

COMMUNICATION COMPLEXITY OF DISJOINTNESS UNDER PRODUCT DISTRIBUTIONS

ZACH HUNTER¹, ALEKSA MILOJEVIĆ¹, BENNY SUDAKOV¹, AND ISTVAN TOMON²

ABSTRACT. Determining the randomized (or distributional) communication complexity of disjointness is a central problem in communication complexity, having roots in the foundational work of Babai, Frankl, and Simon in the 1980s and culminating in the famous works of Kalyanasundaram-Schnitger and Razborov in 1992. However, the question of obtaining tight bounds for product distributions persisted until the more recent work of Bottesch, Gavinsky, and Klauck resolved it. In this note we revisit this classical problem and give a short, streamlined proof of the best bounds, with improved quantitative dependence on the error parameter.

Our approach is based on a simple combinatorial lemma that may be of independent interest: if two sets drawn independently from two distributions are disjoint with non-negligible probability, then one can extract two subfamilies of reasonably large measure that are fully cross-disjoint (equivalently, a large monochromatic rectangle for disjointness).

1. INTRODUCTION

Communication complexity studies how much information two separated parties must exchange in order to compute a function of their joint input. Formally, fix finite sets $\mathcal{A}, \mathcal{B}$ and a Boolean function $f : \mathcal{A} \times \mathcal{B} \rightarrow \{0, 1\}$. Alice receives $A \in \mathcal{A}$ and Bob receives $B \in \mathcal{B}$, and neither party sees the other's input. They communicate by exchanging bits across a two-way channel according to a (deterministic) protocol, at the end of which one of the players outputs a value intended to equal $f(A, B)$. The *cost* of a protocol is the maximum, over all inputs (A, B) , of the total number of bits exchanged on input (A, B) , and the *deterministic communication complexity* of f , denoted by $\text{CC}(f)$, is the minimum cost of a protocol that is correct on every input. We refer to [7, 10] for background, motivation, and applications.

One of the most studied functions is *set disjointness*. Fix a ground set $[n]$ and let $\mathcal{A} = \mathcal{B} = 2^{[n]}$ and

$$\text{DISJ}_n(A, B) = 1 \iff A \cap B = \emptyset.$$

Disjointness is the canonical example of a hard function to compute in this model and it is a starting point for many reductions. Its deterministic communication complexity is $\Theta(n)$, and the lower bound follows from the rank method by noting that $\text{CC}(\text{DISJ}_n) \geq \log_2 \text{rank} M$, where $M_{A,B} = \text{DISJ}_n(A, B)$ is a $2^n \times 2^n$ matrix of full rank which encodes the function DISJ_n (see [11] for details).

In this paper we study disjointness in the *distributional* model. Fix a distribution μ on $\mathcal{A} \times \mathcal{B}$ and an error parameter $\varepsilon \in (0, 1/2)$. In this model, the input (A, B) is sampled according to μ , and $D_\varepsilon^\mu(f)$ denotes the minimum communication cost of a deterministic protocol that computes f correctly on a $(1 - \varepsilon)$ -fraction of pairs of inputs (with respect to μ). Note that the protocol may depend on μ .

By Yao's minimax principle, $\max_\mu D_\varepsilon^\mu(f)$ equals the public-coin randomized communication complexity of f with error ε . Therefore, determining the distributional complexity of disjointness was a problem of great interest, until it was famously resolved by Kalyanasundaram and Schnitger [6] and by Razborov [9] in the 1990s. Razborov proved that disjointness has randomized communication complexity $\Omega(n)$ by constructing a hard distribution μ for which $D_{1/3}^\mu(\text{DISJ}_n) = \Theta(n)$. Thus, in the fully general distributional/randomized setting disjointness remains linearly hard.

¹Department of Mathematics, ETH Zürich, Switzerland. Email: {zach.hunter, aleksa.milojevic, benjamin.sudakov}@math.ethz.ch. Research supported in part by SNSF grant 200021-228014.

²Umeå University, e-mail: istvantomon@gmail.com, Research supported in part by the Swedish Research Council grant VR 2023-03375.

