Communication Complexity Theory: Thirty-Five Years of Set Disjointness

Alexander A. Sherstov*

University of California, Los Angeles
sherstov@cs.ucla.edu

Abstract. The set disjointness problem features k communicating parties and k subsets $S_1, S_2, \ldots, S_k \subseteq \{1, 2, \ldots, n\}$. No single party knows all k subsets, and the objective is to determine with minimal communication whether the k subsets have nonempty intersection. The important special case $k = 2$ corresponds to two parties trying to determine whether their respective sets intersect. The study of the set disjointness problem spans almost four decades and offers a unique perspective on the remarkable evolution of communication complexity theory. We discuss known results on the communication complexity of set disjointness in the deterministic, nondeterministic, randomized, unbounded-error, and multiparty models, emphasizing the variety of mathematical techniques involved.

Keywords: Set disjointness problem, communication complexity, communication lower bounds

1 Introduction

Communication complexity theory, initiated by Andrew Yao [52] thirty-five years ago, is a central branch of theoretical computer science. The theory studies the minimum amount of communication, measured in bits, required in order to compute functions whose arguments are distributed among several parties. In addition to the basic importance of studying communication as a bottleneck resource, the theory has found a vast number of applications to other research areas, including mechanism design, streaming algorithms, machine learning, data structures, pseudorandom generators, and chip layout. Communication complexity theory is an abundant source of fascinating research questions that can be easily explained to a high school graduate but require deep mathematics and decades of collective effort to resolve. Progress in this area over the years has been truly remarkable, both in the depth and volume of research results obtained and in the diversity of techniques invented to obtain them.

* Supported by a National Science Foundation CAREER award and an Alfred P. Sloan Foundation Research Fellowship.

Our survey focuses on a single communication problem, whose study began with the theory's inception in 1979 and actively continues to this day, with much left to discover. This problem is *set disjointness*. Its simplest version features two parties who are each given a subset of $\{1, 2, \ldots, n\}$ and asked to determine with minimal communication whether the two subsets intersect. One can interpret the problem as scheduling a meeting subject to the availability of the two parties—or rather, checking whether such a meeting can be scheduled. The study of set disjointness has had a significant impact on communication complexity theory and has in many ways shaped it. First and foremost, the difficulty of determining the communication requirements of set disjointness in all but the simplest models has fueled a rapid development of the field's techniques. Moreover, set disjointness has acquired special status in communication complexity theory in that it often arises as an extremal example or as a problem separating one communication model from another. In what follows, we survey some of the highlights of this story, from basic models such as nondeterminism to advanced formalisms such as unbounded-error and multiparty communication.

2 Deterministic communication

The simplest model of communication is the two-party deterministic model. Consider a function $f\colon X \times Y \to \{0,1\}$, where X and Y are finite sets. The model features two cooperating parties, traditionally called Alice and Bob. Alice receives an input $x \in X$, Bob receives an input $y \in Y$, and their objective is to compute $f(x,y)$. To this end, Alice and Bob communicate back and forth according to an agreed-upon protocol. The *cost* of a given communication protocol is the maximum number of bits exchanged on any input pair (x,y). The *deterministic communication complexity* of f, denoted $D(f)$, is the least cost of a communication protocol for f. In this formalism, the set disjointness problem corresponds to the function $\mathrm{DISJ}_n\colon \mathscr{P}(\{1,2,\ldots,n\}) \times \mathscr{P}(\{1,2,\ldots,n\}) \to \{0,1\}$ given by

$$\mathrm{DISJ}_n(A,B) = \begin{cases} 1 & \text{if } A \cap B = \varnothing, \\ 0 & \text{otherwise,} \end{cases}$$

where $\mathscr{P}$ refers as usual to the powerset operator.

We start by reviewing some fundamental notions, which are easiest to explain in the deterministic model and become increasingly important in more advanced models. A *combinatorial rectangle* on $X \times Y$ is any subset R of the form $R = A \times B$, where $A \subseteq X$ and $B \subseteq Y$. For brevity, we will refer to such subsets as simply *rectangles*. Given a communication problem $f\colon X \times Y \to \{0,1\}$, a rectangle R is called *f-monochromatic* if f is constant on R. Rectangles play a central role in the study of communication complexity due to the following fact [29], which shows among other things that an efficient deterministic protocol for a given function f partitions the domain into a disjoint union of a small number of f-monochromatic rectangles.

FACT 1. *Let $\Pi\colon X \times Y \to \{0,1\}$ be a deterministic communication protocol of cost at most c. Then there exist pairwise disjoint rectangles $R_1, R_2, R_3, \ldots, R_{2^c}$ such that*

$$\bigcup_{i=1}^{2^c} R_i = X \times Y$$

and Π is constant on each R_i.

Fooling set method. A straightforward technique for proving communication lower bounds is the *fooling set method* [29], which works by identifying a large set of inputs no two of which can occupy the same f-monochromatic rectangle. Formally, a *fooling set* for $f\colon X \times Y \to \{0,1\}$ is any subset $S \subseteq X \times Y$ with the following two properties: (i) f is constant on S; and (ii) if (x,y) and (x',y') are two distinct elements of S, then f is not constant on $\{(x,y),(x,y'),(x',y),(x',y')\}$. A moment's reflection reveals that an f-monochromatic rectangle can contain at most one element of S. Therefore, any partition (or even cover!) of $X \times Y$ by f-monochromatic rectangles must feature a rectangle for each point in the fooling set S, which in light of Fact 1 means that the deterministic communication complexity of f is at least $\log_2 |S|$. We summarize this discussion in the following theorem.

THEOREM 2 (Fooling set method). *Let $f\colon X \times Y \to \{0,1\}$ be a given communication problem. If S is a fooling set for f, then*

$$D(f) \geqslant \log_2 |S|.$$

The fooling set method works perfectly for the set disjointness problem. Indeed, the set $\{(A, \{1,2,\ldots,n\} \setminus A) : A \subseteq \{1,2,\ldots,n\}\}$ is easily seen to be a fooling set for DISJ_n, whence $D(\mathrm{DISJ}_n) \geqslant n$. A somewhat more careful accounting yields the tight bound $D(\mathrm{DISJ}_n) = n + 1$.

Rank bound. A more versatile technique for deterministic communication complexity was pioneered by Mehlhorn and Schmidt [33], who took an algebraic view of the question. These authors associated to every communication problem $f\colon X \times Y \to \{0,1\}$ its *characteristic matrix* $M_f = [f(x,y)]_{x \in X, y \in Y}$ and observed that any partition of $X \times Y$ into 2^c f-monochromatic rectangles gives an upper bound of 2^c on the rank of the characteristic matrix over the reals. In view of Fact 1, this gives the so-called *rank bound* on deterministic communication complexity.

THEOREM 3 (Mehlhorn and Schmidt). *For any $f\colon X \times Y \to \{0,1\}$,*

$$D(f) \geqslant \log_2(\operatorname{rk} M_f).$$

This method, too, works well for set disjointness. Indeed, the characteristic matrix of DISJ_n is the Kronecker product of n matrices

$$\begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} \otimes \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix} \otimes \cdots \otimes \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}$$

and therefore has full rank. As a result, $D(\mathrm{DISJ}_n) \geqslant n$ by the rank bound.

