Exponential Lower Bounds for Semantic Resolution[*]

S. Jukna[†‡]

Abstract

In a semantic resolution proof we operate with clauses only but allow *arbitrary* rules of inference: consistency is the only requirement. We prove a very simple exponential lower bound for the size of bounded fanin semantic resolution proofs of a general *Hitting Set Principle* stating that, for any set system with hitting set number τ, no set of size less than τ can be a hitting set. The pigeonhole principle and known blocking principles for finite (affine and projective) planes are special cases of this general principle.

1. Introduction

The resolution proof system introduced by Blacke [2] and further developed by Davis and Putnam [12] and Robinson [17] is one of the first and simplest in the hierarchy of propositional proof systems. This system operates with clauses and has one rule of inference

$$\frac{C_1 \vee x_i \quad C_2 \vee \neg x_i}{C_1 \vee C_2}$$

called the resolution rule. First exponential lower bound for regular resolution (these are resolution proofs with the additional restriction that along every path every particular variable x_i can be resolved at most once) was proved by Tseitin [18] almost 30 years ago. However, despite its apparent simplicity, the first lower bounds for non-regular resolution were proven only in 1985 by Haken [13]. These bounds were achieved for the pigeonhole principle PHP_n^{n+1} asserting that $n+1$ pigeons can not sit in n holes so that every pigeon is alone

[*] To appear in: *Proc. of DIMACS Workshop on Feasible Arithmetics and Length of Proofs* (April 21-23, 1996), P. Beame and S. Buss eds., AMS-DIMACS Series in Discrete Mathematics and Theoretical Computer Science

[†] Research supported by DFG grant Me 1077/10–1.

[‡] Department of Computer Science, University of Trier, D-54286, Trier, Germany. On leave from Department of Mathematical Logic, Institute of Mathematics, Akademijos 4, Vilnius 2000, Lithuania. E-mail: jukna@ti.uni-trier.de

in his hole. Buss and Turán [8] extended this bound to $\exp(\Omega(n^2/m))$ for more general form PHP_n^m of the pigeonhole principle in which the number of pigeons, m is another parameter. Haken's argument was further refined and applied to other tautologies by Urquhart [19] and Chvátal and Szemerédi [9].

In a recent work [1] Beame and Pitassi have found a direct and elegant proof of Haken's [13] lower bound for PHP_n^{n+1}. In this paper we simplify and generalize the combinatorial part of their argument and show that it in fact works: (i) for other principles than PHP_n^m, and (ii) for proof systems that generalize resolution. Examples of these new principles are, so-called, blocking principles for finite (affine and projective) planes. The model generalizing resolution is that of *semantic resolution*. Like in standard resolution proof, here we operate with clauses only. The main difference is that we allow arbitrary inference rules

$$\frac{C_1,\ldots,C_l}{C}$$

Their *consistency* is the only requirement: every truth assignment satisfying all the hypotheses $C_1,\ldots,C_l$, must also satisfy the conclusion C. The number of hypotheses, l is the *fanin* of that rule. The resolution rule is a very special case of this general rule with $l = 2$. The *size* (or *length*) of a proof is the total number of clauses in it.

We will prove a very simple exponential lower bound for the size of semantic resolution proofs of a general principle, which we call the *Hitting Set Principle*. This principle, $HS(\mathcal{F})$ states that for any set system $\mathcal{F}$ with hitting set number τ, no set A of size less than τ can be a hitting set. The *hitting* (or *blocking*) *set* for $\mathcal{F}$ is a set which hits (i.e. intersects) every edge of $\mathcal{F}$; the *hitting number* $\tau(\mathcal{F})$ is the minimal possible size of such a set. This number is hard-coded into the CNF, and the underlying variables are the variables describing A and other variables may be used to explain that A has the right size. This principle is a generalization of the pigeonhole principle PHP_n^m: here $\mathcal{F}$ consists of m mutually disjoint n-element sets. The sets A of interest are all sets of n points which define partial 1-to-1 mappings from n of the pigeons to the n holes. The pigeonhole principle states precisely that either A is too large, or some edge is not intersected by A (see also Example 1 below).

