Electronic Colloquium on Computational Complexity, Report No. 123 (2021)

ON PUBLIC-COIN ZERO-ERROR RANDOMIZED COMMUNICATION COMPLEXITY

BEN DAVIS, HAMED HATAMI, WILLIAM PIRES, RAN TAO, AND HAMZA USMANI

Abstract. We prove that for every Boolean function, the public-coin zero-error randomized communication complexity and the deterministic communication complexity are polynomially equivalent.

1. Introduction

The field of communication complexity studies the amount of communication required to solve the problem of computing discrete functions when the input is split between two or more parties. In the most commonly studied framework, there are two parties, often called Alice and Bob, and a communication problem is defined by a Boolean function $f : \mathcal{X} \times \mathcal{Y} \to \{0, 1\}$. An input $x \in \mathcal{X}$ is given to Alice, and an input $y \in \mathcal{Y}$ is given to Bob. Together, they should both compute the entry $f(x, y)$ by exchanging bits of information in turn, according to a previously agreed-on protocol. There is no restriction on their computational power; the only measure we care to minimize is the number of exchanged bits.

A deterministic protocol π specifies for each of the two players, the bit to send next, as a function of their input and history of the communication so far. It naturally corresponds to a binary tree as follows. Every internal node is associated with either Alice or Bob. If an internal node v is associated with Alice, then it is labeled with a function $a_v : \mathcal{X} \to \{0, 1\}$, which prescribes the bit sent by Alice at this node as a function of her input. After this bit was sent, the players move to the corresponding child of v: they move to the left child if the bit is 0, and to the right child if the bit is 1. Similarly, Bob's nodes are labeled with Boolean functions on $\mathcal{Y}$. Each leaf is labeled by 0 or 1, which corresponds to the output of the protocol. The *cost* of the protocol on an input (x, y), denoted by $\text{cost}(\pi(x, y))$, is the number of bits exchanged on this input. The cost of the protocol is the maximum of $\text{cost}(\pi(x, y))$ over all inputs (x, y). The *deterministic communication complexity* of f, denoted by $\text{D}(f)$, is the smallest cost of a protocol that computes f correctly on all inputs.

Next we discuss randomized communication complexity. The randomness can be introduced in two different ways: private randomness or public randomness.

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

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

In this article, we are interested in *zero-error* randomized communication protocols. These protocols, which are also called Las Vegas protocols, are not allowed to make any errors, and their

HH was supported by an NSERC Discovery Grant. This project was carried as a summer student research project at McGill University under the supervision of HH.

ISSN 1433-8092

costs are defined as the *expected number* of bits exchanged over the worst input (x, y). The private-coin zero-error randomized communication complexity of a function $f : \mathcal{X} \times \mathcal{Y} \to \{0, 1\}$, denoted by $\mathrm{R}_0^{\mathrm{prv}}(f)$, is the infimum cost over all private-coin randomized protocols that compute f with zero error:

$$\mathrm{R}_0^{\mathrm{prv}}(f) = \inf_{\pi_{R_A,R_B}} \max_{(x,y)} \mathbb{E}_{R_A,R_B} \left[\mathrm{cost}(\pi_{R_A,R_B}(x, y))\right],$$

where the infimum is over all zero-error *private-coin* protocols π_{R_A,R_B}. The public-coin zero-error randomized communication complexity is defined similarly using public randomness and is denoted by $\mathrm{R}_0^{\mathrm{pub}}(f)$. That is

$$\mathrm{R}_0^{\mathrm{pub}}(f) = \inf_{\pi_R} \max_{(x,y)} \mathbb{E}_R \left[\mathrm{cost}(\pi_R(x, y))\right],$$

where the infimum is over all zero-error *public-coin* protocols π_R.

