Electronic Colloquium on Computational Complexity, Report No. 116 (2013)

Lower Bounds for DNF-Refutations of a Relativized Weak Pigeonhole Principle*

Albert Atserias
Universitat Politècnica de Catalunya
Barcelona, Spain.

Moritz Müller
Kurt Gödel Research Center
Vienna, Austria.

Sergi Oliva
Universitat Politècnica de Catalunya
Barcelona, Spain.

August 29, 2013

Abstract

The relativized weak pigeonhole principle states that if at least $2n$ out of n^2 pigeons fly into n holes, then some hole must be doubly occupied. We prove that every DNF-refutation of the CNF encoding of this principle requires size $2^{(\log n)^{3/2-\epsilon}}$ for every $\epsilon > 0$ and every sufficiently large n. By reducing it to the standard weak pigeonhole principle with $2n$ pigeons and n holes, we also show that this lower bound is essentially tight in that there exist DNF-refutations of size $2^{(\log n)^{O(1)}}$ even in R(log). For the lower bound proof we need to discuss the existence of unbalanced low-degree bipartite expanders satisfying a certain robustness condition.

1 Introduction

1.1 Weak pigeonhole principles

The pigeonhole principle PHP_n^m expresses the fact that there is no injection from m pigeons into n holes whenever m is bigger than n. As usual, we formulate PHP_n^m as a contradictory CNF in the propositional variables $P_{u,v}$ with u ranging over an m-element set $[m]$ of pigeons and v ranging over an n-element set $[n]$ of holes. The formula has clauses $\neg P_{u,v} \vee \neg P_{u',v}$ for $u, u' \in [m]$ with $u \neq u'$ and $v \in [n]$ forcing different pigeons to fly to different holes, and clauses $\bigvee_{v\in[n]} P_{u,v}$ for $u \in [m]$ forcing every pigeon to fly to some hole. Estimating the

*A preliminary version of this paper appeared in the Proceedings of 28th Annual Conference on Computational Complexity (CCC), 2013 and as Chapter 5 in the third author's PhD Thesis.

ISSN 1433-8092

refutation-complexity of this set of clauses in various proof systems has a long history in proof complexity dating back to Cook and Reckhow's seminal article [14].

One of the most quoted results of propositional proof complexity is that PHP_n^{n+1} does not have short proofs in the standard propositional proof systems that "lack the ability to count". This is confirmed by the seminal results of Haken [18] for resolution, and Ajtai [1] for standard proof systems manipulating formulas of bounded depth (i.e. AC^0-Frege), followed by the great quantitative improvements by Beame, Impagliazzo and Pitassi [8] and Krajíček, Pudlák and Woods [22] on Ajtai's result. In contrast, short polynomial-size proofs exist as soon as the proof system are allowed formulas that express counting properties, such as arbitrary propositional formulas [12] (i.e. NC^1-Frege), or even threshold formulas of bounded depth (i.e. TC^0-Frege).

From the above, the ability to count looks like an essential ingredient for proving PHP_n^{n+1}. On the other hand, since *approximate counting* is available in AC^0 via explicit polynomial-size formulas [2], one may speculate that *weaker* pigeonhole principles with a much bigger gap between the number of pigeons and the number of holes, such as $\mathrm{PHP}_n^{n^2}$ or PHP_n^{2n}, may have polynomial-size bounded-depth proofs. However, this is a notorious 25-year old open problem [25], the main obstacle being that although the known AC^0-formulas for approximate counting are explicit, their *correctness* seems hard to prove. The only known superpolynomial lower bounds are for resolution in the case of $\mathrm{PHP}_n^{n^2}$ [28, 30], and for proofs manipulating k-DNFs with $k \leq \epsilon \log n / \log\log n$ for some $\epsilon > 0$ in the case of PHP_n^{2n} [6, 34, 29].

Indeed, for those weaker pigeonhole principles some positive results are known: Paris, Wilkie and Woods [25] proved that $\mathrm{PHP}_n^{n^2}$ and PHP_n^{2n} do have quasipolynomial-size bounded-depth proofs, in fact, proofs of barely superpolynomial size (cf. [25, 4]). Their proof does not rely on approximate counting. Instead, they prove $\mathrm{PHP}_n^{n^2}$ by a clever diagonalization argument and employ an amplification argument to reduce PHP_n^{2n} to $\mathrm{PHP}_n^{n^2}$. Analyzing their argument in bounded arithmetic, Krajíček [20, 19] got quasipolynomial-size proofs of the onto-version of PHP_n^{2n} by depth-2 formulas. Indeed, he obtained quasipolynomial size R(log)-refutations (cf. [20]), i.e. refutations by k-DNF formulas for k logarithmic in the size of the proof. This was later improved by Maciel, Pitassi and Woods [23] who gave $n^{O((\log n)^2)}$-size such proofs of the original version. The main question remains open: do $\mathrm{PHP}_n^{n^2}$ or PHP_n^{2n} have polynomial-size bounded-depth refutations? This is an important problem.

From the perspective of mathematical logic, the problem is tightly connected to the question whether $I\Delta_0$ refutes a certain first-order formulation of the principle in the language of arithmetic augmented by a relation symbol for (the graph of) the alleged injection from $[2n]$ into $[n]$. The standard Paris-Wilkie translation (cf. [24]) translates such refutations into polynomial-size bounded-depth refutations of PHP_n^{2n} (see [19]). Conversely, sufficiently uniform such refutations would show that $I\Delta_0$ refutes the first-order formulation. By an argument due to Paris, Wilkie and Woods, this would also imply that the infinitude of primes is provable in $I\Delta_0$ [25], which is another standing open problem. A negative answer would establish a new independence result for a weak fragment of arithmetic and this is of possible interest, e.g. in [26] Pudlák asked for methods to prove such results.

From the perspective of computational complexity theory, a negative answer could have

consequences for our understanding of approximate counting as a computational problem. In short, it would mean that approximate counting cannot be solved by polynomial-size bounded-depth circuits with elementary (i.e. comparably complex) proofs of correctness. We refer to Section 6 for a discussion.

1.2 Our results

With the hope of contributing some progress on these open problems, we study the following modified weak pigeonhole principle: if at least $2n$ out of n^2 pigeons fly into n holes, then some hole must be doubly occupied. Note that, intuitively, the ability to approximately count should still be enough to prove this principle. To formulate this principle we use additional propositional variables R_u for $u \in [n^2]$ intended to express that pigeon u decides to fly. Formally, the relativized weak pigeonhole principle $\mathrm{PHP}_n^{n^2,2n}$ has clauses

$$\begin{aligned} &\neg R_u \vee \neg R_{u'} \vee \neg P_{u,v} \vee \neg P_{u',v} && \text{for } u, u' \in [n^2] \text{ with } u \neq u' \text{ and } v \in [n], \\ &\neg R_u \vee \bigvee\nolimits_{v \in [n]} P_{u,v} && \text{for } u \in [n^2], \end{aligned}$$

together with a set of *threshold* clauses

$$\mathrm{TH}_{2n}(R, X)$$

in the R_u-variables R and some auxiliary variables X. These threshold clauses express that at least $2n$ pigeons decide to fly. More precisely, $\mathrm{TH}_{2n}(R, X)$ is a polynomial-size (in n) set of clauses such that for every assignment α to the variables R the following holds: there exists an assignment ξ to the auxiliary variables X such that $\alpha \cup \xi$ satisfies $\mathrm{TH}_{2n}(R, X)$ if and only if α sets at least $2n$ many variables in R to true.

We are ready to state the main result of this paper:

Theorem 1. *For every real $\epsilon > 0$ and every sufficiently large n, every strongly sound semantic DNF-refutation of $\mathrm{PHP}_n^{n^2,2n}$ has size at least $2^{(\log n)^{3/2-\epsilon}}$.*

We stress the fact that the lower bound holds for *any* choice of $\mathrm{TH}_{2n}(R, X)$ that satisfies the above properties. By a semantic DNF-refutation we mean a sequence of DNFs ending in the empty clause such that each DNF in the sequence either is a clause from $\mathrm{PHP}_n^{n^2,2n}$, or is an elementary tautology $(X \vee \neg X)$ for some variable X, or is logically implied by at most two DNFs appearing earlier in the sequence. Thus semantic refutations abstract away from a particular choice of inference rules. Being strongly sound is a property shared by most common inference rules (cf. [34, 33]).

Second, we show that the lower bound in Theorem 1 is essentially tight, namely, we show that for a particular natural choice of the threshold clauses $\mathrm{TH}_{2n}(R, X)$, the set $\mathrm{PHP}_n^{n^2,2n}$ has quasipolynomial size refutations even in R(log).

Theorem 2. *For a particular choice of threshold clauses $\mathrm{TH}_{2n}(R, X)$ with the properties described above, there exist* R(log)*-refutations of $\mathrm{PHP}_n^{n^2,2n}$ of size $2^{(\log n)^{O(1)}}$.*

This is not hard to prove via bounded arithmetic, here is a sketch of such a proof: we give a first-order formula $\pi_n^{n^2,2n}$ in the language of arithmetic plus some additional symbols whose Paris-Wilkie translation is $\mathrm{PHP}_n^{n^2,2n}$ for a particular choice of threshold clauses. The mentioned quasipolynomial size R(log)-refutations of $\mathrm{PHP}_n^{n^2,2n}$ given by Maciel, Pitassi and Woods are sufficiently uniform to actually establish that Buss' bounded arithmetic theory T_2^2 refutes the first-order formulation π_n^{2n} of PHP_n^{2n} [23]. This refutation is readily transformed into a T_2^2 refutation of our principle $\pi_n^{n^2,2n}$ via a simple (i.e. Δ_1^b-) first-order interpretation. Theorem 2 then follows by standard means applying the Paris-Wilkie translation to this refutation.

Let us also point out that all the results of this paper would hold without any essential modification for $\mathrm{PHP}_n^{n^c,dn}$ for any two constants $d > 1$ and $c > 1$. We focus on $d = c = 2$ for concreteness.

1.3 Relevance of the results

In the previous section we expressed the hope that our study of the relativized weak pigeonhole principle could lead to some progress on the main questions about the standard weak pigeonhole principle. Let us discuss how.

First, the mentioned first-order interpretation of π_n^{2n} in $\pi_n^{n^2,2n}$ also shows that $\mathrm{PHP}_n^{n^2,2n}$ has polynomial-size bounded-depth refutations if and only if so does PHP_n^{2n} (the *if* direction is given by the interpretation and the *only if* direction is even easier). As already said the latter is an important open problem, so our lower bound for depth-2 is of potential relevance for the eventual resolution of this problem. The fact that our methods yield quasipolynomial lower bounds where comparably big upper bounds exist is particularly encouraging. To our knowledge, this is the first natural occurrence of this phenomenon. However, our lower bound falls short of making explicit progress on whether polynomial-size bounded-depth proofs exist. We refer to Section 6 for a discussion.

Second, let us note that Theorem 1 gives the first superpolynomial lower bound for strongly sound semantic DNF-refutations. Previously known lower bounds for refutations handling arbitrary DNFs come from principles exponentially hard for bounded-depth Frege systems, and there are essentially only two such principles known, namely PHP_n^{n+1} and the so-called counting principles (cf. [19]). These lower bounds do not hold for strongly sound semantic refutations but hold only relative to a fixed choice of finitely many inference rules.

Finally, we find it worthwhile to point out that there is no "complexity gap" [31] for R(log), which is to be put in contrast with those known for tree-like systems [31, 15, 21]. Indeed, $\mathrm{PHP}_n^{n^2,2n}$ can be seen as the Paris-Wilkie translation of a suitable first-order sentence without "built-in" arithmetical symbols (cf. [15] for a discussion), and our lower and upper bounds state that $\mathrm{PHP}_n^{n^2,2n}$ has intermediate proof complexity in R(log), i.e. superpolynomial but not exponential.

1.4 Proof outline and comparison to previous work

Our proof follows the random restriction method, so successfully used in previous works in propositional proof complexity, with some additional ideas. The typical skeleton of a proof by the random restriction method goes as follows: Assume a short proof of F is given. Apply a random restriction from a suitable distribution in such a way that, with high probability, every formula in the proof simplifies significantly, but the proved formula F remains hard. Finally argue directly that the restricted F cannot have a short proof with such simple formulas.

For an example, suppose PHP_n^{2n} has polynomial-size resolution refutations. For the random restriction we choose an assignment that describes a 1-1 mapping from $n/2$ randomly chosen pigeons onto $n/2$ randomly chosen holes, and leaves all the other variables unset. With these parameters, the restricted PHP_n^{2n} becomes $\mathrm{PHP}_{0.5\cdot n}^{1.5\cdot n}$, and each *complex* clause of the proof has been made true with high probability. Now a direct prover-adversary argument shows that a proof of $\mathrm{PHP}_{0.5\cdot n}^{1.5\cdot n}$ with non-complex clauses only is impossible.

Trying to apply this argument to DNF-refutations hits several difficulties. First, a random *matching* restriction as above is not likely to simplify an arbitrary DNF formula, even if this formula is small. Indeed, the DNF could be the negation of PHP_n^{2n} itself, and the point of the argument above was precisely that this formula does not simplify much. Here is where our modified version $\mathrm{PHP}_n^{n^2,2n}$ enters the picture. By choosing $2n$ out of n^2 pigeons at random and setting all the variables about the other pigeons completely at random, it is very likely that each DNF in the proof simplifies into one all whose terms mention very few of the $2n$ chosen pigeons. This sort of restriction comes inspired by the so-called Dantchev-Riis restrictions [15], and its analysis for our case requires arguments of the type Furst, Saxe, and Sipser introduced in their seminal work on bounded-depth circuits [17][1].

