Average and Randomized Communication Complexity

ALON ORLITSKY[1] and ABBAS EL GAMAL[2]

Abstract

The communication complexity of a two-variable function $f(x,y)$ is the number of information bits two communicatiors need to exchange in order to compute f when, initially, each knows only one of the variables. There are several communication-complexity measures corresponding to whether (1) worst-case or average number of bits is considered, (2) computation errors are allowed or not, and (3) randomization is allowed or not. Tight bounds are provided for the typical behavior of all bounded-error communication-complexity measures of boolean functions. Let $\mathcal{F}_s$ be the set of functions $f:\{0,..,n-1\}\times\{0,..,n-1\}\to\{0,1\}$ that contain s ones (i.e., $|\{(x,y): f(x,y)=1\}|=s$). It is shown that for every $n\le s\le n^2/2$, the communication-complexity measures fall into two classes:

The $\log n$ Class: The error-free worst-case randomized complexity and, more importantly, the error-free worst-case deterministic complexity of most functions in $\mathcal{F}_s$ are between $\log n-4$ and $\log n+1$ bits. This holds even for $s=n$.

The $\log\frac{s}{n}$ Class: The ϵ-error worst-case randomized complexity and the ϵ-error average randomized complexity of most functions in $\mathcal{F}_s$ are between $(1-2\epsilon)(\log\frac{s}{n}-2\log\log\frac{s}{n})$ and $(1-2\epsilon)(\log\frac{s}{n}+5.3\log\log n)$ bits. More importantly, the error-free average

[1]AT&T Bell Laboratories, Murray Hill, NJ 07974

[2]Information Systems Laboratory, Department of Electrical Engineering, Stanford University, CA 94305

deterministic complexity of *all* functions in $\mathcal{F}_s$ is less than $\log \frac{s}{n} + 8.3 \log\log n$ bits. For most of these functions it is also $\geq \log \frac{s}{n} - 2\log\log \frac{s}{n}$ bits.

The difference between the complexity of the $\log n$ class, and that of the $\log \frac{s}{n}$ class, ranges from a constant (for $s \approx n^2/2$) to exponential (for $s \approx n \log n$). In particular, since most functions have about $n^2/2$ ones, all their complexity measures are around $\log n$ bits.

1 Introduction

Communication complexity was first introduced in [1] and [2]. It is concerned with different aspects of the following problem: n is a positive integer, f is a binary function defined on $\{0,..,n-1\}\times\{0,..,n-1\}$, and $\epsilon \geq 0$. Two *communicators*, $P_{\mathcal{X}}$ having a random integer $X \in \{0,..,n-1\}$ and $P_{\mathcal{Y}}$ having a random integer $Y \in \{0,..,n-1\}$ use a predetermined protocol to compute $f(X,Y)$ with probability of error $\leq \epsilon$. How many bits must they exchange on the average? at worst?

This section is aimed at making these notions more precise. In Subsection 1.1 we formally define the deterministic model. In Subsection 1.2 we describe randomized protocols and compare them to deterministic ones. In Subsection 1.3 we survey previous work and describe the results obtained in this paper. These results are proved in Section 2.

1.1 Deterministic Protocols

Communication is performed over a binary channel that can carry only one bit at a time. Therefore, each communicator, in his turn, transmits a (possibly empty) sequence of bits that we call a *string* or a *message*. Since empty messages are allowed, we lose no generality by assuming that $P_{\mathcal{X}}$ transmits the first message. Communication proceeds according to a predetermined protocol that, at each time, specifies the message transmitted and, at the end of the communication, decides on the computed value that the communicators ($P_{\mathcal{X}}$ and $P_{\mathcal{Y}}$) assume is the correct value of the function. The protocol is *deterministic* if whenever $P_{\mathcal{X}}$ and $P_{\mathcal{Y}}$ have the same random integers, they transmit the same sequence of messages and decide on the same computed value.

With this communication model in mind, we formally define deterministic protocols. The definition is similar to that of [3] and [4]. It differs from [3] in that (1) it views the protocol as a codebook – an approach that facilitates information-theoretic treatment, and (2) the computed value of the function does not have to be transmitted explicitly. An *input* is an ordered pair $(x, y) \in \{0, .., n-1\} \times \{0, .., n-1\}$ thought of as a value assignment for X and Y. If $b_k, \ldots, b_l$ are strings, we let $\langle b_k, \ldots, b_l \rangle$ or, equivalently, $\langle b_j \rangle_{j=k}^{l}$ denote the $(l+1-k)$-element sequence whose ith element is b_{k+i-1}. A *deterministic protocol* ϕ *for* $\{0, .., n-1\} \times \{0, .., n-1\}$ can therefore be defined as a mapping that with each input, associates:

1. a sequence $\langle b_1(x, y), b_2(x, y), \ldots, b_{m(x,y)}(x, y) \rangle$ of transmitted messages, all the odd messages to be transmitted by $P_{\mathcal{X}}$ and all the even messages — by $P_{\mathcal{Y}}$;

2. a *computed value* $v_\phi(x, y)$ that $P_{\mathcal{X}}$ and $P_{\mathcal{Y}}$ assume is equal to $f(x, y)$.

Although the number of messages, $m(x, y)$, is finite for all protocols of interest, we extend the definition of $b_i(x, y)$ indefinitely and let $b_i(x, y) \stackrel{\text{def}}{=} \emptyset$, the empty string, for all $i > m(x, y)$. This simplifies notation and shortens proofs.

Remarks:

1. The messages depend of course on the protocol used; but, to shorten notation, we write $b_i(x, y)$ and $m(x, y)$ without the (sub/super)script ϕ.

2. In general, each of $P_{\mathcal{X}}$ and $P_{\mathcal{Y}}$ could decide on a different computed value. However, any protocol ensuring that a given communicator (say $P_{\mathcal{X}}$) knows $f(X, Y)$ for all inputs, can be augmented by one bit to yield a protocol ensuring that both $P_{\mathcal{X}}$ and $P_{\mathcal{Y}}$ know $f(X, Y)$ for all inputs. (Note that requiring merely that at least one communicator

know $f(X,Y)$ for all inputs trivializes the problem as $P_{\mathcal{X}}$ can always decide that $f(X,Y)=0$ and $P_{\mathcal{Y}}$ that $f(X,Y)=1$).

Not every protocol can be carried out in a distributed environment where, initially, $P_{\mathcal{X}}$ knows only X, $P_{\mathcal{Y}}$ only knows Y, and their only information about each other's random variable derives from previously exchanged messages. To enable communication in such an environment, we restrict consideration to protocols possessing the following properties:

Separate-transmissions property

All that a communicator knows prior to transmitting a message is his random variable and previous transmissions. To guarantee that his message will not depend on the other communicator's variable, we require that for all x, x', y, y':

for odd i, $\langle b_j(x,y)\rangle_{j=1}^{i-1} = \langle b_j(x,y')\rangle_{j=1}^{i-1}$ implies that $b_i(x,y) = b_i(x,y')$,

and,

for even i, $\langle b_j(x,y)\rangle_{j=1}^{i-1} = \langle b_j(x',y)\rangle_{j=1}^{i-1}$ implies that $b_i(x,y) = b_i(x',y)$.

Prefix-free-messages property

Let $\alpha_1,\ldots,\alpha_k$ and $\beta_1,\ldots,\beta_l$ be bits. The string $\alpha_1\cdots\alpha_k$ is a *prefix* of the string $\beta_1\cdots\beta_l$ if $k \le l$ and $\alpha_i = \beta_i$ for $i = 1,\ldots,k$. It is a *proper prefix* if, in addition, $k \ne l$. A set of strings is called *prefix free* if no string in the set is a prefix of another. In order for each person to know when the message he receives terminates, we require that, given his random variable and previous transmissions, the set of all possible messages he can receive is prefix-free. Hence,

for odd i, $\langle b_j(x,y)\rangle_{j=1}^{i-1} = \langle b_j(x',y)\rangle_{j=1}^{i-1}$ implies that neither $b_i(x,y)$ nor $b_i(x',y)$ is a proper prefix of the other, and,

for even i, $\langle b_j(x,y)\rangle_{j=1}^{i-1} = \langle b_j(x,y')\rangle_{j=1}^{i-1}$ implies that neither $b_i(x,y)$ nor $b_i(x,y')$ is a proper prefix of the other.

Separate-decisions property

When a communicator decides on the computed value, he must, again, base his decision on his random variable and the messages exchanged. Therefore,

$\langle b_j(x,y)\rangle_{j=1}^{\infty} = \langle b_j(x,y')\rangle_{j=1}^{\infty}$ implies that $v_\phi(x,y) = v_\phi(x,y')$,

and

$\langle b_j(x,y)\rangle_{j=1}^{\infty} = \langle b_j(x',y)\rangle_{j=1}^{\infty}$ implies that $v_\phi(x,y) = v_\phi(x',y)$.

Remarks:

1. Note that the first two properties must apply to all i's, even those larger than $m(x,y)$.

2. It can be shown that although empty messages are allowed, the prefix-free-messages property does not increase the amount of communication required. Every protocol containing empty messages can be trivially modified to obtain a protocol that transmits the same number of bits but does not contain any empty message.

From here on we assume that all protocols possess these three properties and we turn to the complexity definitions.

Let $l_\phi(x,y) \stackrel{\text{def}}{=} \sum_{i=1}^{m(x,y)} |b_i(x,y)|$ denote the total number of bits transmitted according to the protocol ϕ when $P_{\mathcal{X}}$ knows x and $P_{\mathcal{Y}}$ knows y. The *worst-case complexity of* ϕ is defined as:

$$\hat{L}_\phi \stackrel{\text{def}}{=} \max\{l_\phi(x,y) : (x,y) \in \{0,..,n-1\} \times \{0,..,n-1\}\}\ .$$

The *average complexity of* ϕ is defined as:

$$\bar{L}_\phi \stackrel{\text{def}}{=} \frac{1}{n^2} \sum_{(x,y)\in\{0,..,n-1\}\times\{0,..,n-1\}} l_\phi(x,y) \ .$$

If f is a function defined on $\{0,..,n-1\}\times\{0,..,n-1\}$ and ϕ is a protocol for $\{0,..,n-1\}\times\{0,..,n-1\}$, we say that ϕ is *error-free* for f if $v_\phi(x,y) = f(x,y)$ for all $(x,y) \in \{0,..,n-1\}\times\{0,..,n-1\}$. The *worst-case error-free communication complexity of* f is then defined as:

$$\hat{C}_D(f,0) \stackrel{\text{def}}{=} \min\{\hat{L}_\phi(x,y) : \phi \text{ is an error-free deterministic protocol for } f\} \ .$$

And the *average error-free communication complexity of* f:

$$\bar{C}_D(f,0) \stackrel{\text{def}}{=} \min\{\bar{L}_\phi(x,y) : \phi \text{ is an error-free deterministic protocol for } f\} \ .$$

The subscript $_D$ stands for *deterministic*, whereas the 0 indicates that no errors are allowed (0 probability of error). As might be suspected by the excessive notation, both requirements will soon be relaxed.

1.2 Randomized protocols

Randomized protocols are like deterministic ones, except that $P_\mathcal{X}$ and $P_\mathcal{Y}$ may use "coin flips" to determine their transmissions. Whereas deterministic protocols require that $P_\mathcal{X}$ base his transmissions and his computed value only on X and preceding transmissions (similarly for $P_\mathcal{Y}$), randomized protocols require that $P_\mathcal{X}$ base the bias of the coin flips (which determine his transmissions and his computed value) only on X and preceding transmissions (similarly for $P_\mathcal{Y}$). Still, the set of all messages that have positive probability at any time must be prefix free. And, the computed values must always agree (this requires a simple rejustification).

We denote randomized protocols by Φ (as opposed to ϕ for deterministic protocols). The number of bits transmitted when $P_{\mathcal{X}}$ knows x and $P_{\mathcal{Y}}$ knows y is now a random variable denoted by $L_\Phi(x,y)$. Its expected value is denoted by $\bar{L}_\Phi(x,y)$.

