CSC2429 / MAT1304: Circuit Complexity March 7, 2019

Lecture 8: Monotone lower bounds for CLIQUE

Instructor: Benjamin Rossman *Scribe: Bruno Pasqualotto Cavalar*

1 Introduction

Let $k \geqslant 3$ be a fixed constant. Let k-CLIQUE : $\{0,1\}^{\binom{n}{2}} \to \{0,1\}$ be the Boolean function which, given a graph G represented by an element of $\{0,1\}^{\binom{n}{2}}$, outputs 1 if and only if G contains a clique of k vertices. It is easy to build a (monotone) circuit of size $O(n^k)$ for k-CLIQUE by considering a brute-force algorithm. It is widely conjectured that the non-monotone circuit size of k-CLIQUE is $n^{\Omega(k)}$.

The best algorithm known for k-CLIQUE in the worst-case runs on time $O(n^{\lceil k/3 \rceil \omega})$, where ω is the matrix multiplication exponent, currently known to satisfy $\omega < 2.373$. The algorithm generalizes the observation that a graph G has a triangle if and only if $\mathrm{Trace}(A_G^3) > 0$, where A_G is the adjacency matrix of G.

Proposition 1. *There exists an $O(n^{\lceil k/3 \rceil \omega})$ algorithm computing k-CLIQUE on the worst-case.*

1.1 The average-case

Let $G(n,p)$ the *Erdös-Rnyi* random graph on n vertices, where each edge appears independently with probability p. It is widely known that there exists a probability $p_c^{K_k} = p_c^{K_k}(n) = \Theta(n^{-2/(k-1)})$ such that $\Pr[G(n, p_c^{K_k})$ contains a k-clique$] = 1/2$. It is also widely conjectured that the non-monotone circuit size of k-CLIQUE is $n^{\Omega(k)}$ in the average-case under the $G(n, p_c^{K_k})$ distribution.

1.2 Why k constant?

When k is constant, clearly $O(n^k)$ is a polynomial bound. However, seeing k-CLIQUE as a problem parametrized by k, one readily obtains the following.

Proposition 2. *If there exists $c(k), \varepsilon(k)$ such that $\lim_{k\to\infty} \varepsilon(k) \to \infty$ and $\mathcal{C}(k\text{-CLIQUE}_n) \geqslant c(k)\Omega(n^{\varepsilon(k)})$, then* $\mathbf{P} \neq \mathbf{NP}$.

1.3 The slice case

Let $G(n,m)$ be the random graph chosen uniformly at random from the n-vertex graphs with exactly m edges. Together with Proposition 2, one can easily see from Berkowitz' argument from previous lectures that, if the *monotone* circuit size of k-CLIQUE under $G(n,m)$ is $n^{\Omega(k)}$, then $\mathbf{P} \neq \mathbf{NP}$.

2 A greedy algorithm for the average-case

We now consider a randomized greedy algorithm for finding cliques in the graph $G \sim G(n, p_c^{K_k})$.

First choose a vertex v_1 of G uniformly at random, then choose among the neighbors of v_1 a vertex v_2 uniformly at random, then choose among the common neighbors of v_1, v_2 a vertex v_3

uniformly at random, and keep doing this until you find a maximal clique of size, say, ℓ. This takes time $\ell \cdot n$, which is $O(n)$ with high probability.

We claim that this algorithm produces a clique of size at least $(k-1)/2$ almost surely.

Claim 3. *If $i < (k-1)/2$, then w.h.p there exists a common neighbor between $v_1, v_2, \dots, v_i$.*

Proof. We have

$$\Pr[\nexists \text{ common neighbor}] = (1-p^i)^{n-i} \leqslant e^{n^{-2i/(k-1)}(n-i)} = e^{-n^{\Omega(1)}} = o(1). \qquad \square$$

By the same argument, one can also see that it suffices to run the greedy algorithm $n^{\varepsilon^2 k}$ times to find a clique of size $(1/2+\varepsilon)k$ in G. By taking $\varepsilon = 1/2$, we obtain an algorithm running in time $n^{k/4+O(1)}$ that computes k-CLIQUE on $G(n, p_{\mathrm{c}}^{K_k})$.

