



New bounds on the half-duplex communication complexity

Yuriy Dementiev* Artur Ignatiev† Vyacheslav Sidelnik‡ Alexander Smal§
Mikhail Ushakov¶

August 4, 2020

Abstract

In this work, we continue the research started in [HIMS18b], where the authors suggested to study the half-duplex communication complexity. Unlike the classical model of communication complexity introduced by Yao, in the half-duplex model, Alice and Bob can speak or listen simultaneously, as if they were talking using a walkie-talkie. The motivation for such a communication model comes from the study of the KRW conjecture. Following the open questions formulated in [HIMS18b], we prove lower bounds for the disjointness function in all variants of half-duplex models and an upper bound in the half-duplex model with zero, that separates disjointness from the inner product function in this setting. Next, we prove lower and upper bounds on the half-duplex complexity of the Karchmer-Wigderson games for the counting functions and for the recursive majority function, adapting the ideas used in the classical communication complexity. Finally, we define the non-deterministic half-duplex complexity and establish bounds connecting it with non-deterministic complexity in the classical model.

Keywords: communication complexity · half-duplex communication · Karchmer-Wigderson games.

1 Introduction

1.1 Background

Communication complexity is a powerful tool with applications in algorithms, circuit complexity, proof complexity, and many other areas of theoretical computer science. In the classical model of communication complexity introduced by Yao [Yao79], two players, Alice and Bob, try to compute $f(x, y)$ for some function f , where Alice only knows x and Bob only knows y . Alice and Bob can communicate by sending bits to each other, one bit per round, and at the end of the communication, both players must know the result $f(x, y)$. The essential property of this classical model is that in each round of communication, one player sends a bit, and the other receives it.

*St.Petersburg State University, yuriy.dementiev2000@gmail.com

†St.Petersburg State University, artur.ignatiev23924@gmail.com

‡St.Petersburg State University, slava.sidelnik@gmail.com

§St.Petersburg Department of Steklov Mathematical Institute of Russian Academy of Sciences, smal@pdmi.ras.ru

¶St.Petersburg State University, mikushakov@gmail.com

There are many extensions of this basic two-party communication model, such as randomized communication complexity, non-deterministic communication complexity, various types of multi-party communication models, etc. In [HIMS18b], the authors suggested considering a communication model where the players speak 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 two persons try to speak simultaneously, then they do not hear each other. Formally speaking, every round, each player chooses one of three actions: send 0, send 1, or receive. There are three different types of rounds: a *classical round*, when one player sends some bit while the other receives, a *spent round*, when both players send bits (these bits get lost), and a *silent round*, when both players receive. In [HIMS18b], the authors defined three variations of the half-duplex model based on what happens in silent rounds: half-duplex models with silence, with zero, and with adversary (see Section 2 for more information).

It turned out that the communication complexity in the half-duplex models not only differs from the classical case, but also behaves differently. For example, in the classical case, the equality function, the disjointness function, and the inner product function have complexity $n + 1$, meaning that all three are the hardest functions. In the half-duplex models with silence and with zero, these three functions are of different complexity.

The original motivation to study the half-duplex communication models comes from the study of the Karchmer-Wigderson games [KW88] for the multiplexer relation [EIRS01]. In [MS20], results from the half-duplex communication complexity were used to prove a lower bound on a composition of the universal relation with the Karchmer-Wigderson game for some function. The authors suggest that a better understanding of the half-duplex communication complexity might help to achieve new bound in the study of the KRW conjecture and even prove a supercubic lower bound on the De Morgan formula size of an implicit function.

We continue the research started [HIMS18b], and close some open questions from it regarding the complexity of the disjointness function. We also study the complexity of the Karchmer-Wigderson games for the counting functions and for the recursive majority function. In addition, we define the non-deterministic half-duplex communication complexity and prove bounds connecting it to the classical non-deterministic communication complexity.

1.2 Overview of results

For a communication problem R , let $D_s^{hd}(R)$, $D_0^{hd}(R)$, and $D_a^{hd}(R)$, denote the half-duplex communication complexity of R with silence, with zero, and with adversary, respectively (see Section 2 for formal definitions). Table 1 contains a summary of lower and upper bounds for the communication problems considered in [HIMS18b], the bounds proved in this paper are marked with \star .

Table 1: Lower and upper bounds for the communication problems considered in [HIMS18b].

	EQ_n	IP_n	$DISJ_n$	KW_{MOD2}
D_s^{hd}	$\geq n/\log 5$ $\leq n/\log 5 + o(n)$	$\geq n/1.67$	$\geq n/\log 5$ \star $\leq n/2 + O(1)$	$\geq 1.12 \log n$ \star $\leq 1.262 \log n$ \star
D_0^{hd}	$\geq n/\log 3$ $\leq n/\log 3 + o(n)$	$\geq n/1.234$	$\geq n/\log 3$ \star $\leq 3n/4 + o(n)$ \star	$\geq 1.62 \log n$ \star $\leq 1.893 \log n$ \star
D_a^{hd}	$\geq n/\log 2.5$	$\geq n$	$\geq n/\log 2.5$ \star	$= 2 \log n$

In addition to the bounds in Table 1, we prove the following upper bounds for special cases of the counting function,

$$D_s^{hd}(\text{KW}_{\text{MOD}3}) \leq 1.893 \log n, \quad D_s^{hd}(\text{KW}_{\text{MOD}5}) \leq 2.46 \log n, \quad D_s^{hd}(\text{KW}_{\text{MOD}11}) \leq 3.48 \log n,$$

and for the recursive majority function,

$$D_s^{hd}(\text{KW}_{\text{RecMaj}}) \leq 2 \log_3 n, \quad D_0^{hd}(\text{KW}_{\text{RecMaj}}) \leq 2 \log_3 n.$$

For arbitrary $p \geq 2$, we prove that $D_s^{hd}(\text{KW}_{\text{MOD}p}) \leq (1 + \lceil \log_3 \frac{p}{2} \rceil) / \left(\log \left(\frac{2}{\sqrt{5}-1} \right) \right) \cdot \log n$. We also prove that the lower bounds for $\text{KW}_{\text{MOD}2}$ in Table 1 apply for $\text{KW}_{\text{MOD}p}$ for arbitrary p in all three models.