The definitions of average and worst-case complexities of protocols can now be extended to randomized protocols. The *average complexity* $\bar{L}_\Phi$ *of* Φ is:

$$\bar{L}_\Phi \stackrel{\text{def}}{=} \frac{1}{n^2} \sum_{(x,y)\in\{0,..,n-1\}\times\{0,..,n-1\}} \bar{L}_\Phi(x,y).$$

The *worst-case complexity* $\hat{L}_\Phi$ *of* Φ is:

$$\hat{L}_\Phi \stackrel{\text{def}}{=} \max\{\bar{L}_\Phi(x,y) : (x,y) \in \{0,..,n-1\}\times\{0,..,n-1\}\}.$$

Similarly, the *computed value*, $V_\Phi(x,y)$, which $P_{\mathcal{X}}$ and $P_{\mathcal{Y}}$ assume is the value of $f(x,y)$ is now a 0-1 random variable. We let $\bar{E}_\Phi(x,y) \stackrel{\text{def}}{=} p(V_\Phi(x,y) \neq f(x,y))$ denote the probability that the computed value determined by Φ is wrong for (x,y), and define $\hat{E}_\Phi(f)$, the worst case error incurred by Φ in computing f, as

$$\hat{E}_\Phi(f) \stackrel{\text{def}}{=} \max\{\bar{E}_\Phi(x,y) : (x,y) \in \{0,..,n-1\}\times\{0,..,n-1\}\}.$$

Remarks:

1. The maximum in $\hat{L}_\Phi$ and $\hat{E}_\Phi(f)$ is taken over the inputs. For each input we still average over the "coin flips".

2. For brevity we only consider worst case errors. Average error results (when average errors are appropriately defined) can also be obtained.

Using these quantities we define two more communication-complexity measures. The *worst-case randomized communication complexity of a function f with ϵ error*:

$$\hat{C}_R(f,\hat{\epsilon}) \stackrel{\text{def}}{=} \min\{\hat{L}_\Phi : \hat{E}_\Phi(f) \le \epsilon\},$$

and the *average randomized communication complexity of f with ϵ error*:

$$\bar{C}_R(f,\hat{\epsilon}) \stackrel{\text{def}}{=} \min\{\bar{L}_\Phi : \hat{E}_\Phi(f) \le \epsilon\}.$$

Again, $\hat{C}$ denotes worst-case complexity while $\bar{C}$ stands for average complexity and the subscript $_R$ indicates that randomized protocols are allowed. The hat in $\hat{\epsilon}$ means that ϵ worst-case error is permitted (we will later use $\bar{\epsilon}$ for ϵ average error).

A simple relationship between randomized and deterministic protocols is useful in proving lower bounds on randomized protocols. Let Φ be a randomized protocol. Each transmission is the outcome of a random experiment whose probability distribution is determined by Φ according to the integer known to the transmitter and previous transmissions. Conceptually, each communicator can perform all the random experiments (corresponding to all possible integers he might have and all possible transmissions) prior to the commencement of the communication. Then, given his value and previous transmissions, he consults the appropriate experiment and discards the rest.

These two huge random experiments, each performed by one communicator, are independent of the current input. Every combined outcome of these two random experiments induces a deterministic protocol. Let $p(\phi)$ be the total probability of outcomes that induce the deterministic protocol ϕ. As before, each ϕ determines $l_\phi(x,y)$ and $v_\phi(x,y)$. The expected communication length for (x,y), denoted earlier in this subsection as $\bar{L}_\Phi(x,y)$, and the probability of error $\bar{E}_\Phi(x,y)$ can now be written as:

$$\begin{aligned} \bar{L}_\Phi(x,y) &= \sum_\phi p(\phi) l_\phi(x,y) \;, \\ \bar{E}_\Phi(x,y) &= \sum_{\{\phi : v_\phi(x,y) \neq f(x,y)\}} p(\phi). \end{aligned} \tag{1}$$

These equations clearly imply that:

$$\begin{aligned} \bar{L}_\Phi &= \sum_\phi p(\phi) \bar{L}_\phi \;, \\ \bar{E}_\Phi(f) &= \sum_\phi p(\phi) \bar{E}_\phi(f). \end{aligned} \tag{2}$$

Before proceeding further, we note that there are two possible models.

1. *Shared random sources:* the communicators have access to each other's random variable.

2. *Separate random sources:* the communicators do not have access to each other's random variable.

Clearly, $\hat{C}_R(f,\hat{\epsilon})$ and $\bar{C}_R(f,\hat{\epsilon})$ depend on the model assumed. The bounds we prove hold in a strong sense: all the lower bounds apply even if we combine the two random experiments into one — known to both communicators. All the upper bounds apply even if the experiments are separated, each known to one communicator. Therefore, all bounds apply to both models.

We have so far defined four complexity measures: $\hat{C}_D(f,0)$, $\bar{C}_D(f,0)$, $\hat{C}_R(f,\hat{\epsilon})$, and $\bar{C}_R(f,\hat{\epsilon})$. Of special interest are the randomized complexities when no errors are allowed ($\epsilon = 0$) — the so called *Las Vegas Complexities*. There are two Las Vegas complexity measures: $\hat{C}_R(f,0)$ and $\bar{C}_R(f,0)$ but it can be easily seen that for all functions f, $\bar{C}_R(f,0) = \bar{C}_D(f,0)$. (We omit the hats because 0 worst-case and 0 average errors are the same). This leaves five different measures. In the rest of the paper we try to determine how hard each one is.

Remark: As in the case of randomized complexities, one could define the deterministic measures $\bar{C}_D(f,\hat{\epsilon})$ and $\hat{C}_D(f,\hat{\epsilon})$. However, for deterministic protocols, $\hat{E}_\phi(f)$ is either 0 or 1, thus $\hat{C}_D(f,\hat{\epsilon}) = \hat{C}_D(f,0)$ and $\bar{C}_D(f,\hat{\epsilon}) = \bar{C}_D(f,0)$.

1.3 Previous and New Results

Complexity measure A is *smaller* than complexity measure B if $A(f) \leq B(f)$ for all functions f. For example, $\bar{C}_R(f,\hat{\epsilon}) \leq \hat{C}_R(f,\hat{\epsilon})$ for all functions f. This relation induces a partial order on the five complexity measures which is depicted in Figure 1.

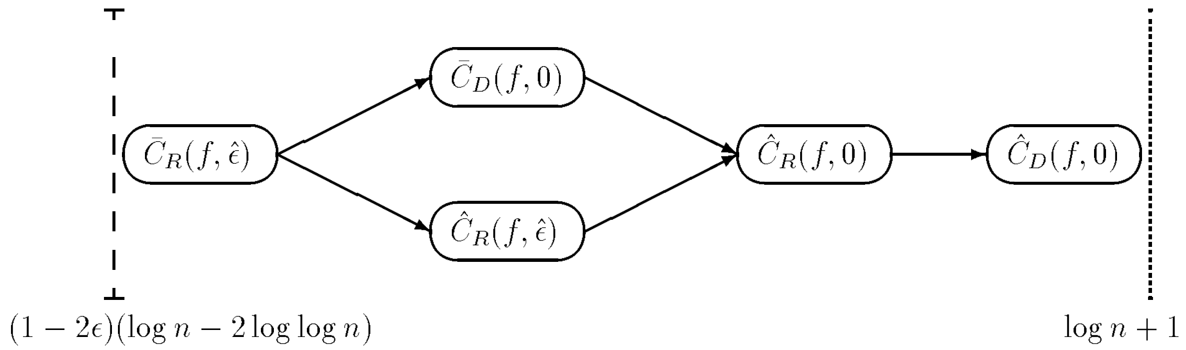

Figure 1: Known relations for most functions

The vertical lines bound the complexity of most functions. The dotted line therefore means that for most functions, all complexity measures are smaller than $\log n + 1$ (in this case, it is trivially true for all functions). The dashed line and other known results are described in the next two paragraphs.

[2] proved that for most functions, $\hat{C}_D(f,0)$ is the worst possible: about $\log n$. The proof can be easily modified to show that the same is true for $\bar{C}_D(f,0)$. For randomized protocols, [2] showed that the equality function (defined in Example 1), has $\hat{C}_D(equ,0) = \log n + 1$, but,

for all $\epsilon > 0$, $\hat{C}_R(equ, \hat{\epsilon}) = \Theta(\log\log n)$. [5] demonstrated a function (defined in example 2) for which $\hat{C}_D(f,0) = N$, but $\hat{C}_R(f,0) = O(\sqrt{N}\log^2 N)$.

More recently, [6] and [7] showed that for most functions $\bar{C}_R(f,\hat{\epsilon}) > (1-2\epsilon)(\log n - \log\log n) - 2$. This result, indicated in Figure 1 by a dashed vertical line, implies that for most functions allowing randomization, ϵ error, and even measuring only the average complexity, cannot help — the complexity remains about the same. In particular, it means that the equality and the function of [5] are more the exception than the rule.

Let $\mathbf{1}_f \stackrel{\text{def}}{=} \{(x,y) : f(x,y) = 1\}$ be the set of inputs for which f is 1, and $\mathbf{0}_f$ — the set of inputs for which f is 0. The vast majority of binary functions have about equal number of ones and zeros in their function table. Therefore, saying that a certain property holds for most functions, is about the same as saying that it holds for most functions with $|\mathbf{1}_f| \approx n^2/2$. A natural question, thus, is whether the complexity is reduced when the number of ones and zeros is not balanced, i.e., $|\mathbf{1}_f| \ll n^2/2$ or $|\mathbf{0}_f| \ll n^2/2$.

This is similar to compression of binary sequences. We know that most sequences are incompressible. But more insight is obtained by noting that most sequences are compressible to their entropy (yet the entropy of most sequences is maximal). Does a similar property hold for communication complexity?

Without loss of generality, we consider only functions with fewer ones than zeros, that is, $\mathbf{1}_f \le n^2/2$. For every integer $s \le n^2/2$, define $\mathcal{F}_s \stackrel{\text{def}}{=} \{f : |\mathbf{1}_f| = s\}$ (the set of functions with s ones). We prove that for all $s \ge n$:

1. [Corollary 4] Most functions in $\mathcal{F}_s$ have worst case error-free (Las Vegas) complexities, $\hat{C}_D(f,0)$ and $\hat{C}_R(f,0)$, at least $\log n - 4$ (even when $s = n$).

2. [Corollary 7] However, *all* functions in $\mathcal{F}_s$ have $\bar{C}_D(f,0) \le \log\frac{s}{n} + 8.3\log\log n + c$ That is, a deterministic protocol that commits no errors and, on the average, exchanges $\le \log\frac{s}{n} + 8.3\log\log n + c$ bits.

3. [Corollary 6] Most of the functions in $\mathcal{F}_s$ have $\hat{C}_R(f,\hat{\epsilon}) \le (1-2\epsilon)(\log\frac{s}{n} + 5.3\log\log n + 2\log\frac{1}{\epsilon})$.

4. [Corollary 8] For most of the functions in $\mathcal{F}_s$, the bounds in (2) and (3) are tight: $\bar{C}_R(f,\hat{\epsilon}) \ge (1-2\epsilon)(\log\frac{s}{n} - 2\log\log\frac{s}{n}) - 9$.

All the bounds proved apply to both models: shared and separate sources of randomness, and they are all tight up to an additive term. They are summarized in Figure 2. As in Figure 1, dotted vertical lines denote upper bounds and dashed lines denote lower bounds.

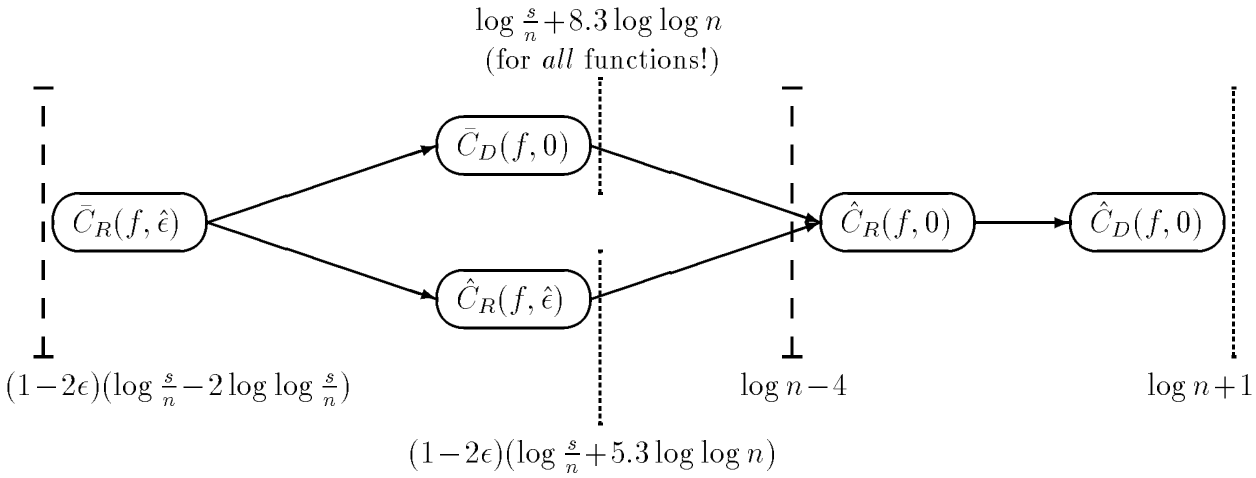

Figure 2: Bounds for most functions in $\mathcal{F}_s$ $\quad (n \le s \le n^2/2)$

Some of the consequences of these results are:

1. All complexity measures fall in two very distinct classes. One (consisting of $\hat{C}_D(f,0)$ and $\hat{C}_R(f,0)$) is at least $\log n - 4$ bits for every s. The other (consisting of $\bar{C}_D(f,0)$, $\hat{C}_R(f,\hat{\epsilon})$, and $\bar{C}_R(f,\hat{\epsilon})$) decreases with s and is about $\log \frac{s}{n}$ bits. For $s \approx n^2/2$, there is a small difference between the complexity measures in the two classes. For $s \approx n$ or $s \approx n \log n$, the complexity measures in the first class are exponentially larger than those in the second.

