Resolution Lower Bounds for Perfect Matching Principles

Alexander A. Razborov[1]

Institute for Advanced Study, Princeton, US and Steklov Mathematical Institute, Moscow, Russia

Abstract

For an arbitrary hypergraph $\mathcal{H}$, let $PM(\mathcal{H})$ be the propositional formula asserting that $\mathcal{H}$ contains a perfect matching. We show that every resolution refutation of $PM(\mathcal{H})$ must have size

$$\exp\left(\Omega\left(\frac{\delta(\mathcal{H})}{\lambda(\mathcal{H})r(\mathcal{H})(\log n(\mathcal{H}))(r(\mathcal{H})+\log n(\mathcal{H}))}\right)\right),$$

where $n(\mathcal{H})$ is the number of vertices, $\delta(\mathcal{H})$ is the minimal degree of a vertex, $r(\mathcal{H})$ is the maximal size of an edge, and $\lambda(\mathcal{H})$ is the maximal number of edges incident to two different vertices.

For ordinary graphs G our general bound considerably simplifies to $\exp\left(\Omega\left(\frac{\delta(G)}{(\log n(G))^2}\right)\right)$ (implying an $\exp(\Omega(\delta(G)^{1/3}))$ lower bound that depends on the minimal degree only). As a direct corollary, every resolution proof of the functional onto version of the pigeonhole principle $onto - FPHP_n^m$ must have size $\exp\left(\Omega\left(\frac{n}{(\log m)^2}\right)\right)$ (which becomes $\exp\left(\Omega(n^{1/3})\right)$ when the number of pigeons m is unbounded). This in turn immediately implies an $\exp(\Omega(t/n^3))$ lower bound on the size of resolution proofs of the principle asserting that the circuit size of the Boolean function f_n in n variables is greater than t. In particular, *Resolution does not possess efficient proofs of* $\mathbf{NP} \not\subseteq \mathbf{P/poly}$.

These results relativize, in a natural way, to a more general principle $M(U|\mathcal{H})$ asserting that $\mathcal{H}$ contains a matching covering all vertices in $U \subseteq V(\mathcal{H})$.

Key words: Proof complexity, Resolution, Pigeonhole principle

Email address: `razborov@ias.edu` (Alexander A. Razborov).
URL: `http://www.mi.ras.ru/~razborov` (Alexander A. Razborov).

[1] Supported by The von Neumann Fund.

Preprint submitted to Elsevier Science *16 February 2004*

1 Introduction

Propositional proof complexity is an area of study that has seen a rapid development over the last decade. It plays as important a role in the theory of feasible proofs as the role played by the complexity of Boolean circuits in the theory of efficient computations. Propositional proof complexity is in a sense complementary to the (non-uniform) computational complexity; moreover, there exist extremely rich and productive relations between the two areas (see e.g. [1,2]).

Many combinatorial principles traditionally considered in the propositional proof complexity naturally appear as statements about graphs or hypergraphs asserting their most basic properties. The most prominent example is probably made by *Tseitin tautologies* [3,4] that are valid for any graph and assert in a way that the sum of degrees of all vertices is even (we will see several more examples below).

This naturally brings about the following general question:

which general combinatorial "hardness conditions" imposed on a (hyper)graph imply hardness of the associated principle with respect to one or another propositional proof system?

In this paper we confine ourselves to Resolution (which is one of the most widely studied proof systems), and for this system some previous work attempting to tackle the question above in this generality was done. Urquhart proved in [4] that Tseitin tautologies are hard for Resolution as long as the underlying graph has sufficiently good expansion properties. [5] introduced the *Hitting Set principle* $HS(\mathcal{H})$ asserting that the hypergraph $\mathcal{H}$ contains a small set of vertices hitting all its edges. He proved that this principle is hard for Resolution whenever $\mathcal{H}$ is a sufficiently good combinatorial design.

Urquhart [6] considered the *Matching principle* $M(G)$ asserting that the bipartite graph G on $U \times V$ has a (multi-valued) matching from U to V. Ben-Sasson and Wigderson [7] considered the same principle $M(G)$ under another name $G - PHP$. They proved that $G - PHP$ is hard for Resolution if G has sufficiently good expansion properties.

Alekhnovich, Ben-Sasson, Razborov and Wigderson [8] introduced the principle $\tau(\mathcal{H}, \vec{g})$ asserting that the Nisan-Wigderson generator based upon the hypergraph $\mathcal{H}$ (treated as a set system) and the Boolean functions $g_1, \dots, g_m$ misses a prescribed point in its image. They proved that if $\mathcal{H}$ has sufficiently good expansion properties and $g_1, \dots, g_m$ are robust with respect to restrictions then $\tau(\mathcal{H}, \vec{g})$ is hard for Resolution, as long as $\mathcal{H}$ does not have too many edges.

The framework from [8] in particular encompasses a natural generalization of Tseitin tautologies to hypergraphs. For the case of bounded vertex degree this generalization was also independently considered by Pudlák and Impagliazzo [9]. They formulated a combinatorial property of the underlying hypergraph implying that the resulting Tseitin tautology is very hard for tree-like Resolution, but this property is by far less natural than those mentioned above.

In this paper we look at the *Perfect Matching principle* $PM(\mathcal{H})$ asserting that the hypergraph $\mathcal{H}$ contains a perfect matching. Our reason to be interested in this principle is at least two-fold. The first motivation is similar to [5,6]: this class unifies in an extremely natural framework such popular combinatorial principles as $onto{-}FPHP_n^m$, $Count_r^n$ and the Mutilated Chessboard Problem.

The second reason is that, in the opposite direction, the Perfect Matching principle $PM(\mathcal{H})$ is a special case of the generator tautologies $\tau(\mathcal{H}, \vec{g})$ from [8] mentioned above. Namely, $PM(\mathcal{H})$ is isomorphic to $\tau(\mathcal{H}^*, \vec{E_1})$, where $\mathcal{H}^*$ is the dual hypergraph, and all g_is are the EXACT-1 functions outputting 1 iff the number of ones in the input string is exactly equal to 1. Thus, the principle $PM(\mathcal{H})$ might as well provide a convenient bridge between these two frameworks.

Our main result is an $\exp\left(\Omega\left(\frac{\delta(\mathcal{H})}{\lambda(\mathcal{H})r(\mathcal{H})(\log n(\mathcal{H}))(r(\mathcal{H})+\log n(\mathcal{H}))}\right)\right)$ lower bound on the size of any resolution refutation of $PM(\mathcal{H})$, where:

- $n(\mathcal{H})$ is the number of vertices;
- $\delta(\mathcal{H})$ is the minimal degree of a vertex;
- $r(\mathcal{H})$ is the maximal size of an edge;
- $\lambda(\mathcal{H})$ is the maximal number of edges incident to two different vertices.

Unlike previous work [4,5,8], our bound involves only the most basic combinatorial parameters of the hypergraph $\mathcal{H}$.

If $\mathcal{H} = G$ is an ordinary graph then $r(G) = 2$, $\lambda(G) = 1$ and this general bound gets simplified to $\exp\left(\Omega\left(\frac{\delta(G)}{(\log n(G))^2}\right)\right)$. Also, our result readily relativizes to the principle $M(U|\mathcal{H})$ asserting that the hypergraph $\mathcal{H}$ contains a matching covering at least all vertices in U; in the resulting bound $n(\mathcal{H})$ and $\delta(\mathcal{H})$ are re-calculated with respect to U.

Since the functional onto version of the pigeonhole principle $onto - FPHP_n^m$ is isomorphic to $PM(K_{m,n})$, we immediately get the bound $\exp\left(\Omega\left(\frac{n}{(\log m)^2}\right)\right)$ on its resolution size complexity (implying an $\exp\left(\Omega(n^{1/3})\right)$ bound when the number of pigeons m is unlimited). This generalizes the same lower bound for its functional version proved in [10] (see also [11–14] for the preceding work). It is worth noting that if we attempt to extract a stand-alone proof of this particular result from our general argument, it will look quite funny (half of the pigeons in its course will change sides and turn into holes and vice versa). This is one additional reason why we prefer to work in the more general framework of arbitrary (hyper)graphs.

As another immediate application of our general result we get an $\exp(\Omega(n/(r^2(\log n)(r+\log n))))$ bound on the resolution size complexity of the counting principle $Count_r^n$. Apparently the only lower bound for this principle that was known for $r > 2$ prior to our work comes from lower bounds for much stronger model of bounded-depth Frege proofs and has the form $\exp(\Omega(n^{\epsilon}))$ (for constant r), where ϵ is a rather small constant (see e.g. [15, Section 12]).

Finally, we show an $\exp(\Omega(t/n^3))$ lower bound on the size of resolution proofs of the principle $\neg Circuit_t(f_n)$ asserting that the circuit size of the Boolean function f_n in n variables is greater than t. In particular, *Resolution does not possess efficient proofs of* $\mathbf{NP} \not\subseteq \mathbf{P/poly}$. Previously this was known only under the existence of one-way functions (easily follows from the efficient interpolation theorem for Resolution), and when the circuits used for computing f_n may have unbounded fan-in [13].

Our proof method to a large extent follows the general pattern laid out in [14,10]. That is, we define an appropriate notion of the pseudo-width and use the "pigeon filter" lemma from [10] for reducing the pseudo-width of any small resolution proof at the expense of introducing certain new axioms (Lemma 17). Lower bounds on pseudo-width (Lemma 18) make the real novelty of this paper. For getting them we use a sort of indirect reduction to find in $\mathcal{H}$ a structure that looks like a restricted version of the *functional* pigeonhole principle. Then we show that the lower bound for the "pure" functional pigeonhole principle from [10] applies to this case with minimal changes.

The paper is organized as follows. In Section 2 we give necessary definitions and preliminaries and formulate our main results. In Section 3 we prove the lower bound for *ordinary* graphs (Theorem 6): its proof is somewhat simpler than the general bound for hypergraphs, while containing almost all essential ideas. The next section 4 shows hardness of $\mathbf{NP} \not\subseteq \mathbf{P/poly}$ for Resolution. In Section 5 we show how to extend the bound from Section 3 to the case of hypergraphs (Theorem 4). We conclude with several open problems in Section 6.

The paper is completely self-contained, although some familiarity with [14,10] may turn out to be helpful for better understanding the proofs.

2 Preliminaries

2.1 Definitions

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 *clause* is a disjunction of literals. The empty clause will be denoted by 0. A clause is *positive* if it contains only positive literals x^1. For two clauses C', C, let $C' \leq C$ mean that every literal appearing in C' also appears in C. A *CNF* is a conjunction of clauses.