Continuing with the sketch of the proof, the application of the Dantchev-Riis restriction to $\mathrm{PHP}_n^{n^2,2n}$ leaves an instance of PHP_n^{2n}. Unfortunately, a term mentioning very few pigeons need not be short itself, which means that we are not yet at a contradiction with the known lower bounds for PHP_n^{2n} in k-DNF resolution for $k \leq \sqrt{\log n / \log\log n}$ from [34] which were later improved to $k \leq \epsilon \log n / \log\log n$ for some $\epsilon > 0$ [29]. Following the ideas in [10], as adapted to k-DNF proofs in [6, 34], this suggests that we restrict the principle further to a low-degree bipartite expander G (with left vertices $[2n]$ and right vertices $[n]$) to get a short proof of $\mathrm{PHP}(G)$. Recall (cf. [10, 33]), this formula is obtained from PHP_n^{2n} by zeroing out all $P_{u,v}$ with (u, v) not an edge of G.

The low-degree condition on G guarantees that whenever a term mentions very few pigeons we can also assume that the term is short, resulting in a k-DNF refutation of $\mathrm{PHP}(G)$ for small k. This would seem to open the door to using the methods in [34].

Unfortunately, the sort of bipartite expanders that are needed for the rest of the argument require degree at least as large as $\log n$, leaving k well above the quantity that a direct application of the methods in [34] can afford. Here comes the second main idea in our proof:

[1] It is interesting to point out that, in view of Theorem 2, the stronger switching-lemma sort of argument (e.g. Håstad-style) cannot work for this because this is the only essential point in the argument where we use that the given proof has quasipolynomial size.

we use a logarithmic degree expander G, but reduce our problem to proving lower bounds for a related formula $\mathrm{BPHP}(G)$ in which the flights of the pigeons along the edges of the graph are encoded in *binary*. This takes us from $k = \Omega(\log n)$ in the unary encoding to $k = O(\log \log n)$ in the binary encoding (at least in the case that we start with polynomial-size proofs), well below the critical $\sqrt{\log n / \log \log n}$.

Now, we would like to call the small restriction switching lemma from [34] to get a restriction turning all formulas into formulas representable by shallow decision trees. Such refutations can be ruled out by an adversary argument. Unfortunately, the move from $\mathrm{PHP}(G)$ to $\mathrm{BPHP}(G)$ increases the depth of the formulas and this blocks a direct application of the switching lemma. This difficulty is sidestepped via a third idea in our proof, which is to use an appropriate weaker notion of representation by decision trees.

Putting all these ideas together into a proper argument requires a fair amount of technical work and this is what the rest of the paper is devoted to. After a few preliminaries in the next section, in Section 3 we discuss the sort of expander graphs we need, and in Section 4 we use them for the proof of the main theorem. Section 5 proves Theorem 2 and Section 6 gives some final discussion.

2 Preliminaries

For a natural $n \in \mathbb{N}$, we write $[n] := \{0, \ldots, n-1\}$ and $|n| := \lceil \log(n+1) \rceil$. All our logarithms are base 2. Note that, for $n > 0$, the natural $|n|$ is the length of the binary representation of n without leading zeros. For $b \in \mathbb{N}$ we write $\mathrm{bit}(b, n)$ for the $(b+1)$-th least significant bit in the binary representation of n; formally, $\mathrm{bit}(b, n) := \lfloor n/2^b \rfloor \mod 2$. Note that if $b \geq |n|$, then $\mathrm{bit}(b, n) = 0$.

2.1 Bipartite graphs

Let $G = (U, V, E)$ with $E \subseteq U \times V$ be a bipartite graph. For a vertex $u \in U \cup V$ let $N_G(u)$ be the set of neighbors of u in G and for a set of vertices $A \subseteq U \cup V$, let $N_G(A) := \bigcup_{u \in A} N_G(u)$. A set $M \subseteq E$ is a *matching (in G)* if no two edges in M share an endpoint. Note that matchings M are bijections and thus have an image $\mathrm{Im}(M)$ and a domain $\mathrm{Dom}(M)$.

We say G is a (U, V, d_L, d_R)-graph if for every $u \in U$ we have that $|N_G(u)| \leq d_L$ and for every $v \in V$ we have that $|N_G(v)| \leq d_R$. With such a graph we associate a bijection ϕ_G with $\mathrm{Dom}(\phi_G) \subseteq U \times [d_L]$ such that for every $u \in U$ and every $v \in N_G(u)$ there is (exactly one) $i \in [d_L]$ such that $(u, i) \in \mathrm{Dom}(\phi_G)$ and $\phi_G(u, i) = v$. For a subset $C \subseteq U \cup V$ we let $G \cap C$ denote the subgraph of G induced by the vertices of C; if ϕ_G is associated to G, then $G \cap C$ is a $(U \cap C, V \cap C, d_L, d_R)$-graph and the map associated to $G \cap C$ is (as a set of pairs) $\phi_{G \cap C} := \phi_G \cap ((C \times [d_L]) \times C)$. We also write $G \setminus C$ for $G \cap ((U \cup V) \setminus C)$.

2.2 Propositional formulas

Propositional variables are also called *atoms*. A *literal* is an atom X or its *negation* $\neg X$. A *formula* is built from literals by means of $\vee$ and $\wedge$. Note that we allow the negation symbol only in front of atoms. The *negation* $\neg F$ of a formula F is defined as the formula obtained from F by interchanging $\wedge$ and $\vee$, and replacing every literal by its complementary literal (i.e. X by $\neg X$ and $\neg X$ by X). If Γ is a set of formulas, we write $\bigwedge\Gamma$ (resp. $\bigvee\Gamma$) for the conjunction (resp. disjunction) of the formulas in Γ; the elements in Γ are the *conjuncts* (resp. *disjuncts*). We allow the empty disjunction 0 and the empty conjunction 1, and refer to them as *constants*. Note that $\neg 1 = 0$ and $\neg 0 = 1$. A *(k-)term* is a conjunction of (at most k many) literals; and a *(k-)clause* is a disjunction of (at most k many) literals. A *(k-)CNF* is a conjunction of (k-)clauses, and a *(k-)DNF* is a disjunction of (k-)terms.

By $|F|$ we denote the *size* of the formula F: literals and constants have size 1, and $|(F \wedge G)| = |(F \vee G)| = 1 + |F| + |G|$. Note that $|F| = |\neg F|$.

2.3 Restrictions and substitutions

An *assignment* is a function mapping all atoms to truth values 0 and 1. A *restriction* ρ is a partial assignment, i.e. a function mapping some atoms into $\{0, 1\}$. For a formula F we let $F \restriction \rho$ denote the formula obtained from F by first "substituting" literals by their truth values under ρ and then "eliminating" constants. To define this we blur the distinction between truth values $\{0, 1\}$ and constants $\{0, 1\}$ and view a restriction or assignment as a substitution mapping atoms to constants. More generally, a *substitution* is a partial function from atoms to formulas. Let S be a substitution. For a formula F let F^S denote the result of substituting in F the formula $S(X)$ for X and $\neg S(X)$ for $\neg X$, simultaneously for all $X \in \mathrm{Dom}(S)$. For a restriction ρ then define

$$F \restriction \rho := \mathrm{E}(F^\rho),$$

where E is the function that eliminates constants and is defined as follows: If F is a literal or a constant, set $\mathrm{E}(F) := F$. For $F = (G \wedge H)$, set $\mathrm{E}(F) := 0$ if $\mathrm{E}(G) = 0$ or $\mathrm{E}(H) = 0$; set $\mathrm{E}(F) := \mathrm{E}(G)$ if $\mathrm{E}(H) = 1$; $\mathrm{E}(F) := \mathrm{E}(H)$ if $\mathrm{E}(G) = 1$; and otherwise set $\mathrm{E}(F) := (\mathrm{E}(G) \wedge \mathrm{E}(H))$. For $F = (G \vee H)$, define $\mathrm{E}(F)$ analoguously with the roles of 0 and 1 switched. If Γ is a set or sequence of formulas we write Γ^S for the result of applying the substitution to every formula in Γ. The notation $\Gamma \restriction \rho$ is defined analogously.

If Γ is a set of formulas and F is a formula, we say that Γ *logically implies* F if every assignment ρ with $G \restriction \rho = 1$ for all G in Γ satisfies $F \restriction \rho = 1$. We say that Γ *strongly implies* F if the same is true for restrictions ρ. In case Γ is a singleton $\{G\}$ we say that G logically or strongly implies F. We say that F and G are *(strongly) equivalent* if they (strongly) imply each other. Of course if F strongly implies G then F logically implies G but the converse need not be true; e.g. $(X \vee \neg X)$ is logically but not strongly implied by 1.

Lemma 3. *Let Γ be a set of formulas, F a formula, ρ a restriction, and S a substitution. Then*

1. *if* Γ *strongly implies* F*, then* $\Gamma \restriction \rho$ *strongly implies* $F \restriction \rho$*, and*
2. *if* Γ *strongly implies* F*, then* Γ^S *strongly implies* F^S*.*

Proof. For restrictions the statement follows directly from the definitions. For substitutions we need a couple of technical observations, both with easy proofs by induction on the structure of formulas:

Claim. For every formula F and every substitution S the following hold:

1. $\mathrm{E}(F^S) = \mathrm{E}(F^{\mathrm{E} \circ S})$,
2. $\mathrm{E}(F^S) = \mathrm{E}(F)^S$ if constants do not appear in $\mathrm{Im}(S)$.

Another observation we need is that, in order to prove the lemma, one can assume that S is defined on all variables appearing in $\Gamma \cup \{F\}$: in case S is undefined on a variable X, extend it to map X to itself.

For a restriction ρ let S^ρ be the substitution mapping $X \in \mathrm{Dom}(S)$ to $S(X)^\rho$, and let $S \restriction \rho$ be the substitution mapping $X \in \mathrm{Dom}(S)$ to $S(X) \restriction \rho$. Observe that $S \restriction \rho$ is the disjoint union of a restriction σ and a substitution S' such that constants do not appear in $\mathrm{Im}(S')$. Then for every $H \in \Gamma \cup \{F\}$ we have

$$H^S \restriction \rho = \mathrm{E}(H^{S^\rho}) = \mathrm{E}(H^{S \restriction \rho}) = \mathrm{E}((H^\sigma)^{S'}) = (H \restriction \sigma)^{S'}. \tag{1}$$

The first equality holds because, by the comment after the claim, S is defined on all variables of H, the second holds by part 1 of the Claim noting $\mathrm{E} \circ S^\rho = S \restriction \rho$, the third holds because $S \restriction \rho$ is the disjoint union of σ and S', and the fourth holds by part 2 of the Claim.

We are ready to show that Γ^S strongly implies F^S. Let ρ be a restriction and assume $G^S \restriction \rho = 1$ for every $G \in \Gamma$. By (1) we have $(G \restriction \sigma)^{S'} = 1$ and hence $G \restriction \sigma = 1$ for every $G \in \Gamma$. Since Γ strongly implies F thus $1 = F \restriction \sigma = (F \restriction \sigma)^{S'} = F^S \restriction \rho$ where the last equality holds again by (1). □

Remark 4. Let F be a CNF and let F' be the DNF that is obtained by distributing the conjunctions over the disjunctions in the straightforward way. In other words, if $F = \bigwedge_i C_i$ where each C_i is a clause, then F' is the disjunction of all the terms that are obtained by conjoining exactly one literal from each C_i. Obviously F and F' are logically equivalent. We point out that they are also strongly equivalent. To see this just note that a restriction that makes F true must make true at least one literal from each clause, and that a restriction that makes F' true must make true all the literals of at least one of its terms.

2.4 Semantic refutations and strongly sound refutations

A *semantic proof of* F *from* Γ is a sequence $F_0, \ldots, F_{m-1}$ of formulas such that $F = F_{m-1}$ and for every $i \in [m]$ either F_i is in Γ, or F_i is an *elementary tautology* $(X \vee \neg X)$ for some variable X, or there is a set $\Delta \subseteq \{F_j \mid j \in [i]\}$ with at most two elements such that Δ logically implies F_i. The proof is *strongly sound* (cf. [34]) if in the third case Δ strongly implies F_i (see the previous section). We speak of a *(k-)DNF-proof* if all F_i with $i \in [m]$ are (k-)DNFs. A *refutation* is a proof of 0. The proof has *size* $\sum_{i \in [m]} |F_i|$ and *length* m.

Remark 5. Length is not a very interesting measure for strongly sound semantic DNF-refutations because if Γ is a contradictory set of clauses, or even a contradictory set of DNFs, then there always is a short strongly sound semantic DNF-refutation of Γ. To see this, suppose that $G_0, \ldots, G_{m-1}$ enumerates Γ where $m = |\Gamma|$. The refutation is $F_0, \ldots, F_m$ with $F_0 := G_0$ and $F_m := 0$, and F_{i+1} a DNF that is strongly equivalent to $(F_i \wedge G_{i+1})$ for $0 < i < m-1$. More precisely, if $F_i = \bigvee_j T_j$ for terms T_j and $G_{i+1} = \bigvee_k T'_k$ for terms T'_k, then $F_{i+1} = \bigvee_{j,k}(T_j \wedge T'_k)$. It is easy to see that F_i and G_{i+1} strongly imply F_{i+1}. The length of this refutation is $m+1$ where m is the cardinality of Γ (but its size is exponential in the size $\sum_{i\in[m]} |G_i|$ of Γ).

