

Exponential Resolution Lower Bounds for Weak Pigeonhole Principle and Perfect Matching Formulas over Sparse Graphs

Received Feb 19, 2024
Revised Dec 12, 2024
Accepted Jan 20, 2025
Published Mar 25, 2025

Key words and phrases
Computational Complexity, Proof Complexity, Resolution, Pigeonhole Principle

Susanna F. de Rezende^a ✉ 

Jakob Nordström^b ✉ 

Kilian Risse^c ✉ 

Dmitry Sokolov^c ✉ 

^a Lund University, Sweden

^b University of Copenhagen, Denmark, and Lund University, Sweden

^c EPFL, Lausanne, Switzerland

ABSTRACT. We show exponential lower bounds on resolution proof length for pigeonhole principle (PHP) formulas and perfect matching formulas over highly unbalanced, sparse expander graphs, thus answering the challenge to establish strong lower bounds in the regime between balanced constant-degree expanders as in [Ben-Sasson and Wigderson '01] and highly unbalanced, dense graphs as in [Raz '04] and [Razborov '03, '04]. We obtain our results by revisiting Razborov's pseudo-width method for PHP formulas over dense graphs and extending it to sparse graphs. This further demonstrates the power of the pseudo-width method, and we believe it could potentially be useful for attacking also other longstanding open problems for resolution and other proof systems.

1. Introduction

In one sentence, proof complexity is the study of efficient certificates of unsatisfiability for formulas in conjunctive normal form (CNF). In its most general form, this is the question of whether coNP can be separated from NP or not, and as such appears out of reach for current

A preliminary version of this article appeared at CCC 2020 [18]. The authors were funded by the Knut and Alice Wallenberg grant KAW 2016.0066. In addition, the second author was supported by the Swedish Research Council grant 2016-00782 and the Independent Research Fund Denmark (DFF) grant 9040-00389B, and the fourth author by the Knut and Alice Wallenberg grant KAW 2016.0433. Most of this work was carried out while the authors were affiliated to KTH Royal Institute of Technology. Some parts were done while visiting the Simons Institute for the Theory of Computing in association with the DIMACS/Simons Collaboration on Lower Bounds in Computational Complexity, which is conducted with support from the National Science Foundation.

techniques. However, if one instead focuses on concrete proof systems, which can be thought of as restricted models of nondeterministic computation, this opens up the view to a rich landscape of results.

One line of research in proof complexity has been to prove superpolynomial lower bounds for stronger and stronger proof systems, as a way of approaching the distant goal of establishing $\text{NP} \neq \text{coNP}$. A perhaps even more fruitful direction, however, has been to study different combinatorial principles and investigate what kind of reasoning is needed to efficiently establish the validity of these principles. In this way, one can quantify the “depth” of different mathematical truths, measured in terms of how strong a proof system is required to prove them.

In this paper, we consider the proof system *resolution* [10], in which one derives new disjunctive clauses from the formula until an explicit contradiction is reached. This is arguably the most well-studied proof system in proof complexity, for which numerous exponential lower bounds on proof size have been shown (starting with [13, 20, 32]). Yet many basic questions about resolution remain stubbornly open. One such set of questions concerns the *pigeonhole principle (PHP)* stating that there is no injective mapping of m pigeons into n holes if $m > n$. This is one of the simplest, and yet most useful, combinatorial principles in mathematics, and it has been topic of extensive study in proof complexity with deep connections to TFNP classes.

When studying the pigeonhole principle, it is convenient to think of it in terms of a bipartite graph $G = (U \dot{\cup} V, E)$ with pigeons $U = [m]$ and holes $V = [n]$ for $m \geq n + 1$. Every pigeon i can fly to its neighbouring pigeonholes $N(i)$ as specified by G , which for now we fix to be the complete bipartite graph $K_{m,n}$ with $N(i) = [n]$ for all $i \in [m]$. Since we wish to study unsatisfiable formulas, we encode the claim that there does in fact exist an injective mapping of pigeons to holes as a CNF formula consisting of *pigeon axioms*

$$P^i = \bigvee_{j \in N(i)} x_{ij} \quad \text{for } i \in [m] \quad (1.1a)$$

and *hole axioms*

$$H_j^{i,i'} = (\bar{x}_{ij} \vee \bar{x}_{i'j}) \quad \text{for } i \neq i' \in [m], j \in N(i) \cap N(i') \quad (1.1b)$$

(where the intended meaning of the variables is that $x_{i,j}$ is true if pigeon i flies to hole j). To rule out multi-valued mappings one can also add *functionality axioms*

$$F_{j,j'}^i = (\bar{x}_{ij} \vee \bar{x}_{ij'}) \quad \text{for } i \in [m], j \neq j' \in N(i), \quad (1.1c)$$

and a further restriction is to include *surjectivity* or *onto axioms*

$$S_j = \bigvee_{i \in N(j)} x_{ij} \quad \text{for } j \in [n] \quad (1.1d)$$

requiring that every hole should get a pigeon. Clearly, the “basic” *pigeonhole principle (PHP) formulas* with clauses (1.1a) and (1.1b) are the least constrained. As one adds clauses (1.1c) to obtain the *functional pigeonhole principle (FPHP)* and also clauses (1.1d) to get the *onto*

functional pigeonhole principle (onto-FPHP), the formulas become more overconstrained and thus (potentially) easier to disprove, meaning that establishing lower bounds becomes harder. A moment of reflection reveals that onto-FPHP formulas are just saying that complete bipartite graphs with m left vertices and n right vertices have perfect matchings, and so these formulas are also referred to as *perfect matching formulas*.

Another way of varying the hardness of PHP formulas is by letting the number of pigeons m grow larger as a function of the number of holes n . What this means is that it is not necessary to count exactly to refute the formulas. Instead, it is sufficient to provide a precise enough estimate to show that $m > n$ must hold (where the hardness of this task depends on how much larger m is than n). Studying the hardness of such so-called *weak PHP formulas* gives a way of measuring how good different proof systems are at approximate counting. A second application of lower bounds for weak PHP formulas is that they can be used to show that proof systems cannot produce efficient proofs of the claim that $\text{NP} \not\subseteq \text{P/poly}$ [26, 29].

Yet another version of more constrained formulas is obtained by restricting what choices the pigeons have for flying into holes, by defining the formulas not over $K_{m,n}$ but sparse bipartite graphs with bounded left degree—such instances are usually called *graph PHP formulas*. Again, this makes the formulas easier to disprove in the sense that pigeons are more constrained, and it also removes the symmetry in the formulas that plays an essential role in many lower bound proofs.

Our work focuses on the most challenging setting in terms of lower bounds, when all of these restrictions apply: the PHP formulas contain both functionality and onto axioms, the number of pigeons m is very large compared to the number of holes n , and the choices of holes are restricted by a sparse graph. But before discussing our contributions, let us review what has been known about resolution and pigeonhole principle formulas. We emphasize that what will follow is a brief and selective overview focusing on resolution only—see Razborov’s beautiful survey paper [27] for a discussion of upper and lower bounds on PHP formulas in other proof systems.

1.1 Previous Work

In a breakthrough result, which served as a strong impetus for further developments in proof complexity, Haken [20] proved a lower bound $\exp(\Omega(n))$ on resolution proof length for $m = n + 1$ pigeons. Haken’s proof was for the basic PHP formulas, but easily extends to onto-FPHP formulas. This result was simplified and improved in a sequence of works [7, 8, 12, 33] to a lower bound of the form $\exp(n^2/m)$, which, unfortunately, does not yield anything nontrivial for $m = \Omega(n^2)$ pigeons.

Buss and Pitassi [11] showed that the pigeonhole principle does in fact get easier for resolution when m becomes sufficiently large: namely, for $m = \exp(\Omega(\sqrt{n \log n}))$ PHP formulas can be refuted in length $\exp(O(\sqrt{n \log n}))$. This is in contrast to what holds for the weaker

subsystem *tree-like resolution*, for which the formulas remain equally hard as the number of pigeons increases, and where the complexity was even sharpened in [9, 11, 15, 17] to an $\exp(\Omega(n \log n))$ length lower bound.

Obtaining lower bounds beyond $m = n^2$ pigeons for non-tree-like resolution turned out to be quite challenging. Haken’s bottleneck counting method fundamentally breaks down when the number of pigeons is quadratic in the number of holes, and the same holds for the celebrated length-width lower bound in [8]. Some progress was made for restricted forms of resolution in [31] and [23], leading up to an $\exp(n^\epsilon)$ lower bound for so-called *regular resolution*. In a technical tour de force, Raz [24] finally proved that general, unrestricted resolution requires length $\exp(n^\epsilon)$ to refute the basic PHP formulas even with arbitrary many pigeons. Razborov followed up on this in three papers where he first simplified and slightly strengthened Raz’s result in [25], then extended it to FPHP formulas in [30] and lastly established an analogous lower bound for onto-FPHP formulas in [29].

More precisely, what Razborov showed is that for any version of the PHP formula with m pigeons and n holes, the minimal proof length required in resolution is $\exp(\Omega(n/\log^2 m))$. It is easy to see that this implies a lower bound $\exp(\Omega(\sqrt[3]{n}))$ for any number of pigeons—for $m = \exp(O(\sqrt[3]{n}))$ we can appeal directly to the bound above, and if a resolution proof would use $\exp(\Omega(\sqrt[3]{n}))$ pigeons, then just mentioning all these different pigeons already requires $\exp(\Omega(\sqrt[3]{n}))$ distinct clauses. It is also clear that considering complexity in terms of the number of holes n is the right measure. Since any formula contains a basic PHP subformula with $n + 1$ pigeons that can be refuted in length $\exp(O(n))$, we can never hope for exponential lower bounds in terms of formula size as the number of pigeons m grows to exponential.

So far we have stated results only for the standard PHP formulas over $K_{m,n}$, where any pigeon can fly to any hole. However, the way Ben-Sasson and Wigderson [8] obtained their result was by considering graph PHP formulas over balanced bipartite expander graphs of constant left degree, from which the lower bound for $K_{m,n}$ easily follows by a restriction argument. It was shown in [21] that an analogous bound holds for onto-FPHP formulas, i.e., perfect matching formulas, on bipartite expanders. In this context is also relevant to mention the exponential lower bounds in [1, 16] on *mutilated chessboard formulas*, which can be viewed as perfect matching formulas on balanced, sparse bipartite graphs with very bad expansion. At the other end of the spectrum, Razborov’s PHP lower bound in [29] for highly unbalanced bipartite graphs also applies in a more general setting than $K_{m,n}$: namely, for any graph where the minimal degree of any left vertex is δ , the minimal length of any resolution proof is $\exp(\Omega(\delta/\log^2 m))$. Thus, for graph PHP formulas we have exponential lower bounds on the one hand [8] for $m \ll n^2$ pigeons, where each pigeon is adjacent to a constant number of holes, and on the other hand [29] for any number of pigeons given that each pigeon is adjacent to a polynomial $n^{\Omega(1)}$ number of holes, but nothing has been known in between these extremes. In [29], Razborov asks whether a “*common generalization*” of the techniques in [8] and [29, 30] can be found

“that would uniformly cover both cases?” Urquhart [34] also discusses Razborov’s lower bound technique, but notes that “the search for a yet more general point of view remains a topic for further research.”

1.2 Our Results

In this work, we give an answer to the questions raised in [29, 34] by presenting a general technique that applies for any number of pigeons m all the way from linear to weakly exponential, and that establishes exponential lower bounds on resolution proof length for all flavours of graph PHP formulas (including perfect matching formulas) even over sparse graphs.

Let us state below three examples of the kind of lower bounds we obtain—the full, formal statements will follow in later sections. Our first theorem is an average-case lower bound for onto-FPHP formulas with slightly superpolynomial number of pigeons.

THEOREM 1.1 (Informal). *Let G be a randomly sampled bipartite graph with n right vertices, $m = n^{o(\log n)}$ left vertices, and left degree $\Theta(\log^2 m)$. Then refuting the onto-FPHP formula (a.k.a. perfect matching formula) over G in resolution requires length $\exp(\Omega(n^{1-o(1)}))$ asymptotically almost surely.*

Note that as the number of pigeons grow larger, it is clear that the left degree also has to grow—otherwise we will get a small number of pigeons constrained to fly to a small number of holes by a birthday paradox argument, yielding a small unsatisfiable subformula that can easily be refuted by brute force.

If the number of pigeons increases further to weakly exponential, then randomly sampled graphs no longer have good enough expansion for our technique to work, but there are explicit constructions of unbalanced expanders for which we can still get lower bounds.

THEOREM 1.2 (Informal). *There are explicitly constructible bipartite graphs G with n right vertices, $m = \exp(O(n^{1/16}))$ left vertices, and left degree $\Theta(\log^4 m)$ such that resolution requires length $\exp(\Omega(n^{1/8-\varepsilon}))$ to refute the perfect matching formula over G .*

Finally, for functional pigeonhole principle formulas we can also prove an exponential lower bound for *constant* left degree even if the number of pigeons is a large polynomial.

THEOREM 1.3 (Informal). *Let G be a randomly sampled bipartite graph with n right vertices, $m = n^k$ left vertices, and left degree $\Theta((k/\varepsilon)^2)$. Then refuting the functional pigeonhole principle formula over G in resolution requires length $\exp(\Omega(n^{1-\varepsilon}))$ asymptotically almost surely.*

1.3 Techniques

At a very high level, what we do in terms of techniques is to revisit the pseudo-width method introduced by Razborov for functional PHP formulas in [30]. We strengthen this method to

work in the setting of sparse graphs by combining it with the closure operation on expander graphs in [2, 4], which is a way to restore expansion after a small set of (potentially adversarially chosen) vertices have been removed. To extend the results further to perfect matching formulas, we apply a “preprocessing step” on the formulas as in [29]. In what remains of this section, we focus on graph FPHP formulas and give an informal overview of the lower bound proof in this setting, which already contains most of the interesting ideas (although the extension to onto-FPHP also raises significant additional challenges).

Let $FPHP(G)$ denote the functional pigeonhole principle formula over the graph G consisting of clauses (1.1a)–(1.1c). A first, quite naive (and incorrect), description of the proof structure is that we start by defining a *pseudo-width* measure on clauses C that counts pigeons i that appear in C in many variables x_{ij} for distinct j . We then show that any short resolution refutation of $FPHP(G)$ can be transformed into a refutation where all clauses have small pseudo-width. By a separate argument, we establish that any refutation of $FPHP(G)$ requires large pseudo-width. Hence, no short refutations can exist, which is precisely what we were aiming to prove.

To fill in the details (and correct) this argument, let us start by making clear what we mean by pseudo-width. Suppose that the graph G has left degree Δ . In what follows, we identify a mapping of pigeon i to a neighbouring hole j with the partial assignment ρ such that $\rho(x_{i,j}) = 1$ and $\rho(x_{i,j'}) = 0$ for all $j' \in N(i) \setminus \{j\}$. We denote by $d_i(C)$ the number of mappings of pigeon i that satisfy C . Note that if C contains at least one negated literal $\bar{x}_{i,j}$, then $d_i(C) \geq \Delta - 1$, and otherwise $d_i(C)$ is the number of positive literals $x_{i,j}$ for $j \in N(i)$. Given a judiciously chosen “filter vector” $\vec{d} = (d_1, \dots, d_m)$ for $d_i \approx \Delta$ and a “slack” $\delta \approx \Delta/\log m$, we say that pigeon i is *heavy* in C if $d_i(C) \geq d_i - \delta$ and *super-heavy* if $d_i(C) \geq d_i$. We define the *pseudo-width* of a clause C to be the number of heavy pigeons in C .

With these definitions in hand, we can give a description of the actual proof:

1. Given any resolution refutation π of $FPHP(G)$ in small length L , we argue that all clauses can be classified as having either low or high pseudo-width, where an important additional guarantee is that the high-width clauses not only have many heavy pigeons but actually many super-heavy pigeons.
2. We replace all clauses C with many super-heavy pigeons with “fake axioms” $C' \subseteq C$ obtained by throwing away literals from C until we have nothing left but a medium number of super-heavy pigeons. By construction, the set \mathcal{A} of such fake axioms is of size $|\mathcal{A}| \leq L$, and after making the replacement we have a resolution refutation π' of $FPHP(G) \cup \mathcal{A}$ in low pseudo-width.
3. However, since \mathcal{A} is not too large, we are able to show that any resolution refutation of $FPHP(G) \cup \mathcal{A}$ must still require large pseudo-width. Hence, L cannot be small, and the lower bound follows.

Part 1 is similar to [30], but with a slight twist. We show that if the length of π is $L < 2^{w_0}$ and if we choose $\delta \leq \varepsilon \Delta \log n / \log m$, then there exists a vector $\vec{d} = (d_1, \dots, d_m)$ such that for all

clauses in π either the number of super-heavy pigeons is at least w_0 or else the number of heavy pigeons is at most $O(w_0 \cdot n^\varepsilon)$. The proof of this is by sampling the coordinates d_i independently from a suitable probability distribution and then applying a union bound argument. Once this has been established, part 2 follows easily: we just replace all clauses with at least w_0 super-heavy pigeons by (stronger) fake axioms. Including all fake axioms \mathcal{A} yields a refutation π' of $FPHP(G) \cup \mathcal{A}$ (since we can add a weakening rule deriving C from $C' \subseteq C$ to resolution without loss of generality) and clearly all clauses in π' have pseudo-width $O(w_0 \cdot n^\varepsilon)$.

Part 3 is where most of the hard work is. Suppose that G is an excellent expander graph, so that for some value r all left vertex sets U' of size $|U'| \leq r$ have at least $(1 - \varepsilon \log n / \log m) \Delta |U'|$ unique neighbours on the right-hand side. We show that, under the assumptions above, refuting $FPHP(G) \cup \mathcal{A}$ requires pseudo-width $\Omega(r \cdot \log n / \log m)$. Tuning the parameters appropriately, this yields a contradiction with part 2.

Before outlining how the proof of part 3 goes, we remark that the requirements we place on the expansion of G are quite severe. Clearly, any left vertex set U can have at most $\Delta |U|$ neighbours in total, and we are asking for all except a vanishingly small fraction of these neighbours to be unique. This is why we can establish Theorem 1.1 but not Theorem 1.2 for randomly sampled graphs. We see no reason to believe that the latter theorem would not hold also for random graphs, but the expansion properties required for our proof are so stringent that they are not satisfied in this parameter regime. This seems to be a fundamental shortcoming of our technique, and it appears that new ideas would be required to circumvent this problem.

In order to argue that refuting $FPHP(G) \cup \mathcal{A}$ in resolution requires large pseudo-width, we want to estimate how much progress the resolution derivation has made up to the point when it derives some clause C . Following Razborov's lead, we measure this by looking at what fraction of partial matchings of all the heavy pigeons in C do not satisfy C (meaning, intuitively, that the derivation has managed to rule out this part of the search space). It is immediate by inspection that all pigeons mentioned in the real axiom clauses (1.1a)–(1.1c) are heavy, and any matching of such pigeons satisfies the clauses. Thus, the original axioms in $FPHP(G)$ do not rule out any matchings. Also, it is easy to show that fake axioms rule out only an exponentially small fraction of matchings, since they contain many super-heavy pigeons and it is hard to match all of these pigeons without satisfying the clause. However, the contradictory empty clause \perp rules out 100% of partial matchings, since it contains no heavy pigeons to match in the first place.

What we would like to prove now is that for any derivation in small pseudo-width it holds that the derived clause cannot rule out any matching other than those already eliminated by the clauses used to derive it. This means that the fake axioms together need to rule out all partial matchings, but since every fake axiom contributes only an exponentially small fraction they are too few to achieve this. Hence, it is not possible to derive contradiction in small pseudo-width, which completes part 3 of our proof outline.

There is one problem, however: the last claim above is not true, and so what is outlined above is only a fake proof. While we have to defer the discussion of what the full proof actually looks like in detail, we conclude this section by attempting to hint at a couple of technical issues and how to resolve them.

Firstly, it does not hold that a derived clause C eliminates only those matchings that are also forbidden by one of the predecessor clauses used to derive C . The issue is that a pigeon i that is heavy in both predecessors might cease to be heavy in C —for instance, if C was derived by a resolution step over a variable $x_{i,j}$. If this is so, then we would need to show that any matching of the heavy pigeons in C can be extended to match also pigeon i to any of its neighbouring holes without satisfying both predecessor clauses. But this will not be true, because a non-heavy pigeon can still have some variable $x_{i,j}$ occurring in both predecessors. The solution to this, introduced in [30], is to do a “lossy counting” of matchings by associating each partial matching with a linear subspace of some suitable vector space, and then to consider the span of all matchings ruled out by C . When we accumulate a “large enough” number of matchings for a pigeon i , then the whole subspace associated to i is spanned and we can stop counting.