It is natural to consider the restricted setting when $\mu = \mu_A \times \mu_B$ is a *product distribution*. This corresponds to sampling the sets $A \sim \mu_A, B \sim \mu_B$ independently. Following the classical work of Babai, Frankl, and Simon [1], we define the *strong distributional complexity* of a function f as

$$\text{SD}_\varepsilon(f) := \sup_{\mu = \mu_A \times \mu_B} D_\varepsilon^\mu(f).$$

Remarkably, restricting to independent inputs decreases the worst-case complexity of disjointness from linear to about $\sqrt{n}$. More precisely, Babai–Frankl–Simon [1] proved in 1986 that for every $\varepsilon < 1/2$,

$$\Omega(\sqrt{n}) \leq \text{SD}_\varepsilon(\text{DISJ}_n) \leq O_\varepsilon(\sqrt{n} \log_2 n).$$

The logarithmic gap between the upper and lower bound persisted for several decades, and was closed only in 2015 in the work of Bottesch, Gavinsky, and Klauck [2], who proved the tight bound $\text{SD}_\varepsilon(\text{DISJ}_n) = \Theta_\varepsilon(\sqrt{n})$.

They also considered this problem in a more general context, which interpolates between the product-distribution regime and the fully general regime by parameterizing distributions according to their mutual information $I(A : B)$. Recall, the *mutual information* between random inputs A and B sampled from the distribution μ with marginals μ_A, μ_B is

$$I(A : B) = \sum_{a \in A} \sum_{b \in B} \mu(a, b) \log_2 \frac{\mu(a, b)}{\mu_A(a) \mu_B(b)}.$$

More generally, they showed that if μ is a distribution with mutual information at most k , then $D_\varepsilon^\mu(\text{DISJ}_n) \leq O(\sqrt{n(k+1)}/\varepsilon^2)$. The proofs of these results rely crucially on the Håstad–Wigderson protocol [4] which tests for disjointness of ℓ -element sets inside an arbitrary ground set using only $O_\varepsilon(\ell)$ bits.

In this paper we revisit this classical problem and give a simple streamlined proof of the result of Bottesch, Gavinsky, and Klauck [2]. Additionally, this proof improves the dependence on the error parameter ε .

Theorem 1. *Let μ be a distribution on $2^{[n]} \times 2^{[n]}$ and let $(A, B) \sim \mu$ satisfy $I(A : B) \leq k$. Then*

$$D_\varepsilon^\mu(\text{DISJ}_n) \leq O\left(\sqrt{\frac{n(k+1)}{\varepsilon}}\right).$$

Interestingly, Bottesch, Gavinsky and Klauck provide a lower bound which shows that Theorem 1 is tight, at least when the mutual information is not very large. Namely, for every $1 \leq k \leq \varepsilon n$ they construct a measure μ with mutual information at most k , for which $D_\varepsilon^\mu(\text{DISJ}_n) \geq \Omega(\sqrt{n(k+1)}/\varepsilon)$, showing that the above bound is tight in this regime.

The special case when $k = 0$ corresponds to the case of product distributions, and there we obtain a very short and completely self-contained proof of the bounds for the strong distributional complexity of disjointness.

Theorem 2. *For any $\varepsilon < 1/2$,*

$$\text{SD}_\varepsilon(\text{DISJ}_n) \leq O\left(\sqrt{n \log_2(1/\varepsilon)}\right).$$

The protocol we construct is deterministic, and it is based on a combinatorial lemma, which we believe may be of independent interest. Roughly speaking, the lemma says if two independent sets $A \sim \mu_A$ and $B \sim \mu_B$ are disjoint with non-negligible probability, then one can extract two subfamilies $\mathcal{A}, \mathcal{B} \subseteq 2^{[n]}$ of reasonably large measure such that $A \cap B = \emptyset$ for every pair $(A, B) \in \mathcal{A} \times \mathcal{B}$ (equivalently, DISJ_n is identically 1 on a large combinatorial rectangle). Using this lemma and a variant of the Nisan–Wigderson reduction, we can use large monochromatic combinatorial rectangles to obtain an efficient protocol for disjointness (see Section 3 for details).

