Lower bounds for k-DNF Resolution on random 3-CNFs

Michael Alekhnovich *

February 24, 2005

Abstract

We prove exponential lower bounds for the refutation of a random 3-CNF with linear number of clauses by k-DNF Resolution for $k \leq \sqrt{\log n / \log \log n}$. For this we design a specially tailored random restrictions that preserve the structure of the input random 3-CNF while mapping every k-DNF with large covering number to 1 with high probability. Next we make use of the switching lemma for small restrictions by Segerlind, Buss and Impagliazzo to prove the lower bound.

This work improves the previously known lower bound for Res(2) system on random 3-CNFs by Atserias, Bonet and Esteban and the result of Segerlind, Buss, Impagliazzo stating that random $O(k^2)$-CNF do not possess short Res(k) refutations.

1 Introduction

Random 3-CNF formulas form the most popular stochastic distribution on the instances of NP-complete problems, the complexity of which has been investigated by many researchers. These formulas are generated as a random subset of Δn clauses over n variables, the goal is to satisfy all the clauses. It is known that for any n there exists a sharp *density threshold* c_n [Fri99], below which the random formulas are almost certainly satisfiable and above which are unsatisfiable w.h.p. It is conjectured (and there is also some empirical evidence) that solving random 3-CNFs becomes very hard when the density is close to the threshold.

Even for the high (constant) density while we know that the formula is most probably unsatisfiable no efficient way to verify it is known. Recently, Feige [Fei02] asked whether the following is true: no polynomial time algorithm may prove w.h.p. unsatisfiability of a random 3-CNF for arbitrary large constant density. The positive resolution of this conjecture would imply several interesting hardness results in the theory of approximation of NP-complete problems.

Random 3-CNFs are also an interesting candidate for proving lower bounds for propositional proof systems, i.e. non-deterministic proving algorithms. No short refutations of these formulas are known even for strong systems as Frege or Extended Frege. While there is little hope to prove lower bounds for such strong systems at the moment, one of more feasible goals is to prove combinatorially the hardness of refuting a random 3-CNF for those proof systems for which lower bounds are known. Random 3-CNFs were first considered in the framework of propositional calculus by V. Chvátal, E. Szemerédi. [CS88], who showed that these formulas are hard to refute for Resolution. Later this result was strengthened for resolution by several researchers [BKPS02], [BW01] and generalized for polynomial calculus [BI99], [AR01].

In the present work we investigate the complexity of random 3-CNFs for Res(k) systems. Introduced by Krajíček [Kra01] k-DNF Resolution is a generalized variant of Resolution which operates

*Institute for Advanced Study, Princeton, US. Supported by CCR grant $\mathcal{N}$CCR-0324906.

with k-DNFs instead of clauses. This is an interesting intermediate system between Resolution and depth two Frege. For Res(k) systems random CNFs were considered by Atserias, Bonet and Esteban [ABE02], who proved the exponential lower bound for Res(2) on random 3-CNFs. Recently, it was shown by Segerlind, Buss and Impagliazzo [SBI02] that random $O(k^2)$-CNFs are hard for Res(k) for $k \leq \sqrt{\log n / \log\log n}$. As one of the open questions they ask whether this bound may be improved for random 3-CNF. In the present work we give a positive answer to this question and prove that random 3CNFs are hard for Res(k) for the same range of k. Our improvement comes from carefully constructed random restrictions. The techniques we use give a general way to analyze the structure of solutions of a linear system over expanding 3-CNF. These techniques were used in a sequence of works [AHI04], present, [ABB+04] and we believe that it may turn out useful in the future. As usual, our result implies an exponential lower bound for a wide class of deterministic heuristics refuting a random 3-CNFs that produce a Res(k) refutation.

As an obvious direction for the further investigation we leave the open problem to strengthen our results for stronger systems. The paper is organized in the following way. Section 2 contains the notation and preliminary definitions. We construct our random restrictions and prove our main technical result (Restriction Lemma) in Section 3. After that we proceed along the lines of [SBI02] and prove the lower bound in Section 4.

2 Preliminaries

Let x be a Boolean variable, i.e. a variable that ranges over the set $\{0,1\}$. A *literal* of x is either x (denoted sometimes as x^1) or $\bar{x}$ (denoted sometimes as x^0). A *term* is a conjunction of literals. A *DNF* is a disjunction of terms. A *k-DNF* is a disjunction in which every term is of size at most k. A *clause* is 1-DNF, i.e. a disjunction of literals. For any Boolean function $f : \{0,1\}^n \to \{0,1\}$, $Vars(f)$ will denote the set of its essential variables. An *assignment to* f is a mapping $\alpha : Vars(f) \to \{0,1\}$. A *restriction of* f is a mapping $\rho : Vars(f) \to \{0,1,\star\}$. We denote by $|\rho|$ the number of assigned variables, $|\rho| \stackrel{\text{def}}{=} |\rho^{-1}(\{0,1\})|$. The *restriction of* f *by* ρ, denoted $f|_\rho$, is the Boolean function obtained from f by setting the value of each $x \in \rho^{-1}(\{0,1\})$ to $\rho(x)$, and leaving each $x \in \rho^{-1}(\star)$ as a variable.

