University of California, Los Angeles
CS 289A Communication Complexity

Instructor: Alexander Sherstov
Scribe: Max Ohlendorf
Date: January 11, 2012

LECTURE
2

Lower Bounds for Deterministic Communication

This lecture builds on the material from our first lecture, providing more tools for studying communications protocols. Specifically, we look more closely at the relationships among *protocols*, *binary trees*, *matrices*, and *rectangles*. These are key notions in communication complexity proofs. We then present techniques for proving lower bounds on a problem's communication complexity. The two techniques covered today include *fooling sets* and *rectangle size bounds*. We use these techniques to derive tight lower bounds for the communication complexity of equality, greater-than, disjointness, and inner product.

2.1 Characteristic Matrices

Given a function $f : X \times Y \to \{0, 1\}$, the characteristic matrix of f is given by:

$$M_f = [f(x, y)]_{x \in X, y \in Y}$$

If $X = \{x_1, x_2, ..., x_m\}$ and $Y = \{y_1, y_2, ..., y_n\}$, this can be pictorially represented as

$$M_f = \begin{bmatrix} f(x_1, y_1) & f(x_1, y_2) & ... & f(x_1, y_n) \\ f(x_2, y_1) & f(x_2, y_2) & ... & f(x_2, y_n) \\ ... & ... & ... & ... \\ f(x_m, y_1) & f(x_m, y_2) & ... & f(x_m, y_n) \end{bmatrix}$$

Note that in the study of communication complexity, we ignore all costs other than those incurred by communication. Because we assume that the system has infinite computational power, *reordering or duplicating the rows and columns of the characteristic matrix does not change the communication complexity*, since we can computationally undo all of these transformations before performing the communication.

2.2 Protocols, Trees, and Rectangles

We talked about the tree representation of a protocol during Lecture 1. Here we show the connection between protocols, trees, and rectangles.

		Y			
		00	01	10	11
X	00	0	0	0	1
	01	0	0	0	1
	10	0	0	0	0
	11	0	1	1	1

FIGURE 2.1: A table mapping out the function values for all possible input values on $\{0,1\}^2 \times \{0,1\}^2$.

For the purpose of illustration, we use the following example based on the algorithm presented in Figure 1.1 of [1]. Let $X = Y = \{00, 01, 10, 11\}$, and let $f : X \times Y \to \{0,1\}$ be given by its characteristic matrix in Figure 2.1. A valid protocol tree for f is illustrated in Figure 2.2. Note that for every input pair (x, y) in the function domain, the protocol reaches a leaf node in the tree that contains the value of $f(x, y)$.

2.3 Partitions

DEFINITION 2.1. Consider a function $f : X \times Y \to \{0,1\}$. In other words, this is a function defined on the Cartesian product of X and Y, and returning a Boolean value. A subset $S \subseteq X \times Y$ is *f-monochromatic* if f is constant over S.

THEOREM 2.2. *The set of inputs $(x, y) \in X \times Y$ that reaches any given leaf of the protocol tree is a rectangle.*

Proof. Consider Figure 2.2. We reach the fourth leaf (in left-to-right order) if and only if $(x, y) \in a_1(x) \cap b_3(y) \cap a_4(x)$. More generally, any path to a leaf node corresponds to the intersection of some $a_i(x)$ and $b_j(y)$ (each viewed as a subset of X or Y, respectively). Note that we can always rearrange the elements of the intersection to put all a_i's in front of any b_j's.

Now, define $A(x) = \bigcap a_i(x)$ and $B(x) = \bigcap b_j(x)$, where the intersection is over those a_i and b_j that are encountered along the path of interest. Then the set of inputs reaching the given leaf corresponds to $A(x) \cap B(y)$, which is a rectangle by definition. ∎

COROLLARY 2.3. *Every deterministic communication protocol for f partitions $X \times Y$ into a disjoint union of f-*monochromatic *rectangles. Furthermore, there are at most 2^C rectangles, where C is the communication cost of the protocol (the height of the tree).*

The upper bound of 2^C in this corollary holds because a binary tree of height C cannot have more than 2^C leaves. To illustrate, Figure 2.3 shows the partition on the matrix in Figure 2.1 corresponding to the protocol tree shown in Figure 2.2.

While every deterministic communication protocol partitions $X \times Y$ into a disjoint union of rectangles, it is not true that every partition of $X \times Y$ into a disjoint union of rectangles

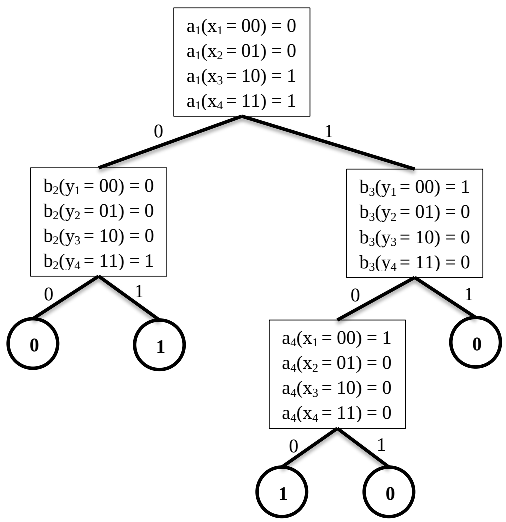

FIGURE 2.2: A protocol tree, where leaf nodes are represented by circles. At each internal node of the protocol, the function determines which branch of the tree to take (0 = left, 1 = right). The communication cost of the protocol is given by the height of the tree.

