Communication Complexity and Applications

Lecturer: Toniann Pitassi

1 The Log Rank Conjecture

The famous logrank conjecture states that the deterministic communication complexity of any function f is polynomially related to the log of the rank of M_f. More formally, let $f : X \times Y \rightarrow \{0,1\}$, and let M_f be the Boolean communication matrix associated with f. Then

$$\mathsf{P}^{\mathsf{cc}}(f) = \log^{O(1)} \mathsf{rank}(M_f)$$

where the rank is taken over the reals.

One direction is quite easy to see: if f has a cost c deterministic protocol, then M_f can be partitioned into at most 2^c monochromatic subrectangles and therefore $\mathsf{rank}(M_f) \leq 2^c$,so the log of the rank is at most c.

The other direction is still wide open. The best that is known is the following upper bound due to Lovett.

$$\mathsf{P}^{\mathsf{cc}}(f) \leq O(\sqrt{\mathsf{rank}(f)} \cdot \log \mathsf{rank}(f)).$$

Nisan and Wigderson studied the following weaker version of the Log Rank Conjecture.

For a communication problem f with associated communication matrix M_f, let $\mathsf{mono}(M_f)$ equal the maximum size of a monochromatic subrectangle of M_f. That is, $\mathsf{mono}(M_f) = max_R \frac{|R|}{|M_f|}$ where the max is taken over all monochromatic subrectangles R of M_f. Since a deterministic protocol Π for f of cost c partitions M_f into monochromatic subrectangles, the average size of the monochromatic subrectangles in the partition is $\frac{|M_f|}{2^c}$, and therefore there must exist at least one monochromatic subrectangle, R of M_f such that $\frac{|R|}{|M_f|} \geq \frac{1}{2^c}$. Therefore, $\mathsf{P}^{\mathsf{cc}}(f) \geq \log \frac{1}{\mathsf{mono}(M_f)}$. Nisan and Wigderson made the following conjecture:

$$\log \frac{1}{\mathsf{mono}(M_f)} \leq \log^{O(1)} \mathsf{rank}(M_f).$$

The following theorem (also due to Nisan and Wigderson) proves that the above conjecture is actually equivalent to the Log Rank Conjecture.

Theorem 1 *Let $r = \mathsf{rank}(M_f)$. For any nondecreasing function δ of r, if $\mathsf{mono}(M_f) \geq \delta(r)$ then $\mathsf{P}^{\mathsf{cc}}(f) \leq O(\log^2 r + \log r \cdot \log \frac{1}{\delta(r)})$.*

If the Nisan-Wigderson conjecture is true, $\delta(r)$ is equal to $2^{-O(\log^k r)}$ for some constant $k > 0$, so by the above Theorem, $\mathsf{P}^{\mathsf{cc}}(f) \leq O(\log^{k+1}(r))$, and thus the Log Rank Conjecture is true.

Proof Let M_f be a matrix of rank r and assume that M_f has a monochromatic subrectangle R of size at least $\frac{\delta(r)}{|M_f|}$. Rearranging rows and columns, we will write M_f as $(RA \mid BC)$ where R the upper left corner of M_f, A is upper right, B is lower left and C is lower right.

Now since R is 1-monochromatic, the rank of RA is at most $\mathsf{rank}(A)$ plus 1, and similarly, the Then by subadditivity of rank ($\mathsf{rank}(A+B) \leq \mathsf{rank}(A) + \mathsf{rank}(B)$) we have: $r = \mathsf{rank}(M_f) = \mathsf{rank}((RA \mid 00) + (00 \mid BC)) \leq \mathsf{rank}(RA) + \mathsf{rank}(BC) \leq \mathsf{rank}(A)$

Assume without loss of generality that $\mathsf{rank}(A) \leq \mathsf{rank}(B)$ and therefore $\mathsf{rank}(A) \leq r/2 + 1$. Thus $\mathsf{rank}(RA) \leq r/2 + 2$.

The protocol is as follows. The row player will communicate one bit to say if their input is in the top half RA or in the bottom half, BC, and then the players will recurse on this half. (If $\mathsf{rank}(B)$ was greater than $\mathsf{rank}(A)$, then the column player would have communicated one bit to say if their input is in the left half or the right half of M_f.)

Let $L(m, r)$ denote the number of leaves of the above protocol where $|M_f| = m$ and $r = \mathsf{rank}(M_f)$. Then we get the following recurrence relation:

$$L(m, r) \leq L(m, \frac{r}{2} + 2) + L(m(1 - \delta(r)), r),$$

since if the row players input x is in the top half, then we can recurse on matrix of rank at most $r/2$, and otherwise when x is in the bottom half, then the rank may still be large but the size of the matrix drops to $m(1 - \delta(r))$.

