

HOW TO COMPRESS INTERACTIVE COMMUNICATION*

BOAZ BARAK[†], MARK BRAVERMAN[‡], XI CHEN[§], AND ANUP RAO[¶]

Abstract. We describe new ways to simulate two-party communication protocols to get protocols with potentially less communication. We show that every communication protocol that communicates C bits and reveals I bits of information about the inputs to the participating parties can be simulated by a new protocol involving at most $\tilde{O}(\sqrt{CI})$ bits of communication. If the protocol reveals I bits of information about the inputs to an observer that watches the communication in the protocol, we show how to carry out the simulation with $\tilde{O}(I)$ bits of communication. These results lead to a direct sum theorem for randomized communication complexity. Ignoring polylogarithmic factors, we show that for worst-case computation, computing n copies of a function requires \sqrt{n} times the communication required for computing one copy of the function. For average case complexity, given any distribution μ on inputs, computing n copies of the function on n inputs sampled independently according to μ requires \sqrt{n} times the communication for computing one copy. If μ is a product distribution, computing n copies on n independent inputs sampled according to μ requires n times the communication required for computing the function. We also study the complexity of computing the sum (or parity) of n evaluations of f , and obtain results analogous to those above. Our results give the first compression schemes for general randomized protocols and the first direct sum results in the general setting of randomized and distributional communication complexity, without requiring bound on the number of rounds in the protocol or that the distribution of inputs is independent.

Key words. interactive information, compression, communication complexity, direct sum

AMS subject classification. 68Q17

DOI. 10.1137/100811969

1. Introduction. In this work, we address two questions: (1) Can we compress the communication of an interactive protocol so it is close to the information conveyed between the parties? (2) Is it harder to compute a function on n independent inputs than to compute it on a single input? In the context of communication complexity, the two questions are related, and our answer to the former will yield an answer to the latter.

Techniques for message compression, first considered by Shannon [31], have had a big impact on computer science, especially with the rise of the Internet and data intensive applications. Today we know how to encode messages so that their length is essentially the same as the amount of information that they carry (see, for example, the textbook [7]). Can we get a similar savings in an interactive setting? A first attempt might be to simply compress each message of the interaction in turn. However, this compression involves at least 1 bit of communication for every message of the

*Received by the editors October 18, 2012; accepted for publication (in revised form) March 4, 2013; published electronically June 25, 2013.

<http://www.siam.org/journals/sicomp/42-3/81196.html>

[†]Microsoft Research New England, Cambridge, MA (b@boazbarak.org).

[‡]Department of Computer Science, Princeton University, Princeton, NJ 08544 (mbraverm@cs.princeton.edu). This author's research was supported in part by a Clay Liffoff Fellowship, an NSERC Discovery Grant, Alfred P. Sloan Fellowship, an NSF CAREER award, and a Turing Centenary Fellowship. Part of this work was done while MB was a postdoctoral researcher with Microsoft Research New England.

[§]Department of Computer Science, Columbia University, New York, NY 10027 (xichen@cs.columbia.edu). This author's research was supported by NSF Grants CCF-0832797 and DMS-0635607.

[¶]Center for Computational Intractability, Princeton University, Princeton, NJ 08544 (anuprao@cs.washington.edu). This author's work was supported by NSF Grant CCF 0832797.

interaction, which can be much larger than the total information conveyed between the parties. In this paper, we show how to compress interactive communication protocols in a way that is independent of the number of rounds of communication, and in some settings, give compressed protocols with communication that has an almost linear dependence on the information conveyed in the original protocol.

The second question is one of the most basic questions of theoretical computer science, called the *direct sum* question, and is closely related to the *direct product* question. A direct product theorem in a particular computational model asserts that the probability of success of performing n independent computational tasks decreases in n . Famous examples of such theorems include Yao's XOR lemma [36] and Raz's parallel repetition theorem [28]. In the context of communication complexity, Shaltiel [30] gave a direct product theorem for the discrepancy of a function, but it remains open to give such a theorem for the success probability of communication tasks. A direct sum theorem asserts that the amount of resources needed to perform n independent tasks grows with n . While the direct sum question for general models such as Boolean circuits has a long history (cf [34, 26, 10]), no general results are known, and indeed they cannot be achieved by the standard reductions used in complexity theory, as a black-box reduction mapping a circuit C performing n tasks into a circuit C' performing a single task will necessarily make C' larger than C , rather than making it smaller. Indeed it is known that at least the most straightforward/optimistic formulation of a direct sum theorem for Boolean circuits is *false*.¹

Nevertheless, direct sum theorems are known to hold in other computational models. For example, an optimal direct sum theorem is easy to prove for decision tree depth. A more interesting model is *communication complexity*, where this question was first raised by Karchmer, Raz, and Wigderson [19] who conjectured a certain direct sum result for deterministic communication complexity of relations, and showed that it would imply that $\mathbf{P} \not\subseteq \mathbf{NC}^1$. Feder et al. [8] gave a direct sum theorem for nondeterministic communication complexity, and deduced from it a somewhat weaker result for deterministic communication complexity—if a single copy of a function f requires C bits of communications, then n copies require $\Omega(\sqrt{Cn})$ bits. Feder et al. also considered the direct sum question for *randomized* communication complexity (see also Open Problem 4.6 in [23]) and showed that the dependence of the communication on the error of the protocol for many copies can be better than that obtained by the naive protocol for many copies.

1.1. External and internal information cost. We follow the *information complexity* approach to proving direct sum theorems. This approach was explicitly introduced by [6] and used by several other works. (See section 1.2 below for more about related works.) The approach is based on a measure, which we call the *external information cost*,² which roughly speaking measures the minimum number of bits of information that an external observer will learn about the inputs of two parties engaged in any protocol to compute a particular function. The formal definition is as follows.

¹The example comes from fast matrix multiplication. By a counting argument, there exists an $n \times n$ matrix A over $\text{GF}(2)$ such that the map $x \mapsto Ax$ requires a circuit of $\Omega(n^2/\log n)$ size. But the map $(x_1, \dots, x_n) \mapsto (Ax_1, \dots, Ax_n)$ is just the product of the matrices A and X (whose columns are x_1, \dots, x_n) and hence can be carried out by a circuit of $O(n^{2.38}) \ll n \cdot (n^2/\log n)$. See Shaltiel's paper [30] for more on this question.

²This notion was called simply *information cost* or *information complexity* by prior works. We use the name “external” to distinguish it from the notion of *internal* information cost defined below, which was not explicitly named by prior works.

DEFINITION 1.1 (external information cost). *Given a distribution μ on inputs X, Y , and protocol π , denoting by $\pi(X, Y)$ the public randomness and messages exchanged during the protocol, the external information cost $\text{IC}_\mu^\circ(\pi)$ is defined to be the mutual information between the inputs and $\pi(X, Y)$:*

$$\text{IC}_\mu^\circ(\pi) \stackrel{\text{def}}{=} I(XY; \pi(X, Y)).$$

For a function f , the external information cost of f with respect to μ is defined to be the infimum of $\text{IC}_\mu^\circ(\pi)$ over all protocols π that compute f on inputs from π with probability larger than $2/3$.

Clearly, the external information cost is always smaller than the communication complexity, and if the inputs to the parties are independent of each other (i.e., X is independent of Y), an optimal direct sum theorem can be proved for this measure of complexity. That is, if the external information cost of f is c and we choose independently n pairs of inputs $(X_1, Y_1), \dots, (X_n, Y_n)$, then any protocol for computing the tuple $f(X_1, Y_1) \cdots f(X_n, Y_n)$ must reveal cn bits of information about these input tuples. Thus the problem of proving direct sum theorems for independent inputs reduces to the problem of simulating a protocol τ with small external information cost with a protocol ρ that has small communication. That is, the direct sum question reduces to the problem of *protocol compression*. Previous works (see section 1.2) obtained restricted direct sum theorems by compressing every message of the protocol individually. Our stronger method of compression allows us to get new direct sum theorems that are independent of the number of rounds of communication.

Internal information cost. In this work we use also a second measure of the information complexity of a communication protocol, which we call the *internal information cost* of a protocol. This is the information that the *parties* in the protocol learn about each other's inputs through observing the messages and public randomness of the protocol. Formally, the definition is as follows.

DEFINITION 1.2 (internal information cost). *Given a distribution μ on inputs X, Y , and protocol π , denoting by $\pi(X, Y)$ the public randomness and messages exchanged during the protocol, the internal information cost $\text{IC}_\mu^i(\pi)$ is defined to be*

$$\text{IC}_\mu^i(\pi) \stackrel{\text{def}}{=} I(X; \pi(X, Y)|Y) + I(Y; \pi(X, Y)|X).$$

Since each party knows her own input, the protocol can only reveal less additional information to her than to an independent observer. Indeed, it can be shown that the internal information cost is never larger than the external information cost and that the two are equal if the inputs X, Y are independent of each other. However, in the case that the inputs are dependent, the internal information cost may be significantly smaller—for example, if μ is a distribution where $X = Y$ always, then the internal information cost is always 0, though the external information cost may be arbitrarily large. It is also easy to check that if π is deterministic, then the internal information cost is simply the sum of the entropies $\text{IC}_\mu^i(\pi) = H(\pi(X, Y)|Y) + H(\pi(X, Y)|X)$, which is the same as $H(\pi(X, Y))$ if X, Y are independent.

The notion of internal information cost was used implicitly by Bar-Yossef et al. [3], and a direct sum theorem for this notion (using the techniques originating from Razborov [27] and Raz [28]) is implicit in their work. In contrast to the analogous result for *external* information cost, the direct sum theorem for *internal* information

cost holds whether or not the inputs to the parties are independent of each other. This means that a method to compress a protocol to have communication related to its *internal* information cost can be used to obtain a direct sum theorem even in the case that the inputs to both parties are not independent. This is crucial for our application to a direct sum theorem for (worst-case) randomized communication complexity.

Our most important contributions are two new *protocol compression* methods that reduce the communication of protocols in terms of their information costs. We give two methods to compress protocols. Our first method shows that one can always simulate a protocol of internal information cost I and communication complexity C using an expected number of $\tilde{O}(\sqrt{IC})$ communication bits. The second method shows how to simulate a protocol of external information cost I with $\tilde{O}(I)$ communication. Note that in both cases the communication complexity of the simulation is *independent* of the number of rounds. Indeed, these are the first compression schemes that do true *protocol* compression, as opposed to compressing each round at a time, and are the first results that succeed for randomized protocols even when the inputs are not independent.

As a result, we obtain the first nontrivial direct sum theorem for randomized communication complexity. Loosely speaking, letting f^n be the function that outputs the concatenation of n invocations of f on independent inputs, and letting f^{+n} be the function that outputs the XOR of n such invocations, we show that (a) the randomized communication complexity of both f^n and f^{+n} is up to logarithmic factors \sqrt{n} times the communication complexity of f , and (b) the distributional complexity of both f^n and f^{+n} over the distribution μ^n , where μ is a product distribution over individual input pairs, is n times the distributional complexity of f .³

1.2. Related works. There has been a vast amount of work on the question of direct sum and direct product theorems for various computation and communication models. We now discuss some of the works that are most relevant to our own, focusing on the direct sum question for randomized and distributional communication complexity in the classical (i.e., nonquantum) setting. Chakrabarti et al. [6] explicitly introduced the notion of (external) information cost, variants of which were implicitly used by other works as well [2, 1, 29]. Chakrabarti et al. focused on the case of the uniform distribution on inputs, which in particular means the inputs to both parties are independent. They showed for this case a direct sum theorem for the external information cost, as well as a (slightly restricted) compression result in the simultaneous messages model (in which the protocol consists of each party sending one message that depends only on its input, and not on the other party's message). They thus derived some direct sum theorems for this model. Bar-Yossef et al. [3] extended the definition of external information cost to arbitrary distributions. They also defined a notion which they called “conditional information cost” which is different but related to our notion of internal information cost, and gave a direct sum theorem for this notion, whose proof we adapt to our direct sum for internal information cost. They then combined this direct sum with information cost lower bounds for concrete simple functions such as AND to obtain improved lower bounds for the set disjointness and L_p approximation problems. In a sequence of works, Jain, Radhakrishnan, and Sen [13, 14] and Harsha et al. [11] improved the message compression result of [6],

³In both (a) and (b), there is a loss of a constant additive factor in the actual statement of the result for f^{+n} . This accounts for the fact that if, say, f is the XOR function itself, then clearly there is no direct sum theorem. See Remark 2.10.

showing that one can compress each message of a protocol to roughly its contribution to the external information cost plus some constant overhead. This results in an essentially optimal direct sum theorem for distributional communication complexity for protocols with a bounded number of rounds with respect to any product distribution over the inputs.

Related problems have also been considered in the information theory literature. Stated in our language, the now classical Slepian–Wolf theorem [33] (see also [35]) is about compressing one message deterministic protocols according to the *internal* information cost of the protocol. In this case, the internal information cost is simply $I(X; M|Y) = H(M|Y)$, where here X, Y are the inputs and M is the lone message transmitted by the first party. The theorem shows how to transmit n independent messages M_1, \dots, M_n to a receiver that knows the corresponding inputs Y_1, \dots, Y_n , using amortized (one way) communication that is close to the internal information cost. In contrast, the results in our work are about compressing a single interaction, which is a harder problem since we cannot enjoy the “economies of scale” of amortized analysis. A caveat is that our compression schemes yield *interactive* communication protocols, as opposed to one way communication.

1.3. Subsequent work. Subsequent papers further developed the notion of internal information complexity and its application to direct sum and direct product theorems. We mention some of the works here. In [5] it has been shown that in the limit, the per-copy randomized communication complexity of solving n copies of a problem is equal to its internal information complexity. Thus progress towards stronger direct sum theorems is tightly linked to progress in obtaining better compression schemes. Unfortunately, no new general compression schemes have emerged. Still, it has been shown that a large class of communication complexity lower bound techniques also yield information complexity lower bounds [20]—thus yielding the direct sum theorem for the randomized communication complexity of a large class of problems. The notion of information complexity has also been extended to the setting without a prior distribution. In the prior-free case, the information complexity of a problem corresponds to its amortized randomized (worst-case) communication complexity.

Closely related to the direct sum problem is the *direct product* problem. While direct sum theorems aim to show that performing n copies of a task in parallel with the same success probability as the original task requires n times as much communication, direct product theorems assert that expending a substantially lower amount of communication will result in an exponentially smaller success probability. Thus direct product theorems are stronger than direct sum theorems. Strong direct product theorems are known in several special cases: for example, when f is the disjointness function [21], or f is known to have small discrepancy [30, 24, 32], or have a smooth rectangle bound [16]. Since we know that the amortized communication complexity of a function is equal to its information complexity, the best direct product theorem one can hope for is that computing n copies of f using communication substantially lower than n times the information complexity of f will only succeed with an exponentially small probability. Unfortunately, such a general result remains elusive. However, a recent result [4] shows that the direct sum theorems in the present paper can in fact be strengthened to direct product theorems.

2. Our results. We give two new *protocol compression* algorithms, that take a protocol π whose information cost is small and transform it into a protocol τ of

small *communication complexity*.⁴ Below we denote the communication complexity of a protocol τ by $\text{CC}(\tau)$.

THEOREM 2.1. *For every distribution μ , every protocol π , and every $\epsilon > 0$, there exist functions π_x, π_y , and a protocol τ such that $|\pi_x(X, \tau(X, Y)) - \pi(X, Y)| < \epsilon$, $\Pr[\pi_x(X, \tau(X, Y)) \neq \pi_y(Y, \tau(X, Y))] < \epsilon$, and*

$$\text{CC}(\tau) \leq O\left(\sqrt{\text{CC}(\pi) \cdot \text{IC}_{\mu}^i(\pi)} \frac{\log(\text{CC}(\pi)/\epsilon)}{\epsilon}\right).$$

If the players want to obtain the results of running the protocol π , they can run τ instead and then use the functions π_x, π_y to reconstruct the effects of running π . The condition $|\pi_x(X, \tau(X, Y)) - \pi(X, Y)| < \epsilon$ ensures that the transcript of τ specifies a unique leaf in the protocol tree for π in such a way that this leaf is ϵ -close in statistical distance to the leaf sampled by π . The condition that $\Pr[\pi_x(X, \tau(X, Y)) \neq \pi_y(Y, \tau(X, Y))] < \epsilon$ guarantees that with high probability both players achieve a consensus on what the sampled leaf was. Thus, the triple τ, π_x, π_y specifies a new protocol that is a compression of π .

