6.841 Advanced Complexity Theory February 23, 2005

Lecture 6

Lecturer: Madhu Sudan *Scribe: Mihai Pătraşcu*

1 Communication Complexity

Communication complexity is interesting from a complexity-theoretic perspective because it offers ways to prove lower bounds in various models of computation. For example, there are several known relations to circuits, one of which we discuss in this lecture. This connection would enable use to prove $\omega(\lg n)$ bounds on circuit depth if we could prove a good lower bound on a certain comunication game.

Communication complexity was introduced by Yao [4], and has been studied intensively since then. Good references are the book by Kushilevitz and Nisan [1] and a series of lecture notes by Ran Raz [3].

The most simple setup is a two-player communication game. We have two players, Alice and Bob, who get inputs x and y, respectively, chosen from $\{0,1\}^n$. The two players want to compute a function $f(x,y)$ by communicating a minimum number of bits. For now, let us restrict the range of f to be boolean. Observe the difference from information theory: while information theory concentrates on sending a message (the data), here we are not interested in the actual data, just a function of the players' inputs. This allows the communication complexity to be lower than sending the entire inputs in many cases.

Formally, the solution to a communication game is a protocol. The protocol is a tuple of functions (s, m_a, m_b, f_a, f_b):

- f takes as argument the communication history and returns which player speaks next, or that the protocol is over.
- m_a takes as arguments the communication history and Alice's input, and decides the next bit sent by Alice; m_b is similar for Bob.
- f_a takes as arguments the communication history and Alice's input, and decides the result of the protocol; f_b is similar for Bob. The protocol is correct if both f_a and f_b are equal to $f(x,y)$ at the end of the protocol.

The complexity of the protocol is the maximum number of bits exchanged over any pair of inputs (x,y). The communication complexity of the problem, $CC(f)$ is the minimum over all protocols for f of the complexity of the protocol.

The are many variations on this model, which are useful in different setting. For example, one can consider: randomization (with zero or bounded error), nondeterminism, message complexity (bound the number of alternations between Alice speaking and Bob speaking), multiplayer games (the most interesting variation being the number-on-the-forehead model, where player i gets all x_j's except x_i) etc.

2 Techniques for Communication Lower Bounds

In the model which we defined, it is clear that $CC(f) \leq n+1$ for any f, because Alice can just send her input, and Bob can reply with $f(x, y)$. Note that the model is nonuniform, and does not care about computational power, so Bob can compute any function once he knows enough information. Lower bounds of $n+1$ can be shown easily (see below), even for very simple explicit functions. The interesting questions are showing bounds for some specific problems, such as the one we discuss in the next section, which relates to circuit complexity.

Two of the most important techniques for communication lower bounds originate in Yao's seminal paper [4]. The first one is the fooling-set method. Consider a matrix A^f of size $2^n \times 2^n$ where $A^f_{x,y} = f(x, y)$. The crucial observation is that communication divides this matrix into disjoint *rectangles*. A rectangle is a set $X \times Y$, with $X \subset \{0,1\}^n, Y \subset \{0,1\}^n$ (since X, Y are arbitrary, we care about combinatorial rectangles, not geometric rectangles). The claim is that for any communication history, the set of inputs (x, y) for which this history materializes is a rectangle. Indeed, assume (x_1, y_1) and (x_2, y_2) generate the same communication history. Now (x_1, y_2) and (x_2, y_1) must still generate the same history. This can be seen by induction of the number of bits sent. Say Alice speaks next; she decides her bit based on her input, and the communication history. Both for y_1 and y_2 as Bob's input, the communication history so far is the same (by the induction hypothesis), so she outputs the same bit.

A correct protocol must announce $f(x, y)$ at the end. It follows that at the end of the protocol, every communication history leads to a monochromatic rectangle. Assume the rectangle contained both zeros and ones. Then, there is either a row or a column containing both zeros and ones. Say it's a row; then, if Alice gets the input corresponding to that row, at the end of the communication she cannot tell whether $f(x, y)$ is zero or one, because she cannot differentiate between Bob's inputs leading to the rectangle.

Now consider the equality function. The matrix A^f has zeros everywhere except the diagonal. It is immediate that the number of one-rectangles is 2^n, because two ones cannot appear in the same rectangle. Similarly, it is seen that there are at least 2^n zero-rectangles. Then, the communication complexity is $\geq n + 1$, because each new bit doubles the number of rectangles.

