Electronic Colloquium on Computational Complexity, Report No. 12 (2004)

The Resolution Complexity of Random Graph k-Colorability

Paul Beame*
Computer Science and Engineering
University of Washington
Box 352350
Seattle, WA 98195
beame@cs.washington.edu

Joseph Culberson†
Department of Computer Science
University of Alberta
Edmonton, Alberta
Canada
joe@cs.ualberta.ca

David Mitchell‡
Department of Computer Science
Simon Fraser University
Burnaby, British Columbia
Canada
mitchell@cs.sfu.ca

Cristopher Moore§
Computer Science Department
University of New Mexico
Albuquerque, NM 87131
moore@cs.unm.edu

December 17, 2003

Abstract

We consider the resolution proof complexity of propositional formulas which encode random instances of graph k-colorability. We obtain a tradeoff between the graph density and the resolution proof complexity. For random graphs with linearly many edges we obtain linear-exponential lower bounds on the length of resolution refutations. For any $\epsilon > 0$, we obtain sub-exponential lower bounds of the form 2^{n^δ} for some $\delta > 0$ for non-k-colorability proofs of graphs with n vertices and $O(n^{\frac{3}{2}-\frac{1}{k}-\epsilon})$ edges. We obtain sharper lower bounds for Davis-Putnam-DPLL proofs and for proofs in a system considered by McDiarmid.

We also show that very simple algorithms achieve qualitatively similar running times.

1 Introduction

The problem of coloring graphs has a long history. The decision problem, k-colorability, being an important NP-complete problem, has generated significant interest for random graphs as well. Since the transcript of any complete algorithm for k-coloring on input G also provides a proof of non-k-colorability in the case that G is not k-colorable, the study of the proof complexity of k-colorability also yields bounds on the running time of such algorithms. We consider here the question of the proof complexity of non-k-colorability for random graphs.

*Research supported by NSF grants CCR-9800124 and CCR-0098066
†Research supported by NSERC grant OGP8053
‡Research supported by NSERC grant RGPIN238987
§Research supported by NSF grant PHY-0200909

ISSN 1433-8092

With sufficiently many edges, a random graph is not k-colorable almost certainly. McDiarmid [McD84] studied the complexity of deciding non-k-colorability in random graphs using a particular natural class of proof procedures that emulates a variety of coloring algorithms. Here we consider natural encodings of the k-colorability of random graphs as CNF formulas and examine the resolution complexity of these formulas.

It is easy to see how many natural algorithms for k-coloring can be simulated as resolution proof procedures. Indeed, many natural algorithms, including the most efficient current procedures can naturally be phrased in terms of resolution procedures in which the graph of resolution inferences is a tree. This is a subclass of resolution proofs known collectively as the Davis-Putnam or DPLL procedure. (Resolution itself comes with no underlying search procedure for proofs whereas DPLL procedures include a search strategy as well.) In our analysis we build on the ideas that have been used to analyze the resolution proof complexity of random k-CNF formulas together with results about the colorability properties of random graphs.

For any graph G, we denote by $\chi(G, k)$ the formula for the "natural" encoding of the statement "G is k-colorable", as a propositional formula. For any CNF formula ϕ, we denote by $Res(\phi)$ the length of the shortest resolution refutation of ϕ. By convention, we let $Res(\phi) = \infty$ if ϕ is satisfiable.

The general scheme of our lower bounds proof is similar to that of the proofs of resolution lower bounds for random k-SAT [CS88, BP96, BKPS02, BSW01]. We make particular use of the recent relationship between clause width and proof length for resolution proofs shown by Ben-Sasson and Wigderson [BSW01]. The argument is as follows.

1. Almost surely, every small sub-graph of G has low degree, and thus is k-colorable. From this, it can be shown that in any refutation of $\chi(G, k)$ there is a "complex clause" that cannot be derived from $\chi(H, k)$, for any H which is a very small subgraph of G.

2. Almost surely, every small sub-graph of G has many vertices of degree less than k. From this it can be shown that any "complex clause" contains many literals. The theorem of Ben-Sasson and Wigderson relating minimum refutation width to minimum refutation length then implies that any refutation of $\chi(G, k)$ must be long.

The class of procedures considered by McDiarmid produces tree-like proofs but these are not DPLL procedures. Nonetheless it is possible to show using arguments similar to those of [CEI96, BSW01] that if these procedures produce short proofs then they have efficient simulations by resolution proofs of small width. Thus, all our lower bounds apply to these procedures as well.

2 Preliminaries

For a set of boolean variables $V = \{v_1, \ldots, v_n\}$, a literal is any variable $x \in V$, or its the negation $\neg x$, a clause is a set of literals, and a formula is a set of clauses. The interpretation of a formula is as a Conjunctive Normal Form (CNF) formula of propositional calculus, that is, a conjunction of disjunctions. A truth assignment for V is a function $\tau : V \mapsto \{T, F\}$. Assignment τ satisfies a clause C iff $\tau(l) = T$ for at least one literal $l \in C$, and τ satisfies a set of clauses ϕ, written $\tau \models \phi$, iff it satisfies every clause in ϕ. A formula is satisfiable iff there is an assignment to its variables which satisfies it. For any clause or formula A, *vars*(A) denotes the set of variables which appear (negated or otherwise) in A. The width of a clause C, denoted $w(C)$ is the number of variables occurring in C. We will consider only clauses which are not tautologies, thus $w(C) \equiv |vars(C)|$.

2.1 k-Colorability Formulas

For any graph $G = \langle V, E \rangle$, we let $n \stackrel{def}{=} |V|$. A *k-coloring* of G is a function col which maps each vertex onto an integer in $[k] = \{1, \ldots, k\}$. A *proper coloring* is a coloring for which $(u, v) \in E \Rightarrow col(u) \neq col(v)$, and we say that G is *k-colorable* iff there is a proper k-coloring of G. We define $\phi = \chi(G, k)$ to be the formula on kn variables with n positive clauses of width k, and $k|E|$ negative clauses of width 2 as follows.

1. For each $v \in V$, $vars(\phi)$ has k propositional variables, $\{x_{v,1}, \ldots, x_{v,k}\}$, and ϕ has one positive clause of width k,
$$\left(\bigvee_{j \in [k]} x_{v,j} \right) \in \phi.$$
2. For each $v \in V$, ϕ has $\binom{k}{2}$ negative clauses of width 2, $\forall i < j \in [k] \;\; (\neg x_{v,i} \vee \neg x_{v,j}) \in \phi$.
3. For each edge $(u, v) \in E$, ϕ has k negative clauses of width 2,
$$\forall l \in [k] \;\; (\neg x_{u,l} \vee \neg x_{v,l}) \in \phi.$$

Clearly, $\chi(G, k)$ is satisfiable iff G is k-colorable. Corresponding to each coloring of G is a unique truth assignment for $\chi(G, k)$. (The reverse is not the case, since some assignments give a vertex no color or multiple colors.) We sometimes fail to distinguish between colorings and assignments. For example, we may say that a coloring makes a clause true, with the obvious meaning.

2.2 Random Graphs

In the study of random graphs there are three natural models one could consider. The most commonly considered models are $\mathcal{G}(n, p)$, where each of the $\binom{n}{2}$ edges is chosen independently with probability p, and $\mathcal{G}_{n,m}$, where a set of precisely m distinct edges is chosen uniformly at random. We find it most convenient to express our proofs in terms of a third distribution, $\hat{\mathcal{G}}_{n,m}$, where m edges are chosen independently with replacement (and duplicates are ignored). As shown, for example, in [AV79] if $p = m/\binom{n}{2}$ then, when considering properties that are monotone (or anti-monotone), the almost certain properties under all three distributions are the same up to a change from m to $m \pm o(m)$. Our results will therefore apply to $\mathcal{G}(n, p)$ and $\mathcal{G}_{n,m}$ as well.

We will consider the dependence of our results on a graph density parameter $\Delta = \Delta(n) = m/n$, the ratio of edges to vertices.

2.3 Resolution Complexity

Resolution is a rule of inference for clauses, which allows one to derive the clause $C \cup D$ from two clauses $C \cup l$ and $D \cup \neg l$. A Resolution derivation from a set of clauses ϕ is a sequence of clauses, $\pi = C_0, \ldots C_m$, where each clause C_i is either an element of ϕ or is derived by the resolution rule from two clauses C_j, E_k occurring in π, for $j, k < i$. The derivation is of length m. A resolution derivation of the empty clause (denoted Λ) from ϕ is called a *refutation* of ϕ. The fundamental property of resolution is that there is a refutation of a set ϕ of clauses if and only if ϕ is unsatisfiable. The resolution complexity of ϕ, here denoted $Res(\phi)$, is the length of the shortest refutation of ϕ.

A related method for CNF formula satisfiability is the Davis-Putnam-DPLL procedure [DLL62]. Such a procedure can be described recursively as follows: First check whether F is trivially satisfiable (has no clauses) or is trivially unsatisfiable (contains an empty clause) and if so stop. Otherwise, select a literal l and search for a satisfying assignment for the formula $F\lceil_{l=1}$ obtained by setting l to true in F (eliminating all clauses containing l and removing $\neg l$ from those clauses that contain it). Otherwise, repeat the search with the formula $F\lceil_{\neg l=1}$. If neither of these searches finds a satisfying assignment then F is not satisfiable. DPLL algorithms will typically select a literal appearing in a clause of length one, called a *unit clause*, if one exists since that literal must be set to true to satisfy the formula. In the case that there are no unit clauses there are many heuristics, called splitting rules, for the selection of the next literal l that have been used in the literature. It is not hard to show that any DPLL algorithm actually produces a resolution refutation and moreover that the form of this refutation is *tree-like* in that the graph of inferences forms a binary tree. Let $DPLL(\phi)$ denote the size of the shortest DPLL refutation of ϕ. The following key relationship between the proof size and resolution width was shown by Ben-Sasson and Wigderson.

