Pr[\mathsf{Ch}(\boldsymbol{m}) = (\boldsymbol{m'})] \coloneqq \Pr\left[\mathsf{Ch}_1(m_1 \mid \emptyset) = m'_1\right] \times \Pr\left[\mathsf{Ch}_2\left(m_2 \mid (m_1, m'_1)\right) = m'_2\right] \cdots \\ \times \Pr\left[\mathsf{Ch}_n\left(m_n \mid (m_1, m'_1)(m_2, m'_2) \cdots (m_{n-1} m'_{n-1})\right) = m'_n\right] \quad (1)$$

Again we assume that the parties alternately send messages. For a specific instance of the protocol, we denote the *transcript* $\boldsymbol{t} = (t^s, t^r) = (m_1 m_2 \cdots, m'_1 m'_2 \cdots)$ as the messages observed for that instance (where t^s denotes the messages the parties send, and t^r the massages they receive). We denote by $t^s_A = m_1 m_3 m_5 \cdots$ the messages sent by A and similarly $t^s_B = m_2 m_4 m_6 \cdots$ is the messages sent by B. The messages received by each party, t^r_A, t^r_B are defined in a similar way.

A protocol $\pi = (\pi_A, \pi_B)$ over a noisy channel Ch is defined in a similar way to the noiseless case: the protocol defines for each party, at each round the next message to send as a function of its inputs, randomness and all the messages *received* up to that round. Explicitly, A's messages on odd rounds are $m_{2i-1} = \pi_A(x, R_A, m'_2 m'_4 \cdots m'_{2i-2})$, and B's messages on even rounds are $m_{2i} = \pi_B(y, R_B, m'_1 m'_3 \cdots m'_{2i-1})$, with $\boldsymbol{m'} = \mathsf{Ch}(\boldsymbol{m})$.

For some protocol π over a channel Ch and some input (x,y) we can ask what is the set of possible transcripts generated by π for that input, and what is the probability to observe each such transcript. Formally, let $T^s(x,y,t^r)$ be the random variable describing $t^s = m_1 m_2 \cdots$ defined by π on inputs (x,y) given that $t^r = m'_1 m'_2 \cdots$ is received by the parties. Then, the probability that an instance of π sending messages over a channel Ch on inputs (x,y) produces $\mathbf{t} = (t^s, t^r)$ is given by

$$\Pr[\mathbf{t} \mid x, y, \mathsf{Ch}] \coloneqq P_{T^s(x,y,t^r)}(t^s) \cdot \Pr[\mathsf{Ch}(t^s) = t^r]. \quad (2)$$

The probability can be generalized to any prefix of $\mathbf{t}$ in the standard way, quantifying the probability that π over Ch produces that prefix on (x,y). The probabilities to have a specific *sent-messages* transcript t^s, and a specific *received-messages* transcript t^r are given by the marginal probabilities,

$$\Pr[t^s \mid x, y, \mathsf{Ch}] = \sum_{t^r} P_{T^s(x,y,t^r)}(t^s) \cdot \Pr\left[\mathsf{Ch}(t^s) = t^r\right], \quad \text{and} \tag{3}$$

$$\Pr[t^r \mid x, y, \mathsf{Ch}] = \sum_{t^s} P_{T^s(x,y,t^r)}(t^s) \cdot \Pr\left[\mathsf{Ch}(t^s) = t^r\right]. \tag{4}$$

For a given input $(x, y) \in X \times Y$ the execution of π with inputs (x, y) over channel Ch induces a distribution over the set of all transcripts according to the above equations. Let $\mathbf{u}$ be a fixed partial transcript. For a given string $\mathbf{u}'$ such that $\mathbf{u} \circ \mathbf{u}'$ is a complete transcript, we can compute the probability that the protocol π (with inputs (x, y) and channel Ch) will produce $\mathbf{u}'$ *when executed with the history* $\mathbf{u}'$. We will use $\pi^{\mathsf{Ch}}(x, y; \mathbf{u})$ to refer to this distribution over completions $\mathbf{u}'$. We will use $\tau^{\mathsf{Ch}}(x, y) \coloneqq \left\{\mathsf{t} \leftarrow \pi^{\mathsf{Ch}}(x, y)\right\}$ to denote the set of all possible (complete) transcripts generated by π on inputs (x, y) assuming messages are being sent over the channel Ch. The output function for A is defined as above, $\mathsf{out}^{\pi}_{A}(x, R_A, t^r_A)$ where t^r_A are the even indices of t^r, i.e., the messages received by A. Similarly, the output for B is $\mathsf{out}^{\pi}_{B}(y, R_B, t^r_B)$.

Error Rate. For $i = 1, 2, \ldots$, by $t^r[i]$, we denote the i-th message in t^r, and similarly for t^s, and let $\mathbf{t}[i] = (t^s[i], t^r[i])$. We define $t^r[0]$ and $t^s[0]$ to be the empty string. By $|t^r|$ (resp. $|t^s|$), we denote the number of messages in transcript t^r (resp. t^s). We denote $|\boldsymbol{t}|$ as the number of pairs in $\boldsymbol{t}$ so that $|\boldsymbol{t}| = |t^r| = |t^s|$. If no noise is added by the channel, then for all $i = 1, 2, \ldots$, we have $t^r[i] = t^s[i]$. Define the *error weight* $\eta(\boldsymbol{t})$ of a transcript $\boldsymbol{t} = (t^r, t^s)$ as

$$\eta(\mathsf{t}) \coloneqq |\{\, i \mid t^r[i] \neq t^s[i] \,\}|\,. \tag{5}$$

Definition 2.4 (Error Rate of Ch). *We say that a channel* Ch *has error rate* $\mu \in [0, 1]$*, if for any inputs* (x, y) *and any transcript* $\boldsymbol{t} \in \tau^{\mathsf{Ch}}(x, y)$*, it holds that*

$$\eta(\mathsf{t}) \leq \mu|\boldsymbol{t}|.$$

Correctness and Privacy. We now augment the privacy and correctness definitions to the noisy case. The correctness definition is straightforward.

Definition 2.5 (Perfect Correctness). *A protocol* π *is* perfectly *correct with respect to some function* f *over* Ch *if for any inputs* x, y*,*

$$\Pr_{\substack{R_A, R_B, R_{\mathsf{Ch}} \\ (t^s, t^r) \leftarrow \pi^{\mathsf{Ch}}(x,y)}}\left[\mathsf{out}^{\pi}_{A}(x, t^r_A) = \mathsf{out}^{\pi}_{B}(y, t^r_B) = f(x, y)\right] = 1.$$

We further say a protocol π *is* perfectly correct with rate μ *if* π *is perfectly correct over every channel* Ch *with error rate* μ*.*

We define privacy in a similar way to the noiseless case, where the transcript doesn't leak more than what is trivially conveyed by the function's output.

Definition 2.6 (Perfect Privacy). *A protocol* π *for* f *is perfectly private over* Ch *with respect to* A *if for any inputs such that* $f(x, y) = f(x', y)$ *it holds that for any* t*,*

$$\Pr[\mathsf{t} \mid x, y, \mathsf{Ch}] = \Pr[\mathsf{t} \mid x', y, \mathsf{Ch}].$$

Similarly, perfect privacy with respect to B requires that whenever $f(x,y) = f(x,y')$ it also holds that $\Pr[\mathbf{t} \mid x, y, \mathsf{Ch}] = \Pr[\mathbf{t} \mid x, y', \mathsf{Ch}]$.
Furthermore, We say a protocol π is perfectly private with rate μ *if it is perfectly private (for both players) over every channel* Ch *with rate μ.*

We mention the following difference between our setting and the one used in [Bea91, Kus92, MPR09]. One of the assumptions used by these works is that the last message of π is the output $f(x,y)$ (this doesn't restrict the generality of those results). However in our noisy model, this assumption doesn't make much sense as the channel can always "violate" the correctness of the protocol by changing this single message. To avoid this inconvenience we assume the output is given by the functions $\mathsf{out}_A, \mathsf{out}_B$ and is not part of the transcript. The next lemma shows that this change has no effect on the distribution of transcripts, namely, that different outputs imply different transcripts, regardless of the parties' randomness (which is trivially the case when the output is part of the transcript).

Lemma 2.7. *Let π be a perfectly correct and perfectly private protocol for a function f, and let x, y, y' be inputs such that $f(x,y) \neq f(x,y')$. Then, for any channel* Ch *and for any transcript* $\mathbf{t}$*, if* $\Pr[\mathbf{t} \mid x, y, \mathsf{Ch}] > 0$*, then* $\Pr[\mathbf{t} \mid x, y', \mathsf{Ch}] = 0$.

Proof. We say A's random tape R_A *is consistent with* transcript $\mathbf{t}$ and input x if for $1 \leq i \leq |\mathbf{t}|$,

$$\pi_A(x, R_A, t^r[0 \dots i{-}1]) = t^s[i].$$

Let $\mathbf{t}^*$ be a transcript such that both $\Pr[\mathbf{t}^* \mid x, y, \mathsf{Ch}] > 0$ and $\Pr[\mathbf{t}^* \mid x, y', \mathsf{Ch}] > 0$. Consider the set,

$$S_{\mathbf{t}^*} = \{\, R_A \mid R_A \textit{ is consistent with } \mathbf{t}^* \textit{ and } x \,\}$$

As $\Pr[\mathbf{t}^* \mid x, y, \mathsf{Ch}] > 0$, by perfect correctness, for all $R_A \in S_{\mathbf{t}^*}$, we have $\mathsf{out}_A(x, R_A, \mathbf{t}^*) = f(x,y)$. Similarly, as $\Pr[\mathbf{t}^* \mid x, y', \mathsf{Ch}] > 0$, for all $R_A \in S_{\mathbf{t}^*}$, we have that $\mathsf{out}_A(x, R_A, \mathbf{t}^*) = f(x,y')$. But these statements are mutually contradictory as $f(x,y) \neq f(x,y')$. Therefore, it must be that $S_{\mathbf{t}^*} = \emptyset$, and either $\Pr[\mathbf{t}^* \mid x, y, \mathsf{Ch}] = 0$ or $\Pr[\mathbf{t}^* \mid x, y', \mathsf{Ch}] = 0$. □

3 Properties of Private Protocols: The Perfect Case

The goal of this section is to show that a private protocol reveals information about the parties' inputs in a particular well-defined order. This section extends the results in [MPR09] to the noisy channel setting. In particular, we analyze the transcripts of the protocol and show that they proceed through well-defined stages that depend upon the partitioning of the function. At each such stage, one of the parties necessarily reveals a particular partition of its input space in which its current input lies. This is formalized by the concept of "frontiers"(see Definition 3.7 for the formal definition). A frontier is a prefix-free set of partial transcripts which is maximal in the sense that every transcript has a prefix in the frontier. The main result of this section is Theorem 3.11, where we construct a set of frontiers which correspond to the stages where the transcripts reveal the partitions.

We focus on functions that have a certain structure, which we call *binary-uniquely decomposable*: their domain is maximally row-partitionable into 2 partitions and is not column-partitionable. Each one of the two restricted domains is either constant or (maximally, binary) column-partitionable but not row-partitionable. This continues recursively in an alternating manner (i.e., one level row-partitionable and the next level is column-partitionable). Although the proofs could be extended to a more general case, this simpler class of functions is enough for our impossibility result, as we are only required to provide *one* function that has no perfectly-correct and private protocol over a noisy channel (of certain parameters). A simple example of such a function is the following Vickrey auction function [Vic61] (a variant of the min function), which is very useful for performing private 2^{nd}-price (Vickrey) auctions,

$$f(x,y) = \begin{cases} 2x & x \leq y \\ 2y+1 & x > y \end{cases}$$

where the inputs x and y are integers less than a certain bound N.

Our first step is to identify certain inputs that form a *minor*. These will be associated with each level of f's decomposition.

Definition 3.1 (Minor). *We say that x, x', y, y' form a ⊟-minor (resp., a ◫-minor) if*

$$\begin{array}{ccc} f(x,y) & = & f(x,y') \\ \neq & & \neq \\ f(x',y) & \neq & f(x',y') \end{array} \qquad \textit{(resp., if} \qquad \begin{array}{ccc} f(x,y) & \neq & f(x,y') \\ = & & \neq \\ f(x',y) & \neq & f(x',y') \end{array} \quad \textit{.)}$$

Lemma 3.2. *Assume a binary-uniquely decomposable function f on $X \times Y$ which is row-partitionable into P, Q, but not column-partitionable. Then, exactly one of the following happens.*

1. *Both $f_{Q\times Y}$ and $f_{P\times Y}$ are constant.*
2. *Only one of $f_{Q\times Y}, f_{P\times Y}$ is constant, and there exists a ⊟-minor (x, x', y, y') with x and x' in different sets of the partition (specifically x is in the constant partition). Furthermore, y, y' are in different sets of the column-partitioning of the non-constant restricted function.*
3. *Both $f_{Q\times Y}$ and $f_{P\times Y}$ are non constants, and there are at least two "overlapping" minors: if $f_{Q\times Y}$ is partitioned into Y_{Q_0}, Y_{Q_1} and $f_{P\times Y}$ is partitioned into Y_{P_0}, Y_{P_1}, there exist minors (x, x', y, y') and (x, x', w, w') where $x \in P$ and $x' \in Q$. And if both y, y' are in the same partition Y_{P_i} and both w, w' are in Y_{Q_j} then $Y_{P_i} \cap Y_{Q_j}$ contains one of w, w' and one of y, y'.*

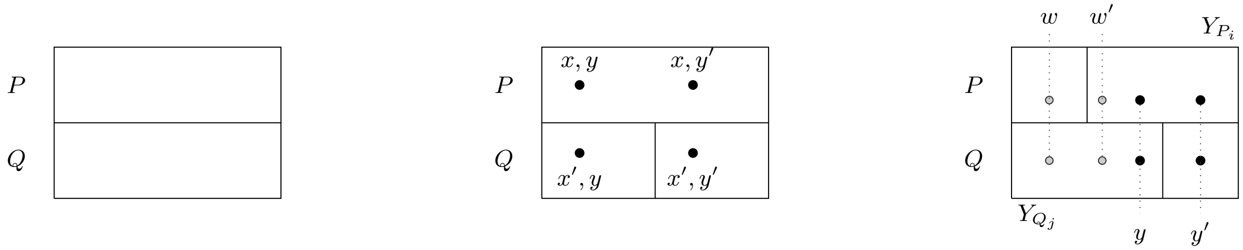

Fig. 1. A simplified illustration of the three cases; A closed area means constant value of f within this area

Proof. We note that cases 1 and 2 are simple, and prove case 3. First, we mention that Y_{P_0}, Y_{P_1} is *not* a valid partition for $f_{Q\times Y}$, since if it was, then f would be column-partitionable in contradiction to its definition. The same applies to Y_{Q_0}, Y_{Q_1} and $f_{P\times Y}$.

