Simplified and Improved Resolution Lower Bounds

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

Toniann Pitassi†
Computer Science Department
University of Arizona
Tucson, AZ 85721
toni@cs.arizona.edu

Abstract

We give simple new lower bounds on the lengths of Resolution proofs for the pigeonhole principle and for randomly generated formulas. For random formulas, our bounds significantly extend the range of formula sizes for which non-trivial lower bounds are known. For example, we show that with probability approaching 1, any Resolution refutation of a randomly chosen 3-CNF formula with at most $n^{6/5-\epsilon}$ clauses requires exponential size. Previous bounds applied only when the number of clauses was at most linear in the number of variables. For the pigeonhole principle our bound is a small improvement over previous bounds. Our proofs are more elementary than previous arguments, and establish a connection between Resolution proof size and maximum clause size.

1 Introduction

The importance of the satisfiability problem permeates all areas of computer science. In the last three decades, there has been a tremendous amount of research in trying to understand the mathematical structure of the satisfiability problem, and in developing algorithms for satisfiability testing, and the complementary problem of propositional theorem proving. The most well-studied and oldest class of algorithms for satisfiability testing is Resolution-based. (The Davis-Putnam procedure is a typical example.) By this, we mean that a particular deterministic implementation of Resolution is used for satisfiability testing in the following way: search for a proof of unsatisfiability; either the search procedure will get stuck and in this case we will get a satisfying assignment as a witness of satisfiability, or the search procedure will succeed, and in this case we will get a Resolution proof as a witness of unsatisfiability.

It is an empirical observation that rather straightforward implementations of Resolution work well as satisfiability testers for many 3-SAT problems. This is well-explained by some important theoretical results. Namely, it is shown in [FS] that if we generate a 3-CNF formula at random with at most $3.003n$ clauses, then with high probability f is satisfiable. (See [CRe, CF, BFU] for previous satisfiability results.) Moreover, to prove their results, [FS], as well as earlier papers, actually exhibit probabilistic, linear-time Resolution-based algorithms that will find a satisfying assignment almost certainly.

Many researchers also use rather straightforward implementations of Resolution as theorem provers, although in this case the theoretical justification is not so evident. One important test case is to consider randomly generated k-CNF formulas with larger numbers of clauses. It is not hard to show [CS] that if there are more than $2^k n \ln 2$ clauses then a random k-CNF formula is almost certainly unsatisfiable (and for 3-CNF formulas this has been improved in [FP, KMPS] with the latter showing that this holds for as little as $4.758n$ clauses.) In a beautiful paper, Chvátal and Szemerédi [CS] showed that Resolution-based theorem-provers must perform badly on such random formulas provided the number of clauses is not too large. In particular, they show that almost certainly any 3-CNF formula with $O(n)$ clauses requires an exponential length Resolution refutation (even ignoring the complexity of searching for it.) Fu [Fu] recently extended this lower bound to apply when the number of clauses is larger but for 3-CNF formulas it is no improvement on [CS].

In this paper, we give new lower bounds for Resolution refutations that are notably simpler than previous proofs. We show, among other things, that random k-

* Research supported by NSF grant CCR-9303017
† Research supported by NSF grant CCR-9457782

CNF formulas are hard for Resolution well outside the range of clause/variable ratios implied by the results of [CS]. In particular, we show that for any $\epsilon > 0$, when $m \leq n^{6/5-\epsilon}$, with high probability a random 3-CNF formula with m clauses and n variables requires an exponential size Resolution refutation.

The inspiration for our lower bounds are recent deterministic simulations of Resolution, due to Clegg, Edmonds and Impagliazzo [CEI]. These simulations are, in a sense, universal search procedures for Resolution proofs in that they work reasonably well whenever *any* Resolution-based procedure would succeed. They are obtained in two steps: First it is shown how to simulate Resolution by small-degree Gröbner proofs; secondly it is shown how to deterministically find a small degree Gröbner proof. In this paper we give direct deterministic simulations for Resolution and tree-like Resolution that can be obtained by studying their simulations. This inspires our lower bounds since proving a lower bound for Resolution becomes a matter of showing that this particular simulation cannot terminate very quickly.

The original lower bounds for general Resolution use the "bottleneck counting" argument of Haken [H]. This was introduced in [H] to show exponential lower bounds for Resolution refutations of the pigeonhole principle and further developed in [U] and [CS] to give more general bounds on Resolution refutations. (Most recently it was used to prove monotone circuit lower bounds [H2].) The bottleneck counting idea is fundamentally very simple. Each clause in the proof is viewed as allowing certain truth assignments to flow through it, namely those that it falsifies. One shows that every truth assignment must flow through some "complex" or "large" clause that only permits a small number of truth assignments to pass. Therefore the number of clauses in the refutation must be big.

However, in previous arguments this clean idea does not suffice and a more complicated form of counting is used. In our method we first apply a random restriction to kill all of the large clauses in the proof. Using simple counting, one can show that almost every small restriction will kill off all large clauses if the proof is short. Then we complete the proof with a direct argument that the remaining restricted proof cannot exist because there are no large clauses in it. Thus our argument simplifies the bottleneck method. To emphasize its simplicity we present an elementary proof of Haken's original lower bound for the propositional pigeonhole principle. As a bonus, the bound we derive is slightly better. We then apply our technique to give simpler and more widely applicable lower bounds for random k-CNF formulas.

