Computational Complexity II Course Instructor: V. Arvind

Monotone Circuit Lower Bounds

Lecturer: V. Arvind *Scribe: Ramprasad Saptharishi*

1 Motivation

The holy grail for computer science has been trying to somehow show that $P \neq NP$. And another problem that is equally intriguing is to show that $NP \not\subseteq P/poly$, trying to find circuit (over $\wedge, \vee, \neg$) lower bounds for problems in NP.

But suppose we were able to drop the $\neg$ gate from the basis, we would be able to compute only monotone functions, but can we show some monotone circuit lower bounds for some hard problem in NP? In this lecture we shall see that the $CLIQUE$ function requires super-polynomial sized monotone circuits.

2 Preliminaries

Definition 1. *A boolean function f is said to be a monotone boolean function if $x \leq y \implies f(x) \leq f(y)$ where $x \leq y$ is by the bit-wise partial order on strings.*

A circuit is said to be monotone if it consists of only $\vee$ and $\wedge$ gates.

Claim 2. *A boolean function is monotone if and only if it can be computed by a monotone circuit*

Proof. One direction is clear, if a function can be evaluated by a monotone circuit, it has to be monotone. The other direction as well is fairly obvious.

Clearly $f(O^n) = 0$ and $f(1^n) = 1$ (otherwise it has to be a constant function, which has a trivial monotone circuit). Hence on every path from O^n to 1^n on the boolean hypercube, take the first string x such that $f(x) = 1$. We now just need to have an $\wedge$ gate with all the 1 positions of x as input, and this has to be $\vee$-ed over the x on every path. □

Definition 3.

$$CLIQUE_{k,n} = \{(x_1, \cdots, x_{\binom{n}{2}}) | \textit{the graph induced by them has a } k \textit{ clique}\}$$

Theorem 4 (Razbarov)**.** *$CLIQUE_{k,n}$ requires monotone circuits of size $O(n^{\sqrt{k}})$ for $k \leq n^{\frac{1}{4}}$*

This could have been an attack on the $NP \nsubseteq P/poly$ problem if the monotone complexity and actual complexity of any function was polynomially related. But unfortunately this isn't true for general functions, one could tweak the perfect matching problem to have a problem in P to have exponential monotone circuit complexity.

However these are true for the so called "Slice Problems" where the functional value is defined by a singe slice on the hyper cube. That is, there exists a k such that for any string with weight more than k has 1 as the output, and any string with weight less than k has 0 as the output.

The proof of Razborov's theorem will be achieved by approximating circuits by "clique indicators", and showing that the approximated circuit will have to make a lot of errors and also that every gate as such makes small errors (thereby meaning that there should be a lot of gates that collectively make a lot of error).

3 Proof of Razborov' Theorem

Definition 5. *A clique indicator I_X is a simple circuit such that I_X is 1 on graphs that have the subset $X \subseteq [n]$ as a clique.*

A (m, l) approximator is $\vee_{i=1}^{r} I_{X_i}$ such that $|X_i| \leq l$ for every i and that $r \leq m$, (i.e) it is the OR of at most m clique indicators each of whose size is bounded by l.

We shall approximate the circuit for $CLIQUE_{k,n}$ with (m, l) indicators and measure the errors on specific types of instances:

- Positive Tests: Graphs on n vertices with just a k-clique. There are $\binom{n}{k}$ such graphs.
- Negative tests: $(k-1)$ colourable graphs or $(k-1)$-partite graphs. If these are counted with the colouring, then there are $(k-1)^n$ of them.

3.1 Building the Approximate Circuit

The circuit shall be built bottom up, each input x_i is already a clique indicator so there is nothing to be done there.

Suppose at some internal node of the circuit, say an $\vee$ gate, has two inputs A and B whose approximators have already been defined.

$$A = \bigvee_{i=1}^{r} I_{X_i}, \quad |X_i| \leq l, r \leq m$$

$$B = \bigvee_{j=1}^{s} I_{Y_j}, \quad |Y_j| \leq l, s \leq m$$

We could approximate this $\vee$ gate by just $A \vee B$ but that could potentially give us a $(l, 2m)$ approximator, while want to stay at (l, m). The following lemma by Erdös and Rado comes to our rescue.

