Electronic Colloquium on Computational Complexity, Report No. 139 (2015)

Lower bound for communication complexity with no public randomness

Eli Ben-Sasson* Gal Maor*

August 26, 2015

Abstract

We give a self contained proof of a logarithmic lower bound on the communication complexity of any non redundant function, given that there is no access to shared randomness. This bound was first stated in Yao's seminal paper [STOC 1979], but no full proof appears in the literature. Our proof uses the method of Babai and Kimmel [Computational Complexity 1997], introduced there in the context of the simultaneous messages model, applying it to the more general and standard communication model of Yao.

1 Introduction

The *communication complexity* of a task is the minimal number of bits that should be communicated to complete the task, when several parties are trying to perform it (cf. [KN97]). The measures $R_\epsilon^{pub}(f), R_\epsilon^{priv}(f)$ represent the randomized communication complexity of a function $f : X \times Y \to Z$, when the parties are using public or private random strings, respectively, and are allowed to err with probability at most ϵ on every input pair.

Observe that if there exist two inputs $x_1, x_2 \in X$ such that for every y, $f(x_1, y) = f(x_2, y)$, then regarding them as the same input would not change the communication complexity of the function. If such a pair (or the symmetric case) exists, we say that the function is *redundant*. Note that every function could be reduced to its non-redundant equivalent, by reducing the input space.

The main theorem we prove here, first stated in [Yao79] without a proof, and later given in [HW07] with a proof sketch (that differs from our method), is the following:

Theorem 1. *[Yao79] For every $\epsilon < \frac{1}{2}$, there exists a universal constant C_ϵ such that for every non redundant function $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$, $R_\epsilon^{priv}(f) \geq C_\epsilon \log(n)$.*

*Department of Computer Science, Technion — Israel Institute of Technology, Haifa, Israel. This research has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement number 240258.

ISSN 1433-8092

This means, that no non-redundant function could be computed with less than logarithmic communication, unless we allow the use of shared randomness. We know, that for several functions (including the equality function), $R_\epsilon^{pub}(f) = O(1)$. Newman's theorem (see [New91]) states that the difference between the two measures can be at most an additive logarithmic term, namely $R_{\epsilon+\delta}^{priv}(f) \leq O(R_\epsilon^{pub}(f) + \log \frac{n}{\delta})$. Theorem 1, combined with Newman's theorem, shows that for every function for which the CC with shared randomness is at most logarithmic, the CC without shared randomness is exactly logarithmic (in the input size).
We use ideas from [BK97], where a square root lower bound on communication complexity was proved for the *simultaneous messages* (SM) model, in which each party is allowed to send only one message to a *referee*, who decides the output. Our generalization provides a weaker bound (logarithmic here, as opposed to square root in the SM model), since every message in our more general model may depend on previous messages sent.

2 Definitions and notation

Let π be a communication protocol with private random strings r_A, r_B. We say that π computes f with error probability ϵ if for every pair (x, y), π outputs the value $f(x, y)$ with probability at least $1-\epsilon$. We denote by $CC(\pi)$ the worst case length of the protocol. The ϵ-error randomized communication complexity of f, $R_\epsilon(f)$, is the minimum of $CC(\pi)$ taken over all private-coin protocols computing f with error probability ϵ.
Assuming $CC(\pi) = c$, we can view the protocol as a binary tree of depth c: At each vertex of the tree, the party owning the vertex decides the child vertex to go to according to its input and random string. Thus, we can think of Alice as sampling a random string, and then according to the string and the input fixing a *strategy* $\phi_A = \phi_A(x, r_A)$, that determines how to proceed from each owned vertex (conditioned on reaching it). Bob picks a strategy $\phi_B = \phi_B(y, r_B)$ likewise. Two such strategies ϕ_A, ϕ_B define a unique leaf that the protocol reaches. We abuse notation and denote by $\pi(\phi_A, \phi_B)$ the output of the relevant leaf. Denote by Φ_A, Φ_B the sets of possible strategies for Alice and Bob, respectively. Denote by $G(\phi_B) \subseteq \Phi_A$ the set of Alice's strategies for which $\pi(\phi_A, \phi_B) = 1$.

3 Proving the lower bound

We will require a few steps before we prove the main theorem.

Fact 1. $|\Phi_A|, |\Phi_B| \leq 2^{2^c}$

Proof. A strategy of Alice could be thought of as a function $s_A : V_A \to \{0, 1\}$, where V_A is the set of vertices owned by Alice. Since $|V_A| < 2^c$, the number of strategies is at most 2^{2^c}. □

Let ν_x^A be the probability measure over Alice's strategies Φ_A given that Alice's input is x, and define ν_y^B similarly.

