Interactive Compression for Product Distributions

Gillat Kol *

Institute for Advanced Study, Princeton, NJ, USA

gillat.kol@gmail.com

ABSTRACT

We study the interactive compression problem: Given a two-party communication protocol with small information cost, can it be compressed so that the total number of bits communicated is also small? We consider the case where the parties have inputs that are independent of each other, and give a simulation protocol that communicates $I^2 \cdot \text{polylog}(I)$ bits, where I is the information cost of the original protocol. Our protocol is the first simulation protocol whose communication complexity is bounded by a polynomial in the information cost of the original protocol.

Categories and Subject Descriptors

F.2

[Analysis of Algorithms and Problem Complexity]:
Miscellaneous

Keywords

communication complexity; interactive compression;
information theory; information complexity

1. INTRODUCTION

In seminal works, Shannon, Fano and Huffman consider the *data compression problem*: Alice wants to send a message x to Bob. How many bits does she need to send, so that Bob will be able to retrieve x with high probability? The answer given in [Sha48, Fan49, Huf52] is that she needs to send only $\lceil \mathbb{H}(x) \rceil$ bits, in expectation, where $\mathbb{H}$ denotes Shannon's entropy function. Roughly speaking, this means that every message can be compressed to its information content.

*Research supported by the National Science Foundation grant No. CCF-1412958 and the Weizmann Institute of Science National Postdoctoral Award Program for Advancing Women in Science.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

STOC'16, June 19–21, 2016, Cambridge, MA, USA

ACM. 978-1-4503-4132-5/16/06...\$15.00

<http://dx.doi.org/10.1145/2897518.2897537>

Over the last decades, interactive communication protocols were studied extensively. The *interactive compression problem* [BBCR10] is the analog of the data compression problem in the interactive setting. It asks whether the transcript of any interactive protocol can be compressed to its information content.

The interactive compression problem is formalized in the setting of communication complexity. In the two-party distributional communication complexity model, each player gets an input, where the inputs are sampled from a joint distribution μ that is known to both players. The players' engage in an interactive communication protocol in order to perform some communication task that depends on both inputs. They may use both common and private random strings. The players communicate in rounds, where in each round one of the players sends a message to the other player. The communication complexity of the protocol is the total number of bits communicated by the two players.

The notion of information cost, introduced by [CSWY01, BYJKS04, BBCR10], measures the amount of information that the players need to reveal about their inputs in order to solve a communication task. This notion was motivated by the fundamental information theoretic question of compressing communication, as well as by fascinating relations to communication complexity, and in particular to the direct sum problem in communication complexity. Roughly speaking, the *(internal) information cost* of a protocol π over the distribution μ , denoted $\text{IC}_\mu(\pi)$, is the number of information bits that the players learn about each other's input, when running the protocol π with inputs sampled from μ .

Using the notions of communication complexity and information cost, the interactive compression problem is stated as follows: Given a communication protocol π , and a distribution μ over the inputs for π , does there exist a protocol π' that simulates π over μ and has communication complexity close to $\text{IC}_\mu(\pi)$ (say, polynomial in $\text{IC}_\mu(\pi)$)? By "simulates" we mean that π' performs the same task as π , except with some small error probability, where the probability is over μ and over the randomness used by the players.

1.1 Previous Works

The interactive compression problem received a lot of attention over the last years, and clever compression protocols were suggested for several cases [HJMR10, BBCR10, BR11, Bra12, BBK⁺13, BMY15, RR15]. Most relevant to our work is a beautiful result by Barak, Braverman, Chen and Rao, showing that over a product distribution μ , any protocol with information cost I and communication complexity C , can be compressed to a protocol with communication complexity $I \cdot \text{polylog}(C)$ [BBCR10]. (We mention that their result is even stronger and gives an $I' \cdot \text{polylog}(C)$ compression for any protocol over any distribution, where I' is the *external* information cost). In [BMY15], a simulation protocol communicating $O(I^2 \cdot \log \log(C))$ bits is shown for the case where the original protocol does not use private randomness. The general case, compressing any protocol over any distribution, is considered in [Bra12], where a $2^{O(I)}$ compression is given, and in [BBCR10], where a $\tilde{O}(\sqrt{C \cdot I})$ compression is given. A compression to $\tilde{O}(I + r)$ is given by [BR11], where r is the number of rounds of the original protocol.

Another line of works studies the limitations of compression protocols. In [GKR14, GKR15b, RS15, GKR15a], exponential separations between communication complexity and information cost are shown. More precisely, for any sufficiently large I , the above results give examples of communication tasks with communication complexity at least 2^I , for which there exist protocols with information cost $O(I)$. Therefore, over general distributions μ , communication protocols cannot always be compressed to their information cost. We mention that in all the above results, the protocol with low information cost has communication complexity C , that is at least double exponential in I . Thus, the $I \cdot \text{polylog}(C)$ compression by [BBCR10], leaves open the possibility that a similar separation can also be obtained with the underlying distribution being a product distribution.

1.2 Our Result

We consider the interactive compression problem over product distributions $\mu = \mu^A \times \mu^B$ (i.e., Alice's input and Bob's input are chosen independently), and give a simulation protocol that communicates $I^2 \cdot \text{polylog}(I)$ bits, where I is the information cost of the original protocol over μ . Our protocol is the first simulation protocol whose communication complexity is polynomial in the information cost of the original protocol, and does not depend on the communication complexity of the original protocol. We mention that our result implies that an exponential separation between information cost and communication complexity, like the ones given by [GKR14, GKR15b, RS15, GKR15a], cannot be obtained for product distributions. Our result is incomparable to the $I \cdot \text{polylog}(C)$ compression protocol of [BBCR10].

Our main result is given in Theorem 1 below. The theorem uses the following notation: Let ξ be a randomized communication protocol between two parties, Alice and Bob. Let x be an input for Alice, and y be an input for Bob.

We denote by $\xi(x, y)$ the random variable representing ξ 's transcript (that is, the concatenation of all the messages exchanged between the players during the execution of ξ), when it is run with the input (x, y) . We also view $\xi(x, y)$ as a distribution over the set of all binary strings. For a pair of distributions D, D' , we denote by $\|D - D'\|$ the statistical distance between D and D' .

Theorem 1. *Let $\varepsilon > 0$. Let π be a randomized protocol that may use private and public coins. Let $\mu = \mu^A \times \mu^B$ be a product distribution over the inputs for π . Then, there exist a public coin protocol τ (that takes the same inputs as π), and a pair of “transcript reconstruction” functions g^A, g^B such that: The function g^A takes as inputs $x \in \text{supp}(\mu^A)$ and a possible transcript of τ , and returns a possible transcript of π . The function g^B takes as inputs $y \in \text{supp}(\mu^B)$ and a possible transcript of τ , and returns a possible transcript of π . In addition, the followings hold:*

1. *The worst case communication complexity of τ is $\text{IC}_\mu^2(\pi) \cdot \text{polylog}(\text{IC}_\mu(\pi)) / \varepsilon^5$.*
2. $\forall (x, y) \in \text{supp}(\mu) : \Pr [g^A(x, \tau(x, y)) \neq g^B(y, \tau(x, y))] \leq \varepsilon$, *where the probability is over the random coins of the protocol τ .*
3. $\mathbf{E}_{(x, y) \leftarrow \mu} [\|g^A(x, \tau(x, y)) - \pi(x, y)\|] \leq \varepsilon$.

2. PROOF SKETCH

Let π be a communication protocol between two players, Alice and Bob. Alice has a private input x and Bob has a private input y , where (x, y) is chosen according to some publicly known joint distribution μ . For the rest of this sketch, $\mu = \mu^A \times \mu^B$ is a product distribution. We assume, without loss of generality, that π does not use public randomness (but may use private randomness), as the public randomness can always be replaced by private randomness without increasing the information cost. We also assume that the players take turns sending bits to one another. Alice sends bits in odd rounds and Bob in the even rounds.

The *external information cost* of a protocol over a distribution μ is the number of information bits that an external observer, who watches the execution of the protocol, learns about the players' inputs, when the inputs are sampled from μ . We next present a sketch of the public coin protocol τ that simulates π , such that the communication complexity of τ is $I^2 \cdot \text{polylog}(I)$, where I is the external information of π over μ . This proves Theorem 1, as over a product distribution μ , the internal and external information costs of a protocol π are the same (see Fact 3). Intuitively, this is true since x conveys no information about y , thus Alice and an external observer who doesn't know x, y have the same information about y at any point in the execution of the protocol. The same is true for Bob. Our protocol builds over the $I \cdot \text{polylog}(C)$ compression by [BBCR10], and parts of this sketch follow their description.

2.1 Divergence Tree

Consider the (directed) binary tree associated with π , where each vertex v of the binary tree corresponds to a possible transcript of π . The two edges going out of v are labeled by 0 and 1, corresponding to the next bit to be transmitted. We think of Alice as owning the non-leaf vertices in the odd layers, and of Bob as owning the non-leaf vertices in the even layers. The protocol π proceeds as follows: Starting from the root, when π reaches a non-leaf vertex v , the player who owns v sends a bit to the other player. The players follow the edge indicated by the sent bit and reach a new vertex.