Expanding the second term we get:

$$L(m, r) \quad \leq \quad L(m, \frac{r}{2} + 2) + L(m(1 - \delta(r)), \frac{r}{2} + 2) + L(m(1 - \delta(r))^2, \frac{r}{2} + 2)$$

Continuing to expand on the second term until every term has rank below approximately $r/2$, using the fact that $(1 - 1/x)^x$ is approximately $1/2$, after $\frac{\log m}{\delta(r)}$ expansions, m drops to a constant. Using the fact that $\delta(r)$ is nondecreasing as r increases (so we can replace $\delta(r')$ by $\delta(r)$ for $r' < r$), and since $L(1, r) = O(1)$ we get

$$L(m, r) \leq O(\frac{\log m}{\delta(r)} \cdot L(m, \frac{r}{2} + 2).$$

Now continuing $\log r$ times until r drops to a constant, and since $L(m, 1) = O(1)$, altogether we have $L(m, r) \leq (\frac{\log m}{\delta(r)})^{\log r}$. Now since protocols can always be balanced, this implies that there is a protocol of cost $\log(L(m, r))$. Thus:

$$\mathsf{P}^{\mathsf{cc}}(f) \leq O(\log r[\log m + \log(\frac{1}{\delta(r)})]) \leq O(\log^2 r + \log r \cdot \log(\frac{1}{\delta(r)}),$$

where the last inequality holds since $m \leq 2^r$.

1.1 Log Rank Conjecture and Zero Communication Protocols

Building on the above Nisan-Wigderson results, Gavinsky and Lovett showed that a seemingly much weaker conjecture actually is equivalent to the LRC. In order to explain this conjecture, we need to define another type of communication complexity measure, called PostBPP.

Definition ($\mathsf{PostBPP}^{\mathsf{cc}}$) Consider a randomized protocol Π for f that can output either 0 or 1 or a third value, $\perp$, that stands for "don't know" or "abort". Let $\frac{1}{2^k}$ be the probability that Π does not output abort. The PostBPP correctness condition requires that conditioned on not outputting abort, we have $Pr[\Pi(x,y) \neq f(x,y)] \leq 1/3$. The cost of such a protocol is $k+c$, where c is (worst-case) the protocol depth.

It is not hard to see that PostBPP protocols can be efficiently amplified, so we can replace the 1/3 with any other constant ϵ. That is, $\mathsf{PostBPP}^{\mathsf{cc}}_{\epsilon}(f) \leq O(\mathsf{PostBPP}^{\mathsf{cc}}_{1/3}(f) \cdot \log(\frac{1}{\epsilon})$.

An alternative and equivalent definition of the PostBPP communication complexity of f is the restriction of the above definition but where the players essentially do not communicate at all. More formally, let $\mathcal{R} = \{(R,b)\}$ be a set of labelled rectangles (so each R is a subrectangle, and $b \in \{0,1\}$ is the label). Given $\mathcal{R}$, on input (x,y), the players randomly sample (R,b) from $\mathcal{R}$. Then Alice (the row player who holds x) sends 1 iff $x \in rows(R)$ and similarly Bob sends 1 iff $y \in cols(R)$. If Alice and Bob both send 1's then the protocol outputs b and otherwise the protocol aborts.

Gavinsky and Lovett proved the following theorem.

Theorem 2

$$\mathsf{P}^{\mathsf{cc}}(f) \leq O(\mathsf{PostBPP}^{\mathsf{cc}}(f) \cdot \log^2 \mathsf{rank}(M_f)).$$

This implies the following corollary, stating that if the cost of $\mathsf{PostBPP}^{\mathsf{cc}}$ protocols (which are much more powerful then deterministic protocols) is bounded by a polynomial in the log of the rank of M_f, then the LRC is true.

