ELSEVIER Annals of Pure and Applied Logic 80 (1996) 195–228

ANNALS OF PURE AND APPLIED LOGIC

An exponential separation between the parity principle and the pigeonhole principle

Paul Beame [a,*], Toniann Pitassi [b,1]

[a] *Computer Science and Engineering, University of Washington, Box 352350, Seattle, WA 98195, USA*
[b] *Department of Mathematics, Thackeray Hall, University of Pittsburgh, Pittsburgh, PA 15260, USA*

Received July 1993; revised July 1995
Communicated by Y. Gurevich

Abstract

The combinatorial parity principle states that there is no perfect matching on an odd number of vertices. This principle generalizes the pigeonhole principle, which states that for a fixed bipartition of the vertices, there is no perfect matching between them. Therefore, it follows from recent lower bounds for the pigeonhole principle that the parity principle requires exponential-size bounded-depth Frege proofs. Ajtai (1990) previously showed that the parity principle does not have polynomial-size bounded-depth Frege proofs even with the pigeonhole principle as an axiom schema. His proof utilizes nonstandard model theory and is nonconstructive. We improve Ajtai's lower bound from barely superpolynomial to exponential and eliminate the nonstandard model theory.

Our lower bound is also related to the inherent complexity of particular search classes (see Papadimitriou, 1991). In particular, oracle separations between the complexity classes *PPA* and *PPAD*, and between *PPA* and *PPP* also follow from our techniques (Beame et al., 1995).

1. Introduction

A fundamental question in propositional proof theory is: how strong is a particular proof system. In particular, how large does a proof of a particular tautology have to be, as a function of the size of the tautology? It is believed that for any conceivable proof system, there exist tautologies (of size n) with no proofs of size polynomial in n, for sufficiently large n. However, proving this for every conceivable proof system is equivalent to proving that NP $\neq$ coNP, a fundamental question

* Corresponding author. E-mail: beame@cs.washington.edu. Research supported by NSF grants CCR-8858799 and CCR-8907960.

[1] Research supported by an NSF postdoctoral fellowship.

in complexity theory. Within the last twenty years, much research has been aimed at proving the existence of tautologies with no polynomial-size proofs for specific, natural classes of proof systems. The first unrestricted lower bound was an exponential lower bound for Resolution proofs of the propositional pigeonhole principle [10]. More recently, it has been shown that the propositional pigeonhole principle requires exponential-size, bounded-depth Frege proofs [4, 12, 15]. This is a major improvement over the Resolution lower bound, since Resolution can be viewed as a depth-1 Frege system.

The most outstanding open problem in this area is to extend these lower bounds to Frege systems and Extended Frege systems. Buss [8] has shown that the propositional pigeonhole principle has polynomial-size Frege proofs, and thus it is even a challenge to find a combinatorial principle requiring large Frege proofs. In a polynomial-size Frege proof, one has the ability to reason algebraically using counting arguments. Thus, intuitively, in order to prove lower bounds for Frege systems, one must come up with a tautology that cannot be proven through a simple counting argument, and moreover, develop techniques to prove this. As a step in this direction, one can add limited counting abilities (in the form of axiom schemas) to a bounded-depth Frege system, and first try to prove lower bounds here. The propositional pigeonhole principle is perhaps the most basic counting axiom, as it asserts that the cardinality n is not equal to the cardinality $n+1$. In this paper, we study what new theorems can and cannot be proven in polynomial-size, and bounded-depth, when we allow the pigeonhole principle as an axiom schema.

The propositional pigeonhole principle can be expressed by a family of propositional formulas, $\{PHP_m : m \geqslant 0\}$, where PHP_m asserts that there is no 1–1 mapping from a set D_0 of size $m+1$ to a set D_1 of size m. A related, but more general principle is the *parity principle*, PAR_n, which states that no graph on $2n+1$ nodes consists of a perfect matching. We encode PAR_n using $\binom{2n+1}{2}$ *matching* variables, $P_{\{i,j\}}$, $i,j \leqslant 2n+1$. Using these variables, PAR_n can be written as the disjunction of the following *matching clauses*:

$$\bigwedge\{\neg P_{\{i,j\}} : j \leqslant 2n+1,\ j \neq i\},\ i \leqslant 2n+1;$$

$$\bigwedge\{P_{\{i,k\}}, P_{\{j,k\}}\}, i \neq j \neq k,\ i,j,\ k \leqslant 2n+1.$$

It is not too hard to see that if there are short, bounded-depth Frege proofs of PAR_n, then there are also short, bounded-depth Frege proofs of the onto version of the pigeonhole principle. Expressed propositionally, the onto version of PHP has additional terms which imply that the function is also surjective. The recent exponential size lower bound for bounded-depth Frege proofs of the onto version of PHP [4] thus also establishes an exponential size lower bound for bounded-depth Frege proofs of PAR_n. (A proof of Urquhart [17] shows that PAR_n, like PHP_n, does have a polynomial size Frege proof of logarithmic depth.)

This suggests the following question: is the parity principle strictly stronger than the pigeonhole principle? Ajtai [2] was the first to show that, in a precise sense, the parity

principle is stronger than the pigeonhole principle. One can generalize the pigeonhole principle by allowing each variable in the *PHP* formula to represent an arbitrary *formula* over some underlying set of propositional variables. Now consider a bounded-depth Frege proof system, with underlying *matching* variables $P_{\{i,j\}}$, where the system is strengthened by allowing all bounded-depth instances of the *PHP* as axioms. In [2], Ajtai showed that PAR_n does not have polynomial-sized, bounded-depth Frege proofs, even in this stronger system.

The structure of Ajtai's argument extends the proof technique in his superpolynomial lower bound for the *PHP* [1]. He first sketches a restriction lemma giving small 'covering sets' for formulas over the matching variables. Then, in the novel part of the paper, he shows how the restrictions cannot have falsified the *PHP* axioms. This is done by first showing that all the information about a given pigeon or hole in a *PHP* axiom can be determined by the values of the matching variables touching a small covering set and then showing that it is not possible to have the information about pigeons or holes locally appear to describe a 1–1 function and yet be globally consistent. This last piece forms the bulk of the paper and uses a somewhat involved counting argument.

In this paper, we present a new proof and we improve the lower bound from superpolynomial to exponential. This result uses the proof-theoretic methods from [4, 6, 12] and a modification of the switching lemma from [4, 15]. The most difficult new part of this proof is showing that each restricted *PHP* axiom is converted to an approximation of a true formula after the various conversions are made. The structure of this argument is similar to Ajtai's: As in [2] we use a bit-wise encoding of the *PHP* formulas to obtain small descriptions of what happens to each pigeon or hole. In our case, rather than small covering sets we use small height matching decision trees along the lines of [4, 15]. This difference is fortuitous because it turns out that this permits a much simpler counting argument to show that it is impossible for the converted subformulas of the *PHP* axiom to locally describe a 1–1 function and be globally consistent.

The lemma which shows this latter result, is of independent interest. In particular, in [2a], we use it to demonstrate oracle separations between certain complexity classes of search problems: between classes *PPA* and *PPAD* and between classes *PPA* and *PPP*. These complexity classes, which characterize the complexity of many interesting problems, were defined by Papadimitriou [13] and lie between the function versions of P and NP.

Lastly, our result is related to questions about the ability to count in various systems of bounded arithmetic. $S_2(R)$ is the relativized system of bounded arithmetic, as defined by Buss [7], and $PHP(R)$ is the pigeonhole principle for the relation R. It follows from our result that $S_2(R) + PHP(R)$ cannot prove the parity principle. For connections between bounded-depth Frege systems and bounded arithmetic, see [16, 14].

We note that, independent of this work, Soren Riis (private communication) has shown similar results using methods of nonstandard model theory.

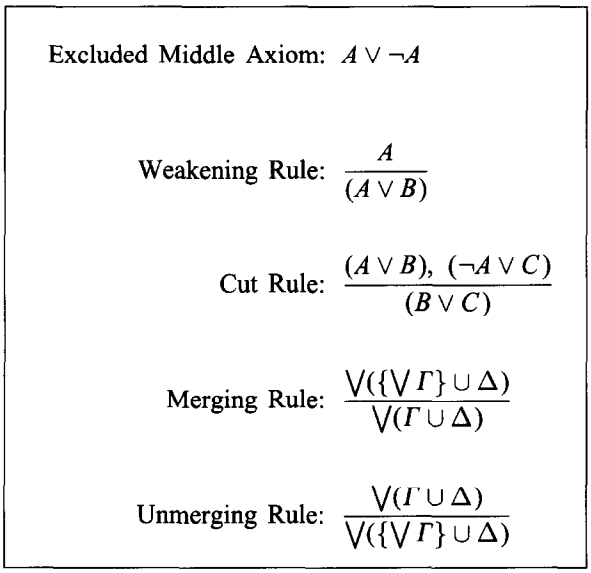

Fig. 1. Rules of the system H.

2. The proof system $H + PHP_b(F)$

The proof system we use is based on the Frege proof system H from [6, 15, 4] augmented with an axiom schema $PHP_b(F)$ which we describe below. H is a Frege system with NOT gates and unbounded fan-in OR gates (see Fig. 1). Excluded middle is the only axiom of H, and the cut rule is the main rule of inference, with extra merging and unmerging rules to manipulate the unbounded fan-in OR's. (By methods of Cook and Reckhow [9] the exact choice of the constant depth Frege system over unbounded fan-in $\wedge$, $\vee$, and $\neg$ we use is not crucial.)

A *proof* of a formula f in H is a sequence of formulas such that the final formula is f, and all intermediate formulas are either instances of the axiom, or follow from previous formulas by an inference rule. The *size* of a formula is the number of occurrences of $\vee$ and $\neg$ in the formula; the size of a Frege proof is the sum of the sizes of the formulas occurring as lines in the proof. Each formula as described above can be written as an unbounded fan-in boolean tree. The *depth* of a formula is the depth of the boolean tree which represents the formula. The depth of a Frege proof is the maximum depth of the formulas in the proof.

We now describe the axiom schema $PHP_b(F)$. Let $F = \{F(i,j) \mid i \leqslant m+1, j \leqslant m\}$ be a set of bounded-depth formulas over the propositional variables $P_{\{i,j\}}$, $i,j \leqslant 2n+1$. The natural form of the pigeonhole principle using $\vee$ and $\neg$ based on F, $PHP(F)$, is an OR of the following subformulas: $C1(F,x) = \neg(F(x,1) \vee F(x,2) \vee \cdots \vee F(x,m))$ for each $x \leqslant m+1$, which expresses the fact that x is not mapped to any hole; $C2(F,x_1,x_2,y) = \neg(\neg F(x_1,y) \vee \neg F(x_2,y))$ for each $y \in m$ and $x_1,x_2 \leqslant m+1$ with $x_1 \neq x_2$, which expresses the fact that hole y has pigeons x_1 and x_2 mapped to it; and

finally $C3(F,x,y_1,y_2) = \neg(\neg F(x,y_1) \vee \neg F(x,y_2))$ for each $x \leqslant m+1$ and $y_1, y_2 \leqslant m$ with $y_1 \neq y_2$ which expresses the fact that pigeon x is mapped to holes y_1 and y_2:

Unfortunately $PHP(F)$ is not a convenient form of the pigeonhole principle for our purposes. For $x \leqslant m+1$ let x_i denote the ith bit of x in binary notation. For each $F(x,y)$, we can express F in "left bitwise" notation by the formula $F_L(x,y)$:

$$F_L(x,y) = \neg \bigvee_{i=1}^{\lceil \log(m+1) \rceil} \neg F^L_{i,y_i}(x) = \bigwedge_{i=1}^{\lceil \log(m+1) \rceil} F^L_{i,y_i}(x),$$

where

$$F^L_{i,b}(x) = \bigvee_{z \leqslant m,\ z_i = b} F(x,z).$$

Similarly, we can also express F in the "right bitwise" notation by the formula $F_R(x,y)$:

$$F_R(x,y) = \neg \bigvee_{i=1}^{\lceil \log m \rceil} \neg F^R_{i,x_i}(y) = \bigwedge_{i=1}^{\lceil \log m \rceil} F^R_{i,x_i}(y),$$

where

$$F^R_{i,b}(y) = \bigvee_{z \leqslant m+1,\ z_i = b} F(z,y).$$

Let $PHP_b(F)$ denote the pigeonhole principle expressed as an OR of the following subformulas:

$$\{ C1_b(F,x) \mid x \leqslant m+1 \}$$

$$\cup \{ C2(F_R, x_1, x_2, y) \mid x_1, x_2 \leqslant m+1,\ x_1 \neq x_2,\ y \leqslant m \}$$

$$\cup \{ C3(F_L, x, y_1, y_2) \mid y_1, y_2 \leqslant m,\ y_1 \neq y_2,\ x \leqslant m+1 \}\},$$

where

$$C1_b(F,x) = \neg \bigvee \{ F^L_{i,b}(x) \mid 1 \leqslant i \leqslant \lceil \log m \rceil, b = 0,1 \},$$

$$C2(F_R, x_1, x_2, y) = \neg(\neg F_R(x_1,y) \vee \neg F_R(x_2,y)),$$

$$C3(F_L, x, y_1, y_2) = \neg(\neg F_L(x,y_1) \vee \neg F_L(x,y_2))$$

The proof of the following lemma is straightforward.

Lemma 1. *Let $F = \{F(x,y) \mid x \leqslant m+1, y \leqslant m\}$, where each $F(x,y)$ is a bounded-depth formula over the matching variables. Then there exists a bounded-depth, polynomial-size Frege proof of $PHP(F)$ from $PHP_b(F)$.*

Let $F_b(x,*)$ denote the set of formulas $\{F_L(x,y) \mid y \leqslant m\}$. Similarly, let $F_b(*,x)$ denote the set of formulas $\{F_R(y,x) \mid y \leqslant m+1\}$. Note that for each x, all formulas

in $F_b(x,*) \cup F_b(*,x)$ are boolean formula over the $O(\log m)$ subformulas, $\{F^L_{i,b}(x) | i \leqslant \lceil \log m \rceil,\ b \in (0,1)\}$, and $\{F^R_{i,b}(x) \mid i \leqslant \lceil \log(m+1) \rceil,\ b \in (0,1)\}$.

