INSTITUTE OF MATHEMATICS

THE CZECH ACADEMY OF SCIENCES

Random resolution refutations

Pavel Pudlák
Neil Thapen

Preprint No. 40-2016

PRAHA 2016

Random resolution refutations

Pavel Pudlák and Neil Thapen*

August 5, 2016

Abstract

We study the *random resolution* proof system defined in [5]. This attempts to capture the notion of a resolution refutation that may make mistakes but is correct most of the time. We give some equivalent formalizations of this idea. We prove some upper and lower bounds, and in particular show separations in both directions between polylogarithmic width random resolution and quasipolynomial size resolution, which solves the problem stated in [5].

1 Introduction

The following system for refuting propositional CNFs was introduced in [5]. Let F be a CNF in variables $x_1, \dots, x_n$ and let $0 < \varepsilon < 1$.

Definition 1.1 *An* ε-random resolution distribution*, or* ε-RR distribution*, of F is a probability distribution $\mathcal{D}$ on pairs $(B_i, \Pi_i)_{i \sim \mathcal{D}}$ such that*

1. *for each $i \in \mathcal{D}$, B_i is a CNF in variables $x_1, \dots, x_n$ and Π_i is a resolution refutation of $F \wedge B_i$*
2. *for every $\alpha \in \{0,1\}^n$, $\Pr_{i \sim \mathcal{D}}[B_i$ is satisfied by $\alpha] \geq 1 - \varepsilon$.*

The size *and the* width *of $\mathcal{D}$ are defined respectively as the maximum size and maximum width of the refutations Π_i (if these maxima exist).*

The definition is attributed in [5] to Stefan Dantchev. Its appearance in [5] is ultimately motivated by an open problem in bounded arithmetic, which we will explain in a moment, after we mention some basic properties and equivalent formulations.

It is sound as a refutational system, in the sense that if F has an ε-RR distribution then F is unsatisfiable. To see this, consider any assignment $\alpha \in \{0,1\}^n$. Since $\varepsilon < 1$, there is at least one pair (B_i, Π_i) such that α satisfies B_i and Π_i is a resolution refutation of $F \wedge B_i$.

*Institute of Mathematics, Czech Academy of Sciences. The authors were supported by the ERC Advanced Grant 339691 (FEALORA)

So α cannot also satisfy F, by the soundness of resolution. The system is also complete, since resolution is complete and we can take $\mathcal{D}$ to consist of a single pair (B, Π) where B is any tautology and Π is a (possibly exponential sized) resolution refutation of F.

On the other hand, as defined it is not a propositional proof system in the sense of of Cook and Reckhow [9], because it is defined by a semantic condition that presumably cannot be tested in polynomial time. Nevertheless it makes perfect sense to compare the complexity of proofs in it with proofs in the standard proof systems, in particular with resolution and bounded depth Frege. We prove some results in this direction in this work. Note also that the definition is particular to resolution, and we must take care if we try to generalize it. For example, if we instead define a random Frege distribution system, in which B and Π can contain arbitrary formulas, then we can trivially refute any unsatisfiable F by setting $B = \neg F$.

As with some concepts of probabilistic computation studied in computation complexity theory, one can use the linear programming duality to give an equivalent definition of the system based on probability distributions over inputs rather than over proofs (see Definition 2.5). This is very useful if one needs to prove lower bounds. Another essentially equivalent formulation is in terms of semantic resolution derivations. This means, roughly speaking, that instead of having an auxiliary formula that is satisfied with high probability, we consider semantic derivations with respect to a large subset of inputs, where lines in the proof are clauses. In a sense, this captures better the intuitive idea of a proof with errors.

Let us also mention that while we tend to think of the error ε as something small, there is a simple amplification lemma that allows us to shrink the error at some cost in proof size. Thus, for the questions we are interested in, without loss of generality we can take $\varepsilon = \frac{1}{2}$.

We now turn to connections with bounded arithmetic. One of the central problems there is to show, in the relativized setting where an undefined relation symbol has been added to the language, that the set $\forall\Sigma^b_1(T^i_2)$, consisting of all $\forall\Sigma^b_1$ consequences of the bounded arithmetic theory T^i_2, gets strictly bigger as i increases. It is known that $\forall\Sigma^b_1(T^2_2)$ is bigger than $\forall\Sigma^b_1(T^1_2)$ [7], but there are no higher separations known, and it cannot be ruled out that $\forall\Sigma^b_1(T^2_2)$ is already the same as $\forall\Sigma^b_1(T_2)$.

There is a more-or-less equivalent question in propositional proof complexity: show that the set of polylogarithmic width CNFs with quasipolynomial size refutations in depth i Frege strictly increases as i increases.[1] Here the system $R(\log)$ [13], which can be thought of as depth $\frac{1}{2}$ Frege, corresponds to the theory T^2_2 and corresponding separations are known. That is, we can separate $R(\log)$ from weaker fragments of constant depth Frege, but not from stronger ones (see for example [17]).

Since the problem of separating $\forall\Sigma^b_1(T^2_2)$ from $\forall\Sigma^b_1(T_2)$ has been notoriously open for many years, it was proposed in [5] to consider, in place of T^2_2, a theory of similar strength but of a rather different nature, namely Jeřábek's theory of approximate counting [11]. This theory, called APC_2 in [5], consists of $T^1_2 + \mathrm{sWPHP}(\mathrm{PV}_2)$ where PV_2 is a formalization of the

[1] We emphasize that we are interested in this question for quasipolynomial size proofs. This matches the natural question in bounded arithmetic, and a separation for polynomial size is known [10], using a padded pigeonhole principle $\mathrm{PHP}_{(\log n)^k}$ which has short proofs in some depth i, but is such that the exponential size lower bound for PHP in depth $i+1$ gives a quasipolynomial lower bound for the padded version.

$\mathrm{FP}^{\mathrm{NP}}$ functions and $\mathrm{sWPHP}(\mathrm{PV}_2)$ expresses that no such function is a surjection from $[a]$ onto $[2a]$, for any a.

A separation of $\forall\Sigma^b_1(\mathrm{APC}_2)$ from $\forall\Sigma^b_1(T_2)$ is still open, but [5] does show separations for subtheories of APC_2. However, they leave the following as an open problem: show an $\forall\Sigma^b_1$ separation between $T^1_2 + \mathrm{sWPHP}(\mathrm{PV}_1)$ and T_2, where $\mathrm{sWPHP}(\mathrm{PV}_1)$ is the surjective weak pigeonhole principle only for FP functions.

A reason why this problem was interesting was the following proposition, which reduces it to a natural-looking question about the complexity of propositional proofs.

Proposition 1.2 ([5]) *Suppose $T^1_2 + \mathrm{sWPHP}(\mathrm{PV}_1) \vdash \forall n \exists y < t(n)\, \theta(n, y)$, where θ is sharply bounded, and everything is relativized. Then the propositional translation $\langle \forall y < t(n)\, \neg\theta(n, y)\rangle$ has a polylogarithmic width $1/q$-RR distribution, for q any quasipolynomial in n.*

Hence one way to solve the open problem mentioned above would be to prove width lower bounds on random resolution, for any CNF which has quasipolynomial size constant depth Frege refutations.

However, the problem specifically about $T^1_2 + \mathrm{sWPHP}(\mathrm{PV}_1)$ was solved in [2], without using Proposition 1.2 or proving a lower bound on random resolution — instead the proof used the properties of the formula $\mathrm{sWPHP}(\mathrm{PV}_1)$ in an essential way.

In this paper we solve the more general problem about random resolution, by proving that the propositional translation of the *coloured polynomial local search* principle CPLS [15], which has polynomial size resolution refutations, does not have narrow $1/2$-random resolution distributions. Previously, lower bounds have only been known for treelike random resolution refutations [5] or for relatively small errors ε [14].

The proof is based on a lemma that looks like a rudimentary version of the switching lemmas used in propositional proof complexity (see the discussion at the start of Section 4). Although this does not give a new separation in bounded arithmetic, we believe that the result is interesting for other reasons. It has been conjectured that in order to separate higher fragments of T_2, we only need switching lemmas for certain more complicated tautologies, similar to CPLS (see for example [18]). Nevertheless, all attempts in this direction have failed so far because of the complexity of the associated combinatorial problems. Our proof gives us some hope that eventually it will be possible to prove such lemmas.

The paper is organized as follows. In Section 2 we fix our notation, prove some basic facts about random resolution distributions and present some equivalent or almost equivalent definitions, in terms of refutations with respect to a distribution over assignments, semantic resolution refutations, and refutations with random extension clauses.

In Section 3 we prove some upper bounds. We show that random 3-CNFs (Proposition 3.1) and the retraction weak pigeonhole principle (Proposition 3.3) have narrow, and thus quasipolynomial size, $1/2$-RR distributions, while they require exponential sized refutations in standard resolution. We also discuss random resolution refutations with respect to the the uniform distribution on assignments (defined below) and observe that this system is unreasonably strong.

In Section 4 we describe our approach to lower bounds, and then use it first to prove that the bit pigeonhole principle requires exponential size 1/2-RR distributions (Theorem 4.3).

In Section 5, using the same approach, we show our main result, that there is a family of logarithmic width CNFs which have polynomial size resolution refutations, but do not have polylogarithmic width 1/2-RR distributions (Theorem 5.10). This answers the open problem about random resolution posed in [5].

In Section 6 we prove a form of feasible interpolation for treelike random resolution (Theorem 6.1) and derive size lower bounds on treelike RR distributions for a family of 3-CNFs (Corollary 6.2).

Finally in Section 7 we show that the usual formalization of the pigeonhole principle requires exponential size 1/2-RR distributions (Theorem 7.1).

Acknowledgments. We would like to thank Pavel Hrubeš and Jan Krajíček for discussions and valuable suggestions. In particular, the proof of Proposition 3.4 generalizes an idea of Hrubeš.

2 Basic properties and alternative definitions

We first introduce some notation. We identify CNF formulas with sets of clauses. We will use 0 (false) and 1 (true) to represent truth values. For a formula F and an assignment α of truth values to its variables, we denote by $F[\alpha]$ the truth value to which the formula is evaluated by α. If ρ is a partial assignment, we denote by F^ρ the formula obtained by substituting ρ into F and simplifying the formula (that is, replacing a conjunction by 0 if one conjunct is 0, etc.).

The *width* of a clause is the number of literals it contains. The *width* and *size* of a refutation are respectively the width of its widest clause and the total number of clauses. A k-CNF is a CNF in which every clause has width at most k.

We will often use the notation

$$p_1 \wedge \cdots \wedge p_r \to q_1 \vee \cdots \vee q_s$$

to stand for the clause $\neg p_1 \vee \cdots \vee \neg p_r \vee q_1 \vee \cdots \vee q_s$, where $p_1, \ldots, p_r, q_1, \ldots, q_r$ can be any literals. In this notation the resolution rule can, for example, have the form: from $A \wedge p \to C$ and $A \wedge \neg p \to D$ conclude $A \to C \vee D$, where p is a literal, A is a conjunction of literals and C and D are clauses.