Corollary 3 *If for every f, $\mathsf{PostBPP}^{\mathsf{cc}}(f) \leq \log^{O(1)} \mathsf{rank}(M_f)$, then the Log Rank Conjecture is true.*

To prove the above theorem, by Nisan-Wigderson, it suffices to prove that any PostBPP protocol with sufficiently small error has a large monochromatic rectangle:

Lemma 4 *Let $r = \mathsf{rank}(M_f)$ and let $\mathsf{PostBPP}^{\mathsf{cc}}_{\epsilon}(f) = c$ where $\epsilon = \frac{1}{8r}$. Then there exists a monochromatic rectangle R such that $|R| \geq \frac{1}{16}2^{-c}|M_r|$.*

Given the above Lemma, the proof of the above theorem follows by setting $\delta(r) = \frac{1}{16}2^{-c}$ and applying Nisan-Wigderson's theorem. It follows that

$$\begin{aligned}
\mathsf{P}^{\mathsf{cc}}(f) &\leq O(\log^2 r + \log r \cdot c) \\
&\leq O(\log^2 r + \log r(\mathsf{PostBPP}^{\mathsf{cc}}(f) \cdot \log(\frac{1}{\epsilon}))) \\
&\leq O(\log^2 r + \log^2 r \cdot \mathsf{PostBPP}^{\mathsf{cc}}(f))
\end{aligned}$$

It is left to prove the above Lemma. First we will argue (using the alternative definition of PostBPP protocols) that there exists a large rectangle R' that is nearly monochromatic. Then we will prove that since M_f has rank at most r, that R' must contain a large, fully monochromatic subrectangle R. For the first step, using the alternative definition of PostBPP, there exists a collection $\mathcal{R}$ of labelled rectangles such that: (1) conditioned on (x,y) lying within at least one of rectangles of $\mathcal{R}$, the probability that the protocol errs is at most ϵ. and (2) the probability that (x, y) lies within one of the rectangles of $\mathcal{R}$ is at least 2^{-c}. Thus by averaging, there must exist some labelled rectangle (R', b) in $\mathcal{R}$ that is large and nearly b-monochomatic. That is, there must exist $(R', b) \in \mathcal{R}$ satisfying: (1) the size of R' is at least $\frac{1}{2} \cdot 2^{-c}|M_f|$ and (2) the fraction of elements $(x, y) \in R'$ such that $f(x, y) \neq b$ is at most $\frac{1}{2\epsilon}$.

Now it is left to prove the second step, that R' must contain a large b-monochromatic subrectangle. To find the monochromatic subrectangle of R' we will first remove from R' any rows that make significantly more errors than average. Let $rows(R')$ be the rows of R'. Call a row $x \in rows(R')$ bad if number of inputs (x, y), $y \in cols(R')$ that are not b-monochromatic is greater than 2 times the expected value. By Markov, The probability that x is bad is at most $\frac{1}{2}$. Now let A be the set of all $x \in rows(R')$ that are not bad; from the above we know that $|A| \geq |rows(R')|/2$.

Now consider the submatrix, Q, of R' restricted to the rows A: $Q = A \times cols(R')$. Since M_r has rank at most r, Q also has rank at most r. Let $x_1, \ldots, x_r$ be a basis for Q – that is, the x_i's are in A and span Q. For each basis element x_i, we define the following set of bad inputs corresponding to x_i:

$$B_i = \{y \in cols(R') \mid f(x_i, y) \neq b\}.$$

Since all x_i's are good, $|B_i| \leq 4\epsilon|cols(R')|$ for all $i = 1, \ldots, r$. Define $B = cols(R') - \cup_{i=1}^{r} B_i$. Thus (using the fact that $\epsilon = \frac{1}{8r}$),

$$|B| \geq |cols(R')|(1 - 4\epsilon r) = |cols(R')|/2.$$

Therefore, we have shown that every row of the matrix $A \times B$ is either 0-monochromatic or 1-monochromatic. Pick $b \in \{0, 1\}$ to be the more popular monochromatic value, and let R be the set of all rows of $A \times B$ that are b-monochromatic. Thus R is a monochromatic subrectangle of R' of size at least

$$\frac{1}{2}|A| \times |B| \geq \frac{1}{2}(\frac{1}{2} \cdot |rows(R')|) \cdot (\frac{1}{2} \cdot |cols(R')|) = \frac{1}{8}|R'| \geq \frac{1}{16}2^{-c}|M_f|.$$

Thus to prove the LRC it suffices to prove the much weaker statement that any M_f of low rank has a small PostBPP protocol. Note that PostBPP protocols are quite powerful as we have:

$$\mathsf{P}^{\mathsf{cc}}(f) \geq \mathsf{BPP}^{\mathsf{cc}}(f) \geq \mathsf{PostBPP}^{\mathsf{cc}}(f).$$

If on the other hand the LRC conjecture is false, then there are low rank matrices that require high communication cost even in an extremely strong model of computation, which would be highly interesting and would show that rank is a very unstable measure, not always correlated with other reasonable properties of random versus structured matrices/functions.