Electronic Colloquium on Computational Complexity, Report No. 112 (2026)

Improved Bounds on the Half-Duplex Communication Complexity

Yuriy Dementiev[1], Tatiana Gladysh[2], Artur Ignatiev[1], Anna Kogan[3], Ivan Mihajlin[4], Timofey Moskalenko[1], Varvara Prozorova[1], Anastasiia Salimova[5], Lev Shpraidun[6], and Alexander V. Smal[4]

[1]Saint Petersburg State University
[2]Steklov Mathematical Institute of Russian Academy of Sciences
[3]Moscow Center for Advanced Studies
[4]JetBrains Research
[5]HSE University
[6]ITMO University

June 30, 2026

Abstract

We continue the study of half-duplex communication complexity, a model introduced in [HIMS18b] and further studied in [DIS+21], in which each player can either send a bit or listen in each round, similarly to communication over a walkie-talkie. We prove improved upper bounds for the Inner Product function in the half-duplex models with silence and with zero, as well as an improved upper bound for the Disjointness function in the model with zero. We also determine the half-duplex communication complexity of a uniformly random Boolean function up to an additive constant: with high probability, it is $n/\log_2 3 + O(1)$ in the model with silence and $n + O(1)$ in the models with zero and with an adversary.

Keywords: communication complexity, half-duplex communication, equality oracle

1 Introduction

Communication complexity is a powerful tool with applications across many areas of computer science, including algorithms, circuit complexity, proof complexity, and more. In the classical model of communication complexity introduced by Yao [Yao79], two players, Alice and Bob, aim to compute $f(x, y)$ for some function f, where Alice knows only x, and Bob knows only y. The players can communicate by exchanging bits, one bit per round, until both parties determine the value of $f(x, y)$. There are many generalizations of this model, such as randomized, non-deterministic, and multiparty communication complexity.

ISSN 1433-8092

The essential property of the classical model is that in each round of communication, one player sends a bit, and the other receives it. In [HIMS18b], the authors proposed a communication model in which the players communicate over a *half-duplex channel*. A well-known example of half-duplex communication is talking using a walkie-talkie: one has to hold a "push-to-talk" button to speak to another person, and the other has to keep it released to listen. If both parties try to speak simultaneously, they cannot hear each other.

Formally speaking, in every round, each player chooses one of three actions: "`send 0`", "`send 1`", or "`receive`". This induces three types of rounds: a *classical round*, when one player sends a bit while the other receives; a *wasted round*, when both players send, and both messages are lost; and a *silent round*, when both players receive. In [HIMS18b], the authors considered three variations of the half-duplex model based on what happens in silent rounds: half-duplex models *with silence*, *with zero*, and *with adversary*.

The motivation for half-duplex communication models stems from the study of the Karchmer-Wigderson games [KW88] for the multiplexer relation [EIRS01]. The framework of half-duplex communication complexity proved useful and was used in a series of papers [MS21, IMS22, Wu23, Mei25] on the variants of the KRW conjecture [KRW95], which offers a promising approach to proving $P \neq NC^1$. In these papers, half-duplex communication models play an essential role: we can prove the existence of a hard function from a lower bound on the half-duplex communication complexity of the multiplexer relation. We do not know how to show it directly for the classical communication complexity. Given that the desired result is usually formulated in terms of classical communication complexity (as a model corresponding to De Morgan formulas), the resulting lower bound must be translated from one model to the other. Since we do not fully understand the relationship between classical and half-duplex communication complexity, translating between the two models degrades the lower bound. Every classical protocol is a valid half-duplex protocol, whereas any half-duplex protocol of depth d can be simulated classically in at most $2d$ rounds. Hence, exact linear coefficients are essential: a function whose half-duplex complexity is only half of its classical complexity provides no gain when the resulting lower bound is translated back to the classical model. This motivates further study of half-duplex communication complexity and of its relationship with classical communication complexity. A better understanding of half-duplex communication may help us make further progress on the KRW conjecture.

The half-duplex communication complexity of functions $\{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ was studied in the original paper [HIMS18b] and later developed in [DIS+21]. It was shown that the communication complexity in half-duplex models not only differs from that of the classical model but also behaves differently. For example, in the classical communication model, the Equality function, the Disjointness function, and the Inner Product function have complexity $n+1$, meaning that all three are the hardest functions among all functions from $\{0,1\}^n \times \{0,1\}^n \to \{0,1\}$. In the half-duplex models with silence and with zero, the Equality is less complex than the other two.

Although our upper-bound results concern only the models with silence and with zero, the model with an adversary will reappear in Section 4, where we transfer a lower bound from communication with a single-bit equality oracle.

1.1 Overview of Results

In this paper, we continue the line of research initiated in [HIMS18b, DIS+21] and present a series of improved bounds. We improve half-duplex communication complexity bounds for the Inner Product function (IP_n) and the Disjointness function (DISJ_n) (see Section 3), and determine the half-duplex communication complexity of a uniformly random Boolean function up to an additive constant in the models with silence, with zero, and with an adversary (see Section 4).

The upper bound for the Inner Product in the model with silence (see Theorem 3) is optimal up to an additive constant, so the question of the complexity of IP in this model is closed up to an additive constant. It is also worth mentioning that obtaining this result made it possible to find a flaw in a published result (see Theorem 4). For several years, it was unknown whether IP_n could be computed in fewer than n rounds in the model with zero. We answer this question affirmatively by proving the upper bound $\mathrm{D}_0^{\mathrm{hd}}(\mathsf{IP}_n) \leqslant \lceil 7n/8 \rceil + 4$ (see Theorem 5). Moreover, Theorem 10 shows that a uniformly random function has complexity at least $n-1$ with high probability, so IP_n is not among the hardest functions in this model. We show a similar lower bound for the model with silence: in this model, neither the Inner Product nor the Disjointness is among the hardest (see Theorem 11). The upper bound for the Disjointness in the model with zero (see Theorem 6) introduces a new idea of recursive reduction that allows us to significantly improve the upper bound. For the comparison with the previously known results, please see Table 1.

