


SIGACT News Complexity Theory Column, March 2024

Structure in Communication Complexity and Constant-Cost Complexity Classes

Hamed Hatami¹ Pooya Hatami²

Abstract

Several theorems and conjectures in communication complexity state or speculate that the complexity of a matrix in a given communication model is controlled by a related analytic or algebraic matrix parameter, e.g., rank, sign-rank, discrepancy, etc. The forward direction is typically easy as the structural implications of small complexity often imply a bound on some matrix parameter. The challenge lies in establishing the reverse direction, which requires understanding the structure of Boolean matrices for which a given matrix parameter is small or large. We will discuss several research directions that align with this overarching theme.

1 Introduction

In 1979, Yao [Yao79] introduced an abstract model for analyzing communication. It quickly became apparent that the applications of this elegant paradigm go far beyond the concept of communication. Many results in communication complexity have equivalent formulations in other fields that are equally natural, and the techniques developed within this framework have proven to be powerful tools applicable across various domains. Today, communication complexity is a vibrant research area with many connections across theoretical computer science and mathematics: in learning theory, circuit design, pseudorandomness, data streaming, data structures, computational complexity, computer networks, time-space trade-offs, discrepancy theory, and property testing.

In this article, we focus on the most standard framework where a communication problem is simply a *Boolean* matrix. Formally, there are two parties, often called Alice and Bob, and a communication problem is defined by a matrix $F \in \{0, 1\}^{\mathcal{X} \times \mathcal{Y}}$. Alice receives a row index $x \in \mathcal{X}$, and Bob receives a column index $y \in \mathcal{Y}$. Together, they should compute the entry $F(x, y)$ by exchanging bits of information according to a previously agreed-on protocol tailored to F . There is no restriction on their computational power; the only measure we care to minimize is the number of exchanged bits.

Many questions in communication complexity concern the basic structural properties of Boolean matrices and are relevant to any field that requires an in-depth analysis of them. Mathematicians,

¹McGill University, Montreal, QC, Canada . hatami@cs.mcgill.ca. Supported by an NSERC grant.

²The Ohio State University, Columbus, OH, USA. pooyahat@gmail.com. Supported by NSF grant CCF-1947546.

of course, have studied matrices for centuries through the lenses of linear algebra, geometry, and analysis. They have produced an extensive collection of tools and theories that apply to any field dealing with these mathematical objects. However, the assumption of Booleanity introduces a novel angle and unveils new problems and challenges.

Consider the example of rank. Elementary linear algebra provides several satisfactory structural descriptions of small-rank matrices. For example, one could construct a small-rank matrix by summing a few rank-one matrices. In contrast, the structure of small-rank *Boolean matrices* is the subject of the most well-known conjecture in communication complexity, *the log-rank conjecture*.

There are many theorems and conjectures in communication complexity that fall into a similar paradigm as the log-rank conjecture: They state or speculate that the complexity of a matrix in a given communication model is essentially determined by a related analytic or algebraic matrix parameter, e.g., rank, sign-rank, discrepancy, trace norm, approximate trace norm, γ_2 -factorization norm, approximate factorization norm. The forward direction is typically easy as the structural implications of small complexity often imply a bound on some matrix parameter. The challenge lies in establishing the reverse direction, which requires understanding *the structure of Boolean matrices* for which a given matrix parameter is small or large.

In this article, we discuss some new research directions and open problems, as well as some classical ones, that align with this overarching theme.

Notation: All logarithms are in base 2. Sometimes, we use $a \lesssim b$ to denote $a = O(b)$. Let \mathbf{I}_m denote the $m \times m$ identity matrix and $\mathbf{J}_{\mathcal{X} \times \mathcal{Y}}$ denote the $\mathcal{X} \times \mathcal{Y}$ all-1 matrix. For a positive integer k , we denote $[k] = \{1, \dots, k\}$. We often identify a Boolean matrix $F_{\mathcal{X} \times \mathcal{Y}}$ with the corresponding function $F : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ defined as $F : (x, y) \mapsto F(x, y)$. We define the *complement* of a communication problem $F = F_{\mathcal{X} \times \mathcal{Y}}$ as $\neg F = \mathbf{J}_{\mathcal{X} \times \mathcal{Y}} - F$. For two $\mathcal{X} \times \mathcal{Y}$ matrices A and B , we denote their entry-wise product (i.e. *Schur product*) by $A \circ B$.

Communication complexity classes: We measure the communication complexity of a matrix $F_{\mathcal{X} \times \mathcal{Y}}$ in relation to the number of input bits $n(F) := \lceil \log \max(|\mathcal{X}|, |\mathcal{Y}|) \rceil$. We often consider $2^n \times 2^n$ matrices $F_n : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, where the inputs of Alice and Bob are n -bit strings. Moreover, similar to computational complexity, the goal is to understand the asymptotic complexity, and therefore, a communication problem typically refers to an infinite family of Boolean matrices rather than a single matrix.

For example, EQUALITY is the family of matrices $\text{EQ}_n : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ with $\text{EQ}(x, y) = 1$ iff $x = y$. Equivalently, EQ_n is the $2^n \times 2^n$ identity matrix, where the rows and columns are labelled with n -bit strings.

In the theory of Turing machines, a polynomial complexity is considered efficient, but this is not a suitable criterion for communication since even in the deterministic model, communication complexity is at most $n + 1$. In an influential paper, Babai, Frankl, and Simon [BFS86] proposed polylogarithmic complexity as the criteria for efficiency and used it to define the communication classes P^{cc} , NP^{cc} , RP^{cc} , BPP^{cc} , PP^{cc} , UPP^{cc} , analogous to the classical computational complexity classes. The definition of communication classes by [BFS86] provides a formal paradigm to compare the power of different communication models.

Since in this article, we only study communication classes, to make the notation less cumbersome, from this point on, we will drop the superscript cc and denote these communication classes simply as P, NP, RP, BPP, PP, UPP.

Constant-cost communication classes: For many communication problems, it may not be completely natural to allow the communication complexity to depend on the matrix size. For example, consider the geometric problem where Alice receives a point $x \in \mathbb{R}^2$, and Bob receives a closed half-space $H \subseteq \mathbb{R}^2$, and they wish to know whether $x \in H$. Alternatively, consider the problem where Alice and Bob receive points $x, y \in \mathbb{R}^2$, respectively, and they wish to know whether the distance between these points is at most 1. In these examples, the number of possible inputs is infinite, and even if we artificially restrict the number of possible inputs to a finite set, the most natural way to represent the inputs is by vectors in \mathbb{R}^2 . Therefore, it is more natural to consider communication protocols that solve the problem (in a certain communication model or under some promised guarantees about the inputs) using a bounded number of communicated bits, independent of the number of possible inputs.

Furthermore, many parameters studied in communication complexity are also relevant to other areas of theoretical computer science, where the Boolean matrices are either infinite or their size is of lesser significance than other features. For example, in learning theory, notions such as VC dimension, Littlestone dimension, margin complexity, and Threshold dimension of binary concept classes are of primary interest rather than the class size. Viewing binary concept classes as Boolean matrices, these parameters are directly related to notions such as stability, discrepancy, and sign-rank that play a central role in communication complexity.

The definition of communication classes by [BFS86], no doubt, is natural and has successfully led to many fruitful lines of research seeking to prove separations between different communication classes. However, if one wishes to view communication complexity in these broader contexts, it becomes essential also to analyze communication problems that have uniformly bounded complexity. Indeed, there is already a rapidly growing body of work dedicated to studying constant-cost communication [LMSS07, LS09a, BCK14, HHL20, Har20, HWZ22, HHP⁺22, EHK22, HHH23, CHHS23, ACHS23, HZ24].

These considerations motivate the definition of the *constant-cost* analogues of the communication classes of [BFS86]. Here, the criterion of effectiveness is a $O(1)$ complexity, independent of the input size n . We will denote these classes by adding a 0 subscript to the corresponding polylogarithmic class, where 0 refers to the power d in the polylogarithmic complexity $O(\log^d(n))$. For example, a family of Boolean matrices F_n is in P_0 if their deterministic communication complexity is uniformly bounded by a constant $C = O(1)$.

We will formally define and discuss the classes $P_0, NP_0, RP_0, BPP_0, PP_0, UPP_0$ later, but for now, let us mention that, unlike in the poly-logarithmic communication classes where $P \subsetneq RP \subsetneq NP$ and $BPP \subsetneq PP$, here we have $P_0 = NP_0 \subsetneq RP_0$ and $BPP_0 = PP_0$. We also believe that $BPP_0 \not\subseteq UPP_0$, while Newman's lemma [New91] implies that $BPP \subseteq UPP$.

Monochromatic rectangles: A *monochromatic rectangle* in a Boolean matrix $F_{\mathcal{X} \times \mathcal{Y}}$ refers to a submatrix $F_{S \times T}$, where all entries share the same value. Monochromatic rectangles are the building blocks of every communication protocol, and many results and problems in communication complexity focus on describing the structure of Boolean matrices in terms of these rectangles or showing the existence of large monochromatic rectangles in them.

Define the *rectangle ratio* of a Boolean matrix $F_{\mathcal{X} \times \mathcal{Y}}$ as

$$\text{rect}(F) := \max_R \frac{|R|}{|\mathcal{X} \times \mathcal{Y}|},$$

where the maximum is over monochromatic rectangles in F .

Concerning communication complexity, the parameter $\text{rect}(\cdot)$ has a few drawbacks: While adding new rows and columns to a matrix can only increase the communication complexity, it can deteriorate $\text{rect}(F)$ by creating large monochromatic rectangles. Moreover, repeating a row or column of F does not change its communication complexity, but it can affect $\text{rect}(F)$. Therefore, it is more natural to consider the following version of rect that considers weighted monochromatic rectangles, first studied by Impagliazzo and Williams [IW10].

Definition 1.1 (Weighted Rectangle Ratio). *For every Boolean matrix F , define*

$$\text{wrect}(F) := \inf_{\mu} \max_R \mu(R),$$

where the infimum is over all product probability measures $\mu = \mu_X \times \mu_Y$ on $\mathcal{X} \times \mathcal{Y}$, and the maximum is taken over all monochromatic rectangles in F .

Both $\text{rect}(\cdot)$ and $\text{wrect}(\cdot)$ are natural parameters that measure the existence of structure in Boolean matrices. As we will discuss throughout the paper, they play a crucial role in several applications in communication complexity, and some basic and consequential questions regarding these quantities remain open.

2 Deterministic communication complexity and rank

A *deterministic communication protocol* is defined by a binary tree, where every internal node v specifies which player speaks at that node and what bit they must send. For example, if an internal node v is associated with Alice, it is labelled with a function $a_v : \mathcal{X} \rightarrow \{0, 1\}$, which prescribes the bit sent by Alice at this node. After this bit is sent, the players move to the corresponding child of v : they move to the left child if the bit is 0 and to the right child if the bit is 1. They continue this process until they reach a leaf. Every leaf is labelled with 0 or 1, which corresponds to the output of the protocol. The *cost* of the protocol is the height of the tree, which is equal to the maximum number of bits exchanged on any input. The *deterministic communication complexity* of F , denoted by $D(F)$, is the smallest cost of a protocol that computes F correctly on all inputs. Let \mathbf{P} and \mathbf{P}_0 be, respectively, the class of problems with poly-logarithmic and $O(1)$ deterministic communication complexity.

