The Monotone Complexity of k-Clique on Random Graphs

Benjamin Rossman
Massachusetts Institute of Technology

November 5, 2009

Abstract

Understanding the average-case complexity of natural problems on natural distributions is an important challenge for computational complexity. In this paper, we consider the average-case monotone complexity of the k-clique problem (for constant k) on Erdős-Rényi random graphs $G(n,p)$. Our main result is a lower bound of $\omega(n^{k/4})$ on the size of monotone circuits which solve the k-clique problem asymptotically almost surely on $G(n,p)$ for all functions $p : \mathbb{N} \longrightarrow [0,1]$ (or even just for two sufficiently far-apart threshold functions, such as $n^{-2/(k-1)}$ and $2n^{-2/(k-1)}$). While stronger lower bounds of $\tilde{\Omega}(n^k)$ are known [17] for the worst-case monotone complexity of k-clique, our result is the first average-case monotone lower bound. This lower bound also supports the intuition that random graphs at the threshold are a source of hard instances for the k-clique problem.

A further result of this paper is a nearly matching upper bound of $n^{k/4+O(1)}$, obtained by monotonizing a construction of constant-depth circuits due to Amano [3]. This upper bound points out a gap between the worst-case and average-case monotone complexity of the k-clique problem.

Similar bounds on the average-case complexity of k-clique for non-monotone constant-depth (AC^0) circuits were previously obtained by the author [18] ($\omega(n^{k/4})$ lower bound) and Amano [3] ($n^{k/4+O(1)}$ upper bound). We remark that the monotone lower bound of the present paper uses entirely different techniques from the AC^0 lower bound of [18]. In particular, we introduce a new variant of *sunflowers* and prove an analogue of the sunflower lemma that may be of independent interest.

1 Introduction

This paper considers the average-case complexity of the k-clique problem (for constant k) on monotone circuits (composed of AND and OR gates only, that is, without NOT gates). The k-clique problem asks whether a simple graph contains a k-clique, that is, a set of k vertices with an edge between every pair. This is one of the most fundamental problems studied in complexity theory. A worst-case lower bound of $\Omega(n^{f(k)})$ for the k-clique problem on general boolean circuits where f is any increasing function would imply $\mathrm{P} \neq \mathrm{NP}$ (in fact, $\mathrm{NP} \not\subseteq \mathrm{P/Poly}$). This fact motivates studying the k-clique problem on restricted classes of circuits, such as AC^0 and monotone circuits, where techniques for unconditional lower bounds have been developed.

In a seminal work, Razborov [17] proved a lower bound of $\tilde{\Omega}(n^k)$ on the size of monotone circuits which solve k-clique in the worst-case. (In the setting of super-constant $k = k(n)$, this lower bound was improved by Alon and Boppana [1] and Amano [5].) More generally, Razborov's result applies to monotone circuits which have value 1 almost surely on a k-clique uniformly planted among a set of n vertices (with no other edges) and value 0 almost surely on a uniform random $(k-1)$-partite graph.

Of course, distinguishing isolated k-cliques from $(k-1)$-partite graphs is an easy task for non-monotone circuits (even for relatively simple constant-depth circuits). The unviability of "isolated k-cliques vs. $(k-1)$-partite graphs" as a premise for non-monotone circuit lower bounds is one good reason to investigate monotone lower bounds for natural distributions of graphs, where it is reasonable to conjecture that even non-monotone circuits cannot solve k-clique efficiently. One very natural class of distributions in this context are Erdős-Rényi random graphs $G(n,p)$, defined as n-vertex graphs in which each potential edge is independently included with probability p (where $p = p(n)$ is a function of n). For every monotone graph property (like the existence of a k-clique), there is a *threshold function* $p(n)$ such that the property holds on $G(n,p)$ with probability bounded away from 0 and 1, while for functions $q(n)$ which are $o(p(n))$ or $\omega(p(n))$, the property respectively holds or does not hold almost surely on $G(n,q)$. In particular, $\Theta(n^{-2/(k-1)})$ is precisely the class of threshold functions for the existence a k-clique.

The present paper is based on the intuition that random graphs $G(n,p)$ for threshold functions $p(n) \in \Theta(n^{-2/(k-1)})$ are a source of hard instances for the k-clique problem. (Similar beliefs about random sat at the threshold are common in the research in statistical physics.) For monotone circuits, we show that this intuition is valid. Specifically, we prove that any monotone circuit which solves the k-clique problem almost surely on $G(n,p)$ for two sufficiently far-apart threshold functions $p(n)$ (such as, for instance, $n^{-2/(k-1)}$ and $2n^{-2/(k-1)}$) has size $\omega(n^{k/4})$. Moreover, we prove a nearly matching upper bound by exhibiting monotone circuits of size $n^{k/4+O(1)}$ which solve the k-clique problem almost surely on $G(n,p)$ for every threshold function $p(n)$. (These results are formally stated in §3.) We remark that while our lower bound does not improve the worst-case lower bound $\tilde{\Omega}(n^k)$, it is the first average-case monotone lower bound for the k-clique problem, while our upper bound points out a gap between the worst-case and average-case complexity.

The intuition that k-clique is hard on random graphs at the threshold has a precedent in previous work of the author [18], which proves a lower bound of $\omega(n^{k/4})$ on the size of non-monotone constant-depth (AC^0) circuits that solve the k-clique problem almost surely on $G(n,p)$ for a single threshold function $p(n)$. (This result in fact improved the best-known worst-case lower bound of $\omega(n^{k/89d^2})$ for depth-d circuits due to Beame [6], in particular by shifting dependence on d from the exponent of n to a lower-order term.) The exponent $k/4$ in this result was shown to be optimal (up to a constant) by Amano [3], whose construction of constant-depth circuits (which solve k-clique almost surely at a threshold) is "monotonized" for the upper bound in the present paper.

Although the AC^0 lower bound of [18] and the monotone lower bound of the present paper proceed from the same intuition, the two proofs are significantly different as far as their technical details. For one thing, the AC^0 result relies on Håstad's Switching Lemma [12], which is no help in the setting of unbounded-depth monotone circuits. While the present monotone lower bound fits the general "approximation method" framework of Razborov [17], it makes essential use of a combinatorial lemma concerning a novel variant of sunflowers called (p,q)-*sunflowers* (defined in §4). This notion, which generalizes an essential property of sunflowers, may be of independent interest.

2 Definitions

Let k be a fixed constant $\geqslant 5$. Let $\mathbb{N} = \{0, 1, 2, \dots\}$. For $n \in \mathbb{N}$, let $[n] = \{1, \dots, n\}$. For a set X and $j \in \mathbb{N}$, let $\binom{X}{j} = \{j\text{-element subsets of } X\}$ and $\binom{X}{<j} = \bigcup_{i=0}^{j-1} \binom{X}{i}$.

"a.a.s." abbreviates "asymptotically almost surely" and means "with probability that tends to 1 as $n \to \infty$".

Graphs: *Graphs* are simple graphs, that is, $G = (V_G, E_G)$ where V_G is a set and $E_G \subseteq \binom{V_G}{2}$. By default, $V_G = [n]$ (where n is an arbitrary integer).

We say that H is a *subgraph* of G and write $H \subseteq G$ if $V_H \subseteq V_G$ and $E_H \subseteq E_G$; we say that H is a *proper subgraph* and write $H \subset G$ if $H \subseteq G$ and $H \neq G$.

For a graph G, let $\mathrm{supp}(G)$ denote the number of non-isolated vertices in G (i.e., the size of the support of G), and let $\mathcal{S}(G)$ denote the set of subgraphs of G.

Cliques: For $A \subseteq [n]$, (by mild abuse of notation) let K_A $(= K_A([n]))$ denote the graph with vertex set $[n]$ and edge set $\binom{A}{2}$; (by mild abuse of terminology) we refer to graphs K_A as $|A|$-*cliques*. The *random k-clique* is the random graph K_A where A is uniformly distributed in $\binom{[n]}{k}$.

For a graph G with vertex set $[n]$, let $\omega_k(G)$ denote the number of k-cliques in G, that is, $\omega_k(G) = |\{A \in \binom{[n]}{k} : K_A \subseteq G\}|$.

Monotone Graph Functions: A *monotone graph function* is a function f from $\{n\text{-vertex graphs}\}$ to $\{0,1\}$ (or a sequence $f = (f_n)$ of functions) satisfying $H \subseteq G \Longrightarrow f(H) \leqslant f(G)$.

A graph H is a *minterm* of f if $f(H) = 1$ and $f(H_0) = 0$ for all $H_0 \subset H$. The set of minterms of f is denoted $\mathcal{M}(f)$.

For a graph H, the H-*indicator* is the monotone graph function Ind_H defined by $\mathrm{Ind}_H(G) = 1$ iff $H \subseteq G$ (so $\mathcal{M}(\mathrm{Ind}_H) = \{H\}$).