Lemma 6 (Sunflower Lemma)**.** *Suppose $\mathfrak{F} \subseteq 2^{[n]}$ such that for all $S \in \mathfrak{F}, |S| \leq l$. Then if $|\mathfrak{F}| > (p-1)^l \cdot l!$, then there exists a "p-petalled sunflower", that is, there exists $Z_1, \cdots Z_p \in \mathfrak{F}$ such that $Z_i \cap Z_j = Z$ for all $i \neq j$.*

Proof. The proof will just be an induction on l.

When $l = 1$, then the core is the null set and we have p disjoint sets as the p petals.

As for the inductive step, pick a maximal disjoint collection of subsets from $\mathfrak{F}$ say $Z_1, \cdots, Z_m$. If $m > p$ there is nothing to be done since we already have a p-petalled sunflower with an empty core.

The other case is when $m < p$. Let $Z = \bigcup Z_i$, and by maximality of the collection of sets, every set in $\mathfrak{F}$ must intersect with Z and $|Z| \leq m \cdot l \leq (p-1) \cdot l$. Hence an element of Z is contained on an average in $\frac{|\mathfrak{F}|}{|Z|} > \frac{(p-1)^l \cdot l!}{(p-1) \cdot l} = (p-1)^{l-1} \cdot (l-1)!$ many sets of $\mathfrak{F}$, and hence there exists an element $x \in Z$ such that it is contained in more than $(p-1)^{l-1} \cdot (l-1)!$ sets of $\mathfrak{F}$. Now consider the collection of those sets that contain x and remove x from it. By induction you have a sunflower of p petals, put back x into the core and we have got the required sunflower $\square$

We shall have $\mathfrak{F} = \{X_1, X_2, \cdots, X_r, Y_1, \cdots, Y_s\}$ and choose m such that $m = (p-1)^l \cdot l!$, hence if we have more than m clique indicators, we "pluck" the sunflower. If $X_1, \cdots, X_p$ form the p-petalled sunflower, replace all of them by the common intersection $X_1 \cap X_2$. And hence, with every plucking, the number of clique indicators go down by 1, repeat this process until you have at most m clique indicators.

Thus we have approximated an $\vee$ gate, we shall call it $A \sqcup B$.

Approximating an $\wedge$ gate requires a little more work. $A \wedge B$ looks like

$$\begin{aligned} A \wedge B &= \left(\bigvee_{i=1}^{r} I_{X_i} \right) \wedge \left(\bigvee_{j=1}^{s} I_{Y_j} \right) \\ &= \bigvee_{i=1}^{r} \bigvee_{j=1}^{s} \left(I_{X_i} \wedge I_{Y_j} \right) \end{aligned}$$

We now have two problems, firstly $I_{X_i} \wedge I_{Y_j}$ isn't even a clique indicator. This can be tackled by just replacing $I_{X_i} \wedge I_{Y_j}$ by $I_{X_i \vee Y_j}$. If $|X_i + Y_j| > l$, just drop this clique indicator. As for our second problem of having a union of $2m$ clique indicators, apply the plucking as before to reduce this number to m. This approximator will be denoted by $A \sqcap B$

The approximator of the output gate will be our approximate circuit C'.

3.2 Approximated circuit makes lots of errors

Lemma 7. *Either C' is 0 on all positive tests or the number of negative tests on which C' is 1 is at least*

$$\left(1 - \frac{\binom{l}{2}}{k-1} \right) (k-1)^n$$

Proof. The only way C' is 0 on all positive instance is when our $A \sqcap B$ approximator throws out all clique indicators because their sizes are greater than l. Hence, if it is non-zero on positive test cases,

$$C' = \bigvee_{i=1}^{r} I_{X_i}$$

where $0 \neq r \leq m$ and $|X_i| \leq l$. Pick a negative instance at random among the $(k-1)^n$ choices. The probability that the approximator outputs a 1 is clearly upper bounded by the probability that I_{X_1} outputs a 1, which is equal to $1 - \Pr[I_{X_1} = 0]$. I_{X_1} outputs a 0 if and only if two vertices of X_1 lie in the same partition of the random negative test, and this happens with probability

$$\frac{\binom{|X_i|}{2}}{k-1} \leq \frac{\binom{l}{2}}{k-1}$$

Hence, the number of errors made on negative tests is lower bounded as claimed in the lemma. $\square$

