Electronic Colloquium on Computational Complexity, Revision 1 of Report No. 89 (2018)

Half-duplex communication complexity

Kenneth Hoover* Russell Impagliazzo† Ivan Mihajlin‡ Alexander V. Smal§

May 16, 2018

Abstract

Suppose Alice and Bob are communicating bits to each other in order to compute some function f, but instead of a classical communication channel they have a pair of walkie-talkie devices. They can use some classical communication protocol for f where each round one player sends bit and the other one receives it. The question is whether talking via walkie-talkie gives them more power? Using walkie-talkie instead of a classical communication channel allows players two extra possibilities: to speak simultaneously (but in this case they do not hear each other) and to listen at the same time (but in this case they do not transfer any bits). We show that for some definitions this non-classical communication model is, in fact, more powerful than the classical one as it allows to compute some functions in a smaller number of rounds. We introduce round elimination technique for proving lower bounds in this setting and use it to prove lower bounds for some Boolean functions. We also apply information theoretic methods to prove better lower bounds for one of the models.

1 Introduction

In the classical communication complexity introduced by Yao [9] there are two players, Alice and Bob, that are trying to compute $f(x, y)$, for some function f, where x is given to Alice and y is given to Bob. Alice and Bob can communicate by sending bits to each other, one bit per round. The essential property of this *classical* model is that in every round of communication one player sends some bit and the other one receives it.

We define three new communication models that generalize the classical one and resemble communication over so-called *half-duplex channels*. A well-known example of half-duplex communication is talking via walkie-talkie: you have to hold a "push-to-talk" button in order to speak to other person, and you have to release it when you want to listen. If by accident two persons try to speak simultaneously then they do not hear each other. We consider communication models where players are allowed to speak simultaneously. Every round each player chooses one of three actions: send 0, send 1, or receive. There are three different types of rounds. If one player sends some bit and the other one receives in a round then communication works as in the classical case, we call such rounds *normal*. If both players send bits in a round then these bits get lost (the same happens if two persons try to speak via walkie-talkie simultaneously), we call these rounds *spent*. If both

*University of California San Diego, khooveri@eng.ucsd.edu
†University of California San Diego, russell@cs.ucsd.edu
‡University of California San Diego, ivmihajlin@gmail.com
§St. Petersburg Department of Steklov Mathematical Institute of Russian Academy of Sciences, smal@pdmi.ras.ru

ISSN 1433-8092

players receive in a round, we call these rounds *silent*. We distinguish three possible models, based on what the players receive in silent rounds:

1. both players receive *nothing*, i.e., it is possible for both players to distinguish a silent round from a normal one, we call this model *half-duplex communication with silence*;
2. both players receive 0, i.e., players cannot distinguish a silent round from a normal round where the other player sends 0, we call this model *half-duplex communication with zero*;
3. each player receives some arbitrary bit, not necessary the same as the other player, we call this model *half-duplex communication with adversary*.

In this paper we study communication complexity of Boolean functions that are hard in the classical communication model.

1.1 Motivation

The original motivation to study these kinds of communication models arose from the question of the complexity of Karchmer-Wigderson games [6] for multiplexers. The *Karchmer-Wigderson game for a function* $f : \{0,1\}^n \to \{0,1\}$ is a (classical) communication problem where Alice is given $x \in f^{-1}(0)$, Bob is given $y \in f^{-1}(1)$, and they want to find $i \in [n]$ such that $x_i \neq y_i$. A *multiplexer* (or *indexing function*) is a function $M_n : \{0,1\}^{2^n} \times \{0,1\}^n \to \{0,1\}$, such that $M_n(t,i) = t[i]$, i.e., M_n interprets the first part of its input as the truth table of some function $f : \{0,1\}^n \to \{0,1\}$ and the second part as an input x to the function, and outputs $f(x)$. There is a lot of work has been done studying Karchmer-Wigderson games for compositions including *universal relations* [4, 2, 3] as a part of Karchmer-Raz-Wigderson program [5] for proving super-logarithmic formula depth lower bounds. Multiplexers are similar to universal relations in the sense that there is a natural reduction from a Karchmer-Wigderson game for some function $f : \{0,1\}^n \to \{0,1\}$ to a Karchmer-Wigderson game for multiplexer M_n: if Alice and Bob are given x and y in the game for f we give them $(tt(f), x)$ and $(tt(f), y)$, respectively, in the game for M_n, where $tt(f)$ is a truth table of function f. On the other hand multiplexers are functions, universal relations are not, so proving analogous results for multiplexers would be one step toward the goal of Karchmer-Raz-Wigderson program. Unfortunately all the techniques that were used for universal relations cannot be applied directly to multiplexers because it is impossible to give Alice and Bob the same input string (all these techniques exploited a symmetry of universal relations that allows to give players the same input string, but this is impossible for functions because inputs of Alice and Bob come from disjoint sets).

Suppose now that Alice and Bob are playing Karchmer-Wigderson game for multiplexer M_n: Alice is given $(tt(f), x)$, $x \in f^{-1}(0)$, and Bob is given $(tt(g), y)$, $y \in g^{-1}(1)$. If the players are also given a promise that $f = g$ then they can use a protocol for Karchmer-Wigderson game for f. But what if this promise is broken? Alice can try to act according to the protocol for f, Bob at the same time can try to act according to a protocol for g, but in some round of this "mixed" protocol they might both want to send or both want to receive at the same time. Such protocol "mixing" can not be done in the classical model. To make it possible we extend the communication model by allowing players to speak or listen simultaneously. How does it affect the communication complexity? As a first step toward answering this question we study half-duplex communication complexity of Boolean functions $\{0,1\}^n \times \{0,1\}^n \to \{0,1\}$.