The proof of this lemma is self-contained and is inspired by ideas developed in our recent work on the Singer–Sudan conjecture concerning families with many disjoint pairs [3]. Namely, inspired

by finding monochromatic rectangles and the log-rank conjecture, Singer and Sudan [12] asked the following question: given families $\mathcal{A}, \mathcal{B} \subseteq 2^{[n]}$ with the property that at least $\varepsilon|\mathcal{A}||\mathcal{B}|$ pairs $(A, B) \in \mathcal{A} \times \mathcal{B}$ are disjoint, how large families $\mathcal{F} \subseteq \mathcal{A}, \mathcal{G} \subseteq \mathcal{B}$ can we find with the property that every pair of sets $(F, G) \in \mathcal{F} \times \mathcal{G}$ is disjoint? In [3], we showed one can take $|\mathcal{F}||\mathcal{G}| \geq 2^{-O(\sqrt{n \log_2 1/\varepsilon})}|\mathcal{A}||\mathcal{B}|$. Our main lemma is a generalization of this statement, where set families are replaced by distributions over sets.

Paper organization. In Section 2, we state and prove our main lemma. Then, in Section 3 we use it to prove Theorem 2. Finally, we use Theorem 2 in Section 4 to derive Theorem 1.

2. DISTRIBUTIONS WITH MANY DISJOINT SETS

Lemma 3. *Let μ_A, μ_B be probability distributions on $2^{[n]}$, and let $\varepsilon \in (2^{-n}, 1/2)$ be a parameter such that*

$$\mathbb{P}_{A \sim \mu_A, B \sim \mu_B}[A \cap B = \emptyset] \geq \varepsilon.$$

Then, there exist collections $\mathcal{A}, \mathcal{B} \subseteq 2^{[n]}$ such that $\mu_A(\mathcal{A}), \mu_B(\mathcal{B}) \geq 2^{-O(\sqrt{n \log_2 1/\varepsilon})}$, and such that all pairs $(A, B) \in \mathcal{A} \times \mathcal{B}$ are disjoint.

Proof. The main idea of the proof is to pick $\ell = \lfloor \sqrt{n/\log_2 \varepsilon^{-1}} \rfloor$ sets $A_1, \dots, A_\ell$ randomly according to μ_A and consider their union $U = \bigcup_{i=1}^\ell A_i$. Then set $\mathcal{A}$ to be the family of sets in $2^{[n]}$ contained in U and set $\mathcal{B}$ to be the family of sets in $2^{[n]}$ contained in $[n] \setminus U$. This construction ensures that for each pair of sets $(A, B) \in \mathcal{A} \times \mathcal{B}$, A and B are disjoint, and therefore our main goal is to show that with positive probability, both $\mu_A(\mathcal{A})$ and $\mu_B(\mathcal{B})$ are sufficiently large.

First, we show that

$$\mathbb{P}[\mu_A(\mathcal{A}) \geq 2^{-3n/\ell}] \geq 1 - 2^{-2n}.$$

Say that a set U is *bad* if $\mathbb{P}_{A \sim \mu_A}[A \subseteq U] = \mu_A(2^U) \leq 2^{-3n/\ell}$. Then the probability that ℓ randomly chosen sets from μ_A land inside a bad set U is at most $(2^{-3n/\ell})^\ell = 2^{-3n}$. Hence, taking the union bound over at most 2^n bad sets, the probability that $U = \bigcup_{i=1}^\ell A_i$ is bad can be bounded by 2^{-2n} . Therefore, with probability at least $1 - 2^{-2n}$, we have $\mu_A(\mathcal{A}) \geq 2^{-3n/\ell}$.

Next, we show that

$$\mathbb{P}[\mu_B(\mathcal{B}) \geq \varepsilon^\ell/2] > 2^{-n}.$$

To do so, we shall compute the expected measure of $\mathcal{B}$ by summing over all sets $B \subseteq [n]$ the probability that B lands in $\mathcal{B}$ times the weight of B under μ_B . By observing that

$$\mathbb{P}[B \in \mathcal{B}] = \mathbb{P}\left[B \cap \bigcup_{i=1}^\ell A_i = \emptyset\right] = \mathbb{P}[A_i \subseteq [n] \setminus B \text{ for all } i \in \{1, \dots, \ell\}] = \mu_A(2^{[n] \setminus B})^\ell,$$

we obtain

$$\mathbb{E}[\mu_B(\mathcal{B})] = \sum_{B \subseteq [n]} \mu_B(B) \cdot \mathbb{P}[B \in \mathcal{B}] = \sum_{B \subseteq [n]} \mu_B(B) \cdot \mu_A(2^{[n] \setminus B})^\ell.$$

