Randomized Simultaneous Messages: Solution of a Problem of Yao in Communication Complexity

László Babai
Peter G. Kimmel
Department of Computer Science
The University of Chicago
1100 East 58th Street
Chicago, Illinois 60637

Abstract

We solve a 17 year old problem of Yao (FOCS 79).

In the two-player communication model introduced by Yao in 1979, Alice and Bob wish to collaboratively evaluate a function $f(x, y)$ in which Alice knows only input x and Bob knows only input y. Both players have unlimited computational power. The objective is to minimize the amount of communication.

Yao (FOCS 79) also introduced an oblivious *version of this communication game which we call the* simultaneous messages *(SM) model. The difference is that in the SM model, Alice and Bob don't communicate with each other. Instead, they simultaneously send messages to a referee, who sees none of the input. The referee then announces the function value.*

The deterministic two-player SM complexity of any function is straightforward to determine. Yao suggested the randomized version of this model, where each player has access to private coin flips.

Our main result is that the order of magnitude of the randomized SM complexity of any function f is at least the square root of the deterministic SM complexity of f. We found this result in February 1996, independently but subsequently to I. Newman and M. Szegedy (STOC 96) who obtain this lower bound for the special case of the "equality" function. Our proof is entirely different from and much simpler than the Newman-Szegedy solution, and it is stronger in that it gives a lower bound not only for the "equality" function but for all functions. A proof in a similar spirit was also found by J. Bourgain and A. Wigderson simultaneously to us (unpublished).

The quadratic reduction actually does occur for the "equality" function (A. Ambainis [2], M. Naor [9], and I. Newman [10] (cf. [7]).) We give a new proof of this fact. This result, combined with our main result, settles Yao's question (FOCS 79), asking the exact randomized SM complexity of the equality function. The lower bound proof uses the probabilistic method; the upper bound uses linear algebra.

We also give a constructive proof that $O(\log n)$ public coins reduce the complexity of "equality" to constant.

1 Introduction

In 1979, Yao [14] introduced the following communication game: Let $f : X \times Y \to \{0, 1\}$ be a boolean function. There are two players, Alice and Bob, who wish to collaboratively compute the value of f on input $(x, y) \in X \times Y$. However, Alice sees only the input x, and Bob sees only the input y. Both Alice and Bob have unlimited computational power. They communicate with each other by writing on a blackboard. The last bit written on the board must be the function value. The *cost* of a communication protocol is the number of bits written on the board for the worst case input. The *communication complexity* of f, denoted $C(f)$, is the minimum cost of a protocol computing f.

Yao [14] also proposed an *oblivious* version of this model which we call the *simultaneous messages* (SM) model: Let $f : X \times Y \to \{0, 1\}$ be a boolean function. Alice is given an input $x \in X$, and Bob is given an input $y \in Y$. Alice, Bob, and a referee wish to collaboratively evaluate $f(x, y)$. Alice sees only input x, Bob sees only input y, and the referee sees none of the inputs. Both players simultaneously pass a message of fixed length to the referee, after which the referee announces the function value. Each player (including the referee) is a function of the arguments it "knows."

Definition 1.1 A *simultaneous messages (SM) protocol P for f* consists of two players along with a referee that correctly computes f on all inputs. The *cost* of an SM protocol

for f is the length of the longer message sent to the referee. The *SM complexity* of f, denoted $C_0(f)$, is the minimum cost of an SM protocol computing f.

This quantity is straightforward to determine. Let $M(f)$ be the *communication matrix* corresponding to f, that is, $M(f)$ is the $|X| \times |Y|$ matrix with entry $f(x,y)$ in the corresponding cell. Let $nrow(f)$ and $ncol(f)$ denote the number of *distinct* rows and columns of $M(f)$, respectively. Then it is easy to show that $C_0(f) = \max\{\lceil \log nrow(f) \rceil, \lceil \log ncol(f) \rceil\}$.

The analagous quantity for several players is very hard to estimate. The SM model with several players is considered in [12], [13], [5], and [3]. (Most of the authors use the term "oblivious communication complexity.")

Yao actually introduces the SM model for randomizing players who use private coins, and calls this a "situation that deserves special attention." He specifically asks the randomized SM complexity of the "equality" function [14, Concl. Rem. D, p. 213]. In this paper we resolve this 17-year old question. In doing so, we use simple but appealing techniques from probabilistic combinatorics and linear algebra. We show in Section 4 that this question is closely related to an *extremal problem in graph theory* ("maximum number of densely connected independent sets"), a fact also established by [11].

Definition 1.2 A *two-sided ϵ-error randomized SM protocol* P for f is an SM protocol in which Alice, Bob, and the referee are allowed to randomize, and for all $(x,y) \in X \times Y$, the referee outputs the correct value of $f(x,y)$ with probability at least $1 - \epsilon$. We define a *one-sided ϵ-error* randomized SM protocol in the same way with the exception that for all $(x,y) \in X \times Y$ such that $f(x,y) = 1$, the referee must always output 1. In the *private-coin* model, each player, including the referee, flips private coins. In the *public-coin* model, Alice and Bob are given the same random bits, but they do not see the referee's random bits, nor does the referee see Alice's and Bob's random bits.