In Section 5, we introduce non-deterministic half-duplex communication complexity. Let $N_s^{hd}(f)$, $N_0^{hd}(f)$, and $N_a^{hd}(f)$ denotes the non-deterministic half-duplex communication complexity of f with silence, with zero, and with adversary, respectively. For any function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, we show that

$$N_s^{hd}(f) = N(f)/\log 5 + \Theta(\log N(f)), \quad N_0^{hd}(f) = N(f)/\log 3 + \Theta(\log N(f)), \quad N_a^{hd}(f) \geq N(f)/\log 3,$$

where $N(f)$ denotes the classical non-deterministic communication complexity of f .

1.3 Organization of this paper

In Section 2, we review the half-duplex communication models. In Section 3, we prove lower and upper bounds for the disjointness function in various half-duplex models. Then, in Section 4, we prove lower and upper bounds on the Karchmer-Wigderson games for the counting functions and for the recursive majority function. In Section 5, we define the non-deterministic half-duplex communication complexity and prove bounds connecting it to the classical non-deterministic complexity. Section 6 contains open problems.

2 Half-duplex Communication Complexity

We expect that the reader is familiar with the standard definitions of communication complexity that can be found in [KN97]. It will be necessary to understand that a communication protocol can be described by a binary tree, where every node has an associated combinatorial rectangle of all input pairs (x, y) such that if the players are given x and y , then the communication goes through this node. Finally, we expect the reader to understand why the rectangles associated with the leaves of the protocol tree are monochromatic.

Let's assume that the players have some synchronizing mechanism, e.g., synchronized clocks, that allows them to understand when each round begins. In the half-duplex communication model, every round, each player chooses one of three actions: send 0, send 1, or receive. So, there are three different types of rounds.

- *A classical round:* one player sends some bit and the other one receives it.
- *A spent round:* both players send bits, and these bits get lost.
- *A silent round:* both players receive.

In [HIMS18b], the authors considered three variations of this model based on what happens in silent rounds.

- *The half-duplex communication model with silence.* In a silent round, both players receive a special symbol `silence`, so it is possible for both players to distinguish a silent round from a classical one.
- *The half-duplex communication model with zero.* In a silent round, both players receive 0, i.e., players can not distinguish a silent round from a classical round where the other player sends 0.
- *The half-duplex communication model with adversary.* In a silent round, each player receives some arbitrary bit, not necessarily the same as the other player.

Note that in the half-duplex model with zero there is no need to send zeros — a player can choose to receive instead and the other player will not notice the difference.

Unlike the classical case, in the half-duplex communication models a player does not always know what the other player’s action was — the information about it can be “lost”, i.e., in spent rounds a player do not know what the other player’s action was. It means that a player might not know what node of the protocol corresponds to the current state of communication. Instead of a one protocol tree, like in the classical case, in the half-duplex case, a protocol is described by two trees — one for each player. The protocol trees have arity 5 in the half-duplex model with silence, arity 3 in the model with zero, and arity 4 in the model with adversary. For the formal definition of the half-duplex communication protocols see [HIMS18b]. It should be noted that despite the differences, every node of a half-duplex protocol tree has an associated rectangle, and every leaf rectangle is monochromatic.

The minimal number of rounds that is enough to solve a communication problem R on all inputs defines the communication complexity of R . For the classical communication model, we denote it by $D(R)$, for the half-duplex models with silence, with zero, and with adversary, we denote it by $D_s^{hd}(R)$, $D_0^{hd}(R)$, and $D_a^{hd}(R)$, respectively.

3 Bounds for the disjointness function

The disjointness function $\text{DISJ}_n : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, such that $\text{DISJ}_n(x, y) = 1 \iff \forall i : x_i \neq 1 \vee y_i \neq 1$, is one of the hardest functions in the classical communication model — it has communication complexity exactly $n + 1$. In other words, the trivial protocol where Alice sends all her bits to Bob, Bob computes the result and then sends it back to Alice, is optimal. This is not the case in the half-duplex models. In [HIMS18b], the authors prove an upper bound $n/2 + O(1)$ on DISJ_n in the half-duplex model with silence (Theorem 16 in [HIMS18b]). Note that in this model the inner product function requires a protocol of complexity at least $n/1.67$, so the disjointness is not the hardest function in this setting. In this section we prove lower bounds for DISJ_n in all three half-duplex models, and an upper bound in the model with zero showing that the disjointness is not the hardest function there as well.

We start with the lower bounds.

Theorem 1. For all $n \in \mathbb{N}$, $D_s^{hd}(\text{DISJ}_n) \geq n/\log 5$, $D_0^{hd}(\text{DISJ}_n) \geq n/\log 3$, $D_a^{hd}(\text{DISJ}_n) \geq n/\log 2.5$.

To prove this theorem, we will need the following folklore property of communication protocols for DISJ_n . For a Boolean vector $x \in \{0, 1\}^n$, let \bar{x} denotes its complement, i.e. $x_i = 1 - \bar{x}_i$ for all $i \in [n]$.

Lemma 2 ([KN97]). *Let $x, y \in \{0, 1\}^n$ and $x \neq y$. The pair of inputs (x, \bar{x}) and (y, \bar{y}) do not belong to the same monochromatic rectangle of DISJ_n .*

Proof. Note that $\text{DISJ}_n(x, \bar{x}) = \text{DISJ}_n(y, \bar{y}) = 1$. Consider a rectangle containing (x, \bar{x}) and (y, \bar{y}) . Due to combinatorial rectangle properties it necessarily contains also (x, \bar{y}) and (y, \bar{x}) . If $y \subset x$ then $\text{DISJ}_n(x, \bar{y}) = 0$; otherwise, $\text{DISJ}_n(\bar{x}, y) = 0$. Hence this rectangle is not monochromatic. \square

This property of DISJ_n allows us to define a sub-additive measure $\mu(R)$ that is equal to the number of pairs (x, \bar{x}) in rectangle R . In [HIMS18b], a special framework was developed specifically for such measures. The following lemma allows to get a lower bound by showing that for every protocol the measure of the root rectangle is large while the measure of all leaf rectangles is small.