For most functions, $|\mathbf{1}_f| \approx \frac{n^2}{2}$. Therefore, all complexity measures coincide and are about $\log n$ bits. However, for most sparse functions ($|\mathbf{1}_f| \ll \frac{n^2}{2}$):

2. In error-free computation, considerably fewer bits are needed on the average than at worst (most interesting, of course, is the difference between $\hat{C}_D(f,0)$ and $\bar{C}_D(f,0)$).

3. In error-free computation, randomization does not help reduce the worst-case complexity. Even for $s \approx n$, most functions in $\mathcal{F}_s$ have $\hat{C}_R(f,0) \geq \log n - 4$. Hence, $\hat{C}_R(f,0)$ is not much smaller than $\hat{C}_D(f,0)$, and the function described in [5] (see Example 2) is an exception.

4. However, if errors are allowed, randomization does reduce the worst-case complexity. In fact, in this case, it is close to average error-free deterministic complexity. Thus the behavior of the equality function is typical of functions in its class: $\mathcal{F}_n$.

2 Proofs

In this section we prove the results claimed. In the first subsection, we prepare the groundwork by exploring some information theoretic properties of protocols. In subsection 2.2

we show that for most functions, $\hat{C}_R(f,0) \geq \log n - 4$ (Corollary 4). In subsection 2.3 we prove the separation between the classes by showing that for all functions $\bar{C}_D(f,0) \leq \log \frac{|\mathbf{1}_f|}{n} + 8.3 \log\log n$ and that, for most functions, $\hat{C}_R(f,\hat{\epsilon}) \leq (1-2\epsilon)(\log \frac{|\mathbf{1}_f|}{n} + 5.3 \log\log n) + 2$ (Corollaries 6 and 7). These results are deduced from a stronger bound that takes into consideration the distribution of ones in the function table. Last, in subsection 2.4 Corollary 8, we show that these bounds are tight up to a negligible additive term by proving that for most functions $\bar{C}_R(f,\hat{\epsilon}) \geq (1-2\epsilon)(\log \frac{|\mathbf{1}_f|}{n} - 2\log\log \frac{|\mathbf{1}_f|}{n}) - 9$.

2.1 Preliminary Results

If $b_k, \ldots, b_l$ are strings, we let $[b_k \ldots b_l]$ or, equivalently, $[b_j]_{j=k}^{l}$ denote the string obtained by their concatenation. We say that the sequence $\langle b_1, \ldots, b_l \rangle$ *parses* the string b if $b = [b_1 \ldots b_l]$. In general, there are many ways to parse a string.

For every (x,y), we let $b_\phi(x,y) \stackrel{\text{def}}{=} [b_i(x,y)]_{i=1}^{m(x,y)}$, *the codeword associated with* (x,y), be the concatenation of all messages transmitted according to ϕ when (x,y) is the input. The next theorem shows that only one parsing of $b_\phi(x,y)$ complies with the separate-transmissions and the prefix-free-messages properties.

Theorem 1 If either $b_\phi(x,y)$ or $b_\phi(x',y')$ is a prefix of the other, then

$$\langle b_j(x,y)\rangle_{j=1}^{\infty} = \langle b_j(x',y)\rangle_{j=1}^{\infty} = \langle b_j(x,y')\rangle_{j=1}^{\infty} = \langle b_j(x',y')\rangle_{j=1}^{\infty}.$$

Proof We show by induction on i that $\langle b_j(x,y)\rangle_{j=1}^{i} = \langle b_j(x',y)\rangle_{j=1}^{i} = \langle b_j(x,y')\rangle_{j=1}^{i} = \langle b_j(x',y')\rangle_{j=1}^{i}$. The induction hypothesis is clearly true for $i = 0$. Assume it is true for $i-1$; then, $[b_j(x,y)]_{j=1}^{i-1} = [b_j(x',y')]_{j=1}^{i-1}$. Hence, one of $[b_j(x,y)]_{j=i}^{\infty}$ and $[b_j(x',y')]_{j=i}^{\infty}$ is a prefix of the other, which implies that one of $b_i(x,y)$ and $b_i(x',y')$ is a prefix of the other. If i is odd, then, by the separate-transmissions property and the induction hypothesis, $b_i(x, y') = b_i(x, y)$. By the prefix-free-messages property, neither $b_i(x, y')$ nor $b_i(x', y')$ can be a proper prefix of the other and hence they must be the same. Using the same argument for (x', y), we get $b_i(x, y) = b_i(x', y) = b_i(x, y') = b_i(x', y')$. The symmetric argument works for even i's, which completes the proof of the induction. $\square$

Besides showing that there is a unique way to parse $b_\phi(x, y)$ into the individual messages, the theorem has three other implications. They are used extensively later so we restate them as individual corollaries. To do so, we need the following definitions.

A *rectangle* is a set of the form $R_\mathcal{X} \times R_\mathcal{Y}$ where $R_\mathcal{X}, R_\mathcal{Y} \subseteq \{0, .., n-1\}$. If S is a subset of $\{0, .., n-1\} \times \{0, .., n-1\}$, we define its *closure* to be $\mathcal{C}\ell(S) \stackrel{\text{def}}{=} \{(x, y) :$ for some u and v in $\{0, .., n-1\}$, both (x, u) and (v, y) are in $S\}$. It can be shown that $\mathcal{C}\ell(S)$ is the intersection of all rectangles that contain S and is therefore a rectangle. The theorem implies that if $b_\phi(x, y) = b_\phi(x', y')$, then the "corner points" (x, y') and (x', y) have the same codeword. Therefore, if a protocol assigns the same codeword to all inputs in a set S, it assigns the same codeword to all inputs in $\mathcal{C}\ell(S)$.

Let $R_\phi(x, y) \stackrel{\text{def}}{=} \{(u, v) : b_\phi(u, v) = b_\phi(x, y)\}$ be the set of inputs having the same codeword as (x, y). The first corollary of Theorem 1 states that $R_\phi(x, y)$ can assume only special forms.

Corollary 1 For all $(x, y) \in \{0, .., n-1\} \times \{0, .., n-1\}$, $R_\phi(x, y)$ is a rectangle.

Proof By the theorem, $(u, v) \in \mathcal{C}\ell(R_\phi(x, y))$ implies that $b_\phi(u, v) = b_\phi(x, y)$, i.e., $(u, v) \in R_\phi(x, y)$. Thus, $\mathcal{C}\ell(R_\phi(x, y)) \subseteq R_\phi(x, y)$. But, $R_\phi(x, y)$ is always contained in $\mathcal{C}\ell(R_\phi(x, y))$. Hence, $R_\phi(x, y) = \mathcal{C}\ell(R_\phi(x, y))$ which is a rectangle. $\square$

This corollary has an intuitive "near-proof". Consider the set S of inputs that can result in a given sequence of messages. By the separate-transmissions property, whenever $P_{\mathcal{X}}$ transmits a message, he partitions this set into *horizontal slices* (sets of the form $S \cap (A_i \times \{0,..,n-1\})$ where the A_is partition S) – one slice corresponding to each message. (Note that this would not have been the case without the separate-transmissions property). Similarly, when $P_{\mathcal{Y}}$ transmits a message, he partitions S into *vertical slices* (sets of the form $S \cap (\{0,..,n-1\} \times A_{\mathcal{Y}})$). Since any horizontal (or vertical) slice of a rectangle is again a rectangle, it is clear, by induction, that $R_\phi(x,y)$ is a rectangle. The missing part in this "proof" is showing that different rectangles cannot correspond to the same concatenation of messages. This is the essence of Theorem 1.

The second corollary states that the codewords associated with the rectangles form a prefix-free set. Let $\mathcal{R}_\phi \stackrel{\text{def}}{=} \{R_\phi(x,y) : (x,y) \in \{0,..,n–1\}\times\{0,..,n–1\}\}$ be the set of rectangles over which ϕ has a fixed codeword. Any input (x,y) belongs to the rectangle $R_\phi(x,y)$, and any two intersecting rectangles in $\mathcal{R}_\phi$ coincide. Therefore, $\mathcal{R}_\phi$ *partitions* $\{0,..,n–1\}\times\{0,..,n–1\}$ into $|\mathcal{R}_\phi|$ rectangles where $|S|$ denotes the cardinality of the set S. For each rectangle $R \in \mathcal{R}_\phi$, define $b_\phi(R)$ to be the unique codeword that ϕ assigns to all inputs in R. Then Theorem 1 implies:

Corollary 2 $\{b_\phi(R) : R \in \mathcal{R}_\phi\}$ is prefix free and its cardinality is $|\mathcal{R}_\phi|$. □

The third corollary combines the theorem with the separate-decisions property to show:

Corollary 3 For all $R \in \mathcal{R}_\phi$, v_ϕ is constant over R.

Proof Let (x,y), (x',y') be two inputs in a rectangle $R \in \mathcal{R}_\phi$. From the theorem, $\langle b_j(x,y)\rangle_{j=1}^{\infty} = \langle b_j(x',y)\rangle_{j=1}^{\infty} = \langle b_j(x',y')\rangle_{j=1}^{\infty}$. By the separate-decisions property, $v_\phi(x,y) =$

$v_\phi(x', y) = v_\phi(x', y')$. $\square$

In summary, the three corollaries show that every deterministic protocol ϕ determines:

1. A partition $\mathcal{R}_\phi$ of $\{0, .., n-1\} \times \{0, .., n-1\}$ into rectangles. We let $R_\phi(x, y)$ denote the unique rectangle in $\mathcal{R}_\phi$ that contains (x, y).

2. A mapping $b_\phi : \mathcal{R}_\phi \to \{0, 1\}^*$ where $b_\phi(R)$ for $R \in \mathcal{R}_\phi$ is the unique codeword that ϕ associates with all inputs in R. That is, for all $(x, y) \in \{0, .., n-1\} \times \{0, .., n-1\}$, $b_\phi(x, y) = b_\phi(R_\phi(x, y))$. The mapping b_ϕ must satisfy several requirements. Most germane to our purposes is that $\{b_\phi(R) : R \in \mathcal{R}_\phi\}$ is prefix free with cardinality $|\mathcal{R}_\phi|$. The length of $b_\phi(R)$ is denoted by $l_\phi(R)$.

3. A mapping $v_\phi : \mathcal{R}_\phi \to \{0, 1\}$ where $v_\phi(R)$, for $R \in \mathcal{R}_\phi$, is the unique value that ϕ assigns to all the inputs (x, y) in R. That is, $v_\phi(x, y) = v_\phi(R_\phi(x, y))$ for all $(x, y) \in \{0, .., n-1\} \times \{0, .., n-1\}$.

Before proceeding to prove the bounds, we need a lemma from information theory. A probability distribution is a set of non-negative reals that add to one. To allow multiple equal probabilities, we regard probability distributions as multisets, but for lack of alternative notation we denote them, as sequences, by angled brackets. The *entropy* of a probability distribution $\langle p(x) : x \in S \rangle$ is defined as

$$H \langle p(x) : x \in S \rangle \stackrel{\text{def}}{=} \sum_{x \in S} p(x) \log \frac{1}{p(x)}.$$

The next information-theoretic lemma says that the number of bits needed to describe the outcome of a random variable is at least the entropy of the underlying probability distribution.

Lemma 1 [8] Let $\langle p_j : j = 1, \ldots, J\rangle$ be a probability distribution and $\langle l_j : j = 1, \ldots, J\rangle$ be the lengths of codewords in a prefix-free code[3]. Then,

$$\sum_{j=1}^{J} p_j l_j \geq H\left\langle p_j : j = 1, \ldots, J\right\rangle \ . \quad \Box$$

2.2 Lower bound on $\hat{C}_R(f, 0)$

In this section we prove that for most functions, however sparse, Las Vegas is not much better than determinism when worst-case complexity is the measure. Slightly more precisely, for all $s \geq n$, most functions in $\mathcal{F}_s$ have $\hat{C}_R(f, 0) \geq \log n - 4$.

