Bounds on Tradeoffs between Randomness and Communication Complexity

Ran Canetti* Oded Goldreich†

August 6, 1990

Abstract

Known results concerning the power of randomness are qualitative, in the sense that they only show that solutions exist or can be improved *if randomness is allowed.* We initiate a quantitative investigation of the power of randomness, in the context of communication complexity.

We prove general lower bounds on the length of the random input of parties computing a function f, depending on the number of bits communicated and the deterministic communication complexity of f. Four standard models for Communication Complexity are considered: the random input of the parties may be shared or local, and the communication may be one-way or two-way.

The bounds are shown to be tight. Namely, we demonstrate functions and protocols for these functions which meet the above bounds up to a constant factor. We do this for all the models, for all values of the deterministic communication complexity, and for all possible quantities of bits exchanged. Furthermore, we use an idea of [BN] to show that it is possible to reduce the number of random bits required by *any* protocol, without increasing the number of bits exchanged (up to a limit depending on the advantage achieved by the protocol).

* *bitnet:*canetti@techunix.bitnet

† *bitnet:* oded@techsel.bitnet. Supported by grant No. 86-00301 from the United States - Israel Binational Science Foundation (BSF), Jerusalem, Israel.

1 Introduction

The power of randomness in computation is a major issue in all aspects of computer-science, and is yet to be fully understood. There are many cases in which there are tremendous gaps between the complexities, or even possibilities of deterministic and randomized computations (e.g. routing [BH,V], Byzantine agreements [FLP,B,FL,FM], Communication Complexity [Y2,PS,MS]).

A method of 'smoothing' these gaps is to measure the 'quantity of randomization' of an algorithm, thus substituting the qualitative question "Is the algorithm deterministic or randomized?" by the quantitative question "How much randomization does the algorithm use?". A standard method for quantifying randomization is measuring the size of the sample-space, or in other words the length of the random input. We initiate a quantitative study of randomness in a computationally simple model: Communication Complexity.

The communication complexity of a function f, as defined by Yao [Y2], measures the minimum number of bits that have to be transfered between two parties in order to compute $f(x, y)$, when one party has x and the other has y.

Although for most functions randomization does not help [AFR], a tremendous gap between the two models of computation exists for some functions. For instance, Yao [Y2] showed that n bits of communication are needed in order to deterministically compute the identity function ($ID(x,y) = 1$ iff $x = y$), and Paturi and Simon [PS] showed a randomized protocol that uses only *two* bits of communication and computes ID (with probability greater than $\frac{1}{2}$).

In view of these gaps between the communication complexities of deterministic and randomized protocols, this model seems to be a convenient test-field for a quantitative investigation of the power of randomness in computation.

We show a tradeoff between the amount of randomness required by a protocol for computing a function, and the number of bits exchanged by the parties while executing the protocol. This tradeoff may be interpreted in two alternative ways. One interpretation is as a lower bound on the number of bits exchanged by parties using a specific number of coin-tosses (namely as a lower bound on the communication complexity of a function, depending on the length of the random input). The other interpretation (which is used in this paper) is as a lower bound on the number of coin-tosses used by a protocol as a function of the number of bits exchanged. It can be seen that the lower bound on the number of coin-tosses used in the protocol increases gradually from zero up to a maximum value (which is at most n), as the given number of bits exchanged decreases from the deterministic communication complexity of the function to the randomized one.

We consider the following variants of the model: the communication may be *one-way,* or *two-way* (with any number of rounds), and the random input may be shared by both parties, or split into two parts, each available to one party only. Other parameters considered are the deterministic communication complexity of the function being computed, and the advantage over $\frac{1}{2}$ achieved by the protocol. The bounds hold for any protocol that computes a function with probability greater than $\frac{1}{2}$ and does not depend heavily on the advantage over $\frac{1}{2}$ achieved by the protocol.

The tradeoff is tight in all these models and for all possible values of the deterministic communication complexity. We construct a sequence of functions that cover the range of all possible deterministic communication complexities. For each function, computation model and a given number of bits to be exchanged, we show a protocol that meets the corresponding lower bound up to a multiplicative constant.

A tradeoff between randomness and communication complexity was independently investigated by [FJM]. They consider the expected communication complexity of 'Las Vegas' protocols (i.e., no error allowed, as opposed to our 'Monte Carlo' model), in the two-way, local coins model. In their setting, they show a tight tradeoff similar to the one presented in this paper[1].

Note that a quantitative study of randomness was carried out in the context of oblivious routing [KR,KPU], and for cashing algorithms [RS].

Organization. In section 2 we define the models and the parameters to be discussed. Section 3 contains the lower bounds for the four models. In section 4 we present functions and protocols for these functions that demonstrate the tightness of the bounds. Finally, we use an observation of Babai and Newman [BN] to show that the number of coin-tosses used in *any* protocol can be reduced up to a value depending only on the advantage over $\frac{1}{2}$ achieved by the protocol.

