

Clique Is Hard on Average for Regular Resolution

Albert Atserias

Universitat Politècnica de Catalunya
Department of Computer Science
Barcelona, Spain
atserias@cs.upc.edu

Ilario Bonacina

Universitat Politècnica de Catalunya
Department of Computer Science
Barcelona, Spain
bonacina@cs.upc.edu

Susanna F. de Rezende

KTH Royal Institute of Technology
School of Electrical Engineering and
Computer Science
Stockholm, Sweden
sfd@kth.se

Massimo Lauria

Sapienza Università di Roma
Department of Statistical Sciences
Rome, Italy
massimo.lauria@uniroma1.it

Jakob Nordström

KTH Royal Institute of Technology
School of Electrical Engineering and
Computer Science
Stockholm, Sweden
jakobn@kth.se

Alexander Razborov

University of Chicago
Chicago, USA
razborov@math.uchicago.edu
Steklov Mathematical Institute
Moscow, Russia
razborov@mi.ras.ru

ABSTRACT

We prove that for $k \ll \sqrt[n]{n}$ regular resolution requires length $n^{\Omega(k)}$ to establish that an Erdős-Rényi graph with appropriately chosen edge density does not contain a k -clique. This lower bound is optimal up to the multiplicative constant in the exponent, and also implies unconditional $n^{\Omega(k)}$ lower bounds on running time for several state-of-the-art algorithms for finding maximum cliques in graphs.

CCS CONCEPTS

• **Theory of computation** → **Proof complexity**; • **Mathematics of computing** → **Random graphs**;

KEYWORDS

Proof complexity, regular resolution, k -clique, Erdős-Rényi random graphs

ACM Reference Format:

Albert Atserias, Ilario Bonacina, Susanna F. de Rezende, Massimo Lauria, Jakob Nordström, and Alexander Razborov. 2018. Clique Is Hard on Average for Regular Resolution. In *Proceedings of 50th Annual ACM SIGACT Symposium on the Theory of Computing (STOC'18)*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3188745.3188856>

1 INTRODUCTION

Deciding whether a graph has a k -clique is one of the most basic computational problems on graphs, and has been extensively studied in computational complexity theory ever since it appeared in Karp's list of 21 NP-complete problems [15]. Not only is this problem widely believed to be infeasible to solve exactly—unless $P = NP$

there does not even exist any polynomial-time algorithm for approximating the maximal size of a clique to within a factor $n^{1-\epsilon}$ for any constant $\epsilon > 0$, where n is the number of vertices in the graph [13, 34]. Furthermore, the problem appears to be hard not only in the worst case but also on average in the Erdős-Rényi random graph model—we know of no efficient algorithms for finding cliques of maximum size asymptotically almost surely on random graphs with appropriate edge densities [16, 31].

In terms of upper bounds, the k -clique problem can clearly be solved in time roughly n^k simply by checking if any of the $\binom{n}{k}$ many sets of vertices of size k forms a clique, which is polynomial if k is constant. This can be improved slightly to $O(n^{\omega k/3})$ using algebraic techniques [26], where $\omega \leq 2.373$ is the matrix multiplication exponent, although in practice such algebraic algorithms are outperformed by combinatorial ones [33].

The motivating problem behind this work is to determine the exact time complexity of the clique problem when k is given as a parameter. As noted above, all known algorithms require time $n^{\Omega(k)}$. It appears quite likely that some dependence on k is needed in the exponent, since otherwise we have the parameterized complexity collapse $FPT = W[1]$ [11]. Even more can be said if we are willing to believe the Exponential Time Hypothesis (ETH) [14]—then the exponent has to depend linearly on k [8], so that the trivial upper bound is essentially tight.

Obtaining such a lower bound unconditionally would, in particular, imply $P \neq NP$, and so currently seems completely out of reach. But is it possible to prove $n^{\Omega(k)}$ lower bounds in restricted but nontrivial models of computation? For circuit complexity, this challenge has been met for circuits that are of bounded depth [30] or are monotone [32]. In this paper we focus on computational models that are powerful enough to capture algorithms that are used in practice.

When analysing such algorithms, it is convenient to view the execution trace as a proof establishing the maximal clique size for the input graph. In particular, if this graph does not have a k -clique, then the trace provides an efficiently verifiable proof of the statement that the graph is k -clique-free. If one can establish a lower bound on the length of such proofs, then this implies a lower

Definitive Version in the ACM Digital Library:
[https://dl.acm.org/citation.cfm?
doid=3188745.3188856](https://dl.acm.org/citation.cfm?doid=3188745.3188856)

bound on the running time of the algorithm, and this lower bound holds even if the algorithm is a non-deterministic heuristic that somehow magically gets to make all the right choices. This brings us to the topic of *proof complexity* [9], which can be viewed as the study of upper and lower bounds in restricted nondeterministic computational models.

Using a standard reduction from k -clique to SAT, we can translate the problem of k -cliques in graphs to that of satisfiability of formulas in conjunctive normal form (CNF). If an algorithm for finding k -cliques is run on a graph G that is k -clique-free, then we can extract a proof of the unsatisfiability of the corresponding CNF formula—the k -clique formula on G —from the execution trace of the algorithm. Is it possible to show any non-trivial lower bound on the length of such proofs? Specifically, does the *resolution* proof system—the method of reasoning underlying state-of-the-art SAT solvers [2, 23, 25]—require length $n^{\Omega(k)}$, or at least $n^{\omega_k(1)}$, to prove the absence of k -cliques in a graph? This question was asked in, e.g., [7] and remains open.

The hardness of k -clique formulas for resolution is also a problem of intrinsic interest in proof complexity, since these formulas escape known methods of proving resolution lower bounds for a range of interesting values of k including $k = O(1)$. In particular, the interpolation technique [18, 28], the random restriction method [4], and the size-width lower bound [5] all seem to fail.

To make this more precise, we should mention that some previous works do use the size-width method, but only for very large k . It was shown in [3] that for $n^{5/6} \ll k \leq n/3$ resolution requires length $\exp(n^{\Omega(1)})$ to certify that a dense enough Erdős-Rényi random graph is k -clique-free. The constant hidden in the $\Omega(1)$ increases with the density of the graph and, in particular, for very dense graphs and $k = n/3$ the length required is $2^{\Omega(n)}$. Also, for a specially tailored CNF encoding, where the i th member of the claimed k -clique is encoded in binary by $\log n$ variables, a lower bound of $n^{\Omega(k)}$ for $k \leq \log n$ can be extracted from a careful reading of [21]. However, in the more natural unary encodings, where indicator variables specify whether a vertex is in the clique, the size-width method cannot yield more than a $2^{\Omega(k^2/n)}$ lower bound since there are resolution proofs of width $O(k)$. This bound becomes trivial when $k \leq \sqrt{n}$.

In the restricted subsystem of *tree-like resolution*, optimal $n^{\Omega(k)}$ length lower bounds were established in [6] for k -clique formulas on complete $(k-1)$ -partite as well as on average for Erdős-Rényi random graphs of appropriate edge density. There is no hope to get hard instances for general resolution from complete $(k-1)$ -partite graphs, however—in the same paper it was shown that all instances from the more general class of $(k-1)$ -colourable graphs are easy for resolution. A closer study of these resolution proofs reveals that they are *regular*, meaning that if the proof is viewed as a directed acyclic graph (DAG), then no variable is eliminated more than once on any source-to-sink path.

More generally, regular resolution is an interesting and non-trivial model to analyse for the k -clique problem since it captures the reasoning used in many state-of-the-art algorithms used in practice (for a survey, see, e.g., [24, 27]). Nonetheless, it has remained consistent with state-of-the-art knowledge that for $k \leq n^{5/6}$ regular

resolution might be able to certify k -clique-freeness in polynomial length independent of the value of k .

Our contribution. We prove optimal $n^{\Omega(k)}$ average-case lower bounds for regular resolution proofs of unsatisfiability for k -clique formulas on Erdős-Rényi random graphs.

THEOREM 1.1 (INFORMAL). *For any integer $k \ll \sqrt[4]{n}$, given an n -vertex graph G sampled at random from the Erdős-Rényi model with the appropriate edge density, regular resolution asymptotically almost surely requires length $n^{\Omega(k)}$ to certify that G does not contain a k -clique.*

In order to make this formal, we need to define how the problem is encoded: depending on the formula considered, the exact statement of what we can prove differs. In this conference paper we consider the simpler encoding for which we can prove an $n^{\Omega(k)}$ lower bound for $k \ll \sqrt{n}$. For a stronger encoding, which in particular captures this simpler one, we prove the above result in the full-length version of this paper.

At a high level, the proof is based on a bottleneck counting argument in the style of [12] with a slight twist that was introduced in [29]. In its classical form, such a proof takes four steps. First, one defines a distribution of random source-to-sink paths on the DAG representation of the proof. Second, a subset of the vertices of the DAG is identified—the set of *bottleneck nodes*—such that any random path must necessarily pass through at least one such node. Third, for any fixed bottleneck node, one shows that it is very unlikely that a random path passes through this particular node. Given this, a final union bound argument yields the conclusion that the DAG must have many bottleneck nodes, and so the resolution proof must be long.

The twist in our argument is that, instead of single bottleneck nodes, we need to define *bottleneck pairs* of nodes. We then argue that any random path passes through at least one such pair but that few random paths pass through any fixed pair; the latter part is based on Markov chain-type reasoning similar to [29, Theorems 3.2, 3.5]. Furthermore, it crucially relies on that the graph satisfies a certain combinatorial property, which captures the idea that the common neighbourhood of a small set of vertices is well distributed across the graph. Identifying this combinatorial property is a key contribution of our work. In a separate argument (that, surprisingly, turned out to be much more elaborate than most arguments of this kind) we then establish that Erdős-Rényi random graphs of the appropriate edge density satisfy this property asymptotically almost surely. Combining these two facts yields our average-case lower bound.

Another contribution of this paper is a relatively simple observation that not only is regular resolution powerful enough to distinguish graphs that contain k -cliques from $(k-1)$ -colourable graphs [6], but it can also distinguish them from graphs that have a homomorphism to any fixed graph H with no k -cliques.

