



Half-duplex communication complexity

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

July 7, 2021

Abstract

Suppose Alice and Bob are communicating in order to compute some function f , but instead of a conventional communication channel, they have a pair of walkie-talkie devices. They can use some classical communication protocol for f , where every round one of the players sends a bit, and the other receives it. The question is whether talking via walkie-talkie gives them more power? Using walkie-talkies 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). The motivation for this kind of communication model comes from the study of the KRW conjecture. 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 also prove lower bounds for these models using both combinatorial and information-theoretic methods.

1 Introduction

In the classical communication complexity model introduced by Yao [12] two players, Alice and Bob, are trying to compute $f(x, y)$, for some function f , where Alice knows only x and Bob knows only y . 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: one has to hold a “push-to-talk” button to speak to another person, and one has to release it to listen. If 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 then communication works like in the classical case, we call such rounds *normal*. If both players send bits during the 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 players receive, we call these

*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

rounds *silent*. We distinguish three possible models, based on what happens in silent rounds. If in silent rounds 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*. Similar models were studied in [7] (without “canceling” spent rounds) and in [3] (for multi-party communication with the noisy broadcast channel). Two other models, we will define later.

In this paper, we study the communication complexity of Boolean functions that are hard in the classical case. It is important to note that we care about multiplicative constants. Every classical communication can be viewed as half-duplex communication with zero and every half-duplex communication with zero can be simulated with classical communication doubling the number of rounds (see Theorem 6 and 7). So the complexity of half-duplex communication is sandwiched between the complexity of the classical case and a half of it. The task of this study is to improve these bounds.

1.1 Motivation

The original motivation to study these kinds of communication models arose from the question of the complexity of the Karchmer-Wigderson games [9] for multiplexers. *The Karchmer-Wigderson game for a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ (KW game)* 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 an $i \in [n]$ such that $x_i \neq y_i$. Let $D(KW(f))$ be a minimal number of rounds that is enough to solve the KW game for f on any pair of possible inputs.

Conjecture 1 (The KRW conjecture [8]). *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and $g : \{0, 1\}^m \rightarrow \{0, 1\}$ be Boolean non-constant functions. Then $D(KW(g \circ f)) \approx D(KW(g)) + D(KW(f))$, where $g \circ f$ denotes a composition $g \circ f : (\{0, 1\}^n)^m \rightarrow \{0, 1\}$ is defined by $(g \circ f)(x_1, \dots, x_m) = g(f(x_1), \dots, f(x_m))$ where $x_1, \dots, x_m \in \{0, 1\}^n$.*

This conjecture implies a super-logarithmic formula depth lower bound (and hence a super-polynomial size lower bound): we can start with a maximally hard function on $\log n$ variables that requires $\log n$ depth and construct a formula on n variables that requires super-logarithmic depth. In attempt to prove it a lot of work has been done studying KW games where one or both functions are replaced with *universal relations* [5, 2, 4]. Another approach to resolving the conjecture lies in examining multiplexer functions. A *multiplexer* (or *indexing function*) is a function $M_n : \{0, 1\}^{2^n} \times \{0, 1\}^n \rightarrow \{0, 1\}$, 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 $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and the second part as an input x to the function, and outputs $f(x)$. Multiplexers are similar to universal relations in the sense that there is a natural reduction from a KW game for some function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ to a KW 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, not relations, so proving analogous results for multiplexers would be one step toward proving the KRW conjecture. 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 the symmetry of universal relations that allows giving 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 solving the KW 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 $f = g$ (note that f and g are parts players inputs, so Alice and Bob plays KW game for M_n on a subset of inputs) then they can use a protocol for KW game for f . However, what if they do not have such a promise (i.e., all inputs are possible, in particular, such that $f \neq g$)? Alice can still try to act as if she plays KW game for f , Bob at the same time can try to act as if he plays KW game for g , but if in fact $f \neq g$ then 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” is impossible 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? When answering this question we care about multiplicative constants — if in this model all (hard) functions become two times easier than in the classical case then this model is useless for proving the KRW conjecture. As a first step toward answering this question, we study the half-duplex communication complexity of Boolean functions $\{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$.

1.2 Organization of this paper

In Section 2, we give definitions for the new communication models. Then, in Section 3, we prove trivial upper and lower bounds that follow immediately from the definitions. Next, in Section 4, we discuss methods 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.