We have the trivial relations

$$\mathrm{R}_0^{\mathrm{pub}}(f) \le \mathrm{R}_0^{\mathrm{prv}}(f) \le \mathrm{D}(f).$$

A fundamental fact in communication complexity is that the leaves of every deterministic protocol computing f partitions f into rectangles (a rectangle is a set $S \times T \subseteq \mathcal{X} \times \mathcal{Y}$) such that the value of f is constant on each rectangle. Such rectangles are called *monochromatic*. Hence every deterministic protocol of cost c provides a partition of $\mathcal{X} \times \mathcal{Y}$ into at most 2^c monochromatic rectangles. Since the rank (as a real matrix) of a monochromatic rectangle is at most one, we obtain the classical lower-bound

$$\mathrm{D}(f) \ge \log \mathrm{rk}(f), \tag{1}$$

where $\mathrm{rk}(f)$ denotes the rank of f as a matrix over the reals.

Let $\mathrm{N}(f)$ denote the logarithm of the minimum number of monochromatic rectangles required to *cover* $\mathcal{X} \times \mathcal{Y}$. It follows from the above discussion that $\mathrm{N}(f) \le \mathrm{D}(f)$, and furthermore as it is shown in [AUY83] the gap can be at most quadratic:

$$\mathrm{D}(f) = O(\mathrm{N}(f)^2). \tag{2}$$

The quantity $\mathrm{N}(f)$ can also be used to lower-bound $\mathrm{R}_0^{\mathrm{prv}}(f)$. Indeed, a private-coin protocol can be interpreted as a deterministic protocol where Alice's input is (x, R_A) and Bob's input is (y, R_B). Hence the leaves of such a protocol provide a *partition* of $(\mathcal{X} \times \Omega_A) \times (\mathcal{Y} \times \Omega_B)$ into combinatorial rectangles, where Ω_A and Ω_B are the supports of the random strings R_A and R_B, respectively. If the protocol is of zero-error, then every rectangle $S \times T \subseteq (\mathcal{X} \times \Omega_A) \times (\mathcal{Y} \times \Omega_B)$ in this partition corresponds to a *monochromatic* rectangle $\widehat{S} \times \widehat{T} \subseteq \mathcal{X} \times \mathcal{Y}$ defined as

$$\widehat{S} = \{x \in \mathcal{X} \ : \ \exists r_A \in \Omega_A,\ (x, r_A) \in S\}, \qquad \widehat{T} = \{y \in \mathcal{Y} \ : \ \exists r_B \in \Omega_B,\ (y, r_B) \in T\}.$$

Furthermore if the protocol has cost c, then for every (x, y), we have $\mathbb{E}[\mathrm{cost}(\pi_{R_A,R_B}(x, y))] \le c$, and thus for every (x, y), there exists a choice of r_A and r_B such that (x, r_A, y, r_B) leads to a leaf of depth at most c. Hence the set of monochromatic rectangles $\widehat{S} \times \widehat{T}$ that correspond to the leaves of depth at most c provide a cover of $\mathcal{X} \times \mathcal{Y}$. There are at most 2^c such leaves, and thus $\mathrm{N}(f) \le \log 2^c = \mathrm{R}_0^{\mathrm{prv}}(f)$. Combining this with Equation (2) yields

$$\mathrm{R}_0^{\mathrm{prv}}(f) = \Omega(\sqrt{\mathrm{D}(f)}).$$

Furthermore as it is shown by Fürer [Für87] there are examples for which this quadratic gap is necessary.

Perhaps the most famous result relating public-coin and private-coin models is Newman's lemma [New91], which in the zero-error case (see [KN97, Exercise 3.15]) states

$$\mathrm{R}_0^{\mathrm{prv}}(f) \le O(\mathrm{R}_0^{\mathrm{pub}}(f) + \log\log|\mathcal{X} \times \mathcal{Y}|). \tag{3}$$

Note that since $\log\log|\mathcal{X}\times\mathcal{Y}|$ can be the dominant term in Equation (3), Newman's lemma does not provide a polynomial relation between the public-coin and private-coin zero-error communication complexities. To the best of our knowledge, it was unknown whether $\mathrm{R}_0^{\mathrm{pub}}(f)$ is polynomially equivalent to $\mathrm{R}_0^{\mathrm{prv}}(f)$ and $\mathrm{D}(f)$. The aim of this article is to establish such a relation.