But this leads to a second problem: when studying matchings of the heavy pigeons in C we might already have assigned pigeons i'_1, \dots, i'_w to occupy holes where pigeon i might want to fly. For standard PHP formulas over complete bipartite graphs this is not a problem, since at least $n - w$ holes are still available and this number is “large enough” in the sense described above. But for a sparse graph it will typically be the case that $w \gg \Delta$, and so it might well be the case that pigeons i'_1, \dots, i'_w are already occupying all the Δ holes available for pigeon i according to G . Although it is perhaps hard to see from our (admittedly somewhat informal) discussion, this turns out to be a very serious problem, and indeed it is one of the main technical challenges we need to overcome.

To address this problem we consider not only the heavy pigeons in C , but also any other pigeons in G that risk becoming far too constrained when the heavy pigeons of C are matched. Inspired by [2, 4], we define the *closure* to be a superset S of the heavy pigeons such that when S and the neighbouring holes of S are removed it holds that the residual graph is still guaranteed to be a good expander. Provided that G is an excellent expander to begin with, and that the number of heavy pigeons in C is not too large, it can then be shown that an analogue of the original argument outlined above goes through.

1.4 Outline of This Paper

We review the necessary preliminaries in Section 2 and introduce two crucial technical tools in Section 3. The lower bounds for weak graph FPHP formulas are then presented in Section 4, after which the perfect matching lower bounds follow in Section 5. We conclude with a discussion of questions for future research in Section 6.

2. Preliminaries

We denote natural logarithms (base e) by \ln , and base 2 logarithms by \log . For positive integers $n \in \mathbb{N}^+$ we write $[n] = \{1, \dots, n\}$.

2.1 Proof Complexity

A *literal* over a Boolean variable x is either the variable x itself (a *positive literal*) or its negation \bar{x} (a *negative literal*). A *clause* $C = \ell_1 \vee \dots \vee \ell_w$ is a disjunction of literals. We write \perp to denote the empty clause without any literals. A *CNF formula* $F = C_1 \wedge \dots \wedge C_m$ is a conjunction of clauses. We think of clauses and CNF formulas as sets: order is irrelevant and there are no repetitions. We let $\text{Vars}(F)$ denote the set of variables of F .

A *resolution refutation* π of an unsatisfiable CNF formula F , or *resolution proof* for (the unsatisfiability of) F , is an ordered sequence of clauses $\pi = (D_1, \dots, D_L)$ such that $D_L = \perp$ and for each $i \in [L]$ either D_i is a clause in F (an *axiom*) or there exist $j < i$ and $k < i$ such that D_i is derived from D_j and D_k by the *resolution rule*

$$\frac{B \vee x \quad C \vee \bar{x}}{B \vee C} . \quad (2.1)$$

We refer to $B \vee C$ as the *resolvent* of $B \vee x$ and $C \vee \bar{x}$ over x , and to x as the *resolved variable*. For technical reasons it is sometimes convenient to also allow clauses to be derived by the *weakening rule*

$$\frac{C}{D} [C \subseteq D] \quad (2.2)$$

(and for two clauses $C \subseteq D$ we will sometimes refer to C as a *strengthening* of D).

The *length* $L(\pi)$ of a refutation $\pi = (D_1, \dots, D_L)$ is L and the length of refuting F is $\min_{\pi: F \vdash \perp} \{L(\pi)\}$, where the minimum is taken over all resolution refutations π of F . It is easy to show that removing the weakening rule (2.2) does not increase the refutation length.

A *partial assignment* or a *restriction* on a formula F is a partial function $\rho : \text{Vars}(F) \rightarrow \{0, 1\}$. The clause C *restricted by* ρ , denoted $C \upharpoonright_\rho$, is the trivial 1-clause if any of the literals in C is satisfied by ρ and otherwise it is C with all falsified literals removed. We extend this definition to CNF formulas in the obvious way by taking unions. For a variable $x \in \text{Vars}(F)$ we write $\rho(x) = *$ if $x \notin \text{dom}(\rho)$, i.e., if ρ does not assign a value to x .

2.2 Concentration Inequalities

We often denote random variables in boldface and write $\mathbf{X} \sim \mathcal{D}$ to denote that \mathbf{X} is sampled from the distribution \mathcal{D} . We will use the following standard forms of the multiplicative Chernoff bounds.

THEOREM 2.1. *Let \mathbf{S} be the sum of independent 0-1 random variables (not necessarily equidistributed) with expectation $\mu = \mathbb{E}[\mathbf{S}]$. Then for $\delta \geq 0$ it holds that*

$$\Pr[\mu - \mathbf{S} \geq \delta] \leq \exp\left(-\frac{\delta^2}{2\mu}\right) \text{ and } \Pr[\mathbf{S} - \mu \geq \delta] \leq \exp\left(-\frac{\delta^2}{2\mu + \delta}\right).$$

2.3 Graph Theory

We write $G = (V, E)$ to denote a graph with vertices V and edges E , where G is always undirected and without loops or multiple edges. Moreover, for bipartite graphs we write $G = (U \dot{\cup} V, E)$, where edges in E have one endpoint in the left vertex set U and the other in the right vertex set V . A *partial matching* φ in G is a subset of edges that are vertex-disjoint. Let $V(\varphi) = \{v \mid \exists e \in \varphi : v \in e\}$ be the vertices of φ and for $v \in V(\varphi)$ denote by φ_v the unique vertex u such that $\{u, v\} \in \varphi$. A vertex v is *covered* by φ if $v \in V(\varphi)$. If φ is a partial matching in a bipartite graph $G = (U \dot{\cup} V, E)$, we identify it with a partial mapping of U to V . When referring to the pigeonhole formula, this mapping will also be identified with an assignment ρ_φ to the variables defined by

$$\rho_\varphi(x_{i,j}) = \begin{cases} * & \text{if } i \notin \text{dom}(\varphi), \\ 0 & \text{if } i \in \text{dom}(\varphi) \text{ and } \varphi(i) \neq j, \\ 1 & \text{if } i \in \text{dom}(\varphi) \text{ and } \varphi(i) = j. \end{cases} \quad (2.3)$$

For simplicity we interchangeably think of partial matchings as partial functions and as Boolean assignments as defined in (2.3).

Given a vertex $v \in V(G)$, we write $N_G(v)$ to denote the set of *neighbours of v* in the graph G and $\Delta_G(v) = |N_G(v)|$ to denote the degree of v . We extend this notion to sets and denote by $N_G(S) = \{v \mid \exists (u, v) \in E \text{ for } u \in S\}$ the *neighbourhood* of a set of vertices $S \subseteq V$. The *boundary*, or *unique neighbourhood*, $\partial_G(S) = \{v \in V \setminus S : |N_G(v) \cap S| = 1\}$ of a set of vertices $S \subseteq V$ contains all vertices in $V \setminus S$ that have a single neighbour in S . If the graph is bipartite, there is of course no need to subtract S from the neighbour set. We will sometimes drop the subscript G when the graph is clear from context. For a set $U \subseteq V$ we denote by $G \setminus U$ the subgraph of G induced by the vertex set $V \setminus U$.

A graph $G = (V, E)$ is an (r, Δ, c) -*expander* if all vertices $v \in V$ have degree at most Δ and for all sets $S \subseteq V$, $|S| \leq r$, it holds that $|N(S) \setminus S| \geq c \cdot |S|$. Similarly, $G = (V, E)$ is an (r, Δ, c) -*boundary expander* if all vertices $v \in V$ have degree at most Δ and for all sets $S \subseteq V$, $|S| \leq r$, it holds that $|\partial(S)| \geq c \cdot |S|$. For bipartite graphs, the degree and expansion requirements only apply to the left vertex set: $G = (U \dot{\cup} V, E)$ is an (r, Δ, c) -*bipartite expander* if all vertices $u \in U$ have degree at most Δ and for all sets $S \subseteq U$, $|S| \leq r$, it holds that $|N(S)| \geq c \cdot |S|$, and an (r, Δ, c) -*bipartite boundary expander* if for all sets $S \subseteq U$, $|S| \leq r$, it holds that $|\partial(S)| \geq c \cdot |S|$. For bipartite graphs we will only ever be interested in bipartite notions of expansions, and so which kind of

expansion is meant will always be clear from context. A simple but useful observation is that

$$|N(S) \setminus S| \leq |\partial(S)| + \frac{\Delta|S| - |\partial(S)|}{2} = \frac{\Delta|S| + |\partial(S)|}{2}, \quad (2.4)$$

since all non-unique neighbours in $N(S) \setminus S$ have at least two incident edges. This implies that if a graph G is an $(r, \Delta, (1 - \xi)\Delta)$ -expander then it is also an $(r, \Delta, (1 - 2\xi)\Delta)$ -boundary expander.

For $n, m, \Delta \in \mathbb{N}$, we denote by $\mathcal{G}(m, n, \Delta)$ the distribution over bipartite graphs with disjoint vertex sets $U = \{u_1, \dots, u_m\}$ and $V = \{v_1, \dots, v_n\}$ where the neighbourhood of a vertex $u \in U$ is chosen by sampling a subset of size Δ uniformly at random from V . A property is said to hold *asymptotically almost surely* on $\mathcal{G}(f(n), n, \Delta)$ if it holds with probability that approaches 1 as n approaches infinity.

For the right parameters, a randomly sampled graph $G \sim \mathcal{G}(m, n, \Delta)$ is asymptotically almost surely a good boundary expander as stated next.

LEMMA 2.2. *Let m, n and Δ be large enough integers such that $m > n \geq \Delta$. Let $\xi, \chi \in \mathbb{R}^+$ be such that $\xi < 1/2$, $\xi \ln \chi \geq 2$ and $\xi \Delta \ln \chi \geq 4 \ln m$. Then for $r = n/(\Delta \cdot \chi)$ and $c = (1 - 2\xi)\Delta$ it holds asymptotically almost surely for a randomly sampled graph $G \sim \mathcal{G}(m, n, \Delta)$ that G is an (r, Δ, c) -boundary expander.*

PROOF. Let $G = (U \dot{\cup} V, E)$. We first estimate the probability that a set $S \subseteq U$ of size at most r violates the boundary expansion. For brevity, let us write $s = |S|$ and $c' = (1 - \xi)\Delta$. In view of (2.4), the probability that S violates the boundary expansion can be bounded by

$$\Pr[|\partial(S)| < cs] \leq \Pr\left[|N(S)| < \frac{\Delta s + cs}{2}\right] \quad (2.5a)$$

$$= \Pr[|N(S)| < c's] \quad (2.5b)$$

$$\leq \binom{n}{c's} \cdot \left(\frac{\binom{c's}{\Delta}}{\binom{n}{\Delta}}\right)^s \quad (2.5c)$$

$$\leq \binom{n}{c's} \cdot \left(\frac{c's}{n}\right)^{\Delta s} \quad (2.5d)$$

$$\leq \left[\left(\frac{en}{c's}\right)^{c'} \cdot \left(\frac{c's}{n}\right)^\Delta\right]^s \quad (2.5e)$$

$$= \left[e^{(1-\xi)\Delta} \cdot \left(\frac{n}{c's}\right)^{-\xi\Delta}\right]^s \quad (2.5f)$$

$$\leq \exp\left(\Delta s \left(1 - \xi \ln\left(\frac{n}{c's}\right)\right)\right) \quad (2.5g)$$

$$\leq \exp\left(\Delta s \left(1 - \xi \ln\left(\frac{\chi}{1-\xi}\right)\right)\right) \quad (2.5h)$$

$$\leq \exp(\Delta s(1 - \xi \ln \chi)) \quad (2.5i)$$

$$\leq \exp(-(\Delta s \xi \ln \chi)/2), \quad (2.5j)$$

where (2.5h) holds since $s \leq r \leq n/(\Delta\chi)$ and (2.5j) holds since $\xi \ln \chi \geq 2$. Hence, the probability that G is not a boundary expander can be bounded by

$$\begin{aligned} \Pr[G \text{ is not an expander}] &\leq \sum_{s \in [r]} \binom{m}{s} \exp(-(\Delta s \xi \ln \chi)/2) \\ &\leq \sum_{s \in [r]} \exp(-s((\xi \Delta \ln \chi)/2 - \ln m)) \\ &\leq \sum_{s \in [r]} \exp(-s \ln m) \leq \frac{1}{m-1}, \end{aligned} \quad (2.6)$$

where the second-to-last inequality holds since $\xi \Delta \ln \chi \geq 4 \ln m$. ■

We will also consider some parameter settings where randomly sampled graphs do not have strong enough expansion for our purposes, but where we can resort to explicit constructions as follows.

THEOREM 2.3 ([19]). *For all positive integers m , $r \leq m$, all $\xi > 0$, and all constant $\nu > 0$, there is an explicit $(r, \Delta, (1 - \xi)\Delta)$ -expander $G = (U \dot{\cup} V, E)$, with $|U| = m$, $|V| = n$, $\Delta = O((\log m)(\log r)/\xi)^{1+1/\nu}$ and $n \leq \Delta^2 \cdot r^{1+\nu}$.*

COROLLARY 2.4. *Let κ, ε, ν be positive constants, $\kappa < \frac{1}{8}$, and let n be a large enough integer. Then there is an explicit graph $G = (U \dot{\cup} V, E)$, with $|U| = m = 2^{\Omega(n^\kappa)}$ and $|V| \leq n$, that is an $(n^{\frac{1}{1+\nu} - \frac{4\kappa}{\nu}}, \Delta, (1 - 2\xi)\Delta)$ -boundary expander for $\xi = \frac{\varepsilon \log n}{\log m}$ and $\Delta = O(\log^{2(1+1/\nu)} m)$.*

PROOF. Let G be the expander from Theorem 2.3 for the parameters $m = 2^{\varepsilon' n^\kappa}$, $r = n^{\frac{1}{1+\nu} - \frac{4\kappa}{\nu}}$, and $\xi = \frac{\varepsilon \log n}{\log m}$, where ε' is chosen to be a small enough constant so that $\Delta^2 \cdot r^{1+\nu} \leq n$. Such a graph G is an $(r, \Delta, (1 - \xi)\Delta)$ -expander for Δ as in the Corollary. By (2.4) it follows that an (r, Δ, c) -expander is an $(r, \Delta, 2c - \Delta)$ -boundary expander, and hence G is an $(r, \Delta, (1 - 2\xi)\Delta)$ -boundary expander. Note that Theorem 2.3 guarantees that the right side of G has size at most $\Delta^2 \cdot r^{1+\nu} \leq n$. ■

3. Two Key Technical Tools

In this section we review two crucial technical ingredients of the resolution lower bound proofs.

3.1 Pigeon Filtering

The following lemma is a generalization of [30, Lemma 6] and states that for every small collection $\vec{r}(1), \dots, \vec{r}(L)$ of integer vectors there is a vector \vec{r} such that each vector $\vec{r}(\ell)$ has either many coordinates i smaller than \vec{r} (i.e., $r_i(\ell) \leq r_i$), or has few coordinates i which are by 1 larger than \vec{r} , that is, there are few coordinates i such that $r_i(\ell) \leq r_i + 1$.

The proof follows by sampling each coordinate of \vec{r} independently from a geometric probability distribution and applying a union bound over the vectors $\vec{r}(1), \dots, \vec{r}(L)$. The difference

to [30] is that our geometric distribution depends on an additional parameter α (which is implicitly fixed to $\alpha = 2$ in [30]) that allows us to get a better upper bound on the numbers r_i . This turns out to be crucial for us—we discuss this in more detail in Section 4 where we explain how the lemma relates to the weak pigeonhole principle.

LEMMA 3.1 (Filter lemma). *Let $m, L \in \mathbb{N}^+$ and suppose that $w_0, \alpha \in [m]$ are such that $w_0 > \ln L$ and $w_0 \geq \alpha^2 \geq 4$. Further, let $\vec{r}(1), \dots, \vec{r}(L)$ be integer vectors, each of the form $\vec{r}(\ell) = (r_1(\ell), \dots, r_m(\ell))$. Then there exists a vector $\vec{r} = (r_1, \dots, r_m)$ of positive integers $r_i \leq \lfloor \frac{\log m}{\log \alpha} \rfloor - 1$ such that for all $\ell \in [L]$ at least one of the following holds:*

1. $|\{i \in [m] : r_i(\ell) \leq r_i\}| \geq w_0$,
2. $|\{i \in [m] : r_i(\ell) \leq r_i + 1\}| \leq O(\alpha \cdot w_0)$.

PROOF. We first define a weight function $W(\vec{r})$ for vectors $\vec{r} = (r_1, \dots, r_m)$ as

$$W(\vec{r}) = \sum_{i \in [m]} \alpha^{-r_i} . \quad (3.1)$$

In order to establish the lemma, it is sufficient to show that there exist constants γ and γ' and a vector $r = (r_1, \dots, r_m)$ such that for all $\ell \in [L]$ the implications

$$W(\vec{r}(\ell)) \geq \frac{\gamma' w_0}{\alpha} \Rightarrow |\{i \in [m] \mid r_i(\ell) \leq r_i\}| \geq w_0 , \quad (3.2a)$$

$$W(\vec{r}(\ell)) \leq \frac{\gamma' w_0}{\alpha} \Rightarrow |\{i \in [m] \mid r_i(\ell) \leq r_i + 1\}| \leq \gamma \alpha w_0 \quad (3.2b)$$

hold. Let $t = \lfloor \frac{\log m}{\log \alpha} \rfloor - 1$ and let μ be a probability distribution on $[t]$ given by $\Pr[\mathbf{r} = i] = \beta \cdot \alpha^{-i}$ for all $i \in [t]$, where $\beta = \frac{\alpha-1}{1-\alpha^{-t}}$. Note that

$$\beta \sum_{i \in [t]} \alpha^{-i} = \frac{\alpha-1}{1-\alpha^{-t}} \left(\frac{1-\alpha^{-t}}{\alpha-1} \right) = 1 \quad (3.3)$$

and thus μ is a valid distribution. Let us write $\vec{\mathbf{r}} = (\mathbf{r}_1, \dots, \mathbf{r}_m)$ to denote a random vector with coordinates sampled independently according to μ . We claim that for every $\ell \in [L]$ the implications (3.2a) and (3.2b) are true asymptotically almost surely. Let us proceed to verify this.

1. Suppose that $W(\vec{r}(\ell)) \geq \frac{\gamma' w_0}{\alpha}$. We wish to show that $|\{i \in [m] : r_i \geq r_i(\ell)\}| \geq w_0$. Observe that coordinates larger than t contribute only

$$\sum_{r_i(\ell) > t} \alpha^{-r_i(\ell)} \leq m \cdot \alpha^{-t-1} < \alpha \quad (3.4)$$

to $W(\vec{r}(\ell))$, and hence the weight function truncated at t is

$$\sum_{r_i(\ell) \leq t} \alpha^{-r_i(\ell)} \geq \frac{\gamma' w_0}{\alpha} - \alpha \geq (\gamma' - 1) \frac{w_0}{\alpha} , \quad (3.5)$$

since $w_0 \geq \alpha^2$. Note that for every coordinate i with $r_i(\ell) \leq t$ we have that $\Pr[\mathbf{r}_i \geq r_i(\ell)] \geq \beta \cdot \alpha^{-r_i(\ell)}$. Consider the random set $P_{\vec{r}}(\ell) = \{i \in [m] \mid r_i(\ell) \leq t \text{ and } \mathbf{r}_i \geq r_i(\ell)\}$. We can appeal to (3.5) to derive that

$$\begin{aligned} \mathbb{E}[|P_{\vec{r}}(\ell)|] &= \sum_{r_i(\ell) \leq t} \Pr[\mathbf{r}_i \geq r_i(\ell)] \\ &\geq \sum_{r_i(\ell) \leq t} \beta \alpha^{-r_i(\ell)} \\ &\geq \beta(\gamma' - 1) \frac{w_0}{\alpha} \geq \frac{\gamma' - 1}{2} w_0 \end{aligned} \quad (3.6)$$

is a lower bound on the expected size of $P_{\vec{r}}(\ell)$. As the events $\mathbf{r}_i \geq r_i(\ell)$ are independent, by the multiplicative Chernoff bound, that is, Theorem 2.1, we get that

$$\Pr[|P_{\vec{r}}(\ell)| < w_0] \leq \Pr[|P_{\vec{r}}(\ell)| - \mathbb{E}[|P_{\vec{r}}(\ell)|] \leq w_0 - \mathbb{E}[|P_{\vec{r}}(\ell)|]] \quad (3.7a)$$

$$= \Pr[\mathbb{E}[|P_{\vec{r}}(\ell)|] - |P_{\vec{r}}(\ell)| \geq \mathbb{E}[|P_{\vec{r}}(\ell)|] - w_0] \quad (3.7b)$$

$$\leq \exp\left(-\frac{(\mathbb{E}[|P_{\vec{r}}(\ell)|] - w_0)^2}{2 \mathbb{E}[|P_{\vec{r}}(\ell)|]}\right) \quad (3.7c)$$

$$= \exp\left(-\frac{\mathbb{E}[|P_{\vec{r}}(\ell)|]^2 - 2 \mathbb{E}[|P_{\vec{r}}(\ell)|] w_0 + w_0^2}{2 \mathbb{E}[|P_{\vec{r}}(\ell)|]}\right) \quad (3.7d)$$

$$\leq \exp\left(-\frac{\mathbb{E}[|P_{\vec{r}}(\ell)|] - 2w_0}{2}\right) \quad (3.7e)$$

$$\leq \exp\left(-\frac{(\gamma' - 5)}{4} w_0\right) \quad (3.7f)$$

$$\leq \exp(-2w_0) \quad (3.7g)$$

$$\leq L^{-2}, \quad (3.7h)$$

where the second to last inequality holds for $\gamma' \geq 13$.

