A sharp threshold in proof complexity

Dimitris Achlioptas
Microsoft Research
One Microsoft Way
Redmond, WA 98052
optas@microsoft.com

Paul Beame*
Computer Science and Engineering
University of Washington
Seattle, WA 98195-2350
beame@cs.washington.edu

Michael Molloy†
Department of Computer Science
University of Toronto
Toronto, Ontario M5S 1A4
molloy@cs.toronto.edu

ABSTRACT

We give the first example of a sharp threshold in proof complexity. More precisely, we show that for any $\epsilon > 0$ and $\Delta > 2.28$, random formulas consisting of $(1-\epsilon)n$ 2-clauses and Δn 3-clauses, which are known to be unsatisfiable almost certainly, almost certainly require resolution and Davis-Putnam proofs of unsatisfiability of exponential size, whereas it is easily seen that random formulas with $(1+\epsilon)n$ 2-clauses (and Δn 3-clauses) have linear size proofs of unsatisfiability almost certainly.

A consequence of our result also yields the first proof that typical random 3-CNF formulas at ratios below the generally accepted satisfiability threshold (and thus expected to be satisfiable almost certainly) cause natural Davis-Putnam algorithms to take exponential time to find satisfying assignments.

1. INTRODUCTION

The satisfiability problem has received a great deal of study as the canonical NP-complete problem. In the last several years the very universality and flexibility of the satisfiability problem that made it a natural starting point for NP-completeness has also made it the basis for significant progress in the solution of a variety of practical problems including problems in constraint satisfaction [32], planning [22, 24], and symbolic model checking [10]. The basic tools for these advances are some very tight and efficient implementations of satisfiability algorithms using back-tracking search based on the Davis-Putnam/DLL procedure [16, 15] and using heuristic search based on hill-climbing and random walks [32, 31]. In a sense, these satisfiability algorithms have become the hammer and there is now a small industry in turning computational problems into nails.

One folklore property of these satisfiability algorithms is that they work extremely well at handling typical satisfiable formulas, although outlying satisfiable formulas can cause very bad behavior for any specific algorithm. (One highly-successful suggested strategy is to use randomized Davis-Putnam/DLL algorithms and re-start them with different random bits if they begin to take too long [21, 20].) It has previously been shown that Davis-Putnam/DLL algorithms are exponentially inefficient for proving unsatisfiability for typical unsatisfiable formulas [13, 7, 9, 6, 8] and thus there must be rare satisfiable formulas close to these unsatisfiable formulas that cause such algorithms to have exponential behavior. Nonetheless, their relative rarity makes it hard for us to lay hands on any single such formula. Thus, prior to this work, there was no known method to prove that any specific satisfiable formulas force Davis-Putnam/DLL algorithms to take exponential time.

We give a method to prove the first lower bounds on the running times of Davis-Putnam/DLL algorithms for typical *satisfiable* CNF formulas. We show that almost all random formulas from a large class of formulas, for which there is strong empirical evidence of satisfiability, force natural Davis-Putnam/DLL algorithms to take exponential time. In particular, we show that natural Davis-Putnam algorithms when applied to random 3-CNF formulas with, for example, $3.81n$ 3-clauses will almost certainly generate subproblems consisting of unsatisfiable mixed formulas of the types that require exponential proofs of unsatisfiability.

Let $\mathcal{F}^n_{\epsilon,\Delta}$ be the distribution where $m_2 = (1-\epsilon)n$ and $m_3 = \Delta n$, for some arbitrary *constants* $\Delta, \epsilon > 0$.

THEOREM 1. *For every $\Delta, \epsilon > 0$, if $F \sim \mathcal{F}^n_{\epsilon,\Delta}$, then w.h.p., $res(F) = 2^{\Omega(n)}$ and $DLL(F) = 2^{\Omega(n)}$.*

The key to proving this result is a proof complexity lower bound for mixed CNF formulas consisting of clauses of length 2 and length 3. This bound is interesting in its own right since it is the first example of a sharp threshold for the proof complexity of unsatisfiable formulas. More precisely, we prove that random formulas consisting of $(1-\epsilon)n$

* Research supported by NSF grant CCR-9800124.
† Research supported by NSERC and by a Sloan Fellowship

STOC'01, July 6-8, 2001, Hersonissos, Crete, Greece.

2-clauses and Δn 3-clauses almost certainly require resolution and Davis-Putnam/DLL proofs of *unsatisfiability* of exponential size. For $\Delta > 2.28$ these formulas are known to be unsatisfiable almost certainly (see below). This is quite surprising since slightly increasing the number of 2-clauses from $(1-\epsilon)n$ to $(1+\epsilon)n$ makes the 2-CNF sub-formula alone unsatisfiable almost certainly and this unsatisfiability problem is trivial for Davis-Putnam/DLL proofs.

This is not the first time that the special properties of mixed formulas consisting of 2- and 3-clauses have been considered. A frequent observation about converting problems from other domains into satisfiability problems is that they typically become mixed CNF formulas with a substantial number of clauses of length 2 along with clauses of length 3. The other interest in such mixed formulas derives from the fact that, as Davis-Putnam/DLL algorithms run, they simplify the current CNF formula by removing and reducing the size of the clauses based on the current partial assignment. Since a clause of length 1 (a unit clause) in a satisfiable CNF formula must be satisfied, these algorithms require no branching to determine what to do if such a clause exists. The case in Davis-Putnam/DLL algorithm design that is most critical to performance is when the reduced formulas are mixtures of clauses of length at least 2.

Before we outline this work, which is highly relevant for our current paper, we first discuss randomly-chosen k-CNF formulas, which are an important test case for satisfiability algorithms. A fundamental conjecture about such formulas says that there is a constant α_k, the *satisfiability threshold*, such that a random k-CNF formula of clause-variable ratio Δ is almost certainly satisfiable for $\Delta < \alpha_k$ and almost certainly unsatisfiable if $\Delta > \alpha_k$. It is known that 1 is a threshold for random 2-SAT [12, 19] and empirically, it appears that there is a threshold for random 3-SAT around 4.2 [33, 25]. (The best proven bounds show that random 3-SAT formulas are almost certainly satisfiable for $\Delta < 3.26$ [5] and almost certainly unsatisfiable for $\Delta > 4.598$ [23] and it is known that the 'phase' transition from satisfiable to unsatisfiable formulas is asymptotically sharp [17].) The formulas we consider are at densities bounded well below the empirical 4.2 threshold and thus should be almost certainly satisfiable.

The proven lower bounds on the satisfiability thresholds have all been derived by analyzing specific Davis-Putnam-like algorithms without back-tracking ('card-type/myopic algorithms' in the terminology of [1, 5]) and showing that they succeed with positive constant probability. A widely-cited and repeated observation of Selman, Mitchell, and Levesque [33], based on experiments with a back-tracking version of a simple card-type algorithm on small problems, was that random 3-SAT is easy in the satisfiable region up to the 4.2 threshold, becomes sharply much harder at the threshold and quickly becomes easy again at larger densities in the unsatisfiable region.

The upper end of this 'easy-hard-easy' characterization is somewhat misleading since it is known that at any ratio above the threshold any Davis-Putnam/DLL or resolution algorithm almost certainly requires exponential size proofs of unsatisfiability [13] and the rate of decline in complexity with increasing ratio has been analyzed as well [6]. Our new results show that the lower end of this characterization is also somewhat misleading; in fact, our results show that the exponentially hard region for the Unit Clause (UC) algorithm studied in [11] (and the basis for the back-tracking algorithm [33]) begins at least at ratio 3.81, well before ratio 4.2. (This concurs with recent experimental evidence that even the best of current Davis-Putnam implementations seem to have bad behavior below the threshold [14].)

An extension of random 3-SAT problems to mixed formulas has led to the study of the so-called $(2+p)$-SAT problem. Here one looks at randomly-generated formulas on n variables with a mix of clauses of length 2 and 3 where the fraction of clauses of length 3 is p. Using a mixture of empirical results and heuristic, non-rigorous techniques borrowed from statistical physics Kirkpatrick et. al. [27, 28, 29] gave evidence that one could add up to roughly $2n/3$ random 3-clauses (corresponding to p around 0.4) and have no impact on typical satisfiability problems in that the point at which the threshold is reached is dependent solely on the number of 2-clauses. Furthermore they suggested that the phase transition for the $(2+p)$-SAT problem itself changed character from so-called a type 1 transition (discontinuous "order parameter") representative of 2-SAT to a type 2 transition (continuous "order parameter") believed to be representative of 3-SAT at this point and suggested that this type 2 phase transition somehow characterizes NP-hard problems. (This latter conjecture has been put in further perspective by [3].)

Achlioptas et. al. [4] proved two rigorous results about $(2+p)$-SAT : first, adding $2n/3$ randomly chosen 3-clauses to a random formula with $(1-\epsilon)n$ 2-clauses yields a formula that is almost surely satisfiable (and this assignment can be found by a simple card-type algorithm), and second, adding $2.28n$ random 3-clauses to a formula with $(1-\epsilon)n$ 2-clauses yields an almost certainly unsatisfiable formula. The cut-off is surely sharp [4] and empirical evidence suggests that the $2n/3$ bound is much closer to the truth [29]. Furthermore $2n/3$ is the provable limit for showing satisfiability by card-type algorithms [1] and in [2], it was conjectured that it is tight for all algorithms.