For a natural $k \geq 1$, an R(k)*-refutation of* Γ is a refutation of Γ such that every formula is a k-DNF that either is in Γ, or is an elementary tautology $(X \vee \neg X)$ for some variable X, or is obtained from earlier formulas by an application of the cut rule or the weakening rule or the $\wedge$-introduction rule (cf. [20] for a definition). An R(log)*-refutation* is an R(k)-refutation with $k \leq \log s$, where s is the size of the proof. Note that R(k)-refutations, and hence R(log)-refutations are strongly sound.

2.5 Decision trees

A *decision tree* is a finite, rooted, ordered tree whose inner vertices are labeled by atoms, whose leafs are labeled by 0 or 1, and such that no atom occurs twice in a branch (i.e. a path from the root to some leaf); we say the tree *queries* an atom if the atom occurs as a label. Each inner vertex has two successors (i.e. immediate successors on a branch). Since the tree is ordered we can distinguish between a *left* and a *right* successor of an inner vertex. By a *0-branch* (*1-branch*) we mean a branch leading to a leaf labeled 0 (labeled 1). Every path π from the root to some vertex corresponds to the following restriction that we also denote by π: if an atom occurs as a label of a vertex p in the path π, then the restriction sets this atom to 0 if the left successor of p is in π and to 1 if the right successor of p is in π; if π contains no successor of p, then the restriction does not evaluate the atom. When we say that a restriction or a branch *extends* the path π we mean extension as restrictions.

A decision tree T *represents* a formula F if $F \restriction \pi = b$ for every $b \in \{0,1\}$ and every b-branch π of T. We let $h(F)$ denote the minimal height of a decision tree representing F. A decision tree could represent a formula but query variables that do not even appear in the formula. However there is not much point in doing that:

Lemma 6. *Every formula F is represented by a decision tree of height at most $h(F)$ that queries only variables appearing in F.*

The following lemma is also easy to verify.

Lemma 7. *Let T_0 and T_1 be decision trees of heights h_0 and h_1 respectively. Then there exists a decision tree T of height at most $h_0 + h_1$ such that*

(a) every 0-branch of T extends a 0-branch of T_0 or it extends a 0-branch of T_1; and

(b) every 1-branch of T extends both a 1-branch of T_0 and a 1-branch of T_1.

Although we do not actually use it, we remark that if T_0 and T_1 represent the formulas F_0 and F_1, then the decision tree T of Lemma 7 represents $F_0 \wedge F_1$.

3 Resilient expanders

In this section we discuss the sort of expander graphs that we need. In short, these are unbalanced low-degree bipartite expanders that satisfy an additional robustness condition: for at least half the subsets of vertices of some fixed size on the right-hand side, the graph remains an expander if these vertices are removed. Let us note that a similar definition was implicit in [6] which was later revisited in [34]. However, both these concepts were very tied to their specific application to proof complexity. Here we provide a more systematic and general treatment.

3.1 Definition and some basic properties

Let $G = (U, V, E)$ be a bipartite graph with $|U| = t$ and $|V| = n$ where $t \geq n$. Let b be a positive real and let q and r be naturals such that $0 \leq q \leq n/(1+b)$ and $0 \leq r \leq n$. Recall that G is a (q, b)-expander if $|N_G(S)| \geq (1+b)|S|$ for every q-element subset $S \subseteq U$. We say that G is a (q, b, r)-*resilient expander* if for a random r-element subset $\mathbf{B} \subseteq V$ we have that $G \setminus \mathbf{B}$ is a (q, b)-expander with probability bigger than $1/2$.

The choice of $1/2$ here is arbitrary; any constant in the open interval $(0, 1)$ would serve our purposes. However, observe that if we were to require that $G \setminus \mathbf{B}$ is a (q, b)-expander with probability 1 over the choice of $\mathbf{B}$, then the minimum degree of G would have to exceed r. Later we will see that for the less demanding requirement of probability strictly smaller than 1 we can afford a much smaller degree.

A first property to note is that if G is a (q, b, r)-resilient expander, then $G \cap C$ is also a (q, b, r)-resilient expander for every $C \subseteq U$. In other words, the property is hereditary under taking subsets of the left-hand side. Similarly, if it is a (q, b, r)-resilient expander then it also is a (q', b', r')-resilient expander for all $q' \leq q$, all positive $b' \leq b$, and all $r' \leq r$. The next lemma proves the only non-trivial case of this statement.

Lemma 8. *If G is a (q, b, r)-resilient expander, then G is a (q, b, s)-resilient expander for all $s \leq r$.*

Proof. Fix $s \leq r$. Call a set $B \subseteq V$ good if $G \setminus B$ is a (q, b)-expander. Observe that any subset of a good set is good. Assume at least half the r-element subsets of V are good. Each good r-element set contains exactly $\binom{r}{s}$ many good s-element sets, and each such s-element set appears in at most $\binom{n-s}{r-s}$ many good r-element sets. Therefore, the number of good s-element sets is at least $\frac{1}{2}\binom{n}{r}\binom{r}{s}/\binom{n-s}{r-s}$. Expanding the binomials, one sees this is precisely $\frac{1}{2}\binom{n}{s}$. □

3.2 Existence

We prove that random bipartite graphs with the appropriate parameters are resilient expanders. For naturals t, n and d, let $\mathbf{G} = \mathbf{G}(t, n, d)$ be the random bipartite graph (U, V, E) with $U = [t]$ and $V = [n]$ defined by the following random experiment: for each $u \in U$ choose a d-element subset N_u of V uniformly and independently at random, and declare each $v \in N_u$ a neighbor of u.

Lemma 9. *Let ε and b be positive reals, let t, n, q, r and d be naturals such that $t \geq n > 1 + 2/\varepsilon$, $q \leq n/12(1+b)$, $r \leq n/12$, and $n \geq d \geq (\log t + (3+b)\log n)/(\log n - \log(3(1+b)q + 3r))$, and let $\mathbf{G} = \mathbf{G}(t, n, d)$. Then*

$$\mathbb{P}[\ \mathbf{G} \text{ is a } (q, b, r)\text{-resilient expander}\ ] > 1 - \varepsilon.$$

Before we prove this, let us look at some special cases to illustrate the complicated expressions in the hypothesis. Think of ε and b as positive constants and think of all other parameters as functions of n. If $t = O(n)$, $q = \Omega(n)$ and $r = \Omega(n)$, then the required lower bound on the degree d is $O(\log n)$. On the other hand, if still $t = O(n)$ but $q = n^{1-\Omega(1)}$ and $r = n^{1-\Omega(1)}$, then the required lower bound on the degree is only $O(1)$. For our application we will have $t = 2n$, $q = n^{1-\Omega(1)}$ and $r = \Theta(n/\log n)$, in which case the required lower bound on the degree is $O(\log n/\log\log n)$.

To prove Lemma 9 we rely on the following probabilistic fact. Let X be a random variable that takes all of its values x with positive probability. Given an event $\mathcal{E}$, recall that $\mathbb{P}[\ \mathcal{E} \mid X\ ]$ is the random variable $f \circ X$ where f is the function defined by $f(x) = \mathbb{P}[\ \mathcal{E} \mid X = x\ ]$ for every value x of X.

Lemma 10. *Let p be a real such that $0 < p < 1$, let $\mathcal{E}$ be an event and let X be a random variable. Then*

$$\mathbb{P}[\ \mathbb{P}[\ \mathcal{E} \mid X\ ] > p\ ] \geq \frac{1}{1-p} \cdot (\mathbb{P}[\ \mathcal{E}\ ] - p)\,.$$

Proof. Since $\mathbb{P}[\ \mathcal{E} \mid X\ ]$ takes values in $[0, 1]$ we have

$$\mathbb{E}[\ \mathbb{P}[\ \mathcal{E} \mid X\ ]\ ] \leq \mathbb{P}[\ \mathbb{P}[\ \mathcal{E} \mid X\ ] > p\ ] \cdot 1 + (1 - \mathbb{P}[\ \mathbb{P}[\ \mathcal{E} \mid X\ ] > p\ ]) \cdot p.$$

On the other hand, direct calculation shows $\mathbb{E}[\ \mathbb{P}[\ \mathcal{E} \mid X\ ]\ ] = \mathbb{P}[\ \mathcal{E}\ ]$. This implies the lemma. □

Proof of Lemma 9. Let $\mathbf{B}$ be an r-element subset of V chosen uniformly at random and independently from $\mathbf{G}$. In the following we let B range over values of $\mathbf{B}$. Let $\mathcal{E}$ be the event that $\mathbf{G} \setminus \mathbf{B}$ is a (q, b)-expander. Observe that the event that $\mathbf{G}$ is a (q, b, r)-resilient expander equals the event that $\mathbb{P}[\ \mathcal{E} \mid \mathbf{G}\ ] > 1/2$. By Lemma 10 it thus suffices to show that

$$\mathbb{P}[\ \mathcal{E}\ ] > 1 - \frac{\varepsilon}{2}. \tag{2}$$

Fix B and let $\mathcal{E}^B$ denote the event that $\mathbf{G} \setminus B$ is a (q, b)-expander. Further, fix two sets $S \subseteq U$ and $T \subseteq V \setminus B$ of cardinalities $i \leq q$ and $j < (1+b)i$ respectively. Recall that $N_{\mathbf{G}}(S)$ denotes the neighbors of S in the random graph $\mathbf{G}$. Then

$$\mathbb{P}[\ N_{\mathbf{G}}(S) \subseteq T \cup B\ ] \leq \left(\frac{\binom{j+r}{d}}{\binom{n}{d}}\right)^i \leq \left(\frac{(j+r)e}{n}\right)^{di};$$

here we use $\binom{j+r}{d} \leq ((j+r)e/d)^d$ and $\binom{n}{d} \geq (n/d)^d$. By the union bound over (non-empty) $S \subseteq U$ and $T \subseteq V \setminus B$ of the appropriate cardinalities we have

$$\mathbb{P}[\ \overline{\mathcal{E}^B}\ ] \leq \sum_{i=1}^{q} \binom{t}{i} \sum_{j=1}^{\lfloor (1+b)i \rfloor} \binom{n}{j} \left(\frac{(j+r)e}{n}\right)^{di}. \tag{3}$$

The term $\binom{n}{j} \cdot ((j+r)e/n)^{di}$ in the internal sum in (3) is bounded by $n^j \cdot ((j+r)e/n)^{di}$, which is an increasing function of j. Plugging in the largest possible j and multiplying by the number of terms, the internal sum in (3) is at most

$$(1+b)i \cdot n^{(1+b)i} \cdot \left(\frac{(1+b)ie + re}{n}\right)^{di} \leq \left(n^{2+b} \cdot \left(\frac{3(1+b)q + 3r}{n}\right)^d\right)^i.$$

Here we use $1 \leq i \leq q$ and $q \leq n/12(1+b)$ so that $(1+b)i \leq n$ and $(1+b)i \cdot n^{(1+b)i} \leq n^{(2+b)i}$. Crudely bounding $\binom{t}{i}$ by t^i, we conclude that (3) is bounded by

$$\sum_{i=1}^{q} \left(t \cdot n^{2+b} \cdot \left(\frac{3(1+b)q + 3r}{n}\right)^d\right)^i.$$

From $q \leq n/12(1+b)$ and $r \leq n/12$ we conclude that the fraction is bounded by $1/2$ and hence is strictly smaller than 1. From $d \geq (\log t + (3+b)\log n)/(\log n - \log(3(1+b)q + 3r))$ we conclude that (3) is bounded by

$$\sum_{i=1}^{\infty} \left(\frac{1}{n}\right)^i = \frac{1}{n-1}.$$

At this point we proved that $\mathbb{P}[\ \overline{\mathcal{E}^B}\ ] \leq 1/(n-1)$ for every B. This implies (2), because

$$\mathbb{P}[\ \overline{\mathcal{E}}\ ] = \sum_B \mathbb{P}[\ \overline{\mathcal{E}^B} \text{ and } \mathbf{B} = B\ ] = \sum_B \mathbb{P}[\ \overline{\mathcal{E}^B}\ ] \cdot \mathbb{P}[\ \mathbf{B} = B\ ] \leq \frac{1}{n-1} < \frac{\varepsilon}{2}.$$

Here, the second displayed equality is due to the independence of the events $\overline{\mathcal{E}^B}$ and $\mathbf{B} = B$, and the last inequality is due to $n > 1 + 2/\varepsilon$. □

3.3 Left and right degrees

Besides being a resilient-expander, we often need our graph to have low right-degree. This is guaranteed in a random graph by the following easy calculation:

Lemma 11. *Let ε be a positive real, let t, n, d and d' be naturals satisfying $t \geq n \geq d$ and $n(tde/nd')^{d'} < \varepsilon$, and let $\mathbf{G} = \mathbf{G}(t, n, d)$. Then*

$$\mathbb{P}[\ \mathbf{G} \text{ has right-degree smaller than } d'\ ] > 1 - \varepsilon.$$

Proof. For fixed vertices $u \in U$ and $v \in V$, the probability that (u, v) is an edge in $\mathbf{G}$ is $\binom{n-1}{d-1}/\binom{n}{d} = d/n$. Moreover, for fixed $v \in V$, these events are mutually independent as u ranges over U. By the union bound over all d'-element subsets of U, this means that the probability that the degree of v is at least d' is bounded by $\binom{t}{d'}(d/n)^{d'}$. By the union bound over v, the probability that the right-degree is at least d' is bounded by $n\binom{t}{d'}(d/n)^{d'}$. The lemma follows from the bound $\binom{t}{d'} \leq (te/d')^{d'}$ and the hypothesis that $n(tde/nd')^{d'} < \varepsilon$. □

As mentioned earlier, in our application of Lemma 9 we will have $b = O(1)$, $t = 2n$, $q = n^{1-\Omega(1)}$ and $r = \Theta(n/\log n)$, in which case the required lower bound on d is $O(\log n/\log\log n)$. Setting $d = \lceil \log n \rceil$ satisfies this lower bound and Lemma 11 gives right-degree $d' = O(\log n)$. Therefore, for the setting of parameters b, t, q and r of our interest, there exists a (q, b, r)-resilient expander with left-degree $O(\log n)$ and right-degree $O(\log n)$. Let us argue now that having a (q, b, r)-resilient expander with right-degree $O(\log n)$ but left-degree $o(\log n/\log\log n)$ is impossible.