Consider $f_{P\times Y}$. There must exist $x \in P$ and $y \in Y_{Q_0}, y' \in Y_{Q_1}$ such that $f(x,y) = f(x,y')$, or otherwise Y_{Q_0}, Y_{Q_1} is a valid column partition of $f_{P\times Y}$. It follows that y, y' must be in the same partition Y_{P_i}. Note that for any $x' \in Q$, (x, x', y, y') is a minor.

Let Y_{Q_j} be the partition that has non-empty intersection with both Y_{P_0}, Y_{P_1} (at least one of Y_{Q_0}, Y_{Q_1} must intersect both Y_{P_0}, Y_{P_1}, or otherwise f is column-partitionable). There must exist $u' \in Q$ and $w \in Y_{Q_j} \cap Y_{P_0}$, $w' \in Y_{Q_j} \cap Y_{P_1}$ such that $f(u', w) = f(u', w')$, or otherwise Y_{P_0}, Y_{P_1} is a valid partition for $f_{Q\times Y}$. Then, for any $u \in P$ (and specifically for x), (u', u, w, w') is a ⊞-minor. Finally, note that we are allowed to choose any $x' \in Q, u \in P$; we choose $x' = u'$ and $u = x$ and complete the proof. □

With this setting we can start exploring the sequential revelation of information in private protocols. Consider inputs (x, x', y, y') which form a ⊞-minor. The following lemma shows that transcripts produced by inputs (x, y) and those produced by inputs (x', y) must differ at some point. Further, the earliest point where they differ must be a point where party A speaks. This round will symbolize the point where A begins to reveal the row-decomposition of this level.

Lemma 3.3. *For a perfectly correct π for f over Ch, the following holds. If $f(x,y) \neq f(x',y)$, then for any $\mathbf{t} \in \tau^{\mathsf{Ch}}(x,y) \cup \tau^{\mathsf{Ch}}(x',y)$ there exists an odd round ρ such that $\Pr[\mathbf{t}_{[1..\rho]} \mid x, y, \mathsf{Ch}] \neq \Pr[\mathbf{t}_{[1..\rho]} \mid x', y, \mathsf{Ch}]$, yet for any $\ell < \rho$ it holds that $\Pr[\mathbf{t}_{[1..\ell]} \mid x, y, \mathsf{Ch}] = \Pr[\mathbf{t}_{[1..\ell]} \mid x', y, \mathsf{Ch}]$.*

Proof. Let $\mathbf{t} = (t^s, t^r) \in \tau^{\mathsf{Ch}}(x,y) \cup \tau^{\mathsf{Ch}}(x',y)$. Since $f(x,y) \neq f(x',y)$ and the protocol is perfectly-correct, Lemma 2.7 suggests that there must exist some round k for which $\Pr[\mathbf{t}_{[1..k]} \mid x, y, \mathsf{Ch}] \neq \Pr[\mathbf{t}_{[1..k]} \mid x', y, \mathsf{Ch}]$. Let ρ be the minimal round that satisfies this condition.

Assume towards contradiction that ρ is even, i.e., the next one to send a message is party B. Then,

$$\begin{aligned}
&\Pr\left[\mathbf{t}_{[1..\rho]} \mid x, y, \mathsf{Ch}\right] = \\
&\qquad = \Pr\left[\pi(x, y, t^r{}_{[1..\rho]}) = t^s{}_{[1..\rho]}\right] \Pr\left[\mathsf{Ch}(t^s{}_{[1..\rho]}) = t^r{}_{[1..\rho]}\right] \\
&\qquad = \Pr\left[\pi_B(y, t^r{}_{[1..\rho-1]}) = t^s[\rho]\right] \Pr\left[\pi(x, y, t^r{}_{[1..\rho-1]}) = t^s{}_{[1..\rho-1]}\right] \\
&\qquad\quad \times \Pr\left[\mathsf{Ch}(t^s[\rho] \mid \mathbf{t}_{[1..\rho-1]}) = t^r[\rho]\right] \Pr\left[\mathsf{Ch}(t^s{}_{[1..\rho-1]}) = t^r{}_{[1..\rho-1]}\right] \\
&\qquad = \Pr\left[\pi_B(y, t^r{}_{[1..\rho-1]}) = t^s[\rho]\right] \Pr\left[(\mathsf{Ch}(t^s[\rho] \mid \mathbf{t}_{[1..\rho-1]}) = t^r[\rho]\right] \times \Pr\left[\mathbf{t}_{[1..\rho-1]} \mid x', y, \mathsf{Ch}\right] \\
&\qquad = \Pr[\mathbf{t}_{[1..\rho]} \mid x', y, \mathsf{Ch}]
\end{aligned}$$

which leads to a contradiction. The third transition is due the fact that for any round $k < \rho$, $\Pr[\mathbf{t}_{[1..k]} \mid x, y, \mathsf{Ch}] = \Pr[\mathbf{t}_{[1..k]} \mid x', y, \mathsf{Ch}]$, by the way we define ρ. □

The next two lemmas show that when considering two consecutive levels of the decomposition (say, row-decomposition followed by a column-decomposition; those are associated with some ⊞-minor (x, x', y, y')), if party B reveals some information about whether its input is y or y' (i.e. the column-decomposition), it must be the case, that party A has completely revealed whether his input is x or x', or more precisely, whether his input is in P or Q for the matching row-partition.

Lemma 3.4. *For a perfectly correct π for f over Ch, the following holds. For any x, x', y, y' that form a ⊞-minor, let $\mathbf{t} \in \tau^{\mathsf{Ch}}(x', y) \cup \tau^{\mathsf{Ch}}(x', y')$, and let ρ be the first round where $\Pr[\mathbf{t}_{[1..\rho]} \mid x', y] \neq \Pr[\mathbf{t}_{[1..\rho]} \mid x', y']$. Then,*

$$\Pr\left[\mathbf{t}_{[1..\rho-1]} \mid x, y, \mathsf{Ch}\right] = \Pr\left[\mathbf{t}_{[1..\rho-1]} \mid x, y', \mathsf{Ch}\right] = 0.$$

Proof. First, note that ρ must be even (this follows from Lemma 3.3). Let $\mathbf{t} \in \tau^{\mathsf{Ch}}(x', y) \cup \tau^{\mathsf{Ch}}(x', y')$, and assume (towards contradiction) that $\Pr[\mathbf{t}_{[1..\rho-1]} \mid x, y', \mathsf{Ch}] > 0$. Now,

$$\begin{aligned}\Pr[\mathbf{t}_{[1..\rho]} \mid x, y, \mathsf{Ch}] &= \Pr\left[\pi_B\left(y, t^r_B{}_{[1..\rho-1]}\right) = t^s[\rho]\right] \Pr\left[\mathsf{Ch}\left(t^s[\rho] \mid \mathbf{t}_{[1..\rho-1]}\right) = t^r[\rho]\right] \cdot \Pr[\mathbf{t}_{[1..\rho-1]} \mid x, y, \mathsf{Ch}] \\ &= \Pr\left[\mathbf{t}[\rho] \mid y, \mathbf{t}_{[1..\rho-1]}, \mathsf{Ch}\right] \Pr[\mathbf{t}_{[1..\rho-1]} \mid x, y', \mathsf{Ch}].\end{aligned}$$

In addition, due to B's privacy and the fact that $f(x, y) = f(x, y')$, for any round ρ' and specifically for the above ρ we have,

$$\begin{aligned}\Pr[\mathbf{t}_{[1..\rho]} \mid x, y, \mathsf{Ch}] &= \Pr[\mathbf{t}_{[1..\rho]} \mid x, y', \mathsf{Ch}] \\ &= \Pr\left[\mathbf{t}[\rho] \mid y', \mathbf{t}_{[1..\rho-1]}, \mathsf{Ch}\right] \Pr[\mathbf{t}_{[1..\rho-1]} \mid x, y', \mathsf{Ch}],\end{aligned}$$

and since $\Pr[\mathbf{t}_{[1..\rho-1]} \mid x, y'] \neq 0$, it must be that

$$\Pr\left[\mathbf{t}[\rho] \mid y, \mathbf{t}_{[1..\rho-1]}, \mathsf{Ch}\right] = \Pr\left[\mathbf{t}[\rho] \mid y', \mathbf{t}_{[1..\rho-1]}, \mathsf{Ch}\right]. \tag{6}$$

On the other hand, we know that ρ is the first round such that $\Pr\left[\mathbf{t}_{[1..\rho]} \mid x', y\right] \neq \Pr\left[\mathbf{t}_{[1..\rho]} \mid x', y'\right]$, thus,

$$\begin{aligned}\Pr[\mathbf{t}_{[1..\rho]} \mid x', y, \mathsf{Ch}] =&\ \Pr\left[\mathbf{t}[\rho] \mid y, \mathbf{t}_{[1..\rho-1]}, \mathsf{Ch}\right] \Pr[\mathbf{t}_{[1..\rho-1]} \mid x', y, \mathsf{Ch}] \\ =&\ \Pr\left[\mathbf{t}[\rho] \mid y, \mathbf{t}_{[1..\rho-1]}, \mathsf{Ch}\right] \Pr[\mathbf{t}_{[1..\rho-1]} \mid x', y', \mathsf{Ch}] \\ &\neq \\ &\Pr\left[\mathbf{t}[\rho] \mid y', \mathbf{t}_{[1..\rho-1]}, \mathsf{Ch}\right] \Pr[\mathbf{t}_{[1..\rho-1]} \mid x', y', \mathsf{Ch}] = \Pr[\mathbf{t}_{[1..\rho]} \mid x', y', \mathsf{Ch}]\end{aligned}$$

and since $\mathbf{t} \in \tau^{\mathsf{Ch}}(x', y) \cup \tau^{\mathsf{Ch}}(x', y')$ we know that $\Pr[\mathbf{t}_{[1..\rho-1]} \mid x', y', \mathsf{Ch}] > 0$ and we conclude that

$$\Pr\left[\mathbf{t}[\rho] \mid y, \mathbf{t}_{[1..\rho-1]}, \mathsf{Ch}\right] \neq \Pr\left[\mathbf{t}[\rho] \mid y', \mathbf{t}_{[1..\rho-1]}, \mathsf{Ch}\right], \tag{7}$$

contradicting Eq. (6). A similar proof works for (x, y). □

Lemma 3.5. *Let π be a perfectly-correct private protocol for f over Ch. For any x, x', y, y' that form a ⊟-minor, let $\mathbf{t} \in \tau^{\mathsf{Ch}}(x, y) \cup \tau^{\mathsf{Ch}}(x, y')$. Then, there exists a round ρ such that*

$$\Pr[\mathbf{t}_{[1..\rho-1]} \mid x', y, \mathsf{Ch}] = 0, \quad \Pr[\mathbf{t}_{[1..\rho-1]} \mid x', y', \mathsf{Ch}] = 0$$

Moreover, for any $\mathbf{t}^* \in \tau^{\mathsf{Ch}}(x', y) \cup \tau^{\mathsf{Ch}}(x', y')$ such that $\mathbf{t}_{[1..\rho']} = \mathbf{t}^*_{[1..\rho']}$ for some ρ', it holds that $\rho' < \rho$.

Proof. The existence of ρ is trivially guaranteed by (perfect-)correctness. Clearly $\rho > \rho'$ since at least one of $\Pr[\mathbf{t}_{[1..\rho']} \mid x', y, \mathsf{Ch}]$, $\Pr[\mathbf{t}_{[1..\rho']} \mid x', y', \mathsf{Ch}]$ must be non zero, by the way we pick $\mathbf{t}^*$ and the fact that $\mathbf{t}_{[1..\rho']} = \mathbf{t}^*_{[1..\rho']}$. □

We conclude with the following Theorem, which is a simple corollary of the above lemmas, and formalizes the fact that there is a round in which we fully know the row-partition but nothing about the next-level column-decomposition.