2 Deterministic simulation of Resolution

If $A \vee l$ and $B \vee \neg\, l$ are clauses, then the clause $A \vee B$ may be inferred by the resolution rule, resolving on the literal l. A Resolution refutation of a CNF formula $C = C_1 \wedge C_2 \wedge \ldots \wedge C_m$ is a derivation of the empty clause from the clauses of C using the resolution rule. Refutations can be represented as directed acyclic graphs of in-degree two whose sources are labelled by input clauses and whose sole sink is labelled by the empty clause. Often these proofs do not reuse derived clauses and such proofs are called tree-like Resolution proofs, since the graph representation for such a proof is a tree. The size of a Resolution refutation is the number of clauses in the derivation.

Recently Clegg, Edmonds, and Impagliazzo [CEI] showed how to search deterministically for a Resolution proof of size S in time $2^{O(\sqrt{n \log S} \log n)}$. They proved this result by showing first that any Resolution proof can be simulated by a degree $O(\sqrt{n \log S})$ Gröbner proof, and then by showing how to deterministically find a degree d Gröbner proof in time $n^{O(d)}$. (They also give an $O(\log S)$ degree proof in the tree-like case.) It is possible to use this intuition to come up with a more direct deterministic simulation of Resolution.

In the tree-like case, the deterministic simulation is as follows. We first search for a literal such that there is a Refutation refutation of the clauses with that literal added in size $S/2$ (half the original size). Once one is found, we set that variable on the other side, and this gives the recurrence equation: $2nF(S/2, n) + F(S, n-1)$, where $F(S, n)$ is the time that it takes to search for a size S Resolution refutation with n underlying variables. Since the base case is 1, we get time roughly $n^{\log S}$.

Fu [Fu] has shown that any 3-CNF formula with more than $2n^2/3$ clauses has a linear size Resolution refutation. Because his refutation is tree-like, the above simulation shows that there is a quasipolynomial time algorithm that almost certainly finds refutations of 3-CNF formulas with more than $2n^2/3$ clauses.

In the general case we remove large clauses rather than high degree terms as was in done in [CEI]. We want to search for a literal x such that there is a Resolution refutation of x with at most $(1-k/2n)S$ lines of size larger than k, where k is approximately $\sqrt{n \log S}$. By an averaging argument, it is clear that such a refutation exists assuming that there is a size S proof. Once we find this literal, set it and continue. The recurrence is bounded by: $2nF((1-k/2n)S, n-1, k) + F(S, n-1, k)$, where now $F(S,n,k)$ is the time that it takes to search for a Resolution refutation with at most S lines larger than k, and n underlying variables. This is roughly $n^{\sqrt{n \log S}} F(0, n, \sqrt{n \log S})$. The base case is not 1, but instead $n^{O(\sqrt{n \log S})}$, since we simply write down all clauses of size at most $\sqrt{n \log S}$ to see if there is a proof. Thus, we end up with time $2^{O(\sqrt{n \log S} \log n)}$ to see if a size S proof exists.

The above algorithms are very simple, and they suggest a method for proving lower bounds on the size of Resolution refutations. Namely, once we have a relatively efficient, deterministic simulation for Resolution (a "universal algorithm"), in order to prove Resolution lower bounds for a particular family of formulas, we only need to show that our universal algorithm does not produce a proof very quickly. Of course, because the deterministic simulation is not tight, the lower bound achieved by this method may not be optimal.

The remainder of this paper shows that by applying this method it is possible to simplify and extend several previous lower bounds on the lengths of Resolution refutations. In order to get better lower bounds, we will not argue directly about the universal algorithm, but nonetheless, understanding why the universal algorithm cannot produce a proof very quickly is the intuition behind the new lower bounds. The overall strategy is to first apply a small restriction to eliminate all large clauses in the proof and yet without severely reducing the difficulty of refuting the restricted formula, and secondly to argue that any refutation of the restricted formula must have a large clause.

3 Lower bounds for the pigeonhole principle

We now give new lower bounds for the pigeonhole principle, $\neg PHP^n_{n-1}$. which was the first example proven hard for Resolution. In general, $\neg PHP^m_n$, with m pigeons and n holes ($m > n$) is expressed propositionally with underlying variables are $P_{i,j}$, $i \leq m$, $j \leq n$. Its clauses are: (1) $P_{i,1} \vee P_{i,2} \vee ... \vee P_{i,n}$, for each $i \leq m$; (2) $\neg P_{i,k} \vee \neg P_{j,k}$, for each $i, j \leq m$, $k \leq n$, $i \neq j$. Note that the number of clauses in $\neg PHP^m_n$ is $m + \binom{m}{2} n \leq m^3$.

As in the lower bound proof of Haken [H], a truth assignment to the underlying variables $P_{i,j}$ is *critical* if it defines a one-to-one, onto map from $n-1$ pigeons to $n-1$ holes, with the remaining pigeon not mapped to any hole. A critical assignment where i is the pigeon left out is called *i-critical.* In what follows we will only be interested in critical truth assignments. We will say that two clauses c_1 and c_2 are *equivalent with respect to critical assignments* if for every critical assignment α, $c_1(\alpha) = c_2(\alpha)$. Similarly, we will say that an inference, c_1 and c_2 imply c_3, is *sound with respect to critical assignments* if for every critical assignment α such that $c_1(\alpha) = c_2(\alpha) = 1$, it is also the case that $c_3(\alpha) = 1$.

