More on the Relative Strength of Counting Principles

Paul Beame Søren Riis

January 18, 1997

1 Introduction

Over the last several years, substantial progress has been made in the study of the complexity of propositional proof systems. A particularly noteworthy development in this effort has been the significant cross-fertilization between research on circuit complexity and research on propositional proof complexity. One area of application of circuit-complexity techniques in proof complexity has been in the study of constant-depth Frege systems, proof systems in the conventional axiom schema/inference rule format each of whose constituent propositional formulas has constant depth.

Ajtai [2] introduced circuit-complexity techniques to the study of constant-depth Frege systems and provided the first bound showing that the propositional pigeonhole principle did not have efficient constant-depth proofs. His arguments were simplified in [9] and the complexity lower bound improved to exponential in [7, 15, 13].

Given these results, a natural question to ask is: What is the power of the proof system if it is augmented by axiom schemas for some family of tautologies that does not have efficient proofs? Tautologies for several combinatorial principles have been studied. Ajtai [3] showed that if the pigeonhole principle is added as an axiom schema then the $Count_2$ tautologies still do not have efficient proofs, where the $Count_p$ tautologies express the fact that there is no perfect partition of a set M into blocks of size p, if $|M| \not\equiv 0 \pmod{p}$. Again, this bound was improved to exponential in [8, 16] and can be extended to arbitrary values of p. These arguments used the usual circuit-complexity techniques augmented by specialized combinatorial techniques to handle the new axiom schemas.

A natural next question (asked originally in [14]) was to examine the relative strength of the $Count_p$ principles for different values of p. Ajtai [1] first showed that, when p and q are distinct primes, proofs of $Count_q$ are not efficient even when given $Count_p$ as axiom schemas. To handle the $Count_p$ axioms Ajtai used a reworking of several ideas from the theory of representations of the symmetric group.

A different approach for handling the $Count_p$ axioms was taken by two other groups of researchers independently. This approach was a natural extension of the methods in

[8, 16]. Riis [16] laid out a framework involving showing that certain 'exceptional forests' of decision trees do not exist. This problem was left unsolved in [16]. Working along similar lines, Beame, Impagliazzo, Krajíček, Pitassi, and Pudlák [6] introduced the notion of a Nullstellensatz proof system and reduced the existence question for objects similar to Riis' exceptional forests to the degree required for certain proofs in this system. They also showed lower bounds on this degree using Ramsey theory, and thus extended Ajtai's results to a somewhat wider class of p, q combinations. Riis [17] applied similar Ramsey theory arguments directly to the forests themselves.

One important contribution of [6] was to show that exponential lower bounds for $Count_q$ given $Count_p$ would follow from improved degree bounds for the Nullstellensatz proofs. These improved degree bounds were shown by Buss, Impagliazzo, Krajíček, Pudlák, Razborov, and Sgall [10] who introduced a nice inductive method for producing such bounds.

The present paper uses similar techniques to give a further refinement of our understanding of the strength of these combinatorial principles. Thus far, we have referred to 'the' pigeonhole principle. However, there are a number of variations of the pigeonhole principle depending on the sizes of the domain and range of the map and on whether or not the map is required to be 'onto' (which is a weaker version). The lower bounds mentioned above have applied to either version equally and have assumed that the domain is one element larger than the range. We show that the onto version of the pigeonhole principle from $n+p^{\lfloor \epsilon \log n \rfloor}$ points to n points, $\textit{onto-PHP}_n^{n+p^{\lfloor \epsilon \log n \rfloor}}$, requires exponential size constant-depth proofs even given $Count_p$ as axiom schemas. The key feature of our argument is a new degree lower bound for Nullstellensatz proofs of $\textit{onto-PHP}_n^{n+p^k}$. (Our results strengthen the results in [18, 19] and are based on a substantially different presentation.)