One of the simplest and the most widely studied propositional proof systems is *Resolution* which operates with clauses and has one rule of inference called *resolution rule*:

$$\frac{C_0 \vee x \qquad C_1 \vee \bar{x}}{C} \ (C_0 \vee C_1 \leq C). \tag{1}$$

A *resolution refutation* of a CNF τ is a resolution proof of the empty clause 0 from the clauses appearing in τ. The *size* $S_R(P)$ of a resolution proof P is the overall number of clauses in it. For a CNF τ, $S_R(\tau)$ is the minimal size of its resolution refutation, and ∞ if no such refutation exists (i.e., τ is satisfiable).

For n, a non-negative integer let $[n] \stackrel{\text{def}}{=} \{1, 2, \ldots, n\}$, and for $\ell \leq n$ let $[n]^\ell \stackrel{\text{def}}{=} \{I \subseteq [n] \mid |I| = \ell\}$.

A *hypergraph* $\mathcal{H}$ is a pair $\mathcal{H} = (V, \mathcal{E})$, where V is a finite set of *vertices*, and $\mathcal{E} \subseteq \mathcal{P}(V)$ is the set of *edges* (thus, in hypergraphs we do allow empty edges and loops but disallow multiple edges). The hypergraph is a *graph* if all its edges have cardinality 2 (thus, in graphs we disallow both multiple edges and loops). In the case of graphs we scale the notation one level down and denote by E the set of all edges, whereas individual edges are denoted by small letters e. A *matching* in a hypergraph $\mathcal{H}$ is any collection of pairwise disjoint edges. The matching is *perfect* if every vertex is covered by (exactly) one edge from the matching.

Definition 1 *For a hypergraph $\mathcal{H} = (V, \mathcal{E})$, the* Perfect Matching principle *$PM(\mathcal{H})$ is the CNF in the variables $\{x_E \mid E \in \mathcal{E}\}$ that is the conjunction of the following clauses:*

$$Q_v \stackrel{\text{def}}{=} \bigvee_{E \ni v} x_E \ (v \in V);$$

$$Q_{E_1,E_2} \stackrel{\text{def}}{=} \bar{x}_{E_1} \vee \bar{x}_{E_2} \ (E_1 \neq E_2 \in \mathcal{E};\ E_1 \cap E_2 \neq \emptyset).$$

Clearly, $PM(\mathcal{H})$ is satisfiable if and only if $\mathcal{H}$ contains a perfect matching.

Example 2 The (negation of the) *functional onto pigeonhole principle* is the unsatisfiable CNF in the variables $\{ x_{ij} \mid i \in [m],\ j \in [n] \}$ denoted by $\neg onto - FPHP^m_n$ that is the conjunction of the following clauses:

$$\begin{aligned}
&Q_i \overset{\text{def}}{=} \bigvee_{j=1}^{n} x_{ij}\ (i \in [m]);\\
&Q_{i_1,i_2;j} \overset{\text{def}}{=} (\bar{x}_{i_1j} \vee \bar{x}_{i_2j})\ (i_1 \neq i_2 \in [m],\ j \in [n]);\\
&Q_{i;j_1,j_2} \overset{\text{def}}{=} (\bar{x}_{ij_1} \vee \bar{x}_{ij_2})\ (i \in [m],\ j_1 \neq j_2 \in [n]);\\
&Q_j \overset{\text{def}}{=} \bigvee_{i=1}^{m} x_{ij}\ (j \in [n])
\end{aligned}$$

(the *basic pigeonhole principle* PHP^m_n consists of the first two groups of axioms, and the *functional pigeonhole principle* $FPHP^m_n$ – of the first three groups). Clearly, $\neg onto - FPHP^m_n$ is identical to $PM(K_{m,n})$, where $K_{m,n}$ is the complete bipartite graph. More generally, [6,7] proposed to consider the principle $G - PHP$ (G a bipartite graph on $[m] \times [n]$) which is a naturally defined restriction of PHP^m_n onto G. Denoting its obvious analogue for the functional onto version by $onto - G - FPHP$, we see that $\neg onto - G - FPHP$ is identical to $PM(G)$.

Example 3 If $\mathcal{H} = ([n], [n]^r)$ is the complete r-hypergraph on n vertices and $r \nmid n$ then $PM(\mathcal{H})$ coincides with the *counting principle* $Count^n_r$.

Given a hypergraph $\mathcal{H} = (V, \mathcal{E})$, let $n(\mathcal{H}) \overset{\text{def}}{=} |V|$ is the number of its vertices. The *star* of a vertex v is $S_{\mathcal{H}}(v) \overset{\text{def}}{=} \{E \in \mathcal{E} \mid v \in E\}$. The *degree* of a vertex v is $\deg_{\mathcal{H}}(v) \overset{\text{def}}{=} |S_{\mathcal{H}}(v)|$. The *minimal degree* of $\mathcal{H}$ is defined as $\delta(\mathcal{H}) \overset{\text{def}}{=} \min_{v \in V} \deg_{\mathcal{H}}(v)$.

r-uniform hypergraphs are characterized as those in which all edges have cardinality r. From this concept we need only the upper bound on the size of an edge so we let $r(\mathcal{H}) \overset{\text{def}}{=} \max_{E \in \mathcal{E}} |E|$.

Pairwise balanced designs with index λ are characterized as those $(V, \mathcal{E})$ for which $|S_{\mathcal{H}}(v) \cap S_{\mathcal{H}}(v')| = \lambda$ for any two different vertices v, v'. From this definition we will also need only the upper bound, so we let $\lambda(\mathcal{H}) \overset{\text{def}}{=} \max_{v \neq v' \in V} |S_{\mathcal{H}}(v) \cap S_{\mathcal{H}}(v')|$.

2.2 Results

The main result of this paper is the following

Theorem 4 $S_R(PM(\mathcal{H})) \geq \exp\left(\Omega\left(\frac{\delta(\mathcal{H})}{\lambda(\mathcal{H})r(\mathcal{H})(\log n(\mathcal{H}))(r(\mathcal{H})+\log n(\mathcal{H}))}\right)\right)$.

This theorem will be fully proved only in Section 5.

If $\mathcal{H} = ([n],[n]^r)$, then $n(\mathcal{H}) = n$, $\delta(\mathcal{H}) = \binom{n-1}{r-1}$, $r(\mathcal{H}) = r$ and $\lambda(\mathcal{H}) = \binom{n-2}{r-2}$, and we immediately get

Corollary 5 $S_R(Count_r^n) \geq \exp(\Omega(n/(r^2(\log n)(r+\log n))))$.

Note that in this corollary r need not be a constant and may arbitrarily depend on n.

For an ordinary graph G, $r(G) = 2$ and $\lambda(G) = 1$. Thus, the following result is a special case of Theorem 4:

Theorem 6 *For an arbitrary graph G,*

$$S_R(PM(G)) \geq \exp\left(\Omega\left(\frac{\delta(G)}{(\log n(G))^2}\right)\right).$$

However, in the next section 3 we will give its independent proof which is a little bit simpler than the proof of Theorem 4.

Applying Theorem 6 to the bipartite graph $K_{m,n}$ with $m > n$, we get

Corollary 7 *For $m > n$, $S_R(\neg onto - FPHP_n^m) \geq \exp\left(\Omega\left(\frac{n}{(\log m)^2}\right)\right)$.*

Corollary 8 *For every $m > n$, $S_R(\neg onto - FPHP_n^m) \geq \exp\left(\Omega(n^{1/3})\right)$.*

Proof of Corollary 8 from Corollary 7. Let $S_R(\neg onto - FPHP_n^m) = S$, and let P be a size S refutation of $\neg onto - FPHP_n^m$. P can use at most S axioms from $\{Q_1, \ldots, Q_m\}$, and it must use at least $(n+1)$ such axioms (otherwise, all axioms occurring in P could have been simultaneously satisfied). Apply to P the restriction that sends to 0 all those x_{ij} for which $Q_i \notin P$. This will show $S_R(\neg onto - FPHP_n^{m'}) \leq S$ for some m' with $n < m' \leq S$. Now the required bound $S \geq \exp\left(\Omega(n^{1/3})\right)$ immediately follows from Corollary 7.$\blacksquare$

Remark 9 If we try to generalize Corollary 8 to arbitrary graphs G, then we immediately face the difficulty that after restricting the graph G, its minimal degree $\delta(G)$ may in general drop. One natural way of circumventing this is to relativize the whole argument to an arbitrary set of "active" vertices U. Namely, for $U \subseteq V(\mathcal{H})$ let $M(U|\mathcal{H})$ be defined in the same way as $PM(\mathcal{H})$, with the exception that the axioms Q_v are allowed only for $v \in U$. Respectively, let $\delta(U|\mathcal{H}) \stackrel{\mathrm{def}}{=} \min_{v\in U} \deg_{\mathcal{H}}(v)$. Then we can generalize our Theorem 4 to

$$S_R(M(U|\mathcal{H})) \geq \exp\left(\Omega\left(\frac{\delta(U|\mathcal{H})}{\lambda(\mathcal{H})r(\mathcal{H})(\log|U|)(r(\mathcal{H})+\log|U|)}\right)\right)$$

and then

$$S_R(M(U|G)) \geq \exp\left(\Omega\left(\frac{\delta(U|G)}{|U|^2}\right)\right) \geq \exp\left(\Omega\left(\frac{\delta(U)}{|U|^2}\right)\right).$$

Applying now the same reasoning as in the proof of Corollary 8, we get

Corollary 10 *For an arbitrary graph G,*

$$S_R(PM(G)) \geq \exp(\Omega(\delta(G)^{1/3})).$$

As another application, for the principle $G - FPHP$ we get the following:

Theorem 11 *For every bipartite graph G on $[m] \times [n]$, $S_R(\neg G - FPHP) \geq \exp\left(\Omega\left(\frac{\min_{i\in[m]}\deg_G(i)}{(\log m)^2}\right)\right)$.*

It is much easier, however, to prove this theorem by using the machinery from [10] in more direct way. Since we are not aware of any other interesting applications of the principle $M(U|\mathcal{H})$ where potentially $\delta(U|\mathcal{H}) \gg \delta(\mathcal{H})$ and/or $|U| \ll V(\mathcal{H})$ (and, likewise, are not aware of interesting graphs for which Corollary 10 can not be replaced by Theorem 6), we will concentrate only on the absolute version $PM(\mathcal{H})$, confining ourselves to a few remarks in appropriate places as to this possibility of relativization.

2.3 Positive calculus

Like in virtually all previous work on the subject ([16,11,5,6,14,10]), it will be convenient to get rid of negations once and for all by using the following normal form for refutations of $PM(\mathcal{H})$.

