1 The Razborov monotone circuit lower bound

Theorem 1 (Razborov-Alon-Boppana)**.** *Let $3 \leq k \leq n^{1/4}$, then the circuit complexity for $Clique_{n,k}$ is $n^{\Omega(\sqrt{k})}$.*

Note: If $k = n^{1/4}$, then we have $n^{\Omega(n^{1/8})} \geq 2^{n^{\frac{1}{8}}}$. In other words, this is an exponential lowerbound.

1.1 Parameters

Throughout we have the following parameters (not all terms have been defined yet):

- n, the number of nodes in an undirected graph
- $\binom{n}{2}$, the number of inputs to our circuit, $x_{\{i,j\}}$. Think of these as answering True/False whether or not there is an edge connecting i and j in a graph.
- k, the size of a clique
- l, the maximum size of a clique indicator
- m the maximum number of clique indicators in an approximator
- p, the number of petals in a sunflower

Definition 2. *A* positive test graph *is a graph that is a k-clique, and no other edges are present. A* negative test graph *is a graph formed by a $(k-1)$-coloring adding edges between all vertices of different colors.*

Our goal is to show that "too-small" monotone circuits will either assign 0 to many positie test graphs, or assign 1 to many negative test graphs. To show this, we will work with approximations of monotone circuits. The approximation is performed by sequentially replacing gates in an $\wedge$ and $\vee$ circuit, with approximate gates $\sqcap$ and $\sqcup$.

Definition 3. *A* clique indicator $\lceil X \rceil$ *is a Boolean function given by a set $X \subseteq \{1, \ldots, n\}$, $|X| \leq l$.*

$$\lceil X \rceil (x_1, \ldots, x_{\binom{n}{2}}) := \begin{cases} 1 & \textit{if } X \textit{ is a clique in this graph} \\ 0 & \textit{otherwise.} \end{cases}$$

Definition 4. *A is an approximator if it is a disjunction of no more than m clique indicators.*

$$A = \bigvee_{i=1}^{m'} \lceil X_i \rceil$$

for $m' \le m$ and $|X_i| \le l$.

1.2 Construction

Let C be a monotone circuit. The approximator $\widetilde{C}$ for C is defined inductively on the size of C. To define the construction, we will make use of plucking sunflowers.

Definition 5. *A* sunflower *is a collection of sets, $z_1, \ldots, z_p$ such that pairwise intersections are constant. In other words, there is some z_0 such that if $i \neq j$, then $z_i \cap z_j = z_0$. This z_0 is called the* center*, and p is the number of petals. To pluck a sunflower is to replace $z_1, \ldots, z_p$ with z_0.*

The construction is as follows.

Base case: if C has 0 gates, then C is just $x_{\{i,j\}}$, so let $X := \{i, j\}$, and $\widetilde{C}$ is defined to be $\lceil X \rceil$. Notice that this approximator gets the thing it is approximating right.

Induction step ($\vee$): In this case $C = C_1 \vee C_2$. Let $A = \bigvee_{i=1}^{m_1} \lceil X_i \rceil$ and $B = \bigvee_{i=1}^{m_2} \lceil Y_i \rceil$ be the approximators for C_1 and C_2. We would like our approximator $A \sqcup B$ to be

$$\bigvee_{i=1}^{m_1} \lceil X_i \rceil \vee \bigvee_{i=1}^{m_2} \lceil Y_i \rceil.$$

If $m_1 + m_2 \le m$, then this is our approximator. Otherwise, we need to pluck sunflowers until we have no more than m clique indicators remaining.

Induction step ($\wedge$) In this case, $C = C_1 \wedge C_2$. Let $A = \bigvee_{i=1}^{m_1} \lceil X_i \rceil$ and $B = \bigvee_{i=1}^{m_2} \lceil Y_i \rceil$ be the approximators for C_1 and C_2, as before. We will form $A \sqcap B$ in stages:

1. Start by considering:
$$\bigvee_{i=1}^{m_1} \bigvee_{j=1}^{m_2} \lceil X_i \rceil \wedge \lceil Y_i \rceil.$$
2. Replace this with:
$$\bigvee_{i=1}^{m_1} \bigvee_{j=1}^{m_2} \lceil X_i \cup Y_i \rceil.$$

3. Discard any $\lceil X_i \cup Y_j \rceil$ with $|X_i \cup Y_j| > l$.

4. There are up to m^2 many clique indicators left, so repeatedly pluck sunflowers until fewer than m are left.

In order to use the above construction, we need a guarantee that sunflowers exist. This is achieved by the next lemma.