1.2 Organization of this paper

In Section 2, we give definitions for considered communication models. Then, in Section 3, we prove trivial upper and lower bounds that follows immediately from the definitions. Next, in Section 4, we discuss combinatorial rectangles of input pairs that can be associated with communication protocol, and their application for proving communication complexity lower bounds. In Sections 5, 6 and 7, we present our main results, upper and lower bounds for proposed communication models. Finally, in Section 8, we state several open questions.

2 Definitions

Definition 1. Let X, Y and Z be some finite sets. We say that two players, Alice and Bob, are solving *half-duplex communication problem* for relation $R \subseteq X \times Y \times Z$ if sets X, Y, Z, and relation R are known by both players, Alice is given some $x \in X$, Bob is given some $y \in Y$, and players want to find some $z \in Z$, $(x, y, z) \in R$, by communicating to each other. The communication is organized in rounds. At every round, each player decides (depending only on its own input and previous communication) to do one of three available actions: send 0, send 1 or receive. If one player sends some bit $b \in \{0, 1\}$ and the other one receives then the latter gets bit b, we call such rounds *normal*. If both players send bits at the same time then these bits get lost, we call such rounds *spent* (it is important that the player that is sending can not distinguish whether this round is normal or spent). If both players receive at the same time, we call such rounds *silent*. There are three variants of half-duplex communication problem depending on how silent rounds work.

- In a silent round both players receive *nothing*, so it is possible for both players to distinguish a silent round from a normal one, the corresponding problem is called *half-duplex communication problem with silence*.
- In a silent round both players receive 0, i.e., players cannot distinguish a silent round from a normal round where the other player sends 0, the corresponding problem is called *half-duplex communication problem with zero*;
- In a silent round each player receives some arbitrary bit, not necessarily the same as the other player; the corresponding problem is called *half-duplex communication problem with adversary*.

We say that half-duplex communication problem for R is *solved* if at the end of communication both players know some z, such that $(x, y, z) \in R$.

Next we define a notion of *communication protocol*. In the classical case a protocol is a binary rooted tree that describes communication of players on all possible inputs: every internal node corresponds to a state of communication and defines which of players is sending this round. Unlike the classical case in half-duplex communication player does not always know what was the other's player action — the information about it can be "lost" i.e., in spent rounds player do not know what was that other's player action. It means that a player might not know what node of the protocol corresponds to the current state of communication. Note also that solving half-duplex communication problem with zero there is no need to send zeros — player can receive instead and the other player will not notice the difference. Keeping all this in mind, we give the following definition of half-duplex protocol.

Definition 2. *Half-duplex communication protocol with silence* that solves a relation $R \subset X \times Y \times Z$ is a pair (T_A, T_B) of rooted trees that describe how Alice and Bob communicate on all possible inputs $(x, y) \in X \times Y$. Every node of T_A corresponds to a state of Alice, every node of T_B — to a state of Bob. Every leaf l is labeled with $z_l \in Z$. Let $\mathcal{A} = \{\text{send } 0, \text{send } 1, \text{receive}\}$ be the set of possible actions, and $\mathcal{E} = \{\text{send } 0, \text{send } 1, \text{receive } 0, \text{receive } 1, \text{silence}\}$ be the set of all possible events. Every node v of T_A and (of T_B) is labeled with two functions $g_v : X \to \mathcal{A}$ ($g_v : Y \to \mathcal{A}$) and $h_v : \mathcal{E} \to C(v)$, where $C(v)$ is a set of child nodes of v. Root nodes of T_A and T_B correspond, respectively, to the initial states of Alice and Bob. If Alice (Bob) is in a state that corresponds to node $v \in T_A$ ($v \in T_B$), then she does action $g_v(x)$ (he does action $g_v(y)$). Events of both players are defined in a natural way by their actions in this round. The next node of the protocol is defined by the function h. When players reach leaves communication ends (they always reach leaves simultaneously). The protocol is correct if for every input pair $(x, y) \in X \times Y$ communication ends in a pair of leaves labeled with the same $z \in Z$ such that $(x, y, z) \in R$.

Half-duplex communication protocol with zero is defined in the same way with the only difference that set of possible events does not include "send 0", i.e. $\mathcal{E} = \{\text{send } 1, \text{receive } 0, \text{receive } 1\}$.

Half-duplex communication protocol with adversary that solves a relation $R \subset X \times Y \times Z$ is a pair (T_A, T_B) of rooted trees that describe how Alice and Bob communicate on all possible inputs $(x, y) \in X \times Y$ and for any strategy of adversary $w \in \{0, 1\}^*$. The structure of the protocol is the same as in half-duplex communication protocol with zero, but with $\mathcal{E} = \{\text{send } 0, \text{send } 1, \text{receive } 0, \text{receive } 1\}$. If both players decide to receive in round i then Alice and Bob receive bits w_{2i-1} and w_{2i} respectively. The protocol is correct if for every input pair $(x, y) \in X \times Y$ and any strategy of adversary $w \in \{0, 1\}^*$ communication ends in two leaves labeled with the same $z \in Z$ such that $(x, y, z) \in R$.

For each of these models, a *partial transcript after k rounds* is a pair (π_a, π_b) of length-k sequences over $\mathcal{E}$ that lists the events observed by Alice and Bob, respectively, after running some protocol on a pair of inputs for k rounds.

The cardinality of set $\mathcal{E}$ upper bounds arity of trees T_A and T_B: arity is 5 for half-duplex communication with silence, 3 for half-duplex communication with zero, and 4 for half-duplex communication with adversary.

Definition 3. We say that half-duplex communication protocol *solves* a communication problem for function $f : X \times Y \to Z$ if it solves a relation $R(f) = \{(x, y, f(x, y)) \mid x \in X, y \in Y\}$.

The classical communication complexity of a communication problem for function f, $D(f)$, is defined in terms of the minimal depth of a protocol solving it. Analogously, we define communication complexity for half-duplex communication problems.