Fix a hypergraph $\mathcal{H} = (V, \mathcal{E})$. For $\mathcal{E}_0 \subseteq \mathcal{E}$, let $X_{\mathcal{E}_0} \stackrel{\mathrm{def}}{=} \bigvee_{E\in\mathcal{E}_0} x_E$; these are exactly all positive clauses in the variables $\{x_E \mid E \in \mathcal{E}\}$. For $\mathcal{E}_0, \mathcal{E}_1 \subseteq \mathcal{E}$, let $\mathcal{E}_0 \perp \mathcal{E}_1 \equiv (\mathcal{E}_0 \cap \mathcal{E}_1 = \emptyset \,\wedge\, (\forall E_0 \in \mathcal{E}_0)(\forall E_1 \in \mathcal{E}_1)(E_0 \cap E_1 \neq \emptyset))$ (intuitively, $\mathcal{E}_0$ and $\mathcal{E}_1$ are inconsistent).

Definition 12 *The* positive calculus *operates with positive clauses in the variables $\{x_E \mid E \in \mathcal{E}\}$ and has one inference rule which is the following* positive rule*:*

$$\frac{C_0 \vee X_{\mathcal{E}_0} \qquad C_1 \vee X_{\mathcal{E}_1}}{C}\ (C_0 \vee C_1 \leq C;\ \mathcal{E}_0 \perp \mathcal{E}_1). \tag{2}$$

A *positive calculus refutation* of a set of positive clauses $\mathcal{A}$ is a positive calculus proof of 0 from $\mathcal{A}$, and the *size* $S(P)$ of a positive calculus proof is the overall number of clauses in it. Finally, let $S_P(PM(\mathcal{H}))$ be the minimal possible size of a positive calculus refutation of the set of axioms $\{Q_v \mid v \in V\}$.

Lemma 13 $S_P(PM(\mathcal{H})) \leq S_R(PM(\mathcal{H})) \leq O(S_P(PM(\mathcal{H})) \cdot |\mathcal{E}|^2)$.

Proof. Suppose that we have a resolution refutation of $PM(\mathcal{H})$. Apply to every line in it the transformation θ that replaces every negated literal $\bar{x}_E$ by the positive clause $X_{\{E' \mid E' \neq E,\ E' \cap E \neq \emptyset\}}$. Clearly, $\theta(Q_v) = Q_v$ and $\theta(Q_{E_1,E_2})$ contains Q_v for an arbitrary $v \in E_1 \cap E_2$. It is also easy to see that θ takes an instance of the resolution rule (1) to an instance of the positive rule; therefore, θ maps P to a positive calculus refutation of the same size.

In the opposite direction, it is straightforward to check that in the presence of the axioms Q_{E_1,E_2} the positive rule is simulated by an $O(|\mathcal{E}|^2)$-sized resolution proof.■

Remark 14 The relativized version of Lemma 13 is also true: if we define $S_P(M(U|\mathcal{H}))$ as the minimal possible size of a positive calculus refutation of the set of axioms $\{Q_v \mid v \in U\}$, then we still have $S_P(M(U|\mathcal{H})) \leq S_R(M(U|\mathcal{H})) \leq O(S_P(M(U|\mathcal{H})) \cdot |\mathcal{E}|^2)$. We, however, should work a little bit harder for establishing the first inequality in this case (cf. [16]). Namely, instead of applying the mapping θ to the axioms Q_{E_1,E_2}, we look at the first time they get resolved with another clause:

$$\frac{C \vee x_{E_1} \qquad Q_{E_1,E_2}}{C \vee \bar{x}_{E_2}},$$

and observe that $\theta(C \vee x_{E_1}) \leq \theta(C \vee \bar{x}_{E_2})$. Thus, these axioms can be eliminated from $\theta(P)$ directly.

2.4 Filter lemma

We will need the following general combinatorial statement proved in [10].

Proposition 15 ([10]) *Suppose that we are given S integer vectors $r^1, r^2, \ldots, r^S$ of length m each: $r^\nu = (r^\nu_1, \ldots, r^\nu_m)$, and let w_0 be an arbitrary integer parameter. Then there exists an integer vector $(r_1, \ldots, r_m)$ such that $r_i < \lfloor \log_2 m \rfloor$*

for all $i \in [m]$ and for every $\nu \in [S]$ at least one of the following two events happens:

(1) $|\{i \in [m] \mid r_i^\nu \le r_i\}| \ge w_0$*;*
(2) $|\{i \in [m] \mid r_i^\nu \le r_i + 1\}| \le O(w_0 + \log S)$*.*

For the sake of completeness, we include its complete proof.

Proof of Proposition 15. We use an easy probabilistic argument. For $r = (r_1, \ldots, r_m)$, let $W(r) \overset{\text{def}}{=} \sum_{i=1}^m 2^{-r_i}$, and let $C > 0$ be a sufficiently large constant. It suffices to prove the existence of a vector r such that for every $\nu \in [S]$ we have:

$$W(r^\nu) \ge C(w_0 + \log_2 S) \Longrightarrow |\{i \in [m] \mid r_i \ge r_i^\nu\}| \ge w_0; \tag{3}$$

$$W(r^\nu) \le C(w_0 + \log_2 S) \qquad \Longrightarrow |\{i \in [m] \mid r_i \ge r_i^\nu - 1\}| \le O(w_0 + \log S). \tag{4}$$

Let $t \overset{\text{def}}{=} \lfloor \log_2 m \rfloor - 1$ and R be the distribution on $[t]$ given by $p_r \overset{\text{def}}{=} 2^{-r}$ $(1 \le r \le t-1)$, $p_t \overset{\text{def}}{=} 2^{1-t}$. Pick independent random variables $\boldsymbol{r_1}, \ldots, \boldsymbol{r_m}$ according to this distribution. Let us check that for any individual $\nu \in [S]$ the related condition (3), (4) is satisfied with high probability.

Case 1. $W(r^\nu) \ge C(w_0 + \log_2 S)$.
Note that $\sum\limits_{r_i^\nu > t} 2^{-r_i^\nu} \le m \cdot 2^{-t-1} \le 2$, therefore $\sum\limits_{r_i^\nu \le t} 2^{-r_i^\nu} \ge C(w_0 + \log_2 S) - 2$. On the other hand, for every i with $r_i^\nu \le t$ we have $\mathbf{P}[\boldsymbol{r_i} \ge r_i^\nu] \ge 2^{-r_i^\nu}$, hence $\mathbf{E}[|\{i \in [m] \mid r_i^\nu \le t \wedge \boldsymbol{r_i} \ge r_i^\nu\}|] \ge C(w_0 + \log_2 S) - 2$. Since the events $\boldsymbol{r}_i \ge r_i^\nu$ are independent, we may apply Chernoff's bound and conclude that $\mathbf{P}[|\{i \in [m] \mid r_i^\nu \le t \wedge \boldsymbol{r}_i \ge r_i^\nu\}| < w_0] \le S^{-2}$ if the constant C is large enough.

Case 2. $W(r^\nu) \le C(w_0 + \log_2 S)$.
In this case $\mathbf{P}[\boldsymbol{r_i} \ge r_i^\nu - 1] \le 2^{2-r_i^\nu}$ and, therefore,

$$\mathbf{E}[|\{i \in [m] \mid \boldsymbol{r_i} \ge r_i^\nu - 1\}|] \le 4W(r^\nu) \le 4C(w_0 + \log_2 S).$$

Applying once more Chernoff's bound, we conclude that

$$\mathbf{P}[|\{i \in [m] \mid \boldsymbol{r_i} \ge r_i^\nu - 1\}| \ge C'(w_0 + \log S)] \le S^{-2}$$

for any sufficiently large constant $C' \gg C$.

So, for every individual $\nu \in [S]$ the probability that the related property (3), (4) fails is at most S^{-2}. Therefore, for at least one choice of $\boldsymbol{r_1}, \ldots, \boldsymbol{r_m}$ they will be satisfied for all $\nu \in [S]$. This completes the proof of Proposition 15.■

3 Proof of the main result for ordinary graphs

In this section we prove Theorem 6. Fix a graph $G = (V, E)$. Given Lemma 13, we may assume that we have a positive calculus refutation P of $\{ Q_v \mid v \in V \}$, and we should lower bound its size $S(P)$. Let $N_G(v)$ be the set of all vertices adjacent to v in G. For a positive clause C in the variables $\{ x_e \mid e \in E \}$, let

$$N_C(v) \stackrel{\rm def}{=} \left\{ w \in N_G(v) \,\middle|\, x_{(v,w)} \in C \right\}$$

and

$$\deg_C(v) \stackrel{\rm def}{=} |N_C(v)|.$$

For analyzing the refutation P we are going to allow certain positive clauses as new axioms. Our allowance criterium will be determined by a fixed integer vector $d = (d_v \mid v \in V)$ ("filter"), and a positive clause C will be allowed as a new axiom if and only if sufficiently many vertices v satisfy $\deg_C(v) \geq d_v$ ("get stuck" at the filter d). In this way we will be able to simplify the refutation P by "filtering out" of it all clauses C with this property and declaring them as new axioms.

Our first task (Section 3.1) will be to show that if the thresholds d_v are chosen properly, then *in every clause C passing the filter d, almost all vertices pass it safely*, i.e. $\deg_C(v)$ is *well* below d_v. This part almost immediately follows from Proposition 15 and is practically identical to [10, Lemma 3.3].

The *pseudo-width* of a clause C will be defined as the number of vertices that *narrowly* pass the filter d. Our second task (Section 3.2) will be to get lower bounds on the pseudo-width of any small positive calculus refutation in the presence of the new axioms described above. It will be performed in two steps. During the first step we use a simple probabilistic argument to identify within G a structure that "looks like" $G' - FPHP$ (where G' is a bipartite subgraph of G) and behaves well with respect to any positive clause in the prospective refutation (Claim 19). Then we complete the proof by sorting out the edges of G according to this structure and evaluating the result in a linear matroid; this part being a relatively easy adaption of the argument in [10, Lemma 3.4] for the "pure" $FPHP^m_n$.

Now we begin the formal proof.

