

Communication Complexity of Set-Disjointness for All Probabilities*

Mika Göös and Thomas Watson

Department of Computer Science, University of Toronto
Toronto, ON, Canada
{mgoos, thomasw}@cs.toronto.edu

Abstract

We study set-disjointness in a generalized model of randomized two-party communication where the probability of acceptance must be at least $\alpha(n)$ on yes-inputs and at most $\beta(n)$ on no-inputs, for some functions $\alpha(n) > \beta(n)$. Our main result is a complete characterization of the private-coin communication complexity of set-disjointness for all functions α and β , and a near-complete characterization for public-coin protocols. In particular, we obtain a simple proof of a theorem of Braverman and Moitra (STOC 2013), who studied the case where $\alpha = 1/2 + \epsilon(n)$ and $\beta = 1/2 - \epsilon(n)$. The following contributions play a crucial role in our characterization and are interesting in their own right.

1. We introduce two communication analogues of the classical complexity class that captures *small bounded-error* computations: we define a “restricted” class SBP (which lies between MA and AM) and an “unrestricted” class USBP. The distinction between them is analogous to the distinction between the well-known communication classes PP and UPP.
2. We show that the SBP communication complexity is precisely captured by the classical *corruption* lower bound method. This sharpens a theorem of Klauck (CCC 2003).
3. We use information complexity arguments to prove a linear lower bound on the USBP complexity of set-disjointness.

1998 ACM Subject Classification F.1.3 Complexity Measures and Classes

Keywords and phrases Communication Complexity, Set-Disjointness, All Probabilities

Digital Object Identifier 10.4230/LIPIcs.APPROX-RANDOM.2014.721

1 Introduction

In the *set-disjointness* problem, Alice is given an $x \subseteq [n]$, Bob is given a $y \subseteq [n]$, and their task is to decide whether $x \cap y = \emptyset$. Equivalently, viewing x and y as binary strings, we define

$$\text{DISJ}(x, y) := \neg \bigvee_{i \in [n]} (x_i \wedge y_i).$$

Set-disjointness is the preeminent coNP-complete problem in communication complexity [2, 13]. A fundamental result of Kalyanasundaram and Schnitger [23] (with alternative proofs given by [31, 4]) states that every randomized protocol for set-disjointness requires $\Omega(n)$ bits of communication to achieve a constant error probability that is bounded away from $1/2$. These lower bounds have been extremely useful in applications of communication complexity to other areas of theoretical computer science, including circuit complexity,

* This work was supported by funding from NSERC.



© Mika Göös and Thomas Watson;

licensed under Creative Commons License CC-BY

17th Int'l Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX'14) /
18th Int'l Workshop on Randomization and Computation (RANDOM'14).

Editors: Klaus Jansen, José Rolim, Nikhil Devanur, and Cristopher Moore; pp. 721–736



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

distributed computing, streaming, data structures, combinatorial optimization, and more; see [28, 22, 13].

In this work, we study set-disjointness in a generalized setting where the probability of acceptance must be at least $\alpha(n)$ on yes-inputs and at most $\beta(n)$ on no-inputs, for any prescribed functions $\alpha(n) > \beta(n)$.

1.1 Main Result

Our main result is a complete characterization of the private-coin communication complexity of set-disjointness for all functions α and β , and a near-complete characterization for public-coin protocols. Roughly speaking, we prove that the randomized complexity is

$$\Theta(n \cdot (1 - \beta/\alpha))$$

for typical functions α and β ; see subsection 1.4 for the statement of the exact bounds.

As a special case, we obtain a simple proof of a result of Braverman and Moitra [6]. They showed that the communication complexity of set-disjointness is $\Theta(\epsilon n)$ in case $\alpha = 1/2 + \epsilon(n)$ and $\beta = 1/2 - \epsilon(n)$. While this special case might suggest that the complexity is determined by the additive gap $\alpha - \beta$, our characterization reveals that, in fact:

Central tenet: *It is not the additive gap between α and β that determines the complexity of set-disjointness; what matters is the multiplicative gap.*

Our proof follows this ideology: we show that in order to understand the communication complexity for all α and β it suffices to understand the *small bounded-error* case where α is tiny (e.g., exponentially small in n) and $\beta = \alpha/2$.

1.2 SBP: Small Bounded-error Probabilities

In classical time-bounded (i.e., poly-time Turing machine) complexity theory, small bounded-error acceptance probabilities are captured by a counting class called SBP, which was introduced by Böhler, Glaßer, and Meister [5] and has also been studied in [38]. In particular, [5] observed that SBP is sandwiched between the Arthur–Merlin classes MA and AM [3].

In this work, we introduce two communication complexity analogues of SBP: a *restricted* class called SBP, and an *unrestricted* class called USBP. These classes are natural and interesting in their own right. Most importantly, they serve to structure our argument.

Randomized Communication Complexity. In what follows, we assume familiarity with basic definitions of communication complexity [28, 22]. Fix a two-party function $f: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ where on input (x, y) Alice is given x and Bob is given y . We say (x, y) is a b -input if $(x, y) \in f^{-1}(b)$. We let $R_{\alpha, \beta}^{\text{pub}}(f)$, respectively $R_{\alpha, \beta}^{\text{priv}}(f)$, denote the minimum communication complexity (as a function of n) of a public-coin, respectively private-coin, protocol for f such that the probability of acceptance is at least $\alpha(n)$ on all 1-inputs and at most $\beta(n)$ on all 0-inputs. As is customary [2], for any communication measure $C(f)$ we often let C stand for the class of functions f with $C(f) = \text{polylog}(n)$.

PP and UPP. To motivate our upcoming definitions for SBP, we take a little detour and recall the communication classes associated with the standard complexity class PP. There

are in fact two distinct measures—*restricted* and *unrestricted*—as introduced in [2, 30]:

$$\begin{aligned} \text{PP}(f) &:= \min_{\epsilon(n) > 0} R_{1/2 + \epsilon, 1/2 - \epsilon}^{\text{pub}}(f) + \log(1/\epsilon), \\ \text{UPP}(f) &:= \min_{\epsilon(n) > 0} R_{1/2 + \epsilon, 1/2 - \epsilon}^{\text{priv}}(f). \end{aligned}$$

In the (restricted) public-coin model, one needs to charge the additional $\log(1/\epsilon)$ term in order for the measure to be well-behaved when ϵ is tiny. (For example, note that $R_{1/2 + \epsilon, 1/2 - \epsilon}^{\text{pub}}(f) \leq 2$ for $\epsilon = 2^{-n-1}$.) The original definition of $\text{PP}(f)$ given in [2] actually charged for the number of public coin flips *instead* of the $+\log(1/\epsilon)$; however, by standard sparsification techniques (see [29] and [28, Theorem 3.14]) the two versions are essentially equivalent—they are within a constant factor plus $O(\log n)$ —and the definition we have stated is much more prevalent in recent literature. It also follows from standard sparsification that we may convert any PP protocol into a UPP protocol of comparable cost: $\text{UPP}(f) \leq O(\text{PP}(f) + \log n)$. In the converse direction, an exponential separation between UPP and PP is known [8, 33, 34].

SBP and USBP. Analogously to the above, we define

$$\begin{aligned} \text{SBP}(f) &:= \min_{\alpha(n) > 0} R_{\alpha, \alpha/2}^{\text{pub}}(f) + \log(1/\alpha), \\ \text{USBP}(f) &:= \min_{\alpha(n) > 0} R_{\alpha, \alpha/2}^{\text{priv}}(f). \end{aligned}$$

Here the constant factor $1/2 = \beta/\alpha$ can be replaced by any positive constant less than 1 while affecting the complexity measures by only a constant factor: if we run a protocol ℓ times and accept iff all iterations accept, then β/α gets raised to the power ℓ while the communication and the $\log(1/\alpha)$ term each get multiplied by ℓ . We call this procedure *and-amplification* (in contrast to the usual *majority-amplification*). We also note that by standard sparsification, $\text{USBP}(f) \leq O(\text{SBP}(f) + \log n)$ holds for all f . In the converse direction, we do not know whether USBP is significantly more powerful than SBP (though a small separation is witnessed by the greater-than function, which has constant USBP complexity but $\Theta(\log n)$ SBP and PP complexity [7]).