Log-rank conjecture. It is an instructive exercise [29] to prove that the rank bound subsumes the fooling set method, in that every communication lower bound obtained using the fooling set method can be rederived up to a constant factor using the rank bound. As a matter of fact, the *log-rank conjecture* due to Lovász and Saks [32] asserts that the rank bound is approximately tight for every function:

$$D(f) \leqslant (\log_2(\operatorname{rk} M_f))^{c_1} + c_2$$

for some universal constants $c_1, c_2 > 0$ and all f. This conjecture remains one of the most intriguing open questions in the area. An earlier, stronger version of the log-rank conjecture with $c_1 = 1$ has been disproved. One counterexample, due to Nisan and Wigderson [37], is a function $f\colon \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ with $D(f) = \Omega(n)$ but $\log_2(\operatorname{rk} M_f) = O(n^{0.631\cdots})$. As the reader might have guessed from the title of our survey, Nisan and Wigderson's construction crucially uses results [40] on the communication complexity of the set disjointness function!

3 Nondeterminism

Nondeterminism plays an important role in the study of communication, both as a natural model in its own right and as a useful intermediate formalism. In a nondeterministic protocol for a given function $f\colon X \times Y \to \{0,1\}$, Alice and Bob start by guessing a bit string, visible to them both. From then on, they communicate deterministically. A nondeterministic protocol for f must output the correct answer for *at least one* guess string when $f(x,y) = 1$ and for *all* guess strings when $f(x,y) = 0$. The cost of a nondeterministic protocol is defined as the worst-case length of the guess string, plus the worst-case cost of the deterministic phase. The *nondeterministic communication complexity* of f, denoted $N(f)$, is the least cost of a nondeterministic protocol for f. As usual, the *co-nondeterministic communication complexity* of f is the quantity $N(\neg f)$.

The nondeterministic communication complexity of a given function f is essentially characterized by the *cover number* of f, which is the smallest number of f-monochromatic rectangles whose union is $f^{-1}(1)$. Indeed, it follows easily from Fact 1 that any nondeterministic protocol of cost c gives rise to such a collection of size at most 2^c. Conversely, any size-2^c collection of f-monochromatic rectangles whose union is $f^{-1}(1)$ gives rise to a nondeterministic protocol for f of cost $c+2$, in which Alice and Bob guess one of the rectangles and check with two bits of deterministic communication whether it contains their input pair.

Fooling set method. The fooling set method, reviewed in the previous section, generalizes to the nondeterministic model. Indeed, as we have already observed, no two points of a fooling set $S \subseteq f^{-1}(1)$ can reside in the same f-monochromatic rectangle, which means that any cover of $f^{-1}(1)$ by f-monochromatic rectangles must contain at least $|S|$ rectangles. In the language of nondeterministic communication complexity, we arrive at the following statement.

THEOREM 4 (Fooling set method). *Let $f\colon X \times Y \to \{0,1\}$ be a given communication problem. If $S \subseteq f^{-1}(1)$ is a fooling set for f, then*

$$N(f) \geqslant \log_2 |S|.$$

Since the set disjointness function has a fooling set $S \subseteq \mathrm{DISJ}_n^{-1}(1)$ of size 2^n, we obtain $N(\mathrm{DISJ}_n) \geqslant n$. Set disjointness should be contrasted in this regard with its complement $\neg\mathrm{DISJ}_n$, known as *set intersection*, whose nondeterministic communication complexity is a mere $\log_2 n + O(1)$. Indeed, Alice and Bob need only guess an element $i \in \{1, 2, \ldots, n\}$ and verify with two bits of communication that it belongs to their respective sets.

Rectangle size bound. The most powerful method for lower bounds on nondeterministic communication complexity is the following beautiful technique, known as the *rectangle size bound* [29].

THEOREM 5 (Rectangle size bound). *Let $f\colon X \times Y \to \{0,1\}$ be a given function. Then for every probability distribution μ on $f^{-1}(1)$,*

$$N(f) \geqslant \log_2 \left(\frac{1}{\max_R \mu(R)} \right),$$

where the maximum is over all rectangles $R \subseteq f^{-1}(1)$.

The rectangle size bound is a generalization of the fooling set method. Indeed, letting μ be the uniform distribution over a given fooling set $S \subseteq f^{-1}(1)$, we immediately recover Theorem 4. The proof of Theorem 5 is straightforward: any cover of $f^{-1}(1)$ by f-monochromatic rectangles must cover a set of μ-measure 1, which means that the total number of rectangles in the cover must be no less than the reciprocal of the largest μ-measure of a rectangle $R \subseteq f^{-1}(1)$. Theorem 5 is of interest in two ways. First of all, it characterizes nondeterministic communication complexity up to a small additive term [29]. Second, as we will see in the next section, ideas analogous to the rectangle size bound play a key role in the study of randomized communication complexity.

It is instructive to rederive the nondeterministic lower bound for set disjointness using the rectangle size bound. One approach is to simply consider the uniform distribution over the fooling set $\{(A, \{1, 2, \ldots, n\} \setminus A) : A \subseteq \{1, 2, \ldots, n\}\}$, which gives $N(\mathrm{DISJ}_n) \geqslant n$. A more revealing choice [29] is to let μ be the uniform distribution over $\mathrm{DISJ}_n^{-1}(1)$, so that

$$\mu(A, B) = \begin{cases} 3^{-n} & \text{if } A \cap B = \varnothing, \\ 0 & \text{otherwise.} \end{cases}$$

Now, let $R = \mathscr{A} \times \mathscr{B}$ be any rectangle in $\mathrm{DISJ}_n^{-1}(1)$. Then it is straightforward to check that the larger rectangle $\mathscr{P}(S) \times \mathscr{P}(T)$, where $S = \bigcup_{A \in \mathscr{A}} A$ and $T = \bigcup_{B \in \mathscr{B}} B$, must also be contained in $\mathrm{DISJ}_n^{-1}(1)$. It follows that $S \cap T = \varnothing$ and therefore $|R| \leqslant |\mathscr{P}(S) \times \mathscr{P}(T)| \leqslant 2^n$. In summary, we have shown that every rectangle $R \subseteq \mathrm{DISJ}_n^{-1}(1)$ contains at most 2^n inputs, whence $\mu(R) \leqslant (2/3)^n$. Applying the rectangle size bound, we arrive at $N(\mathrm{DISJ}_n) \geqslant n \log_2(3/2)$. While this bound is weaker than the previous bound $N(\mathrm{DISJ}_n) \geqslant n$, the analysis just given is more broadly applicable and is good preparation for the next section on randomized communication.

Complexity classes. In a seminal paper, Babai, Frankl, and Simon [5] initiated a systematic study of communication from the standpoint of complexity classes. Analogous to computational complexity, the focus here is on the asymptotic communication requirements of a *family* of functions, one function for each input size. Specifically, one considers families $\{f_n\}_{n=1}^{\infty}$ where $f_n \colon \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$. Among the complexity classes defined in [5] are $\mathsf{P}^{\mathsf{cc}}, \mathsf{NP}^{\mathsf{cc}}$, and $\mathsf{coNP}^{\mathsf{cc}}$, corresponding to function families with efficient deterministic, nondeterministic, and co-nondeterministic protocols, respectively. Formally, P^{cc} is the class of all families $\{f_n\}_{n=1}^{\infty}$ for which $D(f_n) \leqslant \log^c n + c$ for some constant $c > 0$ and all n. The classes $\mathsf{NP}^{\mathsf{cc}}$ and $\mathsf{coNP}^{\mathsf{cc}}$ are defined analogously with respect to the requirements $N(f_n) \leqslant \log^c n + c$ and $N(\neg f_n) \leqslant \log^c n + c$. Set disjointness is helpful in characterizing the relations among these classes. Indeed, recall that $D(\mathrm{DISJ}_n) \geqslant N(\mathrm{DISJ}_n) \geqslant n$ and $N(\neg\mathrm{DISJ}_n) \leqslant \log_2 n + O(1)$. An immediate consequence is that $\mathsf{P}^{\mathsf{cc}} \subsetneq \mathsf{NP}^{\mathsf{cc}}$, with an exponential gap between deterministic and nondeterministic complexity achieved for $\neg\mathrm{DISJ}_n$. One analogously obtains $\mathsf{NP}^{\mathsf{cc}} \neq \mathsf{coNP}^{\mathsf{cc}}$, with an exponential gap for DISJ_n. The significance of set disjointness in the study of nondeterminism is no accident: Babai et al. show that it is a complete problem for the class $\mathsf{coNP}^{\mathsf{cc}}$.