3.1 Pseudo-width and its reduction

Suppose that we are given an integer vector $d = (d_v \mid v \in V)$ indexed by vertices of the graph G. For a positive clause C let

$$V_d(C) \stackrel{\text{def}}{=} \{v \in V \mid \deg_C(v) \geq d_v\}.$$

Fix for the rest of Section 3 the parameters δ_v as follows[2] :

$$\delta_v \stackrel{\text{def}}{=} \frac{\deg_G(v)}{2\log|V|}, \tag{5}$$

and let

$$V_d^\sharp(C) \stackrel{\text{def}}{=} \{v \in V \mid \deg_C(v) \geq d_v - \delta_v\}.$$

Define the *pseudo-width of the clause* C as

$$w_d(C) \stackrel{\text{def}}{=} |V_d^\sharp(C)|.$$

The *pseudo-width of a positive calculus proof* P is naturally defined as

$$w_d(P) \stackrel{\text{def}}{=} \max\{w_d(C) \mid C \in P\}.$$

A (w_0, d)*-axiom* is a positive clause C such that $|V_d(C)| \geq w_0$.

Remark 16 For the relativized version $M(U|G)$ (that is, when we only have the axioms $\{Q_v \mid v \in U\}$ for some $U \subset V$), the vectors d_v, δ_v are defined only for $v \in U$. (5) will have $\log|U|$ in the denominator, $V_d(C), V_d^\sharp(C)$ will be subsets of U etc.

Lemma 17 *Suppose that there exists a positive calculus refutation P of $\{Q_v \mid v \in V\}$, and let $w_0 \leq \frac{\delta(G)}{4}$ be an arbitrary integer parameter. Then there exists an integer vector $d = (d_v \mid v \in V)$ with $\delta_v < d_v \leq \deg_G(v)$ for all $v \in V$, a set of (w_0, d)-axioms $\mathcal{A}$ and a positive calculus refutation P' of $\{Q_v \mid v \in V\} \cup \mathcal{A}$ such that $S(P') \leq S(P)$ and*

$$w_d(P') \leq O(w_0 + \log S(P)). \tag{6}$$

Proof. As we already mentioned above, this lemma is very similar to [10, Lemma 3.3], and for this reason we will be rather concise here. Fix a positive calculus refutation P of $\{Q_v \mid v \in V\}$, and let $S \stackrel{\text{def}}{=} S(P)$. For $C \in P$ define

$$r_v(C) \stackrel{\text{def}}{=} \lfloor \frac{\deg_G(v) - \deg_C(v)}{\delta_v} \rfloor + 1.$$

[2] All logarithms in this paper are base 2

Let $m \stackrel{\rm def}{=} |V|$ and $r(C) \stackrel{\rm def}{=} (r_v(C) \mid v \in V)$ be the integer vector of length m. We apply Proposition 15 to the vectors $\{r(C) \mid C \in P\}$, and let $(r_v \mid v \in V)$ satisfy the conclusion of that proposition.

Set $d_v \stackrel{\rm def}{=} \lfloor \deg_G(v) - \delta_v r_v \rfloor + 1$ (so that d_v is the minimal integer with the property $\lfloor \frac{\deg_G(v)-d_v}{\delta_v} \rfloor + 1 \leq r_v$). Note that since $r_v < \lfloor \log_2 m \rfloor$, $w_0 \leq \frac{\delta(G)}{4}$ and also $\delta_v \leq \frac{\delta(G)}{2 \log |V|}$ (by (5)), we have $d_v > \frac{\delta(G)}{2} \geq \delta_v + w_0$.

Consider now an arbitrary $C \in P$. If for the vector $r(C)$ the first case in Proposition 15 takes place, then $\lfloor \frac{\deg_G(v)-\deg_C(v)}{\delta_v} \rfloor + 1 \leq r_v$ for at least w_0 different vertices $v \in V$; therefore every such C is an (w_0, d)-axiom. Choose arbitrarily w_0 vertices in $V_d(C)$, and remove from C all those x_e for which e is not incident to at least one of the chosen vertices. The resulting clause C' will still be an (w_0, d)-axiom and $\deg_{C'}(v) \leq w_0$ for every vertex v that has not been chosen. Hence, due to the inequality $d_v > \delta_v + w_0$, no such vertex may belong to $V_d^\sharp(C')$ which implies $w_d(C') = w_0$. Replace C by C', and put the latter into $\mathcal{A}$.

In the second case, $\left| \left\{ v \in V \,\middle|\, \lfloor \frac{\deg_G(v)-\deg_C(v)}{\delta_v} \rfloor \leq r_v \right\} \right| \leq O(w_0 + \log S)$. Since $v \in V_d^\sharp(C)$ implies the inequality $\lfloor \frac{\deg_G(v)-\deg_C(v)}{\delta_v} \rfloor \leq r_v$, for all such C we have $w_d(C) \leq O(w_0 + \log S)$.

This completes the proof of Lemma 17.■

3.2 Lower bounds on pseudo-width

Given Lemma 17, we must now show that for every choice of the vector d, there is no small size small pseudo-width positive calculus refutation of $\{Q_v \mid v \in V\} \cup \mathcal{A}$, where $\mathcal{A}$ is any set of (w_0, d)-axioms. Before we begin the formal proof, let us try to convey some intuition toward it.

As we already mentioned above, our overall strategy will be to find inside G a "well-behaving" (with respect to the refutation) structure which sufficiently resembles $G' - FPHP$ for some bipartite subgraph G'. For this purpose we randomly divide the vertices V into *pigeon vertices* V_P and *hole vertices* V_H. If our prospective refutation P is small enough, then we may expect that this partition will look random to every clause $C \in P$.

The partition (V_P, V_H) induces a classification of all edges into pigeon-pigeon edges, pigeon-hole edges and hole-hole edges. Pigeon-pigeon edges are of no importance and are removed immediately.

Pigeon-hole edges are the most crucial, they form the subgraph G' and they are used to simulate $G' - FPHP$. The fact that our partition is random enough with respect to every $C \in P$ implies that there are sufficiently many pigeon-hole edges, and that when everything is restricted to them, degrees are scaled down by almost exactly a factor of two, the sets $V_d(C)$ and $V_d^\sharp(C)$ also behave in an expected manner etc. This ensures us that we can easily adopt the algebraic argument for the functional pigeonhole principle [10, Lemma 3.4].

One remaining problem is that in its original form this argument seems to be inherently incapable of taking care of the axioms $\{Q_v\}$ with v a hole vertex (missing in the functional version of PHP^m_n), and this is exactly what the hole-hole edges are used for. More specifically, our algebraic invariant is preserved under adding or deleting such edges, and when we need to "force" an axiom Q_v with $v \in V_H$ (see the proof of Claim 21), we do it simply by appending any legitimate hole-hole edge (v, w) to the current matching b.

Let us now proceed to the rigorous proof. Recall that δ_v are given by (5). At this point let us also define

$$S_0 \stackrel{\text{def}}{=} \exp\left(\frac{\epsilon^2\delta(G)}{(\log|V|)^2}\right) \tag{7}$$

and

$$w_0 \stackrel{\text{def}}{=} \exp\left(\frac{\epsilon\delta(G)}{(\log|V|)}\right), \tag{8}$$

where $\epsilon < 0$ is a sufficiently small constant. For technical reasons we also need to assume

$$|V| \leq S_0 \tag{9}$$

(or, in other words, $\delta(G) \geq \frac{1}{\epsilon^2} \cdot (\log|V|)^3$); at the end of Section 3 we will show how to get rid of this restriction.

Our lower bound on the pseudo-width looks like this:

Lemma 18 *Let $d = (d_v \mid v \in V)$ be an integer vector such that $\delta_v < d_v \leq \deg_G(v)$ for all $v \in V$, and P be a positive calculus refutation of $\{Q_v \mid v \in V\} \cup \mathcal{A}$, where $\mathcal{A}$ is an arbitrary set of (w_0, d)-axioms, such that $S(P) \leq S_0$. Then $w_d(P) \geq \frac{\delta(G)}{200\log|V|}$.*

Proof. Fix $d = (d_v \mid v \in V)$, $\mathcal{A}$ and P satisfying these assumptions. We need the following easy claim (the analogue of this claim for hypergraphs, however, will be by far less transparent).

Claim 19 *There exists a vertex partition $V = V_P \dot\cup V_H$ such that the following two properties are satisfied:*

(1) for every $A \in \mathcal{A}$, $|V_d(A) \cap V_P| \geq w_0/3$;
(2) for every $C \in P \cup \{X_E\}$ and every $v \in V$,

$$\left| |N_C(v) \cap V_H| - \frac{1}{2}\deg_C(v) \right| \leq \frac{\delta_v}{10}$$

(recall that $X_E = \bigvee_{e \in E} x_e$).

Proof. Uniformly pick a random partition $V = \boldsymbol{V_P} \dot\cup \boldsymbol{V_H}$. For estimating the probabilities that it satisfies the required properties, it will be convenient to use the following special case of Bernstein's inequality (see e.g. [17, page 205]) that, in a convenient way, generalizes both Chernoff's and variance bounds for the sum of independent Poisson trials.

Proposition 20 *Let $\boldsymbol{S}$ be the sum of independent 0-1 variables (not necessarily equidistributed), and let E be its expectation. Then $\mathbf{P}[|\boldsymbol{S} - E| \geq \delta] \leq \exp\left(-\Omega\left(\frac{\delta^2}{\delta+E}\right)\right)$.*

In particular, for the property 1 we have that for every individual $A \in \mathcal{A}$, $\mathbf{P}[|V_d(A) \cap \boldsymbol{V_P}| \leq w_0/3] \leq \exp(-\Omega(w_0)) \leq S_0^{-2}$. For property 2, given any individual positive clause C and $v \in V$,

$$\begin{aligned}\mathbf{P}&\left[\left||N_C(v) \cap \boldsymbol{V_H}| - \frac{1}{2}\deg_C(v)\right| \geq \frac{\delta_v}{10}\right] \\ &\leq \exp\left(-\Omega\left(\frac{\delta_v^2}{\delta_v + \deg_C(v)}\right)\right) \leq \exp\left(-\Omega\left(\frac{\delta_v^2}{\deg_G(v)}\right)\right) \\ &\leq \exp\left(-\Omega\left(\frac{\deg_G(v)}{(\log|V|)^2}\right)\right) \leq \exp\left(-\Omega\left(\frac{\delta(G)}{(\log|V|)^2}\right)\right) \leq S_0^{-3}\end{aligned}$$

provided the constant ϵ in (7) is small enough. Given our assumption (9), Claim 19 now follows by the union bound.■

We return to the proof of Lemma 18. Fix an arbitrary partition $V = V_P \dot\cup V_H$ satisfying properties 1, 2 of Claim 19. Let D be the set of all (partial) matchings in G. We will sometimes identify matchings $a \in D$ with their characteristic functions, i.e., with Boolean assignments to the variables $\{x_e \mid e \in E\}$. Let $\mathrm{dom}(a)$ be the set of all vertices in V incident to an edge in a.