If p is a literal, we will sometimes write $p = 1$ instead of the literal p and $p = 0$ instead of the literal $\neg p$. Similarly we will write $p \neq 1$ or $p \neq 0$ to mean respectively $\neg p$ or p. If $p_1, \ldots, p_r$ are literals and $\beta \in \{0,1\}^r$ we write $\bar{p} = \beta$ to stand for the conjunction $\bigwedge_{1 \le i \le k} p_i = \beta_i$ where each conjunct is formally either p_i or its negation, as above; and $\bar{p} \neq \beta$ to stand for the disjunction $\bigvee_{1 \le i \le k} p_i \neq \beta_i$. The following observation will be useful:

Lemma 2.1 *The CNF $\bigwedge\{\bar{p} \neq \beta : \beta \in \{0,1\}^r\}$ has a refutation of width r, using $2^r - 1$ resolution steps.*

We write $[n]$ for $\{0, \ldots, n-1\}$. When we formalize combinatorial principles as CNFs, if the principle involves a function $f : [n] \to [m]$ we will often formalize f by introducing variables for its "bit-graph". That is, for each $x < n$ we introduce $\log m$ variables $(f(x))_0, \ldots, (f(x))_{\log m - 1}$ representing the value of $f(x)$ in binary. For the sake of simplicity, in this situation we will assume that m is a power of 2. For $y < m$ we will write $f(x) = y$ to stand for the conjunction $\bigwedge_i (f(x))_i = \beta_i$, where $\beta \in \{0,1\}^{\log m}$ is y written in binary, and we will write $f(x) \neq y$ for the disjunction $\bigvee_i (f(x))_i \neq \beta_i$.

A standard example here is the bit pigeonhole principle BPHP_n (see Section 4 below). Suppose $n = 2^k$. Then BPHP_n is a contradictory propositional CNF asserting that f is an injection from $[n+1]$ to $[n]$. In our notation, it consists of clauses

$$f(x) \neq y \vee f(x') \neq y$$

for all $x < x' < n+1$ and all $y < m$. Each clause has width $2k$.

2.1 Random resolution distributions

In the rest of this section, let F be a CNF in variables $x_1, \ldots, x_n$ and let $0 < \varepsilon < 1$. Our definition of the size of an ε-RR distribution above does not take into account the size of the sample space. We show now that the size of the sample space can be bounded, at the cost of slightly increasing the error ε.

Lemma 2.2 *If F has an ε-RR distribution, then it also has 2ε-RR distribution where the sample space has size $O(n/\varepsilon)$.*

Proof. Given the probability distribution $(B_i, \Pi_i)_{i \sim \mathcal{D}}$, consider m random samples $i_1, \ldots, i_m$. By the Chernoff bound, for each α the probability that $B_{i_j}[\alpha] \neq 1$ for more than a fraction 2ε of the m samples is at most $\mathrm{e}^{-\varepsilon m/3}$. The probability that this happens for *some* α is at most 2^n times larger. Hence if $m > \frac{3 \ln 2}{\varepsilon} n$, there is nonzero probability that this does not happen for any α. We pick a set of samples for which it does not happen, and let the new distribution $\mathcal{D}'$ be given by the uniform distribution on $\{i_1, \ldots, i_m\}$. ∎

We next show an "amplification" result, that we can decrease the error ε at the cost of increasing the width and size.

Lemma 2.3 *Suppose F has an ε-RR distribution of width w and size s. Then for every $k \geq 1$ it also has an ε^k-RR distribution of width at most kw and size $O(s^k)$.*

Proof. Given $(B_i, \Pi_i)_{i \sim \mathcal{D}}$, take the distribution $(B_{i_1 \ldots i_k}, \Pi_{i_1 \ldots i_k})_{(i_1 \ldots i_k) \sim \mathcal{D}^k}$, where $B_{i_1 \ldots i_k}$ is the CNF formula obtained from $\bigvee_{\ell=1}^k B_{i_\ell}$ by applying the distributive law and $\Pi_{i_1 \ldots i_k}$ is a refutation composed of the refutations $\Pi_{i_1}, \ldots, \Pi_{i_k}$. In more detail, every clause C of $B_{i_1 \ldots i_k}$ has the form $C_1 \vee \cdots \vee C_k$ where each C_ℓ is a clause of B_{i_ℓ}. For every fixed tuple $C_2, \ldots, C_k$ we can use Π_{i_1} to derive $C' := C_2 \vee \cdots \vee C_k$ from C and F. Then for every fixed tuple $C_3, \ldots, C_k$ we can use Π_{i_2} to derive $C_3 \vee \cdots \vee C_k$ from C' and F, and so on until we derive the empty clause. ∎

Corollary 2.4 *Let "small" mean either "polylogarithmic width" or "quasipolynomial size". Then for any quasipolynomial $q(n)$, F has a small $1/2$-RR distribution if and only if F has a small $1/q(n)$-RR distribution.*

The system we are most interested in is *narrow random resolution*. This is 1/2-RR with polylogarithmic width (and hence quasipolynomial size) distributions.

2.2 Random resolution refutations

Definition 2.5 *Let Δ be a probability distribution on $\{0,1\}^n$. An* (ε, Δ)-random resolution refutation*, or* (ε, Δ)-RR refutation*, of F is a pair (B, Π) such that*

1. *B is a CNF in variables $x_1, \ldots, x_n$ and Π is a resolution refutation of $F \wedge B$*
2. *$\Pr_{\alpha \sim \Delta}[B[\alpha] = 1] \geq 1 - \varepsilon$.*

Note that this definition is, in general, not sound. In particular, let F be any (nonempty) CNF whatsoever. Let C be any clause from F and let α be an assignment which falsifies C. Let Δ be the distribution that puts all its weight on the single assignment α, and let B be the CNF $\bigwedge_i x_i = x_i[\alpha]$. Then B is true with probability 1 over Δ, and we can easily derive the empty clause from $F \wedge B$, since B contains the negation of each literal in C as a singleton clause. See also Section 3.3 below.

However, if an (ε, Δ)-RR refutation exists for *all* distributions Δ, then this is equivalent to the existence of an ε-RR distribution, as follows.

Proposition 2.6 *The following are equivalent.*

1. *F has an ε-RR distribution of width w and size s.*
2. *F has an (ε, Δ)-RR refutation of width w and size s for every distribution Δ on $\{0,1\}^n$.*

Proof. This is an immediate consequence of the minimax theorem. Consider a zero-sum game between two players, called the Prover and the Adversary, in which the Prover picks a pair (B, Π) such that B is a CNF and Π is a refutation of $F \wedge B$ of width w and size s, and the Adversary picks an assignment α. The payoff is $B[\alpha]$, that is, the Prover gets 1 if α satisfies B and 0 otherwise.

Then condition 1 says that the Prover has a mixed strategy to achieve a payoff of at least $1 - \varepsilon$, and condition 2 says that the Adversary does not have a mixed strategy to achieve a payoff less than $1 - \varepsilon$. By the minimax theorem these statements are equivalent. ∎

Our main use of RR refutations will be to prove lower bounds on RR distributions, by carefully choosing a suitable distribution Δ. To help with this we extend the notion of RR refutations to distributions over partial assignments, as these will appear in resolution lower bounds.

Definition 2.7 *Let $\mathcal{R}$ be a distribution of partial assignments to the variables $x_1, \ldots, x_n$. An* $(\varepsilon, \mathcal{R})$-RR refutation *of F is a pair (B, Π) such that*

1. *B is a CNF and Π is a resolution refutation of $F \wedge B$*
2. *$\Pr_{\rho \sim \mathcal{R}}[B^\rho = 0] \leq \varepsilon$.*

Proposition 2.8 *The following are equivalent.*

1. *F has an ε-RR distribution of width w and size s.*
2. *F has an $(\varepsilon, \mathcal{R})$-RR refutation of width w and size s for every distribution $\mathcal{R}$ over partial assignments.*

Proof. Suppose condition 1 holds and let $\mathcal{R}$ be any distribution over partial assignments. Define a distribution Δ over total assignments as follows: choose $\rho \sim \mathcal{R}$ at random, then extend it to a total assignment by setting all unset variables to 0. Let (B, Π) be the (ε, Δ)-RR refutation of F given by Proposition 2.6, so that $\Pr_{\alpha \sim \Delta}[B[\alpha] = 0] \leq \varepsilon$. By the construction of Δ it follows that $\Pr_{\rho \sim \mathcal{R}}[B^\rho = 0] \leq \varepsilon$, and thus (B, Π) is also a $(\varepsilon, \mathcal{R})$-RR refutation of F.

The other direction is immediate from Proposition 2.6. ∎

2.3 Semantic resolution refutations

Definition 2.9 *Let $\mathcal{A} \subseteq \{0,1\}^n$ be a nonempty set of truth assignments. We say that a formula C is a* semantic consequence over $\mathcal{A}$ *of formulas $C_1, \ldots, C_r$, written $C_1, \ldots, C_r \vDash^{\mathcal{A}} C$, if every assignment in $\mathcal{A}$ that satisfies $C_1, \ldots, C_r$ also satisfies C.*

A degree d semantic resolution refutation of F over $\mathcal{A}$ *is a sequence* Π *of* clauses*, ending with the empty clause, in which every clause either belongs to F or is a semantic consequence over $\mathcal{A}$ of at most d earlier clauses.*

For the sake of simplicity we will only consider degree 2 semantic resolution refutations, which we will simply call *semantic resolution refutations.*

Definition 2.10 *Let Δ be a probability distribution on $\{0,1\}^n$. An* (ε, Δ)-semantic refutation *of F is a pair $(\mathcal{A}, \Pi)$ such that*

1. *Π is a semantic refutation of F over $\mathcal{A}$, and*
2. *$\Pr_{\alpha \sim \Delta}[\alpha \in \mathcal{A}] \geq 1 - \varepsilon$.*

One can also define an *ε-semantic resolution distribution* and prove that it is equivalent to the existence of (ε, Δ)-semantic resolution refutations for all distributions Δ, in the same manner as our proof in Proposition 2.6 that an ε-RR distribution is equivalent to the existence of (ε, Δ)-RR refutations for all distributions Δ. We leave this easy exercise to the reader.

Proposition 2.11 *If F has an (ε, Δ)-RR refutation of width w and size s, then it also has an (ε, Δ)-semantic resolution refutation of width $\leq w$ and size $\leq s$.*

In the opposite direction, if F has has an (ε, Δ)-semantic refutation of width w and size s, then it also has an (ε, Δ)-RR refutation of width $O(w)$ and size at most $O(sn^2)$.