The algorithms that have been analyzed for 3-SAT, UC (unit clauses first), GUC, etc., all card-type algorithms, have the property that they make irrevocable choices to the partial assignment and at every point in their execution the partial formula that remains is an unbiased random formula characterized by the number of clauses of each size. (The severely limited backtracking extension of these algorithms analyzed by Frieze and Suen [18] also has this property. Note that it merely increases the probability of success without changing for range of parameters for which they work.) Since interesting points are when these algorithms have exhausted all unit clauses, they are characterized by the sequence of $(c_2 n, c_3 n)$ of random subformula types they pass through. The algorithms will succeed almost certainly if they exhaust the 3-clauses without c_2 ever reaching 1. In fact, if they reach $((1-\epsilon)n, 2n/3)$ without c_2 every reaching 1 they will succeed with positive constant probability. For card-type algorithms there are differential equations that can be used to describe this sequence.

Our lower bounds for arbitrary back-tracking extensions of card-type algorithms follow by showing that with appropriate initial clause ratio these algorithms pass through a sequence $((1-\epsilon)n, c_3 n)$ where c_3 is at least 2.28. This formula is almost certainly unsatisfiable but this almost certainly requires an exponential size Davis-Putnam proof and thus, once this node has been explored in the Davis-Putnam search, the search cannot leave the sub-tree for an exponentially long time. Our results show that this happens for UC at $3.81n$ clauses and for GUC at $3.98n$ clauses. Furthermore, these bounds will be improved with any improvement in the $2.28n$ unsatisfiability limit for $(2+p)$-SAT . In fact, if the upper bound is reduced to $2n/3$ to match the lower bound then we will have the following remarkably sharp behavior: every card-type algorithm A is such that it operates in linear time with positive constant probability up to some threshold β_A but any extension of A requires exponential time for all ratios larger than β_A!

Our lower bound on the proof complexity of mixed formulas is similar in general spirit to other lower bounds for resolution algorithms but requires considerably more subtlety. We first prove a number of detailed combinatorial properties of random 2-CNF formulas with $(1-\epsilon)n$ clauses. We need non-trivial properties of extensions, that we call *bomponents,* of the notion of connected component to the natural directed graphs on literals associated with each 2-CNF formula. By carefully bounding the number of vertices in bomponents of each size we show that the 2-CNF formula almost certainly has properties that assure that almost all extensions by linear-sized 3-CNF formulas yield exponential size resolution (and Davis-Putnam) proofs.

Since we have already sketched the general argument for the lower bounds for Davis-Putnam search on satisfiable formulas, for the remainder of this extended abstract we concentrate entirely on the proof complexity result. We leave the details of the proof of the application to Davis-Putnam algorithms to the full version of the paper.

2. BOUNDING RESOLUTION REFUTATION SIZE

The resolution rule allows one to derive a clause $(A \vee B)$ from two clauses $(A \vee x)$ and $(B \vee \bar{x})$. A resolution refutation of an unsatisfiable CNF formula F begins with the clauses of F and by a sequence of inferences using the resolution rule derives the empty clause. The size of a resolution refutation is the number of clauses appearing in the proof. Given F, let $res(F)$ be the length of the shortest resolution refutation of F. The Davis-Putnam/DLL algorithm on a CNF formula F performs a backtracking search for a satisfying assignment of F by extending partial assignments until they either reach a satisfying assignment or violate a clause of F. It is well known that for an unsatisfiable formula F, the tree of nodes explored by any DLL algorithm can be converted to a resolution refutation of F where the pattern of inferences forms a tree. Let $DLL(F)$ be the size of the smallest such refutation, i.e. the size of the smallest DLL tree associated with F.

DEFINITION 2. *Let F be an an arbitrary CNF formula.*

- *Given a set F of clauses, a literal x is* pure *in F if and only if x appears in F but $\neg x$ does not appear in F.*
- *$V(F)$ denotes the set of variables of F.*
- *$|F|$ denotes the number of clauses in F.*
- *The* degree *of a variable v, $\deg(v)$, is the number of clauses of F containing one of $v, \bar{v}$ (analogously for literals). The* average degree *of F is $\left(\sum_v \deg(v)\right)/|V(F)|$.*
- *If F is an arbitrary 2-CNF formula on variables $\{x_1, \ldots, x_n\}$ we associate the following directed graph $\vec{D}(F)$ with F:*
  - *The vertex set is $\{x_1, \ldots, x_n, \bar{x}_1, \ldots, \bar{x}_n\}$.*
  - *The edge set is $\{x \to y : (\bar{x} \vee y)$ is a clause in $F\}$.*
- *We will say that a literal ℓ is* near-pure *in F if $\deg(\bar{\ell}) = 1$.*
- *We will say that a directed cycle $C = \ell_1 \to \ell_2 \to \cdots \to \ell_q \to \ell_1$ in $\vec{D}(F_2)$ is* pure *if all of $\ell_1, \ldots, \ell_q$ are near pure in F. (Note that each literal can appear in at most one pure cycle and that any pure cycle forms a directed cycle in $\vec{D}(F_2)$).*
- *F is μ-*pure *if it has at most $\mu|V(F)|$ pure items, i.e. pure literals and cycles.*

The *boundary* of F, $b(F)$, is the set of pure items in F. Note that this generalizes the definition of boundary from [9] where the boundary was defined to be those literals whose variables appeared in precisely one clause of F. Define the *satisfiability threshold for F*, $s(F)$, to be the size of the minimal unsatisfiable subset of G.

Given a set of clauses F we say that $F \Rightarrow_{Res} C$ if and only if there is a resolution derivation of C from F that uses all clauses in F. Define the *sub-critical expansion of F*, $e(F) = \max_{s \le s(F)} \min\{|b(G)| : s/2 \le |G| < s\}$. The following are extensions of [9].

LEMMA 3. *If $F \Rightarrow_{Res} C$ then C contains at least $|b(F)|$ literals.*

PROOF. This is trivial for pure literals since, once they appear in the derivation, they cannot be removed. For pure cycles, observe that once a literal from the pure cycle appears in a clause in the derivation, all clauses derived from it will also contain a literal from that cycle. The argument is by induction on the proof in topological order: Suppose that a literal x from the cycle appears in clause A and does not appear in clause D where D is derived from A and some clause B by the resolution rule. Therefore x must appear negatively in B and the resolution rule must use x. The only way that $\bar{x}$ can be in B is if B is derived using the one negative clause $\bar{x} \vee y$ that appears in the pure cycle containing x. Therefore by the induction hypothesis B contains some literal in this pure cycle and this literal will appear in D. Since each literal is in at most one pure cycle, the lemma follows. □

Using Lemma 3 the following is immediate from [9].

PROPOSITION 4. *If all clauses in a formula F have size at most k, then $res(F) = 2^{\Omega([e(F)-k]^2/n)}$ and $DLL(F) \ge 2^{e(F)-k}$.*

For each n, let us fix a canonical set of n variables and for each fixed $k \ge 0$, let $C_k(n)$ denote the set containing all $2^k\binom{n}{k}$ non-trivial k-clauses on our canonical set of variables. We will consider a random formula F on n variables formed by selecting uniformly, independently and with replacement $m_2 = m_2(n)$ clauses from $C_2(n)$ and $m_3 = m_3(n)$ clauses from $C_3(n)$. We will say that a sequence of random events $\mathcal{E}_n$ occurs with high probability (w.h.p.) if $\lim_{n\to\infty}\Pr[\mathcal{E}_n]=1$.

Let $\mathcal{F}^n_{\epsilon,\Delta}$ be the distribution where $m_2=(1-\epsilon)n$ and $m_3=\Delta n$, for some arbitrary *constants* $\Delta,\epsilon>0$.

LEMMA 5. *For every $\Delta,\epsilon>0$ there exist $\zeta=\zeta(\Delta,\epsilon)>0$ and $\mu=\mu(\Delta,\epsilon)>0$ such that*

1. *W.h.p. F contains no unsatisfiable subformula on $v\le\zeta n$ variables.*
2. *W.h.p. F contains no μ-pure subformula on $\frac{1}{2}\zeta n\le v\le\zeta n$ variables.*

COROLLARY 6. *For $F\sim\mathcal{F}^n_{\epsilon,\Delta}$, w.h.p., $res(F)=2^{\Omega(n)}$ and $DLL(F)=2^{\Omega(n)}$.*

The overall proof strategy will be to analyze almost certain properties of the 2-clauses of F and show that such subformulas must have large boundaries and then to show that the addition of 3-clauses does not significantly reduce the size of the boundary. (Since the 2-clause subformula of F alone is almost certainly satisfiable, the fact that the two clauses alone have a large boundary does not contradict the fact that 2SAT is easy for resolution.) This argument is subtle because the 2-clauses of F are so close to being unsatisfiable themselves and because we need to handle all possible subformulas among the 2-clauses. The latter requirement necessitates the introduction of a somewhat unusual graph-theoretic concept in the di-graph associated with the 2-clauses of F.