Lemma 6 (Sunflower Lemma)**.** *Let $p \geq 2$. Suppose Z is a collection of sets of size no bigger than l. If $|Z| > (p-1)^l l!$, then Z contains a sunflower, $z_1, \ldots, z_p \in Z$ with p petals.*

Proof. We induct on l.

For $l = 1$, $|Z| > (p-1)^l l! \geq p-1$. Take $z_1, \ldots, z_p$ to be distinct elements of Z. They form a sunflower with empty center.

For $l \geq 2$: Choose a maximal set $z_1, \ldots z_t$ of Z with empty center. If $t \geq p$, then we're done (we have a sunflower with empty center). So assume $t < p$. Thus we have,

$$\left| \bigcup_{i=1}^{t} z_i \right| \leq tl \leq (p-1)l.$$

We know that every element $z \in Z$ intersects this union (because $z_1, \ldots, z_p$ are maximal pairwise disjoint). So some $x \in \bigcup_{i=1}^{t} z_i$ intersects fraction $\frac{1}{(p-1)l}$ members of Z.

Fix some such x. Let $F := \{z \in Z : \ x \in z\}$. Then $|F| \geq \frac{|Z|}{(p-1)l} > (p-1)^{l-1}(l-1)!$. Let $G := \{z - \{x\} : \ z \in F\}$. G satisfies the hypotheses of lemma for $l-1$, and has more than $(p-1)^{l-1}(l-1)!$ members each of size $\leq l-1$. By the induction hypothesis, there is a p petal sunflower, $T_1, \ldots, T_p$ of G with center T_0. But then $T_1 \cup \{x\}, T_2 \cup \{x\}, \ldots, T_p \cup \{x\}$ is a p petal sunflower in Z. $\square$

1.3 The approximators do what we want

At this point we have two subgoals:

Subgoal 1: We want to show that any monotone circuit C that is "too small" is closely approximated by its approximator $\widetilde{C}$ on both positive and negative test graphs.

Subgoal 2: We want to show that any approximator either makes a lot of errors (outputs 0) on positive test graphs, or makes a lot of errors (outputs 1) on negative test graphs.

Lemma 7 (Subgoal 2). *Let $A = \bigvee_{i=1}^{m'} \lceil X_i \rceil$ be an approximator. Then either (i) A outputs 0 on all graphs or (ii) A outputs 1 (accepts) more than*

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

negative test graphs.

Note: there are $(k-1)^n$ many negative test graphs, because we describe a negative test graph by a $(k-1)$ coloring (so two negatives test graphs that are isomorphic but colored differently are counted as distinct).

Proof. If $m' = 0$, then A is the empty disjunction (in other words, 0). If $m' \geq 1$, consider just some $\lceil X_j \rceil$, $|X_j| \leq l$. We claim any $\lceil X_j \rceil$ accepts a randomly chosen negative test graph with probability at least

$$1 - \frac{\binom{l}{2}}{k-1}.$$

In other words, that $\lceil X_j \rceil$ rejects a negative test graph with probability less than $\binom{l}{2}/(k-1)$. But the probability that $\lceil X_j \rceil$ rejects a negative test graph is the same as the probability that there exists an x, y in X_j and x and y are assigned the same color. The probability that any two nodes receive the same color is $1/(k-1)$, so by the union bound we have

$$\Pr[\exists x, y \in X_j \text{ and } x, y \text{ are same color}] \leq \frac{\binom{l}{2}}{k-1}$$

which completes the proof. □

Lemma 8 (Subgoal 1, part i). *The number of positive test graphs which are accepted by C, but rejected by its approximator, $\widetilde{C}$ is less than or equal to size$(C)m^2\binom{n-l-1}{k-l-1}$.*

Proof. We'll show that each gate of C contributes no more than $n^2\binom{n-l-1}{k-l-1}$ such positive test graphs. We argue by induction. For the base case, (in other words on the inputs) approximators are the same as the circuits they approximate, so these contribute no positive test graphs rejected by $\widetilde{C}$. For the $\vee$ induction step: We want to count the number of positive test graphs which are accepted by $\widetilde{C}_1 \vee \widetilde{C}_2$ but are rejected by $\widetilde{C}_1 \sqcup \widetilde{C}_2$. But any test graph accepted by $\widetilde{C}_1 \vee \widetilde{C}_2$ is also accepted by $\widetilde{C}_1 \sqcup \widetilde{C}_2$. This is because these are the same, accepted in the approximator, we replace a certain number of sunflowers by their centers. This is okay, though, because a disjunction of a sunflower has the same truth value as the center on positive test graphs.