For n, a non-negative integer let $[n] \stackrel{\text{def}}{=} \{1,2,\ldots,n\}$. For a vector $v = (v_1, ..., v_m)$ and index set $I \subseteq [m]$ we denote by v_I the subvector with coordinates chosen according to I. For a matrix A and a set of rows $I \subseteq [m]$ we use the notation A_I for the submatrix of A corresponding to these rows. In particular, we denote the ith row of A by A_i and identify it with the set $\{j \mid A_{ij} = 1\}$. The cardinality of this set is denoted by $|A_i|$.

2.1 Resolution and k-DNF Resolution

Resolution is a simple propositional proof system that operates with clauses and has one rule of inference called *resolution rule*:

$$\frac{A \vee x, B \vee \bar{x}}{A \vee B}.$$

A *resolution refutation* of a CNF formula τ is a resolution proof of the empty clause from the clauses appearing in τ. The *size* of a resolution proof is the number of different clauses in it. The *width* $w(C)$ *of a clause* C is the number of literals in C. The *width* $w(\tau)$ *of a set of clauses* τ (in particular, the width of a resolution proof) is the maximal width of the clauses appearing in this set. For an unsatisfiable set of clauses τ denote by $S_R(\tau)$ the size of minimal refutation in Resolution. Denote by $w_R(\tau)$ the minimal refutation width over all possible proofs of τ.

k-DNF Resolution (or Res(k)) is a generalization of Resolution that operates with k-DNFs and has the following inference rules:

1. Weakening: $\frac{F}{F \vee \ell}$
2. AND-introduction: $\frac{F \vee \ell_1, \ldots, F \vee \ell_w}{F \vee (\bigwedge_{i=1}^{w} \ell_i)}$
3. Cut: $\frac{F \vee (\bigwedge_{i=1}^{w} \ell_i), G \vee (\bigvee_{i=1}^{w} \bar{\ell}_i)}{F \vee G}$
4. AND-elimination: $\frac{F \vee (\bigwedge_{i=1}^{w} \ell_i)}{F \vee \ell_i}$

Res(k) refutation of unsatisfiable CNF τ is the inference of the empty clause from the clauses in τ using the rules 1-4. Similarly to Resolution, define $S_{R(k)}(\tau)$ as the size of the smallest Res(k) refutation of CNF τ. We will need the following results proved in [SBI02]:

Definition 2.1 ([SBI02], Definition 3.1). A *decision tree* is a rooted binary tree such that every internal node is labeled with a variable, the edges leaving this node correspond to whether the variable is set to 0 or 1, and the leaves are labeled with either 0 or 1. Every path from the root to a leaf may be viewed as a partial assignment. For a decision tree T and $v \in \{0, 1\}$, we write the set of paths that lead from the root to a leaf labeled v as $Br_v(T)$. We say that a decision tree T *strongly represents* a DNF F if for every $\pi \in Br_0(T)$ and for all $t \in F$, $t|_\pi = 0$ and for every $\pi \in Br_1(T)$, there exists $t \in F$ such that $t|_\pi = 1$. Let the representation height of F, $h(F)$ be the minimum height of a decision tree strongly representing F.

Definition 2.2 ([SBI02], Definition 3.2). Let F be a DNF and let S be a set of variables. If every term of F contains a variable from S then we say that S is a *cover* of F. The *covering number* of F, $c(F)$ is the minimum cardinality of a cover of F.

Lemma 2.1 ([SBI02],Corollary 3.4). Let k, s and d be positive integers, let γ and δ be real numbers from the range $(0, 1]$ and let $\mathcal{D}$ be a distribution on partial assignments so that for every k-DNF G, $\Pr[G|_\rho \neq 1] \leq d2^{-\delta(c(G))^\gamma}$. For every k-DNF F:

$$\Pr_{\rho \in \mathcal{D}}[h(F|_\rho) > 2s] \leq dk2^{-\delta' s^{\gamma'}},$$

where $\delta' = 2(\delta/4)^k$ and $\gamma' = \gamma^k$.

Lemma 2.2 ([SBI02], Theorem 5.1). Let τ be a set of clauses of width $\leq h$. If τ has a Res(k) refutation so that for each line F of the refutation $h(F) \leq h$, then $w_R(\tau) \leq kh$.

2.2 Expanders

A boundary expander is a bounded-degree graph that has many neighbors for every subset of its nodes. Similarly to [ABRW00], we use a more general notion of expander as an $m \times n$ matrix. We start with describing our notation.