Proof. The first part follows immediately by letting $\mathcal{A}$ be the set of assignments that satisfy the auxiliary CNF B from the RR refutation. For the second part, let $(\mathcal{A}, \Pi)$ be the (ε, Δ)-semantic refutation of F. We may assume, by adding dummy premises as necessary, that every semantic consequence step in Π has exactly two premises.

Suppose C_1, C_2 and C are clauses such that $C_1, C_2 \vDash^{\mathcal{A}} C$. We claim that for all literals $x \in C_1$ and $y \in C_2$ we have $\vDash^{\mathcal{A}} x \wedge y \to C$. For otherwise there would exist some $\alpha \in \mathcal{A}$ with $C \cup \{\neg x, \neg y\}[\alpha] = 0$, which implies that $C[\alpha] = 0$ and $x[\alpha] = y[\alpha] = 1$, whence $C_1[\alpha] = C_2[\alpha] = 1$, which contradicts the assumption. Let $B_{C_1,C_2,C}$ be the CNF

$$\bigwedge \{x \wedge y \to C : x \in C_1,\, y \in C_r\}$$

(the clauses of this CNF may contain repeated variables).

Let B be the conjunction of the CNFs $B_{C_1,C_2,C}$ over all semantic consequence steps $C_1, C_2 \vDash^{\mathcal{A}} C$ in the semantic refutation Π. By the claim, $B[\alpha] = 1$ for every $\alpha \in \mathcal{A}$. Hence $\Pr_{\alpha \sim \Delta}[B[\alpha] = 1] \geq \Pr_{\alpha \sim \Delta}[\alpha \in \mathcal{A}] \geq 1 - \varepsilon$.

It remains to construct a small resolution refutation of $F \wedge B$. We can derive C from C_1, C_2 and $B_{C_1,C_2,C}$ as follows. For each $y \in C_2$, we derive $y \to C$ by resolving C_1 with all clauses in the set $\{x \wedge y \to C : x \in C_1\}$ in turn. After $O(n^2)$ steps, we have derived every clause in the set $\{y \to C : y \in C_2\}$. Resolving all these clauses with C_2 in turn, another $O(n)$ steps, gives us C. We replace every semantic consequence step in Π with a derivation of this form. ∎

2.4 Random extension clauses

The paper [4] considered the problem of proving lower bounds on bounded depth Frege proofs with connectives defined using counting modulo a prime, and reduced it to a problem about proving lower bounds on Nullstellensatz proofs containing certain specific low degree *extension polynomials*. These polynomials use additional variables r_j and have the following property: for every fixed assignment α to the original variables x_i, the extension polynomials are zero with high probability, if we fix the new variables r_j randomly.

Unfortunately, this property alone cannot be used for proving lower bounds. In [6] the authors showed that for every unsatisfiable set of low degree polynomials, there exist some extension polynomials that have the above property but are such that using them, one can derive a contradiction with a low degree proof.

In contrast to this, resolution with random extension clauses is a nontrivial concept. We will show that it is essentially equivalent to RR distributions.

Definition 2.12 *A* resolution refutation with ε-random extension clauses *is a pair* (B, Π) *such that*

1. *B is a CNF containing additional variables $r_1, \ldots, r_\ell$ not appearing in F*

2. *Π is a resolution refutation of $F \wedge B$*

3. *for every $\alpha \in \{0,1\}^n$, $\Pr_\beta[B[\alpha, \beta] = 1] \geq 1 - \varepsilon$ where the probability is with respect to an assignment β to the variables $\bar{r}$ chosen uniformly from $\{0,1\}^\ell$.*

Proposition 2.13 *If F has a resolution refutation with ε-random extension clauses of width w and size s, then it has an ε-RR distribution of width w and size s.*

In the other direction, if F has an ε-RR distribution of width w and size s, then it also has a resolution refutation with 3ε-random extension clauses of width $w + \log(n/\varepsilon^2) + O(1)$ and size $O(sn/\varepsilon^2)$.

Proof. The first part follows by substituting a random assignment to the variables $\bar{r}$ into B.

For the second part, let $\ell = \log(n/\varepsilon^2) + c$, where c is a constant which we will specify later. By Lemma 2.2, there is a 2ε-RR distribution of F with width w and size s and with a sample space containing $m = O(n/\varepsilon)$ pairs (B_i, Π_i). Let the probabilities of the samples be $p_1, \ldots, p_m$. We will approximate these numbers by multiples of $2^{-\ell}$. Choose $Q_1, \ldots, Q_m$ so that each Q_i is $2^\ell p_i$ rounded down, or up, to an integer in such a way that $\sum_i Q_i = 2^\ell$. Let $p'_i = Q_i/2^\ell$. Then for every subset $I \subseteq [n]$ we have $\left|\sum_{i \in I} p_i - \sum_{i \in I} p'_i\right| < m/2^\ell \leq \varepsilon$, where the last inequality holds if we take c sufficiently large.

Distribute the strings $\beta \in \{0,1\}^\ell$ among the formulas B_i so that the fraction of strings assigned to B_i is p'_i. More formally, take a mapping $\iota : \{0,1\}^\ell \to [m]$ such that the preimage of each i has size $p'_i 2^\ell$. For $\beta \in \{0,1\}^\ell$ let B_β be the CNF

$$\bigwedge \{r_1 = \beta_1 \wedge \cdots \wedge r_\ell = \beta_\ell \to C : \ C \text{ is a clause in } B_{\iota(\beta)}\}.$$

This is logically equivalent to $\bar{r} = \beta \to B_{\iota(\beta)}$. Let B be the CNF $\bigwedge_{\beta \in \{0,1\}^\ell} B_\beta$.

Clearly, any substitution β to the variables $\bar{r}$ in B produces the formula $B_{\iota(\beta)}$. Hence for every assignment α to $\bar{x}$ the probability that a random assignment β to $\bar{r}$ falsifies F is, by construction, at most 3ε. Hence B satisfies the third part of Definition 2.12.

For the second part of Definition 2.12, we must construct a small refutation of $F \wedge B$. First, for each β, we use the refutation $\Pi_{\iota(\beta)}$ to derive the clause $\bar{r} \neq \beta$ from B_β. We then apply Lemma 2.1. The resulting refutation has size $O(s2^\ell) = O(sn/\varepsilon^2)$ and width $w + \ell = w + \log(n/\varepsilon^2) + O(1)$, as required. ∎

3 Upper bounds

3.1 Random 3-CNFs

We will show that random 3-CNFs with sufficiently high density have small RR distributions, while as is well-known, they only have exponentially large resolution refutations [8].

Proposition 3.1 *A random 3-CNF with n variables and $64n$ clauses has a $7/8$-RR distribution of width 3 and linear size, with probability $1-o(1)$. It follows by Lemma 2.3 that such 3-CNFs have $1/2$-RR distributions of constant width and polynomial size.*

Proof. Let $m = 64n$. Let $C_1, \ldots, C_m$ be randomly chosen 3-clauses. Let α be a fixed truth assignment. Let N_α denote the number of clauses that are not satisfied by α. The expectation of N_α is $m/8$. By the Chernoff bound,

$$\Pr[N_\alpha \leq \tfrac{m}{16}] \leq \mathrm{e}^{-(\frac{1}{2})^2 \cdot \frac{1}{2} \cdot \frac{m}{8}} = \mathrm{e}^{-\frac{m}{64}}.$$

The probability that there is any α for which $N_\alpha \leq \frac{m}{16}$ is, by the union bound, at most $2^n \cdot \mathrm{e}^{-\frac{m}{64}} = \mathrm{e}^{(\ln 2 - 1)n}$, which is exponentially small.

Let $F := C_1 \wedge \ldots \wedge C_m$ be a fixed formula such that for every α the number of clauses that are satisfied by α is at most $\frac{15}{16}m$. We define the distribution of auxiliary CNFs B_i as follows. Choose $i \in \{1, \ldots, m\}$ uniformly at random. The clause C_i is a disjunction $y_1 \vee y_2 \vee y_3$ of three literals; we set B_i simply to be the conjunction $\neg y_1 \wedge \neg y_2 \wedge \neg y_3$. Then we can derive the empty clause from C_i and B_i in just three resolution steps, and for any α the probability that B_i is true is $1/8$. ∎

3.2 The retraction weak pigeonhole principle

We can also separate narrow random resolution from resolution using an explicit sequence of CNFs. The *retraction weak pigeonhole principle* (see [11]), which we denote rWPHP_n, asserts that there is no pair of functions $f : [2n] \to [n]$ and $g : [n] \to [2n]$ such that $g(f(x)) = x$ for all $x < n$. In particular, if rWPHP fails for f and g, then f is an injection and g is a surjection, so both the injective and surjective forms of the usual weak pigeonhole principle fail.

We formalize this as a propositional contradiction, which talks about f and g via their bit-graphs. So suppose $n = 2^k$. Then rWPHP_n is a CNF with variables $(f(x))_i$ for $x < 2n$ and $i < k$, for the ith bit of the value of $f(x)$, and variables $(g(y))_i$ for $y < n$ and $i < k+1$, for the ith bit of the value of $g(y)$. It consists of the clauses

$$f(x) = y \to (g(y))_i = x_i \quad \text{for } x < 2n,\ y < n \text{ and } i < k+1$$

where x_i is the ith bit of x. This is logically equivalent to $f(x) = y \to g(y) = x$ for every x and y.

Proposition 3.2 *Every resolution refutation of rWPHP_n requires width at least n and exponential size.*

Proof. It is easy to show the width lower bound by a Prover-Adversary argument, in which the Adversary maintains a partial matching between $[2n]$ and $[n]$ of size up to n.

Let m be the total number of variables in rWPHP_n and let ℓ be its initial width. Then $m \leq O(n \log n)$ and $\ell \leq O(\log n)$. By a well-known result of Ben-Sasson and Wigderson [3], if rWPHP_n has a refutation of size s then it has one of width $\ell + \sqrt{m \log s}$. A straightforward calculation shows that our width lower bound implies an exponential size lower bound. ∎

Proposition 3.3 *The formulas* rWPHP_n *have narrow* $1/2$*-RR distributions.*

Proof. Choose $x < 2n$ uniformly at random and let B_x be the $(k+1)$-CNF

$$\bigwedge_{y<n} g(y) \neq x.$$

In any total assignment this is true with probability at least $1/2$.

For the narrow refutation, first for each $y < n$ resolve the clause $g(y) \neq x$ from B_x with the clauses $f(x) \neq y \vee (g(y))_i = x_i$ for $i = 0, \ldots, k$ from rWPHP_n. The result is the clause $f(x) \neq y$. Once we have all clauses $f(x) \neq y$ we can derive the empty clause with a narrow resolution derivation, using Lemma 2.1. ∎

3.3 RR refutations over the uniform distribution

As discussed in Section 2.2, the definition of (ε, Δ)-RR refutations is unreasonably strong if we are able to fix the distribution Δ. We show in this section that the distribution does not have to be unnatural for this to happen.

