University of California, Los Angeles
CS 289A Communication Complexity

Instructor: Alexander Sherstov
Scribe: Dong Wang
Date: January 23, 2012

LECTURE
4

Partitions and Covers

In previous lectures, we saw that every communication protocol induces a partition of the domain into monochromatic rectangles, and learned two lower bound techniques for the number of rectangles in such a partition. The techniques presented so far take into account only the number of monochromatic rectangles in partitions of the space $X \times Y$ and ignore the additional restriction that these partitions should correspond to some protocol. A natural question is how tight these lower bound techniques are.

In addition, we are also interested in relaxing the need to partition the space $X \times Y$ into f-monochromatic rectangles by allowing overlaps among rectangles. In other words, we are interested in covering $X \times Y$ by monochromatic rectangles. Coverings are more natural combinatorial objects than partitions and can sometimes be more efficient. This raises the question of how much more efficient a cover can be than a partition. In this lecture, we study the relations among these combinatorial measures and their relation to communication complexity.

4.1 Key definitions

Definition 4.1. Let $f : X \times Y \to \{0, 1\}$ be a function. Define $C^P(f)$ to be the smallest number of leaves in any protocol tree for f. Define $C^D(f)$ to be the smallest number of f-monochromatic rectangles in any partition of $X \times Y$. Define $C(f)$ to be the smallest number of f-monochromatic rectangles in any cover of $X \times Y$.

The following relationships hold among these quantities:

Theorem 4.2. *For all $f : X \times Y \to \{0, 1\}$:*

$$2^{\Theta(D(f))} = C^P(f) \geq C^D(f) \geq C(f) \geq 2^{\Theta(\sqrt{D(f)})}$$

The first and last inequalities are fundamental results and will be proved in this lecture. The other inequalities are immediate from the definitions. In particular, we infer that covers and partitions cannot be significantly more efficient than partitions that arise from communication protocols.

4.2 Communication complexity versus protocol cover

THEOREM 4.3. *For all $f : X \times Y \to \{0,1\}$,*

$$D(f) = \Theta(\log C^P(f))$$

Proof. The maximum number of leaves in a binary tree with depth $D(f)$ is $2^{D(f)}$. Thus, $C^P(f) \leq 2^{D(f)}$, on equivalently $D(f) \geq \log_2 C^P(f)$. For the upper bound on $D(f)$, consider a protocol for f with l leaves.

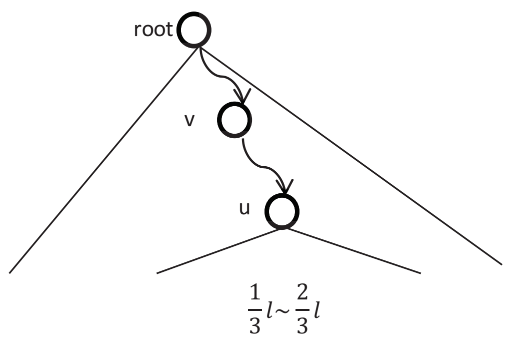

FIGURE 4.1: Figure for finding target u.

Consider Figure 4.1. Alice and Bob identify a node u in the tree such that the subtree rooted at u has between $\frac{l}{3}$ and $\frac{2l}{3}$ leaves. To find such a node, we start at the root of the tree and keep going to the child that has more than $\frac{2l}{3}$ leaves in its subtree. This process stops at some node v. By assumption, the subtrees rooted at v's children each have at most $\frac{2l}{3}$ leaves, and furthermore one of them must have at least $\frac{l}{3}$ leaves (since the subtree rooted at v has more than $\frac{2l}{3}$ leaves). Thus, one of v's children is the desired node u. Note that finding u requires no communication.