The fact that the bit communicated at every node depends only on one of the inputs implies the most useful fact in communication complexity: for every node v , the set of all inputs that lead the protocol to v is a *combinatorial rectangle*, which is a product set $R_v = S_v \times T_v$ with $S_v \subseteq \mathcal{X}$ and $T_v \subseteq \mathcal{Y}$.

For every $b \in \{0, 1\}$, a combinatorial rectangle R is called a *b -monochromatic rectangle* if $F(x, y) = b$ for every entry $(x, y) \in R$. Note that for every b -labelled leaf ℓ , the set R_ℓ must be a b -monochromatic rectangle. Therefore, the leaves of a deterministic protocol partition the matrix into monochromatic rectangles.

A Boolean matrix F with small $D(F)$ is highly structured, as it consists of at most $2^{D(F)}$ monochromatic rectangles. This partition of F into at most $2^{D(F)}$ disjoint monochromatic rectangles implies

$$D(F) \geq \log(1/\text{rect}(F)). \tag{1}$$

Moreover, since the real rank of a monochromatic rectangle is at most 1, we have $\text{rk}(F) \leq 2^{D(F)}$, where $\text{rk}(F)$ denotes the real rank of F . Combined with the easy fact $D(F) \leq \text{rk}(F) + 1$, we have

$$\log \text{rk}(F) \leq D(F) \leq \text{rk}(F) + 1. \tag{2}$$

Therefore, P_0 coincides with the class of problems with bounded rank. However, the exponential gap between the lower bound and the upper bound is too wide to provide a characterization of P in terms of rank. The *log-rank conjecture* speculates that the lower bound is essentially sharp and $D(F)$ and $\log \text{rk}(F)$ are polynomially equivalent.

Conjecture 1 (The log-rank conjecture [LS88]). *There exists a universal constant $C > 0$ such that for every Boolean matrix F ,*

$$D(F) \leq C (\log \text{rk}(F))^C.$$

The log-rank conjecture, if true, provides a structural description for Boolean matrices of small rank. Namely, they must have the same tree-like partition into monochromatic rectangles as matrices with small deterministic communication complexity. Nisan and Wigderson [NW95] showed that to prove the log-rank conjecture, it suffices to show the existence of one large monochromatic rectangle.

Conjecture 2 (The log-rank conjecture - equivalent formulation). *There exists a universal constant $C > 0$ such that every Boolean matrix F satisfies*

$$\text{rect}^{-1}(F) \leq C \log^C \text{rk}(F).$$

To date, the strongest known bound is still exponentially far from that in [Conjecture 1](#). Concretely, Lovett [Lov16] showed that $D(f) \leq O(\sqrt{\text{rk}(f)} \log \text{rk}(f))$, which was recently improved by Sudakov and Tomon [ST23] to $O(\sqrt{\text{rk}(f)})$. On the lower bound side, a construction of [GPW18a] shows that the constant C cannot be strictly less than 2. We refer to surveys [Lov14, LS23] for a detailed discussion of the log-rank conjecture.

3 Equality oracles and γ_2 -Factorization norm

The results discussed in this section mostly stem from joint work by Lianna Hambardzumyan and the authors [HHH23]. Regarding communication complexity, we will focus on the deterministic model with access to an EQUALITY oracle. We will show that this model has compelling ties to harmonic analysis and operator theory. In particular, we will discuss a conjecture in communication complexity with profound links to the characterization of the idempotents of the algebra of Schur multipliers and Cohen’s celebrated idempotent theorem, a well-known and notable theorem in harmonic analysis.

3.1 An analytic version of the rank lower-bound

The rank lower bound on $D(F)$ follows from the observation that every Boolean matrix F can be written as $F = \sum_{i=1}^{2^{D(F)}} B_i$, where each B_i is a rank-one *Boolean* matrix. If instead of counting the number of rank-one *Boolean* matrices in this sum, we focus on the sum of the coefficients, we can establish an alternative lower bound on $D(F)$ —a lower bound with an analytical flavour.

Define the μ -norm of a *real* matrix A as

$$\|A\|_\mu = \inf_{\lambda_i, B_i} \left\{ \sum_i |\lambda_i| : A = \sum_i \lambda_i B_i \right\},$$

where each B_i is a rank-one *Boolean* matrix and each λ_i is a real number. Note that $\|\cdot\|_\mu$ satisfies all the axioms of a norm, and we have

$$\log \|F\|_\mu \leq D(F). \quad (3)$$

The μ -norm is not well-known in functional analysis, and its definition is tailored to its purpose as a lower bound in communication complexity. Fortunately, the Grothendieck inequality shows that the μ -norm is equivalent to the well-studied γ_2 -factorization norm.

Definition 3.1. *The γ_2 -norm of a real matrix $A_{\mathcal{X} \times \mathcal{Y}}$, denoted by $\|A\|_{\gamma_2}$, is the infimum of c such that there exists a positive integer d and vectors $u_x, v_y \in \mathbb{R}^d$ for $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ with $\langle u_x, v_y \rangle = A(x, y)$ and $\|u_x\|_2 \cdot \|v_y\|_2 \leq c$ for all x, y .*

A key property of the γ_2 -norm is that for every two $\mathcal{X} \times \mathcal{Y}$ real matrices A_1 and A_2 , we have

$$\|A_1 \circ A_2\|_{\gamma_2} \leq \|A_1\|_{\gamma_2} \|A_2\|_{\gamma_2}, \quad (4)$$

where we recall that $A_1 \circ A_2$ is the entry-wise product of A_1 and A_2 .

The Grothendieck inequality implies that for every real matrix A , we have

$$\|A\|_{\gamma_2} \leq \|A\|_\mu \leq 4K \|A\|_{\gamma_2},$$

where $K \leq \frac{\pi}{2 \ln(1+\sqrt{2})} \approx 1.7822$ is the so-called Grothendieck constant for real numbers. In light of this equivalence, we will rephrase the lower bound [Equation \(3\)](#) in terms of the γ_2 -norm:

$$\log \|F\|_{\gamma_2} \leq D(F). \quad (5)$$

If we compare this lower bound to [Equation \(2\)](#), it is natural to wonder whether we can bound $D(F)$ from above by a function of $\|F\|_{\gamma_2}$. The answer is negative. The γ_2 -norm of every identity matrix is 1, but identity matrices can be of arbitrarily large rank and, therefore, of arbitrarily large deterministic communication complexity.

Proposition 3.2. *For every m , the γ_2 -norm of the $m \times m$ identity matrix \mathbf{I}_m is 1,*

Proof. Note that since the standard basis $e_1, \dots, e_m \in \mathbb{R}^m$ satisfies

$$\langle e_i, e_j \rangle = \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases},$$

by [Definition 3.1](#), \mathbf{I}_m satisfies $\|\mathbf{I}_m\|_{\gamma_2} \leq 1$. Moreover, it is clear from the definition of the γ_2 -norm that for every matrix A , we have $\|A\|_{\gamma_2} \geq \|A\|_\infty := \max_{x,y} |A(x, y)|$, and therefore, $\|\mathbf{I}_m\|_{\gamma_2} \geq 1$. \square

What are the Boolean matrices with small γ_2 -norm then? It follows from $\|A\|_{\gamma_2} \geq \|A\|_\infty$ that the γ_2 -norm of every non-zero Boolean matrix is at least 1. Let us first study the Boolean matrices whose γ_2 -norm is exactly 1.

[Proposition 3.2](#) shows that all identity matrices have γ_2 -norm 1. Note that the proof of [Proposition 3.2](#) generalizes to a larger class of matrices, which we call *blocky matrices*.

Definition 3.3 (Blocky matrices). *A Boolean matrix $F_{\mathcal{X} \times \mathcal{Y}}$ is blocky if there exist disjoint sets $\mathcal{X}_i \subseteq \mathcal{X}$ and disjoint sets $\mathcal{Y}_i \subseteq \mathcal{Y}$ such that the support of F is exactly $\bigcup_i \mathcal{X}_i \times \mathcal{Y}_i$.*

Let *Blocky* denote the set of all blocky matrices. It turns out that *Blocky* is precisely the set of Boolean matrices with γ_2 -norm 1.

Proposition 3.4 (Livshits [Liv95]). *A Boolean matrix F satisfies $\|F\|_{\gamma_2} = 1$ iff $F \in \text{Blocky}$.*

Proof. It is observed in [Liv95] that

$$\left\| \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \right\|_{\gamma_2} = \frac{2}{\sqrt{3}} > 1.$$

Since $\|\cdot\|_{\gamma_2}$ norm is invariant under row and column permutations, a Boolean matrix F with $\|F\|_{\gamma_2} = 1$ cannot have any 2×2 submatrices with exactly 3 ones. It is straightforward to verify that a Boolean matrix satisfying this property must be blocky. \square

Since every Boolean matrix with γ_2 -norm 1 is blocky, it is natural to ask whether Boolean matrices of bounded γ_2 -norm can be characterized through blocky matrices.

We first show that if F is generated by entry-wise operations from a few blocky matrices, it must have a small γ_2 -norm.

Proposition 3.5. *Consider $\mathcal{X} \times \mathcal{Y}$ blocky matrices B_1, \dots, B_r and a combining function $\Gamma : \{0, 1\}^r \rightarrow \{0, 1\}$. The Boolean matrix $F_{\mathcal{X} \times \mathcal{Y}}$ defined as*

$$F(x, y) := \Gamma(B_1(x, y), \dots, B_r(x, y)) \quad (6)$$

satisfies $\|F\|_{\gamma_2} \leq 3^r$.

Proof. We prove the statement by induction on r . The base case $r = 0$ is trivial. Next, consider a Boolean matrix F satisfying Equation (6), and write $F = (B_1 \circ F_1) + (\mathbf{J} - B_1) \circ F_2$, where the entries of F_1 and F_2 depend only on B_2, \dots, B_r and \mathbf{J} is the all-1 matrix. Note that \mathbf{J} is a blocky matrix and satisfies $\|\mathbf{J}\|_{\gamma_2} = 1$. We have

$$\begin{aligned} \|F\|_{\gamma_2} &\leq \|B_1 \circ F_1\|_{\gamma_2} + \|(\mathbf{J} - B_1) \circ F_2\|_{\gamma_2} \\ &\leq \|F_1\|_{\gamma_2} + (\|\mathbf{J}\|_{\gamma_2} + \|B_1\|_{\gamma_2})\|F_2\|_{\gamma_2} \leq 3^{r-1} + 2 \cdot 3^{r-1} = 3^r. \end{aligned} \quad \square$$

In [HHH23], we conjectured that Boolean matrices of small γ_2 -norm are precisely those of the form Equation (6).

Conjecture 3. *Suppose that F is a Boolean matrix with $\|F\|_{\gamma_2} \leq c$. Then we may write*

$$F = \sum_{i=1}^L \pm B_i, \quad (7)$$

where B_i are blocky matrices and $L \leq \ell(c)$ for some integer $\ell(c)$ depending only on c .

Conjecture 3 is inspired by Cohen's idempotent theorem, and it is known to be true for a large class of Boolean matrices, including the XOR-lifts of Boolean functions.

Recall that the XOR-lift of a function $f : \mathbb{Z}_2^n \rightarrow \{0, 1\}$ is the matrix $f^\oplus : \mathbb{Z}_2^n \times \mathbb{Z}_2^n \rightarrow \{0, 1\}$ defined as $f^\oplus(x, y) = f(x + y)$.

