

Lecture 5: Deterministic Communication Complexity

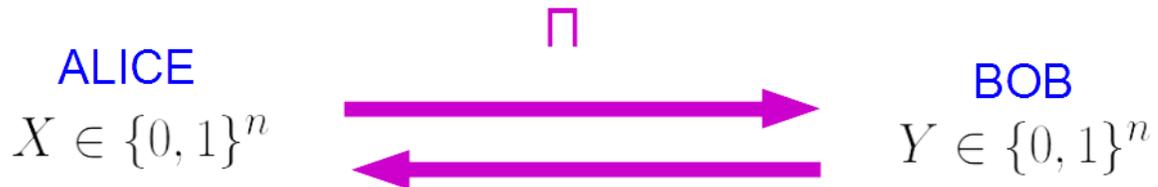
Lecturer: Omri Weinstein

Scribes: Natania Wolansky

5.1 Motivation

Before, we dealt with just one-way communication – Alice is trying to communicate something to Bob. the goal was to communicate message X in a lossless way. The question was always, "how much does Alice need to communicate?"

We now consider a reciprocal communication structure, described by a protocol Π , shown below:



Now we would like to compute some function (loosely defined) of the two messages: $f(X, Y) : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$. Now, the question to ask is not how to transmit message X , but WHAT to transmit so that either Bob or Alice (or both) can compute $f\{X, Y\}$ with as little information as possible. For now, we assume that our protocol Π needs to be able to work for ANY given $\{X, Y\}$ pair.

Of course, in any situation like this, the naive protocol is simply for Alice to send her entire message to Bob, and to have him compute the function – so we can upper bound any protocol by n bits of communication.

But in many situations, we will need fewer bits than that. Consider three examples below.

5.1.1 Example 1: XOR

Consider the example where our function is $f\{X, Y\} := \bigoplus_{i=1}^n (x_i \oplus y_i)$ – that is, we xor each bit in X with each bit in Y and then xor all these bits together.

This surprisingly requires only 1 bit of communication – Alice can simply compute $\bigoplus_{i=1}^n (x_i)$ and send this to Bob, where he can then do the same, and simply compare these two numbers.

5.1.2 Example 2: Equality

Consider the example where $f(X, Y)$ is the equality function: 1 if $X = Y$, 0 otherwise. This example needs $n + 1$ bits, since we need to check if every single bit matches. (This is for the deterministic case, when we

cannot be wrong for ANY combination of X and Y . In reality/practice, we would use hash functions, where there is a small probability of collision, or error in computing the equality function. In fact, we can use multiple hashes, and it is proven by some wise mathematicians that with $O(\log(\frac{1}{\epsilon}))$ random hashes, we will be correct (non-colliding) with a probability $\geq 1 - \epsilon$.

5.1.3 Example 3: Greater Than

Consider the example where $f(X, Y) = 1$ if $X \geq Y$, 0 otherwise – where X can be considered the binary encoding of a single number. In this case, the deterministic communication complexity is n – since it's possible that we'd have to compare every bit to see the first one that differs (perhaps the numbers are only 1 away from each other). But again, in the randomized case, we'd need only $O(\log^2(n))$ bits, because we could do a binary search for the first differing bit.

5.2 Formal Definition of Protocol

Now let's define formally what we mean by a process for Alice and Bob to send information to each other.

Definition 5.2.1. We define a Communication Protocol Π (basically the formalization of a conversation) as a pair of mappings:

$$\begin{aligned} f_A : (\Pi_{<i}, X) &\rightarrow \{0, 1\}, \text{ even } i \\ f_B : (\Pi_{<i}, Y) &\rightarrow \{0, 1\}, \text{ odd } i \end{aligned}$$

By convention, we assume Alice always speaks first.

What do we mean by $\Pi_{<i}$? Essentially, Alice has access to all the communication in the past (what she told Bob, and what Bob told her) plus her string, X . Bob has access to the previous information exchanged in the conversation, plus Y , but NOT X .