Our second result gives a simulation for protocols with small external information cost.

THEOREM 2.2. *For every distribution μ , every protocol π , and every $\alpha > 0$, there exist functions π_x, π_y , and a protocol τ such that $|\pi_x(X, \tau(X, Y)) - \pi(X, Y)| < \alpha$, $\Pr[\pi_x(X, \tau(X, Y)) \neq \pi_y(Y, \tau(X, Y))] < \alpha$, and*

$$\text{CC}(\tau) \leq O\left(\text{IC}_{\mu}^o(\pi) \frac{\log(\text{CC}(\pi)/\alpha)}{\alpha^2}\right).$$

Our results can be viewed as a generalization of the traditional notion of string compression, a notion that applies only to the restricted case of deterministic one way protocols. In the above theorems, our compressed protocols may use public randomness that can be large (though still bounded in terms of the communication complexity of the original protocol). However, we note that by the results of Newman [25], any protocol that achieves some functionality can be converted into another protocol that achieves the same functionality and uses few public random bits. Thus our compression schemes are useful even when public randomness is expensive.

2.1. Direct sum theorems. Given a function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{Z}$, we define the function $f^n : \mathcal{X}^n \times \mathcal{Y}^n \rightarrow \mathcal{Z}^n$ to be the concatenation of the evaluations:

$$f^n(x_1, \dots, x_n, y_1, \dots, y_n) \stackrel{\text{def}}{=} (f(x_1, y_1), f(x_2, y_2), \dots, f(x_n, y_n)).$$

Denote by $R_{\rho}(f)$ the communication complexity of the best randomized public coin protocol for computing f that errs with probability at most ρ . In this paper we show the following theorem.

THEOREM 2.3 (direct sum for randomized communication complexity). *For every $\alpha > 0$,*

$$R_{\rho}(f^n) \cdot \log(R_{\rho}(f^n)/\alpha) \geq \Omega(R_{\rho+\alpha}(f)\alpha\sqrt{n}).$$

⁴ We note that this is in the communication complexity model, and hence these algorithms are not necessarily computationally efficient. Even for single message compression, there are distributions with small entropy that cannot be efficiently compressed (e.g., pseudorandom distributions).

Theorem 2.3 is obtained using Yao’s min-max principle from an analogous theorem for *distributional* communication complexity. For a distribution μ on the inputs $\mathcal{X} \times \mathcal{Y}$, we write $D_\rho^\mu(f)$ to denote the communication complexity of the best protocol (randomized or deterministic) that computes f with probability of error at most ρ when the inputs are sampled according to μ . We write μ^n to denote the distribution on n inputs, where each is sampled according to μ independently.

We first state the direct sum theorem for information content that is implicit in the work of [3].

THEOREM 2.4. *For every μ, f, ρ there exists a protocol τ computing f on inputs drawn from μ with probability of error at most ρ and communication at most $D_\rho^{\mu^n}(f^n)$ such that $IC_\mu(\tau) \leq \frac{2D_\rho^{\mu^n}(f^n)}{n}$.*

Compressing protocol τ above using Theorem 2.1 reduces the communication of this protocol to $\tilde{O}(\sqrt{IC_\mu(\tau)D_\rho^{\mu^n}(f^n)}) = \tilde{O}(D_\rho^{\mu^n}(f^n)/\sqrt{n})$. Formally, we prove the following theorem.

THEOREM 2.5 (direct sum for distributional communication complexity). *For every $\alpha > 0$,*

$$D_\rho^{\mu^n}(f^n) \cdot \log \left(D_\rho^{\mu^n}(f^n)/\alpha \right) \geq \Omega \left(D_{\rho+\alpha}^\mu(f) \alpha \sqrt{n} \right).$$

The communication complexity bound of Theorem 2.5 only grows as the square root of the number of repetitions. However, in the case that the distribution on inputs is a product distribution, we use our stronger compression (Theorem 2.2) to obtain a direct sum theorem that is optimal up to logarithmic factor.

THEOREM 2.6 (direct sum for product distributions). *If μ is a product distribution, then for every $\alpha > 0$*

$$D_\rho^{\mu^n}(f^n) \cdot \log \left(D_\rho^{\mu^n}(f^n)/\alpha \right) \geq \Omega \left(D_{\rho+\alpha}^\mu(f) \alpha^2 n \right).$$

2.1.1. XOR lemmas for communication complexity. When n is very large in terms of the other quantities, the above theorems can be superseded by trivial arguments, since f^n must require at least n bits of communication just to describe the output. Our next set of theorems show that almost the same bounds apply to the complexity of the XOR (or more generally sum modulo K) of n copies of f , where the trivial arguments do not hold. Assume that the output of the function f is in the group \mathbb{Z}_K for some integer K , and define

$$f^{+n}(x_1, \dots, x_n, y_1, \dots, y_n) \stackrel{\text{def}}{=} \sum_{i=1}^n f(x_i, y_i).$$

We have the following results for the complexity of f^{+n} .

THEOREM 2.7 (XOR lemma for randomized communication complexity). *For every $\alpha > 0$,*

$$R_\rho(f^{+n}) \cdot \log \left(R_\rho(f^{+n})/\alpha \right) \geq \Omega \left((R_{\rho+\alpha}(f) - 2 \log K) \alpha \sqrt{n} \right).$$

THEOREM 2.8 (XOR lemma for distributional communication complexity). *For every $\alpha > 0$,*

$$D_\rho^{\mu^n}(f^{+n}) \cdot \log \left(D_\rho^{\mu^n}(f^{+n})/\alpha \right) \geq \Omega \left((D_{\rho+\alpha}^\mu(f) - 2 \log K) \alpha \sqrt{n} \right).$$

THEOREM 2.9 (XOR lemma for product distributions). *If μ is a product distribution, then for every $\alpha > 0$,*

$$D_{\rho}^{\mu^n}(f^{+n}) \cdot \log \left(D_{\rho}^{\mu^n}(f^{+n}) / \alpha \right) \geq \Omega \left((D_{\rho+\alpha}^{\mu}(f) - 2 \log K) \alpha^2 n \right).$$

Remark 2.10. If $f : \mathbb{Z}_K \times \mathbb{Z}_K \rightarrow \mathbb{Z}_K$ is itself the sum function, then the communication complexity of f^{+n} does not grow at all, since there is a simple protocol to compute $\sum_i (x_i + y_i) = \sum_i x_i + \sum_j y_j$ using $2 \log K$ bits. This suggests that some kind of additive loss (like the $2 \log K$ term above) is necessary in the above theorems.

3. Our techniques. We now give an informal overview of our compression algorithms. Our direct sum results are obtained in section 5 by combining these with the direct sum for information content proven in section 6. Full descriptions of the compression algorithms are given in section 7 (for the general case) and section 8 (for the product distribution case).

The goal of our compression algorithms is to take a protocol that uses large amounts of communication and conveys little information, and convert it into a protocol that makes better use of the communication to achieve better communication complexity. (Such algorithms need not be necessarily computationally efficient; see footnote 4.)

Note that generic message compression can be fit into this context by considering a deterministic one way protocol, where player X needs to send a message to player Y . In this classical setting it is well known that protocol compression (i.e., simple data compression) can be achieved. In principle, one could try to apply round-by-round message compression to compress entire protocols. This approach suffers from the following fatal flaw: individual messages may (and are even likely) to contain $\ll 1$ bit of information. The communication cost of ≥ 1 bit per round would thus be \gg information content of the round. Thus any attempt to implement the compression on a round-by-round basis, as opposed to an entire-protocol basis may work when the number of rounds is bounded, but is doomed to fail in general.

An instructive example on conveying a subconstant amount of information that we will use later in this exposition is the following. Suppose that player X gets n independent random bits x_1, \dots, x_n and Y has no information about them. X then computes the majority $m = \text{MAJ}(x_1, \dots, x_n)$ and sends it to Y . With a perfectly random prior, the bit m is perfectly balanced, and thus in total X conveys one bit of information to Y . Suppose that in the protocol Y only really cared about the value of x_5 . How much information did X convey about the input x_5 ? By symmetry and independence of the inputs, X conveys $1/n$ bits of information about x_5 . After the bit m (suppose $m = 1$) is received by Y , her estimate of $P[x_5 = 1]$ changes from $1/2$ to $1/2 + \Theta(1/\sqrt{n})$. The fact that changing the probability from $1/2$ to $1/2 + \epsilon$ only costs ϵ^2 bits of information is the cause for the suboptimality of our general compression algorithm.

There are several challenges that need to be overcome to compress an arbitrary protocol. An interesting case to consider is a protocol where the players alternate sending each other messages, and each transmitted message is just a bit with information content $\epsilon \ll 1$. In this case, we cannot afford to even transmit one bit to simulate each of the messages, since that would incur an overhead of $1/\epsilon$, which would be too large for our application. This barrier was one of the big stumbling blocks for earlier works, which is why their results applied only when the number of rounds in the protocols was forced to be small.




FIG. 1. An illustration of the protocol tree for π . The round nodes are owned by X and the square nodes are owned by Y . On each edge the “correct” probability is indicated. The “approximate” probability that is estimated by the player who does not own the node is shown in the brackets.

We give two simulation protocols to solve this problem. The first solution works for all distributions, achieving suboptimal parameters, while the second works only for product input distributions and achieves optimal parameters up to polylogarithmic factors. In both solutions, the players simulate the original protocol π using shared randomness. The intuition is that if a message contains a small amount of information, then we do not need to communicate it, and can sample it using shared randomness instead.

It will be convenient to think of a protocol in terms of its protocol tree, after fixing the shared randomness (there may still be private randomness that is not fixed). This is a binary tree where every node v belongs to one of the parties in the protocol, and specifies the probability of sending 1 or 0 as the next bit. We then define the tree of probabilities illustrated in Figure 1 as follows. For each node v_x of the protocol tree that is owned by the player X (i.e., it is his turn to speak), player X knows the “correct” probabilities $O_{v_x,x}(0)$ and $O_{v_x,x}(1)$ of the bit that she is about to send. Player Y does not know these probabilities, but she has estimates $O_{v_x,y}(0)$ and $O_{v_x,y}(1)$ for them based on her input Y (formally these estimates are simply the probability of seeing a 0 or 1 conditioned on the protocol reaching v_x and conditioned on y). In the case where the input distribution μ is a product distribution $\mu = \mu_x \times \mu_y$, the X player can also compute the estimates $O_{v_x,y}(0)$ and $O_{v_x,y}(1)$, since they are independent of the input y given the node v_x . The goal is to simulate the protocol according to the “correct” distributions.

3.1. Compression according to the internal information cost. In our first compression protocol, the players use shared randomness to sample the bit at every node of the protocol tree for $\pi(x, y)$. In other words, for every prefix v of messages, each player samples the next bit of the interaction according to the best guess that they have for how this bit is distributed, even if the next bit is actually transmitted by the other player in the original protocol. The players do this using shared randomness, in a way that guarantees that if their guesses are close to the correct distribution, then the probability that they sample the same bit is high. More precisely, the players share a random number $\kappa_v \in [0, 1]$ for every node v in the tree, and each player guesses the next bit following v to be 1, if the player’s estimated probability for the message being 1 is at least κ_v . Note that the player that owns v samples the next bit with the correct probability. It is not hard to see that the probability of getting inconsistent samples at the node v is at most $|O_{v,x} - O_{v,y}| \stackrel{def}{=} (|O_{v,x}(0) - O_{v,y}(0)| + |O_{v,x}(1) - O_{v,y}(1)|)/2$. Once they have each sampled from the possible interactions, we shall argue that there




FIG. 2. An illustration of the compression protocol for nonproduct distributions. The circle nodes are owned by player X and the square nodes are owned by Y . The figure illustrates the states of the protocol trees after all the bits have been sampled. The players then proceed to resolve their disagreements. The disagreement at node u is resolved in favor of X since he owns the node. The protocol proceeds to node v where the disagreement is resolved in favor of Y . The final computation path in this case is $u - v - w$, the output is 0, and the total number of disagreements along the path is 2.

is a *correct leaf* in the protocol tree, whose distribution is exactly the same as the leaf in the original protocol. This is the leaf that is obtained by starting at the root and repeatedly taking the edge that was sampled by the owner of the node. We then show how the players can use hashing and binary search to communicate a polylogarithmic number of bits with each other to resolve the inconsistencies in their samples and find this correct path with high probability. In this way, the final outcome will be statistically close to the distribution of the original protocol. An example run for this protocol is illustrated in Figure 2. The additional interaction cost scales according to the expected number of inconsistencies on the path to the correct leaf, which we show can be bounded by $\sqrt{I \cdot C}$, where I is the information content and C is the communication cost of the original protocol.

Recall from the majority example above that ϵ information can mean that $|O_{v,x} - O_{v,y}| \approx \sqrt{\epsilon}$. In fact, the “worst-case” example for us is when in each round I/C information is conveyed, leading to a per-round error of $\sqrt{I/C}$ and a total expected number of mistakes of $\sqrt{I/C} \cdot C = \sqrt{I \cdot C}$.

3.2. Compression according to the external information cost. Our more efficient solution gives a protocol with communication complexity within polylogarithmic factors of the external information cost. It is illustrated on Figure 3. The idea in this case is to simulate chunks of the protocol that convey a constant amount of information each. If we can simulate a portion of the protocol that conveys a constant (or even $1/\text{polylog}$) amount of information using a polylogarithmic number of bits of communication, then we can simulate the entire protocol using $\tilde{O}(I)$ bits of communication.

The advantage the players have in the case that we are measuring information from the viewpoint of an observer is that for each node in the tree, the player who owns that node knows not only the correct distribution for the next bit, but also knows what the distribution that the observer has in mind is. They can use this shared knowledge to sample entire paths according to the distribution that is common knowledge at every step. In general, the distribution of the sampled path can deviate quite a bit from the correct distribution. However, we argue that if the information




FIG. 3. An illustration of the compression protocol for product distributions. The gray layer represents the “frontier” of nodes where some fixed amount of information is conveyed in the original protocol, and which is simulated in one iteration of the compressed protocol. Once the players agree on a node u , they compute a new frontier, illustrated here by the black layer.

conveyed on a path is small ($1/\text{polylog}$ bit), then the difference between the correct and the approximate probability is constant. After sampling the approximate bits for an appropriate number of steps so as to cover $1/\text{polylog}$ information, the players can communicate to estimate the correct probability with which this node was supposed to occur. The players can then either accept the sequence or resample a new sequence in order to get a final sample that behaves in a way that is close to the distribution of the original protocol. The method of sampling a message by repeatedly sampling from shared randomness is widely known as *rejection sampling*, and was used by several prior works about compressing protocols [11, 15] and by others [22, 12] in the context of theoretical computer science.

There are several technical challenges involved in getting this to work. The fact that the inputs of the players are independent is important for the players to decide how many messages the players should try to sample at once to get to the frontier where $1/\text{polylog}$ bits of information have been revealed. When the players’ inputs are dependent, they cannot estimate how many messages they should sample before the information content becomes too high, and we are unable to make this approach work.

4. Preliminaries.

Notation. We reserve capital letters for random variables and distributions, calligraphic letters for sets, and small letters for elements of sets. Throughout this paper, we often use the notation $|b$ to denote conditioning on the event $B = b$. Thus $A|b$ is shorthand for $A|B = b$. Given a sequence of symbols $A = A_1, A_2, \dots, A_k$, we use $A_{\leq j}$ to denote the prefix of length j .

We use the standard notion of *statistical/total variation* distance between two distributions.