The sum of the absolute values of the Fourier coefficients of a function $f : \mathbb{F}_2^n \rightarrow \mathbb{R}$ is called the *Fourier algebra norm* or *spectral norm* of f and is denoted by

$$\|f\|_A := \|\widehat{f}\|_1 = \sum_{\chi \in \widehat{G}} |\widehat{f}(\chi)|.$$

The following identity relating the Fourier algebra norm of f to the γ_2 -norm of its XOR lift is due to [LS09b, Lemma 36].

$$\|f\|_A = \|f^\oplus\|_{\gamma_2}. \quad (8)$$

Therefore, in the case of XOR-lifts, the assumption $\|f^\oplus\|_{\gamma_2} \leq c$ of [Conjecture 3](#) is equivalent $\|f\|_A \leq c$. The structure of Boolean functions $f : \mathbb{Z}_2^n \rightarrow \{0, 1\}$ with $\|f\|_A \leq c$ is characterized by the quantitative version of the so-called Cohen's idempotent theorem for \mathbb{Z}_2^n .

Theorem 3.6 (Quantitative Cohen's theorem for \mathbb{Z}_2^n [Coh60, GS08]). *If $S \subseteq \mathbb{Z}_2^n$ satisfies $\|\mathbf{1}_S\|_A \leq c$, we may write*

$$\mathbf{1}_S = \sum_{i=1}^L \pm \mathbf{1}_{H_i + a_i}, \quad (9)$$

where $H_i + a_i$ are cosets and $L \leq \ell(c)$ for some integer $\ell(c)$ depending only on c .

Note that the XOR-lift of every $\mathbf{1}_{H_i + a_i}$ is a blocky matrix, and therefore, [Theorem 3.6](#) verifies [Conjecture 3](#) for XOR-lifts. In fact, these results extend to every finite group. Given a finite group G , we can generalize the notion of XOR-lifts to *group-lifts*, where we define $F(x, y) = f(y^{-1}x)$ for $f : G \rightarrow \mathbb{C}$. The notion of algebra norm also generalizes to other finite groups. We will not give the original definition of the algebra norm, but for the purposes of this paper, it suffices to know that similar to [Equation \(8\)](#), we have $\|f\|_A = \|F\|_{\gamma_2}$. This identity was observed in [HHH23] based on a result of Davidson and Donsig [DD07].

Theorem 3.7 (Quantitative Cohen's theorem for finite groups [San11]). *There exists a function $\ell : \mathbb{N} \rightarrow \mathbb{N}$ such that the following holds. For every finite group G and every $S \subseteq G$, if $F_{G \times G}(x, y) = \mathbf{1}_S(y^{-1}x)$ satisfies $\|F\|_{\gamma_2} \leq c$, we may write*

$$\mathbf{1}_S = \sum_{i=1}^L \pm \mathbf{1}_{H_i a_i}, \quad (10)$$

where $H_i a_i$ are cosets and $L \leq \ell(c)$. In particular, $F = \sum_{i=1}^L \pm B_i$, where B_i are blocky matrices.

[Theorem 3.7](#) verifies [Conjecture 3](#) for the adjacency matrix of every Cayley (directed) graph. On the other hand, for general Boolean matrices, it is not even known whether $\|F\|_{\gamma_2} \leq c$ implies $\text{rect}(F) \geq \kappa(c)$ for some $\kappa(c) > 0$, which would be an easy consequence of [Conjecture 3](#).

Regarding the bound in [Theorem 3.7](#), for general finite groups, one can take $\ell(c) = A(6, O(c))$, where A is the Ackerman function [San11]. For the case of finite Abelian groups, a better bound of $\ell(c) = 2^{O(c^4 \text{polylog}(c))}$ is due to [San20].

In the special case of \mathbb{Z}_2^n which corresponds to the XOR-lifts, the best bound [San19] that appears in the literature is $\ell(c) \leq 2^{O(c^3 \text{polylog}(c))}$. However, recently, Gowers, Green, Manners, and Tao [GGMT23] announced a proof for Morton's conjecture (aka polynomial Freiman–Ruzsa conjecture). Substituting this result in Sanders proof for [San19, Proposition 2] shows that in the case of \mathbb{Z}_2^n , one may take $\ell(c) = 2^{O(c \text{polylog}(c))}$.

3.2 Equality Oracle Protocols

EQUALITY is the canonical problem with the strongest possible separation between deterministic and randomized communication complexities. We have $D(\text{EQ}_n) = n+1$, which is the largest possible value for any n -bit communication problem. On the other hand, as we will discuss in [Section 5.1](#), the randomized communication complexity of EQ_n is only $O(1)$.

We know that the deterministic model cannot solve EQUALITY efficiently. What if we augment the model with an equality oracle? Does this result in a significantly stronger model? Can this model efficiently solve every problem with small randomized communication complexity?

Formally, in the deterministic communication model with access to an EQUALITY oracle, a protocol for a Boolean matrix $F_{\mathcal{X} \times \mathcal{Y}}$ corresponds to a binary tree. Each non-leaf node v in the tree is labelled with two functions $a_v : \mathcal{X} \rightarrow \{0, 1\}^m$ and $b_v : \mathcal{Y} \rightarrow \{0, 1\}^m$ for some m . On this node, the players map their inputs to strings $a_v(x)$ and $b_v(y)$, respectively, and the oracle will broadcast the value of $\text{EQ}_m(a_v(x), b_v(y))$ to both players. This will contribute only 1 to the cost of the protocol. Note that the oracle queries can simulate sending one-bit messages from each party to the other one. For example, if it is Alice's turn to send a bit a , the query $\text{EQ}_1(a, 1)$ can transmit it to Bob. Hence, in this model, we can assume that all the communication is through oracle queries.

Let $D^{\text{EQ}}(F)$ denote the smallest cost of a deterministic protocol with equality oracle for the matrix F , and define \mathcal{P}^{EQ} and $\mathcal{P}_0^{\text{EQ}}$ to be, respectively, the class of problems with poly-logarithmic and constant communication costs in this model.

In the same way that combinatorial rectangles are the building blocks of deterministic communication protocols, blocky matrices serve as the foundational components of equality oracle protocols. Indeed, every node v of an equality oracle protocol for computing $F(x, y)$ corresponds to $B_v(x, y) = \text{EQ}(a_v(x), b_v(y))$ where B_v is a blocky matrix.

[Equation \(2\)](#) characterizes \mathcal{P}_0 as the set of problems with $O(1)$ rank. Can we obtain a similar characterization for $\mathcal{P}_0^{\text{EQ}}$ via the blocky matrices? To this end, let us define a notion of rank based on blocky matrices.

Definition 3.8 (Blocky Rank). *The blocky rank of a real matrix A , denoted $\text{rk}_{\text{blocky}}(A)$, is the smallest integer r such that A is a real linear combination of r blocky matrices.*

Blocky rank has interesting connections to circuit and communication complexity theory [[AY22](#), [HHH23](#)]. The following proposition shows analogous bounds to [Equation \(5\)](#) on $D^{\text{EQ}}(F)$, and implies that a matrix family $\{F_n\}$ is in $\mathcal{P}_0^{\text{EQ}}$ iff $\text{rk}_{\text{blocky}}(F_n) = O(1)$.

Proposition 3.9 ([[HHH23](#)]). *For every Boolean matrix $F_{\mathcal{X} \times \mathcal{Y}}$, we have*

$$\frac{1}{2} \log \text{rk}_{\text{blocky}}(F) \leq D^{\text{EQ}}(F) \leq \text{rk}_{\text{blocky}}(F)$$

and

$$\log \|F\|_{\gamma_2} \leq 2 \cdot D^{\text{EQ}}(F). \tag{11}$$

Proof. We first prove $D^{\text{EQ}}(F) \leq \text{rk}_{\text{blocky}}(F)$. Let $k = \text{rk}_{\text{blocky}}(F)$. We construct an EQ-oracle protocol for F . In advance, Alice and Bob agree on a decomposition $F = \sum_{i=1}^k \lambda_i B_i$, where B_i is a blocky matrix and $\lambda_i \in \mathbb{R}$ for $i \in [k]$. Since each blocky matrix B_i corresponds to an EQ query, for an input (x, y) , Alice and Bob make k queries to the oracle to determine $F(x, y)$.

For the lower bounds, let $d = D^{\text{EQ}}(M)$. Consider a leaf ℓ in the EQ-oracle protocol tree computing F and let P_ℓ denote the path of length $k_\ell \leq d$ from the root to ℓ . Note that each non-leaf node

v in the tree corresponds to a query to the equality oracle, and each such query corresponds to a blocky matrix B_v . Define $B_v^1 = B_v$ and $B_v^0 = \neg B_v = J_{\mathcal{X} \times \mathcal{Y}} - B_v$.

Suppose $P_\ell = v_1, v_2, \dots, v_{k_\ell}, \ell$, and consider the matrix

$$F_{P_\ell} := B_{v_1}^{\sigma_{v_1}} \circ B_{v_2}^{\sigma_{v_2}} \circ \dots \circ B_{v_{k_\ell}}^{\sigma_{v_{k_\ell}}},$$

where $\sigma_{v_i} \in \{0, 1\}$ and $\sigma_{v_i} = 1$ iff the edge (v_{i-1}, v_i) is labeled by 1. Hence, after simplification, F_{P_ℓ} can be written as a sum of at most 2^d summands with ± 1 coefficients, where each summand is a Schur product of at most k_ℓ blocky matrices. Observe that the Schur product of two blocky matrices is a blocky matrix. Thus, F_{P_ℓ} is a sum of at most 2^d blocky matrices with ± 1 coefficients.

Summing over all the leaves that are labelled by 1, we get $F = \sum_{\ell \text{ is a 1-leaf}} F_{P_\ell}$. As the number of leaves is bounded by 2^d , and each F_{P_ℓ} is a ± 1 linear combination of at most 2^d blocky matrices, we have $\text{rk}_{\text{blocky}}(F) \leq 2^{2d}$ and $\|F\|_{\gamma_2} \leq 2^{2d}$. \square

3.3 Analogue of the log-rank conjecture for blocky rank is false

The log-rank conjecture speculates that the deterministic communication complexity is polynomially equivalent to the logarithm of the rank. In light of [Proposition 3.9](#) it is natural to ask a similar question for D^{Eq} and $\log \text{rk}_{\text{blocky}}$. Arkadev Chattopadhyay³ observed that the recent counter-example to the so-called log-approximate-rank conjecture by Chattopadhyay, Mande, and Sherif [[CMS19](#)] implies that the answer is negative.

Recall that a node in a directed graph is called a sink if all of its adjacent edges are incoming. Define a function $\text{SINK}_m : \{0, 1\}^{\binom{m}{2}} \rightarrow \{0, 1\}$ where the input of length $\binom{m}{2}$ specifies the orientation of the edges of the complete graph on m vertices. The function outputs 1 if there is a vertex that is a sink in the given orientation of edges and 0 otherwise.