As a first step, we will replace each clause C in the Resolution refutation by a totally monotone clause by replacing each occurrence of a negative literal $P_{i,k}$ by the set of literals $\{P_{l,k} \mid l \neq i\}$. It is easy to check that the set of monotone clauses is equivalent to the original clauses with respect to critical truth assignments and therefore, inferences using these clauses are still sound with respect to critical truth assignments. Thus in what follows, we will show that the totally monotone proof cannot be a sound refutation for all of the critical truth assignments. In our monotone proof of $\neg PHP^n_{n-1}$, define a large clause to be any clause with at least $n^2/10$ (positive) literals, i.e. one that includes at least 1/10-th of all the variables. The transformation to a monotone proof (due to Sam Buss) is not essential, but will make our argument slightly cleaner.

Assume that we have a size S (monotone) refutation of $\neg PHP^n_{n-1}$, $S < 2^{n/20}$. Then the maximal number of large clauses is S. Thus, on average, setting a single $P_{i,j}$ to 1 will set at least $S/10$ many large clauses to 1. Choose a particular $P_{i,j}$ that achieves at least the average, and set it to 1. In addition to setting $P_{i,j}$ to 1, set $P_{i,l}$, $P_{l',j}$ to zero, for all $l \neq j$, $l' \neq i$. Applying this restriction to the entire Resolution proof, leaves us with a new Resolution refutation of $\neg PHP^{n-1}_{n-2}$, where the number of large clauses is at most $9S/10$. Continue in this fashion until we have set all large clauses to 1. Applying this argument iteratively $\log_{10/9} S$ many times, we are guaranteed to have knocked out all large clauses. Thus, we are left with a Resolution refutation of $\neg PHP^{n'}_{n'-1}$, where

$$n' \geq n - \log_{10/9} S = (1 - (log_{10/9} 2)/20)n > 0.671n,$$

and where no clause in the refutation is large. But this contradicts the following lemma (originally due to Haken [H]) which states that such a refutation must have a clause of size $2(n')^2/9 > 2(0.45n^2)/9 = n^2/10$.

Lemma 1: Any Resolution refutation of $\neg PHP^n_{n-1}$

must have a clause with $2n^2/9$ literals.

Proof: Let P be a refutation of $\neg PHP^n_{n-1}$. If S is a set of clauses, we will say that S implies C on all critical assignments if: every critical assignment satisfying every clause in S also satisfies C. For each clause C, let the *complexity* of C be the minimum number of clauses in $\neg PHP^n_{n-1}$ that implies C on all critical truth assignments. Since we are considering only critical truth assignments, only the "pigeon" clauses saying that some pigeon i must be mapped to a hole will be included in a minimal set. Note that the complexity of the initial "pigeon" clauses is 1, and the complexity of the final false clause is n. By soundness, the complexity of a resolvent is at most the sum of the complexities of the two clauses from which it was derived, and therefore there must exist a clause C in the proof with $n/3 < \mathit{complexity}(C) \leq 2n/3$. We will show that C contains a large number of variables.

Let S be a minimal set of pigeon clauses in $\neg PHP^n_{n-1}$ that implies C, and let $|S| = m$. We will now show that C has at least $(n-m)m \geq 2n^2/9$ distinct literals mentioned. Fix some $i \in S$, and let α be an i-critical truth assignment falsifying C. For each $j \notin S$, consider the j-critical assignment, α', obtained from α by replacing i by j. This assignment satisfies C, and differs from α only in one place: if α mapped j to l, then α' maps i to l. Since C is monotone, it must contain the variable $P_{i,l}$. Running over all $n-m$ j's not in S (using the same α), it follows that C must contain at least $n-m$ distinct variables $P_{i,l}$, $l \leq n$. Repeating the argument for all $i \in S$ shows that C contains at least $(n-m)m$ positive literals. □

Theorem 2: For sufficiently large n, any Resolution proof of $\neg PHP^n_{n-1}$ requires size $2^{n/20}$.

We note that this improves somewhat upon Haken's bound of $2^{n/577}$ although our major interest is in its simpler proof rather than in the better size bound. Buss and Turán [BT] extend Haken's argument to show that $\neg PHP^m_n$ requires superpolynomial size Resolution lower bounds as long as $m < n^2/\log n$. The argument presented here can be modified to re-prove their result.

4 Lower Bounds for Random k-CNF formulas

Chvátal and Szemerédi applied the bottleneck counting technique of Haken [H] as formalized by Urquhart [U] to prove that random k-CNF formulas with at most cn clauses, for any constant c, require exponential size Resolution proofs. Their proof technique first proves that almost certainly a hypergraph defined by a random k-CNF formula satisfies a certain sparseness property with two different choices of parameters. Using the property with these parameters they then define a mapping between certain "special pairs", each consisting of a set of variables and a restriction defined over a large portion of that set, and "complex clauses" in the proof. By computing a lower bound on the number of special pairs and an upper bound on the number of pairs associated with each complex clause they determine a lower bound on the number of complex clauses in the proof and therefore on the total size of the proof. Xudong Fu [Fu] used the same technique and extended the range of number of clauses for which exponential lower bounds apply to $m \leq n^{(k-1)/4}$. (For $k \leq 5$ this is no improvement.)