3. Restrictions, decision trees and the switching lemma

The overall idea of the proof is to apply a *restriction* (or partial 0–1 setting) to the underlying variables in the proof in such a way that: (1) after applying the restriction we are left with a (sub)proof of the same principle, over the subdomain of unset variables; and (2) the proof that remains is greatly simplified. In order for (1) to hold, we must choose a restriction that defines a partial matching between some of the vertices. The definitions for partial matching restrictions that we use are entirely analogous to those in [4, 15] but we include them for definiteness. The *variables over* D are $\{P_{\{i,j\}} : i \neq j \in D\}$. The i and j will be called the *endpoints* of $P_{\{i,j\}}$. For convenience we will write both P_{ij} and P_{ji} to represent variable $P_{\{i,j\}}$. A *map over* D is defined to be a conjunction of the form $\bigwedge \Gamma$, where Γ is a set of variables over D such that distinct variables in Γ have distinct endpoints. Maps describe partial matchings on the set D. The *size* of a map $\bigwedge \Gamma$ is $|\Gamma|$. An OR of maps is called a *map disjunction*. The *mapsize* of a map disjunction is the size of the largest map in the disjunction; if all the maps are of size at most t, then it will be called a *t-disjunction*. A map σ' *extends* map σ if $\sigma = \bigwedge \Gamma$ and $\sigma' = \bigwedge \Gamma'$ such that $\Gamma \subseteq \Gamma'$. We say that two maps σ and τ are *compatible* if there is some map π that extends both σ and τ and we denote the smallest such map π by $\sigma\tau$. A *truth assignment* φ over D is any total assignment of $\{0,1\}$ to the variables over D. Let $D' \subseteq D$. A truth assignment φ over D is a *matching on* D' if for all $i \in D'$ there is a unique $j \in D$ such that $P_{ij} = 1$.

If Y is a map or a set of variables, then $v(Y)$ denotes the set of endpoints of variables in Y.

We will now define a probability space of partial matchings on D, where $|D| = 2n + 1$ The probability space $\mathcal{M}^D_p$ is the set of all pairs $\rho = \langle \pi, \pi_* \rangle$ where π is a random matching of n edges in D and π_* is a random subset of the edges of π where each edge of π_* is chosen independently at random with probability p.

Every $\rho = \langle \pi, \pi_* \rangle$ in $\mathcal{M}^D_p$ determines a unique *restriction*, r, of the variables over D as follows.

$$r(P_{ij}) = \begin{cases} 1 \text{ if } \{i,j\} \in \pi \setminus \pi_*, \\ 0 \text{ if there is a } k \text{ such that } \{i,k\} \\ \quad \text{or } \{j,k\} \in \pi \setminus \pi_*, \\ * \text{ otherwise.} \end{cases}$$

In this way, the distribution $\mathcal{M}^D_p$ defines a probability distribution of restrictions. If r is a random restriction obtained by choosing a random ρ according to $\mathcal{M}^D_p$, we will refer

to both the restriction and the random partial matching by ρ. For a Boolean formula F and an element $\rho \in \mathcal{M}_p^D$, F restricted by ρ will be denoted by $F\restriction_\rho$.

3.1. Matching decision trees

In this subsection, we define a combinatorial structure called a matching decision tree. A matching-decision tree can be thought of as a simple method for describing a function on truth assignments that are almost total matchings. Our eventual goal is to approximate each formula in the original proof by a small-depth decision tree.

A *matching decision tree* over domain D is defined as follows. It is a rooted tree where each interior node v is labelled by a query $i \in D$ and each edge is labelled by some pair $\{i,j\}$ where $j \neq i \in D$. Leaves are labelled with either "0" or "1". For each interior node v labelled by $i \in D$, there is exactly one out-edge labelled $\{i,j\}$ for each $j \in D \setminus \{i\}$ that does not appear in any edge label on the path from the root to v. The label of an interior node v may not appear in any edge label on the path from the root to v. Thus the set of edge labels on any path defines a map. A matching decision tree where all of the leaves are labelled "1" will be called a *1-tree*. A matching decision tree T' *extends* a matching decision tree T if, for any root to leaf path p' in T', there is a unique path p in T such that the map σ' defined by p' extends the map σ defined by p. (Note that the leaf labels are not required to be related in this definition.)

A matching decision tree T over D *represents* a function f over domain D if for all leaf nodes $v \in T$, if we let σ be the map defined by the path in T from the root to v then for all truth assignments α over D that are matchings on $v(\sigma)$ and satisfy σ, $f(\alpha)$ is equal to the label of v. For a boolean function f over domain D, we define $d_D(f)$ to be the minimum height of all matching decision trees computing f. (Note that the empty matching decision tree has only one node, labelled by either "1" or "0", and represents the function "true" or "false", respectively.)

Let T be a matching decision tree. In the remainder of this paper, the function represented by T is defined to be the map-disjunction, $maps(T)$, consisting of the labels of all of the paths in T that end in leaves labelled 1. Note that if T has height t, then the function computed by T is a t-disjunction. Furthermore note that for any partial matching restriction ρ over D, $maps(T\restriction_\rho) = maps(T)\restriction_\rho$.

Extending this definition, if f is a tree with intermediate nodes labelled by *OR* and *NOT* gates, and leaf nodes labelled by matching decision trees, then the function computed by f is obtained by iteratively computing the functions evaluated by the subtrees of f.

If ρ is a partial matching restriction over D and T is a matching decision tree over D, then define $T\restriction_\rho$ to be the decision tree obtained from T by removing all paths which have a label that has been set to "0" by ρ, and contracting all edges whose labels are set to "1" by ρ.

Lemma 2. *Let f be a boolean function over D and let T be a matching decision tree representing f over D. If ρ is a partial matching restriction over D, then $T\restriction_\rho$ is a matching decision tree for $f\restriction_\rho$ over $D\restriction_\rho$.*

Note that if T represents f over D then the tree T^c obtained by switching the 1's and 0's labelling the leaves of T represents $\neg f$. The lemma in the next section actually is a switching lemma in the spirit of [11] because it will allow us to obtain a map disjunction that approximates the negation of f by representing f by a matching decision tree T and then taking $maps(T^c)$.

Where it is convenient, we shall assume that an ordering is given on D. Whenever we write a real number where an integer is required, we mean the integer part of the real number (floor). If f is a map disjunction defined over a set D and ρ is a restriction on D then we will use the notation $\delta(f\restriction_\rho)$ for $d_{D\restriction_\rho}(f\restriction_\rho)$. We now state the main combinatorial lemma.

Lemma 3 (Switching Lemma). *Let f be an r-disjunction over $D \subseteq D^n$, where D^n is a set of size $2n+1$. Choose ρ at random from $\mathcal{M}_p^D$. If $s \geqslant 0$ and $pn \geqslant (r+s)(2r+2s+1)$ then*

$$Pr[\delta(f\restriction_\rho) \geqslant s] < \alpha^s,$$

where $\alpha > 0$ satisfies $(1 + 225p^4n^3/\alpha^2)^r \leqslant 2$.

The inequality $(1 + 225p^4n^3/\alpha^2)^r \leqslant 2$ holds when $\alpha = 19p^2n^{3/2}r^{1/2}$. This can be seen by taking the natural logarithm of both sides and the applying the inequality $\ln(1+x) \leqslant x$.

The proof of the Switching Lemma is given in Section 6.

4. Exponential lower bounds

The overall structure of the exponential lower bound argument is very similar to the argument in [12] and in [4]. Given an alleged proof, P, of depth d and size S, a series of d restrictions are applied and after each one the proof is converted using a switching lemma to reduce the depth, until we end up with a sequence of formulas, each of which can be represented by matching decision trees. We will show that if P had size no greater than S, then after the d conversions, each matching decision tree is a 1-tree. But on the other hand, the final formula in the proof is the converted PAR formula, which becomes a matching decision tree which is not a 1-tree, and hence we have reached the contradiction.

The new part of the argument is showing that each instance of a PHP axiom schema gets converted into a 1-tree. This is the subject of Section 5.

There are some formal and technical differences from [4, 15] in how we apply the depth reduction in our argument. First of all, for convenience, we maintain a

proof with small height matching decision trees at the leaves rather than formulas and finish reduction when each formula is a small height matching decision tree. More importantly, in order to preserve the formulas in the bitwise version of the $PHP_b(F)$ axiom schema we do not always apply the switching lemma to $\vee$'s of decision trees. If the $\vee$ has fan-in at most $\log S$ then we simply "stack" the decision trees one on top of the other in the natural way creating a new deeper decision tree that evaluates all of the $\log S$ trees along each branch.

Using the switching lemma one can easily maintain via induction that after i levels of conversions have been applied the height of the decision trees at the leaves of the formulas in the proof is at most $2\log^i S$. The domain, D, of the matching variables declines by a fixed fractional power at each step.

In this section we will prove the following theorem.

Theorem 4. *Any proof of PAR_n in $H+PHP_b(F)$ of depth d must have size S at least* $\exp[n^{\Omega(1/d4^d)}]$; *more precisely,* $S \geqslant \exp[n^{2/(5d4^d)}/3]$.

Corollary 5. *Any proof of PAR_n in $H+PHP_b(F)$ of polynomial size must have depth* $\Omega(\log\log n)$.

4.1. The conversion process

The conversion proceeds in rounds where each round reduces the depth of the formulas by 1. In each round, subformulas lying just above the leaves are converted into decision trees. A certain set of these are converted using the switching lemma, others are converted by simpler means. Based on the set of subformulas for which the switching lemma is to be applied, a restriction σ is chosen to keep the heights of all the resulting decision trees small. Then the conversions themselves are done using the method previously decided upon for each subformula. A given subformula will, in general, appear several times throughout the proof. Each time it appears, the same conversion is applied.

More formally, after σ is applied, if f is $\neg T$ for some decision tree T then the conversion of f, $\mathscr{C}[f] = T^c$ and if f is $\bigvee_{i=1}^{q} T_i$ then

(a) if $q > \log S$, then $\mathscr{C}[f]$ is the result of applying the switching lemma argument to $\bigvee_{i=1}^{q} maps(T_i)$ as described below, and

(b) if $q \leqslant \log S$, then $\mathscr{C}[f]$ is obtained by stacking the decision trees T_i such that a leaf at the end of path p is labelled "1" if p forces some T_i to 1. (One stacks decision trees T_1 and T_2 by replacing each leaf of T_1 by a copy of T_2, deleting incompatible paths, contracting redundant queries, and labelling the leaves of this copy of T_2 by the OR of the original leaf value and the value of the leaf of T_1 that this copy of T_2 replaced.)

Note that if f is $\neg T$ or f is $\bigvee_{i=1}^{q} T_i$ for $q \leqslant \log S$, i.e. the stacking method (b) is used to produce $\mathscr{C}[f]$, then the application of any restriction ρ commutes with the

conversion process on f, i.e. $\mathscr{C}[f]\restriction_\rho = \mathscr{C}[f\restriction_\rho]$, although this is not true in general if the switching lemma is used.

4.2. *Proof of Theorem 4*

Let P be an alleged proof of PAR_n over D, $|D| = 2n+1$, of size at most S, and depth d (in $H + PHP_b(F)$). We will first show that there exists a sequence of good restrictions which allows us to convert the formulas in P into small-depth decision trees. Recall that each formula in P consists of d levels of *OR*s and *NOT*s, followed by the bottom level, which are depth-1 decision trees.

Let $t_0 = 2$, and $t_i = t_0 \log^i S$ for $i > 0$. Define $\lambda(n) = n^{1/4}/(8\log^{d/4} S)$, and $p_i = \lambda(n_i)/n_i$. If $\lambda^{(i)}$ is the i-fold composition of λ with itself, then it can be shown that $\lambda^{(i)}(n) \geqslant n^{4^{-i}}/(16\log^{d/3} S)$.

We will show that in the conversion process each formula is converted into a decision tree of height t_d which will be much smaller than the size of the universe which is $2\lambda^{(d)}(n)+1$. We will argue that this is impossible and thus we must have $\log S \geqslant (n^{4^{-d}}/144)^{2/(5d)} \geqslant n^{2/5d4^d}/3$ since $d \geqslant 2$, i.e. $S \geqslant \exp(n^{\Omega(4^{-d}/d)})$ as required.

Lemma 6. *If* $\log S < (n^{4^{-d}}/144)^{2/(5d)}$ *then after* d *applications of the conversion process above to depth* d *proof* P *of size* S, *all the formulas in the proof are converted to matching decision trees over a domain of size* $2\lambda^{(d)}(n)+1$, *where each tree has depth at most* $t_d = 2\log^d S \ll \lambda^{(d)}(n)$.

Proof. We will let $D^0 = D$, $P^0 = P$ and define $\rho^1, \ldots, \rho^d$ as a sequence of restrictions such that for all $1 \leqslant k \leqslant d$, ρ^k leaves all variables over D^k unset, $|D^k| = 2n_k + 1$. We will use $P^1, \ldots, P^d$ to denote the sequence of proofs generated, where P^k will be P^{k-1} converted by ρ^k. We will show that for all $i < d$, if each formula in P^i has depth $d-i$, with decision trees of height t_i at its leaves, and total size S, then P^i converted by ρ^{i+1} yields a new sequence of formulas, P^{i+1}, of depth $d-(i+1)$, decision trees of height t_{i+1} at its leaves and total size S.

Since $\log S < (n^{4^{-d}}/144)^{2/(5d)}$, it follows that $n^{4^{-d}} > 144\log^{5d/2} S$. Therefore,

$$\begin{aligned}
\lambda^{(d)}(n) &\geqslant n^{4^{-d}}/(16\log^{d/3} S) \\
&> \frac{9\log^{5d/2} S}{\log^{d/3} S} \\
&= 9\log^{13d/6} S \\
&\gg 9\log^{2d} S \\
&> (t_{d-1} + t_d)(2t_{d-1} + 2t_d + 1) > t_d
\end{aligned}$$

for n sufficiently large.

Let D^i be the domain of the formulas in P^i and let $i < d$. We will argue by induction that the leaves of P^i are decision trees of height at most t_i and that P^i is defined on variables over $2n_i + 1$ vertices where $n_i \geq \lambda^{(i)}(n)$. When $i = 0$ the claim is clearly true.