As mentioned in section 1.2, every randomized protocol can be regarded as a collection of deterministic protocols; the randomness is confined to the choice of the protocol. Given any error-free deterministic protocol ϕ and any input (x, y), there exists an error-free deterministic protocol ϕ' that exchanges only two bits when the input is (x, y) and at most two bits more than ϕ for every other input. It is therefore not a priori clear that there must be an input with a large expected communication ($\bar{L}_\Phi(x, y)$). We prove the lower bound by demonstrating a *set* S of inputs over which every deterministic protocol must exchange a large number of bits *on the average*.

We define the average complexity of a protocol Φ over a set $S \subseteq \{0, .., n-1\} \times \{0, .., n-1\}$ to be:

$$\bar{L}_\Phi^S \stackrel{\text{def}}{=} \frac{1}{|S|} \sum_{(x,y) \in S} \bar{L}_\Phi(x, y).$$

The definition applies of course also to deterministic protocols. If ϕ is deterministic, $\mathcal{R}_\phi$ (defined in section 2.1) induces a partition $S_1, \ldots, S_J$ of S such that each S_j is the intersection

[3] A set of strings is traditionally called *a codebook* or a *code*. Its elements are called codewords.

of S with some rectangle $R_j \in \mathcal{R}_\phi$. All inputs in each S_j have the same codeword — $b_\phi(R_j)$ — whose length we denote by l_j. The next lemma relates $\bar{L}_\phi^S$ to the sizes of the S_j's.

Lemma 2 If S is partitioned by $\mathcal{R}_\phi$ into $S_1, \ldots, S_J$ then

$$\bar{L}_\phi^S \geq H \langle |S_j|/|S| : j = 1, \ldots, J \rangle \ .$$

Proof

$$\begin{aligned} \bar{L}_\phi^S &\stackrel{\text{def}}{=} \sum_{(x,y)\in S} \frac{1}{|S|} l_\phi(x,y) \\ &= \sum_{R\in\mathcal{R}_\phi} \sum_{(x,y)\in R\cap S} \frac{1}{|S|} l_\phi(x,y) \\ &= \sum_{j=1}^{J} \sum_{(x,y)\in S_j} \frac{1}{|S|} l_j \\ &= \sum_{j=1}^{J} l_j \frac{|S_j|}{|S|} \end{aligned}$$

But, from Lemma 2, $l_1, \ldots, l_J$ are lengths of codewords in a prefix free code. Hence, by Lemma 1

$$\sum_{j=1}^{J} l_j \frac{|S_j|}{|S|} \geq H \langle |S_j|/|S| : j = 1, \ldots, J \rangle \ . \quad \square$$

The rest of the proof applies this lemma to most functions in $\mathcal{F}_s$, for every $n \leq s \leq n^2/2$. We divide the range $n \leq s \leq n^2/2$ into three intervals: $n \leq s < 2n \log n$, $2n \log n \leq s < n^2/8$, and $n^2/8 \leq s \leq n^2/2$. The larger s, the more likely are functions in $\mathcal{F}_s$ to have high complexity. Therefore, the most interesting interval to prove is $n \leq s \leq 2n \log n$. We only prove this interval here and mention the stronger results pertaining to the other intervals of s in Corollary 4.

The closure of a set T was defined in section 2.1 to be:

$\mathcal{C}\ell(T) \stackrel{\text{def}}{=} \{(x,y) :$ for some u and v in $\{0,..,n-1\}$, both (x,u) and (v,y) are in $T\}$.

If $S \subseteq \{0,..,n-1\} \times \{0,..,n-1\}$, we let λ_f^S be the size of the largest subset T of S such that f is constant over $\mathcal{C}\ell(T)$:

$$\lambda_f^S \stackrel{\text{def}}{=} \max\{|T| : T \subseteq S, \text{ and } (x_1, y_1), (x_2, y_2) \in C\ell(T) \Rightarrow f(x_1, y_1) = f(x_2, y_2)\} .$$

(note that $\mathcal{C}\ell(T)$ does not have to be contained in S). The next theorem uses λ_f^S to upper bound the size of each S_j; then, it uses Lemma 2 to lower bound $\hat{C}_R(f, 0)$.

Theorem 2 $\hat{C}_R(f, 0) \geq \log \frac{|S|}{\lambda_f^S}$ for all $S \subseteq \{0,..,n-1\} \times \{0,..,n-1\}$.

Proof Let S be any set in $\{0,..,n-1\} \times \{0,..,n-1\}$, Φ be any error-free randomized protocol, and ϕ — any deterministic protocol $\in \Phi$ that occurs with non-zero probability.

Consider the partition $S_1, \ldots, S_J$ of S induced by $\mathcal{R}_\phi$. By definition, each S_j is contained in one rectangle of $\mathcal{R}_\phi$, hence, $\mathcal{C}\ell(S_j)$ is also contained in that rectangle. From Corollary 3, $v_\phi(x, y)$ must be constant over $\mathcal{C}\ell(S_j)$. But, since Φ is error-free and ϕ occurs with positive probability, ϕ must be error-free too. Hence, $f(x, y)$ must be constant over $\mathcal{C}\ell(S_j)$. This implies that $|S_j| \leq \lambda_f^S$. A simple calculation can show that if $|S_j| \leq \lambda_f^S$ for all j, then $H\langle |S_j|/|S| : j = 1, \ldots, J\rangle \geq \log \frac{|S|}{\lambda_f^S}$. Therefore, by the last lemma, $\bar{L}_\phi^S \geq \log \frac{|S|}{\lambda_f^S}$. This is true for all positive-probability protocols $\phi \in \Phi$ therefore, $\bar{L}_\Phi^S \geq \log \frac{|S|}{\lambda_f^S}$. By definition of $\bar{L}_\Phi^S$ as an average, there must be an input $(x, y) \in S$ such that $\bar{L}_\Phi(x, y) \geq \log \frac{|S|}{\lambda_f^S}$. Therefore, for every error-free protocol Φ, $\hat{L}_\Phi \geq \log \frac{|S|}{\lambda_f^S}$ and the theorem follows. $\square$

The next lemma uses a counting argument to show that for s close to n, most functions in $\mathcal{F}_s$ contain a large set S with $\lambda_f^S \leq 2$.

Lemma 3 For all $s \geq n \geq 8$, a fraction larger than $1 - 2\frac{s^9}{n^{12}}$ of the functions in $\mathcal{F}_s$ contain a set $S \subseteq \mathbf{1}_f$ of size $n/8$ such that for all subsets T of S, $|T| \geq 3$ implies that $\mathcal{C}\ell(T) \not\subseteq \mathbf{1}_f$.

Proof We define a *square* to be a product set of the form $A_{\mathcal{X}} \times A_{\mathcal{Y}}$ where $A_{\mathcal{X}}, A_{\mathcal{Y}} \subseteq \{0,..,n-1\}$ are of equal size. We call $|A_{\mathcal{X}}|$ (equivalently, $|A_{\mathcal{Y}}|$) — the *side* of the square. The proof proceeds with two claims:

Claim 1: The fraction of functions f in $\mathcal{F}_s$, for which $\mathbf{1}_f$ contains a square of side 3, is at most $\frac{s^9}{n^{12}}$.

Proof of Claim 1: Fix a 3×3 square in $\{0,..,n-1\} \times \{0,..,n-1\}$. The number of sets of size s (thought of as $\mathbf{1}_f$) that contain this square, is $\binom{n^2-9}{s-9}$. Hence, the number of sets of size s that contain any 3×3 square is at most $\binom{n}{3}\binom{n}{3}\binom{n^2-9}{s-9}$. The fraction of functions in $\mathcal{F}_s$ such that $\mathbf{1}_f$ contains a 3×3 rectangle is, therefore, at most $\binom{n}{3}\binom{n}{3}\binom{n^2-9}{s-9}/\binom{n^2}{s} \le s^9/n^{12}$.

Claim 2: For all $s \ge n$, the fraction of functions f in $\mathcal{F}_s$ such that $\mathbf{0}_f$ contains a square of side $\ge 7n/8$ is at most $1/2^n$.

Proof of Claim 2: As before, fix a square of side $7n/8$. The number of sets of size $n^2 - s$ (thought of as $\mathbf{0}_f$) that contain this square is $\binom{n^2-(\frac{7}{8}n)^2}{s}$. Therefore, the number of sets of size n^2-s that contain any square of side $7n/8$ is at most $\binom{n^2-(\frac{7}{8}n)^2}{s}\binom{n}{\frac{7}{8}n}^2$. Hence, the fraction of functions in $\mathcal{F}_s$ such that $\mathbf{0}_f$ contains a square of side $\ge 7n/8$ is at most,

$$\frac{\binom{n^2-(\frac{7}{8}n)^2}{s}\binom{n}{\frac{7}{8}n}^2}{\binom{n^2}{s}} \le \left(\frac{n^2-(\frac{7}{8}n)^2}{n^2}\right)^s \binom{n}{\frac{7}{8}n}^2 \le \left(\frac{15}{64}\right)^s 1.458^{2n} < \frac{1}{2^n}\ .$$

Combined, the two claims imply that for a fraction of at least $1 - .5^n - \frac{s^9}{n^{12}} \ge 1 - 2\frac{s^9}{n^{12}}$ of the functions f in $\mathcal{F}_s$, the set $\mathbf{0}_f$ doesn't contain a square of side $7n/8$, and the set $\mathbf{1}_f$ doesn't contain a square of side 3.

To complete the proof, we need one more definition. Call a set S *diagonal* if no two of its elements are in the same row or column. (That is, $(x, y), (x', y') \in S$ implies $x \neq x'$ and $y \neq y'$). It is easy to see that if $\mathbf{0}_f$ does not contain a square of side k $(1 \leq k \leq n)$, then $\mathbf{1}_f$ contains a diagonal set of size $n - k + 1$.

We can therefore deduce that for a fraction of at least $1 - 2\frac{s^9}{n^{12}}$ of the functions f in $\mathcal{F}_s$, the set $\mathbf{1}_f$ contains a diagonal set D_f of size $n/8$ but doesn't contain a square of side 3. D_f is clearly the desired set because if $T \subseteq D_f$ and $|T| \geq 3$ then $\mathcal{C}\ell(T)$ contains a square of side 3. This implies that $\mathcal{C}\ell(T) \not\subseteq \mathbf{1}_f$. $\square$

The claimed bounds can now be easily proved:

Corollary 4 For any $s \geq n$, a fraction of at least $1 - 2\frac{(2\log n)^9}{n^3}$ of the functions in $\mathcal{F}_s$ have $\hat{C}_R(f, 0) \geq \log n - 4$.

Proof If $n \leq 16$ the claim is trivially true. Otherwise, as mentioned before, we divide the proof into three intervals of s: $n \leq s < 2n \log n$, $2n \log n \leq s < n^2/8$, and $n^2/8 \leq s \leq n^2/2$. We prove the result for the first interval and state the (stronger) results that can be proved for the others.

<u>$n \leq s < 2n \log n$</u>

A fraction of at least $1 - 2\frac{(2\log n)^9}{n^3}$ of the functions in $\mathcal{F}_s$ have $\hat{C}_R(f, 0) \geq \log n - 4$.

Proof: According to Lemma 3, a fraction of at least $1 - 2\frac{(2\log n)^9}{n^3}$ of the functions in $\mathcal{F}_s$ contain a set S of size $n/8$ with $\lambda_f^S \leq 2$. Theorem 2 ensures that for all these functions, $\hat{C}_R(f, 0) \geq \log \frac{n/8}{2} = \log n - 4$.

<u>$2n \log n < s \leq n^2/8$</u>

A fraction of at least $1 - 2/n^3$ of the functions in $\mathcal{F}_s$ have $\hat{C}_R(f, 0) \geq \log n - 2$.

$n^2/8 \le s \le n^2/2$

A fraction of at least $1 - 0.365^n \cdot n^8$ of the functions in $\mathcal{F}_s$ have $\hat{C}_R(f, 0) \ge \log n - 2$.