Let $b \in \{0, 1\}$. For every non-leaf vertex w , we denote by $O_w(b)$ the probability of transmitting the bit b at w , conditioned on reaching that vertex (without taking into consideration the actual values of the inputs x, y). We denote by $O_{wx}(b)$, $O_{wy}(b)$ the probabilities of transmitting the bit b at w , conditioned on a particular fixing of x or y (respectively), and conditioned on the event of reaching w during the run of the protocol. We denote by $O_{wxy}(b)$ the probability of transmitting the bit b at w , conditioned on a particular fixing of both x, y , and conditioned on the event of reaching w . We view $O_{wxy} = (O_{wxy}(0), O_{wxy}(1))$ as the “true” probability distribution at w . Observe that Alice can compute the distributions O_{wxy} and O_{wy} for vertices w that she owns: Since Alice is the one deciding on the bit to be sent at w , it holds that $O_{wxy} = O_{wx}$. Since μ is a product distribution, it also holds that $O_{wy} = O_w$. Similarly, Bob can compute O_{wxy} and O_{wx} for vertices w that he owns.

We define the divergence at w with respect to x , denoted by $\mathbb{D}_{wx}$, as $\mathbb{D}(O_{wx} \| O_w)$, where $\mathbb{D}((p, 1-p) \| (q, 1-q)) = p \log(p/q) + (1-p) \log((1-p)/(1-q))$ is the *divergence* (also known as *relative entropy* or the *Kullback-Leibler distance*) between the distributions $(p, 1-p)$ and $(q, 1-q)$. Observe that for a vertex w owned by Bob, $\mathbb{D}_{wx} = \mathbb{D}(O_{wx} \| O_w) = \mathbb{D}(O_w \| O_w) = 0$. Similarly, the divergence at w with respect to y , denoted by $\mathbb{D}_{wy}$, is defined as $\mathbb{D}(O_{wy} \| O_w)$.

2.2 Frontiers

Let v be a vertex and let $\mathcal{C}$ be a subset of descendants of v . The set $\mathcal{C}$ is a *frontier* (or a *cut*) with respect to v , if every path from v to a leaf that is a descendant of v , contains exactly one element of $\mathcal{C}$.

Let $\beta = 1/\text{polylog}(I)$. Given a vertex v , we define $\mathcal{C}_{vx}$ to be the set of descendants w of v such that if we sum up $\mathbb{D}_{w'x}$ for all intermediate vertices w' on the path from v to w we get a total $< \beta$, but adding $\mathbb{D}_{wx}$ makes the total at least β , or w is a leaf. Intuitively, vertices in $\mathcal{C}_{vx}$ correspond to the shortest transcripts for which Alice reveals β additional bits of information about x (in addition to the information revealed at v). The set $\mathcal{C}_{vy}$ is defined similarly. Observe that $\mathcal{C}_{vx}$ and $\mathcal{C}_{vy}$ are frontiers with respect to v . In addition, Alice knows the frontier $\mathcal{C}_{vx}$ and Bob knows the frontier $\mathcal{C}_{vy}$.

Let w be a vertex. We define $D_{vxy}(w)$ to be the probability that π reaches w , when the inputs are x, y , conditioned on it reaching v . Note that $D_{vxy}(w)$ is obtained by multiplying the probabilities $O_{w'xy}(b_{w'})$ for vertices w' along the path

from v to w , where $b_{w'} = 0$ if the path from v to w passes through the edge going out of w' that is labeled by 0, and $b_{w'} = 1$ otherwise. In other words, following the probabilities $O_{w'xy} = O_{w'x}$ for vertices owned by Alice, and $O_{w'xy} = O_{w'y}$ for vertices owned by Bob.

We define $D_{vx}(w)$ to be the “best estimate” of $D_{vxy}(w)$ by Alice. That is, $D_{vx}(w)$ is induced by following the probabilities $O_{w'x}$ for vertices owned by Alice, and $O_{w'x} = O_{w'}$ for vertices owned by Bob. The value $D_{vy}(w)$ is defined similarly. We define $D_v(w)$ to be the “best estimate” of $D_{vxy}(w)$ by an observer who doesn’t know x, y . As before, $D_v(w)$ is induced by following the probabilities $O_{w'}$ along the path. Note that $D_v(w)$ is known to both players, as well as to an external observer. Also note that by restricting each of the functions $D_{vxy}, D_{vx}, D_{vy}, D_v$ (defined over the set of all descendants of v) to any frontier with respect to v , a probability distribution is obtained. We denote by $\tilde{D}_{vxy}, \tilde{D}_{vx}, \tilde{D}_{vy}, \tilde{D}_v$ the distributions obtained by restricting the functions $D_{vxy}, D_{vx}, D_{vy}, D_v$ (respectively) to the leaves that are descendants of v . When we omit the vertex v from the notation (e.g., $D_{xy}, \tilde{D}$ etc.), we mean that v is the root of the tree.

2.3 The Simulation Protocol τ

The simulation protocol proceeds as follows: Initially v is set to the root of the tree, t below is some large constant, $r = \text{poly}(I)$, and $\eta = 1/\text{poly}(I)$.

1. **Selecting a Leader:** Players sample a bit p from the public randomness. If $p = 1$ (Bob leads), they run Steps 2–4 as written. Otherwise, if $p = 0$ (Alice leads), they run Steps 2–4 with the roles of Alice and Bob, and the roles of x and y , switched.
2. **Correlated Sampling:** Players jointly sample a leaf u according to $\tilde{D}_{vy}$. (Recall that Bob knows the distribution $\tilde{D}_{vy}$, while Alice’s best estimate of $\tilde{D}_{vy}$ is $\tilde{D}_v$). This is done by running the correlated sampling protocol of [BR11] with $P = \tilde{D}_{vy}$, $Q = \tilde{D}_v$, and error probability η , and requires the exchange of $\approx \mathbb{D}(\tilde{D}_{vy} \| \tilde{D}_v) + \log(1/\eta)$ bits (see Lemma 14).
3. **Finding a Separation:** Consider the unique vertex w' in the intersection of $\mathcal{C}_{vx}$ and the path from v to u . We define $c_{vx}(u)$ to be the index (number) of w' on the path from v to u . We define $c_{vy}(u)$ similarly. The goal of this step is for the players to agree on an index k such that $\min\{c_{vx}(u), c_{vy}(u)\} \leq k \leq \max\{c_{vx}(u), c_{vy}(u)\}$. This is done with error $1/r$ by sending $O(\log(r))$ bits, as follows:
  - (a) Let $c_1 \leq \dots \leq c_r \in \mathbb{N}$ be such that c_i is the first index satisfying $\Pr_{x'}[c_{vx'}(u) \leq c_i] \geq i/r$, where the probability is over x' that is sampled according to the current distribution over Alice’s inputs (i.e., the original distribution μ^A conditioned on the current transcript of τ). Observe that this distribution is still a product distribution and that the indices c_i are known to both players.

- (b) Alice sends the first index $a \in [r]$ such that $c_{vx}(u) \leq c_a$. Bob sends the first index $b \in [r]$ such that $c_{vy}(u) \leq c_b$.
- (c) Return $k = \min\{c_a, c_b\}$. Let w be the k^{th} vertex on the path from v to u .

4. Rejection Sampling: Alice sends a bit a' that equals 0 (reject) if w has an ancestor in $\mathcal{C}_{vx}$ (intuitively, the transcript w reveals $\geq \beta$ additional bits of information about x). Otherwise, a' equals 1 (accept) with probability $\min\{1, D_{vx}(w)/tD_v(w)\}$. If $a' = 1$, then the players set $v = w$. If v is a leaf they end the protocol, otherwise they go back to Step 1.

2.4 Protocol Analysis

To get a rough idea of why this protocol works, fix x, y and the transcript of τ so far, and assume that vertex v was reached. We first consider the case where Step 1 selects $p = 1$ (Bob leads).

Consider a particular execution of Step 3, and let k be the returned index. Let E be the event that in this execution of Step 3, $c_a = c_b$ and $c_a \neq c_{vx}(u)$. We claim that $\min\{c_{vx}(u), c_{vy}(u)\} \leq k \leq \max\{c_{vx}(u), c_{vy}(u)\}$, unless E occurs:

1. If $c_a > c_b$ then $k = \min\{c_a, c_b\} = c_b$. By the definition of c_a as the first index for which $c_{vx}(u) \leq c_a$, it holds that $k = c_b < c_{vx}(u)$. Therefore, $c_{vy}(u) \leq c_b = k < c_{vx}(u)$, and we are done. The case where $c_b > c_a$ is similar.
2. If $c_a = c_b$ and $c_a = c_{vx}(u)$, then $c_{vx}(u) = c_a = k = c_b \geq c_{vy}(u)$, and we are also done.

Fix v, u . The current distribution over the players' inputs is obtained by conditioning the original distribution μ on the current transcript of the protocol τ . Observe that this is still a product distribution, as it is obtained by conditioning a product distribution on a transcript of a protocol. We will prove that since this distribution is a product distribution, the event E occurs with probability at most $1/r$ (over the selection of inputs). For intuition, consider the case where the index a is (almost) uniformly distributed in $[r]$ (this always happens when the indices c_i are all distinct). Since Alice's input and Bob's input are still independent, we have that $c_b = c_a$ with probability $1/r$, and thus E occurs with probability at most $1/r$. Another extreme case is when there exists $c \in [r]$, such that it is always the case that $c_a = c$ (this always happens when all the indices c_i coincide). Then, except with probability $1/r$ over the selection of x , we have that $c_{vx}(u) = c = c_a$, and thus E does not occur. The actual proof follows from these two intuitions.