DEFINITION 4.1. Let D and F be two random variables taking values in a set \mathcal{S} . Their statistical distance is

$$|D - F| \stackrel{\text{def}}{=} \max_{\mathcal{T} \subseteq \mathcal{S}} (|\Pr[D \in \mathcal{T}] - \Pr[F \in \mathcal{T}]|) = \frac{1}{2} \sum_{s \in \mathcal{S}} |\Pr[D = s] - \Pr[F = s]|.$$

If $|D - F| \leq \epsilon$ we shall say that D is ϵ -close to F . We shall also use the notation $D \stackrel{\epsilon}{\approx} F$ to mean D is ϵ -close to F .

4.1. Information theory.

DEFINITION 4.2 (entropy). The entropy of a random variable X is

$$H(X) \stackrel{\text{def}}{=} \sum_x \Pr[X = x] \log(1/\Pr[X = x]).$$

The conditional entropy $H(X|Y)$ is defined to be $\mathbb{E}_{y \in_R Y} [H(X|Y = y)]$.

FACT 4.3. $H(AB) = H(A) + H(B|A)$.

DEFINITION 4.4 (mutual information). The mutual information between two random variables A and B , denoted $I(A; B)$, is defined to be the quantity $H(A) - H(A|B) = H(B) - H(B|A)$. The conditional mutual information $I(A; B|C)$ is $H(A|C) - H(A|BC)$.

In analogy with the fact that $H(AB) = H(A) + H(B|A)$, we have the next proposition.

PROPOSITION 4.5. Let C_1, C_2, D, B be random variables. Then

$$I(C_1 C_2; B|D) = I(C_1; B|D) + I(C_2; B|C_1 D).$$

The previous proposition immediately implies the following proposition.

PROPOSITION 4.6 (superadditivity of mutual information). Let C_1, C_2, D, B be random variables such that for every fixing of D , C_1 and C_2 are independent. Then

$$I(C_1; B|D) + I(C_2; B|D) \leq I(C_1 C_2; B|D).$$

We also use the notion of *divergence*, which is a different way to measure the distance between two distributions.

DEFINITION 4.7 (divergence). The informational divergence between two distributions is $\mathbb{D}(A||B) \stackrel{\text{def}}{=} \sum_x A(x) \log(A(x)/B(x))$.

For example, if B is the uniform distribution on $\{0, 1\}^n$, then $\mathbb{D}(A||B) = n - H(A)$.

PROPOSITION 4.8. $\mathbb{D}(A||B) \geq |A - B|^2$.

PROPOSITION 4.9. Let A, B, C be random variables in the same probability space. For every a in the support of A and c in the support of C , let B_a denote $B|A = a$ and B_{ac} denote $B|A = a, C = c$. Then $I(A; B|C) = \mathbb{E}_{a, c \in_R A, C} [\mathbb{D}(B_{ac}||B_c)]$.

The above facts imply the following easy proposition.

PROPOSITION 4.10. With notation as in Proposition 4.9, for any random variables A, B ,

$$\mathbb{E}_{a \in_R A} [|(B_a) - B|] \leq \sqrt{I(A; B)}.$$

Proof.

$$\begin{aligned} \mathbb{E}_{a \in_R A} [|(B_a) - B|] &\leq \mathbb{E}_{a \in_R A} \left[\sqrt{\mathbb{D}(B_a||B)} \right] \\ &\leq \sqrt{\mathbb{E}_{a \in_R A} [\mathbb{D}(B_a||B)]} && \text{(by convexity)} \\ &= \sqrt{I(A; B)} && \text{(by Proposition 4.9).} \quad \square \end{aligned}$$

Generic Communication Protocol
<ol style="list-style-type: none"> 1. Set v to be the root of the protocol tree. 2. If v is a leaf, the protocol ends and Alice and Bob output the corresponding values $A_v(x)$ and $B_v(x)$. Otherwise, the player owning v samples a child of v according to the distribution associated with her input for v and sends a bit to the other player to indicate which child was sampled. 3. Set v to be the newly sampled node and return to the previous step.

FIG. 4. A communication protocol.

4.2. Communication complexity. Let \mathcal{X}, \mathcal{Y} denote the set of possible inputs to the two players, who we name P_x, P_y . In this paper, we view a *private coin protocol* for computing a function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{Z}_K$ as a binary tree with the following structure:

- each node is *owned* by P_x or by P_y ;
- for every $x \in \mathcal{X}$, each internal node v owned by P_x is associated with a distribution $O_{v,x}$ supported on the children of v . Similarly, for every $y \in \mathcal{Y}$, each internal node v owned by P_y is associated with a distribution $O_{v,y}$ supported on the children of v ;
- each leaf ℓ of the protocol is labeled with functions $A_\ell : \mathcal{X} \rightarrow \mathbb{Z}_K$ and $B_\ell : \mathcal{Y} \rightarrow \mathbb{Z}_K$ that allow Alice and Bob to compute the output.

On input x, y , the protocol π is executed as in Figure 4. The protocol is said to *succeed* in computing f if it terminates on a leaf ℓ such that $A_\ell(x) = B_\ell(y) = f(x, y)$. Note that in general Alice and Bob do not have to give the same output at the end of the protocol's execution, but giving divergent outputs results in an automatic failure.

A public coin protocol is a distribution on private coin protocols, run by first using shared randomness to sample an index r and then running the corresponding private coin protocol π_r . Every private coin protocol is thus a public coin protocol. The protocol is called *deterministic* if all distributions labeling the nodes have support size 1.

DEFINITION 4.11. *The communication complexity of a public coin protocol π , denoted $CC(\pi)$, is the maximum depth of the protocol trees in the support of π .*

Given a protocol π , $\pi(x, y)$ denotes the concatenation of the public randomness with all the messages that are sent during the execution of π . We call this the *transcript* of the protocol. We shall use the notation $\pi(x, y)_j$ to refer to the j th transmitted bit in the protocol. We write $\pi(x, y)_{\leq j}$ to denote the concatenation of the public randomness in the protocol with the first j message bits that were transmitted in the protocol. Given a transcript, or a prefix of the transcript, v , we write $CC(v)$ to denote the number of message bits in v (i.e., the length of the communication).

We often assume that every leaf in the protocol is at the same depth. We can do this since if some leaf is at depth less than the maximum, we can modify the protocol by adding dummy nodes which are always picked with probability 1, until all leaves are at the same depth. This does not change the communication complexity.

DEFINITION 4.12 (communication complexity notation). *For a function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{Z}_K$, a distribution μ supported on $\mathcal{X} \times \mathcal{Y}$, and a parameter $\rho > 0$, $D_\rho^\mu(f)$ denotes the communication complexity of the cheapest deterministic protocol for computing f on inputs sampled according to μ with error ρ . $R_\rho(f)$ denotes the cost of the best randomized public coin protocol for computing f with error at most ρ on every input.*

We shall use the following fact, first observed by Yao.

FACT 4.13 (Yao's min-max). $R_\rho(f) = \max_\mu D_\rho^\mu(f)$.

4.3. Finding differences in inputs. We use the following lemma of Feige et al. [9].

LEMMA 4.14 (see [9]). *There is a randomized public coin protocol τ with communication complexity $O(\log(k/\epsilon))$ such that on input of two k -bit strings x, y , it outputs the first index $i \in [k]$ such that $x_i \neq y_i$ with probability at least $1 - \epsilon$, if such an i exists.*

For completeness, we include the proof (based on hashing) in Appendix C.

4.4. Measures of information complexity. Here we briefly discuss the two measures of information cost defined in the introduction.

Let R be the public randomness, and X, Y be the inputs to the protocol π . By the chain rule for mutual information,

$$\begin{aligned} \text{IC}_\mu^\circ(\pi) &= I(XY; \pi(X, Y)) \\ &= I(XY; R) + \sum_{i=1}^{\text{CC}(\pi)} I(XY; \pi(X, Y)_i | \pi(X, Y)_{<i}) \\ &= 0 + \sum_{i=1}^{\text{CC}(\pi)} I(XY; \pi(X, Y)_i | \pi(X, Y)_{<i}), \\ \text{IC}_\mu^i(\pi) &= I(X; \pi(X, Y) | Y) + I(Y; \pi(X, Y) | X) \\ &= I(X; R | Y) + I(Y; R | X) + \sum_{i=1}^{\text{CC}(\pi)} I(X; \pi(X, Y)_i | Y \pi(X, Y)_{<i}) \\ &\quad + I(Y; \pi(X, Y)_i | X \pi(X, Y)_{<i}) \\ &= 0 + \sum_{i=1}^{\text{CC}(\pi)} I(X; \pi(X, Y)_i | Y \pi(X, Y)_{<i}) + I(Y; \pi(X, Y)_i | X \pi(X, Y)_{<i}). \end{aligned}$$

Let w be any fixed prefix of the transcript of length $i - 1$. If it is the X player's turn to speak in the protocol, $I(Y; \pi(X, Y)_i | X, \pi(X, Y)_{\leq i-1} = w) = 0$. If it is the Y player's turn to speak, then $I(X; \pi(X, Y)_i | Y, \pi(X, Y)_{\leq i-1} = w) = 0$. On the other hand $I(XY; \pi(X, Y)_i | \pi(X, Y)_{\leq i-1} = w) \geq \max\{I(X; \pi(X, Y)_i | Y \pi(X, Y)_{\leq i-1} = w), I(Y; \pi(X, Y)_i | X \pi(X, Y)_{\leq i-1} = w)\}$ by the chain rule.

Thus, we get the following fact.

FACT 4.15. $\text{IC}_\mu^\circ(\pi) \geq \text{IC}_\mu^i(\pi)$.

If μ is a product distribution,

$$\begin{aligned} &I(XY; \pi(X, Y)_i | \pi(X, Y)_{\leq i-1} = w) \\ &= I(X; \pi(X, Y)_i | \pi(X, Y)_{\leq i-1} = w) + I(Y; \pi(X, Y)_i | X \pi(X, Y)_{\leq i-1} = w) \\ &= I(X; \pi(X, Y)_i | Y \pi(X, Y)_{\leq i-1} = w) + I(Y; \pi(X, Y)_i | X \pi(X, Y)_{\leq i-1} = w). \end{aligned}$$

So we can conclude with the following.

FACT 4.16. *If μ is a product distribution, $\text{IC}_\mu^i(\pi) = \text{IC}_\mu^\circ(\pi)$.*

We note that $\text{IC}_\mu^i(\pi)$ and $\text{IC}_\mu^\circ(\pi)$ can be arbitrarily far apart, for example, if μ is such that $\Pr[X = Y] = 1$, then $\text{IC}_\mu^i(\pi) = 0$, even though $\text{IC}_\mu^\circ(\pi)$ may be arbitrarily large.

A remark on the role of public and private randomness. Public randomness is considered part of the protocol's transcript. But even if the randomness is short compared to the overall communication complexity, making it public can have a dramatic

effect on the information content of the protocol. (As an example, consider a protocol where one party sends a message of $x \oplus r$, where x is its input and r is random. If the randomness r is private, then this message has zero information content. If the randomness is public, then the message completely reveals the input. This protocol may seem trivial since its communication complexity is larger than the input length, but in fact we will be dealing with exactly such protocols, as our goal will be to “compress” communication of protocols that have very large communication complexity, but very small information content.)

5. Proof of the direct sum theorems. In this section, we prove our direct sum theorems. By Yao’s min-max principle, for every function f , $R_\rho(f) = \max_\mu D_\rho^\mu(f)$. Thus Theorem 2.5 implies Theorem 2.3 and Theorem 2.8 implies Theorem 2.7. So we shall focus on proving Theorems 2.5 and 2.6, and the XOR lemmas Theorems 2.8 and 2.9.

By Theorem 2.4, the main step to establish Theorem 2.5 is to give an efficient simulation of a protocol with small information content by a protocol with small communication complexity. We shall use our two results on compression, that we restate here.

THEOREM 2.1 1 (restated). *For every distribution μ , every protocol π , and every $\epsilon > 0$, there exist functions π_x, π_y , and a protocol τ such that $|\pi_x(X, \tau(X, Y)) - \pi(X, Y)| < \epsilon$, $\Pr[\pi_x(X, \tau(X, Y)) \neq \pi_y(Y, \tau(X, Y))] < \epsilon$, and*

$$\text{CC}(\tau) \leq O\left(\sqrt{\text{CC}(\pi) \cdot \text{IC}_\mu^i(\pi)} \frac{\log(\text{CC}(\pi)/\epsilon)}{\epsilon}\right).$$

THEOREM 2.2 1 (restated). *For every distribution μ , every protocol π , and every $\alpha > 0$, there exist functions π_x, π_y , and a protocol τ such that $|\pi_x(X, \tau(X, Y)) - \pi(X, Y)| < \alpha$, $\Pr[\pi_x(X, \tau(X, Y)) \neq \pi_y(Y, \tau(X, Y))] < \alpha$, and*

$$\text{CC}(\tau) \leq O\left(\text{IC}_\mu^o(\pi) \frac{\log(\text{CC}(\pi)/\alpha)}{\alpha^2}\right).$$

Proof of Theorem 2.5 from Theorem 2.1. Let π be any protocol computing f^n on inputs drawn from μ^n with probability of error less than ρ . Then by Theorem 2.4, there exists a protocol τ_1 computing f on inputs drawn from μ with error at most ρ with $\text{CC}(\tau_1) \leq \text{CC}(\pi)$ and $\text{IC}_\mu^i(\tau_1) \leq 2\text{CC}(\pi)/n$. Next, applying Theorem 2.1 to the protocol τ_1 gives that there must exist a protocol τ_2 computing f on inputs drawn from μ with error at most $\rho + \alpha$ and

$$\begin{aligned} \text{CC}(\tau_2) &\leq O\left(\sqrt{\text{CC}(\tau_1) \text{IC}_\mu^i(\tau_1)} \log(\text{CC}(\tau_1)/\alpha)/\alpha\right) \\ &= O\left(\sqrt{\text{CC}(\pi) \text{CC}(\pi)/n} \log(\text{CC}(\pi)/\alpha)/\alpha\right) \\ &= O\left(\frac{\text{CC}(\pi) \log(\text{CC}(\pi)/\alpha)}{\sqrt{n}}\right). \end{aligned}$$

This proves Theorem 2.5. \square

Proof of Theorem 2.6 from Theorem 2.2. Let π be any protocol computing f^n on inputs drawn from μ^n with probability of error less than ρ . Then by Theorem 2.4, there exists a protocol τ_1 computing f on inputs drawn from μ with error at most ρ with $\text{CC}(\tau_1) \leq \text{CC}(\pi)$ and $\text{IC}_\mu^i(\tau_1) \leq 2\text{CC}(\pi)/n$. Since μ is a product distribution, we have that $\text{IC}_\mu^o(\pi) = \text{IC}_\mu^i(\pi)$. Next, applying Theorem 2.2 to the protocol τ_1 gives that there must exist a protocol τ_2 computing f on inputs drawn from μ with error

at most $\rho + \alpha$ and

$$\begin{aligned} \text{CC}(\tau_2) &\leq O\left(\text{IC}_\mu^\circ(\tau_1) \log(\text{CC}(\tau_1)/\alpha)/\alpha^2\right) \\ &= O\left(\frac{\text{CC}(\pi) \log(\text{CC}(\pi)/\alpha)}{n\alpha^2}\right) \end{aligned}$$

This proves Theorem 2.6. \square

Proof of the XOR lemma. The proof for Theorem 2.8 (XOR lemma for distributional complexity) is very similar. First, we show an XOR analogue of Theorem 2.4 in section 6.