Theorem 3.6. *Let π be perfectly-correct private protocol for f over Ch. Let x, x', y, y' form a ⊟-minor. Then, for any transcript $\mathbf{t}$, there is a round ρ such that at least one of the following is satisfied.*

1. $\Pr[\mathbf{t}_{[1..\rho]} \mid x, y, \mathsf{Ch}] = \Pr[\mathbf{t}_{[1..\rho]} \mid x, y', \mathsf{Ch}] = 0$, *and* $\Pr[\mathbf{t}_{[1..\rho]} \mid x', y, \mathsf{Ch}] = \Pr[\mathbf{t}_{[1..\rho]} \mid x', y', \mathsf{Ch}]$.
2. $\Pr[\mathbf{t}_{[1..\rho]} \mid x', y, \mathsf{Ch}] = \Pr[\mathbf{t}_{[1..\rho]} \mid x', y', \mathsf{Ch}] = 0$, *and* $\Pr[\mathbf{t}_{[1..\rho]} \mid x, y, \mathsf{Ch}] = \Pr[\mathbf{t}_{[1..\rho]} \mid x, y', \mathsf{Ch}]$.

It is important to note at this point that analogous claims for the above Lemma 3.2 through Theorem 3.6 can similarly be shown for a binary-uniquely decomposable f which is column-partitionable but not row-partitionable, and for (x, x', y, y') being a ◫-minor instead of a ⊟-minor, etc.

3.1 Frontiers: dealing with more than a single transcript

While above we have analyzed one transcript at a time, we now extend the definitions to treat many transcripts (e.g., the entire $\tau^{\mathsf{Ch}}(x,y)$) at the same time. In the spirit of [MPR09] we define *frontiers* as a means to deal with several transcripts rather than a single transcript. We begin with a few definitions.

Definition 3.7 (Frontiers). *A set F of partial transcripts is called a* frontier *if,*

1. *F is prefix-free*[4]*.*
2. *F is* maximal*, i.e.,* $\Pr[F \mid x, y, \mathsf{Ch}] = 1$ *for all inputs x, y, where*

$$\Pr[F \mid x, y, \mathsf{Ch}] \coloneqq \sum_{\mathbf{u} \in F} \Pr[\mathbf{u} \mid x, y, \mathsf{Ch}].$$

Informally, we say that a partial transcript $\mathbf{u}$ has reached frontier F if $\mathbf{u} \in F$. For frontiers F and G, we will also be interested in the probability that F 'precedes' G; that is, the probability that the protocol reaches frontier F before it reaches frontier G. Formally, we have the following definition.

Definition 3.8. *For any frontiers F and G, define*

$$\Pr[F \le G \mid x, y, \mathsf{Ch}] = \sum_{\{\mathbf{u} \in F \mid \mathbf{u} \text{ is a prefix of some } \mathbf{u}' \in G\}} \Pr[\mathbf{u} \mid x, y, \mathsf{Ch}]$$

the weighted probability of all the transcripts in F that are prefixes of transcripts in G. In the same manner define $\Pr[F < G \mid x, y, \mathsf{Ch}]$ *for transcripts which are strict prefixes.*

For notational conciseness, for fixed x, y, Ch, we will use the statement "$F \le G$" as a shorthand for the statement "$\Pr[F \le G \mid x, y, \mathsf{Ch}] = 1$". The same holds for the statement "$F < G$".

Given a minor (x, x', y, y') we define the set $F(x, x', y, y')$ as the collection of transcripts prefixes up to the point where the "partition" related with this minor happens via Theorem 3.6. Note that the length of each prefix can be different.

Definition 3.9. *For a minor (x, x', y, y') define*

$$F(x, x', y, y') \coloneqq \left\{ \mathbf{t}_{[1..\rho]} \mid \mathbf{t} \textit{ is a transcript and } \rho = \rho(\mathbf{t}) \textit{ is the minimal satisfying Theorem 3.6.} \right\}$$

Lemma 3.10. *Given a minor (x, x', y, y'), the set $F(x, x', y, y')$ is a frontier.*

Proof. First, we show that $F = F(x, x', y, y')$ is prefix free. Assume that $\mathbf{u}, \mathbf{v} \in F$ such that $\mathbf{u}$ is a strict prefix of $\mathbf{v}$. Let $\mathbf{t}$ be any transcript with prefix $\mathbf{v}$. The minimal ρ that satisfies Theorem 3.6 for $\mathbf{t}$ is $|\mathbf{u}|$, therefore it cannot be that $\mathbf{v} \in F$ because of $\mathbf{t}$. However, this claim holds for any $\mathbf{t}$ whose prefix is $\mathbf{v}$, thus $\mathbf{v} \notin F$. Second, observe that F is maximal since we consider all possible transcripts $\mathbf{t}$. □

[4] That is, no string in F is a proper prefix of another string in F. Formally, for any $\mathbf{t}, \mathbf{t}' \in F$ where $|\mathbf{t}| \le |\mathbf{t}'|$, it holds that $\mathbf{t} \neq \mathbf{t}'_{[1..|\mathbf{t}|]}$.

We are now ready for the main theorem of this section. Let f be a binary-uniquely decomposable function on $X \times Y$ and let π be a perfectly correct and private protocol for f. We define a sequence of domains $X_1 \times Y_1 \subseteq \cdots \subseteq X_d \times Y_d$, according to the partitioning of f, where $X_d \times Y_d = X \times Y$. For any such sequence we define a sequence of frontiers $F_i = F(X_i \times Y_i)$ where F_i represents one decomposition step of the protocol (i.e., the decomposition of $X_i \times Y_i$ into $X_{i-1} \times Y_{i-1}$ and $(X_i \setminus X_{i-1}) \times Y_{i-1}$, assuming that the i-th decomposition level is a row-partition). Furthermore, it holds that $F_{i+1} \leq F_i$, that is, the protocol reaches these frontiers exactly in their order (i.e., the information about f's decomposition is revealed exactly in this order).

Theorem 3.11. *Suppose that f is binary-uniquely decomposable and assume that the maximal number of decomposition in f is $d \geq 1$. Let π be a perfectly-correct private protocol for f over Ch. Let $X_1 \times Y_1 \subseteq \cdots \subseteq X_d \times Y_d$ be a sequence of restricted domains such that $X_d \times Y_d = X \times Y$, and $X_i \times Y_i$ is one of the partitions of $X_{i+1} \times Y_{i+1}$. Then, for any $X_i \times Y_i$ there exists a frontier $F_i = F(X_i \times Y_i)$ such that (assuming $f_{X_i \times Y_i}$ is row-partitionable; a similar claim holds if $f_{X_i \times Y_i}$ is column-partitionable.), for any $x, x' \in X_i$ and $y, y' \in Y_i$, and any $\mathbf{u} \in F_i$,*

1. *if x, x' are in the same partition,* $\Pr[\mathbf{u} \mid x, y, \mathsf{Ch}] = \Pr[\mathbf{u} \mid x', y', \mathsf{Ch}]$.
2. *if x, x' belong to different partitions, at least one of* $\Pr[\mathbf{u} \mid x, y, \mathsf{Ch}]$, $\Pr[\mathbf{u} \mid x', y', \mathsf{Ch}]$ *is zero.*

Moreover, for any restricted domain $X_j \times Y_j$ such that $X_i \times Y_i \subset X_j \times Y_j$, it holds that $F_j \leq F_i$.

Proof. Assume $f_{X_i \times Y_i}$ is row-partitionable into $P, Q \subset X$. We prove the claim by induction on $i = 1 \ldots d$. We split the proof into cases according to Lemma 3.2.

Case I: $f_{P \times Y_i}, f_{Q \times Y_i}$ are constant. Let F_i be the set of complete transcripts $F_i = \{\mathbf{t} \in \tau^{\mathsf{Ch}}(x, y) \mid x, y \in X_i \times Y_i\}$, and the claim trivially follows from privacy, and Lemma 2.7. Also note that this must be the case for the induction's base case $X_1 \times Y_1$.

Case II: $f_{P \times Y_i}$ is constant but $f_{Q \times Y_i}$ is not. By Lemma 3.2 we know that there exists a minor $(x_\circ, x'_\circ, y_\circ, y'_\circ)$ where $x_\circ \in P$ and $x'_\circ \in Q$. Define $F_i = F(x_\circ, x'_\circ, y_\circ, y'_\circ)$ as given by Definition 3.9. We begin with proving the first property.

If $x, x' \in P$ the first property trivially holds from the privacy of (the constant) $f_{P \times Y_i}$.

If $x, x' \in Q$ then, assume $f_{Q \times Y_i}$ is column-partitionable into Y_i^0, Y_i^1, and recall that Lemma 3.2 suggests that $y_\circ$ and $y'_\circ$ are in different partitions, say $y_\circ \in Y_i^0$ and $y'_\circ \in Y_i^1$. Now, if y, y' are in the same partition, then the first property holds due to the induction hypothesis on $Q \times Y_i$. Namely, $Q \times Y_i \subset X_i \times Y_i$ and the induction implies that any string $\mathbf{u'}$ in $F_{i-1} = F(Q \times Y_i)$ satisfies property (1.); since $F_i \leq F_{i-1}$, then any $\mathbf{u} \in F_i$ is a prefix of some $u' \in F_{i-1}$ and must satisfy the same property.

If $y \in Y_i^0$ and $y' \in Y_i^1$ then, again using the induction hypothesis we know that for any $\mathbf{u} \in F_i$ and for any $x, x' \in Q$,

$$\Pr[\mathbf{u} \mid x, y, \mathsf{Ch}] = \Pr[\mathbf{u} \mid x'_\circ, y_\circ, \mathsf{Ch}] \quad \text{and} \quad \Pr[\mathbf{u} \mid x', y', \mathsf{Ch}] = \Pr[\mathbf{u} \mid x'_\circ, y'_\circ, \mathsf{Ch}].$$

The claim follows since $\Pr[\mathbf{u} \mid x'_\circ, y_\circ, \mathsf{Ch}] = \Pr[\mathbf{u} \mid x'_\circ, y'_\circ, \mathsf{Ch}]$ as given by Theorem 3.6, and the way we have defined F_i.

We now prove the second property. Let $x \in P, x' \in Q$ and $y, y' \in Y_i$. Note that due to Theorem 3.6, for any $\mathbf{u} \in F_i$ it holds that at least one of $\Pr[\mathbf{u} \mid x'_\circ, y_\circ, \mathsf{Ch}]$ and $\Pr[\mathbf{u} \mid x_\circ, y'_\circ, \mathsf{Ch}]$ is zero. Also note that $\Pr[\mathbf{u} \mid x, y, \mathsf{Ch}] = \Pr[\mathbf{u} \mid x_\circ, y, \mathsf{Ch}]$ and $\Pr[\mathbf{u} \mid x', y', \mathsf{Ch}] = \Pr[\mathbf{u} \mid x'_\circ, y', \mathsf{Ch}]$, the former due to privacy (recall that $f_{P \times Y_i}$ is constant) and the latter due to the induction hypothesis on $Q \times Y_i$. This completes the proof of property (2.) for this case.

Case III: both $f_{P\times Y_i}$ and $f_{Q\times Y_i}$ are non constant. Assume $f_{P\times Y_i}$ is column-partitionable to Y_p^0, Y_p^1 and that $f_{Q\times Y_i}$ to Y_q^0, Y_q^1. Let $(x_\circ, x'_\circ, y_\circ, y'_\circ)$ and $(x_\circ, x'_\circ, w_\circ, w'_\circ)$ be the minors guaranteed by Lemma 3.2. Define F_i to be $F(x_\circ, x'_\circ, y_\circ, y'_\circ)$, and note that $F(x_\circ, x'_\circ, y_\circ, y'_\circ) = F(x_\circ, x'_\circ, w_\circ, w'_\circ)$: by using the induction hypothesis on each of the domains $Q \times Y_{Q_0}$, $Q \times Y_{Q_1}$, $P \times Y_{P_0}$,$P \times Y_{P_1}$ we get that $\Pr[\mathbf{u} \mid x, y, \mathsf{Ch}] = \Pr[\mathbf{u} \mid x, y', \mathsf{Ch}] = \Pr[\mathbf{u} \mid x, w', \mathsf{Ch}]$ and $\Pr[\mathbf{u} \mid x', w, \mathsf{Ch}] = \Pr[\mathbf{u} \mid x', w', \mathsf{Ch}] = \Pr[\mathbf{u} \mid x', y, \mathsf{Ch}]$ (refer to Figure 1). By using the same reasoning as in Case II, the first term also equals $\Pr[\mathbf{u} \mid x, w, \mathsf{Ch}]$ and the second also equals $\Pr[\mathbf{u}' \mid x', y', \mathsf{Ch}]$ which implies these frontiers are in fact the same.

It is easy to verify that the induction holds for this case as well: the first property holds due to the same reason as the case where $x, x' \in Q$ in Case II; when $x, x' \in P$ we use the minor $(x_\circ, x'_\circ, w_\circ, w'_\circ)$ and when $x, x' \in Q$ we use the minor $(x_\circ, x'_\circ, y_\circ, y'_\circ)$.

The second property also follows the same way as in Case II.