Fix x, y and the transcript of τ so far, and assume that vertex v was reached. Consider the set $\mathcal{C} = \mathcal{C}_{vxy}$ of vertices $w = w_u$ obtained by Step 3 of the protocol τ for the different leaves u that are descendants of v . We claim that $\mathcal{C}$ is "close" to being a frontier with respect to v : Consider the protocol τ' that operates the same as τ , except that Step 3 is changed so whenever E occurs, the protocol τ' has

Alice and Bob exchanging $c_{vx}(u)$ and $c_{vy}(u)$, and returns $k = \min\{c_{vx}(u), c_{vy}(u)\}$. It can be shown that the set $\mathcal{C}' = \mathcal{C}'_{vxy}$ of vertices $w = w_u$ obtained by Step 3 of the protocol τ' for the different leaves u that are descendants of v , is an actual frontier with respect to v . We also note that $\mathcal{C}'$ is "always between" $\mathcal{C}_{vx}$ and $\mathcal{C}_{vy}$: As claimed before, if E does not occur, the returned value k satisfies $\min\{c_{vx}(u), c_{vy}(u)\} \leq k \leq \max\{c_{vx}(u), c_{vy}(u)\}$. If E occurs then Step 3 of τ' returns $k = \min\{c_{vx}(u), c_{vy}(u)\}$. That is, a vertex in $\mathcal{C}'$ has either an ancestor in $\mathcal{C}_{vx}$ and a descendant in $\mathcal{C}_{vy}$, or has an ancestor in $\mathcal{C}_{vy}$ and a descendant in $\mathcal{C}_{vx}$.

Assume for simplicity that $\mathcal{C} = \mathcal{C}'$, and, in particular, that $\mathcal{C}$ is a frontier that is "always between" $\mathcal{C}_{vx}$ and $\mathcal{C}_{vy}$ ($\mathcal{C}$ and $\mathcal{C}'$ are "close" anyway, as E is an event that occurs with probability $\leq 1/r$). Also assume that all the vertices in $\mathcal{C}$ are two levels below v , with the first vertex (i.e., v) owned by Alice, and the vertices in the intermediate level owned by Bob. Let $w \in \mathcal{C}$. Recall that $D_{vxy}(w)$ is the true probability of arriving at w conditioned on reaching v , and that $D_v(w)$ is the best estimate of $D_{vxy}(w)$ by an observer who does not know x, y . Fixing w , we write $D_{vxy}(w) = D_1 D_2$ and $D_v(w) = D'_1 D'_2$, where D_i denotes the true probability that step i is taken according to x, y , and D'_i denotes this probability as estimated by an observer who does not know x, y . Observe that $D_{vx}(w) = D_1 D'_2$, $D_{vy}(w) = D'_1 D_2$. Also note that the probability that w is selected by Step 3 is $D_{vy}(w) = D'_1 D_2$, as the correlated sampling of Step 2 outputs a leaf u distributed according to $\tilde{D}_{vy}$.

We first consider the case where the frontier $\mathcal{C}_{vy}$ is "always above" $\mathcal{C}_{vx}$. That is, every vertex in $\mathcal{C}_{vx}$ is a descendant of some vertex in $\mathcal{C}_{vy}$. Intuitively, this means that upon reaching v , Bob always gives β bits of information about y before Alice gives β bits of information about x . Since we assume that $\mathcal{C}_{vy}$ is "always above" $\mathcal{C}_{vx}$, and since $\mathcal{C}$ is "always between" $\mathcal{C}_{vy}$ and $\mathcal{C}_{vx}$, it holds that w has a descendant in $\mathcal{C}_{vx}$. This intuitively means that the transcript w reveals $< \beta$ additional bits of information about x , thus $D_{vx}(w)$ and $D_v(w)$ (which is the estimation of $D_{vx}(w)$ by an observer who doesn't know x) tend to be close. Now assume that the threshold t is set high enough so that $tD_v(w) \geq D_{vx}(w)$ with high probability. In this case the probability that w is accepted by Step 4 equals

$$\Pr[a' = 1] = \frac{D_{vx}(w)}{tD_v(w)} = \frac{D_1 D'_2}{tD'_1 D'_2} = \frac{D_1}{tD'_1}.$$

This implies that the total probability that w is selected by Step 4 to be the new v , is $D'_1 D_2$ times D_1/tD'_1 , which is exactly its correct probability $D_1 D_2$ divided by t . Hence, we get an overhead of t steps, but output the right distribution over w .

Now, it may be the case that $\mathcal{C}_{vy}$ is not always above $\mathcal{C}_{vx}$. Actually, it may even be the case that $\mathcal{C}_{vy}$ is not always above $\mathcal{C}_{vx}$ and not always below it, but rather, some vertices in $\mathcal{C}_{vy}$ have descendants in $\mathcal{C}_{vx}$, and others have ancestors in $\mathcal{C}_{vx}$. This means that after reaching v , Alice may give $\geq \beta$ additional bits of information about her input before Bob does so for his input, or it may be the other way around. Step 1 randomly chooses one of the players in every iteration

and assumes that this player is going to be the first to give β additional bits of information about his input at the vertex w that will be selected by Step 3.

Assume that Bob was selected by Step 1 ($p = 1$). If w has an ancestor in $\mathcal{C}_{vy}$ (equivalently, w has a descendant in $\mathcal{C}_{vx}$), then Bob gives β additional bits of information before Alice does, and the assumption of Step 1 is true. In this case, Alice just “corrects” the probability of reaching w according to her own view, by accepting it with probability $D_{vx}(w)/tD_v(w)$. Observe that Bob does not need to correct w ’s probability, as it was already sampled according to D_{vy} . Otherwise, if w has an ancestor in $\mathcal{C}_{vx}$, then Alice gives β additional bits of information before Bob does. Thus, the assumption of Step 1 is false, and w gets rejected by Alice. This case, where Alice gives information first, will be handled when Step 1 selects $p = 0$ (Alice leads). Since the protocol τ cannot predict a priori (before u is selected by Step 2) which player is going to be the first to give β information, Step 1 randomly selects the leader.

2.5 The Communication Complexity of the Protocol

Let w be the vertex selected by Step 3. Since w is between $\mathcal{C}_{vx}$ and $\mathcal{C}_{vy}$, if we sum up the term $\mathbb{D}_{w'x} + \mathbb{D}_{w'y}$ for all intermediate vertices w' on the path from v to w we get a total $\geq \beta$ (unless w is a leaf). Intuitively, this means that $\geq \beta$ bits of information about x, y are revealed to the external observer when we go from the v to w . Since the external information of π over μ is only I , and since we set $\beta = 1/\text{polylog}(I)$, this means that the vertex v gets updated by Step 4 at most $m \approx I/\beta = I \cdot \text{polylog}(I)$ times, on average. Assume again that Bob was selected by Step 1 ($p = 1$) and that $\mathcal{C}_{vy}$ is always above $\mathcal{C}_{vx}$. As we claimed above, the probability that w is selected by Step 4 to be the new v , is roughly $D_{vxy}(w)/t$. Therefore, summing over all possible vertices w , in each iteration of the protocol τ , the vertex v gets updated by Step 4 with probability roughly $1/t$. This implies that the protocol τ runs for at most $m' \approx tm = I \cdot \text{polylog}(I)$ iterations, on average.

Each iteration consists of the following steps: Step 2 has error η . Since we run this step at most m' times, in order for the total error introduced by this step to be small, we set $\eta = 1/\text{poly}(I) \ll 1/m'$. Step 2 communicates $\approx \mathbb{D}(\tilde{D}_{vy} \parallel \tilde{D}_v) + \log(1/\eta)$ bits if $p = 1$, and $\approx \mathbb{D}(\tilde{D}_{vx} \parallel \tilde{D}_v) + \log(1/\eta)$ bits if $p = 0$. Proposition 9 shows that for every x and y ,

$$\mathbb{D}(\tilde{D}_x \parallel \tilde{D}) \leq \mathbf{E}_{y \leftarrow Y} [\mathbb{D}(\tilde{D}_{xy} \parallel \tilde{D})].$$

$$\mathbb{D}(\tilde{D}_y \parallel \tilde{D}) \leq \mathbf{E}_{x \leftarrow X} [\mathbb{D}(\tilde{D}_{xy} \parallel \tilde{D})].$$

Loosely stated, the first equation above says that the distribution $\tilde{D}_x$ is “closer” to $\tilde{D}$ than $\tilde{D}_{xy}$. This is true as both $\tilde{D}_x$ and $\tilde{D}_{xy}$ are induced by following the probabilities $O_{w'x}$ for vertices owned by Alice. However, for vertices owned by Bob, $\tilde{D}_x$ follows the probability $O_{w'}$, while $\tilde{D}_{xy}$ follows $O_{w'y}$. The distribution $\tilde{D}$ always follows $O_{w'}$.

Proposition 8 gives an alternative definition of external information, by showing that I , the external information of π over μ , satisfies

$$I = \mathbf{E}_{(x,y) \leftarrow \mu} [\mathbb{D}(\tilde{D}_{xy} \parallel \tilde{D})].$$

Therefore, Step 2 communicates $O(I)$ bits, in expectation.

Step 3 has error $1/r$. Since we run this step at most m' times, in order for the total error introduced by this step to be small, we set $r = \text{poly}(I) \gg m'$. The communication required by this step is then $O(\log(r)) = O(\log(I))$ bits, as the players only exchange indices $a, b \in [r]$. Step 4 requires 1 bit of communication. Hence, the total communication per round is $O(I)$, in expectation, and the protocol communicates $O(m'I) = I^2 \cdot \text{polylog}(I)$ bits.

3. PRELIMINARIES

The proofs are omitted from this section and can be found in the full version of this paper.

3.1 Notation

Let π be a randomized communication protocol. Let (x, y) be a possible input for π . We denote by $\pi(x, y)$ the random variable representing π ’s transcript (messages exchanged during π ’s execution), when it is run with the input (x, y) . We sometimes confuse random variables and their distribution. For example, we often view $\pi(x, y)$ as a distribution.