Fix m , and for $i \in [m]$, define $\psi_i : \{0, 1\}^{\binom{m}{2}} \rightarrow \{0, 1\}$ to be the indicator function of whether i is a sink in the orientation given by ψ_i . Note that

$$\psi_i(x) = 1 \Leftrightarrow x_{j,i} = 1 \ \forall j \neq i, \tag{12}$$

where $x_{j,i} = 1$ indicates that the edge between i and j is oriented towards i . Since no orientation of the complete graph has more than one sink, we have

$$\text{SINK}_m(x) = \sum_{i=1}^m \psi_i(x).$$

We will consider the family of XOR-lifts of sink functions. Recall that the XOR-lift of a function $f : \{0, 1\}^n \rightarrow \mathbb{R}$ is $f^\oplus : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \mathbb{R}$ with $f^\oplus(x, y) := f(x \oplus y)$. We have

$$\text{SINK}_m^\oplus = \sum_{i=1}^m \psi_i^\oplus.$$

It follows from [Equation \(12\)](#) that each ψ_i^\oplus is a blocky matrix, and therefore $\text{rk}_{\text{blocky}}(\text{SINK}_m^\oplus) \leq m$. On the other hand, Chattopadhyay, Mande, and Sherif [[CMS19](#)] prove that the $\text{R}(\text{SINK}_m^\oplus) = \Theta(m)$. Since R provides a lower bound on D^{Eq} (see [Equation \(15\)](#)), we obtain the following theorem.

Theorem 3.10 (Chattopadhyay, Mande, and Sherif [[CMS19](#)]). *For the family of Boolean matrices $F_m = \text{SINK}_m^\oplus$, we have $\text{D}^{\text{Eq}}(F_m) = \tilde{\Omega}(m)$ and $\text{rk}_{\text{blocky}}(F_m) \leq m$.*

³private communication, no pun intended!

3.4 Blocky matrices and Idempotents of Schur Multipliers

Let \mathcal{X} and \mathcal{Y} be countable sets, and let $B(\mathcal{Y}, \mathcal{X})$ denote the space of bounded linear operators $A : \ell_2(\mathcal{Y}) \rightarrow \ell_2(\mathcal{X})$ endowed with the operator norm:

$$\|A\| = \sup_{x \in \ell_2(\mathcal{Y}): \|x\|_2=1} \|Ax\|_2.$$

A matrix $M_{\mathcal{X} \times \mathcal{Y}}$ is called a *Schur multiplier* if, for every $A \in B(\mathcal{Y}, \mathcal{X})$, we have $M \circ A \in B(\mathcal{Y}, \mathcal{X})$. In other words, $\|M \circ A\| < \infty$ for every $A = A_{\mathcal{X} \times \mathcal{Y}}$ with $\|A\| < \infty$. Note that Schur multipliers form an algebra with addition and Schur product: If M_1 and M_2 are Schur multipliers, then $M_1 + M_2$ and $M_1 \circ M_2$ are both Schur multipliers.

Every Schur multiplier M defines a map $B(\mathcal{Y}, \mathcal{X}) \rightarrow B(\mathcal{Y}, \mathcal{X})$ via $A \mapsto M \circ A$, which assigns an operator norm to it:

$$\|M\|_{\mathbf{m}} := \|M\|_{B(\mathcal{Y}, \mathcal{X}) \rightarrow B(\mathcal{Y}, \mathcal{X})} = \sup_{\substack{A \in B(\mathcal{Y}, \mathcal{X}) \\ \|A\|=1}} \|M \circ A\|.$$

Note that $\|\cdot\|_{\mathbf{m}}$ is an algebra norm as for every M_1 and M_2 , we have

$$\|M_1 \circ M_2\|_{\mathbf{m}} \leq \|M_1\|_{\mathbf{m}} \|M_2\|_{\mathbf{m}}.$$

In other words, the algebra of Schur multipliers endowed with the norm $\|\cdot\|_{\mathbf{m}}$ is a Banach algebra. A classical result, due to Grothendieck, shows that the multiplier norm coincides with the γ_2 -norm.

Proposition 3.11 (See [Pis96, Theorem 5.1]). *For every matrix A , we have $\|A\|_{\mathbf{m}} = \|A\|_{\gamma_2}$.*

An element a of a Banach algebra is said to be an *idempotent* (aka *projection*) if $a^2 = a$. The following question arises naturally.

What are the *idempotents* of the algebra of Schur multipliers?

Every idempotent F of this algebra must satisfy $F = F \circ F$ and, therefore, is a Boolean matrix. However, not every (infinite) Boolean matrix is a bounded Schur multiplier, as it is possible to have $\|F\|_{\mathbf{m}} = \infty$ for a Boolean matrix F . **Proposition 3.4** shows that blocky matrices are precisely the set of all *contractive* idempotents. In other words, an idempotent Schur multiplier satisfies $\|F\|_{\mathbf{m}} \leq 1$ iff it is a blocky matrix.

Question 4. *Are the idempotent Schur multipliers precisely those Boolean matrices that can be written as a ± 1 -linear combination of finitely many contractive idempotents (equivalently, blocky matrices)?*

A simple compactness argument, as outlined in [HHH23], shows that this problem is equivalent to **Conjecture 3**. Therefore, a positive to **Conjecture 3** would characterize idempotents of Schur multipliers, analogous to Cohen's [Coh60] characterization of the idempotents of the Fourier–Stieltjes algebra.

4 Nondeterministic Model and P^{NP}

In a nondeterministic protocol π for a problem $F_{\mathcal{X} \times \mathcal{Y}}$, the parties receive a shared advice string a and use it in a standard deterministic protocol π_a . We say that a protocol computes F if

$$F(x, y) = 1 \Leftrightarrow \exists a, \pi_a(x, y) = 1.$$

The cost of the protocol is the bit-length of a plus the maximum cost of $\pi_a(x, y)$ over all choices of a, x, y . The nondeterministic communication complexity of F , denoted by $\mathsf{N}(F)$, is the minimum cost of such a protocol for F . A matrix family is in the class NP if they have poly-logarithmic nondeterministic communication complexity.

Unlike in the Turing-Machine complexity, in the communication framework, it is known that $\mathsf{P} = \mathsf{NP} \cap \text{coNP}$, which follows from $\mathsf{D}(F) = O(\mathsf{N}(F) \cdot \mathsf{N}(\neg F))$; see [KN97, Theorem 2.11]. However, nondeterministic protocols are provably more powerful than deterministic ones, as can be demonstrated by the important example of the set intersection problem.

The SET-INT problem, INT_n , is defined by $\text{INT}_n(x, y) = 1$ if there exists a coordinate i such that $x_i = y_i = 1$. Since the players can use their nondeterminism to guess the intersecting coordinate i , we have $\mathsf{N}(\text{INT}_n) = O(\log n)$. However, it is easy to see that $\mathsf{D}(\text{INT}_n) = n + 1$. In fact, [BFS86] already in the 1980s proved that SET-INT does not belong to P^{EQ} .

The structural properties of NP : The nondeterministic communication complexity of a problem is fully captured by its *monochromatic rectangle covering number*. Let $\mathsf{C}^1(F)$ denote the minimum number of 1-monochromatic rectangles required to cover the 1 entries of F . It is easy to see [KN97] that

$$\mathsf{N}(F) = \log(\mathsf{C}^1(F)) + O(1). \tag{13}$$

Combined with $\mathsf{D}(F) = O(\mathsf{C}^1(F))$, we have

$$\mathsf{D}(F) \leq O(2^{\mathsf{N}(F)}). \tag{14}$$

Therefore, $\mathsf{P}_0 = \mathsf{NP}_0 = \text{coNP}_0$. The following proposition shows that nondeterministic protocols, while more powerful than deterministic ones, satisfy the same quantitative bound on $\text{wrect}(\cdot)$.

Proposition 4.1. *For every Boolean matrix F ,*

$$\mathsf{N}(F) \gtrsim \log(1/\text{wrect}(F)).$$

Proof. Let $F_{\mathcal{X} \times \mathcal{Y}}$ be a Boolean matrix, and $c = \mathsf{C}^1(F) = O(2^{\mathsf{N}(F)})$. Let $\mu_{\mathcal{X}} \times \mu_{\mathcal{Y}}$ be a product probability measure on $\mathcal{X} \times \mathcal{Y}$, and let $S_1 \times T_1, \dots, S_c \times T_c$ be a 1-monochromatic rectangle covering of F . The case $c \leq 1$ is trivial, so assume $c > 1$.

If there exists i with $\mu_{\mathcal{X}}(S_i) \cdot \mu_{\mathcal{Y}}(T_i) \geq 1/4c^2$, then we are done. So, assume otherwise that for every i , we have $\mu_{\mathcal{X}}(S_i) \cdot \mu_{\mathcal{Y}}(T_i) < 1/4c^2$. Let I to be the set of indices i such that $\mu_{\mathcal{X}}(S_i) < 1/2c$. Note that, if $i \notin I$, then $\mu_{\mathcal{Y}}(T_i) < 1/2c$. Now define, $A = \mathcal{X} \setminus \cup_{i \in I} S_i$ and $B = \mathcal{Y} \setminus \cup_{j \notin I} T_j$. It is easy to see that $A \times B$ is a 0-monochromatic rectangle of F and $\mu_{\mathcal{X}}(A) \cdot \mu_{\mathcal{Y}}(B) > 1/4$. \square

Impagliazzo and Williams [IW10] extended the bound in Proposition 4.1 to the more powerful model of deterministic communication with access to NP oracles. Let us first define this model formally.

An *oracle communication protocol* for a communication problem F is a protocol where each node v is either a regular communication node or it is labelled with a triple (P_v, a_v, b_v) where P_v is Boolean matrix, and $P_v(a_v(x), b_v(y))$ is used to decide whether to travel to the left or the right child of v .

The complexity class P^{NP} . The D^{NP} cost of an oracle communication protocol is the largest cost of a path from the root to a leaf, which is the sum of the communicated bits plus the sum of $\mathsf{N}(P_v)$ for every v on the path.

Define $\mathsf{D}^{\text{NP}}(F)$ to be the smallest D^{NP} cost of an oracle communication protocol for F . The complexity class P^{NP} is the class of problems $\{F_n\}$ with $\mathsf{D}^{\text{NP}}(F_n) = \text{polylog } n$.

Among the extensive list of complexity classes detailed in Göös, Pitassi, and Watson’s article [GPW18b], titled “*the landscape of communication complexity classes*”, P^{NP} is the largest non-probabilistic class for which an explicit lower bound is known. For example, consider the inner product problem $\text{IP}_n : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ defined as $\text{IP}_n(x, y) = x_1y_1 + \dots + x_ny_n \pmod 2$. A simple argument, based on dimension and orthogonality (see [RY20, Claim 1.17]), shows that every monochromatic rectangle in IP_n is of size at most 2^n , and therefore, $\text{wrect}(\text{IP}_n)^{-1} \geq 2^{\Omega(n)}$. The following theorem of [IW10] shows that $\mathsf{D}^{\text{NP}}(\text{IP}_n) = \Omega(n)$.

Theorem 4.2 (Impagliazzo and Williams [IW10]). *For every Boolean matrix F , we have*

$$\mathsf{D}^{\text{NP}}(F) \gtrsim \log(\text{wrect}(F)^{-1}).$$

One might ask whether $\log(\text{wrect}(F)^{-1})$ and $\mathsf{D}^{\text{NP}}(F)$ are polynomially equivalent. The answer is negative as [GKPW17] constructs an explicit family of Boolean matrices exhibiting a large gap between the two quantities.

Theorem 4.3 (Göös, Kamath, Pitassi, and Watson [GKPW17]). *There exists a sequence of $2^n \times 2^n$ Boolean matrices F_n satisfying $\mathsf{D}^{\text{NP}}(F_n) \geq n^{\Omega(1)}$ and $\log(\text{wrect}(F_n)^{-1}) \leq \log^{O(1)}(n)$.*

5 Probabilistic Communication Models

Next, we discuss probabilistic communication protocols where the players can act in a randomized fashion. Randomness can be introduced in two different ways: private randomness and public randomness.