Our main result (Theorem 1) says that as long as the set system $\mathcal{F}$ satisfies certain combinatorial conditions, then any CNF formula formalizing the hitting set principle for $\mathcal{F}$ requires an exponentially long semantic resolution proof. Our argument is essentially the same argument employed by Beame and Pitassi [1] in the case of PHP_n^{n+1}. Using simple "greedy" algorithm, we first show that some small restriction will kill off all long 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 long clauses in it.

2. The lower bound

In this section we state our main result and describe several its applications. We first need to setup some notation. A *hypergraph* (or *set system*) over a set X is simply a family $\mathcal{F}$ of its subsets; elements of X are *points*, and sets in $\mathcal{F}$ are *edges*. We will be interested in the size of semantic resolution proofs for the hitting set principles $HS(\mathcal{F})$. In any such proof we have $|\mathcal{F}|$ leaves labeled by (positive) clauses $C_E = \bigvee_{i \in E} x_i$, one for each edge $E \in \mathcal{F}$. We call these leaves *primary*. All other leaves are *secondary* and may be labeled by arbitrary clauses. In particular, besides the x-variables (corresponding to points of $\mathcal{F}$) these clauses may contain any other variables. We require only that the conjunction of all these secondary clauses must be satisfiable on any set of size less than $\tau(\mathcal{F})$. For example, one can take as secondary the following set of clauses:

$$y_{i1} \vee \cdots \vee y_{im}, \quad \neg y_{ik} \vee \neg y_{jk}, \quad \neg y_{ik} \vee \neg x_k$$

where $1 \leq i \neq j \leq m - \tau(\mathcal{F}) + 1$, $1 \leq k \leq m$ and $m = |X|$ is the total number of points in the hypergraph $\mathcal{F}$. It is easy to see that, given an assignment to x-variables, this set of clauses is satisfiable if and only if the number of ones in that assignment is less than $\tau(\mathcal{F})$. For now, let us point out that our lower bounds argument *does not* depend on the actual form of these secondary clauses: important will be only combinatorial properties of primary clauses, corresponding to the edges of $\mathcal{F}$.

Simple (but useful for the rest of the paper) observation is that, due to consistency of inference rules, any semantic resolution proof for $HS(\mathcal{F})$ is in fact a nondeterministic algorithm for the following *search problem*: Given a set A of size less than $\tau(\mathcal{F})$, find an edge $E \in \mathcal{F}$ such that $A \cap E = \emptyset$. To see this, associate with each such set A an assignment u_A to the remaining variables so that $f(v_A, u_A) = 1$, where f is the conjunction of all secondary leaves and v_A is the incidence vector of A. Fix this injection $A \mapsto u_A$, and traverse the proof starting from the last clause of the proof (which is empty) by always choosing that of hypotheses C, for which $C(v_A, u_A) = 0$. Consistency ensures that proceeding this way we will necessarily reach a leaf. The fact that $C(v_A, u_A) = 1$ for all the secondary leaves C, ensures that the reached leaf is primary, i.e. has the form $C_E = \bigvee_{i \in E} x_i$ with $E \in \mathcal{F}$. Since $C_E(v_A, u_A) = 0$, the edge E avoids our set A, and we are done: the desired edge is found.

This observation allows us to concentrate on the lower bounds problem on the length of semantic resolution proofs, solving the search problem for particular hypergraphs $\mathcal{F}$: any such bound is immediately a lower bound on the length of a shortest semantic resolution proof of any CNF, describing the hitting set principle for $\mathcal{F}$.

We will be particularly interested in special k-partite hypergraphs. Let $S_1, \ldots, S_k$ be mutually disjoint subsets of X, called *blocks*. A *partial transversal* is a set $B \subseteq X$ which intersects each block in at most one point; B is a *transversal* if $|B| = k$ (in this case B intersects each block in exactly one point).

Definition. We call a hypergraph $\mathcal{F}$ a (k, b, λ, d)-*design* if there exist k mutually disjoint blocks $S_1, \ldots, S_k$ such that:

1. Every edge of $\mathcal{F}$ is a transversal for $S_1, \ldots, S_k$;
2. $|S_i| \le b$ for all $i = 1, \ldots, k$;
3. $|E \cap F| \le \lambda$ for all edges $E \ne F \in \mathcal{F}$;
4. Every point belongs to at most d edges of $\mathcal{F}$.