Definition 4. The minimal depth of a communication protocol solving half-duplex communication problem for function f with silence, with zero, with adversary, define *half-duplex communication complexity of function f* with silence, denoted $D_s^{hd}(f)$, with zero, denoted $D_0^{hd}(f)$, with adversary, denoted $D_a^{hd}(f)$, respectively. Analogously, we define *half-duplex communication complexity of relation R* with silence, $D_s^{hd}(R)$, with zero, $D_0^{hd}(R)$, and with adversary, $D_a^{hd}(R)$.

In this paper we study half-duplex communication complexity for a special case of Boolean functions $\{0, 1\}^n \times \{0, 1\}^n \to \{0, 1\}$ (i.e., $X = Y = \{0, 1\}^n$, $Z = \{0, 1\}$).

3 Trivial bounds

As far as half-duplex communication generalizes classical communication the following upper bound is immediate.

Theorem 1. *For every function* $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^k$, $D_s^{hd}(f) \leq D_0^{hd}(f) \leq D_a^{hd}(f) \leq D(f)$.

Proof. Every classical communication protocol can be embedded in half-duplex communication protocol that does not use spent and silent rounds. □

Next theorem shows that every half-duplex protocol with zero or with adversary can be transformed in a classical communication protocol of double depth.

Theorem 2. *For every function* $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^k$, $\frac{D(f)}{2} \leq D_0^{hd}(f) \leq D_a^{hd}(f)$.

Proof. Every t-round half-duplex communication protocol with silence or with adversary can be transformed into $2t$-round classical communication protocol. Every round of the original protocol corresponds to two consecutive rounds of the new one: at first round Alice sends a bit she was sending in the original protocol or sends 0 if she was receiving, at second round Bob does the same thing. □

As we will see later, half-duplex protocols with silence can use silent rounds as an additional third symbol and hence not every t-round half-duplex protocol with silence can be embedded in $2t$ classical protocol. The following theorem shows that instead we can embed every such protocol in a classical protocol with $3t$ rounds.

Theorem 3. *For every function* $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^k$, $D_s^{hd}(f) \geq \frac{D(f)}{3}$.

Proof. Every t-round half-duplex communication protocol with silence can be transformed into $3t$-round classical communication protocol. Every round of the original protocol corresponds to three consecutive rounds of the new one: at first round Alice sends 1 to indicate if she was sending a bit in the original protocol, or sends 0 otherwise, at second round Bob does the same thing symmetrically. After that they are both aware of the intentions of each other. If they were both planning to send, they can skip the third round. If they were both planning to receive, then they can just assume that they heard silence. If one player was planning to send and the other one was planning to receive they can perform such an action on third round. □

4 Rectangles

Many lower bounds on classical communication complexity were proved by considering combinatorial rectangles that are associated with the nodes of communication protocol [8]: it's easy to see that every node v of the (classical) protocol corresponds to a combinatorial rectangle $R_v = X_v \times Y_v$, where $X_v \subseteq X$, $Y_v \subseteq Y$, such that if Alice and Bob are given an input from R_v then their communication will necessarily pass through node v. This implies that the rectangles associated with the child nodes of v define a subdivision of R_v.

There is a general technique [8] for proving lower bounds using associated combinational rectangles in: if for some sub-additive measure μ defined on combinatorial rectangles we show both a lower bound on the measure of $X \times Y$, the rectangle in the root node, i.e., $\mu(X \times Y) \geq \mu_r$ for some μ_r, and an upper bound on the measure of rectangles in leaves, i.e., for every leaf l the measure of the corresponding rectangle R_l is at most μ_ℓ for some μ_ℓ, then we can claim lower bound of $\log_2(\mu_r/\mu_\ell)$ on the depth of the protocol.

One of the most studied sub-additive measure on rectangles is $\mu_M(R)$ that is equal to the minimal number of *monochromatic* rectangles that covers R. Rectangle R is z*-monochromatic* in respect to function f for some $z \in Z$ if for all $(x, y) \in R$, $f(x, y) = z$. As far as both players have to come up with the same answer at the end of communication every rectangle in leaves is monochromatic, thus for this measure $\mu_\ell = 1$.

Almost the same technique can be used for half-duplex protocols. There are some technical differences that we have to keep in mind. First of all, we can apply this idea for both trees T_A and T_B. We should also note that trees T_A and T_B are non-binary. Secondly, we should be careful while defining associated combinatorial rectangles for half-duplex protocols with adversary — in case of silent rounds the next node of the protocol depends also on a strategy w of adversary, so we have to formally consider w it as a part of input. This leads to the following lower bound for equality function $\mathrm{EQ}_n : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$, such that $\mathrm{EQ}_n(x, y) = [x = y]$.

Theorem 4.

- $D^{hd}_s(\mathrm{EQ}_n) \geq \log_5 2^n = n/\log 5$,
- $D^{hd}_0(\mathrm{EQ}_n) \geq \log_3 2^n = n/\log 3$,
- $D^{hd}_a(\mathrm{EQ}_n) \geq \log_4 2^n = n/2$.

Proof. Let $\mu = \mu_M$. All rectangle in leaves are monochromatic, $\mu_\ell = 1$. Every 1-monochromatic rectangle is of size one: if some rectangle contains two elements, say (x, x) and (x', x'), then it also contains (x, x') and (x', x), so it is not 1-monochromatic. Thus, the root rectangle has measure at least $\mu_r = 2^n + 1$ [8]. □

Surprisingly, as we will see later, first two result are sharp up to additive logarithmic term. We can get better bound if we improve this technique using *round elimination*.

4.1 Round elimination

Let us fix a protocol for some half-duplex communication problem and consider the first round. Let $R_c = X \times Y$ be the corresponding rectangle of all possible inputs. We can subdivide R_c in nine rectangles, one for each possible combination of actions.