1.3 Changes from previous versions

In the previous versions on this paper including the conference version [6] we claimed better lower bounds for IP_n in all tree half-duplex models based on an information theoretic approach. Unfortunately, proofs of these results (Theorems 18, 21, and 24 in [6]) contain a critical flaw (Lemma 13 in [6] is not true). Moreover, as far as we know, this flaw can not be fixed — there is an upper bound of $n/2 + 2$ on half-duplex complexity of IP_n with silence (proven by Tatiana Gladyshevskaya) that matches the lower bound proven here up to an additive constant 2. This upper bound contradicts the lower bound claimed in the previous versions of the paper. All these theorems were removed from the paper preserving numeration of other theorems.

2 Definitions

Definition 1. Let X , Y , and Z be some finite sets. We say that two players, Alice and Bob, are solving the *half-duplex communication problem* for a relation $R \subseteq X \times Y \times Z$ if sets X , Y , Z , and the 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 $z \in Z$ such that $(x, y, z) \in R$, by communicating to each other via a half-duplex channel. The communication is organized into rounds. At each round, both players decide (depending only on their inputs 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 crucial that the player that is sending cannot 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 a special symbol `silence`, so it is possible for both players to distinguish a silent round from a normal one, the corresponding problem is called *a 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 *a 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 *a 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 the 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 a player does not always know what the other's player 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. Note also that solving half-duplex communication problem with zero there is no need to send zeros — a 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 (of T_B) is labeled with two functions $g_v : X \rightarrow \mathcal{A}$ ($g_v : Y \rightarrow \mathcal{A}$) and $h_v : \mathcal{E} \rightarrow C(v)$, 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 the players reach a leaf they stop (they always reach a leaf 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 a different set of possible events $\mathcal{E} = \{\text{send}(1), \text{receive}(0), \text{receive}(1)\}$, i.e it does not include `send(0)`.

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 the adversary.

Definition 3. Half-duplex communication protocol *solves* a communication problem for function $f : X \times Y \rightarrow Z$ 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, defines *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 $\{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ (i.e., $X = Y = \{0, 1\}^n$, $Z = \{0, 1\}$).

Definition 5.

- *Equality function* $\text{EQ}_n : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, such that $\text{EQ}_n(x, y) = 1 \iff x = y$.
- *Inner product function* $\text{IP}_n : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, such that $\text{IP}_n(x, y) = \bigoplus_{i \in [n]} x_i y_i$.
- *Disjointness function* $\text{DISJ}_n : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, such that $\text{DISJ}_n(x, y) = 1 \iff \forall i : x_i \neq 1 \vee y_i \neq 1$.

All these function require n bits of communication in the classical model.

3 Trivial bounds

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

Theorem 6. For every function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, $D_s^{hd}(f) \leq D_0^{hd}(f) \leq D_a^{hd}(f) \leq D(f)$.

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

Next theorem shows that one can always transform half-duplex protocol with zero or with the adversary into a classical communication protocol of double depth.

Theorem 7. For every function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, $\frac{D(f)}{2} \leq D_0^{hd}(f) \leq D_a^{hd}(f)$.

Proof. Every t -round half-duplex communication protocol with zero or with the 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: on the 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. \square

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

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: on the 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 could skip the third round. If they were both planning to receive, then they can assume that they heard silence. If one player was planning to send and the other one was planning to receive they can perform such action on the third round. \square

Remark 1. Theorems 6, 7, and 8 hold also for $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^k$.

4 Methods for lower bounds

4.1 Rectangles

Many lower bounds on classical communication complexity were proved by considering combinatorial rectangles associated with the nodes of communication protocol [11]: it is 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 [11] for proving lower bounds using associated combinatorial 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 $\mu_r > 0$, 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 $\mu_\ell > 0$, 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 $\mu_\ell = 1$.

We can use almost the same technique for half-duplex protocols. There are some technical differences that we have to keep in mind. First of all, we can apply this idea to both trees T_A and

T_B . We should also note that trees T_A and T_B are non-binary; hence arity became the base of the logarithm. 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 as a part of input. This leads to the following lower bound for equality.

Theorem 9.

- $D_s^{hd}(\text{EQ}_n) \geq \log_5 2^n = n/\log 5$,
- $D_0^{hd}(\text{EQ}_n) \geq \log_3 2^n = n/\log 3$,
- $D_a^{hd}(\text{EQ}_n) \geq \log_4 2^n = n/2$.