Such a design $\mathcal{F}$ is *large* if every transversal of $S_1, \ldots, S_k$ avoids at least one edge of $\mathcal{F}$. The corresponding *edge-search problem* for $\mathcal{F}$ is to find such an edge. Note that any design, with more than kd edges, is large, but there also are large designs with smaller number of edges.

Our main result is the following general lower bound for semantic resolution.

Theorem 1. *Let $\mathcal{F}$ be a large (k, b, λ, d)-design, and G be a semantic resolution proof of fanin at most l. Let s and t be integers satisfying*

$$ls \le \min\{|\mathcal{F}| - dt, k - t\} \quad \textit{and} \quad t \ge k/2 \tag{1}$$

If G solves the edge-search problem for $\mathcal{F}$ then

$$\mathrm{size}(G) \ge 2^{M/b} \quad \textit{where} \quad M = \frac{s(k - t - ls + 1)^2}{k + \lambda \cdot (s - 1)} \tag{2}$$

In particular, if $|\mathcal{F}| \ge k(d + 1)/2$ then

$$\mathrm{size}(G) \ge \exp\left(\Omega\left(\frac{k^2}{b(l + \lambda)}\right)\right). \tag{3}$$

Few words about the parameters. The bound (3) follows from (2) by taking $t = \lceil k/2 \rceil$ and $s = \lceil k/(4l) \rceil$. The bound itself becomes trivial if the block size b is near to the square of their number k. But this is inherent weakness of our argument (as well as all previous lower bound proofs for PHP_n^m): in order to eliminate a single variable we are forced to do this simultaneously for all the variables in whole block (see also Remark 2 in Section 5).

To motivate the rest of the paper, let us mention several applications of Theorem 1.

Example 1. (Pigeonhole principle). The generalized pigeonhole principle PHP_n^m $(m \ge n + 1)$ says that if each of n holes may be occupied by only one of m pigeons then at least one pigeon must have no hole. The corresponding search problem is to find such a pigeon. More exactly, given an $n \times m$ $(0, 1)$-matrix M with $m > n$ and exactly one 1 in each row, the problem is to find an all-0 column. In this case we have a hypergraph $\mathcal{F}$ with m edges, corresponding to columns, and n blocks, corresponding to rows. Since $|\mathcal{F}| = m > n$, this hypergraph is a large (k, b, λ, d)-design with $k = n$, $b = m$, $\lambda = 0$ and $d = 1$. Since $|\mathcal{F}| = m > n = k(d+1)/2$, we can apply (3), which yields the lower bound $2^{\Omega\left(n^2/(ml)\right)}$. Recall that $2^{\Omega(n^2/m)}$ is the best known lower bound for the minimal length of a resolution refutation proof of PHP_n^m [13, 19, 8, 10]. So, the reason why PHP_n^m is hard for resolution, seems to lie not in the weakness of the resolution rule itself, but rather in the impossibility to keep enough information about possible outcomes, using small (up to l) sets of clauses.

Example 2. (Affine planes). Take an affine plane $\mathrm{AG}(2,q)$ of order q. Every point lies on $q+1$ lines, and there are $q(q+1)$ lines, each two of which intersect in at most one point. It is known (see [14, 4]) that every set of less than $2q-1$ points misses at least one line of $\mathrm{AG}(2,q)$. This result leads to the following *line search problem for* $\mathrm{AG}(2,q)$. We have $n = q(q+1)$ variables $x_1, \ldots, x_n$ corresponding to points, and n leaves, labeled by clauses $C_L = \bigvee_{i \in L} x_i$, corresponding to lines L. Given a set of at most $2(q-1)$ points, the problem is to find a line with no point in this set. By the result, mentioned above, this problem is well defined. Any semantic resolution proof for this problem solves the edge-search problem for the following design $\mathcal{F}$. Take any set $L_1, \ldots, L_q$ of q parallel (i.e. mutually disjoint) lines, and consider the hypergraph $\mathcal{F}$, the edges of which are all the remaining q^2 lines. Since every such line intersects each of the lines $L_1, \ldots, L_q$ in exactly one point, the hypergraph $\mathcal{F}$ is a (k, b, λ, d)-design with $k = b = d = q$ and $\lambda = 1$. To verify the largeness of this design, let B be a transversal of $L_1, \ldots, L_q$. If B would intersect all the lines in $\mathcal{F}$ then it would intersect *all* the lines of $\mathrm{AG}(2,q)$, which is impossible because $|B| = q < 2q-1$. Thus, every transversal B avoids at least one line of $\mathcal{F}$, and hence, $\mathcal{F}$ is large. Since $|\mathcal{F}| = q^2 > q(q+1)/2 = k(d+1)/2$, we we can apply (3), which yields the lower bound $2^{\Omega(q/l)} = \exp\left(\sqrt{|\mathcal{F}|}/l\right)$ on the size of any semantic resolution proof of fanin at most l, solving the edge-search problem for $\mathcal{F}$, and hence, for any such proof solving the line search problem for $\mathrm{AG}(2,q)$.