By applying Jensen's inequality to the convex function $x \mapsto x^\ell$ and the measure μ_B we obtain

$$\mathbb{E}[\mu_B(\mathcal{B})] = \sum_{B \subseteq [n]} \mu_B(B) \cdot \mu_A(2^{[n] \setminus B})^\ell \geq \left(\sum_{B \subseteq [n]} \mu_B(B) \cdot \mu_A(2^{[n] \setminus B}) \right)^\ell \geq \varepsilon^\ell.$$

The last inequality follows from the fact $\sum_{B \subseteq [n]} \mu_B(B) \cdot \mu_A(2^{[n] \setminus B}) = \mathbb{P}_{A \sim \mu_A, B \sim \mu_B}[A \cap B = \emptyset] \geq \varepsilon$. Using that $0 \leq \mu_B(\mathcal{B}) \leq 1$, we can write

$$(1) \quad \varepsilon^\ell \leq \mathbb{E}[\mu_B(\mathcal{B})] \leq 1 \cdot \mathbb{P}\left[\mu_B(\mathcal{B}) \geq \varepsilon^\ell/2\right] + \frac{\varepsilon^\ell}{2} \cdot \mathbb{P}\left[\mu_B(\mathcal{B}) < \varepsilon^\ell/2\right] < \mathbb{P}\left[\mu_B(\mathcal{B}) \geq \varepsilon^\ell/2\right] + \frac{\varepsilon^\ell}{2}.$$

Note that $\varepsilon^\ell/2 \geq \varepsilon^{\sqrt{n/\log_2 \varepsilon^{-1}}}/2 = 2^{-\sqrt{n \log_2 \varepsilon^{-1}}-1} > 2^{-n-1}$, where the last inequality holds by our assumption $\varepsilon > 2^{-n}$. Therefore, comparing the two sides of Eq. 1 gives $\mathbb{P}[\mu_B(\mathcal{B}) \geq \varepsilon^\ell/2] > \varepsilon^\ell/2 \geq 2^{-n-1}$.

Since $\mathbb{P}[\mu_A(\mathcal{A}) \geq 2^{-3n/\ell}] + \mathbb{P}[\mu_B(\mathcal{B}) \geq \varepsilon^\ell/2] > 1$, there exists a choice of $A_1, \dots, A_\ell$ for which

$$\begin{aligned}\mu_A(\mathcal{A}) &\geq 2^{-3n/\ell} \geq 2^{-O(\sqrt{n \log_2 \varepsilon^{-1}})}, \\ \mu_B(\mathcal{B}) &\geq \varepsilon^\ell/2 = 2^{-\ell \log_2 \varepsilon^{-1} - 1} \geq 2^{-O(\sqrt{n \log_2 \varepsilon^{-1}})},\end{aligned}$$

where we have used that $\ell = \lfloor \sqrt{n/\log_2 \varepsilon^{-1}} \rfloor$. This completes the proof. $\square$

3. DISTRIBUTIONAL COMPLEXITY FOR PRODUCT DISTRIBUTIONS

In this section, we use a variant of the Nisan-Wigderson reduction [8] to prove Theorem 2, i.e. to prove that $D_\varepsilon^{\mu_A \times \mu_B}(\text{DISJ}_n) \leq O(\sqrt{n \log_2 1/\varepsilon})$. We say that a measure ν is *almost-empty* if $\mathbb{P}_{(A,B) \sim \nu}[A \cap B = \emptyset] \leq \varepsilon/2$. Also, we say that a combinatorial rectangle $R = \mathcal{F} \times \mathcal{G} \subseteq 2^{[n]} \times 2^{[n]}$ is *full* if all pairs $(F, G) \in \mathcal{F} \times \mathcal{G}$ are disjoint.

Proof of Theorem 2. We present an efficient protocol for computing DISJ_n correctly on a $(1 - \varepsilon)$ -fraction of the inputs under the given measure μ . Let $A, B \in 2^{[n]}$ be the inputs known to Alice and Bob.