Remark 1.3 In view of amplification by repetition, all positive constants $\epsilon < 1/2$ are equivalent for two-sided error randomized protocols, and all positive constants $\epsilon < 1$ are equivalent for one-sided error randomized protocols.

Recently, the randomized SM and one-way communication models for two players have been studied in [7] in connection with the VC dimension and the problem of computing the inner product of two real vectors.

We shall briefly discuss the public-coin model at the end of this paper, but our main concern is the private-coin model. Our main result is the following:

Theorem 1.4 *Let $f : X \times Y \to \{0,1\}$ be any boolean function. Any private-coin two-sided error randomized SM protocol for f has cost $\Omega(\sqrt{C_0(f)})$.*

We also show using simple facts from linear algebra that the quadratic reduction of cost can be achieved for the "equality" function $EQ_n : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$, where $EQ_n(x,y) = 1$ iff $x = y$.

Theorem 1.5 *There exists a private-coin one-sided error randomized SM protocol for EQ_n of cost $\Theta(\sqrt{n}) = \Theta(\sqrt{C_0(EQ_n)})$.*

Corollary 1.6 *The private-coin randomized SM complexity of EQ_n is $\Theta(\sqrt{n})$.* ■

This answers Yao's question [14, Concl. Rem. D, p. 213].

Acknowledgment: We wish to thank Avi Wigderson for bringing to our attention the history of this problem. In particular, our main result, Theorem 1.4, was recently proved by Ilan Newman and Mario Szegedy [11] for the specific case of the "equality" function. After Szegedy's presentation of the [11] results at the I.A.S., J. Bourgain and A. Wigderson found a greatly simplified proof of the general statement of Theorem 1.4 (all functions) [6]. Our work was done independently of [11] and [6] and roughly simultaneously to the latter. Our work was completed by February, 1996. Although the [6] proof is different from ours, both proofs are quite simple, use the Chernoff bounds, and are closely related in spirit.

Avi Wigderson has also communicated to us that the $\sqrt{n}$ upper bound of Theorem 1.5 was also found previously by A. Ambainis [2], M. Naor [9], and I. Newman [10] (cf. [7]) using protocols different from ours.

2 The Lower Bound

Let $f : X \times Y \to \{0,1\}$ be any boolean function. Let $M(f)$ be the communication matrix corresponding to f. Without loss of generality, we assume that $M(f)$ has no identical rows or columns. Let P be a private-coin two-sided ϵ-error randomized SM protocol for f. Without loss of generality, we may assume $\epsilon = 0.01$. Let $a(n)$ and $b(n)$ be the number of bits that Alice and Bob send respectively to the referee.

Theorem 2.1 $a(n) \cdot b(n) \geq \Omega(C_0(f))$.

Note that Theorem 1.4 is an immediate consequence. ■

Proof of Theorem 2.1:

Let Ω and Φ be the set of messages of Alice and Bob respectively. Let $a(n) = \log|\Omega|$ and $b(n) = \log|\Phi|$. For $x \in X$, let μ_x be Alice's probability measure on Ω given input x. For $y \in Y$, let ν_y be the Bob's probability measure on Φ given input y.

The idea behind the proof is that for each μ_x and ν_y we pick a logarithmic size sample space and prove, using

a Chernoff bound, that these small sets uniquely correspond to x and y respectively.

For $y \in Y$ and subset $S \subseteq \Phi$, let $\nu_y(S) = \sum_{\varphi \in S} \nu_y(\varphi)$. For $x \in X$ and multiset $T = \{\omega_1, \ldots, \omega_t\}$, $\omega_i \in \Omega$, let $\mu_x(T) = \sum_{i=1}^{t} \mu_x(\omega_i)$. For $\omega \in \Omega$, $\varphi \in \Phi$, let $\rho(\omega, \varphi)$ be the probability that the referee outputs 1 on message pair (ω, φ).

Remark 2.2 One could also proceed by assuming that the referee is deterministic: From a protocol P with a randomizing referee, create a protocol P' in which the referee outputs 1 for every (ω, φ) such that $\rho(\omega, \varphi) \geq 0.5$, and 0 otherwise. It is not hard to see that this increases the error by at most a factor of 2. This is how the lower bound proof for the equality function proceeds in [11].

For $x \in X$ and $\varphi \in \Phi$, we call the quantity $F(x, \varphi) := \sum_{\omega \in \Omega} \mu_x(\omega)\rho(\omega, \varphi)$ the *strength* of φ for x. For $0 \leq \alpha \leq 1$, we say $\varphi \in \Phi$ is *α-strong* for $x \in X$ if $F(x, \varphi) \geq \alpha$; otherwise, φ is α-weak for x. For $x \in X$, let $S_x = \{\varphi \in \Phi : \varphi \text{ is } 0.9\text{-strong for } x\}$ and let $W_x = \{\varphi \in \Phi : \varphi \text{ is } 0.1\text{-weak for } x\}$.