Relationship to Arthur–Merlin Classes. Klauck [25, 27] and Aaronson and Wigderson [1] took up the study of communication complexity analogues of Arthur–Merlin games. Their results have already found applications in data streaming [9, 10, 19]. We do not define the communication models MA and AM here, but we note that the classical inclusions continue to hold in the communication setting (for the same reasons):

$$\text{MA} \subseteq \text{SBP} \subseteq \text{AM}.$$

Indeed, if $\text{MA}(f) = m$ then by majority-amplification and by absorbing Merlin’s nondeterminism into the randomness we obtain $R_{2^{-m-1}, 2^{-m-2}}^{\text{pub}}(f) \leq O(m^2)$. Thus $\text{SBP}(f) \leq O(\text{MA}(f)^2)$ (and the quadratic blow-up is necessary for “black-box” simulations [15]). On the other hand, $\text{AM}(f) \leq O(\text{SBP}(f) + \log n)$ holds by sparsifying the randomness and using the Goldwasser–Sipser protocol [18].

1.3 Results for SBP and USBP

We prove that SBP communication complexity is exactly characterized by the well-known *corruption* lower bound method (also known as the *rectangle bound* or *one-sided discrepancy*).

The definition of the corruption bound $\text{Corr}(f)$ is given in section 2, but for now, we note that $\text{Corr}(f)$ essentially depends on the size of the largest approximately 1-monochromatic rectangle in the communication matrix of f . (For an extensive discussion of the different lower bound methods in communication complexity, see [20].) Previously, Klauck [25] showed that $\text{Corr}(f)$ lies somewhere between the MA and AM communication complexities of f ; namely $\Omega(\text{AM}(f)) \leq \text{Corr}(f) \leq O(\text{MA}(f)^2)$. Klauck also gave a combinatorial near-characterization of $\text{Corr}(f)$ (tight up to logarithmic factors) using so-called *one-sided uniform threshold covers*. The following theorem sharpens these results by pinpointing precisely the class between MA and AM that is characterized by corruption.

► **Theorem 1.** $\text{SBP}(f) = \Theta(\text{Corr}(f))$ for all f .

One way to frame Theorem 1 is as follows. A lot of effort (e.g., [25, 20, 24, 21, 17]) has been spent on comparing the relative strengths of different lower bound methods in communication complexity with the goal of finding a natural method that captures the bounded-error randomized communication complexity of every function. Theorem 1 can be viewed as achieving a diametrically opposite goal: we start with a historically important lower bound method (i.e., corruption) and find a natural communication measure that it captures. Theorem 1 is also somewhat analogous, in content and proof, to another result of Klauck [26] showing that the discrepancy bound captures PP.

Razborov [31] famously proved that $\text{Corr}(\text{DISJ}) = \Theta(n)$. (The first linear lower bound for set-disjointness [23] did not use corruption.) By the results of [25], this implies that $\text{MA}(\text{DISJ}) \geq \Omega(\sqrt{n})$. We immediately have a stronger corollary.

► **Corollary 2.** $\text{SBP}(\text{DISJ}) = \Theta(n)$.

While the classical rectangle-based methods suffice to analyze SBP protocols, these techniques are not well-suited for handling acceptance probabilities $\alpha(n)$ that are arbitrarily small functions of n (e.g., doubly exponentially small in n). To obtain lower bounds for USBP we pursue a different avenue and show that the *information complexity* framework, as formulated by Bar-Yossef, Jayram, Kumar, and Sivakumar [4] (see also [12]), can be adapted to suit our purposes. The main technical result of this work is the following, proved in section 3.

► **Theorem 3.** $\text{USBP}(\text{DISJ}) = \Theta(n)$.

We note that the statement of Theorem 3 is similar in spirit to Forster's theorem [16] stating that the UPP complexity of the inner product function is $\Theta(n)$. Note also that Theorem 2 is of course a corollary of Theorem 3, too, but the corruption-based proof via Theorem 1 is arguably more elementary than the proof of Theorem 3. Finally, we note that the well-studied gap-Hamming-distance promise problem [11, 37, 36] (where 1-inputs have distance $\geq \frac{n}{2} + \sqrt{n}$ and 0-inputs have distance $\leq \frac{n}{2} - \sqrt{n}$) has SBP and USBP complexities $\Theta(\sqrt{n})$, where the lower bound follows by Theorem 3 and a standard reduction from DISJ, and the upper bound follows by and-amplification of the trivial protocol that checks inequality at a random bit position.

1.4 Characterization for All α and β

Using our results for SBP and USBP in a black-box manner we derive the following (near) complete characterization for the randomized communication complexity of set-disjointness in section 4.

► **Theorem 4** (Private-coin). *For all $\alpha(n) > \beta(n)$,*

$$R_{\alpha, \beta}^{\text{priv}}(\text{DISJ}) = \Theta(n \cdot (1 - \beta/\alpha) + \log n).$$

► **Theorem 5** (Public-coin). *There is a universal constant $C > 0$ such that for all $\alpha(n) > \beta(n)$,*

$$R_{\alpha, \beta}^{\text{pub}}(\text{DISJ}) = \begin{cases} \Theta(n \cdot (1 - \beta/\alpha)) & \text{when } \log(1/\alpha) \leq C \cdot n \cdot (1 - \beta/\alpha), \\ 2 & \text{when } \log(1/\alpha) \geq \lceil n \cdot (1 - \beta/\alpha) \rceil. \end{cases}$$

We stress that for the public-coin characterization (and in particular, the result of [6] as a corollary), it suffices to rely only on Razborov’s corruption lemma (via Theorem 2), and not on any information complexity techniques. Braverman and Moitra [6] observed that $R_{1/2 + \epsilon, 1/2 - \epsilon}^{\text{pub}}(\text{DISJ}) \geq \Omega(\epsilon^2 n)$ follows from the standard bounded-error lower bound by majority-amplification, and they obtained the tight $\Omega(\epsilon n)$ bound by developing information complexity techniques tailored to this setting. Our idea is that and-amplification imposes only an ϵ factor loss (rather than the ϵ^2 factor loss imposed by majority-amplification) while still reducing to a case where the corruption method applies.

We also note that for public-coin protocols there remains a small gap in the parameters around the threshold $\log(1/\alpha) = \Theta(n \cdot (1 - \beta/\alpha))$ that is not covered by our theorem. As we discuss in section 4, the power of the public coins kicks in at this threshold.

Finally, we mention that all the set-disjointness lower bounds in this paper continue to hold under the *unique-intersection* promise where the inputs are either disjoint or intersect in exactly one coordinate: for Theorem 2 this property is inherited from Razborov’s proof; for Theorem 3 this property is implicit in our proof.

2 SBP is Characterized by Corruption

In this section we prove Theorem 1, which states that $\text{SBP}(f) = \Theta(\text{Corr}(f))$ for all f . We start by defining the corruption bound. We say a distribution μ over inputs is *balanced* (with respect to f) if $\mu(f^{-1}(1)) = \mu(f^{-1}(0)) = 1/2$. We say a rectangle R is *1-biased* (with respect to f and μ) if $\mu(R \cap f^{-1}(0)) \leq \mu(R)/8$. The corruption bound is defined as

$$\text{Corr}(f) := \max_{\text{balanced } \mu} \min_{\text{1-biased } R} \log \left(\frac{1}{\mu(R)} \right).$$

It was proved in [25] that the constant factor of $1/8$ (in the definition of 1-biased) can be replaced by any positive constant at most $1/8$ while affecting the corruption bound by only a constant factor. It was also proved in [25] that the bound is robust with respect to the balance condition on μ .