Since $\textit{onto-PHP}_n^{n+p^{\lfloor \epsilon \log n \rfloor}}$ does follow efficiently from axiom schemas for the general PHP_n^{n+1}, and additional axiom schemas for $Count_p$ do yield efficient proofs of $\textit{onto-PHP}_n^{n+1}$, it follows that PHP_n^{n+1} requires exponential proofs given axiom schemas for $\textit{onto-PHP}_n^{n+1}$. Also, since additional axiom schemas for $Count_q$, for appropriate $q \neq p$ do give short proofs of $\textit{onto-PHP}_n^{n+p^{\lfloor \epsilon \log n \rfloor}}$, the exponential separations between $Count_q$ and $Count_p$ principles are also corollaries of our results. The idea of examining the relationship between $\textit{onto-PHP}_n^{n+p^{\lfloor \epsilon \log n \rfloor}}$ and $Count_p$ was originally used implicitly by Riis as the basis of the approach in [16] towards proving a separation between $Count_p$ and $Count_q$.

There has been a substantial improvement in the precision and presentation of the methods for proving lower bounds on constant-depth Frege systems with additional axiom schemas and the papers above do not give entirely self-contained explanations of the best of current techniques. In this paper we attempt to give as complete a presentation as possible.

We now outline the structure of the argument, giving references for the key techniques. We use the notion of a k-evaluation due to Krajíček, Pudlák, and Woods [13], incorporating the matching decision trees of Pitassi, Beame, Impagliazzo [15], and built for any small Frege proof using a switching lemma proved with the methods of Beame [4]. Then, as in the argument of Riis [16] and Beame and Pitassi [8], we show that having a k-evaluation implies the existence of a certain forest of matching decision trees. Following this we show, using a reduction analogous to that of Beame, Impagliazzo, Krajíček, Pitassi, and Pudlák [6], that the existence of such a forest implies a small degree Nullstellensatz refutation of an associated family of polynomials. Finally, the proof that such a small degree refutation does not exist is analogous to that of Buss, Impagliazzo, Krajíček, Pudlák, Razborov, and Sgall [10]. This last is the main new technical contribution and the reader who is familiar with the other aspects of this paper may wish to skip directly to section 8. In section 9 we combine the arguments from the previous sections to show our main results.

2 Frege Proofs and Counting Principles

A *Frege system* is a sound and implicationally complete propositional proof system with a finite number of axiom schemas and inference rules. The *size* of a proof in a Frege system is the total number of subformulas appearing in the proof. We consider formulas over the basis $\vee$ (binary) and $\neg$ with propositional variables and the constants 0 (false) and 1 (true) as atoms. We use $F \wedge G$ as a shorthand for $\neg(\neg F \vee \neg G)$ and $\bigvee_{i=1}^{r} F_i$ as a shorthand for an arbitrarily parenthesized tree of binary $\vee$'s with $F_1, \ldots, F_r$ at the leaves.

The *depth* of a formula F is the maximum number of runs of $\vee$ connectives on any path from an atom of F to the main connective of F. A *depth d Frege system* is a restriction of a Frege system to proofs all of whose formulas have depth at most d.

(One can define Frege systems over any basis of binary connectives and the sizes of proofs in these systems are polynomially related to each other. For constant depth d Frege systems this is also true, given the depth measure above, provided that one excludes connectives $\oplus$ and $\leftrightarrow$. We concentrate on connectives $\neg$ and $\vee$ for ease of presentation.)

Pigeonhole Principles Let D and R be disjoint sets and consider propositional variables P_{ij}, $i \in D$, $j \in R$. There are 3 natural variations of the Pigeonhole Principle:

"There is no total injective relation on $D \times R$"

$$PHP_R^D = (\bigvee_{i\in D} \neg \bigvee_{j \in R} P_{ij}) \vee \bigvee_{j\in R} \bigvee_{i\neq i' \in D} (P_{ij} \wedge P_{i'j}).$$

"There is no total 1-1 function from D to R:"

$$\textit{fun-}PHP_R^D = PHP_R^D \vee \bigvee_{i\in D} \bigvee_{j\neq j'\in R} (P_{ij} \wedge P_{ij'}).$$

These two are polynomially equivalent as may be seen by imposing an order on R and setting:

$$P'_{ij} = P_{ij} \wedge \bigvee_{j'<j} \neg P_{ij'}.$$

However, the following variant, as we will see, is strictly weaker than PHP^D_R:

"There is no 1-1 onto function from D to R:"

$$\textit{onto-PHP}^D_R = \textit{fun-PHP}^D_R \vee (\bigvee_{j\in R} \neg \bigvee_{i\in D} \neg P_{ij}).$$

When $|D| > |R|$ all the variations are tautologies and for the purposes of this paper we will assume that $|D| \geq |R|$. Clearly, each pigeonhole principle variation only depends on the sizes of D and R so we will usually refer to $PHP^{|D|}_{|R|}$, etc.