Proof. Let $\mu = \mu_M$. All leaf rectangles 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 $\mu_r = 2^n + 1$ (see [11] for more information). \square

Surprisingly, as we will see later, first two result are sharp up to additive logarithmic term. We developed an extension of this technique that we call *round elimination*.

4.2 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., the 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 decides 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 is necessary to proceed to the next round. Lets call a rectangle R *good for (partial) function 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 10. *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 , $\mu(R_l) \leq \mu_\ell$. Fix a protocol \mathcal{P} . If for any rectangle R appearing in the protocol there is 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 $\log_{1/\alpha} \frac{\mu_r}{\mu_\ell}$.*

Proof. We start with $R = X \times Y$. Every round we show that $f \upharpoonright R$ can be restricted to some good $R_{good} \subset R$ such that $\mu(R_{good}) \geq \alpha \cdot \mu(R)$, let R to be R_{good} , and proceed to the next round until we reach a leaf. Thus there are at least $\log_{1/\alpha}(\mu_r/\mu_\ell)$ rounds. \square

4.3 Upper bound on internal information

Another useful tool for proving lower bounds on the communication complexity of problems in the classical model is upper bounding the information Alice and Bob have learned about the other's inputs, as a function of the number of rounds.

Let \mathcal{P} be a half-duplex communication protocol for relation $R \subset X \times Y \times Z$, and consider an arbitrary distribution \mathcal{D} over the pair of inputs of Alice and Bob. Let \mathcal{X} , and \mathcal{Y} be the marginal distributions over inputs to Alice and Bob, also, let Π_A and Π_B be the marginal distributions over Alice and Bob's transcripts induced by \mathcal{D} .

Definition 11. An *internal information cost of protocol \mathcal{P}* is

$$\text{IC}_{\mathcal{D}}(\mathcal{P}) = I(\mathcal{X} : \Pi_B \mid \mathcal{Y}) + I(\mathcal{Y} : \Pi_A \mid \mathcal{X}).$$

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} . An *internal information cost of first k rounds of \mathcal{P}* is

$$\text{IC}_{\mathcal{D}}^k(\mathcal{P}) = I(\mathcal{X} : \Pi_B^k \mid \mathcal{Y}) + I(\mathcal{Y} : \Pi_A^k \mid \mathcal{X}).$$

For more details on information theory, we refer to [1, 4].

Lemma 12. Suppose that $H(\mathcal{X} \mid \mathcal{Y}) + H(\mathcal{Y} \mid \mathcal{X}) \geq \rho$. If for any half-duplex communication protocol with silence/zero/adversary \mathcal{P} computing R ,

- $H(\mathcal{X} \mid \mathcal{Y}, \Pi_A) + H(\mathcal{Y} \mid \mathcal{X}, \Pi_B) \leq \lambda$, and
- for every k , $\text{IC}_{\mathcal{D}}^k(\mathcal{P}) \leq \alpha k$,

then half-duplex complexity of R with silence/zero/adversary is at least $(\rho - \lambda)/\alpha$.

Proof. We know that before any communication Alice and Bob have total uncertainty at least ρ about $(\mathcal{X}, \mathcal{Y})$. After the protocol ends, Alice and Bob have total uncertainty at most λ about $(\mathcal{X}, \mathcal{Y})$. Thus, during the protocol Alice and Bob have to learn at least $\rho - \lambda$ bits of information. In k rounds the players can learn at most αk bits. Hence, the number of rounds is at least $(\rho - \lambda)/\alpha$. \square

5 Half-duplex communication with silence

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

Theorem 14. For every $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, $D_s^{hd}(f) \leq \lceil n/\log 3 \rceil + 1$.

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 evaluates $f(x, y)$ and sends the result to Alice. \square

Using the idea of non-binary encoding, we prove a better upper bound for equality.

Theorem 15. $D_s^{hd}(\text{EQ}_n) \leq \lceil n/\log 5 \rceil + \lceil \log n/\log 3 \rceil + 2.$

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 i th 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 i th round is normal then one player can check whether i th symbols are different. If i th round is silent then again one player knows if i th 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. To check it, Alice sends the number of normal rounds she was receiving 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

Using almost the same ideas we can show an upper bound for disjointness.

Theorem 16. $D_s^{hd}(\text{DISJ}_n) \leq \lceil n/2 \rceil + 2.$

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

Using round elimination we get the following lower bound for the inner product function.

Theorem 17. $D_s^{hd}(\text{IP}_n) \geq n/2.$

To prove this theorem we need the following property of IP_n .