Suppose G is an (t, n, d_L, d_R)-graph that is a (q, b, r)-resilient expander where b, t, q and r are as above and $d_R = O(\log n)$. Then there exist at least $t/(d_L \cdot d_R)$ vertices in U with pairwise disjoint neighborhoods in V. Let $\tilde{\mathbf{B}}$ be a random subset of V obtained by placing each vertex in it independently with probability r/n. For a fixed vertex $u \in U$, the probability that $\tilde{\mathbf{B}}$ contains all the neighbors of u is at least $(r/n)^{d_L}$. Moreover, these events are mutually independent for vertices from U that have pairwise disjoint neighborhoods in V. Therefore, the probability that $\tilde{\mathbf{B}}$ does not contain all the neighbors of any vertex in U is bounded by

$$\left(1 - \left(\frac{r}{n}\right)^{d_L}\right)^{\frac{t}{d_L \cdot d_R}} \leq \exp\left(-\left(\frac{r}{n}\right)^{d_L} \cdot \frac{t}{d_L \cdot d_R}\right).$$

The probability of this event for a random r-element subset $\mathbf{B} \subseteq V$ is at most a multiplicative factor $3\sqrt{r}$ bigger (see equation (6) in Section 4). Since G is a (q, b, r)-resilient expander, the probability of this event for $\mathbf{B}$ is at least $1/2$. But since $t \geq n$, $r = \Omega(n/\log n)$ and $d_R = O(\log n)$, this is possible only if d_L is $\Omega(\log n/\log\log n)$.

4 Lower bound

In this section we develop the proof of Theorem 1 as outlined in the introduction.

4.1 Killing large conjunctions

Let t be a natural such that $n < t < m$. Let $\boldsymbol{\rho} = \boldsymbol{\rho}(t)$ be the random restriction[2] on the variables of $\mathrm{PHP}_n^{m,t}$ defined by the following random experiment:

1. choose a subset $\mathbf{A} \subseteq [m]$ uniformly at random among all t-element subsets of $[m]$;
2. let $\boldsymbol{\rho}$ be the restriction that maps every R_u to 1 if $u \in \mathbf{A}$ and to 0 otherwise;
3. extend $\boldsymbol{\rho}$ to the auxiliary variables X such that $\mathrm{TH}_t(R, X)$ is satisfied;
4. extend $\boldsymbol{\rho}$ by mapping every $P_{u,v}$ with $u \in [m] \setminus \mathbf{A}$ and $v \in [n]$ to 1 independently with probability $1/2$ and to 0 otherwise.

In the following, by a *pigeon variable* we mean a variable $P_{u,v}$ for $u \in [m]$ and $v \in [n]$; we say $P_{u,v}$ *mentions* pigeon u; a formula *mentions* a pigeon if so does some variable occuring in it. For later use, note that if ρ is a realization of $\boldsymbol{\rho}$, then $\mathrm{PHP}_n^{m,t} \restriction \rho$ and PHP_n^t are the same formula up to renaming of pigeons.

Lemma 12. *Let p be a natural such that $p < t$ and $p < m - t$, and T be a term that mentions at least p many pigeons. Then*

$$\mathbb{P}\big[\ T \restriction \boldsymbol{\rho} \neq 0\ \big] \leq \left(\frac{1}{2} + \frac{t}{m-p}\right)^p.$$

Proof. Choose p literals in T mentioning pairwise different pigeons. Let P be the set of pigeons mentioned by these literals, and for every $u \in P$ let ℓ_u be the literal chosen for pigeon u. Consider the events $\mathcal{E} :=$ "$\boldsymbol{\rho}(\ell_u) \neq 0$ for all $u \in P \setminus \mathbf{A}$", and $\mathcal{F}_i :=$ "$|P \setminus \mathbf{A}| = i$", where $i \in \{0, \ldots, p\}$. Note that $\mathbb{P}[\ T \restriction \boldsymbol{\rho} \neq 0\ ] \leq \mathbb{P}[\ \mathcal{E}\ ]$ and

$$\mathbb{P}[\ \mathcal{E}\ ] = \sum_{i=0}^{p} \mathbb{P}[\ \mathcal{E} \mid \mathcal{F}_i\ ] \cdot \mathbb{P}[\ \mathcal{F}_i\ ] = \sum_{i=0}^{p} \frac{1}{2^i} \cdot \frac{\binom{p}{i}\binom{m-p}{t-p+i}}{\binom{m}{t}}.$$

For naturals $m \geq k$ we write $m^{\underline{k}}$ for the falling factorial $m^{\underline{k}} := m \cdot (m-1) \cdots (m-k+1)$. Note that our assumptions on p ensure $m - p > t - p + i > 0$. Using $0 \leq i \leq p$ and noting $m^{\underline{p}} = m^{\underline{i}} \cdot (m-i)^{\underline{p-i}}$, we have

$$\frac{\binom{m-p}{t-p+i}}{\binom{m}{t}} = \frac{(m-t)^{\underline{i}}}{m^{\underline{i}}} \cdot \frac{t^{\underline{p-i}}}{(m-i)^{\underline{p-i}}} \leq \frac{t^{\underline{p-i}}}{(m-i)^{\underline{p-i}}} \leq \left(\frac{t}{m-p}\right)^{p-i}.$$

Replacing, and using the binomial formula, the probability we want is bounded by

$$\sum_{i=0}^{p} \binom{p}{i} \left(\frac{1}{2}\right)^i \left(\frac{t}{m-p}\right)^{p-i} = \left(\frac{1}{2} + \frac{t}{m-p}\right)^p.$$

□

[2]Of course, by a random restriction we mean a random variable whose values are restrictions.

Lemma 13. *Let p and s be naturals such that $s < p < t$, and T be a term that mentions at most p many pigeons. Then*

$$\mathbb{P}\big[\ T \restriction \boldsymbol{\rho} \text{ mentions more than } s \text{ many pigeons } \big] \leq \binom{p}{s+1}\left(\frac{t}{m}\right)^{s+1}.$$

Proof. For any $s+1$ pigeon variables in T mentioning pairwise different pigeons, the probability that they all remain unset by $\boldsymbol{\rho}$ is

$$\frac{\binom{m-s-1}{t-s-1}}{\binom{m}{t}} = \frac{t^{\underline{s+1}}}{m^{\underline{s+1}}} \leq \left(\frac{t}{m}\right)^{s+1}.$$

The claim thus follows by the union bound. □

4.2 Restriction to a graph and binary encoding

Let t be a natural such that $n < t < m$ and let $G = (U, V, E)$ be a bipartite graph with $U = [t]$ and $V = [n]$. Consider the following restriction θ_G: it maps every variable $P_{u,v}$ to 0 if $(u, v) \notin E$ and is undefined on all other variables. Then $\mathrm{PHP}^t_n \restriction \theta_G$ is the CNF with clauses (1 and)

$$\begin{array}{ll} \bigvee_{v \in N_G(u)} P_{u,v} & \text{for } u \in U, \\ \neg P_{u,v} \vee \neg P_{u',v} & \text{for } (u,v),(u',v) \in E \text{ with } u \neq u'. \end{array}$$

This formula is commonly denoted by $\mathrm{PHP}(G)$ (cf. [10, 33]).

Now assume that G is a (U, V, d_L, d_R)-graph with associated function ϕ_G. Write

$$\ell := |d_L - 1|$$

for the length of the binary representation of the largest number in $[d_L]$. We introduce a new set of *binary pigeon variables* $P_{u;b}$ for $u \in U$ and $b \in [\ell]$. Again, we say that $P_{u;b}$ mentions pigeon u, and that a formula mentions the pigeons mentioned by some atom occuring in it. The intuitive meaning of an assignment to the binary pigeon variables is that pigeon u flies to hole $\phi_G(u, j)$, where j is the number whose binary representation is given by the truth values $P_{u;\ell-1}, \ldots, P_{u;0}$, where the truth value of $P_{u;0}$ is the least-significant bit. The formula $\mathrm{BPHP}(G)$ has *domain clauses* and *collision clauses*:

$$\begin{array}{ll} \bigvee_{b \in [\ell]} \neg^{\mathrm{bit}(b,j)} P_{u;b} & \text{for } (u,j) \in U \times [2^\ell] \text{ such that } (u,j) \notin \mathrm{Dom}(\phi_G), \\ \bigvee_{b \in [\ell]} \neg^{\mathrm{bit}(b,j)} P_{u;b} \vee \bigvee_{b \in [\ell]} \neg^{\mathrm{bit}(b,j')} P_{u';b} & \text{for } (u,j) \in \mathrm{Dom}(\phi_G) \text{ and } (u',j') \in \mathrm{Dom}(\phi_G) \\ & \text{such that } u \neq u' \text{ and } \phi_G(u,j) = \phi_G(u',j'). \end{array}$$

Here, for a variable X we write $\neg^0 X := X$ and $\neg^1 X := \neg X$.

The unary encoding $\mathrm{PHP}(G)$ and the binary encoding $\mathrm{BPHP}(G)$ are closely related. Indeed, the formula obtained from $\mathrm{PHP}(G)$ by substituting every variable $P_{u,v}$ by the term

$\bigwedge_{b\in[\ell]} \neg^{1-\mathrm{bit}(b,j)} P_{u;b}$, where $j \in [2^\ell]$ is such that $\phi(u,j) = v$, is the conjunction of the collision clauses of BPHP(G) and *sporadic axioms*:

$$\bigvee_{j\in J_G(u)} \bigwedge_{b\in[\ell]} \neg^{1-\mathrm{bit}(b,j)} P_{u;b} \qquad \text{for } u \in U \text{ with } J_G(u) := \{j \in [2^\ell] \mid (u,j) \in \mathrm{Dom}(\phi_G)\}.$$

The following lemma states that these sporadic axioms are redundant:

Lemma 14. *Every sporadic axiom has a strongly sound semantic DNF-proof from the domain clauses of* BPHP(G)*. The length of the proof is at most* $3 \cdot 2^\ell$ *and every term appearing in the proof mentions only one pigeon.*

Proof. Fix $u \in U$. For $1 \le i \le \ell$, let F_i be the DNF formula $\bigvee_{j\in[2^i]} \bigwedge_{b\in[i]} \neg^{1-\mathrm{bit}(b,j)} P_{u;b}$. Then $F_1 = (P_{u;0} \vee \neg P_{u;0})$ and, for $1 \le i \le \ell - 1$, the formula F_{i+1} is strongly implied by F_i and the elementary tautology $(P_{u;i} \vee \neg P_{u;i})$. It follows that F_ℓ has a strongly sound proof of length $2\ell - 1$.

The sporadic axiom is obtained from F_ℓ by eliminating the terms for j such that $(u,j) \notin \mathrm{Dom}(\phi_G)$ one after the other. Note that a DNF with a term $\bigwedge_{b\in[\ell]} \neg^{1-\mathrm{bit}(b,j)} P_{u;b}$ such that $(u,j) \notin \mathrm{Dom}(\phi_G)$ together with the domain clause for (u,j) strongly implies the DNF obtained by deleting the term. In total this needs at most another $2 \cdot 2^\ell$ steps. □

We note also that the elementary tautologies $(P_{u,v} \vee \neg P_{u,v})$ for $(u,v) \in E$ become DNFs that we call *assignment tautologies*:

$$(\bigwedge_{b\in[\ell]} \neg^{1-\mathrm{bit}(b,j)} P_{u;b}) \vee (\bigvee_{b\in[\ell]} \neg^{\mathrm{bit}(b,j)} P_{u;b}) \qquad \text{for } (u,j) \in \mathrm{Dom}(\phi_G).$$

These assignment tautologies are also redundant:

Lemma 15. *Every assignment tautology has a strongly sound semantic DNF-proof from no assumptions. The length of the proof is* $2\ell - 1$ *and every term appearing in the proof mentions only one pigeon.*

Proof. Fix $(u,j) \in \mathrm{Dom}(\phi_G)$. For $1 \le i \le \ell$, let G_i be the formula $(\bigwedge_{b\in[i]} \neg^{1-\mathrm{bit}(b,j)} P_{u;b}) \vee (\bigvee_{b\in[i]} \neg^{\mathrm{bit}(b,j)} P_{u;b})$. Then $G_1 = (P_{u;0} \vee \neg P_{u;0})$ and G_ℓ is the assignment tautology corresponding to (u,j). Moreover, for $1 \le i \le \ell - 1$, the formula G_{i+1} is strongly implied by G_i and the elementary tautology $(P_{u;i} \vee \neg P_{u;i})$. It follows that G_ℓ has a strongly sound proof of length $2\ell - 1$. □

4.3 Killing large disjunctions

Let t be a natural such that $n < t < m$ and let $G = (U, V, E)$ be a (t, n, d_L, d_R)-graph with associated function ϕ_G. Let r be a natural such that $1 \le r \le n$. We define a random restriction $\boldsymbol{\mu} = \boldsymbol{\mu}(G, r)$ on the variables of BPHP(G) by the following random experiment:

1. independently for every $v \in V$, choose a pigeon $\mathbf{Q}_v \in N_G(v)$ uniformly at random;
2. independently, choose an r-element subset $\mathbf{B} \subseteq V$ uniformly at random;

3. let $\mathbf{M} := \{(\mathbf{Q}_v, v) \mid v \in \mathbf{B} \text{ and } \mathbf{Q}_v \neq \mathbf{Q}_{v'} \text{ for all } v' \in \mathbf{B} \setminus \{v\}\}$;
4. let $\boldsymbol{\mu}$ be the restriction associated with the matching $\mathbf{M}$.