Paper outline. The rest of this paper is organized as follows. Section 2 presents some preliminaries. We show that some nontrivial k -clique instances are easy for regular resolution in Section 3. Section 4 contains the formal statement of the lower bounds we prove for Erdős-Rényi random graphs. In Section 5 we define a combinatorial property of graphs and show that clique formulas on such

graphs are hard for regular resolution, and the proof that Erdős-Rényi random graphs satisfy this property asymptotically almost surely is in Section 6. We conclude in Section 7 with a discussion of open problems.

2 PRELIMINARIES

We write $G = (V, E)$ to denote a graph with vertices V and edges E , where G is always undirected, without loops and multiple edges. Given a vertex $v \in V$, we write $N(v) = \{u \mid \text{there exists } v \in V \text{ such that } \{u, v\} \in E\}$ to denote the set of *neighbours of* v . For a set of vertices $R \subseteq V$ we write $\widehat{N}(R) = \bigcap_{v \in R} N(v)$ to denote the set of *common neighbours of* R . For two sets of vertices $R \subseteq V$ and $W \subseteq V$ we write $\widehat{N}_W(R) = \widehat{N}(R) \cap W$ to denote the set of *common neighbours of* R *inside* W . For a set $U \subseteq V$ we denote by $G[U]$ the subgraph of G induced by the set U . For $n \in \mathbb{N}^+$ we write $[n] = \{1, \dots, n\}$. We say that $V_1 \dot{\cup} V_2 \dot{\cup} \dots \dot{\cup} V_k = V$ is a *balanced* k -partition of V if for all $i, j \in [k]$ it holds that $|V_i| \leq |V_j| + 1$. All logarithms are natural (base e) if not specified otherwise.

Probability and Erdős-Rényi random graphs. We denote random variables in boldface and write $X \sim \mathcal{D}$ to denote that X is sampled from the distribution \mathcal{D} . A p -biased coin, or a *Bernoulli variable*, is the outcome of a coin flip that yields 1 with probability p and 0 with probability $1 - p$. We use the special case of Markov's inequality saying that if X is non-negative, then $\Pr[X \geq 1] \leq \mathbb{E}[X]$. We also need the following special case of the multiplicative Chernoff bound: if X is a binomial random variable (i.e., the sum of i.i.d. Bernoulli variables) with expectation $\mu = \mathbb{E}[X]$, then $\Pr[X \leq \mu/2] \leq e^{-\mu/8}$.

We consider the Erdős-Rényi distribution $\mathcal{G}(n, p)$ of random graphs on a fixed set V of n vertices. A random graph sampled from $\mathcal{G}(n, p)$ is produced by placing each potential edge $\{u, v\}$ independently with probability p , $0 \leq p \leq 1$ (the edge probability p may be a function of n). A property of graphs is said to hold *asymptotically almost surely* on $\mathcal{G}(n, p(n))$ if it holds with probability that approaches 1 as n approaches infinity.

For a positive integer k , let X_k be the random variable that counts the number of k -cliques in a random graph from $\mathcal{G}(n, p)$. It follows from Markov's inequality that asymptotically almost surely there are no k -cliques in $\mathcal{G}(n, p)$ whenever p and k are such that $\mathbb{E}[X_k] = p^{\binom{k}{2}} \binom{n}{k}$ approaches 0 as n approaches infinity. This is the case, for example, if $p = n^{-2\eta/(k-1)}$ for $k \geq 2$ and $\eta > 1$.

CNF formulas and resolution. A *literal* over a Boolean variable x is either the variable x itself (a *positive literal*) or its negation $\neg x$ (a *negative literal*). A *clause* $C = a_1 \vee \dots \vee a_w$ is a disjunction of literals; we say that the *width* of C is w . The empty clause will be denoted by \perp . A *CNF formula* $F = C_1 \wedge \dots \wedge C_m$ is a conjunction of clauses. We think of clauses as sets of literals and of CNF formulas as sets of clauses, so that order is irrelevant and there are no repetitions. For a formula F we denote by $\text{Vars}(F)$ the set of variables of F .

A *resolution derivation* from a CNF formula F is as an ordered sequence of clauses $\pi = (D_1, \dots, D_L)$ such that for each $i \in [L]$ either D_i is a clause in F or there exist $j < i$ and $k < i$ such that D_i is derived from D_j and D_k by the *resolution rule*

$$\frac{B \vee x \quad C \vee \neg x}{B \vee C}, \quad (1)$$

$D_i = B \vee C$, $D_j = B \vee x$, $D_k = C \vee \neg x$. We refer to $B \vee C$ as the *resolvent* of $B \vee x$ and $C \vee \neg x$ over x , and to x as the *resolved variable*. The *length* (or *size*) of a resolution derivation $\pi = (D_1, \dots, D_L)$ is L and it is denoted by $|\pi|$. A *resolution refutation* of F , or *resolution proof* for (the unsatisfiability of) F , is a resolution derivation from F that ends in the empty clause \perp .

A resolution derivation $\pi = (D_1, \dots, D_L)$ can also be viewed as a labelled DAG with set of nodes $\{1, \dots, L\}$ and edges $(j, i), (k, i)$ for each application of the resolution rule deriving D_i from D_j and D_k . Each node i in this DAG is labelled by its associated clause D_i , and each non-source node is also labelled by the resolved variable in its associated derivation step in the refutation. A resolution refutation is called *regular* if along any source-to-sink path in its associated DAG every variable is resolved at most once.

For a partial assignment ρ we say that a clause C *restricted by* ρ , denoted $C \upharpoonright_\rho$, is the trivial 1-clause if any of the literals in C is satisfied by ρ or otherwise is C with all falsified literals removed. We extend this definition to CNFs in the obvious way: $(C_1 \wedge \dots \wedge C_m) \upharpoonright_\rho = C_1 \upharpoonright_\rho \wedge \dots \wedge C_m \upharpoonright_\rho$. Applying a restriction preserves (regular) resolution derivations. To see this, observe that in every application of the resolution rule the restricted consequence is either killed (becomes identically 1) or obtained, as before, by resolving the two restricted premises or it is a copy of one of them. Thus, we have:

FACT 2.1. *Let π be a (regular) resolution refutation of a CNF formula F . For any partial assignment ρ to the variables of F there is an efficiently constructible (regular) resolution refutation $\pi \upharpoonright_\rho$ of the CNF formula $F \upharpoonright_\rho$, so that the length of $\pi \upharpoonright_\rho$ is at most the length of π .*

Branching programs. A branching program on variables x_1, \dots, x_n is a DAG that has one source node and where every non-sink node is labelled by one of the variables x_1, \dots, x_n and has exactly two outgoing edges labelled 0 and 1. The size of a branching program is the total number of nodes in the graph. In a *read-once branching program* it holds in addition that along every path every variable appears as a node label at most once.

For each node a in a branching program, let $X(a)$ denote the variable that labels a , and let a^0 and a^1 be the nodes that are reached from a through the edges labelled 0 and 1, respectively. A truth-value assignment $\sigma : \{x_1, \dots, x_n\} \rightarrow \{0, 1\}$ determines a path in a branching program in the following way. The path starts at the source node. At an internal node a , the path is extended along the edge labelled $\sigma(X(a))$ so that the next node in the path is $a^{\sigma(X(a))}$. The path ends when it reaches a sink. We write $\text{path}(\sigma)$ for the path determined by σ . When extending the path from a node a to the node $a^{\sigma(X(a))}$, we say that the *answer to the query* $X(a)$ at a is $\sigma(X(a))$ and that the path *sets* the variable $X(a)$ to the value $\sigma(X(a))$. For each node a of the branching program, let $\beta(a)$ be the maximal partial assignment that is contained in any assignment σ such that $\text{path}(\sigma)$ passes through a . Equivalently, this is the set of all those $\sigma(x_i) = \gamma$ for which the query x_i is made, and answered by γ , along every consistent path from the source to a . If the program is read-once, the consistency condition becomes redundant.

The *falsified clause search problem* for an unsatisfiable CNF formula F is the task of finding a clause $C \in F$ that is falsified by a given truth value assignment σ . A branching program P on the variables $\text{Vars}(F)$ *solves* the falsified clause search problem for F

if each sink is labelled by a clause of F such that for every assignment σ , the clause that labels the sink reached by $\text{path}(\sigma)$ is falsified by σ . The minimal size of any regular resolution refutation of an unsatisfiable CNF formula F is exactly the same as the minimal size of any read-once branching program solving the falsified clause search problem for F . This can be seen by taking the refutation DAG and reversing the edges to get a branching program or vice versa. For a formal proof see, e.g., [19, Theorem 4.3].

The k -clique formula. In order to analyse the complexity of resolution proofs that establish that a given graph does not contain a k -clique we must formulate the problem as a propositional formula in conjunctive normal form (CNF). We consider two distinct encodings for the clique problem originally defined in [3].

The first propositional encoding we present, $\text{Clique}(G, k)$, is based on mapping of vertices to clique members. This formula is defined over variables $x_{v,i}$ ($v \in V, i \in [k]$) and consists of the following set of clauses:

$$\neg x_{u,i} \vee \neg x_{v,j} \quad i, j \in [k], i \neq j, u, v \in V, \{u, v\} \notin E, \quad (2a)$$

$$\bigvee_{v \in V} x_{v,i} \quad i \in [k], \quad (2b)$$

$$\neg x_{u,i} \vee \neg x_{v,i} \quad i \in [k], u, v \in V, u \neq v, \quad (2c)$$

We refer to (2a) as *edge axioms*, (2b) as *clique axioms* and (2c) as *functionality axioms*. Note that $\text{Clique}(G, k)$ is satisfiable if and only if G contains a k -clique, and that this is true even if clauses (2c) are omitted—we write $\text{Clique}^*(G, k)$ to denote this formula with only clauses (2a) and (2b).