Suppose it is true for some $i < d$. We can apply the Matching Switching Lemma for ρ^{i+1} drawn at random from $\mathscr{M}_{p_i}^{D^i}$ to each distinct map disjunction representing a formula at one level above the leaves in P^i since the maps in each leaf decision tree are of size at most t_i and

$$\begin{aligned} p_i n_i = \lambda(n_i) &\geq \lambda^{(i+1)}(n) \geq \lambda^{(d)}(n) \\ &\geq (t_{d-1} + t_d)(2t_{d-1} + 2t_d + 1) \\ &\geq (t_i + t_{i+1})(2t_i + 2t_{i+1} + 1). \end{aligned}$$

For each map disjunction, f, corresponding to one of these formulas in P^i, for a randomly chosen $\rho \in \mathscr{M}_{p_i}^{D^i}$, the probability that $f\restriction_\rho$ cannot be represented by a matching decision tree over D^{i+1} of depth at most t_{i+1} is most $\alpha^{t_{i+1}}$, where $0 < \alpha < 19 p_i^2 n_i^{3/2} t_i^{1/2}$. Since $p_i = \lambda(n_i)/n_i$ and $t_i = 2\log^i S < 2\log^d S$,

$$\alpha < \frac{19 p_i^2 n_i^2 t_i^{1/2}}{n_i^{1/2}} = \frac{19\lambda(n_i)^2 t_i^{1/2}}{n_i^{1/2}} < \frac{19 n_i^{1/2}\sqrt{2}\log^{d/2} S}{(64\log^{d/2} S) n_i^{1/2}} = \frac{19\sqrt{2}}{64} < \frac{1}{2}.$$

Because the size of P^i is at most S, there are at most S map disjunctions in P^i, and therefore, for a randomly chosen ρ, the probability that every formula of P^i at a level one above the leaves cannot be represented by a depth-t_{i+1} matching decision tree over D^{i+1} is at most $S\alpha^{t_{i+1}} \leq S\alpha^{2\log S} < 1/S < 1/6$ for n sufficiently large.

The expected number of edges in the random matching defining ρ that are starred is $n_i p_i = \lambda(n_i)$. Since the number of stars is binomially distributed, for sufficiently large n_0, a random ρ leaves at least the expected number of stars with probability greater than 1/3. (See, for example, Lemma 4.1 of [3]). In this case the number of vertices in the resulting domain D is at least $2\lambda(n_i) + 1$. Thus, there exists a restriction, ρ^{i+1}, leaving $2n_{i+1} + 1$ domain vertices, $n_{i+1} \geq \lambda(n_i) \geq \lambda^{(i)}(n)$, such that each decision tree in D^{i+1} has depth at most t_{i+1}.

By induction the claim is true for d and we obtain a sequence $\mathscr{D}$ of decision trees over a smaller domain of size $2\lambda^{(d)}(n) + 1$, where each tree has depth at most $t_d = 2\log^d S \ll \lambda^{(d)}(n)$. □

From now on we assume that S satisfies the conditions of Lemma 6. Let F be a subformula in proof P of PAR_n. Define T_F to be the decision tree $\mathscr{C}[F]\restriction_\rho$ where $\rho = \rho_1\rho_2\cdots\rho_d$ as defined in the proof of Lemma 6. Note that $\mathscr{D}$ consists of the sequence $\{T_{F_i}\}$ where $\{F_i\}$ is the sequence of formulas in proof P.

Lemma 7. *Let F be a subformula of one of the formulas in proof P. Then*

(a) *if F is $\neg G$ then T_F is T_G^c,*

(b) *if F is $\bigvee_{i=1}^q G_i$ for $q \leq \log S$ then T_F may be obtained by stacking $T_{G_1}, \ldots, T_{G_q}$, and*

(c) *if* F *is* $\bigvee_{i=1}^{q} G_i$ *for* $q > \log S$ *then* T_F *represents* $\bigvee_{i=1}^{q} maps(T_{G_i})$ *over* D^d.

Proof. Suppose that F is converted in round k of the conversion process. By the rules of the conversion, each of these hold for the decision trees corresponding to F and the G's at the time subformula F is converted. Let $\sigma = \rho_{k+1} \cdots \rho_d$. The first two cases follow since in these cases conversion commutes with application of restriction σ. Since $maps(T\restriction_\sigma) = maps(T)\restriction_\sigma$, the third case follows as well. □

To prove Theorem 4 we will need the following theorem, which will be proven in the subsequent section.

Theorem 8 (PHP Axiom Soundness). *Let* $PHP_b(F)$ *be an instance of the* PHP_b *axiom schema in proof* P. *Then* $T_{PHP_b(F)}$ *is a* 1*-tree.*

The remainder of the proof follows by first showing that under our assumption about S, each decision tree in $\mathscr{D}$ is a 1-tree. Then we will derive a contradiction by showing that the PAR_n formula cannot be converted into a 1-tree by the conversion process.

Lemma 9. *Let* $\mathscr{D}$ *be the result of applying the conversion process to proof* P. *Every decision tree in* $\mathscr{D}$ *is a* 1*-tree.*

Proof. The proof proceeds by induction on the sequence of trees in $\mathscr{D}$, or equivalently on the sequence of formulas in P. Now every formula in P is either an instantiation of an axiom or follows from previous formulas via some inference rule.

The only axioms are instances of the PHP axiom schema or the excluded middle axiom. By Theorem 8, any instance of the PHP axiom schema is converted into a 1-tree. Suppose it is an instance of the excluded middle axiom, say $A \vee \neg A$. By Lemma 7, the decision tree representing $\neg A$ is the tree T_A, but with the opposite leaf labelling. Therefore, the decision tree for $A \vee \neg A$ is a 1-tree.

There are four different rules of inference to deal with. The more difficult cases are those involving unbounded fan-in OR gates – i.e. the merging and unmerging rules. We will first give the proof when the inference is an application of the cut rule, and then when the inference is an application of unmerging. The other rules are analogous.

Suppose that the inference is the cut-rule, and let A be the formula $X \vee Y$, let B be the formula $\neg X \vee Z$, and let C be the formula $Y \vee Z$. We want to show that if T_A and T_B are 1-trees, then so is T_C. By Lemma 7, T_A is obtained by stacking the decision trees T_X and T_Y. Similarly, T_B is obtained by stacking the decision trees T_X^c and T_Z, and T_C is obtained by stacking the decision trees T_Y and T_Z. Suppose, for sake of contradiction, that a path, π, of T_C has leaf label 0. Thus, there are compatible subpaths π_Y in T_Y and π_Z in T_Z that both have leaf labels 0. Since both T_C and T_X have height much smaller than the universe size there is some path ρ in T_X (and thus also in T_X^c) compatible with π. By construction of T_A, $\sigma = \rho\pi_Y$ labels some path in T_A and by construction of T_B, $\tau = \rho\pi_Z$ labels some path in T_B. Now, in either T_X or

T_X^c, the path ρ has leaf label 0. Thus either σ in T_A or τ in T_B has leaf label 0, a contradiction.

Intuitively, the above argument holds because, in the case of the cut rule, the OR gate and the negations involved in the inference are not approximated and therefore, since both antecedent formulas are 1, the derived formula should also be 1. Now consider the unmerging rule. Let A be the formula $\bigvee\{X1,\ldots,Xn,Y1,\ldots,Ym\}$, and let B be the formula $X \vee Y$ where $X = \bigvee\{X1,\ldots,Xn\}$, $Y = \bigvee\{Y1,\ldots,Ym\}$, and B follows from A by the unmerging rule. Assume that T_A is a 1-tree. By Lemma 7 T_A represents $A' = \bigvee\{maps(T_{X1}),\ldots,maps(T_{Xn}),maps(T_{Y1}),\ldots,maps(T_{Ym})\}$, T_X represents $\bigvee\{maps(T_{X1}),\ldots,maps(T_{Xn})\}$, T_Y represents $\bigvee\{maps(T_{Y1}),\ldots,maps(T_{Ym})\}$, and T_B is obtained by stacking T_X and T_Y.

Fix a path, π labelling T_B. We want to show that π has leaf label 1. Since both T_A and T_B are decision trees of height much smaller than the universe size, there is some path ρ in T_A compatible with π. By assumption, the path ρ in T_A has leaf label 1. Since T_A represents A', it follows that $A'(\rho) = 1$. Therefore,

$$\bigvee\{maps(T_{X1}),\ldots,maps(T_{Xn}),maps(T_{Y1}),\ldots,maps(T_{Ym})\}(\rho) = 1.$$

Thus, either $\bigvee\{maps(T_{X1}),\ldots,maps(T_{Xn})\}(\rho) = 1$ or $\bigvee\{maps(T_{Y1}),\ldots,maps(T_{Ym})\}(\rho) = 1$, say the former. Since ρ is compatible with the path π_X in T_X that is a subpath of π, and T_X represents $\bigvee\{maps(T_{X1}),\ldots,maps(T_{Xn})\}$, this path π_X of T_X also has leaf label 1. Finally, because T_B is obtained by stacking the decision trees T_X and T_Y, the path of T_B labelled by π must have leaf label 1. The cases of the other rules are similar and the claim follows by induction on the number of steps in the proof. □

We now obtain a contradiction by showing that PAR_n converts into a 0-tree.

Lemma 10. *The result of applying the conversion process for proof P to PAR_n, T_{PAR_n}, is a 0-tree.*

Proof. By Lemma 7, T_{PAR_n} represents the OR of the maps in the decision trees representing its clauses:

$$T_{\neg\bigvee\{P_{ij}:\ j\leqslant 2n+1,\ j\neq i\}} \quad \text{for } i\leqslant 2n+1;$$
$$T_{\neg(\neg P_{ik}\vee\neg P_{jk})} \quad \text{for } i\neq j\neq k,\ i,j,k\leqslant 2n+1.$$

Consider any path in T_{PAR_n} and let π be the map labeling this path. We will show that its leaf label must be 0.

Suppose instead that its leaf label is 1. Therefore, we know that there is some leaf labelled 1 in some decision tree respresenting a clause of PAR_n that must be reached by π. Now it is easy to see that any $T_{\neg(\neg P_{ik}\vee\neg P_{jk})}$ must be a 0-tree since no partial matching is consistent with matching both i and j to k and so the satisfied map cannot be from one of these clauses. Therefore, the leaf labelled 1 is in some decision tree of the form $T_{\neg\bigvee\{P_{ij}:\ j\leqslant 2n+1,j\neq i\}}$. In this case, π must not match any

element to i. However, since π leaves at least one other element j of D^d unmatched there is a truth assignment α that extends π in which i is matched to j and the clause is falsified by α. This contradicts the fact that π reaches a leaf labelled 1 in the converted clause.

Thus, the final formula PAR_n converts into a 0-tree. □

Theorem 4 follows immediately.

5. The soundness of PHP_b

In this section we will prove Theorem 8. Since the decision trees produced are all of small height, in order to show that the tree produced by converting a pigeonhole axiom is a 1-tree, it suffices to show that it is impossible to force this tree to 0 by a small partial matching restriction.

Proof of Theorem 8. Consider an instance of the PHP_b axiom schema, $PHP_b(F)$, in the original proof, P, such that $PHP_b(F)$ is the pigeonhole principle on $m(m+1)$ subformulas $F(x, y)$. We wish to show that $T = T_{PHP_b(F)}$ is a 1-tree.

Suppose that T is not a 1-tree and thus it has a leaf labelled 0. Since the height of T is at most t_d there is some map σ of size $\leqslant t_d$ such that $T\restriction_\sigma = 0$.

For each subformula G of $PHP_b(F)$, we will use the notation G' to denote $T_G\restriction_\sigma$.

By definition, $T\restriction_\sigma = 0$ if and only if $PHP_b(F)'$ is identically 0. The argument that the latter holds is based on the properties of the decision trees obtained for the various subformulas of $PHP_b(F)$. The easy cases are when some single decision tree expresses the fact that either the function is undefined on one of the $m+1$ points or that the function is not 1–1. The difficult case is when these decision trees do not obviously contradict the pigeonhole principle. That is, each one appears to define a part of a one-to-one function from $m+1$ to m.

If there were some partial matching, α, that extended some path in every tree, then it is easy to see that in this case $PHP_b(F)'$ would not be identically 0 and we would be finished. Unfortunately, this ideal situation may not hold because a particular partial matching may not extend any path in a given decision tree. For example, if the root node of a tree queries i, then all matchings where i is unmatched will not extend any path of the tree. However, since we have required that the matching decision trees are not too deep, we will still be able to show that $PHP_b(F)'$ cannot be identically 0.

Theorem 11. *$PHP_b(F)'$ is not identically 0.*

Let the size of the universe remaining after σ is applied be $2n'+1$ and call the resulting domain D'. By construction, the tree $T_{PHP_b(F)}$ represents the $\bigvee$ of the maps in the various trees $T_{C1_b(F,x)}$, $T_{C2(F_R,x_1,x_2,y)}$ and $T_{C3(F_L,x,y_1,y_2)}$ over D^d. We can apply the restriction σ to all of these trees and conclude that $PHP_b(F)'$ represents the $\vee$ of the

maps in the various $C1_b(F,x)'$, $C2(F_R,x_1,x_2,y)'$, and $C3(F_L,x,y_1,y_2)'$ over D'. Thus, it suffices to show that at least one of these trees has a branch with leaf label 1.

The various subformulas $C1_b$, $C2$, and $C3$ are defined in terms of subformulas $F_L(x,y)$ and $F_R(x,y)$ and thus on $F^L_{i,b}(x)$ and $F^R_{i,b}(y)$ for $x \leqslant m+1$ and $y \leqslant m$ which are in turn based on the subformulas $F(x,y)$. The largest OR used in these definitions is of size $\lceil \log(m+1) \rceil$. Since $\lceil \log(m+1) \rceil \leqslant \log S$, by Lemma 7, trees T_{C1_b}, T_{C2}, and T_{C3} are obtained by stacking the bitwise trees $T_{F^L_{i,b}(x)}$ and $T_{F^R_{i,b}(y)}$ and setting their leaf labels appropriately. The same property holds for the way trees $C1'_b$, $C2'$, and $C3'$ may be obtained from $F'^L_{i,b}(x)$ and $F'^R_{i,b}(y)$ for $x \leqslant m+1$ and $y \leqslant m$, since restrictions commute with stacking.

By Lemma 7, $F_L(x,y)$ becomes the tree, $F'_L(x,y)$, obtained by stacking the $\lceil \log m \rceil$ trees, $F'^L_{i,y_j}(x)$. Similarly $F'_R(x,y)$ is obtained by stacking the $\lceil \log(m+1) \rceil$ trees $F'^R_{i,x_j}(y)$. For each x, $1 \leqslant x \leqslant m+1$ we can define $L_x = \{F'^L_{i,b}(x) \mid i \leqslant \lceil \log m \rceil, b \in \{0,1\}\}$, and similarly each tree $F'_R(y,x)$, $y \leqslant m$ is the stacking of $\lceil \log(m+1) \rceil$ trees from the set of $2\lceil \log(m+1) \rceil$ trees $R_x = \{F'^R_{i,b}(Y) \mid i \leqslant \lceil \log(m+1) \rceil, b \in \{0,1\}\}$.