Counting Principles Let M be any set and p be a positive integer. Let $M^{(p)}$ denote the set of all p-element subsets of M. The mod p counting principle over M is defined on the set of propositional variables Y_e, $e \in M^{(p)}$.

"There is no perfect p-partition of M."

$$Count^M_p = (\bigvee_{v\in M} \bigwedge_{v\in e\in M^{(p)}} \neg Y_e) \vee \bigvee_{e,f\in M^{(p)},\ e\perp f} (Y_e \wedge Y_f)$$

where we write $e \perp f$ if $e \neq f$ and $e \cap f \neq \emptyset$. If $|M| \not\equiv 0(\mod p)$ then $Count^M_p$ is a tautology. Again, the tautology really only depends on $|M|$, so we can refer to $Count^m_p$. We will use $Count_p$ to refer to the family of tautologies $Count^m_p$ with $m \not\equiv 0$ ($\mod p$).

3 Restrictions and Matching Decision Trees

The main argument in this paper is a lower bound for the lengths of bounded-depth Frege proofs of *onto-PHP*, given *Count* formulas as axiom schemas. Therefore the propositional variables with which will primarily be concerned are those that appear in the *onto-PHP* formula. We introduce some notation for discussing formulas involving these variables.

Let $\mathcal{M}_{D\times R}$ be the set of all partial matchings on $D \times R$. A *matching term* A is $\bigwedge_{\langle i,j\rangle\in\pi} P_{ij}$ for some matching $\pi \in \mathcal{M}_{D\times R}$. A *matching disjunction* F is $\bigvee_i A_i$ where A_i are matching terms.

We say that i and j are the *endpoints* of P_{ij}. If Y is a term or a set of variables, then $v(Y)$ denotes the set of endpoints of variables in Y. We use $\mathrm{dom}(Y) = v(Y) \cap D$ and $\mathrm{range}(Y) = v(Y) \cap R$.

For $|R| = n$ and $|D| = n+m$, define $\mathcal{M}^{\ell}_{D\times R}$ to be the set of all partial matchings on $D \times R$, ρ which match all but ℓ nodes of R.

Every ρ in $\mathcal{M}^{\ell}_{D\times R}$ determines a unique partial assignment or *restriction*, r,

$$r(P_{ij}) = \begin{cases} 1 & \text{if } \langle i,j\rangle \in \rho \\ 0 & \text{if there is an } e' \in \rho \text{ such that } |v(e') \cap \{i,j\}| = 1 \\ * & \text{otherwise} \end{cases}$$

where $*$ indicates that P_{ij} is not assigned a value. If r is the restriction obtained from ρ, we will refer to both the restriction and the partial matching by ρ. For a Boolean formula F in the variables over $D \times R$ and a partial matching ρ, F restricted by ρ will be the formula in the variables unset by ρ that remains after assigning values to the set by ρ; we denote this by $F\restriction_\rho$. Given a set $S \subseteq D \cup R$, let $S\restriction_\rho$ denote $S \setminus v(\rho)$.

Our key interest in the set of restrictions given by the partial matchings $\rho \in \mathcal{M}^{\ell}_{D\times R}$, is that for any such ρ, *onto-PHP*$^D_R\restriction_\rho =$ *onto-PHP*$^{D\restriction_\rho}_{R\restriction_\rho}$.

We say that two partial matchings σ and τ are *compatible* if $\sigma \cup \tau$ is also a partial matching. When viewed as restrictions, we use the notation $\sigma\tau$ to denote the restriction defined by the partial matching $\sigma \cup \tau$.

Definition 3.1: A *matching decision tree over* $D \times R$ is a rooted directed tree T whose internal nodes are labelled by elements of $D \cup R$, and whose leaves may be labelled by elements of some label set L so that:

1. (a) If the root of T is labelled by $i \in D$ then for each $j \in R$ there is one out-edge from the root labelled $\langle i,j\rangle$.

   (b) If the root of T is labelled by $j \in R$ then for each $i \in D$ there is one out-edge from the root labelled $\langle i,j\rangle$.