Here, the restriction μ associated with a matching M of G maps for every $(u, v) \in M$ and $b \in [\ell]$ the variable $P_{u;b}$ to $\mathrm{bit}(b, j)$ where j is such that $\phi_G(u, j) = v$; it is undefined on all other variables. Call a formula F *matching-satisfiable (in G)* if $F \restriction \mu = 1$ for some such restriction μ.

Two formulas F and F' are *very disjoint (in G)* if $N_G(P)$ and $N_G(P')$ are disjoint, where $P \subseteq U$ and $P' \subseteq U$ are the sets of pigeons mentioned by F and F' respectively.

Lemma 16. *Let s and w be naturals such that $r \geq s \geq 1$ and $w \geq 1$. Further, let $F = \bigvee \Gamma$ where Γ contains at least w matching-satisfiable, pairwise very disjoint formulas each mentioning at most s pigeons. Then*

$$\mathbb{P}\big[\ F \restriction \boldsymbol{\mu} \neq 1\ \big] \leq 3\sqrt{r} \cdot \exp\left(-w \cdot \left(\frac{r}{d_R \cdot n}\right)^s \cdot \left(1 - \frac{r}{n}\right)^{d_L \cdot s}\right).$$

Proof. Define the random variables $\tilde{\mathbf{B}}$, $(\tilde{\mathbf{Q}}_v)_{v \in V}$, $\tilde{\mathbf{M}}$, $\tilde{\boldsymbol{\mu}}$ similarly as above but letting $\tilde{\mathbf{B}}$ be the random subset of V that contains every $v \in V$ independently with probability r/n. Let $\tilde{\mathbf{B}}_v$ denote the indicator variable for the event that $v \in \tilde{\mathbf{B}}$; note that the indicator variables are independent.

Fix a matching-satisfiable formula $F' \in \Gamma$ mentioning at most s pigeons. Choose a minimal matching M such that $F' \restriction \mu = 1$ where μ is the restriction associated with M. Write $M_0 := \mathrm{Dom}(M)$ and $M_1 := \mathrm{Im}(M)$. Then, by minimality of M, the domain M_0 is included in the set of pigeons $P \subseteq U$ mentioned by F'. Observe that the event that $F' \restriction \tilde{\boldsymbol{\mu}} = 1$ is implied by the event that $M \subseteq \tilde{\mathbf{M}}$. The latter event is implied by the intersection of

$\mathcal{E}_1 :=$ "$\tilde{\mathbf{B}}_v = 1$ for every $v \in M_1$", and
$\mathcal{E}_2 :=$ "$\tilde{\mathbf{Q}}_v = M^{-1}(v)$ for every $v \in M_1$"

and the event that $\tilde{\mathbf{Q}}_v \notin M_0$ for every $v \in \tilde{\mathbf{B}} \setminus M_1$. Thus it is implied by the intersection of $\mathcal{E}_1$, $\mathcal{E}_2$ and

$\mathcal{E}_3 :=$ "$\tilde{\mathbf{B}}_v = 0$ for every $v \in N_G(M_0) \setminus M_1$".

Now, the probability of $\mathcal{E}_1$ is at least $(r/n)^s$, the probability of $\mathcal{E}_2$ is at least $(1/d_R)^s$, and the probability of $\mathcal{E}_3$ is at least $(1 - r/n)^{d_L \cdot s}$, the last because $|N_G(M_0) \setminus M_1| \leq d_L \cdot s$. These three events are independent. Hence

$$\mathbb{P}[\ \mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3\ ] \geq \left(\frac{r}{n}\right)^s \cdot \left(\frac{1}{d_R}\right)^s \cdot \left(1 - \frac{r}{n}\right)^{d_L \cdot s} =: p.$$

The event $\mathcal{E}_1 \cap \mathcal{E}_2 \cap \mathcal{E}_3$ depends only on the variables $\tilde{\mathbf{Q}}_v$ and $\tilde{\mathbf{B}}_v$ with $v \in N_G(M_0) \subseteq N_G(P)$. Thus, for a family of pairwise very disjoint formulas in Γ, the events are independent. Using the assumption of the lemma,

$$\mathbb{P}[\ F \restriction \tilde{\boldsymbol{\mu}} \neq 1\ ] \leq (1 - p)^w \leq \exp(-wp). \tag{4}$$

Writing $B(m,q)(k) = \binom{m}{k} q^k (1-q)^{m-k}$ for the binomial distribution, we have

$$\mathbb{P}[\ F \restriction \tilde{\boldsymbol{\mu}} \neq 1\ ] \geq \mathbb{P}[\ |\tilde{\mathbf{B}}| = r\ ] \cdot \mathbb{P}[\ F \restriction \tilde{\boldsymbol{\mu}} \neq 1 \mid |\tilde{\mathbf{B}}| = r\ ] = B\left(n, \frac{r}{n}\right)(r) \cdot \mathbb{P}[\ F \restriction \boldsymbol{\mu}\ ]. \quad (5)$$

Using Robbins' [32] version of Stirling's formula, one can derive the following bound (see also [11, p.4, Eq. (1.5)]):

$$B\left(n, \frac{r}{n}\right)(r) \geq \frac{1}{e^{1/6}} \cdot \frac{1}{\sqrt{2\pi}} \cdot \left(\frac{n}{r(n-r)}\right)^{1/2} \geq \frac{1}{3}\frac{1}{\sqrt{r}}. \quad (6)$$

Combining (4), (5) and (6) yields the lemma. □

4.4 Switching lemma

Associate with a DNF F the hypergraph $\mathcal{H}(F)$ which has as universe the set of variables of F and which has for each term T in F a hyperedge consisting in the variables of T. The *covering number* $\mathrm{cv}(F)$ of F is the size of a smallest hitting set of $\mathcal{H}(F)$.

Lemma 17. *Let F be a k-DNF in the binary pigeon variables. Then F contains at least $\frac{\mathrm{cv}(F)}{\ell \cdot k \cdot d_L \cdot d_R}$ many pairwise very disjoint terms.*

Proof. Let $\mathcal{T}$ be a maximal family of very disjoint terms in F. Let P be the set of pigeons mentioned by $\bigvee \mathcal{T}$. Then the set of all pigeon variables mentioning pigeons in $N_G(N_G(P))$ is a hitting set of $\mathcal{H}(F)$. Noting that $N_G(N_G(P))$ has cardinality at most $|P| \cdot d_L \cdot d_R$ and $|P| \leq |\mathcal{T}| \cdot k$ we get

$$\mathrm{cv}(F) \leq \ell \cdot |N_G(N_G(P))| \leq |\mathcal{T}| \cdot \ell \cdot k \cdot d_L \cdot d_R$$

and the lemma follows. □

Interest in the covering number stems from the following lemma proved by Segerlind, Buss and Impagliazzo [34] (see also the survey [33, Corollary 9.3]).

Lemma 18 ([34])**.** *Let k, h and c be positive naturals and let γ be a positive real. Let Γ be a set of k-DNFs that is closed under restrictions and assume that $\boldsymbol{\sigma}$ is a random restriction such that $\mathbb{P}[\ F \restriction \boldsymbol{\sigma} \neq 1\ ] \leq c \cdot 2^{-\gamma \cdot \mathrm{cv}(F)}$ for every F in Γ. Then*

$$\mathbb{P}\big[\ h(F \restriction \boldsymbol{\sigma}) > h\ \big] \leq c \cdot k \cdot 2^{-(\gamma/4)^k \cdot h}$$

for every F in Γ.

In the above statement, recall $h(F)$ denotes the minimal height of a decision tree representing the formula F. For what follows, we need to define a generalization of the notion of representation by decision trees, and the corresponding notion of height.

Let F be a formula, let $\mathcal{C}$ be a set of restrictions, and let T be a decision tree. We say that T $\mathcal{C}$-represents F if it queries only variables from F and

1. for every 1-branch π of T and every $\mu \in \mathcal{C}$ extending π we have $F \restriction \mu = 1$,
2. for every 0-branch π of T and every $\mu \in \mathcal{C}$ extending π we have $F \restriction \mu \neq 1$.

The minimal height of a decision tree that $\mathcal{C}$-represents F is denoted $h(F;\mathcal{C})$.

Remark 19. If $\mathcal{C}' \subseteq \mathcal{C}$ then every tree that $\mathcal{C}$-represents F also $\mathcal{C}'$-represents F; in particular, $h(F;\mathcal{C}') \leq h(F;\mathcal{C})$. For all $\mathcal{C}$ we have $h(F;\mathcal{C}) \leq h(F)$ (by Lemma 6). If $\mathcal{C} = \emptyset$, then $h(F;\mathcal{C}) = 0$. A formula without variables is $\mathcal{C}$-represented by and only by the one node tree labeled with its truth value.

We need some notation. For a matching M in G let μ_M denote the restriction associated with M (cf. Section 4.3); the set of restrictions of this form is $\mathcal{M}(G)$. If G is clear from context and M is a matching in G, we write

$$\mathcal{M}(M) := \mathcal{M}(G \setminus (\mathrm{Dom}(M) \cup \mathrm{Im}(M))).$$

Observe that μ belongs to $\mathcal{M}(M)$ if and only if $\mu \cup \mu_M$ belongs to $\mathcal{M}(G)$.

For a DNF-formula F in the variables of BPHP(G), let $F_{\mathcal{M}(G)}$ be the disjunction of all terms T of F for which there exists a μ in $\mathcal{M}(G)$ such that $T \restriction \mu = 1$. Observe that the terms of $F_{\mathcal{M}(G)}$ are precisely the matching-satisfiable (in G) terms of F. We need the following lemma:

Lemma 20. *Let M be a matching in G and F a DNF. Every decision tree that $\mathcal{M}(M)$-represents $F_{\mathcal{M}(G)} \restriction \mu_M$ also $\mathcal{M}(M)$-represents $F \restriction \mu_M$.*

Proof. Let T $\mathcal{M}(M)$-represent $F_{\mathcal{M}(G)} \restriction \mu_M$.

Let π be a 1-branch of T and let μ_N in $\mathcal{M}(M)$ be such that $\mu_N \supseteq \pi$. Since T $\mathcal{M}(M)$-represents $F_{\mathcal{M}(G)} \restriction \mu_M$ we have that $(F_{\mathcal{M}(G)} \restriction \mu_M) \restriction \mu_N = 1$. Since $F_{\mathcal{M}(G)}$ is obtained by deleting some terms from F we get also $(F \restriction \mu_M) \restriction \mu_N = 1$.

Let π be a 0-branch of T and let μ_N in $\mathcal{M}(M)$ be such that $\mu_N \supseteq \pi$. Assume for contradiction that $(F \restriction \mu_M) \restriction \mu_N = 1$. Then there is a term T in F such that $(T \restriction \mu_M) \restriction \mu_N = 1$. But $M \cup N$ is a matching in G and $(T \restriction \mu_M) \restriction \mu_N = T \restriction \mu_{M \cup N}$. Hence T is matching-satisfiable (in G) and thus appears in $F_{\mathcal{M}(G)}$. Hence $(F_{\mathcal{M}(G)} \restriction \mu_M) \restriction \mu_N = 1$ and this contradicts the fact that T $\mathcal{M}(M)$-represents $F_{\mathcal{M}(G)} \restriction \mu_M$. □

4.5 Matching game

In the next section we show that if G is a good expander, then all the refutations of BPHP(G) involve some formula that cannot be represented by a shallow decision tree. For its proof we use the *matching games* from [9] later simplified in [5]. Here we provide even cleaner proofs.

Let G be a (U, V, d_L, d_R)-graph. For $S \subseteq U$ and $T \subseteq V$, we say that S *is matchable into* T if there exists a matching M of G with $S \subseteq \mathrm{Dom}(M)$ and $\mathrm{Im}(M) \subseteq T$. If S is not matchable into T but every proper subset of S is, we call it *minimally non-matchable*. For a matching M and a natural $q > 0$, we say that M is *q-extendible* if every $S \subseteq U \setminus \mathrm{Dom}(M)$ of cardinality at most q is matchable into $V \setminus \mathrm{Im}(M)$.