Therefore, for each x, the tree obtained by stacking the trees in $L_x \cup R_x$ is an extension of all of the trees $F'_L(x,*)$, and $F'_R(*,x)$. We define $\mathcal{T}_x$ to be this single matching tree over D', $|D'| = 2n'+1$, which simultaneously extends all of the trees $F'_L(x,*)$ and $F'_R(*,x)$, with the further modification that all the root-leaf paths are extended to some fixed length $k \ll n'$. This is accomplished by adding queries of other matching variables to any paths that are too short. Note that $\mathcal{T}_x$ still extends all of the trees in $F'_R(*,x)$ and $F'_L(x,*)$. The leaves of $\mathcal{T}_x$ are labelled with pairs $\{x \to u_1, \ldots, x \to u_k, v_1 \to x, \ldots, v_l \to x\}$ where the pair $x \to u_i$ is a label of some path, p if and only if $F'_L(x,u_i)\restriction_p = 1$. Similarly, the pair $v_j \to x$ is a label of p if and only if $F'_R(v_j,x)\restriction_p = 1$. Note that since $F_R(*,m+1)$ is not defined, no leaf label of $\mathcal{T}_{m+1}$ will contain a pair $u \to m+1$, for any $u \leqslant m+1$. We now let $\mathcal{T} = \{\mathcal{T}_x \mid x \leqslant m+1\}$.

Definition. $\mathcal{T}$ is a *local function* if and only if: $\forall x \leqslant m+1$, $\forall$ paths p of $\mathcal{T}_x$, there exists some $z \leqslant m$ such that the leaf label of p contains the pair $x \to z$. In other words, if the map defined by p is ρ then there exists $z \leqslant m$ such that $F'_L(x,z)\restriction_\rho = 1$.

Definition. $\mathcal{T}$ is *locally* 1–1 if for all x and for all paths, p, in $\mathcal{T}_x$, the leaf associated with p has at most one label of the form $x \to z_1$ and at most one label of the form $z_2 \to x$.

Definition. $\mathcal{T}$ is *consistent* if for all $x, y \leqslant m+1$ and for all pairs of paths, p_x in $\mathcal{T}_x$ and p_y in $\mathcal{T}_y$, if the maps they define, ρ_x and ρ_y, respectively, are compatible then $x \to y$ labels the leaf of p_x if and only if $x \to y$ labels the leaf of p_y.

We first show that if either $\mathcal{T}$ is not locally 1–1 or not a local function then $PHP_b(F)'$ is not identically 0. Then we will argue that one of these cases must be true. We do this by showing, using the way that $\mathcal{T}$ is constructed, that if $\mathcal{T}$ is locally 1–1 then it is also consistent and then showing that it is impossible for $\mathcal{T}$ to be

consistent as well as both a local function and locally 1–1. This latter proof requires a combinatorial argument.

Lemma 12. *If $\mathcal{T}$ is not locally* 1–1, *then $PHP_b(F)'$ is not identically* 0.

Proof. Assume that $\mathcal{T}$ is not locally 1–1. Then there exists an $x \leqslant m+1$, and a path p in $\mathcal{T}_x$ such that leaf label associated with p contains either (1) $x \to z_1$ and $x \to z_2$, $z_1 \neq z_2$, or (2) $z_1 \to x$ and $z_2 \to x$, $z_1 \neq z_2$. Consider the first case. Let ρ be the map defined by p. Since $\mathcal{T}_x$ extends both $F'_L(x,z_1)$ and $F'_L(x,z_2)$,

$$F'_L(x,z_1)\restriction_\rho = F'_L(x,z_2)\restriction_\rho = 1$$

and thus $C2(F_L,x,z_1,z_2)'\restriction_\rho = 1$ which means that $PHP_b(F)'$ is not identically 0. The second case is handled similarly. □

Lemma 13. *If $\mathcal{T}$ is not a local function, then $PHP_b(F)'$ is not identically zero.*

Proof. If $\mathcal{T}$ is not a local function, then for some $x \leqslant m+1$, there exists a path, p, of $\mathcal{T}_x$, whose leaf label does not contain $x \to y$, for any $y \leqslant m$. Let ρ be the map defined by p. Since $\mathcal{T}_x$ extends all $F'^L_{i,b}(x)$, for every $i \leqslant \lceil \log m \rceil$ and $b = 0$ or 1 we have $F'^L_{i,b}(x)\restriction_\rho = 0$ and thus $C1_b(F,x)'\restriction_\rho = 1$ so $PHP_b(F)'$ is not identically 0. □

Lemma 14. *If $\mathcal{T}$ is locally* 1–1 *then $\mathcal{T}$ is consistent.*

Proof. We will prove the contrapositive. Suppose that $\mathcal{T}$ is not consistent. Then there exists $x, y \leqslant m+1$ and compatible maps, ρ_x labelling path p_x in $\mathcal{T}_x$, and ρ_y labelling path p_y in $\mathcal{T}_y$, such that either $F'_L(x,y)\restriction_{\rho_x} = 1$ and $F'_R(x,y)\restriction_{\rho_y} = 0$ or vice versa. We will assume that the former case occurs (in which case we also know that $y \leqslant m$); the latter case is completely analogous.

We now sketch the remainder of the argument. Since $F_L(x,y)$ and $F_R(x,y)$ are constructed from the bitwise versions of F, this inconsistency occurs exactly if x is mapped to at least two different z's in F, at least one of which agrees with y in each bit position (in effect the left bitwise version sees a phantom edge not really present in F.) Thus the underlying F is not 1–1 and this is easily translated upward to show that $\mathcal{T}$ is not locally 1–1. The formal argument follows.

Recall that

$$F'_L(x,y)\restriction_{\rho_x} = \mathscr{C}[\neg \bigvee_{i=1,\ldots,\lceil \log m \rceil} \neg F'^L_{i,y_i}(x)]\restriction_{\rho_x}.$$

Because of the rules for conversion, the switching lemma is not used in producing $F'_L(x,y)$ from the various $F'^L_{i,y_i}(x)$ and so $F'_L(x,y)\restriction_{\rho_x} = 1$ implies that for all $i \leqslant \lceil \log m \rceil$, $F'^L_{i,y_i}(x)\restriction_{\rho_x} = 1$. By similar reasoning, $F'_R(x,y)\restriction_{\rho_y} = 0$ implies that there exists a j such that $F'^R_{j,x_j}(y)\restriction_{\rho_y} = 0$.

Now $F'^{L}_{i,y_i}(x) = T_{F^{L}_{i,y_i}(x)}\restriction_{\sigma}$ and by Lemma 7 $T_{F^{L}_{i,y_i}(x)}$ represents

$$\bigvee_{z \leqslant m, z_i = y_i} maps(T_{F(x,z)})$$

over D^d and thus $F'^{L}_{i,y_i}(x)$ represents

$$A^{L}_{i,y_i} = \bigvee_{z \leqslant m, z_i = y_i} maps(F'(x,z))$$

over D'. Similarly, $F'^{R}_{j,x_j}(y)$ represents

$$\bigvee_{w \leqslant m+1, w_j = x_j} maps(F'(w,y))$$

over D'.

Since $F'^{R}_{j,x_j}(y)\restriction_{\rho_y} = 0$ this implies in particular that $F'(x,y)\restriction_{\rho_y} = 0$. Also, since for each i, $F'^{L}_{i,y_i}(x)\restriction_{\rho_x} = 1$ this representation implies that there must be $z^1,\ldots,z^{\lceil \log m \rceil} \leqslant m$ such that for each i, $z^i_i = y_i$ and $F'(x,z^i)\restriction_{\rho_x} = 1$. Now, because ρ_x and ρ_y are compatible, $F'(x,y)\restriction_{\rho_y} = 0$ implies that $F'(x,y)\restriction_{\rho_x} \neq 1$. Thus, for each i, $y \neq z^i$, and so there must be at least two different values $u \neq v$ among the z^i such that

$$F'(x,u)\restriction_{\rho_x} = F'(x,v)\restriction_{\rho_x} = 1.$$

We will now use this to show that $\mathscr{T}$ is not locally 1–1. Since $F'(x,v)\restriction_{\rho_x} = F'(x,u)\restriction_{\rho_x} = 1$, we have for each $i \leqslant \lceil \log m \rceil$,

$$A^{L}_{i,u_i}(x)\restriction_{\rho_x} = \left(\bigvee_{z \leqslant m, z_i = u_i} maps(F'(x,z)) \right)\restriction_{\rho_x} = 1$$

and

$$A^{L}_{i,v_i}(x)\restriction_{\rho_x} = \left(\bigvee_{z \leqslant m, z_i = v_i} maps(F'(x,z)) \right)\restriction_{\rho_x} = 1.$$

Since $\mathscr{T}_x$ extends every $F'^{L}_{i,b}(x)$, ρ_x fixes the value of every $F'^{L}_{i,b}(x)$, in particular of every $F'^{L}_{i,u_i}(x)$ and $F'^{L}_{i,v_i}(x)$. Because $F'^{L}_{i,u_i}(x)$ represents $A^{L}_{i,u_i}(x)$ over D', and $F'^{L}_{i,v_i}(x)$ represents $A^{L}_{i,v_i}(x)$ over D', for every i $F'^{L}_{i,u_i}(x)\restriction_{\rho_x} = A^{L}_{i,u_i}(x)\restriction_{\rho_x} = 1$ and $F'^{L}_{i,v_i}(x)\restriction_{\rho_x} = A^{L}_{i,v_i}(x)\restriction_{\rho_x} = 1$. Therefore, by construction, $F'_{L}(x,u)\restriction_{\rho_x} = 1$ and $F'_{L}(x,v)\restriction_{\rho_x} = 1$ and thus $x \to u$ and $x \to v$ both label the leaf followed by ρ_x in $\mathscr{T}_x$. Since $u \neq v$ this shows that $\mathscr{T}$ is not locally 1–1. □

Lemma 15. *Let* $\{\mathscr{T}_x \mid 1 \leqslant x \leqslant m+1\}$ *be matching decision trees, as described above. Then it is impossible for* $\mathscr{T}$ *to be at the same time a local function, locally* 1–1, *and consistent.*

Proof. Assume for sake of contradiction that $\mathscr{T}$ is locally 1–1, consistent, and a local function. By definition of $\mathscr{T}$, we also know that no leaf label of $\mathscr{T}_{m+1}$ contains $(x, m+1)$, for any x, $1 \leqslant x \leqslant m+1$. We will show that this leads to a contradiction.

Let U, V be maps over D of size exactly k. Let $\mathcal{T}_x$, $\mathcal{T}_y$ be complete matching decision trees over D. Then we have the following definitions.

(1) $r(U) = r_L(U) - r_R(U)$, where $r_L(U) = \#\{(x, y) \mid U$ labels a path in $\mathcal{T}_x$ mapping x to $y\}$, and $r_R(U) = \#\{(x, y) \mid U$ labels a path in $\mathcal{T}_y$ mapping x to $y\}$.

(2) $d(U, V) = \#\{(x, y) \mid U$ labels a branch in $\mathcal{T}_x$ mapping x to $y\}$ and V labels a branch in $\mathcal{T}_y$ mapping x to y and U is compatible with V.

(3) Let $a(N, k)$ be the number of leaves in a complete, matching decision tree of height k over D, $|D| = N$.

(4) Let $b(N, r, k)$ be the number of leaves in a complete matching decision tree of height k over D, $|D| = N$, that lie below a given node of height r.

We will write $r_L(U)$ as $\sum_{(x,y)} r_L(U, x, y)$, where $r_L(U, x, y) = 1$ if U labels a path in $\mathcal{T}_x$ with leaf value (x, y), and otherwise $r_L(U, x, y) = 0$. Analogously, $r_R(U) = \sum_{(x,y)} r_R(U, x, y)$, where $r_R(U, x, y) = 1$ if U labels a path in $\mathcal{T}_y$ with leaf value (x, y). Similarly, we will write $d(U, V)$ as $\sum_{x,y} d(U, V, x, y)$, where $d(U, V, x, y)$ is 1 if: U is compatible with V; U labels a path in $\mathcal{T}_x$ with leaf value (x, y); and V labels a path in $\mathcal{T}_y$ with leaf value (x, y).

Lemma 16. *Given the quantities defined above,*

(a) $a(N - 2k, k) \cdot r_L(U) = \sum_V d(U, V) \cdot b(N - 2k, |U \cap V|, k)$.

(b) $a(N - 2k, k) \cdot r_R(U) = \sum_V d(V, U) \cdot b(N - 2k, |U \cap V|, k)$.

Proof. We give the proof of part (a). The proof of part (b) is analogous. Rewriting the left- and right-hand sides, we want to show

$$\sum_{x,y} r_L(U, x, y) \cdot a(N - 2k, k) = \sum_{x,y} \sum_V d(U, V, x, y) \cdot b(N - 2k, |U \cap V|, k).$$

Fix U, x, y. Then we will show that $r_L(U, x, y) \cdot a(N - 2k, k) = \sum_V d(U, V, x, y) \cdot b(N - 2k, |U \cap V|, k)$. If $r_L(U, x, y) = 0$, then $d(U, V, x, y) = 0$ for all V, and therefore the above equality holds for these choices of U, x, y.

The other case is when $r_L(U, x, y) = 1$. Recall that U labels a path of $\mathcal{T}_x$ with leaf label $x \to y$ if and only if $r_L(U, x, y) = 1$. Let $\mathcal{T}' = \mathcal{T}_y \restriction_U$. We claim that the number of paths in $\mathcal{T}'$ equals $\sum_V d(U, V, x, y)$. To see that each path of $\mathcal{T}'$ contributes 1 to the quantity $\sum_V d(U, V, x, y)$, notice that if p is a path of $\mathcal{T}'$ labelled by V', then V' is compatible with U, and can be extended to a map, V, which labels a path of $\mathcal{T}_y$. Because the decision trees are consistent, since there is a path in $\mathcal{T}_x$ consistent with V and with leaf label $x \to y$, it must be the case that $x \to y$ is also a leaf label of the path labelled by V in $\mathcal{T}_y$, and therefore $d(U, V, x, y) = 1$. In the other direction, if $d(U, V, x, y) = 1$, then V is consistent with U, and V labels a path of $\mathcal{T}_y$ with leaf value $x \to y$, and therefore the restricted path, labelled by $V \restriction_U$, will be a path of $\mathcal{T}'$.