We conclude with yet another use of set disjointness. A fundamental result due to Aho, Ullman, and Yannakakis [2] states that $D(f) \leqslant cN(f)N(\neg f)$ for some absolute constant $c > 0$ and every function f. In particular, one obtains the surprising equality $\mathsf{P}^{\mathsf{cc}} = \mathsf{NP}^{\mathsf{cc}} \cap \mathsf{coNP}^{\mathsf{cc}}$. A variant of the set disjointness problem, known as *k-set disjointness* [39], shows that the upper bound of Aho et al. is tight up to a constant factor.

4 Randomized communication

In many ways, the randomized model is the most realistic abstraction of two-party communication. As usual, consider a function $f \colon X \times Y \to \{0,1\}$, where X and Y are finite sets. Alice receives an input $x \in X$, Bob receives an input $y \in Y$, and their objective is to compute $f(x,y)$ by communicating back and forth according to an agreed-upon protocol. In addition, Alice and Bob share an unlimited supply of uniformly random bits, which they can use in deciding what messages to send. The *cost* of a randomized protocol is the maximum number of bits exchanged on any input pair (x,y). Since the random bits are shared,

they do not count toward the protocol cost. A protocol is said to *compute f with error ϵ* if on every input pair (x, y), the output of the protocol is correct with probability at least $1 - \epsilon$. The *ϵ-error randomized communication complexity* of f, denoted $R_\epsilon(f)$, is the least cost of a randomized protocol that computes f with error ϵ. The canonical quantity to study is $R_{1/3}(f)$. This setting of the error parameter is without loss of generality. Indeed, for any constants $\epsilon, \epsilon' \in (0, 1/2)$, the error of a communication protocol can be reduced from ϵ to ϵ' by running the protocol constantly many times and outputting the majority answer.

There are several other ways to formalize randomized communication, all of which turn out to be equivalent [29]. Most notably, one can consider a model where Alice and Bob each have a private source of random bits, known as the *private-coin model*. A fundamental theorem due to Newman [34] shows that whether the random bits are shared or private affects the communication complexity of any given function $f\colon \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ by at most an additive term of $O(\log n)$.

Corruption bound. The randomized communication complexity of a function can be vastly smaller than its deterministic or nondeterministic complexity. For example, the problem of checking two n-bit strings for equality has randomized communication complexity $O(1)$, in contrast to its $\Omega(n)$ complexity in the deterministic and nondeterministic models. The fact that randomized protocols can be so powerful means that proving lower bounds in this model is correspondingly more difficult. The most common method for lower bounds on randomized communication complexity the *corruption bound* due to Yao [53]. As we shall soon see, this technique is strong enough to yield the celebrated $\Omega(n)$ lower bound for set disjointness.

THEOREM 6 (Corruption bound). *Let $f\colon X \times Y \to \{0,1\}$ be a given function, $\alpha, \beta > 0$ given parameters. Let μ be a probability distribution on $X \times Y$ such that every rectangle R obeys*

$$\mu(R \cap f^{-1}(0)) \geqslant \alpha\mu(R) - \beta.$$

Then for all $\epsilon > 0$,

$$R_\epsilon(f) \geqslant \log_2 \left(\frac{\alpha\mu(f^{-1}(1)) - \epsilon}{\beta} \right).$$

The technical details of this theorem are somewhat tedious, but the intuition is entirely straightforward. Fix a probability distribution μ on the domain of the given communication problem f. The hypothesis of the theorem states that with respect to μ, the "0" entries make up at least an α fraction of any rectangle—except for particularly small rectangles, with measure on the order of β. As a result, any cover of $f^{-1}(1)$ by rectangles that are "almost" f-monochromatic requires roughly $\mu(f^{-1}(1))/\beta$ rectangles, for a communication cost of roughly $\log_2(\mu(f^{-1}(1))/\beta)$. It is not too difficult to turn this informal discussion into a rigorous proof of the corruption bound, by using Fact 1 and viewing a randomized

protocol of a given cost as a probability distribution on deterministic protocols of the same cost.

An $\Omega(\sqrt{n})$ lower bound. The randomized communication complexity of set disjointness has been extensively studied. A variety of proof techniques have been brought to bear on this question, including Kolmogorov complexity, information theory, matrix analysis, and approximation theory. The first strong result in this line of work is an $\Omega(\sqrt{n})$ lower bound on the randomized communication complexity of DISJ_n, due to Babai, Frankl, and Simon [5]. Their proof, presented below in its entirety, uses nothing but basic combinatorics and is exceedingly elegant.

THEOREM 7 (Babai, Frankl, and Simon). $R_{1/3}(\mathrm{DISJ}_n) \geqslant \Omega(\sqrt{n})$.

Proof. Without loss of generality, we may assume that n is a perfect square divisible by 12. We will work with a restriction of the set disjointness problem, in which Alice and Bob's inputs are sets of size exactly $\sqrt{n}$. Let μ denote the uniform probability distribution over all such inputs. Then

$$\mu(\mathrm{DISJ}_n^{-1}(1)) = \frac{\binom{n-\sqrt{n}}{\sqrt{n}}}{\binom{n}{\sqrt{n}}} = \Omega(1).$$

The crux of the proof is the following purely combinatorial fact:

CLAIM. *Let $\mathscr{R} = \mathscr{A} \times \mathscr{B}$ be any rectangle with $\mathbf{P}_{(A,B)\in\mathscr{R}}[A \cap B = \varnothing] \geqslant 1 - \alpha$ and $|\mathscr{A}| \geqslant 2^{-\delta\sqrt{n}}\binom{n}{\sqrt{n}}$, where $\alpha > 0$ and $\delta > 0$ are sufficiently small absolute constants. Then*

$$|\mathscr{B}| \leqslant 2^{-\delta\sqrt{n}}\binom{n}{\sqrt{n}}.$$

Let us finish the proof of the theorem before moving on to the claim itself. The claim is logically equivalent to the following statement: there exist absolute constants $\alpha > 0$ and $\delta > 0$ such that any rectangle $\mathscr{R}$ with $\mu(\mathscr{R}) \geqslant 2^{-\delta\sqrt{n}}$ satisfies

$$\mu(\mathscr{R} \cap \mathrm{DISJ}_n^{-1}(0)) > \alpha\mu(\mathscr{R}).$$

Applying the corruption bound (Theorem 6) with $\beta = 2^{-\delta\sqrt{n}}$, we conclude that $R_\epsilon(\mathrm{DISJ}_n) = \Omega(\sqrt{n})$ for sufficiently small $\epsilon = \epsilon(\alpha, \delta) > 0$. By error reduction, this implies the conclusion of the theorem. □

Proof of Claim. Consider the matrix $M = [\mathrm{DISJ}_n(A, B)]_{A\in\mathscr{A}, B\in\mathscr{B}}$. By hypothesis, the "0" entries make up at most an α fraction of M. Without loss of generality, we may assume that the fraction of "0" entries is at most 2α in every row of M (if not, simply remove the offending rows, which reduces the size of $\mathscr{A}$ by at most a factor of 2). Now, abbreviate $k = \sqrt{n}/3$ and inductively find sets $A_1, A_2, \ldots, A_k \in \mathscr{A}$ that are *well separated*, in the sense that for all i,

$$|A_i \setminus (A_1 \cup A_2 \cup \cdots \cup A_{i-1})| \geqslant \frac{\sqrt{n}}{2}.$$

That such sets must exist is a straightforward exercise in counting, with $\alpha > 0$ small enough.