Let $\mathbf{U}_n$ be the uniform distribution on $\{0,1\}^n$. We show first that, for every constant k, every unsatisfiable k-CNF has a small size and width random resolution refutation with respect to $\mathbf{U}_n$. In fact we prove something slightly more general, that this is true even for $k-$CNFs which are satisfied with some small probability. Note that it follows that the $(\varepsilon, \mathbf{U}_n)$-RR refutation system is not sound.

Proposition 3.4 *For every* $k \in \mathbb{N}$ *and* $\varepsilon > 0$*, there exist* $s \in \mathbb{N}$ *and* $\delta > 0$ *such that for every* k*-CNF* F *that is satisfied with probability* $\leq \delta$ *in* $\mathbf{U}_n$*, there exists an* $(\varepsilon, \mathbf{U}_n)$*-RR refutation of size* $\leq s$*.*

Proof. The proof is by induction on k. First suppose $k = 1$. We put $\delta = \varepsilon$. The 1-CNF F is just a conjunction of literals. If F has two complementary literals we can derive the empty clause in one step. So suppose that F has m literals, with no pair of complementary ones, and is satisfied with probability $\leq \delta$. Then $m \geq \lceil \log_2 \delta \rceil$. Let the auxiliary CNF B be a single clause consisting of $\lceil \log_2 \delta \rceil$ negated literals from F. Clearly, B has the desired properties.

Now suppose that the proposition is true for k and let ε be given. Let ℓ be the largest integer such that $(1 - 2^{-(k+1)})^\ell \leq \varepsilon$ and let $r := (\ell - 1)(k+1)$. Let $\delta > 0$ be the constant given by the inductive assumption for k and $\varepsilon 2^{-r}$. Let F be a $(k+1)$-CNF that is satisfied with probability $\leq \delta 2^{-r}$. Now there are two cases.

First, suppose that F has ℓ disjoint clauses (meaning that no two clauses share a common variable, negated or not negated). Then let Γ be a conjunction of such a set of clauses. Then Γ is satisfied with probability exactly $(1 - 2^{-(k+1)})^\ell \leq \varepsilon$. Let B be $\neg\Gamma$ written as a CNF. Then B has at most $(k+1)^\ell$ clauses, each of size ℓ, and B is satisfied with probability $\geq 1-\varepsilon$. Since ℓ is a constant, we also have a constant size refutation of $\Gamma \wedge B$, and hence of $F \wedge B$, since Γ is a subset of the clauses of F.

Otherwise, there exists a set of variables X of size r such that every clause of F contains a variable from X. Consider any assignment σ to the variables X. Then F^σ is a k-CNF and by the assumption of the proposition, F^σ is satisfied with probability at most δ. By the inductive hypothesis there exists a formula B_σ such that $F^\sigma \wedge B_\sigma$ has a constant size refutation and B_σ is satisfied with probability $\geq 1 - \varepsilon 2^{-r}$. Hence $F \wedge \bigwedge_\sigma B_\sigma$ has a constant size refutation and $\bigwedge_\sigma B_\sigma$ is satisfied with probability at least $1 - \varepsilon$, where the conjunction is over all assignments σ to the variables X. ∎

One can prove a similar proposition for other parameters, either by applying Lemma 2.3 or by modifying the proof above. For example, for every constant k, every unsatisfiable k-CNF F has an $(\varepsilon, \mathbf{U}_n)$-RR refutation with polylogarithmic width and ε^{-1} quasipolynomial.

We next show that any narrow CNF with a small resolution refutation has a narrow RR refutation with respect to $\mathbf{U}_n$. When proving size lower bounds, the first step often is to reduce the width of a refutation by applying a random restriction. The proposition shows that in the case of RR refutations there is another possibility, if the distribution is uniform or close to uniform: we can reduce the width by increasing the error. We do something like this in the proof of Theorem 4.3 below.

Proposition 3.5 *Suppose F has width v and there is a resolution refutation of F of size s. Then for every $k \geq 1$, F also has a $(2^{-k}, \mathbf{U}_n)$-RR refutation of size $2s$ and width $\leq 2\max\{v, k + \lceil \log s \rceil\}$.*

Proof. Let Π be a refutation of F of size s. Let $w = \max\{v, k + \lceil \log s \rceil\}$ and say that a clause is *wide* if it has width at least w. Each wide clause is falsified by at most $2^{n-w} \geq 2^{n-k}/s$ total assignments. Let $\mathcal{A}$ be the set of assignments which falsify any wide clause in Π. Then $|\mathcal{A}| \leq 2^{n-k}$. Delete all wide clauses from Π and let Σ be the resulting sequence of clauses. We can view Σ as a semantic resolution refutation over $\{0,1\}^n \setminus \mathcal{A}$. Hence, by Proposition 2.11, there exists a $(2^{-k}, \mathbf{U}_n)$-RR refutation of F of size $2s$ and width $2w$. ∎

It is clear that this proposition can be generalized in at least two ways. First, we may weaken the assumption that F has a small resolution refutation to the assumption that F has an ε-RR refutation, with some error ε that will be added to the parameter in the narrow proof. Second, we can use an arbitrary distribution Δ instead of $\mathbf{U}_n$, with a suitably modified concept of the width. Namely, we could define the width of a clause C with respect to Δ by

$$\mathrm{width}_\Delta(C) := \log(\Pr_{\alpha \sim \Delta}[C[\alpha] = 0])^{-1}.$$

4 Lower bounds for the bit pigeonhole principle

We present the first of our three main lower bounds on RR distributions. Before going into details, we outline the basic structure that the proofs will follow.

In this section and in Section 7 we prove size lower bounds on versions of the pigeonhole principle, by first using a random restriction in to reduce this to a width lower bound. In Section 5, on the coloured polynomial local search principle, we only prove a width lower bound, as the principle already has small resolution refutations.

To prove width lower bounds on a 1/2-RR distribution, we use Proposition 2.8 to convert the distribution into a $(1/2, \mathcal{R})$-RR refutation (B, Π) with respect to a distribution $\mathcal{R}$ on partial assignments (we will use the terms "restriction" and "partial assignment" interchangeably). The crucial thing is to choose the distribution $\mathcal{R}$ carefully.

The ideal would be that there are many restrictions ρ from $\mathcal{R}$ which make the auxiliary formula B true, thus making it vanish and leaving us with a resolution refutation for which we already have a lower bound. To this end we use a sort of rudimentary version of the switching lemma, which we call a *fixing lemma* (a different lemma in each case, because it depends on the formula F). Intuitively this shows that, with reasonably high probability, ρ fixes the value of B to either true or false. From the definition of a $(1/2, \mathcal{R})$-RR refutation we know that $B^\rho = 0$ with probability at most a half, so we can conclude that many restrictions ρ make B true.

However, in practice it is not possible to achieve the ideal that ρ makes B true. Instead we only ask that the restricted formula B^ρ cannot be falsified by any "legal" extension $\sigma \supseteq \rho$. What counts as a legal extension depends on F, and the definition is chosen so so that we can both prove the fixing lemma and then prove a width lower bound on Π by an adversary argument, in which the adversary only works with legal extensions of ρ.

The proof of a fixing lemma should, in principle, be a special case of a proof of a switching lemma, since we are essentially switching a CNF to a decision tree of height 0, or to a trivial DNF. However in the one case we consider in which a switching lemma is known, for the (non-bit) pigeonhole principle, we do not use it directly, but rather prove our own fixing lemma. One reason is that the usual lemma works with *syntactic* transformations of formulas and does not seem to guarantee that our *semantic* condition on B, that B is satisfied with high probability, is preserved. For the CPLS formula in the next section, there is unlikely to be any traditional switching lemma. This is because, understood very broadly, such a lemma would imply strong size lower bounds on CPLS in constant depth Frege, while we know that CPLS already has polynomial size refutations in resolution.

We continue with our lower bound proof for BPHP_n. Let $n = 2^k$. As already described, this is a contradictary CNF asserting that a function f is an injection from $[n+1]$ to $[n]$. It has variables $(f(x))_j$ for each $x < n+1$ and $j < k$, for the jth bit of the value of $f(x)$, and consists of clauses

$$f(x) \neq y \vee f(x') \neq y$$

for all $x < x' < n+1$ and all $y < n$.

In our proof, we will only consider partial assignments in which, for every x, either all or none of the variables $(f(x))_j$ are set. We identify such assignments with the corresponding partial functions from $n+1$ pigeons to n holes.

Given a probability p, define the distribution $\mathcal{R}_p$ of partial injections ρ from $[n+1]$ into $[n]$ as follows: choose the range of ρ by putting each hole into the range independently

at random with probability $1-p$, then choose uniformly at random from all possible partial injections onto this range. For the rest of the proof, set $p = n^{-1/2}$ and $w = n^{1/3}$.

Lemma 4.1 (fixing lemma) *Let n be sufficiently large. Suppose B is a w-CNF such that $\Pr[B^\rho = 0] \leq 1/2$. Then*

$$\Pr[\textit{there exists a partial injection } \sigma \supseteq \rho \textit{ with } B^\sigma = 0] \leq 3/4.$$

Proof. Let S be the set of $\rho \in \mathcal{R}_p$ for which there exists a partial injection $\sigma \supseteq \rho$ which falsifies B. Partition S into the set S^0 of restrictions which falsify B, and the set S^1 of restrictions which do not falsify B themselves, but which have an extension to a partial injection which falsifies B. We know $\Pr[S^0] \leq 1/2$, so it remains to bound the size of S^1.

Consider any $\rho \in S^1$. No clause in B is falsified by ρ, but there must be at least one clause which is falsified in some partial injection $\sigma \supseteq \rho$. Let C be the first such clause and let σ be such an extension of ρ falsifying it. Let x be the first pigeon mentioned in C which is not in the domain of ρ, and let $i < w$ be the position in C at which the first variable from pigeon x appears. Let σ' be σ restricted to the pigeons in the domain of ρ together with pigeon x.

Define a function θ on S^1 by $\theta : \rho \mapsto (\sigma', i)$, where σ' and i are chosen as above. Then θ is an injection, because we can first recover C from $\theta(\rho)$ as the first clause of B which is falsified in some extension of σ' to a partial injection; then we can recover x as the pigeon associated with the variable at position i in C; and finally we can recover ρ from σ' by unsetting pigeon x.