We now give a simpler proof of the lower bound for random k-CNF formulas. Our overall strategy is the same as in the previous section. That is, we will first choose a restriction to remove all large clauses, and then argue that the restricted formula is still random enough that any proof of it must still contain a large clause, hence a contradiction. The latter property is proven by modifying the sparseness property of Chvátal and Szemerédi . A nice side-effect of our simplification is that we obtain exponential lower bounds for a much greater range of values of m.

DEFINITION 4.1: A CNF formula F is n'-sparse if every set of $s \leq n'$ variables contains at most s clauses of F.

Proposition 3: If CNF formula F is n'-sparse then every subset of up to n' clauses from F is simultaneously satisfiable.

Proof: Let T be a set of clauses of F with $|T| = n'$. By the definition of n'-sparsity, every set of clauses $S \subseteq T$ contains at least $|S|$ different variables. By Hall's Theorem we can choose a system of distinct representative variables, one for each clause of T. We satisfy the clauses of T by setting the representative variable of each clause to satisfy the clause. (Note that the n'-sparsity of F implies that there is no clause of size 0.) □

DEFINITION 4.2: Let $n' < n''$. A CNF formula F in n variables is (n', n'', y)-sparse if every set of s variables, $n' < s \leq n''$, contains at most ys clauses.

The *boundary*, $b(S)$, of a set S of clauses is the set of variables that appear in only one clause of S. The following proposition is essentially from Chvátal and Szemerédi .

Proposition 4: Let F be a CNF formula with clause size at most k and suppose that F is $(n'(k+\epsilon)/2, n''(k+\epsilon)/2, 2/(k+\epsilon))$-sparse. Then every set S of ℓ clauses of F, with $n' < \ell \leq n''$, has a boundary of size at least $\epsilon\ell$.

Proof: Let S be a set of ℓ clauses from F, $n' < \ell \leq n''$. Suppose that S has a boundary of size less than $\epsilon\ell$. There are at most $k\ell$ occurrences of variables among the clauses of F. Then the maximum number of different variables appearing in S is less than $\epsilon\ell + (k\ell - \epsilon\ell)/2 \leq k\ell/2 + \epsilon\ell/2 \leq (k+\epsilon)\ell/2$ since each boundary variable occurs once and every one of the remaining variables occurs at least twice. However this contradicts the assumption that F is $(n'(k+\epsilon)/2, n''(k+\epsilon)/2, 2/(k+\epsilon))$-sparse. □

Lemma 5: [Complex Clause Lemma] Let $n' \leq n$ and F be an unsatisfiable CNF formula in n variables with clauses of size at most k that is both n'-sparse and $(n'(k+\epsilon)/4, n'(k+\epsilon)/2, 2/(k+\epsilon))$-sparse. Then any Resolution proof P of the unsatisfiability of F must include a clause of length at least $\epsilon n'/2$.

Proof: Let F be a CNF formula satisfying the conditions of the lemma and let P be a Resolution refutation of F. If S is a set of clauses, we say that S implies clause C if every truth assignment satisfying the conjunction of clauses in S also satisfies C. For each clause C in P, let the *complexity* of C be the minimum number of clauses of F that implies C.

Since F is n'-sparse, by Proposition 3 any subset of at most n' clauses of F is satisfiable. Therefore the complexity of the empty clause is $> n'$. Since the complexity of a resolvent is at most the sum of the complexities of the two clauses from which it is derived, there must exist a clause C in the proof with $n'/2 < complexity(C) \leq n'$. We will show that C contains a large number of variables.

Let S, $n'/2 < |S| \leq n'$, be a set of clauses of F witnessing the complexity of C. By Proposition 4 and the fact that F is $(n'(k+\epsilon)/4, n'(k+\epsilon)/2, 2/(k+\epsilon))$-sparse, S has a boundary $b(S)$ of size at least $\epsilon|S| > \epsilon n'/2$. It suffices to prove that C contains all the variables in $b(S)$.

Let x be an element of $b(S)$ and let C' be the unique clause of S containing x. By definition of S, the clauses in $S - \{C'\}$ does not imply C but S does imply C. Therefore there is some assignment to the variables of S and C such that all clauses in $S - \{C'\}$ are true but C' and C are false. If we modify this assignment by toggling the truth value of x in order to satisfy C' then we have an assignment that satisfies all clauses of S and therefore satisfies C by definition. We have only modified the truth value of x and have changed the truth value of C. Therefore C contains x. □

Lemma 6: Let P be a Resolution refutation of f of size S. The large clauses of P are those clauses mentioning more than an distinct variables. With probability greater than $1 - 2^{1-at/4}|S|$, a random restriction of size t sets all large clauses in S to 1.