In a *private-coin randomized protocol*, each player has access to their own independent random bits and can use them to decide which bit to send next. More precisely, Alice and Bob have access to random strings R_A and R_B , respectively. These two strings are chosen independently, each according to some probability distribution described by the protocol. The bit sent by Alice at a node v is determined by a function a_v of both x and R_A . Similarly, the bits sent by Bob are determined by functions of y and R_B .

In the *public-coin* model, the players have access to a shared source of randomness. In other words, Alice and Bob both receive the same random string R . The public-coin model is stronger than the private-coin model as the former can simulate the latter by setting $R = (R_A, R_B)$.

The cost of a randomized protocol is the maximum number of communicated bits over all inputs and all choices of random strings. A probabilistic protocol is allowed to make errors. It is common to consider three types of errors:

- **Two-sided error (BPP):** For every x, y , the probability that the protocol makes an error on (x, y) is at most ϵ for some $\epsilon < 1/2$. When ϵ is a fixed constant strictly less than $1/2$, the protocol is called a bounded-error protocol. The particular choice of ϵ is unimportant as a simple error reduction shows that it affects the complexity by only a constant factor. Therefore, as it is common, we will fix the error parameter to $\epsilon = 1/3$.
- **One-sided error (RP):** In this setting, the protocol can only make an error if $F(x, y) = 1$. In other words, for every x, y with $F(x, y) = 0$, the protocol must always correctly output 0, but for every x, y with $F(x, y) = 1$, it might output a wrong answer with probability at most ϵ for some fixed $\epsilon < 1$. We will fix the error parameter to $\epsilon = 1/3$.
- **Zero-error (ZPP):** In this case, the output of a protocol is 0, 1, or \perp , where \perp indicates a failure to compute $F(x, y)$. The protocol must never output 0 or 1 erroneously; however, on every input, it is allowed to output \perp with probability at most $\frac{1}{2}$.

A classical result in communication complexity, called Newman’s lemma, states that in the two-sided error, one-sided error, and zero-error settings, the following is true. The difference between public-coin and private-coin randomized communication complexities of any n -bit communication problem is $O(\log(n))$.

Newman’s lemma shows that when defining the poly-logarithmic communication complexity classes BPP, RP, ZPP, it is unimportant whether we use shared randomness or private randomness. However, to define constant-cost classes $\text{BPP}_0, \text{RP}_0, \text{ZPP}_0$, we need to make a choice. It turns out that in the setting of private-coin, all these classes collapse to P_0 . Therefore, we shall define these classes in the public-coin model.

Regarding the zero-error protocols, the following theorem shows that even in the public-coin model, the zero-error randomized communication complexity of a matrix F is polynomially equivalent to $\text{D}(F)$ and in particular $\text{ZPP}_0 = \text{P}_0$.

Theorem 5.1 ([DHP⁺22, Theorem 2.1]). *The public-coin zero-error randomized communication complexity of every Boolean matrix F is at least $\Omega(\text{D}(F)^{1/4})$.*

It is an open problem whether this bound can be improved to $\Omega(\sqrt{\text{D}(F)})$, which, if true, would be sharp. We will not further discuss the zero-error model and refer the interested reader to [DHP⁺22] for further reading.

5.1 The power of randomness: BPP

The *randomized communication complexity* of a Boolean matrix F , denoted by $\text{R}(F)$, is the minimum cost of a public-coin randomized protocol with two-sided error $\leq 1/3$. Let BPP and BPP_0 be, respectively, the class of problems with poly-logarithmic and $O(1)$ randomized communication complexities.

Are probabilistic protocols more powerful than deterministic protocols? The example of EQUALITY shows that randomness can provide a significant advantage. To test whether $x \neq y$, Alice and Bob can use their shared randomness to jointly sample a random subset $S \subseteq \{0, 1\}^n$ at no cost and then, by exchanging two bits of information, indicate to each other whether their inputs belong to S . If they see a disparity, they can conclude confidently that $x \neq y$. They can run this test twice, and if they do not detect $x \neq y$, they declare $x = y$. Note that the probability of error is $\leq 1/4$.

Therefore, $R(\text{EQ}_n) = O(1)$, and $\text{EQUALITY} \in \text{BPP}_0$. In particular, we have the relations $\text{P} \subsetneq \text{BPP}$ and $\text{BPP}_0 \not\subseteq \text{P}$. Also note that $R(\text{EQ}_n) = O(1)$ implies via standard error-reduction that

$$R(F) \lesssim D^{\text{Eq}}(F) \log D^{\text{Eq}}(F), \quad (15)$$

which establishes $\text{P}^{\text{Eq}} \subseteq \text{BPP}$ and $\text{P}_0^{\text{Eq}} \subseteq \text{BPP}_0$.

Which problems have efficient randomized protocols? A substantial portion of the literature in communication complexity is dedicated to lower-bound techniques against randomized communication complexity, and many celebrated results establish such lower bounds for important concrete problems, such as Set Disjointness [Raz92], Gap Hamming Distance [CR12], and Halfspace [She08]. While it is possible to write a voluminous book about the lower bounds against randomized communication complexity, we know very little about what is inside BPP and BPP_0 . In fact, until recently, it was not known whether there is any problem in BPP that is not in P^{Eq} . Let us list some classical problems in BPP .

- **GREATER-THAN** is the family of communication problems $\text{GT}_n : [2^n] \times [2^n] \rightarrow \{0, 1\}$ where $\text{GT}_n(x, y) = 1$ iff $x \leq y$. It is known [Nis93, Vio15, RS15] that $R(\text{GT}_n) = \Theta(\log(n))$, and therefore,

$$\text{GREATER-THAN} \in \text{BPP} \setminus \text{BPP}_0.$$

- **HYPERCUBE** is the family of communication problems $\text{Q}_n : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ where $\text{Q}_n(x, y) = 1$ iff x and y differ in exactly one coordinate. Given $x, y \in \{0, 1\}^n$, Alice and Bob can pick a uniform partition of $[n]$ into 8 sets S_1, \dots, S_8 and accept if for exactly one $i \in [8]$, it holds that $(\bigoplus_{j \in S_i} x_j) \oplus (\bigoplus_{j \in S_i} y_j) = 1$. It is easy to see that the communication cost of this protocol is constant and that the error probability is at most $1/3$. Therefore, $R(\text{Q}_n) = O(1)$. On the other hand, [HHH23, Lemma 2.15] and Equation (8) shows that $\|\text{Q}_n\|_{\gamma_2} \geq \Omega(\sqrt{n})$ and therefore, $D^{\text{Eq}}(\text{Q}_n) \geq \Omega(\log n)$. We have

$$\text{HYPERCUBE} \in \text{BPP}_0 \setminus \text{P}_0^{\text{Eq}};$$

see also [HWZ22] for a different proof of this fact.

- More generally, let $\ell(n) < n/2$ be an integer, and $S_n \subseteq \{0, \dots, \ell(n)\} \cup \{n - \ell(n), \dots, n\}$ and denote the hamming weight of an $x \in \{0, 1\}^n$ by $|x|$. If $\ell(n) = \text{polylog}(n)$, then the family of the XOR-lifts $\mathbf{1}_{S_n}^{\oplus} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ defined as $\mathbf{1}_{S_n}^{\oplus}(x, y) = \mathbf{1}_{S_n}(|x \oplus y|)$ is in BPP . If $\ell(n) = O(1)$, then this family is in BPP_0 [Yao03]. Note that **HYPERCUBE** corresponds to $S_n = \{1\}$.
- **INTEGER INNER PRODUCT**: Given a fixed positive integer t , the communication problem IIP_t is the family of functions $\text{IIP}_{t,n} : [-2^n, 2^n]^t \times [-2^n, 2^n]^t \rightarrow \{0, 1\}$ with $\text{IIP}_{t,n}([x_1, \dots, x_t], [y_1, \dots, y_t]) = 1$ iff $x_1 y_1 + \dots + x_t y_t = 0$. To check the validity of the equation, Alice and Bob can choose a random prime $p \approx \log(n)$ and exchange $x_i \bmod p$ and $y_i \bmod p$ for $i = 1, \dots, t$. This leads to a randomized protocol for IIP_t with cost $O(\log(n))$, which shows $\text{IIP}_t \in \text{BPP}$. On the other hand, for any fixed $t > 2$, it was shown in [CLV19] that $D^{\text{Eq}}(\text{IIP}_{t,n}) \geq \Omega(n)$. In fact, as [CHHS23] shows, one even has $\|\text{IIP}_{t,n}\|_{\gamma_2} \geq 2^{\Omega(n)}$. Hence, for $t > 2$,

$$\text{IIP}_t \in \text{BPP} \setminus \text{P}^{\text{Eq}}.$$

We are unaware of any examples in BPP that fundamentally differ from those listed above. In fact, IIP_t , which was introduced by Chattopadhyay, Lovett, and Vinyals [CLV19], is the only known example of a communication problem in BPP that is not in P^{EQ} (see also [PSS23, CHHS23]). Let us mention a conjecture about IIP_t before proceeding further. We do not know how to prove any $\omega(1)$ lower bound for $\text{R}(\text{IIP}_{t,n})$.

Conjecture 5 (See [CHHS23, Conjecture 6.4]). *For $t > 2$, $\text{IIP}_t \notin \text{BPP}_0$.*

Note that disproving [Conjecture 5](#) would imply that $\text{BPP}_0 \not\subseteq \text{P}^{\text{EQ}}$.

The communication problems in P^{EQ} are highly structured as they are linear combinations of a few blocky matrices. On the other hand, the only known example in $\text{BPP} \setminus \text{P}^{\text{EQ}}$ is the INTEGER INNER PRODUCT, which has a low-dimensional geometric representation (i.e. bounded sign-rank) and enjoys nice structural properties. All these known examples contain large monochromatic rectangles. Does every Boolean matrix with an efficient randomized communication protocol contain a large monochromatic rectangle? More specifically, Gös, Kamath, Pitassi, and Watson [GKPW17] asked the following question.

Question 6. *Is it the case that for every family of n -bit communication problems F_n in BPP, there exists $c > 0$ such that $\text{rect}(F_n) \geq 2^{-c(\log n)^c}$?*

By [Theorem 4.2](#), a negative answer to [Question 6](#) would imply that $\text{BPP} \not\subseteq \text{P}^{\text{NP}}$, a relation that remains unknown.

Conjecture 7. $\text{BPP} \not\subseteq \text{P}^{\text{NP}}$.

Independently from [GKPW17], and also motivated by [Conjecture 7](#), [CLV19] asked whether there exists a $c > 0$ such that every communication problem F satisfies $\text{rect}(F) \geq 2^{-c\text{R}(F)^c}$. In fact, we do not know whether there is a uniform lower bound on $\text{rect}(F)$ depending only on $\text{R}(F)$.

Question 8 ([CLV19, HHH23]). *Is there a function $\kappa : \mathbb{N} \rightarrow (0, 1)$ such that $\text{rect}(F) \geq \kappa(\text{R}(F))$?*

As we shall see in [Theorem 5.6](#), a negative answer to [Question 8](#) would imply that $\text{BPP}_0 \not\subseteq \text{UPP}_0$, which is currently unknown.