We first define a subprotocol, with the following goal. Assume that there is a set $X \subset [n]$, known to both Alice and Bob, such that $A \setminus X$ and $B \setminus X$ are disjoint, or equivalently, $A \cap B = \emptyset \Leftrightarrow (A \cap X) \cap (B \cap X) = \emptyset$. At the end of the subprotocol, either Alice and Bob declare whether A and B are disjoint, or they will agree on a set Y of size $|Y| \leq |X|/2$ such that $A \setminus Y$ and $B \setminus Y$ are disjoint, effectively halving the size of the ground set. Additionally, if T is the set of pairs of inputs where the subprotocol outputs an answer whether $A \cap B = \emptyset$ and S is the set of pairs of inputs where the answer was incorrect, we will guarantee $\mu(S) \leq \frac{\varepsilon}{2}(\mu(T) + \frac{1}{|X|})$.

We claim that such a subprotocol can be executed with $O(\sqrt{|X| \log_2 1/\varepsilon})$ bits of communication. This claim implies our theorem, since we can start with $X = [n]$, iterate the subprotocol until the size of the ground set reduces to $|X| \leq 4 \log_2 1/\varepsilon$, halving the size of the ground set in every step. If up to this point Alice and Bob did not decide whether $A \cap B = \emptyset$, Alice can communicate all elements of $A \cap X$ using $|X| = O(\sqrt{|X| \log_2 1/\varepsilon})$ bits, thus solving the problem. The total number of bits communicated in this protocol is at most

$$\sum_{i=0}^{\lfloor \log_2 n \rfloor - 1} O\left(\sqrt{\frac{n}{2^i} \log_2 1/\varepsilon}\right) = O(\sqrt{n \log_2 1/\varepsilon}).$$

Moreover, the measure of the set S where the protocol makes an incorrect answer is at most

$$\frac{\varepsilon}{2} \left(1 + \sum_{i=0}^{\lfloor \log_2 n \rfloor - 1} \frac{1}{n/2^i} \right) \leq \frac{\varepsilon}{2} \cdot 2 = \varepsilon.$$

Subprotocol specification: Let M_0 be the disjointness matrix of the subsets of X , that is, for any $A_0, B_0 \subseteq X$ we have $(M_0)_{A_0, B_0} = 1 \Leftrightarrow A_0 \cap B_0 = \emptyset$. Also, let $A' = X \cap A$, $B' = X \cap B$, $\mathcal{A}_0 = \mathcal{B}_0 = 2^{X'}$, and let $\nu = \nu_A \times \nu_B$ be the rescaling of the measure μ on $2^{X'} \times 2^{X'}$ such that ν is a probability measure.

If $\nu_A(A') \leq \frac{\varepsilon}{2^{|X'|}}$, then Alice communicates this to Bob, or if $\nu_B(B') \leq \frac{\varepsilon}{2^{|X'|}}$, Bob communicates this to Alice using a single bit. If either of these two events happen, Alice and Bob may declare that the sets are, say, not disjoint. The measure of inputs where they make a mistake in this step is at most $\frac{\varepsilon}{2^{|X'| - 1}} \leq \frac{\varepsilon}{2^{|X'|}}$, since the total measure of pairs (A', B') with $\nu_A(A') \leq \frac{\varepsilon}{2^{|X'|}}$ or $\nu_B(B') \leq \frac{\varepsilon}{2^{|X'|}}$ is at most $2 \cdot 2^{|X'|} \cdot \varepsilon/2^{|X'|} = \varepsilon/2^{|X'| - 1}$.

If ν is almost-empty, then Alice and Bob declare that their sets are not disjoint and end the protocol. Then their answer is correct for at least $(1 - \varepsilon/2)$ measure of the inputs in $\mathcal{A}_0 \times \mathcal{B}_0$, and thus $\mu(S) \leq \frac{\varepsilon}{2}\mu(T) + \frac{\varepsilon}{2^{|X'|}}$ is satisfied.

If ν is not almost-empty, let $R = \mathcal{F} \times \mathcal{G}$ be the full rectangle in M_0 of largest measure. Since $\varepsilon/2 \geq 2^{-|X'|}$ (due to the stopping condition on X), by Lemma 3, we have $\nu(R) \geq 2^{-O(\sqrt{|X'| \log_2 1/\varepsilon})}$, which

also implies $\nu_A(\mathcal{F}), \nu_B(\mathcal{G}) \geq 2^{-O(\sqrt{|X|\log_2 1/\varepsilon})}$. Since all sets $F \in \mathcal{F}$ are disjoint from all sets $G \in \mathcal{G}$, we have that their unions $U = \bigcup_{F \in \mathcal{F}} F$ and $V = \bigcup_{G \in \mathcal{G}} G$ are also disjoint. Hence, $|U| + |V| \leq |X|$ and so one of U, V has size at most $|X|/2$, say U (otherwise, we may swap the roles of the players). If $A' \in \mathcal{F}$, then $A' \subseteq U$, so $A' \setminus U$ and $B' \setminus U$ are disjoint. In this case, Alice and Bob set $Y := U$ and stop the subprotocol. If $A' \notin \mathcal{F}$, then define $\mathcal{A}_1 := \mathcal{A}_0 \setminus \mathcal{F}$ and $\mathcal{B}_1 := \mathcal{B}_0$.