Let $s \in \{0, 1\}^*$ be a binary string. We denote by $|s| \in \mathbb{N} \cup \{0\}$ the length of s . For example, $|\pi(x, y)|$ is the length of the transcript represented by the random variable $\pi(x, y)$. For $i \in [|s|]$, we denote by $s_i \in \{0, 1\}$ the i^{th} bit of s , and by $s_{\leq i} \in \{0, 1\}^i$ the prefix of length i of s . For $i \leq 0$, set $s_{\leq i} = \phi$, where ϕ denotes the empty string (the string of length 0).

For a pair of distributions D, D' over the same domain Ω , we denote by $\|D - D'\|$ the *statistical/total variation* distance between D and D' , given by $\|D - D'\| = \max_{S \subseteq \Omega} \{|D(S) - D'(S)|\}$. We sometimes consider the statistical distance $\|D - D'\|$, where D is a distribution over a set $\mathcal{S}$ of binary strings, and D' is a distribution over a set $\mathcal{S}'$ of binary strings (e.g., $\|\pi(x, y) - \pi'(x, y)\|$, where π, π' are protocols and (x, y) is an input for both π and π'). For this statistical distance to be defined, we view both D and D' as distributions over the set of all binary strings. For a function $D : \Omega \rightarrow \mathbb{R}$ (not necessarily a distribution) and a set $\mathcal{S} \subseteq \Omega$, we define $D(\mathcal{S}) = \sum_{x \in \mathcal{S}} D(x)$.

For a pair of random variables X, Y , we denote by $\mathbb{I}(X; Y)$ the *mutual information* between X and Y . For a pair of distributions D, D' over the same domain, we denote by $\mathbb{D}(D \parallel D')$ the *relative entropy* between D and D' , given by

$$\mathbb{D}(D \parallel D') = \mathbf{E}_{x \leftarrow D} \left[\log \left(\frac{D(x)}{D'(x)} \right) \right].$$

3.2 Information Cost

Definition 1 (Internal Information Cost). *The internal information cost of a (private coin) protocol π over random*

inputs (X, Y) that are drawn according to a joint distribution μ , is defined as

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

Definition 2 (External Information Cost). *The external information cost of a (private coin) protocol π over random inputs (X, Y) that are drawn according to a joint distribution μ , is defined as*

$$\text{Ext}_\mu(\pi) = \mathbb{I}(\pi(X, Y); X, Y).$$

Fact 2. *Let π be a protocol and let μ be a distribution over the inputs for π . Then, $\text{IC}_\mu(\pi) \leq \text{Ext}_\mu(\pi)$.*

Fact 3 (Fact 4.16 in [BBCR10]). *Let π be a protocol and let $\mu = \mu^A \times \mu^B$ be a product distribution over the inputs for π . Then, $\text{IC}_\mu(\pi) = \text{Ext}_\mu(\pi)$.*

3.3 Divergence Tree

Transcript tree.

Let π be a communication protocol between two players, Alice and Bob. Alice has a private input x and Bob has a private input y , where (x, y) is chosen according to some publicly known joint distribution μ . In this work, we consider the case where $\mu = \mu^A \times \mu^B$ is a product distribution. We assume, without loss of generality, that π does not use public randomness (but may use private randomness), as the public randomness can always be replaced by private randomness without increasing the (internal or external) information cost. We also assume that the players take turns sending bits to one another. Alice send bits in odd rounds and Bob in the even rounds. We further assume, without loss of generality, that the players communicate the same number of bits in every execution of the protocol.

We denote by $\mathcal{T}$ the (directed) binary tree associated with the communication protocol π . That is, every vertex v of $\mathcal{T}$ corresponds to a possible transcript of π . The two edges going out of v are labeled by 0 and 1, corresponding to the next bit to be transmitted. We assume that $\mathcal{T}$ is a complete binary tree. We denote by $\mathcal{V}$ the set of vertices of $\mathcal{T}$, by v_0 the root of $\mathcal{T}$, and by $\mathcal{L}$ the set of leaves of $\mathcal{T}$. Since a vertex $v \in \mathcal{V}$ corresponds to a possible transcript of π , we often think of v as a string. That is, we view v as the binary string induced by the labels of the edges on the path from the root to v .

We think of Alice as owning the non-leaf vertices in the odd layers of $\mathcal{T}$, and of Bob as owning the non-leaf vertices in the even layers of $\mathcal{T}$. We denote by $\mathcal{V}^A \subseteq \mathcal{V} \setminus \mathcal{L}$ the set of vertices owned by Alice and by $\mathcal{V}^B \subseteq \mathcal{V} \setminus \mathcal{L}$ the set of vertices owned by Bob. The protocol π proceeds as follows: Starting from the root, when the protocol π reaches a non-leaf vertex v , the player who owns v sends a bit to the other player. The players follow the edge indicated by the sent bit and reach a new vertex.

For a vertex $v \in \mathcal{V}$, we denote by $\mathcal{V}(v) \subseteq \mathcal{V}$ the set of vertices in the subtree rooted in v (including v itself), and by $\mathcal{L}(v)$ the set of leaves in the subtree rooted in v (including v itself). Let $v, w \in \mathcal{V}$. We say that w is a *descendant* of v , if w is in the subtree rooted in v (in particular, v is a descendant

of itself). We say that w is a *strict descendant* of v , if w is in the subtree rooted in v and $w \neq v$. We say that v is an *ancestor* of w , if w is a descendant of v . We say that v is a *strict ancestor* of w , if w is a strict descendant of v .

Distributions associated with a vertex.

Let (X, Y) be a pair of random variables distributed according to μ , representing the players' inputs. Since μ is a product distribution, X and Y are independent. We view $\pi(X, Y)$ as the random variable representing the leaf of $\mathcal{T}$ reached by the protocol π .

Let $v \in \mathcal{V}$. We assume that for every $(x, y) \in \text{supp}(\mu)$ it is possible for the protocol π to reach the vertex v . That is, $\Pr[\pi(X, Y)_{\leq |v|} = v \mid X = x, Y = y] > 0$. This can be assumed without loss of generality: Consider the private coin protocol π^* , in which Alice follows π with probability $1 - \delta$, and, with probability δ , she sends a random bit for each vertex that she owns, for some sufficiently small $\delta > 0$. Bob acts similarly. Observe that π^* has the required property and $\|\pi(X, Y) - \pi^*(X, Y)\| \leq 2\delta$.

Let $v \in \mathcal{V} \setminus \mathcal{L}$ and $(x, y) \in \text{supp}(\mu)$. We define the following distributions over the two edges going out of v : O_{vxy} is the distribution over the edges going out of v according to the protocol π . We think of O_{vxy} as the “real” distribution at v . The distribution O_{vx} is Alice's best estimate O_{vxy} , and the distribution O_{vy} is Bob's best estimate O_{vxy} . The distribution O_v is the best estimate of O_{vxy} by an external observer who doesn't know neither x nor y . Formally, for $b \in \{0, 1\}$,

$$O_{vxy}(b) =$$

$$\Pr[\pi(X, Y)_{|v|+1} = b \mid \pi(X, Y)_{\leq |v|} = v, X = x, Y = y],$$

$$O_{vx}(b) = \Pr[\pi(X, Y)_{|v|+1} = b \mid \pi(X, Y)_{\leq |v|} = v, X = x],$$

$$O_{vy}(b) = \Pr[\pi(X, Y)_{|v|+1} = b \mid \pi(X, Y)_{\leq |v|} = v, Y = y],$$

$$O_v(b) = \Pr[\pi(X, Y)_{|v|+1} = b \mid \pi(X, Y)_{\leq |v|} = v].$$

We note that

$$\forall v \in \mathcal{V}^B : O_{vxy}(b) = O_{vy}(b); \quad \forall v \in \mathcal{V}^A : O_{vxy}(b) = O_{vx}(b). \quad (1)$$

Furthermore, since X and Y are independent even conditioned on any transcript v of the protocol π , we get

$$\forall v \in \mathcal{V}^B : O_{vx}(b) = O_v(b); \quad \forall v \in \mathcal{V}^A : O_{vy}(b) = O_v(b). \quad (2)$$

Probabilities of reaching a vertex.

Let $v \in \mathcal{V}$ and $(x, y) \in \text{supp}(\mu)$. We define the functions $D_{vxy}, D_{vx}, D_{vy}, D_v : \mathcal{V}(v) \rightarrow [0, 1]$ by:

$$D_{vxy}(v) = D_{vx}(v) = D_{vy}(v) = D_v(v) = 1,$$

and for $w \in \mathcal{V}(v) \setminus \{v\}$,

$$D_{vxy}(w) = \prod_{i \in \{|v|, \dots, |w|-1\}} O_{w_{\leq i}xy}(w_{i+1}),$$

$$D_{vx}(w) = \prod_{i \in \{|v|, \dots, |w|-1\}} O_{w \leq_i x}(w_{i+1}),$$

$$D_{vy}(w) = \prod_{i \in \{|v|, \dots, |w|-1\}} O_{w \leq_i y}(w_{i+1}),$$

$$D_v(w) = \prod_{i \in \{|v|, \dots, |w|-1\}} O_{w \leq_i}(w_{i+1}).$$

Let $\tilde{D}_{vxy}, \tilde{D}_{vx}, \tilde{D}_{vy}, \tilde{D}_v : \mathcal{L}(v) \rightarrow [0, 1]$ be the restrictions of the functions $D_{vxy}, D_{vx}, D_{vy}, D_v$ to $\mathcal{L}(v)$ (respectively). Observe that $\tilde{D}_{vxy}, \tilde{D}_{vx}, \tilde{D}_{vy}, \tilde{D}_v$ are distributions over $\mathcal{L}(v)$. Whenever we omit the subscript v , we mean that v is v_0 , the root of $\mathcal{T}$. E.g., $D_{xy} = D_{v_0xy}$, $\tilde{D}_{xy} = \tilde{D}_{v_0xy}$.