2. Suppose that $W(\vec{r}(\ell)) \leq \frac{\gamma' w_0}{\alpha}$. Now we need to show that $|\{i \in [m] : r_i \geq r_i(\ell) - 1\}| \leq \gamma \alpha w_0$ holds asymptotically almost surely. Note that

$$\Pr[\mathbf{r}_i \geq r_i(\ell) - 1] = \beta \sum_{j=r_i(\ell)-1}^t \alpha^{-j} \quad (3.8a)$$

$$= \frac{\alpha - 1}{1 - \alpha^{-t}} \left(\frac{\alpha^{-r_i(\ell)+2} - \alpha^{-t}}{\alpha - 1} \right) \quad (3.8b)$$

$$= \frac{\alpha^{-r_i(\ell)+2} - \alpha^{-t}}{1 - \alpha^{-t}} \quad (3.8c)$$

$$= \frac{\alpha^{t-r_i(\ell)+2} - 1}{\alpha^t - 1} \quad (3.8d)$$

$$\leq \frac{\alpha^{t-r_i(\ell)+2}}{\alpha^t/2} = 2\alpha^{2-r_i(\ell)}. \quad (3.8e)$$

Similar to the previous case, let $Q_{\vec{r}}(\ell) = \{i \in [m] \mid \mathbf{r}_i \geq r_i(\ell) - 1\}$. We can upper-bound the expected cardinality of $Q_{\vec{r}}(\ell)$ by

$$\mathbb{E}[|Q_{\vec{r}}(\ell)|] = \sum_{i \in [m]} \Pr[\mathbf{r}_i \geq r_i(\ell) - 1] \leq 2\alpha^2 \mathcal{W}(\vec{r}(\ell)) \leq 2\gamma' \alpha w_0. \quad (3.9)$$

Again, we apply the Chernoff bound in Theorem 2.1 and conclude that

$$\begin{aligned} \Pr[|Q_{\vec{r}}(\ell)| \geq \gamma \alpha w_0] &\leq \Pr[|Q_{\vec{r}}(\ell)| - \mathbb{E}[|Q_{\vec{r}}(\ell)|] \geq \gamma \alpha w_0 - 2\gamma' \alpha w_0] \\ &\leq \exp\left(-\frac{(\gamma - 2\gamma')^2 (\alpha w_0)^2}{4\gamma' \alpha w_0 + (\gamma - 2\gamma') \alpha w_0}\right) \\ &\leq \exp(-\alpha w_0) \\ &\leq L^{-2} \end{aligned} \quad (3.10)$$

where the second to last inequality holds for γ sufficiently larger than γ' , say $\gamma \geq 5\gamma'$.

A union bound argument over all vectors in $\{\vec{r}(\ell) : \ell \in [L]\}$ for both cases shows that for $\gamma' \geq 13$ and $\gamma \geq 5\gamma'$ there exists a choice of $\vec{r} = (r_1, \dots, r_m)$ such that both implications (3.2a) and (3.2b) hold. ■

3.2 Graph Closure

A key concept in our work is that of a *closure* of a vertex set, which seems to have originated in [2, 4]. In later arguments we will want to remove a set T of vertices along with its neighbourhood $N(T)$ from an expanding graph G while maintaining some guarantee of expansion. The closure of T is a suitably small set $S \supseteq T$ such that the graph $G \setminus (S \cup N(S))$ is a reasonably good expander; see Lemma 3.4 for the formal guarantees. For intuition, it is useful to think of S as being constructed by an iterative process $T = S_0 \subset S_1 \subset \dots \subset S_t = S$ where $S_i = S_{i-1} \cup S'_i$ for S'_i any small set of vertices in $G \setminus (S_{i-1} \cup N(S_{i-1}))$ that is not sufficiently expanding. It turns out to be more convenient to work with the somewhat technical definition that follows.

In order to have a definition that makes sense for both expanders and bipartite expanders, we define $V_{\text{exp}}(G)$ to be the set of vertices of G that expand, that is, if $G = (V, E)$ is an expander then $V_{\text{exp}}(G) = V$, and if $G = (U \dot{\cup} V, E)$ is a bipartite expander then $V_{\text{exp}}(G) = U$.

DEFINITION 3.2 (Closure). For an expander graph G and vertex sets $S \subseteq V_{\text{exp}}(G)$ and $U \subseteq V(G)$, we say that the set S is (U, r, ν) -contained if $|S| \leq r$ and $|\partial(S) \setminus U| < \nu \cdot |S|$.

For any expander graph G and any set $T \subseteq V_{\text{exp}}(G)$ of size $|T| \leq r$, we let $\text{closure}_{r,\nu}(T)$ denote an arbitrary but fixed maximal set such that $T \subseteq \text{closure}_{r,\nu}(T) \subseteq V_{\text{exp}}(G)$ and $\text{closure}_{r,\nu}(T)$ is $(N(T), r, \nu)$ -contained.

Note that the closure of any set T of size $|T| \leq r$ as defined above does indeed exist, since T itself is $(N(T), r, \nu)$ -contained.

LEMMA 3.3. *Suppose that G is an (r, Δ, c) -boundary expander and that $T \subseteq V_{\text{exp}}(G)$ has size $|T| \leq k \leq r$. Then $|\text{closure}_{r,v}(T)| < \frac{k\Delta}{c-v}$.*

PROOF. By definition we have that $|\partial(\text{closure}_{r,v}(T)) \setminus N(T)| < v \cdot |\text{closure}_{r,v}(T)|$. Furthermore, since $|\text{closure}_{r,v}(T)| \leq r$ by definition, we can use the expansion property of the graph to derive the inequality $|\partial(\text{closure}_{r,v}(T)) \setminus N(T)| \geq |\partial(\text{closure}_{r,v}(T))| - |N(T)| \geq c \cdot |\text{closure}_{r,v}(T)| - k\Delta$. Note that we also use the fact that the neighbourhood of T is of size at most $k\Delta$. The conclusion follows by combining both statements. ■

Suppose G is an excellent boundary expander and that $T \subseteq V_{\text{exp}}(G)$ is not too large. Then Lemma 3.3 shows that the closure of T is not much larger. And if the closure is not too large, then after removing the closure and its neighbourhood from the graph we are still left with a decent expander, a fact which will play a key role in the technical arguments in later sections. The following lemma makes this intuition precise.

LEMMA 3.4. *For G an (r, Δ, c) -boundary expander, let $T \subseteq V_{\text{exp}}(G)$ be such that $|T| \leq r$ and $|\text{closure}_{r,v}(T)| \leq r/2$, let $G' = G \setminus (\text{closure}_{r,v}(T) \cup N(\text{closure}_{r,v}(T)))$ and $V_{\text{exp}}(G') = V_{\text{exp}}(G) \cap V(G')$. Then any set $S \subseteq V_{\text{exp}}(G')$ of size $|S| \leq r/2$ satisfies $|\partial_{G'}(S)| \geq v|S|$.*

PROOF. Suppose the set $S \subseteq V_{\text{exp}}(G')$ is of size $|S| \leq r/2$ and does not satisfy $|\partial_{G'}(S)| \geq v|S|$. Since $\text{closure}_{r,v}(T)$ is also of size at most $r/2$, we have that the set $(\text{closure}_{r,v}(T) \cup S)$ is $(N(T), r, v)$ -contained in the graph G . But this contradicts the maximality of $\text{closure}_{r,v}(T)$. ■

4. Lower Bounds for Weak Graph FPHP Formulas

We now proceed to establish lower bounds on the length of resolution refutations of functional pigeonhole principle formulas defined over bipartite graphs. We write $G = (V_P \dot{\cup} V_H, E)$ to denote the graph over which the formulas are defined and \mathcal{M} to denote the set of partial matchings on G (also viewed as partial mappings of V_P to V_H). Let us start by making more precise some of the technical notions discussed in the introduction (which were originally defined in [25]).

For a clause C and a pigeon i we denote the set of holes j with the property that C is satisfied if i is matched to j by

$$N_C(i) = \{j \in V_H \mid e = \{i, j\} \in E \text{ and } \rho_{\{e\}}(C) = 1\} \quad (4.1)$$

and we define the *i th pigeon degree* $\Delta_C(i)$ of C as

$$\Delta_C(i) = |N_C(i)| . \quad (4.2)$$

We think of a pigeon i with large $\Delta_C(i)$ as a pigeon on which the derivation has not made any significant progress up to the point of deriving C , since the clause rules out very few holes.

The pigeons with high enough pigeon degree in a clause are the *heavy pigeons* of the clause as defined next.

DEFINITION 4.1 (Pigeon weight, pseudo-width and (w_0, \vec{d}) -axioms). Let C be a clause and let $\vec{d} = (d_1, \dots, d_m)$ and $\vec{\delta} = (\delta_1, \dots, \delta_m)$ be two vectors of positive integers such that \vec{d} is elementwise greater than $\vec{\delta}$. We say that pigeon i is \vec{d} -*super-heavy* for C if $\Delta_C(i) \geq d_i$ and that pigeon i is $(\vec{d}, \vec{\delta})$ -*heavy* for C if $\Delta_C(i) \geq d_i - \delta_i$. When \vec{d} and $\vec{\delta}$ are understood from context, which is most often the case, we omit the parameters and just refer to *super-heavy* and *heavy* pigeons. Pigeons that are not heavy are referred to as *light pigeons*. The set of pigeons that are super-heavy for C is denoted by

$$P_{\vec{d}}(C) = \{i \in [m] \mid \Delta_C(i) \geq d_i\}$$

and the set of pigeons that are heavy for C is denoted by

$$P_{\vec{d}, \vec{\delta}}(C) = \{i \in [m] \mid \Delta_C(i) \geq d_i - \delta_i\} .$$

The *pseudo-width* of C is the number of heavy pigeons in C and the pseudo-width of a resolution refutation π , denoted by $w_{\vec{d}, \vec{\delta}}(\pi)$, is $\max_{C \in \pi} w_{\vec{d}, \vec{\delta}}(C)$. Finally, we will refer to clauses C with precisely w_0 super-heavy pigeons, i.e., such that $|P_{\vec{d}}(C)| = w_0$, as (w_0, \vec{d}) -*axioms*.

Note that according to Definition 4.1 super-heavy pigeons are also heavy. Making the connection back to our informal discussion in the introduction, the “fake axioms” mentioned there are nothing other than (w_0, \vec{d}) -axioms.

Now that we have all the notions needed, let us give a detailed proof outline. Given a short resolution refutation π of the formula $FPHP(G)$, we use the Filter lemma (Lemma 3.1) to get a filter vector $\vec{d} = (d_1, \dots, d_m)$ such that each clause either has many super-heavy pigeons or there are not too many heavy pigeons (for an appropriately chosen vector $\vec{\delta}$). Clearly, clauses that fall into the second case of the filter lemma have bounded pseudo-width. On the other hand, clauses in the first case may have very large pseudo-width. In order to obtain a proof of low pseudo-width, these clauses are strengthened to (w_0, \vec{d}) -axioms and added to a special set \mathcal{A} . This then gives a refutation π' that refutes the formula $FPHP(G) \cup \mathcal{A}$ in bounded pseudo-width. The following lemma summarizes the upper bound on pseudo-width that we obtain.

LEMMA 4.2. *Let $G = (V_P \dot{\cup} V_H, E)$ be a bipartite graph with $|V_P| = m$ and $|V_H| = n$; let π be a resolution refutation of $FPHP(G)$; let $w_0, \alpha \in [m]$ be such that $w_0 > \log L(\pi)$ and $w_0 \geq \alpha^2 \geq 4$, and let $\vec{\delta} = (\delta_1, \dots, \delta_m)$ be defined by $\delta_i = \frac{\Delta_G(i) \log \alpha}{\log m}$. Then there exists an integer vector $\vec{d} = (d_1, \dots, d_m)$, with $\delta_i < d_i \leq \Delta_G(i)$ for all $i \in V_P$, a set of (w_0, \vec{d}) -axioms \mathcal{A} with $|\mathcal{A}| \leq L(\pi)$, and a resolution refutation π' of $FPHP(G) \cup \mathcal{A}$ such that $w_{\vec{d}, \vec{\delta}}(\pi') = O(\alpha \cdot w_0)$.*

As mentioned above, this upper bound is a straightforward application of Lemma 3.1. We defer the formal proof to Section 4.2. What we will need from Lemma 4.2 is that a resolution

refutation of $FPHP(G)$ in length less than 2^{w_0} can be transformed into a refutation of $FPHP(G) \cup \mathcal{A}$ in pseudo-width at most $O(\alpha \cdot w_0)$.

The second step in the proof is to show that any resolution refutation π of $FPHP(G) \cup \mathcal{A}$ requires large pseudo-width. The high-level idea is to define a progress measure on clauses $C \in \pi$ by counting the number of matchings on $P_{\vec{d}, \vec{\delta}}(C)$ that do not satisfy C . We then show that in order to increase this progress measure we need large pseudo-width. The following lemma states the pseudo-width lower bound.

LEMMA 4.3. *Let $\xi \leq 1/4$ and $m, n, r, \Delta \in \mathbb{N}$; let $G = (V_P \dot{\cup} V_H, E)$ with $|V_P| = m$ and $|V_H| = n$ be an $(r, \Delta, (1 - 2\xi)\Delta)$ -boundary expander, and let $\vec{\delta} = (\delta_1, \dots, \delta_m)$ be defined by $\delta_i = 4\Delta_G(i)\xi$. Suppose that $\vec{d} = (d_1, \dots, d_m)$ is an integer vector such that $\delta_i < d_i \leq \Delta_G(i)$ for all $i \in V_P$. Let w_0 be an arbitrary parameter and \mathcal{A} be an arbitrary set of (w_0, \vec{d}) -axioms with $|\mathcal{A}| \leq (1 + \xi)^{w_0}$. Then every resolution refutation π of $FPHP(G) \cup \mathcal{A}$ must satisfy $w_{\vec{d}, \vec{\delta}}(\pi) \geq r\xi/4$.*

In one sentence, the lemma states that if the set of “fake axioms” \mathcal{A} is not too large, then resolution requires large pseudo-width to refute $FPHP(G) \cup \mathcal{A}$. Note that this lemma holds for any filter vector and not just for the one obtained from Lemma 4.2.

In order to prove Lemma 4.3, we wish to define a progress measure on clauses that indicates how close the derivation is to refuting the formula (i.e., it should be small for axiom clauses but large for contradiction). A first attempt would be to define the progress of a clause C as the number of ruled-out matchings (i.e., matchings that do not satisfy C) on the pigeons mentioned by C . This definition does not quite work since we will want to keep track of *which* matchings are ruled out—unless there are more than $\Delta_G(i) - d_i + \delta_i/4$ holes to which a pigeon can be mapped without satisfying C . Then we will want to think of C as ruling out *all holes* for this pigeon. Since the pigeon degree of a light pigeon i is at most $d_i - \delta_i$, such a pigeon will certainly have at least $\Delta_G(i) - d_i + \delta_i \geq \Delta_G(i) - d_i + \delta_i/4$ holes to which it can be mapped, and the “lossy counting” will thus ensure that all holes are considered as ruled out.

We realize this “lossy counting” through a linear space Λ , in which each partial matching φ is associated with a subspace $\lambda(\varphi)$. Roughly speaking, the progress $\lambda(C)$ of a clause C is then defined to be the span of all partial matchings that are ruled out by C . We design the association between matchings and subspaces so that the contradictory empty clause \perp has $\lambda(\perp) = \Lambda$ but so that the span of all the axioms $\text{span}(\{\lambda(A) \mid A \in FPHP(G) \cup \mathcal{A}\})$ is a proper subspace of Λ . For this argument we crucially rely on the gap (r_i vs $r_i + 1$) between the two cases of the filter lemma (Lemma 3.1). This implies that in a refutation π of $FPHP(G) \cup \mathcal{A}$ there must exist a resolution step deriving a clause C from clauses C_0 and C_1 such that the linear space of the resolvent $\lambda(C)$ is not contained in $\text{span}(\lambda(C_0), \lambda(C_1))$. But the main technical lemma of this section (Lemma 4.10) says that for any derivation in low pseudo-width the linear space of the resolvent is contained in the span of the linear spaces of the clauses being resolved. Hence,

in order for π to be a refutation it must contain a clause with large pseudo-width, and this establishes Lemma 4.3.

So far our argument follows that of Razborov very closely, but it turns out we cannot realize this proof idea if we only keep track of heavy and light pigeons. Let us attempt a proof of the claim in Lemma 4.10 that low-width resolution steps cannot increase the span to illustrate what the problem is. The interesting case is when there is a pigeon i that is heavy for C_0 or C_1 but not for their resolvent C . Then, following Razborov, for any matching φ on the heavy pigeons of C that fails to satisfy C , we need to be able to extend φ in at least $\Delta_G(i) - d_i + \delta_i/4$ different ways to a matching including also pigeon i that falsifies either C_0 or C_1 . If this can be done, then we think of C_0 and C_1 as together ruling out (essentially) all holes for i , and the linear space associated with C will be contained in the span of the spaces for C_0 and C_1 . However, the problem is that φ can send all heavy pigeons to the neighbourhood of pigeon i . In this scenario, there might be very few holes, or even no holes, to which i can be mapped when extending φ , and even our lossy counting will not be able to pick up enough holes for the argument to go through. We resolve this problem by not only considering the heavy pigeons but a larger set of *relevant* pigeons including all pigeons i' that can become overly constrained when some matching on the heavy pigeons shrinks the neighbourhood of i' too much. Formally, the *closure* of the set of heavy pigeons, as defined in Definition 3.2, is the notion we need.

4.1 Formal Statements of Graph FPHP Formula Lower Bounds

Deferring the proofs of all technical lemmas for now, let us state our lower bounds for graph FPHP formulas and see how they follow from Lemmas 4.2 and 4.3 above.

THEOREM 4.4. *Let $m = |U|$ and $n = |V|$ and suppose that $G = (U \dot{\cup} V, E)$ is an $(r, \Delta, (1 - \frac{\log \alpha}{2 \log m})\Delta)$ -boundary expander for $\alpha \in [m]$ such that $8 \leq \frac{\alpha^3}{\log \alpha} = o(\frac{r}{\log m})$. Then resolution requires length $\exp\left(\Omega\left(\frac{r \log^2 \alpha}{\alpha \log^2 m}\right)\right)$ to refute FPHP(G).*

As promised in Section 3, let us briefly discuss the parameter α . Note that, on the one hand, the larger α is, the more relaxed we can be with respect to the expansion requirements, and hence the set of formulas to which the lower bound applies becomes larger. On the other hand, the strength of the lower bound deteriorates quickly with α . Hence, we need to choose α carefully to find a good compromise between these two concerns.