Finally, we complete the induction by showing that $F_i \le F_{i-1}$ (assuming the induction holds for all F_j with $j < i$). Suppose not, and let $\mathbf{u} \in F_i$ be a transcript that violates this condition, that is, there exists $\mathbf{u}' \in F_{i-1}$ such that $\mathbf{u}'$ is a strict prefix of $\mathbf{u}$. Due to the way we construct the frontiers, either F_i contains only complete transcripts (and then any frontier F_j satisfies $F_j \le F_i$) or otherwise F_{i-1} is defined via Definition 3.9 by some minor x, x', y, y' where $x, x' \in X_{i-1}$ and $y, y' \in Y_{i-1}$ (wlog, a ⊞-minor); additionally, let v, v', w, w' be the (⊞-)minor that defines F_i, where $v, v' \in X_i$ and $w, w' \in Y_i$.

By Theorem 3.6 we know that for any $\mathbf{u}' \in F_{i-1}$, either $\Pr[\mathbf{u}' \mid x, y, \mathsf{Ch}] = 0$ or $\Pr[\mathbf{u}' \mid x, y', \mathsf{Ch}] = 0$. If both are 0 then $\mathbf{u}'$ (or prefix or it) is in F_i by its construction and it can't be that $\mathbf{u} \in F_i$ for any $\mathbf{u}$ that has a proper prefix $\mathbf{u}'$.

Otherwise, since $\mathbf{u}'$ is a prefix of $\mathbf{u}$, it follows that also either $\Pr[\mathbf{u} \mid x, y, \mathsf{Ch}] = 0$ or $\Pr[\mathbf{u} \mid x', y, \mathsf{Ch}] = 0$. However, X_{i-1} and $X_i \setminus X_{i-1}$ is a valid partitioning of X_i, thus x, x' are both in the same partition after the decomposition step of F_i. By the way we construct F_i, it holds that $\Pr[\mathbf{u} \mid v', w, \mathsf{Ch}] = \Pr[\mathbf{u} \mid v', w', \mathsf{Ch}]$. Furthermore, using the induction hypothesis on the two partitions of F_{i-1}, we get that $\Pr[\mathbf{u} \mid x, y, \mathsf{Ch}] = \Pr[\mathbf{u} \mid x', y, \mathsf{Ch}] = \Pr[\mathbf{u} \mid v', w, \mathsf{Ch}]$ as well as $\Pr[\mathbf{u} \mid x, y', \mathsf{Ch}] = \Pr[\mathbf{u} \mid x', y', \mathsf{Ch}] = \Pr[\mathbf{u} \mid v', w', \mathsf{Ch}]$. It follows that $\Pr[\mathbf{u} \mid x, y', \mathsf{Ch}] = \Pr[\mathbf{u} \mid x, y, \mathsf{Ch}] = 0$ and we reached a contradiction. □

4 Impossibility of Constant-Rate Coding for Private Protocols

Intuitively, Theorem 3.11 implies that the information is leaked by order: first A speaks and tells where, in the first decomposition his input lies (and until then, B doesn't say anything meaningful); then it's B's turn to reveal where his input lies in the next decomposition, etc. It is not allowed to give information beyond the current point of decomposition (this will damage privacy), and moreover, if the other side does not acknowledge the correct partition of the decomposition, he might violate his privacy (or at least the protocol outputs a wrong output). This gives a strict bound on the error rate any such protocol can withstand.

Before we prove our impossibility result, we discuss a subtlety regarding the definition of error rate in protocols of varying length. Consider a (deterministic) protocol so that for any input x, y we can match a single transcript $t_{(x,y)}$. Assume that for any $(x, y) \neq (x', y')$ it holds that $|t_{(x,y)}| \neq |t_{(x',y')}|$. In this setting, Definition 2.4 has no practical meaning. Indeed, if the channel is allowed to corrupt a fraction $c \in (0, 1)$ of the transmission, and the protocol runs on the input whose transcript t is the longest, the protocol might terminate before round $|t|$, and the ‘effective’ fraction of errors in the observed transcript will be higher than c.

To overcome this issue, we restrict our discussion only to protocols whose output length is the same for every transcript. This doesn’t restrict the power of the protocol: after executing $\mathsf{out}_A, \mathsf{out}_B$ the parties can send “0” as many times as needed, ignoring any received messages. On the other hand, the channel is potentially stronger, as it is allowed a higher proportion of errors for transcripts that prematurely terminate (if any exists). We also mention that protocols in which both parties always have consensus about whose turn it is to speak and whether or not the protocol ended, *must* be of the same length for any input; see Claim 9 in [BR11].

Theorem 4.1 (main). *Let f be a function that is binary-uniquely decomposable and has maximal decomposition depth d. Then, no private protocol for f is perfectly correct and perfectly private with error rate $1/d$. Namely, for every perfectly correct protocol π for f, there exists a channel Ch^* with error rate $1/d$, such that π^{Ch^*} is not private.*

We begin with an intuitive outline of the proof. For any protocol π, we will construct a channel Ch^* which violates party B’s privacy. In particular, we will show there exist inputs $x \in X, y, y' \in Y$, and a transcript $\mathbf{t}$, such that $f(x, y) = f(x, y')$, but $\Pr[\mathbf{t} \mid x, y, \mathsf{Ch}^*] \neq \Pr[\mathbf{t} \mid x, y', \mathsf{Ch}^*]$.

We start with inputs $x_0, x_1 \in X, y_0, y_1 \in Y$ which form a ⊞-minor, and analyze protocol π with the noiseless channel Ch^0. Theorem 3.11 tells us precisely how information about the inputs (x_0, x_1, y_0, y_1) is revealed by π: *first* A reveals whether its input is in the partition corresponding to x_0 or x_1, *then*, if A’s input is in the partition corresponding to x_1, B reveals whether its input is in the partition corresponding to y_0 or y_1. More precisely, there exists $1 \leq i \leq d$, such that,

- till frontier F_i, party A does not reveal any information that can distinguish the partitions corresponding to x_0 from x_1,
- at frontier F_{i+1}, party A reveals whether its input is in the partition corresponding to x_0 or x_1,
- if A’s input was in partition corresponding to x_1, B reveals the partition of its own input at frontier F_{i+2}.

We use the above observations to design the adversarial channel Ch^*. Consider an execution of π with inputs x_0, y_0. The channel lets the messages go back and forth unmodified till the protocol reaches frontier F_i. At this stage, it samples a random tape R_A for A consistent with the transcript so far, *and input* x_1. Thereafter, the channel replaces A’s messages with what it would have sent had its input been x_1 and random tape R_A. Note that as transcripts for x_0 and x_1 are identically distributed till frontier F_i, B’s view is consistent with the execution of π with inputs x_1, y_0 with the *noiseless channel.* The channel Ch^* continues modifying A’s messages this way until in B’s view, the protocol reaches frontier F_{i+2}. At this point party B has revealed information about the partition of its input, which violates its privacy.

The crucial missing piece in the above is the rate of the channel Ch^*. Recall that our channel is allowed to introduce only a $1/d$ fraction of errors. If we pick an arbitrary minor (x_0, x_1, y_0, y_1) and launch the above attack, it is possible that the frontiers F_{i+2} and F_i are “far apart”, and thus the channel has to introduce large number of errors. To avoid this, we start by picking inputs x_1, y_0, and a transcript $\mathbf{t}$ in $\tau^{\mathsf{Ch}}(x_1, y_0)$. As there are d frontiers, by an averaging argument, there exists an i such that the fraction of messages in $\mathbf{t}$ between F_i and F_{i+2} is $2/d$. Further, there must exist x_0, y_1 such that (x_0, x_1, y_0, y_1) forms a ⊞-minor which corresponds to the frontier F_i. The channel Ch^* uses this minor for the attack described above. The formal proof follows.

Proof. (**Theorem 4.1**) Let Ch^0 be the 'noiseless channel', that is, for any $\ell \in \mathbb{N}$ and any sequence of messages $\boldsymbol{m} = m_1 \cdots m_\ell$,

$$\Pr[\mathsf{Ch}^0(\boldsymbol{m}) = \boldsymbol{m}] = 1.$$

Suppose $d > 2$ (the theorem trivially holds for $d \leq 2$) and let $X_\circ \times Y_\circ$ be a subdomain obtained by recursively performing d decomposition of f's domain (a restricted domain of depth d). Let $F_1, \ldots, F_d$ be the frontiers given by Theorem 3.11, and let F_0 be the empty set. We re-index the frontiers in a reverse order so that $F_d = F(X_\circ \times Y_\circ)$ and $F_1 \leq \cdots \leq F_d$. Let $x, y \in X_\circ \times Y_\circ$ and fix $\mathbf{t} \in \tau^{\mathsf{Ch}^0}(x, y)$. For $1 \leq j \leq d$, let $\lambda_j(\mathbf{t})$ be the length of the longest prefix of $\mathbf{t}$ in F_j. Then for $0 \leq j < k \leq d$, we can define the *number of messages in* $\mathbf{t}$ *between* F_j *and* F_k, $\lambda_{j,k}(\mathbf{t})$, as

$$\lambda_{j,k}(\mathbf{t}) \coloneqq \lambda_k(\mathbf{t}) - \lambda_j(\mathbf{t}).$$

As $\sum_{j=1}^{d} \lambda_{j-1,j}(\mathbf{t}) = |\mathbf{t}|$, there exists i such that $\lambda \coloneqq \lambda_{i,i+2}(\mathbf{t}) \leq 2|\mathbf{t}|/d$.

For any partial transcript $\mathbf{u}$, we define the *active domain at* $\mathbf{u}$ as the set of inputs that have a non-zero probability of producing a transcript with prefix $\mathbf{u}$. Let $X_i \times Y_i$ be the active domain at the longest prefix of $\mathbf{t}$ in F_i (clearly $(x, y) \in X_i \times Y_i)$). Let the next decomposition step be a row-decomposition, with row partitions $X_i^0, X_i^1 \subset X_i$. Let us rename the partitions so that $x \in X_i^1$. First, as the decomposition depth of (x, y) is $d \geq i+2$, it can not be the case that both $f_{X_i^0, Y_i}$ is constant, and $f_{X_i^1, Y_i}$ is constant. In particular, as $x \in X_i^1$, $f_{X_i^1, Y_i}$ is not constant. Let $Y_i^0, Y_i^1 \subset Y_i$ be the next column-partition of $X_i^1 \times Y_i$. Then, by Lemma 3.2, there exist $x_0 \in X_i^0, x_1 \in X_i^1, y_0 \in Y_i^0, y_1 \in Y_i^1$ such that (x_0, x_1, y_0, y_1) form a ⊞-minor. As $f(x, y_0) \neq f(x, y_1)$, inputs (x_0, x, y_0, y_1) also form a ⊞-minor.

Now we will construct a channel Ch^* that violates B's privacy. Let $\mathbf{u}$ be the longest prefix of $\mathbf{t}$ in F_i. Fix, $\mathbf{t}' \in \tau^{\mathsf{Ch}^0}(x_0, y_0)$ that also has $\mathbf{u}$ as a prefix. The existence of $\mathbf{t}'$ is guaranteed by Theorem 3.11. Define the channel Ch^* as follows:

- Ch^* does not introduce any errors till $|\mathbf{u}|$ rounds. If the partial transcript so far is not identical to $\mathbf{u}$, the channel acts as the noiseless channel Ch^0 for the entire protocol.
- If the partial transcript up to round $|\mathbf{u}|$ is identical to $\mathbf{u}$, then the channel modifies messages as follows: sample a random tape R_A for party A from the conditional distribution over random tapes given the transcript so far and the input $x \in X_i^1$. For the next $\lambda = \lambda_{i,i+2}$ rounds, send B's messages to A without change. Change A's sent messages to what A would have sent when using input x and the random tape sampled earlier. That is, for every message m, every odd round j such that $|\mathbf{u}| < j \leq |\mathbf{u}| + \lambda$,

$$\Pr\left[\mathsf{Ch}^*\left(t'^s[j] \mid \mathbf{t}'[1..(j-1)]\right) = m\right] = \Pr\left[\pi_A\left(x, R_A, t^s[1..(j-1)]\right) = m\right].$$

The channel continues with the remaining protocol without modifying any messages.

First, note that Ch^* introduces only $1/d$ fraction of errors. This is because it only changes half of the messages between frontiers F_i and F_{i+2}, of which there are only a fraction $2/d$ of the length of transcript $\mathbf{t}$. As all transcripts are of the same length, the channel makes only $1/d$ fraction errors.

Next, we argue that channel Ch^* violates B's privacy. Fix a transcript $\mathbf{t}^* \in \tau^{\mathsf{Ch}^*}(x_0, y_0)$ which has $\mathbf{u}$ as a prefix. Again, the existence of $\mathbf{t}^*$ is guaranteed by Theorem 3.11. We will show that $\Pr[\mathbf{t}^* \mid x_0, y_1, \mathsf{Ch}^*] = 0$.

We have,

$$\Pr[(t^{*s}, t^{*r}) \mid x_0, y_1, \mathsf{Ch}^*] = \Pr[\pi_A(x_0, t^{*r}_A) = t^{*s}_A] \Pr[\mathsf{Ch}^*(t^{*s}) = t^{*r}] \Pr[\pi_B(y_1, t^{*r}_B) = t^{*s}_B]. \qquad (8)$$

To maintain readability of the probability calculations below, we take two notational shortcuts: (1) all probabilities below are conditioned on observing the prefix $\mathbf{u}$, but we do not explicitly mention this conditioning, and (2) we analyze $\mathbf{t}^*$ only up to round $|\mathbf{u}| + \lambda$, so from now on we will take $\mathbf{t}^*$ to mean the prefix $\mathbf{t}^*{}_{[1..(|\mathbf{u}|+\lambda)]}$ (same for t^{*s}, t^{*r}).