Table 1: Comparison of the previously known bounds and the bounds from this paper.

Model	Bound type	IP_n	DISJ_n
with silence	previous upper bound	$n+1$	$n/2+2$
	upper bound in this paper	$\lceil n/2 \rceil + 2$	—
	best known lower bound	$n/2$	$n \log_5 2 \approx 0.43068n$
with zero	previous upper bound	$n+1$	$3n/4 + o(n)$
	upper bound in this paper	$\lceil 7n/8 \rceil + 4$	$0.6615n + O(\log^2 n)$
	best known lower bound	$n \log_{\frac{2}{3-\sqrt{5}}} 2 \approx 0.72021n$	$n \log_3 2 \approx 0.63092n$

2 Half-duplex communication complexity

In this section, we give a formal definition of two half-duplex communication complexity models: *with silence* and *with zero*.

Consider the following game with two players. Alice and Bob aim to compute a function $f\colon X \times Y \to Z$ on an input $(x, y) \in X \times Y$. Throughout this paper, we use $X = Y = \{0,1\}^n$ and $Z = \{0,1\}$. The function f is known to both players, but the input is distributed between them: Alice receives x, whereas Bob receives y. They communicate in order to determine $f(x, y)$. Alice and Bob can exchange bits using a half-duplex communication channel together with a synchronizing mechanism (e.g., synchronized clock), which allows them to organize communication in rounds. In every round, each player chooses one of the *actions*: "`send 0`", "`send 1`", and "`receive`". Note that in the classical communication complexity model, whenever one player sends a bit, the other always receives it; the information is never lost, and both players always know about the other's action. In the half-duplex communication model, there might be three different situations:

- *classical round*: one player chooses to send some bit b while the other player chooses to receive. In this case, the receiving player receives bit b.
- *wasted round:* both players choose to send some bits. In this case, the messages are lost, and players do not learn any information (i.e., a player knows nothing about the other player's action).
- *silent round:* both players choose to receive.

  In [HIMS18b], the authors suggested three ways to define what happens in silent rounds.

  1. *half-duplex model with silence:* the players receive a special symbol `silence` meaning that there is no message in the channel.
  2. *half-duplex model with zero:* both players receive 0.

     (Note that a player who received 0 cannot distinguish a silent round from a classical round in which the other player sent 0.)
  3. *half-duplex model with an adversary:* the players receive arbitrary bits, not necessarily the same.

The formal definitions and the upper-bound results in this paper concern the first two models. The model with an adversary is used only in Section 4 for transferring a lower bound. For its formal definition, we refer to [HIMS18a, DV23].

To be able to compute f for all possible inputs $(x, y) \in X \times Y$, Alice and Bob have to follow a *communication protocol* that defines their actions on all inputs.

Definition 1. A *half-duplex communication protocol with silence* for a function $f\colon X \times Y \to Z$ is a pair of rooted trees (T_A, T_B) of arity at most five such that:

- every leaf l is labeled with some $z_l \in Z$,
- every internal node $v \in T_A$ is labeled with a pair of functions $g_v : X \to \mathcal{A}$ and $h_v : \mathcal{E}_s \to D(v)$, where $D(v)$ denotes the set of children of v,

  $$\mathcal{A} = \{\texttt{send}\,0, \texttt{send}\,1, \texttt{receive}\}$$

  is the set of *actions*,

  $$\mathcal{E}_s = \{\texttt{sent}\,0, \texttt{sent}\,1, \texttt{received}\,0, \texttt{received}\,1, \texttt{silence}\}$$

  is the set of *events*;
- similarly, every internal node $u \in T_B$ is labeled with a pair of functions $g_u : Y \to \mathcal{A}$ and $h_u : \mathcal{E}_s \to D(u)$.

For any $(x, y) \in X \times Y$, a half-duplex communication protocol with silence defines *Alice's and Bob's traces* $\pi_A(x, y)$ and $\pi_B(x, y)$ as the sequences of nodes $v_1, \ldots, v_t \in T_A$ and $u_1, \ldots, u_t \in T_B$ such that:

- v_1 and u_1 are the roots of T_A and T_B,
- v_t and u_t are leaves,
- for every $i \in \{1, \ldots, t-1\}$, it holds that

$$v_{i+1} = h_{v_i}\big(\phi\big(g_{v_i}(x), g_{u_i}(y)\big)\big),$$
$$u_{i+1} = h_{u_i}\big(\phi\big(g_{u_i}(y), g_{v_i}(x)\big)\big),$$

where $\phi : \mathcal{A} \times \mathcal{A} \to \mathcal{E}_s$ and the value of $\phi(a_1, a_2)$ is defined by the following table

a_1 \ a_2	`send 0`	`send 1`	`receive`
`send 0`	`sent 0`	`sent 0`	`sent 0`
`send 1`	`sent 1`	`sent 1`	`sent 1`
`receive`	`received 0`	`received 1`	`silence`

(i.e., function ϕ maps actions of the players into events).

A protocol *computes* a function f if for any $(x, y) \in X \times Y$, both players' traces end in leaves with the same label $z = f(x, y)$. The depth of a protocol is the maximum, over all inputs (x, y), of the number of rounds in the corresponding execution. The *half-duplex communication complexity with silence*, denoted by $\mathrm{D}^{\mathrm{hd}}_{\mathrm{s}}(f)$, is the minimum depth of a protocol computing f.