PROOF OF THEOREM 4.4. Let $\xi = \frac{\log \alpha}{4 \log m}$ and let $w_0 = \frac{\varepsilon_0 r \xi}{\alpha}$ for some small enough $\varepsilon_0 > 0$. We note that the choice of parameters and the condition on α ensure that $4 \leq \alpha^2 \leq w_0$. Furthermore, in terms of ξ , the graph G is an $(r, \Delta, (1 - 2\xi)\Delta)$ -boundary expander.

We proceed by contradiction. Suppose π is a resolution refutation with $L(\pi) < 2^{\varepsilon' w_0 \xi}$ for a small enough constant $\varepsilon' > 0$. Applying Lemma 4.2 we get a set of (w_0, \vec{d}) -axioms \mathcal{A} with

$|\mathcal{A}| \leq L(\pi)$ and a resolution refutation π' of $FPHP(G) \cup \mathcal{A}$ such that $w_{\vec{a}, \vec{\delta}}(\pi') \leq K\alpha w_0$ for some large enough constant K .

Note that $|\mathcal{A}| \leq L(\pi) < 2^{\varepsilon' w_0 \xi} \leq (1 + \xi)^{w_0}$ for $\varepsilon' < 1/2$. Applying Lemma 4.3 to π' yields a pseudo-width lower bound of $r\xi/4$. We conclude that

$$r\xi/4 \leq w_{\vec{a}, \vec{\delta}}(\pi') \leq K\alpha w_0 = \varepsilon_0 K r \xi . \quad (4.3)$$

Choosing $\varepsilon_0 < \frac{1}{4K}$ yields a contradiction. ■

The following corollary summarizes our claims for random graphs.

COROLLARY 4.5. *Let m and n be positive integers and let $\Delta : \mathbb{N}^+ \rightarrow \mathbb{N}^+$ and $\varepsilon : \mathbb{N}^+ \rightarrow [0, 1]$ be any monotone functions of n such that $n < m \leq n^{(\varepsilon/16)^2 \log n}$ and $n \geq \Delta \geq \left(\frac{16 \log m}{\varepsilon \log n}\right)^2$. Then asymptotically almost surely resolution requires length $\exp(\Omega(n^{1-\varepsilon}))$ to refute $FPHP(G)$ for $G \sim \mathcal{G}(m, n, \Delta)$.*

PROOF. Let us assume that $n^{(\varepsilon/16)^2 \log n}$ and $((16 \log m)/(\varepsilon \log n))^2$ are integers. Observe that if $G \sim \mathcal{G}(m, n, \Delta)$ for $\Delta > ((16 \log m)/(\varepsilon \log n))^2$, then we can sample a graph G' from the distribution $\mathcal{G}(m, n, ((16 \log m)/(\varepsilon \log n))^2)$ by choosing a random subset of appropriate size of each neighbourhood of a left vertex of G (and applying a restriction zeroing out the other edges). Hence, we can restrict our attention to the case where $\Delta = ((16 \log m)/(\varepsilon \log n))^2$. Also, it is sufficient to prove the claim for $m = n^{(\varepsilon/16)^2 \log n}$, since choosing m smaller can only make the formula less constrained and hence makes the lower bound easier to obtain.

We want to apply Lemma 2.2 for $\chi = \alpha = n^{\varepsilon/4}$ and $\xi = \frac{\log \alpha}{4 \log m}$. In order to do so, we need to verify the inequalities

$$\xi < 1/2 , \quad (4.4a)$$

$$\xi \ln \chi \geq 2 , \quad (4.4b)$$

$$\xi \Delta \ln \chi \geq 4 \ln m . \quad (4.4c)$$

For (4.4a) we observe that $\xi = \frac{16}{\varepsilon \log n}$ and since $n < n^{(\varepsilon/16)^2 \log n}$ we see that $\frac{1}{\log n} < \left(\frac{\varepsilon}{16}\right)^2$. Hence, the first condition holds for n large enough. To check (4.4b), we compute

$$\xi \ln \chi = \frac{16}{\varepsilon \log n} \frac{\varepsilon \ln n}{4} \geq 2 . \quad (4.5)$$

For (4.4c), we observe that $\Delta = \log m$ and hence

$$\xi \Delta \ln \chi = \frac{4}{\log e} \log m = 4 \ln m . \quad (4.6)$$

We conclude that asymptotically almost surely, $G \sim \mathcal{G}(m, n, \Delta)$ is an $(n^{1-\varepsilon/2}, \Delta, (1 - 2\xi)\Delta)$ -boundary expander. Theorem 4.4 then gives a length lower bound of $\exp(\Omega(n^{1-\varepsilon}))$, as required. ■

The following two corollaries are simple consequences of Corollary 4.5, optimizing for different parameters. The first corollary gives the strongest lower bounds, while the second minimizes the degree.

COROLLARY 4.6. *Let m, n be such that $m \leq n^{o(\log n)}$. Then asymptotically almost surely resolution requires length $\exp(\Omega(n^{1-o(1)}))$ to refute $\text{FPHP}(G)$ for $G \sim \mathcal{G}(m, n, \log m)$.*

PROOF. Let $m = n^{f(n)}$, where $f(n) = o(\log n)$. Applying Corollary 4.5 for $\varepsilon = 16\sqrt{\frac{f(n)}{\log n}} = o(1)$ we get the desired statement. ■

COROLLARY 4.7 (Restatement of Theorem 1.3). *Let k and n be positive integers and let $m = n^k$ and $\varepsilon \in \mathbb{R}^+$. Then asymptotically almost surely resolution requires length $\exp(\Omega(n^{1-\varepsilon}))$ to refute $\text{FPHP}(G)$ for $G \sim \mathcal{G}\left(m, n, \left(\frac{16k}{\varepsilon}\right)^2\right)$.*

PROOF. We appeal to Corollary 4.5 with $\Delta = \left(\frac{16k}{\varepsilon}\right)^2$, $m = n^k$ and ε constant. A short calculation shows that all conditions are met. ■

Our final corollary shows that we can get meaningful lower bounds even for a weakly exponential number of pigeons. Unfortunately, the statement does not hold for random graphs.

COROLLARY 4.8. *Let $\kappa < 3/2 - \sqrt{2}$ and $\varepsilon > 0$ be constant and n be integer. Then there is a family of explicitly constructible graphs G with $m = 2^{\Omega(n^\kappa)}$ and left degree $O(\log^{1/\sqrt{\kappa}}(m))$ such that resolution requires length $\exp(\Omega(n^{1-2\sqrt{\kappa}(2-\sqrt{\kappa})-\varepsilon}))$ to refute $\text{FPHP}(G)$.*

PROOF. Let G be the graph from Corollary 2.4 with $v = \frac{2\sqrt{\kappa}}{1-2\sqrt{\kappa}}$. An appeal to Theorem 4.4 using the graph G yields the desired lower bound. ■

4.2 A Pseudo-Width Upper Bound for Graph FPHP Formulas with Extra Axioms

Let us now prove Lemma 4.2. For this proof, let us identify V_P with $[m]$. For every clause C in the refutation π , let $\vec{r}(C) = (r_1(C), \dots, r_m(C))$ be the vector where each coordinate is given by

$$r_i(C) = \left\lfloor \frac{\Delta_G(i) - \Delta_C(i)}{\delta_i} \right\rfloor + 1, \quad (4.7)$$

where $\delta_i = \frac{\Delta_G(i) \log \alpha}{\log m}$ as in Lemma 4.2 and $\Delta_C(i)$ denotes the i th pigeon degree of C , that is, the number of holes pigeon i may go to such that the clause C is satisfied under the corresponding partial assignment (see (4.2)). We apply the filter lemma (Lemma 3.1) to the set of vectors $\{\vec{r}(C) \mid C \in \pi\}$. Denote by $\vec{r} = (r_1, \dots, r_m)$ a vector as guaranteed to exist by Lemma 3.1. Let

$$d_i = \Delta_G(i) - \lceil \delta_i r_i \rceil + 1. \quad (4.8)$$

A short calculation establishes that d_i is the smallest integer such that $\left\lfloor \frac{\Delta_G(i) - d_i}{\delta_i} \right\rfloor + 1 \leq r_i$.

Note that every pigeon $i \in [m]$ such that $r_i(C) \leq r_i$ is super-heavy for C . Also, every heavy pigeon of a clause C satisfies that $r_i(C) \leq r_i + 1$.

To obtain a refutation π' that satisfies the conclusions of the lemma, we consider every clause $C \in \pi$ and either there is a (w_0, \vec{d}) -axiom which is a strengthening of C that can be added to \mathcal{A} or the pseudo-width of C is small enough and the clause can thus remain in π' . More concretely, we make a case distinction whether $\vec{r}(C)$ satisfies case 1 of Lemma 3.1 or only case 2. In the former case there is a (w_0, \vec{d}) -axiom which is a strengthening of C , while in the other the pseudo-width of C is bounded:

1. C satisfies $|\{i \in [m] \mid r_i(C) \leq r_i\}| \geq w_0$: As every pigeon $i \in [m]$ with $r_i(C) \leq r_i$ also satisfies $\Delta_C(i) \geq d_i$, there is a (w_0, \vec{d}) -axiom which is a strengthening of C that can be added to \mathcal{A} . This reduces the pseudo-width of this clause to w_0 .
2. C satisfies $|\{i \in [m] \mid r_i(C) \leq r_i + 1\}| \leq O(\alpha \cdot w_0)$: As every heavy pigeon always satisfies $r_i(C) \leq r_i + 1$, the pseudo-width of C is $O(\alpha \cdot w_0)$.

This concludes the proof as $|\mathcal{A}| \leq L(\pi)$ and the pseudo-width of π' is $O(\alpha \cdot w_0)$ by construction.

4.3 A Pseudo-Width Lower Bound for Graph FPHP Formulas with Extra Axioms

We continue to the proof of Lemma 4.3. For convenience, we restate the lemma here.

LEMMA 4.3. (Restated) *Let $\xi \leq 1/4$ and $m, n, r, \Delta \in \mathbb{N}$; let $G = (V_P \dot{\cup} V_H, E)$ with $|V_P| = m$ and $|V_H| = n$ be an $(r, \Delta, (1 - 2\xi)\Delta)$ -boundary expander; and let $\vec{\delta} = (\delta_1, \dots, \delta_m)$ be defined by $\delta_i = 4\Delta_G(i)\xi$. Suppose that $\vec{d} = (d_1, \dots, d_m)$ is an integer vector such that $\delta_i < d_i \leq \Delta_G(i)$ for all $i \in V_P$. Let w_0 be an arbitrary parameter and \mathcal{A} be an arbitrary set of (w_0, \vec{d}) -axioms with $|\mathcal{A}| \leq (1 + \xi)^{w_0}$. Then every resolution refutation π of $\text{FPHP}(G) \cup \mathcal{A}$ must satisfy $w_{\vec{d}, \vec{\delta}}(\pi) \geq r\xi/4$.*

Using Definition 3.2, we define the set of *relevant* pigeons of a clause C as

$$\text{closure}(C) = \text{closure}_{r, (1-3\xi)\Delta}(P_{\vec{d}, \vec{\delta}}(C)) , \quad (4.9)$$

where $P_{\vec{d}, \vec{\delta}}(C)$ denotes the set of $(\vec{d}, \vec{\delta})$ -heavy pigeons for C as defined in Definition 4.1. By definition, the closure of a set T contains T itself but is only defined if $|T| \leq r$. However, if $|P_{\vec{d}, \vec{\delta}}(C)| \geq r \geq r\xi/4$ then we already have the lower bound claimed in the lemma, and so we may assume that the closure is well-defined for all clauses in the refutation π . This implies, in particular, that for every clause $C \in \pi$ we have $P_{\vec{d}, \vec{\delta}}(C) \subseteq \text{closure}(C)$.

Let us next construct the linear space Λ and describe how matchings are mapped into it. Fix a field \mathbb{F} of characteristic 0 and for each pigeon $i \in V_P$ let Λ_i be a linear space over \mathbb{F} of dimension $\Delta_G(i) - d_i + \delta_i/4$. Let Λ be the tensor product $\Lambda = \bigotimes_{i \in V_P} \Lambda_i$ and denote by $\lambda_i : V_H \mapsto \Lambda_i$ a function with the property¹ that any subset of holes $J \subseteq V_H$ of size at least

¹ It is readily seen that such functions exist by, e.g., considering mappings that send each hole $j \in V_H$ to a uniformly sampled unit vector.

$\dim(\Lambda_i)$ spans Λ_i . In other words, for J as above we have that $\Lambda_i = \text{span}(\lambda_i(j) : j \in J)$. This is how we will realize the idea of “lossy counting”. For $J \subseteq V_H$ such that $|J| \leq \dim(\Lambda_i)$ we have exact counting $\dim(\text{span}(\{\lambda_i(j) \mid j \in J\})) = |J|$ as otherwise we can construct a set of size $\dim(\Lambda_i)$ that does not span the entire space, but when $|J| > \dim(\Lambda_i)$ gets large enough we have $\dim(\text{span}(\{\lambda_i(j) \mid j \in J\})) = \dim(\Lambda_i)$. We crucially rely on this notion of “lossy counting” in the proof of Lemma 4.10.

In order to map functions $V_P \mapsto V_H$ into Λ , we define $\lambda : V_H^{V_P} \mapsto \Lambda$ by $\lambda(j_1, \dots, j_m) = \bigotimes_{i \in V_P} \lambda_i(j_i)$, where we will abuse notions slightly in that we identify a vector with the 1-dimensional space spanned by this vector. For a partial function $\varphi : V_P \mapsto V_H$, we let $\lambda(\varphi)$ be the span of all total extensions of φ (not necessarily matchings), or equivalently

$$\lambda(\varphi) = \bigotimes_{i \in \text{dom}(\varphi)} \lambda_i(\varphi_i) \otimes \bigotimes_{i \notin \text{dom}(\varphi)} \Lambda_i . \quad (4.10)$$

Recall that \mathcal{M} is the set of all partial matchings on the graph G and that we interchangeably think of partial matchings as partial functions $\varphi : V_P \rightarrow V_H$ or as Boolean assignments as defined in (2.3). For each clause C , we are interested in the partial matchings $\varphi \in \mathcal{M}$ with domain $\text{dom}(\varphi) = \text{closure}(C)$ such that φ does not satisfy C . We refer to the set of such matchings as the *zero space* of C and denote it by

$$Z(C) = \{\varphi \in \mathcal{M} \mid \text{dom}(\varphi) = \text{closure}(C) \wedge C \upharpoonright_{\varphi} \neq 1\} . \quad (4.11)$$

We associate C with the linear space

$$\lambda(C) = \text{span}(\{\lambda(\varphi) \mid \varphi \in Z(C)\}) . \quad (4.12)$$

Note that contradiction is mapped to Λ , i.e., $\lambda(\perp) = \Lambda$.

We assert that the span of the axioms $\text{span}(\{\lambda(A) \mid A \in \text{FPHP}(G) \cup \mathcal{A}\})$ is a proper subspace of Λ .

LEMMA 4.9. *If $|\mathcal{A}| \leq (1 + \xi)^{w_0}$, then $\text{span}(\{\lambda(A) \mid A \in \text{FPHP}(G) \cup \mathcal{A}\}) \subsetneq \Lambda$.*

Accepting this claim without proof for now, this implies that in π there is some resolution step deriving C from C_0 and C_1 where the subspace of the resolvent is not contained in the span of the subspaces of the premises, or in other words $\lambda(C) \not\subseteq \text{span}(\lambda(C_0), \lambda(C_1))$. Our next lemma, which is the heart of the argument, says that this cannot happen as long as the closures of the clauses are small.

LEMMA 4.10. *Let C be a clause derived from clauses C_0 and C_1 by the resolution rule. If it holds that*

$$\max\{|\text{closure}(C_0)|, |\text{closure}(C_1)|, |\text{closure}(C)|\} \leq r/4 ,$$

then $\lambda(C) \subseteq \text{span}(\lambda(C_0), \lambda(C_1))$.

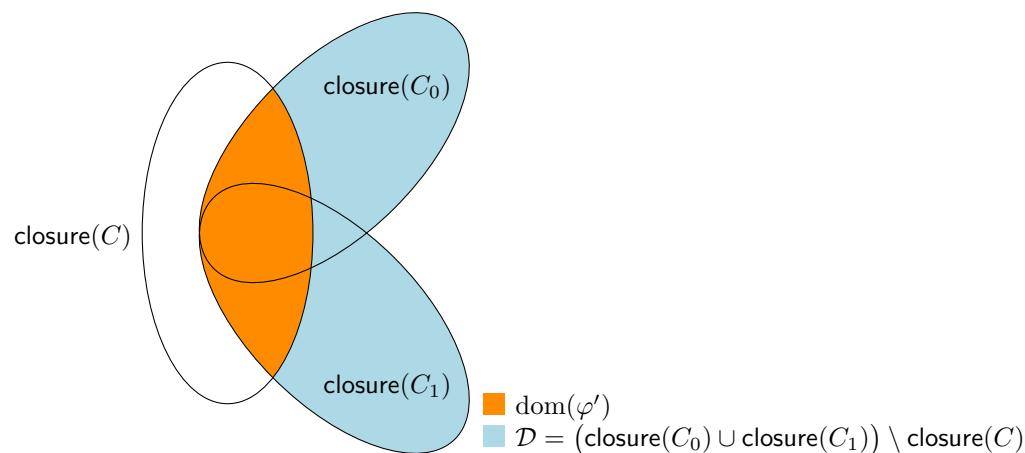


Figure 1. Depiction of relations between $\text{closure}(C)$, $\text{closure}(C_i)$, $i = 1, 2$, $\text{dom}(\varphi')$ and \mathcal{D} in proof of Lemma 4.10.

Since contradiction cannot be derived while the closure is of size at most $r/4$, any refutation π must contain a clause C with $|\text{closure}(C)| > r/4$. But then Lemma 3.3 implies that C has pseudo-width at least $r\xi/4$, and Lemma 4.3 follows. All that remains for us is to establish Lemmas 4.9 and 4.10.

PROOF OF LEMMA 4.9. We need to show that the axioms $FPHP(G) \cup \mathcal{A}$ do not span all of Λ . We start with the axioms in $FPHP(G)$.

Let A be pigeon axiom P^i as in (1.1a) or a functionality axiom $F_{j,j'}^i$ as in (1.1c). Note that i is a heavy pigeon for A . Clearly, there are no pigeon-to-hole assignments for pigeon i that do not satisfy A . Thus there are no matchings on $\text{closure}(A)$ that do not satisfy A . We conclude that $\lambda(A) = \emptyset$. If instead A is a hole axiom $H_j^{i,i'}$ as in (1.1b), then we can observe that $\Delta_G(i) - 1 \geq d_i - \delta_i$ since $\delta_i = 4\xi\Delta_G(i) \geq 2\xi\Delta \geq 1$ (by boundary expansion). This implies that A has two heavy pigeons. Observe that there are no matchings on these two pigeons that do not satisfy A . Thus $Z(A) = \emptyset$ and we conclude that $\lambda(A) = \emptyset$.

Now consider the (w_0, \vec{d}) -axioms in \mathcal{A} . We wish to show that any $A \in \mathcal{A}$ can only span a very small fraction of Λ . We can estimate the number of dimensions $\lambda(A)$ spans by

$$\dim \lambda(A) \leq \prod_{i \notin P_{\vec{d}}(A)} \dim \Lambda_i \prod_{i \in P_{\vec{d}}(A)} (\Delta_G(i) - \Delta_A(i)) \leq \prod_{i \notin P_{\vec{d}}(A)} \dim \Lambda_i \prod_{i \in P_{\vec{d}}(A)} (\Delta_G(i) - d_i), \quad (4.13)$$

since all pigeons $i \in P_{\vec{d}}(A)$ are super-heavy, that is, it holds that $\Delta_A(i) \geq d_i$. Hence the fraction of the space Λ that A may span is bounded by

$$\frac{\dim \lambda(A)}{\dim \Lambda} \leq \prod_{i \in P_{\vec{d}}(A)} \frac{\Delta_G(i) - d_i}{\Delta_G(i) - d_i + \delta_i/4} \leq (1 - \xi)^{w_0}. \quad (4.14)$$

As $|\mathcal{A}| \leq (1 + \xi)^{w_0}$ we can conclude that not all of Λ is spanned by the axioms. ■

PROOF OF LEMMA 4.10. Let us write $S_{01} = \text{closure}(C_0) \cup \text{closure}(C_1)$ and $S = \text{closure}(C)$. In order to establish the lemma, we need to show for all $\varphi \in Z(C)$ that

$$\lambda(\varphi) \subseteq \text{span}(\lambda(C_0), \lambda(C_1)) . \quad (4.15)$$

Denote by φ' the restriction of φ to the domain $S \cap S_{01}$ and note that C is not satisfied under the Boolean assignment associated with φ' . Also, observe that if a matching η extends a matching η' , then $\lambda(\eta)$ is contained in $\lambda(\eta')$. This is so since for any pigeon $i \in \text{dom}(\eta) \setminus \text{dom}(\eta')$ we have from (4.10) that η' picks up the whole subspace Λ_i while η only gets a single vector. Thus, if we can show that $\lambda(\varphi') \subseteq \text{span}(\lambda(C_0), \lambda(C_1))$, then we are done as φ extends φ' and hence $\lambda(\varphi) \subseteq \lambda(\varphi')$. For the following argument it may be helpful to refer to the illustration in Figure 1.