Lemma 3 (Lemma 10 in [HIMS18b]). *Let μ be some sub-additive measure on rectangles such that $\mu(X \times Y) \geq \mu_r$ and for any leaf rectangle R_l , $\mu(R_l) \leq \mu_\ell$. If for any rectangle R there is always a good subrectangle for function $f \upharpoonright R$ of measure at least $\alpha \cdot \mu(R)$ then the depth of the protocol is at least $\log_{1/\alpha} \frac{\mu_r}{\mu_\ell}$.*

Now we are ready to prove Theorem 1.

Proof of Theorem 1. For the first two lower bounds it is enough to use *the fooling set method* [KN97]. There are 2^n pairs (x, \bar{x}) , and hence there are at least 2^n 1-monochromatic rectangles, that gives a lower bound on the number of leaves in the protocol. It remains to note that the protocol trees in the half-duplex models with silence and with zero have arities 5 and 3, respectively.

The third lower bound requires a little bit more effort. The same argument would prove only $n/2$ lower bound which is trivial. Instead, we use *the rectangle elimination technique* introduced in [HIMS18b]. The proof uses the fact that any rectangle R can be covered by a set of five *good rectangles* that cover the entire rectangle twice, and hence one of rectangles has measure at least $2/5 \cdot \mu(R)$ (see [HIMS18a, Theorem 13] for more information). Application of Lemma 3 for μ and $\alpha = 2/5$ concludes the proof. \square

Now we proceed to the upper bound for DISJ_n in the half-duplex model with zero. In order to show a protocol with $\frac{3}{4}n + o(n)$ rounds, we start with a less efficient protocol and then improve it. Let us remind that in the half-duplex model with zero there is no need to send zeros, so the players never do it.

Theorem 4. *For all $n \in \mathbb{N}$, $D_0^{hd}(\text{DISJ}_n) \leq 5n/6 + O(\log n)$.*

Proof. W.l.o.g., we assume that n is even. Consider the following protocol. Alice and Bob split their input strings into blocks of size 2, so each player has $n/2$ such blocks. Let $\#(ab)$ denotes the number of blocks ab among the blocks of Bob. Note that the following three cases are possible:

1. $\#(00) \geq n/6$, then $\#(01) + \#(10) + \#(11) \leq n/2 - n/6 = n/3$,
2. $\#(01) \geq n/6$, then $\#(00) + \#(10) + \#(11) \leq n/2 - n/6 = n/3$,
3. $\#(00) + \#(01) < n/3$.

At the beginning Bob must check which of the cases applies and tell Alice using two bits of communication. Further communication will depend on it. We will show that in all the cases the players can solve the problem using at most $5n/6 + O(\log n)$ rounds. The communication will be divided into two stages. In the first stage, Alice and Bob process their input considering one pair of corresponding blocks per round. Each round they act as it is described in Table 2. Alice and Bob need to be able to distinguish the situations when the corresponding blocks intersect, that is, to distinguish block pairs (01, 01), (01, 11), (11, 01), (11, 10), (11, 11), (10, 10), (10, 11) from others.

Table 2: The first stage of communication of the half-duplex protocol for DISJ_n with zero.

Case 1			Case 2			Case 3		
Block	Alice	Bob	Block	Alice	Bob	Block	Alice	Bob
00	receive	send(1)	00	receive	receive	00	send(1)	receive
01	send(1)	receive	01	send(1)	send(1)	01	receive	receive
10	receive	receive	10	receive	receive	10	receive	send(1)
11	receive	receive	11	receive	receive	11	receive	receive

Case 1. After $n/2$ rounds, Bob tells Alice whether he ever has received 1 while processing a block 11 or 01. If he has, then the corresponding block of Alice was 01, and hence their strings intersect. This corresponds to identifying (01, 01) and (01, 11). Further, Alice tells Bob whether she ever has a silent round while processing 11. If she has, then the corresponding block of Bob was one of 01, 10 or 11, and hence their strings intersect, so the players identified block pairs (11, 01), (11, 10), and (11, 11). If any intersecting blocks have been found, the players stop the communication and output 0. Otherwise, they proceed to the second stage.

In the second stage, to identify the remaining two block pairs (10, 10) and (10, 11), Alice needs to distinguish zeros she received when Bob was processing 10 or 11, and zeroes she received when Bob was processing 01. Bob sequentially (starting from his first silent round on 01, 10 or 11) goes through all his blocks 01, 10, 11 corresponding to silent rounds. When he processes 10 and 11, he sends 1, and when he processes 01, he sends 0. Alice knows how many bits Bob will send her, since he sends one bit for every silent round. Alice simultaneously processes her blocks corresponding to the silent rounds. Now for every such blocks she knows whether Bob has 10 or 11, and 01. If she receives 1 having 10 then they identified (10, 10) and (10, 11). At the end of the protocol, Alice needs one more round to tell Bob whether she has found an intersection.

The complexity of the first stage is $n/2 + O(1)$, the complexity of the second stage is $n/3 + O(1)$, so the total complexity is $5n/6 + O(1)$.

Case 2. After $n/2$ rounds, Bob tells Alice whether he has ever received 1 while processing 11. If he has, then Alice had 01 in the corresponding block, and hence their strings intersect, so the players identified (01, 11). After that, Alice tells Bob whether she has ever received 1 while processing 11. If she has, then they identified (11, 01). Next, Bob tells the number of rounds in which he sent 1. Alice compares it with the number of 1 she received. If these two numbers are equal, then there were no spent rounds in the first stage, otherwise the players identified a block pair (01, 01). If any intersecting blocks have been found, the players stop the communication and output 0. Otherwise, they proceed to the second stage.

In the second stage, to identify the remaining four block pairs (10, 10), (10, 11), (11, 10) and (11, 11), Alice needs to distinguish zeroes received when Bob was processing 10 or 11 from zeros received when Bob was processing 00. Bob sequentially (starting from his first silent round on 00, 10 or 11) goes through all his blocks 00, 10, 11 corresponding to silent rounds. When he processes 10 and 11, he sends 1, and when he processes 00, he sends 0. Similarly to the previous case, if she receives 1 having 10 or 11 then they identified one of the desired pairs of blocks. At the end of the protocol Alice needs one more round to tell Bob whether she has found an intersection.