2.1 Karp's conjecture

It can be seen by a calculation that the random graph $G(n, 1/2)$ contains a clique of size $\sim 2\log n$ with high probability. Therefore, the greedy algorithm finds a clique of size $\sim \log n$ with high probability. In light of this, Karp made the following conjecture.

Conjecture 4 (Karp's conjecture)**.** *For every $\varepsilon > 0$, there does not exists a polynomial-time algorithm which finds a clique of size $(1+\varepsilon)\log n$ in $G(n, 1/2)$ w.h.p.*

One can "scale down" this conjecture to the $G(n, p_{\mathrm{c}}^{K_k})$ case in the following way.

Conjecture 5. *For every $0 < \varepsilon < 1/2$, there does not exist an algorithm running in time $O(n^{\varepsilon^2 k})$ which finds a clique of size $(1/2+\varepsilon)k$ in $G(n, p_{\mathrm{c}}^{K_k})$ w.h.p.*

3 Razborov's lower bound

Theorem 6 (Razborov '85)**.** *For every constant k, we have*

$$\mathcal{C}(k\text{-CLIQUE}) = \Omega\left(\frac{n}{(\log n)^2}\right)^k.$$

The proof uses the "sunflower-plucking approximation method", which will be unfolded in the next lectures. Observe that this lower bound is bigger than the upper bound given in Proposition 1, which evidences a gap between monotone and non-monotone complexity.

Let $\mathcal{Y}$ be the distribution on n-vertex graphs which chooses a set $K \in \binom{[n]}{k}$ uniformly at random and outputs a graph with a clique in K and no other edges. Let $\mathcal{N}$ be the distribution on n-vertex graphs which chooses a function $c : [n] \to [k-1]$ uniformly at random and outputs a graph with edges between vertices $i, j \in [n]$ if and only if $c(i) \neq c(j)$. Clearly, graphs in the support of $\mathcal{N}$ are k-clique-free. Razborov's lower bound can be interpreted as considering the distribution which is half of the time $\mathcal{Y}$ and half of the time $\mathcal{N}$, and proving a $1/2 + o(1)$ correlation bound with k-CLIQUE for small circuits.

3.1 The approximation method

Razborov's lower bound introduces a technique known as the "approximation method", which we outline here. Let $\mathcal{B}^+$ denote the set of all monotone function $\mathcal{G} \to \{0,1\}$. For $f, g \in \mathcal{B}^+$, we write $f \leqslant g$ iff $f(X) \leqslant g(X)$ for all $X \in \mathcal{G}$. Note that $\mathcal{B}^+$ is a lattice under the partial order $\leqslant$: the least upper bound and greatest lower bound of f and g are given by $f \vee g$ and $f \wedge g$.

The idea of the approximation method is to replace $\mathcal{B}^+$ with a subset $\mathcal{A} \subseteq \mathcal{B}^+$ which is also a lattice under $\leqslant$. We call functions in $\mathcal{A}$ *approximators*, and we denote by $f \sqcup g$ and $f \sqcap g$ the l.u.b. and g.l.b. of two approximators f and g. (Note: For general $f, g \in \mathcal{A}$, functions $f \sqcup g, f \sqcap g \in \mathcal{A}$ need not coincide with $f \vee g$ and $f \wedge g$.) We also require $\mathcal{A}$ to contain all coordinate functions $x \mapsto x_i$ (that is, edge-indicator functions in the case of k-CLIQUE).