Proposition 4. *Let $v \in \mathcal{V}$, $w \in \mathcal{V}(v)$, and $(x, y) \in \text{supp}(\mu)$. Then,*

$$\frac{D_{vy}(w) \cdot D_{vx}(w)}{D_v(w)} = D_{vxy}(w).$$

Let $v \in \mathcal{V}$, $w \in \mathcal{V}(v)$, and $(x, y) \in \text{supp}(\mu)$. If $w = v$, we define $\mathbb{D}_{vwx} = \mathbb{D}_{vwx} = \mathbb{D}_{vwy} = 0$. If $w \neq v$, we define

$$\begin{aligned} \mathbb{D}_{vwx} &= \sum_{i \in \{|v|, \dots, |w|-1\}} \mathbb{D}(O_{w \leq_i xy} \| O_{w \leq_i}), \\ \mathbb{D}_{vwy} &= \sum_{i \in \{|v|, \dots, |w|-1\}} \mathbb{D}(O_{w \leq_i xy} \| O_{w \leq_i}), \\ \mathbb{D}_{vwx} &= \sum_{i \in \{|v|, \dots, |w|-1\}} \mathbb{D}(O_{w \leq_i xy} \| O_{w \leq_i}), \end{aligned}$$

By Equations (1) and (2) it holds that $\mathbb{D}_{vwx} = \mathbb{D}_{vwx} + \mathbb{D}_{vwy}$.

Proposition 5 (Corollary 6 in [Gan12]). *Let $v \in \mathcal{V}$ and $(x, y) \in \text{supp}(\mu)$. Then,*

$$\begin{aligned} \mathbb{D}(\tilde{D}_{vxy} \| \tilde{D}_v) &= \sum_{w \in \mathcal{V}(v) \setminus \mathcal{L}(v)} D_{vxy}(w) \cdot \mathbb{D}(O_{wx} \| O_w) \\ &= \sum_{u \in \mathcal{L}(v)} \tilde{D}_{vxy}(u) \cdot \mathbb{D}_{vux}. \end{aligned}$$

$$\begin{aligned} \mathbb{D}(\tilde{D}_{vx} \| \tilde{D}_v) &= \sum_{w \in \mathcal{V}(v) \setminus \mathcal{L}(v)} D_{vx}(w) \cdot \mathbb{D}(O_{wx} \| O_w) \\ &= \sum_{u \in \mathcal{L}(v)} \tilde{D}_{vx}(u) \cdot \mathbb{D}_{vux}. \end{aligned}$$

$$\begin{aligned} \mathbb{D}(\tilde{D}_{vy} \| \tilde{D}_v) &= \sum_{w \in \mathcal{V}(v) \setminus \mathcal{L}(v)} D_{vy}(w) \cdot \mathbb{D}(O_{wy} \| O_w) \\ &= \sum_{u \in \mathcal{L}(v)} \tilde{D}_{vy}(u) \cdot \mathbb{D}_{vuy}. \end{aligned}$$

Proposition 6. *For every $(x, y) \in \text{supp}(\mu)$ and $w \in \mathcal{V}$,*

$$D_{xy}(w) = \Pr[\pi(X, Y)_{\leq |w|} = w \mid X = x, Y = y],$$

$$D_x(w) = \Pr[\pi(X, Y)_{\leq |w|} = w \mid X = x],$$

$$D_y(w) = \Pr[\pi(X, Y)_{\leq |w|} = w \mid Y = y],$$

$$D(w) = \Pr[\pi(X, Y)_{\leq |w|} = w].$$

Proposition 7. *For every $x \in \text{supp}(\mu^A)$, $y \in \text{supp}(\mu^B)$, and $w \in \mathcal{V}$,*

$$D_x(w) = \Pr_{y \leftarrow Y}[D_{xy}(w)],$$

$$D_y(w) = \Pr_{x \leftarrow X}[D_{xy}(w)].$$

The following claim gives a useful alternative definition of external information.

Proposition 8. *It holds that*

$$\text{Ext}_\mu(\pi) = \mathbf{E}_{(x, y) \leftarrow \mu} [\mathbb{D}(\tilde{D}_{xy} \| \tilde{D})].$$

Proposition 9. *For every $x \in \text{supp}(\mu^A)$ and $y \in \text{supp}(\mu^B)$,*

$$\mathbb{D}(\tilde{D}_x \| \tilde{D}) \leq \mathbf{E}_{y \leftarrow Y} [\mathbb{D}(\tilde{D}_{xy} \| \tilde{D})],$$

$$\mathbb{D}(\tilde{D}_y \| \tilde{D}) \leq \mathbf{E}_{x \leftarrow X} [\mathbb{D}(\tilde{D}_{xy} \| \tilde{D})].$$

3.4 Frontiers

Definition 3 (Frontier). *Let $v \in \mathcal{V}$ and $\mathcal{C} \subseteq \mathcal{V}(v)$. The set $\mathcal{C}$ is a frontier with respect to v , if every path from v to a leaf $u \in \mathcal{L}(v)$ contains exactly one element of $\mathcal{C}$.*

Let $(x, y) \in \text{supp}(\mu)$. Let $v \in \mathcal{V}$ and let $\mathcal{C} \subseteq \mathcal{V}(v)$ be a frontier with respect to v . Observe that when restricting each of the functions $D_{vxy}, D_{vx}, D_{vy}, D_v$ to the frontier $\mathcal{C}$, we get a distribution. In particular, since $\mathcal{L}(v)$ is a frontier with respect to v , $\tilde{D}_{vxy}, \tilde{D}_{vx}, \tilde{D}_{vy}, \tilde{D}_v$ are distributions.

Proposition 10. *Let $x \in \text{supp}(\mu^A)$ and $y \in \text{supp}(\mu^B)$. Let $\mathcal{C}$ be a frontier with respect to the root v_0 . Then,*

$$\mathbb{D}(\tilde{D}_x \| \tilde{D}) \geq \mathbf{E}_{v \leftarrow \tilde{D}_x|_{\mathcal{C}}} [\mathbb{D}(\tilde{D}_{vx} \| \tilde{D}_v)] = \sum_{v \in \mathcal{C}} D_x(v) \cdot \mathbb{D}(\tilde{D}_{vx} \| \tilde{D}_v),$$

$$\mathbb{D}(\tilde{D}_y \| \tilde{D}) \geq \mathbf{E}_{v \leftarrow \tilde{D}_y|_{\mathcal{C}}} [\mathbb{D}(\tilde{D}_{vy} \| \tilde{D}_v)] = \sum_{v \in \mathcal{C}} D_y(v) \cdot \mathbb{D}(\tilde{D}_{vy} \| \tilde{D}_v),$$

where $D_x|_{\mathcal{C}}$ and $D_y|_{\mathcal{C}}$ denote the distributions obtained by restricting the functions D_x and D_y (respectively) to the frontier $\mathcal{C}$.

The sets $\mathcal{V}_{vx}, \mathcal{V}_{vy}$ and the frontiers $\mathcal{C}_{vx}, \mathcal{C}_{vy}$.

Assume that the leaves of $\mathcal{T}$ are (all) in level d , and that d is odd. Let $\mathcal{L}^B$ be set of vertices in the level $d-1$, the last level before the leaves. Note that level $d-1$ is owned by Bob. Let $\mathcal{L}^A$ be set of vertices in the level $d-2$, the second to last level before the leaves. Note that level $d-2$ is owned by Alice. Let $\mathcal{L}^+ = \mathcal{L} \cup \mathcal{L}^A \cup \mathcal{L}^B$.

We use a parameter $\beta > 0$ that will be set later. Let $v \in \mathcal{V} \setminus \mathcal{L}^+$, $x \in \text{supp}(\mu^A)$ and $y \in \text{supp}(\mu^B)$. We define the set $\mathcal{V}_{vx}$ as the set of all $w \in \mathcal{V}(v) \setminus \mathcal{L}$ satisfying $\mathbb{D}_{vw \leq |w|-1} x < \beta$. That is, we include w in $\mathcal{V}_{vx}$ if $\mathbb{D}_{vw} x < \beta$, or if $\mathbb{D}_{vw} x \geq \beta$ and w is the first vertex on the path from v to w for which $\mathbb{D}_{vw} x \geq \beta$. Similarly, we define the set $\mathcal{V}_{vy}$ as the set of all $w \in \mathcal{V}(v) \setminus (\mathcal{L} \cup \mathcal{L}^B)$ satisfying $\mathbb{D}_{vw \leq |w|-1} y < \beta$. Note that Alice knows the set $\mathcal{V}_{vx}$ and Bob knows the set $\mathcal{V}_{vy}$.

We define the set $\mathcal{C}_{vx} \subseteq \mathcal{V}_{vx}$ as the set of all $w \in \mathcal{V}_{vx}$ such that w 's children are not in $\mathcal{V}_{vx}$ (observe that if one child is not in $\mathcal{V}_{vx}$, then the other child is not in $\mathcal{V}_{vx}$ either). In other

words, the set $\mathcal{C}_{vx}$ is the “border” of the set $\mathcal{V}_{vx}$. Similarly, we define the set $\mathcal{C}_{vy} \subseteq \mathcal{V}_{vy}$ as the set of all $w \in \mathcal{V}_{vy}$ such that w ’s children are not in $\mathcal{V}_{vy}$. Observe that both $\mathcal{C}_{vx}$ and $\mathcal{C}_{vy}$ are frontiers with respect to v . Note that Alice knows the frontier $\mathcal{C}_{vx}$ and Bob knows the frontier $\mathcal{C}_{vy}$. In addition, $\mathcal{C}_{vx} \subseteq \mathcal{V}^B$ and $\mathcal{C}_{vy} \subseteq \mathcal{V}^A$, thus

$$\mathcal{C}_{vx} \cap \mathcal{C}_{vy} = \phi. \quad (3)$$

The indices $c_{vx}(u), c_{vy}(u)$.