5.2 Is two-sided error necessary?

Define the one-sided randomized communication complexity $\text{R}^1(F)$ and its corresponding complexity classes RP and RP_0 analogous to the two-sided error counterparts $\text{R}(F)$, BPP , and BPP_0 .

Are two-sided error protocols genuinely more powerful than one-sided error protocols?

One could give an affirmative answer to this question by referring to EQUALITY, which satisfies $\text{R}(\text{EQ}_n) = O(1)$ while $\text{R}^1(\text{EQ}_n) = \Omega(n)$. This, however, is not a fully satisfactory separation. Indeed, since $\text{R}^1(\neg\text{EQ}_n) = O(1)$, we can solve EQUALITY with a single oracle query to the NONEQUALITY problem which belongs to RP_0 . In other words, $\text{EQUALITY} \in \text{coRP}_0$.

If we examine all the known examples in BPP, we realize that they all essentially boil down to solving problems with one-sided error in the sense that they are either in $\text{RP} \cup \text{coRP}$, or they are composed of a few components, each belonging to $\text{RP} \cup \text{coRP}$. We find this surprising, as we are not aware of any evident reasons as to why a two-sided error protocol might be simulated by a series

of steps that can be performed by efficient one-sided error protocols. We suspect this phenomenon to be due to our limited knowledge of examples in BPP.

Define the class P^{RP} similarly to P^{NP} , except that the protocol is now charged $R^1(P_v)$ for its oracle queries P_v at a node v . The simple inclusions of $\mathsf{RP} \subseteq \mathsf{BPP}$, $\mathsf{P}^{\mathsf{RP}} \subseteq \mathsf{BPP}$ and $\mathsf{P}_0^{\mathsf{RP}} \subseteq \mathsf{BPP}_0$ are immediate from the definitions.

Question 9. *Is it true that $\mathsf{BPP} = \mathsf{P}^{\mathsf{RP}}$?*

It is known that nondeterministic protocols can simulate one-sided protocols with a logarithmic loss,

$$N(F) \leq R^1(F) + O(\log n). \quad (16)$$

This shows that $\mathsf{RP} \subseteq \mathsf{NP}$ and $\mathsf{P}^{\mathsf{RP}} \subseteq \mathsf{P}^{\mathsf{NP}}$. In particular, [Conjecture 7](#) would imply a negative answer to [Question 9](#).

It is interesting to ask the above questions in the constant-cost setting.

Question 10. *Is it true that $\mathsf{BPP}_0 = \mathsf{P}_0^{\mathsf{RP}}$?*

[Theorem 5.2](#) below implies that for every communication problem in $\mathsf{P}_0^{\mathsf{RP}}$, we have $\text{wrect}(F) = \Omega(1)$. In particular, a negative answer to [Question 8](#) would imply $\mathsf{BPP}_0 \subsetneq \mathsf{P}_0^{\mathsf{RP}}$.

Theorem 5.2 ([\[HHH23, Theorem 3.8\]](#)). *For every communication problem F ,*

$$\text{wrect}(F) \geq 2^{-O(R^1(F))}.$$

5.3 Sign-rank and UPP

The *unbounded-error communication complexity* of F , denoted by $U(F)$, is the smallest communication cost of a *private-coin* randomized protocol π that satisfies

$$\Pr[\pi(x, y) \neq f(x, y)] < \frac{1}{2} \quad \forall x, y.$$

In other words, the protocol is only required to outperform a random guess. The complexity classes corresponding to this measure are UPP and UPP₀.

It is crucial that in this communication model, the players have only access to private randomness. Otherwise, given access to shared randomness, they could jointly sample a random input (x_0, y_0) at no cost and use two bits of communication to verify whether $(x, y) = (x_0, y_0)$. If $(x, y) = (x_0, y_0)$, then they know the output $F(x, y)$, and if it is not, they can output a random bit. This protocol has an error probability strictly less than $1/2$.

Paturi and Simon [\[PS86\]](#) proved that the unbounded-error communication complexity is precisely determined by an elegant matrix parameter called *sign-rank*.

To discuss sign-rank, it is more convenient to switch from Boolean matrices to *sign matrices*, which are matrices with ± 1 entries. The *sign-rank* $\text{rk}_{\pm}(F)$ of a sign matrix $F_{\mathcal{X} \times \mathcal{Y}}$ is the smallest rank of a real matrix $A_{\mathcal{X} \times \mathcal{Y}}$ such that the entries of A are nonzero and have the same signs as their corresponding entries in F . Geometrically, sign-rank corresponds to the smallest dimension where we can represent F as points and homogeneous half-spaces.

We can reformulate the definition of sign-rank as follows.

Definition 5.3 (Sign-rank). *The sign-rank of a sign matrix $F_{\mathcal{X} \times \mathcal{Y}}$ is the smallest d such that there exist vectors $u_x, v_y \in \mathbb{R}^d$ with $F(x, y) = \text{sgn}(\langle u_x, v_y \rangle)$ for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$.*

Recall that the log-rank conjecture speculates that for deterministic protocols, the communication complexity is polynomially related to the logarithm of the rank of the corresponding matrix. Paturi and Simon proved that a similar and tighter connection is true for unbounded-error protocols, except that rank is replaced by sign-rank.

Theorem 5.4 (Paturi and Simon [PS86]). *For every sign-matrix F , we have*

$$U(F) = \log \text{rk}_{\pm}(F) \pm O(1).$$

In light of [Theorem 5.4](#), to study $U(F)$, one can set aside the intricacies of communication and focus on the geometric notion of sign-rank.

Number of matrices of small sign-rank: Shortly after the introduction of sign-rank in [PS86], Alon, Frankl, and Rödl [AFR85] used results of [Mil64, Tho65, War68] on the number of connected components of real algebraic varieties and obtained a linear lower bound on the sign-rank of random sign matrices. This argument was later refined in [AMY16, Lemma 24] to the following bound on the number of low sign-rank matrices.

Lemma 5.5 (See [AMY16, Lemma 24]). *For $d \leq \frac{m}{2}$, the number of $m \times m$ sign matrices of sign-rank at most d does not exceed $(O(m/d))^{2dm} \leq 2^{O(dm \log(m))}$.*

[Lemma 5.5](#) shows that there are very few matrices with small sign-rank and that a typical $m \times m$ sign matrix has sign-rank $\Omega(m)$. This scarcity of small-sign-rank matrices suggests that they might possess strong structural properties.

Large monochromatic rectangles: [APP⁺05] used the geometric properties of sign-rank to prove that every $\mathcal{X} \times \mathcal{Y}$ sign matrix of sign-rank d contains an $\frac{|\mathcal{X}|}{2^{d+1}} \times \frac{|\mathcal{Y}|}{2^{d+1}}$ monochromatic rectangle. Their result uses a theorem of Yao and Yao [YY85], which is based on the Borsuk-Ulam theorem, a result in topology. Slightly different bounds are also obtained in [FPS16] using the cutting lemma of Chazelle [Cha93]. In our notation, we have the following relation between sign-rank and $\text{wrect}(F)$.

Theorem 5.6 (See [APP⁺05, Theorem 1.3]). *For every sign-matrix F , we have*

$$\text{rk}_{\pm}(F) \gtrsim \log(\text{wrect}(F)^{-1}). \tag{17}$$

On the other hand, [HHP⁺22] used a counting argument to show that there are matrices with $\text{wrect}(A)^{-1} = O(1)$ and very large sign-rank.

Theorem 5.7 ([HHP⁺22, Theorem 3.2]). *There exists $m \times m$ sign matrices A such that*

$$\text{wrect}(A)^{-1} \leq 2^{15}, \quad \text{while} \quad \text{rk}_{\pm}(A) = \Omega\left(\frac{m^{1/3}}{\log(m)}\right).$$

It is known that $\text{P}^{\text{NP}} \not\subseteq \text{UPP}$ (see [GPW18b]). Therefore, [Theorem 5.7](#) is stronger than the separation of [Theorem 4.3](#), as it shows the existence of communication problems with $\text{wrect}(F)^{-1} = O(1)$ that are not in UPP. This answers an open problem by Göös, Kamath, Pitassi, and Watason [GKPW17].

The caveat of [Theorem 5.7](#) is that its existential proof does not provide any explicit construction. In fact, regarding explicit examples that separate sign-rank and $\text{wrect}(\cdot)^{-1}$, our knowledge is embarrassingly limited. The following problem is open.

Problem 11. Construct an explicit sequence of matrices F_n such that $\text{wrect}(F_n)^{-1} = O(1)$ and

$$\lim_{n \rightarrow \infty} \text{rk}_{\pm}(F_n) = \infty.$$

By [Theorem 5.7](#), we know such matrices exist, and in fact, with very large sign-ranks. On the other hand, none of the known lower bound techniques can directly imply a solution to [Problem 11](#). Indeed, in addition to the monochromatic rectangle lower bound, there are only two other known methods for proving lower bounds on the sign-rank of explicit matrices: (i) Sign-rank is at least the VC-dimension: $\text{rk}_{\pm}(A) \geq \text{VC}(A)$; (ii) Forster’s method, which states that sign-rank is at least the inverse of the largest possible average margin among the representations of the matrix by points and half-spaces: $\text{rk}_{\pm}(F) \geq m^{\text{avg}}(F)^{-1}$. We refer the reader to [\[HHP⁺22\]](#) for the definition of average margin and a thorough discussion of these facts.

Qualitatively, [Equation \(17\)](#) is the strongest known method for proving lower bounds on the sign-rank of an explicit matrix. If it fails to provide a super-constant lower bound for the sign-rank of a matrix, then the other two methods will also fail. More precisely, we have

$$\sqrt{\text{VC}(A)} \leq m^{\text{avg}}(A)^{-1} \leq \text{wrect}^{-1}(A). \tag{18}$$

In this sense, [Problem 11](#) captures the limitation of the currently known lower bound techniques for sign-rank.

Sign-rank of hypercubes and BPP_0 vs UPP_0 : Linial, Mendelson, Schechtman, and Shraibman [\[LMSS07\]](#) asked whether sign-rank can be bounded from above by a function of the so-called margin complexity. The relation between margin complexity, discrepancy, and randomized communication complexity, which were discovered later, allows us to rephrase their question as follows.

Question 12. *Is it true that $\text{BPP}_0 \subseteq \text{UPP}_0$? Equivalently, is it possible to upper bound $\text{rk}_{\pm}(F)$ by a function of $R(F)$?*

We believe the answer to this question to be negative. Consider the sign version of the HYPERCUBE problem, that is, let \mathbf{Q}_n be the sign matrix whose rows and columns are indexed with the elements of $\{0, 1\}^n$, and $\mathbf{Q}_n(x, y) = -1$ if x and y differ in exactly one coordinate. As we discussed earlier, $R(\mathbf{Q}_n) = O(1)$. We conjecture that the sign-rank of \mathbf{Q}_n tends to infinity as n grows, which, if true, would imply $\text{BPP}_0 \not\subseteq \text{UPP}_0$.

Conjecture 13 (Sign-rank of HYPERCUBE [\[HHP⁺22\]](#)). *We have*

$$\lim_{n \rightarrow \infty} \text{rk}_{\pm}(\mathbf{Q}_n) = \infty.$$

It is worth pointing out that proving [Conjecture 13](#) in the positive would likely require some new lower bound techniques, as $\text{wrect}(\mathbf{Q}_n)^{-1} = O(1)$, shown in [\[HHP⁺22\]](#). Note that, a positive answer to [Conjecture 13](#) would also solve [Problem 11](#).

