Title
Exponential Separation of Res(k) and Res(k+1)

Authors
Buss, Sam
Impagliazzo, Russell
Segerlind, Nathan

Publication Date
2002-01-11

Peer reviewed

Exponential separation of Res(k) and Res($k+1$)

Sam Buss and Russell Impagliazzo and Nathan Segerlind

January 10, 2002

Abstract

For each $k \geq 1$, we give a famiily of unsatisfiable sets of clauses which have polynomial size Res($k+1$) refutations, but which require Res(k) refutations of $2^{n^{\epsilon_k}}$. This improves the superpolynomial- separation between resolution and Res(2) given by Bonet and Galesi to exponential. As a corollary, we obtain an exponential separation between depth 0 Frege and depth 1 Frege, improving upon the superpolynomial separation given by the weak pigeonhole principle.

1 Introduction

The central question of propositional proof complexity is "Given a propositional proof system, can we show that there are tautologies which require large proofs in that system?" This problem has consequences for both algorithm analysis and complexity theory. Many algorithms for problems such as graph coloring and independent set work by implicitly generating derivations in certain proof systems, see [21] and [4]. Lower bounds for their corresponding proof systems yield lower bounds for the running times of these algorithms. In the seminal work of Cook and Reckhow, [8], it is shown that there is a proof system which has a polynomial size proof for every tautology if and only if NP equals coNP.

It seems that before we are able to prove lower bounds for arbitrary systems, we should be able to prove lower bounds for particular, familiar systems such as the elementary textbook propositional proof systems: the Frege systems. Unfortunately, Frege systems yet seem resistant to our efforts for proving lower bounds, so we narrow our attention to Frege proofs in which the lines of the proofs are restricted to a weak circuit class.

The popular resolution proof system can be viewed in this way, it is essentially a Frege refutation system in which every line of the proof is a clause. The first resolution bounds, for the propositional pigeonhole principle, were shown by Haken, [11], and improved by Buss and Turán in [7]. More recent advances in this technique have used the methods of relationships between random restriction and width, [3], as well as and size-width trade-offs, [5].

There has also been progress for bounded-depth Frege systems: the method of random restrictions used in circuit complexity to prove lower bounds for bounded depth circuits with unbounded fan-in AND and OR gates ([9], [12], [2]) can also prove size lower bounds for proof systems which use such bounded-depth formulas ([20], [13]). However, the techniques for studying the fine structure of bounded-depth circuits have not been as successful for studying the fine structure of bounded-depth Frege proofs.

Despite the random restriction techniques of [12] which established an exponentials separation between depth d circuits and depth $d+1$ circuits, researchers have yet been able only to establish only superpolynomial, not exponential, separation between depth d and depth $d+1$

Frege systems. Moreover, there has been little understanding of how bottom fan-in affects the abilities of a proof system. Little has been known even for augmentations of resolution which allow the formation of conjunctions of a small number of literals.

Extensions of resolution which allow the formation of small conjuncts, called Res(k) systems, have been studied implicitly algorithmic work and explicitly in proof theoretic work. The graph 3-coloring algorithm of Beigel and Eppstein, [4] can be viewed as generating Res(2) derivations because it branches on conjunctions of two variables. Res(k) systems have also received theoretical attention in the work of [1], [15].

Our main theorem is an exponential separation between resolution which allows k-conjuncts and resolution which allows $(k+1)$-conjuncts.

The sets of clauses we use to obtain the separations are based on the following combinatorial fact: For a finite, undirected graph G, it is impossible to put a partial order on the vertices of G so that every vertex is preceded by one of its neighbors. These are modified versions of the GT_n tautologies used by Goerdt,[10], and Bonet and Galesi, [6]. The differences are that we do not require the order to be linear, and that we require each node to have a predecessor which is adjacent to it in the given graph. The important properties about these principles, which we call "graph ordering principles" and write as "GOP(G)", where G is a fixed graph, are that for certain graphs they have polynomial size resolution refutations but require linear positive width.

From this starting point, we replace the variables of the $GOP(G)$ principles with k-conjunctions of new variables and obtain principles we call $GOP^k(G)$. If the degree of G is low, it turns out that the $GOP(G)$ principle will have low width and the $GOP^k(G)$ principle will therefore have small size. However, refuting it while allowing only width $k-1$ conjuncts requires exponential size.

Theorem 1 *For every k, there exists a family of graphs so that for sufficiently large n, the sets of clauses $GOP^{k+1}(G)$ have size $n^{O(1)}$, require size $2^{n^{\Omega(1)}}$ refutations in Res(k), but have size $n^{O(1)}$ refutations in Res($k+1$).*

As far as the authors know, this is the first separation between Res(k) and Res($k+1$) for $k \geq 2$, and it improves upon the previous separation of [1] between Res(1) and Res(2), which was superpolynomial but not exponential.

As a consequence of our separation between Res(1) and Res(2) we obtain an *exponential* separation between depth 0 versus depth 1 Frege.

Theorem 2 *The principles $GOP^2(G)$ require exponential size depth-0 Frege proofs, but have polynomial size depth-1 Frege proofs.*

The separation given by the weak pigeonhole principle, cn into n, is not of exponential quality. The weak pigeonhole principle for cn into n has size $n^{O(\log n)}$ proofs in depth 1 Frege, [19], [16], but it requires size $2^{\Omega(n)}$ proofs in depth 0 Frege, [11], [7], [3]. These bounds are consistent with the possibility that depth 0 Frege might be able to simulate a size S depth 1 Frege proof with size $2^{2^{O(\sqrt{\log S})}}$.

A likely consequence of our separation is the possibility of improving the superpolynomial separation between depth d and depth $d+1$ Frege to an exponential separation. In [14], Krajicek shows that the weak-pigeonhole principle for depth d Sipser functions requires requires size $2^{n^{\Omega(1)}}$ to refute in detph d Frege, but has size $n^{O(\log n)}$ proofs in depth $d+1$ Frege. As in the preceding paragraph, this separation is not of exponential quality. An investigation of the proof reveals that the cause of this shortcoming is the quasi-polynomial upper bound for the weak pigeonhole principle.

It is our belief that the techniques of [14] and [12] can be adopted to establish that the GOP^2 principles in which every variable is replaced by a depth d Sipser function provide a truly exponential separation between depth d and depth $d+1$ Frege.

Conjecture: There are tautologies τ_n of size $n^{O(1)}$ which have $n^{O(1)}$ size proofs in depth $d+1$ Frege but which require $2^{n^{\Omega(1)}}$ size proofs in depth d Frege.

Section 2 defines the proof systems we are studying and our notational conventions. Section 3 defines the CNFs for which we prove the separation, section 4 gives the upper bounds and section 5 gives the lower bounds.

The proof of the lower bound builds upon the familiar idea of "prove a width lower bound, apply a random restriction and, if the proof is small, obtain a narrow proof" used in [3]. For clauses, it usually works out that if the clause contains a large number of variables, then it is satisfied with very high probability because the events of each literal being satisfied are independent. For k-DNFs, this is not the case. Distinct terms may share variables, and therefore the probability of satisfying a large k-DNF may be bounded away from 1. However, such correlation allows us to make use of the distributive rule to simplify the k-DNF into a relatively small number of $(k-1)$-DNFs. This relation between the likelihood of satisfaction and the number of $(k-1)$-DNFs in a formulas expansion is the crux of our proof because it allows to determine either a DNF was very likely satisfied or its expansion into clauses did not make the proof very much larger. With this knowledge in hand, we can expand the Res(k) proof into a resolution proof, apply a random restriction and obtain a narrow resolution proof, contradicting the width lower bound for resolution.