Alice\Bob	send 0	send 1	receive
send 0	R_{00}	R_{01}	R_{0r}
send 1	R_{10}	R_{11}	R_{1r}
receive	R_{r0}	R_{r1}	R_{rr}

Consider two rectangles: $R_{good} = R_{00} \cup R_{01} \cup R_{0r}$ and $R_{bad} = R_{0r} \cup R_{1r}$. If we restrict f to be a partial function defined only on R_{good}, i.e., players will always get some $(x, y) \in R_{good}$, then there is no need in the first round — the information the players get about the other part of the input is fixed: Alice does not get any information, Bob can receive 0 if he decide to receive. On the other hand if we restrict f to R_{bad} then the first round is still needed: Bob can receive both 0 and 1 and this information in necessary to proceed to the next round. Lets call a rectangle R *good for* *functions* f if restricting f to R makes the first round unnecessary (i.e., protocol without the first round is correct for all $(x, y) \in R$). The idea of this method is to consider some covering of R_c with a set of *good* rectangles and prove that there is always a good rectangle of large enough measure. If we can show that there is always a rectangle of measure at least $\alpha \cdot \mu(R_c)$ then we can iterate this idea and claim that protocol depth is at least $\log_{1/\alpha}(\mu_r/\mu_\ell)$, where μ_r is a lower bound on the measure of the root rectangle and μ_ℓ is an upper bound on the measure of leaf rectangles.

Lemma 1. *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 $f \restriction R$ of measure at least $\alpha \cdot \mu(R)$ then the depth of the protocol is at least $\log_{1/\alpha} \frac{\mu_r}{\mu_\ell}$.*

Proof. We start with $R = X \times Y$. Every round restrict f to some good $R_{good} \subseteq R$ such that $\mu(R_{good}) \geq \alpha \cdot \mu(R)$, let R to be R_{good}, and proceed to the next round. At the end we will reach some leaf. Thus there is at least $\log_{1/\alpha}(\mu_r/\mu_\ell)$ rounds. $\square$

5 Half-duplex communication with silence

The main advantage of this model over the other models we consider is that whenever players have silent round, they learn about it. In some sense they have a third symbol in the alphabet — receiving player can get either 0/1 or a special symbol corresponding to "silence". Next theorem shows how players can take the advantage of silence to transfer data.

Theorem 5. *For every $f : \{0, 1\}^n \times \{0, 1\}^n \to \{0, 1\}$, $D_s^{hd}(f) \leq \lceil n/\log 3\rceil + 1$.*

Proof. Alice encodes x in ternary alphabet $\{0, 1, 2\}$ and sends it to Bob: in order to send 0 or 1 Alice sends the corresponding bit, sending 2 is emulated by receiving (keeping silence). This requires $\lceil \log_3 2^n\rceil = \lceil n/\log 3\rceil$ bits. At the last round Bob computes $f(x, y)$ and sends it back to Alice. $\square$

Using the idea of encoding in a non-binary alphabet, we managed to prove a better upper bound for equality function.

Theorem 6. $D_s^{hd}(\mathrm{EQ}_n) \leq \lceil n/\log 5\rceil + \lceil \log n/\log 3\rceil + 2$.

Proof. Alice and Bob encode their inputs in alphabet of size five $\{0, 1, 2, 3, 4\}$. Then they process their inputs symbol by symbol sequentially in $\lceil n/\log 5\rceil$ rounds. At round i they process ith symbol in the following manner.

Symbol	**Alice**	**Bob**
0	send 0	receive
1	send 1	receive
2	receive	send 0
3	receive	send 1
4	receive	receive

If ith round is normal then one player can check whether ith symbols are different. If ith round is silent then again one player knows if ith symbols are different. If after $\lceil n/\log 5\rceil$ rounds one of the players has already learned that the answer is 0, then he or she sends 0. If this round is not silent, then both players know that the answer is 0. Otherwise, Alice and Bob have to make sure that there were no spent rounds. In order to check it, Alice sends the number normal rounds she was receiving in encoded in ternary, that requires $\lceil \log n / \log 3 \rceil$ rounds. Bob checks whether this number is equal to the number of rounds he was sending in. If so, inputs are equal. In the last round, Bob sends the answer back to Alice. $\square$

The next theorem shows better than $n/\log 3$ upper bound for disjointness function $\mathrm{DISJ}_n : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$, such that $\mathrm{DISJ}_n(x,y) = \bigwedge_{i\in[n]} \neg(x_i \wedge y_i)$, which in classical case is one of the hardest functions of this type.

Theorem 7. $D_s^{hd}(\mathrm{DISJ}_n) \leq \lceil n/2 \rceil + 2$.

Proof. Alice and Bob process their inputs two bits per round, $\lceil n/\log 2 \rceil$ rounds. At round i they process symbols $2i-1$ and $2i$ in the following manner.

Symbols	**Alice**	**Bob**
00	send 0	receive
01	receive	send 0
10	receive	send 1
11	receive	receive

At the end of communication Bob tells Alice whether there was a silent round in which Bob's input was 11 (i.e., inputs are not disjoint). Alice tells Bob whether she ever received 0 having 01 or 11, or received 1 having 10 or 11 (again, inputs are not disjoint). $\square$

The next function we have results for is the inner product function $\mathrm{IP}_n : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$, such that $\mathrm{IP}_n(x,y) = \bigoplus_{i\in[n]} x_i y_i$. In the classical model, this function is one of the harder ones. This might also be the case for half-duplex models as the same time we do not know efficient protocols for it, and this is the function we can prove the best lower bounds for. On the other hand, the best lower bound we can prove for it in this model is $n/2$.

Theorem 8. $D_s^{hd}(\mathrm{IP}_n) \geq n/2$.

For this theorem we need the following fact about inner product function.

Lemma 2. *Every leaf rectangle of a protocol solving communication problem for* IP_n *has size at most* 2^n.