2.1 SBP is Lower Bounded by Corruption

Here we show the lower bound $\text{SBP}(f) \geq \Omega(\text{Corr}(f))$. The intuition is as follows. The first step is to fix the public randomness of an SBP protocol in such a way that the average-case behavior of the resulting deterministic protocol mimics the worst-case behavior of the original protocol. Typically, this sort of thing is done by invoking the distributive law (linearity of expectation), but here we need a more elaborate calculation due to the asymmetric nature of SBP. Then, the rest of the argument follows along similar lines as the proof in [25] (that $\text{Corr}(f) \leq O(\text{MA}(f)^2)$), showing that the 1-inputs are mostly covered by “small” transcript rectangles (of our average-case protocol), hence many such rectangles are needed.

We proceed with the formal proof. Let Π be an $R_{\alpha, \alpha/32}^{\text{pub}}$ protocol for f ; recall that by and-amplification we may assume $\beta = \alpha/32$ rather than $\beta = \alpha/2$ in the definition of SBP. Assuming $\log(1/\alpha) < \text{Corr}(f)/2$, we show that Π uses $\Omega(\text{Corr}(f))$ bits of communication. To this end, fix a balanced distribution μ such that for all 1-biased rectangles R , $\mu(R) \leq 2^{-\text{Corr}(f)}$.

Identify the possible outcomes of public randomness with $\{1, \dots, m\}$, and let Π_i denote Π running with public randomness i . Let p_i be the probability the public randomness is i (so $p_i = 1/m$ if the public randomness is uniformly distributed). Let q_i be the probability over μ that Π_i accepts, conditioned on the input being a 1-input. Let r_i be the same but conditioned on a 0-input. Now

$$\sum_i p_i q_i = \Pr_{i, (x,y) \sim \mu} [\Pi_i(x,y) \text{ accepts} \mid f(x,y) = 1] \geq \alpha, \quad (1)$$

$$\sum_i p_i r_i = \Pr_{i, (x,y) \sim \mu} [\Pi_i(x,y) \text{ accepts} \mid f(x,y) = 0] \leq \alpha/32. \quad (2)$$

► **Claim 6.** There exists an i^* such that $q_{i^*} \geq \alpha/2$ and $r_{i^*} \leq q_{i^*}/16$.

Proof of claim. Suppose for contradiction that for all i either $q_i < \alpha/2$ or $r_i > q_i/16$. Let $S \subseteq \{1, \dots, m\}$ be such that for all $i \in S$, $q_i < \alpha/2$, and for all $i \in \bar{S}$, $r_i > q_i/16$. Then

$$\begin{aligned} \sum_i p_i r_i &\geq \sum_{i \in \bar{S}} p_i r_i \geq \sum_{i \in \bar{S}} p_i q_i / 16 \\ &= \frac{1}{16} \left(\sum_i p_i r_i - \sum_{i \in S} p_i r_i \right) \geq \frac{1}{16} \left(\alpha - \left(\sum_{i \in S} p_i \right) \alpha / 2 \right) \geq \alpha / 32. \end{aligned}$$

Furthermore, at least one of the inequalities must be strict, contradicting (2). ◀

Fix an i^* guaranteed by Claim 6. Using i^* as the public randomness in Π , we can now apply the usual corruption argument. Consider the 1-rectangles that correspond to accepting transcripts of Π_{i^*} . Call a 1-rectangle R *large* if $\mu(R) > 2^{-\text{Corr}(f)}$ and *small* otherwise. Recall that by our assumption on μ , no large 1-rectangle is 1-biased: for every large 1-rectangle R we have $\mu(R \cap f^{-1}(0)) > \mu(R)/8$. Under μ , the total measure of large 1-rectangles is at most half the total measure of all 1-rectangles, since otherwise

$$\begin{aligned} r_{i^*} &= 2 \Pr_{(x,y) \sim \mu} [\Pi_{i^*}(x,y) \text{ accepts and } f(x,y) = 0] \\ &= 2 \sum_{\text{1-rectangles } R} \mu(R \cap f^{-1}(0)) \\ &\geq 2 \sum_{\text{large 1-rectangles } R} \mu(R \cap f^{-1}(0)) \\ &> \frac{1}{4} \sum_{\text{large 1-rectangles } R} \mu(R) \\ &> \frac{1}{4} \cdot \frac{1}{2} \sum_{\text{1-rectangles } R} \mu(R) \\ &\geq \frac{1}{8} \sum_{\text{1-rectangles } R} \mu(R \cap f^{-1}(1)) \\ &= \frac{1}{8} \Pr_{(x,y) \sim \mu} [\Pi_{i^*}(x,y) \text{ accepts and } f(x,y) = 1] \\ &= \frac{1}{8} \cdot q_{i^*} / 2 \\ &= q_{i^*} / 16. \end{aligned}$$

Therefore $\sum_{\text{small 1-rectangles } R} \mu(R) \geq \frac{1}{2} \sum_{\text{1-rectangles } R} \mu(R) \geq \frac{1}{2} \cdot q_{i^*} / 2 \geq \alpha/8 > 2^{-\text{Corr}(f)/2-3}$. Thus there are at least $2^{-\text{Corr}(f)/2-3} / 2^{-\text{Corr}(f)} = 2^{\text{Corr}(f)/2-3}$ small 1-rectangles, which implies that Π uses at least $\text{Corr}(f)/2 - 3$ bits of communication.

2.2 SBP is Upper Bounded by Corruption

Here we show the upper bound $\text{SBP}(f) \leq O(\text{Corr}(f))$. The intuition is as follows. If the corruption bound is small, that means for every balanced distribution over inputs there exists a rectangle (which can be viewed as a 2-bit protocol) that exhibits average-case SBP-like behavior—accepting a random 1-input with not-too-small probability, and accepting a random 0-input with constant-factor-smaller probability. We use the minimax theorem to convert this property into a distribution over rectangles, with a worst-case SBP guarantee. Several technical issues arise with this argument. One is the asymmetry between 1-inputs and 0-inputs, but this can be massaged away using a linear transformation of probabilities before invoking minimax. Another is that the corruption bound can yield an average-case SBP rectangle with a different “ α ” for different balanced distributions, whereas the minimax application requires a single α to work uniformly for all balanced distributions. This issue is fixed by passing to an appropriate subrectangle to decrease the α if necessary, for any given balanced distribution.

We proceed with the formal proof. For notational convenience we let $\mathbf{0}$ and $\mathbf{1}$ stand for the events $f^{-1}(0)$ and $f^{-1}(1)$, respectively. For example, $\mu(\mathbf{0} | R) = \mu(R \cap f^{-1}(0)) / \mu(R)$ and $\mu(R | \mathbf{0}) = \mu(R \cap f^{-1}(0)) / \mu(f^{-1}(0))$.

Define $\alpha = 2^{-\text{Corr}(f)}$.

► **Claim 7.** For every balanced μ there exists a rectangle R with $\mu(R | \mathbf{1}) \geq \alpha$ and $\mu(R | \mathbf{0}) \leq \alpha/2$.