Proposition 1 ([BSW01]). *Let $w^*(F)$ is the minimum over all resolution refutations Π of F of the largest width of a clause in Π. Then $DPLL(F) \geq 2^{w^*(F)-w(F)}$ and $Res(F) \geq 2^{c(w^*(F)-w(F))^2/n}$ for $c = 1/(9 \ln 2)$.*

The class of algorithms considered by McDiarmid is similar in spirit to DPLL algorithms, except that instead of trying assignments to a particular Boolean variable (akin to choosing the specific color for a vertex), one chooses whether or not two non-adjacent vertices will be colored the same or differently. This is represented by graph operations that either identify the non-adjacent vertices or add an edge between them.

More formally, McDiarmid's proof system for non-k-colorability has as its objects graphs H derived from the input graph G. The axioms of the proof system are the k'-cliques for any $k' > k$. Given a graph H with two non-adjacent vertices $u, v \in H$, then H follows from $H \cup \{(u, v)\}$ and H_{uv} where H_{uv} is the graph obtained by identifying u and v and naming the resulting vertex v. H can also follow from any H' such that H' is a subgraph of H. McDiarmid only considered proofs whose inference graph forms a tree.

Lemma 1. *Let $k \geq 2$. If the non-k-colorability of a graph G can be proven by a size S tree-like proof in McDiarmid's proof system then there is a resolution refutation of $\chi(G, k)$ of width at most $k(k + 1) + 2k \log_2 S$.*

Proof. The proof follows a general argument due to Russell Impagliazzo (personal communication) that extends the width-size relationship for tree-like resolution in [BSW01] to decisions involving bounded numbers of variables. Write $F \vdash_w F'$ if and only if there is a resolution derivation of F' from F, each of whose clauses has width at most w.

The proof is by induction on S. We begin with the base case of $S = 1$. Clearly for $k' > k$, by considering the clauses of $\chi(K_{k'}, k)$ that only involve the variables for the first $k + 1$ vertices of $K_{k'}$, $\chi(K_{k'}, k) \vdash_{k(k+1)} \Lambda$.

Now consider the last inference of the tree-like McDiarmid proof and suppose that the claim is true for all strictly smaller McDiarmid proofs. If that last inference derived H from a subgraph H' then we note that $\chi(H, k) \vdash_k \chi(H', k)$ since $\chi(H', k)$ is a subformula of $\chi(H, k)$ and each of its clauses has size at most k; therefore the size bounds follows by the inductive hypothesis for H'. Alternatively, the last inference derived H from $H_0 = H \cup \{(u, v)\}$ and $H_1 = H_{uv}$ where $(u, v) \notin H$. One of the proofs that these graphs are not k-colorable has size at most $S/2$ and the other has size at most S.

Let $T_0 = \text{Unequal}(u,v) = \bigwedge_{l\in[k]}(\neg x_{u,l} \vee \neg x_{v,l})$, $T_1 = \text{Equal}(u,v) = \bigwedge_{l\in[k]}((\neg x_{u,l} \vee x_{v,l}) \wedge (\neg x_{v,l} \vee x_{u,l}))$, and $\text{Colored}(u) = (x_{u,1} \vee \ldots \vee x_{u,k})$. Observe that $(\text{Colored}(u) \wedge \text{Colored}(v)) \rightarrow (T_0 \vee T_1)$ is a tautology involving $2k$ variables and that for $u, v \in H$, $\text{Colored}(u)$ and $\text{Colored}(v)$ are clauses of $\chi(H,k)$.

We show that for $b = 0, 1$, $\chi(H,k) \wedge T_b \vdash_k \chi(H_b,k)$. For $b = 0$, note that $\chi(H_0,k) = \chi(H \cup \{(u,v)\},k) = \chi(H,k) \wedge \text{Unequal}(u,v) = \chi(H,k) \wedge T_0$ and it follows trivially. For $b = 1$, observe that the only clauses of $\chi(H_{uv},k)$ that are not already in $\chi(H,k)$ are those of the form $(\neg x_{v,l} \vee \neg x_{w,l})$ where $(\neg x_{u,l} \vee \neg x_{w,l})$ is in $\chi(H,k)$. Each such clause follows easily from $\chi(H,k) \wedge \text{Equal}(u,v)$ by resolving $(\neg x_{u,l} \vee \neg x_{w,l})$ with $(\neg x_{v,l} \vee x_{u,l})$. Therefore $\chi(H,k) \wedge T_1 = \chi(H,k) \wedge \text{Equal}(u,v) \vdash_k \chi(H_{uv},k) = \chi(H_1,k)$.

By the inductive hypothesis there are resolution refutations of $\chi(H_0,k)$, $\chi(H_1,k)$ such that one has width at most $w = k(k+1) + 2k\log_2 S$ and the other one has width at most $w - 2k$. Assume that the refutation of H_0 is narrower (wlog). Since $\chi(H,k) \wedge T_b \vdash_k \chi(H_b,k)$ for $b = 0, 1$ and $w \geq k(k+1) \geq 3k$, we have $\chi(H,k) \wedge T_0 \vdash_{w-2k} \Lambda$ and $\chi(H,k) \wedge T_1 \vdash_w \Lambda$.

We now come to the key point in the argument. We can convert the resolution refutation witnessing $\chi(H,k) \wedge T_0 \vdash_{w-2k} \Lambda$ into a refutation witnessing $\chi(H,k) \vdash_w T_1$. For every assignment σ to the $2k$ variables in T_0 that satisfies T_0 apply σ as a restriction to the proof. The result is a derivation $\chi(H,k)\lceil_\sigma \vdash_{w-2k} \Lambda$. By Lemma 3.1 of [BSW01] one can add back literals of $\bar{\sigma}$ as needed and apply weakening to derive $\chi(H,k) \vdash_w \bar{\sigma}$. Doing this for all such choices of σ we obtain $\chi(H,k) \vdash_w \bar{T_0}$ where $\bar{T_0}$ is the canonical CNF formula for the truth table of $\neg T_0$. Now using the clauses $\text{Colored}(u)$ and $\text{Colored}(v)$ from $\chi(H,k)$ together with $\bar{T_0}$ we derive T_1 which is logically implied. This last derivation requires width only $2k$. Finally, we apply the proof witnessing $\chi(H,k) \wedge T_1 \vdash_w \Lambda$. The overall width is at most w as required. □

2.4 k-Colorability of Random Graphs

As proved by Achlioptas and Friedgut [AF99], for every integer $k \geq 2$, there is a function $c_k(n)$ bounded by a constant such that for $\Delta > c_k(n)$ for $G \sim \mathcal{G}(n, 2\Delta/(n-1))$, the probability G is k-colorable goes to 0 and for $\Delta < c_k(n)$, this probability goes to 1. Let $c_k^+ = \limsup_{n\to\infty} c_k(n)$ and $c_k^- = \liminf_{n\to\infty} c_k(n)$. By a result of Łuczak [Łuc91, JŁR00], both c_k^+ and c_k^- are $k\ln k + O(k\ln\ln k)$; Achlioptas and Naor [AN] have recently shown even tighter results that at every density the chromatic number almost certainly takes on one of at most two values. Further, for $k = 3$ by results of Achlioptas, Molloy, and Moore [AM99, AM97, AM02], $c_3^+ < 2.522$ and $c_3^- > 2.01$.

3 Proof of Lower Bounds

3.1 Subgraph Boundary Size, Expansion, Width, and Length

We define the *k-boundary* of a graph G, denoted $\beta_k(G)$, to be the set of vertices in G of degree between 1 and $k-1$.

For a subgraph $H < G$, let $E_k(H)$ denote the conjunction of the edge (negative) clauses of $\chi(G,k)$ corresponding to the edges of H. We say that H *implies* a clause C if and only if $E_k(H) \rightarrow C$ is true at α for all truth assignments α corresponding to total (but not necessarily proper) colorings of G.

Lemma 2. *Let C be clause in the variables of $\chi(G,k)$. If $H < G$ is a minimal induced subgraph of G that implies C then*

- *H has no isolated vertices, and*
- $w(C) \geq |\beta_k(H)|$.

Proof. First observe that if H has an isolated vertex u then $E_k(H) = E_k(H - u)$ and thus $H - u$ also implies C contradicting the assumption that H was minimal.

Now consider a vertex v in H of degree between 1 and $k-1$. We derive the lower bound on the size of C by showing that there is some variable $x_{v,i}$ that appears in C. By the minimality of H there is some truth assignment α corresponding to a total coloring of G such that $E_k(H - v) \rightarrow C$ is false at α; i.e. α satisfies $E_k(H - v)$ but not C. Since α satisfies $E_k(H - v)$, it is proper with respect to all the edges of $H - v$. Since the degree of v is at most $k - 1$, we can extend this proper coloring on $H - v$ to one that is proper on all of H by changing the color of vertex v to get a new assignment α'. Now $E_k(H)$ is satisfied by α' and since H implies C, C is satisfied by α'.