We can visualize a deterministic protocol Π as a binary tree τ_Π :

As we can see, it's the identity matrix. Let's look at what happens on the communication matrix as we go through the protocol Π . In the first step, Alice sends one bit to Bob, either a 0 or 1. Perhaps she sends the first bit of her message. She therefore cuts the communication matrix in half, rowwise. For instance, if she sends a 0 then Bob knows her message is either 00 or 01 – so it lies in the top half of the communication matrix.

Now, Bob sends a message back. Perhaps he also sends the first bit of his string. Let's say Bob has 10, so he sends a 1 – then Alice knows that the answer lies in the right half of the top half of the matrix – So Bob has now further cut the submatrix in half, columnwise. With each new piece of communication, the submatrix Alice and Bob currently live in is split into two, either rowwise or columnwise depending on who is speaking. They can both stop when the submatrix is monochromatic (same value) – meaning there is only one possible value for $f(X, Y)$. So in our case, the top right half quarter of the matrix is monochromatic in 0, so Alice and Bob can stop communicating and output 0. Note, they do NOT need to know EXACTLY which string the other has – they are only trying to calculate the function.

Definition 5.3.1. A subset $R \subset \{0, 1\}^n \times \{0, 1\}^n$ is a combinatorial rectangle IFF:

$$\begin{aligned} R &= A \times B \\ &= \{(x, y) | x \in A, y \in B\} \end{aligned}$$

That is, we take some subset A of the rows and some subset B of the columns, and we have a combinatorial rectangle iff every combination of these subsets is in our subset R .

An alternate definition is: R is a combinatorial rectangle IFF:

$$(x, y), (x', y') \in R \implies (x, y'), (x', y) \in R$$

meaning, if x shows up with some y in the rectangle, then it must show up with EVERY y in the rectangle, and vice versa with y .

5.3.2 Protocol Trees - Combinatorial Rectangles

Our claim is that each node of the tree described above can be equated with a specific combinatorial rectangle in the communication matrix.

Claim 5.3.2. $\forall v_{\leq i}$ (some node in level i of the protocol tree) let $S_{v_{\leq i}}$ be the set of possible values of X and Y at $v_{\leq i}$ – formally:

$$S_{v_{\leq i}} := \{(x, y) | \Pi(x, y)_{\leq i} = v_{\leq i}\}$$

The claim is that $S_{v_{\leq i}}$ is a rectangle.

Proof. We will do a proof by induction on i . For the root of the tree, the claim obviously holds true. That takes care of the base case. Now, we assume that the claim holds for level i . Suppose i is even, WLOG, and suppose $S_{v_{\leq i}}$ is a rectangle. Then by the definition of rectangles, $S_{v_{\leq i}} = A_v \times B_v$. Say, WLOG, that Alice at this point chooses to communicate 1 to Bob. Then at level $i + 1$, we have a new set $S_{v_{\leq i+1}}$:

$$S_{v_{\leq i+1}} = (A_i \cap \{x | \Pi_{v_{\leq i}}(x) = 1\}) \times B_v.$$

which is also a rectangle. Thus we've proven this connection. \square

This also holds for the leaves of the tree. So finally, we can say that all S_l (leaves of tau_{Π}) are f -monochromatic (monochromatic under the function f that we are trying to compute) combinatorial rectangles. So The number of communication steps, and thus the deterministic communication complexity of the function f , is the depth of the tree – the max number of steps to reach the leaves of the tree. formally,

Definition 5.3.3. For all valid protocols – that is, for all $\Pi : \Pi(x, y) = f(x, y)$ for every (x, y) in the domain of the function f , the deterministic communication complexity is defined as:

$$D(f) = \min_{\Pi} (|\Pi|)$$

And by this connection, we can go the other way: any deterministic communication complexity protocol Π partitions M_f into (at most $2^{|\Pi|}$) f -monochromatic rectangles.

5.4 Lower Bound on Deterministic Communication Complexity

Suppose we are given a monochromatic partitioning – does this correspond to a protocol? No! Because every intermediate step also needs to be a rectangle. We can't assume that a monochromatic partitioning is also a protocol – that's why we can't naively just find the minimum coloring of the matrix and call it a day. Let the partition number of M_f , $P(M_f)$ be this minimum number of partition rectangles. Then we know that:

$$\log(P(M_f)) \leq D(f)$$

but we can't turn that inequality into an equality.