Theorem 1.1 (Main Theorem)**.** *For every $f:\mathcal{X}\times\mathcal{Y}\to\{0,1\}$ we have*

$$\mathrm{R}_0^{\mathrm{pub}}(f)=\Omega(\mathrm{D}(f)^{1/4}).$$

2. Proof of Theorem 1.1

The key insight for proving Theorem 1.1 is to consider the relation between $\mathrm{rk}(f)$ and $\mathrm{D}(f)$. If $\log\mathrm{rk}(f)$ is small, then by applying a classical result of Nisan and Widgerson [NW95], one can deduce that $\mathrm{D}(f)$ cannot be much larger than $\mathrm{R}_0^{\mathrm{pub}}(f)$.

On the other hand, if $\log\mathrm{rk}(f)$ is large, then by Equation (1) one can consider a full-rank submatrix of f whose deterministic communication complexity is large as well. Since this submatrix is of full-rank, its dimensions are equal to its rank, and thus one can successfully apply Newman's lemma (the $\log\log(\cdot)$ term will be small in Equation (3)) to obtain a strong lower-bound on its public-coin zero-error communication complexity.

We start by stating the result of Nisan and Widgerson [NW95]. In the following statement, for a rectangle $A\subseteq\mathcal{X}\times\mathcal{Y}$, $|A|$ denotes its set-theoretic cardinality.

Lemma 2.1. [NW95] *Let $f:\mathcal{X}\times\mathcal{Y}\to\{0,1\}$ be a Boolean function, and $\delta>0$ a real number. Assume that for every rectangle $A_1\subseteq\mathcal{X}\times\mathcal{Y}$, there exists a sub-rectangle $A_2\subseteq A_1$ such that A_2 is monochromatic and $|A_2|\geq\delta|A_1|$. Then*

$$\mathrm{D}(f)\leq O(\log(1/\delta)\log\mathrm{rk}(f)+(\log^2\mathrm{rk}(f))).$$

In the following simple corollary, we observe that f satisfies the assumption of Lemma 2.1 with $\delta=2^{-2\,\mathrm{R}_0^{\mathrm{pub}}(f)-1}$.

Corollary 2.2. *There exists a universal constant $c\geq1$ such that every $f:\mathcal{X}\times\mathcal{Y}\to\{0,1\}$ satisfies*

$$\mathrm{D}(f)\leq c\,\mathrm{R}_0^{\mathrm{pub}}(f)\log\mathrm{rk}(f)+c\log^2\mathrm{rk}(f).$$

Proof. Let π_R be a zero-error public-coin communication protocol of cost $\mathrm{R}_0^{\mathrm{pub}}(f)$ for f. Here R denotes the public randomness, and any fixation of R to a string r leads to a deterministic protocol π_r. Consider a rectangle $A_1\subseteq\mathcal{X}\times\mathcal{Y}$. We have

$$\mathbb{E}_R\,\mathbb{E}_{(x,y)\in A_1}\,\mathrm{cost}(\pi_R(x,y))\leq\mathrm{R}_0^{\mathrm{pub}}(f),$$

which shows that there exists a fixation of R to r such that

$$\mathbb{E}_{(x,y)\in A_1}\,\mathrm{cost}(\pi_r(x,y))\leq\mathrm{R}_0^{\mathrm{pub}}(f).$$

Consequently,

$$\Pr_{(x,y)\in A_1}\left[\mathrm{cost}(\pi_r(x,y))>2\,\mathrm{R}_0^{\mathrm{pub}}(f)\right]\leq\frac{1}{2}.$$

This shows that the leaves of the protocol π_r that are in depth at most $2\,\mathrm{R}_0^{\mathrm{pub}}(f)$ cover at least half of the points in A_1. Since each leaf corresponds to a monochromatic sub-rectangle, one of these leaves must correspond to a monochromatic sub-rectangle of size at least $2^{-2\,\mathrm{R}_0^{\mathrm{pub}}(f)-1}|A_1|$. □

To optimize our bound, instead of applying Newman's lemma, we apply the next lemma, whose proof is almost identical to the proof of Newman's lemma.