Therefore $C(\alpha) \neq C(\alpha')$ and since α and α' differ only on the assignments to two variables, $x_{v,i}$ and $x_{v,i'}$ where $i \neq i'$ are the old and new colors of v, C must contain one of them. □

Definition: Given a clause C over the variables from $\chi(G, k)$, denote $\mu_{G,k}(C)$ to be the minimal number of vertices in an induced subgraph $H < G$ that implies C.

Clearly $\mu_{G,k}$ has two key properties:

Lemma 3. *Let G be a graph and $k \geq 2$ be an integer.*

(a) $\mu_{G,k}(\Lambda)$ is the number of vertices in the smallest k-uncolorable sub-graph of G.

(b) If D is a resolvent of B and C then $\mu_{G,k}(D) \leq \mu_{G,k}(B) + \mu_{G,k}(C)$.

Definition: Let $s + 1$ be the minimum number of vertices in a sub-graph of G that is not k-colorable. The *sub-critical k-expansion*, $e_k(G)$, of a graph G is defined to be the maximum over all t, $2 \leq t \leq s$, of the minimum k-boundary size of any induced subgraph H of G that has no isolated vertices and has between $t/2$ and t vertices.

Lemma 4. *For $k \geq 3$, any resolution refutation of $\chi(G, k)$ must contain a clause of width at least $e_k(G)$.*

Proof. Let $s \geq k \geq 3$ and $t \leq s$ be chosen as in the definition of $e_k(G)$. Let π be a resolution refutation of $\chi(G, k)$. By Lemma 3(a), $\mu_{G,k}(\Lambda) = s + 1$. Further, any clause C of $\chi(G, k)$ has $\mu_{G,k}(C) \leq 2$. Therefore there is a clause D in π such that $\mu_{G,k}(D) > t \geq 2$ and no ancestor of D has $\mu_{G,k}$ greater than t. Since $\mu_{G,k}(D) > 2$ there must be two parent clauses B and C in π such that D is the resolvent of B and C. By Lemma 3(b), at least one of these clauses, say B, must have $\mu_{G,k}$ between $t/2$ and t. If $H < G$ witnesses the value of $\mu_{G,k}(B)$ then by Lemma 2, H has no isolated vertices and $w(B) \geq |\beta_k(H)|$. Thus, by definition of $e_k(G)$, $w(B) \geq e_k(G)$ as required. □

Corollary 1. *If G is a graph and $k \geq 3$ is an integer then for $c = 1/(9 \ln 2)$,*

$$Res(\chi(G, k)) \geq 2^{c(e_k(G)-k)^2/n}$$

and

$$DPLL(\chi(G, k)) \geq 2^{e_k(G)-k}.$$

Proof. Clearly $w(\chi(G,k)) = k$ and Lemma 4 implies that $w^*(\chi(G,k)) \geq e_k(G)$. Applying Proposition 1 yields the claimed results. □

We obtain a similar result for tree-like proofs in McDiarmid's proof system using Lemma 1.

Corollary 2. *If G is a graph and $k \geq 3$ is an integer then any tree-like proof of non-k-colorability of G in McDiarmid's proof system requires size at least $2^{(e_k(G)-k(k+1))/2k}$.*

3.2 Lower bounding sub-critical k-expansion

We now prove lower bounds on $e_k(G)$ for most $G \sim \hat{\mathcal{G}}_{n,m}$. We show these bounds by first showing that such a G is almost certainly locally sparse.

We say that a graph G is (r,q)-*dense* if some subset of r vertices of G contains at least q edges of G.

Lemma 5. *Let $G \sim \hat{\mathcal{G}}_{n,m}$. For $r, q \geq 1$,*

$$\Pr[G \text{ is } (r,q)\text{-dense}] \leq \left(\frac{ne}{r}\right)^r \left(\frac{emr^2}{qn^2}\right)^q.$$

Proof. Let R be a set of vertices with $|R| = r$. Let $p = \binom{r}{2}/\binom{n}{2} \leq (r/n)^2$ denote the probability that a randomly chosen edge on n vertices is contained in R. For $G \sim \hat{\mathcal{G}}_{n,m}$, the number of edges of G contained in R has the binomial distribution, $\mathrm{Bin}(m,p)$. The probability that at least q edges of G are contained in R is bounded above by:

$$\Pr[\mathrm{Bin}(m,p) \geq q] \leq \binom{m}{q} p^q \leq \left(\frac{emr^2}{qn^2}\right)^q, \tag{1}$$

Summing this over the $\binom{n}{r} \leq (en/r)^r$ r-subsets of the set of vertices of G we obtain

$$\Pr[G \text{ is } (r,q)\text{-dense}] \leq \left(\frac{ne}{r}\right)^r \left(\frac{emr^2}{qn^2}\right)^q.$$

□

Lemma 6. *For each integer $k \geq 3$ there is a constant C_k such that the following holds. Let m, n be integers with $m = \Delta n$. If $s = C_k n/\Delta^{k/(k-2)}$ then the probability that $G \sim \hat{\mathcal{G}}_{n,m}$ contains a subgraph of size at most s that is not k-colorable is $o(1)$ in s.*

Proof. The probability that G contains a k-uncolorable subgraph of size at most s is the probability that there is some *minimally* k-uncolorable graph $H < G$ with $r \leq s$ vertices. Observe that such an H must have $r \geq k+1$ and have minimum degree at least k since a vertex of degree at most $k-1$ can always be colored by one of the k colors so that none of its incident edges is mono-chromatic. In particular, this implies that H must have average degree at least k and thus contain at least $kr/2$ edges.

Thus, the probability that G contains such a subgraph H is at most

$$\sum_{r=k+1}^{s} \Pr[G \text{ is } (r, kr/2)\text{-dense}].$$

By Lemma 5 we have $\Pr[G \text{ is } (r, kr/2)\text{-dense}] \leq D(r)$ where

$$\begin{aligned} D(r) &= \left(\frac{ne}{r}\right)^r \left(\frac{2emr^2}{krn^2}\right)^{kr/2} \\ &= \left(ne(2em/kn^2)^{k/2} r^{(k-2)/2}\right)^r \\ &= (Q(k,m,n) r^{(k-2)/2})^r \end{aligned}$$

for $Q(k,m,n) = ne(2em/kn^2)^{k/2}$. Now

$$\begin{aligned} \frac{D(r+1)}{D(r)} &= \frac{(Q(k,m,n)(r+1)^{(k-2)/2})^{r+1}}{(Q(k,m,n)r^{(k-2)/2})^r} \\ &= Q(k,m,n)(r+1)^{(k-2)/2} \left(\frac{r+1}{r}\right)^{r(k-2)/2} \\ &\leq Q(k,m,n)(e(r+1))^{(k-2)/2} \\ &= ne(2em/kn^2)^{k/2}(e(r+1))^{(k-2)/2} \\ &= (2e^2 m/kn)^{k/2}((r+1)/n)^{(k-2)/2} \\ &= (2e^2\Delta/k)^{k/2}((r+1)/n)^{(k-2)/2} \\ &\leq 1/2 \end{aligned}$$

for $1 \leq r \leq C_k n/\Delta^{k/(k-2)}$, where $C_k > 0$ depends only on k. Let $s = C_k n/\Delta^{k/(k-2)}$. Therefore the probability that G contains such an k-uncolorable subgraph is a geometric series in r and is at most twice its largest term which is less than

$$D(1) = ne(2em/kn^2)^{k/2} = e(2e\Delta/k)^{k/2}/n^{(k-2)/2} = c_k(1/s)^{(k-2)/2}$$

for some constant c_k. Thus it is $o(1)$ in s as required. □

Lemma 7. *For each integer $k \geq 3$ and ϵ with $1 - 1/(k-1) > \epsilon > 0$, there is a constant $c_{\epsilon,k} > 0$ such that the following holds. Let m, n be integers with $m = \Delta n$. If $t \leq c_{\epsilon,k} n/\Delta^{\frac{k-(k-1)\epsilon}{k-(k-1)\epsilon-2}}$ then the probability that $G \sim \hat{\mathcal{G}}_{n,m}$ contains a subgraph on r vertices, $t/2 < r \leq t$, that has no isolated vertices and at most ϵr vertices of degree $< k$ is $o(1)$ in t.*

Proof. Fix $k \geq 3$, $\epsilon > 0$ and m, n with $m = \Delta n$ and $G \sim \hat{\mathcal{G}}_{n,m}$. If there is a subgraph $H < G$ on r vertices that has no isolated vertices and at most ϵr vertices of degree $< k$, then H has at least $r - \epsilon r$ vertices of degree at least k and the remaining vertices of degree at least 1. Therefore H contains at least $[k(r - \epsilon r) + \epsilon r]/2 = r(k - (k-1)\epsilon)/2$ edges. Thus by Lemma 5 the probability that such an H exists with $t/2 \leq r \leq t$ is at most

$$\begin{aligned} &\sum_{r=t/2}^{t} \Pr[G \text{ is } (r, r(k-(k-1)\epsilon)/2)\text{-dense}] \\ &\leq \sum_{r=t/2}^{t} \left(\frac{ne}{r}\right)^r \left(\frac{2emr^2}{r(k-(k-1)\epsilon)n^2}\right)^{r(k-(k-1)\epsilon)/2} \\ &< \sum_{r=t/2}^{t} \left(\left(\frac{2e^2 m}{(k-(k-1)\epsilon)n}\right)^{(k-(k-1)\epsilon)/2} \left(\frac{r}{n}\right)^{(k-(k-1)\epsilon-2)/2}\right)^r \end{aligned}$$

$$= \sum_{r=t/2}^{t} \left(\left(\frac{2e^2\Delta}{(k-(k-1)\epsilon)} \right)^{(k-(k-1)\epsilon)/2} \left(\frac{r}{n}\right)^{(k-(k-1)\epsilon-2)/2} \right)^r$$