Lemma 21. *Let $q > 0$ be a natural. If M is a q-extendible matching and (u, v) is an edge in M, then $M \setminus \{(u, v)\}$ is a q-extendible matching.*

Proof. Write $M_0 := \mathrm{Dom}(M)$ and $M_1 := \mathrm{Im}(M)$ and note that $u \in M_0$ and $v \in M_1$. Let S' be a subset of $U \setminus (M_0 \setminus \{u\})$ of cardinality at most q. We need to show that S' is matchable into $V \setminus (M_1 \setminus \{v\})$. We consider two cases: $u \in S'$ and $u \notin S'$. In case $u \in S'$, using that $u \in M_0$, we have that $S' \setminus \{u\}$ is a subset of $U \setminus M_0$ of cardinality at most q. Since M is q-extendible, $S' \setminus \{u\}$ is matchable into $V \setminus M_1$. But then, using that $v \in M_1$, the set S' is also matchable into $V \setminus (M_1 \setminus \{v\})$ by adding (u, v) to the matching that witnesses this. In case $u \notin S'$ then S' is a subset of $U \setminus M_0$ of cardinality at most q. Since M is q-extendible we conclude that S' is matchable into $V \setminus M_1$, and hence into $V \setminus (M_1 \setminus \{v\})$. □

For a natural $q > 0$ and a real $b > 0$, the graph G is a (q, b)-*expander* if $|N_G(S)| \geq (1+b)|S|$ for every $S \subseteq U$ of cardinality at most q.

Lemma 22. *Let $q > 0$ be a natural and $b > 0$ a real. If G is a (q, b)-expander, M is a q-extendible matching with $|M| < \lfloor qb/d_L \rfloor$ and $u \in U \setminus \mathrm{Dom}(M)$, then there exists $v \in N_G(u) \setminus \mathrm{Im}(M)$ such that $M \cup \{(u, v)\}$ is a q-extendible matching.*

Proof. Again write $M_0 := \mathrm{Dom}(M)$ and $M_1 := \mathrm{Im}(M)$. Let $v_1, \ldots, v_l$ be an enumeration of $N_G(u) \setminus M_1$. Since M is q-extendible and $q \geq 1$, we have that $\{u\}$ is matchable into $V \setminus M_1$, so $l \geq 1$. Clearly, $M \cup \{(u, v_i)\}$ is a matching for every $i \in \{1, \ldots, l\}$. Assume for contradiction that $M \cup \{(u, v_i)\}$ is not q-extendible for any $i \in \{1, \ldots, l\}$. For every $i \in \{1, \ldots, l\}$ let S_i be a subset of $U \setminus (M_0 \cup \{u\})$ of cardinality at most q that is minimally non-matchable into $V \setminus (M_1 \cup \{v_i\})$. By Hall's Theorem and the minimality of S_i we have $|N_G(S_i) \setminus (M_1 \cup \{v_i\})| < |S_i|$, and hence $|N_G(S_i)| < |S_i| + (qb/d_L - 1) + 1$. On the other hand $|S_i| \leq q$, and hence $|N_G(S_i)| \geq (1 + b)|S_i|$ by expansion of G. These together imply $|S_i| < q/d_L$ and hence $|S_i| < q/l$ because $1 \leq l \leq d_L$. Since this holds for every $i \in \{1, \ldots, l\}$ we get $|S| \leq q$ for $S := \bigcup_{i=1}^{l} S_i \cup \{u\}$. Since M is q-extendible and $S \subseteq U \setminus M_0$ we conclude that S is matchable into $V \setminus M_1$. A matching M' witnessing this matches u to v_i for some $i \in \{1, \ldots, l\}$. As M' matches S_i into $V \setminus M_1$ while S_i is non-matchable into $V \setminus (M_1 \cup \{v_i\})$, necessarily M' matches some $u_i \in S_i$ to v_i. But this contradicts M' to be a matching because $u_i \neq u$ as $u \notin S_i$. □

4.6 Adversary argument

Let G be a (U, V, d_L, d_R)-graph. We derive a lower bound on the height of formulas in a refutation of $\mathrm{BPHP}(G)$ provided G is suitably expanding. This is done by an adversary argument (cf. [27]) based on Lemmas 21 and 22.

Recall that $\mathcal{M}(G)$ denotes the set of restrictions μ_M associated with matchings M in G (cf. Section 4.4).

Lemma 23. *Let $q > 1$ be a natural and $b > 0$ a real and let G be a (q, b)-expander. Assume $F_0, \ldots, F_{s-1}$ is a strongly sound refutation of $\mathrm{BPHP}(G)$ such that every F_i with $i \in [s]$ is a*

formula in the variables of BPHP(G). *Then there exists* $i \in [s]$ *such that*

$$h(F_i; \mathcal{M}(G)) > \tfrac{1}{3}\lfloor qb/d_L \rfloor.$$

Proof. For the sake of contradiction assume that for every $i \in [s]$ there exists a decision tree T_i of height at most $\frac{1}{3}\lfloor qb/d_L \rfloor$ that $\mathcal{M}(G)$-represents F_i. Since $F_{s-1} = 0$, the tree T_{s-1} is the tree with one node labeled 0.

Claim. Let $i \in [s]$ be such that F_i is neither a clause in BPHP(G) nor an elementary tautology of the form $(P_{u;b} \vee \neg P_{u;b})$ for a $u \in U$ and $b \in [\ell]$. Further assume that M is a q-extendible matching of size $|M| \leq \frac{1}{3}\lfloor qb/d_L \rfloor$ such that μ_M extends some 0-branch of T_i. Then there exists $i' \in [i]$ and a q-extendible matching M' of size $|M'| \leq \frac{1}{3}\lfloor qb/d_L \rfloor$ such that $\mu_{M'}$ extends some 0-branch of $T_{i'}$.

Proof of the Claim. Let i and M accord the assumption. In particular $F_i \restriction \mu_M \neq 1$. Then F_i is not strongly implied by $\emptyset$. Hence there exist (not necessarily distinct) j and k in $[i]$ such that F_i is strongly implied by $\{F_j, F_k\}$. Choose T for T_j and T_k according Lemma 7. In particular, T has height at most $\frac{2}{3}\lfloor qb/d_L \rfloor$. Given a path π in T starting at the root, call a matching *appropriate for* π if it is q-extendible, contains M, its associated restriction extends π, and its domain is $\mathrm{Dom}(M) \cup P(\pi)$, where $P(\pi)$ is the set of pigeons mentioned by some variable queried in π.

Subclaim. There exists a branch π of T and a matching M_π appropriate for π.

The Subclaim implies the Claim: by $M_\pi \supseteq M$ we have $\mu_{M_\pi} \supseteq \mu_M$, so μ_{M_π} extends a 0-branch of T_i (as μ_M does). Since T_i $\mathcal{M}(G)$-represents F_i, we have $F_i \restriction \mu_{M_\pi} \neq 1$. As F_i is strongly implied by $\{F_j, F_k\}$ there is $i'' \in \{j, k\}$ such that $F_{i''} \restriction \mu_{M_\pi} \neq 1$. Then μ_{M_π} and hence π does not extend a 1-branch of $T_{i''}$. By choice of T (Lemma 7 (a)) thus π is a 0-branch. Then there is $i' \in \{j, k\}$ such that π extends a 0-branch π' of $T_{i'}$ (Lemma 7 (b)). Let M' be the restriction of M_π to $P(\pi')$. Then M' is a q-extendible matching (by Lemma 21), $|M'| \leq \frac{1}{3}\lfloor qb/d_L \rfloor$ (as $|P(\pi')| \leq \frac{1}{3}\lfloor qb/d \rfloor$), and clearly $\mu_{M'}$ extends π', that is, M' and i' satisfy the Claim.

We are left to prove the Subclaim. Observe that M is an appropriate matching for the path π consisting only in the root of T. To prove the subclaim it thus suffices to show that if we have a path π with appropriate matching M_π such that π does not lead to a leaf of T then we can extend π by one node t such that there is an appropriate matching $M_{\pi t}$ for the path πt. So let π and M_π be as stated, say, π leads to an inner node t of T querying the variable $P_{u;b}$. We distinguish two cases. In case $u \in \mathrm{Dom}(M_\pi)$ then μ_{M_π} evaluates $P_{u;b}$; in this case we prolongue π by the corresponding successor t' of t and let $M_{\pi t'} := M_\pi$. In case $u \notin \mathrm{Dom}(M_\pi)$ we look for some v such that $M_\pi \cup \{(u, v)\}$ is a q-extendible matching and then proceed as in the first case. Such a v can be found because $\mathrm{Dom}(M_\pi) = \mathrm{Dom}(M) \cup P(\pi)$ has cardinality at most

$$|\mathrm{Dom}(M)| + |P(\pi)| \leq \tfrac{1}{3}\lfloor qb/d_L \rfloor + \tfrac{2}{3}\lfloor qb/d_L \rfloor - 1 < \lfloor qb/d_L \rfloor,$$

and Lemma 22 applies. Here we use that $|P(\pi)|$ is bounded by the length of π, and this is at most $\frac{2}{3}\lfloor qb/d_L \rfloor - 1$ because π leads to an internal node of T and T has height at most $\frac{2}{3}\lfloor qb/d_L \rfloor$. $\dashv$

We observe that $M := \emptyset$ and $i := s-1$ satisfy the assumptions of the Claim: $F_{s-1} = 0$ is neither a clause in BPHP(G) nor an elementary tautology, T_{s-1} is the tree with one node labeled 0, so $\mu_\emptyset = \emptyset$ extends a 0-branch of T_{s-1}, and obviously $|\emptyset| = 0 \leq \frac{1}{3}\lfloor qb/d_L \rfloor$; finally, that $\emptyset$ is q-extendible follows from G being a (q,b)-expander and Hall's Theorem.

The Claim implies that there exist $i \in [s]$ and M such that F_i is a clause in BPHP(G) or an elementary tautology of the form $(P_{u;b} \vee \neg P_{u;b})$ for a $u \in U$ and $b \in [\ell]$, and M is a q-extendible matching such that μ_M extends a 0-branch of T_i. By $\mathcal{M}(G)$-representation there is no matching $M' \supseteq M$ such that $F_i \restriction \mu_{M'} = 1$. We get the desired contradiction by showing that such M' indeed exists. We have three cases.

Case 1. F_i is a domain clause, say for $(u,j) \notin \mathrm{Dom}(\phi_G)$. Since M is q- and hence 1-extendible there exists a matching $M' \supseteq M$ such that $u \in \mathrm{Dom}(M')$. Then there is j' such that $\phi(u,j') = M'(u)$ and in particular $(u,j') \in \mathrm{Dom}(\phi_G)$. Hence $j \neq j'$ and there is $b \in [\ell]$ such that $\mathrm{bit}(b,j) \neq \mathrm{bit}(b,j')$. Then $\mu_{M'}$ evaluates $P_{u;b}$ to $\mathrm{bit}(b,j')$, and hence $\neg^{\mathrm{bit}(b,j)} P_{u;b} \restriction \mu_{M'} = 1$. Then $F_i \restriction \mu_{M'} = 1$.

Case 2. F_i is a collision clause, say for u, u', j, j' with $u \neq u'$ and $\phi_G(u,j) = \phi_G(u,j')$. Since M is q- and hence 2-extendible there exists a matching $M' \supseteq M$ such that $u, u' \in \mathrm{Dom}(M')$. Since M' is a matching, $M'(u) \neq \phi_G(u,j)$ or $M'(u') \neq \phi_G(u',j')$. Assume the first and choose j'' such that $M'(u) = \phi_G(u,j'')$. Then $j \neq j''$, so $\mathrm{bit}(b,j) \neq \mathrm{bit}(b,j'')$ for some $b \in [\ell]$. As above, this implies $\neg^{\mathrm{bit}(b,j)} P_{u;b} \restriction \mu_{M'} = 1$, so $F_i \restriction \mu_{M'} = 1$.

Case 3. F_i is an elementary tautology of the form $(P_{u;b} \vee \neg P_{u;b})$ for $u \in U$ and $b \in [\ell]$. Since M is q- and hence 1-extendible there exists a matching $M' \supseteq M$ such that $u \in \mathrm{Dom}(M')$. Then $\mu_{M'}$ is defined on $P_{u;b}$ and $F_i \restriction \mu_{M'} = 1$ follows. □

4.7 Proof size lower bound

We prove Theorem 1. Let $\epsilon > 0$ be arbitrary and write

$$m := n^2,\ t := 2n,\ s := (\log n)^{1/2-\epsilon}.$$

Assume for the sake of contradiction that there exists infinitely many n such that $\mathrm{PHP}_n^{m,t}$ has a strongly sound semantic DNF-refutation $R = R_n$ of size at most n^s. For the next claim recall the random restriction $\boldsymbol{\rho} = \boldsymbol{\rho}(t)$ from Section 4.1.

Claim 1. There exists a realization ρ of $\boldsymbol{\rho}$ such that every term in every DNF in $R \restriction \rho$ mentions at most s pigeons.

Proof of Claim 1: Call a term *long* if it mentions more than $p := 2s\log(n)$ pigeons, and *short* otherwise. By Lemma 12, a long term T does not restrict to 0 (under $\boldsymbol{\rho}$) with probability at most

$$\left(\frac{1}{2} + \frac{t}{m-p}\right)^p \leq \frac{1}{2^p} \cdot e^{\frac{tp}{2(m-p)}}.$$

But this is smaller than $n^{-s} \cdot 1/2$ noting $\frac{tp}{2(m-p)} \approx 0$ for large enough n. By the union bound, with probability bigger than $1/2$ every long term of R restricts under $\boldsymbol{\rho}$ to 0.