The complexity of the first stage is $n/2 + O(1)$, the complexity of the second stage is $n/3 + O(\log n)$, so the total complexity is $5n/6 + O(\log n)$.

Case 3. After $n/2$ rounds, Bob tells Alice if there was a silent round corresponding to his block 11. If there was such a round, the Alice had 01, 10 or 11 in the corresponding block, so the players identified block pairs (01, 11), (10, 11), and (11, 11). Next, Alice tells Bob if she has ever received 1 while processing 10 or 11. If she has, then Bob had 10, and hence they identified a block (10, 10) or (11, 10). If any intersecting blocks have been found, the players stop the communication and output 0. Otherwise, they proceed to the second stage.

In the second stage, to identify the remaining two block pairs (01, 01) and (11, 01), Alice needs to distinguish zeroes received while Bob was processing a block 00 and zeroes received while Bob was processing a block 01. Bob sequentially (starting from his first silent round on 00 or 01) goes through all his blocks 00 and 01 corresponding to silent rounds. When he processes a block 01, he sends 1, and when he processes 00, he sends 0. Similarly to the previous cases, if Alice receives 1 while processing 01 or 11, she identifies one of the desired block pairs.

Similarly to the other cases, the complexity of the protocol in this case is $5n/6 + O(1)$. \square

The protocol proposed in Theorem 4 can be improved if we notice that reiterating over a part of the blocks that happens in the second stage can be reduced to solving disjointness problem on smaller inputs.

Theorem 5. For all $n \in \mathbb{N}$, $D_0^{hd}(\text{DISJ}_n) \leq 3n/4 + o(n)$.

Proof. We are going to modify the protocol from the proof of Theorem 4. In the modified protocol, the players consider the same three cases, and they have the same first stages in all cases. The second stage is different. Instead of reiterating all blocks corresponding to silent rounds, Alice and Bob reduce this problem to solving disjointness on strings of size at most $n/3$, and then run the same protocol for disjointness recursively. Thus, we get the following bound

$$D_0^{hd}(\text{DISJ}_n) \leq \sum_{i=0}^{\lceil \log_3(n) \rceil} \frac{n}{2 \cdot 3^i} + o(n) \leq \sum_{i=0}^{\infty} \frac{n}{2 \cdot 3^i} + o(n) = \frac{3n}{4} + o(n).$$

It remains to understand how the second stage works in each of the cases.

Case 1. To identify two block pairs (10, 10) and (10, 11), Bob writes down a string x' that has one bit for every silent rounds: 1 for a block 10 or 11, and 0 for 01. Similarly, Alice writes down a string y' that has one bit for every silent round: 1 for a block 10, and 0 for other blocks. It is not hard to see, that $\text{DISJ}(x', y') = 0$ if and only if there was a silent round corresponding to (10, 10) or (10, 11).

Case 2. To identify four block pairs (10, 10), (10, 11), (11, 10) and (11, 11), Bob writes down a string x' that has one bit for every silent rounds: 1 for a block 10 or 11, and 0 for 00. Similarly, Alice writes down a string y' using the same rules. It is not hard to see, that $\text{DISJ}(x', y') = 0$ if and only if there was a silent round corresponding to one of the desired block pairs.

Case 3. To identify the remaining two block pairs (01, 01) and (11, 01), Bob writes down a string x' that has one bit for every silent rounds: 1 for a block 01, and 0 for others (at this point, he already knows that there were no silent rounds corresponding to 11). Similarly, Alice writes down a string y' that has one bit for every silent round: 1 for a block 01 or 11, and 0 for 10. And again, it is not hard to see, that $\text{DISJ}(x', y') = 0$ if and only if there was a silent round corresponding (01, 01) and (11, 01). \square

It is worth noting that this theorem shows that $D_0^{hd}(\text{DISJ}_n) < D_a^{hd}(\text{DISJ}_n)$. Moreover, we separate DISJ_n and IP_n in the half-duplex model with zero, $D_0^{hd}(\text{DISJ}_n) < D_0^{hd}(\text{IP}_n)$, showing that disjointness is not the hardest function in this model.

4 Bounds on the Karchmer-Wigderson games

The seminal work of Karchmer and Wigderson [KW88] established a correspondence between De Morgan formulas for non-constant Boolean function f and communication protocols for the Karchmer-Wigderson game for f . *The Karchmer-Wigderson game* for Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is the following communication problem: Alice gets an input $x \in \{0, 1\}^n$ such that $f(x) = 0$, and Bob gets as input $y \in \{0, 1\}^n$ such that $f(y) = 1$. Their goal is to find a coordinate $i \in [n]$ such that $x_i \neq y_i$. The Karchmer-Wigderson game for f can be considered as a communication problem for *the Karchmer-Wigderson relation for f* :

$$\text{KW}_f = \{(x, y, i) \mid x, y \in \{0, 1\}^n, i \in [n], f(x) = 0, f(y) = 1, x_i \neq y_i\}.$$

Karchmer and Wigderson showed that the communication complexity of KW_f is exactly equal to the depth formula complexity of f . This correspondence allows us to use communication complexity methods for proving formula depth lower bounds.

In this section, we prove bounds on the half-duplex communication complexity of the Karchmer-Wigderson games for *the counting function* $\text{MOD}_p : \{0, 1\}^n \rightarrow \{0, 1\}$, such that

$$\text{MOD}_p(x) = 0 \iff x_1 + \dots + x_n = 0 \pmod{p},$$

and for *the recursive majority function* $\text{RecMaj}_n : \{0, 1\}^n \rightarrow \{0, 1\}$, such that

$$\text{RecMaj}_n(x_1, \dots, x_n) = \text{Maj}_3(\text{RecMaj}_{\frac{n}{3}}(x_1, \dots, x_{\frac{n}{3}}), \text{RecMaj}_{\frac{n}{3}}(x_{\frac{n}{3}+1}, \dots, x_{\frac{2n}{3}}), \text{RecMaj}_{\frac{n}{3}}(x_{\frac{2n}{3}+1}, \dots, x_n)),$$

where $\text{Maj}_3 : \{0, 1\}^3 \rightarrow \{0, 1\}$ is majority of three bits, and n is a power of three. This does not immediately imply any bounds for De Morgan formulas, however, similar results can be potentially used to prove bounds in the classical case as it was done in [MS20].