For $t/2 \leq r \leq t$, and for some constant $c_{\epsilon,k} > 0$, if $t \leq c_{\epsilon,k} n/\Delta^{\frac{k-(k-1)\epsilon}{k-(k-1)\epsilon-2}}$, each term in the sum is at most 2^{-r} and thus the sum is less than $2^{1-t/2}$ which is $o(1)$ in t. □

Corollary 3. *For each integer $k \geq 3$ and ϵ with $1 - 1/(k-1) > \epsilon > 0$, there is a constant $c'_{\epsilon,k} > 0$ such that the following holds. Let m, n be integers with $m = \Delta n$. Let $W = n/\Delta^{\frac{k-(k-1)\epsilon}{k-(k-1)\epsilon-2}}$. The probability that $G \sim \hat{\mathcal{G}}_{n,m}$ has $e_k(G) < c'_{\epsilon,k} W$ is $o(1)$ in W.*

Proof. Let $c_{\epsilon,k} > 0$ be the constant from Lemma 7 and C_k be the constant from Lemma 6. Let $t = \min(C_k, c_{\epsilon,k})W$. By Lemma 6, if $G \sim \hat{\mathcal{G}}_{n,m}$ and $s = C_k n/\Delta^{k/(k-2)} \geq t$ then the probability that a subgraph $H < G$ of size at most s is not k-colorable is $o(1)$ in s and thus $o(1)$ in W. Also, by Lemma 7 the probability that $G \sim \hat{\mathcal{G}}_{n,m}$ contains a subgraph on r vertices, $t/2 < r \leq t$, that has no isolated vertices and at most ϵr vertices of degree $< k$ is $o(1)$ in t, and thus $o(1)$ in W. Thus every induced subgraph H on r vertices with no isolated vertices and $t/2 < r \leq t \leq s$ has k-boundary of size at least $\epsilon r \geq \epsilon t/2 = \epsilon \min(C_k, c_{\epsilon,k})W/2$. Letting $c'_{\epsilon,k} = \epsilon \min(C_k, c_{\epsilon,k})/2$ yields the lower bound on $e_k(G)$. □

3.3 Lower bound theorems

Theorem 1. *For each integer $k \geq 3$ and $\epsilon > 0$ there are constants $C_{\epsilon,k}, C'_{\epsilon,k} > 0$ such that if m,n are integers with $m = \Delta n$ and $G \sim \hat{\mathcal{G}}_{n,m}$, then with probability $1-o(1)$ in n, $Res(\chi(G,k)) \geq \exp(C_{\epsilon,k} n/\Delta^{2+\frac{4}{k-2}+\epsilon})$ and $DPLL(\chi(G,k)) \geq \exp(C'_{\epsilon,k} n/\Delta^{1+\frac{2}{k-2}+\epsilon})$.*

Proof. Let $\epsilon' = \epsilon(k-2)^2/[(k-1)(4+\epsilon(k-2))]$. Clearly $0 < \epsilon' < (k-2)/(k-1) = 1 - 1/(k-1)$. By Corollary 1 the resolution complexity of $\chi(G,k)$ is at least $2^{(e_k(G)-k)^2/n}$. By Corollary 3, there is a $c'_{\epsilon',k}$ such that with probability $1 - o(1)$ in $W = n/\Delta^{(k-(k-1)\epsilon')/(k-(k-1)\epsilon'-2)}$, $e_k(G) \geq c'_{\epsilon',k} W$. Therefore

$$\begin{aligned} e_k^2(G)/n &\geq (c'_{\epsilon',k})^2 W^2/n \\ &= (c'_{\epsilon',k})^2 n/\Delta^{2(1+2/(k-(k-1)\epsilon'-2))} \\ &= (c'_{\epsilon',k})^2 n/\Delta^{2+\frac{4}{k-2}+\epsilon} \end{aligned}$$

by our choice of ϵ'. Now, if this quantity is at least 1, then clearly W is $\omega(1)$ in n since W is larger than it by a factor of $c'_{\epsilon',k}\Delta^{1+2/(k-(k-1)\epsilon'-2)}$. Therefore with probability $1 - o(1)$ in n, $e_k(G)^2/n \geq (c'_{\epsilon',k})^2 n/\Delta^{2+\frac{4}{k-2}+\epsilon}$. Clearly, we can choose $C_{\epsilon,k}$ and can absorb the $-k$ in the constant $C_{\epsilon,k}$ to obtain the desired result for resolution. For DPLL procedures, the result follows even more directly. □

It is worth noting that we can obtain a lower bound on the size of a proof of non-k-colorability in the system considered by McDiarmid [McD84] that is similar to the DPLL proof size lower bound proof size since their bound on the proof width as a function of proof size only differs by a factor of $2k$. Thus the same bound as the DPLL bound above holds for McDiarmid's system with a slightly different value of the constant $C'_{\epsilon,k}$.

Corollary 4. *For each integer $k \geq 3$, $\epsilon > 0$, and $\Delta > 0$ there is a constant $C'_{\epsilon,k} > 0$ such that*

- *if* $p = 2\Delta/(n-1)$ *and* $G \sim \mathcal{G}_{n,p}$*, then with probability* $1-o(1)$ *in* n*,* $Res(\chi(G,k)) \geq \exp(C'_{\epsilon,k} n/\Delta^{2+\frac{4}{k-2}+\epsilon})$.
- *if* $m = \Delta n$ *and* $G \sim \mathcal{G}_{n,m}$*, then with probability* $1-o(1)$ *in* n*,* $Res(\chi(G,k)) \geq \exp(C'_{\epsilon,k} n/\Delta^{2+\frac{4}{k-2}+\epsilon})$.

Corollary 5. *For each integer* $k \geq 3$*,* $\epsilon > 0$*, there is a* $\delta > 0$ *such that if* $m \leq n^{\frac{3k-2}{2k}-\epsilon}$ *and* $G \sim \mathcal{G}_{n,m}$ *or* $G \sim \hat{\mathcal{G}}_{n,m}$ *then, with probability* $1 - o(1)$ *in* n*,* $Res(\chi(G,k)) \geq 2^{n^\delta}$.

Proof. For this range of m, $\Delta \leq n^{\frac{k-2}{2k}-\epsilon}$. Applying the Theorem 1 with a suitably small value of ϵ' in place of ϵ yields the desired result. □

Corollary 6. *For each integer* $k \geq 3$*,* $\epsilon > 0$*, there is a* $\delta > 0$ *such that if* $m \leq n^{2-\frac{2}{k}-\epsilon}$ *and* $G \sim \mathcal{G}_{n,m}$ *or* $G \sim \hat{\mathcal{G}}_{n,m}$ *then, with probability* $1 - o(1)$ *in* n*,* $DPLL(\chi(G,k)) \geq 2^{n^\delta}$.

4 Upper Bounds

A very simple brute-force procedure achieves a $2^{O(n/\Delta)}$ upper bound for proving non-k-colorability, based on the following observation.

Lemma 8. *Let* $p = 2\Delta/(n-1)$ *and* $G \sim \mathcal{G}(n,p)$*. Let* G_r *be the subgraph of* G *induced on the first* r *vertices of* G*. If* $r - 1 > c_k^+(n-1)/\Delta$ *then with probability* $1 - o(1)$ *in* r*,* G_r *is not* k*-colorable.*

Proof. Clearly, if $G \sim \mathcal{G}(m,p)$ then the induced graph $G_r \sim \mathcal{G}(r,p)$. By our assumption on r, $p(r-1)/2 > c_k^+$. The conclusion follows by the definition of c_k^+. □

Therefore the simple algorithm that searches through all possible k-colorings of the first $c_k^+(n-1)/\Delta+1$ vertices of a $G \sim \mathcal{G}(n,p)$ will almost certainly find a witness to the non-k-colorability of G. Such an algorithm can easily be phrased as a simple DPLL search procedure, called the *ordered DPLL procedure* in [BKPS02], that always splits on the first unset variable. In fact, [Wil84, BW85] analyze essentially the same simple backtracking procedure for k-coloring, although it is not described as being based on DPLL. (In their algorithm, vertices are listed in a fixed order and all colors of a vertex compatible with previously assigned vertices are tried recursively.) For this algorithm, with $p = 2\Delta/(n-1)$, $\Delta \in o(n)$ and $G \sim \mathcal{G}(n,p)$, they show that the log of the expected number of nodes in the search tree is $k(\log k)^2 n/4\Delta$ plus lower order terms for $pn \to \infty$. For completeness, we state and prove the following simpler version of such a theorem.

Theorem 2. *Let* $k > 1$*,* $p = 2\Delta/(n-1)$ *and* $G \sim \mathcal{G}(n,p)$*. With probability* $1 - o(1)$ *in* n*, the ordered DPLL procedure where all* k *variables associated with each vertex are numbered consecutively witnesses the fact that* $Res(\chi(G,k)) \leq 2^{O(k \log^2 k\, n/\Delta)} n^{O(1)}$.