3. BOUNDING PURITY

By Lemma 3, any minimal unsatisfiable formula must have no pure items. Therefore, it will be convenient to restate Lemma 5 in terms of the following definition.

DEFINITION 7. A subformula H of F on v variables is (μ,ζ)-*pure* if: i) $v<\frac{1}{2}\zeta n$ and H is 0-pure, or, ii) $\frac{1}{2}\zeta n\le v\le\zeta n$ and H is μ-pure.

Hence, Lemma 5 can be restated as

LEMMA 8. *For every $\Delta,\epsilon>0$, there exists $\zeta=\zeta(\Delta,\epsilon)>0$, $\mu=\mu(\Delta,\epsilon)$ such that w.h.p. F contains no (μ,ζ)-pure subformula.*

To prove Lemma 5 we will need need to consider two possibilities for the fraction of 3-clauses appearing in a potential (μ,ζ)-pure subformula. In particular, it will be relatively easy to prove that w.h.p. every subformula rich in 3-clauses is not (μ,ζ)-pure (Lemma 11). Proving the same assertion for subformulas where the fraction of 3-clauses is arbitrarily small (Lemma 10) will be harder, comprising Sections 4 and 5. More precisely, let us say that

DEFINITION 9. *A formula F is* δ-rich *if it has at least $\delta|V(F)|$ 3-clauses, and* δ-poor *otherwise.*

Lemma 8 will follow readily from the following two lemmas.

LEMMA 10. *For every $\Delta,\epsilon>0$ and all $\zeta\le\zeta_1(\Delta,\epsilon)$, there exist $\delta_1=\delta_1(\Delta,\epsilon,\zeta)>0$ and $\mu_1=\mu_1(\Delta,\epsilon,\zeta)$ such that w.h.p. $\mathcal{F}^n_{\epsilon,\Delta}$ contains no δ_1-poor, (μ_1,ζ)-pure subformula.*

LEMMA 11. *For every $\Delta,\epsilon,\delta>0$, and all $\zeta\le\zeta_2(\Delta,\epsilon,\delta)$, there exists $\mu_2=\mu_2(\Delta,\epsilon,\zeta)$ such that w.h.p. $\mathcal{F}^n_{\epsilon,\Delta}$ contains no δ-rich, (μ_2,ζ)-pure subformula.*

PROOF OF LEMMA 8. Given $\Delta,\epsilon>0$ we first apply Lemma 10 with $\zeta=\zeta_1(\Delta,\epsilon)$. This determines a value $\delta^*=\delta_1(\Delta,\epsilon,\zeta_1)>0$. We then apply Lemma 11 with $\delta=\delta^*$ and $\zeta=\zeta_2(\Delta,\epsilon,\delta^*)>0$. As a result, we get that for $\zeta=\min\{\zeta_1(\Delta,\epsilon),\zeta_2(\Delta,\epsilon,\delta^*)\}>0$ and $\mu=\max\{\mu_1(\Delta,\epsilon,\zeta),\mu_2(\Delta,\epsilon,\zeta)\}>0$, w.h.p. there are no δ^*-poor, (μ,ζ)-pure subformulas and no δ^*-rich, (μ,ζ)-pure subformulas. □

To prove Lemma 11 we will need the following

LEMMA 12. *For every $\Delta,\epsilon,\beta>0$ there exists $\zeta=\zeta(\Delta,\epsilon,\beta)>0$ such that w.h.p. no subformula of $F\sim\mathcal{F}^n_{\epsilon,\Delta}$ on $v\le\zeta n$ variables has average degree at least $2-\beta/2$ and at least βv 3-clauses.*

PROOF. For a given $\beta>0$ let α be such that $2\alpha+3\beta=2-\beta/2$, implying $\alpha+2\beta-1=\beta/4$. We will bound the expected number of subformulas on v variables having αv 2-clauses and βv 3-clauses. Using this bound we will show that there exists $\zeta=\zeta(\Delta,\epsilon,\beta)>0$ such that the total expected number of such subformulas for $1\le v\le\zeta n$ variables is $o(1)$.

For a fixed v, we can bound the expected number, $Q_\beta(v)$, of such subformulas by

$$\binom{n}{v}\binom{(1-\epsilon)n}{\alpha v}\binom{\Delta n}{\beta v}\left(\frac{v}{n}\right)^{2\alpha v+3\beta v}\qquad(1)$$

$$\le\left(\frac{en}{v}\right)^v\left(\frac{e(1-\epsilon)n}{\alpha v}\right)^{\alpha v}\left(\frac{e\Delta n}{\beta v}\right)^{\beta v}\left(\frac{v}{n}\right)^{2\alpha v+3\beta v}$$

$$=\left(\frac{(1-\epsilon)^\alpha e^{1+\alpha+\beta}\Delta^\beta}{\alpha^\alpha\beta^\beta}\times\left(\frac{v}{n}\right)^{\alpha+2\beta-1}\right)^v$$

$$=\left(K\times\left(\frac{v}{n}\right)^{\beta/4}\right)^v,\qquad(2)$$

for some (constant) $K=K(\beta,\Delta,\epsilon)$. It is easy to see that for $\zeta=\zeta(\beta,K)$ sufficiently small and all n sufficiently large, if $v\le\zeta n$ the right hand side of (2) is decreasing with v. Therefore,

$$\sum_{1\le v\le\zeta n}Q_\beta(v)\le\sum_{v=1}^{8/\beta}Kn^{-\beta/4}+\sum_{v>8/\beta}K\times\frac{(8/\beta)^{\beta/4}}{n^2}=o(1).$$

□

PROOF OF LEMMA 11. We first observe that any μ-pure subformula H on v variables must contain at most μv pure literals and therefore have average degree at least $2(1-\mu)$. Moreover, if H is δ-rich, it contains at least δv 3-clauses. We will take $\mu_2=\delta/4$. Thus, any δ-rich, μ_2-pure subformula must have average degree at least $2-\delta/2$ and contain at least δv 3-clauses. But from Lemma 12, there exists $\zeta_2=\zeta(\Delta,\epsilon,\delta)>0$ such that w.h.p. F contains no such subformula on $v\le\zeta_2 n$ variables. □

To prove Lemma 10 we will in fact prove a stronger lemma. In particular, rather than proving the lemma's assertion for $\mathcal{F}^n_{\epsilon,\Delta}$, we will prove it for an arbitrary formula on n variables formed by starting with a 2-CNF formula satisfying certain properties and adding to it $m_3=\Delta n$ random 3-clauses. To complete the proof, in appendix A, we prove that F_2 satisfies these properties w.h.p.

4. THE δ-POOR CASE

As we mentioned earlier, we will prove the assertion of Lemma 10 for formulas formed by starting with a 2-CNF formula satisfying certain properties and adding m_3 random 3-clauses. To describe these properties we need to introduce the following definitions.

DEFINITION 13. Let F be an arbitrary 2-CNF formula. For literals x, y appearing in F let us write $x \rightsquigarrow_F y$ iff $x = y$ or there exists a directed path in $\vec{D}(F)$ from x to y.

- For each literal x we let $\mathrm{In}_F(x) = \{y : y \rightsquigarrow_F x\}$.
- For a set of literals S let $G(S) = G_F(S)$ be the undirected graph formed by considering the subgraph of $\vec{D}(F)$ induced by the vertices corresponding to S and ignoring the direction of arcs.
  - We will say that $\mathrm{In}_F(x)$ is *tree-like* if $G(\mathrm{In}_F(x))$ contains *no* cycle.
  - We will say that $\mathrm{In}_F(x)$ is *simple* if $G(\mathrm{In}_F(x))$ contains *at most one* cycle.
- For each literal x in F, the *bomponent* of x, $\mathrm{Bomp}_F(x) = \mathrm{In}_F(x) \cup \bigcup_{y \in \mathrm{In}_F(x)} \mathrm{In}_F(\bar{y})$.
- $T_i(F) = |\{x : |\mathrm{Bomp}_F(x)| = i\}|$.

Lemma 10 will follow readily from the following two lemmas.

LEMMA 14. *Fix $\rho \in (0, 1)$ and $\Delta > 0$. Let F^* be a formula formed by taking*

- *Any set of clauses from $C_2(n)$ such that the resulting formula F_2^* satisfies:*
  1. *For every literal ℓ, $\mathrm{In}_{F_2^*}(\ell)$ is simple.*
  2. *There are at most $\log n$ literals ℓ, such that $\mathrm{In}_{F_2^*}(\ell)$ is not tree-like.*
  3. *For all i, $T_i(F_2^*) \le 2(1 - \rho)^i\, n$.*
- *No more than Δn clauses from $C_3(n)$, chosen uniformly independently and with replacement.*

For all $\zeta \le \zeta_0(\Delta, \rho)$, there exist $\delta = \delta_0(\Delta, \rho, \zeta) > 0$, and $\mu = \mu_0(\Delta, \rho, \zeta) > 0$ such that, w.h.p. F^* contains no δ-poor, (μ, ζ)-pure subformula.