Definition 2.3. For a set of rows $I \subseteq [m]$ of an $m \times n$ matrix A, we define its *boundary* $\partial_A I$ (or just ∂I) as the set of all $j \in [n]$ (called *boundary elements*) such that there exists exactly one row $i \in I$ that contains j. We say that A is an (r, c)-*boundary expander* if

$\forall I \subseteq [m]$ $(|I| \leq r \Rightarrow |\partial I| \geq c \cdot |I|)$.

2.3 Random 3CNFs and linear systems

In this paper we use the following model of generating a random 3CNF.

Definition 2.4. Denote by $\phi_{n,\Delta}$ the random 3-CNF with Δn clauses and n variables, in which every clause is chosen independently from the set of all $2^3\binom{n}{3}$ clauses.

Every 3CNF can be translated to a system of linear equations modulo 2, such that every clause of the CNF is a semantical corollary of some linear equation of the system (in particular every unsatisfiable 3CNF is translated into an unsatisfiable linear system):

Definition 2.5. For every value of the random variable $\phi_{n,\Delta}$ we denote a matrix $A_{n,\Delta}(\phi_{n,\Delta})$ and a vector $b(\phi_{n,\Delta})$ in the following way. Let $A_{n,\Delta}$ be $\Delta n \times n$ matrix in which $(A_{n,\Delta})_{ij} = 1$ iff the ith clause of $\phi_{n,\Delta}$ contains (either negatively or positively) the variable x_j. Define a vector $b \in \{0,1\}^{\Delta n}$ so that for the ith clause $(x_{j_1}^{\epsilon_1} \vee x_{j_2}^{\epsilon_2} \vee x_{j_3}^{\epsilon_3})$ $b_i = \epsilon_1 + \epsilon_2 + \epsilon_3$.

Instead of proving a lower bound on the size of the refutation of $\phi_{n,\Delta}$ we will prove a stronger bound on the refutation of the system $A_{n,\Delta}x = b$. Note that $A_{n,\Delta}$ is a random matrix in which each row contain 3 ones. We need the following well known fact (first noticed in [CS88]) saying that the random matrix is a good expander.

Fact 2.1. For all constants $\Delta > 0$, $c < 1$ there exists some constant δ s.t. with probability $1 - o(1)$ $A_{n,\Delta}$ is $(\delta n, c)$-boundary expander.

2.4 Various closures of a set of columns

We make use of two different operations of taking closures. The first was defined in [AR01].

Definition 2.6. Let $A \in \{0,1\}^{m\times n}$. For a set of columns $J \subseteq [n]$ define the following inference relation $\vdash_J$ on the set $[m]$ of rows of A:

$$I \vdash_J I_1 \iff |I_1| \le r/2 \ \wedge\ \partial_A(I_1) \subseteq \left[\bigcup_{i\in I} A_i \cup J\right]. \tag{1}$$

Let the *closure* $\mathrm{Cl}(J)$ of J be the set of all rows which can be inferred via $\vdash_J$ from the empty set (i.e. the minimal set of rows I such that $\forall I'\ I \vdash_J I' \rightarrow I' \subseteq I$).

The following lemma was proved in [AR01] (Lemma 3.16 in the cited work). For the completeness we include the proof in the appendix.

Lemma 2.3. For any set J with $|J| \le (cr/2)$, $|\mathrm{Cl}(J)| \le c^{-1}|J|$.

We also need another (stronger) closure operation the intuitive sense of which is to extract a good expander out of a given matrix by removing rows and columns. The following definition has appeared in [AHI04].

Definition 2.7. For an $A \in \{0,1\}^{m\times n}$ and a subset of its columns $J \subseteq [n]$ we define an inference relation $\vdash_J^e$ on the set of $[m]$ rows of A:

$$I \vdash_J^e I_1 \iff |I_1| \le r/2 \ \wedge\ \left|\partial_A(I_1) \setminus \left[\bigcup_{i\in I} A_i \cup J\right]\right| < (c/2)|I_1| \tag{2}$$

Given a set of rows I and a set of columns J consider the following *cleaning* step:

- If there exists a nonempty subset of rows I_1 such that $I \vdash^e_J I_1$, then
  - Add I_1 to I.
  - Remove all rows corresponding to I_1 from A.

Repeat the cleaning step as long as it is applicable. Fix any ordering to exclude ambiguity, initialize $I = \emptyset$ and denote the resulting content of I at the end by $\mathrm{Cl}^e(J)$.

Lemma 2.4. Assume that A is an arbitrary matrix and J is a set of its columns. Let $I' = \mathrm{Cl}^e(J)$, $J' = \bigcup_{i \in \mathrm{Cl}^e(J)} A_i$. Denote by $\hat{A}$ the matrix that results from A by removing the rows corresponding to I' and columns to J'. If $\hat{A}$ is non-empty than it is $(r/2, c/2)$-boundary expander.