2 Preliminaries

Let $f : \{0,1\}^n \times \{0,1\}^n \rightarrow \{0,1\}$ and let P_1 and P_2 be two parties having inputs $x, y \in \{0,1\}^n$ respectively, and communicating according to

[1] In fact, the lower bounds of [FJM] can be derived, up to an additive constant, from the bound presented in part (b) of Theorem 2 below. (They consider the range of more than $\sqrt{n}$ bits exchanged, where n is the input length.)

a randomized protocol π in order to compute $f(x,y)$. Denote the output of the protocol on input x,y as $\pi(x,y)$. (Namely, $\pi(x,y)$ is a random variable determined by the coin tosses of π.)

Consider the following two parameters.

- The communication may be *one-way* or *two-way*. In the one-way model, party P_1 sends a message to P_2, and party P_2 decides on the output. In the two-way model, the parties take turns on sending messages, until a party sends a special symbol *'halt'* and outputs the protocol answer. (We assume that the messages transfered between the parties are prefix-free.)

- The coins tosses used in the protocol may be *local* or *shared*. The outcome of a local coin is known only to the party tossing it, while the outcome of a shared coin is known to both parties without need of communication.

Clearly, local coins can be emulated by shared coins (and a one-way protocol is a special case of a two-way protocol).

In the shared coins models, let r_π denote the number of coin-tosses used during the execution of π on the worst input pair:

$$r_\pi \stackrel{\text{def}}{=} \max_{x,y\in\{0,1\}^n} r_\pi(x,y),$$

where $r_\pi(x,y)$ is the number of coin tosses in π on input (x,y). In the local coins models let r^1_π (r^2_π) denote the number of coin-tosses used by P_1 (P_2) during the execution of π on the worst input pair.

A protocol π *computes* a function f with advantage ϵ_π if $\epsilon_\pi > 0$, where

$$\epsilon_\pi = \min_{x,y\in\{0,1\}^n} \text{Prob}(\pi(x,y) = f(x,y)) - \frac{1}{2}.$$

(The probability is taken over the coin tosses of π.)

Let

$$m_\pi \stackrel{\text{def}}{=} \max_{x,y\in\{0,1\}^n, r\in\{0,1\}^{r_\pi(x,y)}} m_\pi(x,y,r),$$

where $m_\pi(x,y,r)$ is the number of bits transfered in protocol π on input x,y and coin-tosses r.

In order to measure the tightness of the bounds we use the following notation. Let $R_{m,\epsilon}(f)$ ($R^i_{m,\epsilon}(f)$, for $i \in \{1,2\}$) denote the minimum of r_π (r^i_π) over all protocols π that compute f with advantage at least ϵ, using up to m communication-bits (namely $m_\pi \leq m$, and $\epsilon_\pi \geq \epsilon$).

Define the one-way (two-way) deterministic communication complexity of a function f, denoted as $C_D^{1\to 2}(f)$ ($C_D^{1\rightleftharpoons 2}(f)$), as the minimum of m_π over all deterministic one-way (two-way) protocols π that compute f. (Note that $C_D^{1\to 2}(f)$ is the logarithm of the number of distinct rows in the matrix representation of f.[2])

A function is *non-degenerate* if all the rows (columns) in the matrix representation of f are distinct. For simplicity of presentation we consider only *non-degenerate* functions. However, the following discussion can be easily extended to all f.

In the sequel $\hat{x}$ will denote the integer that the binary representation of which is x, and $e \in_R D$ will denote that element e is chosen at random from domain D, with uniform probability distribution.

3 Lower bounds

We show lower bounds on $R_{m,\epsilon}(f)$, in the different models of computation. The proofs of the bounds use either combinatorial arguments (counting the number of vectors of a particular type), or simulations (of a randomized protocol by a deterministic one).

Table 1 contains the bounds for the four models.

[2] All the logarithms in this paper are of base 2.

Table 1: A summary of the bounds

	one-way	two-way
local coins	$R^1_{m,\epsilon}(f) \geq \frac{n}{2^m} - 1$ $R^1_{m,\epsilon}(f) \geq \log\left(\frac{n}{(1-\epsilon/2)m}\right)$	$R^i_{m,\epsilon}(f) \geq \frac{n}{2^m} - 1,\ i = 1,2$ $R^1_{m,\epsilon}(f) + R^2_{m,\epsilon}(f) \geq \log\left(\frac{\frac{C_D^{1\rightleftharpoons 2}(f)}{m} - 1}{1-2\epsilon}\right)$
shared coins	$R_{m,\epsilon}(f) \geq \log\left(\frac{\frac{n}{m}-1}{1-2\epsilon}\right)$	$R_{m,\epsilon}(f) \geq \log\left(\frac{\frac{C_D^{1\rightleftharpoons 2}(f)}{m} - 1}{1-2\epsilon}\right)$ $R_{m,\epsilon}(f) \geq \log\left(\frac{n}{2^m - \epsilon m}\right)$