By Lemma 13, a short term restricts to one mentioning more than s many pigeons with probability at most

$$\binom{p}{s+1} \cdot \left(\frac{t}{m}\right)^{s+1} \leq \left(\frac{pt}{m}\right)^{s+1}.$$

But this is smaller than $n^{-s} \cdot 1/2$ for sufficiently large n. By the union bound, with probability bigger than $1/2$ every short term of R restricts to one mentioning at most s pigeons. The claim follows. $\dashv$

Choose ρ according Claim 1. We already observed in Section 4.1 that, up to some renaming of pigeons, $R \restriction \rho$ is a DNF-refutation of PHP_n^t of size at most n^s.

Set

$$b := 1,\ q := \lceil \sqrt{n} \rceil,\ r := \lceil n/\log n \rceil,\ d_L := \lceil \log n \rceil,\ d_R := 7\lceil \log n \rceil.$$

Recall for later use that $\ell := |d_L - 1|$ and therefore ℓ is $O(\log\log n)$. Assuming n is sufficiently large the hypotheses of Lemmas 9 and 11 are satisfied for $\varepsilon := 1/2$ and imply the existence of a (U, V, d_L, d_R)-graph G that is a (q, b, r)-resilient expander where $U = [t]$ and $V = [n]$.

Recall the restriction θ_G from Section 4.2. There we observed that $\mathrm{PHP}_n^t \restriction \theta_G$ is $\mathrm{PHP}(G)$, so $(R \restriction \rho) \restriction \theta_G$ is a strongly sound semantic refutation of $\mathrm{PHP}(G)$ of size and hence length at most n^s (Lemma 3 (1)). Let ϕ_G be a map associated with G as in Section 2.1. To $(R \restriction \rho) \restriction \theta_G$ apply the substitution mapping $P_{u,v}$ to the ℓ-term $\bigwedge_{b \in [\ell]} \neg^{1-\mathrm{bit}(b,j)} P_{u;b}$, where j is such that $\phi_G(u, j) = v$. By the discussions just before Lemmas 14 and 15, the substitution turns the clauses in $\mathrm{PHP}(G)$ into collision clauses of $\mathrm{BPHP}(G)$ and sporadic axioms, and the elementary tautologies $(P_{u,v} \vee \neg P_{u,v})$ for $(u, v) \in E$ into assignment tautologies. By Lemmas 14 and 15 we can add proofs of the sporadic axioms from the domain clauses of $\mathrm{BPHP}(G)$, and of the assignment tautologies from no assumptions; this way we get a semantic refutation R' of $\mathrm{BPHP}(G)$ of length at most $n^{c_1 \cdot s}$ for some constant c_1. By Lemma 3 (2) and since the added proofs are strongly sound, this refutation R' is again strongly sound.

Every term in every DNF in $(R \restriction \rho) \restriction \theta_G$ mentions at most s pigeons and becomes, after the substitution, an ℓ-CNF mentioning at most s pigeons. The additional proofs added for the sporadic axioms and the assignment tautologies mention only one pigeon. Hence, all the formulas in R' are disjunctions of ℓ-CNFs each mentioning at most s pigeons. By Remark 4, each such formula is strongly equivalent to a DNF with terms that mention at most s pigeons. Since there are at most $s \cdot \ell$ binary pigeon variables mentioning some fixed set of at most s pigeons, this DNF is an $\lfloor s \cdot \ell \rfloor$-DNF whose terms still mention at most s pigeons. Let R'' be the strongly sound semantic refutation that results from replacing each formula in R' by its strongly equivalent $\lfloor s \cdot \ell \rfloor$-DNF whose terms mention at most s pigeons. Note this does not increase the length, so R'' has length at most $n^{c_1 \cdot s}$.

For the next claim, let $\mathbf{B}$ and $\boldsymbol{\mu}$ be random variables defined for G as in Section 4.3.

Claim 2. There exists a realization (B, μ) of $(\mathbf{B}, \boldsymbol{\mu})$ such that

(a) $h(F_{\mathcal{M}(G)} \restriction \mu) \leq \frac{1}{3}\lfloor qb/d_L \rfloor$ for all F in R'', and

(b) $G \setminus B$ is a $(q, 1)$-expander.

Proof of Claim 2. Note a random $\mathbf{B}$ satisfies (b) with probability bigger than $1/2$ because G is (q, b, r)-resilient. It thus suffices to show that for any F in R'' we have

$$n^{c_1 \cdot s} \cdot \mathbb{P}[\ h(F_{\mathcal{M}(G)} \restriction \boldsymbol{\mu}) > \tfrac{1}{3}\lfloor qb/d_L \rfloor\ ] \leq \tfrac{1}{2}. \tag{7}$$

To prove this we intend to apply Lemma 18 taking the random restriction $\boldsymbol{\mu}$ for $\boldsymbol{\sigma}$ and taking as Γ the set of k-DNFs in the variables of BPHP(G) all of whose terms are matching-satisfiable (in G) and mention at most s pigeons. Observe that if F is a k-DNF all of whose terms mention at most s pigeons, then $F_{\mathcal{M}(G)}$ belongs to Γ. By Lemmas 16 and 17, the assumptions of Lemma 18 are satisfied for

$$k := \lfloor s \cdot \ell \rfloor,\ h := \lfloor \tfrac{1}{3} qb/d_L \rfloor,\ c := \lceil 3\sqrt{r} \rceil,$$

and

$$\gamma := \left(\frac{r}{d_R \cdot n}\right)^s \cdot \left(1 - \frac{r}{n}\right)^{d_L \cdot s} \cdot \frac{\log(e)}{\ell \cdot k \cdot d_L \cdot d_R}.$$

Thus by Lemma 18 we have that $\mathbb{P}[\ h(F_{\mathcal{M}(G)} \restriction \boldsymbol{\mu}) > \frac{1}{3}\lfloor qb/d_L \rfloor\ ]$ is at most $c \cdot k \cdot 2^{-(\gamma/4)^k \cdot h}$. Note that if n is sufficiently large, then $(1 - r/n)^{d_L \cdot s} \geq (1/e)^{c_2 \cdot s}$ for some constant $c_2 > 0$. It is then easy to see that $\gamma/4 \geq (1/\log n)^{c_3 \cdot s}$, and hence $(\gamma/4)^k \geq (1/\log n)^{c_3 \cdot s^2 \cdot \ell} \geq n^{-1/(\log n)^\epsilon}$ for some other constant $c_3 > 0$. As $h \geq n^{1/3}$ we get $(\gamma/4)^k \cdot h \geq n^{1/4}$ for sufficiently large n. Noting $c \cdot k \leq n$, then (7) follows. $\dashv$

Choose (B, μ) according to Claim 2 and recall $\mu = \mu_M$ for some matching M in G. Let $G' := G \setminus (\mathrm{Dom}(M) \cup \mathrm{Im}(M))$, and let $U' := U \setminus \mathrm{Dom}(M)$ and $V' := V \setminus \mathrm{Im}(M)$. Recall further that R'' is a strongly sound semantic refutation of BPHP(G).

Claim 3. $R'' \restriction \mu_M$ is a refutation of the union of the set of collision clauses of BPHP(G') and the set of disjunctions of at most two domain clauses of BPHP(G').

Proof of Claim 3. We analyze how every clause C of BPHP(G) and every elementary tautology $(P_{u;b} \vee \neg P_{u;b})$ for $u \in U$ and $b \in [\ell]$ restricts under μ_M.

Assume C is a collision clause for $(u, j) \in \mathrm{Dom}(\phi_G)$ and $(u', j') \in \mathrm{Dom}(\phi_G)$ with $u \neq u'$ and $\phi_G(u, j) = \phi_G(u', j') = v$. If v is not in $\mathrm{Im}(M)$ and neither u nor u' are in $\mathrm{Dom}(M)$, then C is a collision clause of BPHP(G') since (u, j) and (u', j') both belong to $\mathrm{Dom}(\phi_{G'})$ and $\phi_{G'}(u, j) = \phi_{G'}(u', j') = v$. This is ensured by the definition of the map associated to a restricted graph (see Section 2.1). If v is in $\mathrm{Im}(M)$, we need to distinguish four cases. If both u and u' are outside $\mathrm{Dom}(M)$, then C is a disjunction of two domain clauses of BPHP(G') since $(u, j) \notin \mathrm{Dom}(\phi_{G'})$ and $(u', j') \notin \mathrm{Dom}(\phi_{G'})$. If exactly one of u and u' is in $\mathrm{Dom}(M)$, say u, and $M(u) \neq v$, then $C \restriction \mu_M = 1$ since if j'' is such that $\phi_G(u, j'') = M(u)$ then the binary representations of j and j'' differ in at least one bit. If exactly one of u and u' is in $\mathrm{Dom}(M)$, say u, and $M(u) = v$, then $C \restriction \mu_M = \bigvee_{b \in [\ell]} \neg^{\mathrm{bit}(b, j')} P_{u';b}$ and this is a domain clause of BPHP(G') again because $(u', j') \notin \mathrm{Dom}(\phi_{G'})$. Finally, if both u and u' are in $\mathrm{Dom}(M)$, then $C \restriction \mu_M = 1$ since $M(u) \neq M(u')$ and if j'' and j''' are such that $\phi_G(u, j'') = M(u)$ and $\phi_G(u, j''') = M(u')$, then the binary representations of j and j'' differ in at least one bit, or the binary representations of j' and j''' differ in at least one bit.

It remains to analyze domain clauses and elementary tautologies. Assume C is a domain clause for $(u, j) \notin \mathrm{Dom}(\phi_G)$. If u is not in $\mathrm{Dom}(M)$, then C is also a domain clause of $\mathrm{Dom}(\phi_{G'})$. If u is in $\mathrm{Dom}(M)$, then $C \restriction \mu_M = 1$ since if j' is such that $\phi_G(u, j') = M(u)$ then the binary representations of j and j' differ in at least one bit. Finally, if C is an elementary tautology $(P_{u;b} \vee \neg P_{u;b})$ for $u \in U$ and $b \in [\ell]$, then if u is not in $\mathrm{Dom}(M)$ then u is in U', and if u is in $\mathrm{Dom}(M)$ then $C \restriction \mu_M = 1$. $\dashv$

The refutation $R'' \restriction \mu_M$ from Claim 3 can be turned into a refutation R''' of $\mathrm{BPHP}(G')$ by deriving the disjunctions of two domain clauses of $\mathrm{BPHP}(G')$ that appear in R'' in one strongly sound step from either of the two domain clauses. Thus R''' is a strongly sound semantic refutation of $\mathrm{BPHP}(G')$ (Lemma 3 (1)).

Since $\mathrm{Im}(M) \subseteq B$, Claim 2 (b) implies that G' is a $(q, 1)$-expander. By Lemma 23 there is a formula F' in R''' such that

$$h(F'; \mathcal{M}(G')) > \tfrac{1}{3}\lfloor qb/d_L \rfloor. \tag{8}$$

In particular F' cannot be a domain clause of $\mathrm{BPHP}(G')$ because otherwise $h(F'; \mathcal{M}(G')) \leq \ell$, and for sufficiently large n, this is smaller than the right-hand side in (8). By the definition of R''', it follows that $F' = F \restriction \mu_M$ for some F in R''. Now, by Remark 19, Lemma 20, and recalling $\mathcal{M}(M) = \mathcal{M}(G')$, we have

$$h(F_{\mathcal{M}(G)} \restriction \mu_M) \geq h(F_{\mathcal{M}(G)} \restriction \mu_M; \mathcal{M}(M)) \geq h(F \restriction \mu_M; \mathcal{M}(M)) = h(F'; \mathcal{M}(G')).$$

Hence, (8) contradicts Claim 2 (a).

5 Upper bound

In this section we prove Theorem 2 as outlined in the introduction. We assume some elementary familiarity with Buss' bounded arithmetic theories (cf. [13, 19]): given a relational language α, its relativized i-th level is denoted by $T^i_2(\alpha)$ and is given by Buss' theory BASIC and the induction scheme for $\Sigma^b_i(\alpha)$-formulas. In the following let P', P, and E be binary relation symbols, and let R be a unary relation symbol.

We define the first-order formula π^m_n in the language of $T^2_2(P')$ as the conjunction of the following bounded formulas with free first-order variables m and n:

$$\forall x < m\ \exists y < n\ P'(x, y),$$
$$\forall x < m\ \forall y < m\ \forall z < n\ (P'(x, z) \wedge P'(y, z) \rightarrow x = y).$$

Next, we define the first-order formula $\pi^{q,m}_n$ in the language of $T^2_2(E, R, P)$ as the conjunction of the following bounded formulas:

$$B^{q,m}_1 := \forall x < m\ \exists y < q\ E(x, y),$$
$$B^{q,m}_2 := \forall x < m\ \forall y < q\ \forall z < q\ (E(x, y) \wedge E(x, z) \rightarrow y = z),$$
$$B^{q,m}_3 := \forall x < m\ \forall y < m\ \forall z < q\ (E(x, z) \wedge E(y, z) \rightarrow x = y),$$

$$B_4^{q,m} := \forall x < m\, \forall y < q\, (E(x,y) \to R(y)),$$
$$B_5^{q,n} := \forall x < q\, (R(x) \to \exists y < n\, P(x,y)),$$
$$B_6^{q,n} := \forall x < q\, \forall y < q\, \forall z < n\, (R(x) \wedge R(y) \wedge P(x,z) \wedge P(y,z) \to x = y).$$

Note that $B_1^{q,m}$, $B_2^{q,m}$, $B_3^{q,m}$ and $B_4^{q,m}$ have m and q as free first-order variables, and that $B_5^{q,n}$ and $B_6^{q,n}$ have q and n as free first-order variables.