Example 3. (Projective planes). Take a projective plane $\mathrm{PG}(2,q)$ of order q. It has the same number $n = q^2+q+1$ of lines and points; each line has $q+1$ points and every point lies in $q+1$ lines; any two lines share exactly one point. It is known (see [5, 6]) that any set of at most $q+\sqrt{q}$ points must either contain a line or must avoid a line. This result leads to the following *line search problem for* $\mathrm{PG}(2,q)$. We have $n = q^2+q+1$ variables $x_1, \ldots, x_n$ corresponding to points, and and $2n$ leaves, labeled by clauses $C_L^+ = \bigvee_{i \in L} x_i$ and $C_L^- = \bigvee_{i \in L} \neg x_i$. Given a set of at most $q+\sqrt{q}$ points, the problem is to find a line which lies entirely either in this set or in its complement. This problem reduces to the line search problem in affine planes. The idea is to use the well-known fact that deletion of any one line L_0 from $\mathrm{PG}(2,q)$ (together with all its points) gives us affine plane $\mathrm{AG}(2,q)$; the lines of this new plane are sets $L \setminus L_0$ where $L \neq L_0$ are lines of the projective plane. Let now G be a fanin-l semantic resolution proof solving the line search problem for $\mathrm{PG}(2,q)$. Fix an arbitrary line L_0 of $\mathrm{PG}(2,q)$ and set to 0 all the variables x_i with $i \in L_0$. This restriction kills (evaluates to 1) all negative leaves of G and deletes (i.e. evaluates to 0) exactly one variable from each positive leaf. The restriction $L_0 \mapsto 0$ corresponds to deletion of L_0 from $\mathrm{PG}(2,q)$, and hence, leads to $\mathrm{AG}(2,q)$. Thus, we obtain a proof which solves the line search problem for $\mathrm{AG}(2,q)$. As shown in the previous example, this proof (and hence the original proof G) must have at least $2^{\Omega(q/l)}$ clauses.

3. Combinatorics

The proof of Theorem 1 consists of three simple steps.

1. Firstly, we replace each clause in the original proof G by a *positive* clause so that the resulting proof G^+ still solves the original edge-search problem for $\mathcal{F}$.
2. The goal of the second 'killing large clauses' step is to show that, if the proof G^+ would have less than $2^{M/b}$ clauses, with M defined by (2), then it would be possible to set some t variables to constants so that all long clauses in G^+ are killed (i.e. are evaluated to 1). Conditions (1) are necessary to ensure that we do not kill too many primary leafs, i.e. that the whole search problem becomes not much easier after this restriction. Our restrictions are *deterministic.* Thus the whole argument avoids randomness.
3. The goal of the final 'forcing large clauses' step is to show that *any* proof solving the desired search problem, must have at least one long clause. Here we essentially use the fact that edges of our design are almost disjoint, i.e. that $|E \cap F| \leq \lambda$ for any $E \neq F \in \mathcal{F}$.

This implies that G could not have less than $2^{M/b}$, as desired.

All the combinatorics we need is accumulated in two easy lemmas: the 'killing' lemma and the 'forcing' lemma.