For a positive clause C, let

$$Z(C) \stackrel{\text{def}}{=} \left\{ a \in D \,\middle|\, \mathrm{dom}(a) \supseteq V_d^\sharp(C) \wedge C(a) = 0 \right\}.$$

Intuitively, $Z(C)$ is the set of all matchings "forcing" C to 0. We are going to keep track of a certain algebraic invariant defined in terms of $Z(C)$ as the refutation P is making progress, and for that purpose we construct a mapping ϕ from D to the set of linear subspaces of a linear space L. A very natural and interesting question (raised in particular by one of the referees) is whether the use of linear algebra is really essential and can not be replaced by a purely combinatorial argument. We will comment on this after we are done with the proof of Lemma 18.

Let E_H consist of those edges $e \in E$ that have *at most one* endpoint in V_P, and let $D_H \overset{\text{def}}{=} \{ a \in D \mid a \subseteq E_H \}$. If $a \notin D_H$, we immediately set $\phi(a) \overset{\text{def}}{=} 0$. Now we show how to define ϕ on D_H. Our construction essentially uses tensor products of linear spaces; we refer to any good textbook in algebra (e.g. [18]) for their definitions and basic properties. In particular (this is what we actually need for our proof), if $L = L_1 \otimes \cdots \otimes L_n$, then for any linear subspaces $L'_1, \ldots, L'_n$ in $L_1, \ldots, L_n$ respectively we can form an uniquely defined subspace in L isomorphic to their tensor product and (for this reason) denoted by $L'_1 \otimes \cdots \otimes L'_n$ with the following two properties.

(1) Denote by $\text{Span}(L_1, \ldots, L_n)$ the linear space spanned by linear subspaces $L_1, \ldots, L_n$ of the same common space L. Then $\otimes$ and Span obey the following distributive law: for any subspaces $L'_1, \ldots, L'_{i-1}, L^1_i, \ldots, L^h_i, L'_{i+1}, \ldots, L'_n$ in the respective $L_1, \ldots, L_n$ we have

$$\begin{aligned} &L'_1 \otimes \cdots \otimes L'_{i-1} \otimes \text{Span}(L^1_i, \ldots, L^h_i) \otimes L'_{i+1} \cdots \otimes L'_n \\ &\quad = \text{Span}(L'_1 \otimes \cdots \otimes L^1_i \otimes \cdots \otimes L'_n, \ldots, L'_1 \otimes \cdots \otimes L^h_i \otimes \cdots \otimes L'_n). \end{aligned}$$

(2) $\dim(L'_1 \otimes \cdots \otimes L'_n) = \prod_{i=1}^n \dim(L'_i)$.

For ease of notation, one-dimensional subspaces L'_i in the expression $L'_1 \otimes \cdots \otimes L'_n$ will be represented by their generating elements.

Fix an arbitrary infinite field k, and for $v \in V_P$ let L_v be an h_v-dimensional linear space over k, where

$$h_v \overset{\text{def}}{=} \left(\frac{\deg_G(v) - d_v}{2} + \frac{\delta_v}{4} \right).$$

Let $L \overset{\text{def}}{=} \bigotimes_{v \in V_P} L_v$. Denote further $N_G(v) \cap V_H$ by $N_H(v)$, and for every $v \in V_P$ fix an arbitrary generic embedding $\phi_v : N_H(v) \longrightarrow L_v$ (so that for every $W \subseteq N_H(v)$ with $|W| = h_v$ the elements $\{ \phi_v(w) \mid w \in W \}$ form a linear basis of L_v).

Next, for $a \in D_H$ we let

$$\phi(a) \overset{\text{def}}{=} \bigotimes_{v \in V_P \setminus \text{dom}(a)} L_v \otimes \bigotimes_{v \in V_P \cap \text{dom}(a)} \phi_v(a_v),$$

where a_v is the uniquely defined vertex in $N_H(v)$ such that $(v, a_v) \in a$. It is important to note that $\phi(a)$ depends only on the set of edges having *exactly* one endpoint in V_P ("pigeon-hole" edges) present in a. Finally, for a positive clause C we let

$$\phi(C) \overset{\text{def}}{=} \text{Span}(\phi(a) | a \in Z(C)).$$

Claim 21 *Suppose that C is obtained from C_0, C_1 via a single application of the positive rule in the refutation P, and assume that $w_d(C_0)$, $w_d(C_1)$ do not exceed $\frac{\delta(G)}{200 \log |V|}$. Then $\phi(C) \subseteq \text{Span}(\phi(C_0), \phi(C_1))$.*

Proof. Fix an arbitrary $a \in Z(C)$; we only need to show that $\phi(a) \subseteq \text{Span}(\phi(C_0), \phi(C_1))$. Let $V' \overset{\text{def}}{=} V_d^\sharp(C_0) \cup V_d^\sharp(C_1)$, and remove from a all edges that are not incident to at least one vertex in V'. Denote the resulting matching by a'. Since the mapping ϕ is anti-monotone w.r.t. inclusion, it is sufficient to show that

$$\phi(a') \subseteq \text{Span}(\phi(C_0), \phi(C_1)). \tag{10}$$

Note that since C is positive, $C(a') = 0$. Let $b \in D_H$ be an extension of a' such that still $C(b) = 0$, and still every $e \in b$ is incident to at least one vertex in V'. Note for the record that the second property implies $|b| \leq w_d(C_0) + w_d(C_1) \leq \frac{\delta(G)}{100 \log |V|}$.

Denote $\pi(b) \overset{\text{def}}{=} |V' \setminus \text{dom}(b)|$. We are going to show by induction on $\pi(b) = 0, 1, \ldots, \pi(a')$ that

$$\phi(b) \subseteq \text{Span}(\phi(C_0), \phi(C_1)). \tag{11}$$

Base $\pi(b) = 0$. Since the positive rule is sound on D, $C(b) = 0$ implies $C_\epsilon(b) = 0$ for some $\epsilon \in \{0, 1\}$. Then $b \in Z(C_\epsilon)$, and (11) follows.

Inductive step. Let $\pi(b) > 0$, and pick an arbitrary $v \in V' \setminus \text{dom}(b)$. Property 2 of Claim 19 (applied to $C = X_E$) implies that $|N_H(v)| \geq \frac{1}{2} \deg_G(v) - \frac{\delta_v}{10}$. Therefore, b has at least

$$|N_H(v)| - 2|b| \geq \frac{1}{2} \deg_G(v) - \frac{\delta_v}{10} - \frac{\delta(G)}{50 \log |V|} \geq \frac{1}{2} \deg_G(v) - \frac{7\delta_v}{50}$$

different extensions $\hat{b} = b \cup \{(v, w)\} \in D_H$ with $w \in H$.

We claim that $v \notin V_d^\sharp(C)$. Indeed, $v \notin \text{dom}(a')$ since $v \notin \text{dom}(b)$ and $a' \subseteq b$. Also, $v \notin \text{dom}(a \setminus a')$ since $v \in V'$ and, therefore, an edge incident to v would not have been removed from a. Hence, $v \notin \text{dom}(a)$ which implies $v \notin V_d^\sharp(C)$ by the definition of $Z(C)$.

This means $\deg_C(v) < d_v - \delta_v$. Applying property 2 of Claim 19 once more, we obtain $|N_C(v) \cap V_H| \leq \frac{1}{2}(d_v - \delta_v) + \frac{\delta_v}{10} = \frac{1}{2}d_v - \frac{2\delta_v}{5}$, and this is the upper bound on the number of extensions $\hat{b}$ of the above form that violate the condition $C(\hat{b}) = 0$. Altogether, we have at least

$$\left(\frac{1}{2}\deg_G(v) - \frac{7\delta_v}{50}\right) - \left(\frac{1}{2}d_v - \frac{2\delta_v}{5}\right) = \frac{1}{2}(\deg_G(v) - d_v) + \frac{13\delta_v}{50} \tag{12}$$

different extensions $\hat{b} = b \cup \{(v, w)\} \in D_H$ with $w \in H$ and such that $C(\hat{b}) = 0$. To every one of these extensions we can apply the inductive hypothesis and conclude that $\phi(\hat{b}) \subseteq \mathrm{Span}(\phi(C_0), \phi(C_1))$.

Now, if $v \in V_H$ then we simply have $\phi(b) = \phi(\hat{b})$ for any such $\hat{b}$ (since b and $\hat{b}$ differ only in one edge (v, w) that is "hole-hole"). If $v \in V_P$ then (12) is greater than h_v; therefore, if $b \cup \{(v, w_1)\}, \ldots, b \cup \{(v, w_t)\}$ is their complete list then $\phi_v(w_1), \ldots, \phi_v(w_t)$ generate L_v. Hence, in this case we also have $\phi(b) \subseteq \mathrm{Span}(\phi_v(w_1), \ldots, \phi_v(w_t))$, and the inductive step follows.

We have completely proved (11). Applying it to $b = a'$, we get (10) which completes the proof of Claim 21.■

Now we complete the proof of Lemma 18 by a simple counting argument. Assume for the sake of contradiction that $w_d(P) < \frac{\delta(G)}{200 \log |V|}$. Note that for every $v \in V$ we have $v \in V_d(Q_v) \subseteq V_d^\sharp(Q_v)$ (since $d_v \leq \deg_G(v)$) and therefore $Z(Q_v) = 0$ (since no partial matching covering v can set Q_v to zero). This implies $\phi(Q_v) = 0$. Also, $V_d^\sharp(0) = \emptyset$ (since $d_v > \delta_v$), the empty matching $\emptyset$ belongs to $Z(0)$ and $\phi(\emptyset) = L$. By iterating Claim 21, we get $\mathrm{Span}(\phi(A) | A \in \mathcal{A}) = L$. Consider an individual $A \in \mathcal{A}$.