Definition 2. A *half-duplex communication protocol with zero* for a function $f : X \times Y \to Z$ is defined similarly to the protocol in Theorem 1, but with a different set of events

$$\mathcal{E}_0 = \{\texttt{sent}\,0, \texttt{sent}\,1, \texttt{received}\,0, \texttt{received}\,1\}$$

that does not contain `silence`, and $\phi(\texttt{receive}, \texttt{receive}) := \texttt{received}\,0$.

The advantage of half-duplex models is that the players can send bits simultaneously and receive simultaneously, which is not possible in the classical model. This allows players in the model with silence to send an additional third symbol: for example, in a given round, Bob can always listen and then, depending on Alice's action, he will receive either 0 or 1, or observe `silence`. This way, Alice can send Bob more information (e.g., Alice can send an n-bit string in $\lceil n/\log_2 3\rceil$ rounds). This explains why the half-duplex model with silence is stronger than the classical one. As for the half-duplex model with zero, its advantages are less obvious. It might seem that the model has no advantage: there is no benefit from wasted rounds, since players get no information in them, and silent rounds are indistinguishable from classical ones. However, it can be shown that this model is strictly stronger than the classical one.

3 Upper bounds

In this section, we prove three new upper bounds on the half-duplex communication complexity — two upper bounds for the Inner Product and one for the Disjointness.

3.1 Inner Product

The Inner Product function $\mathsf{IP}_n : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ is defined by

$$\mathsf{IP}_n(x,y) = \bigoplus_{i=1}^{n} x_i \cdot y_i = \sum_{i=1}^{n} x_i \cdot y_i \pmod 2.$$

This is one of the hardest functions in the classical communication model — it has communication complexity exactly $n+1$. In other words, the trivial protocol in which Alice sends all her bits to Bob, Bob computes the result, and then sends it back to Alice, is optimal. For the half-duplex model with silence, one might expect this function to have complexity roughly $n/\log_2 3$ (as Alice can send her entire input to Bob using this number of rounds). The following theorem shows that the players can compute the Inner Product more efficiently.

Theorem 3. *For any* $n \in \mathbb{N}$, $\mathrm{D}_{\mathrm{s}}^{\mathrm{hd}}(\mathsf{IP}_n) \leqslant \lceil \frac{n}{2} \rceil + 2$.

Remark 4. *The conference version of the paper introducing half-duplex communication complexity [HIMS18b, Theorem 18] claimed the lower bound* $\mathrm{D}_{\mathrm{s}}^{\mathrm{hd}}(\mathsf{IP}_n) \geqslant n/1.67$*, which contradicts Theorem 3. The proof of that claim contains a flaw, and the lower bound does not appear in the updated full version [HIMS18a]. The proof of this upper bound led to the discovery of this issue.*

Proof of Theorem 3. For simplicity, we assume that n is even. In the case of odd n, Alice and Bob can append one zero coordinate to each input without changing the value of the function. During the communication, the players will maintain two *accumulators*: Alice will maintain a bit $a \in \{0,1\}$, and Bob will maintain a bit $b \in \{0,1\}$. Initially $a = b = 0$.

For $i \in \{1, \dots, n/2\}$, we will refer to the pairs of input bits x_{2i-1}, x_{2i} and y_{2i-1}, y_{2i} as *Alice's* i*th block* and *Bob's* i*th block*, respectively. In round i, the players choose their actions based on the values of bits in their ith blocks, and their goal is to compute

$$\delta_i = x_{2i-1}y_{2i-1} \oplus x_{2i}y_{2i}.$$

Consider the following rules for Alice's and Bob's actions. In the round i, Alice chooses her action depending on the values of bits x_{2i-1}, x_{2i}. For simplicity, we will write the values of bits in blocks without a comma.

$$00 \mapsto \texttt{send}\,0, \qquad 01 \mapsto \texttt{send}\,1, \qquad 10 \mapsto \texttt{receive}, \qquad 11 \mapsto \texttt{receive}.$$

Similarly, in the round i, Bob chooses his action depending on the values of bits y_{2i-1}, y_{2i}.

$$00 \mapsto \texttt{send}\,0, \qquad 01 \mapsto \texttt{receive}, \qquad 10 \mapsto \texttt{send}\,1, \qquad 11 \mapsto \texttt{receive}.$$

A player flips their accumulator in round i if either

- the current block is 01 or 10 and the event is `received 1`, or
- the current block is 11 and the event is not `received 0`.

A player with block 00 never flips the accumulator.

Let $a_i, b_i \in \{0, 1\}$ indicate whether Alice's and Bob's accumulators, respectively, are flipped in round i. We now show that $a_i \oplus b_i = \delta_i$.

Indeed, we distinguish the following cases.

1. If one of the blocks is 00, then no accumulator is flipped, so $(a_i, b_i) = (0, 0)$. Also, $\delta_i = 0$.
2. If the blocks are $(01, 10)$ or $(10, 01)$, then neither player receives 1, hence again $(a_i, b_i) = (0, 0)$ and $\delta_i = 0$.
3. If the blocks are $(01, 01)$, then Bob receives 1 and flips his accumulator, while Alice does not flip. Thus $(a_i, b_i) = (0, 1)$ and $\delta_i = 1$.
4. If the blocks are $(10, 10)$, then Alice receives 1 and flips her accumulator, while Bob does not flip. Thus $(a_i, b_i) = (1, 0)$ and $\delta_i = 1$.
5. If exactly one of the blocks is 11 and the other one is 01 or 10, then only the player with block 11 flips the accumulator. Hence (a_i, b_i) is either $(1, 0)$ or $(0, 1)$, and in both cases $\delta_i = 1$.
6. If the blocks are $(11, 11)$, then both players flip their accumulators, so $(a_i, b_i) = (1, 1)$, while $\delta_i = 1 \cdot 1 \oplus 1 \cdot 1 = 0$.