We start with lower bounds for MOD_p functions. In the classical case, a lower bound on the communication complexity of the Karchmer-Wigderson game for some function f corresponds to a lower bound on De Morgan formula depth complexity of f . For MOD_2 , the parity function, we have the tight bound $2 \log n$: the lower bound is due to the famous work of Khrapchenko [Khr71], and

the upper bound is straightforward by implementing binary search. The method of Khrapchenko can also be used to prove $2 \log n - O(1)$ lower bounds for $\text{MOD}p$ for arbitrary p . In [HIMS18b], the authors proved $2 \log n$ lower bound for $\text{KW}_{\text{MOD}2n}$ in the half-duplex model with adversary. We use the same ideas to prove general lower bounds for $\text{MOD}p$ in all the half-duplex models.

Theorem 6. *For any $p \geq 2$,*

$$D_s^{hd}(\text{KW}_{\text{MOD}p}) > 1.12 \log n, \quad D_0^{hd}(\text{KW}_{\text{MOD}p}) > 1.62 \log n, \quad D_a^{hd}(\text{KW}_{\text{MOD}p}) \geq 2 \log n - O(1).$$

We prove this theorem using the information-theoretic approach from [HIMS18b]. For basic definitions of information theory, we refer to [CT06]. The high-level idea of the proof is simple. We show that there is a random distribution over the inputs of Alice and Bob, such that at the beginning of the communication each player has uncertainty roughly $\log n$ bits about the input of the other player. At the end of the communication (for this specific distribution) each player necessarily knows the input of the other player. This means that during the protocol each player learns roughly $\log n$ bits of information. For the classical communication model, this would be enough to show $2 \log n$ lower bound, because in every round only one of the players learns at most one bit of information, so there must be at least $2 \log n$ rounds. In the half-duplex communication, the situation is more complicated — in silent rounds both players might learn some information. To estimate the amount of information the players learn during half-duplex communication, we use the upper bounds proved in [HIMS18b].

Proof. Let $(\mathcal{X}, \mathcal{Y})$ be a pair of jointly distributed random variables where \mathcal{X} is uniformly distributed on a set

$$\{x \in \{0, 1\}^n \mid \text{MOD}p(x) = 0, n/4 < \|x\|_1 < 3n/4\},$$

that corresponds to all Alice's inputs of Hamming weight between $n/4$ and $3n/4$, and \mathcal{Y} is obtained from \mathcal{X} by flipping one 0 bit uniformly at random. Thus, $H(\mathcal{Y} \mid \mathcal{X}) \geq \log(n/4)$ and $H(\mathcal{X} \mid \mathcal{Y}) \geq \log(n/4)$. Before any communication takes place $H(\mathcal{Y} \mid \mathcal{X}) + H(\mathcal{X} \mid \mathcal{Y}) \geq 2 \log n - O(1)$. Given inputs from the distribution $(\mathcal{X}, \mathcal{Y})$ the players have to find the unique bit of difference.

Let \mathcal{P} be a protocol for $\text{MOD}p$. W.l.o.g., we assume that all the leafs of \mathcal{P} are on the same depth. For any natural k , let Π_A^k and Π_B^k be the marginal distributions over Alice's and Bob's partial transcripts after running \mathcal{P} for k rounds induced by $(\mathcal{X}, \mathcal{Y})$. If the protocol \mathcal{P} has depth d , then $H(\mathcal{Y} \mid \mathcal{X}, \Pi_A^d) + H(\mathcal{X} \mid \mathcal{Y}, \Pi_B^d) = 0$. This means that the players have to learn $2 \log n - O(1)$ bits during the communication.

Now we apply upper bounds on the amount of information that can be learned in one round. For the half-duplex model with silence, there exists a distribution on the inputs such that the players can learn at most 1.67 bits per round [HIMS18a, Theorem 18], so we get $D_s^{hd}(\text{KW}_{\text{MOD}p}) > (2 \log n)/1.67 > 1.12 \log n$. For the half-duplex model with zero, there exists a distribution on the inputs such that the players can learn at most 1.234 bits per round [HIMS18a, Theorem 21], hence using the same reasoning, $D_0^{hd}(\text{KW}_{\text{MOD}p}) > (2 \log n)/1.234 > 1.62 \log n$. Finally, for the half-duplex model with adversary, there exists a distribution on the inputs such that the players can learn at most one bit per round [HIMS18a, Theorem 24], that concludes the proof. \square

Now we proceed to the upper bounds. First, we will consider a few special cases of the counting function, and then we will prove the general upper bound for arbitrary $p \geq 2$. The following two theorems establish upper bounds for $\text{MOD}2$ in the half-duplex models with silence and with zero. Both protocols will use the idea of ternary search but in a little bit different manner.

Theorem 7. $D_s^{hd}(\text{KW}_{\text{MOD}2}) \leq 2 \log_3 n + O(1) < 1.262 \log n$.

Proof. Alice and Bob split their input strings into three equal parts and compute MOD2 for all the resulting substrings. There are four possible cases for each player.

- Parities of the Alice's substrings: 1) 000; 2) 011; 3) 101; 4) 110.
- Parities of the Bob's substrings: 1) 111; 2) 100; 3) 010; 4) 001.

Using two rounds of communication the players determine which pair of corresponding substrings have different parities, and then repeat the protocol recursively for these substrings of size $n/3$. That gives the desired bound, $D_s^{hd}(\text{MOD}2) \leq 2 \log_3 n + O(1)$. These two rounds are described in Table 3.

Table 3: The half-duplex protocol for $\text{KW}_{\text{MOD}2}$ with silence.

Round 1			Round 2		
Case	Alice	Bob	Case	Alice	Bob
1	receive	receive	1	send 0	send 1
2	receive	receive	2	send 1	send 0
3	send 0	send 1	3	receive	receive
4	send 1	send 0	4	receive	receive

If the first round or the second round was silent, i.e., the players received the symbol of silence, then they know that the first substrings have different parities. Otherwise, if none of the rounds was silent, then there were no spent rounds (if the first was spent, then the second necessarily was silent, and vice versa, if the second was spent, then the first was silent). In this case, the players know that they have the same parity of the first substrings, and each player have sent one bit that was received by the other player. Moreover, these bits correspond to the parities of the second substrings. So, both Alice and Bob know the parities of the first two substrings of the other player's input, hence both players know the parities of all the substrings. \square

Theorem 8. $D_0^{hd}(\text{KW}_{\text{MOD}2}) \leq 3 \log_3 n + O(1) < 1.893 \log n$.