2 Proof Systems

A *literal* is a variable or its negation. When l is a literal we write $\neg l$ for the opposite literal of l, not the syntactic object obtained by prepending l with the negation symbol. In this interpretation, $\neg\neg X$ is the variable X.

A *term* is a conjunction of literals. We will view these as sets of literals, so that for two terms T_1 and T_2, $T_1 \wedge T_2 = T_1 \cup T_2$. A *DNF* is a disjunction of terms, which we view as a set of terms, so that $F_1 \vee F_2 = F_1 \cup F_2$. A *$k$-DNF* is a DNF whose terms are each of size at most k.

Definition 2.1 *Res(k) is the refutation system whose lines are k-DNFs and whose inference rules are:*

cut:

$$\frac{A \vee l \quad B \vee \neg l}{A \vee B}$$

weakening:

$$\frac{A}{A \vee l}$$

AND-introduction, for $1 \leq j \leq k$:

$$\frac{A_1 \vee l_1 \ \ldots \ A_j \vee l_j}{A_1 \vee \ \ldots \ \vee A_j \ \vee \bigwedge_{i=1}^{j} l_i}$$

AND-elimination, for $1 \leq i \leq j \leq k$:

$$\frac{A \vee \bigwedge_{i=1}^{j} l_i}{A \vee l_i}$$

We define the size of a refutation Γ, $s(\Gamma)$, to be the number of subformulas appearing in it, although this is clearly polynomially related to an optional notion of refutation size, the number of symbols appearing in Γ. The width of a refutation Γ, $w(\Gamma)$, is the maximum number of variables appearing in a line of Γ. We are concerned with this measure primarily for resolution refutations, in which case it is equal to the maximum number of literals in a line of Γ because, without loss of generality, no clause may contain a variable and its negation.

An optional definition of Res(k) would be to allow cuts on conjunctions as follows:

$$\frac{A \vee \bigwedge_{i=1}^{j} l_i \quad B \vee \bigvee_{i=1}^{j} \neg l_i}{A \vee B}$$

We call this proof system Res$^*(k)$, however, such refutations can be simulated by Res(k) with only a $O(k)$ factor increase in the size. To do this, replace each such cut by at most k applications of the AND-elimination rule to obtain the formulas $A \vee l_1, \ \ldots A \vee l_j$, and then cutt these with the formula $B \vee \bigvee_{i=1}^{j} \neg l_i$ to derive $A \vee B$.

3 Graph Ordering Principles

Definition 3.1 *Let G be an undirected graph. For each vertex u of G, let $N(u)$ denote the set of neighbors of u in G. For each ordered pair of vertices $(u, v) \in V(G)^2$ let there be a propositional variable $P_{u,v}$.*

The graph ordering principle on G, $GOP(G)$, *is the following set of clauses:*

- *For all $u, v, w \in V(G)$, if u precedes v and v precedes w then u precedes w:*

$$P_{u,v} \wedge P_{v,w} \rightarrow P_{u,w}$$

- *For each $u \in V(G)$, u does not precede itself*

$$\neg P_{u,u}$$

- *For every $u \in V(G)$, one of u's neighbors in G precedes u:*

$$\bigvee_{v \in N(u)} P_{v,u}$$

Notice that for a graph G on n vertices with maximum degree $d \geq 3$, the principle $GOP(G)$ consists of $O(n^3)$ many clauses each of width at most d.

The principles hard for Res(k) are obtained by starting with the $GOP(G)$ principles and replacing each variable by a conjunction of $k+1$ new variables. These principles can be expressed as CNFs by the application of the distributive rule.

Definition 3.2 *Let C be a clause in the variables $X_1, \ \ldots \ X_N$.*

Let $X_1^1, \ \ldots \ X_N^1, \ \ldots X_1^k, \ \ldots \ X_N^k$ be new variables.

The k-substitution of C, $S_k(C)$ *is defined to be the k-DNF obtained by making the substitution of $\bigwedge_{j=1}^{k} X_i^j$ for each X_j:*

$$S_k(C) = C\left[X_i \leftarrow \bigwedge_{j=1}^{k} X_i^j, \neg X_i \leftarrow \bigvee_{j=1}^{k} \neg X_i^j \mid 1 \leq i \leq N\right]$$

The expanded k-substitution of C, $E_k(C)$ *is the set of clauses obtained by applying the distributive rule to $S_k(C)$ as follows:*

$$E_k(C) = \{\bigvee_{i=1}^{w} l_i \mid S_k(C) = \bigvee_{i=1}^{w} T_i, \ l_i \in T_i\}$$

We write $GOP^k(G)$ for the expanded k-substitution of $GOP(G)$:

$$GOP^k(G) = \bigcup_{C \in GOP(G)} E_k(C)$$

Notice that if G is a graph on n vertices with maximum degree $d \geq 3$, then $GOP^k(G)$ is set of at most $O(k^d n^3)$ many clauses of width at most kd.

4 The Upper Bounds

Theorem 3 *For each k, and every G with n vertices and degree at most $d \geq 3$, $GOP^k(G)$ has a Res(k) refutation of size $\mathrm{O}(kn^3 + k^{d+1})$.*

Proof:

The construction of the proofs has three steps. First, we construct the resolution refutation for $GOP(G)$, this is then used to construct a Res$^*(k)$ refutation of $\{S_k(C) \mid C \in GOP(G)\}$, which can be transformed into a Res$^*(k)$ refutation of $GOP^k(G)$, which in turn can be transformed into a Res(k) refutation of $GOP^k(G)$.

To construct the resolution refutation of $GOP(G)$, we iteratively derive the formulas $\bigvee_{i=l}^{n} P_{i,j}$ for each j, l. The clauses $\bigvee_{i=1}^{n} P_{i,j}$ are weakenings of the hypotheses. Let l range from 1 up to n. At stage l, we have $\bigvee_{i=l}^{n} P_{i,l}$ and the hypothesis $\neg P_{l,l}$, so we can derive $\bigvee_{i=l+1}^{n} P_{i,l}$. For $j \neq l$, we combine $\bigvee_{i=l+1}^{n} P_{i,l}$ with the transitivity axioms to obtain $\neg P_{l,j} \vee \bigvee_{i=l+1}^{n} P_{i,j}$, which is resolved with $\bigvee_{i=l}^{n} P_{i,j}$ to obtain $\bigvee_{i=l+1}^{n} P_{i,j}$. When we finally derive $P_{n,n}$, it is resolved with the hypothesis $\neg P_{n,n}$ to obtain the empty clause. This refutation has size $\mathrm{O}(n^3)$.

Because resolution refutations can have their weakening inferences removed with no ill effect on the size, see lemma 17 in the appendix, there is a weakening-free resolution refutation Γ of $GOP(G)$ of size $\mathrm{O}(n^3)$.

We now show that the set of formulas $S_k(\Gamma) = \{S_k(C) \mid C \in \Gamma\}$ is a Res$^*(k)$ refutation of $\{S_k(C) \mid C \in GOP(G)\}$. Consider resolution inference $\frac{A \vee X_{u,v} \quad B \vee \neg X_{u,v}}{A \vee B}$. The formula $S_k(A \vee X_{u,v})$ has the form $S_k(A) \vee \bigwedge_{i=1}^{k} X_{u,v}^i$, the formula $S_k(B \vee \neg X_{u,v})$ has the form $S_k(B) \vee \bigvee_{i=1}^{k} \neg X_{u,v}^i$ and the formula $S_k(A \vee B)$ has the form $S_k(A) \vee S_k(B)$. The number of lines of $S_k(\Gamma)$ equals the number of lines in Γ, and the substitution increases the size of each line by at most a factor of $\mathrm{O}(k)$, so $S_k(\Gamma)$ has size $\mathrm{O}(n^3)$.

