

The Monotone Complexity of k -Clique on Random Graphs

Benjamin Rossman
MIT

Abstract—It is widely suspected that Erdős-Rényi random graphs are a source of hard instances for clique problems. Giving further evidence for this belief, we prove the first average-case hardness result for the k -clique problem on monotone circuits. Specifically, we show that no monotone circuit of size $O(n^{k/4})$ solves the k -clique problem with high probability on $G(n, p)$ for two sufficiently far-apart threshold functions $p(n)$ (for instance $n^{-2/(k-1)}$ and $2n^{-2/(k-1)}$). Moreover, the exponent $k/4$ in this result is tight up to an additive constant.

One technical contribution of this paper is the introduction of *quasi-sunflowers*, a new relaxation of sunflowers in which petals may overlap slightly on average. A “quasi-sunflower lemma” (à la the Erdős-Rado sunflower lemma) leads to our novel lower bounds within Razborov’s method of approximations.

Keywords-clique; monotone circuits; average-case complexity; quasi-sunflowers

I. INTRODUCTION

In this paper, we study the average-case complexity of k -CLIQUE on monotone circuits. By k -CLIQUE, we mean the decision problem of testing whether an n -vertex graph G contains a clique (= complete subgraph) of fixed constant size k . The *worst-case* monotone complexity of k -CLIQUE has been well-known since the mid-80’s: Razborov [13] proved a lower bound of $\omega(n^k / \log^{2k} n)$ (improved to $\omega(n^k / \log^k n)$ by Alon and Boppana [1]), while the brute-force algorithm (a monotone DNF) gives an upper bound of $O(n^k)$. However, the *average-case* complexity of k -CLIQUE has been a mystery until recently. In a previous work [15], we proved a lower bound of $\omega(n^{k/4})$ for the average-case complexity of k -CLIQUE on AC^0 circuits (= constant-depth polynomial-size circuits with gates of unbounded fan-in). This lower bound was shown to be nearly tight by Amano [3], who gave circuits of size $n^{k/4+O(1)}$ and depth $O(k)$ solving k -CLIQUE in the average-case. In this paper, we prove a similar $\omega(n^{k/4})$ average-case lower bound for k -CLIQUE on monotone circuits (which, moreover, is also tight).

But what is the “average case” for k -CLIQUE? For any monotone graph property \mathcal{P} (such as k -CLIQUE), there is a natural class of distributions to consider: Erdős-Rényi random graphs $G(n, p)$ where $p(n)$ is a threshold function for \mathcal{P} . (That is, $G(n, p)$ is the random graph on n vertices in which edges are independently present with probability $p(n)$, where $\Pr[G(n, p) \in \mathcal{P}]$ is bounded away from 0 and 1.) Our results in [15] and the present paper support an intuition that *random graphs at the threshold are a source*

*of hard instances for clique problems.*¹ This intuition goes back to a question raised by Karp [12] in ‘76. Karp observed that, while the balanced random graph $G(n, 1/2)$ is known to have maximum clique size $\sim 2 \log n$, the greedy algorithm with high probability only finds a clique of size $\sim \log n$. Karp asked whether any polynomial algorithm almost surely finds a clique of size $(1 + \epsilon) \log n$ for some constant $\epsilon > 0$. Despite having received considerable attention over the years, this question remains wide open today.

We point out that Karp’s question “scales down” to $G(n, p)$ where $p(n) = \Theta(n^{-2/(k-1)})$ is a threshold function for k -CLIQUE. Here the maximum clique size is almost surely k or $k - 1$, yet most of the time the greedy algorithm only finds a clique of size $\lfloor k/2 \rfloor$ or $\lceil k/2 \rceil$. Not coincidentally, $G(n, p)$ has the most expected cliques of sizes $\lfloor k/2 \rfloor$ or $\lceil k/2 \rceil$ (see §VI). We remark that for every $\epsilon \in [0, \frac{1}{2}]$, we can find a clique of size $\sim (\frac{1}{2} + \epsilon)k$ with high probability simply by running the greedy algorithm $n^{\epsilon^2 k + O(1)}$ times. In particular, for $\epsilon = \frac{1}{2}$ this gives an upper bound of $n^{k/4 + O(1)}$ on the average-case complexity of k -CLIQUE (effectively by finding all cliques in $G(n, p)$).² As we show, this upper bound can be implemented on monotone circuits. More interesting: our lower bound (stated precisely in §III) shows that this naive algorithm is best possible for monotone circuits which solve k -CLIQUE at two threshold functions such as $n^{-2/(k-1)}$ and $2n^{-2/(k-1)}$.

Outline: In §II we give some basic definitions and state a few background lemmas. We then formally state our results in §III. In §IV we discuss Razborov’s lower bound and the approximation method. In §V we introduce a new relaxation of sunflowers called *quasi-sunflowers* and prove a “Quasi-sunflower Lemma” (similar to the Erdős-Rado Sunflower Lemma), which may be of independent interest. In §VI we divide the subgraphs of K_k into three “sizes”. In §VII we define a closure operation in the lattice of monotone graph functions (in the style of the approximation method). Our main theorems are proved in §VIII and §IX. We state our conclusions in §X.

¹Similar beliefs about random SAT at the threshold are common in statistical physics.

²There are also deterministic algorithms, such as the Bron-Kerbosch algorithm [6], which solve the MAXCLIQUE problem on a graph G in time $\text{poly}(n) \times \#\{\text{cliques in } G\}$ by exhaustively finding all cliques.

II. PRELIMINARIES

Let $k \geq 5$ be an arbitrary, but fixed, integer. Let n be a positive number and let $[n] = \{1, \dots, n\}$. All asymptotic statements and notation ($O(\cdot)$, $o(\cdot)$, $\Omega(\cdot)$, etc.) refer to growing n . The hidden constants in this asymptotic notation are universal (in particular, independent of k). Expressions with *high probability* (w.h.p.) and *almost surely* (a.s.) mean with probability tending to 1 as $n \rightarrow \infty$.

For a set X and integer $t \geq 0$, $\binom{X}{t}$ denotes the set of t -element subsets of X .

$\log(\cdot)$ is the base-2 logarithm and $\ln(\cdot)$ is the natural logarithm.

A. Graphs and patterns

Graphs in this paper are finite simple graphs. Formally, a graph is a pair $G = (V_G, E_G)$ where V_G is a finite set and $E_G \subseteq \binom{V_G}{2}$. We denote by \mathcal{G}^n the set of graphs with vertex set $[n]$. By default *graphs* are elements of \mathcal{G}^n . By distinction, we refer to constant-size graphs with no isolated vertices as *patterns*.

For both graphs and patterns, \cup is the union operation and \subseteq is the subgraph/subpattern relation. For $\ell \in \mathbb{N}$, K_ℓ denotes the complete pattern with vertex set $\{1, \dots, \ell\}$ and edge set $\binom{\{1, \dots, \ell\}}{2}$. An ℓ -clique in a graph G is a set of ℓ vertices with all $\binom{\ell}{2}$ possible edges present between.

B. Monotone functions and minterms

By *graph function* we always mean a function from \mathcal{G}^n to $\{0, 1\}$. A graph function f is *monotone* if $f(G_1) \leq f(G_2)$ whenever G_1 is a subgraph of G_2 .

A graph H is a *minterm* of monotone graph function f if $f(H) = 1$ and $f(H') = 0$ for every proper subgraph $H' \subset H$. For a pattern P , a minterm H is a P -minterm if the induced pattern on the non-isolated vertices of H is isomorphic to P . The set of minterms (resp. P -minterms) of f is denoted $\mathcal{M}(f)$ (resp. $\mathcal{M}(f, P)$).

C. Monotone circuits

A *monotone circuit* on m variables is an acyclic directed graph C with m sources (called *inputs*) and one sink (called the *output*). Non-source nodes in C (called *gates*) have in-degree 2 and are labelled either \wedge or \vee . C computes a monotone function $\{0, 1\}^m \rightarrow \{0, 1\}$ in the natural way.

For $m = \binom{n}{2}$, we view C as computing a monotone graph function (via a natural bijection between \mathcal{G}^n and $\{0, 1\}^{\binom{n}{2}}$). The value of C on a graph G is denoted $C(G)$. $\mathcal{M}(C)$ (resp. $\mathcal{M}(C, P)$) denotes the set of minterms (resp. P -minterms) of the function computed by C .