Proof. Indeed, assume that there exists a set of rows I_1 in $\hat{A}$ of size less than or equal $r/2$ for which $\partial_{\hat{A}}(I_1) \le (c/2)|I_1|$. Then $I' \vdash_{J'} I_1$ and the matrix $\hat{A}$ is not closed. □

The following lemma is proved in [AHI04]. For the sake of completeness we give the proof in the appendix.

Lemma 2.5. If $|J| < cr/4$, then $|\mathrm{Cl}^e(J)| < 2c^{-1}|J|$ no matter in what order have cleaning steps taken place.

For a term t we will write $\mathrm{Cl}(t)$ and $\mathrm{Cl}^e(t)$ for $\mathrm{Cl}(\mathrm{Vars}(t))$ and $\mathrm{Cl}^e(\mathrm{Vars}(t))$ respectively.

2.5 Metric on a set of columns and rows

Definition 2.8. For $A \in \{0,1\}^{m \times n}$ let $G(A)$ be the corresponding bipartite graph between m row vertices and n column vertices with incidence matrix A. For two vertices v_1, v_2 each of which is either a column or a row, the distance $d_A(v_1, v_2)$ is the length of the shortest path between them. For two subsets $J_1, J_2 \subseteq [m] \sqcup [n]$ we let

$$d_A(J_1, J_2) = \min_{v_1 \in J_1, v_2 \in J_2} d_A(v_1, v_2).$$

Proposition 2.1. Assume that every column of A has at most Δ ones for some $\Delta \ge 3$. For any set J there is at most $|J| \cdot \Delta^k$ points v for which $d_A(v, J) \le k$.

□

3 Random restriction lemma

In this section we define the special random restrictions and prove the main technical result (Theorem 3.1) that will be used in the next section to lower bound the size of $\mathrm{Res}(k)$ refutation.

Definition 3.1. Let $A \in \{0,1\}^{m \times n}$ be (r,c)-boundary expander and $b \in \{0,1\}^m$, let $X = \{x_1, ..., x_n\}$ be the set of all variables. Denote by $\rho_{A,b}$ a random restriction on X that results from the following experiment. Choose a random subset $X_1 \subset X$ of size $cr/4$. Denote by $\hat{I} = \mathrm{Cl}^e(X_1)$, $\hat{X} = X_1 \cup \{x_j | \; \exists i \in \hat{I} \; A_{ij} = 1\}$. The restriction $\rho_{A,b}$ assigns a random partial assignment chosen uniformly from all $\hat{x} \in \{0,1\}^{\hat{X}}$ satisfying

$$A_{\hat{I}} \hat{x} = b_{\hat{I}}.$$

The following definition (as well as Lemma 3.1) appeared simultaneously in the present paper and in [AHI04].

Definition 3.2. For $A \in \{0,1\}^{m\times n}, b \in \{0,1\}^m$, a term t is said to be *locally consistent* iff the formula

$$t \wedge [A_I x = b_I]$$

is satisfiable, where $I = \mathrm{Cl}(t)$.

The following statement shows that a term is locally consistent w.r.t. a linear system iff it is consistent w.r.t. each of its subsystems of small size:

Lemma 3.1. Assume that term t is locally consistent. Then for any set I with $|I| < r/2$ the formula

$$t \wedge [A_I x = b_I]$$

is satisfiable.

Proof. Let us regard the conjunction of literals in t as a conjunction of linear equations (having one variable each). Assume for the contradiction that linear equations corresponding to t and $A_I x - b_I$ are inconsistent. From the basic linear algebra this implies the existence of two sets $t' \subseteq t$, and $I' \subseteq I$ s.t.

$$\sum_{i \in I'} (A_i x - b_i) + \sum_{x_j^\epsilon \in t'} (x_j - \epsilon) \equiv 1.$$

This in particular implies $\partial_A(I') \subseteq \mathrm{Vars}(t')$, hence $I' \subseteq \mathrm{Cl}(t)$. This however contradicts to the fact that t is locally consistent. □

Definition 3.3. A DNF ϕ is in normal form w.r.t. A, b if each of its terms is locally consistent.

Theorem 3.1 (Restriction Lemma). Let $A \in \{0,1\}^{m\times n}$ be (r,c)-boundary expander in which every column contains at most Δ ones, $b \in \{0,1\}^m$ be an arbitrary vector and assume that $r = \Omega(n/\Delta)$. For any k-DNF ϕ in normal form holds

$$\Pr[\phi|_\rho \not\equiv 1] < \left(1 - 2^{-k}\right)^{c(\phi)/\Delta^{O(k)}}.$$

Corollary 3.1. There exists an absolute constant D s.t. under the assumptions of the theorem for any normal form k-DNF ϕ

$$\Pr[\phi|_\rho \not\equiv 1] < 2^{-c(\phi)/\Delta^{Dk}}.$$

Proof of Theorem 3.1.