Recall that the "0" entries make up at most a 2α fraction of the entries in $[\mathrm{DISJ}_n(A_i, B)]_{i=1,2,\ldots,k;\ B\in\mathscr{B}}$. In particular, at least half of the sets $B \in \mathscr{B}$ must satisfy

$$\mathop{\mathbf{P}}_{i=1,2,\ldots,k}[A_i \cap B \neq \varnothing] \leqslant 4\alpha.$$

But by the well-separated property of $A_1, A_2, \ldots, A_k$, the number of such sets B is at most $\binom{k}{4\alpha k}\binom{n-(1-4\alpha)k\sqrt{n}/2}{\sqrt{n}}$. (Verify this!) For $\alpha > 0$ and $\delta > 0$ small enough, this estimate does not exceed $\frac{1}{2} \cdot 2^{-\delta\sqrt{n}}\binom{n}{\sqrt{n}}$, which gives the claimed upper bound on $|\mathscr{B}|$. □

This lower bound on the randomized communication complexity of set disjointness has an important implication for communication complexity classes. Analogous to $\mathsf{P}^{\mathsf{cc}}, \mathsf{NP}^{\mathsf{cc}}, \mathsf{coNP}^{\mathsf{cc}}$, Babai et al. [5] defined $\mathsf{BPP}^{\mathsf{cc}}$ as the class of all communication problems $\{f_n\}_{n=1}^{\infty}$ for which $R_{1/3}(f_n) \leqslant \log^c n + c$ for some constant $c > 0$ and all n. Theorem 7 shows that $\{\mathrm{DISJ}_n\}_{n=1}^{\infty} \notin \mathsf{BPP}^{\mathsf{cc}}$, thus separating the classes $\mathsf{NP}^{\mathsf{cc}}$ and $\mathsf{coNP}^{\mathsf{cc}}$ from $\mathsf{BPP}^{\mathsf{cc}}$.

Tight lower bound. The problem of determining the randomized communication complexity of set disjointness remained open for several years after the work of Babai et al. It was finally resolved by Kalyanasundaram and Schnitger [24], who used Kolmogorov complexity to obtain the tight lower bound $R_{1/3}(\mathrm{DISJ}_n) = \Omega(n)$. Shortly thereafter, Razborov [40] gave a celebrated alternate proof of the linear lower bound for set disjointness. In fact, Razborov considered an easier communication problem known as *unique set disjointness*, in which Alice and Bob's input sets $A, B \subseteq \{1, 2, \ldots, n\}$ are either disjoint or intersect in a unique element. He studied the probability distribution μ that places weight $3/4$ on disjoint pairs (A, B) of cardinality $|A| = |B| = \lfloor n/4 \rfloor$, and weight $1/4$ on uniquely intersecting pairs again of cardinality $|A| = |B| = \lfloor n/4 \rfloor$; in both cases, each such pair is equally likely. He proved that $\mu(\mathscr{R} \cap \mathrm{DISJ}^{-1}(0)) \geqslant \alpha\mu(\mathscr{R}) - 2^{-\delta n}$ for some constants $\alpha > 0$ and $\delta > 0$ and every combinatorial rectangle $\mathscr{R}$, from which the tight lower bound $R_{1/3}(\mathrm{DISJ}_n) = \Omega(n)$ follows immediately by Theorem 6. Razborov's analysis is based on an entropy argument along with an ingenious use of conditioning.

Razborov's result as well as his proof inspired much follow-up work. The fact that the lower bound holds even for unique set disjointness was a crucial ingredient in Nisan and Wigderson's counterexample to the "strong" log-rank conjecture (see Section 3). The linear lower bound on the randomized communication complexity of set disjointness has found several surprising applications, including streaming algorithms [4] and combinatorial auctions [35]. In a testament to the mathematical richness of this problem, Bar-Yossef et al. [7] discovered a simpler yet, information-theoretic proof of the linear lower bound. This line of work is still active, with a recent paper by Braverman et al. [13] determining the randomized communication complexity of set disjointness up to lower-order terms.

5 Unbounded-error communication

The *unbounded-error model*, due to Paturi and Simon [38], is a fascinating model of communication with applications to matrix analysis, circuit complexity, and learning theory [38, 3, 11, 19, 20, 28, 31, 43, 46, 42]. Let $f\colon X \times Y \to \{0,1\}$ be a communication problem of interest. As usual, Alice and Bob receive inputs $x \in X$ and $y \in Y$, respectively, and their objective is to compute $f(x,y)$ with minimal communication. They each have an unlimited private source of random bits. Their protocol is said to compute f in the unbounded-error model if on every input (x,y), the output is correct with probability strictly greater than $1/2$. The *unbounded-error communication complexity* of f, denoted $U(f)$, is the least cost of a protocol that computes f.

Observe that the unbounded-error model is the same as the private-coin randomized model discussed in Section 4, with one exception: in the latter case the protocol must produce the correct answer with probability at least $2/3$, whereas in the former case the probability of correctness merely needs to exceed $1/2$, by an arbitrarily small amount. This difference has far-reaching implications. For example, the fact that the parties in the unbounded-error model do not have a *shared* source of random bits is crucial: it is a good exercise to check that allowing shared randomness in the unbounded-error model would make the complexity of every function a constant. This contrasts with the randomized model, where making the randomness public has almost no effect on the complexity of any given function.

There are several reasons why the unbounded-error model occupies a special place in communication complexity theory. To start with, it is vastly more powerful than the deterministic, nondeterministic, randomized, and *quantum* models [42]. Another compelling reason is that unbounded-error communication complexity is closely related to the fundamental matrix-theoretic notion of *sign-rank*, which is defined for a Boolean matrix $M = [M_{ij}]$ as the minimum rank of a real matrix $R = [R_{ij}]$ such that $\operatorname{sgn} R_{ij} = (-1)^{M_{ij}}$ for all i, j. In other words, the sign-rank of a Boolean matrix M is the minimum rank of real matrix R that sign-represents it, with negative and positive entries in R corresponding to the true and false entries in M, respectively. We let $\operatorname{rk}_{\pm} M$ denote the sign-rank of M. Paturi and Simon [38] proved the following beautiful theorem, which shows that unbounded-error communication and sign-rank are equivalent notions.

Fact 8 (Paturi and Simon). *For some absolute constant c and every function $f\colon X \times Y \to \{0,1\}$,*

$$U(f) - c \leqslant \log_2(\operatorname{rk}_{\pm} M_f) \leqslant U(f) + c.$$

Proving lower bounds on sign-rank is difficult. Indeed, obtaining a strong lower bound on the unbounded-error communication complexity of any explicit function was a longstanding problem until the breakthrough work of Forster [19] several years ago. Fortunately, the unbounded-error complexity of set disjointness is easy to analyze. The following two theorems give a complete answer, up to an additive constant.

THEOREM 9 (Folklore). $U(\mathrm{DISJ}_n) \leqslant \log_2 n + O(1)$.

Proof. Consider the following randomized protocol, where A and B denote Alice and Bob's input sets, respectively. The players pick an index $i \in \{1, 2, \ldots, n\}$ uniformly at random and verify with two bits of communication whether $i \in A \cap B$. If so, they output 0. In the complementary case $i \notin A \cap B$, they output 1 with probability $n/(2n-1)$ and 0 otherwise. It is easy to verify that this protocol is correct with probability at least $n/(2n-1) > 1/2$ on every input. Moreover, it clearly has cost at most $\log_2 n + O(1)$ in the private-coin model. □

THEOREM 10 (Paturi and Simon). $U(\mathrm{DISJ}_n) \geqslant \log_2 n - O(1)$.