Proof. Similarly to the proof of Theorem 7, Alice and Bob split their input strings into three equal parts and compute MOD2 for all the resulting substrings. There are the same four possible cases for each player.

- Parities of the Alice's substrings: 1) 000; 2) 011; 3) 101; 4) 110.
- Parities of the Bob's substrings: 1) 111; 2) 100; 3) 010; 4) 001.

The players use three rounds to determine which pair of corresponding substrings have different parities, and then repeat the protocol recursively for these substrings of size $n/3$. That gives the desired bound $D_0^{hd}(\text{MOD}2) \leq 3 \log_3 n + O(1)$.

In the first round, Alice and Bob send 1 if they both have Cases 1, otherwise they receive.

- If the first round was silent (i.e., both players received 0) then they both know that none of them have Case 1. In the second round, Alice sends 1 if her first substring is even (Case 2), otherwise she receives. Bob does the reverse, he sends 1 if his first substring is odd (Case 2), otherwise he receives. The third round is similar, Alice and Bob send 1 in their Cases 3, otherwise they receive.

- If at least one of them sent 1 in the first round, then again both know about it. Alice sends 1 if her first substring is odd, otherwise she receives. Bob does the reverse, he sends 1 if his first substring is even, otherwise he receives. In the third round, the players do the same for the second substrings.

If the second or the third rounds were silent, then, respectively, the first or the second substrings have different parities. Otherwise, the third substrings have different parities. \square

The next theorem considers the MOD3 function. In the classical case, the best known upper bound for it is $2.881 \log n$ [Chi90]. We show a simple upper bound based on the idea of ternary search.

Theorem 9. $D_s^{hd}(\text{KW}_{\text{MOD}3}) \leq 3 \log_3 n + O(1) \leq 1.893 \log n$.

Proof. Similarly to the protocols for MOD2, Alice and Bob split their inputs into three equal parts and for all the resulting substrings compute the number of ones modulo 3. Then they spend three rounds to find a pair of corresponding substrings with different remainders modulo 3, and run the protocol recursively on these substrings. Thus, we get the desired bound. In the first two rounds, Alice sends her remainders for the first and the second substrings naturally encoded in a ternary alphabet $\{0, 1, \text{silence}\}$. Bob compares the received numbers with the remainders of his first two substrings, decides on which part they should proceed, and sends this number to Alice in the third round using the same encoding. \square

The next function, that we are going to consider — the MOD5 function. The best known upper bound in the classical case is $3.475 \log n$ [Chi90]. For this and all the following upper bounds for MOD p functions, we adapt the prefix code technique used in [Chi90].

Theorem 10. $D_s^{hd}(\text{KW}_{\text{MOD}5}) \leq 2.46 \log n$.

Proof. Alice and Bob divide their inputs into two part: the first of length εn , and the second of length $(1 - \varepsilon)n$, and compute *the remainder* for every resulting substring, that is the number of ones modulo 5. During the protocol the players will narrow the search area from the current string to one of its substrings over and over again. At the beginning of the communication, Alice sends the remainder of her first substring to Bob. Then the players speak in turns, starting with Bob, sending each other pairs (r, b) , where r is a remainder and b is a bit flag. Every turn, the speaking player does the following.

- (except for the very first turn) If the bit flag b in the previous was set to 0 then the player narrows the search area to the first substring, otherwise the player narrows the search area to the second substring. Then the player subdivides the new search area into two parts in proportion $\varepsilon : (1 - \varepsilon)$.
- The player choose one of the substrings and narrows the search area to it. If the remainder r received in the previous message is equal to the remainder of the first substring, then the player choose the second substring, otherwise the player choose the first one. The player subdivides the new search area into two parts in proportion $\varepsilon : (1 - \varepsilon)$.
- The player sends message (r, b) , where r is the remainder of the first substring of the current search area, and b is set to 1 if and only if the second substring was chosen in the previous step.

It remains for us to discuss, how exactly the pair (r, b) pair is encoded. The encoding is the key ingredient of this technique. We are going to use the following prefix-free code in ternary alphabet.

$b \setminus r$	0	1	2	3	4
0	00	01	0s	10	11
1	s00	s01	s0s	s10	s11

Note that “s” stands for “silence”. The code is chosen such that every message with $b = 0$ has encoding of length 2, and every message with $b = 1$ has encoding of length 3.

To estimate the number of rounds, we need to solve the following system of recurrent relations, where $T(n)$ stands for the number of rounds.

$$\begin{cases} T(n) = 2 + T(\varepsilon n) \\ T(n) = 3 + T((1 - \varepsilon)n) \end{cases} \implies \begin{cases} T(n) = 2 \log_{\frac{1}{\varepsilon}} n \\ T(n) = 3 \log_{\frac{1}{1-\varepsilon}} n \end{cases}$$

From $\frac{2}{3} = \frac{\log \varepsilon}{\log(1-\varepsilon)}$ we get $\varepsilon < 0.57$, and hence $T(n) < 2.466 \log n$. \square

Following [Chi90], we consider the MOD11 function with the best known upper bound is $4.93 \log n$ in the classical case.

Theorem 11. $D_s^{hd}(\text{KW}_{\text{MOD11}}) \leq 3.48 \log n$.

Proof. The proof is almost identical to proof of Theorem 10. The players use a prefix-free encoding, such that every message with $b = 0$ has encoding of length 3, and every other message has encoding of length 4.