Monotone Circuits on Graphs: A *monotone circuit on graphs* is a directed acyclic graph $\mathscr{C}$ with a unique sink ("output") in which sources ("inputs") are labeled by elements of $\binom{[n]}{2} \cup \{0,1\}$ (i.e., potential edges or constants 0 or 1) and all other nodes ("gates") are labeled either $\wedge$ or $\vee$. For nodes $\nu \in \mathscr{C}$, the value of (the monotone function computed at) ν on a graph G is denoted $\nu(G)$, while the value of $\mathscr{C}$ on G is denoted $\mathscr{C}(G)$. The *fan-in* of G is the maximum in-degree of a gate. (In §5–7 we consider monotone circuits with fan-in 2, while for our upper bound in §8, we consider constant-depth monotone circuits with unbounded fan-in.)

Erdős-Rényi Random Graphs: For a function $p : \mathbb{N} \longrightarrow [0,1]$, $G(n,p)$ denotes the *Erdős-Rényi random graph* with vertex set $[n]$ where potential edges in $\binom{[n]}{2}$ are independently included with probability $p(n)$. Classes $o(n^{-2/(k-1)})$, $\omega(n^{-2/(k-1)})$ and $\Theta(n^{-2/(k-1)})$ are respectively the classes of *subcritical*, *supercritical* and *threshold* functions for the existence of a k-clique in $G(n,p)$. In particular, the random variable $\mathrm{E}[\omega_k(G(n,p))]$ is asymptotically 0 in the subcritical regime, ∞ in the supercritical regime, and Poisson with constant mean in the threshold regime [7] (see also Ch. 4 of [8]).

A boolean function f on graphs is said to *solve k-clique on (Erdős-Rényi) random graphs* if $f(G) = 1 \iff \omega_k(G) \geqslant 1$ a.a.s. for $G = G(n,p)$ for all functions $p : \mathbb{N} \longrightarrow [0,1]$. (Note that all the "action" in this definition is for threshold functions $p(n)$, since constant functions 0 and 1 solve k-clique a.a.s. in the subcritical and supercritical regimes.)

3 Results

Fix sufficiently small $\delta > 0$ (to be determined later, though $\delta = k^{-3}$ suffices) and let

$$p^-(n) = n^{-\frac{2}{k-1}(1+\delta)}, \qquad p^\theta(n) = n^{-\frac{2}{k-1}}, \qquad p^+(n) = n^{-\frac{2}{k-1}(1-\delta)}.$$

Throughout this paper, $G^- = G(n,p^-)$, $G^\theta = G(n,p^\theta)$ and $G^+ = G(n,p^+)$ are (independent) subcritical, threshold and supercritical random graphs.

We first prove a lower bound on the size of monotone circuits which distinguish random k-cliques from subcritical random graphs G^-.

Theorem 1. *Let $\mathscr{C}$ be a monotone circuit of size $O(n^{k/4})$ such that $\mathrm{E}[\mathscr{C}(\text{random } k\text{-clique})] \geqslant \Omega(1)$. Then $\mathrm{E}[\mathscr{C}(G^-)] \geqslant 1 - \exp(-n^{\Omega(1)})$.*

We remark that every monotone graph function f satisfying $\mathrm{E}[f(\text{random } k\text{-clique})] \geqslant \Omega(1)$ also satisfies $\mathrm{E}[f(G^+)] \geqslant 1 - \exp(-n^{\Omega(1)})$. (This can be shown using Janson's inequality (Lemma 22).) Theorem 1 thus implies that monotone circuits of size $O(n^{k/4})$ can either recognize a positive fraction of k-cliques, or distinguish subcritical G^- from supercritical G^+, but not both simultaneously.

One key ingredients in the proof is a combinatorial lemma concerning a new variant of sunflowers (described in §4), leading to an approximation of the nodes of a monotone circuit by a special class of functions (a common theme in monotone circuit lower bounds ever since Razborov[17]). Another key concept in the proof (inspired by a similar idea in the AC^0 lower bound of [18]) is that a circuit on graphs which computes a minterm in the form of some k-clique K_A (where $A \in \binom{[n]}{k}$) must, for some node, produce minterms which are in some "intermediate" class of subgraphs of K_A. In §5 we identify a special class $\mathcal{J}$ of such "intermediate" subgraphs of k-cliques, which plays an important role in the $\omega(n^{k/4})$ lower bounds of the present paper as well as [18].

Using Theorem 1, we derive the lower bound mentioned in the introduction.

Theorem 2. *Monotone circuits solving k-clique on random graphs have size $\omega(n^{k/4})$.*

In fact, the proof gives a lower bound for every $\varepsilon > 0$ of $\Omega\big(n^{\frac{k+1}{4}+\frac{9}{4(k-1)}-\varepsilon}\big)$ for even k and $\Omega\big(n^{\frac{k+1}{4}+\frac{2}{k-1}-\varepsilon}\big)$ for odd k. This already gives a non-trivial lower bound of $\Omega(n^{11/5-\varepsilon})$ for $k = 6$. Moreover, to state Theorem 2 in the strongest form that the proof entails, we can replace "solves k-clique on random graphs" with "solves k-clique a.a.s. on $G(n,p)$ and $G(n,2p)$ where $p \in \Theta(n^{-2/(k-1)})$"; further, we can even replace $2p$ with $p + p^{1+\varepsilon}$ for sufficiently small $\varepsilon > 0$.

We also show that $k/4$ is essentially optimal in Theorem 2 by constructing a monotone version of AC^0 circuits due to Amano [3].

Theorem 3. *There exist constant-depth monotone circuits of size $n^{k/4+O(1)}$ which solve k-clique on random graphs.*

After presenting some technical preliminaries in §4 and §5, Theorems 1, 2, 3 are proved in §6, §7, §8 respectively.

4 $(1/N, \exp(-N^{\varepsilon}))$-Sunflower Lemma

A *hypergraph* on a set X is a family $\mathcal{F} \subseteq \wp(X)$ of subsets of X. $\mathcal{F}$ is *s-uniform* if $|f| = s$ for all $f \in \mathcal{F}$. We say that $\mathcal{F}$ is a *sunflower with core* Y if $f \cap g = Y$ for all distinct $f, g \in \mathcal{F}$; sets $f \setminus Y$ where $f \in \mathcal{F}$ are called *petals* of $\mathcal{F}$. Let N be any positive integer (unrelated to $|X|$).

Lemma 4 (Erdős-Rado Sunflower Lemma [9])**.** *Every s-uniform hypergraph $\mathcal{F}$ of cardinality $>$ $s!(N-1)^s$ contains a sunflower of size N.* □

A *p-biased random subset* W of X is a random $W \subseteq X$ such that events $x \in W$ are mutually independent with probability p. The following definition of *(p, q)-sunflowers* relaxes the notion of sunflower by allowing petals to overlap (but not too much on average). Other variants of sunflowers have been studied in extremal combinatorics (see Ch. 7 of [15]), but this notion appears to be new.

Definition 5 ((p, q)-Sunflower)**.** For $p, q \in [0, 1]$, we say that $\mathcal{F}$ is a *(p, q)-sunflower over* Y if $Y \subseteq \bigcap \mathcal{F}$ and $\Pr_{p\text{-biased } W \subseteq X}(W \cup Y \text{ does not contain a hyperedge of } \mathcal{F}) \leqslant q$.

Note that every s-uniform sunflower of cardinality m is a $(p, (1-p^s)^m)$-sunflower for all $p \in [0, 1]$. Indeed, for many applications of sunflowers in circuit complexity [1, 5, 17], it would be good enough to work with the notion of $(p, (1-p^s)^m)$-sunflowers instead of sunflowers (that is, Definition 5 captures the essential property of sunflowers for these applications).

It follows immediately from Lemma 4 that every s-uniform hypergraph of cardinality $s!N^{s(s+\varepsilon)}$ contains an $(1/N, e^{-N^{\varepsilon}})$-sunflower of size N. Our combinatorial main lemma improves this to $\Omega(s! \ln(3/2)^{-s} N^{s(1+\varepsilon)})$. The proof uses an inductive argument tailored to an application of Janson's inequality (stated as Lemma 22 in Appendix A).

Lemma 6 (Main Lemma)**.** *Every s-uniform hypergraph $\mathcal{F}$ of cardinality* $\Omega(s! \ln(3/2)^{-s} N^{s(1+\varepsilon)})$ *contains an* $(1/N, \exp(-N^{\varepsilon}))$*-sunflower of size N.*

Proof. Define sequence $\ell_0, \ell_1, \ell_2, \dots$ as follows: $\ell_0 = 1$ and $\ell_j = 2 \sum_{i=0}^{j-1} \binom{j}{i} \ell_i$ for all $j \geqslant 1$. By an elementary calculation (omitted), we have $\ell_s \geqslant \Omega(s! \ln(3/2)^{-s})$. Arguing by induction on s, we will show $\mathcal{F}$ that contains an $(1/N, \exp(-N^{\varepsilon}))$-sunflower of size N whenever $|\mathcal{F}| > \ell_s N^{s(1+\varepsilon)}$. Note that the lemma is vacuously true in the base case when $s = 0$ (since in this case $|\mathcal{F}| \leqslant 1 = \ell_0 N^{0(1+\varepsilon)}$).