Let $v \in \mathcal{V} \setminus \mathcal{L}^+$ and $u \in \mathcal{L}(v)$. We denote by $P(v, u)$ the set of vertices on the path from v to u . Let $x \in \text{supp}(\mu^A)$ and $y \in \text{supp}(\mu^B)$. Consider the unique vertex w in the intersection of $\mathcal{C}_{vx}$ and $P(v, u)$. We set $c_{vx}(u)$ to be the index (number) of w on the path from v to u . That is, if the vertices on the path from v to u are $w_1, \dots, w_k$ (where $w_1 = v$ and $w_k = u$), and $w_c \in \mathcal{C}_{vx}$ ($c \in [k]$), then $c_{vx}(u) = c$. Similarly, we consider the unique vertex w' in the intersection of $\mathcal{C}_{vy}$ and $P(v, u)$. We set $c_{vy}(u)$ to be the index (number) of w' on the path from v to u .

The frontiers $\mathcal{C}_{vxy}^{\min}$ and $\mathcal{C}_{vxy}^{\max}$.

Let $v \in \mathcal{V} \setminus \mathcal{L}^+$ and $(x, y) \in \text{supp}(\mu)$. We define a set $\mathcal{C}_{vxy}^{\min} \subseteq \mathcal{C}_{vx} \cup \mathcal{C}_{vy}$ as follows: Include $w \in \mathcal{C}_{vx}$ in $\mathcal{C}_{vxy}^{\min}$ if w has a descendant in $\mathcal{C}_{vy}$. In addition, include $w \in \mathcal{C}_{vy}$ in $\mathcal{C}_{vxy}^{\min}$ if w has a descendant in $\mathcal{C}_{vx}$. We define a set $\mathcal{C}_{vxy}^{\max} \subseteq \mathcal{C}_{vx} \cup \mathcal{C}_{vy}$ as follows: Include $w \in \mathcal{C}_{vx}$ in $\mathcal{C}_{vxy}^{\max}$ if w has an ancestor in $\mathcal{C}_{vy}$. In addition, include $w \in \mathcal{C}_{vy}$ in $\mathcal{C}_{vxy}^{\max}$ if w has an ancestor in $\mathcal{C}_{vx}$. Observe that $\mathcal{C}_{vxy}^{\min}$ and $\mathcal{C}_{vxy}^{\max}$ are frontiers.

Let $v \in \mathcal{V}$ and $(x, y) \in \text{supp}(\mu)$. Let $w \in \mathcal{V}(v)$. We write $\mathcal{C}_{vxy}^{\min} < w \leq \mathcal{C}_{vxy}^{\max}$, if w is a strict descendant of a vertex $w' \in \mathcal{C}_{vxy}^{\min}$, and w is an ancestor of a vertex $w'' \in \mathcal{C}_{vxy}^{\max}$ (note that we allow $w = w''$, but do not allow $w = w'$). Observe that if $\mathcal{C}_{vxy}^{\min} < w \leq \mathcal{C}_{vxy}^{\max}$, then $w \in \mathcal{V}_{vx} \cup \mathcal{V}_{vy}$. For a set $\mathcal{C} \subseteq \mathcal{V}(v)$, we write $\mathcal{C}_{vxy}^{\min} < \mathcal{C} \leq \mathcal{C}_{vxy}^{\max}$, if for every $w \in \mathcal{C}$, it holds that $\mathcal{C}_{vxy}^{\min} < w \leq \mathcal{C}_{vxy}^{\max}$.

Proposition 11. *Let $v \in \mathcal{V} \setminus \mathcal{L}^+$ and $(x, y) \in \text{supp}(\mu)$. Let $w \in \mathcal{V}(v)$ such that $\mathcal{C}_{vxy}^{\min} < w \leq \mathcal{C}_{vxy}^{\max}$. Then, either $w \in \mathcal{V}_{vx} \setminus \mathcal{V}_{vy}$ or $w \in \mathcal{V}_{vy} \setminus \mathcal{V}_{vx}$.*

Proposition 12. *Let $v \in \mathcal{V} \setminus \mathcal{L}^+$ and $(x, y) \in \text{supp}(\mu)$. Let $\mathcal{C}', \mathcal{C}'' \subseteq \mathcal{V}(v)$ be two frontiers with respect to v such that $\mathcal{C}_{vxy}^{\min} < \mathcal{C}', \mathcal{C}'' \leq \mathcal{C}_{vxy}^{\max}$. Let*

$$\mathcal{C} = (\mathcal{C}' \cap \mathcal{V}_{vx}) \cup (\mathcal{C}'' \cap \mathcal{V}_{vy}).$$

Then, $\mathcal{C}$ is a frontier with respect to v . In addition,

$$(\mathcal{C}' \cap \mathcal{V}_{vx}) \cap (\mathcal{C}'' \cap \mathcal{V}_{vy}) = \phi.$$

3.5 Smooth Simulation

Like in [BBCR10], it will be convenient for us to assume that the protocol to be simulated has “smooth” messages, 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 next argue that any protocol can be transformed into a smooth protocol without increasing the external information cost by much.

Definition 4 (Smooth Protocol). *Let $\beta > 0$. The protocol π is β -smooth if for every vertex $v \in \mathcal{V} \setminus \mathcal{L}$, every possible input (x, y) for π , and any $b \in \{0, 1\}$, it holds that*

$$O_{vxy}(b) \in [1/2 - \beta, 1/2 + \beta].$$

Lemma 13 (Smooth Simulation). *Let $\kappa, \beta > 0$. Let π' be a private coin protocol and μ be a distribution over the inputs for π' . Then, there exists a β -smooth private coin protocol π , and a transcript reconstruction function f that takes as an input a possible transcript of π , and returns a possible transcript of π' , such that for every $(x, y) \in \text{supp}(\mu)$,*

$$\|f(\pi(x, y)) - \pi'(x, y)\| \leq \kappa.$$

In addition,

$$\text{Ext}_\mu(\pi) \leq \text{Ext}_\mu(\pi') + \kappa.$$

3.6 Correlated Sampling

Our simulation protocol uses the correlated sampling protocol *CorrelatedSampling*, promised by the following lemma proved in [BR11]:

Lemma 14 (Theorems 2.1 and 4.1 in [BR11]). *Let $\eta > 0$. Suppose that Alice is given a distribution P and Bob is given a distribution Q , both over the same universe U (the distributions are described by the probabilities assigned to each point). There is a public coin protocol, *CorrelatedSampling*(P, Q, η), such that at the end of the protocol:*

- *Alice outputs an element a distributed according to P .*
- *Bob outputs b such that $\Pr[b \neq a] < \eta$.*
- *The expected communication complexity of the protocol is at most*

$$10\mathbb{D}(P\|Q) + 2\log(1/\eta) + 10,$$

where the expectation is over the randomness used by the players.

We mention that we only apply Lemma 14 for the special case where Alice knows both P and Q .

4. THE SIMULATION PROTOCOL

Let π be a β -smooth private coin communication protocol, for β to be specified next. Let μ be a distribution over the inputs for π . In this section we present the protocol τ^{ST} that simulates π over μ . We first present a related protocol τ (Section 4.1). The actual simulation protocol, τ^{ST} , is then easily obtained from τ (Section 4.2).

Parameters.

Fix a proximity parameter $\varepsilon \in (0, 1)$. Let $I = \text{Ext}_\mu(\pi)$. The protocols τ and τ^{ST} use the following parameters that depend on ε and I : Below $c \in \mathbb{N}$ is some sufficiently large constant.

$$\beta = \frac{\varepsilon}{\log^c(I)} = \frac{\varepsilon}{\text{poly}(\log(I))}; \quad t = O(1);$$

$$m = \frac{10^5 I}{\varepsilon^2 \beta} = I \text{poly}(\log(I)/\varepsilon);$$

$$T = \frac{10^{20} t I^2 \log^3(I)}{\varepsilon^5 \beta} = I^2 \text{poly}(\log(I)/\varepsilon);$$

$$\eta = (\varepsilon/10T)^{10} = \text{poly}(\varepsilon/I); \quad r = 10^4 t m^2 / \varepsilon^2 = \text{poly}(I/\varepsilon);$$

$$\alpha = (0.001\varepsilon/m)^2 = \text{poly}(\varepsilon/I).$$

4.1 The Protocol τ

The protocol τ is formally presented in Algorithms 1-4. Below is an informal description of the input and output of each of the (sub)protocols used by τ . For the rest of the paper we consider the tree $\mathcal{T}$ corresponding to π , and use the sets, functions, and distributions defined with respect to π in Section 3.

The protocol τ (Algorithm 1).

At the beginning of the protocol it is assumed that Alice knows the input x and Bob knows the input y . The players' goal is to agree on a leaf of $\mathcal{T}$ distributed according to a distribution that is close (in statistical distance) to $\tilde{D}_{xy}$. The protocol starts at the root of $\mathcal{T}$ and proceeds in a sequence of m iterations. In every iteration a new vertex v is reached. For $i \in [m]$, the vertex v reached at the end of iteration i is a descendant of the vertex v reached at the end of iteration $i-1$. The vertex v reached by iteration m is a leaf with high probability.