Let $V_0 \overset{\text{def}}{=} V_d(A) \cap V_P$. Then, clearly,

$$\phi(A) \subseteq \bigotimes_{v \in V_P \setminus V_0} L_v \otimes \bigotimes_{v \in V_0} \mathrm{Span}(\phi_v(w) | w \in N_H(v) \setminus N_A(v)).$$

By property 1 of Claim 19, $|V_0| \geq w_0/3$. By the definition of $V_d(A)$, $\deg_A(v) \geq d_v$ for every $v \in V_0$, hence, by property 2 of Claim 19, $|N_H(v) \cap N_A(v)| \geq \frac{1}{2}d_v - \frac{\delta_v}{10}$. Also (by the same claim) $|N_H(v)| \leq \frac{1}{2}\deg_G(v) + \frac{\delta_v}{10}$. Therefore, $|N_H(v) \setminus N_A(v)| \leq \frac{1}{2}(\deg_G(v) - d_v) + \frac{\delta_v}{5} = h_v - \frac{\delta_v}{20}$. Putting things together, we get

$$\frac{\dim(\phi(A))}{\dim(L)} \leq \prod_{v \in V_0} \frac{h_v - \delta_v/20}{h_v} \leq \exp\left(-\Omega(w_0/(\log |V|))\right) < S_0^{-1}$$

if the constant ϵ in (7), (8) is small enough. Since $|\mathcal{A}| \leq S_0$, $\{\phi(A) \mid A \in \mathcal{A}\}$ can not generate L, and this contradiction completes the proof of Lemma 18.■

Remark 22 The algebraic argument used in the proof of Lemma 18 looks rather ad hoc, and it also hinders further potential applications of our method. It would be extremely interesting to replace this with some purely combinatorial reasoning; in particular, the author believes that the yet unknown combinatorial machinery for this purpose would very likely suffice to solve a couple of open problems from Section 6. Unfortunately, so far we have not been able to do anything along these lines.

Proof of Theorem 6. Let G be a graph, and define the parameters δ_v, S_0, w_0 by (5), (7), (8). Assume first that (9) is true, and assume, for the sake of contradiction, that $S_R(PM(G)) \leq S_0$. Applying Lemma 13, we get a positive calculus refutation of $\{\, Q_v \mid v \in V \,\}$ that has the same size bound S_0. Applying Lemma 3.1, we get some vector d and another positive calculus refutation with the same size bound S_0 that additionally allows (w_0, d)-axioms, but at the same time obeys the bound (6) on pseudo-width. This bound, however, is in a direct contradiction with Lemma 18, as long as the constant ϵ in (8) is small enough. This contradiction shows that $S_R(PM(G)) \geq S_0$ and proves Theorem 6 in the case $S_0 \geq |V|$.

In order to take care of the "degenerate" case $S_0 \leq |V|$, let P be the minimum size resolution refutation of $PM(G)$. If $S(P) \geq S_0$, we are done so let us assume $S(P) \leq S_0$. Let $V_{\text{active}} \overset{\text{def}}{=} \{v \in V \mid Q_v \in P\}$, then $|V_{\text{active}}| \leq S(P) \leq S_0$ and $S_R(M(V_{\text{active}}|G)) = S(P)$. However, when we relativize the argument in this section (see Remarks 9, 14, 16), the relativized version of (9) will become $S_0 \geq |V_{\text{active}}|$, and that much we already know (note that the value of S_0 can only increase under the relativization!)

Theorem 6 is completely proved.■

We conclude this section with one technical observation that will make things cleaner in Section 4. Let $S_P(\neg onto - FPHP^m_n) \overset{\text{def}}{=} S_P(PM(K_{m,n}))$ be the *positive calculus* complexity of the functional onto version of the pigeonhole principle. Since Theorem 6, by the nature of its proof, readily applies to the positive calculus, we also have

Lemma 23 *For $m > n$, $S_P(\neg onto - FPHP^m_n) \geq \exp\left(\Omega\left(\frac{n}{(\log m)^2}\right)\right)$*

(a straightforward application of Lemma 13 would have resulted in an annoying $(mn)^2$ factor).

4 Unprovability of circuit lower bounds by small resolution proofs

The material of this section is a minor adaptation of [19, Section 5], so we will be rather concise (and, in particular, skip all motivations). Also, it is not used anywhere in Section 5, so the reader interested to see the conclusion of the proof of Theorem 4 may proceed directly to that section.

Let f_n be a Boolean function in n variables, and let $t \leq 2^n$. Denote by $Circuit_t(f_n)$ the following 5-CNF of size $2^{O(n)}$ encoding the description of a size-t Boolean circuit for computing f_n.

First, we list all variables of $Circuit_t(f_n)$ (some of them have peculiar long names like $InputType'_\nu(v)$), along with their intended meaning:

y_{av} $(a \in \{0,1\}^n,\ v \in [t])$ – the Boolean value computed at the node v on the input string a;

$y_{a\nu v}$ $(a \in \{0,1\}^n,\ \nu \in \{1,2\},\ v \in [t])$ – the value brought to v by ν's wire on a;

$Fanin(v)$ – this is 0 if v is NOT-gate and 1 if v is AND-gate or OR-gate;

$Type(v)$ – when $Fanin(v) = 1$, this is 0 if v is AND-gate and 1 if v is OR-gate;

$InputType_\nu(v)$ – this is 0 if ν's input to v is a constant or a variable and 1 if it is one of the previous gates;

$InputType'_\nu(v)$ – when $InputType_\nu(v) = 0$, this is 0 if ν's input to v is a constant, and 1 if it is a variable;

$InputType''_\nu(v)$ – when $InputType_\nu(v) = InputType'_\nu = 0$, this equals the ν's input to v;

$InputVar_\nu(v,i)$ $(i \in [n])$ – when $InputType_\nu(v) = 0$, $InputType'_\nu(v) = 1$, this is 1 iff ν's input to v is x_i;

$INPUTVAR_\nu(v,i)$ – equals $\bigvee_{i' \leq i} InputVar(v,i')$, introduced to keep bottom fan-in bounded;

$InputNode_\nu(v,v')$ $(v' < v)$ – when $InputType_\nu(v) = 1$, this is 1 iff ν's input to v is the previous

gate v';
$INPUTNODE_\nu(v,v')$ – analogously to $INPUTVAR_\nu(v,i)$.

$Circuit_t(f_n)$ is the conjunction of (conjunctive normal forms equivalent to) the following axioms:

$\neg InputType_\nu(v) \wedge \neg InputType'_\nu(v) \longrightarrow (y_{a\nu v} \equiv InputType''_\nu(v))$;
$\neg InputType_\nu(v) \wedge InputType'_\nu(v) \longrightarrow \neg(InputVar_\nu(v,i) \wedge$
$InputVar_\nu(v,i'))$ $(i \neq i')$;
$\neg InputType_\nu(v) \wedge InputType'_\nu(v) \longrightarrow (INPUTVAR_\nu(v,i) \equiv$
$(INPUTVAR_\nu(v,i-1) \vee InputVar_\nu(v,i)))$
$(INPUTVAR_\nu(v,0) \overset{\text{def}}{\equiv} 0)$;
$\neg InputType_\nu(v) \wedge InputType'_\nu(v) \longrightarrow INPUTVAR_\nu(v,n)$;
$\neg InputType_\nu(v) \wedge InputType'_\nu(v) \wedge InputVar_\nu(v,i) \longrightarrow (y_{a\nu v} \equiv a_i)$;
the analogous group of axioms for $InputNode$;
$\neg Fanin(v) \longrightarrow (y_{av} \equiv \neg y_{a1v})$;
$Fanin(v) \wedge \neg Type(v) \longrightarrow (y_{av} \equiv (y_{a1v} \wedge y_{a2v}))$;
$Fanin(v) \wedge Type(v) \longrightarrow (y_{av} \equiv (y_{a1v} \vee y_{a2v}))$;
$y_{at} \equiv f(a)$.

In this section we prove

Theorem 24 $S_R(Circuit_t(f_n)) \geq \exp(\Omega(t/n^3))$.

Proof. [19, Section 5] constructed a reduction from $\neg FPHP^m_t$ to $Circuit_t(f_n)$ which works in the context of the Polynomial Calculus. A closer inspection reveals that it will also work for Resolution, but only if we weaken $FPHP^m_t$ to $onto - FPHP^m_t$ (cf. [13]), and our proof will essentially consist in conducting this inspection.

Definition 25 *$PDNF_t(f_n)$ is the following 3-CNF of size $2^{O(n)}$ encoding the description of a size-t perfect DNF $K_1 \vee \ldots \vee K_t$ (K_j elementary conjunctions of maximal length n) for computing f_n. The variables of $PDNF_t(f_n)$, along with their intended meaning, are:*

y_{ajk} $(a \in \{0,1\}^n,\ j \in [t],\ k \in [n])$ – *a is consistent with the first k literals in K_j;*
y_{aj} $(a \in \{0,1\}^n,\ j \in [t])$ – $K_1(a) \vee \ldots \vee K_j(a) = 1$;
z_{jk} $(j \in [t],\ k \in [n])$ – *the sign with which x_k occurs in K_j.*

The axioms of $PDNF_t(f_n)$ are (the 3-CNF resulting from expanding):

$$\begin{aligned} y_{ajk} &\equiv \left(y_{aj,k-1} \wedge z_{jk}^{a_k}\right) \ \ (\textit{with } y_{aj0} \overset{\text{def}}{=} 1); \\ y_{aj} &\equiv (y_{a,j-1} \vee y_{ajn}) \ \ (\textit{with } y_{a0} \overset{\text{def}}{=} 0); \\ y_{at} &\equiv f(a). \end{aligned}$$

Proposition 26 $S_R(Circuit_t(f_n)) \geq S_R(PDNF_{\lceil t/2n \rceil}(f_n))$

Proof. [19, proof of Corollary 5.2] noticed the existence of a *variable substitution* that takes $Circuit_t(f_n)$ to $PDNF_{\lceil t/2n \rceil}(f_n)$, and variable substitutions work for any reasonable proof system including Resolution.■

Lemma 27 *There exists m with $t+1 \leq m \leq 2^n$ such that*

$$S_R(PDNF_t(f_n)) \geq S_P(\neg onto - FPHP_n^m).$$

Proof. (cf. [19, proof of Theorem 5.1]) Let $m \overset{\text{def}}{=} |f^{-1}(1)|$; we may assume that $m \geq t+1$ since otherwise $PDNF_t(f_n)$ is satisfiable. Identify pigeons $i \in [m]$ with Boolean assignments $a \in f^{-1}(1)$; thus, pigeonhole variables will look like x_{aj} where $f(a) = 1$. Construct the following mapping that takes every *literal* of a variable of $PDNF_t(f_n)$ to a *positive clause* in the pigeonhole variables:

$$\begin{aligned} y_{ajk} &\mapsto \bigvee \{x_{bj} \mid f(b) = 1 \wedge b_1 = a_1 \wedge \ldots \wedge b_k = a_k\} \ \ (a \in \{0,1\}^n); \\ \bar{y}_{ajk} &\mapsto \bigvee \{x_{bj} \mid f(b) = 1 \wedge \neg(b_1 = a_1 \wedge \ldots \wedge b_k = a_k)\} \ \ (a \in \{0,1\}^n); \\ y_{aj}^{\epsilon} &\mapsto \bar{\epsilon} \ (f(a) = 0); \\ y_{aj} &\mapsto \bigvee_{j' \leq j} x_{aj'} \ (f(a) = 1); \\ \bar{y}_{aj} &\mapsto \bigvee_{j' > j} x_{aj'} \ (f(a) = 1); \\ z_{jk}^{\epsilon} &\mapsto \bigvee \{x_{aj} \mid f(a) = 1 \wedge a_k = \epsilon\}. \end{aligned}$$

An easy inspection shows that this mapping takes every resolution refutation of $PDNF_t(f_n)$ into a positive calculus refutation of $\neg onto - FPHP_n^m$. Lemma 27 follows.■

Theorem 24 is now immediately implied by Proposition 26, Lemma 27 and Lemma 23.■

5 Proof of the main result: general case

In this section we show how to adapt the proof from Section 3 to the case of arbitrary hypergraphs and prove Theorem 4. Before we begin, let us pinpoint the main difficulties with the naive generalization (all of them are in one or another way related to Claim 19).