A second lower-bound technique takes an algebraic view of A^f. It can be shown, using a similar argumentation as above, that $CC(f) \geq \lg \operatorname{rank}(A^f)$. It is known that there exist f for which $CC(f) = \omega(\lg \operatorname{rank}(A^f))$. However, it is still an open question whether $CC(f) = O(\operatorname{poly}(\lg \operatorname{rank}(A^f)))$.

3 A Relation to Circuit Depth

We now investigate an interesting connection between communication complexity and circuit depth, which was described by Karchmer and Wigderson [2]. For this, we need to generalize communication complexity to consider relations. Consider a relation $R \subset X_A \times X_B \times D$, where $X_A, X_B \subseteq \{0,1\}^n$ and D is arbitrary. At the beginning, Alice receives $x \in X_A$, and Bob receives $y \in X_B$. At the end of the protocol, each player must decide on a certain i based on his input and the communication history. It must be the case that both players decide the same i, and that $(x, y, i) \in R$.

For a function $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$, we define the following cannonical relation. Players receives inputs from $X_A = f^{-1}(0)$ and $X_B = f^{-1}(1)$; we have $D = \{1, \ldots, n\}$. The players must decide the relation $R_f = \{(x, y, i) \mid x \in X_A, y \in X_B, x_i \neq y_i\}$. Clearly, if $f(x) \neq f(y)$, there must be at least one bit where they differ; this is saying that the communication problem is for the two players to find that bit.

Theorem 1 *Let $d(f)$ be the circuit depth of f on a circuit with basis $(2-AND, 2-OR, NOT)$. Then $CC(R_f) = \Theta(d(f))$.*

Proof We first show $CC(R_f) \leq d(f)$. The players know a circuit which outputs 0, respectively 1, when applied to x and y. Consider the last gate of the circuit, and assume, for instance, that it is an AND. Thus, the output of the circuit is $f(z) = f_0(z) \wedge f_1(z)$. If $f(x) = 0$, at least one of $f_0(x), f_1(x)$ is zero; since $f(x) = 1$, $f_0(y) = f_1(y) = 1$. Now Alice just outputs $f_0(x)$. If this is zero, the players have a circuit of depth $d(f_0) \leq d(f) - 1$ which differentiates x and y; otherwise, f_1 differentiates between x and y and it also has depth $\leq d(f) - 1$. Then, the conclusion follows by induction on the depth. The base case is that a single bit is found (depth zero) which differentiates x and y.

We now show $CC(R_f) \geq d(f)$. This also works by induction on $CC(R_f)$. Assume by induction that for any two disjoint sets $X_A, X_B \subset \{0,1\}^n$ for which there is a d-bit protocol "separating" X_A and X_B, there exists a depth d protocol implementing a function $g : \{0,1\}^n \rightarrow \{0,1\}$ with $g^{-1}(0) \supseteq X_A, g^{-1}(1) \supseteq X_B$. The base case is $d = 0$. Then, Alice and Bob know a bit which is different for inputs with $f(x) = 0$ vs $f(x) = 1$. So the circuit can just look at that bit (this is fixed nonuniformly), and g is the bit or the bit negated, depending on which way the separation goes.

We now prove the claim for $d+1$. Assume by symmetry that Alice speaks, and her bit is $P(x)$ where $P : X_A \rightarrow \{0,1\}$. This bit induces a partition of X_A into $P^{-1}(0)$ and $P^{-1}(1)$. If the first bit is zero, the remaining protocol has complexity $\leq d$ and it separates $P^{-1}(0)$ from X_B. If the first bit is one, we have a protocol that separates $P^{-1}(1)$ from X_B. By induction, these protocols can be converted into circuits C_0 and C_1. Observe that if $x \in X_B$, both $C_0(x)$ and $C_1(x)$ output 1, by induction. If $x \in X_A$, either $x \in P^{-1}(0)$ or $x \in P^{-1}(1)$. Then either $C_0(x)$ or $C_1(x)$ outputs zero. Then, if we take $C = C_0 \wedge C_1$, it outputs one on X_B and zero on X_A. ■

References

[1] Eyal Kushilevitz and Noam Nisan: *Communication Complexity*, Cambridge University Press, 1996.

[2] Mauricio Karchmer, Avi Wigderson: *Monotone Circuits for Connectivity Require Super-logarithmic Depth*, STOC 1988: 539-550.

[3] Ran Raz: *Circuit Complexity and Communication Complexity*, lecture notes series from IAS summer school on Complexity Theory. Park City Mathematical Series, Volume 10, 2000.

[4] Andrew Chi-Chih Yao: *Some Complexity Questions Related to Distributive Computing (Preliminary Report)*, STOC 1979: 209-213.