Since u is a node in the protocol tree, the set of inputs arriving at u is a rectangle; call it R_u. Thus, with 2 bits of communication Alice and Bob are able to determine whether $(x,y) \in R_u$. If $(x,y) \in R_u$, Alice and Bob recursively process the subtree rooted at u, which has between $\frac{l}{3}$ and $\frac{2l}{3}$ leaves. If $(x,y) \notin R_u$, they recursively process the original tree with u replaced by a terminal node 0 (the label actually does not matter).

In either case, the recursive step is applied to a tree with at most $\frac{2l}{3}$ leaves, giving the recurrence $D(l) \leq 2 + D(2l/3)$ where $D(l)$ is the maximum communication complexity over all functions f with $C^P(f) \leq l$. The recurrence solves to $D(l) \leq \Theta(\log l)$ as claimed. ∎

4.3 Communication complexity versus cover size

DEFINITION 4.4. Let $f : X \times Y \to \{0,1\}$ be a function. Define $C^0(f)$ to be the smallest number of f-monochromatic rectangles needed to cover $f^{-1}(0)$. Define $C^1(f)$ to be the smallest number of f-monochromatic rectangles needed to cover $f^{-1}(1)$.

By definition, $C(f) = C^0(f) + C^1(f)$. We will now show a bound on $D(f)$ in terms of $C(f)$.

Theorem 4.5 (Aho, Ullman, Yannakakis [1]).

$$D(f) \leq O(\log_2 C^0(f) \log_2 C^1(f)).$$

In particular,

$$D(f) \leq O(\log^2 C(f)).$$

Proof. The proof relies on a simple property of rectangles. Let R and S denote f-monochromatic rectangles in a cover of $f^{-1}(0)$ and $f^{-1}(1)$, respectively. As Figure 4.2 shows, R and S must be disjoint either in rows or columns or both.

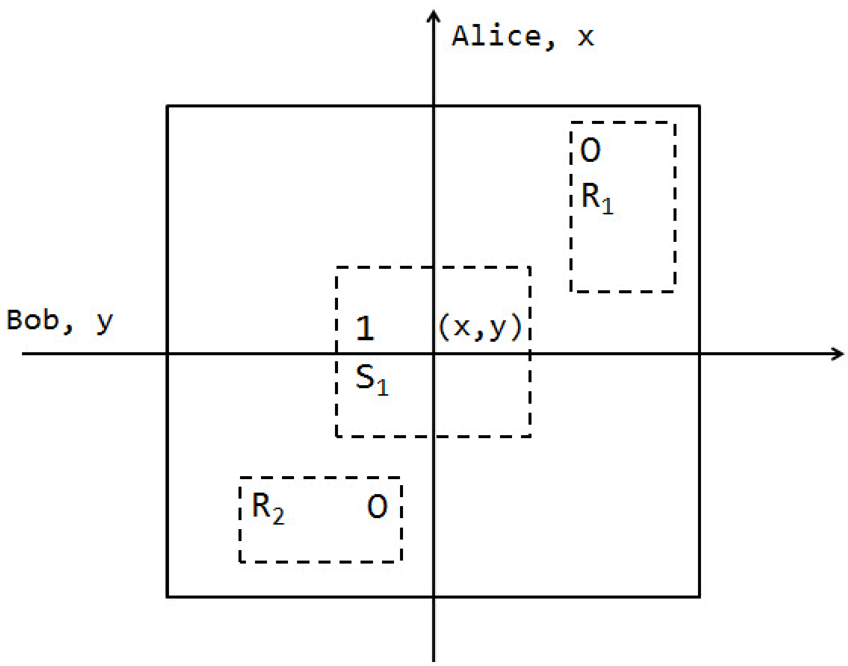

Figure 4.2: Intersections in rows or columns

Fix an optimal cover of $f^{-1}(0)$ and $f^{-1}(1)$ by f-monochromatic rectangles. We now describe a protocol for Alice and Bob. The idea is for them to look for a 1-rectangle that contains the input (x, y). If they fail, the conclude that $f(x, y) = 0$. In each round the players do the following:

1. Alice looks for a 1-rectangle Q that contains column x and is disjoint from at least half of the rectangles in the cover of $f^{-1}(0)$. If she finds it, she sends its index to Bob, using $\log_2 C^1(f)$ bits, and the rectangles in the cover of $f^{-1}(0)$ that are disjoint from Q are discarded.