3.1 The one-way, local coins model

Theorem 1 *Let $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ be a non-degenerate function, and let π be a one-way, local coins protocol that computes f with advantage ϵ. Then,*

$$\frac{\binom{2^{r^1_\pi} + 2^{m_\pi} - 1}{2^{r^1_\pi}}}{\binom{\frac{\epsilon}{2}2^{r^1_\pi} + 2^{m_\pi} - 1}{\frac{\epsilon}{2}2^{r^1_\pi}}} \geq 2^n. \qquad (1)$$

Approximating (1), we derive the following two inequalities, which hold simultaneously:

$$\begin{array}{lrcl} \text{(a)} & r^1_\pi & \geq & \frac{n}{2^{m_\pi}} - 1 \\ \text{(b)} & r^1_\pi & \geq & \log(\frac{n}{(1-\epsilon/2)m_\pi}). \end{array}$$

The first inequality is stronger for $0 < m_\pi < \log n - \log\log n$, and the second for $\log n < m_\pi < n$.

Proof. Consider an enumeration of all possible (up to 2^{m_π}) messages. Let p^x_i denote the probability that party P_1 sends the i-th message (denoted as msg_i) on input x, and q^y_i the probability that P_2, on input y and having received msg_i from P_1, outputs 1. Let $\vec{p^x} \stackrel{\text{def}}{=} (p^x_1 \cdots p^x_{2^{m_\pi}})$ and $\vec{q^y} \stackrel{\text{def}}{=} (q^y_1 \cdots q^y_{2^{m_\pi}})$. Then, $\text{Prob}(\pi(x,y) = 1) = \sum_{i=1}^{2^{m_\pi}} p^x_i \cdot q^y_i$.

Consider two different inputs $x, x' \in \{0,1\}^n$. Since f is non-degenerate, there exists y such that $f(x,y) \neq f(x',y)$. Since π computes f, we have that $|\text{Prob}(\pi(x,y) = 1) - \text{Prob}(\pi(x',y) = 1)| \geq 2\epsilon$. Since $0 \leq p^x_i, q^y_i \leq 1$, we have

$$\vec{p^x} \diamond \vec{p^{x'}} \stackrel{\text{def}}{=} \sum_{i=1}^{2^{m_\pi}} |\, p^x_i - p^{x'}_i \,| \geq$$

$$\geq |\, \sum_{i=1}^{2^{m_\pi}} p^x_i \cdot q^y_i - \sum_{i=1}^{2^{m_\pi}} p^{x'}_i \cdot q^y_i \,| \geq 2\epsilon.$$

Thus, the protocol implies the existence of 2^n distinct vectors $\vec{p^x}$ such that for every x, x', we have $\vec{p^x} \diamond \vec{p^{x'}} \geq 2\epsilon$.

However, since party P_1 tosses only r^1_π coins, each probability p^x_i may be assigned only $2^{r^1_\pi} + 1$ different values: $\{i \cdot 2^{-r^1_\pi} | 0 \leq i \leq 2^{r^1_\pi}\}$. Moreover,we have $\sum_{i=1}^{2^{m_\pi}} p^x_i = 1$. Thus, the number of distinct such vectors is at most the number of possibilities to partition $2^{r^1_\pi}$ elements among 2^{m_π} cells, namely $\binom{2^{r^1_\pi} + 2^{m_\pi} - 1}{2^{r^1_\pi}}$.

We estimate the maximum size of a set S of probability vectors satisfying for each pair $\vec{p} \diamond \vec{p'} \geq 2\epsilon$, in the following way. For every such vector $\vec{p}$, the number of vectors $\vec{p'}$ such that $\vec{p} \diamond \vec{p'} < \epsilon$ is at least the number of possibilities to partition $\frac{\epsilon}{2}2^{r_\pi}$ elements among 2^{m_π} cells, namely $\binom{\frac{\epsilon}{2}2^{r^1_\pi} + 2^{m_\pi} - 1}{\frac{\epsilon}{2}2^{r^1_\pi}}$. Thus, for every probability vector (out of the $\binom{2^{r^1_\pi} + 2^{m_\pi} - 1}{2^{r^1_\pi}}$ possible) appearing in set S, there exist $\binom{\frac{\epsilon}{2}2^{r^1_\pi} + 2^{m_\pi} - 1}{\frac{\epsilon}{2}2^{r^1_\pi}}$ distinct vectors that may not appear in S. The Theorem follows. $\square$

Note that there is no bound on r^2_π, as will become clear in section 4.1.1.