For the $\wedge$ induction step: We need to count the number of positive test graphs for which $\widetilde{C}_1 \wedge \widetilde{C}_2$ is true, but $\widetilde{C}_1 \sqcap \widetilde{C}_2$ is false. At the stage where we pass from $\lceil X_i \rceil \wedge \lceil Y_j \rceil$ to $\lceil X_i \cup Y_j \rceil$, there are no new rejected positive test graphs, because these functions have the same truth value on positive test graphs. But from here, we discard any $\lceil X_i \cup Y_j \rceil$ for which $|X_i \cup Y_j| > l$, and this can cause the approximator to reject a positive test graph. There are no more than $\binom{n-l-1}{k-l-1}$ such positive test graphs for X_i and Y_j, and there are at most m^2 such $\lceil X_i \cup Y_j \rceil$, so we're done.

$\square$

Lemma 9 (Subgoal 1, part ii)**.** *The number of negative test graphs (which are rejected by C) that are accepted by $\widetilde{C}$ is less than or equal to*

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

Proof. The argument is similar to the last lemma. We consider each gate passing to its approximator, and counting the number of erroneous negative test graphs that could be introduced at each gate.

First, consider $C = C_1 \vee C_2$. Say $\widetilde{C}_1 = \bigvee_{i=1}^{m_1} \lceil X_i \rceil$, $\widetilde{C}_2 = \bigvee_{j=1}^{m_2} \lceil Y_j \rceil$. Now we ask: What negative test graphs are rejected by $\widetilde{C}_1 \vee \widetilde{C}_2$ but not by $\widetilde{C}$? Let $z_1, \ldots, z_p$ be a sunflower. We want to bound the number of negative test graphs rejected by $\bigvee_{i=1}^{p} \lceil z_i \rceil$, but accepted by $\lceil z_0 \rceil$, where z_0 is the center of the sunflower.

We'll compute the probability that a randomly chosen negative test graph has this property. Say that a collection of vertices is "properly colored (p.c.)" if all vertices receive different colors. A negative test graph has a clique on a set of vertices exactly when that set of vertices is properly colored. So then we want to bound

$$\Pr[z_0 \text{ is p.c. and } z_1, \ldots, z_p \text{ are each not p.c. }].$$

This is less than or equal to

$$\Pr[z_1, \ldots, z_p \text{ are each not p.c. } \mid z_0 \text{ is p.c. }]$$

By independence, this is less than or equal to

$$\prod_{i=1}^{p} \Pr[z_i \text{ is not p.c. } \mid z_0 \text{ is p.c. }] \leq \prod_{i=1}^{p} \Pr[z_i \text{ is not p.c.}]$$

$$\leq \prod_{i=1}^{p} \binom{l}{2} \frac{1}{k-1} = \left(\binom{l}{2} \frac{1}{k-1} \right)^p$$

For the $\wedge$ case, this is similar. Discarding $\lceil X_i \cup Y_j \rceil$ is ok, because we never accept any new graphs. And there are fewer than m^2 many pluckings, so we use the same bounds as above. $\square$

We are now ready to prove the theorem.

Theorem 10 (Razborov-Alon-Boppana)**.** *Let $3 \leq k \leq n^{1/4}$, then the circuit complexity for Clique$_{n,k}$ is $n^{\Omega(\sqrt{k})}$.*

Proof. Let C be a circuit for Clique$_{n,k}$. Let $l = \sqrt{k}$, $p = \lceil 10\sqrt{k}\log n \rceil$, $m = (p-1)^l l!$. By Lemma 2, we have two possible situations: either (i) $\widetilde{C}$ rejects all positive test graphs or (ii) $\widetilde{C}$ accepts at least $\left(1 - \frac{\binom{l}{2}}{k-1}\right)(k-1)^n$ many negative graphs.

If we are in the first situation, then by Lemma 3 we know that

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

If we are in the second situation, then by Lemma 4, we know that

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

At this point we are done. To see this, we do some algebra.