$b \setminus r$	0	1	2	3	4	5	6	7	8	9	10
0	000	001	010	011	00s	0s0	01s	0s1	0ss	100	101
1	s000	s001	s010	s011	s00s	s0s0	s01s	s0s1	s0ss	s100	s101

That leads to the following system of recurrent relations:

$$\begin{cases} T(n) = 3 + T(\varepsilon n) \\ T(n) = 4 + T((1 - \varepsilon)n) \end{cases} \implies \begin{cases} T(n) = 3 \log_{\frac{1}{\varepsilon}} n \\ T(n) = 4 \log_{\frac{1}{1-\varepsilon}} n \end{cases}$$

From $\frac{3}{4} = \frac{\log \varepsilon}{\log(1-\varepsilon)}$ we get $\varepsilon < 0.55$, and hence $T(n) < 3.48 \log n$. \square

To generalize the previous results, we prove the following upper bound for MOD p for arbitrary $p \geq 2$. Note that this upper bound is not useful for some p , such as 2, 3 or 7. Moreover, in the classical case, the corresponding bound [Chi90] is surpassed by the upper bound on all symmetric functions [BH96] starting with some p . We expect that the similar happens in the half-duplex model with silence.

Theorem 12. For all $p \geq 2$, $D_s^{hd}(\text{KW}_{\text{MOD}p}) \leq \frac{1 + \lceil \log_3 \frac{p}{2} \rceil}{\log(\frac{2}{\sqrt{5}-1})} \log n$.

Proof. The protocol is similar to the protocol from the proof of Theorem 10. The prefix-free code is chosen such that for $b = 0$ the encoding starts with 0 or 1, and for $b = 1$ the encoding starts with \mathbf{s} . The length of the encoding for $b = 0$ is $\lceil \log_3 \frac{3p}{2} \rceil$, for $b = 1$ is $1 + \lceil \log_3 p \rceil$. Thus, we get the following system:

$$\begin{cases} T(n) = 1 + \lceil \log_3 \frac{p}{2} \rceil + T(\varepsilon n) \\ T(n) = 1 + \lceil \log_3 p \rceil + T((1 - \varepsilon)n) \end{cases} \implies \begin{cases} T(n) = (1 + \lceil \log_3 \frac{p}{2} \rceil) \log_{\frac{1}{\varepsilon}} n \\ T(n) = (1 + \lceil \log_3 p \rceil) \log_{\frac{1}{1-\varepsilon}} n \end{cases}$$

$$\frac{1 + \lceil \log_3 \frac{p}{2} \rceil}{1 + \lceil \log_3 p \rceil} = \frac{\log \varepsilon}{\log(1 - \varepsilon)}$$

Note that the right hand side of the last equation decreases with ε , and the left hand side increases with p . So it is easy to see, that $\frac{1 + \lceil \log_3 \frac{p}{2} \rceil}{1 + \lceil \log_3 p \rceil} \in [\frac{1}{2}, 1)$, and thus $\varepsilon \in (\frac{1}{2}, \frac{\sqrt{5}-1}{2}]$. That gives us the desired bound $T(n) \leq \frac{1 + \lceil \log_3 \frac{p}{2} \rceil}{\log(\frac{2}{\sqrt{5}-1})} \log n$. \square

To conclude this series of results, we consider the recursive majority function RecMaj . In the classical case, its communication complexity is known to be in between $2 \log_3 n$ and $3 \log_3 n$ [LLS05]. The structure of this function is ideal for implementing a ternary search. So, given the fact that the half-duplex model with silence allows the players to send messages encoded in ternary, the following theorem is straightforward.

Theorem 13. $D_s^{hd}(\text{KW}_{\text{RecMaj}}) \leq 2 \log_3 n$.

Proof. The players split their input strings into three equal parts and implement a ternary search spending two rounds per iteration. Every iteration, Alice sends the index of a substring with $\text{RecMaj} = 1$ (there is at most one such substring), or 0 if all substrings have $\text{RecMaj} = 0$. Bob chooses a substring with $\text{RecMaj} = 1$ among the other two substrings and sends its number back to Alice. \square

We show that the same upper bound holds in the half-duplex model with zero.

Theorem 14. $D_0^{hd}(\text{KW}_{\text{RecMaj}}) \leq 2 \log_3 n$.

Proof. The players split their input strings into three equal parts and implement a ternary search. Each iteration consists of two rounds. In the first round, Alice is silent if her first substring has $\text{RecMaj} = 0$, otherwise she sends 1, while Bob is silent if his first substring has $\text{RecMaj} = 1$, otherwise he sends 1. In the second round they repeat it for the second substrings. Note that both rounds can not be spent, since has at most one substring with $\text{RecMaj} = 1$. If one of the rounds was silent then the players narrow the search area to the corresponding substrings. Otherwise, they go with the third substrings. \square

5 Non-deterministic half-duplex complexity

The standard definition of the non-deterministic communication complexity does not involve any communication at all, so it is applicable to the half-duplex model without any changes. Let X and Y be non-empty finite sets.

Definition 15. We say that a function $f : X \times Y \rightarrow \{0, 1\}$ has *non-deterministic communication protocol* of complexity d if there are two functions $A : X \times \{0, 1\}^d \rightarrow \{0, 1\}$ and $B : Y \times \{0, 1\}^d \rightarrow \{0, 1\}$, such that

- $\forall (x, y) \in f^{-1}(1) \exists w \in \{0, 1\}^d : A(x, w) = B(y, w) = 1,$
- $\forall (x, y) \in f^{-1}(0) \forall w \in \{0, 1\}^d : A(x, w) \neq 1 \vee B(y, w) \neq 1.$

The *non-deterministic communication complexity* of f , denoted $N(f)$, is the minimal complexity of a non-deterministic communication protocol for f .

There is also an alternative definition of the non-deterministic communication complexity that uses communication between players [MS20].

Definition 16. We say that a function $f : X \times Y \rightarrow \{0, 1\}$ has *privately non-deterministic communication protocol* of complexity d if there is a function $\hat{f} : (X \times \{0, 1\}^d) \times (Y \times \{0, 1\}^d) \rightarrow \{0, 1\}$ of (deterministic) communication complexity at most d such that

- $\forall (x, y) \in f^{-1}(1) \exists w_x, w_y \in \{0, 1\}^d : \hat{f}((x, w_x), (y, w_y)) = 1,$
- $\forall (x, y) \in f^{-1}(0) \forall w_x, w_y \in \{0, 1\}^d : \hat{f}((x, w_x), (y, w_y)) = 0.$

The *privately non-deterministic communication complexity* of f is the minimal depth of a privately non-deterministic communication protocol for f .