Recall the discussion at the beginning of Section 3.2. Given a partition (V_P, V_H), we still must classify every edge with at least two pigeons in it as useless (see the definition of D_H in the proof of Claim 21). As long as $r(\mathcal{H})$ is large, this will result in the unpleasant fact that there are only a few useful (that is, pigeon-hole) edges, and the argument breaks down. We circumvent this by biasing our distribution $(\boldsymbol{V_P}, \boldsymbol{V_H})$ in favour of *holes* (notice the striking difference with the ordinary PHP^m_n), and we have to pay for this an extra $r(\mathcal{H})$ factor in the final bound.

The most serious problem, however, is structural rather than numerical: we no longer have a workable definition of the *vertex* neighbourhood set $N_{\mathcal{H}}(v)$, and we must work with stars $S_{\mathcal{H}}(v)$ instead. This in particular implies that, as long as $\lambda(\mathcal{H}) > 1$, the edges in this star are no longer classified independently of each other, and we are facing difficulties with estimating the probability of large deviation in proving property 2 of Claim 19. We circumvent this by an ad hoc trick, and we will have to pay at least an extra $\lambda(\mathcal{H})$ factor in the final bound for this trick.

Finally, as long as $\mathcal{H}$ is not uniform, the probability that $E \in S_{\mathcal{H}}(v)$ will be classified as (say) pigeon-hole edge also depends on $|E|$. This makes even the expectation in the proof of property 2 of Claim 19 unpredictable in terms of $\deg_C(v)$. The remedy for this, however, is very easy (and comes free of charge): we assign to edges appropriate weights according to their size.

Let us now turn to the formal argument. Fix a hypergraph $\mathcal{H} = (V, \mathcal{E})$. For $E \in \mathcal{E}$, we define its weight as

$$\mu(E) \overset{\text{def}}{=} \left(1 - \frac{1}{r(\mathcal{H})}\right)^{|E|-1}$$

(the reason for this choice of the weight function will become clear in the proof of Claim 29). We adjust all degree-depending notions according to this weight function:

$$\widetilde{\deg}_{\mathcal{H}}(v) \stackrel{\rm def}{=} \sum_{E\in S_{\mathcal{H}}(v)} \mu(E),$$
$$\tilde{\delta}(\mathcal{H}) \stackrel{\rm def}{=} \min_{v\in V} \widetilde{\deg}_{\mathcal{H}}(v),$$

and for a positive clause C in the variables $\{ x_E \mid E \in \mathcal{E} \}$,

$$\widetilde{\deg}_C(v) \stackrel{\rm def}{=} \sum_{E\in S_C(v)} \mu(E),$$

where naturally

$$S_C(v) \stackrel{\rm def}{=} \{ E \in S_{\mathcal{H}}(v) \mid x_E \in C \}.$$

Note for the record that $e^{-1} \leq \mu(E) \leq 1$, hence $\widetilde{\deg}_{\mathcal{H}}(v), \tilde{\delta}(\mathcal{H}), \widetilde{\deg}_C(v)$ differ from their unweighted analogs by at most a constant factor. Also, for ease of comparison with Section 3, note that $\mu(E) = 1/2$ for ordinary graphs G and we have $\widetilde{\deg}_G(v) = \frac{1}{2}\deg_G(v)$, $\tilde{\delta}(G) = \frac{1}{2}\delta(G)$ and $\widetilde{\deg}_C(v) = \frac{1}{2}\deg_C(v)$.

We now adapt the next series of definitions as follows. We let

$$\delta_v \stackrel{\rm def}{=} \frac{\widetilde{\deg}_{\mathcal{H}}(v)}{2\log|V|}.$$

For a vector $d = (d_v \mid v \in V)$, let

$$V_d(C) \stackrel{\rm def}{=} \left\{ v \in V \,\middle|\, \widetilde{\deg}_C(v) \geq d_v \right\}$$

and

$$V_d^\sharp(C) \stackrel{\rm def}{=} \left\{ v \in V \,\middle|\, \widetilde{\deg}_C(v) \geq d_v - \delta_v \right\}.$$

$w_d(C), w_d(P)$ and the notion of an (w_0, d)-axiom are defined on the base of these new $V_d(C), V_d^\sharp(C)$ exactly as before. The adjustment of Lemma 17 is fairly straightforward:

Lemma 28 *Suppose that there exists a positive calculus refutation P of $\{ Q_v \mid v \in V \}$, and let $w_0 \leq \frac{\tilde{\delta}(\mathcal{H})}{4\lambda(\mathcal{H})}$ be an arbitrary integer parameter. Then there exists an integer vector $d = (d_v \mid v \in V)$ with $\delta_v < d_v \leq \widetilde{\deg}_{\mathcal{H}}(v)$ for all $v \in V$, a set of (w_0, d)-axioms $\mathcal{A}$ and a positive calculus refutation P' of $\{ Q_v \mid v \in V \} \cup \mathcal{A}$ such that $S(P') \leq S(P)$ and $w_d(P') \leq O(w_0 + \log S(P))$.*

The only remark which should be made in connection with its proof is this: if $|V_0| = w_0$ and $E \cap V_0 \neq \emptyset$ for every $x_E \in C'$, then $\widetilde{\deg}_{C'}(v) \leq \deg_{C'}(v) \leq w_0\lambda(G)$ for every $v \in V_0$ (this guarantees that after cleaning up any (w_0, d)-axiom its pseudo-width will get reduced to w_0).

Fix the parameters S_0, w_0 as

$$S_0 \stackrel{\rm def}{=} \exp\left(\frac{\epsilon^2\tilde{\delta}(\mathcal{H})}{\lambda(\mathcal{H})r(\mathcal{H})(\log|V|)(r(\mathcal{H})+\log|V|)}\right); \qquad (13)$$

$$w_0 \stackrel{\rm def}{=} \frac{\epsilon\tilde{\delta}(\mathcal{H})}{\lambda(\mathcal{H})r(\mathcal{H})(\log|V|)}. \qquad (14)$$

Instead of (9), we will need the stronger inequality

$$S_0 \geq \max\{|V|, |\mathcal{E}|\}. \qquad (15)$$

Under the assumptions of Lemma 18, we will be proving the lower bound

$$w_d(P) \geq \frac{\tilde{\delta}(\mathcal{H})}{50\lambda(\mathcal{H})r(\mathcal{H})(\log|V|)}. \qquad (16)$$

Fix $d = (d_v \mid v \in V)$, $\mathcal{A}$ and P satisfying those assumptions.

Claim 29 *There exist a partition $V = V_P \dot\cup V_H$ such that the following two properties are satisfied:*

(1) for every $A \in \mathcal{A}$, $|V_d(A) \cap V_P| \geq w_0/(2r(\mathcal{H}))$;
(2) for every $C \in P \cup \{X_{\mathcal{E}}\}$ and every $v \in V$,

$$\Big| |\{E \in S_C(v) \mid E - \{v\} \subseteq V_H\}| - \widetilde{\deg}_C(v) \Big| \leq \frac{\delta_v}{5}.$$

Remark 30 Note that in 2 we have the *real* cardinality of the set $\{E \in S_C(v) \mid E - \{v\} \subseteq V_H\}$, not its weighted version.

Since this claim is most seriously affected by the transition from graphs to hypergraphs, we give its complete proof from scratch.

Proof of Claim 29. Let $(\boldsymbol{V_P} \cup \boldsymbol{V_H})$ be a random partition of V in which for every $v \in V$, $\mathbf{P}[v \in \boldsymbol{V_P}] = \frac{1}{r(\mathcal{H})}$, and these events are independent for different v. Applying Proposition 20 to every individual $A \in \mathcal{A}$, we get $\mathbf{P}[|V_d(A) \cap \boldsymbol{V_P}| \leq w_0/(2r(\mathcal{H}))] \leq \exp(-\Omega(w_0/r(\mathcal{H}))) \leq S_0^{-2}$, as long as the constant ϵ in (13), (14) is small enough.

Fix now an individual positive clause C and $v \in V$. Recall that a set system $S_1, \ldots, S_t$ is called a *sunflower* if all pairwise intersections $S_i \cap S_j$ $(1 \leq i < j \leq t)$ are equal to the same set called *the center of the sunflower*.