For a clause C of width w, we can derive $S_k(C)$ from $E_k(C)$ by a sequence of $\mathrm{O}(k^w)$ many AND-introduction inferences. Because the width of the $GOP(G)$ clauses is at most d, we can derive $S_k(GOP(G))$ from $GOP^k(G)$ with a sequence of $\mathrm{O}(k^d)$ many AND-introduction inferences, and then refute $S_k(GOP(G))$ with a Res$^*(k)$ refutation of size $\mathrm{O}(n^3)$. Therefore, there is a size $\mathrm{O}(n^3 + k^d)$ Res$^*(k)$ refutation of $GOP^k(G)$.

Because Res$^*(k)$ refutations can be converted into Res(k) refutations with at most $\mathrm{O}(k)$ factor of increase in the size, there is a $\mathrm{O}(kn^3 + k^{d+1})$ size Res(k) refutation of $GOP^k(G)$.

■

5 The Lower Bounds

The primary result of this paper is the proof that Res(k) requires exponential size to refute the $GOP^{k+1}(G)$ principles for certain graphs with small maximum degree.

Theorem 4 *Let k be given. There exists a constant $\epsilon_k > 0$, and a family of graphs G on n vertices (for n sufficiently large) with maximum degree $\mathrm{O}(\log n)$ so that Res(k) proofs of $GOP^{k+1}(G)$ require size $2^{n^{\epsilon_k}}$.*

There are three steps to the lower bound proof. First, we show a width lower bound for resolution refutations of $GOP(G)$ for certain graphs. Second, we carefully expand Res(k) refutations into resolution refutations. Finally, we apply a random restriction which satisfies wide lines with very high probability, thereby obtaining a narrow resolution refutation to some smaller instance of the tautology, contradicting the resolution width lower bound. The difficulty is of course that the expansion of a Res(k) refutation into a resolution refutation may increase the size by an exponential factor. We handle this by choosing the expansion and the restriction in such a way that if a formula were to contribute too many formulas upon expansion, then with very high probability the random restriction satisfies it.

The parameter relating the number of $(k-1)$-DNFs needed to express a k-DNF to the likelihood of it being satisfied by our random restrictions is the "k-covering number" of a DNF.

Definition 5.1 *Let F be a k-DNF, and let S be a set of literals. If every width k term of F contains a literal of S, then we say that S* is a k-cover of F.

*The k-*covering number of F, $c_k(F)$, *is the cardinality of the smallest k-cover of F.*

The definition is motivated by the use of the distributive law. Because $(A \wedge B) \vee C$ is equivalent to $(A \vee C) \wedge (B \vee C)$, we can convert a k-DNF into a set of $(k-1)$-DNFs by making the transformation that replaces each $\bigwedge_{i=1}^{k} l_i$ by l_1 or $\bigwedge_{i=1}^{k-1} l_i$. A priori, the size of such a translation is exponential to the number of k-terms in the formula. However, if $c_k(F)$ is small, we can do better. By use of the distributive law, it is easy to see that if F is a k-DNF, then F is equivalent to the conjunction of $2^{c_k(F)}$ many $(k-1)$-DNFs, see definition 5.7 for details.

If F is a k-DNF with a large cover number, it contains many literal-disjoint k-terms. If we were to independently assign a random value to each variable, the term satisfaction events would be independent, and F would be satisfied with very high probability. However, because we wish to reduce the tautology to a smaller instance which is difficult to to refute, we do not assign values to each variable independently, but randomly partition the vertices of the graph and then, for each pair of disconnected vertices u, v, we make an assignment to $X^1_{u,v}, \ldots X^{k+1}_{u,v}$ which sets at least one of the $X^i_{u,v}$'s to 0. Any k-term will have a non-zero chance of being satisfied by such an assignment. Independence between the term satisfaction events is more difficult. If the sets of vertices underlying the terms were disjoint, we would have total independence, but this often does not happen, and we require some machinery from extremal set theory.

The proof of the lower bound appears in four parts. In subsection 5.1 we establish the width lower bound for resolution refutations of the graph-ordering principles. In subsection 5.2, we define a distribution on assignments to the variables of $GOP^{k+1}(G)$ and show that if F has a large k-cover number, then F is very likely satisfied by these restrictions. In subsection 5.3, we show that we can iteratively apply the expansion to transform a Res(k) refutatation into a resolution refutation. Finally, in subsection 5.4, we prove the lower bounds for Res(k) refutations of $\mathrm{GOP}^{k+1}(G)$.

5.1 Width Lower Bound for Resolution

We show that for certain graphs, G any resolution refutation Γ of $GOP(G)$, has $w(\Gamma) = \Omega(n)$.

Our proof follows that of Bonet and Galesi showing that the GT_n tautologies require high width, [6]. A crucial step of their proof suggested the graph property we use to guarantee high width: at one point two vertices are chosen from sufficiently large sets, and because they work with complete graphs, there is an edge joining the two vertices. We cannot work with a complete graph, because $GOP^2(K_n)$ contains more than 2^n many clauses, so we need low-degree graphs with the property that sufficiently large sets are joined by an edge.

Definition 5.2 *Let G be an undirected graph. We say that G is a s-neighborly if between every pair of disjoint sets of vertices, $A, B \subseteq V(G)$ with $|A|, |B| \geq s$, there exists an edge joining A and B.*

We now show that resolution refutations of $GOP(G)$ require large width when G is a connected graph so that even its small sets are neighborly.

Lemma 5 *If G is a connected graph of n vertices that is s-neighborly, then every resolution refutation of $GOP(G)$ contains a clause of width $\frac{n-3s}{6}$.*

Proof:

We begin by defining the "measure" of a clause. A critical truth assignment is an assignment to the variables of $GOP(G)$ which forms a total order on $V(G)$. For each $v \in V(G)$, let $C_v := \bigvee_{u \in N(v)} P_{u,v}$, and for each $I \subseteq V(G)$, $C_I := \bigwedge_{v \in I} C_v$. Let C be a clause. The *measure*

of C, $\mu(C)$, is the minimum cardinality of a set $I \subseteq V(G)$ so that for every cta α, if $\alpha \models C_I$ then $\alpha \models C$.

We can now show that $\mu(\{\}) = n$. Suppose otherwise, and let I be a subset of $V(G)$ with $|I| \leq n-1$. Choose one vertex $v_0 \in V(G) \setminus I$ and let α be a total order which arises by taking a depth-first search of G starting with v_0. Clearly $\alpha \models C_I$ but $\alpha \not\models \{\}$.

Because every clause of $GOP(G)$ has measure either 0 or 1, the empty clause has measure n and the measure is subadditive, we can choose a clause C of Γ so that $\frac{n}{3} \leq \mu(C) \leq \frac{2n}{3}$.

Suppose for the sake of contradiction that $w(C) < \frac{n-3s}{6}$.

Let I be a minimal subset of $V(G)$ so that for every critical truth assignment α, if $\alpha \models C_I$ then $\alpha \models C$. Let $J = V(G) \setminus I$. Notice that $|I|, |J| \geq \frac{n}{3}$.