The *size* of a monotone circuit is the number of gates it contains. The *monotone complexity* of a monotone function $f : \{0, 1\}^m \rightarrow \{0, 1\}$ is the size of the smallest monotone circuit that computes f .

D. Probability

We consistently represent random objects by boldface symbols (\mathbf{G} , \mathbf{W} , etc.). For a set X and $p \in [0, 1]$, notation $\mathbf{W} \subseteq_p X$ expresses that \mathbf{W} is a random subset of X where each $x \in X$ belongs to \mathbf{W} independently with probability p . $\text{Po}(\lambda)$ is the Poisson distribution with mean λ . d_{TV} is total variation distance ($= 1/2$ of the ℓ_1 -distance between two distributions).

For a function $p : \mathbb{N} \rightarrow [0, 1]$, $\mathbf{G} \sim G(n, p)$ is the *Erdős-Rényi random graph* on n vertices in which each element of $\binom{[n]}{2}$ is an edge independently with probability $p(n)$ (i.e., $V_{\mathbf{G}} = [n]$ and $E_{\mathbf{G}} \subseteq_p \binom{[n]}{2}$).

We denote by $\mathbf{K}_k (= \mathbf{K}_k(n))$ the *random planted k -clique* on n vertices (i.e., $V_{\mathbf{K}_k} = [n]$ and $E_{\mathbf{K}_k} = \binom{U}{2}$ where U is uniform random k -element subset of $[n]$).

For background, we state two basic lemmas on k -cliques in $G(n, p)$. (For proofs and additional background, see any of [2], [5], [10].) The first lemma says that $\Theta(n^{-2/(k-1)})$ is precisely the class of threshold functions $p(n)$.

Lemma 1. *If $p(n) = o(n^{-2/(k-1)})$ then w.h.p. $G(n, p)$ is k -clique-free. If $p(n) = \omega(n^{-2/(k-1)})$ then w.h.p. $G(n, p)$ contains a k -clique.*

The second lemma concerns random graphs $G(n, p)$ at threshold functions $p(n) \sim cn^{-2/(k-1)}$. In particular, it includes the fact that the number of k -cliques in $G(n, p)$ is asymptotically Poisson.

Lemma 2. *Denote by $\kappa(G)$ the number of k -cliques in a graph G . Fix $c > 0$ and let $\mathbf{G} \sim G(n, cn^{-2/(k-1)})$ and $\mathbf{K}_k \sim \text{Plant}(n, K_k)$ and $\mathbf{X} \sim \text{Po}(c \binom{n}{k} / k!)$. For $t \in \mathbb{N}$, let $\mathbf{G}_t \sim G(n, cn^{-2/(k-1)})$ conditioned on $\kappa(\mathbf{G}_t) = t$. Then*

$$\begin{aligned} d_{\text{TV}}(\kappa(\mathbf{G}), \mathbf{X}) &= o(1), \\ d_{\text{TV}}(\mathbf{G}_{t+1}, \mathbf{G}_t \cup \mathbf{K}_k) &= o(1), \\ \lim_{n \rightarrow \infty} d_{\text{TV}}(\kappa(\mathbf{G}), \kappa(\mathbf{G} \cup \mathbf{K}_k)) &= d_{\text{TV}}(\mathbf{X}, \mathbf{X} + 1) < 1. \end{aligned}$$

Later on we will also need Janson's inequality (which we state in §V).

III. OUR RESULTS

Let $p = n^{-2/(k-1)}$ (to fix a particular threshold function) and $\delta = k^{-2}$ (just think of δ as a sufficiently small constant). Let \mathbf{G}^- , \mathbf{G} , \mathbf{G}^+ be independent Erdős-Rényi random graphs:

$$\mathbf{G}^- \sim G(n, p^{1+\delta}), \quad \mathbf{G} \sim G(n, p), \quad \mathbf{G}^+ \sim G(n, p^{1-\delta}).$$

By Lemma 1, w.h.p. \mathbf{G}^- is k -clique-free and \mathbf{G}^+ contains a k -clique. That is, with respect to the property of containing a k -clique, \mathbf{G}^- and \mathbf{G}^+ are *subcritical* and *supercritical*.

Our main theorem is a lower bound for monotone circuits which solve k -CLIQUE w.h.p. on both \mathbf{G} and $\mathbf{G} \cup \mathbf{G}^-$.

Theorem 3. *No monotone circuit of size $O(n^{k/4})$ solves k -CLIQUE w.h.p. on both \mathbf{G} and $\mathbf{G} \cup \mathbf{G}^-$.*

Note that $\mathbf{G} \cup \mathbf{G}^-$ is an Erdős-Rényi random graph $G(n, \tilde{p})$ where $\tilde{p} = p + (1-p)p^{1+\delta}$, which is also a threshold function for k -CLIQUE. Moreover, since $\tilde{p} = p + o(p)$, the numbers of k -cliques in \mathbf{G} and $\mathbf{G} \cup \mathbf{G}^-$ are asymptotically equivalent Poisson random variables (by Lemma 2).³

Remark 4. Theorem 3 implies that no monotone circuit of size $O(n^{k/4})$ solves k -CLIQUE w.h.p. on both $G(n, p)$ and $G(n, 2p)$. This follows from the observation that if monotone graph functions f and g agrees w.h.p. on both $G(n, p_1)$ and $G(n, p_2)$ for $p_1, p_2 : \mathbb{N} \rightarrow [0, 1]$ such that $p_1(n) \leq p_2(n)$, then f and g also agree w.h.p. on $G(n, q)$ for every $q : \mathbb{N} \rightarrow [0, 1]$ such that $p_1(n) \leq q(n) \leq p_2(n)$. By the same observation, Theorem 3 may be stated as an average-case hardness result on a single distribution $G(n, \mathbf{q})$ where \mathbf{q} equals p with probability $1/2$ and \tilde{p} with probability (or, alternatively, where \mathbf{q} is uniformly distributed in $[p, \tilde{p}]$).

It would be nice to reduce or eliminate the “gap” of $\tilde{p} - p \sim p^{1+\delta}$ between threshold functions p and \tilde{p} in Theorem 3, in order to get an average-case hardness result for monotone circuits at a single threshold (like the lower bound of [15] for AC^0 circuits). We conjecture that the gap can be eliminated entirely (in §X). However, there is reason to believe that this gap is hard to close, since the single-threshold version of Theorem 3 seems to require techniques that go beyond the approximation method.⁴

Preliminary to Theorem 3, we prove the following lower bound:

Theorem 5. *If C is a monotone circuit of size $O(n^{k/4})$ such that $E[C(\mathbf{K}_k)] = 1 - o(1)$, then $E[C(\mathbf{G}^-)] = 1 - \exp(-\Omega(n^\delta))$.*

Theorem 5 should be compared with the following fact (a consequence of Janson’s inequality).

Fact 6. *If f is a monotone graph function such that $E[f(\mathbf{K}_k)] = 1 - o(1)$, then $E[f(\mathbf{G}^+)] = 1 - \exp(-\Omega(n^\delta))$ (irrespective of the monotone circuit complexity of f).*

Note that subcritical \mathbf{G}^- is replaced by supercritical \mathbf{G}^+ in Fact 6.

The full version of this paper contains two additional results (omitted here due to length constraints):

- We strengthen Theorems 3 and 5 by removing the fan-in 2 restriction on monotone circuits (that is, we get the same $\omega(n^{k/4})$ lower bounds for monotone circuits with \wedge and \vee gates of unbounded fan-in).
- We construct monotone circuits of size $n^{k/4+O(1)}$ and depth $3k$ that solve k -CLIQUE w.h.p. on $G(n, p)$ for all functions $p : \mathbb{N} \rightarrow [0, 1]$. This shows that $k/4$ is tight up to

³Notwithstanding, the total variation distance between random graphs \mathbf{G} and $\mathbf{G} \cup \mathbf{G}^-$ is $1 - o(1)$.

⁴What about the random graph with exactly $\lceil \binom{n}{k} p \rceil$ edges? Note that the monotone complexity of k -CLIQUE on this distribution is polynomially equivalent to the non-monotone complexity, since we are deal with a slice function.

an additive constant in Theorem 3. Moreover, in view of Theorem 7, it demonstrates a gap between the worst-case and average-case monotone complexity of k -CLIQUE.