For the first case,

$$\frac{\binom{n}{k}}{\binom{n-l-1}{k-l-1}} = \frac{n}{k}\frac{n-1}{k-1}\cdots\frac{n-l}{k-l} \geq \left(\frac{n-\sqrt{k}}{k}\right)^{\sqrt{k}}$$

Also,

$$m = (p-1)^l l! = (p-1)^{\sqrt{k}}(\sqrt{k})!$$
$$\leq (10\sqrt{k}\log n)^{\sqrt{k}} \frac{\sqrt{k}^{\sqrt{k}}}{e^{\sqrt{k}}}$$
$$= k^{\sqrt{k}} \left(\frac{10}{e}\log n\right)^{\sqrt{k}}$$

Hence,

$$m^2 \leq k^{2\sqrt{k}} \left(\frac{10}{e} \log n \right)^{2\sqrt{k}}.$$

Therefore,

$$\text{size}(C) \geq \left(\frac{n - \sqrt{k}}{k} \right)^{\sqrt{k}} \frac{1}{k^{2\sqrt{k}} \left(\frac{10}{e} \log n \right)^{2\sqrt{k}}}$$

$$\text{size}(C) \geq \frac{(n - \sqrt{k})^{\sqrt{k}}}{k^{3\sqrt{k}} \left(\frac{10}{e} \log n \right)^{2\sqrt{k}}}$$

Since $k \leq n^{1/4}$, we have

$$\text{size}(C) \geq \frac{(n - \sqrt{k})^{\sqrt{k}}}{n^{3/4\sqrt{k}} \left(\frac{10}{e} \log n \right)^{2\sqrt{k}}}$$

Therefore $\text{size}(C) \geq n^{\Omega(\sqrt{k})}$.

In the other situation, we have

$$\frac{\binom{l}{2}}{k-1} < \frac{1}{2}$$

since $l = \sqrt{k}$. So then,

$$\text{size}(C) m^2 \frac{1}{2^p} \geq \frac{1}{2}$$

Hence,

$$\text{size}(C) \geq \frac{2^{p-1}}{m^2} \geq \frac{2^{10\sqrt{k} \log n}}{k^{2\sqrt{k}} \left(\frac{10}{e} \log n \right)^{2\sqrt{k}}} = \frac{n^{10\sqrt{k}}}{n^{\sqrt{k}/2} \left(\frac{10}{e} \log n \right)^{2\sqrt{k}}}$$

Thus $\text{size}(C) \geq n^{\Omega(\sqrt{k})}$. $\square$

2 Constant Depth Circuits

In this setting we allow unbounded fan-in $\wedge$'s, $\vee$'s, and negations only on variables. Typically we measure complexity of constant depth circuits by counting only the number of $\wedge$'s and $\vee$'s.

We could use other connectives. Other possibilities are, unbounded fan-in parity gates, unbounded fan-in mod k gates (i.e. gates that output 1 if the number of true inputs is 0 mod k), Majority gates, and Threshold gates.

Definition 11 (Π_k- and Σ_k-circuits and formulas)**.** *We definie these inductively. A Π_1-circuit is a conjunction of literals. A Σ_1-circuit is a disjunction of literals. A Π_{k+1}-circuit is a circuit with output gate $\wedge$, and all its inputs are Σ_k-circuits. A Σ_{k+1}-circuit is defined dually. Σ_k-formulas and Π_k-formulas are defined in the same way, replacing everywhere the word circuit with formula above.*

Proposition 12. *Any depth d circuit C of $\bigvee$'s, $\bigwedge$'s, and n literals can be converted into a Π_{d+1}-circuit and Σ_{d+1} circuit of size less than or equal to*

$$d \cdot \mathrm{size}(C) + d \cdot n$$

Proof. The idea is that if we have the output of one gate feed into another, and the gates are of the same type, we merge the inuts of the first gate with the second, and we duplicate gates as needed to make sure the levels work out. □

Theorem 13. *Any n-ary Boolean function f has Σ_2-circuits of size $\leq 2^n$, and dually, Π_2-circuits of the same size.*

Proof. Use the CNF and DNF forms. □

Theorem 14. *If f is an n-ary Boolean function, then f has Σ_3-circuits of size $O(2^n/n)$ (dually, Π_3-circuits of the same size).*

Proof. This follows from the Lupanov construction. □