$\Box$

Example 1 Let $d_H^N(x, y)$ denote the Hamming distance between the two N-bit sequences x and y. For $N \ge 0$ and $0 \le k \le N$, define the function $H_k^N : \{0,1\}^N \times \{0,1\}^N \to \{0,1\}$ by

$$H_k^N(x, y) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } d_H^N(x, y) = k \\ 0 & \text{otherwise.} \end{cases}$$

To determine the worst-case Las Vegas complexity of H_k^N, consider $\mathbf{1}_{H_k^N}$ (the set of inputs for which H_k^N is 1). Its size is $2^N \binom{N}{k}$. Yet, [9] showed that for $0 \le k \le \lceil N/2 - \sqrt{N/4} \rceil$, the largest rectangle contained in $\mathbf{1}_{H_k^N}$ is of size $\binom{N}{k}$. Taking S of Theorem 2 to be the union of $\mathbf{1}_{H_k^N}$ and any $\binom{n}{k}$ element subset of $\mathbf{0}_{H_k^N}$ we obtain: $\hat{C}_R(H_k^N, 0) \ge \lceil \log(2^N + 1) \rceil = N + 1$. This clearly implies that for all k's in the above range, $\hat{C}_R(H_k^N, 0) = N + 1$.

When n is a power of 2, the equality function is equivalent to H_0^N. Therefore, $\hat{C}_R(equ, 0) = N + 1 = \log n + 1$. $\Box$

2.3 Upper bound on $\bar{C}_D(f, 0)$ and $\hat{C}_R(f, \hat{\epsilon})$

The upper bound proved is stronger than described in section 1.3. It bounds the complexity in terms of the distribution of the 1's and 0's (rather than just their number). For "non-regular" distributions it is even lower than $\log(\frac{|\mathbf{1}_f|}{n})$.

For notational convenience we define log 0 to be -1 and define $G(f)$ as follows:

$$G_{Y/X}(f) \stackrel{\text{def}}{=} 1 + \frac{1}{n} \min\Big(\sum_{x=0}^{n-1} \log |\{y : f(x,y) = 1\}|, \sum_{x=0}^{n-1} \log |\{y : f(x,y) = 0\}| \Big)$$

$$G_{X/Y}(f) \stackrel{\text{def}}{=} 1 + \frac{1}{n} \min\Big(\sum_{y=0}^{n-1} \log |\{x : f(x,y) = 1\}|, \sum_{y=0}^{n-1} \log |\{x : f(x,y) = 0\}| \Big)$$

$$G(f) \stackrel{\text{def}}{=} \min\Big(G_{Y/X}(f), G_{X/Y}(f) \Big).$$

We show that for *all* functions, $\bar{C}_R(f,\hat{\epsilon}) \leq (1-2\epsilon)(G(f)+6.3\log\log n + 2\log\frac{1}{\epsilon}+c)$ and use this result to establish the upper bounds claimed in the introduction.

One interpretation of this bound is that each row (column) in the function table contributes about log of the number of 1's (0's) it contains to the complexity of the function. If a function has about the same number of 1's in each row then a protocol that requires about twice that many bits in the worst case can be obtained by generalizing and refining the equality-function protocol of [10] (see [2]).

Lemma 4 Let $n \geq 2$, $m \geq 1$, and $f : \{0,..,n-1\}\times\{0,..,n-1\} \to \{0,1\}$ satisfy: $|\{y : f(x,y)=1\}| \leq m$ for all $x \in \{0,..,n-1\}$. Then, for all $0 < \epsilon \leq 1/2$,

$$\hat{C}_R(f,\hat{\epsilon}) \leq 2(\log m + \log\ln n + \log\tfrac{1}{\epsilon} + 1)\ .$$

Proof The trivial protocol Φ that guesses with probability 2ϵ and transmits all bits otherwise, has $\hat{E}_\Phi(f) \leq \epsilon$ and $\hat{L}_\Phi \leq 1+2\epsilon+(1-2\epsilon)(\log n+2) = 3-2\epsilon+(1-2\epsilon)\log n$. For $n \leq 60{,}000$, this is less than the upper bound for all ϵ. Thus, from here on, we assume that $n > 60{,}000$.

For every $x \in \{0,..,n-1\}$, let $m(x) \stackrel{\text{def}}{=} |\{y : f(x,y)=1\}|$ and for $i=1,\ldots,m(x)$, let $y_i(x)$ be the ith smallest integer in $\{y : f(x,y)=1\}$. Clearly, for all $x \in \{0,..,n-1\}$, $m(x) \leq m$ and $\{y_1(x),\ldots,y_{m(x)}(x)\} = \{y : f(x,y)=1\}$.

Let $\sigma \stackrel{\text{def}}{=} \sqrt{2}\frac{m\ln n}{\epsilon}$. Then, $n > 60{,}000$ implies $\sigma > 31.1$ so the number of primes between σ and 2σ is at least $\frac{\sigma}{\sqrt{2}\ln\sigma}$ and at most $\frac{\sigma}{\ln\sigma}$. That is, $\frac{\sigma}{\sqrt{2}\ln\sigma} < \Pi(2\sigma)-\Pi(\sigma) < \frac{\sigma}{\ln\sigma}$ where $\Pi(x)$ is the number of primes not exceeding x. The number of primes between σ and 2σ that divide any positive integer $\leq n$ is always at most $\log_\sigma n$.

For brevity, let $(x)_\alpha$ denote x mod α. We show that the following protocol, denoted Φ_1, achieves the claimed bound.

1. $P_{\mathcal{Y}}$ picks at random a prime α such that $\sigma < \alpha < 2\sigma$. He transmits α and $(y)_\alpha$ to $P_{\mathcal{X}}$.

2. $P_{\mathcal{X}}$ computes and transmits
$$g(x,\alpha,(y)_\alpha) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } (y)_\alpha = (y_i(x))_\alpha \text{ for some } 1 \le i \le m(x) \\ 0 & \text{otherwise.} \end{cases}$$

3. Both $P_{\mathcal{X}}$ and $P_{\mathcal{Y}}$ accept $g(x,\alpha,(y)_\alpha)$ as the value of $f(x,y)$.

Now, if $f(x,y) = 1$ then $y = y_i(x)$ for some $1 \le i \le m(x)$, thus $g(x,\alpha,(y)_\alpha) = 1$. Hence $P\left(V_{\Phi_1}(x,y) = 0 | (x,y) \in \mathbf{1}_f\right) = 0$. On the other hand,

$$\begin{aligned} &P\Big(V_{\Phi_1}(x,y) = 1 | (x,y) \in \mathbf{0}_f\Big) \\ &\quad = P\Big(\exists i \in \{1,\ldots,m(x)\} \text{ such that } (y_i(x))_\alpha = (y)_\alpha \;\Big|\; y \ne y_j(x) \text{ for all } j \in \{1,\ldots,m(x)\}\Big) \\ &\quad \le \sum_{i=1}^{m(x)} P\Big(y - y_i(x) \equiv 0 \pmod{\alpha} | y - y_i(x) \ne 0\Big) \\ &\quad = \sum_{i=1}^{m(x)} \frac{\text{number of primes between } \sigma \text{ and } 2\sigma \text{ dividing the } \textit{nonzero} \text{ integer } (y - y_i(x))}{\text{number of primes between } \sigma \text{ and } 2\sigma} \\ &\quad \le \sum_{i=1}^{m(x)} \log_\sigma n / \tfrac{\sigma}{\sqrt{2}\ln\sigma} \\ &\quad \le \sqrt{2} m \tfrac{\ln n}{\sigma} = \epsilon\ . \end{aligned}$$

The average number of bits transmitted is upper bounded by:

Transmitting α : $\log\frac{\sigma}{\ln\sigma} \le \log m + \log\ln n + \log\frac{1}{\epsilon} + .5 - \log\ln 31.1$ bits.

Transmitting y mod α: $\log 1.7\sigma \le \log m + \log\ln n + \log\frac{1}{\epsilon} + .5 + .77$ bits.

Rounding off the above: 1 bit.

Transmitting $g(x,\alpha,y \bmod \alpha)$: 1 bit.

Hence the total communication is at most $2(\log m + \log\ln n + \log\frac{1}{\epsilon} + 1)$ bits. $\Box$

In the next example, we use the lemma to improve the bound on a function defined in [5].

Example 2 Let $M \geq 0$, and $x^1, \ldots, x^M, y^1, \ldots, y^M \in \{0,1\}^M$ (we index the sequences with superscripts because subscripts were used earlier to denote a bit in a sequence). If $n = 2^{M^2}$, then each $(x^1, \ldots, x^M)$ can be identified with an integer in $\{0,..,n-1\}$. The *component equality* function is defined as:

$$c.e.\Big((x^1, \ldots, x^M), (y^1, \ldots, y^M)\Big) \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } x^i = y^i \text{ for some } 1 \leq i \leq M \\ 0 & \text{otherwise.} \end{cases}$$

It was shown in [5] that $\hat{C}_R(f, 0) \leq O(N \log^2 N)$. By choosing $\epsilon = \frac{\log N}{N}$ in the last lemma, and proceeding as in the original proof, we get, $\hat{C}_R(f, 0) \leq 5N \log N + c$. $\Box$

However, to meet the lower bound, Lemma 4 needs to be improved even for functions with equally dense rows. The reason being that the protocol uses about as many bits describing the results of the random experiments as it does describing the values. One way around this, is using less randomization. The following lemma, proved in [4], shows that in the initial phases of the communication, deterministic protocols can be efficient.

Lemma 5 Let $n \geq 2$, $m \geq 1$, and $f : \{0,..,n-1\} \times \{0,..,n-1\} \to \{0,1\}$ satisfy:
$|\{y : f(x,y) = 1\}| \leq m$ for all $x \in \{0,..,n-1\}$.
Then, there exists a partition $B_1, \ldots, B_{\lceil cm/(\ln n)^{1.1} \rceil}$ of $\{0, \ldots, n-1\}$ such that:
for $x = 0, \ldots, n-1$ and $j = 1, \ldots, \lceil cm/(\ln n)^{1.1} \rceil$, $\quad |\{y \in B_j : f(x,y) = 1\}| < (\ln n)^{1.1}$
where c is a constant independent of n, m, and f. $\Box$

Combined, the two lemmas yield:

Lemma 6 For all $n \geq 2$, $\quad 0 < \epsilon \leq 1/2$, and $\quad f : \{0,..,n-1\} \times \{0,..,n-1\} \to \{0,1\}$,

$$\bar{C}_R(f, \hat{\epsilon}) \leq G(f) + 4.1 \log \log n + 2 \log \frac{1}{\epsilon} + c$$

where c is a constant independent of n, ϵ, and f.

Proof Without loss of generality assume that $G(f) = G_{Y/X}(f) = \frac{1}{n}\sum_{x=0}^{n-1} \log|\{y : f(x,y) = 1\}|$.

For $i = -1, 0, 1, \ldots, \lceil \log n \rceil$ define $A_i \stackrel{\text{def}}{=} \{x : i-1 < \log m(x) \leq i\}$. Clearly, the sets $A_{-1}, A_0, A_1, \ldots, A_{\lceil \log n \rceil}$ partition $\{0, .., n-1\}$.

By Lemma 5, for $i = \lceil \log(\ln n)^{1.1} \rceil, \ldots, \lceil \log n \rceil$, there exists a partition $\{B_{i,j} : j = 1, \ldots, \lceil c\frac{2^i}{(\ln n)^{1.1}} \rceil\}$ of $\{0, \ldots, n-1\}$ such that for all $i \in A_i$ and all $j \in \{1, \ldots, \lceil c\frac{2^i}{(\ln n)^{1.1}} \rceil\}$, $|\{y \in B_{i,j} : f(x,y) = 1\}| < (\ln n)^{1.1}$ where c is a constant.

$P_{\mathcal{X}}$ and $P_{\mathcal{Y}}$ agree on such a collection of partitions and conduct a protocol Φ_2 defined as follows:

1. $P_{\mathcal{X}}$ transmits $\lceil \log m(x) \rceil$ to $P_{\mathcal{Y}}$ (thereby telling him of the index of the set A_i that x is in).

2. If $\lceil \log m(x) \rceil = -1$, they decide that the value of the function is 0 and stop.

   If $-1 < \lceil \log m(x) \rceil < \lceil \log(\ln n)^{1.1} \rceil$, they move to step 3.

   Otherwise, $P_{\mathcal{Y}}$ transmits the index j of the set $B_{\lceil \log m(x) \rceil, j}$ that contains y.

   From now on, they know that (x, y) is in the generalized rectangle $A_{\lceil \log m(x) \rceil} \times B_{\lceil \log m(x) \rceil, j}$. This rectangle has the property that each of its rows contains at most $(\ln n)^{1.1}$ elements of $\mathbf{1}_f$. That is, $|\{y \in B_{\lceil \log m(x) \rceil, j} : f(x,y) = 1\}| < (\ln n)^{1.1}$ for all x in $A_{\lceil \log m(x) \rceil}$.