Proof. By the bound of Łuczak, $c_k^+ = k \ln k + O(k \ln \ln k)$. Let $r = \max\{\log n, c_k^+(n-1)/\Delta + 1\}$. Apply Lemma 8 to say that with probability $1 - o(1)$ in r and thus $1 - o(1)$ in n since $n \leq 2^r$, the induced graph G_r on the first r vertices of G is not k-colorable.

The ordered DPLL procedure will have a branch for each of k^r different k-colorings of the first r vertices of G. Although there are k Boolean variables associated with each vertex, it is easy to see that the height k tree corresponding to the branches on these k variables has only k non-trivial children. Therefore the ordered DPLL tree has size at most proportional to k^r, the number of k-colorings. Plugging in the value of r yields the desired result. □

More generally, given any upper bound r' on the number of vertices so that random graphs with density Δ almost certainly have a minimal k-uncolorable subgraph of size at most r', one obtains a naive $2^{O(r' \log n)}$ algorithm that does a brute force search for such a subgraph. Such an algorithm can easily be phrased as a resolution proof of non-k-colorability. Lemma 8 simply shows that $r' = O(n/\Delta)$.

We can, however, do much better by using a more careful, if somewhat articifial, splitting rule and which we show w.h.p. proves that $G \sim \mathcal{G}(n, \Delta/n)$ is non-k-colorable in time $\exp(O(n/\Delta^{\alpha_k}))$ where

$$\alpha_k = \frac{k-1}{k-2} \ . \tag{2}$$

Interestingly, this value of α_k coincides with a heuristic bound suggested by calculations of Ein-Dor and Monasson [EDM03] who estimated the running time of DPLL on $G \sim \mathcal{G}(n, \Delta/n)$ by making the (clearly false) assumption that the branches of the search tree are independent.

The splitting rule depends on the current level in the search tree, i.e., the number of vertices that are currently colored. We go through a series of stages; to test G for k-colorability, the algorithm starts in stage k, switches to stage $k-1$ at a certain depth, and so on. The idea is that at stage j, the algorithm is currently checking a subgraph of G for j-colorability, where the vertices in this subgraph are neighbors of vertices colored at previous stages and have at most j allowed colors. This continues until we reach stage 2, at which point we check a subgraph for 2-colorability in polynomial time using unit propagation around an odd cycle. At each stage we use an arbitrary numbering, as in ordered DPLL, to choose among the vertices in the current subgraph. We illustrate this in Figure 1.

We now define this splitting rule, which we call R. To introduce some notation, let T be the current depth, and for $1 \le t \le T$ let v_t be the vertex that was colored at level t of the tree and $c(v_t)$ be its assigned color. For $3 \le i \le k$ let $t_i = a_i n/\Delta^{\alpha_i}$, where α_i is given by (2) and $a_3, \ldots, a_k$ is a set of parameters we will define below. Let $T_j = \sum_{i=j}^{k} t_i$ with $T_{k+1} = 0$, and let $T_2 = n$; then we will say that we are currently in *stage* j if j is the largest integer such that $T < T_j$.

Given that we are in stage j, define a set of vertices V_j as follows. For $j < i \le k$, let $S_i = \{v_t : T_{i+1} \le t < T_i\}$ be the set of vertices colored during stage i, let c_i be the most common color among the vertices in S_i, and let $U_i = \{v \in S_i : c(v) = c_i\}$. Then let V_j be the set of uncolored vertices v such that v has a neighbor $u_i \in U_i$ for all $j < i \le k$. By construction the c_i are all distinct, so vertices in V_j have at most j allowed colors. We take $V_k = V$.

Finally, if $j > 2$, R splits on the vertex in V_j of smallest number. If $j = 2$, R splits on the vertex in V_2 of smallest number, and performs unit propagation whenever a vertex in V_2 exists with only one allowed color. (Note that although it does not take advantage of unit propagation until $j = 2$, up to a polynomial factor, the

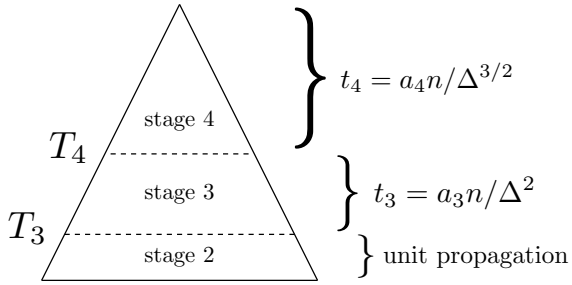

Figure 1: Our depth-dependent splitting rule R is equivalent to a series of stages; during stage j, the algorithm is checking a subgraph V_j for j-colorability. Here $k = 4$.

Algorithm $\mathtt{A}_k(G)$ {
    If $k = 2$, check G for 2-colorability.
    Else {
        Let $t = a_k n/\Delta^{\alpha_k}$.
        If $|G| < t$, fail.
        Else, for all k^t assignments of colors to the t vertices of lowest number, do {
            Let c be the most common color among these t vertices,
            let U be the set of these vertices assigned color c, and
            let G' be the subgraph induced by U's uncolored neighbors.
            Run $\mathtt{A}_{k-1}(G')$.
        }
    }
}

Table 1: A recursive algorithm equivalent to the splitting rule R.

algorithm's running time can only improve if it performs additional unit propagations at an earlier stage.)

To do our analysis, it is convenient to recognize that DPLL with this splitting rule is equivalent to a recursive algorithm $\mathtt{A}_k(G)$ which colors $t = t_k$ vertices and then calls itself on the subgraph G' of uncolored neighbors of the vertices assigned the most common color. We illustrate this algorithm in Table 1.

Theorem 3. *Let $G \sim \mathcal{G}(n, p = \Delta/n)$. For all $k \geq 3$, there exist constants $a_3, \ldots a_k$ and constants b_k, $c_k > 0$, and d_k such that, for all $\Delta > d_k$, DPLL with splitting rule* R*, or equivalently algorithm $\mathtt{A}_k$, refutes $\chi(G, k)$ in time $\exp(b_k n/\Delta^{\alpha_k})\, n^{O(1)}$ with probability at least $1 - \exp(-c_k n/\Delta^{\alpha_k})$, where $\alpha_k = (k-1)/(k-2)$. Therefore we have*

$$DPLL(\chi(G, k)) \leq \exp(b_k n/\Delta^{\alpha_k})\, n^{O(1)}$$

with probability $1 - o(1)$ in n.

Our proof of Theorem 3 is similar to that of Theorem 6.1 of [BKPS02], which establishes an upper bound for DPLL on random k-SAT. The idea there is that setting a certain number of variables creates a large enough density of 2-variable clauses so that an unsatisfiable 2-SAT subformula appears with high probability, which ordered DPLL with unit propagation proves is unsatisfiable in linear time. Similarly, we will show inductively that coloring the first t_k vertices yields a subgraph which $\mathtt{A}_{k-1}$ quickly proves is non-$(k-1)$-colorable, until $k = 2$ and we can check for 2-colorability quickly using our the smallest numbered vertex splitting rule and unit propagation. This induction works if the α_k obey a certain recurrence, yielding (2).

We first prove that if G has some simple properties, which hold with high probability whenever G has sufficiently large degree, then checking for non-2-colorability is extremely fast. (Ordered DPLL is particularly naive because it does not take advantange of the symmetry of the colors. If we used a more sophisticated splitting rule then we could easily derive a linear upper bound for all non-2-colorable graphs.)

Lemma 9. *There is a constant c_0 such that if a graph G has a non-2-colorable connected subgraph on at least $7/8$ of its vertices, then under a random numbering of the vertices of G, the expected time for ordered DPLL to refute $\chi(G, 2)$ is at most $c_0 n$.*

Proof. Observe that once ordered DPLL colors one vertex of a connected component of G, all other vertices of G receive the implied colors by unit propagation. Thus as soon as ordered DPLL chooses some vertex

from the large non-2-colorable connected component it reaches a contradiction and backtracks, tries the other color value, reaches a second contradiction and backtracks again, failing on that component. If such a vertex has the smallest number then a contradiction is reached immediately. However, if a vertex of some 2-colorable component is encountered previously then on backtracking past this component, ordered DPLL will recolor the first vertex of this 2-colorable component and then return to the non-2-colorable large subgraph and again determine failure before finally deriving a contradiction. Thus if the smallest-numbered vertex of the large non-2-colorable component is numbered $b+1$ then the number of times we color first vertices in some component is at most $r = 2^{b+1}$ and the cost of unit propagation on each component is Cn for some constant C. Since the large component has size at least $7n/8$ the probability that this happens is at most $8^{-b} = (2/r)^3$. The expected time is then at most $C\sum_{r=1}^{\infty}(2/r)^3 rn = (8C\sum_{r=1}^{\infty}1/r^2)n = (4C\pi^2/3)n$. □

We now prove that a random graph of sufficiently high, but constant, average degree is overwhelmingly likely to have the large non-2-colorable component called for by Lemma 9, so that our very simple DPLL procedure will determine that it is not 2-colorable in linear expected time. Note that, here and elsewhere, we have made no attempt to optimize the constants (other than α_k) that appear in the exponents.