Therefore, after processing the first k blocks,

$$a \oplus b = \bigoplus_{i=1}^{k} \delta_i.$$

In the last two rounds, Alice sends the value of a, Bob computes $\mathsf{IP}_n(x, y) = a \oplus b$, and sends the result back to Alice. □

The following theorem establishes an upper bound on the half-duplex communication complexity of the Inner Product in the model with zero.

Theorem 5. *For any* $n \in \mathbb{N}$, $\mathrm{D}_0^{\mathrm{hd}}(\mathsf{IP}_n) \leqslant \lceil \frac{7n}{8} \rceil + 4$.

Proof. Similarly to the proof of Theorem 3, Alice and Bob will view their inputs $x, y \in \{0, 1\}^n$ as sequences of $\lceil n/2 \rceil$ blocks of length 2. If n is odd, both players append a zero to their inputs. Consider the matrix defining the values of the function $\mathsf{IP}_2(u, v)$.

u \ v	00	01	10	11
00	0	0	0	0
01	0	1	0	1
10	0	0	1	1
11	0	1	1	0

The crucial idea of the protocol is to decompose IP_2 into the XOR of three functions. The goal is to reduce the number of entries equal to 1 in the function matrix. We flip all entries in the row indexed by 11 and then all entries in the column indexed by 11. This leads to the following decomposition

$$\mathsf{IP}_2(u, v) = \psi(u, v) \oplus [u = 11] \oplus [v = 11].$$

The resulting function $\psi(u, v)$ is given by the following table.

$u \backslash v$	00	01	10	11
00	0	0	0	**1**
01	0	1	0	**0**
10	0	0	1	**0**
11	**1**	**0**	**0**	**0**

Note that this decomposition decreases the number of entries equal to 1 from six to four.

Let u_i and v_i denote Alice's and Bob's ith blocks, respectively.

$$\begin{aligned}\mathsf{IP}_n(x, y) = \bigoplus_{i=1}^{\lceil n/2 \rceil} \mathsf{IP}_2(u_i, v_i) &= \bigoplus_{i=1}^{\lceil n/2 \rceil} \big(\psi(u_i, v_i) \oplus [u_i = 11] \oplus [v_i = 11]\big) \\ &= \left(\bigoplus_{i=1}^{\lceil n/2 \rceil} \psi(u_i, v_i)\right) \oplus \left(\bigoplus_{i=1}^{\lceil n/2 \rceil} [u_i = 11]\right) \oplus \left(\bigoplus_{i=1}^{\lceil n/2 \rceil} [v_i = 11]\right) \\ &= \left(\bigoplus_{i=1}^{\lceil n/2 \rceil} \psi(u_i, v_i)\right) \oplus (\#_x(11) \bmod 2) \oplus (\#_y(11) \bmod 2)\,,\end{aligned}$$

where $\#_x(11)$ and $\#_y(11)$ denote the numbers of blocks 11 in x and y, respectively. So, we reduced computing block-wise IP_2 to computing block-wise ψ and two local correction terms, $\#_x(11) \bmod 2$ and $\#_y(11) \bmod 2$, which Alice and Bob can compute from their respective inputs without communication. We are now ready to describe the protocol.

For blocks $a, b \in \{00, 01, 10, 11\}$, let $\#(a, b) := \big|\{i : (u_i, v_i) = (a, b)\}\big|$.

Preprocessing stage Alice chooses the most frequent block u in x, breaking ties according to a fixed order known to both players, and sends u to Bob in two rounds. Based on the value of u, the players choose their strategies for two main stages of the protocol. In each of Stages 1 and 2, each player uses the corresponding trigger block specified in Table 2.

Table 2: Strategies of Alice and Bob based on the most frequent block in x.

Most frequent block in x	00	01	10	11
Alice's block in Stage 1	00	01	10	11
Bob's block in Stage 1	00	10	01	11
Alice's block in Stage 2	01	00	11	10
Bob's block in Stage 2	10	00	11	01

The quantities learned by the players in the two stages are as follows.

u	Alice, Stage 1	Bob, Stage 1	Alice, Stage 2	Bob, Stage 2
00	$\#(11,00)$	$\#(00,11)$	$\#(10,10)$	$\#(01,01)$
01	$\#(10,10)$	$\#(01,01)$	$\#(11,00)$	$\#(00,11)$
10	$\#(01,01)$	$\#(10,10)$	$\#(00,11)$	$\#(11,00)$
11	$\#(00,11)$	$\#(11,00)$	$\#(01,01)$	$\#(10,10)$

Stage 1 Alice and Bob iterate over all block positions, one position per round. Alice sends 1 if and only if her block is her Stage 1 trigger block from Table 2; otherwise, she receives. Bob acts analogously. By counting the relevant rounds in which they receive 1, Alice and Bob learn the two Stage 1 quantities shown above.

A round is silent if both players receive. Since the only transmitted bit is 1, both players recognize such positions by receiving 0. Since Alice's trigger block is the most frequent block u, the number of silent rounds is at most

$$\frac{3}{4}\left\lceil\frac{n}{2}\right\rceil \leqslant \frac{3(n+1)}{8}.$$

Stage 2 Alice and Bob iterate only over the positions that were silent in Stage 1. In each such position, Alice sends 1 if and only if her block is her Stage 2 trigger block from Table 2; otherwise, she receives. Bob acts analogously.

By counting the relevant rounds in which they receive 1, Alice and Bob learn the two Stage 2 quantities shown above. All pairs counted in Stage 2 are silent in Stage 1. Thus, for every value of u, the two stages together count exactly the four pairs

$$(00,11), \quad (01,01), \quad (10,10), \quad (11,00),$$

which are precisely the entries on which ψ is equal to 1. Stage 2 uses at most $3(n+1)/8$ rounds.