IV. RAZBOROV’S APPROXIMATION METHOD

In a seminal paper [13], Razborov proved the first lower bounds on the monotone complexity of k -CLIQUE.

Theorem 7. *k -CLIQUE has monotone circuit complexity $\omega(n^k / \log^{2k} n)$.*⁵

Razborov in fact shows something stronger. Let \mathbf{H} be the uniform random complete $(k-1)$ -partite graph with vertex set $[n]$ (that is, $E_{\mathbf{H}} = \{\{i, j\} \in \binom{[n]}{2} : \pi(i) \neq \pi(j)\}$ for uniform random function $\pi : [n] \rightarrow \{1, \dots, k\}$). The following result is also from [13] (note the similarity to Theorem 5):

Theorem 8. *If C is a monotone circuit of size $O(n^k / \log^{2k} n)$ such that $E[C(\mathbf{K}_k)] = 1 - o(1)$, then $E[C(\mathbf{H})] = 1 - o(1)$.*⁶

The technique introduced in [13] to prove Theorem 8 is known as the *approximation method*. (Note: The following summary is for background only. Our lower bounds do not explicitly follow this framework.) The idea of the approximation method is to replace the lattice $(\mathfrak{M}, \wedge, \vee)$ of monotone functions $\{0, 1\}^m \rightarrow \{0, 1\}$ with a smaller lattice $(\overline{\mathfrak{M}}, \overline{\wedge}, \overline{\vee})$ where $\overline{\mathfrak{M}} \subset \mathfrak{M}$ such that

- $\overline{\mathfrak{M}}$ contains the function $x \mapsto x_i$ for every $i \in \{1, \dots, m\}$, and
- $\overline{\wedge}$ and $\overline{\vee}$ are the g.l.b. and l.u.b. operations in $\overline{\mathfrak{M}}$ with respect to the natural partial order on functions (i.e., $f \leq g$ iff $f(x) \leq g(x)$ for all $x \in \{0, 1\}^m$).

For every monotone circuit C on m variables, there is a corresponding $\{\overline{\wedge}, \overline{\vee}\}$ -circuit \overline{C} in which the \wedge and \vee gates are replaced by $\overline{\wedge}$ and $\overline{\vee}$ gates. Note that \overline{C} computes a function in $\overline{\mathfrak{M}}$.

Let Δ_0 and Δ_1 be two distributions on $\{0, 1\}^m$ (e.g., the random graphs \mathbf{H} and \mathbf{K}_k from Theorem 8). Suppose our goal is to prove that no monotone circuit C of size S separates Δ_0 and Δ_1 in the sense that $E[C(\Delta_0)] = o(1)$ and $E[C(\Delta_1)] = 1 - o(1)$. Then it suffices to show that:

- 1) no function $f \in \overline{\mathfrak{M}}$ satisfies $E[f(\Delta_0)] = o(1)$ and $E[f(\Delta_1)] = 1 - o(1)$,
- 2) for all $f, g \in \overline{\mathfrak{M}}$,

$$\begin{aligned} E[(f \overline{\vee} g)(\Delta_0)] - E[(f \vee g)(\Delta_0)] &= o(1/S), \\ E[(f \overline{\wedge} g)(\Delta_1)] - E[(f \wedge g)(\Delta_1)] &= o(1/S). \end{aligned}$$

⁵This bound is for constant k . [13] also gives lower bounds for k which depends on n .

⁶Moreover, if C is a monotone circuit of size $n^{k-\Omega(1)}$ such that $E[C(\mathbf{K}_k)] = 1 - o(1)$, then $E[C(\mathbf{H})] = 1 - \exp(-n^{\Omega(1)})$.

By bounding “local errors” in this way, (2) shows that for any C of size S ,

$$\begin{aligned} \mathbb{E}[C(\Delta_0)] &\leq \mathbb{E}[\overline{C}(\Delta_0)] + o(1), \\ \mathbb{E}[C(\Delta_1)] &\geq \mathbb{E}[\overline{C}(\Delta_1)] - o(1). \end{aligned}$$

It follows that C cannot satisfy both $\mathbb{E}[C(\Delta_0)] = o(1)$ and $\mathbb{E}[C(\Delta_1)] = 1 - o(1)$.

Of course, being able to show (1) and (2) for given Δ_0 and Δ_1 depends on a clever choice of the lattice $\overline{\mathfrak{M}}$. To prove Theorem 8, Razborov defines a lattice $\overline{\mathfrak{M}}$ where the l.u.b. operation $\overline{\vee}$ involves “plucking” large sunflowers among the minterms of the function $f \vee g$. (For a full description of $\overline{\mathfrak{M}}$, see [13] or [1].)

Our proof of Theorem 5 does not precisely follow this framework. Rather, we work with a “one-sided version” of the approximation method (via a closure operation $\text{cl} : \mathfrak{M} \rightarrow \overline{\mathfrak{M}}$ defined in §VII). The difference is merely a matter of exposition: our proof could easily be formulated in terms of an approximating lattice $\overline{\mathfrak{M}}$.

V. QUASI-SUNFLOWERS

In this section we introduce a new relaxation of sunflowers called *quasi-sunflowers* (parameterized by $p \in [0, 1]$ and $\gamma \geq 0$). Like sunflowers, quasi-sunflowers are special hypergraphs. Some definitions: A *hypergraph* is a family \mathcal{F} of subsets of a set X (i.e., $\mathcal{F} \subseteq \wp(X)$). Elements of \mathcal{F} are called *hyperedges*. For an integer $s \geq 1$, \mathcal{F} is *s-uniform* if every hyperedge has size s (i.e., $\mathcal{F} \subseteq \binom{X}{s}$).

A *sunflower* is a hypergraph \mathcal{F} such that the intersection of any two distinct hyperedges coincides with the intersection $\bigcap \mathcal{F}$ ($= \bigcap_{U \in \mathcal{F}} U$) of all hyperedges. The set $\bigcap \mathcal{F}$ is called the *core* and sets $U \setminus \bigcap \mathcal{F}$ where $U \in \mathcal{F}$ are called *petals* (note that petals are mutually disjoint). An essential fact about sunflowers is:

Fact 9 (Erdős-Rado Sunflower Lemma [7]). *Every s-uniform hypergraph \mathcal{F} of size $> s!(N-1)^s$ contains a sunflower of size N .*

Quasi-sunflowers are a relaxation of sunflowers in which petals may overlap slightly on average. While other variants of sunflowers are studied in extremal combinatorics (see Ch. 7 of [11]), the following definition appears to be new.

Definition 10. *Let \mathcal{F} be a hypergraph on a set X and let $Y \subseteq \bigcap \mathcal{F}$. For $p \in [0, 1]$ and $\gamma \geq 0$, we say that \mathcal{F} is (p, γ) -quasi-sunflower over Y if for the random set $\mathbf{W} \subseteq_p X$,*

$$\Pr[\mathbf{W} \cup Y \text{ contains a hyperedge of } \mathcal{F}] \geq 1 - e^{-\gamma}.$$

Observation 11. Let $\mathcal{F} \subseteq \binom{X}{s}$ be an s -uniform sunflower of size n . Then \mathcal{F} is a (p, np^s) -quasi-sunflower for every $p \in [0, 1]$. To see this, let $Y = \bigcap \mathcal{F}$ and note that for $\mathbf{W} \subseteq_p X$, the probability that $\mathbf{W} \cup Y$ contains a hyperedge of \mathcal{F} is

$$1 - (1 - p^{s-|Y|})^n \geq 1 - \exp(np^{s-|Y|}) \geq 1 - \exp(np^s).$$

For small p , this $\gamma = np^s$ is nearly tight if $Y = \emptyset$, but (as we will see) is far from tight if $Y \neq \emptyset$.

We suspect that wherever s -uniform sunflowers are used in monotone circuit lower bounds (e.g., [1], [4], [13]), one could just as well work with $(1/2, N/2^s)$ -quasi-sunflowers instead. That is, Definition 10 captures the essential property of sunflowers for these applications. Perhaps one even gets stronger bounds (as we do in this paper) by virtue of the following result.

Theorem 12 (“Quasi-sunflower lemma”). *For all $p \in [0, 1]$ and $\gamma \geq 1$ and $s \geq 1$, every s -uniform hypergraph of size $\geq s!(2.47\gamma/p)^s$ contains a (p, γ) -quasi-sunflower.*

Remark 13. It follows from Fact 9 and Obs. 11 that every s -uniform hypergraph of size $\geq s!(\gamma/p^s)^s$ contains a (p, γ) -quasi-sunflower (namely, a sunflower of size γ/p^s). Theorem 12 is a significant quantitative improvement of this observation.