Equality oracles, $\text{P}_0^{\text{EQ}} \subsetneq \text{UPP}_0$: It is easy to show that the $\text{EQ}_n \in \text{UPP}_0$, as its sign-rank is 3. The following theorem shows that, in fact, $\text{P}_0^{\text{EQ}} \subseteq \text{UPP}_0$. The example of GREATER-THAN shows that this inclusion is strict. It is easy to see that $\text{GREATER-THAN} \in \text{UPP}_0$ as its sign-rank is 2. This fact combined with $R(\text{GT}_n) = \Theta(\log n)$ shows that

$$\text{GREATER-THAN} \in \text{UPP}_0 \setminus \text{BPP}_0^{\text{EQ}} \subseteq \text{UPP}_0 \setminus \text{P}_0^{\text{EQ}}.$$

Theorem 5.8 ([HHP+22]). *For every sign matrix $F_{\mathcal{X} \times \mathcal{Y}}$, we have $\text{rk}_{\pm}(F) \leq 4^{\text{D}^{\text{EQ}}(F)}$. In particular, $\text{U}(F) \leq 2\text{D}^{\text{EQ}}(F) + O(1)$.*

Proof. We proceed by induction on $d := \text{D}^{\text{EQ}}(A)$. When $d = 1$, F corresponds to a blocky matrix, which in fact has $\text{rk}_{\pm}(F) \leq 3$. For larger d , consider a cost d protocol for a sign matrix F and suppose the equality query at the root of the tree is $\text{EQ}(a(x), b(y))$, where here we assume without loss of generality $a(x)$ and $b(y)$ take integer values. Let $S_{\mathcal{X} \times \mathcal{Y}}$ be the matrix with entries $S_{xy} = \mathbf{1}_{a(x)=b(y)}$. We branch according to the output of the first query either to the left or the right subtree of the root, each corresponding to a protocol with cost at most $d-1$. Let the corresponding matrices for these protocols be Π_1 and Π_2 , and note that

$$F = S \circ \Pi_1 + (J - S) \circ \Pi_2,$$

where $J := J_{\mathcal{X} \times \mathcal{Y}}$ is the all-ones matrix. By the induction hypothesis, Π_1 and Π_2 have sign-rank at most $\leq 4^{d-1}$. Let $\tilde{\Pi}_1$ and $\tilde{\Pi}_2$ be real matrices with rank at most 4^{d-1} that satisfy $\text{sgn}(\tilde{\Pi}_1) = \Pi_1$ and $\text{sgn}(\tilde{\Pi}_2) = \Pi_2$. Let $E_{\mathcal{X} \times \mathcal{Y}}$ be the rank-3 matrix with entries $E_{xy} = (a(x) - b(y))^2$. Note that for a sufficiently large k , we have

$$A = \text{sgn}(\tilde{\Pi}_1 + kE \circ \tilde{\Pi}_2).$$

Finally, we have

$$\text{rk}(\tilde{\Pi}_1 + k\tilde{\Pi}_2 \circ E) \leq \text{rk}(\tilde{\Pi}_1) + \text{rk}(\tilde{\Pi}_2) \cdot \text{rk}(E) \leq 4^{d-1} + 3 \cdot 4^{d-1} = 4^d. \quad \square$$

The above theorem combined with [Equation \(15\)](#) shows that $\text{P}_0^{\text{EQ}} \subseteq \text{UPP}_0 \cap \text{BPP}_0$. Both the inclusions $\text{P}_0^{\text{EQ}} \subsetneq \text{UPP}_0$ and $\text{P}_0^{\text{EQ}} \subsetneq \text{BPP}_0$ are strict; the former follows from the example of GREATER-THAN and the latter holds for HYPERCUBE [[HHH23](#), [HWZ22](#)]. The question of whether these separations can be obtained simultaneously was asked recently by [[HZ24](#)].

Question 14 ([[HZ24](#)]). *Is it the case that $\text{P}_0^{\text{EQ}} = \text{UPP}_0 \cap \text{BPP}_0$?*

Since $\text{HYPERCUBE} \notin \text{P}_0^{\text{EQ}}$ and $\text{HYPERCUBE} \in \text{BPP}_0$, a positive answer to [Question 14](#) would imply $\text{HYPERCUBE} \notin \text{UPP}_0$ and solve [Conjecture 13](#).

In the converse direction, [Conjecture 13](#) would imply a positive answer to [Question 14](#) for the special case of XOR-lifts. Indeed, if $\text{HYPERCUBE} \notin \text{UPP}_0$, then the result of [[CHZZ22](#)] would imply that every family of XOR-lift $f_n^{\oplus} \in \text{UPP}_0 \cap \text{BPP}_0$ must satisfy $\|f_n\|_A = O(1)$ and therefore by [Theorem 3.7](#) and [Proposition 3.9](#), we must have $f_n^{\oplus} \in \text{P}_0^{\text{EQ}}$.

5.4 Weakly unbounded-error complexity, PP

The *weakly unbounded-error communication complexity* of a problem $F_{\mathcal{X} \times \mathcal{Y}}$ is defined as

$$\text{PP}(F) := \min_{\epsilon < 1/2} \text{R}_{\epsilon}(F) + \log \left(\frac{1}{\frac{1}{2} - \epsilon} \right),$$

where $\text{R}_{\epsilon}(F)$ is the minimum cost of a public-coin randomized protocol with two-sided error at most ϵ .

Define PP and PP_0 to be the class of families of n -bit problems F_n with $\text{PP}(F_n) = \text{polylog } n$ and $\text{PP}(F_n) = O(1)$, respectively. It is immediate from the definitions and Newman’s lemma [New91] that

$$\text{BPP} \subseteq \text{PP} \subseteq \text{UPP}, \quad \text{and} \quad \text{PP}_0 = \text{BPP}_0.$$

As discussed before, the GREATER-THAN problem separates UPP_0 from $\text{BPP}_0 = \text{PP}_0$. Babai, Frankl, and Simon [BFS86] asked whether $\text{UPP} = \text{PP}$. Their question remained unanswered for over two decades, until [BVdW07, She08] independently showed that there are $2^n \times 2^n$ sign matrices F with $\text{rk}_\pm(F) = n$ but $\text{PP}(F) = n^{\Omega(1)}$. The separation was strengthened in subsequent works to $\text{rk}_\pm(F) = n$ and $\text{PP}(F) = \Omega(n)$ in [She21].

The example proposed in [BVdW07], in fact, belongs to P^{NP} , and therefore also shows that $\text{P}^{\text{NP}} \not\subseteq \text{PP}$. It turns out that the opposite direction of this inclusion is not true either, as follows from an argument involving bounds on the rectangle ratio.

Theorem 5.9. *There exists a family of n -bit communication problems F_n with $\text{PP}(F_n) = O(\log(n))$ and $\text{rect}(F_n)^{-1} = 2^{\Omega(n)}$. In particular, $\text{D}^{\text{NP}}(F_n) = \Omega(n)$ and $\text{PP} \not\subseteq \text{P}^{\text{NP}}$.*

Proof. Let $F_n(x, y) = 1$ iff x and y differ on at least $n/2$ bits. It is known through classical results from combinatorics [FF81] that $\text{rect}(F_n)^{-1} = 2^{\Omega(n)}$. Thus by Theorem 4.2, we get $\text{D}^{\text{NP}}(F_n) = \Omega(n)$.

To see the inclusion in PP , note that the simple protocol where two parties pick a random index i uniformly at random and output 1 iff $x_i \neq y_i$, has cost $O(\log n)$. \square

Finally, recent works have shown that PP does not even contain UPP_0 . Indeed, [HHL20, ACHS23] gave simple constructions of n -bit communication problems F with $\text{rk}_\pm(F) = 3$ and $\text{PP}(F) = \Omega(n)$.

6 Final remarks

We discussed several open problems that indicate significant gaps in our understanding of communication complexity and capture the limitations of the currently available techniques. For example, disproving $\text{BPP} = \text{P}^{\text{RP}}$, $\text{BPP} \subseteq \text{P}^{\text{NP}}$, or giving a negative answer to Question 6 or Question 8, requires constructing a family of matrices in BPP that is fundamentally different from all the currently known examples. Conversely, proving that any of these statements is true would be a major stride toward achieving a structural description of BPP . Similarly, Conjecture 13 and Problem 11 require a new lower-bound technique for sign-rank that can reach beyond the $\log \text{wrect}(\cdot)^{-1}$ bound of Theorem 5.6.

We hope that similar to the introduction of communication classes by [BFS86], the formal paradigm of constant-cost communication classes will catalyze future research—That the efforts to establish separations between these classes will lead to the discovery of new examples and lower-bound techniques and give us a deeper understanding of communication models and their connections to other areas of theoretical computer science.

Due to space limitations, we did not discuss quantum communication models, multi-party models, search problems, and various related query models, most notably parity decision trees. The questions that are being discussed in this article can be asked in a similar way for these models.

We conclude by presenting, in Figure 1 and Figure 2, the known relations and separations among various classes discussed in this article. These figures include a selected list of classes, excluding easier-to-handle classes such as P , RP , NP , $\text{P}_0 = \text{NP}_0$. We define the classes Rect and Rect_0 to consist

of matrix families with $\text{wrect}^{-1}(\cdot)$ bounded from above by $2^{\text{polylog } n}$ and $O(1)$, respectively. The class Rect appears in [GKPW17] with the different name of PM (for Product Method).

	P_0^{EQ}	P_0^{RP}	BPP_0	UPP_0	$Rect_0$	P^{EQ}	P^{RP}	BPP	PP	UPP	P^{NP}	Rect
P_0^{EQ}	=	\subseteq	\subseteq	\subseteq	\subseteq	\subseteq	\subseteq	\subseteq	\subseteq	\subseteq	\subseteq	\subseteq
P_0^{RP}	$\not\subseteq$	=	\subseteq	?	\subseteq	?	\subseteq	\subseteq	\subseteq	\subseteq	\subseteq	\subseteq
BPP_0	$\not\subseteq$?	=	?	?	?	?	\subseteq	\subseteq	\subseteq	?	?
UPP_0	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	=	\subseteq	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	\subseteq	?	\subseteq
$Rect_0$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	=	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	\subseteq
P^{EQ}	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	=	\subseteq	\subseteq	\subseteq	\subseteq	\subseteq	\subseteq
P^{RP}	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	=	\subseteq	\subseteq	\subseteq	\subseteq	\subseteq
BPP	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$?	=	\subseteq	\subseteq	?	?
PP	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	=	\subseteq	$\not\subseteq$	$\not\subseteq$
UPP	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	=	$\not\subseteq$	$\not\subseteq$
P^{NP}	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	\subseteq	=	\subseteq
Rect	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	$\not\subseteq$	=

Figure 1: The entry at a row A and a column B indicates whether $A \subseteq B$ or $A \not\subseteq B$. A question mark indicates that the relationship is unknown. The separations in grey entries follow trivially via padding.




Figure 2: $A \rightarrow B$ indicates $A \subseteq B$.

Acknowledgements. We wish to thank Tsun-Ming Cheung, Mika Göös, Lianna Hambardzumyan, Nathan Harms, Shachar Lovett, and Alexander Razborov for their helpful comments on a draft of this article.