If a restriction ρ sets m pigeons, then the probability of ρ is

$$\Pr[\rho] = (1-p)^m p^{n-m} \frac{(n+1-m)!}{(n+1)!}.$$

Hence $\Pr[\sigma']/\Pr[\rho] = (1-p)/p(n-m+1)$. By the Chernoff bound, $(1-p)/(n-m+1) > 2/3$ with exponentially high probability in n. Let S_{bad} be the set of restrictions for which this bound fails, so that $\Pr[\sigma']/\Pr[\rho] > 2/3p$ for $\rho \in S^1 \setminus S_{\text{bad}}$. Partition $S^1 \setminus S_{\text{bad}}$ into subsets $S_0, \ldots, S_{w-1}$ according to the second component i of θ. On each S_i, the first component θ_1 of θ is an injection from $\mathcal{R}_p$ to $\mathcal{R}_p$ which increases probability by at least $2/3p$. Therefore

$$\Pr[\theta_1[S_\beta]] = \sum_{\rho \in S_i} \Pr[\theta_1(\rho)] > \frac{2}{3p} \sum_{\rho \in S_i} \Pr[\rho] = \frac{2}{3p} \Pr[S_\beta].$$

Since $\Pr[\theta_1[S_i]] \leq 1$ we can conclude that $\Pr[S_i] < 3p/2$, and hence that $\Pr[S^1 \setminus S_{\text{bad}}] < 3pw/2 = 3n^{-1/6}/2$. Since $\Pr[S_{\text{bad}}]$ is also exponentially small, the result follows. ∎

Theorem 4.2 BPHP_n *has no $1/2$-RR distribution of width $w = n^{1/3}$.*

Proof. We will show that BPHP_n has no $(1/2, \mathcal{R}_p)$-RR refutation with this width. Suppose for a contradiction that there is such a refutation (B, Π), where B is the auxiliary w-CNF which is false in $\mathcal{R}_p$ with probability at most $1/2$.

By Lemma 4.1, for a random $\rho \in \mathcal{R}_p$ with probability at least $1/4$ there is no extension of ρ to a partial injection which falsifies any clause from B. Thus by the Chernoff bound we can fix one such restriction ρ which also leaves at least $pn/2 = n^{1/2}/2$ holes free.

Now consider any clause C in the refutation Π. Suppose we have a partial injection $\sigma \supseteq \rho$ that falsifies C, and suppose that C is derived by resolution from clauses $D \vee v$ and $E \vee v$, where v is a variable $(f(x))_j$ for some $x < n+1$ and $j < k$. Since $|C| \leq n^{1/3}$ we may assume without loss of generality that σ sets at most $n^{1/3}$ pigeons not set in ρ. Hence we can find a free hole to assign to pigeon x, thus extending σ to a partial injection which falsifies either $D \vee v$ or $E \vee \neg v$.

In this way, working inductively up through the refutation, we can find a partial injection $\sigma \supseteq \rho$ which falsifies some initial clause. But this is a contradiction, since a partial injection cannot falsify any clause from BPHP_n, and by our choice of ρ a partial injection extending ρ cannot falsify any clause from B. ∎

We show size lower bounds by combining the argument of Theorem 4.2 with a standard application of random restrictions to remove wide clauses from Π.

Theorem 4.3 BPHP_n *has no* $1/2$*-RR distribution of subexponential size.*

Proof. Suppose (B, Π) is a subexponential size $(1/2, \mathcal{R}_p)$-RR refutation of BPHP_n. Let $p = n^{-1/2}$ and $w = n^{1/3}$ as above.

Let C be any clause that mentions at least w pigeons, and let $v_1, \ldots, v_w$ be literals from C such that v_i comes from the ith pigeon mentioned by C. Consider ρ chosen at random from $\mathcal{R}_p$. Then ρ sets almost every pigeon, so for any literal v_i we have that ρ satisfies v_i with probability almost $1/2$. These probabilities are not completely independent for different i, since ρ is constrained to be a partial injection so does not map pigeons independently, but it is not hard to show that the probability that none of these literals is satisfied in ρ is bounded above by $(1/3)^w$. By the union bound, with exponential high probability ρ satisfies every clause in Π that mentions w or more pigeons.

Now the arguments of Lemma 4.1 and Theorem 4.2 go through, with some tweaks. First, they still work if we replace the *width* of a clause C with the *number of pigeons mentioned* in C. In particular Lemma 4.1 works if we only assume that each clause in B mentions no more than w pigeons – we just need to use the index i to record which of those w pigeons σ extends ρ by, rather than the position of the relevant literal. Second, in the proof of Lemma 4.1 we ignored a set S_{bad} of restrictions ρ with a certain undesirable property, and it was safe to do this because the probability of S_{bad} was exponentially small. We now add to S_{bad} the set of restrictions ρ which do not satisfy every clause that mentions w or more pigeons, and can then safely assume that every clause C that we consider in the proof mentions fewer than w pigeons, to give the required bound on i. ∎

5 A separation of resolution from narrow RR

The *coloured polynomial local search* principle (CPLS) was introduced in [15] and the propositional version of it was studied in [19]. We refer to those two papers for more on the principle, and only remark here that it is a good candidate for proving separations of this kind because it is in some sense "complete" among narrow CNFs with short resolution refutations, while at the same time its combinatorial structure is (just) simple enough that we are able to come up with useful random restrictions. We take our definitions from [19].

Consider a leveled directed graph whose nodes consist of all pairs (i, x) from $[a] \times [b]$. We refer to (i, x) as *node x on level i*. If $i < a - 1$, this node has a single neighbour in the graph, node $f_i(x)$ on level $i+1$. Every node in the graph is coloured with some set of colours from $[c]$. CPLS expresses that the following three sentences cannot all be true at once.

1. Node 0 on level 0 has no colours.
2. For every node x on every level $i < a - 1$, if the neighbour $f_i(x)$ of x on level $i + 1$ has any colour y, then x also has colour y.
3. Every node x on the bottom level $a - 1$ has at least one colour, $u(x)$.

We will express this principle as a family of propositional contradictions. Let a be any natural number and let b and c be powers of two. We will define a CNF formula $\text{CPLS}_{a,b,c}$, in the following propositional variables.

- For each $i < a$, $x < b$ and $y < c$, there is a variable $G_i(x, y)$, expressing whether colour y is present at node (i, x).
- For each $i < a$, $x < b$ and $j < \log b$, there is a variable $(f_i(x))_j$, standing for the jth bit of the value of $f_i(x)$.
- For each $x < b$ and $j < \log c$, there is a variable $(u(x))_j$, standing for the jth bit of the value of $u(x)$.

Definition 5.1 *The formula* CPLS *consists of the following three sets of clauses, which we will call Axioms 1, 2 and 3:*

Axiom 1. For each $y < c$, the clause

$$\neg G_0(0, y)$$

Axiom 2. For each $i < a - 1$, each pair $x, x' < b$ and each $y < c$, the clause

$$f_i(x) = x' \wedge G_{i+1}(x', y) \to G_i(x, y)$$

Axiom 3. For each $x < b$ and each $y < c$, the clause

$$u(x) = y \to G_{a-1}(x, y).$$

Note that we do not require f_i to be one-to-one. However, our lower bound proof will also work for the contradiction that additionally contains the clauses $f_i(x) \neq y \vee f_i(x') \neq y$ for all $x \neq x'$ and all y.

Proposition 5.2 $\mathrm{CPLS}_{a,b,c}$ *has polynomial size resolution refutations.*

Proof. For $i < a$, let M_i be the set of clauses $\{\bigvee_{y<c} G_i(x,y) : x < b\}$ expressing that every node at level i has a colour. We can derive M_{a-1} from Axiom 3, using Lemma 2.1. Then repeatedly using Axiom 2 and Lemma 2.1 we can derive M_{a-2}, M_{a-3}, etc. Once we have M_0 we can derive a contradiction from Axiom 1. For more detail see [19]. ∎

We define a class of partial assignments which we will use in our lower bound argument. We first define the important notion of a *path* in a partial assignment.

Definition 5.3 *A* path *in a partial assignment β is a sequence of nodes $(i,x_0),\ldots,(i+k,x_k)$ of maximal length such that $f_{i+j}(x_j) = x_{j+1}$ in β for each $j \in [0,k)$.*

A path may consist of only one node. Thus every node is on some unique path.

Definition 5.4 (legal restriction) *A* legal restriction *is a partial assignment β with the following properties.*

1. *At every node (i,x), either all variables $(f_i(x))_j$ are set, or none are. At every level i, the variables that are set define f_i as a partial injection.*
2. *For every node (i,x) on the path π beginning at $(0,0)$, we have $G_i(x,y) = 0$ for every colour y. Furthermore π does not reach all the way to the bottom level $a-1$. We call π the* zero path*.*
3. *Every other path π is either* starred *or* coloured*, where*
   (a) *if π is starred, then for every node (i,x) on π, no colour $G_i(x,y)$ is set*
   (b) *if π is coloured, then there is some single colour y such that for every node (i,x) on π, $G_i(x,y) = 1$ and $G_i(x,y') = 0$ for all colours $y' \neq y$.*
4. *For every node $(a-1,x)$ on the bottom level, let π be the path containing $(a-1,x)$. If π is starred then all variables $(u(x))_j$ are unset. If π is coloured then $u(x) = y$ where y is the unique colour such that $G_i(x,y) = 1$.*

We state two obvious lemmas, without proof.

Lemma 5.5 *No legal restriction falsifies* $\mathrm{CPLS}_{a,b,c}$.

Lemma 5.6 *Let $\rho \subseteq \sigma$ be legal restrictions. Let z be a variable that is not set by ρ, but is set in σ. Then there exists a unique minimal legal extension $\rho \subset \sigma' \subseteq \sigma$ that sets z and extends ρ in one of the following two ways.*

1. *It changes some starred path into a coloured path.*

2. *It sets some value $f_i(x)$ that is not set in ρ. This necessarily means connecting two paths. Either at least one of these is starred, or they both have the same colouring. The resulting path inherits the colouring of one, or both paths. If they are both starred, then so is the resulting path.*

In either case if a node $(a-1, x)$ which was previously on a starred path is now on a coloured path, then σ' also sets $u(x) = y$ where y is the colour of the path.

We now define a distribution of random restrictions for which we will be able to prove a form of switching lemma.

Definition 5.7 (random restriction) *Fix parameters $0 < p, q < 1$. Let $\mathcal{R}_{p,q}$ be the distribution of random restrictions chosen as follows.*

1. *For each pair $i < a$ and $x < b$, with probability $(1-p)$ include (i, x) in a set S. For each $i < a$, choose f_i uniformly at random from the partial injections from the domain $\{x < b : (i, x) \in S\}$ into b.*

2. *Colour the path beginning at $(0, 0)$ so that $G_i(x, y) = 0$ for all nodes (i, x) on that path.*

3. *For every other path π, with probability $(1-q)$* colour *π randomly with one colour. That is, choose uniformly at random a colour y and, for every node (i, x) on π, set $G_i(x, y) = 1$ and set $G_i(x, y') = 0$ for all $y' \neq y$.*

4. *Finally consider each node $(a-1, x)$ on the bottom level. It is on some path π. If π was coloured at step 3, then set $u(x) = y$ where y is the unique colour assigned to π (that is, $G_{a-1}(x, y) = 1$). Otherwise leave $u(x)$ undefined.*

Abusing notation, we will also use $\mathcal{R}_{p,q}$ to denote the set of restrictions which have nonzero probability in $\mathcal{R}_{p,q}$. Note that this contains all legal restrictions.