In order to prove Theorem 12, we need a probabilistic result known as Janson’s inequality [9] (also see Ch. 2 of [10] and Ch. 8 of [2] for background.)

Lemma 14 (Janson’s Inequality). *Let \mathcal{F} be a hypergraph on a set X . Let \mathbf{W} be a random subset of X such that events $x \in \mathbf{W}$ for $x \in X$ are mutually independent (for example, $\mathbf{W} \subseteq_p X$). Define μ and Δ by*

$$\begin{aligned} \mu &= \sum_{U \in \mathcal{F}} \Pr[U \subseteq \mathbf{W}], \\ \Delta &= \sum_{\substack{U, V \in \mathcal{F}: \\ U \neq V, U \cap V \neq \emptyset}} \Pr[U \cup V \subseteq \mathbf{W}]. \end{aligned}$$

Then $\Pr\left[\bigwedge_{U \in \mathcal{F}} U \not\subseteq \mathbf{W}\right] \leq \exp\left(-\min\left\{\frac{\mu}{2}, \frac{\mu^2}{2\Delta}\right\}\right)$.

Our proof of Theorem 12 uses Janson’s inequality within an inductive argument that resembles proofs of the Erdős-Rado Sunflower Lemma.

Proof of “Quasi-sunflower Lemma” (Theorem 12): Consider the sequence ℓ_1, ℓ_2, \dots defined by $\ell_1 = 1$ and $\ell_s = 2 \sum_{t=1}^{s-1} \binom{s}{t} \ell_t$ for $s \geq 2$. We have $\ell_s \leq s! \ln^{-s}(3/2)$ ($< s!2.47^s$) by induction: for $s \geq 2$, assuming $\ell_t \leq t! \ln^{-t}(3/2)$ for every $t \in \{1, \dots, s-1\}$, we have

$$\begin{aligned} \ell_s &\leq 2 \sum_{t=1}^{s-1} \binom{s}{t} t! \ln^{-t}(3/2) \\ &= 2 \left(\sum_{t=1}^{s-1} \frac{\ln^{s-t}(3/2)}{(s-t)!} \right) s! \ln^{-s}(3/2) \\ &\leq 2 \left(-1 + \sum_{j=0}^{\infty} \frac{\ln^j(3/2)}{j!} \right) s! \ln^{-s}(3/2) \\ &= s! \ln^{-s}(3/2). \end{aligned}$$

Suppose \mathcal{F} is an s -uniform hypergraph of size $\geq \ell_s(\gamma/p)^s$. Arguing by induction on s , we claim that \mathcal{F} contains an (p, γ) -quasi-sunflower (proving the theorem). In the base case where $s = 1$, let $\mathbf{W} \subseteq_p X$ and note that events $U \subseteq \mathbf{W}$ for $U \in \mathcal{F}$ are mutually independent. Therefore,

$$\Pr \left[\bigwedge_{U \in \mathcal{F}} U \not\subseteq \mathbf{W} \right] = (1-p)^{|\mathcal{F}|} \leq (1-p)^{\gamma/p} \leq e^{-\gamma},$$

so \mathcal{F} itself is a (p, γ) -quasi-sunflower over the empty set.

For the induction step, let $s \geq 2$ and assume the claim holds for $t \in \{1, \dots, s-1\}$. For every $A \subseteq X$ with $1 \leq |A| \leq s-1$, let

$$\mathcal{F}_A = \{U \setminus A : U \in \mathcal{F} \text{ such that } A \subseteq U\}.$$

Note that \mathcal{F}_A is an $(s - |A|)$ -uniform hypergraph. We now consider two cases.

First Case: Suppose there exist $t \in \{1, \dots, s-1\}$ and $A \in \binom{X}{t}$ such that $|\mathcal{F}_A| \geq \ell_{s-t}(\gamma/p)^{s-t}$. By the induction hypothesis, \mathcal{F}_A contains a (p, γ) -quasi-sunflower \mathcal{F}' over some $Y' \subseteq \bigcap \mathcal{F}'$. Note that $\{U \cup A : U \in \mathcal{F}'\} \subseteq \mathcal{F}$ is a (p, γ) -quasi-sunflower over $Y' \cup A$.

Second Case: Suppose $|\mathcal{F}_A| \leq \ell_{s-t}(\gamma/p)^{s-t}$ for all $t \in \{1, \dots, s-1\}$ and $A \in \binom{X}{t}$. We will show that \mathcal{F} itself is a (p, γ) -quasi-sunflower over the empty set. Let $\mathbf{W} \subseteq_p X$ and define μ and Δ exactly as in the statement of Janson's inequality (Lemma 14), which says:

$$\Pr \left[\bigwedge_{U \in \mathcal{F}} U \not\subseteq \mathbf{W} \right] \leq \exp \left(-\min \left\{ \frac{\mu}{2}, \frac{\mu^2}{2\Delta} \right\} \right).$$

Thus, it suffices to show that $\min\{\mu/2, \mu^2/2\Delta\} \geq \gamma$.

Clearly $\mu = |\mathcal{F}|p^s$ since $\Pr[U \subseteq \mathbf{W}] = p^s$ for every $U \in \mathcal{F}$. Since $|\mathcal{F}| \geq \ell_s(\gamma/p)^s$ and $\ell_s \geq 2$ (as $s \geq 1$) and $\gamma^s \geq \gamma$ (as $\gamma \geq 1$), we have $\mu/2 \geq \gamma$.

It remains to show that $\mu^2/2\Delta \geq \gamma$. For every $t \in \{1, \dots, s-1\}$, we have $\sum_{A \in \binom{X}{t}} |\mathcal{F}_A| = \binom{s}{t} |\mathcal{F}|$ since each hyperedge in \mathcal{F} is counted $\binom{s}{t}$ times in this summation. Therefore,

$$\begin{aligned} \sum_{A \in \binom{X}{t}} |\mathcal{F}_A|^2 &\leq |\mathcal{F}| \sum_{A \in \binom{X}{t}} |\mathcal{F}_A| \\ &\leq \mu \binom{s}{t} \ell_{s-t} \gamma^{s-t} p^{t-2s} \end{aligned}$$

(using $|\mathcal{F}| = \mu p^{-s}$ and $|\mathcal{F}_A| \leq \ell_{s-t}(\gamma/p)^{s-t}$). Noting that $\Pr[U \cup V \subseteq \mathbf{W}] = p^{2s-|U \cap V|}$ for all $U, V \in \mathcal{F}$, we bound

Δ as follows:

$$\begin{aligned} \Delta &= \sum_{\substack{A \subseteq X : \\ 1 \leq |A| \leq s-1}} \sum_{\substack{U, V \in \mathcal{F} : \\ U \cap V = A}} \Pr[U \cup V \subseteq \mathbf{W}] \\ &\leq \sum_{t=1}^{s-1} \left(\sum_{A \in \binom{X}{t}} |\mathcal{F}_A|^2 \right) p^{2s-t} \\ &\leq \mu \sum_{t=1}^{s-1} \binom{s}{t} \ell_{s-t} \gamma^{s-t} \\ &\leq \mu \gamma^{s-1} \sum_{t=1}^{s-1} \binom{s}{t} \ell_t \quad (\text{using } \gamma^t \leq \gamma^{s-1}) \\ &= \frac{\mu \gamma^{s-1} \ell_s}{2} \quad (\text{by definition of } \ell_s). \end{aligned}$$

Completing the proof, we have

$$\frac{\mu^2}{2\Delta} \geq \frac{\mu}{\gamma^{s-1} \ell_s} = \frac{|\mathcal{F}| p^s}{\gamma^{s-1} \ell_s} \geq \gamma. \quad \blacksquare$$

VI. SMALL, MEDIUM, LARGE

Let $\mathbf{G} \sim \mathbf{G}(n, \Theta(n^{-2/(k-1)}))$ be a random graph at a threshold function for containing k -cliques. It is instructive to calculate the expected number of ℓ -cliques in \mathbf{G} for $\ell \in \{0, \dots, k\}$:

$$\mathbb{E}[\# \text{ of } \ell\text{-cliques in } \mathbf{G}] = \Theta(n^{\ell - \frac{2}{k-1} \binom{\ell}{2}}).$$

Letting $\lambda = \ell/k$, we have

$$\ell - \frac{2}{k-1} \binom{\ell}{2} = \lambda(1-\lambda)k + O(1).$$

Note that $\lambda(1-\lambda)k$ is maximal with value $k/4$ for $\lambda = 1/2$. (Indeed, $\ell - \frac{2}{k-1} \binom{\ell}{2}$ is maximal for $\ell \in \{[k/2], \lceil k/2 \rceil\}$.)