Let S be the set of vertices mentioned by literals of C. By assumption, $|S| < 2\left(\frac{n-3s}{6}\right) = \frac{n-3s}{3}$. Therefore, $|I \setminus S| \geq \frac{n}{3} - \frac{n-3s}{3} = s$. Similarly, $|J \setminus S| \geq s$. Because G is s-neighborly, we may choose $u \in I \setminus S$ and $v \in J \setminus S$ so that $\{u, v\}$ is an edge of G.

Let α be a critical truth assignment so that $\alpha \models C_{I \setminus \{u\}}$ but $\alpha \not\models C_u$ and $\alpha \not\models C$. Let β be the critical truth assignment which arises by moving v to the front of α. For $w \in I$, $w \neq u$, $\beta \models C_w$ because every predecessor of w in α is a predecessor of w in β. For u, $\beta \models C_u$ because $\beta \models P_{v,u}$. However, $\beta \not\models C$ because $\alpha \not\models C$ and no variable mentioning u or v appears in C. Therefore, $\beta \models C_I$ but $\beta \not\models C$, contradiction to the choice of I.

■

5.2 Random Restrictions and Cover Number

We use the following distribution on random partial assignments to satisfy k-DNFs with high cover number. The intended meaning is that we randomly partition the graph into $4k$ many pieces. To do this we randomly color the graph, and then between vertices u and v of distinct color classes, we choose an assignment $X^1_{u,v}, \ldots X^{k+1}_{u,v}$ which makes $\bigwedge_{i=1}^{k+1} X^i_{u,v}$ false.

Definition 5.3 *Let $k \geq 1$ be given. Let G be a graph. The distribution $\mathcal{P}_{k+1}(G)$ on partial assignments ρ to the variables of $GOP^{k+1}(G)$ is defined as follows:*

Select a random coloring of $V(G)$ by $4k$ many colors, $c^\rho : V(G) \to [4k]$.

For each $(u,v) \in V(G)^2$, let $\sigma^\rho_{u,v}$ be chosen uniformly among $0, 1$ assignments to $X^1_{u,v}, \ldots, X^{k+1}_{u,v}$ such that for at least one i, $\sigma^\rho_{u,v}(X^i_{u,v}) = 0$.

The partial assignment ρ, and an auxilliary total assignment, σ^ρ, are defined as follows:

$$\rho = \bigcup_{\substack{(u,v) \in V(G)^2 \\ c^\rho(u) \neq c^\rho(v)}} \sigma^\rho_{u,v}$$

$$\sigma^\rho = \bigcup_{(u,v) \in V(G)^2} \sigma^\rho_{u,v}$$

Our first lemma, which we state without proof, is that if a restriction corresponds to deleting edges of G to obtain G', then applying that restriction to $GOP^{k+1}(G)$ results in $GOP(G')$.

Lemma 6 *Let G be a graph. Let $\rho \in \mathcal{P}_{k+1}$ be given. Let G' be the graph induced by deleting edges of G that are bichromatic under c^ρ.*

$$GOP^{k+1}(G) \restriction_\rho = GOP^{k+1}(G')$$

Formulas with high cover number contain many literal-disjoint terms, but the events of term satisfaction are not necessarily independent. Consider the following example with literals $X^1_{u,v}, X^1_{v,w}, X^1_{u,w}$. If $X^1_{u,v}$ is falsified, and $X^1_{v,w}$ is falsified, then u and w are more likely to receive the same color, so $X^1_{u,w}$ is less likely to be satisfied. To obtain independence, we apply some extremal set theory to the sets of vertices involved with the terms.

Definition 5.4 *Let $X^i_{u,v}$ be a variable. The* underlying pair *of $X^i_{u,v}$ is $\{u,v\}$.*

Let T be a term. The set of vertex pairs *of T, P_T, is defined to be*

$$P_T = \{\{u,v\} \mid \ \{u,v\} \text{ is the underlying pair of a variable in } T\}$$

The set of vertices *of T, S_T, is defined to be*

$$S_T = \bigcup P_T$$

We use the notion of a *sunflower* to obtain statistical independence between the term satisfaction events. The precise formulation we use is a version of the Erdös-Rado lemma that appears in [18].

Definition 5.5 *A (p,l)* sunflower *is a collection of sets $P_1, \ldots P_p$, so that for each i, $1 \leq i \leq p$, with each $|P_i| \leq l$ and for all $1 \leq i < j \leq p$, $1 \leq i' < j' \leq p$, $P_i \cap P_j = P_{i'} \cap P_{j'}$.*

Theorem 7 *Let l be given. Let $\mathcal{Z}$ be a family of M distinct sets, each with cardinality $\leq l$. $\mathcal{Z}$ contains a (p,l) sunflower where $p \geq \left(\frac{M}{l!}\right)^{\frac{1}{l}}$.*

Definition 5.6 *Let $T_1, \ldots T_t$ be a set of terms.*

We say that the terms are sufficiently independent *if the following conditions hold:*

1. *The family of sets $\{S_{T_i} \mid 1 \leq i \leq t\}$ forms a sunflower with core C.*
2. *For each i, each $\{u,v\} \in P_i$, either $u \notin C$ or $v \notin C$*

Notice that the sets P_{T_i}, $1 \leq i \leq t$, are disjoint as a consequnce of this definition.

Lemma 8 *Let F be a k-DNF. F contains a sufficiently independent set of $\left(\frac{c_k(F)}{(2k)!(2(k+1)(2k)(2k-1))^k}\right)^{\frac{1}{2k}} - 16k^2(k+1)$ many terms.*

Proof:

Set $s = c_k(F)$. F contains a set of s many literal-disjoint k-terms, $T_1, \ldots T_s$. It is quite possible that $S_{T_i} = S_{T_j}$ for some $i \neq j$, however, a set of $\leq 2k$ many vertices can be the underlying set of at most $(2(k+1)(2k)(2k-1))^k$ many different k-terms. Therefore, there is a sub-collection of of $\frac{s}{(2(k+1)(2k)(2k-1))^k}$ many terms whose underlying sets of vertices are distinct.

Because the underlying sets of vertices have size at most $2k$, we can apply the sunflower lemma, and find $s' = \left(\frac{s}{(2(k+1)(2k)(2k-1))^k (2k)!}\right)^{\frac{1}{2k}}$ many terms whose sets of underlying vertices form an $(s', 2k)$ sunflower. We rename these terms $T_1, \ldots T_{s'}$.

Let C be the core of the sunflower $S_{T_1}, \ldots S_{T_{s'}}$. Consider pairs of vertices $\{u,v\} \subseteq C$. There are no more than $4k^2$ many pairs in $[C]^2$, and each such pair is the underlying pair of exactly $4(k+1)$ many literals. Therefore, there are at most $4(k+1) \cdot 4k^2 = 16k^2(k+1)$ many literals whose underlying vertices are both in C. The terms $T_1, \ldots T_{s'}$ are literal-disjoint, so each such literal appears in at most one term, and when we remove all terms containing these literals, we obtain a sufficiently independent set of terms of size $s' - 16k^2(k+1) = \left(\frac{s}{(2(k+1)(2k)(2k-1))^k (2k)!}\right)^{\frac{1}{2k}} - 16k^2(k+1)$.

■

Lemma 9 *Let F be a k-DNF which contains contains t sufficiently independent terms.*

$$Pr_\rho[F \restriction_\rho \neq 1] \leq \left(1 - \frac{1}{2^{k^2+3k}}\right)^t$$

Proof: Let $T_1, \ldots T_t$ be the sufficiently independent terms of F. Let C be the core of the sunflower $S_{T_1}, \ldots S_{T_t}$.