Final stage Let p_A be the parity of the sum of the two quantities learned by Alice in Stages 1 and 2, and define p_B analogously for Bob. Then

$$p_A \oplus p_B = \bigoplus_{i=1}^{\lceil n/2\rceil} \psi(u_i, v_i).$$

Alice sends

$$(\#_x(11) + p_A) \bmod 2$$

to Bob. Bob XORs it with $(\#_y(11)+p_B) \bmod 2$, obtaining $\mathsf{IP}_n(x,y)$, and sends the result to Alice in the last round.

The total number of rounds in the whole protocol is at most

$$2 + \frac{n+1}{2} + \frac{3(n+1)}{8} + 2 = \frac{7(n+1)}{8} + 4 < \frac{7n}{8} + 5.$$

Since the number of rounds is an integer, it is at most $\lceil \frac{7n}{8} \rceil + 4$. □

3.2 Disjointness

The Disjointness function $\mathsf{DISJ}_n : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ is defined by

$$\mathsf{DISJ}_n(x,y) = 1 \iff \forall i \in [n] : x_i \neq 1 \vee y_i \neq 1.$$

Classically, Disjointness has communication complexity $n+1$. We show that it can be computed more efficiently in the half-duplex model with zero.

Theorem 6. *For all* $n \in \mathbb{N}$, $\mathrm{D}_0^{\mathrm{hd}}(\mathsf{DISJ}_n) \leqslant 0.6615n + O(\log^2 n)$.

The proof is based on the following more general statement.

Lemma 7. *Let* $\alpha \in [0,1]$, *and define*

$$a(\alpha) := \max\left\{\frac{1}{1+\alpha}, \frac{1}{2} + \frac{\alpha}{2\log_2 3}\right\}.$$

Then, for every $n \in \mathbb{N}$,

$$\mathrm{D}_0^{\mathrm{hd}}(\mathsf{DISJ}_n) \leqslant a(\alpha)n + O(\log^2 n).$$

Theorem 6 follows from the lemma by taking $\alpha = 0.5118$.

Proof of Theorem 6. For $\alpha = 0.5118$,

$$a(\alpha) = \max\left\{\frac{1}{1+\alpha}, \frac{1}{2} + \frac{\alpha}{2\log_2 3}\right\} < 0.6615.$$

The claim follows from Theorem 7. □

Proof of Theorem 7. We describe a recursive protocol and analyze its worst-case number of rounds. Inputs of bounded length are handled by the trivial protocol; this affects only the constant hidden in the $O(\log^2 n)$ term.

Reduction to even length If n is odd, the players exchange x_n and y_n, the last bits of their inputs, using two rounds of communication. If $x_n = y_n = 1$, then both players output 0 and terminate. Otherwise, they discard the last coordinate. In the rest of the proof, we assume that n is even.

Pair Counting Similarly to the previous proofs, Alice and Bob will view their inputs $x, y \in \{0,1\}^n$ as sequences of $n/2$ blocks of length 2.

First, Bob sends Alice the three numbers $\#_y(00)$, $\#_y(01)$, and $\#_y(10)$. Alice then computes $\#_y(11) = n/2 - \#_y(00) - \#_y(01) - \#_y(10)$. Alice sends the analogous three numbers to Bob. Since every counter is at most $n/2$, this stage requires $O(\log_2 n)$ rounds.

Consider the following four inequalities.

$$\begin{aligned} \#_x(01) + \#_y(01) &\geqslant \alpha n/2, & \#_x(00) + \#_x(11) &\geqslant (1-\alpha)n/2, \\ \#_x(10) + \#_y(10) &\geqslant \alpha n/2, & \#_y(00) + \#_y(11) &\geqslant (1-\alpha)n/2. \end{aligned}$$

We claim that at least one of these four inequalities holds. Indeed, if all four inequalities are false simultaneously then

$$(\#_x(00) + \#_x(01) + \#_x(10) + \#_x(11)) + (\#_y(00) + \#_y(01) + \#_y(10) + \#_y(11)) \\ < 2(\alpha n/2 + (1-\alpha)n/2) = n.$$

This contradicts the fact that there are exactly n block occurrences across the two inputs, so at least one of the four inequalities is true. We will consider these cases separately. Since all these counters are known to both players, they can determine which inequalities hold. If several inequalities hold, the players choose the first one in the displayed order.

Type 1 Branch Suppose that $\#_x(01) + \#_y(01) \geqslant \alpha n/2$. In $n/2$ rounds, the players iterate over their blocks, one block per round, and send 1 every time the current block is 01; on all other blocks, the players receive. Each player sets a rejection flag in either of the following cases:

1. the number of received 1's is smaller than the number of 01 blocks in the other player's input; this means that a wasted round with block pair $(01, 01)$ occurred;
2. the player receives 1 while their own block is 11; this means that the block pair is $(11, 01)$ or $(01, 11)$.

In either case, the inputs are not disjoint. The players exchange their rejection flags using two rounds and terminate with output 0 if either flag is set.