Lemma 3.2. Let A be (r,c)-boundary expander. Assume that I is a set of rows of A, $|I| < r/2$ and t is a term such that the formula

$$t \wedge [A_I x = b_I]$$

is satisfiable. Then for any locally consistent term t_1 with $|t_1| \le k$ s.t.

$$d_A(\mathrm{Cl}^e(t), t_1) > 4c^{-1}k$$

the formula

$$t_1 \wedge t \wedge [A_I x = b_I] \tag{3}$$

is also satisfiable.

Proof. Let us regard the conjunction of literals in t, t_1 as a conjunction of linear equations (having one variable each). Assume for the contradiction that (3) is unsatisfiable. This means that the linear equations corresponding to t, t_1 and $A_I x = b_I$ are inconsistent. Just as in the proof of Lemma 3.1 this implies the existence of three sets $t' \subseteq t$, $t'_1 \subseteq t_1$ and $I' \subseteq I$ s.t.

$$\sum_{i \in I'} (A_i x - b_i) + \sum_{x_j^\epsilon \in t'} (x_j - \epsilon) + \sum_{x_l^\epsilon \in t'_1} (x_l - \epsilon) \equiv 1. \tag{4}$$

Definition 3.4. Assume that $\mathcal{L}$ contains linear equations of $Ax = b$ and literals. Denote by $G_{\mathcal{L}}$ the induced subgraph of $G(A)$ that contains the corresponding row vertex for each linear equation in $\mathcal{L}$ and the corresponding column vertex for each variable contained in $\mathcal{L}$ (either in literals or in equations).

Let us consider a subcombination $\mathcal{L} = (I', t', t'_1)$ of (4) with the minimal number of equations that sum up to one. In this case the graph $G_{\mathcal{L}}$ is connected (otherwise one may split the equations in (4) into two smaller sets corresponding to the connected components in $G_{\mathcal{L}}$ which both sum up to a constant, this would contradict to the minimality condition).

It follows from (4) that $\partial_A(I') \subseteq \mathrm{Vars}(t') \cup \mathrm{Vars}(t'_1)$. By the assumption in the statement t is consistent with $A_I x - b_I$, by Lemma 3.1 t_1 is also consistent with $A_I x - b_I$. Thus, the equation (4) can hold only when both t' and t'_1 are non-empty. Note that $\partial(I') \setminus \mathrm{Vars}(t) \subseteq \mathrm{Vars}(t_1)$.

Case 1. $|I' \setminus \mathrm{Cl}^e(t)| > 2c^{-1}k$. In this case

$$|\partial(I') \setminus \left[\bigcup_{i \in \mathrm{Cl}^e(t)} A_i \cup \mathrm{Vars}(t) \right] | \leq k < (c/2)|I' \setminus \mathrm{Cl}^e(t)|,$$

thus $\mathrm{Cl}^e(t)$ is not closed since we may add $I' \setminus \mathrm{Cl}^e(t)$.

Case 2. $|I' \setminus \mathrm{Cl}^e(t)| \leq 2c^{-1}k$. In this case there exists a path of length $4c^{-1}k$ in $G(A)$ that connects $\mathrm{Cl}^e(t)$ and t_1. To see this consider the minimal path in $G_{\mathcal{L}}$ that connects equations corresponding to t with those corresponding to t_1, this path goes along the equations in I'. Given this path one may construct a path of length $2|I' \setminus \mathrm{Cl}^e(t)|$ that connects $\mathrm{Cl}^e(t)$ with t_1 in $G(A)$. This however contradicts to the assumption of the lemma. □

We are ready now to prove Theorem 3.1. Consider some k-DNF ϕ in normal form and the restriction $\rho_{A,b}$. Let $X_1, \hat{X}, \hat{I}$ be the corresponding random variables from Definition 3.1. We may extract at least $C_0 = c(\phi)/k$ variable disjoint terms, each of which will be covered by X_1 with probability at least $(cr/4n)^k$, thus the expected number of covered disjoint terms is $C_0 \left(\frac{cr}{4n}\right)^k$. By Chernoff bound we may assume that there exist $C_1 = C_0/\Delta^{O(k)}$ variable disjoint terms covered by X_1. The following lemma follows from the basics of linear algebra.

Lemma 3.3. Assume that ξ is a random vector uniformly distributed on the affine subspace $\mathcal{A} \subseteq \{0,1\}^X$, $\mathcal{A} = a + \mathcal{L}$, where $\mathcal{L}$ is a linear subspace of $\{0,1\}^X$. For any maximal I s.t. vectors $\{e_i\}_{i \in I}$ are linearly independent modulo $\mathcal{L}$ the restriction ξ_I is a random variable uniformly distributed over $\{0,1\}^I$ and the value of ξ is fully determined by ξ_I.

Corollary 3.2. Let $Y \subset X$ be a set of variables. Assume that b is a partial assignment on Y distributed uniformly on some affine subspace $\mathcal{A} \subseteq \{0,1\}^Y$. For any term t in Y variables either $\Pr[t|_b \equiv 1] = 0$ or $\Pr[t|_b \equiv 1] \geq 2^{-|t|}$.