LEMMA 15. *Fix $\epsilon > 0$ and let F_2 be a random 2-SAT formula formed by selecting uniformly, independently and with replacement $m_2 \le (1-\epsilon)n$ 2-clauses from $C_2(n)$. There exists $\rho = \rho(\epsilon)$, such that w.h.p. F_2 simultaneously satisfies all three conditions of Lemma 14 (where $F_2^* = F_2$).*

PROOF OF LEMMA 10. Given $\Delta, \epsilon > 0$ we first apply Lemma 15 to get $\rho = \rho(\epsilon)$ and then apply lemma 14 with that ρ. Thus, we get that for all $\zeta \le \zeta_0(\Delta, \rho)$, there exist $\delta = \delta_0(\Delta, \rho, \zeta) > 0$, and $\mu = \mu_0(\Delta, \rho, \zeta) > 0$ such that, w.h.p. $\mathcal{F}^n_{\epsilon,\Delta}$ contains no δ-poor, (μ, ζ)-pure subformula. To conclude the proof we set $\zeta_1(\Delta, \epsilon) = \zeta_0(\Delta, \rho(\epsilon))$ and similarly for δ_1, μ_1. $\square$

5. PROOF OF LEMMA 14

Let $\rho, \Delta > 0$, F_2^* be fixed and choose F^* as in the statement of the lemma. For any $\delta, \zeta, \mu > 0$, consider any (candidate) δ-poor, (μ, ζ)-pure subformula H of F^*. Let $v = |V(H)|$ and denote by H_2 the subformula induced by the 2-clauses of H.

The general idea of the argument is as follows. The subformula H_2 has many loose ends, the pure items of H_2, that must be (mostly) covered by the 3-clauses of H in order to for H to have very few pure items. More generally, H may itself create new loose ends but then must cover most of them up again. We show we can cover the literals appearing in H by the bomponents of the loose ends. Further since the bomponent sizes are typically small we get a large number of loose ends, the set $P(H)$ below, each of which must be covered by a different 3-clause literal.

In the case that the set of variables in H is large, we will set parameters so that the contribution of all large bomponents is so small that the δ-poverty of H simply doesn't allow enough 3-clause literals to cover enough loose ends. In the case that the set of variables in H is not so large, we get to use the fact that *every* loose end must be covered but we need a probability argument. Intuitively, it seems unlikely that the random 3-clauses of the formula will exactly cover all loose ends, both those from H_2 and the ones that the 3-clauses themselves generate. One subtlety of the probability analysis is that the new loose ends generated depend on the 3-clauses themselves. In order to make this analysis work, we need stronger bounds on the set of variables in H; we show that, except possibly for the case of the very small number of bomponents that are not tree-like, we get two loose ends per bomponent instead of just one and this is enough to make the probability calculation work. Finally, we use a sharp specialized moment bound to show that the rare large bomponents do not skew the probabilities too much and derive the claimed result.

We now work through the details of the argument. Define the set $P = P(H)$ of literals based on H as follows: P consists of the pure literals of H_2, the smallest numbered literal in each pure cycle of H_2, and every literal on the variables of $V(H) - V(H_2)$. Clearly P contains every pure literal of H and also contains one literal from each pure cycle of H (and since pure cycles are disjoint they are represented by distinct literals).

LEMMA 16. *For any subformula H of F^*, the number of distinct literals in the 3-clauses of H is at least the number of literals in $P(H)$ that are not contained in pure items in H.*

PROOF. We define a one-to-one (but not necessarily onto) mapping from the literals of $P = P(H)$ that are not contained in pure items of H to the literals appearing in the 3-clauses of H. Any literal x in P, that was pure in H_2 or is a literal on $V(H) - V(H_2)$ but is not pure in H, must have $\bar{x}$ in some 3-clause of H and so we map x to $\bar{x}$. The pure cycles of H_2, whose smallest numbered literals form the remainder of P, are disjoint from each other and from the other literals in P. Consider such a cycle C that is pure in H_2 and let $x \in P$ be the smallest numbered literal in C. C will remain pure in H unless there is some y in C such that $\bar{y}$ appears in a 3-clause of H. We map x to $\bar{y}$. The fact that our map is one-to-one follows from the disjointness property of the cycles. $\square$

LEMMA 17. *For any subformula H of F^*, for each literal x on the variables in $V(H)$, there is some literal $y \in P(H)$ such that $x \in \mathrm{Bomp}^*_{F_2}(y)$.*

PROOF. If the literal x is on $V(H) - V(H_2)$ then $x \in P$ so we can take $y = x$. If x is a literal on $V(H_2)$ and $\bar{x}$ appears in H_2, then in the digraph $\vec{D}(H_2)$ walk forward from x as far as possible, since all bomponents contain at most one cycle, either we reach a sink of $\vec{D}(H_2)$, in which case the label of that sink is a pure literal y in P which satisfies $x \in \mathrm{In}_{H_2}(y) \subseteq \mathrm{In}_{F_2}(y)$, or we reach a pure cycle of H_2 in which case the smallest numbered literal y in this cycle satisfies $x \in \mathrm{In}_{H_2}(y) \subseteq \mathrm{In}_{F_2^*}(y)$. Clearly in either case, both x and $\bar{x}$ appear in $\mathrm{Bomp}_{F_2^*}(y)$. $\square$

For convenience throughout the rest of this proof we will write $\mathrm{Bomp}(x)$ for $\mathrm{Bomp}_{F_2^*}(x)$ and for a set T of literals we will write $\mathrm{Bomp}(T) = \bigcup_{x \in T} \mathrm{Bomp}(x)$.

LEMMA 18. *For any $\zeta \le 1$, there are $\mu_0 = \mu_0(\zeta, \rho)$ and $\delta_0 = \delta_0(\zeta, \rho) > 0$ such that if $\mu \le \mu_0$, $\delta \le \delta_0$, and $\zeta n/2 \le v = |V(H)|$ then H is not a μ-pure subformula of F^*.*

PROOF. Suppose that H is a μ-pure subformula of F^* and $\zeta n/2 \le V(H)$. Since H has at most δv 3-clauses and at most μv items remain pure in H, by Lemma 16, $|P| \le (3\delta + \mu)v$. Choose $I = I(\zeta, \rho)$ such that $\sum_{i > I} 2i(1-\rho)^i < \zeta/2$ and let $\mu_0 = \mu_0(\zeta, \rho) = 1/(4I)$ and $\delta_0 = \delta_0(\zeta, \rho) = 1/(4I)$. Therefore $|P| \le v/I$. By Lemma 17, $|\mathrm{Bomp}(P)| \ge 2v$. Let $L_I = \{x \ : \ |\mathrm{Bomp}(x)| > I\}$. Using the bound on the sizes of $T_i(F_2^*)$, $|\mathrm{Bomp}(P \cap L_I)| \le |\mathrm{Bomp}(L_I)| \le \sum_{i > I} 2i(1-\rho)^i n < \zeta n/2 \le v$. Therefore $|\mathrm{Bomp}(P - L_I)| \ge |\mathrm{Bomp}(P) - \mathrm{Bomp}(P \cap L_I)| > v$. However, $|\mathrm{Bomp}(P - L_I)| \le |P| \cdot I \le v$ which is a contradiction. $\square$

It remains then to show that the probability that there is some 0-pure subformula H of F^* with $|V(H)| < \zeta n/2$ is $o(1)$ in n. In this part of the proof we won't even use the fact that we only need to show this for δ-poor subformulas.

Suppose that H is 0-pure and $v = |V(H)| < \zeta n/2$. For any literal x (set of literals T), let $cover(x)$ (resp. $cover(T)$) be the set of literals appearing in $\mathrm{Bomp}(x)$ (resp. $\mathrm{Bomp}(T)$) together with their complements.

LEMMA 19. *For any subformula H of F^* that is 0-pure there exists $P^* = P^*(H) \subseteq P = P(H)$ such that*

1. *$cover(P^*)$ contains every literal appearing in H and*
2. *$|P^*| \le \lfloor \frac{1}{2}(|P| + t_c) \rfloor$ where $t_c = t_c(H)$ is the number of literals $x \in P^*$ such that $\mathrm{In}_{H_2}(x)$ is not tree-like.*

PROOF. Let $\hat{P} \subseteq P$ be the set of literals in P on variables in $V(H_2)$. By definition, for every $x \in P - \hat{P}$, $\bar{x} \in P - \hat{P}$. Let $P_{tree} \subseteq \hat{P}$ be the set of all literals $x \in \hat{P}$ with $\mathrm{In}_{F_2^*}(x)$ tree-like. First we prove that for every $x \in P_{tree}$ there is at least one $y \in \hat{P}$, $y \ne x$ such that $\bar{y} \in \mathrm{In}_{F_2^*}(x)$. For $x \in P_{tree}$, $\mathrm{In}_{H_2}(x)$ is tree-like since $\mathrm{In}_{H_2}(x) \subseteq \mathrm{In}_{F_2^*}(x)$. Therefore there is a vertex $z \in \mathrm{In}_{H_2}(x)$ of in-degree 0 in $\vec{D}(H_2)$ such that $z \ne \bar{x}$. Furthermore, since z appears in H_2, $\bar{x} \in \hat{P}$ so we can take $y = \bar{z} \ne x, \bar{x}$.