Let $\mathcal{D} = S_{01} \setminus S$ and denote by $\mathcal{M}_{\mathcal{D}}$ the set of matchings that extend φ' to the domain \mathcal{D} and do not satisfy C . Since each matching $\psi \in \mathcal{M}_{\mathcal{D}}$ fails to satisfy C , by the soundness of the resolution rule we have that it also fails to satisfy either C_0 or C_1 . Assume without loss of generality that ψ does not satisfy C_0 and denote by ψ' the restriction of ψ to the domain of $\text{closure}(C_0)$. From (4.11) we see that $\psi' \in Z(C_0)$ and therefore $\lambda(\psi) \subseteq \lambda(\psi') \subseteq \lambda(C_0)$.

So far we have argued that for all matchings $\psi \in \mathcal{M}_{\mathcal{D}}$ it holds that $\lambda(\psi)$ is contained in $\text{span}(\lambda(C_0), \lambda(C_1))$. Let $\lambda(\mathcal{M}_{\mathcal{D}}) = \text{span}(\lambda(\psi) \mid \psi \in \mathcal{M}_{\mathcal{D}})$. If we can show that the set of matchings $\mathcal{M}_{\mathcal{D}}$ is large enough for $\lambda(\mathcal{M}_{\mathcal{D}}) = \lambda(\varphi')$ to hold, then the lemma follows. In other words, we want to show that $\lambda(\mathcal{M}_{\mathcal{D}})$ projected to $\Lambda_{\mathcal{D}} = \bigotimes_{i \in \mathcal{D}} \Lambda_i$ spans all of the space $\Lambda_{\mathcal{D}}$.

To argue this, note first that \mathcal{D} is completely outside the $\text{closure}(C)$. Furthermore, by assumption we have $|\text{closure}(C)| \leq r/4$ and $|\mathcal{D}| \leq |S_{01}| \leq r/2$. An application of Lemma 3.4 now tells us that

$$|\partial_{G \setminus (\text{closure}(C) \cup N(\text{closure}(C)))}(\mathcal{D})| \geq (1 - 3\xi)\Delta|\mathcal{D}| . \quad (4.16)$$

By an averaging argument, there must exist a pigeon $i_1 \in \mathcal{D}$ that has more than $(1 - 3\xi)\Delta$ unique neighbours in $\partial_{G \setminus (\text{closure}(C) \cup N(\text{closure}(C)))}(\mathcal{D})$. The same argument applied to $\mathcal{D} \setminus \{i_1\}$ show that some pigeon i_2 has more than $(1 - 3\xi)\Delta$ unique neighbours on top of the neighbours reserved for pigeon i_1 . Iterating this argument, we derive by induction that for each pigeon $i \in \mathcal{D}$ we can find $(1 - 3\xi)\Delta$ distinct holes in $N(\mathcal{D})$. Since all pigeons in \mathcal{D} are light in C , it follows that at most $d_i - \delta_i$ mappings of pigeon i can satisfy the clause C . Hence, there are at least

$$\begin{aligned} (1 - 3\xi)\Delta - (d_i - \delta_i) &\geq (1 - 3\xi)\Delta_G(i) - d_i + 4\xi\Delta_G(i) \\ &\geq \Delta_G(i) - d_i + \delta_i/4 \end{aligned} \quad (4.17)$$

many holes to which each pigeon in \mathcal{D} can be sent, independently of all other pigeons in \mathcal{D} , without satisfying C . As we have that $\dim(\Lambda_i) = \Delta_G(i) - d_i + \delta_i/4$, we conclude that $\lambda(\mathcal{M}_{\mathcal{D}})$ projected to $\Lambda_{\mathcal{D}}$ spans the whole space. This concludes the proof of the lemma. ■

5. Lower Bounds for Perfect Matching Principle Formulas

In this section, we show that the perfect matching principle formulas defined over even highly unbalanced bipartite graphs require exponentially long resolution refutations if the graphs are expanding enough.

Just as in [29], our proof is by an indirect reduction to the FPHP lower bound, and therefore there is a significant overlap in concepts and notation with Section 4. However, since there are also quite a few subtle shifts in meaning, we restate all definitions in full below to make the exposition in this section self-contained and unambiguous.

We first review some useful notions from [25]. Let $G = (V, E)$ denote the graph over which the formulas are defined. For a clause C and a vertex $v \in V(G)$, let the *clause-neighbourhood of v in C* , denoted by $N_C(v)$, be the vertices $u \in V(G)$ with the property that C is satisfied if v is matched to u , that is,

$$N_C(v) = \{u \in V \mid e = \{u, v\} \in E \text{ and } \rho_{\{e\}}(C) = 1\} . \quad (5.1)$$

For a set $V \subseteq V(G)$ let $N_C(V)$ be the union of the clause-neighbourhoods of the vertices in V , i.e., $N_C(V) = \bigcup_{v \in V} N_C(v)$ and let the *v th vertex degree of C* be

$$\Delta_C(v) = |N_C(v)| . \quad (5.2)$$

We think of a vertex v with large degree $\Delta_C(v)$ as a vertex on which the derivation has not made any progress up to the point of deriving C , since the clause rules out very few neighbours. The vertices with high enough vertex degree in a clause are the *heavy vertices* of the clause as defined next.

DEFINITION 5.1 (Vertex weight, pseudo-width and (w_0, \vec{d}) -axioms). Let $\vec{d} = (d_1, \dots, d_{m+n})$ and $\vec{\delta} = (\delta_1, \dots, \delta_{m+n})$ be two vectors such that \vec{d} is elementwise greater than $\vec{\delta}$. We say that a vertex v is *\vec{d} -super-heavy for C* if $\Delta_C(v) \geq d_v$ and that vertex v is *$(\vec{d}, \vec{\delta})$ -heavy for C* if $\Delta_C(v) \geq d_v - \delta_v$. When \vec{d} and $\vec{\delta}$ are understood from context we omit the parameters and just refer to *super-heavy* and *heavy* vertices. Vertices that are not heavy are referred to as *light vertices*. The set of vertices that are super-heavy for C is denoted by

$$V_{\vec{d}}(C) = \{v \in V \mid \Delta_C(v) \geq d_v\} \quad (5.3)$$

and the set of heavy vertices for C is denoted by

$$V_{\vec{d}, \vec{\delta}}(C) = \{v \in V \mid \Delta_C(v) \geq d_v - \delta_v\} . \quad (5.4)$$

The *pseudo-width* $w_{\vec{d}, \vec{\delta}}(C) = |V_{\vec{d}, \vec{\delta}}(C)|$ of a clause C is the number of heavy vertices in it, and the pseudo-width of a resolution refutation π is $w_{\vec{d}, \vec{\delta}}(\pi) = \max_{C \in \pi} w_{\vec{d}, \vec{\delta}}(C)$. We refer to clauses C with precisely w_0 super-heavy vertices as *(w_0, \vec{d}) -axioms*.

To a large extent, the proof of the lower bounds for perfect matching formulas follows the general idea of the proof of Theorem 4.4: given a short refutation we first apply the filter lemma to obtain a refutation of small pseudo-width; we then prove that in small pseudo-width contradiction cannot be derived and can thus conclude that no short refutation exists. In more detail, given a short resolution refutation π , we use the filter lemma (Lemma 3.1) to get a filter vector $\vec{d} = (d_1, \dots, d_{m+n})$ such that each clause either has many super-heavy vertices or not too many heavy vertices (for an appropriately chosen vector $\vec{\delta}$). Clearly, clauses that fall into the second case of the filter lemma have bounded pseudo-width. Clauses in the first case, however, may have very large pseudo-width. In order to obtain a proof of low pseudo-width, these latter clauses are strengthened to (w_0, \vec{d}) -axioms and added to a special set \mathcal{A} . This then gives a refutation π' that refutes the formula $PM(G) \cup \mathcal{A}$ in bounded pseudo-width as stated in the next lemma.

LEMMA 5.2. *Let $G = (V_L \dot{\cup} V_R, E)$ be a bipartite graph with $|V_L| = m$ and $|V_R| = n$; let π be a resolution refutation of $PM(G)$; let $w_0, \alpha \in [m+n]$ be such that $w_0 > \log L(\pi)$ and $w_0 \geq \alpha^2 \geq 4$, and let $\vec{\delta} = (\delta_1, \dots, \delta_{m+n})$ be defined by $\delta_v = \frac{\Delta_G(v) \log \alpha}{\log(m+n)}$ for $v \in V(G)$. Then there exists an integer vector $\vec{d} = (d_1, \dots, d_{m+n})$, with $\delta_v < d_v \leq \Delta_G(v)$ for all $v \in V(G)$, a set of (w_0, \vec{d}) -axioms \mathcal{A} with $|\mathcal{A}| \leq L(\pi)$, and a resolution refutation π' of $PM(G) \cup \mathcal{A}$ such that $L(\pi') \leq L(\pi)$ and $w_{\vec{d}, \vec{\delta}}(\pi') \leq O(\alpha \cdot w_0)$.*

The proof of the above lemma is omitted as it is syntactically equivalent to the proof of Lemma 4.2. Until this point, we have almost mimicked the proof of Theorem 4.4. The main differences will appear in the proof of the counterpart to Lemma 5.2, which states a pseudo-width lower bound.

LEMMA 5.3. *Assume for $\xi \leq 1/64$ and $m, n, r, \Delta \in \mathbb{N}$ that $G = (V_L \dot{\cup} V_R, E)$ is an $(r, \Delta, (1 - 2\xi)\Delta)$ -boundary expander with $|V_L| = m$, $|V_R| = n$, $\Delta \geq \log m/\xi^2$, and $\min\{\Delta_G(v) : v \in V_R\} \geq r/\xi$. Let $\vec{\delta} = (\delta_v \mid v \in V(G))$ be defined by $\delta_v = 64\Delta_G(v)\xi$ and suppose that $\vec{d} = (d_v \mid v \in V(G))$ is an integer vector such that $\delta_v < d_v \leq \Delta_G(v)$ for all $v \in V(G)$. Fix w_0 such that $64 \leq w_0 \leq r\xi - \log n$ and let \mathcal{A} be an arbitrary set of (w_0, \vec{d}) -axioms with $|\mathcal{A}| \leq (1 + 16\xi)^{w_0/8}$. Then every resolution refutation π of $PM(G) \cup \mathcal{A}$ has either length $L(\pi) \geq 2^{w_0/32}$ or pseudo-width $w_{\vec{d}, \vec{\delta}}(\pi) \geq r\xi$.*

The proof of the above lemma is based on a sort of reduction to the $FPHP(G)$ case. The idea, due to Razborov [29], is to first pick a partition of the vertices of G that looks random to every clause in the refutation and then simulate the $FPHP(G)$ lower bound on this partition. In our setting, however, this process gets quite involved. Already implementing the partition idea of Razborov is non-trivial: for a fixed clause C some vertices that are light may be super-heavy with respect to the partition, and we do not have an upper bound on the pseudo-width any longer. The insight needed to solve this issue is to show that by expansion there are not too

many such vertices per clause, and then adapt the closure definition to take these vertices into account.

Another issue we run into is that the span argument from Section 4 cannot be applied to all the vertices in the graph since vertices in V_R are not guaranteed to be expanding. Instead, for the vertices in V_R , we need to resort to the span argument from [30]. Moreover, vertices in the neighbourhood of \mathcal{D} (as defined in the proof of Lemma 4.10) may already be matched, and we are hence unable to attain enough matchings. Our solution is to consider a “lazy” edge removal procedure from the original matching, which with a careful analysis can be shown to circumvent the problem—see Section 5.3 for details.

5.1 Formal Statements of Perfect Matching Formula Lower Bounds

Let us state our lower bounds for the perfect matching formulas and defer the proof of Lemma 5.3 to Section 5.3.

THEOREM 5.4. *Let $G = (U \dot{\cup} V, E)$ be a bipartite graph with $m = |U|$ and $n = |V|$. Suppose that G is an $(r, \Delta, (1 - 2\xi)\Delta)$ -boundary expander for $\Delta \geq \frac{\log(m+n)}{\xi^2}$ and $\xi = \frac{\log \alpha}{64 \log(m+n)}$ where $\alpha \geq 2$ and $\frac{\alpha^3}{\log \alpha} = o\left(\frac{r}{\log(m+n)}\right)$, which furthermore satisfies the degree requirement $\min\{\Delta_G(v) : v \in V\} \geq r/\xi$. Then resolution requires length $\exp\left(\Omega\left(\frac{r \log^2 \alpha}{\alpha \log^2(m+n)}\right)\right)$ to refute the perfect matching formula $PM(G)$ defined over G .*

We remark that this theorem also holds if we replace the minimum degree constraint of V with an expansion guarantee from V to U . We state the theorem in the above form as we want to apply it to the graphs from [19] for which we have no expansion guarantee from V to U .

PROOF OF THEOREM 5.4. Let $w_0 = \frac{\varepsilon_0 r \xi}{\alpha}$, for some small enough $\varepsilon_0 > 0$. Suppose for the sake of contradiction that π is a resolution refutation of $PM(G)$ such that $L(\pi) < (1 + 16\xi)^{w_0/8}$. Since $w_0 > \log L(\pi)$, by Lemma 5.2 we have that there exists an integer vector $\vec{d} = (d_1, \dots, d_{m+n})$, with $\delta_v < d_v \leq \Delta_G(v)$, a set of (w_0, \vec{d}) -axioms \mathcal{A} with $|\mathcal{A}| \leq L(\pi) < (1 + 16\xi)^{w_0/8}$, and a resolution refutation π' of $PM(G) \cup \mathcal{A}$ such that $L(\pi') \leq L(\pi)$ and $w_{\vec{d}, \vec{\delta}}(\pi') \leq K\alpha w_0$ for some large enough constant K . Since $L(\pi') < (1 + 16\xi)^{w_0/8} \leq 2^{w_0/32}$, by Lemma 5.3, we have that $w_{\vec{d}, \vec{\delta}}(\pi') \geq r\xi \geq \alpha w_0 / \varepsilon_0$. Choosing $\varepsilon_0 < 1/K$, we get a contradiction and, thus, $L(\pi) \geq (1 + 16\xi)^{w_0/8} = \exp\left(\Omega\left(\frac{r\xi^2}{\alpha}\right)\right)$. ■

As in Section 4, we have a general statement for random graphs.

COROLLARY 5.5. *Let m and n be positive integers, let $\Delta : \mathbb{N}^+ \rightarrow \mathbb{N}^+$ and $\varepsilon : \mathbb{N}^+ \rightarrow [0, 1]$ be any monotone functions of n such that $n^3 < m \leq n^{(\varepsilon/128)^2 \log n}$ and $n \geq \Delta \geq \log(m+n) \left(\frac{128 \log(m+n)}{\varepsilon \log n}\right)^2$. Then asymptotically almost surely resolution requires length $\exp(\Omega(n^{1-\varepsilon}))$ to refute $PM(G)$ for $G \sim \mathcal{G}(m, n, \Delta)$.*

PROOF. For simplicity, let us assume that $m^+ = n^{(\varepsilon/128)^2 \log n}$ and $\Delta^- = \log(m+n) \cdot ((128 \log(m+n))/(\varepsilon \log n))^2$ are integers. It suffices to prove the claim for $m = m^+$ and $\Delta = \Delta^-$. Indeed, if $G \sim \mathcal{G}(m, n, \Delta)$, for $\Delta > \Delta^-$, we can sample a random subgraph $G' \sim \mathcal{G}(m, n, \Delta^-)$ of G by choosing a random subset of appropriate size of each neighbourhood of a left vertex and applying a restriction zeroing out the other edges. Furthermore, as for smaller m the formula gets less constrained and hence the lower bound is easier to obtain, it suffices to prove it for $m = m^+$.

We want to apply Lemma 2.2 for $\chi = \alpha = n^{\varepsilon/4}$ and $\xi = \frac{\log \alpha}{64 \log m}$, and towards this end we argue that the inequalities

$$\xi < 1/2, \quad (5.5a)$$

$$\xi \ln \chi \geq 2, \quad (5.5b)$$

$$\xi \Delta \ln \chi \geq 4 \ln m \quad (5.5c)$$

all hold. First observe that $\xi = \frac{32}{\varepsilon \log n}$ and $n < n^{(\varepsilon/128)^2 \log n}$, from which we conclude that $\frac{1}{\log n} < (\frac{\varepsilon}{128})^2$. Hence, the first inequality (5.5a) holds for n large enough. A simple calculation

$$\xi \ln \chi = \frac{32}{\varepsilon \log n} \frac{\varepsilon \ln n}{4} \geq 2 \quad (5.6)$$

shows that (5.5b) is also true. Finally, for (5.5c), we observe that $\Delta \geq \log^2 m$ and hence

$$\xi \Delta \ln \chi \geq \frac{8}{\log e} \log^2 m \geq 4 \ln m. \quad (5.7)$$

We conclude that asymptotically almost surely $G \sim \mathcal{G}(m, n, \Delta)$ is an $(n^{1-\varepsilon/2}, \Delta, (1 - 2\xi)\Delta)$ -boundary expander. Furthermore, by the Chernoff inequality (Theorem 2.1) asymptotically almost surely all right vertices have degree at least $n \cdot \frac{64 \log(m+n)}{\varepsilon \log n}$. Thus, Theorem 5.4 gives a length lower bound of $\exp(\Omega(n^{1-\varepsilon}))$ as claimed. ■

The following corollary is a simple consequence of Corollary 5.5, optimizing for the strongest lower bounds.

COROLLARY 5.6 (Restatement of Theorem 1.1). *Let m, n be such that $m \leq n^{o(\log n)}$. Then asymptotically almost surely resolution requires length $\exp(\Omega(n^{1-o(1)}))$ to refute $PM(G)$ for $G \sim \mathcal{G}(m, n, 8 \log^2 m)$.*

PROOF. Let $m = n^{f(n)}$, where $f(n) = o(\log n)$. Applying Corollary 5.5 for $\varepsilon = 128 \sqrt{\frac{f(n)}{\log n}} = o(1)$, we get the desired statement. ■

Our final corollary shows that we even get meaningful lower bounds for highly unbalanced bipartite graphs. As was the case for $FPHP(G)$, the required expansion is too strong to hold for random graphs with such large imbalance, but does hold for explicitly constructed graphs from [19].

COROLLARY 5.7 (Restatement of Theorem 1.2). *Let $\kappa < 3/2 - \sqrt{2}$ and $\varepsilon > 0$ be constants, and let n be an integer. Then there is a family of (explicitly constructible) graphs G with $m = 2^{\Omega(n^\kappa)}$ and left degree $O(\log^{1/\sqrt{\kappa}}(m))$, such that resolution requires length $\exp(\Omega(n^{1-2\sqrt{\kappa}(2-\sqrt{\kappa})-\varepsilon}))$ to refute $PM(G)$.*

PROOF. Let G be the graph from Corollary 2.4 with $\nu = \frac{2\sqrt{\kappa}}{1-2\sqrt{\kappa}}$. In order to apply Theorem 5.4 we need to satisfy the minimum right degree constraint. A simple way of doing this is by adding n^2 edges to G such that each vertex on the right has exactly n incident edges added while each vertex on the left has at most one incident edge added. This will leave us with a graph which has large enough right degree while each left degree increased by at most one. The additional edges may reduce the boundary expansion a bit, but a short calculation shows that by choosing $\xi = \frac{\log \alpha}{128 \log(m+n)}$ in Corollary 2.4, we can still guarantee the needed boundary expansion for Theorem 5.4. The corollary bound follows. ■

5.2 Defining Pigeons and Holes

As stated earlier, we prove the $PM(G)$ lower bound by simulating the $FPHP(G)$ lower bound from Section 4 on a partition $V_P \dot{\cup} V_H$ of the vertices of G . As the notation suggests, we think of the vertices in V_P as pigeons and of the vertices in V_H as holes, that is, we will ignore all edges contained in V_P or in V_H so that pigeons may only fly to holes (formally we restrict the variables corresponding to such edges to 0).