For the induction step, let $s \geqslant 1$ and assume $|\mathcal{F}| > \ell_s N^{s(1+\varepsilon)}$. For all $A \subseteq X$, let $\mathcal{F}_A = \{f \setminus A : A \subseteq f \in \mathcal{F}\}$ and note that $\mathcal{F}_A$ is an $(s - |A|)$-uniform hypergraph on X. We consider two cases.

First Case: Suppose there exist $i \in \{1, \dots, s\}$ and $A \in \binom{X}{i}$ such that $|\mathcal{F}_A| > \ell_{s-i} N^{(s-i)(1+\varepsilon)}$. By the induction hypothesis, $\mathcal{F}_A$ contains a $(1/N, e^{-N^{\varepsilon}})$-sunflower $\mathcal{F}'$ over some Y'. Letting $Y = A \cup Y'$ and $\mathcal{F}_0 = \{f \in \mathcal{F}' : A \subseteq f\}$, we have $\mathcal{F}_0 \subseteq \mathcal{F}$ is a $(1/N, e^{-N^{\varepsilon}})$-sunflower over Y.

Second Case: Suppose $|\mathcal{F}_A| \leqslant \ell_{s-i} N^{(s-i)(1+\varepsilon)}$ for all $i \in \{1, \dots, s\}$ and $A \in \binom{X}{i}$. We will show that $\mathcal{F}$ is itself a $(1/N, e^{-N^{\varepsilon}})$-sunflower over core $\emptyset$. Let W be a $1/N$-biased subset of X and for every hyperedge $f \in \mathcal{F}$, let Λ_f be the event that $f \subseteq W$. Note that $\bigwedge_{f \in \mathcal{F}} \overline{\Lambda_f}$ is precisely the event that W contains no hyperedge of $\mathcal{F}$. Define μ and Δ by

$$\mu = \sum_{f \in \mathcal{F}} \Pr(\Lambda_f), \qquad \Delta = \sum_{f, g \in \mathcal{F} \,:\, f \cap g \neq \emptyset} \Pr(\Lambda_f \wedge \Lambda_g).$$

By Janson's inequality (Lemma 22), we have $\Pr(\bigwedge_{f \in \mathcal{F}} \overline{\Lambda_f}) \leqslant \exp(-\min(\mu/2, \mu^2/2\Delta))$. Therefore, it suffices to show $\mu/2 \geqslant N^{\varepsilon}$ and $\mu^2/2\Delta \geqslant N^{\varepsilon}$. The first of these inequalities follows from

$$\mu = |\mathcal{F}|/N^s > \ell_s N^{s\varepsilon} \geqslant 2N^{\varepsilon}.$$

We now bound Δ. For all $i \in \{1, \dots, s\}$, note that $\sum_{A \in \binom{X}{i}} |\mathcal{F}_A| = \binom{s}{i}|\mathcal{F}|$ since each hyperedge in $\mathcal{F}$ is counted $\binom{s}{i}$ times in this summation. Thus,

$$\sum_{A \in \binom{X}{i}} |\mathcal{F}_A|^2 \leqslant \binom{s}{i} |\mathcal{F}| \ell_{s-i} N^{(s-i)(1+\varepsilon)} = \mu \binom{s}{i} \ell_{s-i} N^{2s-i+(s-i)\varepsilon}.$$

For all $f, g \in \mathcal{F}$ such that $|f \cap g| = i$, note that $\Pr(\Lambda_f \wedge \Lambda_g) = 1/N^{2s-i}$. So we have

$$\Delta = \sum_{\substack{A \subseteq X:\\ 1 \leqslant |A| \leqslant s}} \sum_{\substack{f,g \in \mathcal{F}:\\ f \cap g = A}} \Pr(\Lambda_f \wedge \Lambda_g) \leqslant \sum_{i=1}^{s} \frac{1}{N^{2s-i}} \sum_{A \in \binom{X}{i}} |\mathcal{F}_A|^2 \leqslant \mu \sum_{i=1}^{s} \binom{s}{i} \ell_{s-i} N^{(s-i)\varepsilon} \leqslant \frac{\mu \ell_s N^{(s-1)\varepsilon}}{2}.$$

It follows that

$$\frac{\mu^2}{2\Delta} \geqslant \frac{\mu}{\ell_s N^{(s-1)\varepsilon}} = \frac{|\mathcal{F}|/N^s}{\ell_s N^{(s-1)\varepsilon}} \geqslant \frac{\ell_s N^{s(1+\varepsilon)}/N^s}{\ell_s N^{(s-1)\varepsilon}} = N^{\varepsilon}.$$

This completes the proof, as

$$\begin{aligned}\Pr_{1/N\text{-biased } W \subseteq X}\big(W \text{ contains no hyperedge of } \mathcal{F}\big) &= \Pr\big(\textstyle\bigwedge_{f \in \mathcal{F}} \overline{\Lambda_f}\big) \\ &\leqslant \exp\big(-\min(\mu/2, \mu^2/2\Delta)\big) \leqslant e^{-N^{\varepsilon}}. \qquad \square\end{aligned}$$

5 $\star$-Closure

Define classes $\mathcal{I}$ and $\mathcal{J}$ of n-vertex graphs by

$$\mathcal{I} = \{H : V_H = [n],\ \mathrm{supp}(H) < k/2\}, \qquad \mathcal{J} = \{H_1 \cup H_2 : H_1, H_2 \in \mathcal{I}\} \setminus \mathcal{I}.$$

That is, $\mathcal{I}$ is the class of n-vertex graphs with fewer than $k/2$ non-isolated vertices, and $\mathcal{J}$ is the class of unions of pairs of graphs in $\mathcal{I}$ having at least $k/2$ non-isolated vertices.

Observation 7. For all monotone graph function f and g, note that

$$\mathcal{M}(f \vee g) \subseteq \mathcal{M}(f) \cup \mathcal{M}(g), \qquad \mathcal{M}(f \wedge g) \subseteq \{H_1 \cup H_2 : H_1 \in \mathcal{M}(f),\ H_2 \in \mathcal{M}(g)\}.$$

(That is, every minterm of $f \vee g$ is a minterm of f or a minterm of g, while every minterm of $f \vee g$ is the union of a minterm of f and a minterm of g.) Thus, if $\mathcal{M}(f), \mathcal{M}(g) \subseteq \mathcal{I}$, then $\mathcal{M}(f \vee g) \subseteq \mathcal{I}$ and $\mathcal{M}(f \wedge g) \subseteq \mathcal{I} \cup \mathcal{J}$.

It follows that if $\mathscr{C}$ is a monotone circuit on graphs and $H \in \mathcal{M}(\mathscr{C}) \setminus \mathcal{I}$, then there is a node $\nu \in \mathscr{C}$ such that $\mathcal{M}(\nu) \cap \mathcal{S}(H) \cap \mathcal{J} \neq \emptyset$ (i.e., some subgraph of H is both a minterm of ν and belongs to the class $\mathcal{J}$). We will extend this observation later on in Lemma 14.

5.1 $\star$-Closed Functions

We define a closure operator $f \longmapsto f^{\star}$ in the lattice of monotone graph functions, that is, an operator satisfying $f \leqslant f^{\star}$ and $(f^{\star})^{\star} = f^{\star}$ and $(f \wedge g)^{\star} = f^{\star} \wedge g^{\star}$. (Recall that $G^{-} = G(n, p^{-})$ where $p^{-}(n) = n^{-\frac{2}{k-1}(1+\delta)}$ and $\delta > 0$ is fixed.)

Definition 8 ($\star$-Closed Monotone Graph Functions)**.** A monotone graph function f is $\star$-*closed* if for every $H \in \mathcal{I} \cup \mathcal{J}$,

$$\mathrm{E}[f(G^{-} \cup H)] \geqslant 1 - e^{-n^{\delta}} \implies f(H) = 1.$$

Equivalently, f is $\star$-*closed* if $\mathrm{E}[f(G^- \cup H)] \notin [1 - e^{-n^\delta}, 1)$ for all $H \in \mathcal{I} \cup \mathcal{J}$.

Note that conjunctions of $\star$-closed functions are $\star$-closed. It follows that for every f, there exists a unique minimal $\star$-closed function which is $\geqslant f$. We denote this function by $f^\star$ and called it the $\star$-*closure* of f.

The key property of $\star$-closed functions follows from our combinatorial main lemma (Lemma 6).