Proof of claim. Fix a balanced distribution μ . By definition of corruption, there exists a rectangle S such that $\mu(S) \geq \alpha$ and $\mu(\mathbf{0} | S) \leq 1/8$. Decompose S as the disjoint union $S_1 \cup S_2 \cup \dots \cup S_m$ where the S_i ’s are the individual rows of S , sorted in nondecreasing order of $\mu(\mathbf{0} | S_i)$. Let $R_i = S_1 \cup S_2 \cup \dots \cup S_i$. For every i we know that $\mu(\mathbf{0} | R_i) \leq \mu(\mathbf{0} | S) \leq 1/8$. If there exists an i such that $\alpha \leq \mu(R_i | \mathbf{1}) \leq 2\alpha$ then $R = R_i$ witnesses the claim since

$$\mu(R_i | \mathbf{0}) = \frac{\mu(\mathbf{0} | R_i) \cdot \mu(R_i | \mathbf{1}) \cdot \mu(\mathbf{1})}{\mu(\mathbf{0}) \cdot \mu(\mathbf{1} | R_i)} \leq \frac{(1/8) \cdot 2\alpha \cdot (1/2)}{(1/2) \cdot (7/8)} = 2\alpha/7 \leq \alpha/2.$$

Otherwise, since $\mu(R_m | \mathbf{1}) = \mu(S | \mathbf{1}) = \mu(\mathbf{1} | S) \cdot \mu(S) / \mu(\mathbf{1}) \geq (7/8) \cdot \alpha / (1/2) > \alpha$ and $\mu(R_0 | \mathbf{1}) = 0 < \alpha$, there must exist an i such that $\mu(R_i | \mathbf{1}) > 2\alpha$ and $\mu(R_{i-1} | \mathbf{1}) < \alpha$ and thus $\mu(S_i | \mathbf{1}) > 2\alpha - \alpha = \alpha$. In this case, the rectangle $R = S_i \cap \mathbf{1}$ witnesses the claim since $\mu(S_i \cap \mathbf{1} | \mathbf{1}) = \mu(S_i | \mathbf{1}) > \alpha$ and $\mu(S_i \cap \mathbf{1} | \mathbf{0}) = 0 \leq \alpha/2$. ◀

Let M be the matrix with rows indexed by inputs $(x, y) \in \{0, 1\}^n \times \{0, 1\}^n$ and columns indexed by rectangles $R \subseteq \{0, 1\}^n \times \{0, 1\}^n$ such that

$$M_{(x,y),R} = \begin{cases} 1 & \text{if } f(x,y) = 1 \text{ and } (x,y) \in R \\ 0 & \text{if } f(x,y) = 1 \text{ and } (x,y) \notin R \\ 0 & \text{if } f(x,y) = 0 \text{ and } (x,y) \in R \\ \frac{\alpha}{1-\alpha/2} & \text{if } f(x,y) = 0 \text{ and } (x,y) \notin R \end{cases}.$$

We claim that for every distribution μ over inputs, there exists a rectangle R such that $\mathbb{E}(M_{\mu,R}) \geq \alpha$ (where \mathbb{E} denotes expectation). If $\mu(\mathbf{0}) = 0$ then take $R = \{0, 1\}^n \times \{0, 1\}^n$, and if $\mu(\mathbf{1}) = 0$ then take $R = \emptyset$. Otherwise, let μ' be the balanced version of μ and invoke Claim 7 to find an R such that $\mu'(R | \mathbf{1}) \geq \alpha$ and $\mu'(R | \mathbf{0}) \leq \alpha/2$. Then we have

$$\mathbb{E}(M_{\mu,R}) = \mu(R | \mathbf{1}) \cdot \mu(\mathbf{1}) + \frac{\alpha}{1-\alpha/2} \cdot \mu(\overline{R} | \mathbf{0}) \cdot \mu(\mathbf{0})$$

$$\begin{aligned}
&= \mu'(R|\mathbf{1}) \cdot \mu(\mathbf{1}) + \frac{\alpha}{1-\alpha/2} \cdot \mu'(\overline{R}|\mathbf{0}) \cdot \mu(\mathbf{0}) \\
&\geq \alpha \cdot \mu(\mathbf{1}) + \frac{\alpha}{1-\alpha/2} \cdot (1-\alpha/2) \cdot \mu(\mathbf{0}) \\
&= \alpha.
\end{aligned}$$

Now by the minimax theorem, there exists a distribution D over rectangles such that for every input (x, y) , $\mathbb{E}(M_{(x,y),D}) \geq \alpha$. If $f(x, y) = 1$ this means the probability a random rectangle from D contains (x, y) is at least α . If $f(x, y) = 0$ this means $\frac{\alpha}{1-\alpha/2}$ times the probability a random rectangle from D does not contain (x, y) is at least α , in other words the probability a random rectangle from D contains (x, y) is at most $\alpha/2$. Thus the protocol that picks a random rectangle from D and accepts iff the input is in the rectangle shows that $R_{\alpha, \alpha/2}^{\text{pub}}(f) \leq 2$ and hence $\text{SBP}(f) \leq 2 + \log(1/\alpha) = 2 + \text{Corr}(f)$.

3 USBP Lower Bound

In this section we prove Theorem 3, which states that $\text{USBP}(\text{DISJ}) = \Theta(n)$. We first give an informal overview.

Our proof uses the by-now standard information complexity approach [4, 12]. In this approach, one considers some suitably distributed random input (X, Y) and measures the amount of information that the protocol transcript $\Pi(X, Y)$ (i.e., the concatenation of all messages sent) “leaks” about the input as quantified by the *mutual information* $\mathbb{I}(\Pi(X, Y); X, Y)$. Lower bounding the mutual information has the side effect of lower bounding the *entropy* $\mathbb{H}(\Pi(X, Y))$ of the transcript, which in turn lower bounds the length of the transcript and thereby the communication complexity. It is often useful to involve the *conditional* versions of these information measures, defined by $\mathbb{H}(\Pi|Z) = \mathbb{E}_{z \sim Z} \mathbb{H}(\Pi|Z=z)$ and $\mathbb{I}(\Pi; X, Y|Z) = \mathbb{E}_{z \sim Z} \mathbb{I}(\Pi; X, Y|Z=z)$ where Z is some random variable (jointly distributed with X and Y). We refer the reader to [14] for discussions of these basic information theory concepts.

A key benefit of studying mutual information is that one automatically obtains for it a *direct sum* property (as in [12, 4]), as long as the coordinates (X_i, Y_i) , $i \in [n]$, are mutually independent. This way, the task of proving an $\Omega(n)$ lower bound for the original problem reduces to the task of proving an $\Omega(1)$ information lower bound for some constant-size “gadget”. For set-disjointness $\text{DISJ} = \text{AND}_n \circ \text{NAND}^n$ this gadget is typically NAND .

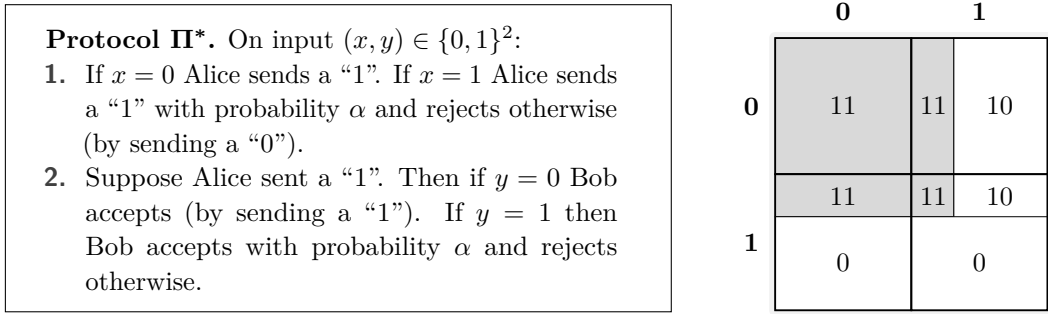
Our proof follows this outline. The reduction to the single-gadget case will be packaged into Theorem 8 and is standard. By contrast, in proving the $\Omega(1)$ information lower bound for the single gadget, we need to overcome the following two new technical issues.