Let ρ be a partial assignment chosen by the distribution $\mathcal{P}_k$.

Notice that for each term T_i, $T_i \restriction_\rho = 1$ if and only if the following two events occur: (i) $T_i \restriction_{\sigma^\rho} = 1$ and (ii) for each $\{u, v\} \in P_{T_i}$, $c^\rho(u) \neq c^\rho(v)$.

Fix a coloring of the vertices in the core $\chi : C \to [4k]$. We will show that

$$\Pr_\rho [F \restriction_\rho \neq 1 \mid c^\rho \restriction_C = \chi] \leq \left(1 - \frac{1}{2^{k^2+3k}}\right)^t$$

First we show that for each T_i, $\Pr_\rho [T_i \restriction_\rho | \, c^\rho \restriction_C = \chi] \geq \frac{1}{2^{k^2+3k}}$. For each $\{u, v\} \in P_{T_i}$, $\{u, v\} \not\subseteq C$, if every $v \in S_{T_i} \setminus C$ with a distinct color outside of Rng χ then every pair of P_{T_i} is bichromatic, and this happens with probability at least $\frac{1}{2^{2k}}$. Because T_i contains at most k literals, $T_i \restriction_{\sigma^\rho} = 1$ occurs with probability at least $\frac{1}{\left(2^{k+1}\right)^k}$. These two events are independent, so we may multiply the probabilities.

Next we show that when we condition on the event that $c^\rho \restriction_C = \chi$, we have that the events $T_i \restriction_\rho = 1$ are totally independent. Because the terms share no underlying pairs of vertices, the events of type (i) are completely independent of one another and the coloring c^ρ. Because no pair $\{u, v\}$ in any P_{T_i} is contained in the core, the events of type (ii) depend only on the values that c^ρ takes on $S_{T_i} \setminus C$, so they are independent of one another and the assignmment σ^ρ.

Therefore, $\Pr_\rho [F \restriction_\rho \neq 1 \mid c^\rho \restriction_C = \chi] \leq \left(1 - \frac{1}{2^{k^2+3k}}\right)^t$.

Because this hold for all colorings of the core, we have that

$$\Pr_\rho [F \restriction_\rho \neq 1 = \chi] \leq \left(1 - \frac{1}{2^{k^2+3k}}\right)^t$$

■

Combining lemmas 8 and 9 we have the following:

Lemma 10 *For each k there exists a positive constant depending only on k, γ_k, so that if F is a k-DNF with $c_k(F) \geq s$, then*

$$Pr_\rho [F \restriction_\rho \neq 1] \leq 2^{-\gamma_k s^{\left(\frac{1}{2k}\right)}}$$

5.3 Expanding Res(k) Refutations into Resolution Refutations

The distributive law of boolean algebra states that $(A \wedge B) \vee C$ is equivalent to $(A \vee C) \wedge (B \vee C)$. This fact allows us to convert k-DNFs into an equivalent set of $(k-1)$-DNFs, and iterating this expansion gives us a conversion into an equivalent set of clauses. In this section we define a method for converting k-DNFs into sets of $(k-1)$-DNFs (and hence clauses) in a manner so that the factor of size increase is at most exponential the k-covering number of the DNF, and demonstrate how this expansion can be used to transform Res(k) refutations into resolution refutations.

Our method of expansion is based on covering sets (recall definition 5.1). For a k-DNF F with k-cover S, we expand F into the set of all possible $(k-1)$-DNFs that arise by choosing, for each literal in the k-covering set to either include it or all of the terms containing that literal, with that literal removed.

Definition 5.7 *Let $k > 1$, let F be a k-DNF and let S be a set of literals which k-covers F.*

For each $l \in S$, set $C_l^0 = \{l\}$ and $C_l^1 = \{T \setminus l \mid T \in F, \ |T| = k, \ l \in T\}$.

Let $\tilde{F}$ be the disjunction of all terms of width $\leq k-1$ appearing in F.

The expansion of F with respect to S, $\mathcal{E}_S(F)$, *is defined as:*

$$\mathcal{E}_S(F) = \{\tilde{F} \cup \bigcup_{l \in S} C_l^{\epsilon_l} \mid \vec{\epsilon} \in \{0,1\}^S\}$$

Clearly, if F is a k-DNF with k-cover S, then $\mathcal{E}_S(F)$ is a set of at most $2^{|S|}$ many k-DNFs. Moreover, F is equivalent to the conjunction of all the formulas in $\mathcal{E}_S(F)$.

We iterate this process to reduce a k-DNFs into sets of clauses. To keep the expansion as small as possible, at each step we expand according to the minimum possible cover set.

Definition 5.8 *Let F be a k-DNF.*
For j, $1 \leq j \leq k$, we define the expansion of F into j-DNFs, $\mathcal{E}_j(F)$, *as follows:*

Set $\mathcal{E}_k(F) = \{F\}$
  For i ranging from k down to 2:
    For each $f \in \mathcal{E}_i(F)$
      let S_f be the first, smallest $(i-1)$-cover of f
    Set $\mathcal{E}_{i-1}(F) = \bigcup_{f \in \mathcal{E}_i(F)} \mathcal{E}_{S_f}(F)$

We now take note of some basic properties of this expansion. The proofs are simple and omitted for space.

Lemma 11
$\mathcal{E}_j(F)$ is a set of j-DNFs whose conjunction is equivalent to F
k-DNF F and $i < j \leq k$, $f \in \mathcal{E}_j(F)$, we have $\mathcal{E}_i(f) \subseteq \mathcal{E}_i(F)$
$|\mathcal{E}_{k-1}(F)| \leq 2^{c_k(F)}$

Unfortunately, the expansion of a Res(k) refutation does not necessarily result in a resolution refutation. Because their respective covering sets could be different, $\mathcal{E}_1(A \vee B)$ need not look very much like $\mathcal{E}_1(A \vee x)$ or $\mathcal{E}_1(B \vee \neg x)$.

Another way to expand k-DNFs into a set of clauses is to simply use the set of all possible clauses which select one literal from each term. This can be very inefficient, but it behaves in a regular way that allows us to connect the expansions of similar formulas. It is used in an auxiliary to help us reason about the more complicated $\mathcal{E}_1$ expansion.

Definition 5.9 *Let $F = \bigvee_{i=1}^{w} T_i$ be a k-DNF. The* brute-force expansion of F, $\mathcal{B}(F)$, *is defined to be the following set of clauses:*

$$\mathcal{B}(F) = \{\bigvee_{i=1}^{w} l_i \mid \forall i,\ 1 \leq i \leq w,\ l_i \in T_i\}$$

Let's take the time to note some properties of the $\mathcal{E}_1$ and $\mathcal{B}$ expansions.

Lemma 12 *For any k-DNF F:*

1. $\mathcal{B}(A \vee B) \subseteq \{f_A \vee f_B \mid f_A \in \mathcal{B}(A),\ f_B \in \mathcal{B}(B)\}$
2. $\mathcal{E}_1(F) \subseteq \mathcal{B}(F)$
3. *For any $f \in \mathcal{E}_1(A \vee l)$ there exists $g \in \mathcal{E}_1(A)$ so that $f = g \vee l$.*
4. *For every $f \in \mathcal{B}(F)$, there exists $f_0 \in \mathcal{E}_1(F)$, $f_0 \subseteq f$.*

Proof:
Proofs of the the first three properties are left to the reader.
The third property is proved by induction, our hypothesis is "For each $1 \leq j \leq k$, if F is a j-DNF, then for each $f \in \mathcal{B}(F)$ there exists $f_0 \in \mathcal{E}_1(F)$ so that $f_0 \subseteq f$".
For the base case when $j = 1$, F is a clause and $\mathcal{B}(F) = \mathcal{E}(F) = \{F\}$.