Lemma 2.3. *Every Boolean function $f:\mathcal{X}\times\mathcal{Y}\to\{0,1\}$ satisfies*

$$\mathrm{D}(f)\leq(2\,\mathrm{R}_0^{\mathrm{pub}}(f)+\log\log|\mathcal{X}||\mathcal{Y}|+1)^2.$$

Proof. Let $f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ be a Boolean function and $\mathrm{R}_0^{\mathrm{pub}}(f) = t$, witnessed by a zero-error randomized protocol π_R. Let π'_R denote the truncation of π_R by restricting to computational paths of length at most $2t$, that is we trim off all branches at the $2t^{\text{th}}$ node and output $\perp$ indicating failure to compute $f(x,y)$ in those cases. By Markov's bound, for every $(x,y) \in \mathcal{X} \times \mathcal{Y}$, we have

$$\Pr_R[\pi'_R(x,y) \neq \perp] \geq \frac{1}{2}.$$

Let $\ell = 2\log(|\mathcal{X}||\mathcal{Y}|)$ so that $2^\ell > |\mathcal{X} \times \mathcal{Y}|$, and consider ℓ independent executions of π'_R by considering ℓ independent copies $R_1, \ldots, R_\ell$ of R. For every (x,y), we have

$$\Pr_{R_1,\ldots,R_\ell}\left[\forall i, \quad \pi'_{R_i}(x,y) = \perp\right] \leq \frac{1}{2^\ell} < \frac{1}{|\mathcal{X} \times \mathcal{Y}|}.$$

In particular there exists a choice of $r_1, \ldots, r_\ell$ such that for every $(x,y) \in \mathcal{X} \times \mathcal{Y}$, there exists at least one $i \in \{1, \ldots, \ell\}$ with $\pi'_{r_i}(x,y) \neq \perp$. Each π'_{r_i} is a deterministic protocol with communication cost at most $2t$, and thus provides a partition of the points (x,y) with $\pi_{r_i}(x,y) \neq \perp$ into at most 2^{2t} monochromatic rectangles. Consequently, there exists a collection of $\ell 2^{2t}$ monochromatic rectangles whose union is all of $\mathcal{X} \times \mathcal{Y}$, or in other words

$$\mathrm{N}(f) \leq \log\left(\ell 2^{2t}\right) = 2\,\mathrm{R}_0^{\mathrm{pub}}(f) + \log\log|\mathcal{X} \times \mathcal{Y}| + 1.$$

Now the result follows form Equation (2). □

Proof of Theorem 1.1. Let $f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ be a Boolean function, and let $c \geq 1$ be the constant from Corollary 2.2. We consider two cases.

Case I: In this case we assume

$$\log \mathrm{rk}(f) \leq \frac{1}{2c}\sqrt{\mathrm{D}(f)}.$$

By Corollary 2.2, we have

$$\mathrm{D}(f) \leq \frac{1}{2}\,\mathrm{R}_0^{\mathrm{pub}}(f)\sqrt{\mathrm{D}(f)} + \frac{\mathrm{D}(f)}{4},$$

which simplifies to

$$\sqrt{\mathrm{D}(f)} \leq \mathrm{R}_0^{\mathrm{pub}}(f).$$

Case II: In this case we assume

$$\log \mathrm{rk}(f) > \frac{1}{2c}\sqrt{\mathrm{D}(f)}.$$

Denote $m = 2^{\frac{1}{2c}\sqrt{\mathrm{D}(f)}}$. By the assumption, f contains a full-rank $m \times m$ submatrix A. By Lemma 2.3, we have

$$\mathrm{D}(A) \leq (2\,\mathrm{R}_0^{\mathrm{pub}}(f) + \log\log\left(m^2\right) + 1)^2 \leq (2\,\mathrm{R}_0^{\mathrm{pub}}(f) + 2\log\log m + 1)^2.$$