3.2 The two-way, local coins model

Theorem 2 *Let $f : \{0,1\}^n \times \{0,1\}^n \rightarrow \{0,1\}$ be a non-degenerate function, and let π be a two-way, local coins protocol that computes f with advantage ϵ. Then the following inequalities hold simultaneously:*

$$\text{(a)} \quad r^i_\pi \;\geq\; \frac{n}{2^{m_\pi}} - 1 \quad \textit{for } i \in \{1,2\}$$

$$\text{(b)} \quad {r^1}_\pi + {r^2}_\pi \;\geq\; \log\left(\frac{\frac{C_D^{1\rightleftharpoons 2}(f)}{m_\pi} - 1}{1-2\epsilon}\right)$$

Proof. Part (a). Consider a conversation, $con = u_1 \circ v_1 \circ \ldots u_k \circ v_k$, between the parties where u_j is the message sent by P_1 in round j of the conversation, v_j is the subsequent message of P_2, and k is the number of rounds. (The parsing is unique since the messages are prefix-free.) Let $p^x_j(con)$ denote the probability that P_1 sends u_j on input x, if the conversation was the corresponding prefix of con, and let $p^x(con) \stackrel{\text{def}}{=} \prod_{j=1}^k p^x_j(con)$. Consider an enumeration of all the (up to 2^{m_π}) conversations, and let $\vec{p^x} \stackrel{\text{def}}{=} (p^x(con_1) \ldots p^x(con_{2^{m_\pi}}))$. Let $q^y(con)$ and $\vec{q^y}$ be similarly defined, with respect to P_2. Let $C_1^{x,y}$ be the set containing all the conversations after which the output of the protocol, on input x, y, is 1. Thus, $\text{Prob}(\pi(x,y) = 1) = \sum_{con \in C_1^{x,y}} p^x(con) \cdot q^y(con)$.

Following the lines of the proof of Theorem 1, we conclude that there exist 2^n distinct such vectors. However, since each element in the vector $\vec{p^x}$ ($\vec{q^y}$) may be assigned only $2^{r^1_\pi}+1$ ($2^{r^2_\pi}+1$) different values, there exist at most $(2^{r^1_\pi}+1)^{2^{m_\pi}}$ distinct such vectors $\vec{p^x}$, and at most $(2^{r^2_\pi}+1)^{2^{m_\pi}}$ distinct vectors $\vec{q^y}$. Part (a) follows.

Part (b). This inequality is derived from the possibility to simulate a randomized protocol π by a deterministic one, going over the coin-tosses of π. Note that it is sufficient to go over more than a $(1-2\epsilon)$-fraction of the coin-tosses of π. $\Box$

3.3 The shared coins models

Theorem 3 *Let $f : \{0,1\}^n \times \{0,1\}^n \rightarrow \{0,1\}$ be a non-degenerate function, and let π be a one-way, shared coins protocol that computes f with advantage ϵ. Then,*

$$r_\pi \geq \log\left(\frac{\frac{n}{m_\pi} - 1}{1-2\epsilon}\right)$$

Proof. The proof uses the same simulation technique as in part (b) of Theorem 2. However, here the resulting deterministic protocol is one-way, and $C_D^{1\rightarrow 2}(f) = n$. $\Box$
(A similar result is achieved using a combinatorial argument.)

Theorem 4 *Let $f : \{0,1\}^n \times \{0,1\}^n \rightarrow \{0,1\}$ be a non-degenerate function, and let π be a two-way, shared coins protocol that computes f with advantage ϵ. Then the following inequalities hold simultaneously.*

$$\text{(a)} \quad r_\pi \geq \log\left(\frac{\frac{C_D^{1\rightleftharpoons 2}(f)}{m_\pi} - 1}{1-2\epsilon}\right)$$

$$\text{(b)} \quad \frac{2^{2^{m_\pi + r_\pi}}}{2^{m_\pi}\epsilon 2^{r_\pi}\binom{2^{r_\pi}}{\epsilon 2^{r_\pi}}} \;\geq\; 2^n.$$

Part (b) yields (after approximation),

$$r_\pi \geq \log\left(\frac{n}{2^{m_\pi} - \epsilon_\pi m_\pi}\right).$$

Note that the first bound is stronger for smaller values of r_π, and the second for larger values of r_π. The crossing point of the bounds is determined by $C_D^{1\rightleftharpoons 2}(f)$.

Proof. Part (a) is identical to part (b) in the proof of Theorem 2.

Part (b). Consider a conversation $con = u_1 \circ v_1 \circ \ldots u_k \circ v_k$, between the parties. Let $p^x_r(con)$ denote if P_1 'is willing to go along' con, on input x and coin-tosses r. Namely, $p^x_r(con) = 1$ if on input x and coin-tosses r and for every j, party P_1 sends u_j if the conversation is the corresponding prefix of con. Otherwise, $p^x_r(con) = 0$. Consider an enumeration of all the possible conversations and coin-tosses and let M^x be the $2^{r_\pi} \times 2^{m_\pi}$ boolean matrix where $M^x_{i,j} = p^x_{r_j}(con_i)$. Clearly, for every two distinct rows x, x' in the function table, the corresponding matrices $M^x, M^{x'}$ differ by at least $2\epsilon 2^{r_\pi}$ rows.