Let $\mathcal{T}''$ be the extension of $\mathcal{T}'$ to a complete, depth-k decision tree over D', $|D'| = N - 2k$. Then the number of branches in the new, extended tree is exactly

$\sum_V d(U,V,x,y)\cdot b(N-2k,|U\cap V|,k)$. Alternatively, the number of branches in $\mathscr{T}''$ is $a(N-2k,k)$, which is equal to $a(N-2k,k)\cdot r_L(U,x,y)$, and thus the lemma holds. □

We are now ready to complete the proof of Lemma 15. Recall that the decision trees $\mathscr{T}$ are over the universe, D' of size $2n'+1$. Let $N=2n'+1$. By the definition of $\mathscr{T}$, we know that for every U that labels a path in $\mathscr{T}_{m+1}$, there is no z such that the leaf label of U contains $z\to m+1$. Therefore, $r(U)>0$ for those U's that label paths in $\mathscr{T}_{m+1}$. Secondly, because we are assuming that $\mathscr{T}$ is both locally 1–1 and a local function, we have that $r(U)\geqslant 0$ for every U. Therefore, $\sum_U r(U)>0$, and thus $\sum_U a(N-2k,k)r(U)>0$ as well.

By Lemma 16, we have $\sum_U\sum_V b(N-2k,|U\cap V|,k)[d(U,V)-d(V,U)]>0$. However,

$$\sum_U\sum_V b(N-2k,|U\cap V|,k)d(U,V)=\sum_V\sum_U b(N-2k,|V\cap U|,k)d(U,V)$$
$$=\sum_U\sum_V b(N-2k,|U\cap V|,k)d(U,V).$$

The first equality follows by swapping the summations and using the commutativity of intersection, and the second equality follows by switching notations for U and V. But this contradicts the inequality above, and therefore the lemma holds. □

Proof of Theorem 11. By Lemmas 14 and 15, if $\mathscr{T}$ is locally 1–1 then it cannot also be a local function. Thus, $\mathscr{T}$ is either not locally 1–1 or not a local function and so, by Lemmas 13 and 12, $PHP_b(F)'$ is not identically 0. □

Theorem 8 now follows as an immediate corollary.

6. The switching lemma

In this section we will assume that D^n is a set with $|D^n|=2n+1$ and the underlying probability distribution will be $\mathscr{M}$ (as defined in Section 3). All other D will be subsets of D^n of odd cardinality.

Let $K\subseteq D$. Then $Proj_D[K]$ is the set of all minimal maps over D which involve all of the elements of K. A map $\sigma\in Proj_D[K]$ induces a restriction; we will refer to σ interchangeably as a restriction and as a map.

We define the *complete matching tree* for $K\subseteq D$ over D inductively as follows. If K consists of a single node $k\in D$, then label the root "k", and create $2n$ edges adjacent to the root, labelled by $\{k,j\}$, for all $j\in D\backslash\{k\}$. Otherwise, $K=K'\cup\{k\}\subseteq D$. Assume that we have created the complete tree for K'; we will now extend it to a complete tree for K. This is done by extending each leaf node v_ℓ as follows. Let p_ℓ be the path from the root to v_ℓ. The edge labellings along p_ℓ define a partial matching involving all elements of K'. If this partial map does not include k, then label v_ℓ by k, and add

new edges leading out of v_ℓ, one for every possible mapping for k that results in a map extending the partial matching along p_ℓ. Otherwise, if k is involved in the partial matching, leave v_ℓ unlabelled. Note that each path of the complete tree over K will be labelled by some $\sigma \in Proj_D[K]$.

For $X \subseteq D$, let $\rho(X) = *$ denote the condition that all vertices in X are unmatched. Also, let $\#(\rho) = k$ denote the condition that exactly k vertices are unmatched by ρ.

Lemma 17. *Let f be a boolean function over the variables P_{ij}, $i \neq j \in D$. For every $K \subseteq D$, there exists a restriction, $\sigma \in Proj_D[K]$ such that $d_D(f) \leqslant |\sigma| + d_{D\restriction_\sigma}(f\restriction_\sigma)$.*

Proof. The proof is very similar to that of Beame and Håstad [3]. Fix $K \subseteq D$. We start with the complete matching tree for K. As noted above the paths of this tree correspond exactly to elements of $Proj_D[K]$. Let v_σ be the leaf node corresponding to the path labelled by $\sigma \in Proj_D[K]$. For each σ, we replace the leaf node, v_σ, by a subtree that is a matching decision tree for $f\restriction_\sigma$ over $D\restriction_\sigma$. The resulting tree clearly represents f over D. The depth of the resulting tree for K is at most $max_\sigma\{|\sigma| + d_{D\restriction_\sigma}(f\restriction_\sigma)\}$. □

If f is a map disjunction defined over a set D and ρ is a restriction on D recall that we use the notation $\delta(f\restriction_\rho)$ for $d_{D\restriction_\rho}(f\restriction_\rho)$. We prove the following lemma.

Lemma 3 (Switching Lemma). *Let f be a r-disjunction over $D \subseteq D^n$. Choose ρ at random from $\mathcal{M}_p^D$. If $s \geqslant 0$ and $pn \geqslant (r+s)(2r+2s+1)$ then*

$$Pr[\delta(f\restriction_\rho) \geqslant s] < \alpha^s,$$

where $\alpha > 0$ satisfies $(1 + 225p^4n^3/\alpha^2)^r = 2$.

The proof of the switching lemma, like that of [11], proceeds by induction on the number of clauses in f. We work along the clauses one by one: if ρ falsifies a particular clause, then we are left with essentially the same problem as before; if ρ does not falsify the clause then, it is much more likely that ρ satisfies the clause (and thus ensures that the whole formula is set to true) than ρ leaves any variable in the clause unset.

There are significant complications however in dealing with partial matching restrictions as opposed to fully independent ones. These complications are similar to those that occurred in [4, 15] where the domain of inputs was bipartite graphs.

We obtain Lemma 3 from the somewhat stronger Lemma 23 by conditioning on some arbitrary function F being forced to 0 and on some arbitrary map Q that is disjoint from f being entirely unset. First we prove several technical lemmas.

Lemma 18. *Let D satisfy $|D| = 2n+1$; let $Q \subseteq D$ with $|Q| = q$ and let $x \in D \setminus Q$. If $p(2n-q-1) \geqslant q+1$ then for ρ chosen at random from $\mathcal{M}_p^D$,*

$$p \leqslant Pr[\rho(x) = * \mid \rho(Q) = *] \leqslant 2p.$$

Proof. Consider $\rho = (\pi, \pi_*) \in \mathcal{M}_p^D$. Let $V = i$ denote the event that exactly i elements of Q are matched outside of Q by the matching, π. Then the probability in question is a weighted average of the probabilities $Pr[\rho(x) = * \mid \rho(Q) = * \wedge V = i]$, for all i, $0 \leqslant i \leqslant q$. (Note that $i + q$ is odd if and only if the unmatched element of D falls in Q.) We will upper bound the above probability for a fixed value of i.

There are two cases based on the way that x is matched by π:

1. x is an endpoint of an edge $e \in \pi$ disjoint from Q. In this case, the probability that x is set to $*$ is the probability that e is set to $*$ which happens with probability p.

2. x is not an endpoint of any matching edge disjoint from Q. In this case, x is certainly set to $*$. However, we now consider the probability that this case occurs. If the unmatched point is in Q then the probability that x is among the i points outside Q that are matched with points in Q is $i/(2n - q + 1)$. If the unmatched point is not in Q then there is an additional probability of $1/(2n - q + 1)$ that x is the unmatched point for a total probability of $i + 1/(2n - q + 1) \leqslant q + 1/(2n - q + 1) \leqslant p$ by hypothesis.

The lower bound follows because the probability that x is set to $*$ in each case is at least p and the upper bound follows by summing the probabilities in the two cases. □

Lemma 19. *Let D satisfy $|D| = 2n + 1$; let Q and R be disjoint subsets of D with $|Q| = q$ and $|R| = r$. If $p(2n - q - r) \geqslant q + r$ then for ρ chosen at random from $\mathcal{M}_p^D$,*

$$p^r \leqslant Pr[\rho(R) = * \mid \rho(Q) = *] \leqslant (2p)^r.$$

Proof. Let $R = \{x_1, x_2, \ldots, x_r\}$. Then the probability that R is set to $*$, given that $\rho(Q) = *$ is equal to

$$Pr[\rho(x_1) = * \mid \rho(Q) = *] \times Pr[\rho(x_2) = * \mid \rho(Q \cup \{x_1\}) = *] \times \cdots \times Pr[\rho(x_r) = * \mid \rho(Q \cup \{x_1, \ldots, x_{r-1}\}) = *].$$

Because each term is of the form $Pr[\rho(x) = * \mid \rho(Q') = *]$, where $q' = |Q'| \leqslant q + r - 1$ satisfies $p(2n - q' - 1) \geqslant q' + 1$, by Lemma 18 each term is between p and $2p$, and therefore the whole quantity is between p^r and $(2p)^r$. □

Lemma 20. *Let $Q \subseteq D$ where $|D| = 2n + 1$ and $|Q| = q$, and let Y be a partial matching over D that is disjoint from Q and $|Y| = k$. If $p(2n - 2k - q) \geqslant 2k + q$, then for ρ chosen at random from $\mathcal{M}_p^D$*

$$Pr[\rho(Y) \neq 0 \mid \rho(Q) = *] \geqslant \left(\frac{1 - p}{2n}\right)^k.$$

Proof. Following along the lines of the previous lemma, the probability $Pr[\rho(Y) \neq 0 \mid \rho(Q) = *]$ can be written as: $Pr[\rho(e_1) \neq 0 \mid \rho(Q) = *] \cdots Pr[\rho(e_k) \neq 0 \mid \rho(Q) = * \wedge \rho(e_1 \cdots e_{k-1}) \neq 0]$ where $Y = e_1 e_2 \cdots e_k$. We will show that if $|Y| = r$ and q satisfy $p(2n - 2(r + 1) - q) \geqslant 2(r + 1) + q$, then for a given map of size one, i.e.

an edge e, disjoint from $v(Y) \cup Q$, $Pr[\rho(e) = 1 \mid \rho(Q) = * \wedge \rho(Y) = 1]$ is at least $(1-p)/2n$. Therefore, since we will only apply this with $r \leqslant k-1$, the probability $Pr[\rho(Y) = 1 \mid \rho(Q) = *]$ is at least $((1-p)/2n)^k$.

Let $W = i$ denote the event that exactly i elements of $v(Y) \cup Q$ are mapped outside of $v(Y) \cup Q$ by π. The probability $Pr[\rho(e) \neq 0 \mid \rho(Q) = * \wedge \rho(Y) \neq 0]$ is a weighted average of the probabilities $Pr[\rho(e) \neq 0 \mid \rho(Q) = * \wedge \rho(Y) \neq 0 \wedge W = i]$. for $0 \leqslant i \leqslant q + 2r$. Note that the conditioning implies that each of the i elements outside $v(Y) \cup Q$ matched with an element of $v(Y) \cup Q$ must be set to $*$ by ρ. Furthermore, if $i + q$ is even then the unmatched point in D lies outside $v(Y) \cup Q$ and is also always set to $*$. If the matching π includes e then certainly $\rho(e) \neq 0$. Since e is disjoint from $v(Y) \cup Q$ the matching π includes e if its smaller endpoint avoids the up to $i+1$ points outside of $v(Y) \cup Q$ mentioned above and π matches this endpoint to the other end of e. This probability is

$$\left(1 - \frac{i+1}{2n - 2r - q + 1}\right) \frac{1}{2n - 2r - q} \geqslant \frac{1-p}{2n}$$

since $i + 1 \leqslant 2r + q + 1 < 2k + q$. □

Lemma 21. *Let F be a boolean formula over D, $|D| = 2n + 1$ and let $R \subseteq D$. Then for all $k \leqslant n$, $Pr[F\restriction_\rho = 0 \mid \#(\rho) = 2(k-1) + 1 \wedge \rho(R) = *] \geqslant Pr[F\restriction_\rho = 0 \mid \#(\rho) = 2k + 1 \wedge \rho(R) = *]$.*

Proof. Note that the distribution of restrictions given $\rho(R) = *$ can be described as follows. Choose a category, i, $0 \leqslant i \leqslant |R|$ from some distribution. Then choose k according to the shifted binomial distribution, $B(n - |R| - i, p) + i$. Choose a random set, S', of size $2k + 1$ from $D \setminus R$, and let $S = S' \cup R$. Choose a random matching, π', on $D \setminus S$. Therefore, the distribution of restrictions given $\rho(R) = *$ and the extra condition that $\#(\rho) = 2k + 1$ can be described by: Choose a random set, S' of size $2k + 1 - |R|$ from $D \setminus R$. Let $S = S' \cup R$ and then choose a random matching, π' on $D \setminus S$.

Let A^k denote the subdistribution of restrictions given that $\rho(R) = *$ and $\#(\rho) = 2k + 1$, and let A^{k-1} denote the subdistribution of restrictions given that $\rho(R) = *$ and $\#(\rho) = 2(k-1) + 1$. We would like to show that the probability that $F\restriction_\rho = 0$ in A^k is no greater than the probability that $F\restriction_\rho = 0$ in A^{k-1}.

Let $\rho^{k-1} = \langle \pi^{k-1}, \pi^{k-1}_* \rangle \in A^{k-1}$, and let $\rho^k = \langle \pi^k, \pi^k_* \rangle \in A^k$. Then we say that ρ^{k-1} and ρ^k *correspond* if: there exists an $(x, y) \in \pi^k_*$, such that $\pi^k_* \cup \langle x, y \rangle = \pi^{k-1}_*$ and $\pi^k \cup \langle x, y \rangle = \pi^{k-1}$. Note that whenever $\rho^k \in A^k$ forces F to 0, so do all of the elements of A^{k-1} which correspond to ρ^k. This is true because for every ρ^{k-1} which corresponds to ρ^k, all underlying variables are the same except for a few variables which are set to * in ρ^k, and set to 0 or 1 in ρ^{k-1}; in other words, ρ^{k-1} is a further restriction of ρ^k. Now, because F is already forced to 0 by ρ^k, it must continue to be 0 as we set more variables. Thus, F is also forced to 0 by ρ^{k-1}.