Lemma 1. *For every $\delta > 0$, there exists an integer $t = 2^{O(c)}$ such that for every $x \in X$, there exists a sequence $T_x = (\phi_1, \phi_2, ..., \phi_t)$ such that for every ϕ_B: $|\frac{1}{t}\sum_{i=1}^{t} \pi(\phi_i, \phi_B) - \nu_x^A(G(\phi_B))| < \delta$.*

Proof. We will use a probabilistic argument. Pick $T = (\phi_1, \phi_2, ..., \phi_t)$, where each ϕ_i is picked uniformly at random according to ν_x^A, with repetitions. It is natural from the definitions that for a specific ϕ_B, the probability that $\pi(\phi_i, \phi_B) = 1$ is exactly $\nu_x^A(G(\phi_B))$. Hence, the random variable $a_i = \pi(\phi_i, \phi_B) - \nu_x^A(G(\phi_B))$ has expected value 0. Using the Chernoff bound we get that

$$\Pr\left[\left|\frac{1}{t}\sum_{i=1}^{t} \pi(\phi_i, \phi_B) - \nu_x^A(G(\phi_B))\right| > \delta\right] < 2e^{\frac{-\delta^2 t}{2}}$$

Choosing $t = \frac{2}{\delta^2}\ln 2|\Phi_B|$, this probability is smaller than $\frac{1}{|\Phi_B|}$, and using a union bound argument, the probability that the set T fails on *any* of Bob's strategies is less than 1. Thus there exists a set T_x satisfying the equation. From fact 1, $t = 2^{O(c)}$. □

We define the protocol π' to be the following: Bob will behave just as he does in π, choosing a strategy according to his input and random string. Alice, given her input x, picks $i \in [t]$ uniformly at random and behaves according to strategy ϕ_i from T_x. Denote the error probability of π' on input (x, y) by $\epsilon'(x, y)$.

Lemma 2. *For every pair (x, y), $\epsilon'(x, y) < \epsilon + \delta$.*

Proof. Note that by definition, the probability for π' to output 1 is

$$\sum_{i=1}^{t} \frac{1}{t} \sum_{\phi_B \in \Phi_B} \nu_y^B(\phi_B)\pi(\phi_i, \phi_B)$$

Assume that $f(x, y) = 1$. Knowing that π errs with probability at most ϵ we get:

$$1 - \epsilon < \Pr[\pi(x, y) = 1] =$$

$$\sum_{\phi_B \in \Phi_B} \nu_y^B(\phi_B)\nu_x^A(G(\phi_B)) < \sum_{\phi_B \in \Phi_B} \nu_y^B(\phi_B)[\frac{1}{t}\sum_{i=1}^{t}\pi(\phi_i, \phi_B) + \delta] =$$

$$\delta + \sum_{i=1}^{t} \frac{1}{t} \sum_{\phi_B \in \Phi_B} \nu_y^B(\phi_B)\pi(\phi_i, \phi_B) = \delta + \Pr[\pi'(x, y) = 1]$$

The case for $f(x, y) = 0$ is similar. □

Corollary 1. *Let ϵ, δ be such that $\epsilon + \delta < \frac{1}{2}$. If f is non redundant, then for every $x_1 \neq x_2 \in X$, $T_{x_1} \neq T_{x_2}$.*

Proof. Since f is non redundant, we get without loss of generality that for $x_1 \neq x_2$ exists some y such that $f(x_1, y) = 1, f(x_2, y) = 0$. Assume $T_{x_1} = T_{x_2}$, then according to the construction of π' the protocol would be the exact same one for both input pairs. According to Lemma 2, since the probability to be correct is at least $1 - \epsilon - \delta$, both the probabilities to output 1 and 0 would be strictly larger than $\frac{1}{2}$, in contradiction. □

Proof of Theorem 1. Assume π has communication complexity c and error probability $\epsilon < \frac{1}{2}$ and pick δ such that $\epsilon + \delta < \frac{1}{2}$. Given Corollary 1 we know that the number of possible x's is bounded by the number of possible choices of T_x. Hence:

$$2^n = |X| \leq |\Phi_A|^t \leq (2^{2^c})^{2^{O(c)}} = 2^{2^{O(c)}}$$

which implies $c = \Omega(\log n)$. □

References

[BK97] László Babai and Peter G Kimmel. Randomized simultaneous messages: Solution of a problem of yao in communication complexity. In *Computational Complexity, 1997. Proceedings., Twelfth Annual IEEE Conference on (Formerly: Structure in Complexity Theory Conference)*, pages 239–246. IEEE, 1997.

[HW07] Johan Håstad and Avi Wigderson. The randomized communication complexity of set disjointness. *Theory of Computing*, 3(1):211–219, 2007.

[KN97] Eyal Kushilevitz and Noam Nisan. *Communication Complexity.* Cambridge University Press, New York, NY, USA, 1997.

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

[Yao79] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (preliminary report). In *Proceedings of the eleventh annual ACM symposium on Theory of computing*, pages 209–213. ACM, 1979.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de