THEOREM 5.1. *For every distribution μ , there exists a protocol τ computing f with probability of error ρ over the distribution μ with $\text{CC}(\tau) \leq D_\rho^{\mu^n}(f^{+n}) + 2 \log K$ such that if τ' is the protocol that is the same as τ but stops running after $D_\rho^{\mu^n}(f^{+n})$ message bits have been sent, then $\text{IC}_\mu^i(\tau') \leq \frac{2D_\rho^{\mu^n}(f^{+n})}{n}$.*

Now let π be any protocol computing f^{+n} on inputs drawn from μ^n with probability of error less than ρ . Then by Theorem 5.1, there exists a protocol τ_1 computing f on inputs drawn from μ with error at most ρ with $\text{CC}(\tau_1) \leq \text{CC}(\pi) + 2 \log K$ and such that if τ'_1 denotes the first $\text{CC}(\pi)$ bits of the message part of the transcript, $\text{IC}_\mu^i(\tau'_1) \leq 2\text{CC}(\pi)/n$. Next, applying Theorem 2.1 to the protocol τ'_1 gives that there must exist a protocol τ'_2 simulating τ'_1 on inputs drawn from μ with error at most $\rho + \alpha$ and

$$\begin{aligned} \text{CC}(\tau'_2) &\leq O\left(\sqrt{\text{CC}(\tau'_1) \text{IC}_\mu^i(\tau'_1)} \log(\text{CC}(\tau'_1)/\alpha)/\alpha\right) \\ &= O\left(\sqrt{\text{CC}(\pi) \text{CC}(\pi)/n} \log(\text{CC}(\pi)/\alpha)/\alpha\right) \\ &= O\left(\frac{\text{CC}(\pi) \log(\text{CC}(\pi)/\alpha)/\alpha}{\sqrt{n}}\right). \end{aligned}$$

Finally we get a protocol for computing f by first running τ'_2 and then running the last $2 \log K$ bits of π . Thus we must have that

$$D_{\rho+\alpha}^\mu(f) \leq O\left(\frac{\text{CC}(\pi) \log(\text{CC}(\pi)/\alpha)/\alpha}{\sqrt{n}}\right) + 2 \log K.$$

This finishes the proof of Theorem 2.8. \square

The proof of Theorem 2.9 is analogous and is omitted here.

6. Reduction to small internal information cost. We now prove Theorems 2.4 and 5.1, showing that the existence of a protocol with communication complexity C for f^n (or f^{+n}) implies a protocol for f with information content roughly C/n .

THEOREM 2.4 1 (restated). *For every μ, f, ρ there exists a protocol τ computing f on inputs drawn from μ with probability of error at most ρ and communication at most $D_\rho^{\mu^n}(f^n)$ such that $\text{IC}_\mu^i(\tau) \leq \frac{2D_\rho^{\mu^n}(f^n)}{n}$.*

THEOREM 5.1 1 (restated). *For every distribution μ , there exists a protocol τ computing f with probability of error ρ over the distribution μ with $\text{CC}(\tau) \leq D_\rho^{\mu^n}(f^{+n}) + 2 \log K$ such that if τ' is the protocol that is the same as τ but stops running after $D_\rho^{\mu^n}(f^{+n})$ message bits have been sent, then $\text{IC}_\mu^i(\tau') \leq \frac{2D_\rho^{\mu^n}(f^{+n})}{n}$.*

The key idea involved in proving the above theorems is a way to split dependencies between the inputs that arose in the study of lower bounds for the communication complexity of disjointness and in the study of parallel repetition [18, 27, 28].

Protocol τ
<p>Public Randomness Phase :</p> <ol style="list-style-type: none"> 1. The players sample $j, w^{-j} \in_{\mathbb{R}} J, W^{-J}$ using public randomness. <p>Private Randomness Phase :</p> <ol style="list-style-type: none"> 1. P_x sets $x_j = x$, P_y sets $y_j = y$. 2. For every $i \neq j$, P_x samples X_i conditioned on the value of w^{-j}. 3. For every $i \neq j$, P_y samples Y_i conditioned on the value of w^{-j}. 4. The players simulate π on the inputs $x_1, \dots, x_n, y_1, \dots, y_n$ and output the jth output of π.

FIG. 5. A protocol simulating π .

Proof. Fix μ, f, n, ρ as in the statement of the theorems. We shall prove Theorem 2.4 first. Theorem 5.1 will easily follow by the nature of our proof. To prove Theorem 2.4, we show how to use the best protocol for computing f^n to get a protocol with small information content computing f . Let π be a deterministic protocol with communication complexity $D_{\rho}^{\mu^n}(f^n)$ computing f^n with probability of error at most ρ .

Let $(X_1, Y_1), \dots, (X_n, Y_n)$ denote random variables distributed according to μ^n . Let $\pi(X^n, Y^n)$ denote the random variable of the transcript (which is just the concatenation of all messages, since this is a deterministic protocol) that is obtained by running the protocol π on inputs $(X_1, Y_1), \dots, (X_n, Y_n)$. We define random variables $W = W_1, \dots, W_n$ where each W_j takes a value in the disjoint union $\mathcal{X} \uplus \mathcal{Y}$ so that each $W_j = X_j$ with probability 1/2 and $W_j = Y_j$ with probability 1/2. Let W^{-j} denote $W_1, \dots, W_{j-1}, W_{j+1}, \dots, W_n$.

Our new protocol τ shall operate as in Figure 5. Note the distinction between *public* and *private* randomness. This distinction make a crucial difference in the definition of information content, as making more of the randomness public reduces the information content of a protocol.

The probability that the protocol τ makes an error on inputs sampled from μ is at most the probability that the protocol π makes an error on inputs sampled from μ^n . It is also immediate that $CC(\tau) = CC(\pi)$. All that remains is to bound the information content $IC_{\mu}^i(\tau)$. We do this by relating it to the communication complexity of π .

To simplify notation, below we will use π to denote $\pi(X, Y)$ when convenient.

$$\begin{aligned}
 D_{\rho}^{\mu^n}(f^n) &\geq CC(\pi) \geq I(X_1 \cdots X_n Y_1 \cdots Y_n; \pi | W) \\
 &\geq \sum_{j=1}^n I(X_j Y_j; \pi | W) = nI(X_J Y_J; \pi | W_J),
 \end{aligned}$$

where the last inequality follows from Proposition 4.6. Next observe that the variables JW^{-J} are independent of X_J, Y_J, W_J . Thus we can write

$$\begin{aligned}
 I(X_J Y_J; \pi | JW) &= I(X_J Y_J; \pi | JW_J W^{-J}) + I(X_J Y_J; JW^{-J} | W_J) \\
 &= I(X_J Y_J; JW^{-J} \pi | W_J) \\
 &= I(X_J Y_J; JW^{-J} \pi | W_J) \\
 &= \frac{I(X_J Y_J; JW^{-J} \pi | X_J) + I(X_J Y_J; JW^{-J} \pi | Y_J)}{2} \\
 &= \frac{I(Y_J; JW^{-J} \pi | X_J) + I(X_J; JW^{-J} \pi | Y_J)}{2},
 \end{aligned}$$

Protocol γ
<p>Public Randomness Phase :</p> <ol style="list-style-type: none"> 1. The players sample $j, w^{-j} \in_{\mathbb{R}} J, W^{-J}$ using public randomness. <p>Private Randomness Phase :</p> <ol style="list-style-type: none"> 1. P_x sets $x_j = x$, P_y sets $y_j = y$. 2. For every $i \neq j$, P_x samples X_i conditioned on the value of w^{-j}. 3. For every $i \neq j$, P_y samples Y_i conditioned on the value of w^{-j}. 4. The players simulate π on the inputs $x_1, \dots, x_n, y_1, \dots, y_n$ to compute $z \in \mathbb{Z}_K$. 5. P_x computes $\sum_{i \neq j, w_i = y_i} f(x_i, w_i)$ and sends this sum to P_y. 6. P_y outputs the value of the function as $z - \sum_{i \neq j, w_i = y_i} f(x_i, w_i) - \sum_{i \neq j, w_i = x_i} f(w_i, y_i)$.

FIG. 6. A protocol simulating π .

where the last equality follows from the fact that X_J determines X and Y_J determines Y . This last quantity is simply the information content of τ . Thus we have shown that $\text{CC}(\pi) \geq (n/2)\text{IC}_{\mu}^i(\tau)$ as required.

Remark 6.1. The analysis above can be easily improved to get the bound $\text{IC}_{\mu}^i(\tau) \leq \text{CC}(\pi)/n$ by taking advantage of the fact that each bit of the transcript gives information about at most one of the players' inputs, but for simplicity we do not prove this here.

This completes the proof for Theorem 2.4. The proof for Theorem 5.1 is very similar. As above, we let π denote the best protocol for computing f^{+n} on inputs sampled according to μ^n . Analogously to τ as above, we define the simulation γ as in Figure 6.

As before, the probability that the protocol γ makes an error on inputs sampled from μ is at most the probability that the protocol π makes an error on inputs sampled from μ^n , since there is an error in γ if and only if there is an error in the computation of z . It is also immediate that $\text{CC}(\gamma) = \text{CC}(\pi) + 2 \log K$.

Let $\gamma'(X, Y)$ denote the concatenation of the public randomness and the messages of γ up to the computation of z . Then, exactly as in the previous case, we have the bound

$$\text{IC}_{\mu}^i(\gamma') \leq 2\text{CC}(\pi)/n.$$

This completes the proof. \square

7. Compression according to the internal information cost. We now prove our main technical theorem, Theorem 2.1.

THEOREM 2.1 2 (restated). *For every distribution μ , every protocol π , and every $\epsilon > 0$, there exist functions π_x, π_y , and a protocol τ such that $|\pi_x(X, \tau(X, Y)) - \pi(X, Y)| < \epsilon$, $\Pr[\pi_x(X, \tau(X, Y)) \neq \pi_y(Y, \tau(X, Y))] < \epsilon$, and*

$$\text{CC}(\tau) \leq O\left(\sqrt{\text{CC}(\pi) \cdot \text{IC}_{\mu}^i(\pi)} \frac{\log(\text{CC}(\pi)/\epsilon)}{\epsilon}\right).$$

7.1. A proof sketch. Here is a high level sketch of the proof. Let μ be a distribution over $\mathcal{X} \times \mathcal{Y}$. Let π be a public coin protocol that does some computation

using the inputs X, Y drawn according to μ . Our goal is to give a protocol τ that simulates π on μ such that⁵

$$CC(\tau) = O\left(\sqrt{CC(\pi) \cdot IC_{\mu}^i(\pi)} \log(CC(\pi))\right).$$

For the sake of simplicity, here we assume that the protocol π has no public randomness. π then specifies a protocol tree which is a binary tree of depth $CC(\pi)$ where every nonleaf node w is owned by one of the players, whose turn it is to speak at this node. Each nonleaf node has a “0 child” and a “1 child.” For every such node w in the tree and every possible message $b \in \{0, 1\}$, the X player gets input x and uses this to define $O_{w,x}(b)$ as the probability in π that, conditioned on reaching the node w and the input being x , the next bit will be b . The Y player defines $O_{w,y}(b)$ analogously. Note that if w is owned by the X player, then $O_{w,x}(b)$ is exactly the correct probability with which b is transmitted in the real protocol.

For every such node w , the players use public randomness to sample a shared random number $\kappa_w \in [0, 1]$ for every nonleaf node w in the tree. The X player uses these numbers to define the child $C_x(w)$ for every node w as follows: if $O_{w,x}(1) < \kappa_w$, $C_x(w)$ is set to the 0 child of w , and is set to the 1 child otherwise. The Y Player does the same using the values $O_{w,y}(1)$ (but the same κ_w) instead.

Now let $v_0, \dots, v_{CC(\pi)}$ be the correct path in the tree. This is the path where every subsequent node was sampled by the player that owned the previous node: for every i ,

$$v_{i+1} = \begin{cases} C_x(v_i) & \text{if X player owns } v_i, \\ C_y(v_i) & \text{if Y player owns } v_i. \end{cases}$$

Then $v_{CC(\pi)}$ has the same distribution as a leaf in π was supposed to have, and the goal of the players will be to identify $v_{CC(\pi)}$ with small communication.

In order to do this, the X player will compute the sequence of nodes $v_0^x, \dots, v_{CC(\pi)}^x$ by setting $v_{i+1}^x = C_x(v_i^x)$. Similarly, the Y player computes the path $v_0^y, \dots, v_{CC(\pi)}^y$ by setting $v_{i+1}^y = C_y(v_i^y)$. Observe that if these two paths agree on the first k nodes, then they must be equal to the correct path up to the first k nodes.

So far, we have not communicated at all. Now the parties communicate to find the first index i for which $v_i^x \neq v_i^y$. If $v_{i-1} = v_{i-1}^x = v_{i-1}^y$ was owned by the X player, the parties reset the i th node in their paths to v_i^x . Similarly, if v_{i-1} was owned by the Y player, the parties reset their i th node to be v_i^y . In this way, they keep fixing their paths until they have computed the correct path.

Thus the communication complexity of the new protocol is bounded by the number of mistakes times the communication complexity of finding a single mistake. Every path in the tree is specified by a $CC(\pi)$ -bit string, and finding the first inconsistency reduces to the problem of finding the first difference in two $CC(\pi)$ -bit strings. A simple protocol of Feige et al. [9] (based on hashing and binary search) gives a protocol for finding this first inconsistency, with communication only $O(\log CC(\pi))$. We describe and analyze this protocol in Appendix C. In section 7.2 we show how to bound the

⁵We identify the communication complexity of the protocols π, τ with their expected communication under μ , as by adding a small error, the two can be related using an easy Markov argument.

Protocol π_r	
Sampling Phase :	
<ol style="list-style-type: none"> 1. For every nonleaf node w in the tree, the player who owns w samples a child according to the distribution given by her input and the public randomness r. This leaves each player with a subtree of the original protocol tree, where each node has out-degree 1 or 0 depending on whether or not it is owned by the player. 	
Path Finding Phase :	
<ol style="list-style-type: none"> 1. Set v to be the root of the tree. 2. If v is a leaf, the computation ends with the value of the node. Else, the player to whom v belongs communicates one bit to the other player to indicate which of the children was sampled. 3. Set v to the sampled child and return to the previous step. 	

FIG. 7. π_r restated.

expected number of mistakes on the correct path in terms of the information content of the protocol. We show that if we are at node v_i in the protocol and the next bit has ϵ information, then the probability that $\Pr[C_x(v_i) \neq C_y(v_i)] \leq \sqrt{\epsilon}$. Since the total information content is $\text{IC}_\mu^i(\pi)$, we can use the Cauchy–Schwarz inequality to bound the expected number of mistakes by $\sqrt{\text{CC}(\pi)\text{IC}_\mu^i(\pi)}$.

7.2. The actual proof. In order to prove Theorem 2.1, we consider the protocol tree \mathcal{T} for π_r , for every fixing of the public randomness r . If R is the random variable for the public randomness used in π , we have the following claim.

CLAIM 7.1. $\text{IC}_\mu^i(\pi) = \mathbb{E}_R [\text{IC}_\mu^i(\pi_R)]$

Proof.

$$\begin{aligned}
 \text{IC}_\mu^i(\pi) &= I(\pi(X, Y); X|Y) + I(\pi(X, Y); Y|X) \\
 &= I(R\pi_R(X, Y); X|Y) + I(R\pi_R(X, Y); Y|X) \\
 &= I(R; X|Y) + I(R; Y|X) + I(\pi_R(X, Y); X|YR) + I(\pi_R(X, Y); Y|XR) \\
 &= I(\pi_R(X, Y); X|YR) + I(\pi_R(X, Y); Y|XR) \\
 &= \mathbb{E}_R [\text{IC}_\mu^i(\pi_R)] \quad \square
 \end{aligned}$$

It will be convenient to describe protocol π_r in a nonstandard, yet equivalent way in Figure 7.

For some error parameters β, γ , we define a randomized protocol $\tau_{\beta, \gamma}$ that will simulate π and use the same protocol tree. The idea behind the simulation is to avoid communicating by guessing what the other player's samples look like. The players shall make many mistakes in doing this, but they shall then use Lemma 4.14 to correct the mistakes and end up with the correct transcript. Our simulation is described in Figure 8.