The second version of clique formulas that we consider is the block encoding $\text{Clique}_{\text{block}}(G, k)$. This formula differs from the previous ones in that it requires a k -clique that has a certain “block-respecting” structure. Let $V_1 \dot{\cup} V_2 \dot{\cup} \dots \dot{\cup} V_k = V$ be a balanced k -partition of V . This formula, defined over variables x_v , encodes the fact that the graph contains a *transversal k -clique*, that is, a k -clique in which each clique member belongs to a different block. Formally, for any positive k and n , the formula $\text{Clique}_{\text{block}}(G, k)$ consists of the following set of clauses:

$$\neg x_u \vee \neg x_v \quad u, v \in V, u \neq v, \{u, v\} \notin E, \quad (3a)$$

$$\bigvee_{v \in V_i} x_v \quad i \in [k], \quad (3b)$$

$$\neg x_u \vee \neg x_v \quad i \in [k], u, v \in V_i, u \neq v. \quad (3c)$$

Note that a graph can contain a k -clique but contain no transversal k -clique for a given partition. Intuitively it is clear that proving that a graph does not contain a transversal k -clique should be easier than proving it does not contain any k -clique, since any proof of the latter fact must in particular establish the former. We make this intuition formal below.

LEMMA 2.2 ([3]). *For any graph G and any $k \in \mathbb{N}^+$, the size of a minimum regular resolution refutation of $\text{Clique}(G, k)$ is bounded from below by the size of a minimum regular resolution refutation of $\text{Clique}_{\text{block}}(G, k)$.*

This lemma was proven in [3] for tree-like and for general resolution via a restriction argument, and it is straightforward to see that the same proof holds for regular resolution.

3 EASY GRAPHS FOR REGULAR RESOLUTION

Before proving our main $n^{\Omega(k)}$ lower bound, in this section we exhibit classes of graphs whose clique formulas have regular resolution refutations of fixed-parameter tractable length, i.e., length $f(k) \cdot n^{O(1)}$ for some function f . This illustrates the strength of regular resolution for the k -clique problem. We note that the upper bounds claimed in this section hold not only for $\text{Clique}(G, k)$ but even for the subformula $\text{Clique}^*(G, k)$ that omits the functionality axioms (2c).

The first example is the class of $(k-1)$ -colourable graphs. Such graphs are hard for tree-like resolution [6], and the known algorithms that distinguish them from graphs that contain k -cliques are highly non-trivial [17, 22]. The second example is the class of graphs that have a homomorphism into a fixed k -clique free graph.

Recall that a homomorphism from a graph $G = (V, E)$ into a graph $G' = (V', E')$ is a mapping $h : V \rightarrow V'$ that maps edges $\{u, v\} \in E$ into edges $\{h(u), h(v)\} \in E'$. A graph is $(k-1)$ -colourable if and only if it has a homomorphism into the $(k-1)$ -clique, which is of course k -clique free. Therefore our second example is a generalization of the first one (but the function $f(k)$ becomes larger).

Both upper bounds follows from a generic procedure, based on Algorithm 1, that builds read-once branching programs for the falsified clause search problem for $\text{Clique}^*(G, k)$.

Given a k -clique free graph G define

$$I(G) = \{G[\widehat{N}(R)] : R \text{ is a clique in } G\}. \quad (4)$$

PROPOSITION 3.1. *There is an efficiently constructible read-once branching program for the falsified clause search problem on formula $\text{Clique}^*(G, k)$ of size at most $|I(G)| \cdot k^2 \cdot |V(G)|^2$.*

PROOF. We build the branching program recursively, following the strategy laid out by Algorithm 1. For the base case $k = 1$, G must be the graph with no vertices. The branching program is a single sink node that outputs the clique axiom of index 1, i.e., the empty clause.

For $k > 1$, fix $n = |V(G)|$ and an ordering v_1, \dots, v_n of the vertices in $V(G)$. We first build a decision tree T by querying the variables $x_{v_1, k}, x_{v_2, k}, \dots$ in order, until we get an answer 1, or until all variables with second index k have been queried. If $x_{v_j, k} = 0$ for all $j \in [n]$ then the k th clique axiom (2b) is falsified by the assignment (see line 14). Otherwise, let v be the first vertex in the order where $x_{v, k} = 1$. The decision tree now queries $x_{w, i}$ for all $w \notin N(v)$ and all $i < k$ to check whether an edge axiom involving v is falsified (lines 4–6). If any of these variables is set to 1 the branching stops and the leaf node is labelled with the corresponding edge axiom $\neg x_{v, k} \vee \neg x_{w, i}$.

The decision tree T built so far has at most kn^2 nodes, and we can identify n “open” leaf nodes $a_{v_1}, a_{v_2}, \dots, a_{v_n}$, where a_{v_i} is the leaf node reached by the path that sets $x_{v_i, k} = 1$ and that does yet determine the answer to the search problem. Let us focus on a specific node a_v for some $v \in V(G)$. The partial assignment $\text{path}(a_v)$ sets v to be the k th member of the clique and no vertex in $V(G) \setminus N(v)$ to be in the clique. Let G_v be the subgraph induced

Algorithm 1 Read-once branching program for the falsified clause search problem on $\text{Clique}^*(G, k)$.

Input $k \in \mathbb{N}^+$, a k -clique free graph G , an assignment $\alpha: \{x_{v,i} \text{ for } v \in V(G), i \in [k]\} \rightarrow \{0, 1\}$
Output A clause of $\text{Clique}^*(G, k)$ falsified by α

- 1: **procedure** SEARCH(G, k, α)
- 2: **for** $v \in V(G)$ **do**
- 3: **if** $\alpha(x_{v,k}) = 1$ **then**
- 4: **for** $w \notin N(v)$ and $i < k$ **do**
- 5: **if** $\alpha(x_{w,i}) = 1$ **then**
- 6: **return** edge axiom $\neg x_{v,k} \vee \neg x_{w,i}$ (2a).
- 7: **end if**
- 8: **end for**
- 9: $G' \leftarrow G[N(v)]$
- 10: $\alpha' \leftarrow \alpha$ restricted to variables $x_{w,j}$ for $w \in V(G')$
and $1 \leq j \leq k-1$
- 11: **return** SEARCH($G', k-1, \alpha'$)
- 12: **end if**
- 13: **end for**
- 14: **return** the k th clique axiom (2b).
- 15: **end procedure**

on G by $N(v)$, let S_v be the set of variables $x_{w,i}$ for $w \in N(v)$ and $i < k$, and let ρ_v be the partial assignment setting $x_{w,i} = 0$ for $w \notin N(v)$ and $i < k$. Clearly $\rho_v \subseteq \text{path}(a_v)$.

By the inductive hypothesis there exists a branching program B_v that solves the search problem on $\text{Clique}^*(G_v, k-1)$ querying only variables in S_v . This corresponds to the recursive call for the subgraph G_v and $k-1$ (lines 9–11). If we attach each B_v to a_v we get a complete branching program for $\text{Clique}^*(G, k)$. This is read-once because B_v only queries variables in S_v and these variables are not in $\text{path}(a_v)$.

To prove that the composed program is correct we consider an assignment σ to the variables in S_v and show that the clause output by B_v on σ is also a valid output for the search problem on $\text{Clique}^*(G, k)$, i.e., it is falsified by the assignment $\text{path}(a_v) \cup \sigma$. Actually we show the stronger claim that it is falsified by $\rho_v \cup \sigma$, which is a subset of $\text{path}(a_v) \cup \sigma$. To this end, note that if the output of B_v on σ is an edge axiom of $\text{Clique}^*(G_v, k-1)$, this must be some $\neg x_{u,i} \vee \neg x_{w,j}$ for $i, j < k$, which is also an edge axiom of $\text{Clique}^*(G, k)$ and is falsified by $\sigma \subseteq \rho_v \cup \sigma$. Now if the output of B_v on σ is the i th clique axiom of $\text{Clique}^*(G_v, k-1)$, then σ falsifies $\bigvee_{w \in N(v)} x_{w,i}$, and therefore $\rho_v \cup \sigma$ falsifies the i th clique axiom in formula $\text{Clique}^*(G, k)$.

The construction so far is correct but produces a very large branching program (in particular, a tree-like one). In order to create a smaller branching program, we observe that if $u, v \in V(G)$ are such that $N(u) = N(v)$ then $G_u = G_v, B_u = B_v$ and $\rho_u = \rho_v$. In this case, we can identify nodes a_u and a_v , resulting in a node we denote a^* , and identify the branching programs B_u and B_v . The correctness of this new program is due to the fact that even after the identification of vertices $\rho_u \subseteq \text{path}(a^*)$ and $\rho_v \subseteq \text{path}(a^*)$. This process leads to having only one subprogram for each distinct induced subgraph at each level of the recursion.

In order to bound the size of this program, we decompose it into k levels. The source is at level zero and corresponds to the graph G . At level i there are nodes corresponding to all subgraphs induced by the common neighbourhood of cliques of size i . Each node in the i th level connects to the nodes of the $(i+1)$ th level by a branching program of size at most kn^2 . Notice that an induced subgraph in $I(G)$ cannot occur twice in the same layers, so the total size of the final branching program is at most $|I(G)| \cdot k^2 n^2$ nodes. \square

We now proceed to prove the upper bounds mentioned previously. A graph G that has a homomorphism into a small k -clique free graph H may still have a large set $I(G)$, making Proposition 3.1 inefficient. The first key observation is that if G has a homomorphism into a graph H then it is a subgraph of a blown up version of H , namely, of a graph obtained by transforming each vertex of H into a “cloud” of vertices where a cloud does not contain any edge, two clouds corresponding to two adjacent vertices in H have all possible edges between them, and two clouds corresponding to two non-adjacent vertices in H have no edges between them. A second crucial point is that if G' is a blown up version of H then it turns out that $|I(G')| = |I(H)|$, making Proposition 3.1 effective for G' . The upper bound then follows from observing that the task of proving that G is k -clique free should not be harder than the same task for a supergraph of G . Indeed Fact 3.2 formalises this intuition. It is interesting to observe that the constructions in Proposition 3.1 and in Fact 3.2 are efficient. The non-constructive part is guessing the homomorphism to H .