Input Compression If neither rejection flag is set, the pair $(01, 01)$ does not occur. Hence, the positions in which Alice has block 01 and those in which Bob has block 01 are disjoint, and their total number is

$$\#_x(01) + \#_y(01) \geqslant \alpha n/2.$$

All these positions are known to both players: a player whose block is not 01 receives 1 exactly when the other player's block is 01, while a player with block 01 already knows that the position must be eliminated. Thus, at most $(1-\alpha)n/2$ positions remain. Now the players construct new inputs x', y' of length $n' \leqslant (1-\alpha)n/2$ by replacing blocks in x, y (ignoring eliminated positions) with individual bits using the following rules:

$$00 \mapsto 0, \quad 10 \mapsto 1, \quad 11 \mapsto 1.$$

For every remaining position, the original blocks intersect if and only if both corresponding compressed bits are 1. Therefore,

$$\mathsf{DISJ}_n(x, y) = \mathsf{DISJ}_{n'}(x', y').$$

We can now apply the inductive hypothesis for the problem of length n'. The case $\#_x(10) + \#_y(10) \geqslant \alpha n/2$ is symmetric after exchanging the first and second coordinates within each block. In this case, the compression rule is

$$00 \mapsto 0, \qquad 01 \mapsto 1, \qquad 11 \mapsto 1.$$

Type 2 Branch Without loss of generality, assume that $\#_x(00)+\#_x(11) \geqslant (1-\alpha)n/2$. The case of $\#_y(00) + \#_y(11) \geqslant (1-\alpha)n/2$ is symmetric.

In $n/2$ rounds, the players iterate over their blocks, one block per round, and send 1 every time the current block is 00; on all other blocks, the players receive. During these rounds, a player sets a rejection flag if their own block is 11 and they receive 0. Since the only transmitted bit is 1, receiving 0 means that the other player also receives and therefore has a block different from 00. Hence, the two blocks intersect. The players exchange their rejection flags using two rounds and terminate with output 0 if either flag is set.

If neither rejection flag is set, every block 11 is paired with a block 00 on the other side. Moreover, both players know all positions in which at least one block is 00: a player whose block is different from 00 receives 1 exactly when the other player's block is 00, while a player with block 00 already knows that the position must be eliminated.

Eliminating these positions also eliminates every position containing a block 11. In particular, all positions in which Alice has block 00 or 11 are eliminated. By the assumed inequality, their number is at least

$$\#_x(00) + \#_x(11) \geqslant (1-\alpha)n/2.$$

Therefore, at most $\alpha n/2$ positions remain, and every remaining block is either 01 or 10.

Now the players construct new inputs x', y' of length $n' \leqslant \alpha n/2$ by replacing blocks in x, y (ignoring eliminated positions) with individual bits using the following rules:

$$\text{Alice:} \quad 01 \mapsto 0, \quad 10 \mapsto 1, \qquad \text{Bob:} \quad 01 \mapsto 1, \quad 10 \mapsto 0.$$

Thus, the remaining instance of the Disjointness reduces to the Equality:

$$\mathsf{DISJ}_n(x,y) = \mathsf{EQ}_{n'}(x',y').$$

The Equality in the half-duplex model with zero can be solved in

$$\frac{n'}{\log_2 3} + O(1) \leqslant \frac{\alpha}{2\log_2 3} n + O(1)$$

rounds [HIMS18a].

Recurrence Analysis Let $T(n)$ be the upper bound on the number of rounds. Then

$$\begin{aligned} &\text{Type 1:} \quad T(n) \leqslant \frac{n}{2} + T(n') + O(\log_2 n), \qquad n' \leqslant \frac{1-\alpha}{2} n, \\ &\text{Type 2:} \quad T(n) \leqslant \left(\frac{1}{2} + \frac{\alpha}{2\log_2 3}\right) n + O(\log_2 n). \end{aligned}$$

Linear Coefficient. For the Type 1 branch, by the definition of $a(\alpha)$,

$$a(\alpha) \geqslant \frac{1}{1+\alpha} \iff \frac{1}{2} + a(\alpha)\frac{1-\alpha}{2} \leqslant a(\alpha).$$

Hence, if the recursive call has size $n' \leqslant (1-\alpha)n/2$, then

$$\frac{n}{2} + a(\alpha)n' \leqslant a(\alpha)n.$$

For the Type 2 branch,

$$\frac{1}{2} + \frac{\alpha}{2\log_2 3} \leqslant a(\alpha).$$

Thus, the total linear contribution is at most $a(\alpha)n$.

Lower-order Terms. Only the Type 1 branch is recursive, and every recursive call reduces the input length by a factor of at most $(1-\alpha)/2 \leqslant 1/2$. Therefore, the recursion depth is $O(\log_2 n)$. Since each level adds $O(\log_2 n)$ rounds, the total lower-order contribution is $O(\log^2 n)$. Hence,

$$T(n) \leqslant a(\alpha)n + O(\log^2 n). \qquad \square$$

4 Complexity of Random Functions

A communication with a single-bit equality oracle (EQ_1) can be viewed as a process where Alice and Bob interact through a third party, Charlie. In each round, they each send one bit to Charlie, who tells them only whether the bits are equal. Based on this information, they update their knowledge and choose the next bits to send. This communication defines a protocol that gradually restricts the set of possible inputs, at each step the current rectangle of inputs is partitioned into up to four subrectangles (depending on the possible pairs of bits sent). See [BFS86, Kra98, CLV19] for background on communication complexity with oracles.

It is also easy to see that the communication complexity with a single-bit equality oracle for a function f (denoted $\mathsf{P}^{\mathsf{EQ}_1}(f)$) is at most the classical communication complexity $\mathrm{D}(f)$. Indeed, each round of a classical protocol can be simulated by a round in the EQ_1 model: one player sends a fixed (constant) bit to Charlie, while the other sends the bit they would have sent in the classical protocol.

We now consider the complexity of random Boolean functions of $2n$ bits with a partitioned input for Alice and Bob.

Theorem 8. *Let $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ be chosen uniformly at random. Then, with probability $1 - o(1)$,*

$$n - 1 \leqslant \mathsf{P}^{\mathsf{EQ}_1}(f) \leqslant n + 1.$$

In particular, $\mathsf{P}^{\mathsf{EQ}_1}(f) \geqslant n - 1$ with probability $1 - o(1)$.

Proof. Let N_d be the number of functions computable by protocols with the EQ_1 oracle of depth at most d. After adding dummy rounds, we may assume that every protocol has depth exactly d.