Define $\pi_x(x, \tau_{\beta, \gamma}(x, y))$ (resp., $\pi_y(y, \tau_{\beta, \gamma}(x, y))$) to be a leaf of the final path computed by P_x (resp., P_y) in the protocol $\tau_{\beta, \gamma}$ (see Figure 8). The definition of the protocol $\tau_{\beta, \gamma}$ implies immediately the following upper bound on its communication

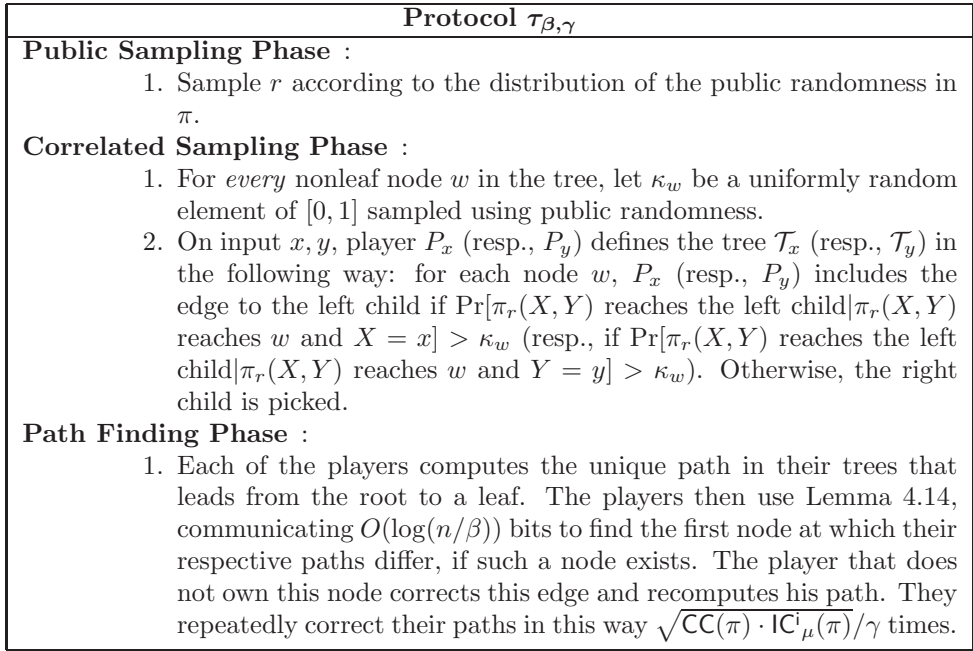


FIG. 8. The simulation of π .

complexity:

$$(7.1) \quad \text{CC}(\tau_{\beta,\gamma}) = O\left(\sqrt{\text{CC}(\pi) \cdot \text{IC}_\mu^i(\pi)} \log(\text{CC}(\pi)/\beta)/\gamma\right).$$

Let $V = V_0, \dots, V_{\text{CC}(\pi)}$ denote the “right path” in the protocol tree of $\tau_{\beta,\gamma}$. That is, every i , $V_{i+1} = 0$ if the left child of $V_{\leq i}$ is sampled by the owner of $V_{\leq i}$ and $V_{i+1} = 1$ otherwise. Observe that this path has the right distribution, since every child is sampled with exactly the right conditional probability by the corresponding owner. That is, we have the following claim.

CLAIM 7.2. *For every x, y, r , the distribution of $V|xyr$ as defined above is the same as the distribution of the sampled transcript in the protocol π .*

This implies in particular, that

$$I(X; V|rY) + I(Y; V|rX) = \text{IC}_\mu^i(\pi_r).$$

Given two fixed trees $\mathcal{T}_x, \mathcal{T}_y$ as in the above protocol, we say there is a *mistake* at level i if the out-edges of V_{i-1} are inconsistent in the trees. We shall first show that the expected number of mistakes that the players make is small.

LEMMA 7.3. $\mathbb{E}[\# \text{ of mistakes in simulating } \pi_r | r] \leq \sqrt{\text{CC}(\pi) \cdot \text{IC}_\mu^i(\pi_r)}$.

Proof. For $i = 1, \dots, \text{CC}(\pi)$, we denote by C_{ir} the indicator random variable for whether or not a mistake occurs at level i in the protocol tree for π_r , so that the number of mistakes is $\sum_{i=1}^{\text{CC}(\pi)} C_{ir}$.

We shall bound $\mathbb{E}[C_{ir}]$ for each i . A mistake occurs at a vertex w at depth i exactly when $\Pr[V_{i+1} = 0 | x \wedge V_{\leq i} = w] \leq \kappa_w < \Pr[V_{i+1} = 0 | y \wedge V_{\leq i} = w]$ or $\Pr[V_{i+1} = 0 | y \wedge V_{\leq i} = w] \leq \kappa_w < \Pr[V_{i+1} = 0 | x \wedge V_{\leq i} = w]$. Thus a mistake occurs at $v_{\leq i}$ with probability at most $|(V_i|xv_{<i}r) - (V_i|yv_{<i}r)|$.

If $v_{<i}$ is owned by P_x , then $V_i|xv_{<i}r$ has the same distribution as $V_i|xyv_{<i}r$; If $v_{<i}$ is owned by P_y , then $V_i|yv_{<i}r$ has the same distribution as $V_i|xyv_{<i}r$. Using Proposition 4.8 and Proposition 4.9, we have

$$\begin{aligned}
 & \mathbb{E}[C_{ir}] \\
 & \leq \mathbb{E}_{xyv_{<i} \in_{\mathbb{R}} XYV_{<i}} [|(V_i|xv_{<i}r) - (V_i|yv_{<i}r)|] \\
 & \leq \mathbb{E}_{xyv_{<i} \in_{\mathbb{R}} XYV_{<i}} [\max\{|(V_i|xyv_{<i}r) - (V_i|yv_{<i}r)|, |(V_i|xyv_{<i}r) - (V_i|xv_{<i}r)|\}] \\
 & \leq \mathbb{E}_{xyv_{<i} \in_{\mathbb{R}} XYV_{<i}} \left[\sqrt{\mathbb{D}(V_i|xyv_{<i}r|V_i|yv_{<i}r) + \mathbb{D}(V_i|xyv_{<i}r|V_i|xv_{<i}r)} \right] \\
 & \hspace{20em} \text{(By Proposition 4.8)} \\
 & \leq \sqrt{\mathbb{E}_{xyv_{<i} \in_{\mathbb{R}} XYV_{<i}} [\mathbb{D}(V_i|xyv_{<i}r|V_i|yv_{<i}r) + \mathbb{D}(V_i|xyv_{<i}r|V_i|xv_{<i}r)]} \\
 & \hspace{20em} \text{(by convexity)} \\
 & = \sqrt{I(X; V_i|YV_{<i}r) + I(Y; V_i|XV_{<i}r)} \\
 & \hspace{20em} \text{(by Proposition 4.9).}
 \end{aligned}$$

Finally we apply the Cauchy–Schwarz inequality to conclude that

$$\begin{aligned}
 \mathbb{E} \left[\sum_{i=1}^{\text{CC}(\pi)} C_{ir} \right] &= \sum_{i=1}^{\text{CC}(\pi)} \mathbb{E}[C_{ir}] \\
 &\leq \sqrt{\text{CC}(\pi) \sum_{i=1}^{\text{CC}(\pi)} \mathbb{E}[C_{ir}]^2} \\
 &\leq \sqrt{\text{CC}(\pi) \sum_{i=1}^{\text{CC}(\pi)} I(X; V_i|YV_{<i}r) + I(Y; V_i|XV_{<i}r)} \\
 &= \sqrt{\text{CC}(\pi) (I(X; V^{\text{CC}(\pi)}|Yr) + I(Y; V^{\text{CC}(\pi)}|Xr))} \\
 &= \sqrt{\text{CC}(\pi) \cdot \text{IC}_{\mu}^i(\pi_r)}. \quad \square
 \end{aligned}$$

We then get that overall the expected number of mistakes is small.

LEMMA 7.4. $\mathbb{E}[\# \text{ of mistakes in simulating } \pi] \leq \sqrt{\text{CC}(\pi) \cdot \text{IC}_{\mu}^i(\pi)}$.

Proof.

$$\begin{aligned}
 \mathbb{E}[\# \text{ of mistakes in simulating } \pi] &= \mathbb{E}_R[\# \text{ of mistakes in simulating } \pi_R] \\
 &\leq \mathbb{E}_R \left[\sqrt{\text{CC}(\pi) \cdot \text{IC}_{\mu}^i(\pi_R)} \right] \\
 &\leq \sqrt{\mathbb{E}_R [\text{CC}(\pi) \cdot \text{IC}_{\mu}^i(\pi_R)]} \\
 &= \sqrt{\text{CC}(\pi) \cdot \text{IC}_{\mu}^i(\pi)}. \quad \square
 \end{aligned}$$

LEMMA 7.5. *The distribution of the leaf sampled by $\tau_{\beta, \gamma}$ is $\gamma + \beta \frac{\sqrt{\text{CC}(\pi) \cdot \text{IC}_{\mu}^i(\pi)}}{\gamma}$ close to the distribution of the leaf sampled by π .*

Proof. We show that in fact the probability that both players do not finish the protocol with the leaf $V_{\text{CC}(\pi)}$ is bounded by $\gamma + \beta \frac{\sqrt{\text{CC}(\pi) \cdot \text{IC}_\mu^i(\pi)}}{\gamma}$. This follows from a simple union bound—the leaf $V_{\text{CC}(\pi)}$ can be missed in two ways: either the number of mistakes on the correct path is larger than $\sqrt{\text{CC}(\pi) \cdot \text{IC}_\mu^i(\pi)}/\gamma$ (probability at most γ by Lemma 7.4 and Markov's inequality) or our protocol fails to detect all mistakes (for each mistake this happens with probability β). \square

We set $\beta = \gamma^2/\text{CC}(\pi)$. Then, since $\text{CC}(\pi) \geq \text{IC}_\mu^i(\pi)$, we get that the protocol errs with probability at most $\rho + 2\gamma$. On the other hand, by (7.1), the communication complexity of the protocol is at most $O(\sqrt{\text{CC}(\pi) \cdot \text{IC}_\mu^i(\pi)} \log(\text{CC}(\pi)/\beta)/\gamma) = O(\sqrt{\text{CC}(\pi) \cdot \text{IC}_\mu^i(\pi)} \log(\text{CC}(\pi)/\gamma)/\gamma)$. Setting $\epsilon = 2\gamma$ proves the theorem. \square

8. Compression according to the external information cost. In this section we argue how to compress protocols according to the information learnt by an observer watching the protocol. We shall prove Theorem 2.6.

8.1. A proof sketch. We start with a rough proof sketch. Given a function f , distribution μ , and protocol π , we want to come up with a protocol τ simulating π such that $\text{CC}(\tau) = \tilde{O}(\text{IC}_\mu^o(\pi))$. We assume that π is a private coin protocol, in this proof sketch, for simplicity. For every nonleaf node w we denote by O_w the probability of transmitting a 1 at the node w conditioned on reaching that node (without taking into consideration the actual values of the inputs to the protocol). We shall write $O_{w,x}$, $O_{w,y}$ to denote this probability conditioned on a particular fixing of x or y , and conditioned on the event of reaching w during the run of the protocol. As a technical condition, we will assume that for every w , $O_{w,x}, O_{w,y} \in 1/2 \pm \beta$ for $\beta = 1/\text{polylog}(\text{CC}(\pi))$. This condition can be achieved for example by re-encoding π so that each party, instead of sending a bit b , sends $\text{polylog}(\text{CC}(\pi))$ random bits such that their majority is b .

For every node w owned by the X player, we define the *divergence at w* , denoted by D_w as $\mathbb{D}(O_{w,x}||O_w)$, where $\mathbb{D}(p||q) = p \log(p/q) + (1-p) \log((1-p)/(1-q))$ equals the divergence (also known as the Kullback–Leibler distance) between the p -biased coin and the q -biased coin. Given a node v , we define \mathcal{B}_v to be the set of descendants w of v such that if we sum up $D_{w'}$ for all intermediate nodes w' on the path from v to w we get a total $< \beta$ but adding D_w makes the total at least β or w is a leaf. We define B_v to be the distribution over \mathcal{B}_v that is induced by following the probabilities $O_{w'}$ along the path. Note that this distribution is known to both parties. We define B_{vx} to be the distribution obtained by assigning the probability of an edge according to $O_{w,x}$ for nodes owned by the x player and O_w for nodes owned by the y player. Similarly we define the distribution B_{vy} .

The protocol proceeds as follows (initially v is set to the root of the tree, t below is some large constant).

1. Both parties use their shared randomness to obtain a random element w according to the distribution B_v . (This involves first sampling a random leaf and then using binary search to find the first location in which the divergence surpasses β .)
2. The X player sends a bit a_1 that equals 1 with probability $\min\{1, B_{vx}(w)/(tB_v(w))\}$.
3. The Y player sends a bit a_2 that equals 1 with probability $\min\{1, B_{vy}(w)/(tB_v(w))\}$.
4. If $a_1 = a_2 = 1$, then they set $v = w$. If v is a leaf they end the protocol, otherwise go back to step 1.

To get a rough idea why this protocol makes sense, consider the case that all the nodes in \mathcal{B}_v are two levels below v , with the first node (i.e., v) owned by the X player, and the nodes in the intermediate level owned by the Y player. For a node $w \in \mathcal{B}_v$, let $B_{vxy}(w)$ be the true probability of arriving at w , and let $\tilde{B}(w) = \tilde{B}_v(w)$ be the estimated probability. Fixing w , we write $B(w) = B_1 B_2$ and $\tilde{B}(w) = \tilde{B}_1 \tilde{B}_2$, where B_i denotes the true probability that step i is taken according to w , and \tilde{B}_i denotes this probability as estimated by an observer who does not know x, y .

The probability that w is output at the end of step 1 is $\tilde{B}_1 \tilde{B}_2$. Now assume that the threshold t is set high enough so that we can assume that $tB_v(w) > B_{vx}(w), B_{vy}(y)$ with high probability. In this case the probability that w is accepted equals

$$(8.1) \quad \Pr[a_1 = 1] \Pr[a_2 = 1] = \left(\frac{B_1 \tilde{B}_2}{t \tilde{B}_1 \tilde{B}_2} \right) \left(\frac{\tilde{B}_1 B_2}{t \tilde{B}_1 \tilde{B}_2} \right) = \frac{B_1 B_2}{t^2 \tilde{B}_1 \tilde{B}_2};$$

thus the total probability that w is output is $\tilde{B}_1 \tilde{B}_2$ times (8.1) which is exactly its correct probability $B_1 B_2$ divided by t^2 , and hence we get an overhead of t^2 steps, but output the right distribution over w .

8.2. The actual proof. We shall compress the protocol into two steps. In the first step, we shall get a protocol simulating π whose messages are *smoothed out* in the sense that every bit in the protocol is relatively close to being unbiased, even conditioned on every fixing of the inputs and the prior transcript. We shall argue that this process does not change the typical divergence of bits in the protocol, a fact that will then let us compress such a protocol.

For every prefix of the transcript w that includes $i \geq 1$ message bits, and every pair of inputs x, y we define the following distributions on prefixes of transcripts that include $i + 1$ message bits:

$$\begin{aligned} O_w(a) &= \Pr[\pi(X, Y)_{\leq i+1} = a | \pi(X, Y)_{\leq i} = w]; \\ O_{w,x}(a) &= \begin{cases} \Pr[\pi(X, Y)_{\leq i+1} = a | \pi(X, Y)_{\leq i} = w, X = x] & \text{if the node } w \text{ is owned by the } x \text{ player,} \\ \Pr[\pi(X, Y)_{\leq i+1} = a | \pi(X, Y)_{\leq i} = w] & \text{else;} \end{cases} \\ O_{w,y}(a) &= \begin{cases} \Pr[\pi(X, Y)_{\leq i+1} = a | \pi(X, Y)_{\leq i} = w, Y = y] & \text{if the node } w \text{ is owned by the } y \text{ player,} \\ \Pr[\pi(X, Y)_{\leq i+1} = a | \pi(X, Y)_{\leq i} = w] & \text{else;} \end{cases} \\ O_{w,x,y}(a) &= \Pr[\pi(X, Y)_{\leq i+1} = a | \pi(X, Y)_{\leq i} = w, X = x, Y = y]. \end{aligned}$$

Next we define the following measures of information.