Proof: Let C be a large clause of P. The expected number of variables of C assigned values by a randomly chosen restriction of size t is $tan/n = at$. Let D be the random variable representing the domain of ρ. By Chernoff-Hoeffding bounds on the tail of the hypergeometric distribution we have

$$\Pr[|C \cap D| \leq at/4] \leq (\sqrt{2}/e^{3/4})^{at} \leq 2^{-at/2}.$$

Also, given that $|C \cap D| = s$, the probability that $C\lceil_\rho$ is not set to 1 is 2^{-s}. Therefore the probability that $C\lceil_\rho$ is not 1 is at most $2^{-at/2} + 2^{-at/4} < 2^{1-at/4}$. Thus, the probability that some large clause of P is not set to 1 is less than $2^{1-at/4}\, S$. □

Lemma 7: Let $x > 0$, $1 \geq y > 1/(k-1)$, and $z \geq 4$. Fix any restriction ρ on $t \leq \min\{xn/2, x^{1-1/y(k-1)}n^{1-1/(k-1)}/z\}$ variables. If F is chosen as a random k-CNF formula in n variables with $m \leq \frac{y}{e^{1+1/y}2^{k+1/y}} x^{1/y-(k-1)} n$ clauses then, with probability at least $1 - 2^{-t} - (2^k+1)/z^{k-1}$, $F\lceil_\rho$ is both $(xn/2, xn, y)$-sparse and xn-sparse.

Proof: The argument is similar to the proof of (x,y)-sparsity of F in [CS] with two exceptions. Firstly, we do not have to worry about small sets for $(xn/2, xn, y)$-sparsity and secondly, we have to take into account effects of the restriction ρ.

Let S be a fixed subset of the n variables of size $s \leq xn$. Let p' be the probability that a randomly chosen k-clause C is such that $C\lceil_\rho \neq 1$ and all variables in $C\lceil_\rho$ are contained in S. In order for this to happen all of the variables of C must lie either in S or in the domain D of ρ. Therefore, in particular, $p' \leq \binom{s+t}{k}/\binom{n}{k} \leq (s+t)^k/n^k$. Let $p = (s+t)^k/n^k$. The m clauses of F are chosen independently. Therefore the distribution of the number of clauses of $F\lceil_\rho$ lying in S is the binomial distribution $B(m,p')$. The probability that more than ys clauses of $F\lceil_\rho$ lie in S is then

$$\Pr[B(m,p') \geq ys] \leq \Pr[B(m,p) \geq ys].$$

By Chernoff bounds on the tail of the binomial distribution this probability is bounded above by

$$(\frac{epm}{ys})^{ys} \leq (\frac{e(s+t)^k m}{ysn^k})^{ys}.$$

There are $\binom{n}{s} \leq (ne/s)^s$ different sets S of size s so the probability that some set S of size s contains more than ys clauses is at most

$$(\frac{ne}{s})^s(\frac{e(s+t)^k m}{ysn^k})^{ys} = (\frac{e^{1+1/y}(s+t)^k m}{ys^{1+1/y}n^{k-1/y}})^{ys}. \qquad (1)$$

Now for $t < s$, $s+t \leq 2s$; therefore, (1) is at most

$$(\frac{e^{1+1/y}2^k s^{k-1-1/y}m}{yn^{k-1/y}})^{ys}.$$

Since $s \leq xn$ this is at most

$$(\frac{e^{1+1/y}2^k x^{k-1-1/y}m}{yn})^{ys} \leq 2^{-s}$$

for $m \leq \frac{y}{e^{1+1/y}2^{k+1/y}}x^{1/y-(k-1)}n$.

Thus the total probability that some set S of size s, $t < s \leq xn$, has more than ys clauses is $< \sum_{s=t+1}^{xn} 2^{-s} < 2^{-t}$. Therefore we have that F is (t, xn, y)-sparse with probability $\geq 1-2^{-t}$. Since $t \leq xn/2$, with at least this probability, F is $(xn/2, xn, y)$-sparse.

Now we need to show that F is xn-sparse. Clearly the fact that F is (t, xn, y)-sparse and $y \leq 1$ implies that no set of size s with $t < s \leq xn$ can contain more than s clauses. It remains to handle sets of size s with $s \leq t$. We consider two cases separately depending on whether or not F contains a clause entirely in the domain D of ρ.

We first assume that no clause of F is entirely contained in D. In this case, the empty clause cannot be generated so the only sets S to worry about are of size s, $0 < s \leq t$. Therefore, the only clauses of F that can become clauses of $F\lceil_\rho$ lying in S must have at least one point in S. There are at most $s\binom{s+t-1}{k-1}$ sets of size k with all points in $S \cup D$ and at least one point in S. Thus the probability that a clause of $F\lceil_\rho$ lies in S in this case is at most

$$\frac{s\binom{s+t-1}{k-1}}{\binom{n}{k}} = \frac{sk\binom{s+t-1}{k-1}}{n\binom{n-1}{k-1}} \leq \frac{sk(s+t)^{k-1}}{n^k}$$

Using this estimate in place of $(s+t)^k/n^k$ in (1) and 1 in place of y, the probability that some set S of size s has more than s clauses of $F\lceil_\rho$ is at most

$$(\frac{e^2k(s+t)^{k-1}m}{sn^{k-1}})^s.$$

For $1 \leq s \leq t$, this is at most

$$(\frac{e^2k2^{k-1}t^{k-1}m}{sn^{k-1}})^s \leq (\frac{e^2k2^{k-1}t^{k-1}m}{n^{k-1}})^s.$$

Now

$$\begin{aligned}\frac{e^2k2^{k-1}t^{k-1}m}{n^{k-1}} &\leq \frac{kt^{k-1}x^{1/y-(k-1)}}{n^{k-2}} \\ &\leq (\frac{2t}{x^{(k-1-1/y)/(k-1)}n^{1-1/(k-1)}})^{k-1}.\end{aligned}$$

The bound on t implies that the total failure probability for all sets of size s, $1 \leq s \leq t$, is at most $\sum_{s=1}^{t}[2/z]^{(k-1)s} < 2(2/z)^{k-1}$ since $z \geq 4$.