Let j, $1 < j \leq K$, be given and assume that the induction hypothesis holds for $(j-1)$ DNFs. Let F be a j-DNF, and let $f \in \mathcal{B}(F)$ be given. We adopt the notation of definition 5.7. Let S be the first, least covering set of F. Let S_0 be the set of literals of S that appear in f, and let S_1 be the set of literals of S that do not appear in f.

Consider the formula F' defined as:

$$F' = \tilde{F} \vee \bigvee_{l \in S_0} l \vee \bigvee_{l \in S_1} C_l^0$$

Clearly, $F' \in \mathcal{E}_{j-1}(F)$ and $f \in \mathcal{B}(F')$. Therefore, by the induction hypothesis, we may choose $f_0 \in \mathcal{E}_1(F')$ so that $f_0 \subseteq f$. However, $\mathcal{E}_1(F') \subseteq \mathcal{E}_1(F)$, so we have found $f_0 \in \mathcal{E}_1(F)$ that weakens f.

■

Lemma 13 *For each $f \in \mathcal{E}_1(A \vee B)$, there is a width $\leq w(f)$ resolution derivation of f from $\mathcal{E}_1(A \vee l)$ and $\mathcal{E}_1(B \vee \neg l)$.*

Proof: By lemma 12, $f \in \mathcal{B}(A \vee B)$, so we may choose $f_A \in \mathcal{B}(A)$, $f_B \in \mathcal{B}(B)$ so that $f = f_A \vee f_B$. The formula $f_A \vee l$ belongs to $\mathcal{B}(A \vee l)$, and $f_B \vee \neg l$ belongs to $\mathcal{B}(B \vee \neg l)$. Therefore, by lemma 12, each of these formulas follows from $\mathcal{E}_1(A \vee l)$ and $\mathcal{E}_1(B \vee \neg l)$ by weakening. We then cut on l to derive $f_A \vee f_B$. ■

Now that we know how to derive $\mathcal{E}_1(A \vee B)$ from $\mathcal{E}_1(A \vee \bigwedge_{i=1}^{j} l_i)$ and $\mathcal{E}_1(B \vee \bigvee_{i=1}^{j} \neg l_i)$, we can expand Res($k$) refutations into resolution refutations by expanding each line into clauses and adding some auxiliary clauses needed for the derivations.

Definition 5.10 *Let Γ be a Res(k) refutation of $\mathcal{C}$. Let F be a formula of Γ, and let $f \in \mathcal{E}_1(F)$ be given.*

If F is hypothesis, or results from a weakening, AND-introduction, or AND-elimination inference, then for each $f \in \mathcal{E}_1(F)$, we set $\mathcal{I}_{\Gamma,F}(f) = \emptyset$.

If F results from a cut inference, then let $\mathcal{I}_{\Gamma,F}(f)$ be formulas of the derivation guaranteed by lemma 13.

$$\mathcal{E}_\Gamma(F) = \mathcal{E}_1(F) \cup \bigcup_{f \in \mathcal{E}_1(F)} \mathcal{I}_{\Gamma,F}(f)$$

Definition 5.11 *Let Γ be a Res(k) refutation of a set of clauses, $\mathcal{C}$. The* resolution expansion *of Γ, $\mathcal{R}(\Gamma)$, is defined by replacing every formula F in Γ by $\mathcal{E}_\Gamma(F)$.*

Lemma 14 *If Γ is a Res(k) refutation of a set of clauses $\mathcal{C}$, then $\mathcal{R}(\Gamma)$ is a resolution refutation of $\mathcal{C}$.*

Proof:

We follow the structure of Γ and show inductively that the set of clauses $\mathcal{E}_\Gamma(F)$ follows by resolution from the clauses of $\mathcal{E}_\Gamma(F')$ where F' precedes F in Γ.

Because every hypothesis $H \in \mathcal{C}$ is a clause, $\mathcal{E}(H) = \{H\}$, and therefore the hypotheses of $\mathcal{E}(\Gamma)$ are the same as the hypotheses of Γ.

Suppose that $F = A \vee B$, and follows in Γ, by a cut of the form $\frac{A \vee l \quad B \vee \neg l}{A \vee B}$. In this case $\mathcal{E}_\Gamma(A \vee B)$ follows from $\mathcal{E}_\Gamma(A \vee l)$ and $\mathcal{E}_\Gamma(B \vee \neg l)$ by lemma 13.

In the remaining cases we will repeatedly use the fact that $\mathcal{E}_\Gamma(F) = \mathcal{E}_1(F)$.

Suppose that $F = A \vee l$ follows by a weakening inference $\frac{A}{A \vee l}$. For each $f \in \mathcal{E}_1(A \vee l)$, there exists $g \in \mathcal{E}_1(A)$ so that $f = g \vee l$. Therefore, for every f in $\mathcal{E}_1(A \vee l)$, f can be inferred by weakening from some $g \in \mathcal{E}_1(A)$.

Suppose that $F = A \vee \bigwedge_{i=1}^{j} l_i$ follows by AND-introduction, $\frac{A \vee l_1 \ \cdots \ A \vee l_j}{A \vee \bigwedge_{i=1}^{j} l_i}$. Because $\mathcal{E}_1(A \vee \bigwedge_{i=1}^{j} l_i) \subseteq \mathcal{B}(A \vee \bigwedge_{i=1}^{j} l_i) \subseteq \bigcup_{i=1}^{j} \mathcal{B}(A \vee l_i)$, for each $f \in \mathcal{E}_1(A \vee \bigwedge_{i=1}^{j} l_i)$, we can choose

$f_0 \in \bigcup_{i=1}^{j} \mathcal{E}_1(A \vee l_i)$ so that $f_0 \subseteq f$. Therefore, every formula of $\mathcal{E}_\Gamma(A \vee \bigwedge_{i=1}^{j} l_i)$ follows by weakening some formula of $\bigcup_{i=1}^{j} \mathcal{E}_1(A \vee l_i) \subseteq \bigcup_{i=1}^{j} \mathcal{E}_\Gamma(A \vee l_i)$.

Suppose that $F = A \vee l$ follows by AND-elimination, $\frac{A \vee l \wedge \bigwedge_{i=1}^{j} l_i}{A \vee l}$. Because $\mathcal{E}_1(A \vee l) \subseteq \mathcal{B}(A \vee l \wedge \bigwedge_{i=1}^{j} l_i)$, for each $f \in \mathcal{E}_\Gamma(A \vee l)$, we can choose $f_0 \in \mathcal{E}_1(A \vee l \wedge \bigwedge_{i=1}^{j} l_i)$ so that $f_0 \subseteq f$. ∎

5.4 The Lower Bound

In this subsection we tie our threads together and prove the lower bound. As has been said before, the proof expands a small Res(k) refutation to a resolution refutation and then restricts it by a random assignment. The crucial observation is that the formulas that contribute a large increase in the size when we expand them, those with high cover number, will be satisfied by the restriction. Therefore, most of the clauses in the expanded proof, and those with high width in particular, will be eliminated from the refutation upon restriction, and thus we obtain a narrow resolution refutation for an instance of the GOP principle.