DEFINITION 8.1 (conditional divergence). *Given a protocol π , a prefix v of the transcript, and $j \in [\text{CC}(v)]$, we define the j th step divergence cost as*

$$\begin{aligned} \mathbb{D}_{x,j}^\pi(v) &\stackrel{\text{def}}{=} \mathbb{D}(O_{v_{\leq j},x} || O_{v_{\leq j}}), \\ \mathbb{D}_{y,j}^\pi(v) &\stackrel{\text{def}}{=} \mathbb{D}(O_{v_{\leq j},y} || O_{v_{\leq j}}). \end{aligned}$$

We define the divergence cost for the whole prefix as the sum of the step divergence costs:

$$\mathbb{D}_x^\pi(v) \stackrel{\text{def}}{=} \sum_{j=1}^{\text{CC}(v)} \mathbb{D}_{x,j}^\pi(v), \quad \mathbb{D}_y^\pi(v) \stackrel{\text{def}}{=} \sum_{j=1}^{\text{CC}(v)} \mathbb{D}_{y,j}^\pi(v).$$

We denote

$$\mathbb{D}_{xy}^\pi(v) \stackrel{def}{=} \mathbb{D}_x^\pi(v) + \mathbb{D}_y^\pi(v).$$

We have the following lemma.

LEMMA 8.2. *For any interval $[i, i + 1, \dots, j]$ of bits from the transcript, and any prefix v ,*

$$\mathbb{E}_{X,Y,\pi(X,Y)|\pi(X,Y)_{<i}=v} \left[\sum_{r=i}^j \mathbb{D}_{X,r}^\pi(\pi(X,Y)) + \mathbb{D}_{Y,r}^\pi(\pi(X,Y)) \right] \leq I(XY; \pi(X,Y)_{\leq j} | \pi(X,Y)_{<i} = v).$$

Proof. By linearity of expectation, the left-hand side is equal to

$$\sum_{r=i}^j \mathbb{E}_{X,Y,\pi(X,Y)|\pi(X,Y)_{<i}=w} [\mathbb{D}_{X,r}^\pi(\pi(X,Y)) + \mathbb{D}_{Y,r}^\pi(\pi(X,Y))].$$

Now consider any fixing of $\pi(X,Y)_{<r} = w$. Suppose without loss of generality that for this fixing it is the x -player's turn to speak in π . Then

$$\mathbb{E}_{X,Y|\pi(X,Y)_{<r}=w} [\mathbb{D}_{Y,j}^\pi(\pi(X,Y))]$$

is 0, since $O_{w,y}$ and O_w are the same distribution. By Proposition 4.9, the contribution of the other term under this fixing is

$$\begin{aligned} \mathbb{E}_{X,Y,\pi(X,Y)|\pi(X,Y)_{<r}=w} [\mathbb{D}_{X,r}^\pi(\pi(X,Y))] &= I(X; \pi(X,Y)_r | \pi(X,Y)_{<r} = w) \\ &\leq I(XY; \pi(X,Y)_r | \pi(X,Y)_{<r} = w). \end{aligned}$$

Thus the entire sum is bounded by

$$\sum_{r=i}^j I(XY; \pi(X,Y)_r | \pi(X,Y)_{<r}, \pi(X,Y)_{<i} = v) = I(XY; \pi(X,Y)_{\leq j} | \pi(X,Y)_{<i} = v)$$

by the chain rule for mutual information. \square

DEFINITION 8.3 (smooth protocols). *A protocol π is β -smooth if for every x, y, i, v_i ,*

$$\Pr[\pi(x, y)_{i+1} = 1 | \pi(x, y)_{\leq i} = v_i] \in [1/2 - \beta, 1/2 + \beta].$$

We show that every protocol can be simulated by a smooth protocol whose typical divergence cost is similar to the original protocol.

LEMMA 8.4 (smooth simulation). *There exists a constant $\ell > 0$ such that for every protocol π and distribution μ on inputs X, Y , and all $0 < \beta, \gamma < 1$ there exists a β -smooth protocol τ such that*

- $|\tau(X, Y) - \pi(X, Y)| < \gamma$,
- $\text{CC}(\tau) = \ell \text{CC}(\pi) \log(\text{CC}(\pi)/\gamma) / \beta^2$, and
- $\Pr_{X,Y} [\mathbb{D}_{XY}^\tau(\tau(X, Y)) > \text{IC}_\mu^\circ(\pi)/\gamma] \leq 2\gamma$.

The main technical part of the proof will be to show how to compress such a smooth protocol. We shall prove the following theorem.

THEOREM 8.5. *There exists a constant k such that for every $\epsilon > 0$, if π is a protocol such that for every x, y, v, i we have that*

$$\Pr[\pi(x, y)_{i+1} = 1 | v_{\leq i}] \in \left[\frac{1}{2} - \frac{1}{k \log(\text{CC}(\pi)/\epsilon)}, \frac{1}{2} + \frac{1}{k \log(\text{CC}(\pi)/\epsilon)} \right],$$

then for every distribution μ on inputs X, Y , there exists a protocol τ with communication complexity Q and a function p such that for every x, y , the expected statistical distance

$$\mathbb{E}_{X,Y} [|p(\tau(X, Y)) - \pi(X, Y)|] \leq \Pr \left[\mathbb{D}_{XY}^{\pi}(\pi(X, Y)) > \frac{\epsilon Q}{k \log(\text{CC}(\pi)/\epsilon)} \right] + k\epsilon.$$

Before proving Lemma 8.4 and Theorem 8.5, we show how to use them to prove Theorem 2.2.

Proof of Theorem 2.2. We set $\gamma = \epsilon = \alpha/8k$, $\beta = \frac{1}{k \log(\text{CC}(\pi)/\gamma)}$. Lemma 8.4 gives a β -smooth simulation τ_1 of π with communication complexity $\text{CC}(\pi) \cdot \text{polylog}(\text{CC}(\pi)/\alpha)$, that is γ close to the correct simulation. Next set $Q = \frac{\text{IC}^{\circ}_{\mu}(\pi) \cdot k \log(\text{CC}(\pi)/\gamma)}{\gamma^2}$. Then the probability that the divergence cost of a transcript of τ_1 exceeds $\frac{\gamma Q}{k \log(\text{CC}(\pi)/\gamma)}$ is at most 2γ . Thus, we can apply Theorem 8.5 to get a new simulation of π with total error $2\gamma + k\gamma \leq \alpha$. The communication complexity of the final simulation is $O(\text{IC}^{\circ}_{\mu}(\pi) \cdot \log(\text{CC}(\pi)/\alpha)/\alpha^2)$. \square

Next we prove the lemma.

Proof of Lemma 8.4. Every time a player wants to send a bit in π , she instead sends $k = \ell \frac{\log(\text{CC}(\pi)/\gamma)}{\beta^2}$ bits which are each independently and privately chosen to be the correct value with probability $1/2 + \beta$. The receiving player takes the majority of the bits sent to reconstruct the intended transmission. The players then proceed assuming that the majority of the bits was the real sampled transmission.

By the Chernoff bound, we can set ℓ to be large enough so that the probability that any transmission is received incorrectly is at most $\frac{\gamma}{\text{CC}(\pi)}$. By the union bound applied to each of the $\text{CC}(\pi)$ transmissions, we have that except with probability γ , all transmissions are correctly received. Thus the distribution of the simulated transcript is γ -close to the correct distribution. All that remains is to bound the probability of having a large divergence cost.

We denote by R the public randomness used in both protocols, by V_i the i th intended transmission in the protocol τ , and by W_i the block of k bits used to simulate the transmission of V_i . We use M_i to denote the majority of the bits W_i (which is the actual transmission). For every i , let G_i denote the event that $V_1, \dots, V_i = M_1, \dots, M_i$, namely, that the first i intended transmissions occurred as intended (set G_0 to be the event that is true always). Then for each i , conditioned on the event G_{i-1} , we have that $X, Y, V_i, R, M_1, \dots, M_{i-1}$ have the same distribution as $X, Y, \pi(X, Y)_i, R, \pi(X, Y)_1, \dots, \pi(X, Y)_{i-1}$. In particular, this implies that

$$(8.2) \quad I(XY; V_i | rm_1, \dots, m_{i-1} G_{i-1}) = I(XY; \pi(X, Y)_i | \pi(X, Y)_{\leq i-1} = rm_1, \dots, m_{i-1}).$$

On the other hand,

$$(8.3) \quad \begin{aligned} & I(XY; W_i | rm_1, \dots, m_{i-1} G_{i-1}) \\ & \leq I(XY; W_i V_i | rm_1, \dots, m_{i-1} G_{i-1}) \\ & = I(XY; V_i | rm_1, \dots, m_{i-1} G_{i-1}) + I(XY; W_i | rm_1, \dots, m_{i-1} V_i G_{i-1}) \\ & = I(XY; V_i | rm_1, \dots, m_{i-1} G_{i-1}). \end{aligned}$$

In the event G_{i-1} , after fixing $V_i, R, M_1, \dots, M_{i-1}$, we have that W_i is independent of the inputs XY . Thus (8.2) and (8.3) imply that

$$(8.4) \quad I(XY; \pi(X, Y)_i | \pi(X, Y)_{\leq i-1} = rm_1, \dots, m_{i-1}) \geq I(XY; W_i | rm_1, \dots, m_{i-1} G_{i-1}).$$

Let D_1, D_2, \dots , be random variables defined as follows. For each i , set

$$D_i = \begin{cases} 0 & \text{if the event } G_{i-1} \text{ does not hold,} \\ \sum_{j=1}^k \mathbb{D}_{X,ik+j}^\tau(\tau(X, Y)) + \mathbb{D}_{Y,ik+j}^\tau(\tau(X, Y)) & \text{otherwise.} \end{cases}$$

Thus we have that conditioned on the event $G_{CC(\pi)}$, $\sum_{i=1}^{CC(\pi)} D_i$ is equal to $\mathbb{D}_{XY}^\tau(\tau(X, Y))$. On the other hand, $\mathbb{E}[D_i] = \Pr[G_i] \mathbb{E}[D_i|G_i] \leq \mathbb{E}[D_i|G_i]$. But by (8.4) and Lemma 8.2, $\mathbb{E}[D_i|G_i]$ can be bounded by $\mathbb{E}_{X,Y}[\mathbb{D}_{X,i}^\pi(\pi(X, Y)) + \mathbb{D}_{Y,i}^\pi(\pi(X, Y))]$. So by linearity of expectation, $\mathbb{E}[\sum_{i=1}^{CC(\pi)} D_i] \leq \mathbb{E}_{X,Y}[\mathbb{D}_{XY}^\pi(\pi(X, Y))] = I(XY; \pi(X, Y)) = IC^\circ_\mu(\pi)$.

Thus, by the union bound,

$$\Pr_{X,Y}[\mathbb{D}_{XY}^\tau(\tau(X, Y)) > IC^\circ_\mu(\pi)/\gamma] \leq (1 - \Pr[G_{CC(\pi)}]) + \Pr\left[\sum_{i=1}^{CC(\pi)} D_i > IC^\circ_\mu(\pi)/\gamma\right].$$

We bound the first term by γ as above, and the second term by γ using Markov's inequality. \square

8.3. Proof of Theorem 8.5. It only remains to prove Theorem 8.5. Set $\beta = 1/k \log(CC(\pi)/\epsilon)$. We use the fact that the bits in our protocol are close to uniform to show that the step divergence is at most $O(\beta)$ for each step.

PROPOSITION 8.6. *For every j , $\mathbb{D}_{x,j}^\pi(v)$ and $\mathbb{D}_{y,j}^\pi(v)$ are bounded by $O(\beta)$.*

Proof. This follows from the fact that all probabilities for each step lie in $[1/2 - \beta, 1/2 + \beta]$. The worst the divergence between two distributions that lie in this range can be is clearly $\log(\frac{1/2+\beta}{1/2-\beta}) = \log(1 + O(\beta)) = O(\beta)$. \square

Next, for every prefix v of the transcript, and inputs x, y , we define a subset of the prefixes of potential transcripts that start with v , \mathcal{B}_{vxy} in the following way: we include w in \mathcal{B}_{vxy} if and only if for every w' that is a strict prefix of w ,

$$\max \left\{ \sum_{j=CC(v)+1}^{CC(w')} \mathbb{D}_{x,j}^\pi(w'), \sum_{j=CC(v)+1}^{\|w'\|} \mathbb{D}_{y,j}^\pi(w') \right\} < \beta,$$

and we have that w itself is either a leaf or satisfies

$$\max \left\{ \sum_{j=CC(v)+1}^{CC(w)} \mathbb{D}_{x,j}^\pi(w), \sum_{j=CC(v)+1}^{\|w\|} \mathbb{D}_{y,j}^\pi(w) \right\} \geq \beta.$$

The set \mathcal{B}_{vxy} has the property that every path from v to a leaf of the protocol tree must intersect exactly one element of \mathcal{B}_{vxy} , i.e., if we cut all paths at the point where they intersect \mathcal{B}_{vxy} , we get a protocol tree that is a subtree of the original tree.

We define the distribution B_{vxy} on the set \mathcal{B}_{vxy} as the distribution on \mathcal{B}_{vxy} induced by the protocol π : namely, we sample from B_{vxy} by sampling each subsequent vertex according to $O_{w,x,y}$ and then taking the unique vertex of \mathcal{B}_{vxy} that the sampled path intersects.

Similarly, we define the distribution B_{vx} on \mathcal{B}_{vxy} , to be the distribution obtained by following the edge out of each subsequent node w according to the distribution $O_{w,x}$, and the distribution B_{vy} by following edges according to $O_{w,y}$, and finally the distribution B_v by sampling edges according to O_w .

Protocol $\tau_{v,t}$	
1.	Both players use public randomness to sample a path according to $\pi(X, Y) v$ and communicate $2 \log \text{CC}(\pi)$ bits to sample an element w of \mathcal{B}_{vxy} according to the distribution B_v .
2.	P_x samples a bit a_1 which is 1 with probability $\min \left\{ \frac{B_{vx}(w)}{tB_v(w)}, 1 \right\}.$
3.	P_y samples a bit a_2 which is 1 with probability $\min \left\{ \frac{B_{vy}(w)}{tB_v(w)}, 1 \right\}.$
4.	If both a_1 and a_2 were 1, they accept w . Else they repeat the protocol.

FIG. 9. The protocol to sample a subsequent part of the transcript.

Observe that for every transcript w , the players can compute the element of \mathcal{B}_{vxy} that intersects the path w by communicating $2 \log \text{CC}(\pi)$ bits (the computation amounts to computing the minimum of two numbers of magnitude at most $\text{CC}(\pi)$). Given these definitions, we are now ready to describe our simulation protocol. The protocol proceeds in rounds. In each round the players shall use rejection sampling to sample some consecutive part of the transcript.

8.3.1. A single round. The first protocol, shown in Figure 9 assumes that we have already sampled the prefix v . We define the protocol for some constant t that we shall set later.

Observe that the distributions we have defined satisfy the equation

$$(8.5) \quad \left(\frac{B_{vx}}{B_v} \right) \left(\frac{B_{vy}}{B_v} \right) = \frac{B_{vxy}}{B_v}.$$

This suggests that our protocol should pick a transcript distributed according to B_{vxy} . We shall argue that the subsequent prefix of the transcript sampled by the protocol in Figure 9 cannot be sampled with much higher probability than what it is sampled with in the real distribution. Let B'_{vxy} denote the distribution of the accepted prefix of $\tau_{v,t}$.

CLAIM 8.7 (no sample gets undue attention). *For every prefix w ,*

$$B'_{vxy}(w)/B_{vxy}(w) \leq 1 + 2 \exp \left(-\Omega \left(\frac{(\log t - O(\beta))^2}{\beta} \right) \right).$$

We shall also show that the expected communication complexity of this protocol is not too high.

CLAIM 8.8 (small number of rounds). *The expected communication complexity of τ_v is at most*

$$\frac{O(t^2)}{1 - \exp \left(-\Omega \left(\frac{(\log t - O(\beta))^2}{\beta} \right) \right)}.$$

Claims 8.7 and 8.8 will follow from the following claim.