Lemma 1. (Killing Lemma) *Let $\mathcal{A}$ be a hypergraph over a set X, and $S_1, \ldots, S_k$ be a partition of X into sets of cardinality at most b. If $|\mathcal{A}| < \left(\frac{k}{k-t}\right)^{r/b}$ and each edge of $\mathcal{A}$ has more than r points then there is a partial transversal T of $S_1, \ldots, S_k$ such that $|T| \leq t+1$ and T intersects all the edges of $\mathcal{A}$.*

Proof. Let $n = |X|$. We construct the set T via the following "greedy" procedure. Let $\mathcal{A}^1 = \mathcal{A}$ and $X^1 = X$. For each i, $1 \leq i \leq t$, include in T the element $x_i \in X^i$ which occurs in the largest number of sets of $\mathcal{A}^i$. Then remove from X^i all the points of that block, which contains x_i, to obtain X^{i+1}, and remove all the sets containing x_i from $\mathcal{A}^i$ to obtain $\mathcal{A}^{i+1}$. Sets deleted after $t+1$ steps intersect the set $\{x_1, \ldots, x_{t+1}\}$. Since $n \leq kb$, the number of remaining sets in $\mathcal{A}$ is bounded from above by $\alpha \cdot |\mathcal{A}|$ where

$$\alpha = \left(1 - \frac{r}{n}\right)\left(1 - \frac{r}{n-b}\right) \cdots \left(1 - \frac{r}{n-bt}\right) \leq \left(\frac{k}{k-t}\right)^{-r/b}.$$

Since $\mathcal{A}$ has less than α^{-1} sets, all the sets of $\mathcal{A}$ are already intersected by T, as desired.■

Let now $\mathcal{F}$ be a large (k, b, λ, d)-design with blocks $S_1, \ldots, S_k$. Fix an arbitrary partial transversal T of these blocks, and consider only those transversals which contain this particular transversal T. Let A be a set of points and $\mathcal{H} \subseteq \mathcal{F}$. We say that $\mathcal{H}$ is a *witness* of A if, for every transversal B containing T, we have that either $B \cap A \neq \emptyset$ or $B \cap E = \emptyset$ for at least one $E \in \mathcal{H}$ (or both). Put otherwise, every extension of T, intersecting all the edges of $\mathcal{H}$, must also intersect the set A. Given a set of points $A \subseteq X$, define its *weight*, $w_T(A)$ to be the minimum number of edges in a witness for A.

Lemma 2. (Forcing Lemma) *Let T be a partial transversal, $t = |T|$, and let $A \subseteq X$ be a set of points of weight $s = w_T(A)$. Then*

$$|A| \geq \frac{s(k-t-s+1)^2}{k + \lambda(s-1)}. \tag{4}$$

Proof. Lemma follows directly from the following two claims. Let $\mathcal{H} = \{E_1, \ldots, E_s\} \subseteq \mathcal{F}$ a minimal set of edges witnessing the weight of A.

Claim 1. The set A intersects every edge of $\mathcal{H}$ in at least $k-t-s+1$ points.

Proof. Take an arbitrary edge $E \in \mathcal{H}$. Since $\mathcal{H}$ is minimal, there must be a transversal $B \supseteq T$ which intersects all the edges of $\mathcal{H}' = \mathcal{H} \setminus \{E\}$ but avoids both sets A and E. For each edge $E' \in \mathcal{H}'$ choose any one point from the intersection $B \cap E'$, and let I be the set of these choosed $\leq |\mathcal{H}'| = s-1$ points. Let $\tilde{E}$ denote the set of all points in E, which belong to no of the blocks intersecting $I \cup T$. Since every edge E is a transversal, every block contains only one point of E, and hence, $|\tilde{E}| \geq |E| - |T| - |I| \geq k-t-s+1$. It remains therefore to prove that $A \supseteq \tilde{E}$ for every edge $E \in \mathcal{H}$.

To prove this, take an edge $E \in \mathcal{H}$ and an arbitrary point $x \in \tilde{E}$. Our goal is to show that x belongs to A. Let S be the (unique) block containing this point x. The fact that point x belongs to $\tilde{E}$ implies that this block S is disjoint from both T and I. Since B is a partial transversal and $B \cap \tilde{E} = \emptyset$, the block S intersects B in some other point $y \neq x$. Remove from B the point y and add the point x. The resulting set $(B \setminus \{y\}) \cup \{x\}$ intersects the edge E. Moreover, $B \setminus \{y\} \supseteq I \cup T$, because $y \in S$ and $S \cap (I \cup T) = \emptyset$.