Upon random restriction by ρ, the graph $G \restriction_\rho$ obtained by deleting the bichromatic edges should have that resolution refutations of $GOP(G \restriction_\rho)$ require large width. By lemma 5.1, it suffices that $G \restriction_\rho$ should consist of large, neighborly connected components with very high probability. We use the notion of "usefulness" to formalize such graphs. Unsurprisingly, random graphs of $\Theta(\log n)$ degree are useful with high probability. This is shown in the appendix.

Definition 5.12 *We say that an n vertex graph G is (m, ϵ, p)* useful *if upon partition into m distinct vertex induced subgraphs, with probability $\geq p$, each subgraph G_i is connected, has size at least $\frac{n}{2m}$, and is $\epsilon|V(G_i)|$-neighborly.*

We now prove theorem 4.

Proof:

Let k be given.

Apply lemma 16 of the appendix and choose G to be an n vertex, degree $c \log n$, $(4k, \frac{1}{6}, \frac{1}{4})$-useful graph.

Choose $s_0 > s_1 > ... > s_k$ satisfying $s_0 = \frac{n}{48k} - 1$ and for each $i > 0$, $s_i = \frac{1}{k}\left(\gamma_k (s_{i-1})^{\frac{1}{2k}} - \ln(2k)\right)$.

Suppose, for the sake of contradiction that Γ is a Res(k) refutation of $\mathrm{GOP}^{k+1}(G)$ with size $\leq 2^{s_k}$.

For each clause $f \in \mathcal{E}_1(\Gamma)$, $g \in \mathcal{E}_j(\Gamma)$ we say that g is a *j-ancestor of f* if $f \in \mathcal{E}_1(g)$. For each clause f of $\mathcal{E}_1(\Gamma)$, and j, $1 \leq j \leq k$, we say that f is j-bad if, for every i, $j < i \leq k$, every i-ancestor g of f has $c_i(g) \leq s_{i-1}$ and there exists a j-ancestor g of f so that $c_j(g) > s_{j-1}$. Call such a g a *j-corruptor*. If f is not j-bad for any j, $1 \leq j \leq k$, then we say that f is good. Notice that good f's are clauses of width $\leq s_0$. Moreover, for each j, $1 \leq j \leq k$, there are at most $2^{\sum_{i=j}^{k} s_i}$ many j-corruptors.

Let ρ be uniformly selected from $\mathcal{P}_{k+1}(G)$. Let j, $1 \leq j < k$, be given. By lemma 10, each j-corruptor in $\mathcal{E}_j(\Gamma)$ survives restriction by ρ with probability at most $2^{-\gamma_k \left(s_{j-1}\right)^{\frac{1}{2k}}}$. Therefore, the probability that there exists a j-corruptor $g \in \mathcal{E}_j(\Gamma)$ so that $g \restriction_\rho \neq 1$ is at most:

$$2^{\sum_{i=j}^{k} s_i - \gamma_k s_{j-1}^{\frac{1}{2k}}} \leq 2^{k s_j - \gamma_k s_{j-1}^{\frac{1}{2k}}} \leq 2^{-\ln(2k)} = \frac{1}{2k}$$

Therefore, with probability at least $\frac{1}{2}$, for each j, $1 \leq j \leq k$, every j-bad $f \in \mathcal{E}_1$ has $f \restriction_\rho = 1$.

Because G is $(4k, \frac{1}{2}, \frac{1}{4})$ useful, and $\frac{1}{2} + \frac{1}{4} < 1$, we can choose $\rho \in \mathcal{P}_{k+1}(G)$ so that for every $f \in \mathcal{E}_1(\Gamma)$, if f is bad then $f \restriction_\rho = 1$ and for each component G_i of $G \restriction_\rho$, G_i has size $\geq \frac{n}{8k}$ and neighborliness $\geq \frac{1}{6}|V(G_i)| \geq \frac{n}{48k}$.

Let Γ' be the refutation obtained by taking $\mathcal{R}(\Gamma) \restriction_\rho$ and removing all clauses that do not reach the empty clause. In particular, for any $F \in \Gamma$, $f \in \mathcal{E}_1(F)$, if $f \restriction_\rho = 1$ then any formulas of $\mathcal{I}_{\Gamma,F}(f)$ are removed.

The only formulas of Γ' are the restrictions of good formulas of $\mathcal{E}_1(\Gamma)$ and their auxiliary formulas, so by lemma 13, the width of Γ' is at most $s_0 = \frac{n}{48k} - 1$.

By lemma 18, there exists a refutation Γ^* with $w(\Gamma^*) \leq w(\Gamma') \leq s_0$ of one of the principles $GOP^{k+1}(G_i)$. However, lemma 5 tells us that $w(\Gamma^*) \geq \frac{1}{6}\left(\frac{n}{8k} - \frac{n}{16k}\right) = \frac{n}{48k}$, contradiction.

■

References

[1] Albert Atserias, Maria Luisa Bonet, and Juan Luis Esteban. Lower bounds for the weak pigeonhole principle beyond resolution. Accepted at Information and Computation, 2001.

[2] Paul Beame. A switching lemma primer. Technical report, Department of Computer Science and Engineering, University of Washington, 1994.

[3] Paul Beame and Toniann Pitassi. Simplified and improved resolution lower bounds. In *37th Annual Symposium on Foundations of Computer Science*, pages 274–282, Burlington, Vermont, 14–16 October 1996. IEEE.

[4] Richard Beigel and David Eppstein. 3-coloring in time $O(1.3446^n)$: a no-MIS algorithm. In *36th IEEE Symposium on Foundations of Computer Science*, pages 444–452, 1995.

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

[6] M. L. Bonet and N. Galesi. A study of proof search algorithms for resolution and polynomial calculus. In IEEE, editor, *40th Annual Symposium on Foundations of Computer Science: October 17–19, 1999, New York City, New York,*, pages 422–431, 1109 Spring Street, Suite 300, Silver Spring, MD 20910, USA, 1999. IEEE Computer Society Press.

[7] Samuel R. Buss and György Turán. Resolution proofs of generalized pigeonhole principles. *Theoretical Computer Science*, 62(3):311–317, December 1988.

[8] Stephen A. Cook and Robert A. Reckhow. The relative efficiency of propositional proof systems. *The Journal of Symbolic Logic*, 44(1):36 – 50, March 1979.

[9] Merrick Furst, James B. Saxe, and Michael Sipser. Parity, circuits, and the polynomial-time hierarchy. *Mathematical Systems Theory*, 17(1):13–27, April 1984.

[10] Andreas Goerdt. Unrestricted resolution versus N-resolution. In Branislav Rovan, editor, *Mathematical Foundations of Computer Science 1990*, volume 452 of *lncs*, pages 300–305, Banská Bystrica, Czechoslovakia, 27–31 August 1990. Springer.

[11] A. Haken. The intractability of resolution. *Theoretical Computer Science*, 39(2-3):297–308, August 1985.

[12] Johan Hastad. Almost optimal lower bounds for small depth circuits. In *Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing*, pages 6–20, Berkeley, California, 28–30 May 1986.

[13] J. Krajicek, P. Pudlak, and A. Woods. An exponential lower bound to the size of bounded depth frege proofs of the pigeonhole principle. *RSA: Random Structures and Algorithms*, 7, 1995.

[14] Jan Krajíček. Lower bounds to the size of constant-depth propositional proofs. *J. of Symbolic Logic*, 59:73–86, 1994.

[15] Jan Krajíček. On the weak pigeonhole principle. *Fudamenta Mathematicae*, 170:123–140, 2001.