Proof. We start with proving it for leaves labeled with 0. Let $R_l = X_l \times Y_l$ be a rectangle of leaf l labeled with 0, i.e., R_l is 0-monochromatic. For every $x \in X_l$ and $y \in Y_l$, $\mathrm{IP}_n(x,y) = 0$, set X_l must be contained in the orthogonal complement for span of Y_l. Thus, $\dim(\{X_l\}) + \dim(\{Y_l\}) \leq n$, and hence, $|R| = |X_l| \times |Y_l| \leq 2^n$.

If leaf is labeled with 1 then for every $x \in X_l$ and $y \in Y_l$, $\mathrm{IP}_n(x,y) = 1$. Let y' be arbitrary element of Y_l. Consider a set $Y_l' = \{y \oplus y' \mid y \in Y_l\}$. It is easy to see that for every $x \in X_l$ and $y \in Y_l'$, $\mathrm{IP}_n(x,y) = 0$, so we can apply the argument above to show that $|X_l| \times |Y_l'| \leq 2^n$. It remains to notice that $|Y_l| = |Y_l'|$. $\square$

Proof of Theorem 8. Let R_c be the rectangle of all possible inputs and $\mu(R) = |R|$. Consider the following set of good rectangles: a rectangle $R_{silent} = R_{rr}$ where round is silent, four rectangles $R_{0*} = R_{00} \cup R_{01} \cup R_{0r}$, $R_{1*} = R_{10} \cup R_{11} \cup R_{1r}$, $R_{*0} = R_{00} \cup R_{10} \cup R_{r0}$, $R_{*1} = R_{01} \cup R_{11} \cup R_{r1}$, where one of players sends some bit, and a rectangle $R_{spent} = R_{00} \cup R_{01} \cup R_{10} \cup R_{11}$, where round is spent. We claim one of these good rectangles has measure at least $\mu(R_c)/4$.

For $\mu(R) = |R|$ we can use the following fact. Let a_0, a_1 and a_r be the probability over all possible inputs that Alice sends 0, sends 1, and receives, respectively. Analogously, we define b_0, b_1 and b_r to be the probability that Bob sends 0, sends 1, and receives. It is easy to see that $a_0 + a_1 + a_r = b_0 + b_1 + b_r = 1$ and for all $\alpha, \beta \in \{0, 1, r\}$, $\mu(R_{\alpha\beta}) = a_\alpha \cdot b_\beta \cdot \mu(R_c)$.

We need to show that $\max\big\{\mu(R_{0*}), \mu(R_{1*}), \mu(R_{*0}), \mu(R_{*1}), \mu(R_{silent}), \mu(R_{spent}))\big\} \geq \mu(R_c)/4$. This is equivalent to showing that $\max\big\{a_1, a_0, b_1, b_0, a_r b_r, (1-a_r)(1-b_r)\big\} \geq 1/4$ for any reals $a_0, a_1, a_r, b_0, b_1, b_r \in [0,1]$, such that $a_0+a_1+a_r = b_0+b_1+b_r = 1$. Let $\bar{a} = (a_1+a_0)/2$, $\bar{b} = (b_1+b_0)/2$. As far as $\max\{a_0, a_1\} \geq \bar{a}$ and $\max\{b_0, b_1\} \geq \bar{b}$,

$$\max\big\{a_1, a_0, b_1, b_0, a_r b_r, (1-a_r)(1-b_r)\big\} \geq \max\big\{\bar{a}, \bar{b}, a_r b_r, (1-a_r)(1-b_r)\big\}.$$

Note that $a_r + 2\bar{a} = 1$, $b_r + 2\bar{b} = 1$. Hence $\bar{a} = (1-a_r)/2$, $\bar{b} = (1-b_r)/2$,

$$\max\big\{\bar{a}, \bar{b}, a_r b_r, (1-a_r)(1-b_r)\big\} = \max\big\{(1-a_r)/2, (1-b_r)/2, a_r b_r, (1-a_r)(1-b_r)\big\}.$$

If $a_r \leq 1/2$ or $b_r \leq 1/2$ then one of first arguments is at least $1/4$. On the other hand if $a_r > 1/2$ and $b_r > 1/2$ then $a_r b_r > 1/4$. Now we apply Lemma 1 for $\mu_r = 4^n$, $\mu_\ell = 2^n$ (Lemma 2), $\alpha = 1/4$, and get the desired bound. □

6 Half-duplex communication with zero

As we have already mentioned before there are only two reasonable actions in this model: send 1 or receive. The following theorem shows that half-duplex communication with zero is more powerful than classical communication, namely, it is possible to solve communication problem for EQ_n in less than n rounds of communication.

Theorem 9. $D_0^{hd}(\mathrm{EQ}_n) \leq \lceil n/\log 3\rceil + 2\lceil \log n\rceil + 1$.

Proof. Alice and Bob encode their inputs in ternary. In the first phase of the protocol, they process their inputs sequentially symbol by symbol in $\lceil n/\log 3\rceil$ rounds. At round i they process ith symbol in the following manner.

Symbol	**Alice**	**Bob**
0	receive	receive
1	send 1	receive
2	receive	send 1

In the next $2\lceil \log n\rceil$ they send each other the number of ones they sent in the first phase. If inputs were different then one of players must have noticed it. At the first phase at round i Alice learns if their corresponding symbols are $(0,2)$, $(2,0)$ or $(2,1)$, Bob learns if their symbols are $(0,1)$ or $(1,0)$. In the second phase, they can learn whether any of $(1,2)$ situation happened in the first phase. The last round players use to notify each other if somebody noticed a mismatch — in this case the player that noticed sends 1. □

Next theorem shows that there are functions of higher complexity than EQ_n.

Theorem 10. $D_0^{hd}(\mathrm{IP}_n) \geq n/\log\frac{2}{3-\sqrt{5}} > n/\log 2.62$.