For the rest of this section we will fix parameters as follows.

$$a = b = n,\ c = \lfloor n^{1/7} \rfloor,\ p = n^{-4/7},\ q = n^{-2/7},\ w = \lfloor n^{1/8} \rfloor$$

where b and c are powers of 2. We will use the well-known Chernoff bound several times to show that the probability of an event is *exponentially small* (or *exponentially high*), by which we mean that the probability is less than $\exp(-n^{\varepsilon})$ (or more than $1 - \exp(-n^{\varepsilon})$) for some constant $\varepsilon > 0$.

Given a restriction σ, we will say that a node $(i+1, x)$ is *free* if it is not in the range of the partial function f_i defined by σ. We will say that a node is zero, starred or coloured if the path containing it is respectively zero, starred or coloured.

Lemma 5.8 *We say that a restriction $\rho \in \mathcal{R}_{p,q}$ is* good *if all of the following hold.*

1. *No path in ρ is longer than $n^{5/7}$. (It follows that ρ is legal.)*
2. *At all levels i, there are at most $2np = 2n^{3/7}$ nodes (i, x) for which $f_i(x)$ is undefined.*
3. *At all levels $i \geq 1$, there are least $npq/2 = n^{1/7}/2$ free, starred nodes.*

Otherwise ρ is bad. *Then the probability that $\rho \sim \mathcal{R}_{p,q}$ is bad is exponentially small.*

Proof. For item 1, for a fixed vertex (i, x) the probability that there is a path of length ℓ starting at (i, x) is at most $(1-p)^\ell$. Using the union bound we can bound the probability, for $\ell = n^{5/7}$, by

$$n^2(1-p)^\ell \approx n^2 \exp(-p\ell) = n^2 \exp(-n^{1/7}).$$

Items 2 and 3 follow immediately from the Chernoff bound. ∎

In the following lemma and proof probabilities, expectations etc. are over $\rho \sim \mathcal{R}_{p,q}$.

Lemma 5.9 (fixing lemma) *Let n be sufficiently large and let B be a w-CNF such that*

$$\Pr[B^\rho = 0] \leq 1/2.$$

Then

$$\Pr[\textit{there exists a legal } \sigma \supseteq \rho \textit{ with } B^\sigma = 0] \leq 3/4.$$

Proof. Let S denote the set of $\rho \in \mathcal{R}_{p,q}$ for which there exists a legal $\sigma \supseteq \rho$ falsifying B. Let $S^0 = \{\rho \in S : \rho \text{ falsifies } B\}$, let $S_{\text{bad}} = \{\rho \in S : \rho \text{ is bad}\}$ and let $S^1 = S \setminus (S^0 \cup S_{\text{bad}})$. By assumption $\Pr[\rho \in S^0] \leq 1/2$ and by Lemma 5.8 we know that $\Pr[\rho \in S_{\text{bad}}]$ is exponentially small. So it is enough to show that $\Pr[\rho \in S^1]$ is small. To estimate this probability, we will construct a mapping $\theta : S^1 \to \mathcal{R}_{p,q}$. For every $\rho \in S^1$, we fix some legal extension σ that falsifies B and define $\theta(\rho)$ as follows.

Let C be the first clause of B that is falsified by σ and let z be the first variable of C that is not fixed by ρ – such a z must exist, because C is not falsified by ρ. Let σ' be the minimal legal extension $\rho \subset \sigma' \subseteq \sigma$ that fixes z, as given by Lemma 5.6. We put $\theta(\rho) := \sigma'$.

Let us compare the probabilities of ρ and $\theta(\rho)$. For $i = 0, \ldots, n-2$ let Z_i be the set of nodes (i, x) for which $f_i(x)$ is defined in ρ. Let $m_i = |Z_i|$ and $m = \sum_i m_i$. Let r be the number of coloured and s the number of starred paths in ρ. Then the probability of ρ is

$$\Pr[\rho] := (1-p)^m p^{n^2-m} \cdot \prod_{i=0}^{n-2} \frac{(n-m_i)!}{n!} \cdot \left(\frac{1-q}{c}\right)^r q^s$$

where the three parts of the product calculate respectively the probability of this choice for the domain of the functions f_i, this choice for the values of the f_i, and this choice of a way of colouring paths. According to Lemma 5.6, there are two possible ways in which σ' extends ρ.

1. Some starred path in ρ is coloured in σ'. Then going from ρ to σ' increases r by one, decreases s by one and leaves the other parameters the same. Hence

$$\frac{\Pr[\theta(\rho)]}{\Pr[\rho]} = \frac{1-q}{c} \cdot \frac{1}{q} \geq \frac{1}{2n^{1/7}} \cdot n^{2/7} = \frac{1}{2} n^{1/7}.$$

2. Some value $f_i(x)$ undefined in ρ is set in σ'. Then m_i and m increase by one. For the other parameters, there are two cases.

   (a) If this connects a starred path to the zero path or a coloured path, or if it connects two starred paths, then s decreases by one and r is unchanged. So

   $$\frac{\Pr[\theta(\rho)]}{\Pr[\rho]} = \frac{1-p}{p} \cdot \frac{1}{n-m_i} \cdot \frac{1}{q} \geq \frac{n^{4/7}}{2} \cdot \frac{1}{2n^{3/7}} \cdot n^{2/7} = \frac{1}{4} n^{3/7}$$

   where we have $n - m_i \leq 2n^{3/7}$ by Lemma 5.8, since ρ is good.

   (b) If it connects two coloured paths, then r decreases by one and s is unchanged. So

   $$\frac{\Pr[\theta(\rho)]}{\Pr[\rho]} = \frac{1-p}{p} \cdot \frac{1}{n-m_i} \cdot \frac{c}{1-q} \geq \frac{n^{4/7}}{2} \cdot \frac{1}{2n^{3/7}} \cdot n^{1/7} = \frac{1}{4} n^{2/7}.$$

The mapping θ is not one-to-one, but is at most $3w$-to-one. This is because we can recover ρ from $\theta(\rho)$ as follows. We find the clause C by taking the first clause in B which is not satisfied by $\theta(\rho)$. Then it suffices to know the position of the literal z in the clause C (a number less than w) and, if $\theta(\rho)$ was obtained by connecting two paths and the resulting path has a colour, to know whether this path inherited the colour from the first part, the second part, or from both parts.

Now partition S^1 as $S^1_0, \ldots, S^1_{3w-1}$ where $S^1_i = \{\rho \in S^1 : \rho \text{ is the } i\text{th preimage of } \theta(\rho)\}$. Then

$$\Pr[S^1_i] = \sum_{\rho \in S^1_i} \Pr[\rho] = \sum_{\rho \in S^1_i} \Pr[\theta(\rho)] \frac{\Pr[\rho]}{\Pr[\theta(\rho)]} \leq 2n^{-1/7} \sum_{\rho \in S^1_i} \Pr[\theta(\rho)] \leq 2n^{-1/7}$$

where we use that $\sum_{\rho \in S^1_i} \Pr[\theta(\rho)] \leq 1$, since θ is injective on S^1_i. It follows that $\Pr[S^1] \leq 6wn^{-1/7} \leq 6n^{1/8-1/7}$, giving the required bound. ∎

We can now prove our main result.

Theorem 5.10 *For all sufficiently large n, the formula $\mathrm{CPLS}_{a,b,c}$ with our parameters $a = b = n$ and $c = \lfloor n^{1/7} \rfloor$ does not have a 1/2-RR distribution of width $n^{1/8}$.*

Proof. Suppose the formula has a 1/2-RR distribution of width w. By Proposition 2.8 it also has $(1/2, \mathcal{R}_{p,q})$-RR refutation (B, Π) of width w. We will show that this implies $w \geq n^{1/8}$.

By the definition of a RR-refutation over partial assignments, $\Pr_{\rho\sim\mathcal{R}_{p,q}}[B^\rho = 0] \leq 1/2$. By Lemmas 5.8 and 5.9 there is a good $\rho \in \mathcal{R}_{p,q}$ such that no legal extension of ρ falsifies B. Fix such a ρ.

We will prove the width lower bound using the well-known Prover-Adversary game. By replacing all clauses with their negations and reversing the direction of the arrows, we can view Π as a strategy for the Prover in the following game: at each turn the Prover either asks the Adversary for the value of a variable, or forgets a variable from memory to free the space for re-use; he wins as soon as the assignment in his memory falsifies either a clause of CPLS or a clause of B (as these were the initial clauses of Π). We will show that to have a winning strategy the Prover must be able to remember at least $n^{1/8}$ variables simultaneously. The width lower bound follows immediately.

So suppose the Prover is limited to remembering at most w variables. The Adversary's strategy is to always have in mind a legal extension σ of ρ which satisfies the conjunction currently known by the Prover, and is small in a certain sense.

Define the *ρ-size* of a legal extension $\sigma \supseteq \rho$ as follows. Let m be the number of nodes (i, x) for which the edge $f_i(x)$ is defined in σ but not in ρ. Let σ^- be the smallest legal extension of ρ which contains all these edges; we could alternatively define σ^- by adding these edges to ρ, suitably colouring any starred paths that are now connected to coloured paths or the zero path, and extending u appropriately. Now σ and σ^- have the same paths. We let r be the number of paths starred in σ^- but coloured in σ. The ρ-size of σ is then $m+r$. The following claim generalises Lemma 5.6.

Claim 1 *If D is a conjunction of size $\ell < c$ and a legal extension $\sigma \supseteq \rho$ satisfies D, then D is also satisfied by a legal extension $\sigma' \supseteq \rho$ with ρ-size at most ℓ.*

Proof. Suppose that D mentions m variables of the form $(f_i(x))_j$ and r variables of the form $G_i(x, y)$ or $(u(x))_j$. We first extend ρ to σ^- by setting every edge $f_i(x)$ mentioned in D the same way it is set in σ and colouring as necessary to make this a legal restriction. This deals with the $(f_i(x))_j$ variables, and every path in σ^- is now a section of a path in σ.

We now extend σ^- to σ' by colouring some paths. For each remaining variable $G_i(x, y)$ or $u(x)$ not already set in σ^-, consider the path on which the corresponding node lies in σ. It cannot be a starred path, or the variable would not be set. If it is coloured, we colour the path it lies on in σ^- the same way as it is coloured in σ. All that remains is a set of variables of the form $G_i(x, y)$ where (i, x) lies on the zero path in σ but not on the (shorter) zero path in σ^-. Since σ satisfies D, these variables must appear negatively in D. Since $\ell < c$, there must be some colour y that is not mentioned in any of these variables. Hence we can colour the paths on which they lie with colour y, and this will satisfy D.