The two-sided error condition on the protocol can be restated as follows. For every $x \in X,\ y \in Y$, the following two conditions hold:

$$f(x, y) = 0 \ \Rightarrow \sum_{\omega, \varphi} \mu_x(\omega) \cdot \nu_y(\varphi) \cdot \rho(\omega, \varphi) \leq 0.01 \quad (1)$$

and

$$f(x, y) = 1 \ \Rightarrow \sum_{\omega, \varphi} \mu_x(\omega) \cdot \nu_y(\varphi) \cdot \rho(\omega, \varphi) \geq 0.99. \quad (2)$$

Observation 2.3 *For every $x \in X,\ y \in Y$ the following two conditions hold:*

$$f(x, y) = 0 \ \text{ implies } \ \nu_y(W_x) \geq 0.9, \text{ and} \quad (3)$$

$$f(x, y) = 1 \ \text{ implies } \ \nu_y(S_x) \geq 0.9. \quad (4)$$

Proof: Consider the case $f(x, y) = 0$. Then

$$\begin{aligned} 1/100 &\geq \sum_{\omega, \varphi} \mu_x(\omega)\nu_y(\varphi)\rho(\omega, \varphi) \\ &\geq \sum_{\varphi \notin W_x} \nu_y(\varphi) \cdot \sum_{\omega} \mu_x(\omega)\rho(\omega, \varphi) \\ &\geq 0.1 \cdot \sum_{\varphi \notin W_x} \nu_y(\varphi) \ = \ 0.1 \cdot (1 - \nu_y(W_x)). \end{aligned}$$

Hence $\nu_y(W_x) \geq 0.9$. The case $f(x, y) = 1$ follows by symmetry.

■ Observation 2.3.

For $x \in X, \varphi \in \Phi$, and multiset $T_x = \{\omega_1, \ldots, \omega_t\}$, $\omega_i \in \Omega$, define independent random variables

$$\xi_i(\varphi) = \begin{cases} 1 & \text{if the referee accepts } (\omega_i, \varphi), \\ 0 & \text{otherwise.} \end{cases}$$

Note that $\xi_i(\varphi)$ depends implicity on ρ. In order for these to be independent, we must ask the referee to perform independent tests on each (ω_i, φ).

Lemma 2.4 *For all $x \in X$, there exists a multiset $T_x = \{\omega_1, \ldots, \omega_t\}$, $\omega_i \in \Omega$, with $t = \lceil 200 \ln(4|\Phi|) \rceil$, such that for all $\varphi \in \Phi$,*

$$\Pr_{\rho}(|\sum_{i=1}^{t} \xi_i(\varphi) - t \cdot F(x, \varphi)| > 0.1 \cdot t) < 1/2, \quad (5)$$

where the probability is taken over the independent tests of the referee.

Proof: Fix $x \in X$. Choose a T_x at random by picking t elements $\omega_1, \ldots, \omega_t \in \Omega$ independently at random according to μ_x. For all $\varphi \in \Phi$, we have

$$E_{T_x, \rho}(\xi_i(\varphi)) = \sum_{\omega \in \Omega} \mu_x(\omega)\rho(\omega, \varphi) = F(x, \varphi),$$

where $E_{T_x, \rho}$ denotes the expectation over the choice of the ω_i and the referee's coins.

For $\varphi \in \Phi$, let

$$\Gamma(\varphi, T_x) \text{ be the event that } |\sum_{i=1}^{t} \xi_i(\varphi) - t \cdot F(x, \varphi)| > 0.1 \cdot t.$$

Note that $\Gamma(\varphi, T_x)$ depends implicitly on ρ.

Fix a $\varphi \in \Phi$. For $1 \leq i \leq t$, define random variables

$$\eta_i(\varphi) = \xi_i(\varphi) - F(x, \varphi).$$

Then the η_i are independent, $|\eta_i(\varphi)| \leq 1$, and $E_{T_x, \rho}(\eta_i(\varphi)) = 0$. Therefore, we can use a Chernoff bound (cf. e. g. [1, Thm. A.16, p. 240]) to obtain

$$\begin{aligned} \Pr_{T_x, \rho}(\Gamma(\varphi, T_x)) \ &= \ \Pr_{T_x, \rho}(|\sum_{i=1}^{t} \eta_i(\varphi)| > 0.1 \cdot t) \\ &< \ 2e^{-(0.1 \cdot t)^2/2t} = 2e^{-t/200} \\ &\leq \ 1/2|\Phi|. \end{aligned}$$

Therefore,

$$\Pr_{T_x, \rho}((\exists \varphi \in \Phi)(\Gamma(\varphi, T_x))) \ < \ 1/2.$$

Thus there exists a choice of T_x such that

$$\Pr_\rho((\exists \varphi \in \Phi)(\Gamma(\varphi, T_x))) \ < \ 1/2.$$

This trivially implies (for the same T_x): for all $\varphi \in \Phi$, we have $\Pr_\rho(\Gamma(\varphi, T_x)) \ < \ 1/2$.