Lemma 18. *Every leaf rectangle of a protocol 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 $y \in Y_l$, $\text{IP}_n(x, y) = 0$, 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, $|R_l| = |X_l| \times |Y_l| \leq 2^n$.

If leaf is labeled with 1 then for every $x \in X_l$ and $y \in Y_l$, $\text{IP}_n(x, y) = 1$. 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 $y \in Y'_l$, $\text{IP}_n(x, y) = 0$, 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

Now we are ready to prove Theorem 17.

Proof of Theorem 17. Let R_c be the rectangle of all possible inputs and $\mu(R) = |R|$. Consider the following set of good rectangles: a rectangle $R_{\text{silent}} = R_{rr}$ 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 the players sends some bit, and a rectangle $R_{\text{spent}} = R_{00} \cup R_{01} \cup R_{10} \cup R_{11}$, 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 $\max\{\mu(R_{0*}), \mu(R_{1*}), \mu(R_{*0}), \mu(R_{*1}), \mu(R_{\text{silent}}), \mu(R_{\text{spent}})\} \geq \mu(R_c)/4$. This is equivalent to showing that $\max\{a_1, a_0, b_1, b_0, a_r b_r, (1 - a_r)(1 - b_r)\} \geq 1/4$ for any reals $a_0, a_1, a_r, b_0, b_1, b_r \in [0, 1]$, such that $a_0 + a_1 + a_r = b_0 + b_1 + b_r = 1$. If $a_0, a_1, b_0, b_1 < 1/4$ then $(1 - a_r)(1 - b_r) > 1/4$. Now we apply Lemma 10 for $\mu_r = 4^n$, $\mu_\ell = 2^n$ (Lemma 18), $\alpha = 1/4$, and get the desired bound. \square

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 compute equality in less than n rounds of communication.

Theorem 19. $D_0^{\text{hd}}(\text{EQ}_n) \leq \lceil n/\log 3 \rceil + 2\lceil \log n \rceil + 1$.

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 i th 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. Depending on values of corresponding inputs, i.e., x_i and y_i , we distinguish 6 types of witnesses of inequality:

(0, 1), (1, 0), (0, 2), (2, 0), (1, 2), (2, 1). If we make sure that each type can be detected by at least one of the players we are done. In the first phase, Alice can detect types (0, 2), (2, 0), (2, 1), while Bob can detect types (1, 0), (0, 1), and (2, 1) (again). This leaves us with detecting witnesses of type (1, 2). Assuming that there are no witnesses of other types, this will be detected in the second phase. \square

The best lower bound for this model is again for IP_n . The following theorem is proved using round elimination.

Theorem 20. $D_0^{hd}(\text{IP}_n) \geq n / \log \frac{2}{3-\sqrt{5}} > n/1.39$.

Proof. Let R_c be the rectangle of all possible inputs and $\mu(R) = |R|$. Consider the following set of good rectangles: $R_{\text{silent}} = R_{rr}$, $R_{1*} = R_{11} \cup R_{1r}$ 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

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

It is equivalent to showing that for any $a, b \in [0, 1]$, $\max\{a, b, (1-a)(1-b)\} \geq \frac{3-\sqrt{5}}{2}$, where a and b denote the probabilities over all possible inputs that, respectively, Alice and Bob sends 1. Now we can argue that minimum value 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 $\max\{a, b, (1-a)(1-b)\} \geq \frac{3-\sqrt{5}}{2}$. Applying Lemma 10 for $\mu_r = 4^n$, $\mu_\ell = 2^n$ (Lemma 18), and $\alpha = \frac{3-\sqrt{5}}{2}$ finishes the proof. \square

7 Half-duplex communication with adversary

The main feature of this model is that receiving player cannot 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. The protocol must be correct for any strategy of the adversary. Our intuition prompts that in this setting silent and spent rounds would be useless. Using combinatorial methods, one can show the following two lower bounds.

Theorem 22. $D_a^{hd}(\text{EQ}_n) \geq n / \log 2.5$.

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

$$\begin{aligned} R_{\bar{1}\bar{1}} &= R_{00} \cup R_{0r} \cup R_{r0} \cup R_{rr}, & R_{\bar{0}\bar{1}} &= R_{10} \cup R_{1r} \cup R_{r0} \cup R_{rr}, \\ R_{\bar{1}\bar{0}} &= R_{01} \cup R_{0r} \cup R_{r1} \cup R_{rr}, & R_{\bar{0}\bar{0}} &= R_{11} \cup R_{1r} \cup R_{r1} \cup R_{rr}, \end{aligned}$$

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 inputs twice, and hence one of it has measure at least $2/5 \cdot \mu(R_c)$. \square

Theorem 23. $D_a^{hd}(\text{IP}_n) \geq n / \log \frac{7}{3}$.