Proof. Let R_c be the rectangle of all possible inputs and $\mu(R) = |R|$. Consider the following set of good rectangles: $R_{slilent} = R_{rr}$, $R_{spent} = R_{11}$, $R_{1*} = R_{11} \cup R_{1r}$ and $R_{*1} = R_{11} \cup R_{r1}$. We claim one of these good rectangles has measure at least $\frac{3-\sqrt{5}}{2} \cdot \mu(R_c)$. We need to show that

$$\max\{\mu(R_{1*}), \mu(R_{*1}), \mu(R_{silent}), \mu(R_{spent}))\} \geq \frac{3-\sqrt{5}}{2} \cdot \mu(R).$$

It is equivalent to showing that for any $a, b \in [0, 1]$, $\max\{a, b, ab, (1-a)(1-b)\} \geq \frac{3-\sqrt{5}}{2}$, where a and b denote the probabilities over all possible inputs that, respectively, Alice and Bob sends 1. It's easy to see minimum value of $\max\{a, b, ab, (1-a)(1-b)\}$ is at most $1/2$, so we can consider only $a \leq 1/2$ and $b \leq 1/2$. Thus,

$$\max\{a, b, ab, (1-a)(1-b)\} = \max\{a, b, (1-a)(1-b)\}.$$

Now we can argue that minimum of this max is achieved when $a = b = (1-a)(1-b)$: indeed, increasing or decreasing a or b increases one of the arguments. Solving corresponding quadratic equation $a = (1-a)^2$ we get $a = \frac{3-\sqrt{5}}{2}$, and hence $\max\{a, b, ab, (1-a)(1-b)\} \geq \frac{3-\sqrt{5}}{2}$. Applying Lemma 1 for $\mu_r = 4^n$, $\mu_\ell = 2^n$, and $\alpha = \frac{3-\sqrt{5}}{2}$ finishes the proof. □

7 Half-duplex communication with adversary

The main feature of this model is that receiving player can not be 100% sure that the received bit if in fact is "real", i.e., this bit originates from the other player, not from an adversary. But the protocol must be correct for any strategy of adversary. Our intuition prompts that in this setting silent and spent rounds would be useless. There is a common obstacle our combinatorial methods faced when we were trying to prove this conjecture — it could be the case that players send different bits in spent rounds. For some reason, our combinatorial methods do not work in this case which is strange because these spend rounds do not transmit any information. If we somehow forbid players to send different bits in spent rounds (e.g., in this case, we immediately terminate the communication and make players output 0) then we can prove that EQ_n requires n rounds of communication. The same bound can be achieved if we allow such spent rounds only on distinct inputs. Using combinatorial methods we managed to show the following two theorems (the proofs are given in Appendix).

Theorem 11. $D^{hd}_a(\mathrm{EQ}_n) \geq n/\log 2.5$.

Theorem 12. $D^{hd}_a(\mathrm{IP}_n) \geq n/\log \frac{7}{3}$.

The better lower bound can be obtained using information-theoretic approach.

7.1 Upper-bound on internal information

A useful tool for proving lower bounds on the communication complexity of problems in the classical model is the upper bound on the information Alice and Bob have learned about the other's inputs, as a function of the number of rounds that have been run. This notion is known as *internal information complexity*. Such tools allow for proving lower bounds, such as the $2\log n$-bit lower bound on the Karchmer-Wigderson game for parity. For more information on information theory we refer to [1].

Theorem 13. *Let f be a partial function and $\mathcal{P}$ a half-duplex communication protocol with adversary computing f, and $\mathcal{D}$ an arbitrary distribution over the range of f. Let $\mathcal{X}$ and $\mathcal{Y}$ be the marginal distributions over inputs to Alice and Bob, and for any k let Π_A^k and Π_B^k be the marginal distributions over Alice and Bob's partial transcripts after running $\mathcal{P}$ for k rounds induced by $\mathcal{D}$, where on silent rounds the adversary picks whether to send 0 or 1 uniformly and independently at random for each player separately. Then for every k,*

$$I(\mathcal{X} : \Pi_B^k \mid \mathcal{Y}) + I(\mathcal{Y} : \Pi_A^k \mid \mathcal{X}) \leq k.$$

Proof. We will induct on k, the number of rounds that have been run. For $k = 0$, there is only one possible partial transcript for either player, the empty transcript, and thus the result is immediate. Now suppose that this is true in round k. Let $\mathcal{E}_A^{k+1}$ and $\mathcal{E}_B^{k+1}$ be the marginal distributions over which event each player will observe. Note that

$$\begin{aligned}
I(\mathcal{X} : \Pi_B^{k+1} \mid \mathcal{Y}) &= H(\mathcal{X} \mid \mathcal{Y}) - H(\mathcal{X} \mid \mathcal{Y}, \Pi_B^{k+1}) \\
&= H(\mathcal{X} \mid \mathcal{Y}) - H(\mathcal{X} \mid \mathcal{Y}, \Pi_B^k) + H(\mathcal{X} \mid \mathcal{Y}, \Pi_B^k) - H(\mathcal{X} \mid \mathcal{Y}, \Pi_B^k, \mathcal{E}_B^{k+1}) \\
&= I(\mathcal{X} : \Pi_B^k \mid \mathcal{Y}) + I(\mathcal{X} : \mathcal{E}_B^{k+1} \mid \mathcal{Y}, \Pi_B^k).
\end{aligned}$$

Thus, it suffices to show that $I(\mathcal{X} : \mathcal{E}_B^{k+1} \mid \mathcal{Y}, \Pi_B^k) + I(\mathcal{Y} : \mathcal{E}_A^{k+1} \mid \mathcal{X}, \Pi_A^k) \leq 1$.