(1) Small Acceptance Probabilities. Since the protocol is only required to succeed with a tiny probability $\alpha(n)$ on 1-inputs, the transcript of Π can be useless most of the time: Imagine a protocol that rejects with probability $1 - \alpha$ at the start (and otherwise does something useful). The entropy of the transcript of such protocols can be as low as $O(\alpha)$.

To address this issue, we do not work with the transcript distribution of Π directly, but rather with the *conditional distribution* given that the protocol accepts. That is, for 1-inputs (x, y) , we consider the random variable

$$T(x, y) := \Pi(x, y) | \Pi(x, y) \text{ is an accepting transcript}$$

and proceed to lower bound $\mathbb{I}(T(X, Y); X, Y)$ instead. One subtlety is that conditioning on acceptance does not “commute” with the reduction to the single-gadget case. We must



■ **Figure 1** Protocol Π^* for NAND. In the illustration on the right, each of the input blocks is further subdivided into rectangles according to the outcomes of the private coins. The rectangles are labeled with the associated transcripts.

consider the distribution of T that arises from first conditioning on acceptance and then doing the reduction, which is generally not the same distribution as if we did the reduction and then conditioned on acceptance. However, this is not a significant technical obstacle.

(2) Large Acceptance Probabilities. The acceptance probability of a protocol Π can vary between α and 1 when run on different 1-inputs. This, together with our conditioning idea above, introduces a new problem: there are USBP protocols for NAND such that the associated T leaks no information about the input!


Indeed, consider the protocol Π^* for NAND given in Figure 1. This protocol accepts the 1-input $(0, 0)$ with probability 1, the 1-inputs $(0, 1)$ and $(1, 0)$ with probability α , and the 0-input $(1, 1)$ with probability α^2 . Choosing α such that $\alpha^2 \leq \alpha/2$ we obtain a USBP protocol for NAND where the associated conditioned-on-acceptance variable T^* is constant (the protocol Π^* has only one accepting transcript, namely “11”).

To avoid this problem, we use a more complicated gadget than NAND; see Figure 2a. The new gadget G contains two instances of NAND: in Figure 2b one instance of NAND corresponds to the pair of edges AB and another one to AC. We show that the bad situation described above cannot happen simultaneously for both of them. One subtlety is that the bad situation—i.e., when a transcript has much higher probability of appearing on the 1-input $(0, 0)$ of NAND than on the other two 1-inputs $(0, 1)$ and $(1, 0)$ of NAND—depends on the transcript, with some transcripts being bad for AB and some being bad for AC, but none being bad for both. We prove an information lower bound (conditioned on acceptance) for whichever instance of NAND behaves better for “most” transcripts.

We note that a similar technical issue arose in the proof of Braverman and Moitra [6] when analyzing the case $\alpha = 1/2 + \epsilon$, $\beta = 1/2 - \epsilon$. Their solution involved applying a certain type of *random-self-reduction* (they called it *smoothing*) to the inputs before invoking the protocol. This approach is highly tailored to their setting and does not seem to be directly helpful to us. Nevertheless, our gadget G was inspired by their analysis.

3.1 Proof of Theorem 3

Define the gadget $G: \{0, 1, 2\} \times \{1, 2\} \rightarrow \{0, 1\}$ as the indicator for non-equality; see Figure 2a. Define the function $F: \{0, 1, 2\}^n \times \{1, 2\}^n \rightarrow \{0, 1\}$ by $F = \text{AND}_n \circ G^n$, i.e., $F(x, y) = 1$ iff $G(x_i, y_i) = 1$ for all $i \in [n]$. Since G reduces to DISJ on 2 bits by the map



■ **Figure 2** (a) Truth table of the gadget G . (b) Distributions Q_A , Q_B , and Q_C are uniform over the endpoints of the edges A , B , and C , respectively.

($0 \mapsto 00$, $1 \mapsto 01$, $2 \mapsto 10$), we find that F reduces to DISJ on $2n$ bits. Hence it suffices to prove that $\text{USBP}(F) \geq \Omega(n)$.

Input Distribution. Define Q_A, Q_B, Q_C to be the following three distributions over $\{0, 1, 2\} \times \{1, 2\}$: Q_A is uniform over $\{(0, 1), (0, 2)\}$, Q_B is uniform over $\{(0, 1), (2, 1)\}$, and Q_C is uniform over $\{(0, 2), (1, 2)\}$; see Figure 2b. For $i \in [n]$, $u \in \{0, 1, 2\}$, $v \in \{1, 2\}$, and z a length- $(n-1)$ string over the alphabet $\{A, B, C\}$ indexed by $[n] \setminus \{i\}$, define $D_{i,u,v,z}$ to be the distribution over pairs $(x, y) \in \{0, 1, 2\}^n \times \{1, 2\}^n$ obtained by setting $x_i = u$, $y_i = v$, and for each $j \neq i$ (independently) sampling (x_j, y_j) from Q_{z_j} . Note that $\text{support}(D_{i,u,v,z}) \subseteq F^{-1}(G(u, v))$ and that x and y are independent when sampled from $D_{i,u,v,z}$.

Reduction to the Single-gadget Case. Let Π be an $R_{\alpha, \alpha/4}^{\text{priv}}$ protocol for F (recall that by and-amplification we may assume $\beta = \alpha/4$ in the definition of USBP). Let $\Pi(x, y)$ denote the transcript of Π on input (x, y) . Thus $\Pi(x, y)$ is a random variable whose outcome depends on the private coins of the protocol. For $(x, y) \in F^{-1}(1)$ define $T(x, y)$ as the random variable whose distribution is that of $\Pi(x, y)$ conditioned on $\Pi(x, y)$ being an accepting transcript.

Suppose for contradiction that the transcripts have length less than $n/2400$. Using the direct sum methodology we will next find a coordinate i (and a string z) such that the protocol leaks very little information about the i -th input (conditioned on the data z). This is formalized in the following lemma whose proof we defer to subsection 3.2 as it is essentially identical to the corresponding argument in [4]. Below, $\|X - Y\|$ denotes the statistical distance between the distributions of the random variables X and Y .

► **Lemma 8.** *If $\gamma > 0$ is such that for all i and z either $\|T(D_{i,0,1,z}) - T(D_{i,0,2,z})\| \geq \gamma$ or $\|T(D_{i,0,1,z}) - T(D_{i,2,1,z})\| \geq \gamma$ or $\|T(D_{i,0,2,z}) - T(D_{i,1,2,z})\| \geq \gamma$, then some transcript has length at least $n\gamma^2/6$.*

Contrapositively, letting $\gamma = 1/20$, Theorem 8 implies that there exists an i and z (which we fix henceforth) such that $\|T(D_{i,0,1,z}) - T(D_{i,0,2,z})\|$, $\|T(D_{i,0,1,z}) - T(D_{i,2,1,z})\|$, and $\|T(D_{i,0,2,z}) - T(D_{i,1,2,z})\|$ are all less than γ .

The Single-gadget Case. Let (u, v) be an input to G and let τ be a transcript. We define

$$\begin{aligned} \pi_{uv}(\tau) &:= \Pr[\Pi(D_{i,u,v,z}) = \tau] \quad \text{for any } (u, v), \\ t_{uv}(\tau) &:= \Pr[T(D_{i,u,v,z}) = \tau] \quad \text{for } (u, v) \in G^{-1}(1). \end{aligned}$$

We henceforth adopt the convention that $0/0 = 0$. Let

$$S := \left\{ \text{accepting } \tau : \frac{\pi_{01}(\tau)}{t_{01}(\tau)} \leq \frac{\pi_{02}(\tau)}{t_{02}(\tau)} \right\}$$

and let $\bar{S} := \{\text{accepting } \tau\} \setminus S$. Since $\|T(D_{i,0,1,z}) - T(D_{i,0,2,z})\| < \gamma$, we have either $\Pr[T(D_{i,0,1,z}) \in S] \geq \frac{1-\gamma}{2}$ or $\Pr[T(D_{i,0,2,z}) \in \bar{S}] \geq \frac{1-\gamma}{2}$. Henceforth assume the former case; a completely analogous argument handles the latter case.