Let us first motivate the properties—captured in Lemma 5.8—that such a partition must satisfy in order for the $FPHP(G)$ simulation to go through. To begin with, recall that in the proof of Lemma 4.9 we show that a (w_0, \vec{d}) -axiom only spans an exponentially small fraction of the linear space Λ . The argument crucially relies on the fact that there are many super-heavy pigeons in every (w_0, \vec{d}) -axiom. To make this work over the partition $V_P \dot{\cup} V_H$, we require that a constant fraction of the super-heavy vertices of every (w_0, \vec{d}) -axiom are in V_P and that super-heavy vertices remain super-heavy with respect to this partition. This first issue is addressed by Property 1 of Lemma 5.8 whereas the second issue is guaranteed by the other properties: Property 2 ensures that for every vertex roughly half of its neighbours are in V_H while Properties 3 and 4 ensure that most clause-neighbourhoods behave in the same manner, i.e., up to a small set of vertices per clause every clause-neighbourhood of a vertex has roughly half of its vertices in V_H . Combining these arguments, we can bound the fraction of the space spanned by a (w_0, \vec{d}) -axiom.

The other main step of the $FPHP(G)$ lower bounds is Lemma 4.10 which states that in low pseudo-width the linear space associated with a resolvent never leaves the span of the premises. This argument relies on the expansion guarantee of the underlying graph and the fact that light pigeons are unconstrained. The required graph expansion (see Lemma 5.10) will follow from

Property 2 and Properties 2–4 are used to argue that light pigeons are also unconstrained with respect to the partition.

LEMMA 5.8. *Let $G = (V_L \dot{\cup} V_R, E)$ be an $(r, \Delta, (1 - 2\xi)\Delta)$ -boundary expander for $\xi \leq 1/4$ and $|V_L| \geq 4$. Fix w_0 such that $64 \leq w_0 \leq r$ and let \mathcal{A} be a set of (w_0, \vec{d}) -axioms of size $|\mathcal{A}| \leq \exp(w_0/32)$. Moreover, suppose that $\Delta \geq \log|V_L|/\xi^2$ and $\min\{\Delta_G(v) : v \in V_R\} \geq (\log|V_R| + w_0)/\xi^2$. If π is a resolution refutation of $PM(G) \cup \mathcal{A}$ with $L(\pi) \leq \exp(w_0/32)$, then there exists a vertex partition $V(G) = V_P \dot{\cup} V_H$ such that*

1. for every $A \in \mathcal{A}$:

$$|V_{\vec{d}}(A) \cap V_P| \geq w_0/4 ,$$

2. for every $v \in V$:

$$||N_G(v) \cap V_H| - 1/2|N_G(v)|| \leq 4\xi|N_G(v)| ,$$

3. for every $C \in \pi$ and for every $v \in V_R$:

$$||N_C(v) \cap V_H| - 1/2|N_C(v)|| \leq 4\xi|N_G(v)| ,$$

4. for every $C \in \pi$ there is a set of vertices $\tilde{V}(C) \subseteq V_L$, satisfying $|\tilde{V}(C)| \leq w_0/8$, such that for every $v \in V_L \setminus \tilde{V}(C)$:

$$||N_C(v) \cap V_H| - 1/2|N_C(v)|| \leq 4\xi\Delta .$$

The analogue of above lemma in [29] is Claim 19. The main difference is that in our setting Property 4 does not always hold for all vertices in the graph while in Razborov's setting the corresponding property always holds.

The following is an auxiliary claim used to argue Property 4, that is, that the error set $\tilde{V}(C)$ is small. The claim states that if G is a good expander and for a fixed clause C there are many vertices $v \in V_L$ such that $|N_C(v) \cap V_H|$ does not behave as expected, then there is a large set of vertices $\tilde{V}^*(C)$ whose clause-neighbourhood in V_H (i.e., the set $N_C(\tilde{V}^*(C)) \cap V_H$) deviates from its expected size.

CLAIM 5.9. *Let $G = (V_L \dot{\cup} V_R, E)$ be an $(r, \Delta, (1 - 2\xi)\Delta)$ -boundary expander. Fix any partition $V(G) = V_P \dot{\cup} V_H$ and any clause C . Let*

$$\tilde{V}(C) = \{v \in V_L : ||N_C(v) \cap V_H| - 1/2|N_C(v)|| > 4\xi\Delta\} .$$

If $|\tilde{V}(C)| > w_0/8$, then there is a set of vertices $\tilde{V}^(C) \subseteq \tilde{V}(C)$, with $|\tilde{V}^*(C)| = w_0/16$, such that*

$$||N_C(\tilde{V}^*(C)) \cap V_H| - 1/2|N_C(\tilde{V}^*(C))|| > 2\xi\Delta|\tilde{V}^*(C)| .$$

PROOF. Denote by $\tilde{V}^+(C)$ ($\tilde{V}^-(C)$ respectively) the vertices in $\tilde{V}(C)$ that have more neighbours (fewer neighbours respectively) in V_H than the expected $1/2|N_C(v)|$. As $|\tilde{V}(C)| > w_0/8$, one of the sets $\tilde{V}^+(C)$ or $\tilde{V}^-(C)$ is of cardinality at least $w_0/16$.

Case 1: Suppose $|\tilde{V}^+(C)| \geq w_0/16$ and let $\tilde{V}^*(C)$ be any subset of $\tilde{V}^+(C)$ of size $w_0/16$. As boundary expansion of G guarantees that $\tilde{V}^*(C)$ has at most $2\xi\Delta|\tilde{V}^*(C)|$ edges to non-unique neighbours in G we derive

$$|N_C(\tilde{V}^*(C)) \cap V_H| \geq \sum_{v \in \tilde{V}^*(C)} |N_C(v) \cap \partial(\tilde{V}^*(C)) \cap V_H| \quad (5.8)$$

$$\geq \sum_{v \in \tilde{V}^*(C)} |N_C(v) \cap V_H| - 2\xi\Delta|\tilde{V}^*(C)| \quad (5.9)$$

$$> \sum_{v \in \tilde{V}^*(C)} (1/2|N_C(v)| + 4\xi\Delta) - 2\xi\Delta|\tilde{V}^*(C)| \quad (5.10)$$

$$\geq 1/2|N_C(\tilde{V}^*(C))| + 2\xi\Delta|\tilde{V}^*(C)|, \quad (5.11)$$

where the strict inequality follows by definition of $\tilde{V}^+(C)$.

Case 2: Suppose $|\tilde{V}^-(C)| \geq w_0/16$ and let $\tilde{V}^*(C)$ be any subset of $\tilde{V}^-(C)$ of size $w_0/16$. Similar to the previous case we can conclude that

$$|N_C(\tilde{V}^*(C)) \cap V_H| \leq \sum_{v \in \tilde{V}^*(C)} |N_C(v) \cap V_H| \quad (5.12)$$

$$< \sum_{v \in \tilde{V}^*(C)} (1/2|N_C(v)| - 4\xi\Delta) \quad (5.13)$$

$$\leq 1/2|N_C(\tilde{V}^*(C))| - 2\xi\Delta|\tilde{V}^*(C)|, \quad (5.14)$$

where the last inequality uses that $\tilde{V}^*(C)$ has at most $2\xi\Delta|\tilde{V}^*(C)|$ edges incident to non-unique neighbours in G .

Combining both cases yields the claim. ■

PROOF OF LEMMA 5.8. Pick a partition $V = \mathbf{V}_P \dot{\cup} \mathbf{V}_H$ uniformly at random. In what follows we show that Property 1 holds with probability at least $3/4$ and Properties 2, 3 and 4 each hold with probability at least $7/8$. Hence by a union bound there exists a partition that satisfies all four properties simultaneously.

For the first property, since $E[|V_{\tilde{d}}(A) \cap \mathbf{V}_P|] = w_0/2$, by the multiplicative Chernoff bound, that is, Theorem 2.1, we have that

$$\Pr[|V_{\tilde{d}}(A) \cap \mathbf{V}_P| \leq w_0/4] \leq \exp(-w_0/16). \quad (5.15)$$

Since $|\mathcal{A}| \leq \exp(w_0/32)$ and $w_0 \geq 64$, a union bound over \mathcal{A} gives us that Property 1 holds except with probability $\exp(-w_0/32) \leq 1/4$.

To analyse Properties 2 and 3, let C either be a clause in π or be the graph G (i.e., the clause that contains all variables) and fix an arbitrary $v \in V(G)$. By Chernoff bound (Theorem 2.1) we get that

$$\begin{aligned} \Pr[||N_C(v) \cap \mathbf{V}_H| - 1/2|N_C(v)|| \geq 4\xi|N_G(v)|] \\ \leq 2 \exp\left(-\frac{(4\xi|N_G(v)|)^2}{|N_C(v)| + 4\xi|N_G(v)|}\right) \\ \leq \exp\left(-8\xi^2|N_G(v)| + 1\right) , \end{aligned} \quad (5.16)$$

where the last inequality holds as $|N_C(v)| \leq |N_G(v)|$ and $\xi \leq 1/4$.

By a union bound argument over the clauses in π and $v \in V_R$, we have that Property 3 holds except with probability $1/8$. For Property 2, we need to analyse vertices in V_L and in V_R separately. On the one hand, since $\min\{\Delta_G(v) : v \in V_L\} \geq (1-2\xi)\Delta \geq \frac{\log|V_L|}{2\xi^2}$ and $|V_L| \geq 4$, a union bound over $v \in V_L$ shows that Property 2 holds for all vertices V_L except with probability $1/16$. On the other, as $\min\{\Delta_G(v) : v \in V_R\} \geq (\log|V_R| + w_0)/\xi^2$, a union bound yields that Property 2 holds for all $v \in V_R$ except with probability $1/16$.

To obtain Property 4, fix a clause C and consider the set $\tilde{\mathbf{V}}(C)$ that contains all vertices $v \in V_L$ satisfying

$$||N_C(v) \cap \mathbf{V}_H| - 1/2|N_C(v)|| > 4\xi\Delta . \quad (5.17)$$

We want to show that it is unlikely that $|\tilde{\mathbf{V}}(C)| \geq w_0/8$. Note that such a large $\tilde{\mathbf{V}}(C)$ implies by Claim 5.9 that there is a set $S \subseteq V_L$ of size $w_0/16$ such that $||N_C(S) \cap \mathbf{V}_H| - 1/2|N_C(S)|| \geq 2\xi\Delta|S|$. By a union bound over all such sets S and applying Chernoff bound (Theorem 2.1) we have that

$$\begin{aligned} \Pr[|\tilde{\mathbf{V}}(C)| \geq w_0/8] \\ \leq \binom{|V_L|}{w_0/16} \max_{\substack{S \subseteq V_L: \\ |S|=w_0/16}} \Pr[||N_C(S) \cap \mathbf{V}_H| - 1/2|N_C(S)|| \geq \xi\Delta w_0/8] \end{aligned} \quad (5.18)$$

$$\leq |V_L|^{w_0/16} \cdot 2 \exp\left(-\frac{(\xi\Delta w_0/8)^2}{\Delta w_0/16 + \xi\Delta w_0/8}\right) \quad (5.19)$$

$$\leq \exp\left(-\xi^2\Delta w_0/8 + 1 + \log|V_L| \cdot w_0/16\right) \quad (5.20)$$

$$\leq \exp\left(-\log|V_L| \cdot w_0/16 + 1\right) , \quad (5.21)$$

where for (5.19) we observe that $|N_C(S)| \leq \Delta|S|$, for (5.20) we need that $\xi \leq 1/4$ and for (5.21) that $\Delta \geq \log|V_L|/\xi^2$. By a union bound over all clauses in π we see that Property 4 holds except with probability $1/8$. ■

Let $V_P \dot{\cup} V_H$ be a partition of $V(G)$ as guaranteed to exist by Lemma 5.8. For an overview of the vertex sets and how they relate we refer to Figure 2. The following lemma shows that the vertices in V_L expand into the set $V_R \cap V_H$. Let $G' = G \setminus (V_R \cap V_P)$ with vertex partition $(V_L \dot{\cup} (V_R \setminus V_P))$.

LEMMA 5.10. *The graph G' is an $(r, (1 + 8\xi)\Delta/2, (1 - 12\xi)\Delta/2)$ -boundary expander.*

PROOF. By Lemma 5.8, Property 2, every vertex in $V_P \cap V_L$ has degree at most $(1 + 8\xi)|N_G(v)|/2$ and at least $(1 - 8\xi)|N_G(v)|/2$. By the expansion guarantee of G , we know that $|N_G(v)| \geq (1 - 2\xi)\Delta$. Therefore all sets of size 1 are good enough boundary expanders. We continue by induction on the size of the set. Let S be a set of vertices of size at most r . In the original graph G , this set S has at least $(1 - 2\xi)\Delta|S|$ many unique neighbours. Thus, there is a vertex v in S that has at least $(1 - 2\xi)\Delta$ many unique neighbours in G . Further, by Lemma 5.8, Property 2, the vertex v has at least $(1 - 8\xi)\Delta/2$ many neighbours in $V_R \cap V_H$. Hence v has at least $(1 - 12\xi)\Delta/2$ many unique neighbours in V_H . From the induction hypothesis on $S \setminus \{v\}$, it follows that S has the required number of unique neighbours in V_H . ■

5.3 Pseudo-Width Lower Bound

We start by setting up the notation we will need to prove Lemma 5.3.

Let C be a clause in π , let $\tilde{V}(C) = \{v \in V_L : ||N_C(v) \cap V_H| - 1/2|N_C(v)|| > 4\xi\Delta\}$ and $\bar{V}(C) = (V_{\vec{d}, \vec{\delta}}(C) \cap V_L) \cup \tilde{V}(C)$. The closure of C is a subset of V_L in the graph G' , defined by

$$\text{closure}(C) = \text{closure}_{r, (1-20\xi)\Delta/2}(\bar{V}(C)) . \quad (5.22)$$

We define the closure only on V_L as we only have an expansion guarantee from V_L into $V_R \cap V_H$. As the concept of closure only makes sense on vertex sets which are expanding, we do not define it on V_R . The set of relevant vertices of a clause C are the vertices in $\text{closure}(C) \cup V_{\vec{d}, \vec{\delta}}(C)$. With this definition at hand we proceed to set up the linear spaces that realize the lossy counting (see Section 4). Let us stress the fact that only vertices in V_P are associated with a linear space.

Fix a field \mathbb{F} of characteristic 0 and for each vertex $v \in V_P$ let Λ_v be a linear space over \mathbb{F} of dimension $1/2(\Delta_G(v) - d_v + \delta_v/2)$. Let $\Lambda = \bigotimes_{v \in V_P} \Lambda_v$ and denote by $\lambda_v : V_H \mapsto \Lambda_v$ a function with the property that any image of a subset $S \subseteq V_H$ of size $|S| \geq \dim(\Lambda_v)$ spans Λ_v , i.e., $\text{span}(\lambda_v(u) : u \in S) = \Lambda_v$.

Let \mathcal{M} be the set of partial matchings in G that contain no edges from $V_P \times V_P$. To map partial matchings $\varphi \in \mathcal{M}$ into Λ , we define $\lambda : \mathcal{M} \mapsto \Lambda$ by

$$\lambda(\varphi) = \bigotimes_{v \in V(\varphi) \cap V_P} \lambda_v(\varphi_v) \otimes \bigotimes_{v \in V_P \setminus V(\varphi)} \Lambda_v . \quad (5.23)$$

Recall that each partial matching $\varphi \in \mathcal{M}$ has an associated partial boolean assignment as defined in (2.3). For each clause C , we are interested in the partial matchings $\varphi \in \mathcal{M}$ that match all of $\text{closure}(C) \cup V_{\vec{d}, \vec{\delta}}(C)$ such that φ does not satisfy C . We refer to the set of such matchings as the *zero space* of C and denote it by

$$Z(C) = \{\varphi \in \mathcal{M} \mid V(\varphi) \supseteq (\text{closure}(C) \cup V_{\vec{d}, \vec{\delta}}(C)) \wedge C \upharpoonright_{\varphi} \neq 1\} . \quad (5.24)$$

We associate C with the linear space

$$\lambda(C) = \text{span}(\lambda(\varphi) \mid \varphi \in Z(C)) . \quad (5.25)$$

Note that contradiction is mapped to Λ , i.e., $\lambda(\perp) = \Lambda$.

The following lemma asserts that the span of the axioms $\text{span}(\{\lambda(A) \mid A \in PM(G) \cup \mathcal{A}\})$ is a proper subspace of Λ .

LEMMA 5.11. *If $|\mathcal{A}| \leq (1 + 16\xi)^{w_0/8}$, then $\text{span}(\{\lambda(A) \mid A \in PM(G) \cup \mathcal{A}\}) \subsetneq \Lambda$.*

Deferring the proof of this lemma for now, note this implies that in the refutation π there is a resolution step deriving C from C_0 and C_1 where the subspace of the resolvent is not contained in the span of the subspaces of the premises, or in other words $\lambda(C) \not\subseteq \text{span}(\lambda(C_0), \lambda(C_1))$. The following lemma, which is the heart of the argument, says that this cannot happen while the sets of relevant vertices of the clauses are small.

LEMMA 5.12. *Let C be derived from C_0 and C_1 . If $\max\{|\text{closure}(C_0) \cup V_{\vec{d}, \vec{\delta}}(C_0)|, |\text{closure}(C_1) \cup V_{\vec{d}, \vec{\delta}}(C_1)|, |\text{closure}(C) \cup V_{\vec{d}, \vec{\delta}}(C)|\} \leq r/4$, then $\lambda(C) \subseteq \text{span}(\lambda(C_0), \lambda(C_1))$.*

Deferring the proof of Lemma 5.12 to Section 5.4, we proceed to show how Lemma 5.3 follows from what we have established so far.

PROOF OF LEMMA 5.3. Lemma 5.11 and Lemma 5.12 imply that contradiction cannot be derived while the set of relevant vertices is of size at most $r/4$ and hence any refutation π must contain a clause C with $|\text{closure}(C) \cup V_{\vec{d}, \vec{\delta}}(C)| \geq r/4$. If for such a clause C it holds that $|V_{\vec{d}, \vec{\delta}}(C)| \geq r\xi$, then Lemma 5.3 follows. Otherwise, recall that $\text{closure}(C) = \text{closure}_{r, \nu}(\bar{V}(C))$, for $\nu = (1 - 20\xi)\Delta/2$, and that G' is an (r, Δ', c) -boundary expander by Lemma 5.10, where $\Delta' = (1 + 8\xi)\Delta/2$ and $c = (1 - 12\xi)\Delta/2$. Thus we can apply Lemma 3.3 to G' and get that $|\bar{V}(C)| \geq \min\{r, (r/4 - r\xi) \cdot (c - \nu) / \Delta'\} \geq 3r\xi/2$. As by definition $\bar{V}(C) = (V_{\vec{d}, \vec{\delta}}(C) \cap V_L) \cup \tilde{V}(C)$ and by property 4 of Lemma 5.8 we have that $|\tilde{V}(C)| \leq w_0/8$, we conclude that

$$w_{\vec{d}, \vec{\delta}}(\pi) \geq |V_{\vec{d}, \vec{\delta}}(C)| \geq |V_{\vec{d}, \vec{\delta}}(C) \cap V_L| \geq |\bar{V}(C)| - |\tilde{V}(C)| \geq 3r\xi/2 - w_0/8 \geq r\xi . \quad (5.26)$$

This completes the proof of Lemma 5.3. ■

PROOF OF LEMMA 5.11. Suppose A is a vertex axiom P^ν or a functionality axiom $F_{w, w'}^\nu$ as in (1.1a) and (1.1c). Observe that ν is a heavy vertex for A . Clearly, there are no matchings on ν that do not satisfy A . We conclude that $\lambda(A) = \emptyset$.