■ Lemma 2.4

Claim 2.5 *For all $x \neq x' \in X$, we have $T_x \neq T_{x'}$.*

Proof: Assume there exist $x \neq x' \in X$ such that $T_x = T_{x'}$. Since the communication matrix $M(f)$ has no identical rows, there exists a $y \in Y$ such that $f(x,y) \neq f(x',y)$. Without loss of generality, let us assume that $f(x,y) = 0$ and $f(x',y) = 1$.

By Observation 2.3, $\nu_y(S_x) \geq 0.9$ and $\nu_y(W_{x'}) \geq 0.9$. Therefore, there exists a $\varphi_0 \in S_x \cap W_{x'}$.

Let A be the event that $|\sum_{i=1}^t \xi_i(\varphi_0) - t \cdot F(x, \varphi_0)| > 0.1 \cdot t)$.

Let B be the event that $|\sum_{i=1}^t \xi_i(\varphi_0) - t \cdot F(x', \varphi_0)| > 0.1 \cdot t)$.

By Lemma 2.4, we have $\Pr(A) \ < \ 1/2$ and $\Pr(B) \ < \ 1/2$. Therefore,

$$\Pr(A \vee B) < 1. \tag{6}$$

By definition of S_x and $W_{x'}$, we know that $F(x, \varphi_0) \geq 0.9$ and $F(x', \varphi_0) < 0.1$. Therefore, any value of $\sum_{i=1}^t \xi_i(\varphi_0)$ forces at least one of the events A or B to happen, so $\Pr(A \vee B) = 1$. This contradicts (6) and concludes the proof.

■ Claim 2.5

Claim 2.5 implies that $|X|$ is bounded from above by the number of possible T_x. This means

$$|X| \ \leq \ |\Omega|^t \leq |\Omega|^{\lceil 200 \ln(4|\Phi|) \rceil} \leq 2^{(\log|\Omega|)(1+200\ln(4|\Phi|))}$$

$$\leq \ 2^{400 \log|\Omega| \log|\Phi|}.$$

Therefore, $400 \log|\Omega| \log|\Phi| \ \geq \ \log|X|$. By symmetrical arguments, $400 \log|\Omega| \log|\Phi| \geq \log|Y|$. Thus,

$$400 \log|\Omega| \log|\Phi| \geq \max\{\log|X|, \log|Y|\} = C_0(f).$$

Since $a(n) = \log|\Omega|$ and $b(n) = \log|\Phi|$, we have $a(n) \cdot b(n) \geq C_0(f)/400$. ■ Theorem 2.1

3 The Upper Bound for Equality

In this section, we prove Theorem 1.5, which shows that the lower bound of Section 2 is tight for the "equality" function $EQ_n : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$, $EQ_n(x,y) = 1$ iff $x = y$.

Proof of Theorem 1.5: Let $\epsilon = 3/4$.

Notation: Let t be the smallest even integer such that $n \ \leq \ t^2/4$. Let Ω be the vector space $GF(2)^t$ with the standard inner product $u \cdot v \ = \ \sum_{i=1}^t u_i v_i$. If $u \cdot v \ = \ 0$ $(u, v \in \Omega)$, then we say that u and v are *perpendicular* and write $u \perp v$. For $S \subseteq \Omega$, let $S^\perp$ denote the subspace $S^\perp = \{v \in \Omega : \ \forall w \in S, \ v \perp w\}$. For $u \in \Omega$, we let $u^\perp$ denote $\{u\}^\perp$.

Fact 3.1 *There are more than $2^{t^2/4}$ subspaces of Ω of dimension $t/2$.* ■

The protocol works as follows. By Fact 3.1, with each $z \in \{0,1\}^n$ we associate a distinct subspace U_z of dimension $t/2$. On input $x \in \{0,1\}^n$, Alice picks a vector u uniformly at random from U_x and sends it to the referee. On input $y \in \{0,1\}^n$, Bob picks a vector v uniformly at random from $U_y^\perp$ and sends it to the referee. The referee outputs 1 if and only if $u \perp v$.

Since $u, v \in \Omega = GF(2)^t$, Alice and Bob each send $t = 2\sqrt{n}(1 + o(1))$ bits to the referee.

If $x = y$, then $U_x = U_y$, and hence for any $u \in U_x$, $v \in U_y^\perp$, we have $u \ \perp \ v$, and thus the referee outputs 1 with probability 1. Suppose now that $x \neq y$.

Observation 3.2 *If U is a subspace of Ω and $u \notin U$, then $u^\perp \not\supseteq U^\perp$.*

Proof: It follows from the well known dimension formula (cf. e.g. [4, Prop. 3.20, p. 53]) that $U = U^{\perp\perp}$. Thus we have $u^\perp \supseteq U^\perp$ iff $u^{\perp\perp} = \text{span}(u) \subseteq U = U^{\perp\perp}$, i.e., $u \in U$.

■ Observation 3.2

Since $x \neq y$, we have $U_x \neq U_y$, and thus $|U_x \cap U_y| \leq |U_x|/2$. Therefore, $\Pr(u \in U_x \setminus U_y) \geq 1/2$.