However, there are only $2^{2^{m_\pi} \cdot 2^{r_\pi}}$ distinct such matrices M. Moreover, the number of distinct matrices that differ from a given matrix by at most $\epsilon 2^{r_\pi}$ rows is at least $2^{m_\pi \epsilon 2^{r_\pi}} \binom{2^{r_\pi}}{\epsilon 2^{r_\pi}}$. Part (b) follows. $\Box$

4 Protocols

We demonstrate the tightness of our lower bounds (summarized in Table 1) by showing protocols for non-degenerate functions. We do this separately for the one-way and the two-way cases.

4.1 Tightness of the bounds for one-way communication

The tightness of the bounds for the one-way cases is demonstrated by showing two families of one-way protocols for the identity function (i.e. $ID(x, y) = 1$ iff $x = y$), one in the local coins model and the other in the shared coins model. Both families consist of protocols parameterized by the maximum number of bits exchanged by the parties (denoted as m_0). The protocols meet the corresponding bounds, up to a multiplicative factor, for all possible values of m_0.

Thus, the tightness of the bounds of Theorems 1 and 3 (the one-way cases) is established.

4.1.1 Local coins

The protocol uses as building blocks two known protocols for ID:

- The following protocol is due to [RY], and will be denoted as π_{RY}. On input $x, y \in \{0,1\}^n$, party P_1 chooses at random a prime p in $[n, 4n]$, and sends $(p, \hat{x} \bmod p)$ to P_2. Party P_2 outputs 1 iff $(\hat{x} \bmod p) = (\hat{y} \bmod p)$.
  It is easy to verify that if $x \neq y$ then $\text{Prob}(\pi(x,y) = 0) \geq \frac{2}{3} - o(1)$. (There are $(3 - o(1))\frac{n}{\log n}$ primes in $[n, 4n]$ and if $x \equiv y$ modulo more than $\frac{n}{\log n}$ of the primes then $x = y$.) If $x = y$ then $\pi(x,y) = 1$. Clearly, $m_{\pi_{RY}} = 2\log n - \log\log n + 4$, $r^1_{\pi_{RY}} = \log n - \log\log n + 2$, and $r^2_{\pi_{RY}} = 0$.
- The following protocol is a slight modification of a protocol presented in [PS], and will be denoted π_{PS}. The parties use 3 different messages. Let p^x_i denote the probability that P_1 sends the ith message on input x, and q^y_i the probability that P_2, on input x and having received the ith message, outputs 1. On inputs $x, y \in \{0,1\}^n$, let
  $$\vec{p^x} = c_x \left(\cos\left(\frac{\hat{x}\pi}{2^{n+1}}\right), \sin\left(\frac{\hat{x}\pi}{2^{n+1}}\right), 1 \right)$$
  where c_x is a normalization constant, and
  $$\vec{q^y} = \tfrac{1}{2}(1 + \cos\left(\tfrac{\hat{y}\pi}{2^{n+1}}\right), 1 + \sin\left(\tfrac{\hat{y}\pi}{2^{n+1}}\right),$$
  $$1 - \cos\left(\tfrac{\pi}{2^{n+2}}\right)).$$
  It is easy to verify that if $x = y$ then $\frac{1}{2} + 2^{-2n-6} \leq \text{Prob}(\pi(x,y) = 1) \leq \frac{1}{2} + 2^{-2n-4}$, and if $x \neq y$ then $\text{Prob}(\pi(x,y) = 0) \geq \frac{1}{2} + 2^{-2n-5}$. (Note that $\frac{1}{3} \leq c_x \leq \frac{1}{2}$ for all x, and that $\frac{\alpha^2}{4} < \cos\alpha < \frac{\alpha^2}{2}$ for $0 < \alpha < \frac{\pi}{2}$.) Note that we may approximate the probabilities by multiples of 2^{-2n-8}, and let $r^1_{\pi_{PS}} = r^2_{\pi_{PS}} = 2n + 8$.

The combined protocol with parameter m_0, operates as follows on inputs $x, y \in \{0,1\}^n$.

- For $m_0 \geq 2\log n - \log\log n + 4$. We partition the input to blocks of length k (to be computed)[3] and execute π_{RY} on each

[3] Assume that k divides n. Otherwise the last block will be shorter.

block, using the same prime for all the blocks. Namely:

Let $x = x_1 \ldots x_{\frac{n}{k}}$, and $y = y_1 \ldots y_{\frac{n}{k}}$, where each x_i, y_i is of length k.
Party P_1**:** Choose a prime $p \in_R [k, 4k]$. compute $x^p = (\hat{x}_1 \bmod p) \circ \ldots \circ (\hat{x}_{\frac{n}{k}} \bmod p)$. Send (p, x^p) to P_2.
Party P_2**:** Compute $y^p = (y_1 \bmod p) \circ \ldots \circ (y_{\frac{n}{k}} \bmod p)$. Output 1 iff $y^p = x^p$.
It can be verified that for these values of m_0,

$$\epsilon_\pi > \tfrac{1}{7}$$
$$m_\pi = (\tfrac{n}{k} + 1)\log 4k - \log\log k$$
$$r^1_\pi = \log\left(\tfrac{4k}{\log k}\right).$$

Setting k to be the largest integer such that $m_\pi \leq m_0$, and using Theorem 1 (part (b)), we get $r^1_\pi - R^1_{m_\pi,\epsilon_\pi}(ID) \leq 3$.