Lemma 10. *There exist constants $c_2, d_2 > 0$ such that, for all $\Delta \geq d_2$, n sufficiently large and $G \sim \mathcal{G}(n, p = \Delta/n)$,*

$$\Pr[G \text{ contains a non-2-colorable connected component of size} \geq 7n/8] \geq 1 - \mathrm{e}^{-c_2\Delta n} \ .$$

Proof. First consider the probability that the largest connected component of G is of size less than $7n/8$. If this is the case then there must be some subset S of vertices with $n/16 \leq |S| < 15n/16$ that is disconnected from $\bar{S}$. (To see this, consider the components $C_1, \ldots, C_l$ of G in an arbitrary order, let j be the smallest integer such that $|C_1 \cup \cdots \cup C_j| \geq n/16$ and set $S = C_1 \cup \cdots \cup C_j$. Since $|C_j| \leq 7n/8$, $n/16 \leq |S| < 15n/16$.) For a fixed set S of this size the probability that there are no edges from S to $\bar{S}$ is at most

$$(1-p)^{|S|\cdot|\bar{S}|} \leq \mathrm{e}^{-p(15n^2/256)} \leq \mathrm{e}^{-15\Delta n/256}$$

Since there fewer than 2^n such sets S,

$$\Pr[G \text{ does not contain a connected component of size} \geq 7n/8] \leq \mathrm{e}^{-(15\Delta/256-\ln 2)n} \ .$$

We now show that with overwhelming probability, no subset of $m \geq n/2$ vertices of G can be bipartite. Fix one such subset B. Let X be the number of 2-colorings of B. There are 2^m 2-color assignments to the vertices of B. For any such assignment, the number of potential edges between vertices of the same color is at least $2\binom{\lfloor m/2 \rfloor}{2}$, which is greater than $m^2/8$ for $n \geq 8$. Therefore,

$$\mathbf{E}[X] \leq 2^m(1-p)^{m^2/8} \leq 2^m \mathrm{e}^{-pm^2/8} = 2^{n/2}\mathrm{e}^{-pn^2/32} \leq \mathrm{e}^{-(\Delta/32-(1/2)\ln 2)n}$$

Since there are fewer than 2^n such sets, the probability that some such set is bipartite is at most $\mathrm{e}^{-(\Delta/32-(3/2)\ln 2)n}$.

Therefore the overall probability that a non-2-colorable component of size at least $(7/8)n$ fails to exist is at most

$$\mathrm{e}^{-(15\Delta/256-\ln 2)n} + \mathrm{e}^{-(\Delta/32-(3/2)\ln 2)n}$$

which is at most $\mathrm{e}^{-c_2\Delta n}$ for $\Delta \geq d_2$ for some constants $c_2, d_2 > 0$. □

Below we will use the following Chernoff bounds on the binomial distribution [ASE92, Appendix A]:

$$\begin{aligned} \Pr[\mathrm{Bin}(m,q) < mq/4] &\leq 2^{-mq/2} && (3) \\ \Pr[\mathrm{Bin}(m,q) > Cmq] &\leq (C/\mathrm{e})^{-Cmq} && (4) \end{aligned}$$

Now, to illustrate the idea and start our induction, we prove Theorem 3 in the case $k = 3$.

Theorem 4. *Let $G \sim \mathcal{G}(n, p = \Delta/n)$. There exist constants a_3, b_3, $c_3 > 0$, and d_3 such that if $\Delta \geq d_3$, then $\mathtt{A}_3$ refutes $\chi(G, 3)$ in time at most $\exp(b_3 n/\Delta^2)\, O(n)$ with probability at least $1 - \exp(-c_3 n/\Delta^2)$.*

Proof. As in our definitions of $\mathtt{R}$ and $\mathtt{A}_3$ above, let G' be the subgraph induced by the uncolored neighbors of the t colored vertices which have been assigned their most common color. We will show that w.h.p. G' satisfies the conditions of Lemma 9 for all 3^t assignments of the first t vertices, then $\mathtt{A}_3$'s expected running time will be at most $3^t c_0 n = \exp((a_3 \ln 3) n/\Delta^2) c_0 n$. Then, setting $b_3 = 2a_3 \ln 3$, by Markov's inequality the probability that the running time exceeds $\exp(b_3 n/\Delta^2) c_0 n$ is at most $\exp(-(a_3 \ln 3) n/\Delta^2)$.

The number n' of vertices in G' is binomially distributed as $\mathrm{Bin}(n - t, q)$ where

$$q \geq 1 - (1-p)^{t/3} \geq 1 - \mathrm{e}^{-pt/3} \geq \frac{pt}{6} = \frac{a_3}{6\Delta}$$

since $\mathrm{e}^{-x} \leq 1 - x/2$ for $0 \leq x \leq 1$. Using the lower Chernoff bound (3) and choosing d_3 large enough so that $t = a_3 n/\Delta^2 < n/2$ for $\Delta \geq d_3$, we have

$$\Pr[n' < a_3 n/(48\Delta)] \leq 2^{-a_3 n/(24\Delta)} \ .$$

Let $a_3 = 48 d_2$ where d_2 is the constant defined by Lemma 10. Then we have

$$\Pr[n' < d_2 n/\Delta] \leq 2^{-2d_2 n/\Delta} \ . \tag{5}$$

Clearly G' is distributed as $\mathcal{G}(n', p)$ if we condition on the value of n'. Let $\Delta' = pn'$ be the mean degree of G'; then if $n' \geq d_2 n/\Delta$ we have $\Delta' \geq d_2$. Combining Lemma 10 with (5) then gives

$$\begin{aligned} \Pr[G' \text{ does not contain a non-2-colorable component of size} \geq 7n'/8] &\leq 2^{-2d_2 n/\Delta} + \mathrm{e}^{-c_2 d_2^2 n/\Delta} \\ &< \exp(-2Cn/\Delta) \end{aligned}$$

for some $C < \min(d_2 \ln 2, c_2 d_2^2/2)$. We wish to bound the probability that G' violates the condition of Lemma 9 for some assignment of the first t vertices. Choose d_3 large enough so that $(a_3 \ln 3)/d_3 < C$; then for all $\Delta \geq d_3$, the union bound gives a total probability of

$$\begin{aligned} 3^t \exp(-2Cn/\Delta) &< \exp\left(\left(\frac{a_3 \ln 3}{\Delta} - 2C\right)\frac{n}{\Delta}\right) \\ &< \exp(-Cn/\Delta) \ . \end{aligned}$$

Combining this with Markov's inequality as described above, the overall probability that the running time of $\mathtt{A}_3$ exceeds $\exp(b_3 n/\Delta^2) c_0 n$ is at most

$$\exp(-(a_3 \ln 3) n/\Delta^2) + \exp(-Cn/\Delta) < \exp(-c_3 n/\Delta^2)$$

for all $\Delta \geq d_3$ and some $c_3 > a_3 \ln 3$. □

We can now prove Theorem 3 for all k.

Proof of Theorem 3. The proof works by induction on k, where each step of the induction parallels that of Theorem 4 and we use $k = 3$ as the base case. Our goal is to set a_k and d_k so that the average degree Δ' of G' is at least d_{k-1}.

First, the number n' of vertices in G' is distributed as $\mathrm{Bin}(n-t, q)$ where

$$q \geq 1 - (1-p)^{t/k} \geq 1 - \mathrm{e}^{-pt/k} \geq \frac{pt}{2k} = \frac{a_k}{2k\Delta^{\alpha_k - 1}}$$

Choosing d_k large enough so that $t < n/2$ for $\Delta \geq d_k$ and setting $a_k = 16kd_{k-1}$, the Chernoff bounds (3), (4) (with $C = 4$) give

$$\Pr[n' < d_{k-1}n/\Delta^{\alpha_k - 1}] \leq 2^{-2d_{k-1}n/\Delta^{\alpha_k - 1}} \tag{6}$$

$$\Pr[n' > 16d_{k-1}n/\Delta^{\alpha_k - 1}] \leq (4/\mathrm{e})^{-16d_{k-1}n/\Delta^{\alpha_k - 1}} \tag{7}$$

Then if $n' \geq d_{k-1}n/\Delta^{\alpha_k - 1}$, we have $\Delta' = pn' \geq d_{k-1}\Delta^{2-\alpha_k} \geq d_{k-1}$, since $\alpha_k \leq 2$ and we assume w.l.o.g. $\Delta > 1$. Therefore, assuming the events of (6) and (7) hold, the running time of $\mathtt{A}_k$ (up to $n^{O(1)}$) is

$$k^t \exp(b_{k-1}n'/(\Delta')^{\alpha_{k-1}}) \leq \exp\left(\left(a_k \ln k + \frac{16b_{k-1}d_{k-1}^{1-\alpha_{k-1}}}{\Delta^{\alpha_{k-1}(2-\alpha_k)-1}}\right)\frac{n}{\Delta^{\alpha_k}}\right) \tag{8}$$