If Alice picks a $u \in U_x \setminus U_y$, then by Observation 3.2, we have $u^\perp \not\supseteq U_y^\perp$, and thus $|U_y^\perp \cap u^\perp| \leq |U_y^\perp|/2$. Therefore, $\Pr(v \notin u^\perp \mid u \notin U_y) \geq 1/2$. Thus we have the $\Pr(\text{referee outputs } 0) = \ \Pr(u \not\perp v) \geq 1/4$. ■ Theorem 1.5.

Remark 3.3 It may seem natural to require that an EQ_n protocol be *symmetric*: Alice and Bob follow the same instructions. This can be accomplished at a cost of a factor of 2 from any (asymmetric) protocol by having Alice and Bob each send what both the old Alice and the old Bob would have sent on their input. The new referee then uses only half of the information from each player. However, we can do slightly better than this factor of 2 in the above protocol for EQ_n if we are a little more careful with the subspaces we associate with the inputs.

Recall that a subspace $U \leq \Omega$ is called *totally isotropic* if $U \subseteq U^\perp$. It is known that there are more than $2^{t(t-2)/8}$ totally isotropic subspaces of Ω of dimension $t/2$. Now we use the above protocol with $n \ \leq \ t(t-2)/8$, and associate a totally isotropic subspace $U_x \leq \Omega$ of dimension $t/2$ with each input x. The protocol is now symmetric because $U_x = U_x^\perp$, and the cost is $t = 2\sqrt{2n}(1 + o(1))$. This is only

a factor of $\sqrt{2}$ more expensive than the original protocol, instead of a factor of 2.

4 A Related Problem in Graph Theory: Densely Connected Independent Sets

The results for the "equality" function give the tight order of magnitude of the logarithm of a graph theoretic extremum, stated in Question 4.3.

Definition 4.1 For a graph $G(V, E)$ and subsets $A, B \subseteq V$, the *density* of G between A and B is

$$|\{(v,w) : v \in A, w \in B, (v,w) \in E\}| \;/\; |A| \cdot |B|.$$

Definition 4.2 Let $0 < \delta < 1$. A family S of subsets of the vertex set of a graph G is called *δ-densely connected* if for all $A, B \in S$, $A \neq B$, the density of G between A and B is at least δ.

Question 4.3 *What is the maximum size of a δ-densely connected family of independent sets of a graph on n vertices?*

It turns out that this question is equivalent to the one-sided error randomized SM complexity of "equality" under a restricted type of protocols. In fact, it was in this context that we arrived at our solutions. Apparently, Newman and Szegedy [11] followed a similar path and obtained the same result prior to our independent work.

Theorem 4.4 *For fixed δ, $0 < \delta < 1$, the maximum size of a δ-densely connected family of independent sets of a graph on n vertices is $n^{\Theta(\log n)}$.*

Definition 4.5 A randomized SM protocol P for a function $f : X \times X \to \{0,1\}$ is called *symmetric* if under P, Alice and Bob follow the same instructions and the referee is symmetric.

Definition 4.6 Let P be a randomized SM protocol for a function $f : X \times Y \to \{0,1\}$, with Ω and Φ as the message sets for Alice and Bob respectively. We call P *uniform* if for each $x \in X$ there is a subset $U_x \subseteq \Omega$ such that on input x, Alice picks a message from U_x uniformly at random, and similarly for each $y \in Y$ there is a subset $W_y \subseteq \Phi$ such that on input y, Bob picks a message from W_y uniformly at random.

Definition 4.7 An *elementary* randomized SM protocol is a uniform, symmetric, one-sided error randomized SM protocol in which the *referee is deterministic*.

Note that the protocol of Remark 3.3 is a $(3/4)$-error elementary randomized SM protocol for EQ_n.

Theorem 4.8 *The following two statements are equivalent:*

1. *There exists a graph G on n vertices with a family of N δ-densely connected independent sets.*
2. *There exists an elementary $(1-\delta)$-error randomized SM protocol for $EQ_{\log N}$ of cost $\log n$.*

We first show how Theorem 4.8, together with our main results, imples Theorem 4.4.

Proof of Theorem 4.4: By Theorem 1.4 and Remark 3.3, there exists an elementary $(3/4)$-error randomized SM protocol for $EQ_{\log N}$ of cost $\log n$ if and only if $\log n = \Omega(\sqrt{\log N})$, or equivalently, $N = n^{O(\log n)}$, with suitable implied constants. By Remark 1.3, the same holds for error $1-\delta$ for any δ, $0 < \delta < 1$. Combined with Theorem 4.8, this proves Theorem 4.4. ■

Proof of Theorem 4.8: Let $G(V,E)$ be a graph on n vertices with a family S of N δ-densely connected independent sets. Let $X = \{0,1\}^{\log N}$. from$G$ and S, we construct a $(1-\delta)$-error randomized SM protocol P for $EQ : X \times X \to \{0,1\}$.

With each input $x \in X$, we associate a distinct independent set $C_x \in S$. The set of possible messages of Alice and Bob will be V. On input $x \in X$, Alice and Bob pick a message v uniformly at random from $C_x \subseteq V$. The referee outputs 0 if and only if $(v,w) \in E$, where v is Alice's message, and w is Bob's.