Note that $\pi_{22}(\tau) \cdot \pi_{01}(\tau) = \pi_{21}(\tau) \cdot \pi_{02}(\tau)$ by the basic rectangular structure of τ . Also note that if $G(u, v) = 1$ and τ is accepting, then the following both hold.

■ We have $\pi_{uv}(\tau) = 0$ iff $t_{uv}(\tau) = 0$, and hence $\pi_{uv}(\tau) = \frac{\pi_{uv}(\tau)}{t_{uv}(\tau)} \cdot t_{uv}(\tau)$.

■ Assuming $\pi_{uv}(\tau)$ and $t_{uv}(\tau)$ are nonzero, we have $\frac{\pi_{uv}(\tau)}{t_{uv}(\tau)} \geq \alpha$ by the correctness of Π .

For $(u, v) \in \{(2, 1), (0, 2)\}$ define $\gamma_{uv}(\tau) = |t_{uv}(\tau) - t_{01}(\tau)|$ and note that

$$\sum_{\text{accepting } \tau} \gamma_{uv}(\tau) = 2\|T(D_{i,u,v,z}) - T(D_{i,0,1,z})\| < 2\gamma. \quad (3)$$

By a case analysis, we have $t_{21}(\tau) \cdot t_{02}(\tau) \geq t_{01}(\tau)^2 - t_{01}(\tau)(\gamma_{21}(\tau) + \gamma_{02}(\tau))$. Recalling our convention that $0/0 = 0$, and considering the case $t_{01}(\tau) = 0$ separately, we find that in all cases

$$\frac{t_{21}(\tau) \cdot t_{02}(\tau)}{t_{01}(\tau)} \geq t_{01}(\tau) - \gamma_{21}(\tau) - \gamma_{02}(\tau). \quad (4)$$

Thus we have

$$\begin{aligned} \Pr[\Pi(D_{i,2,2,z}) \text{ accepts}] &= \sum_{\text{accepting } \tau} \pi_{22}(\tau) \\ &\geq \sum_{\tau \in S} \pi_{22}(\tau) \\ &\geq \sum_{\tau \in S} \frac{\pi_{21}(\tau) \cdot \pi_{02}(\tau)}{\pi_{01}(\tau)} \\ &= \sum_{\tau \in S} \frac{\pi_{21}(\tau)}{t_{21}(\tau)} \cdot \frac{\pi_{02}(\tau)}{t_{02}(\tau)} \cdot \frac{t_{21}(\tau) \cdot t_{02}(\tau)}{t_{01}(\tau)} \\ &\geq \sum_{\tau \in S} \alpha \cdot (t_{01}(\tau) - \gamma_{21}(\tau) - \gamma_{02}(\tau)) \\ &> \alpha \cdot \left(\frac{1-\gamma}{2} - 2\gamma - 2\gamma\right) \\ &> \alpha/4. \end{aligned}$$

To see that the fifth line follows from the fourth, consider each $\tau \in S$: If $t_{21}(\tau) \neq 0$ and $t_{02}(\tau) \neq 0$ then it follows by $\frac{\pi_{21}(\tau)}{t_{21}(\tau)} \geq \alpha$ and $\frac{\pi_{02}(\tau)}{t_{02}(\tau)} / \frac{\pi_{01}(\tau)}{t_{01}(\tau)} \geq 1$ (since $\tau \in S$) and (4). On the other hand, if $t_{21}(\tau) = 0$ or $t_{02}(\tau) = 0$, say, $t_{21}(\tau) = 0$, then it follows since the summand on the fourth line is 0, and $t_{01}(\tau) - \gamma_{21}(\tau) = 0$ so the summand on the fifth line is nonpositive. The sixth line follows from the fifth by $\sum_{\tau \in S} t_{01}(\tau) = \Pr[T(D_{i,0,1,z}) \in S] \geq \frac{1-\gamma}{2}$ and $\sum_{\tau \in S} \gamma_{21}(\tau) \leq \sum_{\text{accepting } \tau} \gamma_{21}(\tau)$ and (3), and similarly for γ_{02} .

We conclude that $\Pr[\Pi(x, y) \text{ accepts}] > \alpha/4$ for some $(x, y) \in \text{support}(D_{i,2,2,z}) \subseteq F^{-1}(0)$, contradicting the correctness of Π . This finishes the proof of Theorem 3.

3.2 Proof of Theorem 8

Define jointly distributed random variables $X = X_1 \cdots X_n \in \{0, 1, 2\}^n$, $Y = Y_1 \cdots Y_n \in \{1, 2\}^n$, and $Z = Z_1 \cdots Z_n \in \{A, B, C\}^n$ as follows: Z is uniform, and given a particular

choice of Z , for each $i \in [n]$ (independently) (X_i, Y_i) is sampled from Q_{Z_i} . Thus the marginal distribution of (X, Y) is that for each i (independently), (X_i, Y_i) has probability $1/3$ for each of $(0, 1)$, $(0, 2)$, and probability $1/6$ for each of $(2, 1)$, $(1, 2)$. Since the support of (X, Y) is in $F^{-1}(1)$, we may also view T as a random variable distributed jointly with (X, Y, Z) . Let Z_{-i} denote $Z_1 \cdots Z_{i-1} Z_{i+1} \cdots Z_n$.

For any $i \in [n]$, $z_i \in \{A, B, C\}$, and $z_{-i} \in \{A, B, C\}^{n-1}$ indexed by $[n] \setminus \{i\}$, we can view $(T(D_{i, Q_{z_i, z_{-i}}}), Q_{z_i})$ as a pair of jointly distributed random variables that is distributed identically to $(T, (X_i, Y_i) \mid Z_i = z_i, Z_{-i} = z_{-i})$. For all i and z_{-i} , by a standard lemma (see [4, Lemma 6.2 and Proposition 6.10]) we have $\mathbb{I}(T(D_{i, Q_{z_i, z_{-i}}}); Q_A) \geq \|T(D_{i, 0, 1, z_{-i}}) - T(D_{i, 0, 2, z_{-i}})\|^2/2$ and similarly for B and C. Therefore

$$\begin{aligned}
\text{Maximum length of transcript} &\geq \mathbb{H}(T \mid Z) \\
&\geq \mathbb{I}(T; X, Y \mid Z) \\
&\geq \sum_i \mathbb{I}(T; X_i, Y_i \mid Z) \\
&= \sum_i \mathbb{E}_{z_{-i}} \frac{1}{3} \sum_{z_i} \mathbb{I}(T; X_i, Y_i \mid Z_i = z_i, Z_{-i} = z_{-i}) \\
&= \sum_i \mathbb{E}_{z_{-i}} \frac{1}{3} \sum_{z_i} \mathbb{I}(T(D_{i, Q_{z_i, z_{-i}}}); Q_{z_i}) \\
&\geq \sum_i \mathbb{E}_{z_{-i}} \frac{1}{3} \cdot \frac{1}{2} \left(\|T(D_{i, 0, 1, z_{-i}}) - T(D_{i, 0, 2, z_{-i}})\|^2 \right. \\
&\quad \left. + \|T(D_{i, 0, 1, z_{-i}}) - T(D_{i, 2, 1, z_{-i}})\|^2 \right. \\
&\quad \left. + \|T(D_{i, 0, 2, z_{-i}}) - T(D_{i, 1, 2, z_{-i}})\|^2 \right) \\
&\geq \sum_i \mathbb{E}_{z_{-i}} (\gamma^2/6) \\
&= n\gamma^2/6.
\end{aligned}$$

where the third line follows by a standard direct sum property for conditional information cost [4].