Therefore, the set $(B \setminus \{y\}) \cup \{x\}$ contains T and intersects all the remaining edges in $\mathcal{H}'$ (since I intersects them). Since $\mathcal{H}$ is a witness for A, we have that $A \cap ((B \setminus \{y\}) \cup \{x\}) \neq \emptyset$. This together with $A \cap B = \emptyset$, implies that $x \in A$.$\blacksquare$

Claim 2. Let $\mathcal{A}$ be a hypergraph with s edges such that $u \leq |E| \leq v$ and $|E \cap F| \leq \lambda$ for all $E \neq F \in \mathcal{A}$. Let $X = \bigcup_{E \in \mathcal{A}} E$ be the underlying set of points. Then $|X| \geq (u^2 s)/(v + (s-1)\lambda)$.

Proof. The proof is a slight modification of a similar counting argument used by K. Corrádi [11] in the case when $u = v$. For a point $x \in X$, let $d(x)$ be the number of sets in $\mathcal{A}$ containing x. Then, for each edge E, $\sum_{x \in E} d(x) = \sum_{F \in \mathcal{A}} |E \cap F| \leq v + (s-1)\lambda$. Summing over all edges we get

$$\sum_{E \in \mathcal{A}} \sum_{x \in E} d(x) = \sum_{x \in X} d(x)^2 \geq \frac{1}{|X|} \left(\sum_{x \in X} d(x) \right)^2 = \frac{1}{|X|} \left(\sum_{E \in \mathcal{A}} |E| \right)^2 \geq \frac{(us)^2}{|X|}.$$

Using the previous estimate we obtain $(us)^2 \leq s \cdot |X| \left(v + (s-1)\lambda \right)$, which gives the desired lower bound on $|X|$.$\blacksquare$

Now we can finish the proof of Forcing Lemma as follows. By Claim 1 there exist s subsets $\tilde{E}_i \subseteq E_i$ such that $A \supseteq \tilde{E}_1 \cup \cdots \cup \tilde{E}_s$ and $u \leq |\tilde{E}_i| \leq v$ with $u = k - t - s + 1$ and $v = k$. Since sets E_i belong to the witness $\mathcal{H}$ (and hence, to the design $\mathcal{F}$), no two of them intersect in more than λ points, and Claim 2 yields the desired lower bound on $|\tilde{E}_1 \cup \cdots \cup \tilde{E}_s|$, and hence, the desired lower bound (4) on $|A|$.$\blacksquare$

4. Proof of Theorem 1

Since G solves the edge-search problem for the design $\mathcal{F}$, it is possible to associate with each transversal B, a truth assignment u_B to the remaining variables so that $f(v_B, u_B) = 1$, where f is the conjunction of all secondary leaves and v_B is the incidence vector of B. Fix this injection $B \mapsto u_B$, and call a truth assignment *legal* if it has a form (v_B, u_B) for some transversal B.

Our first goal is to replace the clauses in G by clauses without negated main variables. The idea of this transformation is similar to that used by Buss [7] in case of the pigeonhole principle. For a point i, let $S(i) = S \setminus \{i\}$ where S is the (unique) block containing this point i. Replace every clause C of G by the clause C^+ which is obtained from C by replacing each negated literal $\neg x_i$ by the set of positive literals $\{x_j : j \in S(i)\}$. Since for any transversal B we have that $i \in B \iff B \cap S(i) = \emptyset$, it follows that $C^+(v_B, u_B) = C(v_B, u_B)$, and hence, the resulting proof G^+ still solves the edge-search problem for $\mathcal{F}$. Moreover, G^+ has at most $\ell = \text{size}(G)$ clauses. Let r be the smallest number for which

$$\ell < \left(\frac{k}{k-t} \right)^{r/b}. \tag{5}$$

By Killing Lemma, there is a partial transversal T of size at most $t+1$ such that, setting to 1 all the variables x_i with $i \in T$, we kill off of all the clauses in G^+ with at least r main variables. Let G' be the resulting proof.