The probability of the second case, i.e. that some clause of F is entirely contained in D, is at most $m\binom{t}{k}/\binom{n}{k} \leq mt^k/n^k < mt^{k-1}/n^{k-1} \leq 1/z^{k-1}$ by the above calculation.

Therefore the total probability that $F\lceil_\rho$ fails to be both $(xn/2, xn, y)$-sparse and xn-sparse is at most $2^{-t} + (2^k+1)/z^{k-1}$. $\Box$

We can now put this all together. Namely, we argue that there is some restriction ρ with the following properties:

(A) Most unsatisfiable formulas with short Resolution refutations have no long clauses in these refutations after ρ is applied to them.

(B) With very high probability, a random formula is satisfiable or requires a refutation with long clauses after ρ is applied.

We then conclude that almost no random formulas can be unsatisfiable and have short Resolution refutations.

Theorem 8: Let $k \geq 3$, $1 > \epsilon > 0$, $y = 2/(k+\epsilon)$, and x, t, z be functions of n such that t and z are $\omega(1)$, and t satisfies the conditions of Lemma 7 for all sufficiently large n. Then with probability approaching 1 as n approaches infinity, a randomly chosen k-CNF formula in n variables with $m \leq \frac{1}{2^{7k/2}}x^{-(k-2-\epsilon)/2}n$ clauses, does not have a Resolution refutation of size $\leq 2^{\frac{\epsilon}{4(k+\epsilon)}xt}/8$.

Proof: Let $S = 2^{\frac{\epsilon}{4(k+\epsilon)}xt}/8$. If a k-CNF formula F is satisfiable then no Resolution refutation for F exists.

Let U be the set of unsatisfiable k-CNF formulas with m clauses in n variables. For each formula $F \in U$ fix some shortest Resolution refutation P_F of F. Consider the set $B \subset U$ of those formulas F that have P_F of size at most S. We will argue that the formulas in B form a negligible fraction of all k-CNF formulas with m clauses on n variables.

By Lemma 6, for any fixed k-CNF formula F in B, the fraction of restrictions ρ which set t variables such that $P_F\lceil_\rho$ contains a clause of length at least $\epsilon xn/(k+\epsilon)$ is at most $\alpha = 2^{1-\frac{\epsilon}{4(k+\epsilon)}xt}S \le 1/4$.

For F in U, call a (ρ, F) pair bad if $P_F\lceil_\rho$ contains a large clause, i.e. one of size $\ge \epsilon xn/(k+\epsilon)$. The total fraction of bad (ρ, F) pairs in B is at most $1/4$. Therefore the fraction of ρ such that (ρ, F) is bad for at least $1/2$ of the F in B is at most $1/2$ by Markov's inequality. Fix some ρ for which less than $1/2$ of the F in B have a clause of length $\ge \epsilon xn/(k+\epsilon)$ in $P_F\lceil_\rho$.

Now since $z(n)$ is $\omega(1)$, for sufficiently large n it is ≥ 4. Also observe that $k-1-1/y = (k-2-\epsilon)/2$ and that $2^{-7k/2} \le y/(e^{1+1/y}2^{k+1/y})$. Therefore for $m \le \frac{1}{2^{7k/2}}x^{-(k-2-\epsilon)/2}n$ all the conditions of Lemma 7 are satisfied for y, z, t, and m. Therefore by Lemma 7 and the fact that both t and z are $\omega(1)$, the probability that for a random k-CNF formula F with m clauses in n variables $F\lceil_\rho$ fails to be both $(xn/2, xn, 2(k+\epsilon))$-sparse and xn-sparse goes to 0 as n goes to infinity. Therefore almost all formulas F are either satisfiable (and not in U) or have these sparseness properties.

Since xn-sparsity implies $2xn/(k+\epsilon)$-sparsity we can apply Lemma 5 with $n' = 2xn/(k+\epsilon)$ to derive that almost all F are either satisfiable (and thus not in U) or have a clause of length at least $\epsilon xn/(k+\epsilon)$ in $P_F\lceil_\rho$. Since $B \subset U$ and at least $1/2$ of the formulas F in B do not have such a large clause in $P_F\lceil_\rho$ the total measure of B is negligibly small. That is, almost all k-CNF formulas in n variables with m clauses do not have Resolution refutations of size at most S. $\square$

Theorem 9: Let $k \ge 3$ and $0 < \epsilon < 1$.

(a) If $v(n) \in o(n^{(k-4+\epsilon)/(k+\epsilon)})$ then a negligible fraction of all k-CNF formulas in n variables with at most
$$\frac{n^{(k+2-\epsilon)/4}}{2^{4k}\ v(n)^{(k-2-\epsilon)/4}}$$
clauses have Resolution refutations of size at most $2^{\frac{\epsilon}{4(k+\epsilon)}v(n)}/8$.