FACT 3.2. *Let $G = (V, E)$ and $G' = (V', E')$ be graphs with no k -clique such that $V \subseteq V'$ and $E \subseteq E' \cap \binom{V}{2}$. If $\text{Clique}^*(G', k)$ has a (regular) refutation of length L , then $\text{Clique}^*(G, k)$ also has a (regular) refutation of length L .*

PROOF. Consider the partial assignment ρ that sets $x_{v,i} = 0$ for every $v \notin V$ and $i \in [k]$. The restricted formula $\text{Clique}^*(G', k)|_\rho$ is isomorphic to $\text{Clique}^*(\tilde{G}, k)$, where $V(\tilde{G}) = V$ and $E(\tilde{G}) = E' \cap \binom{V}{2}$, and thus, by Fact 2.1, has a (regular) refutation π of length at most L . Removing edges from a graph only introduces additional edge axioms (2a) in the corresponding formula, therefore $\text{Clique}^*(\tilde{G}, k) \subseteq \text{Clique}^*(G, k)$ and π is a valid refutation of $\text{Clique}^*(G, k)$ as well. \square

It was shown in [6] that the k -clique formula of a complete $(k-1)$ -partite graph on n vertices has a regular resolution refutation of length $2^k n^{O(1)}$, although the regularity is not stressed in that paper. Since it is instructive to see how this refutation is constructed in this framework, we give a self-contained proof.

PROPOSITION 3.3 ([6, PROPOSITION 5.3]). *If G is a $(k-1)$ -colourable graph on n vertices, then $\text{Clique}^*(G, k)$ has a regular resolution refutation of length at most $2^k k^2 n^2$.*

PROOF. Let $V = V(G)$ and let $V_1 \cup V_2 \cup \dots \cup V_{k-1}$ be a partition of V into colour classes. Define the graph $G' = (V, E')$ where the edge set E' has an edge between any pair of vertices belonging to two different colour classes. Clearly G is a subgraph of G' . Observe that any clique R in G' has at most one vertex in each colour class, and that the common neighbours of R are all the vertices in the colour classes not touched by R .

Therefore, there is a one-to-one correspondence between the members of $I(G')$ and the subsets of $[k-1]$. By Proposition 3.1 there is a read-once branching program for the falsified clause search problem on formula $\text{Clique}^*(G', k)$ of size at most $2^k k^2 n^2$. This read-once branching program corresponds to a regular resolution refutation of $\text{Clique}^*(G', k)$ of the same size. By Fact 3.2 there must be a regular resolution refutation of size at most $2^k k^2 n^2$ for $\text{Clique}^*(G, k)$ as well. \square

Next we generalize Proposition 3.3 to graphs G that have a homomorphism to a k -clique free graph H .

PROPOSITION 3.4. *If G is a graph on n vertices that has a homomorphism into a k -clique free graph H on m vertices, then $\text{Clique}^*(G, k)$ has a regular resolution refutation of length at most $m^k k^2 n^2$.*

PROOF. Fix a homomorphism $h: V(G) \rightarrow V(H)$ and an ordering u_1, \dots, u_m of the vertices of H . Let $V_1 \dot{\cup} V_2 \dot{\cup} \dots \dot{\cup} V_m$ be the partition of $V(G)$ such that V_i is the set of vertices of G mapped to u_i by h . We define the graph $G' = (V, E')$ where

$$E' = \bigcup_{\{u_i, u_j\} \in E(H)} V_i \times V_j, \quad (5)$$

that is, G' is a blown up version of H that contains G as a subgraph. To prove our result we note that, by Proposition 3.1, there is a read-once branching program for the falsified clause search problem on $\text{Clique}^*(G', k)$ —and hence also a regular resolution refutations of the same formula—of size at most $|I(G')| \cdot k^2 n^2$. This implies that, by Fact 3.2, there is a regular resolution refutation of $\text{Clique}^*(G, k)$ of at most the same size.

To conclude the proof it remains only to show that $|I(G')| \leq m^k$. By construction, h maps injectively a clique $R \subseteq V(G')$ into a clique $R_H \subseteq V(H)$ of the same size. Moreover, note that if $U = \widehat{N}(R_H)$, then $\widehat{N}(R) = \cup_{u_i \in U} V_i$. Therefore, for any clique $R' \subseteq V(G')$ that is mapped by h to R_H it holds that $\widehat{N}(R) = \widehat{N}(R')$, i.e., $\widehat{N}(R')$ is completely characterized by the clique in H it is mapped to. Thus $I(G)$ has at most one element for each clique in H and we have that $|I(G')| = |I(H)|$. Finally, note that $|I(H)| \leq m^k$ since, being k -clique free, H cannot have more than m^k cliques. \square

4 RANDOM GRAPHS ARE HARD

The main result of this paper is an average case lower bound of $n^{\Omega(k)}$ for regular resolution for the k -clique problem. As we saw in Section 2, the k -clique problem can be encoded in different ways and depending on the preferred formula the range of k for which we can obtain a lower bound differs. In this section we present a summary of our results for the different encodings.

THEOREM 4.1. *For any real constant $\epsilon > 0$, any sufficiently large integer n , any positive integer $k \leq n^{1/4-\epsilon}$, and any real $\xi > 1$, if $G \sim \mathcal{G}(n, n^{-2\xi/(k-1)})$ is an Erdős-Rényi random graph, then, with probability at least $1 - \exp(-\sqrt{n})$, any regular resolution refutation of $\text{Clique}_{\text{block}}(G, k)$ has length at least $n^{\Omega(k/\xi^2)}$.*

The parameter ξ determines the density of the graph: the larger ξ the sparser the graph and the problem of determining whether G contains a k -clique becomes easier. For constant ξ , where the edge probability is somewhat close to the threshold for containing a

k -clique, the theorem yields a $n^{\Omega(k)}$ lower bound which is tight up to the multiplicative constant in the exponent. The lower bound decreases smoothly with the edge density and is non-trivial for $\xi = o(\sqrt{k})$.

A problem which is closely related to the problem we consider is that of distinguishing a random graph sampled from $\mathcal{G}(n, p)$ from a random graph from the same distribution with a planted k -clique. The most studied setting is when $p = 1/2$. In this scenario the problem can be solved in polynomial time with high probability for $k \approx \sqrt{n}$ [1, 20]. It is still an open problem whether there exists a polynomial time algorithm solving this problem for $\log n \ll k \ll \sqrt{n}$. For $G \sim \mathcal{G}(n, 1/2)$, Theorem 4.1 implies that to refute $\text{Clique}_{\text{block}}(G, k)$ asymptotically almost surely regular resolution requires $n^{\Omega(\log n)}$ size for $k = O(\log n)$ and super-polynomial size for $k = o(\log^2 n)$.

An interesting question is whether Theorem 4.1 holds for larger values of k . We show that for the formula $\text{Clique}(G, k)$ (recall that by Lemma 2.2 this encoding is easier for the purpose of lower bounds) we can prove the lower bound for $k \leq n^{1/2-\epsilon}$ as long as the edge density of the graph is close to the threshold for containing a k -clique.

THEOREM 4.2. *For any real constant $\epsilon > 0$, any sufficiently large integer n , any positive integer k , and any real $\xi > 1$ such that $k\sqrt{\xi} \leq n^{1/2-\epsilon}$, if $G \sim \mathcal{G}(n, n^{-2\xi/(k-1)})$ is an Erdős-Rényi random graph, then, with probability at least $1 - \exp(-\sqrt{n})$, any regular resolution refutation of $\text{Clique}(G, k)$ has length at least $n^{\Omega(k/\xi^2)}$.*

In this extended abstract we prove Theorem 4.2 and we refer to the upcoming full-length version of this paper for the proof of Theorem 4.1. We note, however, that both proofs are very similar and having seen one it is an easy exercise to obtain the other. The proof of Theorem 4.2 is deferred to Section 6 and is based on a general lower bound technique we develop in Section 5.

5 CLIQUE-DENSENESS IMPLIES HARDNESS

In this section we define a combinatorial property of graphs, which we call *clique-denseness*, and prove that if a k -clique-free graph G is clique-dense with the appropriate parameters, then this implies a lower bound $n^{\Omega(k)}$ on the length of any regular resolution refutation of the k -clique formula on G .

In order to argue that regular resolution has a hard time certifying the k -clique-freeness of a graph G , one property that seems useful to have is that for every small enough clique in the graph there are many ways of extending it to a larger clique. In other words, if $R \subseteq V$ forms a clique and R is small, we would like the common neighbourhood $\widehat{N}_V(R)$ to be large. This motivates the following definitions.

Definition 5.1 (Neighbour-dense set). Given a graph $G = (V, E)$ and $q, r \in \mathbb{R}^+$, a set $W \subseteq V$ is *q-neighbour-dense* for $R \subseteq V$ if $|\widehat{N}_W(R)| \geq q$. We say that W is *(r, q)-neighbour-dense* if it is *q-neighbour-dense* for every $R \subseteq V$ of size $|R| \leq r$.

If W is an (r, q) -neighbour-dense set, then we know that any clique of size r can be extended to a clique of size $r+1$ in at least q different ways by adding some vertex of W . Note, however, that

the definition of (r, q) -neighbour-dense is more general than this since R is not required to be a clique.

We next define a more robust notion of neighbour-denseness. For some settings of r and q of interest to us it is too much to hope for a set W which is q -neighbour-dense for every $R \subseteq V$ of size at most r . In this case we would still like to be able to find a “mostly neighbour-dense” set W in the sense that we can “localize” bad sets $R \subseteq V$ of size $|R| \leq r$, i.e., those for which W fails to be q -neighbour-dense.

Definition 5.2 (Mostly neighbour-dense set). Given $G = (V, E)$ and $r', r, q', s \in \mathbb{R}^+$ with $r' \geq r$, a set $W \subseteq V$ is (r', r, q', s) -mostly neighbour-dense if there exists a set $S \subseteq V$ of size $|S| \leq s$ such that for every $R \subseteq V$ with $|R| \leq r'$ for which W is not q' -neighbour-dense, it holds that $|R \cap S| \geq r$.