Claim 31 *There exists a partition $S_C(v) = S_C^1(v) \dot\cup \ldots \dot\cup S_C^t(v)$, where for every $\nu \in [t]$, $S_C^\nu(v)$ is a sunflower with the center $\{v\}$ and $t \leq r(\mathcal{H})\lambda(\mathcal{H})$.*

Proof of Claim 31. Let us construct an auxiliary (ordinary) graph on $S_C(v)$ by connecting E and E' if and only if $E \cap E' \neq \{v\}$. The degree of every vertex

E in this auxiliary graph is bounded by $(r(\mathcal{H})-1)(\lambda(\mathcal{H})-1)$: there are at most $r(\mathcal{H})-1$ choices of $v' \neq v$ in E, and for every such v' at most $\lambda(\mathcal{H})-1$ edges $E' \neq E$ containing both v and v'. Hence, the chromatic number of this auxiliary graph does not exceed $(r(\mathcal{H})-1)(\lambda(\mathcal{H})-1)+1 \leq r(\mathcal{H})\lambda(\mathcal{H})$. It only remains to note that independent sets in this graph are exactly $\{v\}$-centered sunflowers in $\mathcal{H}$.■

Now we are ready to analyze the probability of large deviation for $|\{E \in S_C(v) \mid E - v \subseteq \boldsymbol{V_H}\}|$. Note first that

$$\mathbf{E}[|\{E \in S_C(v) \mid E - \{v\} \subseteq \boldsymbol{V_H}\}|] = \sum_{E \in S_C(v)} \mathbf{P}[E - v \subseteq \boldsymbol{V_H}]$$

$$= \sum_{E \in S_C(v)} \left(1 - \frac{1}{r(\mathcal{H})}\right)^{|E|-1} = \widetilde{\deg}_C(v).$$

Next, fix the partition $S_C(v) = S_C^1(v) \,\dot{\cup}\, \ldots \,\dot{\cup}\, S_C^t(v)$ guaranteed by Claim 31. Note that $\boldsymbol{S_\nu} \stackrel{\text{def}}{=} \{E \in S_C^\nu(v) \mid E - \{v\} \subseteq \boldsymbol{V_H}\}$ is a sum of *independent* 0-1 variables; denote its expectation by E^ν. We have $\sum_{\nu=1}^t E^\nu = \widetilde{\deg}_C(v)$. Applying Proposition 20, we get

$$\mathbf{P}\Big[|\boldsymbol{S_1} + \cdots + \boldsymbol{S_t} - \widetilde{\deg}_C(v)| \geq \delta_v/5\Big]$$
$$\leq \sum_{v=1}^{t} \mathbf{P}\left[|\boldsymbol{S_\nu} - E^\nu| \geq \frac{\delta_v}{10}\left(\frac{1}{t} + \sqrt{\frac{E^\nu}{t \cdot \widetilde{\deg}_C(v)}}\right)\right]$$
$$\leq t \cdot \exp\left(-\Omega\left(\frac{\delta_v^2}{t \cdot \widetilde{\deg}_{\mathcal{H}}(v)}\right)\right) \leq S_0^{-3}$$

as long as the constant ϵ in (13) is small enough (for the last inequality we also need to observe that $t \leq r(\mathcal{H})\lambda(\mathcal{H}) \leq |V| \cdot |\mathcal{E}| \leq S_0^2$ by (15)). Claim 29 now follows by the union bound.■

The definitions of D and $Z(C)$ do not change. Let

$$\mathcal{E}_H \stackrel{\text{def}}{=} \{E \in \mathcal{E} \mid |E \cap V_P| \leq 1\}$$

and, as before, $D_H \stackrel{\text{def}}{=} \{a \in D \mid a \subseteq \mathcal{E}_H\}$. As before, the mapping ϕ vanishes outside of D_H.

For $v \in V_P$ let

$$h_v \stackrel{\text{def}}{=} (\widetilde{\deg}_{\mathcal{H}}(v) - d_v) + \delta_v/2,$$

and let

$$S_H(v) \stackrel{\text{def}}{=} \{E \in S_{\mathcal{H}}(v) \mid E - \{v\} \subseteq V_H\}.$$

We construct generic embeddings $\phi_v : S_H(v) \longrightarrow L_v$, their tensor product ϕ and its action on D_H just in the same way as before. There are no structural changes to the proof of Claim 21, but we need to adjust calculations. Recall that we assume the upper bound (16) on $w_d(C_0), w_d(C_1)$. Then we have $|b| \leq \frac{\tilde{\delta}(\mathcal{H})}{25\lambda(\mathcal{H})r(\mathcal{H})(\log|V|)}$ for every matching b considered in that proof. In particular, any such b covers at most $\frac{\tilde{\delta}(\mathcal{H})}{25\lambda(\mathcal{H})(\log|V|)}$ vertices.

In the inductive step, the lower bound on the number of extensions $\hat{b} = b \cup \{E\} \in D_H$ with $E \in S_H(v)$ becomes $|S_H(v)| - r(\mathcal{H})\lambda(\mathcal{H})|b| \geq \widetilde{\deg}_{\mathcal{H}}(v) - \frac{\delta_v}{5} - \frac{\tilde{\delta}(\mathcal{H})}{25(\log|V|)} \geq \widetilde{\deg}_{\mathcal{H}}(v) - \frac{7\delta_v}{25}$, and the upper bound on the number of these extensions violating $C(\hat{b}) = 0$ becomes $|S_C(v) \cap S_H(v)| \leq (d_v - \delta_v) + \frac{\delta_v}{5} = d_v - \frac{4\delta_v}{5}$. Altogether we have at least $\widetilde{\deg}_{\mathcal{H}}(v) - d_v + \frac{13\delta_v}{25}$ good extensions which is greater than h_v if $v \in V_P$.

The rest of the proof of Theorem 4 (under the assumption (15)) closely follows the pattern in Section 3 (note that the factor of $r(\mathcal{H})$ lost in part 1 of Proposition 29 is accounted for in (13)).

Finally, we show how to get rid of (15). Once more, let P be the minimal size refutation of $PM(\mathcal{H})$ such that $S(P) \leq S_0$ and $V_{\text{active}} \overset{\text{def}}{=} \{v \in V \mid Q_v \in P\}$. Let also $\mathcal{E}_{\text{active}} \overset{\text{def}}{=} \bigcup_{v \in V_{\text{active}}} S_{\mathcal{H}}(v)$. By relativizing the whole argument to V_{active}, the assumption (15) can be relaxed to $S_0 \geq \max\{|V_{\text{active}}|, |\mathcal{E}_{\text{active}}|\}$. It only remains to note that we also have $S(P) \geq |\mathcal{E}_{\text{active}}|$. The reason is that every x_E with $E \in \mathcal{E}_{\text{active}}$ must be resolved somewhere in the refutation P since otherwise for every $v \in E \cap V_{\text{active}}$, there could not be any path from Q_v to the empty clause and, contrary to the minimality of P, we could have removed Q_v from it.

6 Open problems

Currently there are two different techniques for proving lower bounds on $S_R(PM(\mathcal{H}))$. The first method [7] is based on the width-size relation and is applicable only when the minimal degree $\delta(\mathcal{H})$ tends to be small. Our method, on the contrary, can be only applied when $\delta(\mathcal{H})$ is large. Can we find their common generalization that would uniformly cover both cases? For example, is it true that $S_R(G - PHP) \geq \exp\left(n^{\Omega(1)}\right)$ for any bipartite G on $[m] \times [n]$ that has a constant expansion rate, without any assumptions about m and the degrees $\deg_G(i)$? This is true if the number of edges is $\leq n^{2-\Omega(1)}$ [7] *or* $\min_i \deg_G(i) \geq n^{\Omega(1)}$ (Theorem 11).

Can the methods developed in [14,10] and in this paper be applied to the tautologies $\tau(A, \vec{g}), \tau_{\oplus}(A, b)$ introduced in [8] and expressing the hardness of the Nisan-Wigderson generator in the context of propositional proof complexity?

The best known upper bound on $S_R(\neg PHP_n^m)$ is $\exp(O(n \log n)^{1/2})$ [16], and we have shown the lower bound $S_R(\neg onto - FPHP_n^m) \geq \exp(\Omega(n^{1/3}))$. It would be interesting to further narrow this gap. Specifically, what is the value of $\limsup_{n\to\infty} \frac{\log_2 \log_2 S_R(\neg PHP_n^\infty)}{\log_2 n}$?

It appears as if one could hope to get slightly better lower bounds for the counting principle $Count_r^n$ by using ordinary restrictions instead of our machinery. Is it for example true that $S_R(\neg Count_r^n) \geq \exp(\Omega(n))$ for any constant r?

Acknowledgement

I am grateful to both anonymous referees for a great deal of useful remarks.

References

[1] A. Razborov, Lower bounds for propositional proofs and independence results in Bounded Arithmetic, in: F. M. auf der Heide, B. Monien (Eds.), Proceedings of the 23rd ICALP, Lecture Notes in Computer Science, 1099, Springer-Verlag, New York/Berlin, 1996, pp. 48–62.

[2] P. Beame, T. Pitassi, Propositional proof complexity: Past, present and future, Tech. Rep. TR98-067, Electronic Colloquium on Computational Complexity, available at ftp://ftp.eccc.uni-trier.de/pub/eccc/reports/1998/TR98-067/index.html (1998).

[3] G. C. Tseitin, On the complexity of derivations in propositional calculus, in: Studies in constructive mathematics and mathematical logic, Part II, Consultants Bureau, New-York-London, 1968.

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

[5] S. Jukna, Exponential lower bounds for semantic resolution, in: P. Beame, S. Buss (Eds.), Proof Complexity and Feasible Arithmetics: DIMACS workshop, April 21-24, 1996, DIMACS Series in Dicrete Mathematics and Theoretical Computer Science, vol. 39, American Math. Soc., 1997, pp. 163–172.

[6] A. Urquhart, Resolution proofs of matching principles, to appear in *Annals of Mathematics and Artificial Intelligence* (1998).

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

[8] M. Alekhnovich, E. Ben-Sasson, A. Razborov, A. Wigderson, Pseudorandom generators in propositional complexity, in: Proceedings of the 41st IEEE FOCS, 2000, pp. 43–53, journal version to appear in *SIAM Journal on Computing*.

[9] P. Pudlák, R. Impagliazzo, A lower bound for DLL algorithms for k-SAT, in: Proceedings of the 11th Symposium on Discrete Algorithms, 2000, pp. 128–136.

[10] A. Razborov, Resolution lower bounds for the weak functional pigeonhole principle, Theoretical Computer Science 303 (1) (2003) 233–243.

[11] A. Razborov, A. Wigderson, A. Yao, Read-once branching programs, rectangular proofs of the pigeonhole principle and the transversal calculus, in: Proceedings of the 29th ACM Symposium on Theory of Computing, 1997, pp. 739–748.

[12] T. Pitassi, R. Raz, Regular resolution lower bounds for the weak pigeonhole principle, in: Proceedings of the 33rd ACM Symposium on the Theory of Computing, 2001, pp. 347–355.

[13] R. Raz, Resolution lower bounds for the weak pigeonhole principle, in: Proceedings of the 34th ACM Symposium on the Theory of Computing, 2002, pp. 553–562.

[14] A. Razborov, Improved resolution lower bounds for the weak pigeonhole principle, Tech. Rep. TR01-055, Electronic Colloquium on Computational Complexity (2001).

[15] J. Krajíček, Bounded arithmetic, propositional logic and complexity theory, Cambridge University Press, 1995.

[16] S. Buss, T. Pitassi, Resolution and the weak pigeonhole principle, in: Proceedings of the CSL97, Lecture Notes in Computer Science, 1414, Springer-Verlag, New York/Berlin, 1997, pp. 149–156.

[17] J. V. Uspensky, Introduction to mathematical probability, McGraw-Hill Book Company, 1937.

[18] S. Lang, Algebra, 3rd Edition, Springer-Verlag, 2002.

[19] A. Razborov, Lower bounds for the polynomial calculus, Computational Complexity 7 (1998) 291–324.