The oracle returns one bit, so its answer histories form the full binary tree of depth d. For every history $h \in \{0,1\}^{<d}$, the next bits sent by Alice and Bob are Boolean functions of their respective inputs. Since there are $2^d - 1$ such histories, these functions can be chosen in at most $\left(2^{2^n}\right)^{2(2^d-1)}$ ways.

In each round, the ordered pair of sent bits has four possible values. Moreover, each player can reconstruct this pair from their own bit and the oracle's answer. Hence the output can be specified as a function of the full pair transcript $\tau \in (\{0,1\}^2)^d$. There are at most 4^d such transcripts and therefore at most 2^{4^d} choices of outputs. Consequently,

$$N_d \leqslant 2^{2^{n+1}(2^d-1)+4^d}.$$

Set $d = n - 2$. Then

$$\log_2 N_d \leqslant 2^{n+1}(2^{n-2} - 1) + 4^{n-2} \leqslant \frac{9}{16} 4^n.$$

Since there are 2^{4^n} Boolean functions on $\{0,1\}^n \times \{0,1\}^n$,

$$\Pr\left[\mathsf{P}^{\mathsf{EQ}_1}(f) \leqslant n - 2\right] \leqslant 2^{-\frac{7}{16}4^n} = o(1).$$

Thus, $\mathsf{P}^{\mathsf{EQ}_1}(f) \geqslant n - 1$ with probability $1 - o(1)$. Finally, $\mathsf{P}^{\mathsf{EQ}_1}(f) \leqslant \mathrm{D}(f) \leqslant n + 1$ for every f. □

Corollary 9. *For any function f, it holds that $\mathsf{P}^{\mathsf{EQ}_1}(f) \leqslant \mathrm{D}_0^{\mathrm{hd}}(f) \leqslant \mathrm{D}_{\mathrm{a}}^{\mathrm{hd}}(f)$.*

Proof. Given a protocol in the half-duplex model with zero, encode both `receive` and `send 0` by the bit 0, and encode `send 1` by the bit 1. In every round, the players send these encoded bits to the EQ_1 oracle.

If a player receives in the original protocol, the reconstructed bit of the other player is exactly the bit they would receive in the half-duplex model with zero. If a player sends, they ignore the oracle answer and follow the transition corresponding to their own sent bit. Thus, every zero-model protocol can be simulated without increasing its depth, and

$$\mathsf{P}^{\mathsf{EQ}_1}(f) \leqslant \mathrm{D}_0^{\mathrm{hd}}(f).$$

Moreover, every protocol that is correct against arbitrary values supplied in silent rounds is also correct when these values are fixed to zero. Hence,

$$\mathrm{D}_0^{\mathrm{hd}}(f) \leqslant \mathrm{D}_{\mathrm{a}}^{\mathrm{hd}}(f).$$ □

Corollary 10. *Let $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ be chosen uniformly at random. Then, with probability $1 - o(1)$,*

$$n - 1 \leqslant \mathrm{D}_0^{\mathrm{hd}}(f) \leqslant \mathrm{D}_{\mathrm{a}}^{\mathrm{hd}}(f) \leqslant n + 1.$$

In particular, the half-duplex complexities with zero and with an adversary are both at least $n - 1$ with high probability.

Proof. The lower bound follows from Theorem 8 and the preceding corollary. For the upper bound, every classical protocol is a valid protocol in the model with an adversary, since it uses neither silent nor wasted rounds. Therefore,

$$\mathrm{D}_{\mathrm{a}}^{\mathrm{hd}}(f) \leqslant \mathrm{D}(f) \leqslant n + 1.$$ □

Theorem 11. *Let $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ be chosen uniformly at random. Then, with probability $1 - o(1)$,*

$$\left\lfloor \frac{n}{\log_2 3} \right\rfloor \leqslant \mathrm{D}_{\mathrm{s}}^{\mathrm{hd}}(f) \leqslant \left\lceil \frac{n}{\log_2 3} \right\rceil + 1.$$

Proof. Let N_d be the number of functions computable by protocols of depth at most d. Extend every protocol to depth exactly d and embed each of its two protocol trees into the full 5-ary tree of depth d.

For each fixed input of a player, at most 3^t nodes at level t are reachable: sending leads to one child, while listening leads to at most three children, corresponding to `received 0`, `received 1`, or `silence`. Hence, over all internal levels, the number of reachable input-node pairs for each player is at most $2^{n-1}(3^d - 1)$. Since every such pair is assigned one of three actions, the local strategies of both players can be chosen in at most $3^{2^n(3^d-1)}$ ways. Choosing the output labels of the leaves of the two full 5-ary trees additionally gives

$$N_d \leqslant 3^{2^n(3^d-1)} 2^{2\cdot 5^d} = 2^{(\log_2 3)2^n(3^d-1)+2\cdot 5^d}.$$

Set $d = \lfloor n/\log_2 3 \rfloor - 1$. Then $3^d \leqslant 2^n/3$ and $5^d = o(4^n)$, since $\log_3 5 < 2$. Therefore,

$$\log_2 N_d \leqslant \left(\frac{\log_2 3}{3} + o(1)\right) 4^n.$$

Since there are 2^{4^n} Boolean functions on $\{0,1\}^n \times \{0,1\}^n$, we have

$$\Pr\big[\mathrm{D}^{\mathrm{hd}}_{\mathrm{s}}(f) \leqslant d\big] \leqslant 2^{-\left(1-\frac{\log_2 3}{3}-o(1)\right)4^n} = o(1).$$