Definitions.

Consider the protocol τ (Algorithm 1). For $i \in \{0, \dots, m\}$, let $\hat{\mathcal{V}}_i$ be the set of possible transcripts of the first i iterations of the loop in Line 1 of τ . Let $\hat{\mathcal{V}} = \bigcup_{i \in \{0, \dots, m\}} \hat{\mathcal{V}}_i$. In particular, any $\hat{v} \in \hat{\mathcal{V}}$ is a (partial) transcript of τ .

Let $(x, y) \in \text{supp}(\mu)$. As will be justified later (see the last paragraph of Section 4.2), we may assume that after any transcript $\hat{v} \in \hat{\mathcal{V}}_i$ of the first i iterations of τ , both Alice and Bob (each using his private input) know the vertex $v \in \mathcal{V}$ reached by τ at the end of this iteration. Let $\hat{v}$ be a possible (partial) transcript of τ , and assume that $\hat{v}'$ is the longest prefix of $\hat{v}$ contained in $\hat{\mathcal{V}}$. We denote by $\text{msg}_{\pi x}(\hat{v}) \in \mathcal{V}$ the vertex v reached by the protocol τ when the transcript of τ is $\hat{v}'$ and Alice's input is x . We denote by $\text{msg}_{\pi y}(\hat{v}) \in \mathcal{V}$ the vertex v reached by the protocol τ when the (partial) transcript of τ is $\hat{v}'$ and Bob's input is y . Since we assume that for every $(x, y) \in \text{supp}(\mu)$ it is the case that $\text{msg}_{\pi x}(\hat{v}) = \text{msg}_{\pi y}(\hat{v})$ (both Alice and Bob know the vertex $v \in \mathcal{V}$ reached), we often write $\text{msg}_{\pi xy}(\hat{v})$ instead of

---

Algorithm 1: The Protocol $\tau(x, y)$

```

1 for  $i \in [m]$  do
2   Players run  $\text{Chunk}(\hat{v}, x, y)$ , where  $\hat{v}$  is the transcript
   of  $\tau$  so far, to get a vertex  $v$ .
3 return  $v$ .
```

Algorithm 2: The Protocol $\text{Chunk}(\hat{v}, x, y)$

```

1 while no output was given do
2    $p \leftarrow$  random bit sampled from the public randomness.
3   if  $p = 1$  then
4     Players run  $\text{Sample}^B(\hat{v}, x, y)$  to get an output  $w$ .
5   else
6     Players run  $\text{Sample}^A(\hat{v}, x, y)$  to get an output  $w$ .
7   if  $w \neq \text{failure}$  then
8     return  $w$ .
```

Algorithm 3: The Protocol $\text{Sample}^B(\hat{v}, x, y)$

```

1  $v \leftarrow \text{msg}_{\pi xy}(\hat{v})$ .
2 if  $v \in \mathcal{L}^+$  then
3   Players return a leaf sampled according to  $\tilde{D}_{vxy}$ , by
   exchanging one bit for every level in  $\mathcal{L}^+$ .
4 Players run  $\text{CorrelatedSampling}(P = \tilde{D}_{vy}, Q = \tilde{D}_v, \eta)$ 
   to get a leaf  $u$ .
5 Players run  $\text{Separate}^B(\hat{v}, u, x, y)$  to get a vertex  $w$ .
6 Alice samples and sends a bit  $a$ : If  $w \notin \mathcal{V}_{vx}$ , then  $a = 0$ .
   Otherwise,  $a = 1$  with probability
    $\min\{1, D_{vx}(w)/tD_v(w)\}$ .
7 return  $w$  if  $a = 1$  and failure otherwise.
```

Algorithm 4: The Protocol $\text{Separate}^B(\hat{v}, u, x, y)$

```

1  $v \leftarrow \text{msg}_{\pi xy}(\hat{v})$ .
2 Let  $c_1 \leq \dots \leq c_r \in \mathbb{N}$  be s.t.  $c_i$  is the first index
   satisfying  $\Pr_{x' \leftarrow \mu_{\hat{v}}^A}[c_{vx'}(u) \leq c_i] \geq i/r$ .
3 Alice sends the first index  $a \in [r]$  such that  $c_{vx}(u) \leq c_a$ .
4 Alice sends a bit  $a'$  which is 1 iff  $c_{vx}(u) = c_a$ .
5 Bob sends the first index  $b \in [r]$  such that  $c_{vy}(u) \leq c_b$ .
6 if  $a \neq b$  then
7   return vertex number  $\min\{c_a, c_b\} + 1$  on the path
   from  $v$  to  $u$ .
8 else if  $a' = 1$  then
9   return vertex number  $c_a$  on the path from  $v$  to  $u$ .
10 else
11   Alice sends  $c_{vx}(u)$ , Bob sends  $c_{vy}(u)$ .
12   return vertex number  $\max\{c_{vx}(u), c_{vy}(u)\}$  on the
   path from  $v$  to  $u$ .
```

$\text{msg}_{\pi x}(\hat{v}), \text{msg}_{\pi y}(\hat{v})$. We note that $\text{msg}_{\pi xy}(\hat{v})$ is well defined: The vertex v reached by the protocol only depends on the inputs x, y and the transcript $\hat{v}$ (specifically, given $x, y, \hat{v}$, the vertex v is independent of the players' private randomness).

As before, (X, Y) is a pair of random variables distributed according to μ , representing the players' inputs. Recall that $\pi(X, Y)$ is the random variable representing the leaf of $\mathcal{T}$ reached by the protocol π . As explained in Section 3.3, we may assume without loss of generality that for every $(x, y) \in \text{supp}(\mu)$ and $\hat{v} \in \hat{\mathcal{V}}$, the transcript of the protocol τ may be $\hat{v}$. That is, we assume that $\Pr[\tau(X, Y)_{\leq |\hat{v}|} = \hat{v} \mid X = x, Y = y] > 0$.

Let $\mu_{\hat{v}}$ be the distribution μ conditioned on the event that the transcript of τ is $\hat{v}$. That is, $\mu_{\hat{v}} = ((X, Y) \mid \tau(X, Y)_{\leq |\hat{v}|} = \hat{v})$. Let $\mu_{\hat{v}}^A$ and $\mu_{\hat{v}}^B$ be the distributions μ^A and μ^B (respectively), conditioned on the event that the transcript of τ is $\hat{v}$. That is, $\mu_{\hat{v}}^A = (X \mid \tau(X, Y)_{\leq |\hat{v}|} = \hat{v})$ and $\mu_{\hat{v}}^B = (Y \mid \tau(X, Y)_{\leq |\hat{v}|} = \hat{v})$. Since $\hat{v}$ is a possible transcript of a protocol, and since $\mu = \mu^A \times \mu^B$ is a product distribution, $\mu_{\hat{v}}$ is also a product distribution and $\mu_{\hat{v}} = \mu_{\hat{v}}^A \times \mu_{\hat{v}}^B$. Observe that given $\hat{v}$, both players can compute $\mu_{\hat{v}}, \mu_{\hat{v}}^A, \mu_{\hat{v}}^B$.

The Protocol *Chunk* (Algorithm 2).

At the beginning of the protocol it is assumed that Alice knows the input x , Bob knows the input y , and that both players know the transcript $\hat{v}$ of τ so far. Let $v = \text{msg}_{\pi xy}(\hat{v})$. The players' goal is to agree on a vertex $w \in \mathcal{V}(v)$ satisfying $\mathbb{D}_{vwx} \geq \beta$ (if such w exists), while the following properties hold: The set $\mathcal{C}_{\hat{v}xy}$ of all the vertices w returned by the protocol is a frontier. In addition, every vertex $w \in \mathcal{C}_{\hat{v}xy}$ is selected by the players with probability close to $D_{vxy}(w)$.

The Protocol *Sample*^B (Algorithm 3).

At the beginning of the protocol it is assumed that Alice knows the input x , Bob knows the input y , and that both players know the transcript of τ up until the beginning of most recent execution of *Chunk*. This transcript is denoted $\hat{v} \in \hat{\mathcal{V}}$. Let $v = \text{msg}_{\pi xy}(\hat{v})$. The players' goal is to agree on a vertex $w \in \mathcal{V}(v)$ such that $w \in \mathcal{V}_{vx} \setminus \mathcal{V}_{vy}$ (in particular, $\mathbb{D}_{vwy} \geq \beta$ and $\mathbb{D}_{vwx} \leq 2\beta$, unless w is a leaf), while the following properties hold: The set of all the vertices w returned by the protocol is a subset of a frontier. This subset is "maximal" in the sense that one cannot add a new vertex $w \in \mathcal{V}_{vx} \setminus \mathcal{V}_{vy}$ to the set, while keeping it a subset of some frontier. In addition, every vertex w in this subset is selected by the players with probability close to $D_{vxy}(w)$. The protocol may fail.

The Protocol *Sample*^A.

The protocol *Sample*^A is obtained from the protocol *Sample*^B by switching the roles Alice and Bob, switching the roles of x and y , and running *Separate*^A instead of *Separate*^B.

The Protocol *Separate*^B (Algorithm 4).

At the beginning of the protocol it is assumed that Alice knows the input x , Bob knows the input y , and that both players know the transcript of τ up until the beginning of most recent execution of *Chunk*. This transcript is denoted $\hat{v} \in \hat{\mathcal{V}}$. It is also assumed that the players agree on a leaf $u \in \mathcal{L}(v)$, where $v = \text{msg}_{\pi xy}(\hat{v})$. The players' goal is to agree on a vertex $w \in P(v, u)$ such that $w \in (\mathcal{V}_{vx} \setminus \mathcal{V}_{vy}) \cup (\mathcal{V}_{vy} \setminus \mathcal{V}_{vx})$.