We CAN set a (better) lower bound on CC for a given f by proving what the minimum number of monochromatic rectangles can be in M_f , given that these rectangles come from a protocol. Here we give a couple of techniques to do this.

5.4.1 Fooling Sets

Definition 5.4.1. A fooling set (FS) of a communication matrix M_f is a set $R = \{(x_i, y_i)\}_{i=1}^k$ (k entries of M_f) where:

$$f(x_i, y_i) = 1 \forall i \in R \tag{5.1}$$

$$f(x_i, y_{j/i}) \text{ or } f(x_{j/i}, y_i) = 0 \tag{5.2}$$

That is, in this set, the combination of i and i makes 1, but only that – the combination with any other letter in R makes 0. So no two fooling set elements can live in the same monochromatic rectangle.

So, if we can find the largest fooling set FS for some given f , then we know that

$$D(f) \geq \log(|FS(M_f)|)$$

because we'd need at least $|FS(M_f)|$ rectangles (so $\log(|FS(M_f)|)$ levels to have this many leaves in the tree), one for each element in the fooling set.

5.4.1.1 Equality

Claim 5.4.2. Claim: For the equality function, $D(EQ_n) \geq n$

Proof. Let's use a fooling set to show the minimum deterministic CC of the equality function. Since we know from earlier that the communication matrix is the identity matrix, we can take as a fooling set all the elements that lie in the diagonal, or all (x, x) elements. Clearly, this is a fooling set – if we take any (x, x') in the set, the two strings are not equal so they give 0. The size of this fooling set is 2^n (because there are 2^n possible strings of length n) so we have placed a lower bound $D(f) \geq n$. \square

5.4.1.2 Set Disjointness

This is the "queen of communication problems." It's really hard to do and we wish it weren't. Consider the situation in which there are n elements in some shopping list, and Alice and Bob both buy some subset of these elements $X, Y \subset [n]$. The sets X, Y are disjoint if they are non-overlapping – if Alice and Bob didn't buy any of the same stuff. So, $DJS_n(X, Y) = 1 \implies X \cap Y = \emptyset$ (it's 0 if they bought the same thing).

We can create a fooling set FS as all possible subsets of $[n]$ and their compliments – $FS = \{(A, \bar{A})\}_{A \in \mathcal{N}}$. We know that these yield 1 by definition – nothing can overlap it's compliment. But consider comparing set A to some other set A' . Clearly, either A overlaps with A' or A' overlaps with \bar{A} . So either $f(A, A') = 0$ or $f(A', \bar{A}) = 0$. So each set needs its own monochromatic rectangle. thus, $|FS| = 2^n$ (the number of possible subsets of set of size n) and $D(DJS_n) \geq n$. We will return to this later to try to get a better lower bound.

5.5 Rank

5.5.1 Basics

The next method for finding lower bounds on DCC is using M_f directly, through its rank. But to use rank, we have to delve into some linear algebra. Don't be scared, Computer Scientists, you can do it!

Definition 5.5.1. We will use three definitions for the rank under \mathbb{F} of a matrix A :

1. $rank_{\mathbb{F}}(A) := \max$ number of linearly independent rows/columns
2. $rank_{\mathbb{F}}(A) := \min$ number of rows that span the rest of the rows in the matrix
3. $rank_{\mathbb{F}}(A) := r$ s.t. $A = \sum_{i=1}^r A_i, rank_{\mathbb{F}}(A_i) = 1$

The first definition is self-explanatory. The second definition, also, self explanatory. The third is basically saying that if you can split the matrix A into the sum of matrices with rank 1, then the rank of A is the number of these matrices.

We will also use two properties of rank and one fact, all of which we will not prove because this is not a mathematics class. Just trust me about this – they are all correct.