By construction the ρ-size of σ' is at most $m + r = \ell$. ∎

Now suppose the Prover's current memory consists of a conjunction D of size $\ell \leq w = n^{1/8}$, and the Adversary knows a legal extension $\sigma \supseteq \rho$ which satisfies D. By the claim, we may assume that the ρ-size of σ is at most ℓ. There are now three cases in the Adversary's strategy, depending on what the Prover does.

Suppose the Prover forgets a variable. Then the Adversary can apply the claim to shrink the legal extension to one of ρ-size at most $\ell - 1$.

Suppose the Prover queries a variable $(f_i(x))_j$. If this is already set in σ, reply with its value. Otherwise choose a free node $(i+1, x')$ on the next level down. This must exist by item 3 of the definition of goodness and the fact that the ρ-size of σ is less than $n^{1/7}/2$. Extend σ to σ' by setting $f_i(x) = x'$ and extending colourings appropriately. By item 1 of the definition of goodness and the bound on ρ-size, the zero path in σ' cannot reach all the way to the bottom level, so σ' is legal. Reply with the value of the variable in σ'.

Suppose the variable queries a variable $G_i(x, y)$ or $u(x)$. If this is already set in σ, reply with its value. Otherwise, extend σ to σ' by colouring the corresponding path with some arbitrary colour y. Reply with the value of the variable in σ'.

Finally we observe that the Prover cannot win against this strategy, since no legal extension of ρ falsifies any clause from $\mathrm{CPLS}_{a,b,c}$ or from B. ∎

6 Feasible interpolation

In this section we prove a form of feasible interpolation for RR distributions in which the resolution refutations are treelike. Feasible interpolation is one of the two main tools for proving lower bounds in propositional proof complexity (the other being random restrictions). Therefore it is important to fully understand the limitations of this method. It seems that the standard form of feasible interpolation does not hold for random resolution, but as we will show, one can prove lower bounds using randomized communication protocols and obtain lower bounds at least for treelike proofs. The idea of the proof is not new; similar arguments have been used in other papers for different proof systems. Furthermore, in [14] Krajíček proved a lower bound on a stronger type of communication protocol, which enabled him to prove a lower bound for general (that is, daglike) random resolution refutations. However his proof seems only to work for a small error and does not give a nontrivial bound for constant $\varepsilon > 0$.

Theorem 6.1 *Let $\bar{x}, \bar{y}, \bar{z}, \bar{u}$ be disjoint tuples of variables with $\bar{x} = x_1 \dots x_n$ and $\bar{y} = y_1 \dots y_n$. Let $F(\bar{x}, \bar{z})$ and $G(\bar{y}, \bar{u})$ be CNF formulas in the variables shown. Suppose that*

$$F \wedge G \wedge \bigwedge_{j=1}^{n} (\neg x_j \vee \neg y_j)$$

is unsatisfiable and has an ε-RR distribution in which the refutations Π_i are treelike and have size at most s. Then there exists a randomized communication protocol P for two players with the following properties. When the players are given assignments α and β in $\{0,1\}^n$ such that $\exists \bar{z}.F[\alpha, \bar{z}] = 1$ and $\exists \bar{u}.G[\beta, \bar{u}] = 1$, then

1. *with probability $\geq 1 - \varepsilon$ the players find some j such that $\alpha_j = \beta_j = 1$*
2. *they use $O(\log s)$ communication bits.*

Proof. Given α and β such that $\exists\bar{z}.F[\alpha,\bar{z}]=1$ and $\exists\bar{u}.G[\beta,\bar{u}]=1$, each player picks some γ (respectively δ) such that $F[\alpha,\gamma]=1$ ($G[\beta,\delta]=1$). Then jointly they randomly pick some (B_i,Π_i) from the RR distribution. If the auxiliary formula B_i is not satisfied by the assignment $\alpha,\beta,\gamma,\delta$ then the protocol may fail, but this happens with probability at most ε. Otherwise, the following protocol succeeds in finding the bit j. The players pick a clause C in Π_i such that the subtree above C has size between $\frac{1}{3}s$ and $\frac{2}{3}s$. They exchange bits in order to find out whether C is falsified by $\alpha,\beta,\gamma,\delta$. If so, they continue with the subtree above C. Otherwise, they delete all clauses above C from Π_i and continue with the modified tree. Since the resulting stumps are satisfied, as are all clauses of F and G, the players eventually reach a clause $\neg x_j\vee\neg y_j$ that is falsified by α,β. Since the size of the tree decreases by a factor of at least $1/3$ at each step, the number of bits they use is $O(\log n)$. ∎

We will show an application of Theorem 6.1 that gives an exponential size lower bound on treelike RR distributions. Note, however, that such a lower bound (for a different CNF, of nonconstant width) was already been proved in [5].

Recall that the disjointness function is defined for two n-bit strings by $D_n(x,y)=1$ if and only if $x_j\wedge y_j=0$ for all j. The probabilistic communication complexity of this function is $\Omega(n)$ [12]. Reduction of this function to the Karchmer-Wigderson games are known for several partial Boolean functions. For example, Raz and Wigderson [16] showed a reduction to the following problem:

> Let V be a set of $3m$ vertices. Player I is given a partial matching P (a set of independent edges) on V of size m. Player II is given a clique Q on V of size $2m+1$. The goal is to find an edge $e\in P\cap Q$.

They prove that every probabilistic protocol for this game with error at most $1/3$ needs $\Omega(m)$ communication bits.

This problem defines a partial Boolean function where the variables stand for the edges, the minterms are partial matchings of size m, and the maxterms are cliques of size $2m+1$. The fact that every maxterm intersects every minterm can be stated as a tautology in the usual way. The negation of this tautology is an unsatisfiable CNF formula of the form used in the theorem above. The variables $\bar{x}$ and $\bar{y}$ stand for the edges, the variables $\bar{z}$ represent a one-to-one mapping from a set of size m into the set of edges, and the variables $\bar{u}$ represent a one-to-one mapping from an $2m+1$-element set into the set of vertices. Such a formalization gives a formula with clauses of nonconstant size, but using extension variables we can modify it to a 3-CNF. Thus we get:

Corollary 6.2 *There exists a sequence of unsatisfiable* 3*-CNF formulas of polynomial size such that all treelike* $\frac{1}{3}$*-RR distributions have size* $2^{\Omega(\sqrt{n})}$*, where* n *is the number of variables.*

7 Lower bounds for the pigeonhole principle

We now consider the usual formalization of the pigeonhole principle, rather than the bit-graph version. It will be convenient to distinguish pigeons from holes, so let U and V be

disjoint sets of vertices with $|U| = n + 1$ and $|V| = n$. The CNF PHP_n has variables p_{ij} for $i \in U$ and $j \in V$ and consists of clauses

1. $\bigvee_{j \in V} p_{ij}$ for all $i \in U$
2. $\neg p_{ij} \vee \neg p_{i'j}$ for all $i, i' \in U$ with $i \neq i'$ and all $j \in V$.

Theorem 7.1 PHP_n *has no* $1/2$*-RR distribution of size less than* $2^{\Omega(n^{1/12})}$.

Proof. Let $\mathcal{M}_n$ be the set of all total assignments to the PHP variables arising from partial matchings of size n, equipped with the uniform distribution. That is, choose uniformly at random a partial matching M that matches every hole but leaves exactly one pigeon unmatched, and set $p_{ij} = 1$ if $(i, j) \in M$ and $p_{ij} = 0$ otherwise.

By Proposition 2.6, it suffices to prove that every $(1/2, \mathcal{M}_n)$-RR refutation of PHP_n has size $2^{\Omega(n^{1/12})}$. So suppose that we have a resolution refutation Π of $\mathrm{PHP}_n \wedge H$, where H is our auxiliary CNF and satisfies the condition

$$\Pr_{M \sim \mathcal{M}_n}[H[M] = 0] < 1/2.$$

We will call a clause or CNF *positive* if it only contains positive literals. To make H positive, we replace every negative literal $\neg p_{ij}$ in H with the disjunction $\bigvee_{k \neq i} p_{kj}$ asserting that some other pigeon goes to hole j, and call the resulting CNF H'. Since H and H' are semantically equivalent over $\mathcal{M}_n$ the condition above is preserved. Furthermore, every clause C of H can be derived by a short proof from the corresponding clause C' of H' and the axioms $\neg p_{ij} \vee \neg p_{i'j}$ of PHP_n.

Now we apply a random restriction to reduce the width of H'. However, we will use a less standard concept of width, introduced by Ajtai in [1]. For a clause C we define $w_{\mathrm{ec}}(C)$, the *edge covering width* or *ec-width* of C, to be the smallest size of a set $W \subseteq U \cup V$ that covers all edges mentioned in C. Formally,

$$w_{\mathrm{ec}}(C) := \min\{|W| \mid \forall p_{ij} \in C \ (i \in W \vee j \in W)\}.$$

If H is a CNF, then $w_{\mathrm{ec}}(H)$ is the maximum of the ec-widths of its clauses.

We will denote by $\mathcal{R}_m$ the set of partial matchings of size $n - m$ equipped with the uniform distribution. We identify $\rho \in \mathcal{R}_m$ with the following partial assignment:

1. set $p_{ij} = 1$ if $(i, j) \in \rho$
2. set $p_{ij} = 0$ if $(i, j) \notin \rho$ and either i or j is in the domain of ρ
3. the value of p_{ij} is undefined otherwise.

The set of vertices covered by the matching ρ will be called the *support of* ρ and denoted by $\sup(\rho)$.

Lemma 7.2 *There exist constants $c > 0$ and $0 < d < 1$ such that for every positive clause C and every $1 \le \ell \le n^{1/2}$,*

$$\Pr[w_{\mathrm{ec}}(C^\rho) > \ell] \le d^\ell,$$

where the probability is over $\rho \sim \mathcal{R}_{\lfloor cn^{1/4} \rfloor}$.

Proof. We will use the following elementary estimate. Let $X \subseteq [n]$, $|X| = x$ and $y \in \mathbb{N}$ be fixed, and choose $Y \subseteq [n]$ with $|Y| = y$ at random. Then

$$\Pr[|X \cap Y| \ge \ell] \le \left(\frac{\mathrm{e}xy}{n\ell}\right)^\ell. \tag{1}$$

We are given a positive clause C. Let $(E; U, V)$ be the bipartite graph determined by C. Let

$$A := \{i \in U \ : \ \deg_E(i) \ge 2n^{1/2}\ell\}.$$

We consider two cases.

Case 1. Suppose $|A| \ge 2n^{1/2}$. We will show that in this case C is satisfied, and hence has ec-width 0, with high probability. Think of ρ being constructed by first selecting its domain $D \subseteq U$ and then gradually defining $\rho(i)$ for $i \in D$, where we begin with pigeons $i \in A \cap D$. Note that $|A \cap D| \ge 2n^{1/2} - cn^{1/4} \ge n^{1/2}$ provided that n is large enough.