4 The Complexity of Set-Disjointness

We now prove Theorem 4 and Theorem 5 using Theorem 3 and Theorem 2.

4.1 Lower Bounds

Private-coin Lower Bounds. Let Π be an $R_{\alpha, \beta}^{\text{priv}}$ protocol for DISJ. We prove that the cost of Π is both $\Omega(n \cdot (1 - \beta/\alpha))$ and $\Omega(\log n)$, as required for Theorem 4.

First, if we do and-amplification by iterating the protocol $\lceil 1/(1 - \beta/\alpha) \rceil$ times and accepting iff all runs accept, we get an $R_{\alpha', \alpha'/2}^{\text{priv}}$ protocol for DISJ with $\alpha' = \alpha^{\lceil 1/(1 - \beta/\alpha) \rceil}$ (since $(\beta/\alpha)^{\lceil 1/(1 - \beta/\alpha) \rceil} < 1/2$). By Theorem 3 the amplified protocol must use $\Omega(n)$ communication and hence Π must have used $\Omega(n \cdot (1 - \beta/\alpha))$ communication.

Second, Forster's result [16] that the UPP complexity of inner product is $\Omega(n)$ gives us the $\Omega(\log n)$ lower bound for Π . Indeed, the inner product function reduces to DISJ with exponential blow-up (see [35, Proposition 6.5]) and we may convert Π into an UPP protocol by shifting the acceptance threshold near $1/2$.

Public-coin Lower Bounds. Let Π be an $R_{\alpha, \beta}^{\text{pub}}$ protocol for DISJ. We consider the two parts of Theorem 5 separately.

For the first part, suppose $\log(1/\alpha) \leq C \cdot n \cdot (1 - \beta/\alpha)$ for a to-be-specified constant C . We proceed exactly as above: We first and-amplify Π into an $R_{\alpha', \alpha'/2}^{\text{priv}}$ protocol. The parameters satisfy $\log(1/\alpha') = \log(1/\alpha) \cdot \lceil 1/(1 - \beta/\alpha) \rceil \leq C \cdot n \cdot (1 - \beta/\alpha) \cdot \lceil 1/(1 - \beta/\alpha) \rceil \leq 2C \cdot n$. Hence

if C is a sufficiently small universal constant then the $\Omega(n)$ lower bound for the amplified protocol (provided now by Theorem 2) must be coming from the communication cost and not from the $\log(1/\alpha')$ term. We conclude that the original protocol Π must have used $\Omega(n \cdot (1 - \alpha/\beta))$ communication.

For the second part, we do not need any restriction on the parameters. We claim that since DISJ has a 2×2 identity submatrix, we cannot have $R_{\alpha, \beta}^{\text{pub}}(\text{DISJ}) \leq 1$.¹ Suppose for contradiction there is a 1-bit protocol and yet $\text{DISJ}(x, y) = \text{DISJ}(x', y') = 1$ and $\text{DISJ}(x, y') = \text{DISJ}(x', y) = 0$. Say r is the probability Alice declares the output and $1 - r$ is the probability Bob declares the output. Conditioned on Alice declaring the output let $p_x, p_{x'}$ be the acceptance probability for the x and x' rows, and conditioned on Bob declaring the output let $q_y, q_{y'}$ be the acceptance probability for the y and y' columns. Letting $\pi_{xy} = rp_x + (1 - r)q_y$ be the overall acceptance probability on input (x, y) , we have $\alpha - \beta \leq \pi_{xy} - \pi_{x'y} = r(p_x - p_{x'})$ and $\alpha - \beta \leq \pi_{x'y'} - \pi_{xy'} = r(p_{x'} - p_x)$, a contradiction.

4.2 Upper Bounds

Public-coin Protocols. We start with a simple $R_{1, \beta/\alpha}^{\text{pub}}$ protocol for DISJ of cost $\Theta(n \cdot (1 - \beta/\alpha))$.

Basic public-coin protocol Π .

1. Use public randomness to pick a uniformly random $S \subseteq [n]$ of size $\lceil n \cdot (1 - \beta/\alpha) \rceil$.
2. Alice sends the substring $x|_S$ to Bob.
3. Bob outputs $\text{DISJ}(x|_S, y|_S)$.

It is straightforward to check that Π is indeed an $R_{1, \beta/\alpha}^{\text{pub}}$ protocol. To obtain an $R_{\alpha, \beta}^{\text{pub}}$ protocol for the first part of Theorem 5 (without needing any restriction on the parameters), we can reject with probability $1 - \alpha$ at the beginning and otherwise run Π . To obtain a protocol of cost 2 for the second part of Theorem 5, we need to better exploit the power of public coins. If we modify Π so that additional public coins are used to guess $x|_S$, then Alice can just send one bit indicating whether the guess is correct, and Bob can send the output bit (rejecting if the guess was wrong). This yields an $R_{1/2^{|S|}, \beta/\alpha 2^{|S|}}^{\text{pub}}$ protocol which, by the restriction that $\alpha \leq 1/2^{|S|}$, can be adapted into an $R_{\alpha, \beta}^{\text{pub}}$ protocol by automatically rejecting with probability $1 - \alpha 2^{|S|}$.

In fact, the above protocols can be seen as special cases of the following general protocol, which interpolates between them. For simplicity of presentation, let us assume that $\log(1/\alpha)$ is an integer and $\log(1/\alpha) \leq |S|$. In step 2 of the basic protocol Π , Alice can expedite the sending of her message to Bob as follows: Alice and Bob interpret additional public coins as guessing the first $\log(1/\alpha)$ bits of Alice's message. Alice can use one bit of communication to indicate whether this guess is correct, and if so she can send the other $|S| - \log(1/\alpha)$ bits

¹ For the public-coin version of UPP, an equivalence actually holds. For all f , we have $\text{UPP}^{\text{pub}}(f) \leq 2$, and it is not difficult to show that the following are equivalent: (i) $\text{UPP}^{\text{pub}}(f) \leq 1$, (ii) there exist row and column values $p_x, q_y \in [0, 1]$ and $r \in [0, 1]$ such that $|rp_x + (1 - r)q_y - f(x, y)| < 1/2$, (iii) the rows and columns can be permuted so each row and each column is monotonically nondecreasing (0's then 1's), (iv) f does not contain as a submatrix the 2×2 identity (or its complement). To see that (iii) \Rightarrow (ii), take $r = 1/2$, and $p_x =$ fraction of 1's in the x row, and $q_y = (y - 1/2)/(\text{number of columns})$ where y is viewed as a positive integer.

of her message normally. The probability that the public guess is correct is $2^{-\log(1/\alpha)} = \alpha$. Thus, this new protocol ends up working in a familiar way: with probability $1 - \alpha$ the public guess fails (in which case we reject), but otherwise we are able to run Π successfully. This results in an $R_{\alpha, \beta}^{\text{pub}}$ protocol of cost $|S| - \log(1/\alpha) + 2$. Here the $+2$ comes from Alice indicating whether the public guess is correct and Bob sending the final answer.

Private-coin Protocols. By sparsification, we may assume the basic protocol Π uses only $O(\log n)$ bits of public randomness. Thus we have $R_{1, \beta/\alpha}^{\text{priv}}(\text{DISJ}) \leq O(n \cdot (1 - \beta/\alpha) + \log n)$ since Alice can pick S privately and send it to Bob along with $x|_S$. An $R_{\alpha, \beta}^{\text{priv}}$ protocol for Theorem 4 can be obtained as previously: automatically reject with probability $1 - \alpha$ and otherwise run the $R_{1, \beta/\alpha}^{\text{priv}}$ protocol.