References

- [ACHS23] Manasseh Ahmed, Tsun-Ming Cheung, Hamed Hatami, and Kusha Sareen. Communication complexity of half-plane membership. In *Electron. Colloquium Comput. Complex.*, 2023.
- [AFR85] Noga Alon, Peter Frankl, and Vojtech Rödl. Geometrical realization of set systems and probabilistic communication complexity. In *26th Annual Symposium on Foundations of Computer Science, FOCS 1985*, pages 277–280. IEEE Computer Society, 1985.
- [AMY16] Noga Alon, Shay Moran, and Amir Yehudayoff. Sign rank versus VC dimension. In *Proceedings of the 29th Conference on Learning Theory, COLT 2016*, volume 49, pages 47–80, 2016.
- [APP⁺05] Noga Alon, János Pach, Rom Pinchasi, Radoš Radoičić, and Micha Sharir. Crossing patterns of semi-algebraic sets. *Journal of Combinatorial Theory, Series A*, 111(2):310–326, 2005.
- [AY22] Daniel Avraham and Amir Yehudayoff. On blocky ranks of matrices. *Electron. Colloquium Comput. Complex.*, TR22-137, 2022.
- [BCK14] Abhishek Bhrushundi, Sourav Chakraborty, and Raghav Kulkarni. Property testing bounds for linear and quadratic functions via parity decision trees. In *Computer science—theory and applications*, volume 8476 of *Lecture Notes in Comput. Sci.*, pages 97–110. Springer, Cham, 2014.
- [BFS86] László Babai, Péter Frankl, and Janos Simon. Complexity classes in communication complexity theory. In *27th Annual Symposium on Foundations of Computer Science (SFCS 1986)*, pages 337–347. IEEE, 1986.
- [BVdW07] Harry Buhrman, Nikolay Vereshchagin, and Ronald de Wolf. On computation and communication with small bias. In *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC’07)*, pages 24–32. IEEE, 2007.
- [Cha93] Bernard Chazelle. Cutting hyperplanes for divide-and-conquer. *Discrete Comput. Geom.*, 9(2):145–158, 1993.
- [CHHS23] Tsun-Ming Cheung, Hamed Hatami, Kaave Hosseini, and Morgan Shirley. Separation of the factorization norm and randomized communication complexity. In *Proceedings, IEEE Conference on Computational Complexity (CCC)*, 2023.
- [CHZZ22] TsunMing Cheung, Hamed Hatami, Rosie Zhao, and Itai Zilberstein. Boolean functions with small approximate spectral norm. In *Electronic Colloquium on Computational Complexity (ECCC)*, TR22-041, 2022.

- [CLV19] Arkadev Chattopadhyay, Shachar Lovett, and Marc Vinyals. Equality alone does not simulate randomness. In *34th Computational Complexity Conference (CCC 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
- [CMS19] Arkadev Chattopadhyay, Nikhil S. Mande, and Suhail Sherif. The log-approximate-rank conjecture is false. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019*, pages 42–53, New York, NY, USA, 2019. ACM.
- [Coh60] Paul J. Cohen. On a conjecture of Littlewood and idempotent measures. *Amer. J. Math.*, 82:191–212, 1960.
- [CR12] Amit Chakrabarti and Oded Regev. An optimal lower bound on the communication complexity of gap-Hamming-distance. *SIAM J. Comput.*, 41(5):1299–1317, 2012.
- [DD07] Kenneth R. Davidson and Allan P. Donsig. Norms of Schur multipliers. *Illinois J. Math.*, 51(3):743–766, 2007.
- [DHP⁺22] Ben Davis, Hamed Hatami, William Pires, Ran Tao, and Hamza Usmani. On public-coin zero-error randomized communication complexity. *Information Processing Letters*, 178:106293, 2022.
- [EHK22] Louis Esperet, Nathaniel Harms, and Andrey Kupavskii. Sketching distances in monotone graph classes. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, page 1, 2022.
- [FF81] Peter Frankl and Zoltán Füredi. A short proof for a theorem of harper about hamming-spheres. *Discrete Mathematics*, 34(3):311–313, 1981.
- [FPS16] Jacob Fox, János Pach, and Andrew Suk. A polynomial regularity lemma for semi-algebraic hypergraphs and its applications in geometry and property testing. *SIAM Journal on Computing*, 45(6):2199–2223, 2016.
- [GGMT23] W. T. Gowers, Ben Green, Freddie Manners, and Terence Tao. On a conjecture of Marton, 2023.
- [GKPW17] Mika Göös, Pritish Kamath, Toniann Pitassi, and Thomas Watson. Query-to-communication lifting for P^{NP} . *32nd Computational Complexity Conference*, 79:Art. No. 12, 16, 2017.
- [GPW18a] Mika Göös, Toniann Pitassi, and Thomas Watson. Deterministic communication vs. partition number. *SIAM Journal on Computing*, 47(6):2435–2450, 2018.
- [GPW18b] Mika Göös, Toniann Pitassi, and Thomas Watson. The landscape of communication complexity classes. *computational complexity*, 27(2):245–304, 2018.
- [GS08] Ben Green and Tom Sanders. Boolean functions with small spectral norm. *Geometric and Functional Analysis*, 18(1):144–162, 2008.
- [Har20] Nathaniel Harms. Universal communication, universal graphs, and graph labeling. In *11th Innovations in Theoretical Computer Science Conference*, volume 151 of *LIPICs. Leibniz Int. Proc. Inform.*, pages Art. No. 33, 27. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2020.

- [HHH23] Lianna Hambardzumyan, Hamed Hatami, and Pooya Hatami. Dimension-free bounds and structural results in communication complexity. *Israel Journal of Mathematics*, 253(2):555–616, 2023.
- [HHL20] Hamed Hatami, Kaave Hosseini, and Shachar Lovett. Sign rank vs discrepancy. In *35th Computational Complexity Conference*, volume 169 of *LIPICs. Leibniz Int. Proc. Inform.*, pages Art. No. 18, 14. 2020.
- [HHP⁺22] Hamed Hatami, Pooya Hatami, William Pires, Ran Tao, and Rosie Zhao. Lower bound methods for sign-rank and their limitations. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2022)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2022.
- [HWZ22] Nathaniel Harms, Sebastian Wild, and Viktor Zamaraev. Randomized communication and implicit graph representations. In *Proceedings, ACM Symposium on Theory of Computing (STOC)*, 2022.
- [HZ24] Nathaniel Harms and Viktor Zamaraev. Randomized communication and implicit representations for matrices and graphs of small sign-rank. In *Proceedings, ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2024.
- [IW10] Russell Impagliazzo and Ryan Williams. Communication complexity with synchronized clocks. In *2010 IEEE 25th Annual Conference on Computational Complexity*, pages 259–269. IEEE, 2010.
- [KN97] Eyal Kushilevitz and Noam Nisan. *Communication complexity*. Cambridge University Press, Cambridge, 1997.
- [Liv95] Leo Livshits. A note on 0-1 Schur multipliers. *Linear Algebra Appl.*, 222:15–22, 1995.
- [LMSS07] Nati Linial, Shahar Mendelson, Gideon Schechtman, and Adi Shraibman. Complexity measures of sign matrices. *Combinatorica*, 27(4):439–463, 2007.
- [Lov14] Shachar Lovett. Recent advances on the log-rank conjecture in communication complexity. *Bull. Eur. Assoc. Theor. Comput. Sci. EATCS*, 112:18–35, 2014.
- [Lov16] Shachar Lovett. Communication is bounded by root of rank. *Journal of the ACM (JACM)*, 63(1):1–9, 2016.
- [LS88] László Lovász and Michael Saks. Lattices, mobius functions and communications complexity. In *[Proceedings 1988] 29th Annual Symposium on Foundations of Computer Science*, pages 81–90. IEEE Computer Society, 1988.
- [LS09a] Nati Linial and Adi Shraibman. Learning complexity vs. communication complexity. *Combin. Probab. Comput.*, 18(1-2):227–245, 2009.
- [LS09b] Nati Linial and Adi Shraibman. Lower bounds in communication complexity based on factorization norms. *Random Structures & Algorithms*, 34(3):368–394, 2009.
- [LS23] Troy Lee and Adi Shraibman. Around the log-rank conjecture. *Israel Journal of Mathematics*, 256(2):441–477, 2023.

- [Mil64] J. Milnor. On the Betti numbers of real varieties. *Proc. Amer. Math. Soc.*, 15:275–280, 1964.
- [New91] Ilan Newman. Private vs. common random bits in communication complexity. *Inf. Process. Lett.*, 39(2):67–71, July 1991.
- [Nis93] Noam Nisan. The communication complexity of threshold gates. *Combinatorics, Paul Erdős is Eighty*, 1:301–315, 1993.
- [NW95] Noam Nisan and Avi Wigderson. On rank vs. communication complexity. *Combinatorica*, 15(4):557–565, 1995.
- [Pis96] Gilles Pisier. *Similarity problems and completely bounded maps*, volume 1618 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1996.
- [PS86] Ramamohan Paturi and Janos Simon. Probabilistic communication complexity. *Journal of Computer and System Sciences*, 33(1):106–123, 1986.
- [PSS23] Toniann Pitassi, Morgan Shirley, and Adi Shraibman. The strength of equality oracles in communication. In *Proceedings, Innovations in Theoretical Computer Science (ITCS)*, 2023.
- [Raz92] A. A. Razborov. On the distributional complexity of disjointness. *Theoret. Comput. Sci.*, 106(2):385–390, 1992.
- [RS15] Sivaramakrishnan Natarajan Ramamoorthy and Makrand Sinha. On the communication complexity of greater-than. In *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 442–444. IEEE, 2015.
- [RY20] Anup Rao and Amir Yehudayoff. *Communication complexity and applications*. Cambridge University Press, Cambridge, 2020.
- [San11] Tom Sanders. A quantitative version of the non-abelian idempotent theorem. *Geom. Funct. Anal.*, 21(1):141–221, 2011.
- [San19] Tom Sanders. Boolean functions with small spectral norm, revisited. *Mathematical Proceedings of the Cambridge Philosophical Society*, 167(2):335–344, 2019.
- [San20] Tom Sanders. Bounds in Cohen’s idempotent theorem. *Journal of Fourier Analysis and Applications*, 26(2):1–64, 2020.
- [She08] Alexander A Sherstov. Halfspace matrices. *Computational Complexity*, 17:149–178, 2008.
- [She21] Alexander A. Sherstov. The hardest halfspace. *Comput. Complexity*, 30(2):Paper No. 11, 85, 2021.
- [ST23] Benny Sudakov and István Tomon. Matrix discrepancy and the log-rank conjecture. *arXiv preprint arXiv:2311.18524*, 2023.

- [Tho65] René Thom. Sur l'homologie des variétés algébriques réelles. In *Differential and Combinatorial Topology (A Symposium in Honor of Marston Morse)*, pages 255–265. Princeton Univ. Press, Princeton, N.J., 1965.
- [Vio15] Emanuele Viola. The communication complexity of addition. *Combinatorica*, 35(6):703–747, 2015.
- [War68] Hugh E. Warren. Lower bounds for approximation by nonlinear manifolds. *Trans. Amer. Math. Soc.*, 133:167–178, 1968.
- [Yao79] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (preliminary report). In *Proceedings of the eleventh annual ACM symposium on Theory of computing*, pages 209–213, 1979.
- [Yao03] Andrew Chi-Chih Yao. On the power of quantum fingerprinting. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 77–81, 2003.
- [YY85] A C Yao and F F Yao. A general approach to d-dimensional geometric queries. In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, STOC '85, page 163–168. Association for Computing Machinery, 1985.