The fact that \mathbf{G} has many cliques of ‘‘intermediate’’ size $\sim k/2$ and few cliques of size $\leq \varepsilon k$ or $\geq (1-\varepsilon)k$ for small $\varepsilon > 0$ motivates the following definition. (The large number of ‘‘intermediate’’ subgraphs plays an important part in our lower bounds.)

Definition 15. A pattern P is:

- small if $|V_P| < k/2$,
- medium if $|V_P| \geq k/2$ and there exist small patterns P_1 and P_2 such that $P = P_1 \cup P_2$, and
- large otherwise.

A graph is small, medium or large according to the induced pattern on its non-isolated vertices.

A key fact to keep in mind is that the union of two small patterns/graphs is small or medium (but never large). Note that the complete pattern K_ℓ is small if $\ell < k/2$ and large otherwise (but never medium). An important example of medium pattern is

$$P = K_{\lceil k/2 \rceil} - \{\text{a single edge}\}.$$

Note that P is the union of two overlapping copies of the small pattern $K_{\lceil k/2 \rceil - 1}$. In fact, this pattern P gives the optimal bound in the following lemma.

Lemma 16. *For every medium pattern P ,*

$$|V_P| - \frac{2}{k-1}|E_P| \geq \frac{k+1}{4} + \frac{2}{k-1}.$$

Proof: Let P be a medium pattern which minimizes $|V_P| - \frac{2}{k-1}|E_P|$. By definition of medium, P is the union of two small patterns P_1 and P_2 . We can assume that P_1 and P_2 are complete, since we only decrease $|V_{P_1 \cup P_2}| - \frac{2}{k-1}|E_{P_1 \cup P_2}|$ by replacing P_1 and P_2 with the (also small) complete patterns with the same vertices. Let $a = |V_P|$, $b = |V_{P_1}|$, $c = |V_{P_2}|$ and note that

$$|V_P| - \frac{2}{k-1}|E_P| = a - \frac{2}{k-1} \left(\binom{b}{2} + \binom{c}{2} - \binom{b+c-a}{2} \right).$$

First, suppose $k = 2t+1$ is odd. Integers a, b, c satisfy $1 \leq b, c \leq t$ and $t+1 \leq a \leq b+c$. Relaxing integrality, let α, β, γ be reals minimizing $\alpha - \frac{1}{t} \left(\binom{\beta}{2} + \binom{\gamma}{2} - \binom{\beta+\gamma-\alpha}{2} \right)$ subject to $1 \leq \beta, \gamma \leq t$ and $t+1 \leq \alpha \leq \beta+\gamma$. Note that $\beta = \gamma$ since, if not, by replacing β and γ with their mean $(\beta+\gamma)/2$ we reduce the objective function while still satisfying the constraints. Thus, our task becomes minimizing $f(\alpha, \beta) = \alpha + \frac{1}{t} \left(2\binom{\beta-\alpha}{2} - \frac{2}{t} \binom{\beta}{2} \right)$ subject to $1 \leq \beta \leq t$ and $t+1 \leq \alpha \leq 2\beta$. Since $\frac{d}{d\alpha} f(\alpha, \beta) > 0$ and $\frac{d}{d\beta} f(\alpha, \beta) < 0$ for all α, β satisfying these constraints, it follows that $\alpha = t+1$ and $\beta = t$. Therefore,

$$|V_P| - \frac{2}{k-1}|E_P| \geq f(t+1, t) = \frac{t+1}{2} + \frac{1}{t} = \frac{k+1}{4} + \frac{2}{k-1}.$$

In the case where k is even, we get $|V_P| - \frac{2}{k-1}|E_P| \geq \frac{k+1}{4} + \frac{9}{4(k-1)} \geq \frac{k+1}{4} + \frac{2}{k-1}$ by a similar calculation. ■

Remark 17. $k/4$ in Lemma 16 is precisely the $k/4$ that appears in the exponent of $n^{k/4}$ in our main theorems. In fact, Lemma 16 also accounts for the exponent of $n^{k/4}$ in the lower bound from [15] on the average-case complexity of k -CLIQUE on bounded-depth circuits. It is interesting that the same ‘‘bottleneck’’ arises in the distinct settings of bounded-depth circuits and monotone circuits.

VII. THE APPROXIMATION VIA A CLOSURE OPERATION

In this section we define a closure operation in the lattice of monotone graph functions. Closed functions will be combinatorially ‘‘nice’’ in the sense of having few P -minterms for small and medium patterns P (Lemma 27).

Remark 18. This is essentially one half of Razborov’s approximation method. Typically, one also defines a ‘‘truncation’’ operator which cuts out large minterms. Although we find it more natural to work with a one-sided version of the approximation method, our proof can be translated into Razborov’s original framework (as described in §IV).

Recall that we have fixed $p = n^{-2/(k-1)}$ (a threshold function for the existence of k -cliques) and $\delta = k^{-2}$ (just think of δ as ‘‘sufficiently small’’). Also recall that $\mathbf{G} \sim$

$\mathbf{G}(n, p)$ (at the k -clique threshold) and $\mathbf{G}^- \sim \mathbf{G}(n, p^{1+\delta})$ (below the k -clique threshold, i.e., \mathbf{G}^- is almost surely k -clique-free).

Definition 19. *A monotone graph function $f : \mathcal{G}^n \rightarrow \{0, 1\}$ is closed if for every small-or-medium graph H ,*

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

Observation 20. If f and g are both closed, then so is $f \wedge g$.

Definition 21. *For a monotone graph function f , we denote by $\mathbf{cl}(f)$ the unique minimal closed function such that $f \leq \mathbf{cl}(f)$, called the closure of f .*

Note that $\mathbf{cl}(f)$ is well-defined in view of Obs. 20 and the fact that the constant function 1 is closed.

Remark 22. $\mathbf{cl}(\cdot)$, viewed as an operation on the set of monotone graph functions, is a closure operation in the usual sense. That is, it is increasing ($f \leq \mathbf{cl}(f)$), monotone (if $f \leq g$ then $\mathbf{cl}(f) \leq \mathbf{cl}(g)$) and idempotent ($\mathbf{cl}(\mathbf{cl}(f)) = \mathbf{cl}(f)$).

Definition 23. *We denote by ∇ the operation on monotone graph functions defined by $f \nabla g = \mathbf{cl}(f \vee g)$. For a monotone circuit \mathbf{C} , we denote by $\overline{\mathbf{C}}$ denote the corresponding circuit with basis $\{\wedge, \nabla\}$ in which the \vee -gates in \mathbf{C} are replaced by ∇ -gates. For nodes ν in \mathbf{C} , we denote by $\overline{\nu}$ the corresponding node in $\overline{\mathbf{C}}$.*

Note that $\mathbf{cl}(\mathbf{C})$ (i.e., $\mathbf{cl}(f)$ where f is the function computed by \mathbf{C}) is not necessarily the same function as $\overline{\mathbf{C}}$, although $\overline{\mathbf{C}}$ is indeed a closed function satisfying $\mathbf{C} \leq \overline{\mathbf{C}}$ (i.e., $\mathbf{C}(G) \leq \overline{\mathbf{C}}(G)$ for all graphs G).

Lemma 24. *For every monotone graph function f ,*

$$\Pr [f(\mathbf{G}^-) \neq (\mathbf{cl}(f))(\mathbf{G}^-)] \leq 2^{k^2} n^k e^{-n^\delta}.$$

Proof: We claim that there exist $t \in \mathbb{N}$ and small-or-medium graphs H_1, \dots, H_t and monotone functions $f_0, \dots, f_t : \mathcal{G}^n \rightarrow \{0, 1\}$ such that

- $f_0 = f$,
- $\mathbb{E}[f_{i-1}(\mathbf{G}^- \cup H_i)] \in [1 - e^{-n^\delta}, 1)$,
- $f_i = f_{i-1} \vee \text{Ind}_{H_i}$ where $\text{Ind}_{H_i} : \mathcal{G}^n \rightarrow \{0, 1\}$ is the function $\text{Ind}_{H_i}(G) = 1$ iff $H_i \subseteq G$,
- f_t is closed.