Set $b_k = a_k \ln k + 16b_{k-1}d_{k-1}^{1-\alpha_{k-1}}$. Then if $\alpha_{k-1}(2-\alpha_k) - 1 = 0$, or alternately

$$\alpha_k = 2 - \frac{1}{\alpha_{k-1}} \ , \tag{9}$$

then the running time (8) becomes $\exp(b_k n/\Delta^{\alpha_k})$ as stated. Indeed, the solution to the recurrence (9) with the initial condition $\alpha_3 = 2$ is

$$\alpha_k = \frac{k-1}{k-2} \ .$$

We now bound the probability that, for some assignment of the first t vertices, $\mathtt{A}_{k-1}$ fails because $|G'| < t_{k-1}$ or its running time exceeds $\exp(b_k n/\Delta^{\alpha_k})$. First, note that since $\alpha_k - 1 < \alpha_{k-1}$, for sufficiently large Δ we have $d_{k-1}n/\Delta^{\alpha_k - 1} > t_{k-1}$, so the probability that $|G'| < t_{k-1}$ is bounded by (6). Then, combining (6) and (7) with our inductive assumption gives

$$\begin{aligned}
\Pr[\mathtt{A}_{k-1} \text{ fails or takes too long on } G'] &\leq 2 \cdot 2^{-2d_{k-1}n/\Delta^{\alpha_k - 1}} + \exp\left(-\frac{c_{k-1}n'}{\Delta'^{\alpha_{k-1}}}\right) \\
&\leq 2^{1-2d_{k-1}n/\Delta^{\alpha_k - 1}} + \exp\left(-\frac{c_{k-1}d_{k-1}n}{(16d_{k-1})^{\alpha_{k-1}}\Delta^{\alpha_k}}\right) \qquad (10) \\
&< \exp(-2c_k n/\Delta^{\alpha_k})
\end{aligned}$$

for some $c_k > c_{k-1}d_k/(16d_{k-1})^{\alpha_{k-1}}$, where we have set $d_k > c_{k-1}/((2\ln 2)(16d_{k-1})^{\alpha_{k-1}})$ so that the second term in (11) dominates for all $\Delta \geq d_k$. Furthermore, set c_k large enough so that $c_k > a_k \ln k$; then the union bound over the k^t assignments of the first t vertices gives

$$\begin{aligned}
\Pr[\mathtt{A}_k \text{ fails}] &= k^t \exp(-2c_k n/\Delta^{\alpha_k}) \\
&= \exp((a_k \ln k - 2c_k)n/\Delta^{\alpha_k}) \\
&< \exp(-c_k n/\Delta^{\alpha_k})
\end{aligned}$$

which completes the proof. □

Another approach, more closely analogous to [BKPS02], is to note that if $\Theta(t)$ of the initial vertices are assigned each color, then the number of their 2-color neighbors with a particular pair of allowed colors is w.h.p. $n' = \Theta(n(pt)^{k-2}) = \Theta(n/\Delta^{(\alpha_k-1)(k-2)})$. These 2-color vertices induce a graph G' of average degree $\Delta' = pn' = \Theta(\Delta^{1-(\alpha_k-1)(k-2)})$, and Lemma 10 shows that G' is w.h.p. non-2-colorable when $\Delta' \geq d_2$. This happens (with appropriate constants) when $\Delta' = \Theta(1)$, giving $1-(\alpha_k-1)(k-2) = 0$ and so $\alpha_k = (k-1)/(k-2)$.

However, since the colors assigned to the first t vertices are (negatively) correlated, proving that $\Theta(t)$ of them receive each of the k colors then becomes a separate problem. One possible method for this would be to use multitype branching processes as in [AM02].

5 Discussion of algorithms

5.1 Backtracking algorithms

In section 4, we described the behavior of the backtracking algorithm analyzed in [Wil84, BW85] on non-k-colorable graphs as a resolution refutation of the formula $\chi(G,k)$. This is very typical. Many more sophisticated backtracking-based coloring algorithms can be emulated by resolution (or by a tree-like McDiarmid proofs to which our resolution-width bounds also apply). Our DPLL upper bound from Theorem 3 corresponds to one such natural algorithm which has better behavior. For example, most of the coloring algorithms in [JT96] were of this form.

Beigel and Eppstein also used recursive algorithms, based on extensive case analyses of local configurations, to give upper bounds for 3-coloring of $O(1.3446^n)$ [BE95], later improved to $O(1.3289^n)$ [BE00]. The case analysis is in terms of general constraint satisfaction problems (CSPs) with domain size 3 or 4 and binary constraints, of which 3-coloring is a special case where the domain size is 3 and all constraints are not-equal constraints. We are not able, so far, to give strict resolution simulations for these algorithms. However, in [Mit02] a variant of the the algorithm in [BE95] is given which establishes the same bound of $O(1.3446^n)$ for $Res(\chi(G,3))$. In the case of the algorithm given in [BE00], we can do the same for all but one case (among some two dozen).

The execution of the Beigel–Eppstein algorithms on input G may easily be described in terms of the formula $\chi(G,k)$. Consider the following "quasi-DPLL" algorithm scheme. For formula ϕ, select a set $S \subset vars(\phi)$ of variables based on some local property of ϕ, and consider a set A of (possibly partial) assignments to S. For each α in A, either α makes a clause C of ϕ false, or we make a recursive call to solve the restriction of ϕ by α. If the set A of assignments covers all assignments to S (that is, each assignment to S is an extension of some partial assignment in A), then ϕ is unsatisfiable if and only if each restricted formula is unsatisfiable. Any "quasi-DPLL" algorithm of this sort can be efficiently simulated by resolution as follows. For each α in A we may derive a clause which α makes false, and from these derived clauses plus clauses of ϕ that mention variables in S, we may construct a refutation of ϕ. The derived clauses are obtained as follows. If α makes a clause $C \in \phi$ false, we use the clause C, otherwise the recursive call to refute the restriction of ϕ by α returns a clause that α makes false. In the case where ϕ is already a restriction of the input formula by partial assignment τ, we can extract from the refutation of ϕ a clause which τ makes false, and return this clause to the parent invocation. (The refutations produced may not be strictly tree-like, but are nearly tree-like in that non-tree parts are local.)

Most of the cases in the Beigel–Eppstein algorithms are captured fairly easily by this scheme. The bounds are obtained by careful choices for sets S and A, and the corresponding number and sizes of recursive calls. However, in a small number of cases the set A of assignments does not cover all assignments to S.

In these cases, the completeness arguments do not have direct analogs in resolution, and to obtain the bounds in [Mit02] alternate handling or sub-case breakdown was used. In the algorithm of [BE00] one case makes use of polytime bipartite matching, which is a more serious impediment to simulation by resolution. Nonetheless, it is plausible that our lower bounds can be shown to apply to both of the Beigel–Eppstein algorithms.

The connection between k-coloring algorithms and resolution-based algorithms for $\chi(G, k)$ is also borne out in practice. For example, the program used for coloring in [CG01] also uses graph reductions based on removing vertices, or merging vertices. These, too, can be emulated by resolutions, mostly on 2 clauses, but occasionally involving 3-clauses. It also does significant forward pruning by propagating statements of the type "the color of v is i or the color of w is j". These can be expressed as 2-clauses $(x_{v,i} \vee x_{w,j})$ which can be generated by resolution. In extensive head-to-head tests [CG01] of this program on graphs and the back-jumping version of the DPLL-based Tableau program **ntab-back** [CA96] on the associated $\chi(G, k)$ behave statistically alike on random 3-color instances. The coloring program is better for larger values of k, but this is likely due to the fact that Tableau was tuned to 3-SAT instances and the fact that some features (e.g. clustering) are easier to identify knowing the graph structure.

5.2 Other complete algorithms

There are many other algorithms suitable for k-coloring or k-colorability that are not primarily based on backtracking search and therefore are not covered by our resolution bounds. Often they are tuned for use on the random graphs we consider. However most non-backtracking coloring algorithms, such as those in [BMWZ, AK97, FF96, CS02] are *incomplete* in that they may miss some possible k-colorings and therefore do not produce certificates of non-k-colorability. We will only be concerned with complete algorithms.

One class is particularly interesting. At sufficiently large density, any graph is almost certainly not k-colorable, however this fact does not in itself provide a proof of such a graph's non-k-colorability. However, recently, Krivelevich and Vu [KV02] showed that tight concentration bounds on the polynomially-computable Lovasz ϑ function of random graphs can be used to approximate the chromatic number in polynomial expected time for sufficiently large graph density Δ. Coja-Oghlan [CO03] has used the same general technique and sharper concentration bounds for the polynomial-time computable *vector chromatic number* $\bar{\vartheta}_{1/2}(G)$ to derive a polynomial expected time k-colorability algorithm for $k = o(\sqrt{n})$ and $\Delta > ck^2$ for some constant $c > 0$. More precisely, since $\bar{\vartheta}_{1/2}(G) \leq \vartheta(\bar{G}) \leq \chi(G)$ for all graphs G, the algorithm first tests if $\bar{\vartheta}_{1/2}(G) > k$ (using semi-definite programming); if so, then this provides a proof of the non-k-colorability of G. Otherwise, the algorithm then calls a standard worst-case exponential-time k-coloring algorithm such as that of Lawler [Law76]. The concentration results for $\bar{\vartheta}(G)$ are such that it is exponentially unlikely that the algorithm will need to resort to the second stage and thus the algorithm has polynomial expected running time.

Clearly, based on our results, these algorithms are provably superior to resolution and backtracking for proving non-k-colorability of random graphs when the graph density is sufficiently large compared to k. These algorithms use the typical properties of a random input to try to quickly produce a certificate that the input is not k-colorable. Although such a certificate may not be guaranteed to exist, it works sufficiently frequently that it is useful. Similar ideas using spectral methods have also been employed by Goerdt et al. [GK01, FG01] for random k-SAT. It is interesting to note that in the case of k-SAT, although the best current algorithms almost certainly yield efficient certificates in a wider range of densities than does DPLL, in contrast to the situation for k-colorability it is still open whether or not they do the same when compared with general resolution,

Finally, we note that since $\chi(G,k)$ is a k-CNF formula, any complete deterministic algorithm for k-SAT is potentially relevant for k-coloring. Most such algorithms that have been analyzed are themselves resolution-based although there are exceptions such as the $(2-2/(k+1))^n$ worst-case time algorithm of Dantsin et al. [DGH$^+$02] which is not competitive and does improve with improved graph density.