Lemma 9 ($\star$-Closed Functions Have Few Minterms in $\mathcal{I} \cup \mathcal{J}$)**.** *If f is $\star$-closed, then for every $H \in \mathcal{I} \cup \mathcal{J}$ with $|E_H| = s$, we have $|\{H' \in \mathcal{M}(f) : H \cong H'\}| < o(n^{s(\frac{2}{k-1}+2\delta)})$.*

Proof. Suppose f is $\star$-closed and let $H \in \mathcal{I} \cup \mathcal{J}$ with $|E_H| = s$. Assume, for contradiction, that $|\{H' \in \mathcal{M}(f) : H \cong H'\}| \geqslant \Omega(n^{s(\frac{2}{k-1}+2\delta)})$. Let $X = \binom{[n]}{2}$ and define s-uniform hypergraph $\mathcal{F} \subseteq \binom{X}{s}$ by $\mathcal{F} = \{E_{H'} : H' \in \mathcal{M}(f),\ H \cong H'\}$. Let $N = n^{\frac{2}{k-1}(1+\delta)}$ $(= 1/p^-)$ and $\varepsilon = (\frac{2}{k-1}(1+\delta))^{-1}\delta$, so that $N^\varepsilon = n^\delta$. Note that $|\mathcal{F}| > \omega(N^{s(1+\varepsilon)})$ (since $\frac{2}{k-1} + 2\delta > \frac{2}{k-1}(1+\delta)(1+\varepsilon)$, by a simple calculation). By Lemma 6, $\mathcal{F}$ contains a $(1/N, e^{-N^\varepsilon})$-sunflower over some Y. Letting H_0 be the graph $([n], Y)$, this means

$$\mathrm{E}[f(G^- \cup H_0)] \geqslant \Pr(\exists H' \in \mathcal{M}(f) \text{ such that } H \cong H' \text{ and } G^- \cup H_0) \geqslant 1 - e^{-N^\varepsilon} = 1 - e^{-n^\delta}.$$

Note that $|Y| < s$ and hence H_0 is a proper subgraph of some $H' \in \mathcal{M}(f)$ isomorphic to H. Since f is $\star$-closed and $H_0 \in \mathcal{I} \cup \mathcal{J}$ (as $H_0 \subseteq H' \in \mathcal{I} \cup \mathcal{J}$), it follows that that $f(H_0) = 1$. But since H_0 is proper subgraph of H', this contradicts the fact that H' is a minterm of f. $\square$

Lemmas 10, 11 and 12 state some additional properties of $\star$-closed functions.

Lemma 10 ($\star$-Closure Algorithm)**.** *Let f be a monotone graph function and consider the following algorithm. Let $f_0 = f$. Starting with $i = 1$, output f_{i-1} if it is $\star$-closed; otherwise, choose any $H_i \in \mathcal{I} \cup \mathcal{J}$ such that $\mathrm{E}[f_{i-1}(G^- \cup H_i)] \in [1 - e^{-n^\delta}, 1)$, let $f_i = f_{i-1} \vee \mathrm{Ind}_{H_i}$, increment i and repeat. This algorithm terminates after $t = n^{O(1)}$ iterations and outputs $f_t = f^\star$.*

Proof. The algorithm clearly terminates after at most $|\mathcal{I} \cup \mathcal{J}| \leqslant n^{O(1)}$ iterations, since each graph in $\mathcal{I} \cup \mathcal{J}$ can occur as H_i only once. Correctness of the algorithm follows from the fact that $f_t = f \vee (\mathrm{Ind}_{H_1} \vee \cdots \vee \mathrm{Ind}_{H_t})$ and $f^\star(H_i) = 1$ for $i = 1, \ldots, t$ (by induction). $\square$

The next lemma says that $f^\star$ approximates f extremely well on $G^- \cup H$ for every graph H.

Lemma 11. *For every monotone graph function f and graph H (not necessarily in $\mathcal{I} \cup \mathcal{J}$), we have* $\Pr\big(f(G^- \cup H) \neq f^\star(G^- \cup H)\big) \leqslant n^{O(1)} e^{-n^\delta}$.

Proof. Let t and $H_1, \ldots, H_t$ and $f_0, \ldots, f_t$ be as in Lemma 10.

$$\begin{aligned}
\Pr\big(f(G^- \cup H) \neq f^\star(G^- \cup H)\big) &\leqslant t \Pr\big(f_{i-1}(G^- \cup H) \neq f_i(G^- \cup H)\big)\\
&= t \Pr\big(f_{i-1}(G^- \cup H) = 0 \text{ and } H_i \subseteq G^- \cup H\big)\\
&\leqslant t \Pr\big(f_{i-1}(G^- \cup H \cup H_i) = 0\big)\\
&\leqslant t \Pr\big(f_{i-1}(G^- \cup H_i) = 0\big) \quad \text{(by monotonicity)}\\
&\leqslant n^{O(1)} e^{-n^\delta}.
\end{aligned}$$

$\square$

Lemma 12. *For every monotone graph function f, we have $\mathcal{M}(f^\star) \subseteq \mathcal{M}(f) \cup \mathcal{I} \cup \mathcal{J}$.*

Proof. Recall that $\mathcal{M}(f_1 \vee f_2) \subseteq \mathcal{M}(f_1) \cup \mathcal{M}(f_2)$ for all monotone graph functions f_1 and f_2. Thus, $\mathcal{M}(f^\star) = \mathcal{M}(f \vee (\mathrm{Ind}_{H_1} \vee \cdots \vee \mathrm{Ind}_{H_t})) \subseteq \mathcal{M}(f) \cup \{H_1, \ldots, H_t\}$. Since all H_i belong to $\mathcal{I} \cup \mathcal{J}$, we have $\mathcal{M}(f^\star) \subseteq \mathcal{M}(f) \cup \mathcal{I} \cup \mathcal{J}$. $\square$

5.2 $\star$-Closed Approximation

We now extend the notion of $\star$-closure to monotone circuits with fan-in 2. Define an operation $\overline{\vee}$ on monotone graph functions by $f \mathbin{\overline{\vee}} g = (f \vee g)^\star$. If $\mathscr{C}$ is a monotone circuit on graphs with fan-in 2, let $\overline{\mathscr{C}}$ denote the $\{\wedge, \overline{\vee}\}$-circuit obtained from $\mathscr{C}$ simply by replacing $\vee$-gates with $\overline{\vee}$-gates. Note that nodes of $\overline{\mathscr{C}}$ compute a $\star$-closed functions. We call $\overline{\mathscr{C}}$ the *$\star$-closed approximation* of $\mathscr{C}$. For every node $\nu \in \mathscr{C}$, let $\overline{\nu}$ denote the function computed at the corresponding node in $\overline{\mathscr{C}}$. Note that $\overline{\nu}$ is possibly a different from $\nu^\star$, although $\nu(G) \leqslant \nu^\star(G) \leqslant \overline{\nu}(G)$ for all graphs G.

The next lemma says that $\overline{\mathscr{C}}$ approximates $\mathscr{C}$ extremely well on $G^- \cup H$ for all graphs H (so long as $|\mathscr{C}| \leqslant n^{O(1)}$).

Lemma 13. *For every graph H, we have* $\Pr\big(\overline{\mathscr{C}}(G^- \cup H) \neq \mathscr{C}(G^- \cup H)\big) \leqslant |\mathscr{C}| n^{O(1)} e^{-n^\delta}$.

Proof. Note that if $\nu(G^- \cup H) = \nu^\star(G^- \cup H)$ for all $\nu \in \mathscr{C}$, then $\overline{\mathscr{C}}(G^- \cup H) = \mathscr{C}(G^- \cup H)$. Thus,

$$\begin{aligned}\Pr\big(\overline{\mathscr{C}}(G^- \cup H) \neq \mathscr{C}(G^- \cup H)\big) &\leqslant \Pr\big(\exists \nu \in \mathscr{C},\ \nu(G^- \cup H) \neq \nu^\star(G^- \cup H)\big)\\ &\leqslant \textstyle\sum_{\nu \in \mathscr{C}} \Pr\big(\nu(G^- \cup H) \neq \nu^\star(G^- \cup H)\big)\\ &\leqslant |\mathscr{C}| n^{O(1)} e^{-n^\delta} \quad \text{(by Lemma 11).}\end{aligned}$$ $\square$

We now state a key property of minterms of nodes in $\overline{\mathscr{C}}$ along the lines of Observation 7.

Lemma 14. *For every $H \in \mathcal{M}(\overline{\mathscr{C}}) \setminus \mathcal{I}$, there is a node $\nu \in \mathscr{C}$ such that $\mathcal{M}(\overline{\nu}) \cap \mathcal{S}(H) \cap \mathcal{J} \neq \emptyset$.*

Proof. Assume that $\mathcal{M}(\overline{\nu}) \cap \mathcal{S}(H) \cap \mathcal{J} = \emptyset$ for all $\nu \in \mathscr{C}$. We will show, by induction, that $\mathcal{M}(\overline{\nu}) \cap \mathcal{S}(H) \subseteq \mathcal{I}$ for all $\nu \in \mathscr{C}$. This establishes the lemma, since it implies $\mathcal{M}(\overline{\mathscr{C}}) \cap \mathcal{S}(H) \subseteq \mathcal{I}$ (taking ν to be the output node of $\mathscr{C}$).

In the base case where ν is an input node (corresponding to a potential edge e), note that ν and $\overline{\nu}$ compute the same function (testing whether a graph contains the edge e). Clearly, $\mathcal{M}(\nu) \subseteq \mathcal{I}$ since the unique minterm (the graph whose only edge is e) has only 2 non-isolated vertices (and $2 < k/2$ as $k \geqslant 5$). Therefore $\mathcal{M}(\overline{\nu}) \cap \mathcal{S}(H) \subseteq \mathcal{I}$.