Suppose Alice and Bob receive inputs x and y respectively. If $x = y$, then Alice and Bob send vertices from independent set C_x, so the referee outputs 1. If $x \neq y$, then the probability that the referee outputs 0 is exactly the density of G between C_x and C_y, which by hypothesis is δ. It is clear that this protocol is elementary.

Conversely, let P be an elementary $(1-\delta)$-error randomized SM protocol for $EQ : X \times X \to \{0,1\}$, where $X = \{0,1\}^{\log N}$. Let V be the message set of Alice and Bob. Since P is elementary, for every input $x \in X$, there is a subset $U_x \subseteq V$ such that on input x, Alice and Bob pick a message from U_x uniformly at random. from$P$, we define a graph $G(V,E)$. The vertices of G are the messages of Alice and Bob under P. The edges of G and the family S of independent sets are given by

$$E = \{(u,v) \in V \times V : \text{ the referee outputs 0 on} (u,v)\},$$

$$S = \{U_x : x \in X\}.$$

Since P has one-sided error, it follows that the U_x are indeed independent sets of G.

Let $x \neq y \in X$. It is clear that the density of G between U_x and U_y is exactly the probability that the referee outputs 0 when Alice and Bob receive inputs x and y respectively. By the condition of the protocol, this is at least δ.

■ Theorem 4.8.

5 The Public Coin Model

Recall that in the *public-coin* model, Alice and Bob share random bits but do not see the referee's random bits, nor does the referee see Alice's and Bob's random bits. We examine two different public-coin models:

Definition 5.1 In the *public-fee* model, the cost of a randomized SM protocol is the length of the longer message sent to the referee *plus* the number of common random bits used by Alice and Bob. In the *public-no-fee* model, the cost of a randomized SM protocol is simply the length of the longer message sent to the referee.

Yao states the following theorem [14, Thm. 5, p. 212] but omits the proof "because of its complexity:"

Theorem 5.2 ***(Yao)*** *For any f, the two-party (private-coin) randomized communication complexity of f is at least $\Omega(\log\log(\text{nrow}(M(f))) + \log\log(\text{ncol}(M(f))))$, where* nrow *and* ncol *are the number of distinct rows and columns of the communication matrix $M(f)$.*

A consequence of this is the following theorem, for which we give a simple proof.

Theorem 5.3 *Let $f : |X| \times |Y| \to \{0,1\}$ be any boolean function. Any public-fee two-sided error randomized SM protocol for f has cost $\geq \Omega(\log(C_0(f)))$.*

Proof: Let P be a public-fee randomized SM protocol of cost m for $f : |X| \times |Y| \to \{0,1\}$. Let Ω and Φ be the set of messages of Alice and Bob respectively. Let $a(n) = \log|\Omega|$ and $b(n) = \log|\Phi|$. Let $r(n)$ be the number of common random bits viewed by Alice and Bob. For $\omega \in \Omega, \varphi \in \Phi$, let $\rho(\omega, \varphi)$ be the probability that the referee outputs 1 on message pair (ω, φ).

From P, we construct a deterministic SM protocol P' for f with players Alice$'$, Bob$'$, and referee Ref$'$ as follows. On input $x \in X$, Alice$'$ sends $(\omega_1, \omega_2, \ldots, \omega_{2^{r(n)}})$, where ω_i is the message Alice would send under P on input x upon seeing the ith possible random string. Similarly, on input $y \in Y$, Bob$'$ sends $(\varphi_1, \varphi_2, \ldots, \varphi_{2^{r(n)}})$, where φ_i is the message Bob would send under P on input y upon seeing the ith possible random string. Ref$'$ computes

$$\eta(x,y) := \sum_{i=1}^{2^{r(n)}} \rho(\omega_i, \varphi_i)/2^{r(n)}$$

$$= \Pr(\text{referee of } P \text{ outputs 1 on input } (x,y)),$$

and outputs 1 if $\eta(x,y) \geq 1/2$, and 0 otherwise. As P is a two-sided ϵ-error randomized SM protocol for f and $\epsilon < 1/2$, we have that Ref$'$ always outputs the correct answer under P'.

The number of bits sent by Alice$'$ and Bob$'$ respectively is $a(n) \cdot 2^{r(n)}$ and $b(n) \cdot 2^{r(n)}$. Since the public-fee cost of P is $a(n)+b(n)+r(n) = m$, each of $a(n), b(n)$, and $r(n)$ is at most m. Therefore, Alice$'$ and Bob$'$ each send at most $m2^m$ bits. Thus $m2^m \geq C_0(f)$.

■ Theorem 5.3.

The corollary of the next theorem shows that this lower bound is tight for the "equality" function.

Theorem 5.4 *There exists an explicit public-coin one-sided error randomized SM protocol for EQ_n using $O(1)$ bits of communication, $O(\log n)$ public random bits, and no private random bits.*

By "explicit," we mean that Alice and Bob can compute the messages they send to the referee in poly(n) time, where $n = |x| = |y|$. (In fact, the time is nearly linear: the dominant part of the computation is the division of an n-digit integer by a $\log n$-digit integer.)