3. They use the protocol Φ_1 of Lemma 4 to find $f(x, y)$ over this rectangle with probability of error $\leq \epsilon$.

The number of bits transmitted is upper bounded by:

Step 1: $\lceil \log\log n \rceil$.

Step 2: 0 bits if $m(x) < (\ln n)^{1.1}$ and $\lceil \log \lceil c\frac{2^{\lceil \log m(x) \rceil}}{(\ln n)^{1.1}} \rceil \rceil \leq \lceil \log \frac{c \cdot 2m(x)}{(\ln n)^{1.1}} \rceil$ otherwise.

Step 3: 0 bits if $m(x) = 0$, $2(\log m(x) + \log\ln n + \log\frac{1}{\epsilon} + 1)$ if $0 < m(x) < (\ln n)^{1.1}$, and

$2(\log(\ln n)^{1.1} + \log\ln n + \log\frac{1}{\epsilon} + 1)$ otherwise.

Hence for every (x, y), the total number of bits transmitted is at most $\log m(x) + 2\log\frac{1}{\epsilon} + 4.1\log\ln n + c$ (where c is a new constant). The average, therefore, is $\leq \frac{1}{n}\sum_{x=0}^{n-1} \log m(x) + 2\log\frac{1}{\epsilon} + 4.1\log\ln n + c = G(f) + 2\log\frac{1}{\epsilon} + 4.1\log\ln n + c$. $\square$

Decreasing ϵ in the lemma increases the complexity only moderately, so if ϵ is large, some bits can be saved by incurring an error smaller than ϵ most of the time and just guessing the result in the rest, yielding:

Theorem 3 For all $n \geq 2$, $0 < \epsilon \leq \frac{1}{2}$, and all $f : \{0, .., n-1\} \times \{0, .., n-1\} \to \{0, 1\}$,

$$\bar{C}_R(f, \hat{\epsilon}) \leq 2 + (1 - 2\epsilon)\Big(G(f) + 6.3\log\log n + 2\log\frac{1}{\epsilon} + c\Big)$$

where c is a constant independent of n, ϵ and f.

Proof Let $\delta \stackrel{\text{def}}{=} \min(\epsilon, \frac{1}{G(f)}, \frac{1}{101})$ and $\epsilon' \stackrel{\text{def}}{=} \frac{\epsilon - \delta}{1/2 - \delta}$. The protocol Φ_3 is defined as follows:

1. $P_{\mathcal{Y}}$ performs a random experiment that is 1 with probability ϵ' and transmits the outcome to $P_{\mathcal{X}}$.

2. If the outcome of the experiment was 1, $P_{\mathcal{Y}}$ performs an unbiased binary experiment and transmits the result to $P_{\mathcal{X}}$. They both accept the outcome of this experiment as the value of $f(x, y)$. If the outcome of the experiment in the first phase was 0 (probability $1 - \epsilon'$), they use protocol Φ_2 (Lemma 6) to compute $f(x, y)$ with probability of error $\leq \delta$.

The total probability of error is $\leq \epsilon'/2 + (1-\epsilon')\delta = \epsilon'(\frac{1}{2} - \delta) + \delta = \epsilon$. The average communication length is

$$\begin{aligned} \bar{L}_{\Phi_3} &= 1 + \epsilon' + (1-\epsilon')\Big(G(f) + 4.1 \log\log n + 2\log\tfrac{1}{\delta} + c\Big) \\ &\leq 2 + (1-2\epsilon)\Big(G(f) + 6.3\log\log n + 2\log\tfrac{1}{\epsilon} + c\Big)\ . \quad \Box \end{aligned}$$

By the convexity of the log function, $G(f) \leq \log \frac{|\mathbf{1}_f|}{n}$. We have therefore proved the upper bound for the average case:

Corollary 5 For all $n \geq 2$, $\quad 0 < \epsilon \leq \frac{1}{2}$, and all functions $f : \{0,..,n-1\} \times \{0,..,n-1\} \rightarrow \{0,1\}$,

$$\bar{C}_R(f,\hat{\epsilon}) \leq 2 + (1-2\epsilon)\left(\log \tfrac{|\mathbf{1}_f|}{n} + 6.3\log\log n + 2\log\tfrac{1}{\epsilon} + c\right)\ . \quad \Box$$

For all $s \geq n$, most functions in $\mathcal{F}_s$ have about the same number of 1's in every row and column. A quick analysis of Lemma 6 can show that the expected number of bits transmitted for every input is about the same (and, that step 1 can be skipped). Hence

Corollary 6 Let $s \geq n \geq 2$ and $0 < \epsilon \leq \frac{1}{2}$. For most functions f in $\mathcal{F}_s$,

$$\hat{C}_R(f,\hat{\epsilon}) \leq 2 + (1-2\epsilon)\Big(\log \tfrac{|\mathbf{1}_f|}{n} + 5.3\log\log n + 2\log\tfrac{1}{\epsilon} + c\Big)\ . \quad \Box$$

The protocol of Theorem 3, can incur an error only when its computed value is 1. This happens with probability $\leq \epsilon + \frac{|\mathbf{1}_f|}{n^2}$. Verifying the result whenever that is the case, we obtain a randomized Las Vegas protocol with $\bar{L}_\Phi \leq \log\frac{|\mathbf{1}_f|}{n} + 8.3\log\log n + c$. This, in turn, implies the existence of a deterministic protocol with the same average length. Thus,

Corollary 7 For *all* functions $f : \{0,..,n-1\} \times \{0,..,n-1\} \rightarrow \{0,1\}$,

$$\bar{C}_D(f,0) \leq \log \tfrac{|\mathbf{1}_f|}{n} + 8.3\log\log n + c\ .$$

Proof If $\frac{|\mathbf{1}_f|}{n^2} > \frac{1}{(\log n)^{8.3}}$, then $\log\frac{|\mathbf{1}_f|}{n} + 8.3\log\log n + c \geq \log n + c$. For $c \geq 2$, this is more than the complexity of the trivial protocol. Else, let $\epsilon \stackrel{\text{def}}{=} \frac{1}{\log n}$. Corollary 5 says that $\bar{C}_R(f,\hat{\epsilon}) \leq \log\frac{|\mathbf{1}_f|}{n} + 8.3\log\log n + c$. The protocol used to derive this bound (Φ_3) errs only when its computed value is 1. Assuming a uniform probability distribution over the inputs, this happens with probability $\leq \epsilon + \frac{|\mathbf{1}_f|}{n^2} \leq \frac{2}{\log n}$. By verifying the result whenever this is the case, we add at most 2 bits on the average. Therefore, $\bar{C}_D(f,0) = \bar{C}_R(f,0) \leq \log\frac{|\mathbf{1}_f|}{n} + 8.3\log\log n + c$. □

Example 3 Consider H_k^N again. Example 1 showed that for $0 \leq k \leq \lceil N/2 - \sqrt{N/4}\rceil$, $\hat{C}_R(H_k^N,0) \geq \log N$. The following protocol, whose main advantage is amenability to simple analysis, shows that the average complexity of H_k^N is much smaller.

First, $P_{\mathcal{X}}$ transmits $x_1,\ldots,x_k$. Then, starting with $i = k+1$ and continuing with consecutive i's, $P_{\mathcal{X}}$ transmits x_i and $P_{\mathcal{Y}}$ responds by transmitting 1 if the Hamming distance between $(x_1,\ldots,x_i)$ and $(y_1,\ldots,y_i)$ is $k+1$, and transmitting 0 otherwise. To the Nth bit, $P_{\mathcal{Y}}$ responds by transmitting the value of H_k^N. They stop after $P_{\mathcal{Y}}$ transmits a 1 or after he responds to x_N. Let $i_j \stackrel{\text{def}}{=} \min\Big\{\{i : d_H\big((x_1,\ldots,x_i),(y_1,\ldots,y_i)\big) = j\} \cup \{N\}\Big\}$ and let I_j be the random variable whose value is i_j. If L is the random variable denoting the length of the communication then, $L = k + 2(I_{k+1} - k) = 2I_{k+1} - k$ so the expected value of L is $\bar{L} = 2\bar{I}_{k+1} - k$.

The expected value of I_{k+1} can be determined via a tedious calculation or using an old trick. For $j = 1,\ldots,k+1$, let $D_j \stackrel{\text{def}}{=} I_j - I_{j-1}$ (I_0 is defined to be 0). Then each D_j is distributed geometrically with parameter $\frac{1}{2}$. Therefore, the expected value of each D_j is 2, and the expected value of I_{k+1} is $E(I_{k+1}) = E(\sum_{j=1}^{k+1} D_j) = \sum_{j=1}^{k+1} 2 = 2(k+1)$. Hence,

$\bar{L} = 3k + 4$.

Note that this protocol achieves a much smaller complexity than that promised by Corollary 7. The next section, however, shows that this is a rare exception. For most functions, the upper bound of the corollary is quite tight. The reason for the discrepancy here is that $\mathbf{0}_{H_k^N}$ contains unusually large generalized rectangles. $\square$

2.4 Lower bound on $\bar{C}_R(f, \hat{\epsilon})$

In section 2.2 we defined the average complexity of a protocol over a set $S \subseteq \{0, .., n-1\} \times \{0, .., n-1\}$. We now define the *average error of* Φ *over* S to be,

$$\bar{E}_\Phi^S(f) \stackrel{\text{def}}{=} \frac{1}{|S|} \sum_S \bar{E}_\Phi(x, y) = \frac{1}{|S|} \sum_S P(V_\Phi(x, y) \neq f(x, y)).$$

The *average error of* Φ, $\bar{E}_\Phi(f)$, is defined to be the arithmetic average of the error over the zeros and the errors over the ones: $\bar{E}_\Phi(f) \stackrel{\text{def}}{=} \frac{1}{2}\left(\bar{E}_\Phi^{\mathbf{0}_f}(f) + \bar{E}_\Phi^{\mathbf{1}_f}(f)\right)$.

Remark: In general, $\bar{E}_\Phi(f)$ can be taken as the maximum of $\bar{E}_\Phi^{\mathbf{0}_f}(f)$ and $\bar{E}_\Phi^{\mathbf{1}_f}(f)$, or any weighted average of the two. Here we chose half-half weights. Other weights, though possible, yield an inferior lower bound.

We can now define average-error complexity in much the same way we did in the introduction. Two measures prove particularly useful. The *average deterministic complexity with* ϵ *average error* is defined as:

$$\bar{C}_D(f, \bar{\epsilon}) \stackrel{\text{def}}{=} \min\{\bar{L}_\phi : \phi \text{ is deterministic and } \bar{E}_\phi(f) \leq \epsilon\},$$

and the *average randomized complexity with* ϵ *average error* is defined as:

$$\bar{C}_R(f, \bar{\epsilon}) \stackrel{\text{def}}{=} \min\{\bar{L}_\Phi : \bar{E}_\Phi(f) \leq \epsilon\} \ .$$

The next lemma shows that although the average error is defined in a peculiar way, ϵ average error is still easier to achieve than ϵ worst error.

Lemma 7 For all functions f and $\epsilon > 0$, $\quad \bar{C}_R(f,\hat{\epsilon}) \geq \bar{C}_R(f,\bar{\epsilon})$.

Proof : $\quad \bar{E}_\Phi^{\mathbf{0}_f}(f) \stackrel{\text{def}}{=} \frac{1}{|\mathbf{0}_f|} \sum_{(x,y)\in\mathbf{0}_f} \bar{E}_\Phi(x,y) \leq \max\{\bar{E}_\Phi(x,y) : (x,y) \in \mathbf{0}_f\}$, and

$$\bar{E}_\Phi^{\mathbf{1}_f}(f) \stackrel{\text{def}}{=} \frac{1}{|\mathbf{1}_f|} \sum_{(x,y)\in\mathbf{1}_f} \bar{E}_\Phi(x,y) \leq \max\{\bar{E}_\Phi(x,y) : (x,y) \in \mathbf{1}_f\}.$$

Therefore,

$$\begin{aligned}
\bar{E}_\Phi(f) &= \tfrac{1}{2}\left(\bar{E}_\Phi^{\mathbf{0}_f}(f) + \bar{E}_\Phi^{\mathbf{1}_f}(f)\right) \\
&\leq \max\left(\bar{E}_\Phi^{\mathbf{0}_f}(f), \bar{E}_\Phi^{\mathbf{1}_f}(f)\right) \\
&\leq \max\left(\max\{\bar{E}_\Phi(x,y) : (x,y) \in \mathbf{0}_f\}, \max\{\bar{E}_\Phi(x,y) : (x,y) \in \mathbf{1}_f\}\right) \\
&= \max\{\bar{E}_\Phi(x,y) : (x,y) \in \{0,..,n-1\}\times\{0,..,n-1\}\} \\
&= \hat{E}_\Phi(f) \ . \quad \square
\end{aligned}$$

Any lower bound for $\bar{C}_R(f,\bar{\epsilon})$ is therefore also a lower bound for $\bar{C}_R(f,\hat{\epsilon})$. It then suffices to prove the result for average error. We begin by proving a lower bound for the deterministic complexity $\bar{C}_D(f,\bar{\epsilon})$.