In what follows, it might be helpful for the reader to think of r' and r as linear in k and q and s as polynomial in n , where we also have that $s \ll q$.

Now we are ready to define a property of graphs that makes it hard for regular resolution to certify that graphs with this property, but without k -cliques, are indeed k -clique-free.

Definition 5.3 (Clique-dense graph). Given $k \in \mathbb{N}^+$ and $t, s, \varepsilon \in \mathbb{R}^+$, $1 \leq t \leq k$, we say that a graph $G = (V, E)$ is (k, t, s, ε) -clique-dense if there exist $r, q \in \mathbb{R}^+$, $r \geq 4k/t^2$, such that

- (1) V is (tr, tq) -neighbour-dense, and
- (2) every (r, q) -neighbour-dense set $W \subseteq V$ is (tr, r, q', s) -mostly neighbour-dense for $q' = 3\varepsilon ks^{1+\varepsilon} \log s$.

THEOREM 5.4. *Given $k \in \mathbb{N}^+$ and $t, s, \varepsilon \in \mathbb{R}^+$ if the graph G is (k, t, s, ε) -clique-dense, then every regular resolution refutation of the CNF formula $\text{Clique}(G, k)$ has length at least $\frac{1}{\sqrt{2}} s^{\varepsilon k/t^2}$.*

The value of q' in Definition 5.3 is tailored so that Theorem 4.2 holds for $k \ll n^{1/2}$ on graphs with edge density close to the threshold for having a k -clique. Setting $q' = \varepsilon r s^{1+\varepsilon} \log s$ and making the necessary modifications in the proof would yield Theorem 4.2 for a larger range of edge densities but only for $k \ll n^{2/5}$.

We will spend the rest of this section establishing Theorem 5.4. Fix $r, q \in \mathbb{R}^+$ witnessing that G is (k, t, s, ε) -clique-dense as per Definition 5.3. We first note that we can assume that $tr \leq k$ since otherwise, by property 1 of Definition 5.3, G contains a k -clique and the theorem follows immediately.

By the discussion in Section 2 it is sufficient to consider read-once branching programs, since they are equivalent to regular resolution refutations, and so in what follows this is the language in which we will phrase our lower bound. Thus, for the rest of this section let P be an arbitrary, fixed read-once branching program that solves the falsified clause search problem for $\text{Clique}(G, k)$. We will use the convention of referring to “vertices” of the graph G and “nodes” of the branching program P to distinguish between the two.

Recall that for a node a of P , $\beta(a)$ denotes the maximal partial assignment that is contained in any assignment σ such that the path $\text{path}(\sigma)$ passes through a . For any partial assignment β we write β^1 to denote the partial assignment that contains exactly the variables that are set to 1 in β . Clearly, if β falsifies an edge axiom or a functionality axiom, then so does β^1 . Furthermore, for any $\beta' \subseteq \beta^1$, if β' falsifies an edge axiom or a functionality axiom, so does β^1 . We will

use this monotonicity property of partial assignments throughout the proof.

For each node a of P and each index $i \in [k]$ we define two sets of vertices

$$V_i^0(a) = \{v \in V \mid \beta(a) \text{ sets } x_{v,i} \text{ to } 0\} \quad (6a)$$

$$V_i^1(a) = \{v \in V \mid \beta(a) \text{ sets } x_{v,i} \text{ to } 1\} \quad (6b)$$

of G . Observe that for $\beta = \beta(a)$ the set of vertices referenced by variables in β^1 is $\bigcup_i V_i^1(a)$.

Intuitively, one can think of $V_i^0(a)$ and $V_i^1(a)$ as the sets of vertices v for which the variable $x_{v,i}$ is assigned 0 and 1, respectively, that are guaranteed to be “remembered” at the node a (in the language of resolution, they correspond to negative and positive occurrences of variables in the clause D_a associated with the node a). Other assignments to variables $x_{u,i}$ for $u \notin V_i^0(a) \cup V_i^1(a)$ encountered along some path to a have been “forgotten” and may not be queried any more on any path starting at a . Formally, we say that a variable $x_{v,i}$ is *forgotten at a* if there is a path from the source of P to a passing through a node b where $x_{v,i}$ is queried, but v is not in $V_i^0(a)$ nor in $V_i^1(a)$. Furthermore, we say index i is *forgotten at a* if for some vertex v the variable $x_{v,i}$ is forgotten at a . Of utter importance is the fact that these notions are persistent: if a variable or an index is forgotten at a node a , then it will also be the case for any node reachable from a by a path. We say that a path in P ends in the i th clique axiom if the clause that labels its last node is the clique axiom (2b) of $\text{Clique}(G, k)$ with index i . The above observation implies that the index i cannot be forgotten at any node along such a path.

We establish our lower bound via a bottleneck counting argument for paths in P . To this end, let us define a distribution \mathcal{D} over paths in P by the following random process. The path starts at the source and ends whenever it reaches a sink of P . At an internal node a with successor nodes a^0 and a^1 , reached by edges labelled 0 and 1 respectively, the process proceeds as follows.

- (1) If $X(a) = x_{u,i}$ and i is forgotten at a then the path proceeds via the edge labelled 0 to a^0 .
- (2) If $X(a) = x_{u,i}$ and $\beta(a) \cup \{x_{u,i} = 1\}$ falsifies an edge axiom (2a) or a functionality axiom (2c), then the path proceeds to a^0 .
- (3) Otherwise, an independent $(rs^{-(1+\varepsilon)}/2ek)$ -biased coin is tossed with outcome $\gamma \in \{0, 1\}$ and the random path proceeds to a^γ .

We say that in cases (1) and (2) the answer to the query $X(a)$ is *forced*. Note that any path α in the support of \mathcal{D} must end in a clique axiom since α does not falsify any edge or functionality axiom by construction. Moreover, a property that will be absolutely crucial is that only answers 0 can be forced—answers 1 are always the result of a coin flip.

CLAIM 5.5. *Every path in the support of \mathcal{D} sets at most k variables to 1.*

PROOF. Let α be a path in the support of \mathcal{D} . We argue that for each $i \in [k]$ at most one variable with second index i is set to 1 on α . Let a and b be two nodes that appear in this order in α . If for some $i \in [k]$, and for some $u, v \in V$, $x_{u,i}$ is set to 1 by α at node a and $x_{v,i}$ is queried at b , then $v \neq u$ by regularity and, by definition of \mathcal{D} , the answer to query $x_{v,i}$ will be forced to 0, either to avoid

violating a functionality or an edge axiom, or because i is forgotten at b . \square

Let us call a pair (a, b) of nodes of P *useful* if there exists an index i such that $V_i^1(b) = \emptyset$, i is not forgotten at b , and the set $V_i^0(b) \setminus V_i^0(a)$ is (r, q) -neighbour-dense. For each useful pair (a, b) , let $i(a, b)$ be an arbitrary but fixed index witnessing that (a, b) is useful. A path is said to *usefully* traverse a useful pair (a, b) if it goes through a and b in that order and sets at most $\lceil k/t \rceil$ variables to 1 between a and b (with a included and b excluded).

As already mentioned, the proof of Theorem 5.4 is based on a bottleneck counting argument in the spirit of [12], with the twist that we consider pairs of bottleneck nodes. To establish the theorem we make use of the following two lemmas which will be proven subsequently.

LEMMA 5.6. *Every path in the support of \mathcal{D} usefully traverses a useful pair.*

LEMMA 5.7. *For every useful pair (a, b) , the probability that a random α chosen from \mathcal{D} usefully traverses (a, b) is at most $2s^{-\varepsilon r/2}$.*

Combining the above lemmas, it is immediate to prove Theorem 5.4. By Lemma 5.6 the probability that a random path α sampled from \mathcal{D} usefully traverses some useful pair is 1. By Lemma 5.7, for any fixed useful pair (a, b) , the probability that a random α usefully traverses (a, b) is at most $2s^{-\varepsilon r/2}$. By a standard union bound argument, it follows that the number of useful pairs is at least $\frac{1}{2}s^{\varepsilon r/2}$, so the number of nodes in P cannot be smaller than $\frac{1}{\sqrt{2}}s^{\varepsilon r/4} \geq \frac{1}{\sqrt{2}}s^{\varepsilon k/t^2}$.

To conclude the proof it remains only to establish Lemmas 5.6 and 5.7.

PROOF OF LEMMA 5.6. Consider any path in the support of \mathcal{D} . By the definition of our random process this path ends in the i^* th clique axiom for some $i^* \in [k]$. By Claim 5.5, the path sets at most k variables to 1 and hence we can split it into t pieces by nodes a_0, a_1, \dots, a_t (a_0 is the source, a_t the sink) so that between a_j and a_{j+1} at most $\lceil k/t \rceil$ variables are set to 1. It remains to prove that for at least one $j \in [t]$ the set

$$W_j = V_{i^*}^0(a_j) \setminus V_{i^*}^0(a_{j-1}) \quad (7)$$

is (r, q) -neighbour-dense. Note that this will prove Lemma 5.6 since by construction (a_{j-1}, a_j) is then a pair that is usefully traversed by the path.