Remark 5.5 With $O(n)$ public random bits, there is a simple randomized SM protocol for EQ_n in which Alice and Bob send $O(1)$ bits. Let z be chosen uniformly at random from $\{0,1\}^n$. On input x, Alice sends $x \cdot z$ (the inner product modulo 2 of x and z). Similarly, on input y, Bob sends $y \cdot z$. The referee outputs 1 if $x \cdot z = y \cdot z$.

If $x = y$, it is clear the referee outputs 1. If $x \neq y$, then $\Pr_z(x \cdot z = y \cdot z) = 1/2$. This one-sided error randomized SM protocol, along with Theorem 5.3, separates the power of the public-fee and public-no-fee models.

Remark 5.6 Using the above protocol and a standard derandomization argument, it is not hard to show the existence of protocol that uses $O(\log n)$ public coins and $O(1)$ bits of communication. This argument was pointed out to us by Jiří Sgall, who learned it from Noam Nisan.

First, let Alice and Bob use protocol P, in which a random string z of length $3n$ is chosen and then Alice and Bob repeat the protocol of Remark 5.5 three times, so that the probability that P errs is $\leq 1/8$. The number of random strings they chose from is thus 2^{3n}, but we can show that it suffices to draw the random strings from a subset of size cn for some constant $c > 0$, and therefore we only need $O(\log n)$ random bits.

Fix input $(x,y) \in \{0,1\}^{2n}$. Let $c = 144$. Choose a multiset $T = \{z_1, \ldots, z_{cn}\} \subset (\{0,1\}^{3n})^{cn}$ at random by picking cn strings $z_1, \ldots z_{cn}$ independently at random from $\{0,1\}^{3n}$. For $1 \leq i \leq cn$, let

$$\xi_i = \begin{cases} 1 & \text{if } P \text{ errs on } z_i, \\ 0 & \text{otherwise} \end{cases}$$

The ξ_i are independent, with $E(\xi_i) = 1/8$. Let protocol P' be the same as P except that the random string for P' of

length $3n$ is chosen uniformly at random from T. Then the probability that P' errs on input (x,y) is $\sum_{i=1}^{cn} \xi_i/cn$.

For $i \leq i \leq cn$, let

$$\eta_i = \xi_i - 1/8.$$

Then the η_i are independent, and for all i we have $E(\eta_i) = 0$, and $|\eta_i| \leq 1$. Therefore, we can use a Chernoff bound (cf. e.g. [1, Thm. A.16, p. 240]) to obtain

$$\Pr_T(|\sum_{i=1}^{cn} \eta_i| > cn/6) < 2e^{-(cn/6)^2/2cn} = 2e^{-2n}.$$

Therefore, the probability that there exist $x, y \in \{0,1\}^n$ such that $\sum_{i=1}^{cn} \eta_i > cn/6$ is less than $2e^{-2n}2^{2n} < 1$, so there exists a multiset $T = \{z_1, \ldots, z_{cn}\} \subset (\{0,1\}^{3n})^{cn}$ such that for all inputs (x,y) we have $\sum_{i=1}^{cn} \eta_i \leq cn/6$, which implies $\sum_{i=1}^{cn} \xi_i \leq cn(1/6 + 1/8) < cn/3$, and so the probability that P' errs on (x,y) is at most 1/3.

Note that this protocol is *not constructive*.

Proof of Theorem 5.4: Our protocol is motivated by a one-sided error randomized one-way protocol for EQ_n by Rabin, Simon, and Yao (cf. [8, Thm. 6.1, p. 22]): Alice and Bob are given n bits each, interpreted as integers less than 2^n. A random prime is chosen using $O(\log n)$ random bits. Alice and Bob compare the remainder of inputs modulo the prime. This communication takes $O(\log n)$ bits.

Instead of sending the remainders, we will test the two $O(\log n)$ bit remainders for equality by repeating the protocol by choosing a random prime using $O(\log\log n)$ additional random bits. This will give us remainders of $O(\log\log n)$ bits that we want to test for equality. We repeat this process until the size of the remainders is below a certain constant, and then communicate them. The formal proof follows.

Let $x, y \in \{0,1\}^n$.
We denote the ith iterated logarithm of n by $\log^{(i)} n$. Note: $\log^{(0)} n = n$, and $\log^{(-1)} n = 2^n$.
Let $\log^* n$ denote the least integer j such that $\log^{(j)} n \leq 1$.
Let $w(1) = 2$ and $w(i+1) = 2^{w(i)}$.
Let $k = \log^* n - 4$, so

$$16 = w(3) < \log^{(k)} n \leq w(4) = 2^{16}.$$

For $1 \leq i \leq k$, let p_i be a random prime

$$p_i \leq (\log^{(i-1)} n)^2.$$

For convenience, we let $p_0 = (\log^{(i-1)} n)^2 = 2^{2n}$.
Let $x_1 = x, y_1 = y$, and for $1 \leq i \leq k$, let

$$x_{i+1} := x_i \bmod p_i$$

and

$$y_{i+1} = y_i \bmod p_i.$$

Alice computes and sends x_{k+1} to the referee. Bob computes and sends y_{k+1} to the referee. The referee outputs 1 if $x_{k+1} = y_{k+1}$, and 0 otherwise.