The Protocol *Separate*^A.

The protocol *Separate*^A is obtained from the protocol *Separate*^B by switching the roles Alice and Bob, switching the roles of x and y , and using $\mu_{\hat{v}}^B$ instead of $\mu_{\hat{v}}^A$.

4.2 The Protocol τ^{ST}

The protocol τ^{ST} gets the same parameters as τ . It operates the same as τ , expect for the following changes:

1. Add the following line at the beginning of the protocol: "At any point in the execution of this protocol, after the players exchange T bits, τ^{ST} terminates and returns failure".
2. Lines 11 and 12 of the protocols *Separate*^B and *Separate*^A are replaced by the single line " τ^{ST} returns failure". That is, if the "else" part in Line 10 of either *Separate*^B (Algorithm 3) or *Separate*^A is reached, then τ^{ST} terminates immediately and returns failure.

We mention that the protocol *CorrelatedSampling* is run by *Sample*^B and *Sample*^A with error parameter η , in order to sample a leaf $u \in \mathcal{L}(v)$. By Lemma 14, such an execution may result in the players getting two different leaves u^A and u^B with probability η . We ignore this possibility and assume that Alice and Bob always sample the same leaf u . Since *CorrelatedSampling* is run at most T times by τ^{ST} , the probability that Alice and Bob sample a different leaf u in one of the executions of *CorrelatedSampling*, is at most $T\eta$. We select η to be significantly smaller than $1/T$, thus $T\eta$ is negligible.

5. PROTOCOL ANALYSIS

In this section we prove Theorem 1. We restate Theorem 1 with slight change of notation, in order to simplify the notation in the proof (the original protocol is called here π' , and the simulation protocol is τ^{ST}).

Theorem (Theorem 1 restated). *Let $\varepsilon > 0$. Let π' be a randomized protocol that may use private and public coins. Let $\mu = \mu^A \times \mu^B$ be a product distribution over the inputs for π' . Then, there exist a public coin protocol τ^{ST} (that takes the same inputs as π'), and a pair of "transcript reconstruction" functions g^A, g^B such that: The function g^A takes as inputs $x \in \text{supp}(\mu^A)$ and a possible transcript of τ^{ST} , and returns a possible transcript of π' . The function g^B takes as inputs $y \in \text{supp}(\mu^B)$ and a possible transcript of τ^{ST} , and returns a possible transcript of π' . In addition, the followings hold:*

1. The worst case communication complexity of τ is $\text{IC}_\mu^2(\pi') \cdot \text{polylog}(\text{IC}_\mu(\pi')) / \varepsilon^5$.
2. $\forall(x, y) \in \text{supp}(\mu) : \Pr[g^A(x, \tau^{ST}(x, y)) \neq g^B(y, \tau^{ST}(x, y))] \leq \varepsilon$, where the probability is over the random coins of the protocol τ^{ST} .
3. $\mathbf{E}_{(x, y) \leftarrow \mu} [\|g^A(x, \tau^{ST}(x, y)) - \pi'(x, y)\|] \leq \varepsilon$.

Let π' be the given randomized communication protocol, and let μ be a distribution over the inputs for π' . As explained before, we may assume that π' is a private coin protocol, as the public randomness can always be replaced by private randomness without increasing the information cost. Let $I' = \text{Ext}_\mu(\pi')$, and assume that I' is sufficiently large. Let π be the β' -smooth protocol promised by Lemma 13 for π' , when applied with $\kappa = \varepsilon/10$ and $\beta' = \frac{\varepsilon}{\log^c(I' + \varepsilon/10)}$ (the same constant c as in the definition of the parameter β). Let $I = \text{Ext}_\mu(\pi)$. By Lemma 13, $I \leq I' + \varepsilon/10$. Thus, $\beta' \leq \beta = \frac{\varepsilon}{\log^c(I)}$, and, in particular, π is also β -smooth. The simulation protocol, τ^{ST} , is used to simulate the protocol π . Our main lemma is the following:

Lemma 15. *It holds that*

$$\mathbf{E}_{(x, y) \leftarrow \mu} [\|\text{msg}_{\pi x}(\tau^{ST}(x, y)) - \pi(x, y)\|] \leq \varepsilon/2.$$

The proof of Lemma 15 can be found in the full version of this paper. We next show how Theorem 1 follows from Lemma 15.

Proof of Theorem 1. The first item in Theorem 1 is satisfied as the worst case communication complexity of π^{ST} is $T \leq I^2 \text{polylog}(I) / \varepsilon^5 \leq (I')^2 \text{polylog}(I') / \varepsilon^5$, and as, by Fact 3, $I' = \text{Ext}_\mu(\pi') = \text{IC}_\mu(\pi')$.

Let f^A be the function that gets an input $x \in \text{supp}(\mu^A)$ and a possible transcript $\hat{v}$ for τ^{ST} , and returns the transcript $\text{msg}_{\pi x}(\hat{v}) \in \mathcal{V}$ for π . Let f^B be the function that gets an input $y \in \text{supp}(\mu^B)$ and a possible transcript $\hat{v}$ for τ^{ST} , and returns the transcript $\text{msg}_{\pi y}(\hat{v}) \in \mathcal{V}$ for π . By Lemma 13, there exists a function f such that $\|f(\pi(x, y)) - \pi'(x, y)\| \leq \varepsilon/10$. Let $g^A = f \circ f^A$ and $g^B = f \circ f^B$, where ‘ $\circ$ ’ denotes function composition.

The second item in Theorem 1 is satisfied, because, as explained in the last paragraph of Section 4.2, Alice and Bob reach a different vertex of $\mathcal{T}$ with probability at most $T\eta \ll \varepsilon$.

The third item in Theorem 1 is satisfied as by Lemma 15,

$$\begin{aligned} & \mathbf{E}_{(x, y) \leftarrow \mu} [\|g^A(x, \tau^{ST}(x, y)) - \pi'(x, y)\|] \\ & \leq \mathbf{E}_{(x, y) \leftarrow \mu} [\|f(f^A(x, \tau^{ST}(x, y))) - f(\pi(x, y))\| \\ & \quad + \|f(\pi(x, y)) - \pi'(x, y)\|] \\ & \leq \mathbf{E}_{(x, y) \leftarrow \mu} [\|\text{msg}_{\pi x}(\tau^{ST}(x, y)) - \pi(x, y)\| \\ & \quad + \|f(\pi(x, y)) - \pi'(x, y)\|] \\ & \leq \varepsilon/2 + \varepsilon/10 \leq \varepsilon. \end{aligned}$$

□

6. ACKNOWLEDGEMENTS

We thank Avi Wigderson for numerous helpful discussions.

7. REFERENCES

- [BBCR10] Boaz Barak, Mark Braverman, Xi Chen, and Anup Rao. How to compress interactive communication. In *STOC*, pages 67–76, 2010.
- [BBK⁺13] Joshua Brody, Harry Buhrman, Michal Koucký, Bruno Loff, Florian Speelman, and Nikolay K. Vereshchagin. Towards a reverse newman’s theorem in interactive information complexity. In *CCC*, pages 24–33, 2013.
- [BMV15] Balthazar Bauer, Shay Moran, and Amir Yehudayoff. Internal compression of protocols to entropy. In *APPROX/RANDOM*, pages 481–496, 2015.
- [BR11] Mark Braverman and Anup Rao. Information equals amortized communication. In *FOCS*, pages 748–757, 2011.
- [Bra12] Mark Braverman. Interactive information complexity. In *STOC*, pages 505–524, 2012.
- [BYJKS04] Ziv Bar-Yossef, T. S. Jayram, Ravi Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. *J. Comput. Syst. Sci.*, 68(4):702–732, 2004.
- [CSWY01] Amit Chakrabarti, Yaoyun Shi, Anthony Wirth, and Andrew Chi-Chih Yao. Informational complexity and the direct sum problem for simultaneous message complexity. In *FOCS*, pages 270–278, 2001.
- [Fan49] Robert M Fano. *The transmission of information*. Massachusetts Institute of Technology, Research Laboratory of Electronics, 1949.
- [Gan12] Anat Ganor. *Efficient Communication (Master’s Thesis)*. Weizmann Institute of Science, 2012.
- [GKR14] Anat Ganor, Gillat Kol, and Ran Raz. Exponential separation of information and communication. In *FOCS*, pages 176–185, 2014.
- [GKR15a] Anat Ganor, Gillat Kol, and Ran Raz. Exponential separation of communication and external information. *Electronic Colloquium on Computational Complexity (ECCC)*, 22:88, 2015.
- [GKR15b] Anat Ganor, Gillat Kol, and Ran Raz. Exponential separation of information and communication for boolean functions. In *STOC*, pages 557–566, 2015.
- [HJMR10] Prahladh Harsha, Rahul Jain, David A. McAllester, and Jaikumar Radhakrishnan. The communication complexity of correlation. *IEEE Transactions on Information Theory*, 56(1):438–449, 2010.

- [Huf52] David A Huffman. A method for the construction of minimum redundancy codes. *proc. IRE*, 40(9):1098–1101, 1952.
- [RR15] Sivaramakrishnan Natarajan Ramamoorthy and Anup Rao. How to compress asymmetric communication. In *CCC*, pages 102–123, 2015.
- [RS15] Anup Rao and Makrand Sinha. Simplified separation of information and communication. *Electronic Colloquium on Computational Complexity (ECCC)*, 22:57, 2015.
- [Sha48] C. E. Shannon. A mathematical theory of communication. *The Bell Systems Technical Journal*, 27:July 379–423, October 623–656, 1948.