To see this, note that we can generate such a sequence (a priori indefinitely) simply by choosing any suitable H_{i+1} so long as f_i is not closed. This process eventually terminates, since each small or medium graph H appears at most once in the sequence H_1, H_2, \dots . In particular,

$$t \leq |\{\text{small and medium graphs in } \mathcal{G}^n\}| \leq 2^{k^2} n^k.$$

An inductive argument shows that $f_i \leq \mathbf{cl}(f)$ for $i = 1, \dots, t$. In particular $f_t \leq \mathbf{cl}(f)$. Since f_t is closed, this

means that $f_t = \text{cl}(f)$. We now have

$$\begin{aligned} \Pr [f(\mathbf{G}^-) \neq (\text{cl}(f))(\mathbf{G}^-)] &\leq \sum_{i=1}^t \Pr [f_{i-1}(\mathbf{G}^-) \neq f_i(\mathbf{G}^-)] \\ &= \sum_{i=1}^t \Pr [f_{i-1}(\mathbf{G}^-) = 0 \text{ and } H_i \subseteq \mathbf{G}^-] \\ &\leq \sum_{i=1}^t \Pr [f_{i-1}(\mathbf{G}^- \cup H_i) = 0] \\ &\leq 2^{k^2} n^k e^{-n^\delta}. \end{aligned}$$

The next two lemmas follow immediately from Lemma 24. ■

Lemma 25. *For every monotone graph function f , $\mathcal{M}(\text{cl}(f)) \setminus \mathcal{M}(f)$ contains only small and medium graphs.*

Proof: The proof of Lemma 24 shows that there exist small-or-medium graphs H_1, \dots, H_t such that $\text{cl}(f) = f \vee \bigvee_{i=1}^t \text{Ind}_{H_i}$. Thus, $\mathcal{M}(\text{cl}(f)) \subseteq \mathcal{M}(f) \cup \{H_1, \dots, H_t\}$. ■

Lemma 26. *For every monotone circuit C of size $\exp(o(n^\delta))$, $\mathbb{E}[\overline{C}(\mathbf{G}^-)] - \mathbb{E}[C(\mathbf{G}^-)] = \exp(-\Omega(n^\delta))$.*

Proof: For any graph H , note that if $C(H) \neq \overline{C}(H)$ then there exists an \vee -gate ν with children μ_1 and μ_2 in C such that $\overline{\nu}(H) \neq (\overline{\mu_1} \vee \overline{\mu_2})(H)$ (equivalently: $f(H) \neq (\text{cl}(f))(H)$ where f is the function $\overline{\mu_1} \vee \overline{\mu_2}$). It follows that

$$\begin{aligned} \mathbb{E}[\overline{C}(\mathbf{G}^-)] - \mathbb{E}[C(\mathbf{G}^-)] &= \Pr [C(\mathbf{G}^-) \neq \overline{C}(\mathbf{G}^-)] \\ &\leq \sum_{\substack{\vee\text{-gates } \nu \text{ in } C \text{ with} \\ \text{children } \mu_1 \text{ and } \mu_2}} \Pr [\overline{\nu}(\mathbf{G}^-) \neq (\overline{\mu_1} \vee \overline{\mu_2})(\mathbf{G}^-)] \\ &\leq \text{size}(C) 2^{k^2} n^k e^{-n^\delta} \quad (\text{by Lemma 24}) \\ &= \exp(-\Omega(n^\delta)). \end{aligned}$$

The last lemma of this section gives an essential property of closed functions (using Theorem 12 on quasi-sunflowers). ■

Lemma 27. *A closed monotone graph function has at most $k^{k^2} (n^\delta/p^{1+\delta})^{|E_P|}$ P -minterms for every small or medium pattern P .*

Proof: Let f be a closed monotone graph function and let P be a small or medium pattern. Toward a contradiction, assume that $|\mathcal{M}(f, P)| \geq k^{k^2} (n^\delta/p^{1+\delta})^{|E_P|}$. Let $X = \binom{[n]}{X^2}$ and consider the $|E_P|$ -uniform hypergraph $\mathcal{F} \subseteq \binom{X^2}{|E_P|}$ defined by $\mathcal{F} = \{E_F : F \in \mathcal{M}(f, P)\}$. Since $|E_P| \leq k^2/4$ (i.e., no medium pattern has more than $k^2/4$ edges), we have $|E_P|! 2.47^{|E_P|} \leq k^{k^2}$ and hence

$$|\mathcal{F}| = |\mathcal{M}(f, P)| \geq |E_P|! 2.47^{|E_P|} (n^\delta/p^{1+\delta})^{|E_P|}.$$

By Theorem 12, there exists a $(p^{1+\delta}, n^\delta)$ -quasi-sunflower $\mathcal{F}_0 \subseteq \mathcal{F}$ over some $Y \subseteq \bigcap \mathcal{F}$. Let H be the graph with edge set $E_H = Y$. Let $\mathbf{W} \subseteq_{p^{1+\delta}} X$ and note that \mathbf{W} has

the same distribution as $E_{\mathbf{G}^-}$. We have

$$\begin{aligned} \mathbb{E}[f(\mathbf{G}^- \cup H)] &\geq \Pr [\mathbf{G}^- \cup H \text{ contains a } P\text{-minterm of } f] \\ &\geq \Pr [\mathbf{W} \cup Y \text{ contains a hyperedge of } \mathcal{F}_0] \\ &\geq 1 - e^{-n^\delta}. \end{aligned}$$

Since f is closed and H is small or medium, it follows that $f(H) = 1$. Note that H has fewer than $|E_P|$ edges, so in particular H is a proper subgraph of some $F \in \mathcal{M}(f, P)$ such that $E_F \in \mathcal{F}_0$. However, this contradicts the fact that F is a minterm of f . ■

VIII. K vs. \mathbf{G}^-

In the previous section, we defined a closure operation $\text{cl}(\cdot)$ on monotone graph functions and an operation $C \mapsto \overline{C}$ transforming a monotone circuit C into a $\{\wedge, \vee\}$ -circuit \overline{C} . In this section, we prove Theorem 5. We begin by noting a basic fact about minterms.

Observation 28. For all monotone graph functions f and g ,

$$\begin{aligned} \mathcal{M}(f \vee g) &\subseteq \mathcal{M}(f) \cup \mathcal{M}(g), \\ \mathcal{M}(f \wedge g) &\subseteq \{F \cup G : F \in \mathcal{M}(f), G \in \mathcal{M}(g)\}. \end{aligned}$$

That is, every minterm of $f \vee g$ is a minterm of f or a minterm of g and every minterm of $f \wedge g$ is the union of a minterm of f and a minterm of g .

Lemma 29. *For every monotone circuit C and graph $H \in \mathcal{M}(\overline{C}, K_k)$, there exist a gate ν in C and a medium subgraph H' of H such that $H' \in \mathcal{M}(\overline{\nu})$.*

Proof: Suppose $H \in \mathcal{M}(\overline{C}, K_k)$. Let $\mathcal{H} = \{\text{subgraphs of } H\}$ and $\mathcal{A} = \{\text{small graphs}\}$ and $\mathcal{B} = \{\text{medium graphs}\}$. Toward a contradiction, assume that $\mathcal{M}(\overline{\nu}) \cap \mathcal{H} \cap \mathcal{B} = \emptyset$ for every gate ν in C . We will show, by induction on ν , that $\mathcal{M}(\overline{\nu}) \cap \mathcal{H} \subseteq \mathcal{A}$ for every node ν in C . This yields a contradiction, since H belongs to $(\mathcal{M}(\overline{\nu_{\text{out}}}) \cap \mathcal{H}) \setminus \mathcal{A}$ where ν_{out} is the output gate of C .

In the base case when ν is an input node labelled by either 0 or 1 or the indicator function for some edge $e \in \binom{[n]}{2}$, $\mathcal{M}(\nu)$ is respectively either the empty set or $\{\text{the empty graph}\}$ or $\{\text{the graph with only edge } e\}$. In any case, all minterms of ν are small. Since $\overline{\nu} = \nu$, $\mathcal{M}(\overline{\nu}) \cap \mathcal{H} \subseteq \mathcal{A}$.