2. If Alice fails, Bob looks for a 1-rectangle Q that contains row y and is disjoint from at least half of the rectangles in the cover of $f^{-1}(0)$. If he finds it, he sends its index to

Alice using $\log_2 C^1(f)$ bits, and the rectangles in the cover of $f^{-1}(0)$ that are disjoint from Q are discarded.

3. If both Alice and Bob failed, we know that $f(x, y) = 0$ and no further communication is necessary.

4. If they find the desired rectangle, they recursively process the remaining rectangles in the cover of f. When there are no 0-rectangle left, they return 1.

Each round reduces the number of 0-rectangles by a factor of at least 2, so there can be at most $\log_2 C^0(f) + 1$ rounds. Therefore, the protocol costs $O(\log_2 C^0(f) \log_2 C^1(f))$ bits in the worst case. ∎

4.4 Lower bounds for covers and disjoint covers

We now revisit several lower bound techniques for deterministic communication complexity and see which ones of them can directly bound quantities such as $C(f)$ and $C^D(f)$. In the statement to follow, we let $\mathrm{FS}(f)$ denote the size of the largest fooling set for f.

Theorem 4.6. *For any $f : X \times Y \to \{0, 1\}$ and any field $\mathbb{F}$,*

- $C(f) \geq \mathrm{FS}(f)$;
- $C^D(f) \geq \mathrm{rk}_{\mathbb{F}}(M_f)$.

Here, we sketch the proofs of these two statements.

Proof. Fix a fooling set $S \subseteq X \times Y$ for f, where $f(S) = z$. For every pair of $(x_1, y_1), (x_2, y_2) \in S$:

- $f(x_1, y_1) = z$
- $f(x_2, y_2) = z$
- Either $f(x_1, y_2) \neq z$ or $f(x_2, y_1) \neq z$.

Therefore, (x_1, y_1) and (x_2, y_2) cannot occupy the same rectangle in a cover of $f^{-1}(z)$, so that $C(f) \geq C^z(f) \geq |S|$.

For the second part, let $\mathscr{R}$ be a cover of $X \times Y$ by pairwise disjoint f-monochromatic rectangles. Let $\mathscr{R}' \subseteq \mathscr{R}$ be the set of rectangles covering $f^{-1}(1)$. By the pairwise disjointness of the rectangles, $M_f = \sum_{R \in \mathscr{R}'} M_R$, where the matrix M_R is defined by $M_R(x, y) = 1$ for $(x, y) \in R$ and $M_R(x, y) = 0$ for $(x, y) \notin R$. By the subadditivity of rank,

$$\mathrm{rk}\, M_f \leq \sum_{R \in \mathscr{R}'} \mathrm{rk}\, M_R \leq |\mathscr{R}| \leq C^D(f),$$

since the rank of each M_R is at most 1. ∎

4.5 The rectangle size bound characterizes cover size

In this section we prove the rather surprising fact, due to Karchmer, Kushilevitz, and Nisan [2], that the rectangle size bound nearly tightly characterizes the smallest cover size for any given function. Recall that to use the rectangle size bound, we define some probability distribution μ on, say, the 1-inputs of f. We then compute the largest $\mu(R)$, where R ranges over all 1-monochromatic rectangles, and conclude that the reciprocal of that quantity is a lower bound on $C^1(f)$.

Definition 4.7. For $f : X \times Y \to \{0,1\}$, define

$$RS^1(f) = \min_{\mu \text{ on } f^{-1}(1)} \max_R \mu(R),$$

where the minimum ranges over all probability distributions μ on $f^{-1}(1)$ and the maximum over all f-monochromatic rectangles $R \subseteq f^{-1}(1)$. Analogously, let

$$RS^0(f) = \min_{\mu \text{ on } f^{-1}(0)} \max_R \mu(R)$$

where the minimum ranges over all probability distributions μ on $f^{-1}(0)$ and the maximum over all f-monochromatic rectangles $R \subseteq f^{-1}(0)$.