By the channel's definition,

$$\begin{aligned}\Pr[\mathsf{Ch}^*(t^{*s}) = t^{*r}] = &\left(\prod_{j=|\mathbf{u}|+1}^{(|\mathbf{u}|+\lambda)/2} \Pr\left[\mathsf{Ch}^*(t^{*s}[2j-1] \mid \mathbf{t}^*{}_{[|\mathbf{u}|..(2j-2)]}) = t^{*r}[2j-1]\right]\right) \times \\ &\left(\prod_{j=|\mathbf{u}|+1}^{(|\mathbf{u}|+\lambda)/2} \Pr\left[\mathsf{Ch}^*(t^{*s}[2j] \mid \mathbf{t}^*{}_{[|\mathbf{u}|..(2j-1)]}) = t^{*r}[2j]\right]\right).\end{aligned}$$

Note that for every even round j, $t^{*s}[j] = t^{*r}[j]$, as Ch^* does not modify B's messages to A. Thus, the second product on the right hand side above is 1. By construction of Ch^*, for every odd $|\mathbf{u}| + 1 \leq j \leq (|\mathbf{u}| + \lambda)$,

$$\Pr\left[\mathsf{Ch}^*(t^{*s}[2j-1] \mid \mathbf{t}^*{}_{[|\mathbf{u}|..(2j-2)]}) = t^{*r}[2j-1]\right] = \Pr\left[\pi_A(x, t^{*s}{}_{[|\mathbf{u}|..(2j-2)]}) = t^{*r}[2j-1]\right]$$

Going back to Eq. (8), we have,

$$\begin{aligned}\Pr[(t^{*s}, t^{*r}) \mid x_0, y_1, \mathsf{Ch}^*] &= \Pr[\pi_A(x_0, t^{*r}_A) = t^{*s}_A] \Pr[\mathsf{Ch}^*(t^{*s}) = t^{*r}] \Pr[\pi_B(y_1, t^{*r}_B) = t^{*s}_B] \\ &= \Pr[\pi_A(x_0, t^{*r}_A) = t^{*s}_A] \times \\ &\quad \prod_{j=|\mathbf{u}|}^{(|\mathbf{u}|+\lambda)/2} \Big(\Pr\left[\pi_A(x, t^{*s}{}_{[|\mathbf{u}|..(2j-2)]}) = t^{*r}[2j-1]\right] \times \\ &\qquad \Pr\left[\pi_B(y_1, t^{*r}{}_{[|\mathbf{u}|..(2j-1)]}) = t^{*s}[2j]\right]\Big) \\ &= \Pr[\pi_A(x_0, t^{*r}_A) = t^{*s}_A] \Pr[\mathbf{t}^{**} \mid x, y_1, \mathsf{Ch}^0].\end{aligned}$$

where $\mathbf{t}^{**} = (t_1^{**s}, t_1^{**r})$ is the following transcript: t^{**s} consists of all messages sent by the channel to B and all messages sent by B, and $t^{**r} = t^{**s}$. We now prove that $\Pr[\mathbf{t}^{**} \mid x, y_1, \mathsf{Ch}^0] = 0$.

Recall that y_0 and y_1 are in different (column) partitions of $X_i^1 \times Y_i$. Also, by construction of Ch^*, $\Pr[\mathbf{t}^{**} \mid x, y_0, \mathsf{Ch}^0] > 0$. Therefore, by Theorem 3.11, $\Pr[\mathbf{t}^{**} \mid x_0, y_1, \mathsf{Ch}^0] = 0$. Thus,

$$\Pr[\mathbf{t}^* \mid x_0, y_1, \mathsf{Ch}^*] = 0,$$

and although $f(x_0, y_0) = f(x_0, y_1)$ we learn that B holds y_0 rather than y_1, in violation of B's privacy according to Definition 2.5. □

5 Properties of Private Protocol: The Non-Perfect Case

The above results assume a perfectly-correct protocol, which is interesting on its own. Indeed, possibility results for computing non-private functions over adversarial channels [Sch96, BR11] do achieve perfect correctness for error rate less than 1/4. However, from a practical point of view, the question of computing private functions over a noisy channel is interesting also for the non-perfect case, that is, where the protocol is only $(1 - \varepsilon)$-correct and $(1 - \delta)$-private, for some $\varepsilon, \delta > 0$.

Definition 5.1 (Correctness). *A protocol π is $(1-\varepsilon)$-correct with respect to some function f over Ch if for any inputs x, y,*

$$\Pr_{\substack{R_A, R_B, R_{\mathsf{Ch}} \\ (t^s, t^r) \leftarrow \pi^{\mathsf{Ch}}(x,y)}} [\mathsf{out}^{\pi}_A(x, t^r_A) = \mathsf{out}^{\pi}_B(y, t^r_B) = f(x,y)] \geq 1 - \varepsilon.$$

Definition 5.2 (Privacy). *We say that a protocol π is $(1-\delta)$-private with respect to party A over a channel Ch, if for any $x, x' \in X$, $y \in Y$ such that $f(x,y) = f(x',y)$, and for any frontier F,*

$$\mathsf{SD}^{\mathsf{Ch}}_F((x,y),(x',y)) \triangleq \frac{1}{2} \sum_{\mathbf{t} \in F} \left| \Pr[\mathbf{t} \mid x, y, \mathsf{Ch}] - \Pr[\mathbf{t} \mid x', y, \mathsf{Ch}] \right| \leq \delta$$

Finally, we say that π is an (ε, δ)-protocol for f if π is $(1-\varepsilon)$-correct and $(1-\delta)$-private.

This leads to our main result for (ε, δ)-protocols,

Theorem 5.3 (main, (ε, δ)-case). *Let f be a function that is binary-uniquely decomposable and has maximal decomposition depth d. Then, for every protocol π and for any small enough $\varepsilon, \delta > 0$, there exists a channel Ch^* with error rate $1/d$, such that π^{Ch^*} is not (ε, δ)-protocol for f.*

Before we prove the main theorem, we show that (ε, δ)-protocols have the same property of revealing the information in a very specific order. This suggests that the same impossibility result holds also for the non-perfect case. The following two lemmas conceptually extend Theorem 3.6 and Theorem 3.11 to (ε, δ)-protocols. We also note that similar lemmas appear in [MPR09] for the case where the channel is error-free, and with $\epsilon = \delta$.

Lemma 5.4. *Let π be a (ε, δ)-protocol for f over Ch, with $\varepsilon, \delta > 0$. For any $x \neq x'$ there exists a frontier F (intuitively, the frontier that is associated to the partition separating x from x') such that,*

1. *If $f(x,y) \neq f(x',y)$ then $\mathsf{SD}^{\mathsf{Ch}}_F((x,y),(x',y)) > 1 - 5\sqrt{\varepsilon + \delta}$.*
2. *For x, x', y, y' that form a $\boxplus$-minor, $\mathsf{SD}^{\mathsf{Ch}}_F((x',y),(x',y')) < \sqrt{\delta}$ and $\mathsf{SD}^{\mathsf{Ch}}_F((x,y),(x,y')) < \delta$.*

Proof. Recall that in our model parties alternately send messages $m \in \Sigma$, where for odd rounds $m_i = \pi_A(x, m_1 \cdots m_{i-1})$, while for even rounds $m_i = \pi_B(y, m_1 \cdots m_{i-1})$. Given an input (x,y) we expand the probability of a transcript $\mathbf{t}$ according to its prefixes $\mathbf{t}_{[1..\rho]}$.

$$\begin{aligned}
\Pr[\mathbf{t} \mid x, y, \mathsf{Ch}] &= \underbrace{\Pr[\pi_A(x,y,\emptyset) = t^s[1]] \Pr[\mathsf{Ch}(t^s[1]) = t^r[1]]}_{\Pr[\mathbf{t}_{[1..1]} \mid x, y, \emptyset, \mathsf{Ch}]} \times \underbrace{\Pr[\pi_B(x,y,t^r{}_{[1..1]}) = t^s[2]) \Pr[\mathsf{Ch}(t^s[2] \mid \mathbf{t}_{[1..1]}) = t^r[2])}_{\Pr[\mathbf{t}_{[1..2]} \mid x, y, \mathbf{t}_{[1..1]}, \mathsf{Ch}]} \times \cdots \\
&= \left(\prod_{j=1}^{|\mathbf{t}|/2} \Pr\left[\mathbf{t}_{[1..2j-1]} \mid x, y, \mathbf{t}_{[1..2j-2]}, \mathsf{Ch}\right] \right) \left(\prod_{j=1}^{|\mathbf{t}|/2} \Pr[\mathbf{t}_{[1..2j]} \mid x, y, \mathbf{t}_{[1..2j-1]}, \mathsf{Ch}] \right) \\
&= P_A(\mathbf{t}, x, \mathsf{Ch}) P_B(\mathbf{t}, y, \mathsf{Ch})
\end{aligned}$$

where P_A and P_B are defined as the terms in the parenthesis.

For a fixed $\mu < 1$ to be defined later, we say that a transcript (or a prefix of a transcript) $\mathbf{u}$ *distinguishes* between x and x' if

$$\frac{|P_A(\mathbf{u}, x, \mathsf{Ch}) - P_A(\mathbf{u}, x', \mathsf{Ch})|}{P_A(\mathbf{u}, x, \mathsf{Ch}) + P_A(\mathbf{u}, x', \mathsf{Ch})} \geq \mu$$

Next, we define a frontier F as a function of x, x' by considering the frontier induced by the following set

$$\{\mathbf{u} \mid \mathbf{u} \text{ is a complete transcript, or } \mathbf{u} \text{ distinguishes between } x \text{ and } x' \text{ over } \mathsf{Ch}\}$$

where a complete transcript is a transcript of full length (recall all inputs terminate after the same number of rounds) on which the parties terminate the protocol and give output. Note that any string in F either distinguishes x, x' or is a complete transcript, and let $F = F_{dist} \cup F_{comp}$ according to this division.

First, let's bound the weight of transcripts that do not distinguish between x, x', given the input (x, y). For any possible output k we set

$$Out_k = \{\mathbf{u} \in F_{comp} \mid \mathbf{u} \text{ implies output } k \text{ with probability } > 1/2\},$$

and let $Not_k = F_{comp} \setminus Out_k$. Note that $Out_k \cap Out_j = \emptyset$ for $k \neq j$. Due to correctness note that $\sum_{\mathbf{u} \in Not_{f(x,y)}} \Pr[\mathbf{u} \mid x, y, \mathsf{Ch}] \leq 2\varepsilon$ and similarly, $\sum_{\mathbf{u} \in Not_{f(x',y)}} \Pr[\mathbf{u} \mid x', y, \mathsf{Ch}] \leq 2\varepsilon$. Now,

$$\begin{aligned}
\Pr[F_{comp} \mid x, y, \mathsf{Ch}] &\leq \sum_{\mathbf{u} \in Not_{f(x,y)}} P_A(\mathbf{u}, x, \mathsf{Ch}) P_B(\mathbf{u}, y, \mathsf{Ch}) + \sum_{\mathbf{u} \in Not_{f(x',y)}} P_A(\mathbf{u}, x, \mathsf{Ch}) P_B(\mathbf{u}, y, \mathsf{Ch}) \\
&\leq \sum_{\mathbf{u} \in Not_{f(x,y)}} P_A(\mathbf{u}, x, \mathsf{Ch}) P_B(\mathbf{u}, y, \mathsf{Ch}) + \frac{1+\mu}{1-\mu} \sum_{\mathbf{u} \in Not_{f(x',y)}} P_A(\mathbf{u}, x', \mathsf{Ch}) P_B(\mathbf{u}, y, \mathsf{Ch}) \\
&= 2\varepsilon \left(1 + \frac{1+\mu}{1-\mu}\right) = \frac{4\varepsilon}{1-\mu}
\end{aligned}$$

where the second transition is due to the fact that $\mathbf{u}$ is not distinguishing, thus $P_A(\mathbf{u}, x, \mathsf{Ch}) \leq \frac{1+\mu}{1-\mu} P_A(\mathbf{u}, x', \mathsf{Ch})$.