For the induction step, let ν be a gate in C with children μ_1 and μ_2 and assume that $\mathcal{M}(\overline{\mu_i}) \cap \mathcal{H} \subseteq \mathcal{A}$ for $i \in \{1, 2\}$. If ν is an \wedge -gate, then

$$\begin{aligned} \mathcal{M}(\overline{\nu}) \cap \mathcal{H} &= \mathcal{M}(\overline{\mu_1} \wedge \overline{\mu_2}) \cap \mathcal{H} \quad (\text{we now use Obs. 28}) \\ &= \{F_1 \cup F_2 : F_1 \in \mathcal{M}(\overline{\mu_1}), F_2 \in \mathcal{M}(\overline{\mu_2})\} \cap \mathcal{H} \\ &= \{F_1 \cup F_2 : F_1 \in \mathcal{M}(\overline{\mu_1}) \cap \mathcal{H}, F_2 \in \mathcal{M}(\overline{\mu_2}) \cap \mathcal{H}\} \\ &\subseteq \{F_1 \cup F_2 : F_1, F_2 \in \mathcal{A}\} \quad (\text{since } \mathcal{M}(\overline{\mu_i}) \cap \mathcal{H} \subseteq \mathcal{A}) \\ &\subseteq \mathcal{A} \cup \mathcal{B} \quad (\text{union of two smalls cannot be large}) \\ &\subseteq \mathcal{A} \quad (\text{since } \mathcal{M}(\overline{\nu}) \cap \mathcal{H} \cap \mathcal{B} = \emptyset). \end{aligned}$$

Finally, if ν is a \vee -gate, then

$$\begin{aligned}
\mathcal{M}(\bar{\nu}) \cap \mathcal{H} &= \mathcal{M}(\bar{\mu}_1 \bar{\vee} \bar{\mu}_2) \cap \mathcal{H} \\
&= \mathcal{M}(\text{cl}(\bar{\mu}_1 \bar{\vee} \bar{\mu}_2)) \cap \mathcal{H} \\
&\subseteq (\mathcal{M}(\bar{\mu}_1 \bar{\vee} \bar{\mu}_2) \cup \mathcal{A} \cup \mathcal{B}) \cap \mathcal{H} \quad (\text{Lemma 25}) \\
&\subseteq (\mathcal{M}(\bar{\mu}_1) \cup \mathcal{M}(\bar{\mu}_2) \cup \mathcal{A} \cup \mathcal{B}) \cap \mathcal{H} \quad (\text{Obs. 28}) \\
&\subseteq \mathcal{A} \cup \mathcal{B} \quad (\text{since } \mathcal{M}(\bar{\mu}_i) \cap \mathcal{H} \subseteq \mathcal{A} \text{ for } i \in \{1, 2\}) \\
&\subseteq \mathcal{A} \quad (\text{since } \mathcal{M}(\bar{\nu}) \cap \mathcal{H} \cap \mathcal{B} = \emptyset).
\end{aligned}$$

■

Lemma 30. *For every monotone circuit C , there exists a medium pattern P such that*

$$|\mathcal{M}(\bar{C}, K_k)| \leq (2k)^{k^2} n^{k-|V_P|} (n^\delta/p^{1+\delta})^{|E_P|} \text{size}(C).$$

Proof: By Lemma 29, for each $H \in \mathcal{M}(\bar{C}, K_k)$, there exists a gate μ_H in C and a medium subgraph H' of H such that $H' \in \mathcal{M}(\bar{\mu}_H)$. Fix choices of μ_H and H' for all $H \in \mathcal{M}(\bar{C}, K_k)$. For every gate ν in C and medium pattern P , let $t(\nu, P) = |\{H \in \mathcal{M}(\bar{C}, K_k) : \mu_H = \nu, H' \in \mathcal{M}(\bar{\nu}, P)\}|$.

By a simple counting argument, there exist ν and P such that

$$\frac{|\mathcal{M}(\bar{C}, K_k)|}{\text{size}(C) |\{\text{medium patterns up to isom.}\}|} \leq t(\nu, P).$$

For each $H' \in \mathcal{M}(\bar{\nu}, P)$, there are at most $n^{k-|V_P|}$ different $H \in \mathcal{M}(\bar{C}, K_k)$ of which H' is a subgraph. It follows that

$$t(\nu, P) \leq n^{k-|V_P|} |\mathcal{M}(\bar{\nu}, P)|.$$

Since $\bar{\nu}$ is closed and P is medium, Lemma 27 implies

$$|\mathcal{M}(\bar{\nu}, P)| \leq k^{k^2} (n^\delta/p^{1+\delta})^{|E_P|}.$$

The result follows by combining these three inequalities, together with the bound 2^{k^2} on the number of medium patterns up to isomorphism. ■

Onto the main result:

Proof of Theorem 5: Suppose $f : \mathcal{G}^n \rightarrow \{0, 1\}$ is computed by monotone circuits of size $O(n^{k/4})$ and satisfies $E[f(\mathbf{K}_k)] = 1 - o(1)$. We must show that $E[f(\mathbf{G}^-)] = 1 - \exp(-\Omega(n^\delta))$.

Let C be the circuit computing f . By Lemma 26,

$$\begin{aligned}
E[\bar{C}(\mathbf{G}^-)] - E[f(\mathbf{G}^-)] &= \Pr[f(\mathbf{G}^-) \neq \bar{C}(\mathbf{G}^-)] \\
&= \exp(-\Omega(n^\delta)).
\end{aligned}$$

Therefore, it suffices to show that $E[\bar{C}(\mathbf{G}^-)] = 1$. We will assume that $E[\bar{C}(\mathbf{G}^-)] \neq 1$ and derive a contradiction.

We claim that $|\mathcal{M}(\bar{C}, K_k)| = (1 - o(1)) \binom{n}{k}$. To show this, we consider the pattern $Q = K_k - \{\text{single edge}\}$ and let $\mathbf{H} \sim \text{Plant}(n, Q)$. Since $E[\bar{C}(\mathbf{K}_k)] \geq E[f(\mathbf{K}_k)] = 1 - o(1)$, it is enough to show that $E[\bar{C}(\mathbf{H})] = o(1)$ (i.e., these two inequalities imply that almost every planted k -clique is a minterm of \bar{C}). The argument goes as follows: if we assume that $E[\bar{C}(\mathbf{H})] = \Omega(1)$, then $\Pr[\bar{C}(\mathbf{G}^-)] =$

$1 - \exp(-\Omega(n^{1/k})) \geq 1 - \exp(-n^\delta)$ for sufficiently large n (recall that $\delta = k^{-2}$) by a straightforward application of Janson's inequality (Lemma 14); but since \bar{C} is closed, it follows that \bar{C} (the empty graph) = 1 (contradicting $E[\bar{C}(\mathbf{G}^-)] \neq 1$).

We now invoke Lemma 30, which gives us a medium pattern P such that

$$\begin{aligned}
\text{size}(C) &\geq \frac{|\mathcal{M}(\bar{C}, K_k)|}{(2k)^{k^2} n^{k-|V_P|} (n^\delta/p^{1+\delta})^{|E_P|}} \\
&= (1 - o(1)) \binom{n}{k} \frac{n^{|V_P|} (p^{1+\delta}/n^\delta)^{|E_P|}}{n^k (2k)^{k^2}} \\
&= \Omega\left(\frac{n^{|V_P| - (\frac{2}{k-1}(1+\delta) + \delta)|E_P|}}{k^k (2k)^{k^2}}\right)
\end{aligned}$$

(using $p = n^{-2/(k-1)}$). Recall that $\delta = 1/k^2$ and note that $|E_P| < k^2/4$ (obs: among medium patterns, the disjoint union of two $\lfloor \frac{k-1}{2} \rfloor$ -cliques has the most edges). By Lemma 16, $|V_P| - \frac{2}{k-1}|E_P| \geq \frac{k}{4} + \frac{1}{4} + \frac{2}{k-1}$. We have:

$$\begin{aligned}
|V_P| - \left(\frac{2}{k-1}(1+\delta) + \delta\right)|E_P| \\
> |V_P| - \frac{2}{k-1}|E_P| - \frac{1}{4}\left(1 + \frac{2}{k-1}\right) > \frac{k}{4} + \frac{1}{k}.
\end{aligned}$$

Therefore, $\text{size}(C) = \Omega(n^{(k/4)+(1/k)}/k^k (2k)^{k^2})$. But since k is a constant, this contradicts the hypothesis that C has size $O(n^{k/4})$. ■

IX. $\mathbf{G} \cup \mathbf{K}$ vs. $\mathbf{G} \cup \mathbf{G}^-$

In this section, we prove Theorem 3 using Theorem 5 together with the following lemma.