- a Subadditivity: $rank(A + B) \leq rank(A) + rank(B)$
- b Subproductability?: $rank(A \times B) \leq \min\{rank(A), rank(B)\}$
- c The rank of Boolean matrices is maximized over \mathbb{R} (as opposed to over, for example, GF2)

What does this "under \mathbb{F} " thing mean? Basically, if we are working in mod 2, for instance, the rank of a matrix may change – although 0 is not a multiple of 1 in \mathbb{R} , it is in GF2 – it's 2×1 . That's why it makes intuitive sense that the rank would be maximized over the real numbers (when all the entries in the matrix are 0 or 1 – there's no way that adding linear combinations of rows would "spill over" and equal 0 or 1, creating a linear dependence where in the real numbers there would be none).

Lemma 5.5.2. $\forall f, \forall \mathbb{F}, \log(\text{rank}_{\mathbb{F}}(M_f)) \leq D(f) \leq \text{rank}_{\mathbb{F}}(M_f)$

Well, you saw this coming – why would we bring up rank in this section if it had nothing to do with finding minimum DCC? The bottom limit kind of makes sense. But the top limit is somewhat surprising! Let's prove the lower limit first:

Proof. First we prove the lower Limit:

We prove this using the third definition of rank – the one with the sum of matrices. Suppose that \exists some protocol Π where $|\Pi| = C$, where we suppose this is the minimal communication complexity Π , and thus the minimum C for some function f . This induces a partition of M_f into $\leq 2^C$ monochromatic rectangles. Each of these rectangles has rank 1, because they are all the same number (so we can, for example, take the first row and show that this row spans the rest of the rows). We know that we can get our original M_f back by adding together all these rectangles (to see this, expand each rectangle to be the size of M_f , with all entries 0 that are not in the original rectangle. Add them. Voila!) So:

$$M_f = \sum_{i=1}^{2^C} M_i, \text{rank}_{\mathbb{R}}(M_i) = 1 \quad (5.3)$$

$$\text{rank}(M_f) = \text{rank} \left(\sum_{i=1}^{2^C} M_i \right) \quad (5.4)$$

$$\leq \sum_{i=1}^{2^C} \text{rank}(M_i) \quad (5.5)$$

$$= 2^C \quad (5.6)$$

where inequality (5) comes from the subadditive property of rank. Rearranging the following equation gives the inequality we are looking for:

$$\begin{aligned} \text{rank}(M_f) &\leq 2^C \\ \log(\text{rank}(M_f)) &\leq C \end{aligned}$$

So we can say that the minimum communication complexity, C , is at least as big as the log of the rank of the communication matrix M_f . Done! \square

5.5.2 Example – Inner Product

Let the function f be inner product: $P_n(X, Y) := \langle x, y \rangle \text{ mod } 2 = \sum_{i=1}^n x_i y_i$ (where the product in the second definition is done in GF2. So we obviously need to do our math in $\mathbb{F} = GF2$, not \mathbb{R}).

The claim is that $\text{rank}_{GF_2}(M_P) \leq n$. Let's show this. Naively, we might say, "Oh! Let's decompose M_f into matrices $M_i(x, y)$ where the matrix entry is 1 only when x and y are both 1 in the i th bit. We can see that this matrix is the inner product matrix projected on the i th coordinate, and to achieve the total inner product, we just need to add these. That is, $M_P(X, Y) = \langle x, y \rangle_2 = \sum_{i=1}^n x_i y_i$. Great! The rank, in any of these projected inner product matrices, is 1, so using the same process as above, we can say that the total matrix is a sum of n of these matrices (one for each bit in the input X or Y). But then we've shown that $D(f) \geq \log(n)$. Can we do better – get a tighter lower bound?

Yes, we can! let's look at the squared matrix, M_P^2 . The rank of M_P is at least as big as the rank of M_P^2 because of subproductness thingy that we defined as the second characteristic of rank. Let's look at a simple matrix to see what this squared matrix would look like:

$$M_P = \begin{matrix} & 00 & 01 & 10 & 11 \\ \begin{matrix} 00 \\ 01 \\ 10 \\ 11 \end{matrix} & \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 \end{pmatrix} \end{matrix}$$

In the squared matrix, we do simple matrix multiplication. We can represent this algebraically as:

$$M_P^2(X, Y) = \sum_{z \in \{0,1\}^n} \langle x, z \rangle \cdot \langle z, y \rangle$$

In words, the entry for (x, y) is the number of strings z (in the real numbers, now) where both $\langle x, z \rangle$ and $\langle z, y \rangle$ is 1. On the diagonal, where $x = y$, this is exactly half the strings (the number of strings where an odd number of 1 bits match up). So this accounts for 2^{n-1} strings. Where $x \neq y$, not on a diagonal, its a fourth of the strings, so 2^{n-2} .