For any $\mathbf{u} \in F_{dist}$ it holds that

$$\begin{aligned}
\mathsf{SD}^{\mathsf{Ch}}_{F_{dist}}((x, y), (x', y)) &= \tfrac{1}{2} \sum_{\mathbf{u} \in F_{dist}} |P_A(\mathbf{u}, x, \mathsf{Ch}) P_B(\mathbf{u}, y, \mathsf{Ch}) - P_A(\mathbf{u}, x', \mathsf{Ch}) P_B(\mathbf{u}, y, \mathsf{Ch})| \\
&= \tfrac{1}{2} \sum_{\mathbf{u} \in F_{dist}} |P_A(\mathbf{u}, x, \mathsf{Ch}) - P_A(\mathbf{u}, x', \mathsf{Ch})| P_B(\mathbf{u}, y, \mathsf{Ch}) \\
&\geq \mu \cdot \tfrac{1}{2} \sum_{\mathbf{u} \in F_{dist}} (P_A(\mathbf{u}, x, \mathsf{Ch}) + P_A(\mathbf{u}, x', \mathsf{Ch})) P_B(\mathbf{u}, y, \mathsf{Ch}) \\
&\geq \frac{\mu}{2} \left(\Pr[F_{dist} \mid x, y, \mathsf{Ch}] + \Pr[F_{dist} \mid x', y, \mathsf{Ch}]\right)
\end{aligned}$$

Finally, $\Pr[F_{dist} \mid x, y, \mathsf{Ch}] = 1 - \Pr[F_{comp} \mid x, y, \mathsf{Ch}] \geq 1 - \frac{4\epsilon}{1-\mu}$, and similarly for x', y, and we get,

$$\mathsf{SD}^{\mathsf{Ch}}_{F}((x, y), (x', y)) \geq \mathsf{SD}^{\mathsf{Ch}}_{F_{dist}}((x, y), (x', y)) \geq \mu \left(1 - \frac{4\varepsilon}{1-\mu}\right).$$

For the other part, note that $\mathsf{SD}^{\mathsf{Ch}}_{F}((x, y), (x, y')) < \delta$ trivially by the privacy of the protocol. In order to bound $\mathsf{SD}^{\mathsf{Ch}}_{F}((x', y), (x', y'))$ we again split F into $F_{comp} \cup F_{dist}$ as above. We define F^{-}_{dist} as the prefixes obtained by taking all $\mathbf{u} \in F_{dist}$ and removing the last round (a message which was sent by A). Note that $\mathsf{SD}^{\mathsf{Ch}}_{F^{-}_{dist}}((x', y), (x', y')) = \mathsf{SD}^{\mathsf{Ch}}_{F_{dist}}((x', y), (x', y'))$ since A's message depends only on the input x', and will be distributed the same for y and for y'. On the other hand, by the way we have defined F and F_{dist}, note that all the prefixes in F^-_{dist} are not distinguishing between x and x' anymore, and thus any $\mathbf{u} \in F^-_{dist}$ satisfies $P_A(\mathbf{u}, x', \mathsf{Ch}) \leq \frac{1+\mu}{1-\mu} P_A(\mathbf{u}, x, \mathsf{Ch})$. The same holds to any $\mathbf{u} \in F_{comp}$, since it is not distinguishing to begin with. Now,

$$\begin{aligned}
\mathsf{SD}^{\mathsf{Ch}}_F((x', y), (x', y')) &= \mathsf{SD}^{\mathsf{Ch}}_{F_{comp}}((x', y), (x', y')) + \mathsf{SD}^{\mathsf{Ch}}_{F_{dist}}((x', y), (x', y')) \\
&= \mathsf{SD}^{\mathsf{Ch}}_{F_{comp}}((x', y), (x', y')) + \mathsf{SD}^{\mathsf{Ch}}_{F^-_{dist}}((x', y), (x', y')) \\
&= \tfrac{1}{2} \sum_{\mathbf{u} \in F_{comp} \cup F^-_{dist}} P_A(\mathbf{u}, x', \mathsf{Ch}) |P_B(\mathbf{u}, y, \mathsf{Ch}) - P_B(\mathbf{u}, y', \mathsf{Ch})| \\
&\leq \tfrac{1}{2} \frac{1+\mu}{1-\mu} \sum_{u \in F_{comp} \cup F^-_{dist}} P_A(\mathbf{u}, x, \mathsf{Ch}) |P_B(\mathbf{u}, y, \mathsf{Ch}) - P_B(\mathbf{u}, y', \mathsf{Ch})| \\
&\leq \tfrac{1}{2} \frac{1+\mu}{1-\mu} \cdot \mathsf{SD}^{\mathsf{Ch}}_F((x, y), (x, y')) = \tfrac{1}{2} \frac{1+\mu}{1-\mu} \cdot \delta
\end{aligned}$$

Choosing $\mu = 1 - \sqrt{\varepsilon + \delta}$ completes the proof. □

Lemma 5.5. *Assume that f is binary-uniquely decomposable with at most $d \geq 1$ levels, and π is a (ε, δ)-protocol for f. Let $X_1 \times Y_1 \subseteq \cdots \subseteq X_d \times Y_d$ be a sequence of restricted domains such that $X_d \times Y_d = X \times Y$, and $X_i \times Y_i$ is one of the partitions of $X_{i+1} \times Y_{i+1}$. For any restricted domain $X_i \times Y_i$ there exists a frontier $F_i = F(X_i \times Y_i)$ such that (assuming $f_{X_i \times Y_i}$ is row-partitionable; a similar claim holds if $f_{X_i \times Y_i}$ is column-partitionable.), for any $x, x' \in X_i$ and $y, y' \in Y_i$, and any $u \in F_i$,*

1. *if x, x' are in the same partition, $\mathsf{SD}^{\mathsf{Ch}}_{F_i}((x, y), (x', y')) < \mu_i$*
2. *if x, x' belong to different partitions, $\mathsf{SD}^{\mathsf{Ch}}_{F_i}((x, y), (x', y')) > 1 - \mu_i$*

with $\mu_i = 2^{O(i)} \sqrt{\varepsilon + \delta}$.
Moreover, for any restricted domains $X_{i-1} \times Y_{i-1} \subset X_i \times Y_i$, it holds that except with probability μ_i, the protocol reaches $F(X_i \times Y_i)$ before reaching $F(X_{i-1} \times Y_{i-1})$, on any input $(x, y) \in X_{i-1} \times Y_{i-1}$.

Proof. Before we begin the proof, we recall the following technical lemma (Lemma 5 in [MPR09])

Lemma 5.6. *For frontiers F and G, for any x, x', y, y',*

$$\mathsf{SD}^{\mathsf{Ch}}_F((x, y), (x', y')) \leq \mathsf{SD}^{\mathsf{Ch}}_G((x, y), (x', y')) + \frac{1}{2} \big(\Pr[G \leq F \mid x, y, \mathsf{Ch}] + \Pr[G \leq F \mid x', y', \mathsf{Ch}]\big)$$

The proof can be found in [MPR09].

Assume $f_{X_i \times Y_i}$ is row-partitionable into $P, Q \subset X_i$. As in the perfect case, we prove the claim by induction on $i = 1 \ldots d$. We split the proof into cases according to Lemma 3.2.

Case I: $f_{P \times Y_i}, f_{Q \times Y_i}$ are constant. Let F_i be the set of complete transcripts. Property (1) follows from privacy (for $\mu_i = \delta$). Property (2) follows from correctness (for $\mu_i = 2\varepsilon$). Also note that this must be the case for the induction's base case $X_1 \times Y_1$, thus $\mu_1 = \max(2\varepsilon, \delta)$.

Case II: $f_{P \times Y_i}$ is constant and $f_{Q \times Y_i}$ is not. By by Lemma 3.2 we know that there exist a minor $(x_\circ, x'_\circ, y_\circ, y'_\circ)$ where $x_\circ \in P$ and $x'_\circ \in Q$. Define F_i to be the frontier given by Lemma 5.4 for this minor.

We begin with proving the first property.
case (II)a $x, x' \in P$**.** Privacy guarantees us that for any $y, y' \in Y_i$, $\mathsf{SD}_{F_i}((x,y),(x',y')) < \delta$.
case (II)b, $x, x' \in Q$**.** For this case, the proof works by bounding $\mathsf{SD}_{F_i}((x,y),(x',y))$ using the induction hypothesis, and using the minor $(x_\circ, x'_\circ, y_\circ, y'_\circ)$ as intermediate points as needed. Specifically, assume $f_{Q\times Y_i}$ is column-partitionable into Y_i^0, Y_i^1, and recall that $y_\circ$ and $y'_\circ$ are in different partition by the proof of Lemma 3.2, say $y_\circ \in Y_i^0$ and $y'_\circ \in Y_i^1$. If y, y' are both in the same partition, we directly bound the statistical distance between the points $(x,y) \leftrightarrow (x',y')$, via the induction hypothesis on $Q \times Y_i^0$. If y, y' are in different partitions, say $y \in Y_i^0$ and $y' \in Y_i^1$ we bound each two consecutive points in the path $(x,y) \leftrightarrow (x'_\circ, y_\circ) \leftrightarrow (x'_\circ, y'_\circ) \leftrightarrow (x', y')$, and use a triangle inequality to bound the statistical distance of the entire path.

- y, y' **are both in the same partition** (say, Y_i^0). we use the induction hypothesis on $F_{i-1} = F(Q \times Y_i^0)$, and get that $\mathsf{SD}^{\mathsf{Ch}}_{F_{i-1}}((x,y),(x',y')) \le \mu_{i-1}$. However we need to bound also the effect to the statistical distance made from messages between F_{i-1} and F_i.
  Using Lemma 5.6, we get that
  $$\begin{aligned}\mathsf{SD}^{\mathsf{Ch}}_{F_i}((x,y),(x',y')) &\le \mathsf{SD}^{\mathsf{Ch}}_{F_{i-1}}((x,y),(x',y')) \\ &\qquad + \frac{1}{2}\left(\Pr[F_{i-1} \le F_i \mid x, y, \mathsf{Ch}] + \Pr[F_{i-1} \le F_i \mid x', y', \mathsf{Ch}]\right) \\ &\le \mu_{i-1} + (2\sqrt{\delta} + 3\mu_{i-1}),\end{aligned}$$
  where the second transition is due to Lemma 5.7 (which is the second part of this induction proof; it shows that the probability the frontiers are not "in order" is small).
- y, y' **are in different partitions**, say $y \in Y_i^0$ and $y' \in Y_i^1$. Then, using the induction hypothesis, we know that for any $x, x' \in Q$,
  $$\mathsf{SD}^{\mathsf{Ch}}_{F_{i-1}}((x,y),(x'_\circ,y_\circ)) \le \mu_{i-1} \quad ; \quad \mathsf{SD}^{\mathsf{Ch}}_{F_{i-1}}((x',y'),(x'_\circ,y_\circ)) \le \mu_{i-1}$$
  and in a similar way to the above, we can bound the distance gained between F_{i-1} to F_i,
  $$\begin{aligned}\mathsf{SD}^{\mathsf{Ch}}_{F_i}((x\ ,y\ ),(x'_\circ,y_\circ)) &\le 2\sqrt{\delta} + 4\mu_{i-1}, \\ \mathsf{SD}^{\mathsf{Ch}}_{F_i}((x',y'),(x'_\circ,y_\circ)) &\le 2\sqrt{\delta} + 4\mu_{i-1}.\end{aligned}$$
  In addition note that $\mathsf{SD}^{\mathsf{Ch}}_{F_i}((x'_\circ,y_\circ),(x'_\circ,y'_\circ)) \le \sqrt{\delta}$ by Lemma 5.4 and the way we construct F_i. With a triangle inequality we conclude that
  $$\begin{aligned}\mathsf{SD}_{F_i}((x,y),(x',y')) &\le \mathsf{SD}^{\mathsf{Ch}}_{F_i}((x,y),(x'_\circ,y_\circ)) + \mathsf{SD}^{\mathsf{Ch}}_{F_i}((x'_\circ,y_\circ),(x'_\circ,y'_\circ)) + \mathsf{SD}^{\mathsf{Ch}}_{F_i}((x'_\circ,y'_\circ),(x,y')) \\ &\le 8\mu_{i-1} + 5\sqrt{\delta} \triangleq \beta_i.\end{aligned}$$

We now prove the second property. Let $x \in P, x' \in Q$ and $y, y' \in Y$. Lemma 5.4 tells us that for every $y \in Y$, $\mathsf{SD}^{\mathsf{Ch}}_{F_i}((x_\circ,y),(x'_\circ,y)) > 1 - 5\sqrt{\varepsilon+\delta}$. A triangle inequality tells us that
$$\mathsf{SD}^{\mathsf{Ch}}_{F_i}((x_\circ,y),(x'_\circ,y)) < \mathsf{SD}^{\mathsf{Ch}}_{F_i}((x_\circ,y),(x,y)) + \mathsf{SD}^{\mathsf{Ch}}_{F_i}((x,y),(x',y')) + \mathsf{SD}^{\mathsf{Ch}}_{F_i}((x',y'),(x'_\circ,y)),$$
thus,
$$\begin{aligned}\mathsf{SD}^{\mathsf{Ch}}_{F_i}((x,y),(x',y')) &\ge -\mathsf{SD}^{\mathsf{Ch}}_{F_i}((x'_\circ,y),(x,y)) + \mathsf{SD}^{\mathsf{Ch}}_{F_i}((x_\circ,y),(x'_\circ,y)) - \mathsf{SD}^{\mathsf{Ch}}_{F_i}((x',y'),(x'_\circ,y)) \\ &\ge -\beta_i + 1 - 5\sqrt{\varepsilon+\delta} - \beta_i = 1 - 16\mu_{i-1} - 10\sqrt{\delta} - 5\sqrt{\varepsilon+\delta}.\end{aligned}$$
And we can set $\mu_i \triangleq 16\mu_{i-1} + 10\sqrt{\delta} + 5\sqrt{\varepsilon+\delta}$ (this is also consistent with case III which we analyze immediately).

Case III: $f_{P\times Y_i}$ and $f_{Q\times Y_i}$ are non constant. This case is very similar to Case II. Let $(x,y),(x',y')$ be given. Recall that there exist two overlapping minor $(x_\circ,x'_\circ,y_\circ,y'_\circ)$ and $(x_\circ,x'_\circ,w_\circ,w'_\circ)$. Since both minors share the same $x_\circ,x'_\circ$, the frontier defined by Lemma 5.4 is the same for both minors, and the proof of Case II holds for most of the situations using one of the minors (see Figure 1). The only case which is different is when (using the notations of Figure 1) $(x,y)\in P\times(Y\setminus Y_{P_i})$ and $(x',y')\in Q\times(Y\setminus Y_{Q_j})$ (in Figure 1, the first input is in the upper-left rectangle, and the other is in the lower-right rectangle, which seemingly requires using both minors). For this case, we need to show that $\mathsf{SD}^{\mathsf{Ch}}_{F_i}((x,y),(x',y'))$ is large.