Proof. We will give a linear-algebraic proof of this result, as opposed to the geometric argument of Paturi and Simon [38]. By Fact 8, it suffices to show that the characteristic matrix of set disjointness has sign-rank at least n. We will actually prove the claim for the submatrix $M = [x_i]_{x \in \{0,1\}^n, i=1,2,\ldots,n}$, whose rows are the 2^n distinct Boolean vectors of length n.

For the sake of contradiction, assume that $\mathrm{rk}_{\pm} M \leqslant n-1$. Then there are vectors $u_1, u_2, \ldots, u_{n-1} \in \mathbb{R}^n$ such that every $\sigma \in \{-1, +1\}^n$ is the (componentwise) sign of some linear combination of $u_1, u_2, \ldots, u_{n-1}$. Let $w \in \mathbb{R}^n$ be a nonzero vector in the orthogonal complement of $\mathrm{span}\{u_1, u_2, \ldots, u_{n-1}\}$. Define $\sigma \in \{-1, +1\}^n$ by

$$\sigma_i = \begin{cases} \operatorname{sgn} w_i & \text{if } w_i \neq 0, \\ 1 & \text{otherwise.} \end{cases}$$

Then $\sigma = \operatorname{sgn}(\sum_{i=1}^{n-1} \lambda_i u_i)$ for some reals $\lambda_1, \ldots, \lambda_{n-1}$, where the sign function is applied componentwise. In particular, $\langle w, \sum_{i=1}^{n-1} \lambda_i u_i \rangle > 0$. But this is impossible since w was chosen to be orthogonal to $u_1, u_2, \ldots, u_{n-1}$. □

The above theorem was in fact the *first* lower bound on unbounded-error communication complexity.

6 Multiparty communication

We now move on to multiparty communication, a topic that is particularly rewarding in its mathematical depth and its applications to many other areas of theoretical computer science. In this setting, k communicating parties need to compute a Boolean-valued function $f(x_1, x_2, \ldots, x_k)$ with k arguments. Each party knows one or more of the arguments to f, but not all. The more information the parties have available to them, the less communication is required. In the extreme setting known as the *number-on-the-forehead model*, each party knows exactly $k-1$ arguments, namely $x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_k$ in the ith party's case. One can visualize this model by thinking of the k parties as seated in a circle, with $x_1, x_2, \ldots, x_k$ written on the foreheads of parties $1, 2, \ldots, k$, respectively. Any given party sees all the arguments except for the one on the party's own

forehead, hence the terminology. The number-on-the-forehead model, introduced by Chandra, Furst, and Lipton [16], is the most powerful model of multiparty communication and is therefore the standard setting in which to prove communication lower bounds.

In this model, the parties communicate via a broadcast channel, with a bit sent by any party instantly reaching everyone else. They also share an unlimited supply of random bits. Analogous to the two-party case, a multiparty communication protocol *computes f with error ϵ* if on every input $(x_1, x_2, \ldots, x_k)$, it outputs the correct answer $f(x_1, x_2, \ldots, x_k)$ with probability at least $1 - \epsilon$. The *cost* of a protocol is the total number of broadcasts on the worst-case input; as usual, the shared randomness does not count toward the communication cost. The *ϵ-error randomized communication complexity* of f, denoted $R_\epsilon(f)$, is the least cost of an ϵ-error communication protocol for f in this model. Again, the canonical quantity to study is $R_{1/3}(F)$, where the choice of $1/3$ is largely arbitrary and can be replaced by any other constant in $(0, 1/2)$ without affecting the theory in any way.

Multiparty set disjointness. The multiparty set disjointness problem is by far the most studied problem in this line of work. In the k-party setting, the inputs to the problem are sets $S_1, S_2, \ldots, S_k \subseteq \{1, 2, \ldots, n\}$, and the ith party knows all the inputs except for S_i. Their goal is to determine whether the sets have empty intersection: $S_1 \cap S_2 \cap \cdots \cap S_k = \varnothing$. When specialized to $k = 2$, this definition is entirely consistent with the two-party set disjointness problem in Sections 1–4. It is common to represent the input to multiparty set disjointness as a $k \times n$ Boolean matrix $X = [x_{ij}]$, whose rows correspond to the characteristic vectors of the input sets. In this notation, set disjointness is given by the simple formula

$$\mathrm{DISJ}_{k,n}(X) = \bigwedge_{j=1}^{n} \bigvee_{i=1}^{k} \overline{x_{ij}}. \tag{1}$$

Progress on the communication complexity of set disjointness for $k \geqslant 3$ parties is summarized in Table 1. In a surprising result, Grolmusz [22] proved an upper bound of $O(\log^2 n + k^2 n/2^k)$. Proving a strong lower bound, even for $k = 3$, turned out to be difficult. Tesson [51] and Beame et al. [9] obtained a lower bound of $\Omega\big(\frac{1}{k} \log n\big)$ for randomized protocols. Four years later, Lee and Shraibman [30] and Chattopadhyay and Ada [18] gave an improved result. These authors generalized the *pattern matrix method* of [44, 45] to $k \geqslant 3$ parties and thereby obtained a lower bound of $\Omega(n/2^{2^k k})^{1/(k+1)}$ on the randomized communication complexity of set disjointness. Their lower bound was strengthened by Beame and Huynh-Ngoc [8] to $(n^{\Omega(\sqrt{k/\log n})}/2^{k^2})^{1/(k+1)}$, which is an improvement for k large enough. All lower bounds listed up to this point are weaker than $\Omega(n/2^{k^3})^{1/(k+1)}$, which means that they become subpolynomial as soon as the number of parties k starts to grow. Three years later, we obtained [47] a lower bound of $\Omega(n/4^k)^{1/4}$ on the randomized communication complexity of set

Bound	Reference
$O\left(\log^2 n + \frac{k^2 n}{2^k}\right)$	Grolmusz [22]
$\Omega\left(\frac{\log n}{k}\right)$	Tesson [51] Beame, Pitassi, Segerlind, and Wigderson [9]
$\Omega\left(\frac{n}{2^{k^k}}\right)^{\frac{1}{k+1}}$	Lee and Shraibman [30] Chattopadhyay and Ada [18]
$\left(\frac{n^{\Omega(\sqrt{k/\log n})}}{2^{k^2}}\right)^{\frac{1}{k+1}}$	Beame and Huynh-Ngoc [8]
$\Omega\left(\frac{n}{4^k}\right)^{1/4}$	Sherstov [47]
$\Omega\left(\frac{\sqrt{n}}{2^k k}\right)$	Sherstov [49]

Table 1. Communication complexity of k-party set disjointness.

disjointness, which remains polynomial for up to $k \approx \frac{1}{2}\log n$ and comes close to matching Grolmusz's upper bound. Most recently [49], we improved the lower bound quadratically to $\Omega(\sqrt{n}/2^k k)$, which is the strongest bound known. This lower bound also holds for *quantum* multiparty protocols, in which case it is tight. However, it is conceivable that the *classical* randomized communication complexity of set disjointness is $\Omega(n/c^k)$ for some constant $c > 1$. Proving such a lower bound, or showing that it does not hold, is a fascinating open problem.

The lower bound from [49] is too demanding to discuss in this survey. In what follows, we will instead focus on the next best lower bound $\Omega(n/4^k)^{1/4}$.