Each rectangle $R \subseteq \{0,..,n-1\}\times\{0,..,n-1\}$ contains a certain proportion $\rho_f(R)$ of 1's defined as: $\rho_f(R) \stackrel{\text{def}}{=} \frac{|R\cap\mathbf{1}_f|}{|R|} = \frac{|\{(x,y)\in R : f(x,y)=1\}|}{|R|}$ (the number of 1's in the rectangle divided by its size). For deterministic ϕ, define $R_{\phi,f}(\delta) \stackrel{\text{def}}{=} \{(x,y) : (1-\delta)\frac{|\mathbf{1}_f|}{n^2} \leq \rho_f(R_\phi(x,y)) \leq (1+\delta)\frac{|\mathbf{1}_f|}{n^2}\}$ to be the inputs covered by rectangles $R \in \mathcal{R}_\phi$ with $\rho_f(R)$ deviating from $\frac{|\mathbf{1}_f|}{n^2}$ by a factor of at most $1 \pm \delta$.

We assumed in the beginning of section 1.3 that $|\mathbf{1}_f| \leq n^2/2$. Therefore, if $R \in \mathcal{R}_\phi$

satisfies $R \subseteq R_{\phi,f}(\delta)$ then,

$$\begin{aligned} |R \cap \mathbf{0}_f|/|R| &= (|R| - |R \cap \mathbf{1}_f|)/|R| \\ &= 1 - \rho_f(R) \\ &\geq 1 - \tfrac{|\mathbf{1}_f|}{n^2}(1+\delta) \\ &\geq (1-\delta)\tfrac{|\mathbf{0}_f|}{n^2} \ . \end{aligned} \tag{3}$$

A similar calculation can show that $|R \cap \mathbf{0}_f|/|R| \leq (1+\delta)\frac{|\mathbf{0}_f|}{n^2}$.

Intuitively, rectangles with proportion of 1's similar to that of the original function table cannot reduce the uncertainty about the function's value. The next lemma says that if the protocol errs with low probability, then $R_{\phi,f}(\delta)$ must be small.

Lemma 8 For every deterministic protocol ϕ, and all $\delta > 0$,

$$\bar{E}_\phi(f) \geq \frac{1-\delta}{2n^2}|R_{\phi,f}(\delta)| \ .$$

Proof Let $0_\phi \stackrel{\text{def}}{=} R_{\phi,f}(\delta) \bigcap \{(x,y) : V_\phi(x,y) = 0\}$ be the set of inputs in $R_{\phi,f}(\delta)$ that ϕ declares 0, and $1_\phi \stackrel{\text{def}}{=} R_{\phi,f}(\delta) \bigcap \{(x,y) : V_\phi(x,y) = 1\}$ — the set of inputs in $R_{\phi,f}(\delta)$ that ϕ declares 1. Clearly, $0_\phi \cup 1_\phi = R_{\phi,f}(\delta)$, and, since ϕ is deterministic, $\bar{E}^S_\phi(f) = \frac{\left|\{(x,y): V_\phi(x,y) \neq f(x,y)\}\right|}{|S|}$. Therefore,

$$\begin{aligned}
\bar{E}_\phi^{\mathbf{1}_f}(f) &= \tfrac{1}{|\mathbf{1}_f|}\sum_{\mathbf{1}_f}\chi_{\{(x,y):V_\phi(x,y)=0\}}(x,y)\\
&\geq \tfrac{1}{|\mathbf{1}_f|}\sum_{\{R\in\mathcal{R}_\phi:R\subseteq 0_\phi\}}\ \sum_{\mathbf{1}_f\cap R}1\\
&\geq \tfrac{1}{|\mathbf{1}_f|}\sum_{\{R\in\mathcal{R}_\phi:R\subseteq 0_\phi\}}(1-\delta)\tfrac{|\mathbf{1}_f|}{n^2}|R|\\
&= \tfrac{1}{n^2}(1-\delta)|0_\phi|.
\end{aligned}$$

We showed in equation (3) that all rectangles $R\in\mathcal{R}_\phi$ that are contained in $R_{\phi,f}(\delta)$, satisfy $|R\cap\mathbf{0}_f|>|R|(1-\delta)\frac{|\mathbf{0}_f|}{n^2}$. Thus, we also have,

$$\bar{E}_\phi^{\mathbf{0}_f}(f)\geq\frac{1}{n^2}(1-\delta)|1_\phi|.$$

Combining the two,

$$\bar{E}_\phi(f)=\frac{1}{2}\left(\bar{E}_\phi^{\mathbf{0}_f}(f)+\bar{E}_\phi^{\mathbf{1}_f}(f)\right)\geq\frac{1-\delta}{2n^2}(|0_\phi|+|1_\phi|)=\frac{1-\delta}{2n^2}|R_{\phi,f}(\delta)|\ .\quad\square$$

Let $r_f(\delta)\stackrel{\text{def}}{=}\max\{|R|:\left|\rho_f(R)-\frac{|\mathbf{1}_f|}{n^2}\right|>\delta\frac{|\mathbf{1}_f|}{n^2}\}$ be the size of the largest rectangle in $\{0,..,n-1\}\times\{0,..,n-1\}$ with proportion of ones deviating from $\frac{s}{n^2}$ by a factor of more than $1\pm\delta$. ($r_f(\delta)$ is independent of ϕ.) Then, for all deterministic protocols ϕ, if R is a rectangle in $\mathcal{R}_\phi$ and $|R|>r_f(\delta)$, then R must be contained in $R_{\phi,f}(\delta)$. From the last lemma,

$$\sum_{\{R\in\mathcal{R}_\phi:|R|>r_f(\delta)\}}\frac{|R|}{n^2}\leq\frac{|R_{\phi,f}(\delta)|}{n^2}\leq\frac{2\bar{E}_\phi(f)}{1-\delta}\ . \tag{4}$$

The inequality implies that the total area occupied by rectangles of size $>r_f(\delta)$ is small if $\bar{E}_\phi(f)$ is small. (In Lemma 12, we show that for most functions $r_f(\delta)$ is also small.) The next simple lemma says that when this is the case (the larger probabilities occupy a small part of the probability space) the entropy is high.

Lemma 9 If $\langle r_j : j=1,\ldots,J\rangle$ is a probability distribution such that $\sum_{\{j:r_j>r\}} r_j\leq A$, then $H\langle r_j\rangle\geq(1-A)\log(\frac{1}{r})$.

Proof Compute the average number of bits needed to describe the events $r_j \leq r$. $\square$

When Lemma 2 is applied to $S = \{0, .., n–1\} \times \{0, .., n–1\}$, it yields $\bar{L}_\phi \geq H \langle |R|/n^2 : R \in \mathcal{R}_\phi \rangle$. Combining it with equation (4) and Lemma 9, we obtain: for all ϕ and all $\delta > 0$,

$$\bar{L}_\phi \geq H \left\langle |R|/n^2 : R \in \mathcal{R}_\phi \right\rangle \geq \left(1 - \frac{2\bar{E}_\phi(f)}{1-\delta}\right) \log \frac{n^2}{r_f(\delta)} . \tag{5}$$

Next, we need to relate $r_f(\delta)$ to $|\mathbf{1}_f|$. First, we use a basic result from probability theory to reduce the problem of functions in $\mathcal{F}_s$ which behave hypergeometrically to that of Bernoulli random variables which are easier to analyze:

Lemma 10 Let H be distributed hypergeometrically with parameters N, s, and r. That is,

$$P(H = k) = H_{N,s,r}(k) \stackrel{\text{def}}{=} \frac{\binom{r}{k}\binom{N-r}{s-k}}{\binom{N}{s}} = \frac{\binom{s}{k}\binom{N-s}{r-k}}{\binom{N}{r}}$$

and let B be distributed binomially with probability $\frac{s}{N}$ over r variables. That is,

$$P(B = k) = B_{\frac{s}{N},r}(k) \stackrel{\text{def}}{=} \binom{r}{k} \left(\frac{s}{N}\right)^k \left(1 - \frac{s}{N}\right)^{r-k} .$$

Then, $r\frac{s}{N}$ is the expected value of both H and B, and, for $a \leq \lfloor r\frac{s}{N} \rfloor \leq \lceil r\frac{s}{N} \rceil \leq b$,

$$P(a \leq H \leq b) \geq P(a \leq B \leq b) .$$

Proof The lemma follows from a result of Uhlmann [11] showing that for all N, s, and r,

$$P(0 \leq H \leq x) - P(0 \leq B \leq x) \begin{cases} \geq 0 & \text{for } x \geq \frac{s}{N}(r-1)\frac{N+1}{N} \\ \leq 0 & \text{for } x \leq \frac{s}{N}(r-1)\frac{N+1}{N} - \frac{r-1}{N} \end{cases}$$

(the original result was phrased slightly differently). Note that $\frac{s}{N}(r-1)\frac{N+1}{N} < \frac{s}{N}r$, hence

$$P(0 \leq H \leq b) \geq P(0 \leq B \leq b) .$$

Also, $\frac{s}{N}(r-1)\frac{N+1}{N} - \frac{r-1}{N} > \frac{s}{N}r - 1$, hence

$$P(0 \leq H \leq a-1) \leq P(0 \leq B \leq a-1)$$

and the lemma follows. $\square$

To bound the size of "non typical" rectangles, we use Bernstein's lemma:

Lemma 11 [12] Let $X_1, \ldots, X_r$ be independent, Bernoulli-p, random variables, and let $\delta < 1-p$. Then,

$$P\Big(\Big|\frac{\sum_{i=1}^{r} X_i}{r} - p\Big| > \delta p\Big) \leq e^{-\frac{\delta^2 rp}{4.5}} \ . \quad \square$$

Therefore, if f is a random function in $\mathcal{F}_s$, the probability that a given rectangle of size $> r$ will have a proportion of f-ones deviating from $\frac{s}{n^2}$ by a factor of more than $1 \pm \delta$ is $< e^{-\frac{\delta^2 rs}{4.5n^2}}$. The probability that *any* rectangle of size $> r$ will have a proportion of f-ones deviating from $\frac{s}{n^2}$ by more than δ is $< 2^{2n} \cdot e^{-\frac{\delta^2 rs}{4.5n^2}}$. If $r \geq \frac{9n^3}{s \cdot \delta^2}$, this probability tends to zero faster than $(\frac{2}{e})^n$. Therefore, a fraction of at least $1 - (\frac{2}{e})^n$ of the functions in $\mathcal{F}_s$ don't contain any rectangle R of size $\geq \frac{9n^3}{s \cdot \delta^2}$ with a proportion of f-ones deviating from $\frac{s}{n^2}$ by a factor of more than $1 \pm \delta$. We have thus proved:

Lemma 12 For all s and $0 < \delta < 1 - \frac{s}{n^2}$, a fraction of at least $1 - (\frac{2}{e})^n$ of the functions in $\mathcal{F}_s$ have $r_f(\delta) \leq \frac{9n^3}{s \cdot \delta^2}$. $\square$

Substituting in equation (5) we get that for $0 < \delta < 1 - \frac{s}{n^2}$,

$$\bar{L}_\phi \geq \left(1 - \frac{2\bar{E}_\phi(f)}{1-\delta}\right) \log \frac{n^2}{r_f(\delta)} \geq \left(1 - \frac{2\bar{E}_\phi(f)}{1-\delta}\right) \log \frac{s\delta^2}{9n} \ .$$

We want results tight up to an additive term so we can think of δ as being very small (in fact, decreasing to 0). For $\delta < \frac{1}{2}$, $\frac{1}{1-\delta} < 1 + 2\delta$, and, for very small δ's, the difference (always less than δ) is negligible. Therefore, we approximate:

$$\begin{aligned}
\bar{L}_\phi &\geq \left(1-2\bar{E}_\phi(f)\right)\log\tfrac{s\delta^2}{9n} - 4\delta\bar{E}_\phi(f)\log\tfrac{s\delta^2}{9n} \\
&> \left(1-2\bar{E}_\phi(f)\right)\log\tfrac{s\delta^2}{9n} - 4\delta\log\tfrac{s\delta^2}{9n} \\
&> \left(1-2\bar{E}_\phi(f)\right)\log\tfrac{s}{n} + 2\left(1-2\bar{E}_\phi(f)\right)\log\delta - 4\delta\log\tfrac{s}{n} - 4\ .
\end{aligned}$$