(b) If $v(n) \in \Omega(n^{(k-4+\epsilon)/(k+\epsilon)})$ then a negligible fraction of all k-CNF formulas in n variables with
$$o(n^{\frac{k^2-k}{3k-4}-\frac{\epsilon}{2}}/v(n)^{\frac{k-1}{3}})$$
clauses have Resolution refutations of size at most $2^{\frac{\epsilon}{4(k+\epsilon)}v(n)}/8$.

Proof: Let $y = 2/(k+\epsilon)$.

Suppose that $v(n) \in o(n^{(k-4+\epsilon)/(k+\epsilon)})$ and define $x(n) = \sqrt{2v(n)/n}$, and $t(n) = x(n)\cdot n/2 = \sqrt{v(n)\cdot n/2}$. Then
$$\begin{aligned} x(n)^{-1/y} &= (\frac{n}{2v(n)})^{(k+\epsilon)/4} \\ &\in \omega(n^{(k+\epsilon)/4}/n^{(k-4+\epsilon)/4}) \\ &\subseteq \omega(n) \end{aligned}$$
by the condition on $v(n)$. Thus
$$\begin{aligned} x(n)^{1-1/y(k-1)}n^{1-1/(k-1)} &\in \omega(x(n)\cdot n) \\ &\subseteq \omega(t(n)) \end{aligned}$$
and so $t(n)$ satisfies the conditions of Lemma 7 for some function $z(n)$ that is $\omega(1)$. Therefore Theorem 8 implies that for m at most
$$2^{-7k/2}x^{-(k-2-\epsilon)/2}n > 2^{-4k}n^{(k+2-\epsilon)/4}/v(n)^{(k-2-\epsilon)/4}$$
a random k-CNF formula with m clauses almost certainly does not have a Resolution refutation of size at most $2^{\frac{\epsilon}{4(k+\epsilon)}xt}/8 = 2^{\frac{\epsilon}{4(k+\epsilon)}v(n)}/8$ as required for part (a).

Now suppose that $v(n) \in \Omega(n^{(k-4+\epsilon)/(k+\epsilon)})$. and let m be $o(n^{\frac{k^2-k}{3k-4}-\frac{\epsilon}{2}}/v(n)^{\frac{k-1}{3}})$. Observe that, since $\frac{k^2-k-(k-1)\epsilon}{3k-4-\epsilon} \ge \frac{k^2-k}{3k-4} - \frac{\epsilon}{2}$ and $\frac{(k-1)(k-2-\epsilon)}{3k-4-\epsilon} \le \frac{k-1}{3}$, m is $o(q)$ where
$$q(n) = n^{\frac{k^2-k-(k-1)\epsilon}{3k-4-\epsilon}}/v(n)^{\frac{(k-1)(k-2-\epsilon)}{3k-4-\epsilon}}.$$
Define $z(n)$ to be $(2^{-7k/2}q(n)/m)^{3/(k-1)}$ which is $\omega(1)$ because m is $o(q)$.
Now let
$$x(n) = n^{-2(k-2)/(3k-4-\epsilon)}[v(n)z(n)]^{2(k-1)/(3k-4-\epsilon)}$$
and define
$$t(n) = x(n)^{1-1/y(k-1)}n^{1-1/(k-1)}/z(n).$$
Note that
$$\begin{aligned} xt &= x^{1+(k-2-\epsilon)/(2(k-1))}n^{1-1/(k-1)}/z(n) \\ &= x^{(3k-4-\epsilon)/(2(k-1))}n^{1-1/(k-1)}/z(n) \\ &= n^{-(k-2)/(k-1)}[v(n)z(n)]n^{1-1/(k-1)}/z(n) \\ &= v(n). \end{aligned}$$

Also note that the condition on $v(n)$ implies that

$$\begin{aligned} x(n)^{-1/y} &\in o([n^{\frac{2(k-2)}{3k-4-\epsilon}}/v(n)^{\frac{2(k-1)}{3k-4-\epsilon}}]^{\frac{k+\epsilon}{2}}) \\ &\subseteq o(n^{\frac{(k-2)(k+\epsilon)}{3k-4-\epsilon}}/n^{\frac{(k-1)(k-4+\epsilon)}{3k-4-\epsilon}}) \\ &= o(n^{\frac{k^2+(\epsilon-2)k-2\epsilon)}{3k-4-\epsilon}}/n^{\frac{k^2+(\epsilon-5)k+4-\epsilon)}{3k-4-\epsilon}}) \\ &= o(n). \end{aligned}$$

Thus

$$\begin{aligned} t(n) &= x(n)\cdot n\cdot x(n)^{-1/y(k-1)}n^{-1/(k-1)}/z(n) \\ &\in o(x(n)\cdot n) \end{aligned}$$

and $t(n)$ satisfies the conditions of Lemma 7. Observe also that

$$\begin{aligned} &2^{-7k/2}x^{-(k-2-\epsilon)/2}n \\ &= (n^{1+\frac{(k-2)(k-2-\epsilon)}{3k-4-\epsilon}})/(2^{7k/2}[v(n)z(n)]^{\frac{(k-1)(k-2-\epsilon)}{3k-4-\epsilon}}) \\ &= q(n)/(2^{7k/2}z(n)^{\frac{(k-1)(k-2-\epsilon)}{3k-4-\epsilon}}) \\ &> q(n)/(2^{7k/2}z(n)^{(k-1)/3}) \\ &= q(n)/(2^{7k/2}(2^{-7k/2}q(n)/m)) = m. \end{aligned}$$

Thus by Theorem 8, a random k-CNF formula with m clauses almost certainly does not have a Resolution refutation of size at most $2^{\frac{\epsilon}{4(k+\epsilon)}xt}/8 = 2^{\frac{\epsilon}{4(k+\epsilon)}v(n)}/8$ as required for part (b). □

Corollary 10: Let $\epsilon > 0$.