For the induction step, first consider the case that $\nu = \nu_1 \wedge \nu_2$ where (by the induction hypothesis) $\mathcal{M}(\overline{\nu_i}) \cap \mathcal{S}(H) \subseteq \mathcal{I}$ for $i = 1, 2$. Suppose that $F \in \mathcal{M}(\overline{\nu}) \cap \mathcal{S}(H)$. We will show that $F \in \mathcal{I}$ (and thus $\mathcal{M}(\overline{\nu}) \cap \mathcal{S}(H) \subseteq \mathcal{I}$). By Lemma 12,

$$\mathcal{M}(\overline{\nu}) = \mathcal{M}(\overline{\nu_1} \mathbin{\overline{\vee}} \overline{\nu_2}) = \mathcal{M}((\overline{\nu_1} \vee \overline{\nu_2})^\star) \subseteq \mathcal{M}(\overline{\nu_1} \vee \overline{\nu_2}) \cup \mathcal{I} \cup \mathcal{J} \subseteq \mathcal{M}(\overline{\nu_1}) \cup \mathcal{M}(\overline{\nu_2}) \cup \mathcal{I} \cup \mathcal{J}.$$

We now consider four cases depending whether F belongs to $\mathcal{M}(\overline{\nu_1})$, $\mathcal{M}(\overline{\nu_2})$, $\mathcal{I}$ or $\mathcal{J}$. If $F \in \mathcal{I}$, there is nothing to prove. Note that $F \notin \mathcal{J}$ since $F \in \mathcal{S}(H)$ and we assumed that $\mathcal{M}(\overline{\nu}) \cap \mathcal{S}(H) \cap \mathcal{J} = \emptyset$. If $F \in \mathcal{M}(\overline{\nu_i})$ for some $i \in \{1, 2\}$, then since $\mathcal{M}(\overline{\nu_i}) \cap \mathcal{S}(H) \subseteq \mathcal{I}$, we have $F \in \mathcal{I}$. Therefore $\mathcal{M}(\overline{\nu}) \cap \mathcal{S}(H) \subseteq \mathcal{I}$.

Finally, suppose $\nu = \nu_1 \wedge \nu_2$ where (by the induction hypothesis) $\mathcal{M}(\overline{\nu_i}) \cap \mathcal{S}(H) \subseteq \mathcal{I}$ for $i = 1, 2$. Again consider $F \in \mathcal{M}(\overline{\nu}) \cap \mathcal{S}(H)$. Since $\overline{\nu} = \overline{\nu_1} \wedge \overline{\nu_2}$, there exist $F_1 \in \mathcal{M}(\overline{\nu_1}) \cap \mathcal{S}(H)$ and $F_2 \in \mathcal{M}(\overline{\nu_2}) \cap \mathcal{S}(H)$ such that $F = F_1 \cup F_2$. We have $F_1, F_2 \in \mathcal{I}$ and hence $F \in \mathcal{I} \cup \mathcal{J}$. Since $\mathcal{M}(\overline{\nu}) \cap \mathcal{S}(H) \cap \mathcal{J} = \emptyset$, it follows that $F \in \mathcal{I}$. Therefore $\mathcal{M}(\overline{\nu}) \cap \mathcal{S}(H) \subseteq \mathcal{I}$. $\square$

6 Theorem 1: Random k-Cliques Versus Subcritical G^-

Let $\mathscr{C}$ be a polynomial-size monotone circuit on n-vertex graphs and assume that $\mathrm{E}[\mathscr{C}(\text{random } k\text{-clique})] \geqslant \Omega(1)$ and $\mathrm{E}[\mathscr{C}(G^-)] < 1 - e^{-n^{o(1)}}$. To prove Theorem 1, we will show that $|\mathscr{C}| > \omega(n^{k/4})$.

Lemma 15. $|\{A \in \binom{[n]}{k} : K_A \in \mathcal{M}(\overline{\mathscr{C}})\}| \leqslant \Omega(n^k)$.

Proof. Since $\overline{\mathscr{C}} \geqslant \mathscr{C}$, we have $\mathrm{E}[\overline{\mathscr{C}}(\text{random } k\text{-clique})] \geqslant \Omega(1)$. To prove the lemma, it clearly suffices to show that $\mathrm{E}[\overline{\mathscr{C}}(Q)] < o(1)$ where Q is uniformly distributed among n-vertex graphs whose non-isolated part is isomorphic to an k-clique minus an edge.

For contradiction, assume that $\mathrm{E}[\overline{\mathscr{C}}(Q)] \geqslant \Omega(1)$. For sufficiently small δ (precisely: for any small enough δ so that $p^-(n)$ is supercritical for the existence of k-cliques minus an edge), we have

$$\mathrm{E}[\overline{\mathscr{C}}(G^- \cup Q)] \geqslant 1 - e^{-n^{\Omega(1)}}$$

by a standard application of Janson's inequality (using only that fact that $\overline{\mathscr{C}}$ is a monotone function). But by Lemma 13,

$$\mathrm{E}[\overline{\mathscr{C}}(G^- \cup Q)] \leqslant \mathrm{E}[\mathscr{C}(G^- \cup Q)] + O(|\mathscr{C}| n^\ell e^{-n^\delta}) < 1 - e^{-n^{o(1)}} + n^{O(1)} e^{-n^\delta} < 1 - e^{-n^{o(1)}}.$$

This implies that $1 - e^{-n^{\Omega(1)}} < 1 - e^{-n^{o(1)}}$ and hence $\Omega(1) < o(1)$, which is absurd. Therefore, $\mathrm{E}[\overline{\mathscr{C}}(Q)] < o(1)$. $\square$

Lemma 16. *There exist $\nu \in \mathscr{C}$ and $H \in \mathcal{J}$ such that $|\{H' \in \mathcal{M}(\overline{\nu}) : H \cong H'\}| \geqslant \Omega(n^{\mathrm{supp}(H)})/|\mathscr{C}|$.*

Proof. For contradiction, assume $|\{H' \in \mathcal{M}(\overline{\nu}) : H \cong H'\}| < o(n^{\mathrm{supp}(H)})/|\mathscr{C}|$ for all $\nu \in \mathscr{C}$ and $H \in \mathcal{J}$. By Lemma 14, for all $A \in \binom{[n]}{k}$ such that $K_A \in \mathcal{M}(\overline{\mathscr{C}})$, we can fix choices of $\nu^{(A)} \in \mathscr{C}$ such that $H^{(A)} \in \mathcal{M}(\overline{\nu^{(A)}}) \cap \mathcal{S}(K_A) \cap \mathcal{J}$. Choose representatives $F_1, \ldots, F_\ell$ from each isomorphism class of graphs in $\mathcal{J}$. Note that each graph isomorphic to F_i belongs to $\mathcal{S}(K_A)$ for $O(n^{k-\mathrm{supp}(F_i)})$ different $A \in \binom{[n]}{k}$. It follows that for every $\nu \in \mathscr{C}$,

$$\begin{aligned}
\left|\left\{A \in \tbinom{[n]}{k} : K_A \in \mathcal{M}(\overline{\mathscr{C}}),\, \nu^{(A)} = \nu\right\}\right| &= \sum_{i \in \{1,\ldots,\ell\}} \left|\left\{A \in \tbinom{[n]}{k} : K_A \in \mathcal{M}(\overline{\mathscr{C}}),\, \nu^{(A)} = \nu,\, H^{(A)} \cong F_j\right\}\right| \\
&\leqslant \sum_{i \in \{1,\ldots,\ell\}} \left|\left\{(A, H) : A \in \tbinom{[n]}{k},\, H \in \mathcal{M}(\overline{\nu}) \cap \mathcal{S}(K_A),\, H \cong F_j\right\}\right| \\
&< \sum_{i \in \{1,\ldots,\ell\}} O(n^{k-\mathrm{supp}(F_i)}) o(n^{\mathrm{supp}(F_i)})/|\mathscr{C}| \\
&\leqslant o(n^k)/|\mathscr{C}| \quad (\text{since } \ell = O(1)).
\end{aligned}$$

Summing over all $\nu \in \mathscr{C}$, we have $|\{A \in \binom{[n]}{k} : K_A \in \mathcal{M}(\overline{\mathscr{C}})\}| < o(n^k)$. But this contradicts Lemma 15. Therefore, there exist ν and H as in the statement of the lemma. $\square$

We now show that $|\mathscr{C}| \geqslant \omega(n^{k/4})$. By Lemma 16, there exist $\nu \in \mathscr{C}$ and $H \in \mathcal{J}$ such that $|\{H' \in \mathcal{M}(\overline{\mathscr{C}}_\nu) : H \cong H'\}| \geqslant \Omega(n^{\mathrm{supp}(H)})/|\mathscr{C}|$. On the other hand, since $\overline{\mathscr{C}}_\nu$ is $\star$-closed, Lemma 9 implies $|\{H' \in \mathcal{M}(\overline{\mathscr{C}}_\nu) : H \cong H'\}| \leqslant O(n^{|E_H|(\frac{2}{k-1} + 2\delta)})$. Therefore,

$$\log_n |\mathscr{C}| > \mathrm{supp}(H) - |E_H|(\tfrac{2}{k-1} + 2\delta) - o(1).$$