Let C^k denote the elements of A^k which force F to 0. For each ρ^k in A^k, there are $(2k + 1)k$ elements in A^{k-1} which correspond to it, and conversely, for each $\rho^{k-1} \in$

A^{k-1}, there are $n+1-k-|R|$ elements of A^k which correspond to it. The probability that a random ρ^k over A^k forces F to 0 equals $|C^k|/|A^k|$; thus the probability that a random ρ^{k-1} over A^{k-1} forces F to 0 is at least $|C^k| \cdot (2k+1)k/(n+1-k-|R|)|A^{k-1}|$. Since $|A_\pi^{k-1}|$ is equal to $(2k+1)k|A^k|/n+1-k-|R|$, the probability that F is forced to 0 over A^{k-1} is greater than or equal to the probability that F is forced to 0 over A^k which is what we wanted to prove. □

Lemma 22. *Suppose that* $0 \leqslant \alpha_0 \leqslant \alpha_1 \leqslant \cdots \leqslant \alpha_n$, *and for all* $k \leqslant n$, $\sum_{j=k}^n a_j \leqslant \sum_{j=k}^n b_j$. *Then for all* $k \leqslant n$, $\sum_{j=k}^n \alpha_j a_j \leqslant \sum_{j=k}^n \alpha_j b_j$.

Lemma 23 (Stronger Switching Lemma). *Let* D *be an arbitrary set with* $|D| = 2n+1$, *and let* Q *be an arbitrary map over* D *with* $|Q| = q$. *Let* f *be an* r*-disjunction over* $D' = D \setminus \{v(Q)\}$ *and let* F *be an arbitrary function over* D. *Let* ρ *be a random restriction chosen according to* $\mathcal{M}_p^D$. *Then for* $s \geqslant 0$ *and* $pn \geqslant (r+s+q)(2r+2s+2q+1)$ *we have*

$$Pr[\delta(f\restriction_\rho) \geqslant s \mid F\restriction_\rho = 0 \wedge \rho(Q) = *] \leqslant \alpha^s,$$

where $\alpha > 0$ *satisfies* $(1 + 225p^4n^3/\alpha^2)^r = 2$.

Proof. The proof proceeds by induction on the total number of maps in f.

Base Case. There are no maps in f. In this case f is identically 0 and therefore f is represented by the tree consisting of the single node labelled 0. Hence, $\delta(f\restriction_\rho) = 0$ and the lemma holds.

Induction Step. Assume that the lemma holds for all map disjunctions with fewer maps than the map disjunction of f. We will write f as $f_1 \vee f_2 \vee \cdots$ where each f_i is a map of f. We will analyze the probability by considering separately the cases in which ρ does or does not force the map f_1 to be 0. The failure probability, the probability that $\delta(f\restriction_\rho) \geqslant s$, is an average of the failure probabilities of these two cases. Thus,

$$\begin{aligned} &Pr\,[\delta(f\restriction_\rho) \geqslant s \mid F\restriction_\rho = 0 \wedge \rho(Q) = *] \\ &\quad \leqslant \max(Pr[\delta(f\restriction_\rho) \geqslant s \mid F\restriction_\rho = 0 \ \wedge \rho(Q) = * \ \wedge f_1\restriction_\rho = 0], \\ &\qquad Pr[\delta(f\restriction_\rho) \geqslant s \mid F\restriction_\rho = 0 \ \wedge \rho(Q) = * \ \wedge f_1\restriction_\rho \neq 0]). \end{aligned}$$

The first term in the maximum is $Pr[\delta(f\restriction_\rho) \geqslant s \mid (F \vee f_1)\restriction_\rho = 0 \wedge \rho(Q) = *]$. Let f' be f with map f_1 removed; then $Pr[\delta(f\restriction_\rho) \geqslant s \mid (F \vee f_1)\restriction_\rho = 0 \wedge \rho(Q) = *] = Pr[\delta(f'\restriction_\rho) \geqslant s \mid (F \vee f_1)\restriction_\rho = 0 \wedge \rho(Q) = *]$. Because f' has one less map than f, this probability is no greater than α^s, by the inductive hypothesis.

Now we will estimate the second term in the maximum. Let T be the set of variables appearing in the first map, f_1. By hypothesis, $size(T) \leqslant r$. We will analyze the cases based on the subset Y of the variables in T to which ρ assigns $*$; we use the notation

$*(\rho_T) = Y$ to denote the event that the variables in T which are assigned $*$ by ρ_T are exactly those in Y. Then

$$Pr[\delta(f\restriction_\rho) \geqslant s \mid F\restriction_\rho = 0 \wedge \rho(Q) = * \wedge f_1\restriction_\rho \neq 0]$$
$$= \sum_{Y \subseteq T} Pr[\delta(f\restriction_\rho) \wedge *(\rho_T) = Y \mid F\restriction_\rho = 0 \wedge \rho(Q) = * \wedge f_1\restriction_\rho \neq 0].$$

Consider the case in which $Y = \emptyset$. In this case the value of f_1 is forced to 1 by ρ. It follows that f is forced to 1 and hence $\delta(f) = 0$ so the term corresponding to $Y = \emptyset$ has probability 0. The sum then becomes

$$\sum_{\substack{Y \subseteq T, \\ Y \neq \emptyset}} Pr[\delta(f\restriction_\rho) \geqslant s \wedge *(\rho_T) = Y \mid F\restriction_\rho = 0 \wedge \rho(Q) = * \wedge f_1\restriction_\rho \neq 0],$$

which is equal to

$$\sum_{\substack{Y \subseteq T, \\ Y \neq \emptyset}} Pr[\delta(f\restriction_\rho) \geqslant s \mid F\restriction_\rho = 0 \wedge \rho(Q) = * \wedge f_1\restriction_\rho \neq 0 \wedge *(\rho_T) = Y] \tag{1}$$

$$\times\ Pr[*(\rho_T) = Y \mid F\restriction_\rho = 0 \wedge \rho(Q) = * \wedge f_1\restriction_\rho \neq 0]. \tag{2}$$

We will first bound the latter term, (2), in each of these products. Given that $f_1\restriction_\rho \neq 0$, the probability that $*(\rho_T) = Y$ is equal to the probability that $\rho(Y) = * \wedge \rho(T \setminus Y) = 1$. Thus, term (2) is no greater than

$$Pr[\rho(Y) = * \wedge \rho(T \setminus Y) = 1 \mid F\restriction_\rho = 0 \wedge \rho(Q) = * \wedge f_1\restriction_\rho \neq 0]$$
$$\leqslant Pr[\rho(Y) = * \mid F\restriction_\rho = 0 \wedge \rho(Q) = * \wedge \rho(T \setminus Y) = 1 \wedge \rho(Y) \neq 0].$$

Let F' be $F \vee G$ where $G\restriction_\rho = 0$ if and only if ρ sets all variables in $T \setminus Y$ to 1; then the above probability is equal to $Pr[\rho(Y) = * \mid F'\restriction_\rho = 0 \wedge \rho(Q) = * \wedge \rho(Y) \neq 0]$.

Claim A. *Let* $|Y| = k$, $|Q| = q$. *When* $pn \geqslant (k+q)(2k+2q+1)$, $Pr[\rho(Y) = * \mid F\restriction_\rho = 0 \wedge \rho(Q) = * \wedge \rho(Y) \neq 0] \leqslant Pr[\rho(Y) = * \mid \rho(Q) = * \wedge \rho(Y) \neq 0]$.

The proof of this claim is somewhat involved so we postpone it until after the rest of the main line of the argument.

Since $|Y| \leqslant r$ and $pn \geqslant (r+s+q)(2r+2s+2q+1)$ we can apply Claim A to show that the term (2) is at most $Pr[\rho(Y) = * \mid \rho(Q) = * \wedge \rho(Y) \neq 0]$. Since the conditions on the parameters also imply that $p(2n-2r-2q) \geqslant 2r+2q \geqslant 2|Y|+2q$, by Lemma 19,

$$Pr[\rho(Y) = * \mid \rho(Q) = *] \leqslant (2p)^{2|Y|}.$$

Also, by Lemma 20,

$$Pr[\rho(Y) \neq 0 \mid \rho(Q) = *] \geqslant \left(\frac{1-p}{2n}\right)^{|Y|}.$$

Therefore,

$$Pr[\rho(Y) = * \mid \rho(Q) = * \wedge \rho(Y) \neq 0] \leqslant \left[\frac{4p^2(2n)}{1-p}\right]^{|Y|} \leqslant (9p^2n)^{|Y|}.$$

Now we look at the first term, (1), in each product. Suppose that $2|Y| \leqslant s$. For each fixed Y, we will analyze the probability above by applying Lemma 17 with $K = v(Y)$ and $D = D\restriction_\rho$. By this lemma, if $\delta(f\restriction_\rho) \geqslant s$ then there is some $\sigma \in Proj_{D\restriction_\rho}[v(Y)]$, such that $d_{(D\restriction_\rho)\restriction_\sigma}((f\restriction_\rho)\restriction_\sigma) \geqslant s - |\sigma|$. To use this requires that we consider all maps in $Proj_{D\restriction_\rho}[v(Y)]$. One difficulty is that $D\restriction_\rho$ is itself a random variable dependent on ρ. We handle this by considering all maps σ in $Proj_D[v(Y)]$ and including them only if $\rho(\sigma) = *$. For notational convenience, let $P(D, Y) = Proj_D[v(Y)]$. When $\rho(\sigma) = *$, $(f\restriction_\rho)\restriction_\sigma = (f\restriction_\sigma)\restriction_\rho$ and applying the definition of $\delta(f\restriction_\rho)$, the above probability is no greater than

$$\begin{aligned}
&\sum_{\sigma \in P(D,Y)} Pr[\delta((f\restriction_\sigma)\restriction_\rho) \geqslant s - |\sigma| \wedge \rho(\sigma) \\
&\quad = * \mid F\restriction_\rho = 0 \wedge \rho(Q) = * \wedge f_1\restriction_\rho \neq 0 \wedge *(\rho_T) = Y] \\
&\quad \leqslant \sum_{\sigma \in P(D,Y)} Pr[\delta((f\restriction_\sigma)\restriction_\rho) \geqslant s - |\sigma| \mid F\restriction_\rho = 0 \wedge \rho(Q) \\
&\quad = * \wedge f_1\restriction_\rho \neq 0 \wedge *(\rho_T) = Y \wedge \rho(\sigma) = *] \\
&\qquad \times Pr[\rho(\sigma) = * \mid F\restriction_\rho = 0 \wedge \rho(Q) = * \wedge f_1\restriction_\rho \neq 0 \wedge *(\rho_T) = Y] \\
&\quad = \sum_{\sigma \in P(D,Y)} Pr[\delta((f\restriction_\sigma)\restriction_\rho) \geqslant s - 2|Y| \mid F\restriction_\rho = 0 \wedge \rho(Q) = * \wedge \rho(T \setminus Y) \\
&\quad = 1 \wedge \rho(\sigma) = *] \\
&\qquad \times Pr[\rho(\sigma) = * \mid F\restriction_\rho = 0 \wedge \rho(Q) = * \wedge \rho(T \setminus Y) = 1 \wedge \rho(Y) = *].
\end{aligned}$$

The last inequality above holds because $|\sigma| \leqslant 2|Y|$, the events $f_1\restriction_\rho \neq 0 \wedge *(\rho_T) = Y$ are equivalent to the events $\rho(Y) = * \wedge \rho(T \setminus Y) = 1$, and the condition $\rho(Y) = *$ is implied by $\rho(\sigma) = *$. Recall that if Y is a map, $v(Y) \subseteq D$ denotes the set of underlying vertices which are contained in the map. We will split up the map σ into two maps, σ_1 and σ_2, where a variable, $P_{ij} \in \sigma$ is in σ_1 if both $i \in v(Y)$ and $j \in v(Y)$. Otherwise, $P_{ij} \in \sigma_2$. Note that for every $\sigma \in Proj_D[v(Y)]$, $0 \leqslant |\sigma_1| \leqslant |Y|$. We further divide the above probability into sums according to the size of σ_1 to get

$$\begin{aligned}
&\sum_{i=0}^{|Y|} \sum_{\substack{\sigma \in P(D,Y), \\ |\sigma_1| = |Y| - i}} Pr[\delta((f\restriction_\sigma)\restriction_\rho) \geqslant s - 2|Y| \mid F\restriction_\rho = 0 \wedge \rho(Q) = * \wedge \rho(T \setminus Y) \\
&\quad = 1 \wedge \rho(\sigma) = *] \qquad (3) \\
&\qquad \times Pr[\rho(\sigma) = * \mid F\restriction_\rho = 0 \wedge \rho(Q) = * \wedge \rho(T \setminus Y) = 1 \wedge \rho(Y) = *]. \qquad (4)
\end{aligned}$$

For a fixed value of Y and $\sigma \in P(D, Y)$, we estimate the first term. Let f' be f with f_1 removed and consider the different possibilities for σ. Let f' be f with the variables in $T \setminus Y$ set to 1. Let F' be $F \vee G$ where $G\restriction_\rho = 0$ if and only if ρ sets all variables in $T \setminus Y$ to 1. Then the first term is equal to

$$Pr[\delta((f'\restriction_\sigma)\restriction_\rho) \geqslant s - 2|Y| \mid F'\restriction_\rho = 0 \wedge \rho(Q) = * \wedge \rho(\sigma) = *].$$

Now, if $\sigma = Y$ then f_1' is satisfied by σ and $f'\restriction_\sigma$ is the constant 1 and this probability is $0 \leqslant \alpha^{s-2|Y|}$. Otherwise, $\sigma \neq Y$, the map f_1' is falsified by σ, so $f'\restriction_\sigma$ has one fewer map than the original f that we started with. If we let Q' be the map which is the conjunction of Q and σ, then we can rewrite the term above as

$$Pr[\delta((f'\restriction_\sigma)\restriction_\rho) \geqslant s - 2|Y| \mid F'\restriction_\rho = 0 \wedge \rho(Q') = *].$$

Furthermore, letting $s' = s - 2|Y| \geqslant 0$, since $|\sigma| \leqslant 2|Y|$ and $q' = |Q'| = |Q| + |\sigma| \leqslant q + 2|Y|$ we have $pn \geqslant (r + s + q)(2r + 2s + 2q + 1) \geqslant (r + s' + q')(2r + 2s' + 2q' + 1)$ so we can apply the inductive hypothesis with f', F', Q', q', and s' to show that the above quantity is no greater than $\alpha^{s-2|Y|}$.