One can iterate the above procedure, thus defining a sequence of families of subsets $\mathcal{A}_0 \supseteq \mathcal{A}_1 \supseteq \dots \supseteq \mathcal{A}_t$ and $\mathcal{B}_0 \supseteq \mathcal{B}_1 \supseteq \dots \supseteq \mathcal{B}_t$. Observe that for each $1 \leq i \leq t$ we have

$$\nu(\mathcal{A}_i \times \mathcal{B}_i) \leq (1 - 2^{-O(\sqrt{|X|\log_2 1/\varepsilon})}) \nu(\mathcal{A}_{i-1} \times \mathcal{B}_{i-1}).$$

Therefore, if the subprotocol runs for t steps, we have

$$\nu(\mathcal{A}_t \times \mathcal{B}_t) \leq (1 - 2^{-O(\sqrt{|X|\log_2 1/\varepsilon})})^t \leq \exp(-t2^{-O(\sqrt{|X|\log_2 1/\varepsilon})}).$$

Setting $t := 2^C \sqrt{|X|\log_2 1/\varepsilon} (|X| + \log_2 1/\varepsilon) = 2^{O(\sqrt{|X|\log_2 1/\varepsilon})}$ for some large constant C , the right-hand-side is at most $2^{-4|X|-2\log_2 1/\varepsilon} \leq \frac{\varepsilon^2}{2^{4|X|}}$. But this is impossible as then $\nu(A' \times B') \leq \frac{\varepsilon^2}{2^{4|X|}}$, implying that either $\nu_A(A') \leq \frac{\varepsilon}{2^{2|X|}}$ or $\nu_B(B') \leq \frac{\varepsilon}{2^{2|X|}}$. Hence, this subprotocol runs for at most t steps, and it defines a protocol tree with $O(t)$ leaves.

It is well-known (see e.g. Chapter 2, Lemma 2.8 in [7]) that every protocol tree can be rebalanced to be a decision tree with depth logarithmic in the number of leaves. This implies that the subprotocol can be executed with $O(\log_2 t) = O(\sqrt{|X|\log_2 1/\varepsilon})$ bits of communication, verifying our claim. $\square$

4. DISTRIBUTIONS WITH MUTUAL INFORMATION

Next, we derive Theorem 1 from Theorem 2. Bottesch, Gavinsky and Klauck observed that the (classical) Substate theorem of Jain, Radhakrishnan and Sen [5] can be used to translate bounds from the independent setting to the setting of bounded mutual information. However, the main point here is that the bounds of Theorem 2 have a very good dependence on ε , which allows us to apply this reduction without big losses.

To state the Substate theorem, we recall the following two definitions. If μ, ν are two probability distributions on $\mathcal{A} \times \mathcal{B}$, $\|\mu - \nu\|_{\text{TV}}$ stands for the total variation distance between distributions μ and ν , which is defined as $\|\mu - \nu\|_{\text{TV}} = \max_{S \subseteq \mathcal{A} \times \mathcal{B}} |\mu(S) - \nu(S)|$. Further, if μ has marginals μ_A, μ_B , we define $I_\infty(\mu) = \max_{(a,b) \in \mathcal{A} \times \mathcal{B}} \log_2 \frac{\mu(a,b)}{\mu_A(a)\mu_B(b)}$.

Theorem 4 ([5]). *Let $\varepsilon > 0$ be a positive parameter, and let $\mathcal{A}, \mathcal{B}$ be finite sets. For every distribution μ on $\mathcal{A} \times \mathcal{B}$, with mutual information bounded by k , there exists another distribution ν , such that $\|\mu - \nu\|_{\text{TV}} \leq \varepsilon$ and $I_\infty(\nu) \leq 4(k+1)/\varepsilon$.*

Using this theorem, we give a proof of Theorem 1.