Since every primary leaf of G corresponds to an edge of $\mathcal{F}$ and each point belongs to no more than d edges of $\mathcal{F}$, at least $|\mathcal{F}|-dt$ of these leaves survive the restriction. We will use them to weight the clauses of G'. Namely, define the weight, $W(C)$ of a clause C to be the minimum number of primary leaves of G' whose conjunction implies C on all the legal truth assignments (v_B, u_B) with $B \supseteq T$. Note that primary leafs have weight 1, whereas secondary leaves have zero weight (they are satisfied by every legal assignment). On the other hand, the root must have weight larger than $\min\{|\mathcal{F}|-dt, k-t\}$, since we have at least $|\mathcal{F}|-dt$ primary leaves, for any $k-t$ of them we can find a transversal $B \supseteq T$ such that $B \setminus T$ intersects all of them (recall that $|B|=k$ and $|T|=t$). Since (by soundness) the weight of every clause is at most the sum of the weights of the (at most l) clauses from which it is derived, we can find a clause C such that $s \leq W(C) \leq ls$, as long as ls does not exceed the weight of the root, which is ensured by (1). This clause has the form $C = \bigvee_{i\in A} x_i \vee C'$ where C' is the auxiliary part of C. Since primary leaves have only x-variables, the weight $W(C)$ of C is exactly the weight $w_T(A)$ of the corresponding set of points A. By Forcing Lemma this set has at least

$$M = \frac{s(k-t-ls+1)^2}{k+\lambda(s-1)}$$

points. Since all clauses with at least r main variables, are already killed, we have that $M < r$. Since r was minimal for which (5) holds, this means that

$$\text{size}(G) = \ell \geq \left(\frac{k}{k-t}\right)^{M/b},$$

which is $\geq 2^{M/b}$ since $t \geq k/2$, as desired. This completes the proof of Theorem 1. ■

5. Concluding remarks

1. The input size of an edge-search problem for a hypergraph $\mathcal{F}$ is the number $|\mathcal{F}|$ of edges. It is interesting to compare the lower bounds which we obtain for searching problems, resulting from the generalized pigeonhole principle and from the line search problem in finite geometries. By Theorem 1, the general lower bound is exponential in $\Omega(k^2/b)$ if $\mathcal{F}$ is k-partite with block size b. Thus, in case of PHP^b_k, the bound is $\exp(k^2/|\mathcal{F}|)$, which is super-polynomial only if $|\mathcal{F}| = o(k^2/\log n)$, and it is still not known if it remains such for $|\mathcal{F}| \geq k^2$. (Recall that k is the number of holes and $|\mathcal{F}| = b$ is the number of pigeons). In this respect, the lower bound for $\mathrm{AG}(2,q)$ is better: here we have $|\mathcal{F}| = k^2$ (with $k=q$) and the bound is $\exp\left(\sqrt{|\mathcal{F}|}\right)$.

2. The reason, why our argument (as well as previous arguments, based on Haken's "bottlenecks counting" idea [13, 19, 8, 10]) does not work for PHP_k^b with $b \geq k^2$, is that we *a priori* restrict our search domain to transversals only. This makes possible the transformation $G \mapsto G^+$ but binds our hands when trying to kill long clauses, since now our killing set T must be (partial) transversal. Note that without this last restriction, we could replace the bound $\left(\frac{k}{k-t}\right)^{r/b}$ in Killing Lemma by $\left(\frac{n}{n-t}\right)^r$, which does not depend on the block size b at all (!). The overall conclusion is that, in order to get lower bounds for PHP_k^b with $b \geq k^2$, one should learn more on how to force large clauses which are not assumed be positive. Quite recently, Razborov, Wigderson and Yao [16] have made an interesting attempt to overcome this k^2 barrier. Using a novel technique they where able to prove exponential lower bounds (for arbitrarily large b!) on the size of some restricted versions of regular resolution proofs for PHP_k^b.