Since δ can be chosen arbitrarily small,

$$\limsup_{n\to\infty} \log_n |\mathscr{C}| \geqslant \min_{H \in \mathcal{J}} \mathrm{supp}(H) - \tfrac{2}{k-1}|E_H|.$$

Among all graphs $H \in \mathcal{J}$, this quantity turns out to be minimal when H is isomorphic to an $\lceil \frac{k}{2} \rceil$-clique minus an edge (by a straightforward argument). Thus, we have

$$\limsup_{n\to\infty} \log_n |\mathscr{C}| \geqslant \lceil \tfrac{k}{2} \rceil - \tfrac{2}{k-1}\left(\tbinom{\lceil k/2 \rceil}{2} - 1\right) = \begin{cases} \frac{k+1}{4} + \frac{9}{4(k-1)} & \text{if } k \text{ even}, \\ \frac{k+1}{4} + \frac{2}{k-1} & \text{if } k \text{ odd}. \end{cases}$$

Therefore $|\mathscr{C}| > \omega(n^{k/4})$, concluding the proof of Theorem 1. $\square$

7 Theorem 2: $\omega(n^{k/4})$ Lower Bound

We now prove Theorem 2, the $\omega(n^{k/4})$ lower bound on the size of monotone circuits solving k-clique on random graphs. Theorem 2 follows from Theorem 1 via the following lemma. (Recall that $G^\theta = G(n, n^{-2/(k-1)})$ is a random graph at the k-clique threshold.)

Lemma 17. *Suppose f is a boolean function which solves k-clique on random graphs. Then*

$$\mathrm{E}\left[f(G^\theta \cup G^-) \,\middle|\, \omega_k(G^\theta) = 0\right] < o(1), \qquad \mathrm{E}\left[f(G^\theta \cup K_A) \,\middle|\, \omega_k(G^\theta) = 0\right] > 1 - o(1).$$

The proof, which involves showing the statistical indistinguishability of two random graph distributions, is given in Appendix B.

Proof of Theorem 2. Suppose $\mathscr{C}$ is a monotone circuit which solves the k-clique problem on random graphs. For any graph H, form a new monotone circuit $\mathscr{C}^H$ simply by relabeling input nodes in $\mathscr{C}$ corresponding to edges in H by the constant 1. Note that $\mathscr{C}^H(G) = \mathscr{C}(G \cup H)$ for all graphs G.

Because $\mathscr{C}$ solves k-clique on random graphs, Lemma 17 directly implies

$$\mathrm{E}\left[\mathscr{C}^{G^\theta}(G^-) \,\middle|\, \omega_k(G^\theta) = 0\right] < o(1), \qquad \mathrm{E}\left[\mathscr{C}^{G^\theta}(K_A) \,\middle|\, \omega_k(G^\theta) = 0\right] > 1 - o(1).$$

Therefore, there exists a graph H such that $\mathrm{E}[\mathscr{C}^H(K_A)] > 1/2$ and $\mathrm{E}[\mathscr{C}^H(G^-)] < 1/2$. Theorem 1 applied to the circuit $\mathscr{C}^H$ implies $(|\mathscr{C}| =)\ |\mathscr{C}^H| > \omega(n^{k/4})$, completing the proof of Theorem 2. $\square$

8 Theorem 3: $n^{k/4+O(1)}$ Upper Bound

The monotone circuits described in this section are adapted from the AC^0 circuits of Amano [3]. As opposed to previous sections, here we consider constant-depth monotone circuits with unbounded fan-in.

Fix a large constant c and small $\varepsilon > 0$ (to be determined). For $j \in \{1, \dots, k\}$, let $t_j = \min\{1, (j-1)(\alpha - \varepsilon)\}$ and fix any $\mathcal{S}_j$ of $[n]$ into sets of size n^{t_j} (i.e., $\lfloor n^{t_j} \rfloor$ or $\lfloor n^{t_j} \rfloor + 1$).

Following Amano [3], we say that a graph G is *good* if for all $j \in \{2, \dots, k\}$ and every $(j-1)$-clique A in G, it holds that $\left|\left\{b \in X : A \cup \{b\} \text{ is a clique in } G\right\}\right| \leqslant c$ for all sets X in the partition $\mathcal{S}_j$. Note that every k-clique (i.e., K_A for every $A \in \binom{[n]}{k}$) is good.

Lemma 18 (Amano [3])**.** *For sufficiently small ε and large c, the random graph $G(n,p)$ is almost surely good for every function $p : \mathbb{N} \longrightarrow [0,1]$ such that $p(n) < o(p^+(n))$.* $\square$

We now define a monotone circuit $\mathscr{C}$ that solves k-clique on all good graphs. As a first step, we fix a set $\mathcal{F}$ of (hash) functions from $[n]$ to $[\log n]$ (i.e., $\{1, \dots, \lceil \log n \rceil\}$) such that $|\mathcal{F}| = O(\log n)$ and

$$(\dagger) \qquad \forall A \in \tbinom{[n]}{\leqslant c}\ \exists f \in \mathcal{F}\ |f(A)| = |A|.$$

(Such $\mathcal{F}$ exists by a probabilistic argument: simply pick $O(\log n)$ functions $[n] \longrightarrow [\log n]$ uniformly at random.)

For all $j \in \{1, \dots, k\}$ and $(S_1, \dots, S_j) \in \mathcal{S}_1 \times \dots \times \mathcal{S}_j$ and $f_1, \dots, f_j \in \mathcal{F}$ and $w_1, \dots, w_j \in [\log n]$ and $a_j \in S_j$ such that $f_j(a_j) = w_j$, we (inductively) define monotone circuits

$$\mathscr{C}_{a_j}^{S_1,\dots,S_j; f_1,\dots,f_j; w_1,\dots,w_j} = \bigwedge_{i \in \{1,\dots,j-1\}} \bigvee_{a_i \in S_i \,:\, f_i(a_i) = w_i} \mathsf{edge}_{\{a_i,a_j\}} \wedge \mathscr{C}_{a_i}^{S_1,\dots,S_i; f_1,\dots,f_i; w_1,\dots,w_i}.$$

Here $\mathsf{edge}_{\{a_i,a_j\}}$ is the input node (variable) corresponding to the potential edge between vertices a_i and a_j. Note that for $j=1$, the circuit $\mathscr{C}_{a_1}^{S_1;f_1;w_1}$ computes the constant function 1 (since $\bigwedge_{i\in\emptyset}$ is vacuously 1). We now define $\mathscr{C}$ by

$$\mathscr{C} = \bigvee_{S_1\in\mathcal{S}_1} \bigwedge_{f_1\in\mathcal{F}} \bigvee_{w_1\in[\log n]} \cdots \bigvee_{S_k\in\mathcal{S}_k} \bigwedge_{f_k\in\mathcal{F}} \bigvee_{w_k\in[\log n]} \bigvee_{a_k\in S_k \,:\, f_k(a_k)=w_k} \mathscr{C}_{a_k}^{S_1,\dots,S_k;f_1,\dots,f_k;w_1,\dots,w_k}.$$

The next two lemmas are proved in Appendix C.

Lemma 19. *If G contains a k-clique, then $\mathscr{C}(G)=1$.*

Lemma 20. *If G is good and $\mathscr{C}(G)=1$, then G contains a k-clique.*

Lemma 21. *$\mathscr{C}$ has size $O(n^{(k/4)+2})$ for sufficiently small ε.*

Proof. Note that $|\mathscr{C}| \leqslant O\big(|\mathcal{S}_1|\cdots|\mathcal{S}_k||\mathcal{F}|^k(\log n)^k n\big) \leqslant O\big(n(\log n)^{2k} n^{\sum_{j=1}^k 1-t_j}\big)$. We have

$$\begin{aligned}\textstyle\sum_{j=1}^k 1-t_j = \sum_{j=1}^k 1-\min(1,(j-1)(\alpha-\varepsilon)) &< \tbinom{k}{2}\varepsilon + \textstyle\sum_{j=1}^{\lceil\frac{k-1}{2}\rceil} 1-\frac{2(j-1)}{(k-1)} \\ &= \tbinom{k}{2}\varepsilon + \begin{cases}\frac{k}{4}+\frac{1}{2}+\frac{1}{2(k-1)} & \text{if } k \text{ even}\\ \frac{k}{4}+\frac{1}{4} & \text{if } k \text{ odd}\end{cases} \\ &< \tfrac{k}{4}+\tfrac{2}{3} \quad \text{(for sufficiently small } \varepsilon).\end{aligned}$$

Hence, (for sufficiently small ε) we have $|\mathscr{C}| \leqslant O\big(n(\log n)^{2k} n^{\binom{k}{2}\varepsilon+\frac{k}{4}+\frac{2}{3}}\big) \leqslant O(n^{(k/4)+2})$. □

Proof of Theorem 3. By Lemmas 18, 19 and 20, the circuit $\mathscr{C}$ solves k-clique a.a.s. on $G(n,p)$ for every function $p(n)$ which is $o(p^+(n))$. On the other hand, the constant function 1 solves k-clique a.a.s. on $G(n,p)$ for every function $p(n)$ which is $\omega(p^\theta(n))$. Thus, for an appropriate function $m(n)$ (for example, $m(n)=n^2\sqrt{p^+p^\theta}$), $\mathscr{C}\vee\mathsf{Threshold}_{m(n)}$ is a constant-depth monotone of size $n^{k/4+O(1)}$ solving k-clique on random graphs, where $\mathsf{Threshold}_{m(n)}$ is an $O(n\log n)$-size constant-depth monotone circuit which has value 1 on graphs with at least $m(n)$ edges. □

9 Future Directions

One question raised by this work is whether the $\omega(n^{k/4})$ average-case monotone lower bound of Theorem 2 can be sharpened to hold for circuits solving the k-clique problem a.a.s. on $G(n,p)$ for a single threshold function $p(n)$. (This sharper average-case lower bound was shown for AC^0 circuits in [18].)