	Y			
X	00	01	10	11
00	0	0	0	1
01	0	0	0	1
10	0	0	0	0
11	0	1	1	1

FIGURE 2.3: The partition of Figure 2.1 corresponding to the deterministic communication protocol of Figure 2.2.

	Y 00	01	10	11
X 00	0	0	0	1
01	0	0	0	1
10	0	0	0	0
11	0	1	1	1

FIGURE 2.4: The above partition cannot be induced by a protocol, since there is no clean first cut that divides the rectangle into two partitions.

can be induced by a protocol. Consider Figure 2.4, which cannot be induced by any protocol. Each "cut" in the partition corresponds to a node in the binary tree, which cleanly divides its children into either a left subtree or a right subtree. Thus, each cut must cleanly divide an existing submatrix into two parts. In the shown partition, there is no way to make a first cut that cleanly divides the matrix into two parts.

2.4 Fooling Sets

DEFINITION 2.4. Let $f : X \times Y \to \{0,1\}$. A set $S \subseteq X \times Y$ is called a *fooling set* for f if there exists a value $z \in \{0,1\}$ such that:

a) For every $(x,y) \in S$, $f(x,y) = z$.

b) For any distinct (x_1,y_1) and (x_2,y_2) in S, either $f(x_1,y_2) \neq z$ or $f(x_2,y_1) \neq z$.

THEOREM 2.5. *The communication complexity of f obeys $D(f) \geq \log_2 |S|$, where S is any fooling set for f.*

Proof. Immediate from Corollary 2.3, since no two elements of the fooling set can occupy the same rectangle in the partition induced by a valid communication protocol. ∎

EXAMPLE 2.6 (corresponds to Example 1.21 in [1]). Consider the equality function $EQ_n : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$,

$$EQ_n(x,y) = \begin{cases} 1, & \text{if } x = y \\ 0, & \text{otherwise} \end{cases}$$

The characteristic matrix is simply the $2^n \times 2^n$ identity matrix, where the value of EQ_n is 1 where the row and column are equal and 0 otherwise.

$$M_{EQ_n} = \begin{bmatrix} 1 & & & \\ & 1 & & \text{\huge 0} \\ & & \ddots & \\ \text{\huge 0} & & & 1 \end{bmatrix}$$

The fooling set in this case is the set of 1-values. Note that if we were to make a rectangle with any pair of 1-values, the corresponding rectangle would necessarily contain a 0.

We can make use of the fooling set to put a lower bound on the communication complexity: $D(EQ_n) \geq \lceil \log_2(2^n + 1) \rceil = n + 1$, where 2^n comes from counting the number of 1 values, and 1 comes from counting the 0-rectangles.

Example 2.7 (Corresponds to Exercise 1.22 in [1]). Consider the greater-than-equal-to function $GTE_n : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$,

$$GTE_n(x,y) = \begin{cases} 1, & \text{if } x \geq y \text{ (lex order)} \\ 0, & \text{otherwise} \end{cases}$$

The characteristic matrix is the lower-triangular matrix of order 2^n, with zeroes above the diagonal and ones everywhere else.

$$M_{GTE_n} = \begin{bmatrix} 1 & & & \\ & 1 & & \Large 0 \\ & & \ddots & \\ \Large 1 & & & 1 \end{bmatrix}$$

We make use of the same fooling set as before, the diagonal 1-entries. Note that if we were to make a rectangle with any pair of 1-values, the corresponding rectangle would necessarily contain a 0.

As in the last problem: $D(GTE_n) \geq \lceil \log_2(2^n + 1) \rceil = n + 1$, where 2^n comes from counting the number of 1 values, and 1 comes from counting the 0-rectangles.

Example 2.8 (Corresponds to Example 1.23 in [1]). Let Alice and Bob each hold a subset of $\{1, ..., n\}$ (denoted by x and y, respectively). Note that each of these subsets can be represented with an n-bit string. Let the disjointness function $DISJ_n : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ be given by

$$DISJ_n(x,y) = \begin{cases} 1, & \text{if } x \cap y = \emptyset \\ 0, & \text{otherwise} \end{cases}$$

For our fooling set, we take $S = \big\{(A, \bar{A}) | A \subseteq \{1, ..., n\}\big\}$. Then $|S| = 2^n$. As in our previous examples, $D(DISJ_n) \geq \lceil \log_2(2^n + 1) \rceil = n + 1$, where 2^n comes from counting the number of 1 values, and 1 comes from counting the 0-rectangles.