- For $m_0 \leq \log n - \log\log n + 4$. For these values of m_0, the parties execute π_{PS} on (x^p, y^p) instead of having party P_1 send x^p to P_2. Namely:

  Party P_1**:** Choose a prime $p \in_R [k, 4k]$. Compute $x^p = (\hat{x}_1 \bmod p) \circ \ldots \circ (\hat{x}_{\frac{n}{k}} \bmod p)$. Send p to P_2 and execute π_{PS} on x^p.
  Party P_2**:** compute $y^p = (\hat{y}_1 \bmod p) \circ \ldots \circ (\hat{y}_{\frac{n}{k}} \bmod p)$. Execute π_{PS} on y^p.
  Using the lower *and* upper bounds on $\epsilon_{\pi_{PS}}$, it can be verified that for these values of r_0,

  $$\epsilon_\pi \geq 2^{-\frac{52n}{k}\log k}$$
  $$r^1_\pi = (\tfrac{2n}{k} + 1)\log 4k - \log\log k + 8$$
  $$m_\pi = \log k - \log\log k + 4$$

  Setting k to be the *smallest* integer such that $m_\pi \leq m_0$, and using Theorem 1 (part (a)), we get that $r^1_\pi / R^1_{m_\pi,\epsilon_\pi}(ID) \leq 32 + o(1)$, for all k.

- For $\log n - \log\log n + 4 < m_0 < 2\log n - \log\log n + 4$. For these values of m_0, use the previous case with $k = n$ (namely one block). According to Theorem 1 (part (b)) we have $r^1_\pi / R^1_{m_\pi,\epsilon_\pi}(ID) \leq 3 + o(1)$.

4.1.2 Shared coins

First we note that $n+1$ shared coins are sufficient for computing *any* function with advantage 2^{-n}, exchanging one bit. (If the first n coin-tosses are equal to the input of P_1, it sends '1' and P_2 computes the function value. Otherwise P_1 sends '0' and P_2 outputs the value of the remaining coin.) However, this method cannot be modified to use smaller amounts of coin-tosses.

We show a one-way, shared coins protocol for ID. Given a limit of m_0 bits to be exchanged, the parties use $\log(\frac{n}{m_0}) + 2$ coin-tosses, and interpret the first $\log(\frac{n}{m_0})$ coins as an integer $i \in_R (1 \ldots \frac{n}{m_0})$. Let the inputs be partitioned to $\frac{n}{m_0}$ blocks $x = x_1 \ldots x_{\frac{n}{m_0}}$ and $y = y_1 \ldots y_{\frac{n}{m_0}}$, such that $|y_j| = |x_j| = m_0$.
Party P_1**:** Send x_i to P_2.
Party P_2**:** If $x_i \neq y_i$ output 0. Otherwise, if $i = 1$ output 1 with probability $\frac{3}{4}$, and if $i \neq 1$, output 1 with probability $\frac{1}{2}$. These random choices are implemented using the two remaining shared coins.

It can be verified that $\epsilon_\pi \geq \frac{m_0}{4n}$. According to Theorem 3, we have that in this model $r_\pi - R_{m_\pi,\epsilon_\pi}(ID) \leq 2$.

4.2 Tightness of the bounds for two-way communication

The bounds for the two-way models behave differently for $m_\pi < \log n$ and $\log n < m_\pi < n$.

In the first case, the bounds depend only on the input length n and on m_π. The bound of Theorem 2 (local coins) is similar to that of Theorem 1 (the corresponding one-way model), thus its tightness is already established. In order to establish the tightness of the bound of Theorem 4 (shared coins), we define and show a family of protocols for the PO function (see section 4.2.1).

In the second case ($\log n < m_\pi < n$), the bounds depend also on the deterministic communication complexity, $C_D^{1=2}(f)$, of the function being computed. Therefore, the tightness of the bounds is demonstrated by showing a family of functions, one for each possible value of $C_D^{1=2}(f)$; for each function and a given number of bits to be exchanged (denoted as m_0), we show a protocol for this function using a number of coin-tosses that meets the corresponding lower bound (see section 4.2.2).

Note that for functions f such that $C_D^{1=2}(f) = n$ (e.g. ID), the one-way and the two-way bounds are similar up to a multiplicative constant for every value of m_π, both in the local coins and the shared coins models. We therefore get that the one-way protocols for ID meet also the corresponding two-way bounds, for all values of m_π.