Towards contradiction, assume instead that no W_j is (r, q) -neighbour-dense, i.e., that for all $j \in [t]$ there exists a set of vertices $R_j \subseteq V$ with $|R_j| \leq r$ such that $|\widehat{N}_{W_j}(R_j)| \leq q$. Let $R = \bigcup_{j \in [t]} R_j$. Since the path ends in the i^* th clique axiom we have $V_{i^*}^0(a_t) = V$, and since i^* is not forgotten along the path, it holds that $V_{i^*}^0(a_{j-1}) \subseteq V_{i^*}^0(a_j)$ for each $j \in [t]$. It follows that the sets W_1, \dots, W_t in (7) form a partition of V , and therefore

$$|\widehat{N}_V(R)| = \sum_{j \in [t]} |\widehat{N}_{W_j}(R)| \leq \sum_{j \in [t]} |\widehat{N}_{W_j}(R_j)| \leq tq. \quad (8)$$

Since $|R| \leq \sum_{j \in [t]} |R_j| \leq tr$ this contradicts the assumption that V is (tr, tq) -neighbour-dense. Lemma 5.6 follows. \square

PROOF OF LEMMA 5.7. Fix a useful pair (a, b) . Let E denote the event that a random path sampled from \mathcal{D} usefully traverses (a, b) . Let $i^* = i(a, b)$, $V^1(a) = \bigcup_{j \in [k]} V_j^1(a)$, and $W = V_{i^*}^0(b) \setminus V_{i^*}^0(a)$. Notice that W is guaranteed to be (r, q) -neighbour-dense by our definition of $i(a, b)$. Since G is (k, t, s, ε) -clique-dense by assumption, this implies that W is (tr, r, q', s) -mostly neighbour-dense, and we let S be the set that witnesses this as per Definition 5.2. We bound the probability of the event E by a case analysis based on the size of the set $V^1(a)$. We remark that all probabilities in the calculations that follow are over the choice of $\alpha \sim \mathcal{D}$.

Case 1 ($|V^1(a)| > r/2$): In this case, we simply prove that already the probability of reaching a is small. By definition of $|V^1(a)|$, we have that $|\beta^1(a)| = |V^1(a)|$. Recall that every answer 1 is necessarily the result of a $(rs^{-(1+\varepsilon)}/2\epsilon k)$ -biased coin flip, and that all these decisions are irreversible. That is, if a path ever decides to set a variable in $V^1(a)$ to 0, then its case is lost and it is guaranteed to miss a . Thus we can upper bound the probability of the event E by the probability that a random α passes through a , and, in particular, by the probability of setting all variables in $\beta^1(a)$ to 1 as follows:

$$\Pr[E] \leq \Pr[\alpha \text{ passes through } a] \quad (9)$$

$$\leq (rs^{-(1+\varepsilon)}/2\epsilon k)^{|\beta^1(a)|} \quad (10)$$

$$\leq s^{-\varepsilon |\beta^1(a)|} \quad (11)$$

$$= s^{-\varepsilon |V^1(a)|} \quad (12)$$

$$\leq 2s^{-\varepsilon r/2}, \quad (13)$$

where for (11) we use the fact that $r \leq k$, which follows from $tr \leq k$ and $t \geq 1$.

Case 2 ($|V^1(a)| \leq r/2$): For every path α , let $R(\alpha)$ denote the set of vertices u for which the path α sets some variable $x_{u,i}$ to 1 at some node between a and b (with a included and b excluded); note that $R(\alpha) = \emptyset$ if α does not go through a and b , and that $|R(\alpha)| \leq \lceil k/t \rceil$ for all paths α that satisfy the event E . For the sets

$$\mathcal{R}_0 = \{R : |R| \leq \lceil k/t \rceil \text{ and } |\widehat{N}_W(R \cup V^1(a))| < q'\} \quad (14a)$$

$$\mathcal{R}_1 = \{R : |R| \leq \lceil k/t \rceil \text{ and } |\widehat{N}_W(R \cup V^1(a))| \geq q'\} \quad (14b)$$

we have that

$$\Pr[E] = \Pr[E \text{ and } R(\alpha) \in \mathcal{R}_0] + \Pr[E \text{ and } R(\alpha) \in \mathcal{R}_1]. \quad (15)$$

The first term in (15) is bounded from above by the probability of $R(\alpha) \in \mathcal{R}_0$. Note that $|R| \leq \lceil k/t \rceil \leq 2k/t \leq rt/2$ (since $r \geq 4k/t^2$) for $R \in \mathcal{R}_0$. Hence we have $|R \cup V^1(a)| \leq rt/2 + r/2 \leq rt$ and therefore $|(R \cup V^1(a)) \cap S| \geq r$ by the choice of S . Thus, the probability of $R(\alpha) \in \mathcal{R}_0$ is bounded by the probability that $|R(\alpha) \cap S| \geq r/2$ since $|V^1(a)| \leq r/2$. But since S is small, we can now apply the

union bound and conclude that

$$\Pr[E \text{ and } R(\alpha) \in \mathcal{R}_0] \leq \Pr[R(\alpha) \in \mathcal{R}_0] \quad (16)$$

$$\leq \Pr[|R(\alpha) \cap S| \geq r/2] \quad (17)$$

$$\leq \left(\frac{|S|k}{r/2} \right) \left(\frac{rs^{-(1+\varepsilon)}}{2ek} \right)^{r/2} \quad (18)$$

$$\leq \left(\frac{2e|S|k}{r} \right)^{r/2} \left(\frac{rs^{-(1+\varepsilon)}}{2ek} \right)^{r/2} \quad (19)$$

$$\leq s^{-\varepsilon r/2} . \quad (20)$$

We now bound the second term in (15). First note that, by definition of W , if α is a path that passes through a and b in this order, then all variables x_{u,i^*} with $u \in W$ must be set to 0 in α at some node between a and b . For each path in the support of \mathcal{D} that passes through a and b , some of the variables x_{u,i^*} with $u \in W$ will be set to zero as a result of a coin flip and others will be forced choices.

Fix a path α contributing to the second term in (15). We claim that along this path at least q' variables x_{u,i^*} ($u \in W$) are set to 0 as a result of a coin flip.

Indeed, since $V_{i^*}^1(b) = \emptyset$ and i^* is not forgotten at b , by the monotonicity property the same holds for every node along α before b . This implies that the answer to a query of the form x_{u,i^*} ($u \in W$) made along α cannot be forced by neither item (1) (forgetfulness) in the definition of \mathcal{D} nor by a functionality axiom. Moreover, since $V^1(c) \subseteq R(\alpha) \cup V^1(a)$ for any node c on the path α between a and b , it holds that all variables x_{u,i^*} with $u \in \widehat{N}_W(R(\alpha) \cup V^1(a))$ can not be forced to 0 by an edge axiom either. Since there are at least q' of them, this proves the claim.

Now the analysis of the second term in (15) is completed by the same Markov chain argument as in Case 1 above (noting that irreversibility of decisions still takes place):

$$\Pr[E \text{ and } R(\alpha) \in \mathcal{R}_1] \leq \Pr[\alpha \text{ flips } \geq q' \text{ coins and gets all 0s}] \quad (21)$$

$$\leq (1 - rs^{-(1+\varepsilon)}/2ek)^{q'} \quad (22)$$

$$\leq s^{-\varepsilon r/2} . \quad (23)$$

Adding (20) and (23) we obtain the lemma. \square

6 RANDOM GRAPHS ARE CLIQUE-DENSE

In this section we show that asymptotically almost surely an Erdős-Rényi random graph $G \sim \mathcal{G}(n, p)$ is (k, t, s, ε) -clique-dense for the right choice of parameters.

THEOREM 6.1. *For any real constant $\varepsilon \in (0, 1/2)$, any sufficiently large integer n , any positive integer k and any real $\xi > 1$ such that $k\sqrt{\xi} \leq n^{1/2-\varepsilon}$, if $G \sim \mathcal{G}(n, n^{-2\xi/(k-1)})$ is an Erdős-Rényi random graph then with probability at least $1 - \exp(-\sqrt{n})$ it holds that G is (k, t, s, ε) -clique-dense with $t = 64\xi/\varepsilon$ and $s = (n/\xi)^{1/2}$.*

As a corollary of Theorem 5.4 and Theorem 6.1 we obtain Theorem 4.2, the main result of this paper.

PROOF OF THEOREM 4.2. Clearly $t \geq 128 \geq 1$ as required by Definition 5.3. We can also assume w.l.o.g. that $t \leq k$ since otherwise

$k/\xi^2 \leq 64/(\xi\varepsilon) \leq O(1)$ and the bound becomes trivial. By plugging in the parameters given by Theorem 6.1 to Theorem 5.4 we immediately get the stated lower bound on the length of any regular refutation π of $\text{Clique}(G, k)$

$$|\pi| \geq \frac{1}{\sqrt{2}} s^{\varepsilon k/t^2} \geq n^{\Omega(k/\xi^2)} , \quad (24)$$

for which we have to note that $s \geq n^{1/4}$ since $\xi \leq t \leq k \leq n^{1/2}$. \square

We will spend the rest of this section proving Theorem 6.1. Let $\delta = 2\xi/(k-1)$. We show that, with probability at least $1 - e^{-\sqrt{n}}$, the random graph G is (k, t, s, ε) -clique-dense for parameters as in the statement of the theorem, $r = 4k/t^2$ and $q = \frac{n^{1-t\delta r}}{4t}$.

Recall that $q' = 3\varepsilon k s^{1+\varepsilon} \log s$. Let us argue that these parameters satisfy constraints

$$t\delta r \leq \frac{\varepsilon}{6} , \quad (25)$$

$$tr \log n \leq \frac{n^{1-t\delta r}}{32} \cdot \frac{\log n}{n^{1/2}} , \quad (26)$$

$$\frac{qn^{-t\delta r} s}{16tr} \geq \frac{n^{1+2\varepsilon/3}}{2^8} , \quad (27)$$

$$q' \leq \frac{qn^{-t\delta r}}{4} \cdot \frac{3 \cdot 2^9 \log n}{n^{\varepsilon/6}} , \quad (28)$$

$$tr \leq \frac{q}{2} , \quad (29)$$

which will be used further on in the proof.