So our matrix M_P^2 is a matrix with all 2^{n-1} on the diagonals and 2^{n-2} off the diagonals. We can see this as the linear combination of the 1s matrix and the identity matrix: specifically, 2^{n-2} 1 matrices and 2^n identity matrices. Each of these has rank 1 in \mathbb{R} , so total we have $\text{rank}(M_f) \geq \text{rank}(M_f^2) = 2^{n-2} + 2^n$. Thus,

$$D(f) \geq n - 1$$

which is a tighter bound than we had before.

5.5.3 Upper Bound

What about the other side of the inequality – the one claiming $D(f) \leq \text{rank}_{\mathbb{R}}(M_f)$? This is difficult to prove, so it takes a little leap of faith for us. Intuitively, we can see using the span definition of rank that if M_f is spanned by r , then the entire matrix is a linear combination of these r spanning rows. So instead of telling Bob the whole matrix, she can just communicate the coefficients of these linear combinations, with r coefficients describing X . For a Boolean matrix the coefficients should only be 0 or 1 – so the coefficients shouldn't be too large – each can be described with one bit. So Alice could potentially just communicate r bits to Bob to describe exactly what the rows of M_f look like.

Open Problem It is believed, via the Log-Rank Conjecture, that this upper bound can be much tighter: $\forall f, D(f) \leq \log^{O(1)}(\text{rank}(M_f))$. In 2013, Lovett proved $D(f) \leq \sqrt{\text{rank}(M_f)}$ but this right now is the best known limit (arXiv:1306.1877).

5.6 Application of Deterministic CC

Finally, we look at one application of this Deterministic CC for a cute problem, published by one of our very own, Yannakakis, in 1988. Consider the Clique-Independent Set Problem CIS_n .

Consider a publicly-known graph $G = (V, E), |V| = n$. Alice gets a clique C and Bob gets an independent set I . The goal is for them to decide whether their two sets intersect – that is, is there a node that is in both the clique and the independent set?

It is easy to see that if there is, it is just one point. Because if there were two points, they would need to be connected to each other to be in Alice's clique C , which means they couldn't both be in Bob's independent set I .

Theorem 5.6.1. $D(CIS_n) \leq O(\log^2(n))$

Proof. Consider the following protocol: Let $d(v)$ be the degree of any node in the entire graph G . Let S be the "checking set", which is initialized to all nodes, and has size m .

- * Alice looks for a $v \in C, \in S$ such that $d(v) \leq \frac{n}{2}$, tells this node to Bob. Bob checks if it is also in his I . If not, they constrain S to just neighbors of v .
- * Bob looks for a $v \in I, \in S$ such that $d(v) \geq \frac{n}{2}$, tells this node to Alice. Alice checks if it is also in her C . If not, they constrain S to just non-neighbors of v .
- * Alice and Bob stop when either they've found a matching node or run out of nodes to check.

Why does this work? Well, in Bob's case, he always constrains the new checking set to non-neighbors of v , meaning he never tampers with I . Alice constrains the new checking set to neighbors of v , so she doesn't tamper with C . If there is an intersecting node, it will be in both C and I , so it will never be thrown away or tampered with. On the other hand, by splitting nodes into those with high degree and low degree, each time Alice or Bob communicates, they limit the search space by at least a factor of 2. So in at most $\log(n)$ rounds, they will have determined the answer to their question.

Thus there are $\log(n)$ rounds, and each time, $\log(n)$ bits are communicated (to share a specific node, which presumably are indexed somehow in a way that both Alice and Bob know). Thus:

$$D(CIS_n(G)) \leq \log^2(n)$$

□