Anatomy of multiparty protocols. Recall that the building blocks of two-party communication protocols are combinatorial rectangles. The corresponding objects in k-party communication are called *cylinder intersections* [6]. For a k-party problem with domain $X_1 \times X_2 \times \cdots \times X_k$, a cylinder intersection is an arbitrary function $\chi\colon X_1 \times X_2 \times \cdots \times X_k \to \{0,1\}$ of the form

$$\chi(x_1, \ldots, x_k) = \prod_{i=1}^{k} \chi_i(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_k),$$

where $\chi_i\colon X_1 \times \cdots \times X_{i-1} \times X_{i+1} \times \cdots \times X_k \to \{0,1\}$. In other words, a k-dimensional cylinder intersection is the product of k Boolean functions, where the ith function does not depend on the ith coordinate but may depend arbitrarily

on the other $k-1$ coordinates. As one would expect, combinatorial rectangles are cylinder intersections for $k = 2$. The following fundamental result is the multiparty analogue of Fact 1.

FACT 11 (Babai, Nisan, and Szegedy). *Let $\Pi\colon X_1 \times X_2 \times \cdots \times X_k \to \{0,1\}$ be a deterministic k-party communication protocol with cost c. Then there exist cylinder intersections $\chi_1, \ldots, \chi_{2^c}$ with pairwise disjoint support such that*

$$\Pi = \sum_{i=1}^{2^c} \chi_i.$$

By viewing a randomized protocol with cost c as a probability distribution on deterministic protocols of cost at most c, one obtains the following corollary, where $\|\cdot\|_\infty$ denotes as usual the ℓ_∞ norm.

COROLLARY 12. *Let $f\colon X_1 \times X_2 \times \cdots \times X_k \to \{0,1\}$ be a given communication problem. If $R_\epsilon(f) = c$, then there exists a linear combination $\Pi = \sum_\chi a_\chi \chi$ of cylinder intersections with $\sum_\chi |a_\chi| \leqslant 2^c$ such that*

$$\|f - \Pi\|_\infty \leqslant \epsilon.$$

Analytic preliminaries. For the past few years, analytic tools have played an increasingly important role in communication complexity theory. We will need two such tools, the Fourier transform and polynomial approximation theory. Consider the real vector space of functions $\phi\colon \{0,1\}^n \to \mathbb{R}$. For $S \subseteq \{1,2,\ldots,n\}$, define $\chi_S\colon \{0,1\}^n \to \{-1,+1\}$ by $\chi_S(x) = \prod_{i\in S}(-1)^{x_i}$. Then every function $\phi\colon \{0,1\}^n \to \mathbb{R}$ has a unique representation of the form $\phi = \sum_S \hat{\phi}(S)\,\chi_S$, where $\hat{\phi}(S) = 2^{-n}\sum_{x\in\{0,1\}^n} \phi(x)\chi_S(x)$. The reals $\hat{\phi}(S)$ are called the *Fourier coefficients of* ϕ, and the mapping $\phi \mapsto \hat{\phi}$ is the *Fourier transform* of ϕ.

The *ϵ-approximate degree* of a function $\phi\colon \{0,1\}^n \to \mathbb{R}$, denoted $\deg_\epsilon(\phi)$, is the least degree of a multivariate real polynomial p that approximates ϕ within ϵ pointwise: $\|\phi - p\|_\infty \leqslant \epsilon$. We also define $E(\phi, d) = \min_p \|\phi - p\|_\infty$, where the minimum is over multivariate real polynomials p of degree at most d. Thus, $E(\phi, d)$ is the least error to which ϕ can be approximated pointwise by a polynomial of degree at most d. In this notation, $\deg_\epsilon(\phi) = \min\{d : E(\phi, d) \leqslant \epsilon\}$. The approximate degree is an extensively studied complexity measure of Boolean functions. The first result in this line of work is due to Nisan and Szegedy [36], who studied the function $\mathrm{AND}_n(x) = \bigwedge_{i=1}^n x_i$.

THEOREM 13 (Nisan and Szegedy). $\deg_{1/3}(\mathrm{AND}_n) = \Theta(\sqrt{n})$.

The $\Omega(n/4^k)^{1/4}$ lower bound. We are now in a position to present the lower bound on the randomized communication complexity of multiparty set disjointness from [47]. The technical centerpiece of this result is the following lemma, which analyzes the correlation of cylinder intersections with the XOR of several independent copies of set disjointness.

LEMMA 14 (Sherstov). *Let k and r be given parameters. Then there is a probability distribution μ on the domain of $\mathrm{DISJ}_{k,r}$ such that*

$$\mu(\mathrm{DISJ}_{k,r}^{-1}(0)) = \mu(\mathrm{DISJ}_{k,r}^{-1}(1))$$

and

$$\left| \mathop{\mathbf{E}}_{X_1,\ldots,X_n \sim \mu} \left[\chi(X_1,\ldots,X_n) \prod_{i=1}^{n} (-1)^{\mathrm{DISJ}_{k,r}(X_i)} \right] \right| \leqslant \left(\frac{2^{k-1}}{\sqrt{r}} \right)^n$$

for every n and every k-party cylinder intersection χ.

A few general remarks are in order before we delve into the proof of the communication lower bound for set disjointness. The proof is best understood by abstracting away from the set disjointness problem and considering arbitrary composed functions. Specifically, let G be a k-party communication problem, with domain $X = X_1 \times X_2 \times \cdots \times X_k$. We refer to G as a *gadget*. We are interested in the communication complexity of functions of the form $F = f(G, G, \ldots, G)$, where $f\colon \{0,1\}^n \to \{0,1\}$. Thus, F is a k-party communication problem with domain $X^n = X_1^n \times X_2^n \times \cdots \times X_k^n$. The motivation for studying such compositions is clear from the defining equation (1) for multiparty set disjointness, which shows that $\mathrm{DISJ}_{k,nr} = \mathrm{AND}_n(\mathrm{DISJ}_{k,r}, \ldots, \mathrm{DISJ}_{k,r})$. A recent line of research [45, 50, 30, 18, 8, 17, 47, 49] gives communication lower bounds for compositions $f(G, G, \ldots, G)$ in terms of the approximate degree of f. For the purpose of proving communication lower bounds for set disjointness, the gadget G needs to be representable as $G = \mathrm{DISJ}_{k,r}$ with $r = r(n,k)$ as small as possible. This miniaturization challenge quickly becomes hard.

THEOREM 15 (Sherstov). *Let $f\colon \{0,1\}^n \to \{0,1\}$ be given. Consider the k-party communication problem $F = f(\mathrm{DISJ}_{k,r}, \ldots, \mathrm{DISJ}_{k,r})$ Then for all $\epsilon, \delta \geqslant 0$,*

$$2^{R_\epsilon(F)} \geqslant (\delta - \epsilon) \left(\frac{\deg_\delta(f)\sqrt{r}}{2^k e n} \right)^{\deg_\delta(f)}. \tag{2}$$

Proof. Let μ be the probability distribution from Lemma 14. Let μ_0 and μ_1 stand for the probability distributions induced by μ on $\mathrm{DISJ}_{k,r}^{-1}(0)$ and $\mathrm{DISJ}_{k,r}^{-1}(1)$, respectively. Consider the following averaging operator L, which linearly sends real functions χ on $(\{0,1\}^{k \times r})^n$ to real functions on $\{0,1\}^n$:

$$(L\chi)(z) = \mathop{\mathbf{E}}_{X_1 \sim \mu_{z_1}} \cdots \mathop{\mathbf{E}}_{X_n \sim \mu_{z_n}} [\chi(X_1,\ldots,X_n)].$$

Observe that $LF = f$. When χ is a k-party cylinder intersection, the Fourier coefficients of $L\chi$ obey

$$\begin{aligned}
|\widehat{L\chi}(S)| &= \left| \mathop{\mathbf{E}}_{z \in \{0,1\}^n} \mathop{\mathbf{E}}_{X_1 \sim \mu_{z_1}} \cdots \mathop{\mathbf{E}}_{X_n \sim \mu_{z_n}} \left[\chi(X_1, \ldots, X_n) \prod_{i \in S} (-1)^{z_i} \right] \right| \\
&= \left| \mathop{\mathbf{E}}_{X_1, \ldots, X_n \sim \mu} \left[\chi(X_1, \ldots, X_n) \prod_{i \in S} (-1)^{\mathrm{DISJ}_{k,r}(X_i)} \right] \right| \\
&\leqslant \left(\frac{2^{k-1}}{\sqrt{r}} \right)^{|S|}, \qquad (3)
\end{aligned}$$

where the second equality uses the fact that μ places equal weight on $\mathrm{DISJ}_{k,r}^{-1}(0)$ and $\mathrm{DISJ}_{k,r}^{-1}(1)$, and the final step follows by Lemma 14.