Note that $\bar{y} \in \mathrm{In}_{H_2}(x) \subseteq \mathrm{In}_{F_2^*}(x)$ implies $\bar{x} \in \mathrm{In}_{F_2^*}(y)$. Thus we form an undirected graph G with vertex set $\hat{P}$ and an edge $\langle x, y \rangle$ for each pair of literals with $\bar{y} \in \mathrm{In}_{F_2^*}(x)$. Let $P' \supseteq P_{tree}$ be the set of vertices in G of positive degree, consider a spanning forest of the vertices in P', and consider any bipartition of that forest. Let P_1 be the smaller side of that bipartition. Therefore P_1 dominates P', i.e. every vertex in $P' - P_1$ has a neighbour in P_1 and thus $P_1 \cup (\hat{P} - P')$ dominates all of $\hat{P}$. Letting $|\hat{P} - P'| = a$, $|P_1 \cup (\hat{P} - P')| \le a + \lfloor \frac{1}{2}(|\hat{P}| - a) \rfloor \le \lfloor \frac{1}{2}(|\hat{P}| + a) \rfloor$. Adding the positive form of each literal in $P - \hat{P}$ to $P_1 \cup (\hat{P} - P')$ we obtain a set P^* of size at most $\lfloor \frac{1}{2}(|P| + a) \rfloor$. Since $\hat{P} - P' \subseteq P^*$ and $P_{tree} \subseteq P'$, $t_c \ge a$ and P^* satisfies the claimed size condition.

By definition of P, $\hat{P}$, and P^*, P^* contains the positive literal corresponding to each variable in $V(H) - V(H_2)$, so $cover(P^*)$ contains all literals on variables in $V(H) - V(H_2)$. Further for every literal x such that $\bar{x}$ appears in H_2, as shown in the proof of Lemma 17, $x \in \mathrm{In}_{H_2}(y)$ for some literal y in P. By definition of P^* either $y \in P^*$ or there is some $z \in P^*$ such that $\bar{y} \in \mathrm{In}_{H_2}(z)$. Therefore $x \in \mathrm{Bomp}(z)$ and thus both x and $\bar{x}$ are in $cover(z)$. Thus $cover(P^*)$ contains all literals on $V(H_2)$ as well, so the lemma follows. $\square$

We will compute the total probability that some 0-pure subformula H of F^* exists by computing the probability for each set T of literals that some such H exists with $P^*(H) = T$ and sum over all choices of T.

LEMMA 20. *Let T be a set of literals, $t = |T|$, and suppose that H is a 0-pure subformula of F^* with $P^*(H) = T$ and $t_c = t_c(H)$. Then H and thus formula F^* must contain at least $2t/3 - t_c/3$ 3-clauses whose literals are contained in $cover(T)$; further if $t \le 10 t_c$ then there are at least $3t/5$ 3-clauses of F^* whose literals are contained in $cover(T)$.*

PROOF. By Lemma 19, since $P^*(H) = T$, $|P(H)| \ge 2|T| - t_c = 2t - t_c$. By Lemma 16, since H is 0-pure the 3-clauses of H contain at least $|P(H)| = 2t - t_c$ literals and therefore H contains at least $(2t - t_c)/3$ 3-clauses of F^*. Since, by Lemma 19, all literals in these clauses are in $cover(P^*) = cover(T)$. In case $t \ge 10 t_c$ then this is at least $(2/3 - 1/30)t > 3t/5$. $\square$

LEMMA 21. *Let Δ, ρ be given. There is a $K = K(\Delta)$ such that for T a set of literals, $t = |T|$, the probability that a random F^* has a 0-pure subformula H with $P^*(H) = T$ is at most*

1. *$R(T) = (K/(tn^2))^{3t/5} |\mathrm{Bomp}(T)|^{9t/5}$ if $t_c(H) \le t/10$, and at most*
2. *$R'(T, t_c) = (K/n^2)^{2t/3 - t_c/3} |\mathrm{Bomp}(T)|^{2t - t_c}$ if $t_c = t_c(H) > t/10$.*

PROOF. Since F^* has Δn 3-clauses, the probability that at least s of them land entirely in $cover(T)$ is at most

$$\begin{aligned}
& \binom{\Delta n}{s} \left[\frac{|cover(T)|^3}{8\binom{n}{3}} \right]^s \\
\le & \binom{\Delta n}{s} \left[\frac{|\mathrm{Bomp}(T)|^3}{\binom{n}{3}} \right]^s \\
\le & [K'/(sn^2)]^s |\mathrm{Bomp}(T)|^{3s}
\end{aligned}$$

for some constant $K' = K'(\Delta)$. Let $K = 5K'/3$. By Lemma 20, if $t_c(H) \leq t/10$ then we get the probability upper bound in part 1 by setting $s = 3t/5$ and if $t_c(H) > t/10$ then we get the probability upper bound in part 2 by setting $s = 2t/3 - t_c/3$ and observing that $s \geq 1$. $\square$

LEMMA 22. *Let $\Delta, \rho > 0$ be given. The probability that there exists some set T of literals and a 0-pure subformula H of F^* with $P^*(H) = T$, and $|T| \leq 10t_c(H)$ is $o(1)$ in n.*

PROOF. Let $t \leq 10t_c$ and suppose that H is a 0-pure subformula of F^* with $t_c(H) = t_c$ and $t = |P^*(H)|$. By assumption about F_2^*, $t_c \leq \log n$ so $t \leq 10 \log n$ and thus $|\text{Bomp}(P^*(H))| \leq K_0 \log^2 n$ for some constant $K_0 = K_0(\rho)$ because our assumption about the sizes of the $T_i(F_2^*)$ implies that there are no bomponents of size $\omega(\log n)$. For each $t \leq 10 \log n$ and each t_c, $t_c \leq t$, there are at most $\binom{\log n}{t_c}\binom{2n}{t-t_c}$ different sets T with $|T| = t$ containing t_c literals x with non-tree-like $\text{In}_{F_2^*}(x)$. Therefore, by Lemma 21.2, the probability that there is some 0-pure subformula H of F^* with $|P^*(H)| = t$ and $|t_c(H)| = t_c$ is at most

$$\binom{\log n}{t_c}(2n)^{t-t_c}(K/n^2)^{2t/3-t_c/3}(K_0 \log^2 n)^{2t-t_c}$$

which is $< (K'' \log n)^{4t} n^{-t/3}$ for some constant $K'' = K''(\Delta, \rho) > 0$. The total probability that an H satisfying the conditions of the lemma exists is then at most

$$\sum_{t_c=1}^{\log n} \sum_{t=t_c}^{10t_c} ((K'')^2 n^{-1/3} \log^4 n)^t \leq 10(K'')^2 n^{-1/3} \log^6 n$$

for n sufficiently large, which is $o(1)$ in n. $\square$

It will be convenient to rewrite the summations over all possible choices of set $T = P^*(H)$ with $|T| = t$ in terms of a probability calculation involving a uniformly chosen random set of literals, T, of size t.

LEMMA 23. *Let $\rho > 0$ be given. There is a constant $B = B(\rho) > 0$ such that for any $t > 0$ and for T a set of literals with $|T| = t$ chosen uniformly at random, $\mathbf{E}_T(|\text{Bomp}(T)|) \leq Bt$.*

PROOF. Let $B = \sum_{i \geq 1} i(1-\rho)^i$. By assumption, for x chosen uniformly at random from among the $2n$ possible literals, $\mathbf{E}_x(|\text{Bomp}(x)|) \leq \sum_{i \geq 1} i(1-\rho)^i = B$ and therefore $\mathbf{E}_T(|\text{Bomp}(T)|) \leq |T| \mathbf{E}_x(|\text{Bomp}(x)|) \leq Bt$. $\square$

LEMMA 24. *For every $\rho > 0$ there exists $\alpha = \alpha(\rho) > 0$ such that for all $r \geq 0$ we have for T a set of literals with $|T| = t$ chosen uniformly at random,*

$$\mathbf{Pr}_T(|\text{Bomp}(T)| > (r+4)\mathbf{E}_T(|\text{Bomp}(T)|)) < e^{-\alpha r t}.$$

PROOF. See appendix. $\square$

LEMMA 25. *Let $\rho > 0$ be given. There is a constant $K_1 = K_1(\rho)$ such that for any $t > 0$ and for a set of literals T with $|T| = t$ chosen uniformly at random $E_T(|\text{Bomp}(T)|^{9t/5}) \leq (K_1 t)^{9t/5}$.*

PROOF. Fix an integer t and consider choosing T uniformly at random with $|T| = t$. We divide up the range of possible values of $|\text{Bomp}(T)|$ into segments of size $\nu(T) = \mathbf{E}_T(|\text{Bomp}(T)|) \leq Bt$ where $B = B(\rho)$ is the constant from Lemma 23 and use our tail bounds within each segment. Therefore by Lemma 24,

$$\begin{aligned}
&\mathbf{E}_T(|\text{Bomp}(T)|^{9t/5}) \\
&\quad \leq \ 4(\mathbf{E}_T(|\text{Bomp}(T)|))^{9t/5} \\
&\qquad\qquad + \sum_{r \geq 0} \mathbf{Pr}_T(|\text{Bomp}(T)| > (r+4)\nu(T)) \\
&\qquad\qquad\qquad \times [(r+5)\nu(T)]^{9t/5} \\
&\quad \leq \ [\nu(T)]^{9t/5} \times (1 + \sum_{r \geq 0} e^{-\alpha r/2}(r+5)^{9t/5}) \\
&\quad \leq \ (Bt)^{9t/5} \times \left(\frac{K_1}{B}\right)^{9t/5} \\
&\quad \leq \ (K_1 t)^{9t/5},
\end{aligned}$$

for some $K_1 = K_1(\alpha, B) = K_1(\rho)$. $\square$

LEMMA 26. *Let Δ, ρ be given. There is a $\zeta_0 = \zeta_0(\Delta, \rho) > 0$ such that the probability that a random F^* has a 0-pure subformula H with $v = |V(H)| \leq \zeta_0 n/2$ is $o(1)$ in n.*