Although the AC^0 lower bound of [18] and the monotone lower bound of the present paper use very different tools, it would be interesting to find a common underlying principal explaining the similar $\omega(n^{k/4})$ result (perhaps a single proof that neatly generalizes both the AC^0 and monotone results).

On a deeper level, it is tempting to speculate that $n^{k/4+\Theta(1)}$ might be the average-case complexity of the k-clique problem for general boolean circuits. Of course, this would imply $\mathrm{P}\neq\mathrm{NP}$. (We point out that the worst-case complexity of k-clique is known to be $n^{(kw/3)+O(1)}$, currently $\approx n^{0.792k+O(1)}$, where w is the exponent of matrix multiplication [16].)

References

[1] Noga Alon and Ravi B. Boppana. The monotone circuit complexity of boolean functions. *Combinatorica*, 7(1):1–22, 1987.

[2] Noga Alon and Joel Spencer. *The Probabilistic Method, 3rd Edition.* John Wiley, 2008.

[3] Kazuyuki Amano. Subgraph isomorphism on AC^0 circuits. In *IEEE Conference on Computational Complexity*, pages 9–18, 2009.

[4] Kazuyuki Amano and Akira Maruoka. The potential of the approximation method. *SIAM J. Computing*, 33(2):433–447, 2004.

[5] Kazuyuki Amano and Akira Maruoka. A superpolynomial lower bound for a circuit computing the clique function with at most $(1/6)\log\log n$ negation gates. *SIAM J. Comput.*, 35(1):201–215, 2005.

[6] Paul Beame. Lower bounds for recognizing small cliques on CRCW PRAM's. *Discrete Appl. Math.*, 29(1):3–20, 1990.

[7] Béla Bollobás. Threshold functions for small subgraphs. *Math. Proc. Camb. Phil. Soc.*, 90:197–206, 1981.

[8] Béla Bollobás. *Random Graphs (2nd Edition).* Cambridge University Press, 2001.

[9] Paul Erdős and Richard Rado. Intersection theorems for systems of sets. *J. London Math. Soc.*, 35:85–90, 1960.

[10] Armin Haken. Counting bottlenecks to show monotone P $\neq$ NP. In *Proc. 36th FOCS*, pages 36–40, 1995.

[11] Danny Harnik and Ran Raz. Higher lower bounds on monotone size. In *Proc. 32nd STOC*, pages 378–387, 2000.

[12] Johan Håstad. Almost optimal lower bounds for small depth circuits. In *STOC '86: Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing*, pages 6–20, 1986.

[13] Svante Janson. Poisson approximation for large deviations. *Random Struct. Algorithms*, 1(2):221–230, 1990.

[14] Svante Janson, Tomasz Łuczak, and Andrzej Ruciński. *Random Graphs.* John Wiley, 2000.

[15] Stasys Jukna. *Extremal Combinatorics with Applications in Computer Science.* Springer, Heidelberg, 2001.

[16] Jaroslav Nesetril and Svatopluk Poljak. On the complexity of the subgraph problem. *Comment. Math. Univ. Carolinae.*, 26(2):415–419, 1985.

[17] Alexander A. Razborov. Lower bounds on the monotone complexity of some boolean functions. *Doklady Akademii Nauk SSSR*, 281:798–801, 1985. English translation in Soviet Math. Doklady 31 (1985), 354–357.

[18] Benjamin Rossman. On the constant-depth complexity of k-clique. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 721–730, 2008.

A Janson's Inequality

Lemma 22 (Janson's Inequality [13], see also Ch. 8 of [2])**.** *Let X be a finite set and let $\{U_i\}_{i\in I}$ be a family of subsets of X. Let W be a random subset of X such that events $\{x \in W\}_{x\in X}$ are mutually independent. Let Λ_i be the event that $U_i \subseteq W$. Define μ and Δ by*

$$\mu = \textstyle\sum_{i\in I} \Pr(\Lambda_i), \qquad \Delta = \textstyle\sum_{i,j\in I : U_i\cap U_j \neq \emptyset} \Pr(\Lambda_i \wedge \Lambda_j).$$

Then

$$\Pr(\textstyle\bigwedge_{i\in I} \overline{\Lambda_i}) \leqslant \exp(-\mu + (\Delta/2)).$$

If $\Delta \geqslant \mu$, then moreover

$$\Pr(\textstyle\bigwedge_{i\in I} \overline{\Lambda_i}) \leqslant \exp(-\mu^2/2\Delta).$$

So in particular

$$\Pr(\textstyle\bigwedge_{i\in I} \overline{\Lambda_i}) \leqslant \exp(-\min(\mu/2, \mu^2/2\Delta)).$$

B Proof of Lemma 17

We need one preliminary lemma.

Lemma 23. *For $p(n) \in \Theta(n^{-2/(k-1)})$ and $G = G(n,p)$ and uniform random $A \in \binom{[n]}{k}$, distributions*

(i) *G conditioned on $\omega_k(G) = 1$*

(ii) *$G \cup K_A$ conditioned on $\omega_k(G) = 0$*

have total variation distance $o(1)$. That is,

$$\sum_{n\text{-vertex graphs } H} \Big| \Pr\big(G = H \,\big|\, \omega_k(G) = 1\big) - \Pr\big(G \cup K_A = H \,\big|\, \omega_k(G) = 0\big) \Big| < o(1).$$

Proof. The total variation distance between distributions (i) and (ii) is easily seen to equal

$$\Pr\big(\omega_k(G\cup K_A) \geqslant 2 \,\big|\, \omega_k(G) = 0\big) + \sum_{H\,:\,\omega_k(H)=1} \Big| \Pr\big(G = H \,\big|\, \omega_k(G) = 1\big) - \Pr\big(G \cup K_A = H \,\big|\, \omega_k(G) = 0\big) \Big|$$

This is at most

$$\Pr\big(\omega_k(G\cup K_A) \geqslant 2 \,\big|\, \omega_k(G) = 0\big) + \max_{H\,:\,\omega_k(H)=1} \left| 1 - \frac{\Pr\big(G \cup K_A = H \,\big|\, \omega_k(G) = 0\big)}{\Pr\big(G = H \,\big|\, \omega_k(G) = 1\big)} \right|.$$

Note that, for every graph H satisfying $\omega_k(H) = 1$, we have

$$\frac{\Pr\big(G \cup K_A = H \,\big|\, \omega_k(G) = 0\big)}{\Pr\big(G = H \,\big|\, \omega_k(G) = 1\big)} = \frac{\Pr\big(\omega_k(G\cup K_A) = 1\big)}{\Pr\big(\omega_k(G) = 0\big)} \Pr\big(\omega_k(G) = 0 \,\big|\, G \cup K_A = H\big).$$

(To see this, first notice that $\Pr\big(G = H \,\big|\, \omega_k(G) = 1\big) = \Pr\big(G \cup K_A = H \,\big|\, \omega_k(G \cup K_A) = 1\big)$. The rest is just Bayes' Theorem and cancellation.) Therefore, it suffices to show

(a) $\Pr\big(\omega_k(G\cup K_A) \geqslant 2 \,\big|\, \omega_k(G) = 0\big) < o(1)$,

(b) $\dfrac{\Pr\left(\omega_k(G \cup K_A) = 1\right)}{\Pr\left(\omega_k(G) = 0\right)} = \Theta(1)$,

(c) $\Pr\left(\omega_k(G) = 0 \,\middle|\, G \cup K_A = H\right) > 1 - o(1)$ where H is an arbitrary graph satisfying $\omega_k(H) = 1$.