2.5 Rectangle Size Bounds

This is another technique for proving lower bounds on communication complexity. In fact, the fooling set technique is a special case of the rectangle bound technique. Our basic strategy is to prove that the size of every monochromatic rectangle is small, and thus that many monochromatic rectangles are needed to partition $X \times Y$.

Let $f : X \times Y \to \{0,1\}$. We define a probability distribution μ on $f^{-1}(1)$, and argue that $\mu(R)$ is small for any rectangle R consisting of only 1s. Alternatively, we can make the same proof by using rectangles consisting only of 0s.

Here we show that fooling set technique is a special case of the rectangle bound technique, corresponding to Proposition 1.24 of [1]:

Example 2.9. Given a fooling set S where $|S| = t$, define μ to be a uniform distribution on the elements of the fooling set. Thus,

$$\mu(x,y) = \begin{cases} 1/t, & \text{for } (x,y) \in S \\ 0, & \text{otherwise} \end{cases}$$

We have shown that no two elements of the fooling set can occupy the same f-monochromatic rectangle. Thus, every rectangle R has measure $\mu(R) \leq 1/t$, and there must be at least $1/(1/t) = t$ leaves in the protocol tree. This leads to the lower bound of $D(f) \geq \lceil \log_2 t \rceil$.

Example 2.10. Corresponds to Example 1.25 in [1]. Let the inner product function $IP_n : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ be given by

$$IP_n(x,y) = \left(\sum_{i=1}^{n} x_i y_i \right) \bmod 2.$$

Let μ be a uniform distribution on $IP_n^{-1}(0)$. μ is supported on $2^{2n-1} = \frac{4^n}{2}$ points. Let $R = S \times T$ be any 0-monochromatic rectangle. Identify $\{0,1\}$ with GF(2), and analogously $\{0,1\}^n$ with the n-dimensional vector space $\mathrm{GF}(2)^n$. Then $IP_n(x,y) = \langle x, y \rangle$, the inner product being taken in $\mathrm{GF}(2)^n$. Define $S' = \operatorname{span} S$ and $T' = \operatorname{span} T$. Then $R' = S' \times T'$ is still monochromatic, because the inner product satisfies

$$\langle a + a', b + b' \rangle = \langle a, b \rangle + \langle a, b' \rangle + \langle a', b \rangle + \langle a', b' \rangle.$$

But S' and T' are orthogonal subspaces of $\mathrm{GF}(2)^n$, so $\dim S' + \dim T' \leq n$. Thus,

$$|S'| \times |T'| \leq 2^{\dim S'} \times 2^{\dim T'} \leq 2^n$$

Therefore,

$$\mu(R) \leq \mu(R') \leq \tfrac{2^n}{4^n/2} = 2^{-(n-1)}$$

and

$$D(IP_n) \geq \log_2 2^{n-1} = \Omega(n)$$

Example 2.11 (Corresponds to Exercise 1.26 in [1]). Returning to the disjointness function $DISJ_n$, recall that we have already shown $D(DISJ_n) = n+1$. Now we seek to prove the same using the rectangle size bound. Consider two randomly-chosen n-bit strings x and y. For each bit $b \in 1, ...n$, there is a 3/4 probability that the two strings do not both have a value of 1 in bit b. Thus, the probability of having two n-bit strings that are totally disjoint will be $\left(\frac{3}{4}\right)^n$. As a result, $|DISJ_n^{-1}(1)| = 4^n(3/4)^n = 3^n$.

Now, let μ be the uniform distribution over $DISJ_n^{-1}(1)$, supported on 3^n points by the calculation just given. Let $R = \mathcal{S} \times \mathcal{T}$ be a 1-monochromatic rectangle of $DISJ_n$, where $\mathcal{S}, \mathcal{T}$ are collections of sets in $\mathcal{P}(\{1, 2, \ldots, n\})$. Let S and T be the union of all sets in $\mathcal{S}$ and $\mathcal{T}$, respectively. Then S and T are disjoint, and therefore we can enlarge R to the 1-monochromatic rectangle $R' = \mathcal{P}(S) \times \mathcal{P}(T)$. Since S and T are disjoint, we have $|S| + |T| \leq n$, so that

$$|\mathcal{P}(S)| \times |\mathcal{P}(T)| \leq 2^{|S|+|T|} \leq 2^n.$$

As a result, $\mu(R') = |R'|/3^n \leq (2/3)^n$, and $D(DISJ_n) = \Omega(n)$.

In preparation for next lecture, think about the following problem.

PROBLEM 2.12. Let $M \in \{0,1\}^{m \times n}$ be a given matrix. Prove that $\operatorname{rank}_{\mathbb{R}} M \geq \operatorname{rank}_{\mathbb{F}} M$, where $\mathbb{F}$ is any field. *Hint:* You may find it helpful to first prove that $\operatorname{rank}_{\mathbb{F}} M = \operatorname{rank}_{\mathbb{K}} M$ for any given fields $\mathbb{F} \subset \mathbb{K}$ and a matrix $M \in \mathbb{F}^{m \times n}$.

References

[1] E. Kushilevitz and N. Nisan. *Communication complexity.* Cambridge University Press, 2nd edition, 2006.