3. In this paper we have shown that the combinatorics of semantic Resolution is captured by two simple "killing" and "forcing" lemmas. Next logical step could be to understand the combinatorics of *cutting planes* proofs. All the known superpolynomial lower bounds for the length of such proofs follow from the corresponding lower bounds on the size of monotone Boolean circuits via appropriate interpolation theorems (see, e.g., [15] for a survey). Thus, these bounds capture the weakness of corresponding circuits rather than the weakness of cutting planes themselves. Moreover, this approach fails in the situations where the corresponding problems (like all three examples in Section 2) *have* small circuits. To get more insight into their nature of cutting planes proofs, it would be interesting to understand the cutting plane complexity of blocking principles for finite geometries. These geometries have more structure then the pigeonhole principle, and the corresponding principles have very natural formulation in terms of linear inequalities. The Jamison-Brower-Schrijev's theorem [14, 4] for AG$(2, q)$ is given by the system of $2n+1$ inequalities:

$$\sum_{i \in L_j} x_i \geq 1, \quad \sum_{i=1}^{n} x_i \leq 2(q-1), \quad 0 \leq x_i \leq 1 \qquad i, j = 1, \ldots, n$$

Bruen's theorem [5] for PG$(2, q)$ also can be stated as a system of $3n+1$ inequalities:

$$1 \leq \sum_{i \in L_j} x_i \leq q-1, \quad \sum_{i=1}^{n} x_i \leq q + \sqrt{q}, \quad 0 \leq x_i \leq 1 \qquad i, j = 1, \ldots, n.$$

What is the cutting plane complexity of these systems? The "quadratic counting" trick used in Bruen's proof makes plausible the conjecture that this system does *not* have a short cutting planes proof, unless we allow quadratic inequalities and/or multiplication of two inequalities. Both answers - a short cutting planes proof of Bruen's theorem or the absence of such proof - would be interesting.

Acknowledgment

I thank Alexander Razborov for turning my attention to resolution proofs and very interesting discussions.

References

[1] P. Beame and T. Pitassi, *Simplified and improved resolution lower bounds*, Proc. 37th IEEE Sympos. on Foundations of Computer Science, 1996.

[2] A. Blacke, *Canonical expressions in Boolean algebra*, PhD thesis, University of Chicago, 1937.

[3] A. Blokhuis, *On the size of a blocking set in $PG(2, p)$*, Combinatorica **14** (1994), 111–114.

[4] A.E. Brower and A. Schrijev, *The blocking number of an affine space*, J. Comb. Theory (A) **24** (1978), 251–253.

[5] A. A. Bruen, *Baer subplanes and blocking sets*, Bull. Amer. Math. Soc. **76** (1970), 342-344.

[6] A. A. Bruen, *Blocking sets in finite projective planes*, SIAM J. Appl. Math. **21** (1971), 380–392.

[7] S. Buss, *Polynomial size proofs of the propositional pigeonhole principle*, J. Symbolic Logic **52** (1987), 916–927.

[8] S. Buss and G. Turán, *Resolution proofs of generalized pigeonhole principles*, Theor. Comput. Sci. **62** (1988), 311–317.

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

[10] S. Cook and T. Pitassi, *A feasibly constructive lower bound for resolution proofs*, Inform. Process. Lett. **34** (1990), 81–85.

[11] K. Corrádi, *Problem at the Schweitzer competition*, Mat. Lapok **20** (1969), 159–162.

[12] M. Davis and H. Putnam, *A computing procedure for quantification theory*, Journal of the ACM **7**(3) (1960), 210–215.

[13] A. Haken, *The intractability of resolution*, Theor. Comput. Sci. **39** (1985), 297–308.

[14] R. Jamison, *Covering finite fields with cosets of subspaces*, J. Comb. Theory (A) **22** (1977), 253–266.

[15] A. Razborov, *Lower bounds for propositional proofs and independence results in bounded arithmetic*, Proc. 23rd Int. Colloq. Automata, Languages and Programming, ICALP'96 (Paderborn, Germany), 1996.

[16] A. Razborov, A. Wigderson and A. Yao, *Read-once branching programs, rectangular proofs of the pigeonhole principle and the transversal calculus*, manuscript, 1996.

[17] J. A. Robinson, *A machine-oriented logic based on the resolution principle*, Journal of the ACM **12**:1 (1965), 23–41.

[18] G. C. Tseitin, *On the complexity of derivations in propositional calculus*, Studies in mathematics and mathematical logic, Part II, ed. A. O. Slisenko, 1968, pp. 115–125.

[19] A. Urquhart, *Hard examples for resolution*, Journal of the ACM **34**:1 (1987), 209–219.