As a first step note that for $k \geq 4$

$$t\delta r = \frac{8\xi k}{t(k-1)} \leq \frac{\varepsilon}{6} , \quad (30)$$

and hence (25) holds. Equation (26) follows from the chain of inequalities

$$tr \log n = \frac{4k \log n}{t} \leq \frac{n^{1/2-\varepsilon} \log n}{32} \leq \frac{n^{1-t\delta r}}{32} \cdot \frac{\log n}{n^{1/2}} . \quad (31)$$

To obtain (27) observe that

$$\frac{qn^{-t\delta r} s}{16tr} = \frac{n^{1-2t\delta r+1/2}}{2^8 k \xi^{1/2}} \geq \frac{n^{1-2t\delta r+\varepsilon}}{2^8} \geq \frac{n^{1+2\varepsilon/3}}{2^8} . \quad (32)$$

To see that (28) holds, note that

$$q' = 3\varepsilon k s^{1+\varepsilon} \log s \quad (33)$$

$$\leq \frac{3\varepsilon k n^{(1+\varepsilon)/2} \log n}{2\xi^{1/2}} \quad (34)$$

$$= \frac{3 \cdot 2^5 \cdot k \xi^{1/2} n^{(1+\varepsilon)/2} \log n}{t} \quad (35)$$

$$\leq \frac{3 \cdot 2^9 \cdot n^{1-\varepsilon/2} \log n}{16t} \quad (36)$$

$$\leq \frac{qn^{-t\delta r}}{4} \cdot \frac{3 \cdot 2^9 \log n}{n^{\varepsilon/6}} . \quad (37)$$

Finally, for (29), we just observe that

$$tr = \frac{4k}{t} \leq \frac{4k^2}{t^2} \leq \frac{n^{1-2\varepsilon}}{16t} \leq \frac{q}{2} , \quad (38)$$

where we use that $k \geq t$ and $t \geq 64$.

We must now prove that asymptotically almost surely G is (k, t, s, ε) -clique-dense for the chosen parameters; all probabilities in this section are over the choice of G . Let $V = V(G)$.

The fact that asymptotically almost surely V is (tr, tq) -neighbour-dense is quite immediate. First, for any $R \subseteq V$ with $|R| \leq tr$,

$$\mathbb{E}[|\widehat{N}(R)|] = |V \setminus R|n^{-\delta|R|} \quad (39)$$

$$\geq (n - tr)n^{-\delta tr} \quad (40)$$

$$\geq \left(n - \frac{q}{2}\right)n^{-\delta tr} \quad (41)$$

$$\geq \frac{n^{1-\delta tr}}{2}, \quad (42)$$

where (42) follows from (29) and the trivial fact that $q \leq n$. Hence, we can bound the probability that V is not (tr, tq) -neighbour-dense by

$$\Pr\left[\exists R \subseteq V, |R| \leq tr \wedge |\widehat{N}(R)| \leq tq\right] \leq \sum_{j=1}^{tr} \binom{n}{j} \max_R \Pr\left[|\widehat{N}(R)| \leq tq\right] \quad (43)$$

$$\leq n^{tr} \max_R \Pr\left[|\widehat{N}(R)| \leq \frac{n^{1-\delta tr}}{4}\right] \quad (44)$$

$$\leq n^{tr} \exp\left(-\frac{n^{1-\delta tr}}{16}\right) \quad (45)$$

$$\leq \exp\left(-\frac{n^{1-\delta tr}}{32} \cdot \left(2 - \frac{\log n}{n^{1/2}}\right)\right) \quad (46)$$

$$\leq e^{-\sqrt{n}}. \quad (47)$$

We note that (43) is a union bound, (44) follows from the definition of q , (45) is the multiplicative form of Chernoff bound (note that the events $v \in \widehat{N}(R)$ ($v \in V \setminus R$) are mutually independent), (46) follows from (26), and (47) holds for large enough n by (25) and the fact that $\varepsilon < 1/2$.

All that is left to prove is that asymptotically almost surely G satisfies property 2 in Definition 5.3, that is that every (r, q) -neighbour-dense set $W \subseteq V$ is (tr, r, q', s) -mostly neighbour-dense. For shortness let P be the event that G satisfies this property. We wish to show that $\Pr[\neg P] \leq e^{-\Omega(n)}$.

Given an (r, q) -neighbour-dense set $W \subseteq V$ we will define a set S_W which will be a ‘‘candidate witness’’ of the fact that W is (tr, r, q', s) -mostly neighbour-dense. First observe that, since W is (r, q) -neighbour-dense and $q' \leq q$ by (28), any set $R \subseteq V$ with $|R| \leq tr$ and $|\widehat{N}_W(R)| \leq q'$ must be such that $|R| > r$. We will use a sequence of such sets R and construct S_W in a somewhat greedy fashion. To this end, the following definition will be useful. A tuple of sets (R_1, \dots, R_m) is said to be r -disjoint if $|R_i \cap (\bigcup_{j < i} R_j)| \leq r$ for every $i \in [m]$.

Fix an arbitrary ordering of the subsets of V . Define $\vec{R}_W = (R_1, \dots, R_m)$ to be a maximally long tuple such that, for every $i = 1, \dots, m$, the set R_i is the first in the ordering such that $|R_i| \leq tr$, $|\widehat{N}_W(R_i)| \leq q'$ and $|R_i \cap (\bigcup_{j < i} R_j)| \leq r$. Note that \vec{R}_W is r -disjoint. Now let $S_W = \bigcup_{i \leq m} R_i$.

Observe that, by maximality of \vec{R}_W , any set $R \subseteq V$ with $|R| \leq tr$ and $|\widehat{N}_W(R)| \leq q'$ must be such that $|R \cap S| > r$. This implies that if $|S_W| \leq s$ then S_W witnesses the fact that W is (tr, r, q', s) -mostly neighbour-dense. Therefore we have that

$$\Pr[\neg P] \leq \Pr[\exists (r, q)\text{-neighbour-dense } W \subseteq V \text{ with } |S_W| > s]. \quad (48)$$

Let $Q(W)$ denote the event that W is (r, q) -neighbour-dense. Moreover, let \mathcal{W} be the collection of all pairs (W, \vec{R}) such that $W \subseteq V$, $\vec{R} = (R_1, \dots, R_\ell)$ for $\ell = \lceil s/tr \rceil$, $R_j \subseteq V$ and $0 < |R_j| \leq tr$ for each $j \in [\ell]$, and \vec{R} is r -disjoint. Notice that if there exists an (r, q) -neighbour-dense W such that $\vec{R}_W = (R_1, \dots, R_m)$ and $|S_W| > s$, then $m \geq \ell$ and $(W, (R_1, \dots, R_\ell)) \in \mathcal{W}$. Furthermore, by definition of \vec{R}_W , for every $j \in [\ell]$ it holds that $|\widehat{N}_W(R_j)| \leq q'$. Hence we can conclude that

$$\Pr[\neg P] \leq \Pr[\exists (W, \vec{R}) \in \mathcal{W} \ Q(W) \wedge \forall j \in [\ell], |\widehat{N}_W(R_j)| \leq q'] \quad (49)$$

$$\leq 2^n n^{tr\ell} \max_{(W, \vec{R}) \in \mathcal{W}} \Pr[Q(W) \wedge \forall j \in [\ell], |\widehat{N}_W(R_j)| \leq q'] \quad (50)$$

$$\leq 2^n n^s \max_{(W, \vec{R}) \in \mathcal{W}} \Pr[Q(W) \wedge \forall j \in [\ell], |\widehat{N}_W(R_j)| \leq \frac{q}{4} n^{-t\delta r}], \quad (51)$$

where (51) follows for n large enough from the bound in (28).

Now fix $(W, \vec{R}) \in \mathcal{W}$ and let R_j^d (resp. R_j^c) be the subset of R_j disjoint from (resp. contained in) $\bigcup_{j' < j} R_{j'}$. Since $|R_j^c| \leq r$ by definition, it holds that if W is (r, q) -neighbour-dense then $|\widehat{N}_W(R_j^c)| > q$. Let $F(j)$ be the event that $|\widehat{N}_W(R_j^c)| > q$ and $|\widehat{N}_W(R_j)| \leq \frac{q}{4} n^{-t\delta r}$. Note that $\Pr[Q(W) \wedge \forall j \in [\ell], |\widehat{N}_W(R_j)| \leq \frac{q}{4} n^{-t\delta r}]$ is at most $\Pr[\forall j \in [\ell], F(j)]$. Let $F'(j)$ be the event that $F(j')$ holds for all $j' \in [j - 1]$. We have that

$$\Pr[\forall j \in [\ell], F(j)] = \prod_{j \in [\ell]} \Pr[F(j) \mid F'(j)]. \quad (52)$$

We can consider the factors of the previous product separately and bound each one by

$$\Pr[F(j) \mid F'(j)] \leq \sum_{\substack{U \subseteq W \\ |U| \geq q}} \Pr\left[|\widehat{N}_U(R_j^d)| \leq \frac{q}{4} n^{-t\delta r} \mid \widehat{N}_W(R_j^c) = U \wedge F'(j)\right] \cdot \Pr\left[\widehat{N}_W(R_j^c) = U \mid F'(j)\right] \quad (53)$$

$$\leq \sum_{\substack{U \subseteq W \\ |U| \geq q}} \Pr\left[|\widehat{N}_U(R_j^d)| \leq \frac{q}{4} n^{-t\delta r}\right] \cdot \Pr\left[\widehat{N}_W(R_j^c) = U \mid F'(j)\right] \quad (54)$$

$$\leq \sum_{\substack{U \subseteq W \\ |U| \geq q}} \exp\left(-\frac{qn^{-t\delta r}}{16}\right) \cdot \Pr\left[\widehat{N}_W(R_j^c) = U \mid F'(j)\right] \quad (55)$$

$$= \exp\left(-\frac{qn^{-t\delta r}}{16}\right) \cdot \sum_{\substack{U \subseteq W \\ |U| \geq q}} \Pr\left[\widehat{N}_W(R_j^c) = U \mid F'(j)\right] \quad (56)$$

$$\leq \exp\left(-\frac{qn^{-t\delta r}}{16}\right). \quad (57)$$

Equation (54) follows from the independence of any two events that involve disjoint sets of potential edges and (55) follows from the multiplicative Chernoff bound and the fact that

$$\mathbb{E}[|\widehat{N}_U(R_j^d)|] = |U \setminus R_j^d| n^{-\delta |R_j^d|} \quad (58)$$

$$\geq (|U| - tr) n^{-\delta tr} \quad (59)$$

$$\geq \frac{q}{2} n^{-\delta tr} . \quad (60)$$

So, putting everything together, we have that

$$\Pr[\neg P] \leq 2^n n^s \exp\left(-\frac{qn^{-t\delta r} \ell}{16}\right) \quad (61)$$

$$\leq e^{(\log 2)n + \sqrt{n} \log n - (n^{1+2\epsilon/3})/2^8} \quad (62)$$

$$\leq e^{-\Omega(n)} , \quad (63)$$

where the last inequality holds for n large enough, and the second to last inequality follows immediately from the bound in (27). This concludes the proof of Theorem 6.1.