Lemma 31. *Let $\mathbf{G}_0 \sim G(n, p)$ and condition on \mathbf{G}_0 being k -clique-free.*

1) *If f solves k -CLIQUE w.h.p. on \mathbf{G} , then*

$$E[f(\mathbf{G}_0 \cup \mathbf{K}_k)] = 1 - o(1).$$

2) *If f solves k -CLIQUE w.h.p. on $\mathbf{G} \cup \mathbf{G}^-$, then*

$$E[f(\mathbf{G}_0 \cup \mathbf{G}^-)] = o(1).$$

Proof: Denote by $\kappa(G)$ the number of k -cliques in a graph G .

For (1): Suppose f solves k -CLIQUE w.h.p. on \mathbf{G} . This means, in particular, that $E[f(\mathbf{G}) \mid \kappa(\mathbf{G}) = 1] = 1 - o(1)$. Let $\mathbf{G}_1 \sim G(n, p)$ conditioned on $\kappa(\mathbf{G}_1) = 1$. Note that $E[f(\mathbf{G}_1)] = 1 - o(1)$ (using the fact that $\Pr[\kappa(\mathbf{G}) = 1] = \Omega(1)$). By Lemma 2, random graphs $\mathbf{G}_0 \cup \mathbf{K}_k$ and \mathbf{G}_1 have total variation distance $o(1)$. Therefore, w.h.p. $E[f(\mathbf{G}_0 \cup \mathbf{K}_k)] = 1 - o(1)$.

For (2): Suppose f solves k -CLIQUE w.h.p. on $\mathbf{G} \cup \mathbf{G}^-$. In particular,

$$(\dagger) \quad E[f(\mathbf{G} \cup \mathbf{G}^-) \mid \kappa(\mathbf{G} \cup \mathbf{G}^-) = 0] = o(1).$$

Since $\mathbf{G} \sim G(n, p)$ and $\mathbf{G} \cup \mathbf{G}^- \sim G(n, p + o(p))$, random variables $\kappa(\mathbf{G})$ and $\kappa(\mathbf{G} \cup \mathbf{G}^-)$ converge in distribution to

the same Poisson distribution by Lemma 2. In particular, we have

$$(\ddagger) \quad \Pr[\kappa(\mathbf{G}) = 0] = (1 + o(1)) \Pr[\kappa(\mathbf{G} \cup \mathbf{G}^-) = 0].$$

Thus, we have

$$\begin{aligned} \mathbb{E}[f(\mathbf{G}_0 \cup \mathbf{G}^-)] &= \Pr[f(\mathbf{G} \cup \mathbf{G}^-) = 1 \mid \kappa(\mathbf{G}) = 0] \\ &= \frac{\Pr[f(\mathbf{G} \cup \mathbf{G}^-) = 1 \ \& \ \kappa(\mathbf{G}) = 0]}{\Pr[\kappa(\mathbf{G}) = 0]} \\ &\geq \frac{\Pr[f(\mathbf{G} \cup \mathbf{G}^-) = 1 \ \& \ \kappa(\mathbf{G} \cup \mathbf{G}^-) = 0]}{\Pr[\kappa(\mathbf{G}) = 0]} \\ &\stackrel{(\ddagger)}{=} \frac{\Pr[f(\mathbf{G} \cup \mathbf{G}^-) = 1 \ \& \ \kappa(\mathbf{G} \cup \mathbf{G}^-) = 0]}{(1 + o(1)) \Pr[\kappa(\mathbf{G} \cup \mathbf{G}^-) = 0]} \\ &= (1 - o(1)) \Pr[f(\mathbf{G} \cup \mathbf{G}^-) = 1 \mid \kappa(\mathbf{G} \cup \mathbf{G}^-) = 0] \\ &\stackrel{(\ddagger)}{=} 1 - o(1). \end{aligned}$$

(Under the assumption that f is monotone, (2) can also be proved using the Holley inequality.) ■

Proof of Theorem 3: Let C be a monotone circuit of size $O(n^{k/4})$. Toward a contradiction, assume that C solves k -CLIQUE w.h.p. on both \mathbf{G} and $\mathbf{G} \cup \mathbf{G}^-$. For a graph G , let C^G be the circuit obtained from C by substituting 1 for each input corresponding to an edge in G . Note that C^G computes the function $C^G(H) = C(G \cup H)$.

Let $\mathbf{G}_0 \sim G(n, p)$ conditioned on \mathbf{G}_0 being k -clique-free. Lemma 31 implies that for every constant $\varepsilon > 0$,

$$\begin{aligned} \Pr_{\mathbf{G}_0} \left[\mathbb{E}_{\mathbf{G}^-} [C^{\mathbf{G}_0}(\mathbf{K}_k)] \geq 1 - \varepsilon \right] &= o(1), \\ \Pr_{\mathbf{G}_0} \left[\mathbb{E}_{\mathbf{K}_k} [C^{\mathbf{G}_0}(\mathbf{G}^-)] \leq \varepsilon \right] &= 1 - o(1). \end{aligned}$$

It follows that there is a sequence of monotone circuits of size $O(n^{k/4})$ (namely, $C^{\mathbf{G}_0}$ for almost every \mathbf{G}_0) with expected value $1 - o(1)$ on \mathbf{K}_k and $o(1)$ on \mathbf{G}^- . However, Theorem 5 says this is impossible, giving the desired contradiction. ■

X. FUTURE DIRECTIONS

The main question left open by this work is whether the $\omega(n^{k/4})$ lower bound of Theorem 3 applies to monotone circuits which solve k -CLIQUE w.h.p. at a single threshold. We conjecture that it does.

Conjecture 32. *No monotone circuit of size $O(n^{k/4})$ solves k -CLIQUE w.h.p. on $G(n, n^{-2/(k-1)})$.*

Theorem 3 strongly suggests that this conjecture should be true. However, the approximation method seems to break down when the distributions on positive and negative inputs are brought so closely together.

Finally, it would be interesting to find other applications of quasi-sunflowers. Given the many demonstrated uses of sunflowers, it may be that quasi-sunflowers lead to better results in some cases.

ACKNOWLEDGMENTS

I would like to thank Tomoyuki Hayasaka and Koutaro Nakagawa for their careful reading of an early draft of this paper. I also thank the anonymous referees for their helpful suggestions.

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. k -subgraph isomorphism on AC^0 circuits. *Computational Complexity*, 19(2):183–210, 2010.
- [4] 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 Journal on Computing*, 35(1):201–215, 2005.
- [5] Béla Bollobás. *Random Graphs (2nd Edition)*. Cambridge University Press, 2001.
- [6] Coenraad Bron and Joep Kerbosch. Finding all cliques of an undirected graph (algorithm 457). *Commun. ACM*, 16(9):575–576, 1973.
- [7] Paul Erdős and Richard Rado. Intersection theorems for systems of sets. *J. London Math. Soc.*, 35:85–90, 1960.
- [8] Mikael Goldmann and Johan Håstad. A simple lower bound for the depth of monotone circuits computing clique using a communication game. *Information Processing Letters*, 41(4):221–226, 1992.
- [9] Svante Janson. Poisson approximation for large deviations. *Random Structures and Algorithms*, 1(2):221–230, 1990.
- [10] Svante Janson, Tomasz Łuczak, and Andrzej Ruciński. *Random Graphs*. John Wiley, 2000.
- [11] Stasys Jukna. *Extremal Combinatorics with Applications in Computer Science*. Springer, Heidelberg, 2001.
- [12] Richard M. Karp. Probabilistic analysis of some combinatorial search problems. In J. F. Traub, editor, *Algorithms and Complexity: New Directions and Recent Results*, pages 1–19. Academic Press, 1976.
- [13] 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.
- [14] Alexander A. Razborov. On the method of approximations. In *Proceedings of the 21st Annual ACM Symposium on the Theory of Computing*, pages 167–176, 1989.
- [15] 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.