Fix a randomized protocol for F with error ϵ and cost $c = R_\epsilon(F)$. Approximate F as in Corollary 12 by a linear combination of cylinder intersections $\mathit{\Pi} = \sum_\chi a_\chi \chi$, where $\sum_\chi |a_\chi| \leqslant 2^c$. For any positive integer d, the triangle inequality gives

$$E(f, d-1) \leqslant \|f - L\mathit{\Pi}\|_\infty + E(L\mathit{\Pi}, d-1). \qquad (4)$$

We proceed to bound the two terms on the right-hand side of this inequality.

(i) By the linearity of L,

$$\|f - L\mathit{\Pi}\|_\infty = \|L(F - \mathit{\Pi})\|_\infty \leqslant \epsilon, \qquad (5)$$

where the last step uses the bound $\|F - \mathit{\Pi}\|_\infty \leqslant \epsilon$ from Corollary 12.

(ii) Discarding the Fourier coefficients of $L\mathit{\Pi}$ of order d and higher gives

$$\begin{aligned}
E(L\mathit{\Pi}, d-1) &\leqslant \min \left\{ 1, \sum_\chi |a_\chi| \sum_{|S| \geqslant d} |\widehat{L\chi}(S)| \right\} \\
&\leqslant \min \left\{ 1, 2^c \sum_{i=d}^{n} \binom{n}{i} \left(\frac{2^{k-1}}{\sqrt{r}} \right)^i \right\} \\
&\leqslant 2^c \left(\frac{2^k \mathrm{e} n}{d\sqrt{r}} \right)^d, \qquad (6)
\end{aligned}$$

where the second step uses (3).

Substituting the newly obtained estimates (5) and (6) into (4),

$$E(f, d-1) \leqslant \epsilon + 2^c \left(\frac{2^k \mathrm{e} n}{d\sqrt{r}} \right)^d.$$

For $d = \deg_\delta(f)$, the left-hand side must exceed δ, forcing (2). $\square$

As an immediate consequence, we obtain the claimed lower bound on the multiparty communication complexity of set disjointness [47]:

Corollary (Sherstov).

$$R_{1/3}(\mathrm{DISJ}_{k,n}) \geqslant \Omega\left(\frac{n}{4^k}\right)^{1/4}. \tag{7}$$

Proof. Recall that $\mathrm{DISJ}_{k,nr} = \mathrm{AND}_n(\mathrm{DISJ}_{k,r}, \ldots, \mathrm{DISJ}_{k,r})$ for all integers n, r. Theorem 13 guarantees that $\deg_{1/3}(\mathrm{AND}_n) > c\sqrt{n}$ for some constant $c > 0$. Thus, letting $f = \mathrm{AND}_n$, $\delta = 1/3$, $\epsilon = 1/4$, and $r = 4^{k+2}\lceil\sqrt{n}/c\rceil^2$ in Theorem 15 gives

$$\begin{aligned} &R_{1/4}(\mathrm{DISJ}_{k,4^{k+2}n\lceil\sqrt{n}/c\rceil^2}) \\ &\qquad = R_{1/4}(\mathrm{AND}_n(\mathrm{DISJ}_{k,4^{k+2}\lceil\sqrt{n}/c\rceil^2}, \ldots, \mathrm{DISJ}_{k,4^{k+2}\lceil\sqrt{n}/c\rceil^2})) \geqslant \Omega(\sqrt{n}), \end{aligned}$$

which is logically equivalent to (7). □

7 Other gems

We have only focused on a small sample of results on the set disjointness problem. Prominently absent in our survey is the fascinating and influential body of work on the *quantum* communication complexity of set disjointness [14, 41, 1, 45, 50]. Much can also be said about deterministic, nondeterministic, and Merlin-Arthur multiparty protocols [25, 21, 47, 49]. Another compelling topic is the multiparty communication complexity of the set disjointness problem in the *number-in-hand* model [7, 15, 12], where each party sees only one of the input sets $S_1, S_2, \ldots, S_k$ as opposed to all but one. Lower bounds for such multiparty protocols play an important role in the study of streaming algorithms. Finally, we have not discussed XOR lemmas and direct product theorems, which deal with the communication complexity of simultaneously solving several independent copies of set disjointness [27, 9, 10, 23, 26, 48, 47, 49].

Acknowledgments

I am thankful to the organizers of MFCS 2014 and Zoltán Ésik in particular for this opportunity to share my passion for the set disjointness problem.

References

1. Aaronson, S., Ambainis, A.: Quantum search of spatial regions. Theory of Computing 1(1), 47–79 (2005)
2. Aho, A.V., Ullman, J.D., Yannakakis, M.: On notions of information transfer in VLSI circuits. In: *Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing* (STOC). pp. 133–139 (1983)

3. Alon, N., Frankl, P., Rödl, V.: Geometrical realization of set systems and probabilistic communication complexity. In: *Proceedings of the Twenty-Sixth Annual IEEE Symposium on Foundations of Computer Science* (FOCS). pp. 277–280 (1985)
4. Alon, N., Matias, Y., Szegedy, M.: The space complexity of approximating the frequency moments. J. Comput. Syst. Sci. 58(1), 137–147 (1999)
5. Babai, L., Frankl, P., Simon, J.: Complexity classes in communication complexity theory. In: *Proceedings of the Twenty-Seventh Annual IEEE Symposium on Foundations of Computer Science* (FOCS). pp. 337–347 (1986)
6. Babai, L., Nisan, N., Szegedy, M.: Multiparty protocols, pseudorandom generators for logspace, and time-space trade-offs. J. Comput. Syst. Sci. 45(2), 204–232 (1992)
7. Bar-Yossef, Z., Jayram, T.S., Kumar, R., Sivakumar, D.: An information statistics approach to data stream and communication complexity. J. Comput. Syst. Sci. 68(4), 702–732 (2004)
8. Beame, P., Huynh-Ngoc, D.T.: Multiparty communication complexity and threshold circuit complexity of AC^0. In: *Proceedings of the Fiftieth Annual IEEE Symposium on Foundations of Computer Science* (FOCS). pp. 53–62 (2009)
9. Beame, P., Pitassi, T., Segerlind, N., Wigderson, A.: A strong direct product theorem for corruption and the multiparty communication complexity of disjointness. Computational Complexity 15(4), 391–432 (2006)
10. Ben-Aroya, A., Regev, O., de Wolf, R.: A hypercontractive inequality for matrix-valued functions with applications to quantum computing and LDCs. In: *Proceedings of the Forty-Ninth Annual IEEE Symposium on Foundations of Computer Science* (FOCS). pp. 477–486 (2008)
11. Ben-David, S., Eiron, N., Simon, H.U.: Limitations of learning via embeddings in Euclidean half spaces. J. Mach. Learn. Res. 3, 441–461 (2003)
12. Braverman, M., Ellen, F., Oshman, R., Pitassi, T., Vaikuntanathan, V.: A tight bound for set disjointness in the message-passing model. In: *Proceedings of the Fifty-Fourth Annual IEEE Symposium on Foundations of Computer Science* (FOCS). pp. 668–677 (2013)
13. Braverman, M., Garg, A., Pankratov, D., Weinstein, O.: From information to exact communication. In: *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing* (STOC). pp. 151–160 (2013)
14. Buhrman, H., Cleve, R., Wigderson, A.: Quantum vs. classical communication and computation. In: *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing* (STOC). pp. 63–68 (1998)
15. Chakrabarti, A., Khot, S., Sun, X.: Near-optimal lower bounds on the multi-party communication complexity of set disjointness. In: *Proceedings of the Eighteenth Annual IEEE Conference on Computational Complexity* (CCC). pp. 107–117 (2003)
16. Chandra, A.K., Furst, M.L., Lipton, R.J.: Multi-party protocols. In: *Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing* (STOC). pp. 94–99 (1983)
17. Chattopadhyay, A.: Circuits, Communication, and Polynomials. Ph.D. thesis, McGill University (2008)
18. Chattopadhyay, A., Ada, A.: Multiparty communication complexity of disjointness. In: Electronic Colloquium on Computational Complexity (ECCC) (January 2008), report TR08-002
19. Forster, J.: A linear lower bound on the unbounded error probabilistic communication complexity. J. Comput. Syst. Sci. 65(4), 612–625 (2002)