For any monotone DeMorgan circuit C (with binary $\vee$ and $\wedge$ gates and inputs labeled by edge-indicators), we may now consider the "approximator circuit" $C^{\mathcal{A}}$ obtained from C by replacing each $\vee$ (resp. $\wedge$) gate with a $\sqcup$ (resp. $\sqcap$) gate. (Whereas C could be computing an arbitrary monotone function, the function computed by $C^{\mathcal{A}}$ is always an approximator.)

The key to the approximation method is defining an appropriate approximator lattice $\mathcal{A}$. We require the following properties of $\mathcal{A}$:

1. For every approximator $f \in \mathcal{A}$,

$$\Pr[f(\mathcal{Y}) = 1] + \Pr[f(\mathcal{N}) = 0] \leqslant 1 + o(1).$$

   (In contrast, $\Pr[k\text{-CLIQUE}(\mathcal{Y}) = 1] + \Pr[f(\mathcal{N}) = 0] = 2$. It follows that k-CLIQUE has correlation at most $\frac{1}{2} + o(1)$ with any approximator under the distribution $\mathcal{X}$ which is $\mathcal{Y}$ half of the time and $\mathcal{N}$ half of the time.)

2. For all approximators $f, g \in \mathcal{A}$,

$$\begin{aligned}
\Pr[(f \vee g)(\mathcal{Y}) = 1 \text{ and } (f \sqcup g)(\mathcal{Y}) = 0] &\leqslant \delta, \\
\Pr[(f \wedge g)(\mathcal{Y}) = 1 \text{ and } (f \sqcap g)(\mathcal{Y}) = 0] &\leqslant \delta, \\
\Pr[(f \vee g)(\mathcal{N}) = 0 \text{ and } (f \sqcup g)(\mathcal{N}) = 1] &\leqslant \delta, \\
\Pr[(f \wedge g)(\mathcal{N}) = 0 \text{ and } (f \sqcap g)(\mathcal{N}) = 1] &\leqslant \delta,
\end{aligned}$$

   for some small constant δ (which equals $O(\frac{(\log n)^2}{n})^k$ in Razborov's proof, in the regime where k is a fixed constant).

Property 2 bounds the approximation error per gate when we replace $\mathcal{C}$ with $\mathcal{C}^{\mathcal{A}}$: that is,

$$\begin{aligned}
\Pr[C(\mathcal{Y}) = 1 \text{ and } C^{\mathcal{A}}(\mathcal{Y}) = 0] &\leqslant \delta \cdot \operatorname{size}(\mathcal{C}), \\
\Pr[C(\mathcal{N}) = 0 \text{ and } \mathcal{C}^{\mathcal{A}}(\mathcal{N}) = 1] &\leqslant \delta \cdot \operatorname{size}(\mathcal{C}).
\end{aligned}$$

It follows that

$$\begin{aligned}
\Pr[C(\mathcal{Y}) = 1] + \Pr[C(\mathcal{N}) = 0] &\leqslant \Pr[C(\mathcal{Y}) = 1 \text{ and } C^{\mathcal{A}}(\mathcal{Y}) = 0] + \Pr[C^{\mathcal{A}}(\mathcal{Y}) = 1] \\
&\quad + \Pr[C(\mathcal{N}) = 0 \text{ and } \mathcal{C}^{\mathcal{A}}(\mathcal{N}) = 1] + \Pr[C^{\mathcal{A}}(\mathcal{N}) = 0] \\
&\leqslant 1 + o(1) + \delta \cdot \operatorname{size}(\mathcal{C}).
\end{aligned}$$

This implies a lower bound $\Omega(1/\delta)$ on the monotone circuit size of k-CLIQUE (or whatever the target function f, provided the distribution $\mathcal{Y}$ and $\mathcal{N}$ satisfy $\Pr[f(\mathcal{Y}) = 1] + \Pr[f(\mathcal{N}) = 0] \geqslant 1 + \Omega(1)$).

In the next lecture, we will use the approximation method to prove a different lower bound for k-CLIQUE.