Proof. Assume that $\mathcal{A} = a + \mathcal{L}$, where $\mathcal{L}$ is a linear space. Let us order the bits of $t = t_1 \wedge t_2 \wedge \ldots \wedge t_k$, for $k = |t|$ and choose the set $I \subset [k]$ such that

- $\{e_i\}_{i \in I}$ are linearly independent modulo $\mathcal{L}$.

- If $i \in [k]$ but $i \notin I$ then b_i is linearly expressible by $\{b_j\}_{j<i}$.

In this case the event $t|_b \equiv 1$ is dependent only on the bits b_I which are distributed uniformly on $\{0,1\}^I$. Thus, in the probability space induced by b_I the probability of any elementary event is at least $2^{-|I|} \geq 2^{-k}$. □

Let us observe the partial assignment $\hat{x} \in \{0,1\}^{\hat{X}}$ given by $\rho_{A,b}$. This is a random variable distributed uniformly on some affine subspace $\mathcal{A} \subset \{0,1\}^{\hat{X}}$. We define the following experiment. Assume that all bits of $\hat{x}$ are hidden. In the base consider a term t_1 of ϕ. Since it is locally consistent there exists a satisfying assignment for the formula

$$t_1 \wedge \left[A_{\hat{I}} x = b_{\hat{I}}\right].$$

Let us reveal the bits of $\hat{x}$ corresponding to t_1. Denote by E_1 the event that t_1 is satisfied, by Lemma 3.2 $\Pr[E_1] \geq (1/2)^k$. If E_1 occurs then we terminate the process successfully, otherwise we proceed according to the following inductive step.

Assume that at step ℓ we revealed values of terms $t_1, ..., t_\ell$ of ϕ. Let $t^{(\ell)}$ be the term corresponding to the partial assignment of revealed bits of $\hat{x}$ on $\mathrm{Vars}(t_1) \cup ... \cup \mathrm{Vars}(t_\ell)$, thus $\hat{x}_{\mathrm{Vars}(t^{(\ell)})} = t^{(\ell)}$, $|t^{(\ell)}| \leq \ell k$. Consider the set of variables $Y^{(\ell)} \subseteq \hat{X}$ located within the distance $4c^{-1}k$ from $\mathrm{Cl}^e(t^{(\ell)})$. If there is no term $t_{\ell+1}$ in ϕ free of $Y^{(\ell)}$-variables then we terminate the process unsuccessfully. Otherwise, consider the term $t_{\ell+1}$ and reveal the corresponding bits of $\hat{x}$. Let $E_{\ell+1}$ be the event that $t_{\ell+1}$ is satisfied. We apply Lemma 3.2 to show that

$$\Pr[E_{\ell+1} | \hat{x}_{\mathrm{Vars}(t^{(\ell)})} = t^{(\ell)}] > 0$$

and then use Corollary 3.2 to conclude that $\Pr[E_{\ell+1} | \hat{x}_{\mathrm{Vars}(t^{(\ell)})} = t^{(\ell)}] \geq (1/2)^k$.

Let T be the stopping time of the inductive process, i.e. the last index ℓ for which we did the inductive step. If at least one event in the list $E_1, E_2, ..., E_T$ has occurred then ϕ is killed to 1. Thus, all we need is to show that T cannot be small in case of unsuccessful termination. First consider the case when $kT \leq cr/4$. Then by Lemma 2.5 $\mathrm{Cl}^e(t^{(T)}) \leq 2c^{-1}Tk$ and by Proposition 2.1

$$|Y^{(T)}| \leq 2c^{-1}Tk \cdot \Delta^{4c^{-1}k}.$$

Since $Y^{(T)}$ is a hitting set for ϕ we have

$$C_1 \leq |Y^{(T)}| \leq 2c^{-1}Tk \cdot \Delta^{4c^{-1}k}$$

and

$$T \geq C_1/\Delta^{O(k)}.$$

In the other case $T > cr/(4k)$, thus $T \geq C_2$, where $C_2 = \min(C_1/\Delta^{O(k)}, cr/(4k))$. Because $r = \Omega(n)$ and $c(\phi) \leq n$ we infer that $C_2 = c(\phi)/\Delta^{O(k)}$. The theorem follows. □

4 Lower bound for Res(k)

In this section we prove a lower bound for Res(k) based on the random restriction constructed in Section 3. The overall argument follows the strategy of [SBI02].

Theorem 4.1. For any constant Δ with probability $1 - o(1)$ every $\mathrm{Res}(k)$ refutation of $\phi_{n,\Delta}$ for $k < \sqrt{\log n / \log\log n}$ has size $2^{n^{1-o(1)}}$.