Let us consider $A \in \mathcal{A}$. These axioms may span a part of the space Λ but the fraction of the space Λ they span is sufficiently small. We first estimate the dimension of $\lambda(A)$. By definition $\tilde{V}(A) = \{\nu \in V_L : ||N_A(\nu) \cap V_H| - 1/2|N_A(\nu)|| > 4\xi\Delta\}$ and by property 4 of Lemma 5.8 it holds that $|\tilde{V}(A)| \leq w_0/8$. We partition V_P into two sets $U = V_P \setminus (V_{\vec{d}, \vec{\delta}}(A) \setminus \tilde{V}(A))$ and $W =$




Figure 2. Depiction of relations between V_L, V_R, V_P, V_H and the vertex sets in the proof of Lemma 5.12

$V_P \cap (V_{\vec{d}, \vec{\delta}}(A) \setminus \tilde{V}(A))$. Note that all vertices $v \in W$ satisfy that $||N_A(v) \cap V_H| - 1/2|N_A(v)|| \leq 4\xi\Delta$. Using property 2 of Lemma 5.8 we get that

$$\dim \lambda(A) \leq \prod_{v \in U} \dim \Lambda_v \cdot \prod_{v \in W} (|N_G(v) \cap V_H| - |N_A(v) \cap V_H|) \quad (5.27)$$

$$\leq \prod_{v \in U} \dim \Lambda_v \cdot \prod_{v \in W} (1/2|N_G(v)| + 4\xi|N_G(v)| - 1/2|N_A(v)| + 4\xi|N_G(v)|) \quad (5.28)$$

$$= \prod_{v \in U} \dim \Lambda_v \cdot \prod_{v \in W} (1/2(|N_G(v)| - |N_A(v)|) + 8\xi|N_G(v)|) \quad (5.29)$$

$$\leq \prod_{v \in U} \dim \Lambda_v \cdot \prod_{v \in W} (1/2(\Delta_G(v) - d_v) + \delta_v/8) \quad (5.30)$$

$$\leq \prod_{v \in U} \dim \Lambda_v \cdot \prod_{v \in W} (\dim \Lambda_v - \delta_v/8), \quad (5.31)$$

where the second to last inequality follows from the fact that $\delta_v = 64\xi|N_G(v)|$ and the last inequality from the definition of $\dim \Lambda_v$.

Note that by property 1 of Lemma 5.8, $|V_P \cap V_{\vec{d}, \vec{\delta}}(A)| \geq w_0/4$ and hence $|W| \geq w_0/8$. We conclude that the fraction of the space Λ that A spans is bounded by

$$\frac{\dim \lambda(A)}{\dim \Lambda} \leq \prod_{v \in W} \frac{\dim \Lambda_v - \delta_v/8}{\dim \Lambda_v} \leq (1 - 16\xi)^{w_0/8}. \quad (5.32)$$

Along with the assumption on $|\mathcal{A}|$, this shows that not all of Λ is spanned by the axioms. ■

5.4 Proof of Lemma 5.12

For conciseness of notation, let us write

$$S_{01} = (\text{closure}(C_0) \cup \text{closure}(C_1)) \cup (V_{\vec{d},\vec{\delta}}(C_0) \cup V_{\vec{d},\vec{\delta}}(C_1)) \quad (5.33)$$

and $S = \text{closure}(C) \cup V_{\vec{d},\vec{\delta}}(C)$. In order to establish Lemma 5.12, we need to show for all $\varphi \in Z(C)$ that

$$\lambda(\varphi) \subseteq \text{span}(\lambda(C_0), \lambda(C_1)) . \quad (5.34)$$

Denote by φ' the restriction of φ to the edges with at least one vertex in $S \cap S_{01}$ and note that C is not satisfied by the partial assignment associated with φ' . Also, observe that if a matching η extends a matching η' , then $\lambda(\eta)$ is a subspace of $\lambda(\eta')$. This is so since for any vertex $v \in V_P \cap (V(\eta) \setminus V(\eta'))$ we have from (5.23) that η' picks up the whole subspace Λ_v while η only gets a single vector. Thus, if we can show that $\lambda(\varphi') \subseteq \text{span}(\lambda(C_0), \lambda(C_1))$, the statement follows since φ extends φ' and hence $\lambda(\varphi) \subseteq \lambda(\varphi')$. For the following argument it may be helpful to refer to Figure 2.

Let $\mathcal{D} = S_{01} \setminus S$ and for a set of matchings $\mathcal{N} \subseteq \mathcal{M}$ let $\lambda(\mathcal{N}) = \text{span}(\{\lambda(\psi) \mid \psi \in \mathcal{N}\})$. In the following we show that there exists a set of matchings $\mathcal{M}_{\mathcal{D}} \subseteq \mathcal{M}$ that do not satisfy C , that cover S_{01} and such that

$$\lambda(\varphi') \subseteq \lambda(\mathcal{M}_{\mathcal{D}}) . \quad (5.35)$$

Before arguing the existence of such a set $\mathcal{M}_{\mathcal{D}}$ let us argue that this would imply the lemma. Observe that by soundness of resolution, no matching in $\mathcal{M}_{\mathcal{D}}$ can satisfy both C_0 and C_1 simultaneously. Fix $\psi \in \mathcal{M}_{\mathcal{D}}$. Without loss of generality, assume that C_0 is not satisfied. Denote by $\psi' \subseteq \psi$ all edges in ψ with at least one vertex in $\text{closure}(C_0) \cup V_{\vec{d},\vec{\delta}}(C_0)$. Clearly, $\psi' \in Z(C_0)$ and hence $\lambda(\psi) \subseteq \lambda(\psi') \subseteq \lambda(C_0)$. Thus, for all matchings $\psi \in \mathcal{M}_{\mathcal{D}}$ we have that $\lambda(\psi) \subseteq \text{span}(\lambda(C_0), \lambda(C_1))$. Combining with (5.35), we get that

$$\lambda(\varphi') \subseteq \lambda(\mathcal{M}_{\mathcal{D}}) \subseteq \text{span}(\lambda(C_0), \lambda(C_1)) \quad (5.36)$$

and hence the lemma follows.

In the remainder, we show how to construct the set $\mathcal{M}_{\mathcal{D}}$. Observe that all vertices $v \in \mathcal{D}$ are light vertices of C . Using property 3 from Lemma 5.8 we get that for all $v_r \in \mathcal{D} \cap V_R$ there are at most

$$|N_C(v_r) \cap V_H| \leq 1/2|N_C(v_r)| + 4\xi|N_G(v_r)| \leq 1/2(d_{v_r} - \delta_{v_r} + 8\xi|N_G(v_r)|) \quad (5.37)$$

mappings of v_r to a vertex in $N_G(v_r) \cap V_H$ that satisfy the clause C . Similarly, using property 4 from Lemma 5.8 and the fact that $\mathcal{D} \cap \tilde{V}(C) = \emptyset$ we see that for all $v_\ell \in \mathcal{D} \cap V_L$ there are at most

$$|N_C(v_\ell) \cap V_H| \leq 1/2|N_C(v_\ell)| + 4\xi|N_G(v_\ell)| \leq 1/2(d_{v_\ell} - \delta_{v_\ell} + 8\xi\Delta) \quad (5.38)$$

mappings of v_ℓ to a vertex in $N_G(v_\ell) \cap V_H$ that satisfy the clause C .

For a set of vertices $W \subseteq V_P \cup V_H$, let $\Lambda_W = \bigotimes_{w \in W \cap V_P} \Lambda_w$ and for a set $U \subseteq V(G)$ let λ^U be the projection of λ to the space Λ_U or in other words

$$\lambda^U(\eta) = \bigotimes_{v \in V(\eta) \cap V_P \cap U} \lambda_v(\eta_v) \otimes \bigotimes_{v \in (V_P \cap U) \setminus V(\eta)} \Lambda_v . \quad (5.39)$$

We extend the notation to sets of matchings as previously for λ . In order to establish (5.35), we have to argue that $\lambda^{\mathcal{D} \setminus V(\varphi')}(\mathcal{M}_{\mathcal{D}})$ spans the space $\Lambda_{\mathcal{D} \setminus V(\varphi')}$. At this point, we deviate from the *FPHP*(G) proof. Note that we only have expansion for the vertices V_L into V_H but \mathcal{D} may also contain vertices from V_R . Thus we cannot apply the argument from Section 4 to all vertices.

Instead, we split the argument into 2 separate parts. First, by an argument similar to the lower bound proof of the *FPHP*(G) formulas, we show that vertices in $\mathcal{D} \cap V_L$ can be matched in many ways. This will in particular imply that $\lambda^{(\mathcal{D} \cap V_L) \setminus V(\varphi')}(\mathcal{M}_{\mathcal{D}})$ spans all of $\Lambda_{(\mathcal{D} \cap V_L) \setminus V(\varphi')}$. After that we consider the vertices in $\mathcal{D} \cap V_R$. As these vertices have very high degree, there are always enough neighbours they can be matched to and therefore $\lambda^{(\mathcal{D} \cap V_R) \setminus V(\varphi')}(\mathcal{M}_{\mathcal{D}})$ spans all of $\Lambda_{(\mathcal{D} \cap V_R) \setminus V(\varphi')}$. Note that this second argument is essentially the span argument from [30].

Consider the vertex set $\mathcal{D} \cap V_L$. Note that $\mathcal{D} \cap V_L$ is completely outside the closure(C). Since, by assumption, the cardinality of closure(C) is upper bounded by $r/4$ and $|\mathcal{D} \cap V_L| \leq |S_{01}| \leq r/2$, by Lemma 3.4 we get that

$$|\partial_{G' \setminus (\text{closure}(C) \cup N_{G'}(\text{closure}(C)))}(\mathcal{D} \cap V_L)| \geq 1/2(1 - 20\xi)\Delta|\mathcal{D} \cap V_L| . \quad (5.40)$$

By an averaging argument, there is a $v \in \mathcal{D} \cap V_L$ that has at least $(1 - 20\xi)\Delta/2$ unique neighbours in $\partial_{G' \setminus (\text{closure}(C) \cup N_{G'}(\text{closure}(C)))}(\mathcal{D} \cap V_L)$. By iterating this argument on $(\mathcal{D} \cap V_L) \setminus \{v\}$ we get a partition $V_{v_1} \dot{\cup} V_{v_2} \dots \dot{\cup} V_{v_{|\mathcal{D} \cap V_L|}}$ of the neighbourhood $\mathcal{D} \cap V_L$. The key properties of this partition are that every vertex $v_\ell \in \mathcal{D} \cap V_L$ can independently be matched to any vertex in V_{v_ℓ} and each set is of size at least $|V_{v_\ell}| \geq (1 - 20\xi)\Delta/2$. Using (5.38), we have that each vertex $v_\ell \in \mathcal{D} \cap V_L$ can be matched to at least

$$\begin{aligned} 1/2(1 - 20\xi)\Delta - 1/2(d_{v_\ell} - \delta_{v_\ell} + 8\xi\Delta) &= 1/2(\Delta - d_{v_\ell} + \delta_{v_\ell} - 28\xi\Delta) \\ &\geq 1/2(\Delta_G(v_\ell) - d_{v_\ell} + \delta_{v_\ell}/2) \end{aligned} \quad (5.41)$$

many vertices in V_{v_ℓ} without satisfying C . Denote these vertices by V'_{v_ℓ} . As in section Section 4, we would like to conclude that every vertex has many choices of vertices it can independently be mapped to and therefore there are enough matchings to span the space $\Lambda_{\mathcal{D} \cap V_L}$. Unfortunately this argument does not work since vertices in V'_{v_ℓ} can be matched in φ' and are hence not available to be matched to v_ℓ , so there might be too few matchings of v_ℓ to span the whole space Λ_{v_ℓ} .

We could attempt to overcome this problem by removing all edges in φ' with a vertex in one of the sets V'_{v_ℓ} . This allows us to independently match all the vertices in $\mathcal{D} \cap V_L$ to sufficiently many neighbours. Regrettably, this edge removal strategy turns out to be too aggressive: it can

```

1:  procedure ExtendedMatching( $T, \psi, V_{v_1}, V_{v_2}, \dots, V_{v_{|T|}}$ )           ▶ extend  $\psi$  to domain  $T$ 
2:      if  $T \setminus V(\psi) \neq \emptyset$                                        ▶ still need to extend  $\psi$ 
3:      then
4:           $\mathcal{M} \leftarrow \emptyset$ 
5:           $v_\ell \leftarrow_{\text{any}} T \setminus V(\psi)$ 
6:          for  $w \in V_{v_\ell}$                                                    ▶  $v_\ell$  can be matched to  $w$ 
7:          do
8:               $\psi' \leftarrow \psi$ 
9:              if  $\exists w'$  such that  $\{w, w'\} \in \psi$  then
10:                  $\psi' \leftarrow \psi' \setminus \{w, w'\}$                        ▶ remove  $w$  from the matching
11:                  $\psi' \leftarrow \psi' \cup \{v_\ell, w\}$                          ▶ match  $v_\ell$  to  $w$ 
12:                  $\mathcal{M} = \mathcal{M} \cup \text{ExtendedMatching}(T, \psi', V_{v_1}, V_{v_2}, \dots, V_{v_{|T|}})$ 
13:             return  $\mathcal{M}$ 
14:      else
15:          return  $\psi$ 

```

Algorithm 1. Extend Matching

occur that a vertex from $S_{01} \cap V_R$, previously matched by φ' , now has no neighbour available to be matched to. Fortunately, this only happens to vertices that were matched in φ' . The solution that suggests itself is to remove edges from φ' in a “lazy” manner: only remove an edge $\{u, v\}$ from φ' when one of the vertices should be matched to some $v_\ell \in V_L$. This ensures that no vertex in V_R that was previously matched by φ' is suddenly unmatched. This is the main idea of Algorithm 1 which takes care of the necessary edge removals.

Let $\mathcal{M}_{\mathcal{D} \cap V_L} = \text{EXTENDMATCHING}(\mathcal{D} \cap V_L, \varphi', V'_{v_1}, \dots, V'_{v_{|\mathcal{D} \cap V_L|}})$. Note that the algorithm terminates on this input as the sets $V'_{v_1}, V'_{v_2}, \dots, V'_{v_{|\mathcal{D} \cap V_L|}}$ are disjoint. Let us establish some claims regarding $\mathcal{M}_{\mathcal{D} \cap V_L}$.

The first claim states that the algorithm cannot remove edges from φ' with a vertex in $S \cap S_{01} \cap V_L$. This is important as we want to get matchings that are defined on all of $S_{01} \cap V_L$. As the algorithm only tries to match vertices in $\mathcal{D} \cap V_L = (S \setminus S_{01}) \cap V_L$, we must ensure that the edges in φ' with an endpoint in $S \cap S_{01} \cap V_L$ are not erased. Note that all edges that are removed by the algorithm have an endpoint in the neighbourhood of $\mathcal{D} \cap V_L$. Hence it suffices to show that the vertices from $S \cap S_{01} \cap V_L$ are not matched to a vertex in the neighbourhood of $\mathcal{D} \cap V_L$.

CLAIM 5.13. *The matching φ' contains no edge $\{w, w'\}$ such that*

$$w \in N_{G' \setminus (\text{closure}(C) \cup N_{G'}(\text{closure}(C)))}(\mathcal{D} \cap V_L)$$

and $w' \in S \cap S_{01} \cap V_L$.

PROOF. Suppose there is an edge $\{w, w'\} \in \varphi'$ for w, w' as in the lemma statement. As $S \cap S_{01} \cap V_L \subseteq \text{closure}(C)$, we see that $w \in N_{G'}(\text{closure}(C))$. But this is a contradiction since w is not in the graph $G' \setminus (\text{closure}(C) \cup N_{G'}(\text{closure}(C)))$. ■

Next, we consider edges in φ' with a vertex in the set $V_P \cap V_R$. Observe that if the algorithm removed such an edge, then the linear space associated with the new matching would differ from the original space in a non-trivial way. Fortunately, this cannot happen.

CLAIM 5.14. *All matchings $\psi \in \mathcal{M}_{\mathcal{D} \cap V_L}$ cover the set $S_{01} \cap V_L$ and an edge $e \in V_L \times (V_P \cap V_R)$ is contained in ψ if and only if it is contained in φ' . Furthermore, if a vertex $v \in V_R$ is matched in φ' , then it is matched in every $\psi \in \mathcal{M}_{\mathcal{D} \cap V_L}$.*

PROOF. By Claim 5.13, Algorithm 1 never removes edges from φ' that are incident to a vertex in $S_{01} \cap S \cap V_L$. As φ' covers all of $S_{01} \cap S \cap V_L$, it follows that every $\psi \in \mathcal{M}_{\mathcal{D} \cap V_L}$ also covers the set $S \cap S_{01} \cap V_L$. Furthermore, the algorithm ensures that every $\psi \in \mathcal{M}_{\mathcal{D} \cap V_L}$ covers the set $\mathcal{D} \cap V_L = (S_{01} \setminus S) \cap V_L$. Combining these statements we see that every matching $\psi \in \mathcal{M}_{\mathcal{D} \cap V_L}$ covers $S_{01} \cap V_L$.

We observe that all edges in φ' that may be deleted by the algorithm must have an endpoint in one of the sets V'_{v_ℓ} and all these sets are contained in $V_H \cap V_R$. As the graph is bipartite (with bipartition $V_L \dot{\cup} V_R$) and the set \mathcal{M} does not contain matchings with edges from $V_P \times V_P$, we see that vertices from $V_P \cap V_R$ can only be matched to vertices in $V_H \cap V_L$. Therefore the algorithm cannot change edges in φ' with an endpoint in $V_P \cap V_R$. This implies that if an edge $e \in V_L \times (V_P \cap V_R)$ is in φ' , then it is also in ψ . For the other direction, observe that since the algorithm can only add edges to ψ with an endpoint in $V_H \cap V_R$, and since the graph is bipartite, the algorithm does not add an edge from $V_L \times (V_P \cap V_R)$.

Finally, the fact that all matched vertices $v \in V_R$ in φ' are also matched in every $\psi \in \mathcal{M}_{\mathcal{D} \cap V_L}$ follows from the “lazy” removal of edges from φ' . ■

We can now show that our set of matchings spans the appropriate space when projected to V_L . Note that for a matching η it holds that $\lambda(\eta) = \lambda^U(\eta) \otimes \lambda^{V_P \setminus U}(\eta)$ for any set U , but the same does not hold for sets of matchings: span does not commute with tensor.

CLAIM 5.15. $\lambda^{V_L}(\varphi') \subseteq \lambda^{V_L}(\mathcal{M}_{\mathcal{D} \cap V_L})$

PROOF. Let us write

$$\lambda^{V_L}(\varphi') = \lambda^{V(\varphi') \cap V_L}(\varphi') \otimes \Lambda_{V_L \setminus V(\varphi')} \quad (5.42)$$

$$\begin{aligned} &= \lambda^{V(\varphi') \cap \mathcal{D} \cap V_L}(\varphi') \otimes \lambda^{(V(\varphi') \cap V_L) \setminus \mathcal{D}}(\varphi') \otimes \\ &\quad \Lambda_{(\mathcal{D} \cap V_L) \setminus V(\varphi')} \otimes \Lambda_{V_L \setminus (\mathcal{D} \cup V(\varphi'))} . \end{aligned} \quad (5.43)$$

Note that no matching $\psi \in \mathcal{M}_{\mathcal{D} \cap V_L}$ covers any of the vertices in $V_L \setminus (\mathcal{D} \cup V(\varphi'))$. This holds as the algorithm can only add edges from the set $(\mathcal{D} \cap V_L) \times (V_H \cap V_R)$. Hence we can write

$$\lambda^{V_L}(\mathcal{M}_{\mathcal{D} \cap V_L}) = \lambda^{V_L \cap (\mathcal{D} \cup V(\varphi'))}(\mathcal{M}_{\mathcal{D} \cap V_L}) \otimes \Lambda_{V_L \setminus (\mathcal{D} \cup V(\varphi'))} . \quad (5.44)$$

Thus we can ignore the space $\Lambda_{V_L \setminus (\mathcal{D} \cup V(\varphi'))}$ for the remainder of this argument. From the algorithm it should be evident that

$$\lambda^{(\mathcal{D} \cap V_L) \setminus V(\varphi')}(\mathcal{M}_{\mathcal{D} \cap V_L}) = \Lambda_{(\mathcal{D} \cap V_L) \setminus V(\varphi')} \quad (5.45)$$

as every vertex in $v \in (\mathcal{D} \cap V_L) \setminus V(\varphi')$ is independently matched to every vertex in V'_v of size $|V'_v| \geq 1/2(\Delta_G(v) - d_v + \delta_v/2)$. As the dimension of $\dim(\Lambda_v) = 1/2(\Delta_G(v) - d_v + \delta_v/2)$, we conclude that $\Lambda_{(V_L \cap \mathcal{D}) \setminus V(\varphi')}$ is spanned.