Since $s \leq n^2/2$, this inequality is valid for all $\delta < \frac{1}{2}$. The (negative) expression $2(1-2\bar{E}_\phi(f))\log\delta - 4\delta\log\frac{s}{n}$ is maximized for $\delta = \frac{1-2\bar{E}_\phi(f)}{2\ln 2\cdot\log\frac{s}{n}}$. Substituting, we obtain:

$$\bar{L}_\phi > \left(1-2\bar{E}_\phi(f)\right)\left(\log\tfrac{s}{n} - 2\log\log\tfrac{s}{n}\right) - 9\ .$$

This result suffices to bound $\bar{C}_D(f,\bar{\epsilon})$ for most functions. To bound $\bar{C}_R(f,\bar{\epsilon})$, we return to the notations introduced at the beginning of the section. Note that for every randomized protocol Φ,

$$\begin{aligned}
\bar{E}_\Phi^{\mathbf{0}_f}(f) &\stackrel{\text{def}}{=} \tfrac{1}{|\mathbf{0}_f|}\sum_{\mathbf{0}_f}\bar{E}_\Phi(x,y) \\
&\stackrel{\text{def}}{=} \tfrac{1}{|\mathbf{0}_f|}\sum_{\mathbf{0}_f} p\left(V_\Phi(x,y)\neq f(x,y)\right) \\
&= \tfrac{1}{|\mathbf{0}_f|}\sum_{\mathbf{0}_f}\sum_{\phi\in\Phi} p(\phi)\chi_{\{(x,y):V_\phi(x,y)\neq f(x,y)\}}(x,y) \\
&= \sum_{\phi\in\Phi} p(\phi)\tfrac{1}{|\mathbf{0}_f|}\sum_{\mathbf{0}_f}\chi_{\{(x,y):V_\phi(x,y)\neq f(x,y)\}}(x,y) \\
&\stackrel{\text{def}}{=} \sum_{\phi\in\Phi} p(\phi)\bar{E}_\phi^{\mathbf{0}_f}(f).
\end{aligned}$$

Similarly,

$$\bar{E}_\Phi^{\mathbf{1}_f}(f) = \sum_{\phi\in\Phi} p(\phi)\bar{E}_\phi^{\mathbf{1}_f}(f).$$

Therefore,

$$\begin{aligned}
\bar{E}_\Phi(f) &\stackrel{\text{def}}{=} \tfrac{1}{2}\left(\bar{E}_\Phi^{\mathbf{0}_f}(f) + \bar{E}_\Phi^{\mathbf{1}_f}(f)\right) \\
&= \tfrac{1}{2}\Big(\sum_{\phi\in\Phi} p(\phi)\bar{E}_\phi^{\mathbf{0}_f}(f) + \sum_{\phi\in\Phi} p(\phi)\bar{E}_\phi^{\mathbf{1}_f}(f)\Big) \\
&= \sum_{\phi\in\Phi} p(\phi)\tfrac{1}{2}\left(\bar{E}_\phi^{\mathbf{0}_f}(f) + \bar{E}_\phi^{\mathbf{1}_f}(f)\right) \\
&= \sum_{\phi\in\Phi} p(\phi)\bar{E}_\phi(f),
\end{aligned}$$

and

$$\begin{aligned}
\bar{L}_\Phi &\stackrel{\text{def}}{=} \tfrac{1}{n^2} \sum_{(x,y)} \bar{L}_\Phi(x,y) \\
&= \sum_{\phi\in\Phi} p(\phi) \tfrac{1}{n^2} \sum_{(x,y)} \bar{L}_\phi(x,y) \\
&= \sum_{\phi\in\Phi} p(\phi) \bar{L}_\phi \\
&> \sum_{\phi\in\Phi} p(\phi) \left[\left(1 - 2\bar{E}_\phi(f)\right) \log \tfrac{s}{n} - 2 \log\log \tfrac{s}{n} - 9\right] \\
&= \left(1 - 2\bar{E}_\Phi(f)\right) \log \tfrac{s}{n} - 2\log\log \tfrac{s}{n} - 9 \ .
\end{aligned}$$

We have proved:

Corollary 8 For all $0 < \epsilon \leq \frac{1}{2}$, a fraction of at least $1 - (\frac{2}{e})^n$ of the functions in $\mathcal{F}_s$ satisfy:

$$\bar{C}_R(f, \hat{\epsilon}) \geq \bar{C}_R(f, \bar{\epsilon}) > (1 - 2\epsilon) \left(\log \tfrac{s}{n} - 2\log\log \tfrac{s}{n}\right) - 9 \ . \quad \Box$$

This completes the proof of the lower bound. We conclude the paper by applying some of the techniques developed to find a tight lower bound for the equality function.

Example 4 Consider the equality function once more. Example 3 showed that $\bar{C}_D(equ, 0) \leq 4$. Consequently, we can only expect high worst-case lower bounds. In this example, we prove that $\hat{C}_R(equ, \bar{\epsilon}) \geq \lceil \log(\frac{n}{1+2\epsilon(n-1)}) \rceil \approx \log \frac{1}{\epsilon}$. First, we relate $\bar{E}_\phi$ to $\bar{L}_\phi^{\mathbf{1}_{equ}}$ in deterministic protocols.

Let ϕ be any deterministic protocol. Then ϕ induces the partition $\mathcal{R}_\phi$ of $\{0,..,n-1\} \times \{0,..,n-1\}$ described in section 2.1. Let $R_1, \ldots, R_J$ be the rectangles in $\mathcal{R}_\phi$ that have a nonempty intersection with $\mathbf{1}_{equ} = \{(0,0), \ldots, (n-1, n-1)\}$. To minimize the error, all other rectangles should have a computed value of 0, so we concentrate on $R_1, \ldots, R_J$.

For $j = 1, \ldots, J$ define $S_j \stackrel{\text{def}}{=} R_j \cap \mathbf{1}_{equ}$. Then, $\sum_{j=1}^{J} |S_j| = n$ and, since each R_j is a rectangle and S is a diagonal, $|S_j|^2 \le |R_j|$. The average error over $\mathbf{0}_{equ}$ is,

$$\bar{E}_\phi^{\mathbf{0}_{equ}} = \sum_{\{R \in \mathcal{R}_\phi : V_\phi(R)=1\}} \frac{|R \cap \mathbf{0}_{equ}|}{|\mathbf{0}_{equ}|} \ge \sum_{\{j : V_\phi(R_j)=1\}} \frac{|R_j - S_j|}{|\mathbf{0}_{equ}|} \ge \sum_{\{j : V_\phi(R_j)=1\}} \frac{|S_j|^2 - |S_j|}{n(n-1)} .$$

For the average error over $\mathbf{1}_{equ}$ we get,

$$\bar{E}_\phi^{\mathbf{1}_{equ}} = \sum_{\{R \in \mathcal{R}_\phi : V_\phi(R)=0\}} \frac{|R \cap \mathbf{1}_{equ}|}{|\mathbf{1}_{equ}|} \ge \sum_{\{j : V_\phi(R_j)=0\}} \frac{|S_j|}{n} \ge \sum_{\{j : V_\phi(R_j)=0\}} \frac{|S_j|^2 - |S_j|}{n(n-1)} .$$

Incorporating these equations into the average error calculation we get,

$$\bar{E}_\phi = \frac{1}{2}\left(\bar{E}_\phi^{\mathbf{0}_{equ}} + \bar{E}_\phi^{\mathbf{1}_{equ}}\right) \ge \frac{1}{2n(n-1)} \sum_{j=1}^{J} \Big(|S_j|^2 - |S_j|\Big) = \frac{1}{2n(n-1)} \Big(\sum_{j=1}^{J} |S_j|^2 - n\Big) .$$

Hence,

$$\sum_{j=1}^{J} (|S_j|/|\mathbf{1}_{equ}|)^2 \le \Big(2\bar{E}_\phi(n-1)+1\Big)/n .$$

Using the inequality,

$$H\langle p_j : j = 1, \ldots, J\rangle \ge \log \frac{1}{\sum_{j=1}^{J} p_j^2}$$

we get,

$$\bar{L}_\phi^{\mathbf{1}_{equ}} \ge H\langle |S_j|/|\mathbf{1}_{equ}| : j = 1, \ldots, J\rangle \ge \log\Big(\frac{n}{1 + 2\bar{E}_\phi(n-1)}\Big) .$$

By the convexity of the logarithm,

$$\bar{L}_\Phi^{\mathbf{1}_{equ}} \ge \sum_{\phi \in \Phi} p(\phi) \log\Big(\frac{n}{1 + 2\bar{E}_\phi(n-1)}\Big) \ge \log\Big(\frac{n}{1 + 2\sum_{\phi \in \Phi} p(\phi)\bar{E}_\phi(n-1)}\Big) = \log\Big(\frac{n}{1 + 2\bar{E}_\Phi(n-1)}\Big) .$$

Therefore, $\hat{C}_R(equ, \bar{\epsilon}) \ge \log(\frac{n}{1+2\epsilon(n-1)})$. To prove tightness we now describe a protocol whose complexity is two bits more than the lower bound. Note, though, that the complexity of the protocol is computed under the shared random sources model.

Let $\epsilon > 0$. Consider all partitions of $\{0, .., n-1\}$ into $\lceil \frac{n}{1+\lfloor 2\epsilon(n-1)\rfloor} \rceil$ sets of size $\le 1 + 2\epsilon(n-1)$ each. $P_{\mathcal{X}}$ picks one such partition at random (all partitions being equally likely) and transmits the index of the set that contains x. By the model assumption he doesn't need to transmit which partition was picked. $P_{\mathcal{Y}}$ then, transmits 1 if y belongs to the same set in the partition and 0 otherwise.

The number of bits transmitted is $1 + \lceil \log \frac{n}{\lfloor 1+2\epsilon(n-1) \rfloor} \rceil$. The error probability is clearly 0 if $x = y$ and, by symmetry or a more cumbersome argument, at most 2ϵ for $x \neq y$.

As mentioned before, the protocol relies heavily on the assumption that the random experiment is shared by $P_{\mathcal{X}}$ and $P_{\mathcal{Y}}$. There are far more than n partitions of $\{0, .., n-1\}$ into sets of size $1 + 2\epsilon(n-1)$. Thus, if $P_{\mathcal{X}}$ has to specify which one he uses, more than $\log n + 1$ bits must be transmitted. In fact, it is stated in [2] that under the separation of random sources assumption, $\hat{C}_R(equ, \hat{\epsilon}) = \Theta(\log \log n)$. We note that for the equality function worst case and average errors yield similar complexities. $\square$

References

[1] H. Abelson, 'Lower Bounds on Information Transfer in Distributed Computations', *Proc. of the 19th Annual Symposium on Foundations of Computer Science*, 1978.

[2] A.C. Yao, 'Some Complexity Questions Related to Distributive Computing', *Proc. of the 11th Annual ACM Symposium on the Theory of Computing*, 1979.

[3] C.H. Papadimitriou and M. Sipser, 'Communication Complexity', *Proc. of the 14th Annual ACM Symposium on the Theory of Computing*, April 1982.

[4] A. El Gamal and A. Orlitsky, 'Interactive Data Compression', *Proc. of the 25th Annual Symposium on Foundations of Computer Science*, October 1984.

[5] K. Mehlhorn and E.M. Schmidt, 'Las Vegas is better then Determinism in VLSI and Distributed Computing', *Proc. of the 14th Annual ACM Symposium on the Theory of Computing*, April 1982.

[6] A. Orlitsky and A. El Gamal, 'Randomized Communication Complexity', *IEEE International Symposium on Information Theory*, June 1985.

[7] B. Chor, and O. Goldreich, 'Unbiased Bits from Sources of Weak Randomness and Probabilistic Communication Complexity', *Proc. of the 26th Annual Symposium on Foundations of Computer Science*, October 1985.

[8] C.E. Shannon, 'A Mathematical Theory of Communication', *Bell Systems Technical Journal*, July 1948.

[9] K.F. Pang and A. El Gamal, 'Communication Complexity of Computing the Hamming Distance', to appear, *SIAM J. of Computing*.

[10] M.O. Rabin and A. Yao, unpublished manuscript (see [2]).

[11] W. Uhlmann, 'Vergleich der hypergeometrischen mit der Binomial Verteilung', *Metrika*, 10, 1966.

[12] A. Rényi, *Probability Theory*. North Holland, 1970.