4.2.1 The Pointer function

Consider the following function.

$$PO(x,y) = \begin{cases} \text{the } \hat{y}\text{-th bit in } x & \text{if } \hat{y} < n \\ & \text{and } \hat{x} \geq n \\ \text{the } \hat{x}\text{-th bit in } y & \text{if } \hat{x} < n \\ & \text{and } \hat{y} \geq n \\ 0 & \text{otherwise.} \end{cases}$$

It can be seen that[4] $C_D^{1=2}(PO) = 2 + \log n$. A lower bound is due to the minimum number of generalized rectangles in a decomposition of PO (see [Y2] for details), and an upper bound is due to the deterministic special case of the following protocol.

We show a shared coins, two-way protocol for PO. This Protocol demonstrates the tightness of the bound of Theorem 4 for $3 \leq m_0 \leq \log n$, and will be used in the next section to show tightness for $\log n < m_\pi < n$.

Let $x, y \in \{0,1\}^n$ be the inputs, and let $k = \lceil \frac{n}{2^{m_0-3}} \rceil$. The parties use $\log k + 1$ coin-tosses, in the following way. As the function is divided into three regions, the parties first exchange two bits in order to learn the region their inputs are in. Without loss of generality, assume that $\hat{y} < n$ and $\hat{x} \geq n$. Then, party P_1 lets $x = x_1 \ldots x_k$, $|x_i| = \frac{n}{k}$. Let the parties interpret the first $\log k$ shared coins as $i \in_R (0 \ldots k-1)$.
Party P_2: If $\frac{in}{k} \leq \hat{y} < \frac{(i+1)n}{k}$ (namely if the $\hat{y}$-th bit in x is in the i-th block) then send $'1' \circ (\hat{y} - \frac{in}{k})$ to P_1. Otherwise send $'0'$.
Party P_1: If received $'0'$ output the value of the remaining shared coin. Otherwise, output the $(\hat{y} - \frac{in}{k})$-th bit in x_i (i.e. the $\hat{y}$-th bit in x).
Clearly, m_0 communication-bits are sufficient, and it can be verified that $\epsilon_\pi = \frac{1}{2k}$. According to part (b) in Theorem 4, we have $r_\pi - R_{m_\pi, \epsilon_\pi}(PO) \leq 4$.

Note that a similar local coins protocol (P_2 sends i as well) is also tight with the corresponding bound.

4.2.2 The hybrid functions

In order to show the tightness of the bounds for all values of $C_D^{1=2}(f)$, we define a series of hybrid functions, for $\log n \leq i \leq n$:

$$h_i(x,y) = \begin{cases} ID(x,y) & \text{if } \hat{x} < 2^i \text{ and } \hat{y} < 2^i \\ PO(x,y) & \text{otherwise} \end{cases}$$

It can be easily seen that[5] $i \leq C_D^{1=2}(h_i) \leq i + 2$. The main contributor to the communication complexity of each hybrid is the ID block, while the PO block is used to keep the function non-degenerate.

Protocols for h_i, parameterized by m_0 (the number of bits to be exchanged), either in the local coins or the shared coins models, are an immediate combination of the corresponding

[4] A result by [PS] can be extended to show that even the unbounded error communication complexity of PO, in the local coins models, is at least $\log n$.

[5] The bound $r_\pi \geq \frac{n}{2^{m_\pi}} - 1$ implies $C_D^{1=2}(f) \geq \log n$ for all non-degenerate functions f, thus this is the entire range of possible communication complexities.

protocols for ID and PO. Namely, the parties first learn, using two bits of communication, which region their inputs are in. Then the parties execute the appropriate protocol. It can be easily verified that for each h_i and for both the local and shared coins models, these combined protocols achieve the corresponding lower bounds up to a multiplicative constant, in the same manner as the protocols they consist of do.

5 Reducing the number of coin-tosses

Throughout the paper, we discussed the number of coin-tosses used by the parties as a function of the number of bits exchanged. (Namely, we showed lower bounds on r_π as a function of m_π and specific protocols meeting that bound.) The following theorem shows that the number of coin-tosses in *any* protocol can be reduced up to a value that depends only on the advantage achieved by the protocol, *without increasing the number of communication-bits used.*

Theorem 5 *Let $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ and let π be a protocol that computes f with advantage ϵ. Then, there exists a protocol π' (of the same model) such that $m_{\pi'} = m_\pi$ and*

(a) $r_{\pi'} \le \log n + 2\log \epsilon_\pi^{-1} + 4$ and $\epsilon_{\pi'} = \frac{\epsilon_\pi}{2}$ (for the shared coins models)

(b) $r_{\pi'} \le \log n + 2\log \epsilon_\pi^{-1} + 5$ and $\epsilon_{\pi'} = \frac{\epsilon_\pi}{3}$ (for the local coins models)

Remark. This upper bound on r_π as a function of ϵ_π, combined with the lower bounds on r_π as a function of m_π (in all the computation models), yields upper bounds on ϵ_π as a function of m_π. The resulting bound for the local coins models is[6]

$$\epsilon_\pi \le 32\sqrt{n} \cdot 2^{\frac{n}{2^{m_\pi+1}}}.$$

In the shared coins models, the resulting upper bound on ϵ_π is trivial.