CLAIM 8.9.

$$\Pr_{w \in_R \mathcal{B}_{vxy}} \left[\frac{B_{vx}(w)}{B_v(w)} \geq t \right] \leq \exp \left(-\Omega \left(\frac{(\log t - O(\beta))^2}{\beta} \right) \right),$$

$$\Pr_{w \in_R \mathcal{B}_{vxy}} \left[\frac{B_{vy}(w)}{B_v(w)} \geq t \right] \leq \exp \left(-\Omega \left(\frac{(\log t - O(\beta))^2}{\beta} \right) \right).$$

Let us first argue that Claim 8.7 follows from Claim 8.9.

Proof of Claim 8.7. Set a to be the function that maps any $w \in \mathcal{B}_{vxy}$ to $\min\{(1/t) \frac{B_{vx}(w)}{B_v(w)}, 1\} \cdot \min\{(1/t) \frac{B_{vy}(w)}{B_v(w)}, 1\}$. Set $a' = (1/t) \frac{B_{vx}(w)}{B_v(w)} (1/t) \frac{B_{vy}(w)}{B_v(w)}$. Then clearly $a'(w) \geq a(w)$ for every w . Applying (8.5), we get

$$a' = (1/t^2) \left(\frac{B_{vx}}{B_v} \right) \left(\frac{B_{vy}}{B_v} \right) = (1/t^2) \frac{B_{vxy}}{B_v}.$$

Thus $B_{vxy} = \beta a' \cdot B_v$ for some constant β . By Proposition B.3, applied to a', a , and the distributions $D = B_{vxy}, D' = B'_{vxy}$, we have that for every w ,

$$\frac{B'_{vxy}(w)}{B_{vxy}(w)} \leq \frac{1}{1 - \Pr_{w \in_R \mathcal{B}_{vxy}} [a'(w) > a(w)]}.$$

On the other hand, by the union bound and Claim 8.9,

$$\Pr_{w \in_R \mathcal{B}_{vxy}} [a'(w) > a(w)] \leq \Pr_{w \in_R \mathcal{B}_{vxy}} \left[\frac{B_{vx}(w)}{B_v(w)} > t \vee \frac{B_{vy}(w)}{B_v(w)} > t \right]$$

$$\leq 2 \exp \left(-\Omega \left(\frac{(\log t - O(\beta))^2}{\beta} \right) \right).$$

Since $1/(1 - z) \leq 1 + O(z)$ for $z \in (0, 1/10)$, we get Claim 8.7. \square

Now we show Claim 8.8 assuming Claim 8.9.

Proof of Claim 8.8. We shall use Proposition B.4. We need to estimate the probability that the first round of $\tau_{v,t}$ accepts its sample. This probability is exactly

$$\sum_{w \in \mathcal{B}_{vxy}} B_v(w) \min \left\{ (1/t) \frac{B_{vx}(w)}{B_v(w)}, 1 \right\} \cdot \min \left\{ (1/t) \frac{B_{vy}(w)}{B_v(w)}, 1 \right\}.$$

Let $A \subset \mathcal{B}_{vxy}$ denote the set $\{w : \frac{B_{vx}(w)}{B_v(w)} \leq t \wedge \frac{B_{vy}(w)}{B_v(w)} \leq t\}$. Then we see that the above sum can be lower bounded:

$$\sum_{w \in \mathcal{B}_{vxy}} B_v(w) \min \left\{ (1/t) \frac{B_{vx}(w)}{B_v(w)}, 1 \right\} \cdot \min \left\{ (1/t) \frac{B_{vy}(w)}{B_v(w)}, 1 \right\}$$

$$\geq (1/t^2) \sum_{w \in A} B_v(w) \left(\frac{B_{vx}(w)}{B_v(w)} \right) \left(\frac{B_{vy}(w)}{B_v(w)} \right) = (1/t^2) \sum_{w \in A} B_{vxy},$$

where the last equality follows from (8.5).

Finally, we see that Claim 8.9 implies that $\sum_{w \in A} B_{vxy} \geq 1 - \exp(-\Omega(\frac{(\log t - O(\beta))^2}{\beta}))$. Proposition B.4 then gives the bound we need. \square

Next we prove Claim 8.9. To do this we shall need to use a generalization of Azuma's inequality, which we prove in Appendix A.

Proof of Claim 8.9. Let W be a random variable distributed according to B_{vxy} . Set $Z_{\text{CC}(v)+1}, \dots, Z_{\text{CC}(\pi)}$ to be real valued random variables such that if $i \leq \text{CC}(W)$,

$$Z_i = \log \left(\frac{O_{w_{\leq i-1}, x}(w_{\leq i})}{O_{w_{\leq i-1}}(w_{\leq i})} \right).$$

If $i > \text{CC}(W)$, set $Z_i = 0$. Observe that $\mathbb{E}[Z_i | w_{\leq i-1}] = \mathbb{D}_{x,i}^\pi(w)$. We also have that

$$\begin{aligned} \sum_{i=\text{CC}(v)+1}^{\text{CC}(\pi)} Z_i &= \sum_{i=\text{CC}(v)+1}^{\text{CC}(\pi)} \log \left(\frac{O_{w_{\leq i-1}, x}(w_{\leq i})}{O_{w_{\leq i-1}}(w_{\leq i})} \right) \\ (8.6) \qquad &= \log \left(\frac{B_{vx}(w)}{B_v(w)} \right). \end{aligned}$$

Next set $T_i = Z_i - \mathbb{E}[Z_i | Z_{i-1}, \dots, Z_1]$. Note that $\mathbb{E}[T_i | T_{i-1}, \dots, T_1] = 0$ (in fact the stronger condition that $\mathbb{E}[T_i | Z_{i-1}, \dots, Z_1] = 0$ holds). For every $w \in \mathcal{B}_{vxy}$, we have that

$$\begin{aligned} &\sup(T_i | w_{\leq i-1}) \\ &\leq \max \left\{ \log \left(\frac{\Pr[\pi(X, Y)_i = 0 | w_{\leq i-1}x]}{\Pr[\pi(X, Y)_i = 0 | w_{\leq i-1}]} \right), \log \left(\frac{\Pr[\pi(X, Y)_i = 1 | w_{\leq i-1}x]}{\Pr[\pi(X, Y)_i = 1 | w_{\leq i-1}]} \right) \right\}, \\ &\inf(T_i | w_{\leq i-1}) \\ &\geq \min \left\{ \log \left(\frac{\Pr[\pi(X, Y)_i = 0 | w_{\leq i-1}x]}{\Pr[\pi(X, Y)_i = 0 | w_{\leq i-1}]} \right) \right. \\ &\quad \left. - \mathbb{D}_{x,i}^\pi(w), \log \left(\frac{\Pr[\pi(X, Y)_i = 1 | w_{\leq i-1}x]}{\Pr[\pi(X, Y)_i = 1 | w_{\leq i-1}]} \right) - \mathbb{D}_{x,i}^\pi(w) \right\}. \end{aligned}$$

By Proposition 4.8 and using the fact that $\pi(x, Y) = 1 \in [1/2 - \beta, 1/2 + \beta]$ we can bound

$$\begin{aligned} \sup(T_i | w_{\leq i-1}) &\leq \log \left(\frac{1/2 - \beta + \sqrt{\mathbb{D}_{x,i}^\pi(w)}}{1/2 - \beta} \right) \\ &= \log \left(1 + O \left(\sqrt{\mathbb{D}_{x,i}^\pi(w)} \right) \right) \\ (8.7) \qquad &= O \left(\sqrt{\mathbb{D}_{x,i}^\pi(w)} \right), \end{aligned}$$

$$\begin{aligned} \inf(T_i | w_{\leq i-1}) &\geq \log \left(\frac{1/2 - \beta}{1/2 - \beta + \sqrt{\mathbb{D}_{x,i}^\pi(w)}} \right) - \mathbb{D}_{x,i}^\pi(w) \\ &= \log \left(1 - O \left(\sqrt{\mathbb{D}_{x,i}^\pi(w)} \right) \right) \\ (8.8) \qquad &= -O \left(\sqrt{\mathbb{D}_{x,i}^\pi(w)} \right), \end{aligned}$$

as long as $\beta < 1/10$.

Protocol τ_t
<ol style="list-style-type: none"> 1. The players publicly sample the public randomness $v \in_R R$ for π. 2. The players repeatedly run $\tau_{v,t}$ to get a new prefix v. They stop only when they reach a leaf of the protocol tree for π.

FIG. 10. The protocol to sample a subsequent part of the transcript.

Equations (8.7) and (8.8) imply that for $w \in \mathcal{B}_{vxy}$,

$$(8.9) \quad \sum_{i=\text{CC}(v)+1}^{\text{CC}(\pi)} (\sup(T_i) - \inf(T_i)|_{w_{\leq i-1}})^2 \leq \sum_{i=\text{CC}(v)+1}^{\text{CC}(\pi)} O(\mathbb{D}_{x,i}^\pi(w)) = O(\beta).$$

For every w ,

$$(8.10) \quad \begin{aligned} \left(\sum_{i=\text{CC}(v)+1}^{\text{CC}(\pi)} T_i \right) |w &= \left(\sum_{i=\text{CC}(v)+1}^{\text{CC}(\pi)} Z_i \right) |w - \sum_{i=\text{CC}(v)+1}^{\text{CC}(\pi)} \mathbb{E}[Z_i | w_{\leq i-1}] \\ &= \sum_{i=\text{CC}(v)+1}^{\text{CC}(w)} \log \left(\frac{\Pr[\pi(x, Y)_i = w_i | v\pi(x, Y)_{\leq i-1} = w_{\leq i-1}]}{\Pr[\pi(X, Y)_i = w_i | \pi(X, Y)_{\leq i-1} = vw_{\leq i-1}]} \right) \\ &\quad - \sum_{i=\text{CC}(v)+1}^{\text{CC}(w)} \mathbb{D}_{x,i}^\pi(w) \\ &\geq \log \left(\frac{B_{vx}(w)}{B_v(w)} \right) - O(\beta), \end{aligned}$$

where the last inequality follows from the definition of \mathcal{B}_{vxy} , Proposition 8.6, and (8.6).

Thus we can use Theorem A.1 to bound

$$\begin{aligned} &\Pr_{w \in_R \mathcal{B}_{vxy}} \left[\frac{B_{vx}(w)}{B_v(w)} \geq t \right] \\ &\leq \Pr_{w \in_R \mathcal{B}_{vxy}} \left[\log \left(\frac{B_{vx}(w)}{B_v(w)} \right) \geq \log t \right] \\ &\leq \Pr \left[\sum_{i=\text{CC}(v)+1}^{\text{CC}(\pi)} T_i \geq \log t - O(\beta) \right] \tag{by (8.10)} \\ &\leq \exp \left(-\Omega \left(\frac{(\log t - O(\beta))^2}{\sum_{i=\text{CC}(v)+1}^{\text{CC}(\pi)} (\sup(T_i) - \inf(T_i)|_{w_{\leq i-1}})^2} \right) \right) \\ &\leq \exp \left(-\Omega \left(\frac{(\log t - O(\beta))^2}{\beta} \right) \right) \tag{by (8.9)}. \quad \square \end{aligned}$$

8.3.2. The whole protocol. Our final protocol for computing f is shown in Figure 10.

We first argue that our simulation returns the correct answer with decent probability. We shall actually argue that the probability for any returned transcript does not increase by too much. To ease notation, let us set

$$\alpha \stackrel{\text{def}}{=} \exp \left(-\Omega \left(\frac{(\log t - O(\beta))^2}{\beta} \right) \right).$$

Set t to be a large enough constant so that

$$\alpha = \exp(-\Omega(1/\beta)) = \exp(-\Omega(k \log(\text{CC}(\pi)/\epsilon))).$$

Let L denote the random variable of the sampled transcript returned by τ_t . Then by Claim 8.7, we get that for every leaf l ,

$$(8.11) \quad \frac{\Pr[L = l|xy]}{\Pr[\pi(x, y) = l]} \leq (1 + \alpha)^{\text{CC}(\pi)} = \exp(O(\epsilon)),$$

and we can set k to be large enough so that $\frac{\Pr[L=l|xy]}{\Pr[\pi(x,y)=l]} \leq 1 + \epsilon/2$. Thus, the leaf sampled by our protocol $\epsilon/2$ is close in statistical distance to the leaf sampled by π .

Observe that if the protocol accepts a leaf l , then the protocol must have involved $s((\mathbb{D}_x^\pi(l) + \mathbb{D}_y^\pi(l))/\beta)$ rounds, for some constant s .

The expected number of bits communicated in each of these rounds is independent of l by Proposition B.1, and is $\frac{t^2}{1-\alpha}$ by Claim 8.8. Let G be the event that $s(\mathbb{D}_x^\pi(L) + \mathbb{D}_y^\pi(L))/\beta \leq \epsilon Q$. Thus, by Markov's inequality, conditioned on G , the expected communication complexity of the protocol is $\epsilon Q(t^2/(1-\alpha)) \leq d\epsilon Q$, for some constant d . By the union bound, we get that the probability that the communication exceeds Q is bounded by $(1 - \Pr[G]) + d\epsilon \leq \Pr[\mathbb{D}_{XY}^\pi(\pi(X, Y)) > \beta\epsilon Q/s] + d\epsilon$. We get our final protocol by terminating the simulation if the communication exceeds Q . The final output is thus $\Pr[\mathbb{D}_{XY}^\pi(\pi(X, Y)) > \beta\epsilon Q/s] + (d+1)\epsilon$ close to the correct distribution.

This completes the proof of Theorem 8.5.

9. Open problems and final thoughts. The main problem that remains open is whether optimal, or near-optimal compression is possible for protocols in the general setting.

Open problem: Is there a generic way to convert any two-party protocol π over a general distribution μ into a protocol that uses only $\text{IC}_\mu^1(\pi) \text{polylog}(\text{CC}(\pi))$ bits of communication?

An affirmative answer to this problem would immediately yield an optimal direct-sum theorem for randomized communication complexity, showing that the communication complexity of f^n is $\tilde{O}(n)$ times as high as the communication complexity of f . Curiously enough, it turns out [5] that the converse is true as well, and the problem above is *complete* for randomized communication direct sum—one can show that if there is no such compression scheme, then there is a (partial) function U for which a direct-sum theorem fails to hold.⁶ In this function U , each player gets as input a protocol tree, as well as the probabilities for all the nodes he owns, and the output is simply the output of the protocol. Unfortunately, by design, information theoretic techniques seem to be powerless in proving lower bounds for U .

Appendix A. A simple generalization of Azuma's inequality.

We shall need the following theorem, whose proof appears in [17]. For completeness, we reproduce the part of the proof we need here:

THEOREM A.1 (Azuma). *Let T_1, \dots, T_k be real valued random variables such that for every i , we have $\mathbb{E}[T_i|T_{i-1}, \dots, T_i] \leq 0$. Set $A_i = (\sup(T_i) - \inf(T_i)|T_{i-1}, \dots, T_1)^2$.*

⁶In a partial function/promise problem the protocol only needs to compute the function if the pair of inputs comes from some subset. Our results in this paper for information measured according to the viewpoint of the players carry over to promise problems as well.

Then if $\sum_{i=1}^k A_i \leq c$, for every $\alpha > 0$,

$$\Pr \left[\sum_{i=1}^k T_i \geq \alpha \right] \leq \exp(-2\alpha^2/c).$$

To prove the theorem, we need the following lemma appearing as Lemma 2.6 in [17].

LEMMA A.2. *Let X be a real valued random variable with $\mathbb{E}[X] = 0$ and $X \in [a, b]$ almost surely. Then $\mathbb{E}[\exp(X)] \leq \exp(\frac{(b-a)^2}{8})$.*

Proof of Theorem A.1. First, we assume without loss of generality that $\mathbb{E}[T_i|T_{i-1}, \dots, T_1] \leq 0$. We can do this by changing each random variable T_i to $T_i - \mathbb{E}[T_i|T_{i-1}, \dots, T_1]$. This does not change any of the conditions above, and only increases $\Pr[\sum_i T_i \geq \alpha]$.