Let (y, π_B^k) be a particular valid input-transcript pair for Bob. Consider $I(\mathcal{X} : \mathcal{E}_B^{k+1} \mid \mathcal{Y} = y, \Pi_B^k = \pi_B^k)$; note that

$$\begin{aligned}
I(\mathcal{X} : \mathcal{E}_B^{k+1} \mid \mathcal{Y} = y, \Pi_B^k = \pi_B^k) &\leq I(\mathcal{X}, \Pi_A^k : \mathcal{E}_B^{k+1} \mid \mathcal{Y} = y, \Pi_B^k = \pi_B^k) \\
&\leq H(\mathcal{E}_B^{k+1} \mid \mathcal{Y} = y, \Pi_B^k = \pi_B^k) - H(\mathcal{E}_B^{k+1} \mid \mathcal{Y} = y, \Pi_B^k = \pi_B^k, \mathcal{X}, \Pi_A^k).
\end{aligned}$$

Suppose Bob will be receiving in round $k+1$; otherwise

$$H(\mathcal{E}_B^{k+1} \mid \mathcal{Y} = y, \Pi_B^k = \pi_B^k) = H(\mathcal{E}_B^{k+1} \mid \mathcal{Y} = y, \Pi_B^k = \pi_B^k, \mathcal{X}, \Pi_A^k) = 0.$$

Consider each (x, π_A^k) input-transcript pair for Alice consistent with (y, π_B^k). Note that $H(\mathcal{E}_B^{k+1} \mid \mathcal{Y} = y, \Pi_B^k = \pi_B^k, \mathcal{X} = x, \Pi_A^k = \pi_A^k)$ will either be 0, if Alice is sending a bit in round $k+1$, or 1, if she is receiving. The latter is because the adversary will choose whether Bob receives a 0 or 1 in round $k+1$ uniformly at random independent of Alice or Bob's transcripts or inputs. Thus

$$H(\mathcal{E}_B^{k+1} \mid \mathcal{Y} = y, \Pi_B^k = \pi_B^k, \mathcal{X}, \Pi_A^k) = \Pr[\text{Alice receives} \mid \mathcal{Y} = y, \Pi_B^k = \pi_B^k],$$

and thus

$$\begin{aligned}
I(\mathcal{X} : \mathcal{E}_B^{k+1} \mid \mathcal{Y} = y, \Pi_B^k = \pi_B^k) &\leq 1 - \Pr[\text{Alice receives} \mid \mathcal{Y} = y, \Pi_B^k = \pi_B^k] \\
&\leq \Pr[\text{Alice sends} \mid \mathcal{Y} = y, \Pi_B^k = \pi_B^k].
\end{aligned}$$

We then have that

$$\begin{aligned}
I(\mathcal{X} : \mathcal{E}_B^{k+1} \mid \mathcal{Y}, \Pi_B^k) &= \sum_{(y, \pi_B^k)} \Pr[y, \pi_B^k] \cdot I(\mathcal{X} : \mathcal{E}_B^{k+1} \mid \mathcal{Y} = y, \Pi_B^k = \pi_B^k) \\
&\leq \sum_{(y, \pi_B^k)} \Pr[\text{Alice sends}, \mathcal{Y} = y, \Pi_B^k = \pi_B^k] \cdot \mathbf{1}[\text{Bob receives}] \\
&\leq \Pr[\text{Alice sends, Bob receives}].
\end{aligned}$$

A symmetric argument holds for Alice, giving

$$I(\mathcal{X} : \mathcal{E}_B^{k+1} \mid \mathcal{Y}, \Pi_B^k) + I(\mathcal{Y} : \mathcal{E}_A^{k+1} \mid \mathcal{X}, \Pi_A^k) \\ \leq \Pr[\text{Alice sends, Bob receives}] + \Pr[\text{Alice receives, Bob sends}] \leq 1.$$

□

As immediate corollary we obtain a lower bound for IP_n.

Corollary 1. *$D_a^{hd}(\mathrm{IP}_n) \geq n$.*

Proof. By the definition of half-duplex communication protocol it have to solve the corresponding problem for any strategy of adversary, including the adversary from Theorem 13. Take the uniform distribution over all input pairs. Then $H(\mathcal{X} \mid \mathcal{Y}) + H(\mathcal{Y} \mid \mathcal{X}) = 2n$. Each leaf of any correct protocol contains at most 2^n input pairs in its rectangle, thus $H(\mathcal{X} \mid \mathcal{Y}, \Pi_B) + H(\mathcal{Y} \mid \mathcal{X}, \Pi_A) \leq n$. □

Unfortunately if we try to apply the same idea to the half-duplex model with zero it will not prove lower bounds better than $n/2$ which is already known. The problem is that in silent round both players learn some information about the other's player input — they know that the other player did not send 1.

The other immediate corollary of Theorem 13 is $2 \log n$ lower bound on the complexity of Karchmer-Wigderson relation for parity function.

Definition 5. Let $X = f^{-1}(0)$, $Y = f^{-1}(1)$ for some Boolean function $f : \{0,1\}^n \to \{0,1\}$. The *KW relation for function f*, $R_f \subseteq X \times Y \times [n]$, is defined by $R_f = \{(x, y, i) \mid x_i \neq x_i\}$.

It it well known that parity function $\oplus_n : \{0,1\}^n \to \{0,1\}$, $\oplus_n(x) = \bigoplus_{i=1}^n x_i$, requires n^2 formula size [7]. In the classical case it is equivalent to saying that KW relations for parity requires $2 \log n$ rounds of communication. Theorem 13 allows us to prove the following analogue of this result.

Corollary 2. *$D_a^{hd}(R_{\oplus_n}) \geq 2 \log n$.*

Proof. Take the uniform distribution over valid input pairs with a single bit of difference. Then $H(\mathcal{Y} \mid \mathcal{X}) + H(\mathcal{X} \mid \mathcal{Y}) = 2 \log n$ before any communication takes place. On the other hand it is easy to see that $H(\mathcal{Y} \mid \mathcal{X}, \Pi_A) + H(\mathcal{X} \mid \mathcal{Y}, \Pi_B) = 0$ at any leaf. □

8 Open problems

It would be interesting to improve upper and lower bounds for Boolean functions for all three half-duplex communication models. The next step in studying these models would be proving non-trivial lower bound for some Karchmer-Wigderson game, e.g., prove that Karchmer-Wigderson game for parity function requires $2 \log n - o(\log n)$. So we propose the following list of open problems.