(a) For $k \geq 4$, almost all k-CNF formulas in n variables with at most $n^{(k+2)/4-\epsilon}$ clauses do not have Resolution refutations of less than exponential size.

(b) Almost all 3-CNF formulas in n variables with at most $n^{6/5-\epsilon}$ clauses do not have Resolution refutations of less than exponential size.

Proof: These both follow by choosing $v(n) = n^{\epsilon/(k+1)}$ and applying the corresponding case of Theorem 9. □

5 Further Research

An open problem is to prove exponential lower bounds for the weak pigeonhole principle, $\neg PHP_n^m$, where the number of pigeons, m, is large (say n^3). Buss and Pitassi [BP] showed that there exists a Resolution refutation of $\neg PHP_n^m$ when $m \geq 2^{\sqrt{n}\log n}$ of size $2^{\sqrt{n}\log n}$. However, when m is polynomial in n, we conjecture that any Resolution refutation of $\neg PHP_n^m$ requires superpolynomial size.

It also remains to close the gap between the range of the number of clauses where we have nontrivial lower bounds for random k-CNF formulas and those where we have good upper bounds. The best upper bounds for random k-CNF formulas currently known are due to Fu [Fu] who shows that there is a constant c_k such that random k-CNF formulas with at least $c_k n^{k-1}$ clauses have short Resolution proofs almost certainly. It would be interesting to know whether or not this property has a strong threshold behavior.

Lastly, it would be interesting to show that our lower bound method is universal in the sense that *any* formula requiring an exponential size Resolution proof can be proven intractable using our method. This would follow if one could show the following proposition: for any 3CNF formula f, if f has a polynomial-size Resolution refutation, then f also has a Resolution refutation with maximum clause size $\sqrt{n}$. Such a result would also justify the simple and natural deterministic simulation of Resolution whereby we exhaustively search for proofs of maximum clause length i, for increasing i.

6 Acknowledgments

We thank Sam Buss and Uwe Schoening for their helpful comments on this work.

References

[BFU] Broder, A., Frieze, A., and Upfal, E., "On the satisfiability and maximum satisfiability of random 3-CNF formulas," Proceedings of the Fourth ACM-SIAM Symposium on Discrete Algorithms, 1993.

[B] Buss, S. "Polynomial size proofs of the propositional pigeonhole principle," *Journal of Symbolic Logic*, v. 52 (1987), pp. 916-927.

[BP] Buss, S., and Pitassi, T., "Resolution and the weak pigeonhole principle," manuscript, 1996.

[BT] Buss, S., and Turán , G., "Resolution proofs of generalized pigeonhole principles," *Theoretical Computer Science*, 62 (1988) pp. 311-317.

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

[CEI] Clegg, M., Edmonds, J., and Impagliazzo, J., "Using the Gröbner basis algorithm to find proofs of unsatisfiability," Proceedings of the 28th ACM Symposium on Theory of Computing, 1996, pp. 174-183.

[CRe] Chvátal , V., Reed, B., "Mick gets some (the odds are on his side)," Proceedings of the 33rd IEEE

Symposium on Foundations of Computer Science, pp. 620-627, 1992.

[CS] Chvátal , V., Szemerédi , E., "Many hard examples for Resolution," *Journal of the Association for Computing Machinery*, Vol. 35, pp. 759-768, 1988.

[CR] Cook, S., Reckhow, R., "The relative efficiency of propositional proof systems," *Journal of Symbolic Logic*, 44, 1979, pp. 36-50.

[FP] Franco, J., and Pauli, M., "Probabilistic analysis of the Davis Putnam procedure for solving the satisfiability problem," *Discrete Applied Mathematics*, 5, 1983, pp.77-87.

[FS] Frieze, A., and Suen, S., "Analysis of simple heuristics for random instances of 3SAT," Manuscript, 1993.

[Fu] Fu, X. "On the complexity of proof systems," PhD Thesis, University of Toronto, 1995.

[H] Haken, A. "The intractability of Resolution," *Theoretical Computer Science*, 39, 1985. pp. 297-308.

[H2] Haken, A. "Counting bottlenecks to show monotone $P \neq NP$," Proceedings of 36th Symposium on Foundations of Computer Science, 1995, pp. 36-40.

[KMPS] Kamath, A., Motwani, R., Palem, K., and Spirakis, P., "Tail bounds for occupancy and the satisfiability threshold conjecture," Proceedings from the 34th IEEE Symposium on Foundations of Computer Science, 1994, pp. 592-603.

[T] Tseitin, G.S. "On the complexity of derivations in the propositional calculus," *Studies in Mathematics and Mathematical Logic, Part II*, A.O. Slisenko, ed., 1970, pp. 115-125.

[U] Urquhart, A., "Hard examples for Resolution," *Journal of ACM*, Vol. 34, 1987, pp. 209-219.