Proof of Theorem 1. Let us consider the measure ν such that $\|\mu - \nu\|_{\text{TV}} \leq \varepsilon/2$ and $I_\infty(\nu) \leq 8(k+1)/\varepsilon$, which exists due to Theorem 4. If the marginals of ν are ν_A and ν_B , then we consider the product distribution $\nu_A \times \nu_B$ and we run the protocol from the proof of Theorem 2 with the error parameter $\varepsilon' = \varepsilon 2^{-8(k+1)/\varepsilon-1}$ and distribution $\nu_A \times \nu_B$.

The cost of this protocol is $O(\sqrt{n \log_2 1/\varepsilon'}) = O(\sqrt{n(k+1)/\varepsilon})$. Moreover, if $S \subseteq 2^{[n]} \times 2^{[n]}$ is the set of pairs of inputs where it makes the mistake, we have that S has measure at most $(\nu_A \times \nu_B)(S) \leq \varepsilon'$. Since $I_\infty(\nu) \leq 8(k+1)/\varepsilon$, this means that for any pair $(A, B) \in \mathcal{A} \times \mathcal{B}$ we have $\nu(A, B) \leq 2^{8(k+1)/\varepsilon} \nu_A(A) \nu_B(B)$ and so $\nu(S) \leq 2^{8(k+1)/\varepsilon} \varepsilon' \leq \varepsilon/2$.

Finally, since $\|\mu - \nu\|_{\text{TV}} \leq \varepsilon/2$, we have $|\mu(S) - \nu(S)| \leq \varepsilon/2$, allowing us to conclude $\mu(S) \leq \varepsilon$. Hence, we have a protocol of complexity at most $O(\sqrt{n(k+1)/\varepsilon})$ making an error on at most an ε -fraction of inputs, completing the proof. $\square$

Acknowledgment. The authors would like to thank Avi Wigderson for valuable insights and for bringing the paper of Babai, Frankl, and Simon to their attention.

REFERENCES

- [1] L. Babai, P. Frankl, and J. Simon, *Complexity classes in communication complexity theory (preliminary version)*, in *Proceedings of the 27th Annual Symposium on Foundations of Computer Science (FOCS 1986)*, 1986, pp. 337–347.
- [2] R. C. Bottesch, D. Gavinsky, and H. Klauck, *Correlation in hard distributions in communication complexity*, in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2015)*, LIPIcs, vol. 40, 2015, pp. 544–572.
- [3] Z. Hunter, A. Milojević, B. Sudakov, and I. Tomon, *Disjoint pairs in set systems and combinatorics of low rank matrices*, arXiv:2411.13510.
- [4] J. Håstad and A. Wigderson, *The randomized communication complexity of set disjointness*, *Theory of Computing* **3** (2007), 211–219.
- [5] R. Jain, J. Radhakrishnan, and P. Sen, *Privacy and interaction in quantum communication complexity and a theorem about the relative entropy of quantum states*, in *Proceedings of the 43rd Annual Symposium on Foundations of Computer Science (FOCS 2002)*, 2002, pp. 429–438.
- [6] B. Kalyanasundaram and G. Schnitger, *The probabilistic communication complexity of set intersection*. *SIAM J. Discrete Math.* **5** (1992), 545–557.
- [7] E. Kushilevitz and N. Nisan, *Communication Complexity*. Cambridge University Press, 1997.
- [8] N. Nisan and A. Wigderson, *On rank vs. communication complexity*, *Proceedings of the 35th Annual Symposium on Foundations of Computer Science (FOCS 1994)* 1994, pp. 831–836.
- [9] A. A. Razborov, *On the distributional complexity of disjointness*, *Theoretical Computer Science* **106** (1992), no. 2, 385–390.
- [10] A. Rao and A. Yehudayoff, *Communication Complexity and Applications*, Cambridge University Press, 2020.
- [11] A. A. Sherstov, *Communication complexity theory: thirty-five years of set disjointness*, in *Mathematical Foundations of Computer Science 2014 (MFCS 2014), Part I*, Lecture Notes in Computer Science, vol. 8634, Springer, 2014, pp. 24–43.
- [12] N. Singer and M. Sudan, *Point-hyperplane incidence geometry and the log-rank conjecture*, *ACM Transactions on Computation Theory* **14** (2022).