Thus, $\mathrm{D}^{\mathrm{hd}}_{\mathrm{s}}(f) \geqslant \lfloor n/\log_2 3 \rfloor$ with probability $1 - o(1)$. Note that in the half-duplex communication model with silence, we have an upper bound $\mathrm{D}^{\mathrm{hd}}_{\mathrm{s}}(f) \leqslant \lceil n/\log_2 3 \rceil + 1$ for every Boolean f via ternary encoding. □

5 Conclusion

In this paper, we obtained new upper bounds on deterministic half-duplex communication complexity. In particular, we proved that

$$\mathrm{D}^{\mathrm{hd}}_{\mathrm{s}}(\mathsf{IP}_n) \leqslant \left\lceil \frac{n}{2} \right\rceil + 2, \qquad \mathrm{D}^{\mathrm{hd}}_{0}(\mathsf{IP}_n) \leqslant \left\lceil \frac{7n}{8} \right\rceil + 4, \qquad \mathrm{D}^{\mathrm{hd}}_{0}(\mathsf{DISJ}_n) \leqslant 0.6615n + O(\log^2 n).$$

We also determined the half-duplex communication complexity of a uniformly random Boolean function: with high probability, it is $n/\log_2 3 + O(1)$ in the model with silence and $n + O(1)$ in the models with zero and with an adversary.

We conclude with the following open problems:

1. Improve the best known bounds for $\mathrm{D}^{\mathrm{hd}}_{0}(\mathsf{IP}_n)$. In contrast to the model with silence, the gap in the model with zero is still substantial.

2. Narrow the gap between the lower and upper bounds for $\mathrm{D}^{\mathrm{hd}}_{0}(\mathsf{DISJ}_n)$ and $\mathrm{D}^{\mathrm{hd}}_{\mathrm{s}}(\mathsf{DISJ}_n)$.

3. For every half-duplex model, find an explicit Boolean function of the maximum communication complexity.

References

[BFS86] Laszlo Babai, Peter Frankl, and Janos Simon. Complexity classes in communication complexity theory. In *27th Annual Symposium on Foundations of Computer Science (sfcs 1986)*, pages 337–347, 1986.

[CLV19] Arkadev Chattopadhyay, Shachar Lovett, and Marc Vinyals. Equality alone does not simulate randomness. In Amir Shpilka, editor, *34th Computational Complexity Conference, CCC 2019, New Brunswick, NJ, USA, July 18-20, 2019*, LIPIcs, pages 14:1–14:11. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[DIS+21] Yuriy Dementiev, Artur Ignatiev, Vyacheslav Sidelnik, Alexander Smal, and Mikhail Ushakov. New bounds on the half-duplex communication complexity. In *SOFSEM 2021: Theory and Practice of Computer Science - 47th International Conference on Current Trends in Theory and Practice of Computer Science, SOFSEM 2021, Bolzano-Bozen, Italy, January 25-29, 2021, Proceedings*, volume 12607 of *Lecture Notes in Computer Science*, pages 233–248. Springer, 2021.

[DV23] Mikhail Dektiarev and Nikolai K. Vereshchagin. Half-duplex communication complexity with adversary? can be less than the classical communication complexity. *Electron. Colloquium Comput. Complex.*, TR23-011, 2023.

[EIRS01] Jeff Edmonds, Russell Impagliazzo, Steven Rudich, and Jirí Sgall. Communication complexity towards lower bounds on circuit depth. *Computational Complexity*, 10(3):210–246, 2001.

[HIMS18a] Kenneth Hoover, Russell Impagliazzo, Ivan Mihajlin, and Alexander Smal. Half-duplex communication complexity. *Electronic Colloquium on Computational Complexity (ECCC)*, 25:89, 2018.

[HIMS18b] Kenneth Hoover, Russell Impagliazzo, Ivan Mihajlin, and Alexander V. Smal. Half-duplex communication complexity. In Wen-Lian Hsu, Der-Tsai Lee, and Chung-Shou Liao, editors, *29th International Symposium on Algorithms and Computation, ISAAC 2018, December 16-19, 2018, Jiaoxi, Yilan, Taiwan*, volume 123 of *LIPIcs*, pages 10:1–10:12. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.

[IMS22] Artur Ignatiev, Ivan Mihajlin, and Alexander Smal. Super-cubic lower bound for generalized Karchmer–Wigderson games. In Sang Won Bae and Heejin Park, editors, *33rd International Symposium on Algorithms and Computation, ISAAC 2022, Seoul, South Korea, December 19-21, 2022*, volume 248 of *LIPIcs*, pages 66:1–66:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

[Kra98] Jan Kraíček. Interpolation by a game. *Mathematical Logic Quarterly*, 44(4):450–458, 1998.

[KRW95] Mauricio Karchmer, Ran Raz, and Avi Wigderson. Super-logarithmic depth lower bounds via the direct sum in communication complexity. *Computational Complexity*, 5(3/4):191–204, 1995.

[KW88] Mauricio Karchmer and Avi Wigderson. Monotone circuits for connectivity require super-logarithmic depth. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2-4, 1988, Chicago, Illinois, USA*, pages 539–550, 1988.

[Mei25] Or Meir. Toward better depth lower bounds: A KRW-like theorem for strong composition. *SIAM J. Comput.*, 54(5):1193–1240, 2025.

[MS21] Ivan Mihajlin and Alexander Smal. Toward better depth lower bounds: The XOR-KRW conjecture. In Valentine Kabanets, editor, *36th Computational Complexity Conference, CCC 2021, July 20-23, 2021, Toronto, Ontario, Canada (Virtual Conference)*, volume 200 of *LIPIcs*, pages 38:1–38:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.

[Wu23] Hao Wu. An improved composition theorem of a universal relation and most functions via effective restriction. *Electron. Colloquium Comput. Complex.*, TR23-151, 2023.

[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*, STOC '79, pages 209–213, New York, NY, USA, 1979. ACM.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il