Proof outline. The proof is based on an idea of Babai and Newman [BN]. We outline the proof for the shared coins case only. Consider the $2^{2n} \times 2^{r_\pi}$ table where a row corresponds to an input pair, a column to a sequence of coin-tosses, and an entry is 1 iff the output of protocol π on the corresponding input and coin-tosses is $f(x,y)$. Fix a row, and choose k columns at random. Clearly, there are at least $(\frac{1}{2}+\epsilon)2^{r_\pi}$ entries of value 1 in this row. Therefore, the probability of having chosen less than $(\frac{1}{2}+\frac{\epsilon}{2})k$ entries of value 1 is, according to the Hoefding inequality [ES], less than $e^{-\frac{k\epsilon^2}{8}}$. Summing this probability over all the 2^{2n} rows, and requiring that the total probability is less than one, we get that $k = \frac{16n}{\epsilon^2}$ is possible. We conclude that there exists a choice of k coin-toss sequences that yields an advantage $\frac{\epsilon}{2}$. Setting $r_\pi = \log k$, the theorem follows. □

[6] A similar bound can be achieved using known simulation techniques.

Acknowledgments

This work was initiated and inspired by a course on Communication Complexity given by Benny Chor, for whom we are greatly indebted. We also wish to thank Guy Even, Hugo Krawczyk, Eyal Kushilevitz, Benny Pinkas and Avi Wigderson for helpful discussions.

References

[AFR] Alon, N., P.Frankl and V. Rödl, "Geometrical Realization of Set Systems and Probabilistic Communication Complexity", *Proc. of 26th FOCS,* pp. 277-280, 1985.

[B] Ben-or M., "Another Advantage of Free Choice: Complete Asynchronous Agreement Protocols", *Proc. of 2nd PODC,* pp. 27-30, 1982.

[BH] Borodin, A., and J. E. Hopcroft, "Routing, Merging, and Sorting on Parallel Models of Computing", *Journal of Computer and System Science 30,* pp. 130-145, 1985.

[BN] Babai, L., and I. Newman, private communication via A. Wigderson, 1989.

[ES] Erdös P., and J. Spencer, *Probabilistic Methods in Combinatorics,* Academic Press, New York, 1974.

[FM] Feldman P., and S. Micali, "Optimal Algorithms for Byzantine Agreement", *Proc. of 20th STOC,* pp. 148-161, 1988.

[FL] Fischer, M. J., and N. A. Lynch, "A Lower Bound on the Time to Assure Interactive Consistency" *Information Processing Letters,* Vol. 14, No. 4, pp. 183-186, 1982.

[FLP] Fischer, M. J., N. A. Lynch, and N. Paterson, "Impossibility of Distributed Consensus with One Faulty Processor", *2nd Symposium on Principles of Database Systems,* 1983.

[FJM] Fleischer, R., H. Jung, and K. Melhorn, "A Time-Randomness Tradeoff for Communication Complexity", *4th International Workshop on Distributed Algorithms,* 1990.

[KR] Karloff, H. J., and P. Raghavan, "Randomized Algorithms and Pseudorandom Generators", *Proc. of 20th STOC,* 1988.

[KPU] Krizanc, D., D. Peleg, and E. Upfal, "A Time-Randomness Tradeoff for Oblivious Routing", *Proc. of 20th STOC*, pp. 93-102, 1988.

[MS] Mehlhorn, K., and E. Schmidt, "Las-Vegas is better than Determinism in VLSI and Distributed Computing", *Proc. of 14th STOC,* pp. 330-337, 1982.

[PS] Paturi, R., and J. Simon, " Probabilistic Communication Complexity ", *Journal of Computer and System Science,* Vol. 33, 106-123, 1986.

[RY] Rabin, M., and A. C. Yao, private communication via M. Rabin, 1990.

[V] Valiant, L. G., "A Scheme for Fast Parallel Communication", *SIAM Journal of Computing,* Vol.11, No. 2, pp. 350-361, 1982.

[Y1] Yao, A. C., "Probabilistic Complexity: Towards a Unified Measure of Complexity", *Proc. of 18th STOC,* 1977, pp. 222-227.

[RS] Raghavan, P., and M. Snir, "Memory Versus Randomization in On-line Algorithms", *ICALP 1989,* pp. 687-703.

[Y2] Yao, A. C., "Some Complexity Questions Related to Distributive Computing", *Proc. of 11th STOC,* 1979, pp. 209-213.