PROOF. By Lemmas 21.1 and 22, the total probability of this event is at most $\sum_{T, |T|=t \geq 1} R(T)$ plus a term that is $o(1)$ in n. By Lemma 25,

$$\begin{aligned}
\sum_{T, |T|=t} R(T) \ &\leq \ \binom{2n}{t}(K/(tn^2))^{3t/5} E_T(|\text{Bomp}(T)|^{9t/5}) \\
&\leq \ (2en/t)^t (K/(tn^2))^{3t/5}(K_1 t)^{9t/5} \\
&= \ ((2e)^5 K^3 K_1^9 t/n)^{t/5} \\
&\leq \ (K_2 t/n)^{t/5}
\end{aligned}$$

for some constant $K_2 = K_2(\Delta, \rho) > 0$. Now if we let $\zeta_0 = 1/(32K_2)$ then for $v \leq \zeta_0 n/2$, $t \leq v \leq n/(32K_2)$. Therefore by Lemma 21, the total probability that such an H exists is at most

$$\sum_{t=1}^{n/(32K_2)} (K_2 t/n)^{t/5} = \sum_{t=1}^{\log n} (K_2 t/n)^{t/5} + \sum_{t=\log n+1}^{n/(32K_2)} (K_2 t/n)^{t/5}.$$

The first summation totals less than $(K_2 n^{-1} \log^2 n)^{1/5}$ which is $o(1)$ in n and the second summation totals at most $\sum_{t > \log n} 2^{-t} \leq 2/n$ which is $o(1)$ in n. $\square$

Lemma 14 now follows immediately from Lemmas 18 and 26.

6. REFERENCES

[1] D. Achlioptas. A survey of lower bounds for random 3-SAT via differential equations. *Theoret. Comput. Sci., to appear.* Available as http://www.research.microsoft.com/∼optas/lbsurvey.ps.

[2] D. Achlioptas. *Threshold Phenomena in random graph coloring and satisfiability.* PhD thesis, Department of Computer Science, University of Toronto, 1999.

[3] D. Achlioptas, A. Chtcherba, I. Istrate, and C. Moore. The phase transition in 1-in-k SAT and NAE 3-SAT. In *12th Annual ACM-SIAM Symposium on Discrete Algorithms (Washington, DC, 2001)*. ACM, New York, 2001. to appear.

[4] D. Achlioptas, L. M. Kirousis, E. Kranakis, and D. Krizanc. Rigorous results for random $(2+p)$-SAT. *Theoret. Comput. Sci., to appear.*

[5] D. Achlioptas and G. B. Sorkin. Optimal myopic algorithms for random 3-SAT. In *41st Annual Symposium on Foundations of Computer Science (Redondo Beach, CA, 2000)*. IEEE, 2000. To appear.

[6] P. Beame, R. Karp, T. Pitassi, and M. Saks. On the complexity of unsatisfiability of random k-CNF formulas. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, pages 561–571, Dallas, TX, May 1998.

[7] P. W. Beame, R. Impagliazzo, J. Krajíček, T. Pitassi, and P. Pudlák. Lower bounds on Hilbert's Nullstellensatz and propositional proofs. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 794–806, Santa Fe, NM, November 1994. IEEE.

[8] E. Ben-Sasson and R. Impagliazzo. Random CNF's are hard for the polynomial calculus. In *Proceedings 40th Annual Symposium on Foundations of Computer Science*, pages 415–421, New York,NY, October 1999. IEEE.

[9] E. Ben-Sasson and A. Wigderson. Short proofs are narrow – resolution made simple. In *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing*, pages 517–526, Atlanta, GA, May 1999.

[10] A. Biere, A. Cimatti, E. Clarke, and Y. Zhu. Symbolic model checking without BDDs. In *Procceedings, 5th International Conference, TACAS'99*, pages 193–207, Berlin, Germany, 1999. Springer-Verlag.

[11] M.T. Chao and J. Franco. Probabilistic analysis of a generalization of the unit-clause literal selection heuristics. *Information Science*, 51:289–314, 1990.

[12] V. Chvátal and B. Reed. Mick gets some (the odds are on his side). In *Proceedings 33rd Annual Symposium on Foundations of Computer Science*, pages 620–627, Pittsburgh, PA, October 1992. IEEE.

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

[14] C. Coarfa, D. D. Demopoulos, A. San Miguel Aguirre, D. Subramanian, and M. Y. Vardi. Random 3-SAT: The plot thickens. In *Proceedings 6th International Conference on Principles and Practice of Constraint Programming*, Singapore, September 2000.

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

[16] M. Davis and H. Putnam. A computing procedure for quantification theory. *Communications of the ACM*, 7:201–215, 1960.

[17] E. Friedgut. Sharp thresholds of graph properties, and the k-sat problem. *Journal of the American Mathematical Society*, 12:1017–1054, 1999.

[18] A. Frieze and S. Suen. Analysis of two simple heuristics on a random instance of k- SAT. *Journal of Algorithms*, 20(2):312–355, 1996.

[19] A. Goerdt. A threshold for unsatisfiability. *Journal of Computer and System Sciences*, 53:469–486, 1996.

[20] C. Gomes, B. Selman, and H. Kautz. Boosting combinatorial search through randomization. In *Proceedings Fifteenth National Conference on Artificial Intelligence (AAAI-98)*, pages 431–437, 1998.

[21] C. P. Gomes, B. Selman, N. Crato, and H. Kautz. Heavy-tailed phenomena in satisfiability and constraint satisfaction problems. *J. Automat. Reason.*, 24(1-2):67–100, 2000.

[22] Kautz H. and Selman B. Pushing the envelope: planning, propositional logic, and stochastic search. In *Proceedings of the 13th AAAI*, pages 1194–2001, 1996.

[23] S. Janson, Y. C. Stamatiou, and M. Vamvakari. Bounding the unsatisfiability threshold of random 3-SAT. *Random Structures Algorithms*, 17(2):103–116, 2000.

[24] H. Kautz, D. McAllester, and B. Selman. Encoding plans in propositional logic. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifth International Conference (KR'96)*, pages 374–384, 1996.

[25] S. Kirkpatrick and B. Selman. Critical behavior in the satisfiability of random formulas. *Science*, 264:1297–1301, May 1994.

[26] C. J.H. McDiarmid. On the method of bounded differences. In *Surveys in Combinatorics, Proceedings of the 12th British Combinatorial Conference*, pages 148–188. 1989.

[27] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky. Phase transition and search cost in the $(2+p)$-SAT problem. In *4th Workshop on Physics and Computation, (Boston, MA, 1996).*

[28] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky. $2+p$-SAT: relation of typical-case complexity to the nature of the phase transition. *Random Structures Algorithms*, 15(3-4):414–435, 1999.

[29] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky. Determining computational complexity from characteristic phase transitions. *Nature*, 400:133–137, 1999.

[30] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical recipes in FORTRAN*. Cambridge University Press, Cambridge, second edition, 1992. The art of scientific computing, With a separately available computer disk.

[31] B. Selman and H. Kautz. Domain-independent extensions to GSAT: Solving large structured satisfiability problems. In *Proceedings of the 13th IJCAI*, pages 290–295, 1993.

[32] B. Selman, H. Levesque, and D. Mitchell. A new method for solving hard satisfiability problems. In *Proceedings Tenth National Conference on Artificial Intelligence (AAAI-92)*, pages 440–446, 1992.

[33] B. Selman, D. Mitchell, and H. Levesque. Generating hard satisfiability problems. *Artificial Intelligence*, 81:17–29, 1996.

APPENDIX

A. PROPERTIES OF SUBCRITICAL RANDOM 2-CNF FORMULAE

We will now prove that subcritical random 2-CNF formulas satisfy the properties in assumptions 1-3 of Lemma 14 w.h.p.

LEMMA 27. *1. W.h.p. for every literal ℓ, $\mathrm{In}_{F_2}(\ell)$ is simple.*

2. W.h.p. there are at most $\log^2 n$ literals ℓ, such that $\mathrm{In}_{F_2^*}(\ell)$ is not tree-like.

3. W.h.p. There exists $\rho = \rho(\epsilon)$ such that for all i, $T_i(F_2) \le 2(1-\rho)^i\, n$.

PROOF. Let $p = (1-\epsilon)/n$ and consider the random graph $G(n,p)$ on n vertices, where each of the $\binom{n}{2}$ possible edges appears independently with probability p. Fix a vertex v in $G(n,p)$ and consider the random subgraph $C(v)$ that corresponds to the connected component of v. Now, fix a literal ℓ in F_2 and consider the undirected graph $G(\mathrm{In}_{F_2}(\ell))$. It is not hard to show that one can couple the random graphs $C(v)$ and $G(\mathrm{In}_{F_2}(\ell))$ so that $G(\mathrm{In}_{F_2}(\ell)) \subseteq C(v)$ always, i.e. $G(\mathrm{In}_{F_2}(\ell))$ is always a subgraph of $C(v)$. This coupling allows us to exploit a number of well-known facts about the components of subcritical random graphs, i.e. $G(n, p = (1-\epsilon)/n)$.