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 22, and four additional rectangles

$$\begin{aligned} R_{0*} &= R_{00} \cup R_{01} \cup R_{0r}, & R_{*0} &= R_{00} \cup R_{10} \cup R_{r0}, \\ R_{1*} &= R_{10} \cup R_{11} \cup R_{1r}, & R_{*1} &= R_{01} \cup R_{11} \cup R_{r1}, \end{aligned}$$

where one of the players sends some fixed bit.

Now we show that for this set of good rectangles and this specific measure we can prove a better bound. In fact, for all half-duplex protocols with adversary

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

We use the idea we have already seen in the proof of Theorem 17. 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 giving us the desired bound. Application of the Lemma 10 for $\mu_r = 4^n$, $\mu_\ell = 2^n$ (Lemma 18) and $\alpha = 3/7$, finishes the proof. \square

Now we proceed to proving $2 \log n$ lower bound on the half-duplex complexity of Karchmer-Wigderson relation for parity function.

Definition 25. Let $X = f^{-1}(0)$, $Y = f^{-1}(1)$ for some Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$. The *KW relation for function f* , $R_f \subseteq X \times Y \times [n]$, is defined by $R_f = \{(x, y, i) \mid x_i \neq y_i\}$.

It is well known that parity function $\oplus_n : \{0, 1\}^n \rightarrow \{0, 1\}$, $\oplus_n(x) = \sum_{i=1}^n x_i \pmod 2$, requires $\Omega(n^2)$ De Morgan formula size [10]. In the classical communication complexity, it is equivalent to saying that KW relations for parity requires $2 \log n$ rounds of communication [9, 4].

We will prove $2 \log n$ lower bound in the half-duplex model with adversary using information-theoretic approach. We start with proving an upper bound on the information Alice and Bob have learned about the other's inputs, as a function of the number of rounds. Next, we show that there is a distribution on inputs of players such that we can show a lower bound on the amount of information the players have to transmit to each other. Together that gives us a lower bound on the number of rounds.

In the following lemma we show that in the half-duplex model with adversary the players need at least k rounds to transmit k bits of information.

Lemma 26. *For any half-duplex protocol \mathcal{P} with adversary, arbitrary probability distribution \mathcal{D} , and natural k ,*

$$\text{IC}_{\mathcal{D}}^k(\mathcal{P}) \leq k.$$

Proof. It suffices to show that

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

Assume that adversarial bits that players receive in silent rounds are independent uniformly random distributed. Let (y, π_B^k) be a particular valid input-transcript pair for Bob. Consider $I(\mathcal{X} : \mathcal{E}_B^{k+1} | E_{y, \pi_B^k})$ where E_{y, π_B^k} denotes an event “ $\mathcal{Y} = y, \Pi_B^k = \pi_B^k$ ”. Note that

$$I(\mathcal{X} : \mathcal{E}_B^{k+1} | E_{y, \pi_B^k}) \leq I(\mathcal{X}, \Pi_A^k : \mathcal{E}_B^{k+1} | E_{y, \pi_B^k}) = H(\mathcal{E}_B^{k+1} | E_{y, \pi_B^k}) - H(\mathcal{E}_B^{k+1} | E_{y, \pi_B^k}, \mathcal{X}, \Pi_A^k).$$

Suppose that on input y and given a partial transcript π_B^k Bob decides to receive in round $k + 1$; otherwise \mathcal{E}_B^{k+1} depends only on y and π_B^k , and hence $H(\mathcal{E}_B^{k+1} | E_{y, \pi_B^k}) = H(\mathcal{E}_B^{k+1} | E_{y, \pi_B^k}, \mathcal{X}, \Pi_A^k) = 0$. Consider each (x, π_A^k) input-transcript pair for Alice consistent with (y, π_B^k) . Note that

$$H(\mathcal{E}_B^{k+1} | \mathcal{Y} = y, \Pi_B^k = \pi_B^k, \mathcal{X} = x, \Pi_A^k = \pi_A^k)$$

will either be 0, if Alice is sending a bit in round $k + 1$, or 1, if she is receiving (i.e., round $k + 1$ is silent). 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

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

and hence $I(\mathcal{X} : \mathcal{E}_B^{k+1} | \mathcal{Y} = y, \Pi_B^k = \pi_B^k) \leq 1 - \Pr[\text{Alice receives} | E_{y, \pi_B^k}] \leq \Pr[\text{Alice sends} | E_{y, \pi_B^k}]$. We then have that

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

A symmetric argument holds for Alice, giving

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

Theorem 27. $D_a^{hd}(R_{\oplus n}) \geq 2 \log n$.