Since $p_i \leq (\log^{(i-1)} n)^2$, we can choose p_i using $O(\log^{(i)} n)$ random bits. Therefore, the number of common random bits used is

$$\sum_{i=1}^{k} O(\log^{(i)} n) \;=\; O(\log n).$$

The number of bits sent by each of Alice and Bob is

$$\log p_k \;\leq\; \log((\log^{(k-1)} n)^2) \;=\; 2\log^{(k)} n$$

$$\leq\; 2 \cdot w(4) = 2^{17} \;\leq\; O(1).$$

Now we show the correctness of the protocol. If $x = y$, it is clear that the referee outputs 1 with probability 1. Let $x \neq y$. Then $x_1 \neq y_1$. The referee outputs 1 if and only if there exists an i, $1 \leq i \leq k$, such that $x_{i+1} = y_{i+1}$ and $x_i \neq y_i$. Therefore,

$$\Pr(\text{ref outputs 1}) \;\leq\; \sum_{i=1}^{k} \Pr(x_{i+1} = y_{i+1} \,|\, x_i \neq y_i). \quad (7)$$

Let $1 \leq i \leq k$, and assume that $x_i \neq y_i$. Then

$$\Pr(x_{i+1} = y_{i+1}) \;=\; \Pr(p_i \text{ divides } x_i - y_i). \quad (8)$$

Since $x_i, y_i \leq p_{i-1}$, we have $|x_i - y_i| \leq p_{i-1} \leq (\log^{(i-2)} n)^2$. Therefore, the number of prime divisors of $x_i - y_i$ is at most $2\log^{(i-1)} n$. Since p_i is chosen to be a random prime $\leq (\log^{(i-1)} n)^2$ and the number of primes $\leq (\log^{(i-1)} n)^2$ is greater than $(\log^{(i-1)} n)^2 / 2\log^{(i)} n$, we have

$$\Pr(p_i \text{ divides } x_i - y_i \,|\, x_i \neq y_i) \;\leq\; \frac{2\log^{(i-1)} n}{(\log^{(i-1)} n)^2 / 2\log^{(i)} n}$$

$$\leq\; \frac{4\log^{(i)} n}{\log^{(i-1)} n}.$$

From this, (7), and (8), it follows that

$$\Pr(\text{ref outputs 1}) \;\leq\; \sum_{i=1}^{k} \Pr(x_{i+1} = y_{i+1} \,|\, x_i \neq y_i)$$

$$\leq\; 4\sum_{i=1}^{k} \frac{\log^{(i)} n}{\log^{(i-1)} n} \;<\; \frac{8\log^{(k)} n}{\log^{(k-1)} n}$$

$$=\; \frac{8\log^{(k)} n}{2^{\log^{(k)} n}} \leq \frac{8 \cdot 16}{2^{16}} = \frac{1}{2^9}$$

■ Theorem 5.4.

Corollary 5.7 *There exists an explicit public-fee one-sided error randomized SM protocol for EQ_n of cost $O(\log(C_0(EQ_n)))$.* ■

References

[1] N. Alon and J. Spencer. *The Probabilistic Method*. John Wiley & Sons, Inc, 1992.

[2] A. Ambainis. Communication complexity in a 3-computer model. *Algorithmica*, 16(3):298 – 301, 1996.

[3] A. Ambainis. Upper bounds on multiparty communication complexity of shifts. *Proceedings of the 13th Symposium on Theoretical Aspects of Computer Science*, pages 631 – 642, 1996.

[4] L. Babai and P. Frankl. *Linear Algebra Methods in Combinatorics, Part I*. Preliminary Version, 1992.

[5] L. Babai, P. Kimmel, and S. V. Lokam. Simultaneous messages vs. communication. *Proceedings of the 12th Symposium on Theoretical Aspects of Computer Science*, pages 361 – 372, 1995.

[6] J. Bourgain and A. Wigderson. Private communication by Avi Wigderson, March, 1996.

[7] I. Kremer, N. Nisan, and D. Ron. On randomized one-round communication complexity. *Proceedings of the 27th ACM Symposium on the Theory of Computing*, pages 596 – 605, 1995.

[8] L. Lovász. Communication complexity: A survey. *Technical Report CS-TR-204-89, Department of Computer Science, Princeton University*, 1989.

[9] M. Naor. Private communication, cited in [7], 1994.

[10] I. Newman. Private communication, cited in [7], 1994.

[11] I. Newman and M. Szegedy. Public vs private coin flips in one round communication games. *Proceedings of the 28th ACM Symposium on the Theory of Computing*, pages 561 – 570, 1996.

[12] N. Nisan and A. Wigderson. Rounds in communication complexity revisited. *SIAM Journal on Computing*, 22(1):211 – 219, 1993.

[13] P. Pudlák, V. Rödl, and J. Sgall. Boolean circuits, tensor ranks and communication complexity. To apear in *SIAM Journal on Computing*, 26(3), 1996.

[14] A. C.-C. Yao. Some complexity questions related to distributive computing. *Proceedings of the 11th ACM Symposium on the Theory of Computing*, pages 209–213, 1979.