By a triangle inequality (using property (1) of this lemma on inputs in the same partition of F_i),

$$\mathsf{SD}^{\mathsf{Ch}}_{F_i}((x_\circ,y_\circ),(x'_\circ,y_\circ))\le \mathsf{SD}^{\mathsf{Ch}}_{F_i}((x_\circ,y_\circ),(x,y))+\mathsf{SD}^{\mathsf{Ch}}_{F_i}((x,y),(x',y'))+\mathsf{SD}^{\mathsf{Ch}}_{F_i}((x',y'),(x'_\circ,y_\circ))$$

Thus, $\mathsf{SD}^{\mathsf{Ch}}_{F_i}((x,y),(x',y'))\ge 1-5\sqrt{\varepsilon+\delta}-2\beta_i=1-\mu_i$

We are left to prove the last part of the lemma, namely that except with small probability the protocol travels through the frontiers in the right order.

Lemma 5.7. *For any $X_{i-1}\times Y_{i-1}\subset X_i\times Y_i$, on any input $(x,y)\in X_{i-1}\times Y_{i-1}$,*

$$\Pr[F(X_{i-1}\times Y_{i-1})<F(X_i\times Y_i)\mid x,y,\mathsf{Ch}]<2\sqrt{\delta}+3\mu_{i-1}.$$

Proof. As above, this proof is by induction on i, where in the base case we take full transcripts and the claim trivially holds.

Let $(x_\circ,x'_\circ,y_\circ,y'_\circ)$ be the minor that defines $F_i=F(X_i\times Y_i)$. Denote $F_{i-1}=F(X_{i-1}\times Y_{i-1})$, and let $(x,y)\in X_{i-1}\times Y_{i-1}$. Assume that $X_{i-1}\times Y_{i-1}$ is row-partitionable into Y_0 and Y_1 so that $x'_\circ\in X_{i-1}$, $y_\circ\in Y_0$ and $y'_\circ\in Y_1$. (If the next level is not partitionable, then F_{i-1} contains complete transcripts and the claim again trivially holds). By Lemma 5.6,

$$\begin{aligned}&\mathsf{SD}^{\mathsf{Ch}}_{F_{i-1}}((x'_\circ,y_\circ),(x'_\circ,y'_\circ))\\&\qquad\le \mathsf{SD}^{\mathsf{Ch}}_{F_i}((x'_\circ,y_\circ),(x'_\circ,y'_\circ))+\frac{1}{2}\Big(\Pr[F_i<F_{i-1}\mid(x'_\circ,y_\circ),\mathsf{Ch}]+\Pr[F_i<F_{i-1}\mid(x'_\circ,y'_\circ),\mathsf{Ch}]\Big)\end{aligned}$$

Substitue $\Pr[F_{i-1}\le F_i\mid(x'_\circ,y_\circ),\mathsf{Ch}]=1-\Pr[F_i<F_{i-1}\mid(x'_\circ,y_\circ),\mathsf{Ch}]$ (and the equivalent for $x'_\circ,y'_\circ$) to obtain

$$\begin{aligned}\frac{1}{2}\Big(\Pr[F_{i-1}\le F_i\mid(x'_\circ,y_\circ),\mathsf{Ch}]+&\Pr[F_{i-1}\le F_i\mid(x'_\circ,y'_\circ),\mathsf{Ch}]\Big)\\&\le \mathsf{SD}^{\mathsf{Ch}}_{F_i}((x'_\circ,y_\circ),(x'_\circ,y'_\circ))-\mathsf{SD}^{\mathsf{Ch}}_{F_{i-1}}((x'_\circ,y_\circ),(x'_\circ,y'_\circ))+1\\&\le\sqrt{\delta}-(1-\mu_{i-1})+1\le\sqrt{\delta}+\mu_{i-1}\end{aligned}\tag{9}$$

where the last transition is due to Lemma 5.4 and the way we construct F_i, and the induction hypothesis for F_{i-1}.

Next, note that the event $F_i<F_{i-1}$ can be determined by looking on transcripts only until the point they cross F_{i-1}, thus it is a statistical test for distributions on F_{i-1}, and it must hold that for any $(u,v),(u',v')\in X_{i-1}\times Y_{i-1}$

$$\big|\Pr[F_i<F_{i-1}\mid(u,v),\mathsf{Ch}]-\Pr[F_i<F_{i-1}\mid(u',v'),\mathsf{Ch}]\big|\le\mathsf{SD}^{\mathsf{Ch}}_{F_{i-1}}((u,v),(u',v'))\le\mu_{i-1}$$

and specifically for $(x'_\circ, y_\circ)$ and for (x, y), which along with Eq. (9) lead to

$$\begin{aligned}\Pr[F_{i-1} \le F_i \mid x, y, \mathsf{Ch}] &\le \Pr[F_{i-1} \le F_i \mid x'_\circ, y_\circ, \mathsf{Ch}] + \mu_{i-1} \\ &\le 2(\sqrt{\delta} + \mu_{i-1}) + \mu_{i-1}.\end{aligned}$$

Which completes the proof of the Lemma. □

□

We are finally ready to prove Theorem 5.3. To this end, we extend the techniques used in the proof of Theorem 4.1 for perfectly-correct and private protocols, and apply them onto (ε, δ)-protocols. Informally, the proof starts with some restricted domain of depth d and constructs a sequence of frontiers according to Lemma 5.5. With high probability (i.e., for most of the transcripts), the frontiers are well-ordered, $F_1 \le \cdots \le F_d$. For every complete transcript $\mathbf{t}$ there are two consecutive frontiers $F_{i(\mathbf{t})}, F_{i(\mathbf{t})+1}$, such that the number of messages between these frontiers is at most $2|\mathbf{t}|/d$.

We define the channel Ch^* as in the perfect case. The channel makes no errors until the transcript reaches $F_{i(\mathbf{t})}$, then the channel changes the messages between $F_{i(\mathbf{t})}$ and $F_{i(\mathbf{t})+2}$ by simulating party A on input $x'_\circ$, for some $\boxplus$-minor that depends on the prefix of the transcript so far. We then run the protocol over Ch^* for the input $x_\circ, y_\circ$. As in the perfect case, B is led to believe that A has the input $x'_\circ$ and by the round where the channel stops changing messages, B has leaked too much information about his input.

Proof. (**Theorem 5.3**) Let $X_\circ \times Y_\circ$ be a restricted domain of depth $d \ge 3$ (the claim trivially holds for $d \le 2$), and let $\mathcal{F} = F_1, \ldots, F_d$ the sequence of frontiers defined by Theorem 5.5, reindexed such that F_d contains the complete transcripts and F_1 contains possibly shorter prefixes; let $F_0 = \emptyset$.

Let x, y be some input in $X_\circ \times Y_\circ$. For any given transcript $\mathbf{t} \in \tau^{\mathsf{Ch}^0}(x, y)$ we can define the number of messages between any two frontiers. Formally, for $1 \le j \le d$, let $\lambda_j(\mathbf{t})$ be the length of the longest prefix of $\mathbf{t}$ in F_j. Then for $0 \le j < k \le d$, we can define the *number of messages in* $\mathbf{t}$ *between* F_j *and* F_k, $\lambda_{j,k}(\mathbf{t})$, as

$$\lambda_{j,k}(\mathbf{t}) \coloneqq \lambda_k(\mathbf{t}) - \lambda_j(\mathbf{t}).$$

As $\sum_{j=1}^{d} \lambda_{j-1,j}(\mathbf{t}) = |\mathbf{t}|$, there exists i such that $\lambda_{i,i+2}(\mathbf{t}) \le 2|\mathbf{t}|/d$, and let $i(\mathbf{t})$ be the minimal such number i for $\mathbf{t}$, and let $\lambda(\mathbf{t}) \coloneqq \lambda_{i(\mathbf{t}),i(\mathbf{t})+2}(\mathbf{t})$. Note that for any two transcripts $\mathbf{t}_1, \mathbf{t}_2$ that share a prefix $\mathbf{u}$, then either $i(\mathbf{t}_1) = i(\mathbf{t}_2)$ or $i(\mathbf{t}_1), i(\mathbf{t}_2)$ are such that the longest prefix of $\mathbf{t}_1$ in $F_{i(\mathbf{t}_1)+2}$ is longer then $\mathbf{u}$ (i.e., $\mathbf{t}_1, \mathbf{t}_2$ diverge before reaching $F_{i(\mathbf{t}_1)+2}, F_{i(\mathbf{t}_2)+2}$).

Practically, this means that when the protocol runs on inputs x, y the event of reaching $i(\mathbf{t})$ is well defined and *causal.* Indicator of this event can be obtained from observing the partial transcript of $\mathbf{t}$ at up to the point when it reaches $F_{i(\mathbf{t})}$, that is, one does not need to know the complete $\mathbf{t}$ in advance but only the prefix $u \in F_{i(\mathbf{t})}$ — any transcript $\mathbf{t}'$ that has $\mathbf{u}$ as a prefix satisfies $i(\mathbf{t}') = i(\mathbf{t})$.

By $\mathbf{t}_{|F_i}$ we denote the partial $\mathbf{v} \in F_i$ such that $\mathbf{v}$ is a prefix of $\mathbf{t}$ until the point it crosses F_i. Next, we wish to fix only a single round i – since different i's mean different frontiers, it is somewhat meaningless to discuss different i's. For each $i \in \mathbb{N}$ define

$$w(i) = \sum_{\mathbf{t} \in \tau^{\mathsf{Ch}^0}(x,y) \text{ s.t. } i(\mathbf{t}) = i} \Pr[\mathbf{t}_{|F_{i(\mathbf{t})}} \mid x, y, \mathsf{Ch}^0]$$

and we fix $i_\circ$ to be the one that maximizes $w(i)$. Obviously, $w(i_\circ) > 1/2d$.

Define $(x_\circ, x'_\circ, y_\circ, y'_\circ)$ as the minor associated with $F_{i_\circ+1}$, and assume wlog it is a ⊞-minor. This allows us to define the channel Ch^* that fails π in the following way. Suppose the complete transcript in this instance is $\mathbf{t}$. The channel doesn't make noise until the prefix it observes reaches $F_{i_\circ}$. Again, this point is well defined without knowing the complete transcript $\mathbf{t}$.

Let $\mathbf{u} := \mathbf{t}_{|F_{i_\circ}}$ be the transcript observed so far. The channel changes any message sent to B by simulating π_A on input $x'_\circ$ given the prefix $\mathbf{u}$. Formally, for a given $\mathbf{u}$, the channel samples a randomness tape R_A that is consistent with the transcript so far, conditioned on the input $x'_\circ$. If no such randomness exists, the channel aborts its attack (i.e., behaves like Ch^0). It is easy to verify that the channel is able find a consistent randomness except with probability at most μ_i. The channel continues by changing A's sent messages to what A would have sent when using input $x'_\circ$ and the random tape sampled earlier. The channel delivers B's messages to A without any change. The channel stops the interference after $2|\mathbf{t}|/d$ rounds, after which it delivers all the messages intact.

To prove our theorem, we will show that if the channel does not abort the attack (that is, if there exists a consistent random tape R_A for A as described above), then $\mathsf{SD}^{\mathsf{Ch}^*}((x_\circ, y_\circ), (x_\circ, y'_\circ))$ is large. Thus, the channel violates B's privacy with probability $1 - \mu_i$. Intuitively, our proof proceeds as follows:

- show that B's 'view' in $\pi^{\mathsf{Ch}^*}(x_\circ, y_\circ)$ is statistically close to B's view in $\pi^{\mathsf{Ch}^0}(x'_\circ, y_\circ)$ up to a certain point in the protocol,
- show that B's 'view' in $\pi^{\mathsf{Ch}^*}(x_\circ, y'_\circ)$ is statistically close to B's view in $\pi^{\mathsf{Ch}^0}(x'_\circ, y'_\circ)$, up to the same point in the protocol,
- use the fact that $\pi^{\mathsf{Ch}^0}(x'_\circ, y_\circ)$ and $\pi^{\mathsf{Ch}^0}(x'_\circ, y'_\circ)$ are statistically far apart at that point to conclude the argument.

Before we begin with the formal details for the above steps, we remark that for our proof we only need to consider partial transcripts of particular lengths. To this end, we introduce the following notation. For a transcript $\mathbf{t}$, define,

$$\mathbf{t}_{|i_\circ+2} \triangleq \mathbf{t}\left[1 \ldots (\lambda_{i_\circ}(\mathbf{t}) + 2|\mathbf{t}|/d)\right].$$

As the Ch^* transcripts and Ch^0 transcripts have different structure (i.e., one has noise, while the other is noiseless), to relate B's view under Ch^* with his view in Ch^0, we introduce a transformation T on partial transcripts under Ch^*. We can view a partial transcript $\mathbf{u}$ as four components: $(u^s_A, u^r_A, u^s_B, u^r_B)$, where u^s_A and u^r_A are the messages sent and received by A, and u^s_B and u^r_B are the messages sent and received by B. Define the transformation T as follows,

$$\mathsf{T} : (u^s_A, u^r_A, u^s_B, u^r_B) \mapsto (u^r_B, u^s_B, u^s_B, u^r_B)_{|i_\circ+2}.$$

Further, let F_T be the set of all truncated transcripts as defined above. That is,

$$F_\mathsf{T} \triangleq \left\{ \mathbf{t}_{|i_\circ+2} \mid \mathbf{t} \text{ is a complete transcript.} \right\}.$$

As statistical distance can not be increased by applying T, we have,

$$\mathsf{SD}\left(\pi^{\mathsf{Ch}^*}(x_\circ, y_\circ), \pi^{\mathsf{Ch}^*}(x_\circ, y'_\circ)\right) \geq \mathsf{SD}\left(\mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ, y_\circ)), \mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ, y'_\circ))\right).$$

Thus, in order to show that B's privacy is compromised (according to Definition 5.2) and complete the proof of the theorem, we only need to show that $\mathsf{SD}\left(\mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ, y_\circ)), \mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ, y'_\circ))\right)$ is large enough (non-negligible). This we show in the following lemma.