Suppose $\rho(i)$ has been fixed for the first $k < n^{1/2}$ elements $i \in A \cap D$. We will estimate the probability that $(i, \rho(i)) \in E$ for the next element $i \in A \cap D$. Note that there are at least $2n^{1/2}\ell - k \ge n^{1/2}\ell$ neighbours of i in the graph E that are not in the range of ρ as constructed so far. So the probability is at least $n^{1/2}\ell/n$. Hence the probability that C is *not* satisfied after the value of $\rho(j)$ is decided for the first $\lceil n^{1/2} \rceil$ elements $j \in A \cap D$ is

$$\le \left(1 - \frac{n^{1/2}\ell}{n}\right)^{\lceil n^{1/2} \rceil} = (1 - o(1))\mathrm{e}^{-\ell}.$$

Case 2. Suppose $|A| < 2n^{1/2}$. Let $D \subseteq U$ and $R \subseteq V$ denote respectively the domain and range of ρ. Let B be the set of all E-neighbours of $U \setminus (A \cup D)$. Then $(A \setminus D) \cup (B \setminus R)$ covers all edges of the graph of C^ρ, so for the lemma it is enough to show this set is small.

By applying (1) with $Y = U \setminus D$, we have

$$\Pr[|A \setminus D| \ge \ell/2] \le \left(\frac{\mathrm{e} \cdot 2n^{1/2} \cdot cn^{1/4}}{(n+1)\ell/2}\right)^{l/2} \le c_1^\ell,$$

for some constant $c_1 < 1$.

We have $|B| \le cn^{1/4} \cdot 2n^{1/2}\ell = 2cn^{3/4}\ell$. Using (1) with $Y = V \setminus R$, we have

$$\Pr[|B \setminus R| \ge \ell/2] \le \left(\frac{\mathrm{e} \cdot 2cn^{3/4}\ell \cdot cn^{1/4}}{n\ell/2}\right)^{\ell/2} = (4\mathrm{e}c^2)^{\ell/2},$$

which is exponentially small if $4\mathrm{e}c^2 < 1$. ∎

Lemma 7.3 (fixing lemma) *Let H be a positive CNF such that $w_{\mathrm{ec}}(H) \leq \ell$. For every clause C of H, pick a set W_C witnessing that $w_{\mathrm{ec}}(C) \leq \ell$. Suppose $\Pr_{M\sim\mathcal{M}_n}[H[M] = 1] \geq 1 - \varepsilon$ and let $\rho \sim \mathcal{R}_m$, where $\ell < m < n$. Then the probability of the following event*

(*) for every clause C of H, if $\sigma \supseteq \rho$ is any extension to a partial matching such that $\sup(\sigma) \supseteq W_C$, then $C^\sigma = 1$

is at least

$$1 - \varepsilon - \frac{\ell m(m-1)}{n-m+1}.$$

Proof. Let $S^0 := \{\rho \in \mathcal{R}_m \mid H^\rho = 0\}$. Clearly, $|S^0|/|\mathcal{R}_m| \leq \varepsilon$, because we can view sampling $\mathcal{R}_m$ as follows. First choose a random size-n matching $M \in \mathcal{M}_n$ and then choose randomly $\rho \subseteq M$. If $H^\rho = 0$ then also $H[M] = 0$.

Let S^1 be the set of all $\rho \in \mathcal{R}_m$ such that neither $H^\rho = 0$, nor does ρ satisfy (*). We will upper-bound $|S^1|$ by defining a one-to-one mapping $\theta : S^1 \to \mathcal{R}_{m-1} \times [\ell]$.

Let orderings of the clauses of H, of partial matchings, of elements of $U \cup V$ and elements of $U \times V$ be fixed. Let $\rho \in S^1$ be given. Take the first clause C of β for which (*) is not true. Let $\sigma \supseteq \rho$ be the first partial matching such that $\sigma \supseteq \rho$, $\sup(\rho) \supseteq W_C$ and $C^\sigma \neq 1$. Then, in fact, $C^\sigma = 0$. Since $H^\rho \neq 0$, there exists a pair (i,j) such that $(i,j) \in \sigma \setminus \rho$ and $i \in W_C$ or $j \in W_C$, because already these pairs from σ set C to 0. Let (i,j) be the first such pair and put $\rho' := \rho \cup \{(i,j)\}$. Then define $\theta(\rho) := (\rho', k)$ where k is the order of i in W_C, or the order of j in W_C if i is not in W_C.

We need to show that θ is one-to-one. Given ρ', we can determine C, because it is the first clause of H such that $C^{\rho'}$ either is 0 or can be set to 0 by an extension of ρ' (as σ is such an extension). The number k determines which pair $(x,y) \in \rho'$ with the property "$x \in W_C$ or $y \in W_C$" is (i,j).

Thus $|S^1|$ is at most $\ell \cdot |\mathcal{R}_{m-1}|$. Hence

$$\Pr_{\rho\sim\mathcal{R}_m}[\rho \in S^1] \leq \frac{\ell|\mathcal{R}_{m-1}|}{|\mathcal{R}_m|} = \frac{\ell\binom{n+1}{m-1}\binom{n}{m-2}(n+2-m)!}{\binom{n+1}{m}\binom{n}{m-1}(n+1-m)!} = \frac{\ell m(m-1)}{n-m+1}.$$

■

Now we can finish the proof of the theorem in a similar way as the previous lower bounds. Let $\ell = \delta n^{1/12}$ for a sufficiently small constant $\delta > 0$. We will show that there is no $(1/2, \mathcal{M}_n)$-RR refutation of PHP_n of size $\leq (1/d)^\ell$, where d is the constant from Lemma 7.2. Suppose Π is such a refutation. Transform the clauses of Π into positive ones and apply Lemma 7.2. Then all clauses have ec-width at most ℓ and the domain n of the PHP has been reduced to $n_1 := \lfloor cn^{1/4} \rfloor$. Now apply Lemma 7.3 with $\varepsilon = 1/2$, n_1 instead of n, and $m = \eta n_1^{1/3}$, where $\eta > 0$ is a sufficiently small constant. The constants δ and η can, clearly, be chosen so that $2\ell < m$ and the estimate on the probability in Lemma 7.3 is positive.

Let Π' be Π after transforming clauses to positive ones and applying the two restrictions from Lemmas 7.2 and 7.3. Note that after the restriction from Lemma 7.3, the auxiliary

clauses H of Π have been converted to clauses which are not falsified by any partial matching. Note also that what corresponds to a resolution step in Π' is this:

$$\frac{\Gamma \vee p_{ij} \quad \Delta \vee \bigvee_{k \neq i} p_{kj}}{\Gamma \vee \Delta} \tag{2}$$

We now traverse Π' from the empty clause at the bottom towards the initial clauses. For each clause we consider, we find a partial matching σ such that $\sup(\sigma) \supseteq W_C$ and $C^\sigma = 0$. Without loss of generality we can furthermore assume that $|\sigma| \leq \ell$, because C is falsified by the edges of σ that are incident with W_C. At the beginning we have $\sigma = \emptyset$. When going through a resolution step where C is derived from C_1 and C_2, we extend, if necessary, σ to σ' in an arbitrary way to ensure $W_{C_1}, W_{C_2} \subseteq \sup(\sigma')$. This is possible, because $|W_{C_1}| + |W_{C_2}| \leq 2\ell < m$. One can easily check (see (2)) that one of the clauses C_i must be falsified by σ'. We pick the falsified clause and, if necessary, reduce the size of the restriction to $\leq \ell$.

In this way we never reach an initial clause, because they cannot be falsified by partial matchings. Hence such a refutation does not exist. ∎

References

[1] Miklós Ajtai. *The complexity of the pigeonhole principle.* Proc. 29th Annual Symp. on Foundations of Computer Science, pp. 346-55, 1988.

[2] Albert Atserias and Neil Thapen. *The Ordering Principle in a Fragment of Approximate Counting.* ACM Transactions on Computational Logic 15:4, article 29, 2014.

[3] Eli Ben-Sasson and Avi Wigderson. *Short proofs are narrow - resolution made simple.* Journal of the ACM 48, pp. 149-169, 2001.

[4] Samuel Buss, Russell Impagliazzo, Jan Krajíček, Pavel Pudlák, Alexander Razborov and Jiří Sgall. *Proof complexity in algebraic systems and bounded depth Frege systems with modular counting.* Computational Complexity 6, pp. 256-298, 1996/1997.

[5] Samuel Buss, Leszek Aleksander Kołodziejczyk and Neil Thapen. *Fragments of approximate counting.* Journal of Symbolic Logic 79:2, pp. 496-525, 2014.

[6] Samuel Buss, Leszek Aleksander Kołodziejczyk and Konrad Zdanowski. *Collapsing Modular Counting in Bounded Arithmetic and Constant Depth Propositional Proofs.* Transactions of the American Mathematical Society 367, pp. 7517-7563, 2015.

[7] Mario Chiari and Jan Krajíček. *Witnessing functions in bounded arithmetic and search problems.* Journal of Symbolic Logic 63:3, pp. 1095-1115, 1998.

[8] Vašek Chvátal and Endre Szemerédi. *Many hard examples for resolution.* Journal of the ACM 35:4, pp. 759-768, 1988

[9] Stephen Cook and Robert Reckhow. *The relative efficiency of propositional proof systems.* Journal of Symbolic Logic 44:1, pp. 36-50, 1979.

[10] Russell Impagliazzo and Jan Krajíček. *A note on conservativity relations among bounded arithmetic theories.* Mathematical Logic Quarterly 48:3, pp.375-7, 2002.

[11] Emil Jeřábek. *On independence of variants of the weak pigeonhole principle.* Journal of Logic and Computation 17:3, pp. 587-604, 2007.

[12] Bala Kalyanasundaram and Georg Schnitger. *The Probabilistic Communication Complexity of Set Intersection.* Structure in Complexity Theory Conference, pp. 41-49, 1987.

[13] Jan Krajíček. *On the weak pigeonhole principle.* Fundamenta Mathematicae 170:1-3, pp.123-140, 2001.

[14] Jan Krajíček. *A feasible interpolation for random resolution.* ArXiv preprint, arXiv:1604.06560, 2016.

[15] Jan Krajíček, Alan Skelley and Neil Thapen. *NP search problems in low fragments of bounded arithmetic.* Journal of Symbolic Logic 72:2, pp. 649-672, 2007.

[16] Ran Raz and Avi Wigderson. *Monotone Circuits for Matching Require Linear Depth.* Journal of the ACM 39, pp. 736-744, 1992.

[17] Alexander Razborov. *Pseudorandom generators hard for k-DNF resolution and polynomial calculus resolution.* Annals of Mathematics 181:2, pp. 415-472, 2015.

[18] Alan Skelley and Neil Thapen. *The provably total search problems of bounded arithmetic.* Proceedings of the London Mathematical Society 103:1, pages 106-138, 2011.

[19] Neil Thapen. *A tradeoff between length and width in resolution.* Theory of Computing 12, article 5, 2016.