6 Open Problems

There is a gap between the exponents of Δ in our upper and lower bounds that would be nice to close, particularly for DPLL. Our DPLL upper bound is of the form $\exp(O(n/\Delta^{(k-1)/(k-2)}))$ whereas our lower bound is of the form $\exp(\Omega(n/\Delta^{k/(k-2)+\epsilon}))$.

Classic results [ER60] show that for $\Delta = \omega(n^{(k-2)/k})$ a random $G \sim \mathcal{G}(n, \Delta/n)$ contains a $(k+1)$-clique with probability $1-o(1)$. The presence of such a clique obviously yields a $O(k^2)$ size DPLL proof of non-k-colorability so we should consider densities below this bound. It is interesting that this is essentially the same range where our current lower bound yields nontrivial results. This suggests that our DPLL lower bound (except maybe for the ϵ) is the more likely to be the correct bound.

One obstacle for improving the DPLL upper bound to match the lower bound may be that the lower bound allows an optimal literal selection rule that may not necessarily be obtainable via an efficient DPLL algorithm. It would be interesting to see if one could obtain improved lower bounds for simple literal selection rules or show that simpler selection rules can achieve the same or better upper bounds.

Our bounds apply only above the threshold for k-colorability, but the use of $\chi(G,k)$ in k-coloring algorithms also is suitable in the k-colorable region since satisfying assignments correspond to k-colorings (unlike the ϑ-based algorithms from section 5 which yield no information about a k-coloring when the graph is k-colorable). Jia and Moore [JM] have recently shown exponential behavior of a natural greedy DPLL algorithm on easily k-colorable random graphs below the threshold: this greedy algorithm repeatedly misses constant-size contradictory subproblems. However, such behavior can easily be eliminated with a more sophisticated literal selection rule. Are there other more sophisticated algorithms where such bounds can be shown for satisfiable instances?

Acknowledgements

We would like to thank Dimitris Achlioptas, Jared Saia, and Vishal Sanwalani for helpful references and discussions.

References

[AF99] D. Achlioptas and E. Friedgut. A sharp threshold for k-colorability. *Random Structures and Algorithms*, 14:63–70, 1999.

[AK97] N. Alon and N. Kahale. A spectral technique for coloring random 3-colorable graphs. *SIAM Journal on Computing*, 26(6):1733–1754, 1997.

[AM97] D. Achlioptas and M. Molloy. The analysis of a list-coloring algorithm on a random graph. In *Proceedings 38th Annual Symposium on Foundations of Computer Science*, pages 204–212. IEEE, October 1997.

[AM99] D. Achlioptas and M. Molloy. Almost all graphs with 2.522n edges are not 3-colorable. *Electronic Journal of Combinatorics*, 6(1):R29, 1999.

[AM02] D. Achlioptas and C. Moore. Almost all graphs with average degree 4 are 3-colorable. In *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing*, pages 199–208, Montreal, Quebec, Canada, May 2002.

[AN] D. Achlioptas and A. Naor. The two possible values of the chromatic number of a random graph. Submitted.

[ASE92] N. Alon, J. H. Spencer, and P. Erdös. *The Probabilistic Method*. John Wiley & Sons, 1992.

[AV79] D. Angluin and L. Valiant. Fast probabilistic algorithms for Hamiltonian circuits and matchings. *Journal of Computer and System Sciences*, 18:155–193, 1979.

[BE95] R. Beigel and D. Eppstein. 3-coloring in time $O(1.3446^n)$: a no-MIS algorithm. In *Proceedings 36th Annual Symposium on Foundations of Computer Science*, pages 444–452, Milwaukee, WI, October 1995. IEEE.

[BE00] R. Beigel and D. Eppstein. 3-coloring in time $O(1.3289^n)$. Technical Report eprint cs.DS/0004046, CoRR, June 2000.

[BKPS02] P. Beame, R. Karp, T. Pitassi, and M. Saks. The efficiency of resolution and Davis-Putnam procedures. *SIAM Journal on Computing*, 31(4):1048–1075, 2002.

[BMWZ] A. Braunstein, M. Mezard, M. Weigt, and R. Zecchina. Constraint satisfaction by survey propagation. http://lanl.arxiv.org/ps/cond-mat/0212451.

[BP96] Paul W. Beame and Toniann Pitassi. Simplified and improved resolution lower bounds. In *Proceedings 37th Annual Symposium on Foundations of Computer Science*, pages 274–282, Burlington, VT, October 1996. IEEE.

[BSW01] E. Ben-Sasson and A. Wigderson. Short proofs are narrow – resolution made simple. *Journal of the ACM*, 48(2):149–169, 2001.

[BW85] E.A. Bender and H.S. Wilf. A theoretical analysis of backtracking in the graph coloring problem. *Journal of Algorithms*, 6(2):275–282, 1985.

[CA96] J.M. Crawford and L.D. Auton. Experimental results on the crossover point in random 3SAT. *Artificial Intelligence*, 81:31–57, 1996.

[CEI96] M. Clegg, J. Edmonds, and R. Impagliazzo. Using the Gröbner basis algorithm to find proofs of unsatisfiability. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, pages 174–183, Philadelphia, PA, May 1996.

[CG01] J. Culberson and I.P. Gent. Frozen development in graph coloring. *Theoretical Computer Science*, 265(1-2):227–264, 2001.

[CO03] A. Coja-Oghlan. The Lovász number of random graphs. In *Proc. 7th Intl. Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM 2003)*, 2003.

[CS88] V. Chvátal and Endre Szemerédi. Many hard examples for resolution. *Journal of the ACM*, 35(4):759–768, 1988.

[CS02] M. Chiarandini and T. Stützle. An application of iterated local search to graph coloring. In *Proceedings of the Computational Symposium on Graph Coloring and its Generalizations*, pages 112–125, 2002.

[DGH+02] E. Dantsin, A. Goerdt, E. Hirsch, R. Kannan, J. Kleinberg, C. Papadimitriou, P. Raghavan, and U. Sch"oning. A deterministic $(2 - 2/(k + 1))^n$ algorithm for k-SAT based on local search. *Theoretical Computer Science*, 280:69–83, 2002.

[DLL62] M. Davis, G. Logemann, and D. Loveland. A machine program for theorem proving. *Communications of the ACM*, 5:394–397, 1962.

[EDM03] L. Ein-Dor and R. Monasson. The dynamics of proving uncolourability of large random graphs i: symmetric colouring heuristic. *J. Phys. A: Math. Gen.*, 36:11055–11067, 2003.

[ER60] P. Erdös and A. Renyi. On the evolution of random graphs. *Magyar Tud. Akad. Mat. Kat. Int. Kozl.*, 5:17–61, 1960.

[FF96] C. Fleurent and J.A. Freland. Genetic and hybrid algorithms for graph coloring. *Annals of Operations Research*, 63:437–464, 1996.

[FG01] J. Friedman and A. Goerdt. Recognizing more unsatisfiable random 3-SAT instances efficiently. In J. Orejas, P. G. Spirakis, and J. van Leeuwen, editors, *Automata, Languages, and Programming: 28th International Colloquium*, volume 2076 of *Lecture Notes in Computer Science*, pages 310–321, Heraklion, Crete, July 2001. Springer-Verlag.

[GK01] A. Goerdt and M. Krivelevich. Efficient recognition of random unsatisfiable k-SAT instances by spectral methods. In *(STACS) 2001: 18th Annual Symposium on Theoretical Aspects of Computer Science*, volume 2010 of *Lecture Notes in Computer Science*, pages 294–304, Dresden, Germany, February 2001. Springer-Verlag.

[JŁR00] S. Janson, T. Łuczak, and A. Ruciński. *Random Graphs*. John Wiley & Sons, 2000.

[JM] H. Jia and C. Moore. How much backtracking does it take to color sparse random graphs? Rigorous results on heavy tails. Submitted.

[JT96] D. S. Johnson and M. Trick, editors. *Cliques, Coloring, and Satisfiability: Second Dimacs Implementation Challenge, October 11-13, 1993*, volume 26 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*. American Mathematical Society, 1996.

[KV02] M. Krivelevich and V. Vu. Approximating the independence number and the chromatic number in expected polynomial time. *Journal of Combinatorial Optimization*, 6:143–155, 2002.

[Law76] E. L. Lawler. A note on the complexity of the chromatic number problem. *Information Processing Letters*, pages 66–67, 1976.

[Łuc91] T. Łuczak. The chromatic number of random graphs. *Combinatorica*, 11:45–54, 1991.

[McD84] C. McDiarmid. Colouring random graphs. *Annals of Operations Research*, 1:183–200, 1984.

[Mit02] David G. Mitchell. *The Resolution Complexity of Constraint Satisfaction*. PhD thesis, University of Toronto, 2002.

[Wil84] Herbert S. Wilf. Backtrack: An $O(1)$ expected time algorithm for the graph coloring problem. *Information Processing Letters*, 18(3):119–121, March 1984.

ECCC ISSN 1433-8092
http://www.eccc.uni-trier.de/eccc
ftp://ftp.eccc.uni-trier.de/pub/eccc
ftpmail@ftp.eccc.uni-trier.de, subject 'help eccc'