3.3 Each gate makes few errors on positive tests

Lemma 8. *The number of positive tests on which C' fails is at most*

$$size(C) \cdot m^2 \cdot \binom{n-l-1}{k-l-1}$$

Proof. We shall consider the errors introduced by the approximator at a single gate, and then apply the union bound to get the bound claimed.

If $g = A \vee B$, then our construction for the approximator for g involves taking a simple $\vee$ (which doesn't introduce any error) and then repeatedly plucking until we get down our number of clique indicators. The plucking operation replaces a larger clique indicator I_{X_i} by I_X with $X \subseteq X_i$ and hence will accept only more graphs and cannot make the approximator answer a 0. Hence $A \sqcup B$ introduces no errors on positive tests.

Suppose $g = A \wedge B$, our first step was to replace $I_{X_i} \wedge I_{Y_j}$ by $I_{X_i \cup Y_j}$. These two functions behave the same on positive tests and hence won't introduce any errors. The second step was to drop indicators of size greater than l. Dropping such clique indicators may make the approximator err on those positive graphs that contain a clique on these $l+1$ or more places, and there are at most $\binom{n-l-1}{k-l-1}$ of them and there are at most m^2 clique indicators, giving us the bound claimed. □

3.4 Each gate makes few errors on negative tests

Lemma 9. *The number of negative tests on which C' fails is at most*

$$size(C) \cdot m^2 \cdot \left(\frac{\binom{l}{2}}{k-1} \right)^p \cdot (k-1)^n$$

Proof. Again, we shall analyse the errors introduced at each gate.

If $g = A \vee B$ then the errors have to be introduced at the plucking operations. We need to consider the negative tests that are accepted after plucking which were rejected before. Pick the $(k-1)$-partition randomly among the $(k-1)^n$ partitions, and let G be the resulting negative graph. For any sunflower $Z_1, \cdots, Z_p$, we estimate the probability that I_{Z_i} outputs 0 on G whereas I_Z outputs a 1.

$$\begin{aligned} Pr\left[I_Z = 1 \wedge \left(\bigvee I_{Z_i} = 0\right)\right] &\leq \Pr\left[\bigvee I_{Z_i} = 0 | I_Z = 1\right] \\ &= \prod_{i=1}^{p} \Pr\left[I_{Z_i} = 0 | I_Z = 1\right] \\ &\leq \prod \Pr[I_{Z_i} = 0] \\ &\leq \left(\frac{\binom{l}{2}}{k-1}\right)^p \end{aligned}$$

The first line follows from the definition of conditional probability. The second line is true because the petals of the sunflower are disjoint and hence the probabilities are independent. The third is true because "Z_i is not a clique" is less likely to happen given the fact that Z is a clique[1]. The fourth line follows because Z_i is not a clique if and only if two vertices of it fall in the same partition.

And since there are at most m plucks, the desired bound follows if g is an $\vee$ gate.

Suppose $g = A \wedge B$, the step where we replace $I_{X_i} \wedge I_{Y_j}$ by $I_{X_i \vee Y_j}$ does not cause any graph that was rejected before to be accepted now. The step of dropping large clique indicators also doesn't not cause additional negative graphs to be accepted. It's again the plucking process that creates the trouble, and that goes through the same analysis as given for the $\vee$ gate. And since there are at most m^2 packing's, the claimed bound is obtained. $\square$

3.5 Choosing parameters

If we choose $l = \lfloor\sqrt{k}\rfloor$, $p = \lceil\sqrt{k}\log n\rceil$ and $m = (p-1)^l \cdot l!$, the lemmas would then show that

$$size(C) = n^{\Omega(\sqrt{k})}$$

The details are left to the reader.

[1] if this intuition is not clear, the reader should work out the easy details