(2) It is well-known that there exists $B = B(\epsilon)$ such that $\Pr[C(v)$ contains a cycle$] < B/n$. Thus, $\mathbf{E}(|Q|) < 2B$ and the claim follows from Markov's inequality.

(1) It is well-known that $\Pr[C(v)$ contains more than one cycle$] < n^{-3/2}$. Thus, $\mathbf{E}(|R|) = o(1)$.

(3) We will first establish that there exists $\sigma = \sigma(\epsilon)$ such that for any fixed literal ℓ and all q,

$$\Pr[|\mathrm{Bomp}_{F_2}(\ell)| = q] < (1-\sigma)^q \ .$$

For this we observe that for a fixed literal ℓ one can construct a random variable W that dominates $|\mathrm{Bomp}_{F_2}(\ell)|$ as follows. Let $X \stackrel{D}{=} |C(v)|$, where v is a fixed vertex in $G(n, p = (1-\epsilon)/n)$ and $|C(v)|$ denotes the number of vertices in its connected component $C(v)$. Now, let

$$W = S_1 + S_2 + \cdots + S_X$$

where the S_i are i.i.d. random variables with $S_i \stackrel{D}{=} |C(v)|$. To bound the upper tail of W we will use a standard moment generating function argument. For any $h > 0$ we have

$$\begin{aligned} \Pr[W = q] &\le \Pr[W \ge q] \\ &= \Pr[\exp(hW) \ge \exp(hq)] \\ &\le \exp(-hq)\mathbf{E}(\exp(hW)) \ , \end{aligned}$$

where for the last inequality we used Markov's (or Bernstein's) inequality. Note now that

$$\begin{aligned} \mathbf{E}(\exp(hW)) &= \sum_{k\ge 1} \Pr[X = k]\, \mathbf{E}\left(\prod_{i=1}^{k} \exp(hS_i)\right) && (3) \\ &= \sum_{k\ge 1} \Pr[X = k] \prod_{i=1}^{k} \mathbf{E}\left(\exp(hS_i)\right) && (4) \\ &= \sum_{k\ge 1} \Pr[X = k]\ \left(\mathbf{E}\left(\exp(hS_i)\right)\right)^k \ , && (5) \end{aligned}$$

where passing from (3) to (4) uses that the random variables S_i are independent while passing from (4) to (5) uses that they are identically distributed.

Let $\Pr[S_i = k] = \Pr[X = k] \equiv p_k$. Let $c = 1-\epsilon$. To bound (5) we will use that, asymptotically in n,

$$p_k = (1/c)\left(ce^{-c}\right)^k k^{k-1}/k! \tag{6}$$

and that for all $c < 1$,

$$\sum_{k\ge 1} p_k = 1 \ . \tag{7}$$

Further, let us note that the function $f(x) = xe^{-x}$ is continuous and increasing in $[0,1)$, going from $f(0) = 0$ to $f(1) = e^{-1}$. In particular for all $0 < \epsilon < 1$, if $c = 1-\epsilon$, $d = c + \epsilon^2/3$ and $h = \epsilon^3/3$, then $ce^{-c+h} < de^{-d}$ and $d < 1$. Moreover, $de^{-c} < e^{-1}$ which will be useful in bounding $\mathbf{E}(\exp(hW))$. With all this in mind,

$$\begin{aligned} \mathbf{E}(\exp(hS_i)) &= \sum_{k\ge 1} p_k \exp(hk) \\ &= (1/c)\sum_{k\ge 1} (ce^{-c+h})^k k^{k-1}/k! \\ &< (1/c)\sum_{k\ge 1} (de^{-d})^k k^{k-1}/k! \\ &= d/c \ . \end{aligned}$$

Thus,

$$\begin{aligned} \mathbf{E}(\exp(hW)) &< \sum_{k\ge 1} p_k (d/c)^k \\ &= (1/c)\sum_{k\ge 1} (de^{-c})^k k^{k-1}/k! \ < \ 1/c \ . \end{aligned}$$

Therefore, for all $\epsilon > 0$ there is $\tau = \tau(\epsilon) > 0$ and $q_0 = q_0(\epsilon)$ such that for all $q \ge q_0$,

$$\begin{aligned} \Pr[|\mathrm{Bomp}_{F_2}(\ell)| = q] &\le \Pr[W = q] \\ &< \frac{1}{1-\epsilon} \exp\left(-\frac{\epsilon^3}{3} q\right) \le (1-\tau)^q \ . \end{aligned}$$

Finally, since for any fixed value of q, $\Pr[|\mathrm{Bomp}_{F_2}(\ell)| = q] < 1$, it follows that there exists $\sigma = \sigma(\tau, q_0)$ such that for all $q \ge 1$,

$$\Pr[|\mathrm{Bomp}_{F_2}(\ell)| = q] < (1-\sigma)^q \ . \tag{8}$$

By linearity of expectation and (8) we get that for all i,

$$\mathbf{E}\left(T_i(F_2)\right) < 2n \times (1-\sigma)^i \ . \tag{9}$$

Our next step is to show that for each i there exists $Q_i \le \mathbf{E}\left(T_i(F_2)\right)$ such that w.h.p. for all i,

$$|T_i(F_2) - Q_i| < n^{3/4} \ . \tag{10}$$

Let us first observe that (8) also implies that there exists $C = C(\epsilon)$ such that w.h.p.

$$\text{for all } \ell \text{ in } F_2, \quad |\mathrm{Bomp}_{F_2}(\ell)| < C \log n \ . \tag{11}$$

To prove (10) we will need to do some work before appealing to a large deviations inequality. The reason for this is that, a priori, replacing a single clause in F_2 could change $T_i(F_2)$ dramatically, for some i; luckily, this is an unlikely event. To capture this last fact we will introduce a family of random variables U_i with the following properties: i) w.h.p. $U_i(F_2) = T_i(F_2)$ for all i, and ii) by definition (of the U_i), replacing a clause in F_2 can affect each U_i by at most polylog(n). Thus, appealing to a large deviation inequality for the U_i will yield the desired result.

The random variables U_i are motivated by the following observation.

Observation 28. *If $x \in \mathrm{Bomp}_{F_2}(y)$ then $\bar{y} \in \mathrm{Bomp}_{F_2}(\bar{x})$. Thus, if $B = \max_x |\mathrm{Bomp}_{F_2}(x)|$, then no literal appears in more than B bomponents.*

The above observation suggests that adding/removing a single arc $\vec{e}$ in $\vec{D}(F_2)$ cannot affect the bomponent size of more than $2B$ literals. To see this note that for every literal x, $\mathrm{Bomp}_{F_2}(x)$ can change only if it contains one of the two endpoints of $\vec{e}$ and, by our observation, each endpoint of $\vec{e}$ appears in at most B bomponents.

This leads us to introduce the notion of the *domponent*, $\mathrm{Domp}(\ell)$, of a literal ℓ. The domponents of all literals in a 2-SAT formula F are determined as follows. We first associate with each arc $\vec{e}$ in $\vec{D}(F)$ a *count*$(\vec{e})$ equal to the number of bomponents in which $\vec{e}$ is present ($\vec{e}$ is present in a bomponent if it was followed at least once in determining that bomponent). We then create a subgraph $\vec{D}'(F)$ of $\vec{D}(F)$ by removing all arcs $\vec{e}$ such that *count*$(\vec{e}) \geq C \log n$. The domponent of each literal ℓ is then its bomponent in $\vec{D}'(F)$. If for a literal ℓ, $\mathrm{Domp}(\ell) = \mathrm{Bomp}_{F_2}(\ell)$ then we will say that $\mathrm{Domp}(\ell)$ is *good*. Analogously to T_i we let

$$U_i(F_2) = |\{x : \mathrm{Domp}(x) = i \text{ and } \mathrm{Domp}(x) \text{ is good}\}| \ .$$

Note that for all i, by definition, $U_i(F_2) \leq T_i(F_2)$ and therefore $\mathbf{E}(U_i(F_2)) \leq \mathbf{E}(T_i(F_2))$. Further, note that by (11), w.h.p. $\vec{D}(F) = \vec{D}'(F)$. Therefore, to prove (10) it suffices to take $Q_i = \mathbf{E}(U_i(F_2))$ and prove that w.h.p. for all i, $|U_i(F_2) - \mathbf{E}(U_i(F_2))| < n^{3/4}$.