Combining this with the rank lower-bound (see Equation (1))

$$\mathrm{D}(A) \geq \log \mathrm{rk}(A) \geq \log m,$$

yields

$$\log m \leq (2\,\mathrm{R}_0^{\mathrm{pub}}(f) + 2\log\log m + 1)^2,$$

which simplifies to

$$\sqrt{\log m} - 2\log\log m - 1 \leq 2\,\mathrm{R}_0^{\mathrm{pub}}(f).$$

Substituting $m = 2^{\frac{1}{2c}\sqrt{\mathrm{D}(f)}}$, we obtain

$$\mathrm{R}_0^{\mathrm{pub}}(f) = \Omega(\mathrm{D}(f)^{1/4}).$$

□

3. Concluding remarks and open problems

- We suspect that the bound in Theorem 1.1 can be improved. The example of Fürer [Für87] shows that there are functions f for which
$$\mathrm{R}_0^{\mathrm{pub}}(f) = \Theta(\sqrt{\mathrm{D}(f)}).$$
Is it true that similar to private-coin model, for every function $f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$,
$$\mathrm{R}_0^{\mathrm{pub}}(f) = \Omega(\sqrt{\mathrm{D}(f)})?$$
- We are not aware of any examples separating $\mathrm{N}(f)$, $\mathrm{R}_0^{\mathrm{prv}}(f)$, and $\mathrm{R}_0^{\mathrm{pub}}(f)$. To the best of our knowledge, it could be the case that these parameters are within constant factors of each other. See also [KN97, Problem 3.11].
- In the case where protocols are allowed to make error, the gap between the public-coin and private-coin randomized communication complexities can be arbitrarily large. For example, for the $N \times N$ identity matrix $\mathtt{I}_N$, allowing an error probability of $\frac{1}{3}$, we have
$$\mathrm{R}_{1/3}^{\mathrm{pub}}(\mathtt{I}_N) = O(1), \qquad \mathrm{R}_{1/3}^{\mathrm{prv}}(\mathtt{I}_N) = \Theta(\log N), \qquad \mathrm{D}(\mathtt{I}_N) = \Theta(\log N).$$
- Corollary 2.2 shows that if $\mathrm{rk}(f)$ is small, then $\mathrm{D}(f)$ can be bounded as a function of $\mathrm{R}_0^{\mathrm{pub}}(f)$. Gavinsky and Lovett [GL14] show that this can be generalized to a similar bound for $\mathrm{R}_{1/3}^{\mathrm{pub}}(f)$, albeit with slightly worse parameters:
$$\mathrm{D}(f) = O(\mathrm{R}_{1/3}^{\mathrm{pub}}(f) \log^2 \mathrm{rk}(f)).$$

References

[AUY83] Alfred V. Aho, Jeffrey D. Ullman, and Mihalis Yannakakis. On notions of information transfer in vlsi circuits. In *Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing*, STOC '83, page 133–139, 1983.

[Für87] Martin Fürer. The power of randomness for communication complexity. In *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing*, STOC '87, page 178–181, 1987.

[GL14] Dmitry Gavinsky and Shachar Lovett. En route to the log-rank conjecture: new reductions and equivalent formulations. In *Automata, languages, and programming. Part I*, volume 8572 of *Lecture Notes in Comput. Sci.*, pages 514–524. Springer, Heidelberg, 2014.

[KN97] Eyal Kushilevitz and Noam Nisan. *Communication complexity*. Cambridge University Press, Cambridge, 1997.

[New91] Ilan Newman. Private vs. common random bits in communication complexity. *Inform. Process. Lett.*, 39(2):67–71, 1991.

[NW95] Noam Nisan and Avi Wigderson. On rank vs. communication complexity. *Combinatorica*, 15(4):557–565, 1995.

School of Computer Science, McGill University, Montreal, QC
Email address: benjamin.davis2@mail.mcgill.ca

School of Computer Science, McGill University, Montreal, QC.
Email address: hatami@cs.mcgill.ca

School of Computer Science, McGill University, Montreal, QC.
Email address: william.pires@mail.mcgill.ca

Department of Mathematics and Statistics, McGill University, Montreal, QC.
Email address: ran.tao6@mail.mcgill.ca

School of Computer Science, McGill University, Montreal, QC.
Email address: hamza.usmani@mail.mcgill.ca

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