[16] Maciel, Pitassi, and Woods. A new proof of the weak pigeonhole principle. In *STOC: ACM Symposium on Theory of Computing (STOC)*, 2000.

[17] Edgar M. Palmer. *Graphical Evolution: An Introduction to the Theory of Random Graphs.* Wiley Interscience, 1985.

[18] Christos H. Papadimitriou. *Computatational Complexity.* Addison-Wesley, Reading, Mass., 1994.

[19] Paris, Wilkie, and Woods. Provability of the pigeonhole principle and the existence of infinitely many primes. *JSL: Journal of Symbolic Logic*, 53, 1988.

[20] Toniann Pitassi, Paul Beame, and Russell Impagliazzo. Exponential lower bounds for the pigeonhole principle. *Computational Complexity*, 3(2):97–140, 1993.

[21] J. M. Robson. Algorithms for maximum independent sets. *Journal of Algorithms*, 7(3):425–440, September 1986.

6 Appendix

6.1 Neighborly and Useful Graphs

A quick and crude application of the probabilistic method shows that with very high probility, a random graph of expected degree $\theta(\log n)$ is neighborly.

Lemma 15 *Let $p = p(n) = \frac{c \log n}{n}$. With probability $\leq e^{\left(2\epsilon - c\epsilon^2\right) n \log n}$, $G_{n,p}$ is* not *ϵn-neighborly.*

Proof: There are at most $n^{2\epsilon n} = 2^{2\epsilon n \log n}$ many pairs of disjoint sets of ϵn many vertices. Each such pair has a chance of at most $(1-p)^{\epsilon^2 n^2}$ of being unconnected. However, $(1-p)^{\epsilon^2 n^2} = (1 - \frac{c \log n}{n})^{\epsilon^2 n^2} \leq e^{-c\epsilon^2 n \log n}$, so an application of the union bound proves the lemma.

■

Lemma 16 *For each m, ϵ there exists a constant c, so that for each n, there exists an $(m, \epsilon, \frac{1}{4})$ useful graph G on n vertices with degree* $\mathrm{O}(\log n)$.

Proof:

Let $p = \frac{c \log n}{n}$.

Consider the following experiment: select a graph G according to the distribution $G_{n,p}$, and then randomly partition its vertex set into m sets, removing all cross edges, to form m vertex induced subgraphs, $G_1, \ \ldots \ G_m$.

Let P be the probability that G has a vertex of degree $> 2c \log n$, or that one of the partition classes of G has size $< \frac{n}{2m}$, or that one of the induced subgraphs is disconnected or is not $\frac{\epsilon n}{2m}$-neighborly. We now bound this probability.

Consider the probability that G has a vertex of degree $\geq 2c \log n$. By the Chernoff bounds, the probability of any one vertex having degree in excess of $2p(n-1)$ is no more than $\left(\frac{e}{4}\right)^{p(n-1)} = \left(\frac{e}{4}\right)^{\frac{c(n-1)\log n}{n}}$.

The Chernoff bounds also allow us to bound the probability that any of the G_i's contain too few vertices. The probability that a given piece of the partition fewer than $\frac{n}{2m}$ vertices is bounded by $e^{-\frac{n}{8m}}$.

Once we condition upon all pieces of the partition containing at least $\frac{n}{2m}$ vertices, we can bound the probability that any induced subgraph is disconnected. Consider a fixed set of $s \geq \frac{n}{2m}$ many vertices, and condition upon the event those vertices recieve the same color in the partition. Each edge internal to the set is included with probability $\frac{c \log n}{n} = \frac{(cs/n)\log n}{s} \geq \frac{(c/2m)\log s}{s}$. By a standard result on connectivity of the random graph (see [17] for details), each side is

disconnected with probability bounded by $\mathrm{O}\left(\frac{1}{n^{c/(2m)-1}}\right)$. Choose a constant d so that this probability is bounded by $\frac{d}{n^{c/2m-1}}$.

Finally, we consider the probability that each of the components G_i is $\epsilon|V(G_i)|$-neighborly. For a fixed set of $s \geq \frac{n}{2m}$ vertices, if we condition on the event that that set forms a component after partition, each internal edge is included with probability $\geq \frac{(c/2m)\log s}{s}$. By lemma 15, that means that the component is *not* a $\epsilon s \geq \frac{\epsilon n}{2m}$ expander with probability at most $e^{(2\epsilon-(c/2m)\epsilon^2)s\log s}$. Provided that $\frac{4m}{\epsilon} < c$, this is less than $e^{(2\epsilon-(c/2m)\epsilon^2)\frac{n}{2m}\log\left(\frac{n}{2m}\right)}$.

Therefore,

$$P \leq \left(\frac{e}{4}\right)^{\frac{c(n-1)\log n}{n}} + me^{-\frac{n}{8m}} + \frac{md}{n^{c/2m-1}} + me^{(2\epsilon-(c/2m)\epsilon^2)\frac{n}{2m}\log\left(\frac{n}{2m}\right)}$$

For a sufficiently large constant c, dependent only on m and ϵ, this is below $\frac{1}{2}$.

Therefore, by an averaging argument on the edge choices, there exists a graph G of maximum degree $\leq 2c\log n$ so that upon random partition of its vertices into m disjoint sets, its induced subgraphs are each connected and of size $\geq \frac{n}{2m}$ with probability $\geq \frac{1}{2}$.

■

6.2 Refutations of Variable Disjoint Sets of Clauses

It is a widely known theorem that resolution refutations can omit weakening inferences with no ill-effect on the size or width. As with all results from the folklore, it is hard to find a reference. We prove it here.

Lemma 17 *Let $\Gamma = C_1,\ C_2,\ \ldots\ C_m$ be a resolution refutation of set of initial clauses $\mathcal{C}$. There is a resolution refution $\Gamma' = C_1{}',\ C_2{}',\ \ldots\ C_m{}'$ of $\mathcal{C}$ so that no weakening inference is used, and for each i, $C_i{}' \subseteq C_i$.*

Proof:

If C_i is an initial clause, then $C'_i := C_i$.

If C_i is a weakening of C_j, then let $C'_i = C'_j$.

If C_i is inferred by the resolution of C_j and C_k on the variable x, we have a case analysis depending on how x appears in C'_j and C'_k. If x appears in both C'_j and C'_j, then let C'_i be the result of resolving C'_j and C'_k. If x does not appear in C'_j, then let C'_i be C'_j. If x does appear in C'_j but not C'_k, then let C'_i be C'_k.

■

Lemma 18 *Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be unsatisfiable sets of claues on disjoint sets of variables. If there is a refutation Γ of $\mathcal{C}_1 \cup \mathcal{C}_2$, then there is a refutation Γ' of either $\mathcal{C}_1$ or $\mathcal{C}_2$. Moreover, $s(\Gamma') \leq s(\Gamma)$, $w(\Gamma) \leq w(\Gamma')$ and $p(\Gamma) \leq p(\Gamma')$.*

Proof: Let Γ^* be the weakening-free refutation of $\mathcal{C}_1 \cup \mathcal{C}_2$ as guaranteed by lemma 17. Because clauses that are resolved with one another share a variable, the clauses of Γ^* can be partitioned into the consequences of $\mathcal{C}_1$ and the consequences of $\mathcal{C}_2$. One of these sets of clauses contains the empty clause, and that set is a resolution refutation Γ' which has $s(\Gamma') \leq s(\Gamma^*) \leq s(\Gamma)$ and $w(\Gamma') \leq w(\Gamma^*) \leq w(\Gamma)$. ■