Recall that

$$C^0(f) \geq \frac{1}{RS^0(f)},$$
$$C^1(f) \geq \frac{1}{RS^1(f)}.$$

We now prove matching upper bounds.

Theorem 4.8 (Karchmer, Kushilevitz, and Nisan [2]). *For any function $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$, one has*

$$C^0(f) \leq O\left(\frac{n}{RS^0(f)}\right),$$
$$C^1(f) \leq O\left(\frac{n}{RS^1(f)}\right).$$

Proof. By symmetry, it suffices to prove the first statement. We do so by building a cover for $f^{-1}(0)$ adaptively.

1. Let μ_i be a uniform distribution on those points in $f^{-1}(0)$ uncovered in previous $i-1$ iterations. Thus, μ_1 is the uniform distribution on $f^{-1}(0)$.
2. According to our definition of $RS^0(f)$, there exists a rectangle $R_i \subseteq f^{-1}(0)$ with $\mu_i(R_i) \geq RS^0(f)$.
3. Remove the points covered by R_i.

Trivially, $|f^{-1}(0)| \leq 4^n$. Each iteration covers at least an $RS^0(f)$ fraction of previously uncovered points, so that after k iterations at most $(1-RS^0(f))^k 4^n$ points remain uncovered. It follows that all points will be covered in $O(n/RS^0(f))$ iterations. ∎

These results still leave a small gap between the lower and upper bounds on $C^0(f)$ in terms of $RS^0(f)$, corresponding to the factor of n in the theorem above. As the following example shows, this gap cannot be narrowed further. We use the following observation:

PROPOSITION 4.9. *For all* $f\colon X \times Y \to \{0,1\}$, $\quad D(f) \leq C^0(f) + 1$.

Proof. On input (x, y), Alice sends Bob a single bit for each rectangle in the cover of $f^{-1}(0)$, indicating whether that rectangle intersects the row x. Bob checks the same for his input y. Finally, he outputs 1 if and only if no rectangle in the cover of $f^{-1}(0)$ intersects both the row x and the column y. ∎

EXAMPLE 4.10. Consider the equality function $EQ_n : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$. From previous lectures, we know that $D(EQ) = n + 1$, so that $C^0(f) \geq n$ by the proposition above. On the other hand, we claim that $RS^0(f) \geq 1/4$. Let μ be any distribution on $EQ_n^{-1}(0)$. Consider a 0-monochromatic rectangle for EQ_n chosen at random as follows: choose a random n-bit string $r \in \{0,1\}^n$ and a random bit b, and let

$$R_{r,b} = \{x : \langle x, r\rangle = b\} \times \{y : \langle y, r\rangle \neq b\}.$$

For any fixed (x, y) with $x \neq y$, $Pr[(x, y) \in R_{r,b}] = 1/4$ by the properties of inner product. It follows that $E_\mu Pr_{r,b}[(x, y) \in R_{r,b}] \geq 1/4$ whence

$$E_{r,b}\, \underbrace{Pr_\mu[(x, y) \in R_{r,b}]}_{=\mu(R_{r,b})} \geq 1/4.$$

Thus, there exists at least one 0-monochromatic rectangle $R_{r,b}$ with $\mu(R_{r,b}) \geq 1/4$. Since the choice of μ was arbitrary, we have $RS^0(EQ_n) \geq 1/4$.

References

[1] A. V. Aho, J. D. Ullman, and M. Yannakakis. On notions of information transfer in VLSI circuits. In *Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing* (STOC), pages 133–139, 1983.

[2] M. Karchmer, E. Kushilevitz, and N. Nisan. Fractional covers and communication complexity. *SIAM J. Discrete Math.*, 8(1):76–92, 1995.