Proof. Let \mathcal{D} be the uniform distribution over valid input pairs with a single bit of difference: \mathcal{X} is uniformly distributed over n -bit strings with even number of ‘1’, and $\mathcal{Y} = \mathcal{X} \oplus \mathcal{I}$, where \mathcal{I} is uniformly distributed over n -bit strings with exactly one ‘1’. Thus, before any communication takes place,

$$H(\mathcal{Y} | \mathcal{X}) + H(\mathcal{X} | \mathcal{Y}) = 2H(\mathcal{I}) = 2 \log n.$$

On the other hand, consider some half-duplex protocol \mathcal{P} for $R_{\oplus n}$ with adversary. It is easy to see that $H(\mathcal{Y} | \mathcal{X}, \Pi_A) + H(\mathcal{X} | \mathcal{Y}, \Pi_B) = 0$ at any leaf. So the players have to learn at least $2 \log n$ bits of information, $\text{IC}_D(\mathcal{P}) \geq 2 \log n$. Applying Lemma 12 for $\rho = 2 \log n$, $\lambda = 0$, and $\alpha = 1$ (Lemma 26), we get the desired lower bound on the number of rounds. \square

8 Open problems

The following table lists lower and upper bounds that we managed to prove in this paper.

	EQ_n	IP_n	DISJ_n
D_s^{hd}	$\geq n/\log 5$ $\leq n/\log 5 + o(n)$	$\geq n/2$ $\leq n/\log 3$	$\geq n/3$ $\leq n/2 + O(1)$
D_0^{hd}	$\geq n/\log 3$ $\leq n/\log 3 + o(n)$	$\geq n/\log(2/(3 - \sqrt{5}))$	$\geq n/2$
D_a^{hd}	$\geq n/\log 2.5$	$\geq n/\log(7/3)$	$\geq n/2$

It would be interesting to improve bounds where it is possible to determine the true complexity of these functions in half-duplex models. So we propose the following list of open problems.

1. Prove better upper and lower bounds for the half-duplex model with silence and zero.
2. Is there any $\alpha < 1$ such that for any $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, $D_0^{hd}(f) \leq \alpha n + o(n)$?
3. Is there any $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, such that at the same time $D(f) \geq n - o(n)$ and $D_a^{hd}(f) \leq \alpha n + o(n)$ for some $\alpha < 1$.

Acknowledgement

We thank the anonymous reviewers whose insightful comments and suggestions helped us to make the paper more readable and to simplify the proofs.

References

- [1] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, New York, 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. [doi:10.1007/s00037-001-8195-x](https://doi.org/10.1007/s00037-001-8195-x).
- [3] Klim Efremenko, Gillat Kol, and Raghuvansh Saxena. Interactive coding over the noisy broadcast channel. *Electronic Colloquium on Computational Complexity (ECCC)*, 24:93, 2017. URL: <https://eccc.weizmann.ac.il/report/2017/093>.
- [4] 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. [doi:10.1137/15M1018319](https://doi.org/10.1137/15M1018319).
- [5] 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>.

- [6] 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](https://doi.org/10.4230/LIPICs.ISAAC.2018.10).
- [7] Russell Impagliazzo and Ryan Williams. Communication complexity with synchronized clocks. In *Proceedings of the 25th Annual IEEE Conference on Computational Complexity, CCC 2010, Cambridge, Massachusetts, USA, June 9-12, 2010*, pages 259–269. IEEE Computer Society, 2010. doi:[10.1109/CCC.2010.32](https://doi.org/10.1109/CCC.2010.32).
- [8] 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. doi:[10.1007/BF01206317](https://doi.org/10.1007/BF01206317).
- [9] 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](https://doi.org/10.1145/62212.62265).
- [10] 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.
- [11] Eyal Kushilevitz and Noam Nisan. *Communication complexity*. Cambridge University Press, 1997.
- [12] 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](https://doi.org/10.1145/800135.804414).