Proof. Recall Definition 2.5 of the random matrix $A_{n,\Delta}$. By Fact 2.1, with high probability $A_{n,\Delta}$ is $(r, 0.8)$- boundary expander for some $r = \Omega(n)$. From this and now on, let us fix r, $A_{n,\Delta}$ and assume that it is $(r, 0.8)$-boundary expander. In order to apply Theorem 3.1 we also need our matrix contain finitely many ones in each column. This is not the case for $A_{n,\Delta}$, however we may delete the columns of high weight from it.

Proposition 4.2. For any $b \in \{0,1\}^m$ the empty term is locally consistent w.r.t. $Ax = b$.

Definition 4.1. For a matrix $A_{n,\Delta}$ let J be the set of $0.2r$ columns of the largest hamming weight. Denote $I' = \mathrm{Cl}^e(J)$, $J' = \bigcup_{i\in I'} A_i$. Denote by $\hat{A}_{n,\Delta}$ the matrix $A_{n,\Delta}$ with all columns corresponding to J' removed and all rows corresponding to I' removed. Similarly define $\hat{b} = b_{[\Delta n]\setminus I'}$.

Lemma 4.1. The matrix $\hat{A}_{n,\Delta}$ is $(r/2, 0.4)$-boundary expander, in which every column has weight bounded by some absolute constant $\hat{\Delta}$ that depends on Δ only.

Proof. $\hat{A}_{n,\Delta}$ is $(r/2, 0.4)$-boundary expander by Lemma 2.4. Since the matrix $A_{n,\Delta}$ contains exactly $3\Delta n$ ones, the matrix $\hat{A}_{n,\Delta}$ may contain at most $3\Delta n/(0.4r)$ ones in each column. Since $r = \Omega(n)$ the lemma follows. ◻

Lemma 4.2. Every $\mathrm{Res}(k)$ refutation of $\phi_{n,\Delta}$ can be transformed into $\mathrm{Res}(k)$ refutation of the system

$$\hat{A}_{n,\Delta}x = \hat{b}, \tag{5}$$

in which every line is in normal form with only polynomial increase of the size.

Proof. Clearly every clause in $\phi_{n,\Delta}$ may be easily inferred from the encoding of the system $A_{n,\Delta}x = b$. Let I', J' be defined as in Definition 4.1. By Proposition 4.2 and Lemma 3.1 we may assign values to $x_{J'}$ so that all the equations in $A_{I'}x = b_{I'}$ are satisfied. Then we get a $\mathrm{Res}(k)$ refutation of the system (5). It is left to show that this refutation can be transformed into normal form.

For every term t that is not locally consistent we may infer $\bar{t}$ from $2k$ axioms in polynomial size in Resolution. Thus we may substitute any occurrence of locally inconsistent terms with $\perp$ with the polynomial increase of the proof. The only subtle point are the singletons x_j^{ϵ}, however our expansion condition on $\hat{A}_{n,\Delta}$ implies that $\forall j\ \mathrm{Cl}(\{j\}) = \emptyset$, thus all singletons are locally consistent. ◻

Lemma 4.3 (Ben-Sasson – Wigderson [BW01]). Assume that the matrix $\hat{A}_{n,\Delta}$ is (r, c)-boundary expander, c is a positive constant. Then every resolution refutation of the system (5) requires width ϵr, where $\epsilon > 0$ depends only on c.

We are ready to finish the proof of Theorem 4.1. Consider some outcome of the random variable $\phi_{n,\Delta}$, let $\hat{A}_{n,\Delta}$ and $\hat{b}$ be defined according to Definitions 2.5, 4.1. In the assumption that there exists $\mathrm{Res}(k)$ refutation of $\phi_{n,\Delta}$ of size S we conclude that there also exists $\mathrm{Res}(k)$ refutation $\mathcal{P}$ of the system (5) of size $Sn^{O(1)}$, which is in normal form. Next we apply the restriction $\rho_{\hat{A}_{n,\Delta},b}$ constructed in the previous section and invoke Corollary 3.1. Using Lemmas 2.1, 2.2 plugging in parameters $\gamma = 1$, $\delta = (1/\Delta)^{Dk}$, $s = \epsilon r$ we conclude that the restricted proof $\mathcal{P}|_{\rho_{\hat{A}_{n,\Delta},b}}$ has width greater than ϵr with probability at least

$$1 - S \cdot k 2^{-\epsilon r (1/\Delta)^{2Dk^2}} > 1 - S 2^{-n/2^{O(k^2)}}.$$

On the other hand, by Lemma 4.3 the probability of this event is zero. Indeed, the formula resulting after the restriction still encodes a linear system over $(r/4, 0.2)$-boundary expander matrix, thus the restricted system has always high resolution refutation width. Altogether this implies $S > 2^{n/2^{O(k^2)}}$ and the theorem follows. $\square$

5 Acknowledgments

We are grateful to Alexander Razborov and Nathan Segerlind for helpful discussions.

References

[ABB+04] Michael Alekhnovich, Allan Borodin, Joshua Buresh-Oppenheim, Russell Impagliazzo, Avner Magen and Toniann Pitassi. Toward a Model for Backtracking and Dynamic Programming. To appear in *Proc. 20th Annual Conference on Computational Complexity*, 2005.