To prove that the random variables U_i are concentrated around their expectation we consider the probability space corresponding to the m_2 independent choices of clauses from $C_2(n)$ that determine F_2. We claim that for any possible set of values for these choices (i.e. for any set of clauses), changing the value of any single random variable (i.e. replacing a clause with some other clause) can only affect $\mathrm{Domp}(\ell)$ for at most $4C \log n$ literals. To prove this claim we break-down the replacement of a clause to four steps corresponding to the four arcs that are removed/added in $\vec{D}(F_2)$. The claim then follows by the fact that the removal/addition of each such arc can affect $\mathrm{Domp}(\ell)$ for at most $C \log n$ literals. This last assertion follows trivially from the fact that, by the definition of domponents, the arc $\vec{e}$ being added (removed), cannot be traversed (have been traversed) during the determination of the domponents more than $C \log n$ times.

Given the above claim, we can apply the following inequality of McDiarmid [26] to get that the probability of each D_i deviating by $n^{3/4}$ is bounded by $\exp(-n^{1/5})$. The union bound then implies that w.h.p. no D_i deviates by that much.

Theorem 29 ([26]). *Let $\mathbf{X} = (X_1, X_2, \ldots, X_n)$ be a family of independent random variables with each X_k taking values in a set A_k. Suppose that the real-valued function f defined on $\prod A_k$ satisfies*

$$|f(\mathbf{x}) - f(\mathbf{x}')| \leq c_k$$

whenever the vectors $\mathbf{x}$ and $\mathbf{x}'$ differ only in the k-th coordinate. Let μ be the expected value of the random variable $f(\mathbf{X})$. Then for any $t \geq 0$,

$$\Pr[f(\mathbf{X}) - \mu \geq t] \leq \exp\left(-2t^2 / \sum c_k^2\right) \ .$$

Combinining (9) and (10) we get that there exists $\sigma = \sigma(\epsilon) > 0$ such that w.h.p.

$$T_i \leq 2n \times (1-\sigma)^i + n^{3/4} \ . \tag{12}$$

Further, recall that by (8)

$$\text{w.h.p. } T_i = 0 \text{ for all } i \geq C \log n \ . \tag{13}$$

Let us now choose $\phi < \sigma$ such that $(1-\phi)^{C \log n} \geq n^{-1/4}$. Thus, for all $i < C \log n$

$$2n \times (1-\phi)^i \geq 2n^{3/4} \ . \tag{14}$$

We claim that w.h.p. for all i,

$$T_i \leq (4n) \times (1-\phi)^i \ . \tag{15}$$

If $i \geq C \log n$ then (15) holds by (13). If $i < C \log n$ then by (12), (15) and $\phi < \sigma$, respectively,

$$\begin{aligned} T_i &\leq 2n \times (1-\sigma)^i + n^{3/4} \\ &\leq 2n \times (1-\sigma)^i + 2n \times (1-\phi)^i \\ &\leq 4n \times (1-\phi)^i \ . \end{aligned}$$

By (12) and (15) it follows that there is $\rho < \phi$ such that w.h.p. for all i, $T_i \leq 2n \times (1-\rho)^i$. $\square$

B. PROOF OF SAMPLING LEMMA

Setting: We are given a set B of $2n$ weighted balls, each ball having weight at least 1. We do not quite know the weight distribution of the balls but, instead, we are given the following promise: there exists $\rho > 0$ such that if T_i is the number of balls with weight i, then

$$T_i \leq 2(1-\rho)^i n, \quad \text{for all } i \ . \tag{16}$$

We will pick a random set $R \subseteq B$ of αn balls, where $\alpha \in (0,1)$ is an arbitrary constant. Let

$$W = \sum_{x \in R} \mathrm{weight}(x) \ .$$

We will prove that there exists a positive function f such that for every $\xi > 0$

$$\Pr[W > 4(1+\xi)^2 \mathbf{E}(W)] < \exp(-f(\xi)n) \ ,$$

Proof: We start by considering W to be defined in the following, equivalent, manner. Let S be an infinite sequence of balls formed by choosing balls uniformly, independently and *with replacement* from B. Let W be the sum of the weights of the first αn distinct elements of S.

Let us consider the prefix $P = p_1, p_2, \ldots, p_d$ of S where $d = 2(1+\xi) \times \alpha n$. In particular, let us form a random set $R' \subseteq B$, by scanning P linearly and adding to R' every ball not seen before, until either $|R'| = \alpha n$ or we exhaust P. We let

$$\begin{aligned} W' &= \sum_{x\in R'} \text{weight}(x) \\ Q &= \sum_{i=1}^{d} \text{weight}(p_i) \ . \end{aligned}$$

Then, by (the miracle of) linearity of expectation, $\mathbf{E}(Q) = 2(1+\xi)\mathbf{E}(W)$ and, thus, for any $\xi > 0$

$$\begin{aligned} &\Pr[W > 4(1+\xi)^2\mathbf{E}(W)] \\ &\quad\le \Pr[W' > 4(1+\xi)^2\mathbf{E}(W)] + \Pr[W' \neq W] \\ &\quad\le \Pr[Q > 4(1+\xi)^2\mathbf{E}(W)] + \Pr[W' \neq W] \\ &\quad\le \Pr[Q > (2+\xi)\mathbf{E}(Q)] + \Pr[W' \neq W] \ . \end{aligned}$$

For $W' \neq W$ to occur it must be that we picked $2(1+\xi)\alpha n$ balls out of $2n$ balls with replacement and got fewer than αn distinct balls. **Using standard results for the coupons collector problem, it is not hard to show that there exists $\phi = \phi(\alpha) > 0$ such that the probability of this event is $\exp(-\phi\xi^2 n)$. We will now prove that $\Pr[Q > (2+\xi)\mathbf{E}(Q)] < \exp(-\tau\xi n)$ for some $\tau = \tau(\alpha,\rho) > 0$. Combining these two estimates the result follows.**

First, note that for any $h > 0$,

$$\begin{aligned} &\Pr[Q > (2+\xi)\mathbf{E}(Q)] \\ &\quad= \Pr[\exp(hQ) > \exp((2+\xi)h\mathbf{E}(Q))] \\ &\quad\le \mathbf{E}(\exp(hQ)) \times \exp(-(2+\xi)h\mathbf{E}(Q)) \ . \end{aligned}$$

Let $\{Q_i\}_{i=1}^d$ be i.i.d.r.v. defined by $Q_i = \text{weight}(p_i)$. Thus, $Q = \sum_{i=1}^d Q_i$ and as a result

$$\begin{aligned} \mathbf{E}(\exp(hQ)) &= \mathbf{E}\left(\prod_{i=1}^{d} \exp(hQ_i)\right) \\ &= \prod_{i=1}^{d} \mathbf{E}\left(\exp\left(hQ_i\right)\right) = \left(\mathbf{E}\left(\exp\left(hQ_i\right)\right)\right)^d \ . \end{aligned}$$

To simplify notation let us replace Q_i with T in the rest of the proof.

Let $\mu = \mathbf{E}(T)$. To go from (17) to (18) we use (16). To go from (18) to (19) we require $h < \rho$, which suffices to guarantee the sum's convergence. Finally, to go from (19) to (20) we use that for $h > 0$, $e^{-h} > 1 - h$.

$$\begin{aligned} &\mathbf{E}(\exp(hT)) \\ &\quad= \sum_{i=1}^{\infty} \Pr[T=i]\exp(hi) \\ &\quad= \sum_{i=1}^{\infty} \Pr[T=i]\left(1 + hi + (\exp(hi) - hi - 1)\right) \qquad (17) \\ &\quad\le 1 + h\mu + \sum_{i=1}^{\infty} (1-\rho)^i(\exp(hi) - hi - 1) \qquad (18) \\ &\quad= 1 + h\mu + (1-\rho)\left(\frac{1}{\rho - 1 + \exp(-h)} - \frac{h+\rho}{\rho^2}\right) \qquad (19) \\ &\quad< 1 + h\mu + (1-\rho)\left(\frac{1}{\rho - h} - \frac{h+\rho}{\rho^2}\right) \qquad (20) \\ &\quad= 1 + h\mu + \frac{h^2(1-\rho)}{\rho^2(\rho-h)} \ . \qquad (21) \end{aligned}$$

Now, substituting $h = \rho^3$ in (21) we get (22), while (23) follows from $\mu \ge 1 > (\rho+1)^{-1}$.

$$\begin{aligned} \mathbf{E}(\exp(\rho^3 T)) &< 1 + \rho^3\mu + \frac{\rho^3}{\rho+1} \qquad (22) \\ &< 1 + 2\rho^3\mu \ . \qquad (23) \end{aligned}$$

Recall now that for all $h > 0$,

$$\begin{aligned} &\Pr[Q > (2+\xi)\mathbf{E}(Q)] \\ &\quad\le \left(\frac{\mathbf{E}(\exp(hT))}{\exp((2+\xi)h\mathbf{E}(T))}\right)^{2\alpha n} \\ &\quad= \left(\frac{\mathbf{E}(\exp(hT))}{\exp(2h\mathbf{E}(T))}\right)^{2\alpha n} \times \exp(-2\alpha h\mu\xi n) \ . \qquad (24) \end{aligned}$$

Taking $h = \rho^3$, (23) implies that the ratio in (24) is bounded by 1. Thus, since $\mu \ge 1$, if $\tau = 2\alpha\rho^3$, then

$$\Pr[Q > (2+\xi)\mathbf{E}(Q)] < \exp(-\tau\xi n) \ .$$