5 Open Problems

It would be interesting to separate SBP and USBP, or to separate MA and SBP, even by a promise problem. In the classical world, it is known that MA and SBP can be separated by an oracle [5, 32]. The relationship between USBP and AM is also open (in both directions).

Among the complexity measures PP, UPP, SBP, and USBP, the first three are characterized by discrepancy [26], sign-rank [30], and corruption (Theorem 1), respectively. It is straightforward to show that the fourth is characterized by the log of the smallest nonnegative rank of a matrix such that the minimum value of a 1-input's entry is at least twice the maximum value of a 0-input's entry. It is open to provide a more natural characterization of USBP.

Acknowledgements. We thank Mark Braverman, Tom Gur, Toniann Pitassi, and anonymous reviewers for comments and discussions.

References

- 1 Scott Aaronson and Avi Wigderson. Algebrization: A new barrier in complexity theory. *ACM Transactions on Computation Theory*, 1(1), 2009.
- 2 László Babai, Peter Frankl, and Janos Simon. Complexity classes in communication complexity theory. In *Proceedings of the 27th Symposium on Foundations of Computer Science (FOCS)*, pages 337–347. IEEE, 1986.
- 3 László Babai and Shlomo Moran. Arthur–Merlin games: A randomized proof system, and a hierarchy of complexity classes. *Journal of Computer and System Sciences*, 36(2):254–276, 1988.
- 4 Ziv Bar-Yossef, T.S. Jayram, Ravi Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. *Journal of Computer and System Sciences*, 68(4):702–732, 2004.
- 5 Elmar Böhler, Christian Glaßer, and Daniel Meister. Error-bounded probabilistic computations between MA and AM. *Journal of Computer and System Sciences*, 72(6):1043–1076, 2006.
- 6 Mark Braverman and Ankur Moitra. An information complexity approach to extended formulations. In *Proceedings of the 45th Symposium on Theory of Computing (STOC)*, pages 161–170. ACM, 2013.
- 7 Mark Braverman and Omri Weinstein. A discrepancy lower bound for information complexity. In *Proceedings of the 16th International Workshop on Randomization and Computation (RANDOM)*, pages 459–470. Springer, 2012.

- 8 Harry Buhman, Nikolai Vereshchagin, and Ronald de Wolf. On computation and communication with small bias. In *Proceedings of the 22nd Conference on Computational Complexity (CCC)*, pages 24–32. IEEE, 2007.
- 9 Amit Chakrabarti, Graham Cormode, and Andrew McGregor. Annotations in data streams. In *Proceedings of the 36th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 222–234. Springer, 2009.
- 10 Amit Chakrabarti, Graham Cormode, Andrew McGregor, Justin Thaler, and Suresh Venkatasubramanian. On interactivity in Arthur–Merlin communication and stream computation. Technical Report TR13-180, Electronic Colloquium on Computational Complexity (ECCC), 2013.
- 11 Amit Chakrabarti and Oded Regev. An optimal lower bound on the communication complexity of Gap-Hamming-Distance. *SIAM Journal on Computing*, 41(5):1299–1317, 2012.
- 12 Amit Chakrabarti, Yaoyun Shi, Anthony Wirth, and Andrew Yao. Informational complexity and the direct sum problem for simultaneous message complexity. In *Proceedings of the 42nd Symposium on Foundations of Computer Science (FOCS)*, pages 270–278. IEEE, 2001.
- 13 Arkadev Chattopadhyay and Toniann Pitassi. The story of set disjointness. *SIGACT News*, 41(3):59–85, 2010.
- 14 Thomas Cover and Joy Thomas. *Elements of Information Theory*. Wiley, 2006.
- 15 Scott Diehl. Lower bounds for swapping Arthur and Merlin. In *Proceedings of the 11th International Workshop on Randomization and Computation (RANDOM)*, pages 449–463. Springer, 2007.
- 16 Jürgen Forster. A linear lower bound on the unbounded error probabilistic communication complexity. *Journal of Computer and System Sciences*, 65(4):612–625, 2002.
- 17 Anat Ganor, Gillat Kol, and Ran Raz. Exponential separation of information and communication. Technical Report TR14-049, Electronic Colloquium on Computational Complexity (ECCC), 2014.
- 18 Shafi Goldwasser and Michael Sipser. Private coins versus public coins in interactive proof systems. In *Proceedings of the 18th Symposium on Theory of Computing (STOC)*, pages 59–68. ACM, 1986.
- 19 Tom Gur and Ran Raz. Arthur–Merlin streaming complexity. In *Proceedings of the 40th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 528–539. Springer, 2013.
- 20 Rahul Jain and Hartmut Klauck. The partition bound for classical communication complexity and query complexity. In *Proceedings of the 25th Conference on Computational Complexity (CCC)*, pages 247–258. IEEE, 2010.
- 21 Rahul Jain, Troy Lee, and Nisheeth Vishnoi. A quadratically tight partition bound for classical communication complexity and query complexity. Technical report, arXiv, 2014.
- 22 Stasys Jukna. *Boolean Function Complexity: Advances and Frontiers*, volume 27 of *Algorithms and Combinatorics*. Springer, 2012.
- 23 Bala Kalyanasundaram and Georg Schnitger. The probabilistic communication complexity of set intersection. *SIAM Journal on Discrete Mathematics*, 5(4):545–557, 1992.
- 24 Iordanis Kerenidis, Sophie Laplante, Virginie Lerys, Jérémie Roland, and David Xiao. Lower bounds on information complexity via zero-communication protocols and applications. In *Proceedings of the 53rd Symposium on Foundations of Computer Science (FOCS)*, pages 500–509, 2012.
- 25 Hartmut Klauck. Rectangle size bounds and threshold covers in communication complexity. In *Proceedings of the 18th Conference on Computational Complexity (CCC)*, pages 118–134. IEEE, 2003.

- 26 Hartmut Klauck. Lower bounds for quantum communication complexity. *SIAM Journal on Computing*, 37(1):20–46, 2007.
- 27 Hartmut Klauck. On Arthur Merlin games in communication complexity. In *Proceedings of the 26th Conference on Computational Complexity (CCC)*, pages 189–199. IEEE, 2011.
- 28 Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, 1997.
- 29 Ilan Newman. Private vs. common random bits in communication complexity. *Information Processing Letters*, 39(2):67–71, 1991.
- 30 Ramamohan Paturi and Janos Simon. Probabilistic communication complexity. *Journal of Computer and System Sciences*, 33(1):106–123, 1986.
- 31 Alexander Razborov. On the distributional complexity of disjointness. *Theoretical Computer Science*, 106(2):385–390, 1992.
- 32 Miklos Santha. Relativized Arthur–Merlin versus Merlin–Arthur games. *Information and Computation*, 80(1):44–49, 1989.
- 33 Alexander Sherstov. Halfspace matrices. *Computational Complexity*, 17(2):149–178, 2008.
- 34 Alexander Sherstov. The pattern matrix method. *SIAM Journal on Computing*, 40(6):1969–2000, 2011.
- 35 Alexander Sherstov. The unbounded-error communication complexity of symmetric functions. *Combinatorica*, 31(5):583–614, 2011.
- 36 Alexander Sherstov. The communication complexity of Gap Hamming Distance. *Theory of Computing*, 8(1):197–208, 2012.
- 37 Thomas Vidick. A concentration inequality for the overlap of a vector on a large set, with application to the communication complexity of the Gap-Hamming-Distance problem. *Chicago Journal of Theoretical Computer Science*, 2012(1):1–12, 2012.
- 38 Thomas Watson. The complexity of estimating min-entropy. *Computational Complexity*, 2015. To appear. Preprint: <http://eccc.hpi-web.de/report/2012/070/>.