[ABE02] A. Atserias, M. Bonet, J. Esteban. Lower bounds for the weak pigeonhole principle and random formulas beyond resolution. *Information and Computation, 176:136–152*, 2002.

[ABRW00] M. Alekhnovich, E. Ben-Sasson, A. Razborov, and A. Wigderson. Pseudorandom generators in propositional complexity. *SIAM Journal on Computing* 34(1): 67-88, 2004.

[AHI04] M. Alekhnovich, E. Hirsch and D. Itsykson. Exponential lower bounds for the running time of DPLL algorithms on satisfiable formulas. To appear in 2005 SAT special issue of the *Journal of Automated Reasoning*. See also *Proc. 31st ICALP*: 84-96, 2004.

[AR01] M. Alekhnovich and A. Razborov. Lower bounds for the polynomial calculus: non-binomial case. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*, pages 190–199, 2001.

[BI99] E. Ben-Sasson and R. Impagliazzo. Random CNF's are Hard for the Polynomial Calculus. In *Proceedings of the 40th IEEE FOCS*, pages 415–421, 1999.

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

[BW01] E. Ben-Sasson and A. Wigderson. Short proofs are narrow — resolution made simple. *Journal of ACM*, 48(2):149–169, 2001.

[CS88] V. Chvátal, E. Szemerédi. Many Hard Examples for Resolution. In *Journal of of the ACM*, 35(4):759–768, 1988.

[Fei02] U. Feige. Relations between Average Case Complexity and Approximation Complexity. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, 534–543, 2002.

[Fri99] E. Friedgut. Sharp Thresholds of Graph Properties, and the k-sat Problem. In *J. Amer. Math. Soc. 12 (1999)*, no. 4, 1017–1054.

[Kra01] J. Krajíček On the weak pigeonhole principle. *Fundamenta Mathematicae*, 170:123–140, 2001.

[SBI02] N. Segerlind, S. R. Buss, R. Impagliazzo. A switching lemma for small restrictions and lower bounds for k - DNF Resolution. To appear in *SIAM journal on Computing.* See also *Proc. of the* 43*nd IEEE Symposium on Foundations of Computer Science*, 2002.

A Bounds on taking closure

Here we include the bounds on the size of $\mathrm{Cl}(J)$ due to [AR01] and $\mathrm{Cl}^e(J)$ due to [AHI04].

Lemma 1.1. For a set $J \subseteq [n]$ with $|J| \leq (cr/2)$, $|\mathrm{Cl}(J)| \leq c^{-1}|J|$.

Proof. Assume the contrary, and choose a chain of subsets $I_1, I_2, \dots$ such that $I_1 \cup \ldots \cup I_{\nu-1} \vdash_J I_\nu$ and $|\cup_\nu I_\nu| > c^{-1}|J|$. Let k be the *smallest* index for which $|\bigcup_{\nu=1}^{k} I_\nu| > c^{-1}|J|$. Then, clearly, $|\bigcup_{\nu=1}^{k} I_\nu| \leq r$ (since $|J| \leq cr/2$). Therefore, $|\partial_A(\bigcup_{\nu=1}^{k} I_\nu)| > |J|$. On the other hand, property (1) of Definition 2.6 implies that every new boundary element that results from appending via $\vdash_J$ some set of rows must belong to J, therefore $\partial_A(\bigcup_{\nu=1}^{k} I_\nu) \subseteq J$. This contradiction proves Lemma 1.1. □

Lemma 1.2. If $|J| < cr/4$, then $|\mathrm{Cl}^e(J)| < 2c^{-1}|J|$ no matter in what order have cleaning steps taken place.

Proof. Assume that $|\mathrm{Cl}^e(J)| \geq 2c^{-1}|J|$. Consider the sequence $I_1, I_2, \dots, I_t$ appearing in the cleaning procedure; i.e.,

$$J \cup \bigcup_{i \in I_1 \cup I_2 \cup \ldots \cup I_k} A_i \vdash_A I_{k+1}.$$

Note that $I_i \cap I_{i'} = \emptyset$ for $i \neq i'$, because we remove the implied set of rows from A at each cleaning step. Denote $C_t = \bigcup_{k=1}^{t} I_k$.

Let T be the first t such that $|C_t| > 2c^{-1}|J|$. Note that $|C_T| \leq r$. Because of the expansion properties of A, $\partial C_T > c|C_T|$, which implies

$$|\partial C_T \setminus J| > c|C_T| - |J| \geq c|C_T|/2. \tag{6}$$

On the other hand, every time we add some I_{t+1} to C_t during the cleaning procedure, only $(c/2)|I_{t+1}|$ new elements can be added to $\partial C_t \setminus J$ (of those elements that have never been there before). This implies

$$|\partial C_T \setminus J| \leq c/2|C_T|,$$

which contradicts (6). □