Having thus broken up the problem, we briefly argue each of (a), (b), (c). For (a), we have

$$\begin{aligned}
\Pr\left(\omega_k(G \cup K_A) \geqslant 2 \,\middle|\, \omega_k(G) = 0\right) &\leqslant \sum_{B \in \binom{[n]}{k} \,:\, |A \cap B| \geqslant 2} \Pr\left(K_B \setminus K_A \subseteq G \,\middle|\, \omega_k(G) = 0\right) \\
&\leqslant \sum_{B \in \binom{[n]}{k} \,:\, |A \cap B| \geqslant 2} \Pr\left(K_B \setminus K_A \subseteq G\right) \quad \text{(by Harris' Theorem)} \\
&= \sum_{j=2}^{k-1} \binom{n-k}{k-j} p^{\binom{k}{2} - \binom{j}{2}} \leqslant \sum_{j=2}^{k-1} O\left(n^{k-j-\frac{2}{k-1}(\binom{k}{2} - \binom{j}{2})}\right) \leqslant O(n^{-1}).
\end{aligned}$$

For (b), note that

$$\frac{\Pr\left(\omega_k(G \cup K_A) = 1\right)}{\Pr\left(\omega_k(G) = 0\right)} = \Pr\left(\omega_k(G \cup K_A) = 1 \,\middle|\, \omega_k(G) = 0\right) + \frac{\Pr\left(\omega_k(G \cup K_A) = 1 \text{ and } \omega_k(G) = 1\right)}{\Pr\left(\omega_k(G) = 0\right)}.$$

Using the fact that $\Pr\left(\omega_k(G) = 0\right) \geqslant \Omega(1)$ (since p is a threshold function), it suffices to show

$$\begin{aligned}
&\Pr\left(\omega_k(G \cup K_A) = 1 \,\middle|\, \omega_k(G) = 0\right) > 1 - o(1), \\
&\Pr\left(\omega_k(G \cup K_A) = 1 \text{ and } \omega_k(G) = 1\right) < o(1).
\end{aligned}$$

The first inequality follows immediately from (a); the second inequality follows from similar calculation to (a).

Finally, for (c), letting H be an arbitrary graph with a unique k-clique (call it B), we have

$$\begin{aligned}
\Pr\left(\omega_k(G) = 0 \,\middle|\, G \cup K_A = H\right) &= \Pr\left(\omega_k(G) = 0 \,\middle|\, G \cup K_B = H\right) \\
&= \Pr\left(\omega_k(G) = 0 \,\middle|\, G \setminus K_B = H \setminus K_B\right) \\
&= \Pr\left(K_B \nsubseteq G \,\middle|\, G \setminus K_B = H \setminus K_B\right) \\
&= \Pr\left(K_B \nsubseteq G\right) \quad \text{(by independence)} \\
&= p^{\binom{k}{2}} = O(n^{-k}).
\end{aligned}$$

□

Proof of Lemma 17. To prove the first inequality of Lemma 17, we note that $G^\theta \cup G^-$ has distribution $G(n, p^\theta + p^- - p^\theta p^-)$ and $p^\theta(n) + p^-(n) - p^\theta(n)p^-(n)$ is a threshold function in $\Theta(n^{-2/(k-1)})$. In particular, the event that $\omega_k(G^\theta \cup G^-) = 0$ holds with probability $\geqslant \Omega(1)$. Therefore,

$$\mathrm{E}\left[f(G^\theta \cup G^-) \,\middle|\, \omega_k(G^\theta \cup G^-) = 0\right] < o(1).$$

It now suffices to show that $\omega_k(G^\theta) = 0 \iff \omega_k(G^\theta \cup G^-) = 0$ holds a.a.s. This follows from the fact that $\Pr\left(\omega_k(G^\theta) = 0\right) \sim \Pr\left(\omega_k(G^\theta \cup G^-) = 0\right)$, which is a consequence of $p^\theta(n) \sim p^\theta(n) + p^-(n) - p^\theta(n)p^-(n)$ (by the well-known asymptotics of $\mathrm{E}\left[\omega_k(G(n,p))\right]$, see e.g. [8]).

For the second inequality of Lemma 17, we note that because $p^\theta(n)$ is a threshold function, random graphs

- $G^\theta \cup K_A$ conditioned on $\omega_k(G^\theta) = 0$
- G^θ conditioned on $\omega_k(G^\theta) = 1$

have total variation distance $o(1)$ by Lemma 23. It follows that

$$\mathrm{E}\left[f(G^\theta \cup K_A) \,\middle|\, \omega_k(G^\theta) = 0\right] > \mathrm{E}\left[f(G^\theta) \,\middle|\, \omega_k(G^\theta) = 1\right] - o(1) > 1 - o(1).$$

□

C Proofs of Lemmas 19 and 20

Proof of Lemma 19. Suppose $\{a_1, \ldots, a_k\}$ is a k-clique in G. We will show, even stronger than Lemma 19, that $\mathscr{C}'(G) = 1$ for the circuit $\mathscr{C}' \leqslant \mathscr{C}$ defined by

$$\mathscr{C}' = \bigvee_{S_1 \in \mathcal{S}_1} \cdots \bigvee_{S_k \in \mathcal{S}_k} \bigwedge_{f_1 \in \mathcal{F}} \bigvee_{w_1 \in [\log n]} \cdots \bigwedge_{f_k \in \mathcal{F}} \bigvee_{w_k \in [\log n]} \bigvee_{a_k \in S_k : f_k(a_k) = w_k} \mathscr{C}_{a_k}^{S_1,\ldots,S_k;f_1,\ldots,f_k;w_1,\ldots,w_k}.$$

For all $j \in [k]$, let $S_j \in \mathcal{S}_j$ be such that $a_j \in S_j$. Let f_1 be any element of $\mathcal{F}$ (adversarily chosen). We select $w_1 = f_1(a_1)$ (in order to satisfy $\mathscr{C}'$). Now let f_2 be any element of $\mathcal{F}$ (again adversarily chosen). We select $w_2 = f_2(a_2)$, and so on. After k rounds, it is clear that $\mathscr{C}_{a_j}^{S_1,\ldots,S_j;f_1,\ldots,f_j;w_1,\ldots,w_j}(G) = 1$ for all $j \in [k]$. Hence $\mathscr{C}'(G) = 1$. $\square$

Proof of Lemma 20. Suppose G is good and $\mathscr{C}(G) = 1$. Fix a witness $S_1 = \{a_1\}$ (for the first $\bigvee_{S_1 \in \mathcal{S}_1}$, viewed as an existential quantifier). Let f_1 be any function in $\mathcal{F}$ (adversarily chosen, corresponding to the universal quantifier $\bigwedge_{f_1 \in \mathcal{F}}$). We then fix a witness w_1 with respect to S_1 and f_1. Note that $f_1(a_1) = w_1$ has to hold. In particular, note that a_1 is the *unique* element of $S_1 \cap f_1^{-1}(w_1)$ such that $\mathscr{C}_{a_1}^{S_1;f_1;w_1} = 1$.

Next fix a witness S_2 with respect to S_1, f_1, w_1. Since G is good, there are at most c different $b_2 \in S_2$ such that $\{a_1, b_2\}$ is an edge in G. By definition of $\mathcal{F}$, there is a function $f_2 \in \mathcal{F}$ which takes different values on all elements of the set $\{b_2 \in S_2 : \{a_1, b_2\}$ is an edge in $G\}$. Fix a witness w_2 with respect to S_1, S_2, f_1, f_2, w_1. Note that there exists *unique* $a_2 \in S_2 \cap f_2^{-1}(w_2)$ such that $\mathscr{C}_{a_2}^{S_1,S_2;f_1,f_2;w_1,w_2}(G) = 1$.

We continue in this manner. Assume we have chosen $S_1, \ldots, S_j, f_1, \ldots, f_j, w_1, \ldots, w_j$ such that there exist unique $a_i \in S_i \cap f_i^{-1}(w_i)$ satisfying $\mathscr{C}_{a_i}^{S_1,\ldots,S_i;f_1,\ldots,f_i;w_1,\ldots,w_i}(G) = 1$ for $i \in \{1, \ldots, j\}$. Fix a witness S_{j+1} with respect to $S_1, \ldots, S_j, f_1, \ldots, f_j, w_1, \ldots, w_j$. Since G is good and $\{a_1, \ldots, a_j\}$ is a clique in G, there are at most c different $b_{j+1} \in S_{j+1}$ such that $\{a_1, \ldots, a_j, b_{j+1}\}$ is a clique in G. By definition of $\mathcal{F}$, there is a function $f_{j+1} \in \mathcal{F}$ which takes different values on all elements of the set $\{b_{j+1} \in S_{j+1} : \{a_1, \ldots, a_j, b_{j+1}\}$ is a clique in $G\}$. Let w_{j+1} be a witness with respect to $S_1, \ldots, S_{j+1}$, $f_1, \ldots, f_{j+1}$, $w_1, \ldots, w_j$. Once again there is a unique $a_{j+1} \in S_{j+1} \cap f_{j+1}^{-1}(w_{j+1})$ such that $\mathscr{C}_{a_{j+1}}^{S_1,\ldots,S_{j+1};f_1,\ldots,f_{j+1};w_1,\ldots,w_{j+1}}(G) = 1$.

Now consider the sequence of vertices $a_1, \ldots, a_k$ produced by this argument. To conclude the proof, note that $\{a_1, \ldots, a_k\}$ must be a k-clique in G. $\square$