To continue the argument, we need the following equivalence relation on matchings. Two matchings $\psi, \psi' \in \mathcal{M}_{\mathcal{D} \cap V_L}$ are equivalent on a vertex set V if they match the vertices in V in the same way, that is, for $v \in V$ we have that $\psi_v = \psi'_v$. We denote the equivalence class with respect to the vertex set V over $\mathcal{M}_{\mathcal{D} \cap V_L}$ of a matching $\psi \in \mathcal{M}_{\mathcal{D} \cap V_L}$ by $\{\psi\}_V$.

We want to show that for every $\psi \in \mathcal{M}_{\mathcal{D} \cap V_L}$ it holds that

$$\lambda^{V(\varphi') \cap \mathcal{D} \cap V_L}(\varphi') \subseteq \text{span}(\lambda^{V(\varphi') \cap \mathcal{D} \cap V_L}(\psi') \mid \psi' \in \{\psi\}_{(\mathcal{D} \cap V_L) \setminus V(\varphi')}) . \quad (5.46)$$

Note that in combination with (5.45) we get that

$$\lambda^{\mathcal{D} \cap V_L}(\varphi') = \lambda^{(\mathcal{D} \cap V_L) \setminus V(\varphi')}(\varphi') \otimes \lambda^{V(\varphi') \cap \mathcal{D} \cap V_L}(\varphi') \quad (5.47)$$

$$= \Lambda_{(\mathcal{D} \cap V_L) \setminus V(\varphi')} \otimes \lambda^{V(\varphi') \cap \mathcal{D} \cap V_L}(\varphi') \quad (5.48)$$

$$\subseteq \lambda^{\mathcal{D} \cap V_L}(\mathcal{M}_{\mathcal{D} \cap V_L}) . \quad (5.49)$$

We prove (5.46) by induction on subsets of $V(\varphi') \cap \mathcal{D} \cap V_L$. The statement clearly holds for the empty set. Fix $U \subseteq V(\varphi') \cap \mathcal{D} \cap V_L$ and a vertex $u \in U$. By induction, we may assume that

$$\lambda^{U \setminus \{u\}}(\varphi') \subseteq \text{span}(\lambda^{U \setminus \{u\}}(\psi') \mid \psi' \in \{\psi\}_{(\mathcal{D} \cap V_L) \setminus V(\varphi')}) . \quad (5.50)$$

We want to show that the statement also holds for the set U . Note that $\lambda^U(\varphi') = \lambda^{U \setminus \{u\}}(\varphi') \otimes \lambda_u(\varphi'_u)$. Further,

$$\text{span}(\lambda^U(\psi') \mid \psi' \in \{\psi\}_{(\mathcal{D} \cap V_L) \setminus V(\varphi')}) = \quad (5.51)$$

$$\begin{aligned} & \text{span}(\lambda^{U \setminus \{u\}}(\psi') \otimes \\ & \text{span}(\lambda_u(\eta) \mid \eta \in \{\psi'\}_{((\mathcal{D} \cap V_L) \setminus V(\varphi')) \cup (U \setminus \{u\})}) \mid \psi' \in \{\psi\}_{(\mathcal{D} \cap V_L) \setminus V(\varphi')}) . \end{aligned} \quad (5.52)$$

Suppose that for every $\psi' \in \{\psi\}_{(V_L \cap \mathcal{D}) \setminus V(\varphi')}$ it holds that

$$\lambda_u(\varphi'_u) \subseteq \text{span}(\lambda_u(\eta_u) \mid \eta \in \{\psi'\}_{((\mathcal{D} \cap V_L) \setminus V(\varphi')) \cup (U \setminus \{u\})}) . \quad (5.53)$$

Then, continuing from above, we see that

$$\text{span}(\lambda^U(\psi') \mid \psi' \in \{\psi\}_{(\mathcal{D} \cap V_L) \setminus V(\varphi')}) \quad (5.54)$$

$$\supseteq \text{span}(\lambda^{U \setminus \{u\}}(\psi') \mid \psi' \in \{\psi\}_{(\mathcal{D} \cap V_L) \setminus V(\varphi')}) \otimes \lambda_u(\varphi'_u) \quad (5.55)$$

$$\supseteq \lambda^{U \setminus \{u\}}(\varphi') \otimes \lambda_u(\varphi'_u) \quad (5.56)$$

$$= \lambda^U(\varphi') , \quad (5.57)$$

where the second inclusion holds by the induction hypothesis (5.50). Thus, to show the statement for U we just need to show (5.53). To this end, fix a matching $\psi' \in \{\psi\}_{(\mathcal{D} \cap V_L) \setminus V(\varphi')}$. Note that if there is a matching $\eta \in \{\psi'\}_{((\mathcal{D} \cap V_L) \setminus V(\varphi')) \cup (U \setminus \{u\})}$ such that $\eta_u = \varphi'_u$, then we are done. Otherwise, Algorithm 1 removed the edge that matched the vertex u in φ' . Hence the vertex u is matched by the procedure to at least $|V'_u| \geq 1/2(\Delta_G(v) - d_v + \delta_v/2)$ different vertices. As the dimension of $\Lambda_u = 1/2(\Delta_G(v) - d_v + \delta_v/2)$, we see that all of the space is spanned. We conclude that (5.53) holds.

What remains is to argue that for every $\psi \in \mathcal{M}_{\mathcal{D} \cap V_L}$ it holds that

$$\lambda^{(V(\varphi') \cap V_L) \setminus \mathcal{D}}(\varphi') \subseteq \text{span}(\lambda^{(V(\varphi') \cap V_L) \setminus \mathcal{D}}(\psi') \mid \psi' \in \{\psi\}_{\mathcal{D} \cap V_L}) . \quad (5.58)$$

The argument goes along the same lines as for the vertices in $V(\varphi') \cap \mathcal{D} \cap V_L$ and we thus omit it.

We can then combine (5.47) and (5.58) to conclude the claim. ■

Observe that the matchings in $\mathcal{M}_{\mathcal{D} \cap V_L}$ are not necessarily extensions of φ' . This is not a problem, however, since the matchings only differ in edges that contain vertices which either do not show up in the linear space or for which the whole linear space associated to the vertex is spanned. Furthermore, vertices from $\mathcal{D} \cap V_H \cap V_L$ are matched to many vertices even though a single vertex would have been sufficient.

It remains only to show that every matching $\psi \in \mathcal{M}_{\mathcal{D} \cap V_L}$ can be extended in many ways to the set $\mathcal{D} \cap V_R$. Fix a matching $\psi \in \mathcal{M}_{\mathcal{D} \cap V_L}$ and recall that these are defined on $S_{01} \cap V_L$. Note

that by Lemma 5.8, property 2, each $v \in \mathcal{D} \cap V_R$ has at least

$$|N_G(v) \cap V_H| \geq 1/2|N_G(v)| - 4\xi|N_G(v)| \quad (5.59)$$

many neighbours in V_H . Using (5.37) we can now bound the number of matchings that do not satisfy C .

Note that the matching ψ contains at most $|S_{01}| \leq r/2$ many edges. Since G is bipartite, this implies that for any $v \in V_R$ at most $r/2$ neighbours are already matched. Observe that some vertex $v \in \mathcal{D} \cap V_H \cap V_R$ may have been matched by Algorithm 1. As these vertices are not associated with a linear space, we only need to match these vertices with a single vertex and hence we can just leave them matched as in ψ . Further, by Claim 5.14, we see that the vertices in $\mathcal{D} \cap V_P \cap V_R$ were not matched by Algorithm 1. All these will be matched in many ways as needed: If $v \in \mathcal{D} \cap V_R$ is not matched by ψ , then by (5.59) and (5.37) it can be matched to at least

$$\begin{aligned} & 1/2(|N_G(v)| - 8\xi|N_G(v)| - d_v + \delta_v - 8\xi|N_G(v)| - r) \\ &= 1/2(\Delta_G(v) - d_v + \delta_v - 16\xi\Delta_G(v) - r) \\ &\geq 1/2(\Delta_G(v) - d_v + \delta_v - 17\xi\Delta_G(v)) \\ &\geq 1/2(\Delta_G(v) - d_v + \delta_v/2) \end{aligned} \quad (5.60)$$

many vertices without satisfying the clause C . Note that in (5.60) we used the assumption that $\Delta_G(v) \geq r\xi$ for $v \in V_R$. As we have that $\dim(\Lambda_v) = 1/2(\Delta_G(v) - d_v + \delta_v/2)$, we conclude that the extensions of ψ can span the linear space $\Lambda_{(\mathcal{D} \cap V_R) \setminus V(\varphi')}$. Hence, by extending each $\psi \in \mathcal{M}_{\mathcal{D} \cap V_L}$, we get a set of matchings $\mathcal{M}_{\mathcal{D}}$, which do not satisfy the clause C , are defined on S_{01} and $\lambda(\varphi') \subseteq \lambda(\mathcal{M}_{\mathcal{D}})$. This establishes the lemma.

6. Concluding Remarks

In this work, we extend the pseudo-width method developed by Razborov [29, 30] for proving lower bounds on severely overconstrained CNF formulas in resolution. In particular, we establish that pigeonhole principle formulas and perfect matching formulas over highly unbalanced bipartite graphs remain exponentially hard for resolution even when these graphs are sparse. This resolves an open problem in [29].

The main technical difference in our work compared to [29, 30] goes right to the heart of the proof, where one wants to argue that resolution in small pseudo-width cannot make progress towards a derivation of contradiction. Here Razborov uses the global symmetry properties of the formula, whereas we resort to a local argument based on graph expansion. This argument needs to be carefully combined with a graph closure operation as in [2, 4] to ensure that the residual graph always remains expanding as matched pigeons and their neighbouring holes are removed. It is this change of perspective that allows us to prove lower bounds for sparse

bipartite graphs with the size m of the left-hand side (i.e., the number of pigeons) varying all the way from linear to exponential in the size n of the right-hand size (i.e., the number of pigeonholes), thus covering the full range between [8] on the one hand and [24, 29, 30] on the other.

One shortcoming of our approach is that the sparse expander graphs are required to have very good expansion—for graphs of left degree Δ , the size of the set of unique neighbours of any not too large left vertex set has to scale like $(1 - o(1))\Delta$. We would like to prove that graph PHP formulas are hard also for graphs with constant expansion $(1 - \varepsilon)\Delta$ for some $\varepsilon > 0$, but there appear to be fundamental technical barriers to extending our lower bound proof to this setting—in particular, since we can only obtain a very small gap in the filter lemma.

Another intriguing problem left over from [29] is to determine the true resolution complexity of weak PHP formulas over complete bipartite graphs $K_{m,n}$ as $m \rightarrow \infty$. The best known upper bound from [11] is $\exp(O(\sqrt{n \log n}))$, whereas the lower bound in [29, 30] is $\exp(\Omega(\sqrt[3]{n}))$. It does not seem unreasonable to hypothesize that $\exp(\Omega(\sqrt[2]{n}))$ should be the correct lower bound (ignoring lower-order terms), but establishing such a lower bound again appears to require substantial new ideas.

We believe that one of the main contributions of our work is that it again demonstrates the power of Razborov’s pseudo-width method, and we are currently optimistic that it could be useful for solving other open problems for resolution and other proof systems. Unfortunately pseudo-width is a quite difficult method to apply; we leave it as an open problem to obtain simplifications of the current approach.

For resolution, an interesting question mentioned in [29] is whether pseudo-width can be useful to prove lower bounds for formulas that encode the Nisan–Wigderson generator [2, 28]. Since the clauses in such formulas encode local constraints, we hope that techniques from our paper could be helpful. Another long-standing open problem is to prove lower bounds on proofs in resolution that k -clique free sparse graph do not contain k -cliques, where the expected length lower bound would be $n^{\Omega(k)}$. Here we only know weakly exponential lower bounds for quite dense random graphs [6, 22], although an asymptotically optimal $n^{\Omega(k)}$ lower bound has been established in the sparse regime for the restricted subsystem of regular resolution [5].

Finally, we want to highlight that for the stronger proof system *polynomial calculus* [3, 14] no lower bounds on proof size are known for PHP formulas with $m \geq n^2$ pigeons. It would be very interesting if some kind of “pseudo-degree” method could be developed that would finally lead to progress on this problem.

Acknowledgements

First and foremost, we are most grateful to Alexander Razborov for many discussions about pseudo-width, graph closure, and other mysteries of the universe. We also thank Paul Beame

and Johan Håstad for useful discussions, and Jonah Brown-Cohen for helpful references on expander graphs. We gratefully acknowledge the feedback from participants of the Dagstuhl workshop 19121 *Computational Complexity of Discrete Problems* in March 2019 and would like to thank the referees for their thorough reviews.

References

- [1] Michael Alekhovich. Mutilated chessboard problem is exponentially hard for resolution. *Theoretical Computer Science*, 310(1–3):513–525, January 2004. DOI (4)
- [2] Michael Alekhovich, Eli Ben-Sasson, Alexander A. Razborov, and Avi Wigderson. Pseudorandom generators in propositional proof complexity. *SIAM Journal on Computing*, 34(1):67–88, 2004. Preliminary version in *FOCS '00*. DOI (6, 8, 15, 43, 44)
- [3] Michael Alekhovich, Eli Ben-Sasson, Alexander A. Razborov, and Avi Wigderson. Space complexity in propositional calculus. *SIAM Journal on Computing*, 31(4):1184–1211, April 2002. Preliminary version in *STOC '00*. DOI (44)
- [4] Michael Alekhovich and Alexander A. Razborov. Lower bounds for polynomial calculus: Non-binomial case. *Proceedings of the Steklov Institute of Mathematics*, 242:18–35, 2003. Preliminary version in *FOCS '01*. DOI ePrint (6, 8, 15, 43)
- [5] Albert Atserias, Ilario Bonacina, Susanna F. de Rezende, Massimo Lauria, Jakob Nordström, and Alexander Razborov. Clique is hard on average for regular resolution. *Journal of the ACM*, 68(4):23:1–23:26, June 2021. Preliminary version in *STOC '18*. DOI (44)
- [6] Paul Beame, Russell Impagliazzo, and Ashish Sabharwal. The resolution complexity of independent sets and vertex covers in random graphs. *Computational Complexity*, 16(3):245–297, October 2007. Preliminary version in *CCC '01*. DOI (44)
- [7] Paul Beame and Toniann Pitassi. Simplified and improved resolution lower bounds. *Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science (FOCS '96)*, pages 274–282, October 1996. DOI (3)
- [8] Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow—resolution made simple. *Journal of the ACM*, 48(2):149–169, March 2001. Preliminary version in *STOC '99*. DOI (3, 4, 44)
- [9] Olaf Beyersdorff, Nicola Galesi, and Massimo Lauria. A lower bound for the pigeonhole principle in tree-like resolution by asymmetric prover-delayer games. *Information Processing Letters*, 110(23):1074–1077, 2010. DOI (4)
- [10] Archie Blake. Canonical Expressions in Boolean Algebra. PhD thesis, University of Chicago, 1937. DOI (2)
- [11] Samuel R. Buss and Toniann Pitassi. Resolution and the weak pigeonhole principle. *11th International Workshop on Computer Science Logic (CSL '97), Selected Papers*, volume 1414 of *Lecture Notes in Computer Science*, pages 149–156. Springer, August 1997. DOI (3, 4, 44)
- [12] Samuel R. Buss and György Turán. Resolution proofs of generalized pigeonhole principles. *Theoretical Computer Science*, 62(3):311–317, December 1988. DOI (3)
- [13] Vašek Chvátal and Endre Szemerédi. Many hard examples for resolution. *Journal of the ACM*, 35(4):759–768, October 1988. DOI (2)
- [14] Matthew Clegg, Jeffery Edmonds, and Russell Impagliazzo. Using the Groebner basis algorithm to find proofs of unsatisfiability. *Proceedings of the 28th Annual ACM Symposium on Theory of Computing (STOC '96)*, pages 174–183, May 1996. DOI (44)
- [15] Stefan S. Dantchev. Resolution width-size trade-offs for the pigeon-hole principle. *Proceedings of the 17th Annual IEEE Conference on Computational Complexity (CCC '02)*, pages 39–43, May 2002. DOI (4)
- [16] Stefan S. Dantchev and Søren Riis. “Planar” tautologies hard for resolution. *Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science (FOCS '01)*, pages 220–229, October 2001. DOI (4)
- [17] Stefan S. Dantchev and Søren Riis. Tree resolution proofs of the weak pigeon-hole principle. *Proceedings of the 16th Annual IEEE Conference on Computational Complexity (CCC '01)*, pages 69–75, June 2001. DOI (4)
- [18] Susanna F. de Rezende, Jakob Nordström, Kilian Risse, and Dmitry Sokolov. Exponential resolution lower bounds for weak pigeonhole principle and perfect matching formulas over sparse graphs. *35th Computational Complexity Conference, CCC 2020, July 28–31, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 169 of *LIPICs*, 28:1–28:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. DOI (1)
- [19] Venkatesan Guruswami, Christopher Umans, and Salil Vadhan. Unbalanced expanders and randomness extractors from Parvaresh-Vardy codes. *Journal of the ACM*, 56(4):20:1–20:34, July 2009. Preliminary version in *CCC '07*. DOI (12, 28, 29)

- [20] Armin Haken. The intractability of resolution. *Theoretical Computer Science*, 39(2-3):297–308, August 1985. DOI (2, 3)
- [21] Dmitry Itsykson, Vsevolod Oparin, Mikhail Slabodkin, and Dmitry Sokolov. Tight lower bounds on the resolution complexity of perfect matching principles. *Fundamenta Informaticae*, 145(3):229–242, August 2016. DOI (4)
- [22] Shuo Pang. Large clique is hard on average for resolution. *Proceedings of the 16th International Computer Science Symposium in Russia (CSR '21)*, volume 12730 of *Lecture Notes in Computer Science*, pages 361–380. Springer, June 2021. DOI (44)
- [23] Toniann Pitassi and Ran Raz. Regular resolution lower bounds for the weak pigeonhole principle. *Combinatorica*, 24(3):503–524, 2004. Preliminary version in *STOC '01*. DOI (4)
- [24] Ran Raz. Resolution lower bounds for the weak pigeonhole principle. *Journal of the ACM*, 51(2):115–138, March 2004. Preliminary version in *STOC '02*. DOI (4, 44)
- [25] Alexander A. Razborov. Improved Resolution Lower Bounds for the Weak Pigeonhole Principle. Technical Report TR01-055, Electronic Colloquium on Computational Complexity (ECCC), July 2001. (4, 16, 26)
- [26] Alexander A. Razborov. Lower bounds for the polynomial calculus. *Computational Complexity*, 7(4):291–324, December 1998. DOI (3)
- [27] Alexander A. Razborov. Proof complexity of pigeonhole principles. *5th International Conference on Developments in Language Theory, (DLT '01), Revised Papers*, volume 2295 of *Lecture Notes in Computer Science*, pages 100–116. Springer, July 2002. DOI (3)
- [28] Alexander A. Razborov. Pseudorandom generators hard for k-DNF resolution and polynomial calculus resolution. *Annals of Mathematics*, 181(2):415–472, March 2015. DOI (44)
- [29] Alexander A. Razborov. Resolution lower bounds for perfect matching principles. *Journal of Computer and System Sciences*, 69(1):3–27, August 2004. Preliminary version in *CCC '02*. DOI (3–6, 26, 27, 31, 43, 44)
- [30] Alexander A. Razborov. Resolution lower bounds for the weak functional pigeonhole principle. *Theoretical Computer Science*, 1(303):233–243, June 2003. DOI (4–6, 8, 12, 13, 28, 38, 43, 44)
- [31] Alexander A. Razborov, Avi Wigderson, and Andrew Chi-Chih Yao. Read-once branching programs, rectangular proofs of the pigeonhole principle and the transversal calculus. *Combinatorica*, 22(4):555–574, 2002. Preliminary version in *STOC '97*. DOI (4)
- [32] Alasdair Urquhart. Hard examples for resolution. *Journal of the ACM*, 34(1):209–219, January 1987. DOI (2)
- [33] Alasdair Urquhart. Resolution proofs of matching principles. *Annals of Mathematics and Artificial Intelligence*, 37(3):241–250, March 2003. DOI (3)
- [34] Alasdair Urquhart. Width versus size in resolution proofs. *Theoretical Computer Science*, 384(1):104–110, September 2007. DOI (5)