This alternative definition of non-deterministic communication uses private witnesses instead of a public one, and hence the players do need to communicate. In the classical case, Definition 15 and Definition 16 are equivalent (see [MS20] for more details). We think that this way of defining it is the right way to define the *non-deterministic half-duplex communication complexity*. So, we define it by replacing the communication model in Definition 16 with the half-duplex models. Let $N_s^{hd}(f)$, $N_0^{hd}(f)$, and $N_a^{hd}(f)$ denotes the non-deterministic half-duplex communication complexity of f with silence, with zero, and with adversary, respectively. We are going to prove bounds that connect the classical non-deterministic communication complexity with the non-deterministic half-duplex communication complexity.

Let's start with the lower bounds.

Theorem 17. For any function $f : X \times Y \rightarrow \{0, 1\}$,

$$N_s^{hd}(f) \geq N(f)/\log 5, \quad N_0^{hd}(f) \geq N(f)/\log 3, \quad N_a^{hd}(f) \geq N(f)/\log 3.$$

Proof. In the half-duplex model with silence, the protocol is a pair of trees of arity 5. The following (classical) non-deterministic protocol can simulate any non-deterministic half-duplex protocol Π with silence. Alice and Bob publicly guess a root-to-leaf path π_A in the tree of Π corresponding to Alice. Alice checks that this transcript is a valid transcript for her input. Bob checks that there exists a root-to-leaf path π_B in his tree that is a valid transcript for his input, and at the same time π_B matches π_A (in all rounds where Alice receives in π_A , Bob in π_B does the corresponding action, and in all rounds where Alice sends, Bob receives the corresponding bit or sends some bit). If Π has complexity d , then the length of the description of π_A is $\lceil d \cdot \log 5 \rceil$. This gives us the first lower bound.

Similarly, in the half-duplex model with zero, the protocol is a pair of trees of arity 3. The same reasoning shows that for any non-deterministic half-duplex protocol Π with zero of complexity d , there exists a (classical) non-deterministic protocol of complexity $\lceil d \cdot \log 3 \rceil$.

In the half-duplex model with adversary, we can consider only such transcripts where the players always receive zeroes in silent rounds. Thus, the lower bound for the half-duplex model with zero applies. \square

Now we proceed to the upper bounds. The upper bounds are based on the upper bounds for the equality function. As far as we don't know any non-trivial upper bounds for the equality in the half-duplex model with adversary, we prove upper bounds for two other models only.

Theorem 18. *For any function $f : X \times Y \rightarrow \{0, 1\}$,*

$$N_s^{hd}(f) \leq N(f)/\log 5 + O(\log N(f)), \quad N_0^{hd}(f) \leq N(f)/\log 3 + O(\log N(f)).$$

Proof. The proof is very straightforward. For any (classical) non-deterministic protocol Π , Alice and Bob can privately guess a public witness w and then check that they both guessed the same witness. This requires solving the equality on strings of length $N(f)$. Together with the upper bounds on the equality [HIMS18a, Theorems 15 and 19], this gives us the desired bounds. \square

6 Open Questions

In addition to the open questions in [HIMS18b], we would be interested to know the answers to the following questions.

1. Is it possible to use information-theoretic methods to improve the lower bounds for disjointness?
2. Prove an upper bound for the Karchmer-Wigderson game for MOD_p in the half-duplex model with zero.
3. The protocol for RecMaj in the half-duplex model with silence seems to be trivial. Can we achieve a better upper bound with a smarter protocol?

References

- [BH96] Gerth Stølting Brodal and Thore Husfeldt. *A communication complexity proof that symmetric functions have logarithmic depth*. BRICS, 1996. URL: <https://www.brics.dk/RS/96/1/BRICS-RS-96-1.pdf>.
- [Chi90] Andrew Chin. On the depth complexity of the counting functions. *Inf. Process. Lett.*, 35(6):325–328, 1990. doi:10.1016/0020-0190(90)90036-W.
- [CT06] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, New York, USA, 2006.

- [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. doi:10.1007/s00037-001-8195-x.
- [HIMS18a] Kenneth Hoover, Russell Impagliazzo, Ivan Mihajlin, and Alexander Smal. Half-duplex communication complexity. *Electronic Colloquium on Computational Complexity (ECCC)*, 25:89, 2018. URL: <https://eccc.weizmann.ac.il/report/2018/089>.
- [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. doi:10.4230/LIPICs.ISAAC.2018.10.
- [Khr71] VM Khrapchenko. Complexity of the realization of a linear function in the class of ii-circuits. *Mathematical Notes of the Academy of Sciences of the USSR*, 9(1):21–23, 1971.
- [KN97] Eyal Kushilevitz and Noam Nisan. *Communication complexity*. Cambridge University Press, 1997.
- [KW88] Mauricio Karchmer and Avi Wigderson. Monotone circuits for connectivity require super-logarithmic depth. In Janos Simon, editor, *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2-4, 1988, Chicago, Illinois, USA*, pages 539–550. ACM, 1988. URL: <http://doi.acm.org/10.1145/62212.62265>, doi:10.1145/62212.62265.
- [LLS05] Sophie Laplante, Troy Lee, and Mario Szegedy. The quantum adversary method and classical formula size lower bounds. In *20th Annual IEEE Conference on Computational Complexity (CCC 2005), 11-15 June 2005, San Jose, CA, USA*, pages 76–90. IEEE Computer Society, 2005. doi:10.1109/CCC.2005.29.
- [MS20] Ivan Mihajlin and Alexander Smal. Toward better depth lower bounds: the xor-krw conjecture. *Electronic Colloquium on Computational Complexity (ECCC)*, 27:116, 2020. URL: <https://eccc.weizmann.ac.il/report/2020/116>.
- [Yao79] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (preliminary report). In Michael J. Fischer, Richard A. DeMillo, Nancy A. Lynch, Walter A. Burkhard, and Alfred V. Aho, editors, *Proceedings of the 11th Annual ACM Symposium on Theory of Computing, April 30 - May 2, 1979, Atlanta, Georgia, USA*, pages 209–213. ACM, 1979. doi:10.1145/800135.804414.