Since the above calculation gives the same upper bound for term (3) for all values of σ, we can pull this quantity outside the sum to obtain

$$\alpha^{s-2|Y|} \sum_{i=0}^{|Y|} \sum_{\substack{\sigma \in P(D,Y), \\ |\sigma_1| = |Y| - i}} Pr[\rho(\sigma) = * \mid F\restriction_\rho = 0 \wedge \rho(Q) = * \ \wedge \rho(T \setminus Y) = 1 \wedge \rho(Y) = *]. \tag{5}$$

Now we will estimate the inner summation for a fixed value of i. As above, we replace the condition $F\restriction_\rho = 0 \wedge \rho(T \setminus Y) = 1$ by the single condition $F'\restriction_\rho = 0$. Also, for a particular σ, the event $\rho(\sigma) = *$ is equivalent to the events $\rho(\sigma_1) = * \wedge \rho(\sigma_2) = *$. Because $\rho(\sigma_1) = *$ is implied by $\rho(Y) = *$, the inner summation is equivalent to

$$\sum_{\substack{\sigma \in P(D,Y), \\ |\sigma_1| = |Y| - i}} Pr[\rho(\sigma_2) = * \mid F'\restriction_\rho = 0 \wedge \rho(Q) = * \wedge \rho(Y) = *].$$

If we let Q' be the conjunction of Q and Y then we can rewrite this as

$$\sum_{\substack{\sigma \in P(D,Y), \\ |\sigma_1| = |Y| - i}} Pr[\rho(\sigma_2) = * \mid F'\restriction_\rho = 0 \wedge \rho(Q') = *].$$

We would like to remove the conditioning on $F'\restriction_\rho = 0$ but we cannot do this for each term as in Claim A. We have to consider the terms in this sum in the aggregate rather than individually. Let N_i be the number of σ's such that $|\sigma_1| = |Y| - i$. Then the above probability can be rewritten as

$$N_i \cdot Pr_{(\sigma_2, \rho)}[\rho(\sigma_2) = * \mid F'\restriction_\rho = 0 \wedge \rho(Q') = *],$$

where the above probability is over all pairs (σ_2, ρ), such that $|\sigma_1| = |Y| - i$. For each σ_2, let u be the set of vertices in σ_2 which are not contained in $v(Y)$. Note that

the number of vertices in u equals $2i$. Also note that for σ_2 chosen at random, u is a uniformly distributed set over $D'' = D \setminus v(Q')$ having these properties. Letting V_i be the collection of all sets over D'' of size $2i$, this probability is equal to $N_i \cdot Pr_{(u,\rho),}[\rho(u) = * \mid F' \restriction_\rho = 0 \wedge \rho(Q') = *]$, where the probability is over all pairs (u, ρ), such that $u \in V_i$ and $\rho \in \mathcal{M}_p^D$. This probability can be further divided according to $\#(\rho)$, the exact number of vertices of D that are unmatched by ρ:

$$N_i \cdot \sum_{j=0}^{n} Pr_{(u,\rho)}[\rho(u) = * \mid F' \restriction_\rho = 0 \wedge \rho(Q') = * \wedge \#(\rho) = 2j+1]$$

$$\times Pr_{(u,\rho)}[\#(\rho) = 2j+1 \mid F' \restriction_\rho = 0 \wedge \rho(Q') = *].$$

Given that $\#(\rho) = j$, for a randomly chosen u the event $\rho(u) = *$ is independent of $F' \restriction_\rho = 0$. Thus, the above probability is equal to

$$N_i \cdot \sum_{j=0}^{n} Pr_{(u,\rho)}[\rho(u) = * \mid \rho(Q') = * \wedge \#(\rho) = 2j+1]$$

$$\times Pr[\#(\rho) = 2j+1 \mid F' \restriction_\rho = 0 \wedge \rho(Q') = *],$$

where we have dropped the subscript on the probability in the second factor in each term since this probability only depends on ρ. For all $k \leqslant n$, $\sum_{j \geqslant k} Pr[\#(\rho) = 2j+1 \mid F' \restriction_\rho = 0 \wedge \rho(Q') = *]$ equals $Pr[\#(\rho) \geqslant 2k+1 \mid F' \restriction_\rho = 0 \wedge \rho(Q') = *]$, because the events are disjoint. Similarly, $\sum_{j \geqslant k} Pr[\#(\rho) = 2j+1 \mid \rho(Q') = *]$ equals $Pr[\#(\rho) \geqslant 2k+1 \wedge \rho(Q') = *]$.

Claim B. *For all k, $Pr[\#(\rho) \geqslant k \mid F' \restriction_\rho = 0 \wedge \rho(Q') = *] \leqslant Pr[\#(\rho) \geqslant k \mid \rho(Q') = *]$.*

Using Claim B and noting that $Pr_{(u,\rho)}[\rho(u) = * \mid \#(\rho) = 2j+1 \wedge \rho(Q') = *] \leqslant Pr_{(u,\rho)}[\rho(u) = * \mid \#(\rho) = 2j+3 \wedge \rho(Q') = *]$ for all $j \geqslant 0$, we can apply Lemma 22 with $\alpha_j = Pr_{(u,\rho)}[\rho(u) = * \mid \#(\rho) = 2j+1 \wedge \rho(Q') = *]$, $a_j = Pr[\#(\rho) = 2j+1 \mid F' \restriction_\rho = 0 \wedge \rho(Q') = *]$, and $b_j = Pr[\#(\rho) = 2j+1 \wedge \rho(Q') = *]$ to show that the above probability is no greater than

$$N_i \cdot \sum_{j=0}^{n} Pr_{(u,\rho)}[\rho(u) = * \mid \#(\rho) = 2j+1 \wedge \rho(Q') = *]$$

$$\times Pr[\#(\rho) = 2j+1 \mid \rho(Q') = *]$$

which is equal to $N_i \cdot Pr_{(u,\rho)}[\rho(u) = * \mid \rho(Q') = *]$.

Since for each fixed value of $u \in V_i$, the probability that $\rho(u) = *$ is the same, the above probability is equal to $N_i \cdot Pr[\rho(u) = * \mid \rho(Q') = *]$, where the probability is now over the distribution $\mathcal{M}_p^D$. Using the fact that $|u| + |v(Q')| = 2i + 2|Y| + 2q \leqslant 2r + 2s + 2q$ and the bound on pn, we can apply Lemma 19 to conclude that for $u \in V_i$, $Pr[\rho(u) = * \mid \rho(Q') = *] \leqslant (2p)^{2i}$.

Recall that N_i is equal to the number of σ's such that $|\sigma_1| = |Y| - i$. There are at most

$$\binom{2|Y|}{2i}\frac{(2|Y|-2i)!}{2^{|Y|-i}(|Y|-i)!} < \binom{2|Y|}{2i}(|Y|-i)^{|Y|-i}$$

choices of σ_1 with $|\sigma_1| = |Y| - i$ and for each such σ_1 there are at most $\binom{m}{2i}(2i)! < m^{2i}$ choices of σ_2 where $m = 2n + 1 - |Y| - 2q \leqslant 2n$. Thus, there are a total of at most

$$\binom{2|Y|}{2i}(|Y|-i)^{|Y|-i}(2n)^{2i}$$

choices of $\sigma \in P(D, Y)$ such that $|\sigma_1| = |Y| - i$.

Thus, for all Y such that $2|Y| \leqslant s$, using the expression in (5), we have

$$\begin{aligned}
&Pr\,[\delta(f\restriction_\rho) \geqslant s \mid F\restriction_\rho = 0 \wedge \rho(Q) = * \wedge f_1\restriction_\rho \neq 0 \wedge *(\rho_T) = Y] \\
&\quad < \alpha^{s-2|Y|}\sum_{i=0}^{|Y|}\binom{2|Y|}{2i}2n^{2i}(|Y|-i)^{|Y|-i}(2p)^{2i} \\
&\quad \leqslant \alpha^{s-2|Y|}\sum_{i=0}^{|Y|}\binom{2|Y|}{2i}2n^{2i}(|Y|)^{|Y|-i}(2p)^{2i} \\
&\quad = \alpha^{s-2|Y|}|Y|^{|Y|}\sum_{i=0}^{|Y|}\binom{2|Y|}{2i}\left(\frac{4pn}{\sqrt{|Y|}}\right)^{2i} \\
&\quad < \alpha^{s-2|Y|}|Y|^{|Y|}\sum_{j=0}^{2|Y|}\binom{2|Y|}{j}\left(\frac{4pn}{\sqrt{|Y|}}\right)^{j} \\
&\quad = \alpha^{s-2|Y|}|Y|^{|Y|}\left(\frac{4pn}{\sqrt{|Y|}}+1\right)^{2|Y|} \\
&\quad \leqslant \alpha^{s-2|Y|}|Y|^{|Y|}\left(\frac{5pn}{\sqrt{|Y|}}\right)^{2|Y|} \\
&\quad = \alpha^{s-2|Y|}(5pn)^{2|Y|}.
\end{aligned}$$

For Y such that $2|Y| > s$ we cannot use the expansion in terms of (3) and (4) to estimate this probability. However, in this case, since $\alpha \leqslant 1$ and $5pn \geqslant 1$, $\alpha^{s-2|Y|}(5pn)^{2|Y|} > 1$ so it still is an upper bound on this probability.

Plugging in the bounds we have for the terms (1) and (2) we get

$$\begin{aligned}
&Pr[\delta(f\restriction_\rho) \geqslant s \mid F\restriction_\rho = 0 \wedge \rho(Q) = * \wedge f_1\restriction_\rho \neq 0] \\
&\quad \leqslant \sum_{\substack{Y \subseteq T, \\ Y \neq \emptyset}} \alpha^{s-2|Y|}(5pn)^{2|Y|}(9p^2n)^{|Y|} \\
&\quad = \alpha^s \sum_{\substack{Y \subseteq T, \\ Y \neq \emptyset}} \left(\frac{225p^4n^3}{\alpha^2}\right)^{|Y|} \\
&\quad \leqslant \alpha^s \sum_{i=1}^{r} \binom{r}{i} \left(\frac{225p^4n^3}{\alpha^2}\right)^{i} \\
&\quad = \alpha^s \left[\left(1 + \frac{225p^4n^3}{\alpha^2}\right)^{r} - 1\right] \\
&\quad \leqslant \alpha^s.
\end{aligned}$$

The last inequality holds since α satisfies $(1 + 225p^4n^3/\alpha^2)^r \leqslant 2$. □

We now provide the proof of Claim A above.

Lemma 24 (Claim A). *Let $|D| = 2n + 1$ and Y and Q be disjoint maps over D with $|Y| = k$ and $|Q| = q$. Let F be an arbitrary function defined over the variables on D. When $pn \geqslant (k + q)(2k + 2q + 1)$, $Pr[\rho(Y) = * \mid F\restriction_\rho = 0 \wedge \rho(Q) = * \wedge \rho(Y) \neq 0] \leqslant Pr[\rho(Y) = * \mid \rho(Q) = * \wedge \rho(Y) \neq 0]$.*

Proof. As in previous proofs of switching lemmas, we will prove Claim A by showing that

$$Pr[F\restriction_\rho = 0 \mid \rho(Y) = * \wedge \rho(Q) = *] \leqslant Pr[F\restriction_\rho = 0 \mid \rho(Y) \neq 0 \wedge \rho(Q) = *].$$

This proves the claim because for arbitrary events A, B, and C, $Pr[A \mid B \wedge C] \leqslant Pr[A \mid C] \leftrightarrow Pr[B \mid A \wedge C] \leqslant Pr[B \mid C]$.

Let $V = i$ denote the event that there are exactly $2i+1$ elements in $D \setminus (v(Y) \cup v(Q))$ not matched with points in $D \setminus (v(Y) \cup v(Q))$ by π. Note that this number is always odd since $|D| = 2n + 1$ and $|v(Y) \cup v(Q)|$ is even. Furthermore, since each such point either is matched to a point of $v(Y) \cup v(Q)$ or is the point unmatched by π we have $0 \leqslant i \leqslant k + q$. Therefore, $Pr[F\restriction_\rho = 0 \mid \rho(Y) = * \wedge \rho(Q) = *]$ is equal to

$$\sum_{i=0}^{k+q} Pr[F\restriction_\rho = 0 \mid \rho(Y) = * \wedge \rho(Q) = * \wedge V = i] \cdot Pr[V = i \mid \rho(Y) = * \wedge \rho(Q) = *].$$

Similarly, $Pr[F\restriction_\rho = 0 \mid \rho(Y) \neq 0 \wedge \rho(Q) = *]$ is equal to

$$\sum_{i=0}^{k+q} Pr[F\restriction_\rho = 0 | \rho(Y) \neq 0 \wedge \rho(Q) = * \wedge V = i] \cdot Pr[V = i \mid \rho(Y) \neq 0 \wedge \rho(Q) = *].$$

We will show:

(1) For each i, $0 \leqslant i \leqslant k+q$,

$$Pr[F\restriction_\rho = 0 \mid \rho(Y) = * \wedge \rho(Q) = * \wedge V = i] \leqslant Pr[F\restriction_\rho = 0 \mid \rho(Y) \neq 0 \wedge \rho(Q) = * \wedge V = i].$$

(2) For each i, $0 \leqslant i \leqslant k+q$,

$$Pr[F\restriction_\rho = 0 \mid \rho(Y) = * \wedge \rho(Q) = * \wedge V = i] \geqslant Pr[F\restriction_\rho = 0 \mid \rho(Y) = * \wedge \rho(Q) = * \wedge V = i+1].$$

(3) For all j, $0 \leqslant j \leqslant k+q$,

$$\sum_{i=0}^{j} Pr[V = i \mid \rho(Y) = * \wedge \rho(Q) = *] \leqslant \sum_{i=0}^{j} Pr[V = i \mid \rho(Y) \neq 0 \wedge \rho(Q) = *].$$

Then, by Lemma 22, the claim follows.

We will first prove step (1). We break up the collection of restrictions satisfying $(\rho(Y) \neq 0 \wedge \rho(Q) = * \wedge V = i)$ into equivalence classes as follows. Suppose that ρ is chosen as (π, π_*). Let $D' = D \setminus (v(Y) \cup v(Q))$ and $\pi^{D'}$ and $\pi_*^{D'} \subseteq \pi^{D'}$ be the restrictions of π and π_* respectively to those edges both of whose endpoints lie in D'. The condition that $V = i$ says that there are exactly $2i+1$ points in D' unmatched by π so it simply fixes the size of $\pi^{D'}$.