1. Prove that for any $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$, $D_s^{hd}(f) \leq n/2 + o(n)$, or disprove it by showing that there is f such that $D_s^{hd}(f) \geq \alpha n - o(n)$ for some $\alpha > 1/2$.

2. Show that there is $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ such that $D_0^{hd}(f) \geq \alpha n - o(n)$ for some $\alpha > 1/\log \frac{2}{3-\sqrt{5}}$ (it is hard to believe that this is in fact correct constant).

3. Is there any $\alpha < 1$ such that for any $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$, $D_0^{hd}(f) \leq \alpha n + o(n)$?

4. Prove that for any $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$, $D_s^{hd}(f) \geq n - o(n)$.

5. Prove linear lower bound for KW relation for some explicit function in the new models.

References

[1] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, New York, NY, USA, 2006.

[2] Jeff Edmonds, Russell Impagliazzo, Steven Rudich, and Jirí Sgall. Communication complexity towards lower bounds on circuit depth. *Computational Complexity*, 10(3):210–246, 2001. URL: https://doi.org/10.1007/s00037-001-8195-x, doi:10.1007/s00037-001-8195-x.

[3] Dmitry Gavinsky, Or Meir, Omri Weinstein, and Avi Wigderson. Toward better formula lower bounds: The composition of a function and a universal relation. *SIAM J. Comput.*, 46(1):114–131, 2017. URL: https://doi.org/10.1137/15M1018319, doi:10.1137/15M1018319.

[4] Johan Håstad and Avi Wigderson. Composition of the universal relation. In Jin-Yi Cai, editor, *Advances In Computational Complexity Theory, Proceedings of a DIMACS Workshop, New Jersey, USA, December 3-7, 1990*, volume 13 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 119–134. DIMACS/AMS, 1990. URL: http://dimacs.rutgers.edu/Volumes/Vol13.html.

[5] 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. URL: https://doi.org/10.1007/BF01206317, doi:10.1007/BF01206317.

[6] 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.

[7] V. M. Khrapchenko. A method of obtaining lower bounds for the complexity of π-schemes. *Mathematical Notes Academy of Sciences USSR*, 10:474–479, 1972.

[8] Eyal Kushilevitz and Noam Nisan. *Communication complexity*. Cambridge University Press, 1997.

[9] 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. URL: http://doi.acm.org/10.1145/800135.804414, doi:10.1145/800135.804414.

Appendix

Proof of Theorem 11

Proof. Let R_c be the rectangle of all possible inputs and $\mu(R) = \big|\{(x,x) \in R\}\big|$. Consider the following set of 5 good rectangles: $R_{spent} = R_{00} + R_{01} + R_{10} + R_{11}$, and four rectangles

$$R_{\bar{1}\bar{1}} = R_{00} \cup R_{0r} \cup R_{r0} \cup R_{rr}, \qquad R_{\bar{0}\bar{1}} = R_{10} \cup R_{1r} \cup R_{r0} \cup R_{rr},$$
$$R_{\bar{1}\bar{0}} = R_{01} \cup R_{0r} \cup R_{r1} \cup R_{rr}, \qquad R_{\bar{0}\bar{0}} = R_{11} \cup R_{1r} \cup R_{r1} \cup R_{rr},$$

where Alice does not send α and Bob does not send β some fixed bits α, β.

Now let us observe that together all these good rectangles cover the entire rectangle of possible input twice, and hence one of it has measure at least $2/5 \cdot \mu(R_c)$.

□

Proof of Theorem 12

Proof. Let R_c be the rectangle of all possible inputs and $\mu(R) = |R|$. We use a set of good rectangles consisted of rectangles $R_{spent}, R_{\bar{1}\bar{1}}, R_{\bar{0}\bar{1}}, R_{\bar{1}\bar{0}}, R_{\bar{0}\bar{0}}$ from the proof of Theorem 11 and four additional rectangles

$$R_{0*} = R_{00} \cup R_{01} \cup R_{0r}, \qquad R_{*0} = R_{00} \cup R_{10} \cup R_{r0},$$
$$R_{1*} = R_{10} \cup R_{11} \cup R_{1r}, \qquad R_{*1} = R_{01} \cup R_{11} \cup R_{r1},$$

where one of players sends some fixed bit. The following lemma shows that for this set of good rectangles and this specific measure we can prove a better bound.

Lemma 3. *For all half-duplex protocols with adversary*

$$\max\big\{\mu(R_{spent}), \mu(R_{0*}), \mu(R_{*0}), \mu(R_{1*}), \mu(R_{*1}), \mu(R_{\bar{1}\bar{1}}), \mu(R_{\bar{0}\bar{1}}), \mu(R_{\bar{1}\bar{0}}), \mu(R_{\bar{0}\bar{0}})\big\} \geq \frac{3}{7} \cdot \mu(R_c).$$

Proof. We use the idea we have already seen in the proof of Theorem 8. Let a_0, a_1 and a_r be the probabilities over all possible inputs that Alice sends 0, sends 1 and receives, respectively. Analogously, we define b_0, b_1 and b_r to be the probabilities that Bob sends 0, sends 1 and receives. It is easy to see that $a_0 + a_1 + a_r = b_0 + b_1 + b_r = 1$ and for all $\alpha, \beta \in \{0, 1, r\}$, $\mu(R_{\alpha\beta}) = a_\alpha \cdot b_\beta \cdot \mu(R_c)$ (it is important here that $\mu(R) = |R|$). Minimization of maximum of linear functions with such constraints can be reduced to a semidefinite programming problem. Its solution gives us a desired bound. □

Application of the Lemma 1 for $\mu_r = 4^n$, $\mu_\ell = 2^n$ and $\alpha = 3/7$, finishes the proof. □

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il