Lemma 5.8. $\mathsf{SD}\left(\mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ, y_\circ)), \mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ, y'_\circ))\right) \geq 1/d - 3\mu_{i_\circ} - \mu_{i_\circ+2}$.

Proof. We begin by showing that the random variables $\mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ, y_\circ))$ and $\mathsf{T}(\pi^{\mathsf{Ch}^0}(x'_\circ, y_\circ))$ are close. Indeed,

$$\begin{aligned}
&\mathsf{SD}(\mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ, y_\circ)), \mathsf{T}(\pi^{\mathsf{Ch}^0}(x'_\circ, y_\circ))) \\
&\quad = \frac{1}{2} \sum_{\mathbf{u} \in F_{\mathsf{T}}} \left| \Pr[\mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ, y_\circ)) = \mathbf{u}] - \Pr[\mathsf{T}(\pi^{\mathsf{Ch}^0}(x'_\circ, y_\circ)) = \mathbf{u}] \right| \\
&\quad = \frac{1}{2} \sum_{\mathbf{u}_1 \in F_{i_\circ}} \sum_{\mathbf{u}_2} \Big| \Pr[\mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ, y_\circ; \mathbf{u}_1)) = \mathbf{u}_2] \Pr[\mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ, y_\circ))_{|F_{i_\circ}} = \mathbf{u}_1] \\
&\qquad\qquad - \Pr[\mathsf{T}(\pi^{\mathsf{Ch}^0}(x'_\circ, y_\circ; \mathbf{u}_1)) = \mathbf{u}_2] \Pr[\mathsf{T}(\pi^{\mathsf{Ch}^0}(x'_\circ, y_\circ))_{|F_{i_\circ}} = \mathbf{u}_1] \Big| . \qquad (10)
\end{aligned}$$

Now, by construction, we have that after frontier $F_{i_\circ}$, the channel plays as party A with input $x'_\circ$. Thus, for every $\mathbf{u}_1 \in F_{i_\circ}$ and any completion $\mathbf{u}_2$ of length $2|\mathbf{t}|/d$,

$$\Pr[\mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ, y_\circ; \mathbf{u}_1)) = \mathbf{u}_2] = \Pr[\mathsf{T}(\pi^{\mathsf{Ch}^0}(x'_\circ, y_\circ; \mathbf{u}_1)) = \mathbf{u}_2].$$

Furthermore, again by construction, the channel does not change any messages before frontier $F_{i_\circ}$. Thus,

$$\Pr[\mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ, y_\circ))_{|F_{i_\circ}} = \mathbf{u}_1] = \Pr[\mathsf{T}(\pi^{\mathsf{Ch}^0}(x_\circ, y_\circ))_{|F_{i_\circ}} = \mathbf{u}_1].$$

Therefore, continuing with Eq. (10),

$$\begin{aligned}
&\mathsf{SD}(\mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ, y_\circ)), \mathsf{T}(\pi^{\mathsf{Ch}^0}(x'_\circ, y_\circ))) \\
&\quad = \sum_{\mathbf{u}_1 \in F_{i_\circ}} \left| \Pr[\mathsf{T}(\pi^{\mathsf{Ch}^0}(x_\circ, y_\circ)) = \mathbf{u}_1] - \Pr[\mathsf{T}(\pi^{\mathsf{Ch}^0}(x'_\circ, y_\circ)) = \mathbf{u}_1] \right| \times \sum_{\mathbf{u}_2} \Pr[\mathsf{T}(\pi^{\mathsf{Ch}^0}(x'_\circ, y_\circ) = \mathbf{u}_2 \mid \mathbf{u}_1] \\
&\quad = \mathsf{SD}\left(\mathsf{T}(\pi^{\mathsf{Ch}^0}(x_\circ, y_\circ))_{|F_{i_\circ}}, \mathsf{T}(\pi^{\mathsf{Ch}^0}(x'_\circ, y_\circ))_{|F_{i_\circ}} \right), \qquad (11)
\end{aligned}$$

where the last equality follows from the fact that the second summation in the first equality sums to 1.

For a "noiseless" transcript $\mathbf{u}$, observe that applying transformation T only shortens the length, and does not swap any components of the transcript; that is, for noiseless $\mathbf{u}$, $\mathsf{T}(\mathbf{u})_{|F_{i_\circ}} = \mathbf{u}_{|F_{i_\circ}}$. Thus, we have,

$$\mathsf{SD}(\mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ, y_\circ))_{|F_{i_\circ}}, \mathsf{T}(\pi^{\mathsf{Ch}^0}(x'_\circ, y_\circ))_{|F_{i_\circ}}) = \mathsf{SD}\left(\pi^{\mathsf{Ch}^0}(x_\circ, y_\circ)_{|F_{i_\circ}}, \pi^{\mathsf{Ch}^0}(x'_\circ, y_\circ)_{|F_{i_\circ}} \right)$$

which with Eq. (11) and Lemma 5.4 gives

$$\mathsf{SD}(\mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ, y_\circ)), \mathsf{T}(\pi^{\mathsf{Ch}^0}(x'_\circ, y_\circ)) \leq \mu_{i_\circ}.$$

By a similar argument, we get,

$$\mathsf{SD}(\mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ, y'_\circ)), \mathsf{T}(\pi^{\mathsf{Ch}^0}(x'_\circ, y'_\circ))) \leq \mu_{i_\circ}.$$

Finally, we have,

$$\begin{aligned}
&\mathsf{SD}(\mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ,y_\circ)),\mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ,y'_\circ))) \\
&\quad\geq \mathsf{SD}(\mathsf{T}(\pi^{\mathsf{Ch}^0}(x'_\circ,y_\circ)),\mathsf{T}(\pi^{\mathsf{Ch}^0}(x'_\circ,y'_\circ))) \\
&\qquad - \mathsf{SD}(\mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ,y_\circ)),\mathsf{T}(\pi^{\mathsf{Ch}^0}(x'_\circ,y_\circ))) - \mathsf{SD}(\mathsf{T}(\pi^{\mathsf{Ch}^*}(x_\circ,y'_\circ)),\mathsf{T}(\pi^{\mathsf{Ch}^0}(x'_\circ,y'_\circ))) \qquad (12)\\
&\quad\geq \mathsf{SD}^{\mathsf{Ch}^0}_{F_\mathsf{T}\cap F_{i_\circ+2}}((x'_\circ,y_\circ),(x'_\circ,y'_\circ)) - 2\mu_{i_\circ} \qquad (13)\\
&\quad\geq 1/d - 3\mu_{i_\circ} - \mu_{i_\circ+2}. \qquad (14)
\end{aligned}$$

In the above, Eq. (12) is due a triangle inequality. Eq. (13) holds because the transformation T only truncates the noiseless transcripts without swapping the components. For Eq. (14), we have from Lemma 5.4 that

$$\mathsf{SD}^{\mathsf{Ch}^0}_{F_{i_\circ+2}}((x'_\circ,y_\circ),(x'_\circ,y'_\circ)) \geq 1-\mu_{i_\circ+2}.$$

However, note that the above statistical distance is over the frontier $F_{i_\circ+2}$, while the statistical distance in Eq. (13) is over the set $F_\mathsf{T}\cap F_{i_\circ+2}$. By construction, $F_\mathsf{T}\cap F_{i_\circ+2}$ contains at least $1/d$ fraction of the probability mass (of transcripts in the execution of $\pi^{\mathsf{Ch}^0}(x_\circ,y_\circ)$), that is,

$$\sum_{\mathbf{u}\in F_\mathsf{T}\cap F_{i_\circ+2}} \Pr[\mathbf{u} \text{ is a prefix of } \pi^{\mathsf{Ch}^0}(x_\circ,y_\circ)] \geq 1/d.$$

By Theorem 5.5, the random variables $\pi^{\mathsf{Ch}^0}(x_\circ,y_\circ)$ and $\pi^{\mathsf{Ch}^0}(x'_\circ,y_\circ)$ are statistically close up to frontier $F_{i_\circ}$ (recall, that $(x_\circ,x'_\circ,y_\circ,y_\circ)$ is the minor associated with $F_{i_\circ+1}$). Thus,

$$\sum_{\mathbf{u}\in F_\mathsf{T}\cap F_{i_\circ+2}} \Pr[\mathbf{u} \text{ is a prefix of } \pi^{\mathsf{Ch}^0}(x'_\circ,y_\circ)] \geq 1/d-\mu_{i_\circ}.$$

Therefore, we have,

$$\mathsf{SD}^{\mathsf{Ch}^0}_{F_\mathsf{T}\cap F_{i_\circ+2}}((x'_\circ,y_\circ),(x'_\circ,y'_\circ)) \geq 1/d-\mu_{i_\circ}-\mu_{i_\circ+2},$$

and Eq. (14) follows. □

This completes the proof of the theorem. □

Acknowledgments

We would like to thank Hemanta Maji and Sanjam Garg for useful discussions. Research supported in part from a DARPA/ONR PROCEED award, NSF grants 1228984, 1136174, 1118096, 1065276, 0916574 and 0830803, 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 U.S. Office of Naval Research under Contract N00014-11-1-0389. 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.

References

[Bea91] D. Beaver. Perfect privacy for two-party protocols. *Proceedings of DIMACS Workshop on Distributed Computing and Cryptography*, vol. 2, pp. 65–77. 1991.

[BK12] Z. Brakerski and Y. T. Kalai. Efficient interactive coding against adversarial noise. *Foundations of Computer Science (FOCS), IEEE Annual Symposium on*, pp. 160–166, 2012.

[BN13] Z. Brakerski and M. Naor. Fast algorithms for interactive coding. *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '13, pp. 443–456. 2013.

[BR11] M. Braverman and A. Rao. Towards coding for maximum errors in interactive communication. *Proceedings of the 43rd annual ACM symposium on Theory of computing*, STOC '11, pp. 159–166. ACM, New York, NY, USA, 2011.

[CK91] B. Chor and E. Kushilevitz. A zero-one law for boolean privacy. *SIAM Journal on Discrete Mathematics*, 4(1):36–47, 1991.

[CPT13] K.-M. Chung, R. Pass, and S. Telang. Interactive coding, revisited. Cryptology ePrint Archive, Report 2013/160, 2013. http://eprint.iacr.org/2013/160.

[FGOS12] M. Franklin, R. Gelles, R. Ostrovsky, and L. J. Schulman. Optimal coding for streaming authentication and interactive communication. *Electronic Colloquium on Computational Complexity (ECCC)*, 2012. TR12-104. To appear in CRYPTO 2013.

[GMS11] R. Gelles, A. Moitra, and A. Sahai. Efficient and explicit coding for interactive communication. *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pp. 768–777. 2011.

[GMW87] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game. *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, STOC '87, pp. 218–229. ACM, New York, NY, USA, 1987.

[IKO$^+$11] Y. Ishai, E. Kushilevitz, R. Ostrovsky, M. Prabhakaran, A. Sahai, and J. Wullschleger. Constant-rate oblivious transfer from noisy channels. P. Rogaway, ed., *Advances in Cryptology – CRYPTO 2011*, *LNCS*, vol. 6841, pp. 667–684. Springer Berlin Heidelberg, 2011.

[Kus89] E. Kushilevitz. Privacy and communication complexity. *Proceedings of the 30th Annual Symposium on Foundations of Computer Science*, FOCS '89, pp. 416–421. IEEE Computer Society, Washington, DC, USA, 1989.

[Kus92] E. Kushilevitz. Privacy and communication complexity. *SIAM Journal on Discrete Mathematics*, 5(2):273–284, 1992.

[MPR09] H. K. Maji, M. Prabhakaran, and M. Rosulek. Complexity of multi-party computation problems: The case of 2-party symmetric secure function evaluation. O. Reingold, ed., *Theory of Cryptography*, *LNCS*, vol. 5444, pp. 256–273. Springer Berlin Heidelberg, 2009.

[RS94] S. Rajagopalan and L. Schulman. A coding theorem for distributed computation. *STOC '94: Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, pp. 790–799. ACM, New York, NY, USA, 1994.

[Sch93] L. J. Schulman. Deterministic coding for interactive communication. *STOC '93: Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pp. 747–756. ACM, New York, NY, USA, 1993.

[Sch96] L. J. Schulman. Coding for interactive communication. *IEEE Transactions on Information Theory*, 42(6):1745–1756, 1996.

[Sha48] C. E. Shannon. A mathematical theory of communication. *ACM SIGMOBILE Mobile Computing and Communications Review*, 5(1):3–55, 2001. Originally appeared in *Bell System Tech. J.* 27:379–423, 623–656, 1948.

[Vic61] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.