20. Forster, J., Krause, M., Lokam, S.V., Mubarakzjanov, R., Schmitt, N., Simon, H.U.: Relations between communication complexity, linear arrangements, and computational complexity. In: Proc. of the 21st Conf. on Foundations of Software Technology and Theoretical Computer Science (FST TCS). pp. 171–182 (2001)
21. Gavinsky, D., Sherstov, A.A.: A separation of NP and coNP in multiparty communication complexity. Theory of Computing 6(10), 227–245 (2010)
22. Grolmusz, V.: The BNS lower bound for multi-party protocols in nearly optimal. Inf. Comput. 112(1), 51–54 (1994)
23. Jain, R., Klauck, H., Nayak, A.: Direct product theorems for classical communication complexity via subdistribution bounds. In: *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing* (STOC). pp. 599–608 (2008)
24. Kalyanasundaram, B., Schnitger, G.: The probabilistic communication complexity of set intersection. SIAM J. Discrete Math. 5(4), 545–557 (1992)
25. Klauck, H.: Rectangle size bounds and threshold covers in communication complexity. In: *Proceedings of the Eighteenth Annual IEEE Conference on Computational Complexity* (CCC). pp. 118–134 (2003)
26. Klauck, H.: A strong direct product theorem for disjointness. In: *Proceedings of the Forty-Second Annual ACM Symposium on Theory of Computing* (STOC). pp. 77–86 (2010)
27. Klauck, H., Špalek, R., de Wolf, R.: Quantum and classical strong direct product theorems and optimal time-space tradeoffs. SIAM J. Comput. 36(5), 1472–1493 (2007)
28. Klivans, A.R., Servedio, R.A.: Learning DNF in time $2^{\tilde{O}(n^{1/3})}$. J. Comput. Syst. Sci. 68(2), 303–318 (2004)
29. Kushilevitz, E., Nisan, N.: Communication complexity. Cambridge University Press (1997)
30. Lee, T., Shraibman, A.: Disjointness is hard in the multiparty number-on-the-forehead model. Computational Complexity 18(2), 309–336 (2009)
31. Linial, N., Mendelson, S., Schechtman, G., Shraibman, A.: Complexity measures of sign matrices. Combinatorica 27(4), 439–463 (2007)
32. Lovász, L., Saks, M.E.: Lattices, Möbius functions and communication complexity. In: *Proceedings of the Twenty-Ninth Annual IEEE Symposium on Foundations of Computer Science* (FOCS). pp. 81–90 (1988)
33. Mehlhorn, K., Schmidt, E.M.: Las Vegas is better than determinism in VLSI and distributed computing. In: *Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing* (STOC). pp. 330–337 (1982)
34. Newman, I.: Private vs. common random bits in communication complexity. Inf. Process. Lett. 39(2), 67–71 (1991)
35. Nisan, N., Segal, I.: The communication requirements of efficient allocations and supporting prices. J. Economic Theory 129(1), 192–224 (2006)
36. Nisan, N., Szegedy, M.: On the degree of Boolean functions as real polynomials. Computational Complexity 4, 301–313 (1994)
37. Nisan, N., Wigderson, A.: On rank vs. communication complexity. Combinatorica 15(4), 557–565 (1995)
38. Paturi, R., Simon, J.: Probabilistic communication complexity. J. Comput. Syst. Sci. 33(1), 106–123 (1986)
39. Razborov, A.A.: Applications of matrix methods to the theory of lower bounds in computational complexity. Combinatorica 10(1), 81–93 (1990)
40. Razborov, A.A.: On the distributional complexity of disjointness. Theor. Comput. Sci. 106(2), 385–390 (1992)

41. Razborov, A.A.: Quantum communication complexity of symmetric predicates. Izvestiya: Mathematics 67(1), 145–159 (2003)
42. Razborov, A.A., Sherstov, A.A.: The sign-rank of AC^0. SIAM J. Comput. 39(5), 1833–1855 (2010), preliminary version in *Proceedings of the Forty-Ninth Annual IEEE Symposium on Foundations of Computer Science* (FOCS), 2008
43. Sherstov, A.A.: Halfspace matrices. Computational Complexity 17(2), 149–178 (2008), preliminary version in *Proceedings of the Twenty-Second Annual IEEE Conference on Computational Complexity* (CCC), 2007
44. Sherstov, A.A.: Separating AC^0 from depth-2 majority circuits. SIAM J. Comput. 38(6), 2113–2129 (2009), preliminary version in *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing* (STOC), 2007
45. Sherstov, A.A.: The pattern matrix method. SIAM J. Comput. 40(6), 1969–2000 (2011), preliminary version in *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing* (STOC), 2008
46. Sherstov, A.A.: The unbounded-error communication complexity of symmetric functions. Combinatorica 31(5), 583–614 (2011), preliminary version in *Proceedings of the Forty-Ninth Annual IEEE Symposium on Foundations of Computer Science* (FOCS), 2008
47. Sherstov, A.A.: The multiparty communication complexity of set disjointness. In: *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing* (STOC). pp. 525–544 (2012)
48. Sherstov, A.A.: Strong direct product theorems for quantum communication and query complexity. SIAM J. Comput. 41(5), 1122–1165 (2012), preliminary version in *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing* (STOC), 2011
49. Sherstov, A.A.: Communication lower bounds using directional derivatives. In: *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing* (STOC). pp. 921–930 (2013)
50. Shi, Y., Zhu, Y.: Quantum communication complexity of block-composed functions. Quantum Information & Computation 9(5–6), 444–460 (2009)
51. Tesson, P.: Computational complexity questions related to finite monoids and semigroups. Ph.D. thesis, McGill University (2003)
52. Yao, A.C.C.: Some complexity questions related to distributive computing. In: *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing* (STOC). pp. 209–213 (1979)
53. Yao, A.C.C.: Lower bounds by probabilistic arguments. In: *Proceedings of the Twenty-Fourth Annual IEEE Symposium on Foundations of Computer Science* (FOCS). pp. 420–428 (1983)