We give a Δ_1^b-interpretation of π_n^m in $\pi_n^{q,m}$ as follows. Set

$$\sigma(x,y) := \exists z < q\, (E(x,z) \wedge P(z,y)),$$
$$\gamma(x,y) := \forall z < q\, (E(x,z) \to P(z,y)).$$

These formulas have q, x and y as free first-order variables. It is straightforward to check that

$$\models \pi_n^{q,m} \to \pi_n^m[P'/\sigma], \tag{9}$$
$$\models B_1^{q,m} \wedge B_2^{q,m} \to \forall x < m\, \forall y < n\, (\sigma(x,y) \leftrightarrow \gamma(x,y)), \tag{10}$$

Here, $[P'/\sigma]$ indicates the substitution of atomic subformulas $P'(t,t')$ for terms t and t' by $\sigma(t,t')$. From [23] we have that $T_2^2(P') \models \neg\pi_n^{2n}$. This implies $T_2^2(P')[P'/\sigma] \models \neg\pi_n^{2n}[P'/\sigma]$. By (10), $T_2^2(E,R,P) \cup \{B_1^{q,m}, B_2^{q,m}\} \models T_2^2(P')[P'/\sigma]$, so $T_2^2(E,R,P) \models (B_1^{q,m} \wedge B_2^{q,m} \to \neg\pi_n^{2n}[P'/\sigma])$, and hence by (9)

$$T_2^2(E,R,P) \models \neg\pi_n^{n^2,2n}. \tag{11}$$

Now note that $\neg\pi_n^{n^2,2n}$ is a DNF_1-formula in the sense of Theorem 3.1 from [20]. This theorem and (11) implies that the Paris-Wilkie translation of $\pi_n^{n^2,2n}$ as a set of clauses has R(log)-refutations of size $2^{(\log n)^{O(1)}}$. It is straightforward to check that this translation produces our set $\mathrm{PHP}_n^{n^2,2n}$, the threshold clauses being the Paris-Wilkie translation of $B_1^{q,m}, \ldots, B_4^{q,m}$.

6 Discussion

6.1 Lower bounds for PHP_n^{2n}?

Besides the Δ_1^b-interpretation of π_n^m in $\pi_n^{q,m}$ from the previous section, there is an obvious quantifier-free interpretation in the reverse direction: define $\rho(y) := (y < m)$, $\epsilon(x,y) := (x = y)$, and $\tau(x,y) := P'(x,y)$ and check that

$$\models \pi_n^m \to \pi_n^{q,m}[R/\rho, E/\epsilon, P/\tau]. \tag{12}$$

Note also that all the formulas ρ, ϵ, τ, σ and γ that define the interpretations are Δ_0. So (9), (10) and (12) imply that $I\Delta_0$ proves the equivalence of $\pi_n^{n^2,2n}$ and π_n^{2n}. Similarly, one sees that PHP_n^{2n} has polynomial-size bounded-depth refutations if and only if so does $\mathrm{PHP}_m^{n^2,2n}$. We have a closer look.

Given a refutation of $\mathrm{PHP}_n^{n^2,2n}$ one can apply a suitable restriction to turn it into a refutation of PHP_n^{2n}. This transformation preserves depth and size. For the other direction, assume you have a polynomial size refutation of PHP_n^{2n}, say written in the variables $P'_{u;v}$. Substitute in every formula every positive $P'_{u;v}$ by the 2-CNF translating $\gamma(u,v)$ and every negative $\neg P'_{u;v}$ by the negation of the 2-DNF translating $\sigma(u,v)$. The result can be "filled up" to a polynomial size refutation of $\mathrm{PHP}_n^{n^2,2n}$. If the original PHP_n^{2n}-refutation used only DNFs, i.e. disjunctions of 1-CNFs, then the $\mathrm{PHP}_n^{n^2,2n}$-refutation obtained uses disjunctions of 2-CNFs. Hence a superpolynomial lower bound for such refutations would entail a superpolynomial lower bound for DNF-refutations of PHP_n^{2n}. This would constitute considerable progress on the main open question concerning the existence of polynomial-size bounded-depth refutations of PHP_n^{2n}.

6.2 Approximate counting and WPHPs

Fix a real $r \geq 1$. The problem of r-approximate counting asks, for a given string $x = x_0 \cdots x_{m-1} \in \{0,1\}^m$, to compute $\hat{w} \in [m+1]$ such that $\hat{w}/r \leq \sum_{i \in [m]} x_i \leq \hat{w} \cdot r$. We are looking for a family of polynomial-size bounded-depth circuits $C_1, C_2, \ldots$ solving this problem: C_m has m input gates and $m+1$ output gates and on input x precisley the $\hat{w}$-th output gate gets value 1.

This is a fundamental problem in computational complexity. In 1981, Furst, Saxe and Sipser [16] showed that the case $r = 1$ of exact counting does not admit a solution. Two years later, Stockmeyer [35] found for any $r > 1$ probabilistic circuits doing the job, and these could be derandomized by methods of Ajtai and Ben-Or [3]. Since it relies on the probabilistic method, this derandomization produces a non-uniform circuit family. A decade later, Ajtai [2] found a uniform (even FO- or DLOGTIME-uniform [7]) solution through an impressive though intricate construction. The problem has been repeatedly revisited until quite recently [36, 37].

As pointed out in the introduction, approximate counting seems sufficient to refute PHP_n^{2n}, or even $\mathrm{PHP}_n^{n^2,2n}$. We have a closer look. We focus on PHP_n^{2n} since the discussion for the relativized version is similar. Given a solution $C_1, C_2, \ldots$ to r-approximate counting let F_m^i be a bounded-depth formula expressing that the $(i+1)$-th output bit of C_m is 1. Recall that PHP_n^{2n} has $m = 2n^2$ variables $P_{u;v}$. On the one hand, PHP_n^{2n} implies that at least $2n$ many variables are true, namely for each $u \in [2n]$ at least one of $P_{u;v}$ for $v \in [n]$; thus, for $i_0 := \lceil 2n/r \rceil$,

$$\mathrm{PHP}_n^{2n} \models F_m^{i_0} \vee F_m^{i_0+1} \vee \ldots \vee F_m^m. \tag{13}$$

On the other hand, PHP_n^{2n} implies that at most n variables are true, namely for each $v \in [n]$ most one of $P_{u;v}$ for $u \in [2n]$; thus, for $i_1 := \lceil nr \rceil$,

$$\mathrm{PHP}_n^{2n} \models \neg F_m^{i_1} \wedge \neg F_m^{i_1+1} \wedge \ldots \wedge \neg F_m^m. \tag{14}$$

Specifically, for $r = \sqrt{2}$ we have $i_0 = i_1$ and the two formulas become negations of each other. To get short bounded-depth refutations of PHP_n^{2n} it would thus be sufficient to find a solution to the $\sqrt{2}$-approximate counting problem such that the implications (13) and (14)

have short bounded-depth derivations from PHP_n^{2n}. This amounts to verify a weak form of *correctness* of the solution.

Acknowledgements The first and third authors would like to thank the CICYT for its support through projects TIN2010-20967-C04-04 (TASSAT) and TIN2007-66523 (FORMALISM) respectively. The second author would like to thank the FWF (Austrian Science Fund) for its support through Project P 24654 N25.

References

[1] M. Ajtai. The complexity of the pigeonhole principle. Proceedings of the 29th Annual Symposium on the Foundations of Computer Science (FOCS), 346-355, 1988.

[2] M. Ajtai. Approximate counting with uniform constant-depth circuits. In Advances in computational complexity theory, DIMACS Series in Discrete Mathematics and Theoretical Computer Science 13:1-20, 1993.

[3] M. Ajtai and M. Ben-Or. A Theorem on Probabilistic Constant Depth Computations. In Proceedings of the 16th Annual ACM Symposium on Theory of Computing (STOC), 471-474, 1984.

[4] A. Atserias. Improved Bounds on the Weak Pigeonhole Principle and Infinitely Many Primes from Weaker Axioms. Theoretical Computer Science 295(1-3): 27-39, 2003.

[5] A. Atserias. On sufficient conditions for unsatisfiability of random formulas. Journal of the ACM 51(2): 281-311, 2004.

[6] A. Atserias, M.L. Bonet and J.L. Esteban. Lower bounds for the weak pigeonhole principle and random formulas beyond resolution, Information and Computation 176(2):136-152, 2002.

[7] D. A. M. Barrington, N. Immerman and H. Straubing. On Uniformity within NC. Journal of Computer and System Sciences. 41(3): 274-306, 1990.

[8] P. Beame, R. Impagliazzo and T. Pitassi. Exponential lower bounds for the pigeonhole principle. Computational Complexity 3(2):97-140, 1993.

[9] E. Ben-Sasson and N. Galesi. Space complexity of random formulae in resolution. Random Structures and Algorithms 23(1):92-109, 2003.

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

[11] B. Bollobás. Random Graphs. 2nd edition, Cambridge University Press, 2001.

[12] S. R. Buss. Polynomial size proofs of the propositional pigeonhole principle. Journal of Symbolic Logic 52(4):916-927, 1987.

[13] S. R. Buss. First-Order Proof Theory of Arithmetic. Handbook of Proof Theory, S. R. Buss (ed.), pp. 79-147, Elsevier, 1998.

[14] S. A. Cook, R. A. Reckhow. The Relative Efficiency of Propositional Proof Systems. Journal of Symbolic Logic 44(1):36-50, 1979.

[15] S. Dantchev and S. Riis. On relativisation and complexity gap for resolution-based proof systems, Proceedings of 17th Annual Conference of the European Association for Computer Science Logic (CSL), Lecture Notes in Computer Science 2803:142-154, Springer, 2003.

[16] M. L. Furst, J. B. Saxe and M. Sipser. Parity, Circuits, and the Polynomial-Time Hierarchy. In Proceedings of 22nd Annual Symposium on Foundations of Computer Science (FOCS), 260-270, 1981.

[17] M. L. Furst, J. B. Saxe and M. Sipser. Parity, circuits, and the polynomial-time hierarchy. Theory of Computing Systems 17(1):13-27, 1984.

[18] A. Haken. The intractability of resolution, Theoretical Computer Science 39(2-3):297-308, 1985.

[19] J. Krajíček. Bounded Arithmetic, Propositional Logic, and Complexity Theory. Encyclopedia of Mathematics and its Applications 60, Cambridge University Press, 1995.

[20] J. Krajíček. On the weak pigeonhole principle. Fundamenta Mathematicae, Vol.170(1-3):123-140, 2001.

[21] J. Krajíček. Combinatorics of first order structures and propositional proof systems. Archive for Mathematical Logic 43(4):427-441, 2004.

[22] J. Krajíček, P. Pudlák and A. Woods. An exponential lower bound to the size of bounded depth Frege proofs of the pigeonhole principle. Random Structures & Algorithms 7(1):15-39, 1995.

[23] A. Maciel, T. Pitassi and A. R. Woods. A new proof of the weak pigeonhole principle. Journal of Computer and System Sciences 64(4):843-872, 2002.

[24] J.B. Paris and A.J. Wilkie. Counting problems in bounded arithmetic. Methods in Mathematical Logic, LNM 1130: 317-340, Springer, 1985.

[25] J.B. Paris, A.J. Wilkie and A.R. Woods. Provability of the pigeonhole principle and the existence of infinitely many primes. Journal of Symbolic Logic 53(4):1235-1244, 1988.

[26] P. Pudlák. A bottom-up approach to foundations of mathematics. Proceedings Gödel'96, Logical Foundations of Mathematics, Computer Science and Physics – Kurt Gödel's Legacy, P. Hajek (ed.), Lecture Notes in Logic 6:81-97, Springer, 1996.

[27] P. Pudlák. Proofs as games. American Mathematical Monthly, pp. 541-550, 2000.

[28] R. Raz. Resolution lower bounds for the weak pigeonhole principle. Journal of the ACM 51(2): 115-138, 2004.

[29] A. A. Razborov. Pseudorandom Generators Hard for k-DNF Resolution and Polynomial Calculus. Unpublished, 2003.

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

[31] S. Riis. A complexity gap for tree-resolution. Computational Complexity 10:179-209, 2001.

[32] H. Robbins. A remark on Stirling's formula. The American Mathematical Monthly 62(1):26-29, 1955.

[33] N. Segerlind. The complexity of propositional proofs. The Bulletin of Symbolic Logic 13(4):417-481, 2007.

[34] N. Segerlind, S. R. Buss and R. Impagliazzo. A switching lemma for small restrictions and lower bounds for k-DNF resolution. SIAM Journal on Computing 33(5): 1171-1200, 2004.

[35] L. J. Stockmeyer. The Complexity of Approximate Counting (Preliminary Version). In Proceedings of the 15th Annual ACM Symposium on Theory of Computing (STOC), 118-126, 1983.

[36] E. Viola. On Approximate Majority and Probabilistic Time. Computational Complexity 18(3), 337-375, 2009.

[37] E. Viola. Randomness Buys Depth for Approximate Counting. In Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS), 230-239, 2011.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de