By Markov's inequality, for every positive λ we have

$$\Pr \left[\sum_{i=1}^k T_i \geq \alpha \right] = \Pr \left[\exp \left(\lambda \sum_{i=1}^k T_i \right) > \exp(\lambda\alpha) \right] \leq \mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^k T_i \right) \right] \exp(-\lambda\alpha).$$

Next we show by induction on k that $\mathbb{E}[\exp(\lambda \sum_{i=1}^k T_i)] \leq \sup(\prod_{i=1}^k \mathbb{E}[\exp(\lambda T_i)|T_{i-1}, \dots, T_1])$. The case $k = 1$ is trivial. For general k we compute

$$\begin{aligned} \mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^k T_i \right) \right] &= \mathbb{E} \left[\exp(\lambda T_1) \mathbb{E} \left[\exp \left(\lambda \sum_{i=2}^k T_i \right) | T_1 \right] \right] \\ &\leq \mathbb{E}[\exp(\lambda T_1)] \sup \left(\prod_{i=2}^k \mathbb{E}[\exp(\lambda T_i)|T_{i-1}, \dots, T_1] \right) \\ &\hspace{15em} \text{(by induction)} \\ &= \sup \left(\mathbb{E}[\exp(\lambda T_1)] \prod_{i=2}^k \mathbb{E}[\exp(\lambda T_i)|T_{i-1}, \dots, T_1] \right) \\ &= \sup \left(\prod_{i=1}^k \mathbb{E}[\exp(\lambda T_i)|T_{i-1}, \dots, T_1] \right). \end{aligned}$$

Thus we can bound

$$\begin{aligned} \Pr \left[\sum_{i=1}^k T_i \geq \alpha \right] &\leq \exp(-\lambda\alpha) \sup \left(\prod_{i=1}^k \mathbb{E}[\exp(\lambda T_i)|T_{i-1}, \dots, T_1] \right) \\ &\leq \exp(-\lambda\alpha) \sup \left(\prod_{i=1}^k \exp \left(\frac{\lambda^2 A_i}{8} \right) \right) \hspace{5em} \text{(by Lemma A.2)} \\ &= \exp(-\lambda\alpha) \sup \left(\exp \left(\frac{\sum_{i=1}^k \lambda^2 A_i}{8} \right) \right) \\ &= \exp(-\lambda\alpha) \exp \left(\sup \left(\frac{\sum_{i=1}^k \lambda^2 A_i}{8} \right) \right) \\ &\leq \exp(-\lambda\alpha + \lambda^2 c/8). \end{aligned}$$

Algorithm Rejection Sampling.
1. Sample an element $z \in_{\mathcal{R}} C$.
2. Accept it with probability $a(z)$, else go to the first step.

FIG. 11. *Generic rejection sampling.*

Setting $\lambda = 4\alpha/c$, we get that

$$\Pr \left[\sum_{i=1}^k T_i \geq \alpha \right] \leq \exp(-2\alpha^2/c). \quad \square$$

Appendix B. Analyzing rejection sampling. In this section we give some basic facts about *rejection sampling*. For a distribution C supported on some finite set \mathcal{C} and a function $a : \mathcal{C} \rightarrow [0, 1]$, Figure 11 describes a generic rejection sampling algorithm.

We prove some simple properties of this kind of sampling. Let D' denote the random variable of the sampled element. Let R denote the random variable that counts the number of rounds before the algorithm accepts the sample. Then we see that D' is independent of R , since for any integers c, c' , $D'|R = c$ has the same distribution as $D'|R = c'$.

PROPOSITION B.1. *D' is independent of R .*

We then see that $D'(w) = \Pr[(R = 1) \wedge w \text{ is accepted}] / \Pr[R = 1] = C(w)a(w) / \Pr[R = 1]$. We have shown the following claim.

CLAIM B.2. *For some constant α , $D' = \alpha a \cdot C$.*

Now let $a' : \mathcal{C} \rightarrow [0, 1]$ be a function such that $a'(w) \geq a(w)$ for all $w \in \mathcal{C}$, and let D denote the random variable of the sampled element. Set $b = a' - a$. Then $D = \beta a' \cdot C = \beta C \cdot (a + b)$ for some $\beta > 0$. Thus, by Claim B.2, there exists a distribution D'' such that D' is a convex combination $D = \beta' D'' + (1 - \beta') D'$. In particular, this implies that $\frac{D'(w)}{D(w)} \leq \frac{1}{1 - \beta'}$. We bound $\beta' \leq \Pr[D' \in \text{supp}(D'')] = \Pr_{w \in_{\mathcal{R}} D}[a'(w) > a(w)]$. This gives us the following two bounds.

PROPOSITION B.3. *Let $D = \beta a' \cdot C$ be a distribution such that $a'(w) \geq a(w)$ for every w . Then for every w ,*

$$\frac{D'(w)}{D(w)} \leq \frac{1}{1 - \Pr_{w \in_{\mathcal{R}} D}[a'(w) > a(w)]}.$$

PROPOSITION B.4. *The expected number of rounds that the above protocol runs for is $1 / \Pr[R = 1]$.*

Proof. From the construction, we see that $\mathbb{E}[R] = \Pr[R = 1] + (1 - \Pr[R = 1])(1 + \mathbb{E}[R])$. Rewriting this, we get $\mathbb{E}[R] = 1 / \Pr[R = 1]$. \square

Appendix C. Finding the first difference in inputs.

Proof Sketch for Lemma 4.14. Without loss of generality, we assume that $k = 2^t$ for an integer t (if not, we can always pad the input strings with 0's until the lengths are of this form before running the protocol). For a parameter C , we define a labeled tree of depth $C \log(k/\epsilon) = C(t + \log(1/\epsilon))$ as follows. The root of the tree is labeled by the interval $[1, 2^t]$. For i ranging from 0 to $t - 1$, every node at depth i labeled by $[a, b]$ has two children, corresponding to splitting the interval $[a, b]$ into equal parts. Thus the left one is labeled by the interval $[a, b - 2^{t-i+1}]$ and the right one is labeled

by $[a + 2^{t-i+1}, b]$. Thus at depth t there are 2^t nodes, each labeled by $[a, a]$ for distinct a 's from $[2^t]$. Every node at depth $\geq t$ has exactly one child, labeled the same as the parent.

In the protocol, the players shall try to narrow down where the first difference in their inputs is by taking a walk on the tree. At each step, the players first check that the interval they are on is correct, and then try to narrow down their search. For any integer $a \in [k]$, let x_a denote the prefix of x of length a . To check whether a given interval $[a, b]$ contains the index that they seek, the players will use public randomness to pick random functions $h_1 : \{0, 1\}^a \rightarrow [18]$ and $h_2 : \{0, 1\}^b \rightarrow [18]$ and compare $h_1(x_a)$ with $h_1(y_a)$ and $h_2(x_b)$ with $h_2(y_b)$. The probability of getting an incorrect answer is thus at most $1/9$.

For a parameter C , the protocol works as follows:

1. the players set v to be the root of the tree;
2. the players run the tests described above to check whether the index with the first difference lies in the interval corresponding to v and in those corresponding to v 's children. If the tests are consistent, and indicate that the interval for v does not contain the index, the players set v to be the parent of the old v (or leave it unchanged if v is the root). If the tests are consistent and indicate that the interval of one of the children contains the index, the players set v to be that child. If the tests are inconsistent, the players leave v unchanged;
3. step 2 is repeated $C(t + \log(1/\epsilon))$ times;
4. if the final vertex is labeled by an interval of the form $[a, a]$, output a . Else conclude that the input strings are equal.

To analyze the protocol, fix x and y . Note that if $x = y$, then the protocol never fails. So let us assume that $x \neq y$ and assume that a is the first index at which x, y differ. Then let w denote the vertex in the tree of largest depth that is labeled by $[a, a]$. Next we direct the edges of the tree so that at every vertex, the only outgoing edge points to the neighbor that is closer to w in terms of shortest path distance. Then observe that at every step of our protocol, v is changed to a neighbor that is closer to w with probability at least $2/3$. Further, our protocol succeeds as long as the number of correct steps on the tree exceeds the number of incorrect steps by t . This happens as long as the number of correct steps is at least $C/2(t + \log(1/\epsilon)) + t/2$. Since the expected number of correct steps is $2C/3(t + \log(1/\epsilon))$, we get that the bad event happens only when we deviate from the expected number by $C/6(t + \log(1/\epsilon)) - t/2 > (C/6 - 1/2)(t + \log(1/\epsilon))$. By the Chernoff bound, the probability that this happens is at most $\exp(\Omega((C/6 - 1/2)^2(t + \log(1/\epsilon))))$. Setting C to be a large enough constant makes this error at most ϵ . \square

Acknowledgments. We thank Noga Alon, Emanuel Milman, Alex Samorodnitsky, Avi Wigderson, and Amir Yehudayoff for useful discussions.

REFERENCES

- [1] F. ABLAYEV, *Lower bounds for one-way probabilistic communication complexity*, in Proceedings of the 20th International Colloquium on Automata, Languages, and Programming, Lecture Notes in Comput. Sci. 700, A. Lingas, R. Karlsson, and S. Carlsson, eds., Springer-Verlag, Berlin, 1993, pp. 241–252.
- [2] R. BAR-YEHUDA, B. CHOR, E. KUSHILEVITZ, AND A. ORLITSKY, *Privacy, additional information and communication*, IEEE Trans. Inform. Theory, 39 (1993), pp. 1930–1943.

- [3] Z. BAR-YOSSEF, T. S. JAYRAM, R. KUMAR, AND D. SIVAKUMAR, *An information statistics approach to data stream and communication complexity*, J. Comput. System Sci., 68 (2004), pp. 702–732.
- [4] M. BRAVERMAN, A. RAO, O. WEINSTEIN, AND A. YEHUDAYOFF, *Direct Products In Communication Complexity*, Technical report TR12-143, Electronic Colloquium on Computational Complexity, 2012.
- [5] M. BRAVERMAN AND A. RAO, *Information equals amortized communication*, in FOCS 2011, R. Ostrovsky, ed., IEEE, Piscataway, NJ, 2011, pp. 748–757.
- [6] A. CHAKRABARTI, Y. SHI, A. WIRTH, AND A. YAO, *Informational complexity and the direct sum problem for simultaneous message complexity*, in Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science, B. Werner, ed., IEEE Computer Society, Los Alamitos, CA, 2001, pp. 270–278.
- [7] T. M. COVER AND J. A. THOMAS, *Elements of Information Theory*, Wiley Ser. Telecomm., Wiley, New York, 1991.
- [8] T. FEDER, E. KUSHILEVITZ, M. NAOR, AND N. NISAN, *Amortized communication complexity*, SIAM J. Comput., 24 (1995), pp. 736–750.
- [9] U. FEIGE, P. RAGHAVAN, D. PELEG, AND E. UPFAL, *Computing with noisy information*, SIAM J. Comput., 23 (1994), pp. 1001–1018.
- [10] G. GALBIATI AND M. FISCHER, *On the complexity of 2-output Boolean networks*, Theoret. Comput. Sci., 16 (1981), pp. 177–185.
- [11] P. HARSHA, R. JAIN, D. A. MCALLESTER, AND J. RADHAKRISHNAN, *The communication complexity of correlation*, in IEEE Conference on Computational Complexity, IEEE Computer Society, Los Alamitos, CA, 2007, pp. 10–23.
- [12] T. HOLENSTEIN, *Parallel repetition: Simplifications and the no-signaling case*, in Proceedings of the 39th Annual ACM Symposium on Theory of Computing, ACM, New York, 2007, pp. 411–419.
- [13] R. JAIN, J. RADHAKRISHNAN, AND P. SEN, *A direct sum theorem in communication complexity via message compression*, in ICALP, Lecture Notes in Comput. Sci. 2719, J. C. M. Baeten, J. K. Lenstra, J. Parrow, and G. J. Woeginger, eds., Springer, Berlin, 2003, pp. 300–315.
- [14] R. JAIN, J. RADHAKRISHNAN, AND P. SEN, *Prior entanglement, message compression and privacy in quantum communication*, in IEEE Conference on Computational Complexity, IEEE Computer Society, Los Alamitos, CA, 2005, pp. 285–296.
- [15] R. JAIN, P. SEN, AND J. RADHAKRISHNAN, *Optimal Direct Sum and Privacy Trade-Off Results for Quantum and Classical Communication Complexity*, preprint, ArXiv:0807.1267v.1 [cs.DC], 2008.
- [16] R. JAIN AND P. YAO, *A Strong Direct Product Theorem in Terms of the Smooth Rectangle Bound*, preprint, arXiv:1209.0263v.1 [cs.CC], 2012.
- [17] M. JERRUM, M. HABIB, C. MCDIARMID, J. L. RAMIREZ-ALFONSIN, AND B. REED, *Probabilistic Methods for Algorithmic Discrete Mathematics*, Algorithms Combin. 16, Springer, Berlin, 1998.
- [18] B. KALYANASUNDARAM AND G. SCHNITGER, *The probabilistic communication complexity of set intersection*, SIAM J. Discrete Math., 5 (1992), pp. 545–557.
- [19] M. KARCHMER, R. RAZ, AND A. WIGDERSON, *Super-logarithmic depth lower bounds via the direct sum in communication complexity*, Comput. Complexity, 5 (1995), pp. 191–204.
- [20] I. KERENIDIS, S. LAPLANTE, V. LERAYS, J. ROLAND, AND D. XIAO, *Lower Bounds on Information Complexity via Zero-Communication Protocols and Applications*, Technical report TR12-038, Electronic Colloquium on Computational Complexity, 2012.
- [21] H. KLAUCK, *A strong direct product theorem for disjointness*, in STOC 2010, ACM, New York, 2010, pp. 77–86.
- [22] J. M. KLEINBERG AND É. TARDOS, *Approximation algorithms for classification problems with pairwise relationships: Metric labeling and markov random fields*, J. ACM, 49 (2002), pp. 616–639.
- [23] E. KUSHILEVITZ AND N. NISAN, *Communication Complexity*, Cambridge University Press, Cambridge, 1997.
- [24] T. LEE, A. SHRAIBMAN, AND R. SPALEK, *A direct product theorem for discrepancy*, in Conference in Computational Complexity, IEEE Computer Society, Low Alamitos, CA, 2008, pp. 71–80.
- [25] I. NEWMAN, *Private vs. common random bits in communication complexity*, Inform. Process. Lett., 39 (1991), pp. 67–71.
- [26] W. PAUL, *Realizing Boolean functions on disjoint sets of variables*, Theoret. Comput. Sci., 2 (1976), pp. 383–396.
- [27] RAZBOROV, *On the distributed complexity of disjointness*, Theoret. Comput. Sci., 106 (1992).

- [28] R. RAZ, *A parallel repetition theorem*, SIAM J. Comput., 27 (1998), pp. 763–803.
- [29] M. E. SAKS AND X. SUN, *Space lower bounds for distance approximation in the data stream model*, in STOC 2002, ACM, New York, 2002, pp. 360–369.
- [30] R. SHALTIEL, *Towards proving strong direct product theorems*, Comput. Complexity, 12 (2003), pp. 1–22.
- [31] C. E. SHANNON, *A mathematical theory of communication*, Bell System Tech. J. 27 (1948), pp. 379–423.
- [32] A. A. SHERSTOV, *Strong direct product theorems for quantum communication and query complexity*, in STOC 2011, ACM, New York, 2011, pp. 41–50.
- [33] D. SLEPIAN AND J. K. WOLF, *Noiseless coding of correlated information sources*, IEEE Trans. Inform. Theory, 19 (1973), pp. 471–480.
- [34] D. UHLIG, *On the synthesis of self-correcting schemes from functional elements with a small number of reliable elements*, Mat. Zametki, 15 (1974), pp. 937–944.
- [35] A. D. WYNER AND J. ZIV, *The rate-distortion function for source coding with side information at the decoder*, IEEE Trans. Inform. Theory, 22 (1976), pp. 1–10.
- [36] A. C.-C. YAO, *Theory and applications of trapdoor functions (extended abstract)*, in Symposium on Foundations of Computer Science, IEEE Computer Society, Los Alamitos, CA, 1982, pp. 80–91.