Consider some fixed choice of $\pi^{D'}$ and $\pi_*^{D'}$ and consider all $\rho = (\pi, \pi_*)$ consistent with them. Given this, the condition that $(\rho(Y) = * \wedge \rho(Q) = *)$ completely determines ρ as a restriction so the event $F\restriction_\rho = 0$ is completely determined. If $F\restriction_\rho = 0$ then it will certainly also be forced to 0 if we allow the possibility that ρ does not assign all of Y to $*$ so in particular it is also 0 if $(\rho(Y) \neq 0 \wedge \rho(Q) = *)$. Therefore, it is at least as likely to be forced to 0 if $(\rho(Y) \neq 0 \wedge \rho(Q) = *)$ as if $(\rho(Y) = * \wedge \rho(Q))$. Since this is true for each choice of $\pi^{D'}$ and $\pi_*^{D'}$, it is true over all. This completes the proof of (1).

The intuition behind step (2) is simply that the larger $V = i$ is, the more stars ρ is likely to have, and hence the less likely it is that F is forced to 0. To prove step (2), let R be the conjunction of Q and Y and fix i. Then we want to prove

$$Pr[F\restriction_\rho = 0 \mid \rho(R) = * \wedge V = i] \leqslant Pr[F\restriction_\rho = 0 \mid \rho(R) = * \wedge V = i-1].$$

The probability $Pr[F\restriction_\rho = 0 \mid \rho(R) = * \wedge V = i]$ can be divided according to the exact number of *'s that are assigned by ρ:

$$Pr[F\restriction_\rho = 0 \mid \rho(R) = * \wedge V = i] = \sum_{j=0}^{n} Pr[F\restriction_\rho = 0 \wedge \#(\rho) = 2j+1 \mid \rho(R) = * \wedge V = i].$$

A similar equation holds when $\rho(R) = *$ and $V = i - 1$:

$$Pr[F\restriction_\rho = 0 \mid \rho(R) = * \wedge V = i-1]$$
$$= \sum_{j=0}^{n} Pr[F\restriction_\rho = 0 \wedge \#(\rho) = 2j+1 \mid \rho(R) = * \wedge V = i-1].$$

(When $j < i + |R|$, $Pr[F\restriction_\rho = 0 \wedge \#(\rho) = 2j+1 \mid \rho(R) = * \wedge V = i] = 0$.)

First, note that the conditioning that $V = i-1$, or $V = i$ is irrelevant to the probability that $F\restriction_\rho = 0$, given that $\#(\rho) = 2j+1$ and $\rho(R) = *$, Therefore, we have

$$Pr[F\restriction_\rho = 0 \mid V = i-1 \wedge \rho(R) = * \wedge \#(\rho) = 2j+1]$$
$$= Pr[F\restriction_\rho = 0 \mid V = i \wedge \rho(R) = * \wedge \#(\rho) = 2j+1]$$
$$= Pr[F\restriction_\rho = 0 \mid \rho(R) = * \wedge \#(\rho) = 2j+1].$$

Thus, it is left to show

$$\sum_{j=0}^{n} Pr[F\restriction_\rho = 0 \mid \#(\rho) = 2j+1 \wedge \rho(R) = *]$$
$$\times Pr[\#(\rho) = 2j+1 \mid V = i \wedge \rho(R) = *]$$
$$\leqslant \sum_{j=0}^{n} Pr[F\restriction_\rho = 0 \mid \#(\rho) = 2j+1 \wedge \rho(R) = *]$$
$$\times Pr[\#(D') = j \mid V = i-1 \wedge \rho(R) = *].$$

By Lemma 21 we know that $Pr[F\restriction_\rho = 0 \mid \rho(R) = * \wedge \#(\rho) = 2j+1]$ is a nonincreasing function of j. That is, the larger j is, the less likely that F is forced to 0. Therefore, it suffices to show that the conditioning that $V = i$ makes it more likely that j is larger than the conditioning that $V = i-1$. More explicitly, by Lemma 22 it suffices to show that for all j,

$$Pr[\#(\rho) \leqslant 2j+1 \mid V = i \wedge \rho(R) = *] \leqslant Pr[\#(\rho) \leqslant 2j+1 \mid V = i-1 \wedge \rho(R) = *].$$

Recall that the distribution given $(V = i \wedge \rho(R) = *)$ can be described as follows. First, choose k at random, according to the binomial distribution, shifted by i: $B(n - |R| - i, p) + i$; then choose a random matching π of $n - |R| - k$ edges on $D \setminus v(R)$. The distribution of $\#(\rho)$ given this conditioning is then given by $2[B(n - |R| - i, p) + i + |R|] + 1$. The distribution of $\#(\rho)$ given $(V = i - 1 \wedge \rho(R) = *)$ is then $2[B(n - |R| - i + 1, p) + i - 1 + |R|] + 1$. Therefore, it is clear that $Pr[\#(\rho) \leqslant 2j+1 \mid V = i \wedge \rho(R) = *] \leqslant Pr[\#(\rho) \leqslant 2j+1 \mid V = i-1 \wedge \rho(R) = *]$. By Lemma 22 this completes step (2).

We will now prove Step (3). We want to show that for all j, $0 \leqslant j \leqslant k+q$,

$$Pr[V \leqslant j \mid \rho(Y) = * \wedge \rho(Q) = *] \leqslant Pr[V \leqslant j \mid \rho(Y) \neq 0 \wedge \rho(Q) = *].$$

The right side of this inequality is a weighted sum of $Pr[V \leqslant j \mid \rho(Y') = 1 \wedge \rho(Y \setminus Y') = * \wedge \rho(Q) = *]$, where Y' ranges over all subsets of Y. We want to show that

this probability is smallest when $Y' = \emptyset$. Notice that for all choices of $Y' \subseteq Y$ of a given size the probability is the same. Therefore, it suffices to prove that, for $Y' \subseteq Y$ and $e \in Y \setminus Y'$ and $Y'' = Y' \cup \{e\}$,

$$\begin{aligned} &Pr[V \leqslant j \mid \rho(Y') = 1 \wedge \rho(Y \setminus Y') = * \wedge \rho(Q) = *] \\ &\quad \leqslant Pr[V \leqslant j \mid \rho(Y'') = 1 \wedge \rho(Y \setminus Y'') = * \wedge \rho(Q) = *]. \end{aligned}$$

Letting $R = Q \cup (Y \setminus Y'')$ we can rewrite what we want to show as

$$\begin{aligned} &Pr[V \leqslant j \mid \rho(Y') = 1 \wedge \rho(e) = * \wedge \rho(R) = *] \\ &\quad \leqslant Pr[V \leqslant j \mid \rho(Y') = 1 \wedge \rho(e) = 1 \wedge \rho(R) = *]. \end{aligned}$$

Let $D' = D \setminus v(Y')$ and $m = n - |Y'|$. Note that the events other than $\rho(Y') = 1$ do not involve variables that touch Y', and the conditional distribution of ρ on the variables over D' given that $\rho(Y') = 1$ is the same as a ρ' chosen from $\mathcal{M}_p^{D'}$. Therefore, it is equivalent to show that

$$Pr[V \leqslant j \mid \rho'(e) = * \wedge \rho'(R) = *] \leqslant Pr[V \leqslant j \mid \rho'(e) = 1 \wedge \rho'(R) = *]$$

which in turn is equivalent to

$$\begin{aligned} &\frac{\sum_{i \leqslant j} Pr[V = i \ \wedge \ \rho'(e) = * \ \wedge \ \rho'(R) = *]}{Pr[\rho'(e) = * \ \wedge \ \rho'(R) = *]} \\ &\leqslant \frac{\sum_{i \leqslant j} Pr[V = i \ \wedge \ \rho'(e) = 1 \ \wedge \ \rho'(R) = *]}{Pr[\rho'(e) = 1 \ \wedge \ \rho'(R) = *]}. \end{aligned} \tag{6}$$

Consider the two denominators:

$$\begin{aligned} &Pr[\rho'(e) = * \ \wedge \ \rho'(R) = *] \\ &\quad = Pr[\rho'(e) = * \mid \rho'(R) = *] \cdot Pr[\rho'(R) = *] \geqslant p^2 Pr[\rho'(R) = *] \end{aligned}$$

by Lemma 19,

$$\begin{aligned} Pr[\rho'(e) = 1 \ \wedge \ \rho'(R) = *] &\leqslant Pr[\rho'(e) \neq 0 \ \wedge \ \rho'(R) = *] \\ &= Pr[\rho'(e) \neq 0 \mid \rho'(R) = *] \cdot Pr[\rho'(R) = *] \\ &\leqslant \frac{1-p}{m} Pr[\rho'(R) = *] \end{aligned}$$

by Lemma 20. Therefore,

$$Pr[\rho'(e) = 1 \wedge \rho'(R) = *] \leqslant \frac{1-p}{mp^2} Pr[\rho'(e) = * \wedge \rho'(R) = *].$$

Now let $r = |R| + 1$. The condition (6) is clearly satisfied when $j \geqslant r - 1$ since the probability on the right side is 1 so we can assume that $j \leqslant r-2$. We will show that for $j \leqslant r-2$ each term in the numerator on the right of (6) is at least $2(1-p)/pr(2r+1)$

times the corresponding term on the left. This is sufficient since then the right-hand side of (6) is at least

$$\frac{2(1-p)}{pr(2r+1)}\frac{mp^2}{1-p} = \frac{2pm}{r(2r+1)} \geqslant 1$$

times the left-hand side of (6) since $2pm \geqslant pn \geqslant r(2r+1)$ by assumption since $r \leqslant |Y| + |Q| \leqslant k+q$.

We will actually compare the two probabilities with $V = i$ in the two cases given a fixed choice of the set I of $2i+1$ points outside of $v(R \cup e)$ that correspond to the $V = i$ event and a fixed choice of ρ' outside of $I \cup v(R \cup e)$:

In the case that $\rho'(e) = *$, we can count the number of ways that π' can match the points in $I \cup v(R \cup e)$ as follows: Consider some fixed ordering of the points in I. We first choose an ordering of the $2r$ points of $R \cup e$ along with an extra dummy point. The first $2i+1$ of the points are paired with the $2i+1$ points of I in order. The remaining points of $v(R \cup e)$ are paired up consecutively according to this order. The point paired with the dummy point will be the point unmatched by π; the other pairs will constitute the matching edges of π'. This ordering overcounts the number of choices of the $r-i$ edges occurring inside $v(R \cup e)$ – there are 2 equivalent orderings of the endpoints within each edge and $(r-i)!$ equivalent orderings amongst the edges. Thus, the total count is $(2r+1)!/2^{r-i}(r-i)!$. Finally, all $r+i$ edges of π' in $I \cup v(R \cup e)$ are set to $*$ which happens with probability p^{r+i}.

In the case that $\rho(e') = 1$, e is guaranteed to be matched by π' so the number of choices for π' can be counted as above except that R replaces $R \cup e$ in the count. Therefore the number of choices of π' is $(2r-1)!/2^{r-i-1}(r-i-1)!$. Finally, e must be set to 1 and the remaining $r+i-1$ edges must be set to $*$ which happens with probability $(1-p)p^{r+i-1}$.

Therefore, the case that $\rho'(e) = 1$ is

$$\frac{\frac{(2r-1)!(1-p)p^{r+i-1}}{2^{r-i-1}(r-i-1)!}}{\frac{(2r+1)!p^{r+i}}{2^{r-i}(r-i)!}} = \frac{(1-p)(r-i)}{pr(2r+1)} \geqslant \frac{2(1-p)}{pr(2r+1)}$$

times as likely as the case that $\rho'(e) = *$ which is what we required. □

Acknowledgements

The authors would like to thank Russell Impagliazzo for his helpful comments and encouragement.

References

[1] M. Ajtai, The complexity of the pigeonhole principle, Proc. 29th Ann. Symp. on Foundations of Computer Science, White Plains, New York (IEEE Press, New York, 1988) 346–355.

[2] M. Ajtai, Parity and the pigeonhole principle, in: S.R. Buss and P.J. Scott, eds., Feasible Mathematics, A Mathematical Sciences Institute Workshop, Ithaca, New York (Birkhäuser, Basel, 1990) 1–24.

[2a] Beame, Cook, Edmonds, Impagliazzo and Patalsi, The complexity of NP search porblems, Proc. 27th Ann. ACM Symp. on Theory of Computing, Las Vegas, NV, USA (1995) 303–314.

[3] P.W. Beame and J. Håstad, Optimal bounds for decision problems on the CRCW PRAM, J. ACM 36 (1989) 643–670.

[4] P.W. Beame, R. Impagliazzo, J. Krajíček, T. Pitassi, P. Pudlák and A. Woods, Exponential lower bounds for the pigeonhole principle, Proc. 24th Ann. ACM Symp. on Theory of Computing, Victoria, BC, Canada (1992) 200–220.

[5] P. Beame and T. Pitassi, An exponential separation between the matching principle and the pigeonhole principle, Proc. 8th Ann. IEEE Symp. on Logic in Computer Science, Montreal, Quebec, 1993.

[6] S. Bellantoni, T. Pitassi and A. Urquhart, Approximation and small depth Frege proofs, Proc. Structure in Complexity Theory, 6th Ann. Conf., Chicago, IL (IEEE Press, New York, 1991) 367–391.

[7] S.R. Buss, Bounded Arithmetic, Vol. 3 Studies in Proof Theory (Bibliopolis, Napoli, 1986).

[8] S.R. Buss, Polynomial size proofs of the pigeonhole principle, J. Symbolic Logic 57 (1987) 916–927.

[9] S.A. Cook and R.A. Reckhow, The relative efficiency of propositional proof systems, J. Symbolic Logic 44 (1977) 36–50.

[10] A. Haken, The intractability of resolution, Theoret. Comput. Sci. 39 (1985) 297–305.

[11] J. Håstad, Computational Limitations of Small-Depth Circuits (MIT Press, New York, 1987). ACM Doctoral Dissertation Award Series, 1986.

[12] J. Krajíček, P. Pudlák and A. Woods, Exponential lower bounds to the size of bounded-depth Frege proofs of the pigeonhole principle, Random Structures and Algorithms 7 (1995) 15–39.

[13] C.H. Papadimitriou, On inefficient proofs of existence and complexity classes, Proc. 4th Czechoslovakian Symp. on Combinatorics, 1991.

[14] T. Pitassi, The power of weak formal systems, Ph.D. Thesis, University of Toronto, August 1992.

[15] T. Pitassi, P. Beame and R. Impagliazzo, Exponential lower bounds for the pigeonhole principle, Comput. Complexity 3 (1993) 97–140.

[16] J. Paris and A. Wilkie, Counting problems in bounded arithmetic, in: Methods in Mathematical Logic: Proc. 6th Latin American Symposium on Mathematical Logic 1983, Lecture Notes in Mathematics, Vol. 1130 (Springer, Berlin, 1985) 317–340.

[17] A. Urquhart, Hard examples for resolution, J. ACM 34 (1987) 209–219.