7 CONCLUDING REMARKS

In this paper we prove optimal average-case lower bounds for regular resolution proofs certifying k -clique-freeness of Erdős-Rényi graphs not containing k -cliques. These lower bounds are also strong enough to apply for several state-of-the-art clique algorithms used in practice.

The most immediate and compelling question arising from this work is whether the lower bounds for regular resolution can be strengthened to hold also for general resolution. A closer study of our proof reveals that there are several steps that rely on regularity. However, there is no connection per se between regular resolution and the abstract combinatorial property of graphs that we show to be sufficient to imply regular resolution lower bounds. Thus, it is tempting to speculate that this property, or perhaps some modification of it, might be sufficient to obtain lower bounds also for general resolution. If so, a natural next step would be to try to extend the lower bound further to the polynomial calculus proof system capturing Gröbner basis calculations.

Another interesting question is whether the lower bounds we obtain asymptotically almost surely for random graphs can also be shown to hold deterministically under the weaker assumption that the graph has certain pseudorandom properties. Specifically, is it possible to get an $n^{\Omega(\log n)}$ length lower bound for the class of Ramsey graphs? A graph on n vertices is called *Ramsey* if it has no set of $\lceil 2 \log_2 n \rceil$ vertices forming a clique or independent set. It is known that for sufficiently large n a random graph sampled from $\mathcal{G}(n, 1/2)$ is Ramsey with high probability. Is it true that for a Ramsey graph G on n vertices the formula $\text{Clique}(G, \lceil 2 \log_2 n \rceil)$ requires (regular) resolution refutations of length $n^{\Omega(\log n)}$? Such a lower bound is known for tree-like resolution [21] and proving it for general resolution would have interesting consequences in other areas of proof complexity [10].

ACKNOWLEDGEMENTS

This work has been a long journey, and different subsets of the authors want to acknowledge fruitful and enlightening discussions with different subsets of Christoph Berkholz, Olaf Beyersdorff, Nicola Galesi, Ciaran McCreesh, Toni Pitassi, Pavel Pudlák, Ben Rossman, Navid Talebanfard, and Neil Thapen. A special thanks to Shuo Pang for having pointed out an inaccuracy in the probabilistic argument in Section 6 and having suggested a fix.

The first, second, and fourth authors were supported by the European Research Council under the European Union's Horizon 2020 Research and Innovation Programme / ERC grant agreement no. 648276 AUTAR. The third and fifth authors were supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007–2013) / ERC grant agreement no. 279611 as well as by Swedish Research Council grants 621-2012-5645 and 2016-00782, and the second author did part of this work while at KTH Royal Institute of Technology supported by the same grants. The last author was supported by the Russian Foundation for Basic Research.

REFERENCES

- [1] Noga Alon, Michael Krivelevich, and Benny Sudakov. 1998. Finding a large hidden clique in a random graph. *Random Structures and Algorithms* 13, 3-4 (1998), 457–466.
- [2] Roberto J. Bayardo Jr. and Robert Schrag. 1997. Using CSP Look-Back Techniques to Solve Real-World SAT Instances. In *Proceedings of the 14th National Conference on Artificial Intelligence (AAAI '97)*, 203–208.
- [3] Paul Beame, Russell Impagliazzo, and Ashish Sabharwal. 2007. The Resolution Complexity of Independent Sets and Vertex Covers in Random Graphs. *Computational Complexity* 16, 3 (Oct. 2007), 245–297. Preliminary version in *CCC '01*.
- [4] Paul Beame and Toniann Pitassi. 1996. Simplified and Improved Resolution Lower Bounds. In *Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science (FOCS '96)*, 274–282.
- [5] Eli Ben-Sasson and Avi Wigderson. 2001. Short Proofs are Narrow—Resolution Made Simple. *J. ACM* 48, 2 (March 2001), 149–169. Preliminary version in *STOC '99*.
- [6] Olaf Beyersdorff, Nicola Galesi, and Massimo Lauria. 2013. Parameterized Complexity of DPLL Search Procedures. *ACM Transactions on Computational Logic* 14, 3, Article 20 (Aug. 2013), 21 pages. Preliminary version in *SAT '11*.
- [7] Olaf Beyersdorff, Nicola Galesi, Massimo Lauria, and Alexander A. Razborov. 2012. Parameterized Bounded-Depth Frege Is not Optimal. *ACM Transactions on Computation Theory* 4, 3 (Sept. 2012), 7:1–7:16. Preliminary version in *ICALP '11*.
- [8] Jianer Chen, Xiuzhen Huang, Iyad A. Kanj, and Ge Xia. 2004. Linear FPT reductions and computational lower bounds. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC '04)*, 212–221.
- [9] Stephen A. Cook and Robert Reckhow. 1979. The relative efficiency of propositional proof systems. *Journal of Symbolic Logic* 44, 1 (March 1979), 36–50.
- [10] Stefan S. Dantchev, Barnaby Martin, and Stefan Szeider. 2011. Parameterized Proof Complexity. *Computational Complexity* 20 (March 2011), 51–85. Issue 1. Preliminary version in *FOCS '07*.
- [11] Rodney Downey and Michael R. Fellows. 1995. Fixed-Parameter Tractability and Completeness II: Completeness for W[1]. *Theoretical Computer Science A* 141 (1995), 109–131. Preliminary versions of some of the results of this paper were presented at the 21st Manitoba Conference on Numerical Mathematics and Computation, 1991.
- [12] Armin Haken. 1985. The Intractability of Resolution. *Theoretical Computer Science* 39, 2-3 (Aug. 1985), 297–308.
- [13] Johan Håstad. 1999. Clique is Hard to Approximate within $n^{1-\epsilon}$. *Acta Mathematica* 182 (1999), 105–142. Preliminary version in *FOCS '96*.
- [14] Russell Impagliazzo and Ramamohan Paturi. 2001. On the Complexity of k -SAT. *J. Comput. System Sci.* 62, 2 (March 2001), 367–375. Preliminary version in *CCC '99*.
- [15] Richard M. Karp. 1972. Reducibility among Combinatorial Problems. In *Complexity of Computer Computations*. Springer, 85–103.
- [16] Richard M. Karp. 1976. The probabilistic analysis of some combinatorial search algorithms. In *Algorithms and Complexity: New Directions and Recent Results*. Academic Press, New York, 1–19.
- [17] Donald E. Knuth. 1994. The sandwich theorem. *The Electronic Journal of Combinatorics* 1, A1 (1994), 1–48.

- [18] Jan Krajíček. 1997. Interpolation Theorems, Lower Bounds for Proof Systems, and Independence Results for Bounded Arithmetic. *Journal of Symbolic Logic* 62, 2 (June 1997), 457–486.
- [19] Jan Krajíček. 1995. *Bounded Arithmetic, Propositional Logic, and Complexity Theory*. Cambridge University Press, New York.
- [20] Luděk Kučera. 1995. Expected complexity of graph partitioning problems. *Discrete Applied Mathematics* 57, 2 (1995), 193–212.
- [21] Massimo Lauria, Pavel Pudlák, Vojtěch Rödl, and Neil Thapen. 2017. The Complexity of Proving That a Graph is Ramsey. *Combinatorica* 37, 2 (April 2017), 253–268. Preliminary version in *ICALP '13*.
- [22] László Lovász. 1979. On the Shannon capacity of a graph. *IEEE Transactions on Information theory* 25, 1 (Jan. 1979), 1–7.
- [23] João P. Marques-Silva and Karem A. Sakallah. 1999. GRASP: A Search Algorithm for Propositional Satisfiability. *IEEE Trans. Comput.* 48, 5 (May 1999), 506–521. Preliminary version in *ICCAD '96*.
- [24] Ciaran McCreesh. 2017. *Solving Hard Subgraph Problems in Parallel*. Ph.D. Dissertation. University of Glasgow.
- [25] Matthew W. Moskewicz, Conor F. Madigan, Ying Zhao, Lintao Zhang, and Sharad Malik. 2001. Chaff: Engineering an Efficient SAT Solver. In *Proceedings of the 38th Design Automation Conference (DAC '01)*. 530–535.
- [26] Jaroslav Nešetřil and Svatopluk Poljak. 1985. On the complexity of the subgraph problem. *Commentationes Mathematicae Universitatis Carolinae* 026, 2 (1985), 415–419.
- [27] Patrick Prosser. 2012. Exact Algorithms for Maximum Clique: A Computational Study. *Algorithms* 5, 4 (2012), 545–587.
- [28] Pavel Pudlák. 1997. Lower Bounds for Resolution and Cutting Plane Proofs and Monotone Computations. *Journal of Symbolic Logic* 62, 3 (Sept. 1997), 981–998.
- [29] Alexander Razborov, Avi Wigderson, and Andrew Yao. 2002. Read-once branching programs, rectangular proofs of the pigeonhole principle and the transversal calculus. *Combinatorica* 22, 4 (2002), 555–574.
- [30] Benjamin Rossman. 2008. On the Constant-Depth Complexity of k -Clique. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC '08)*. 721–730.
- [31] Benjamin Rossman. 2010. *Average-Case Complexity of Detecting Cliques*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- [32] Benjamin Rossman. 2014. The Monotone Complexity of k -Clique on Random Graphs. *SIAM J. Comput.* 43, 1 (2014), 256–279. Preliminary version in *FOCS '10*.
- [33] Virginia Vassilevska. 2009. Efficient Algorithms for Clique Problems. *Inform. Process. Lett.* 109, 4 (Jan. 2009), 254–257.
- [34] David Zuckerman. 2007. Linear Degree Extractors and the Inapproximability of Max Clique and Chromatic Number. *Theory of Computing* 3, 6 (Aug. 2007), 103–128. Preliminary version in *STOC '06*.