   (c) There are no other out-edges from the root of T.

2. $T^{\langle i,j\rangle}$ is a matching decision tree over $D' \times R'$ where $D' = D \setminus \{i\}$, $R' = R \setminus \{j\}$, and $T^{\langle i,j\rangle}$ is the tree whose root is the node connected to the root of T by the edge labelled $\langle i,j\rangle$.

Define $\mathrm{Br}(T)$ to be the set of branches (root-leaf paths) in T and $\mathrm{Br}_a(T)$ to be the set of those branches in T with leaf label $a \in L$. The set of edge labels along any branch of T forms a partial matching. We identify a branch with its matching so we view $\mathrm{Br}(T)$ and $\mathrm{Br}_a(T)$ as sets of partial matchings. We say that $T \equiv a$ if and only if $\mathrm{Br}(T) = \mathrm{Br}_a(T)$.

Lemma 1: Let π be a matching and T be a matching decision tree over $D \times R$ such that $|\pi| + \mathrm{height}(T) \le \min(|D|, |R|)$. Then

(i) there is a $\sigma \in \mathrm{Br}(T)$ compatible with π.

(ii) the tree $T\restriction_\pi$ obtained by contracting all edges of T whose label is in π and deleting all edges of T (and their associated subtrees) whose labels are not compatible with π is a matching decision tree over $D\restriction_\pi \times R\restriction_\pi$.

Proof We prove part (ii) first by induction on the height of T: The base case when T is a single labelled vertex is trivial.

If the label of the root of T is touches π in edge $\langle i,j\rangle$ then $T\restriction_\pi = T^{\langle i,j\rangle}\restriction_{\pi'}$ where $\pi = \pi' \cup \langle i,j\rangle$. We apply the inductive hypothesis to $T^{\langle i,j\rangle}$ and π' over $(D\setminus\{i\})\times(R\setminus\{j\})$ to obtain the desired result.

If the label $i \in D$ of the root of T is not touched by π then the tree $T\restriction_\pi$ consists of the root of T with an outedge labelled by $\langle i,j\rangle$ for each $j \in R\restriction_\pi$ and this reaches subtree $T^{\langle i,j\rangle}\restriction_\pi$. Apply the inductive hypothesis to each such $T^{\langle i,j\rangle}$ over $(D\setminus\{i\})\cup(R\setminus\{j\})$ to obtain the desired result.

The case when the root label of T is $j \in R$ and is not touched by π is similar. Part (i) follows by observing that any branch in T that is contracted to a branch in $T\restriction_\pi$ suffices. $\square$

Definition 3.2: For any matching decision tree T with label set $L = \{0,1\}$, let T^c be the same tree as T except that the leaf labels 0 and 1 are reversed, i.e. $\mathrm{Br}_1(T^c) = \mathrm{Br}_0(T)$ and $\mathrm{Br}_0(T^c) = \mathrm{Br}_1(T)$. Matching decision tree T *represents* boolean formula or function f iff:

$$\forall \pi \in \mathrm{Br}(T),\ f\restriction_\pi \equiv \text{the leaf label of } \pi \text{ in } T.$$

Given matching decision tree T, the matching disjunction given by T is

$$Disj(T) = \bigvee_{\pi \in \mathrm{Br}_1(T)} \bigwedge_{\langle i,j\rangle \in \pi} P_{ij}$$

Note that T represents $Disj(T)$ and that if T has height $\leq k$ then $Disj(T)$ has terms of size $\leq k$. Observe that $Disj(T^c)$ is *not* equivalent to the negation of $Disj(T)$ but that if T represents f then the tree T^c does represent $\neg f$.

4 k-Evaluations

Definition 4.1: Let Γ be a set of formulas closed under subformulas. A *k-evaluation*, $\mathbf{T}$, of Γ is an association of a matching decision tree $\mathbf{T}_A$ of height $\leq k$ with each formula $A \in \Gamma$ such that

(1) T_0 and T_1 are single nodes labelled 0 and 1, respectively, and $T_{P_{ij}}$ is the unique tree of height 1 querying i that represents P_{ij},

(2) $T_{\neg A} = T_A^c$.

(3) If the major connective of A is $\vee$ then write $A = \bigvee_{i \in I} A_i$ where the major connective of each A_i is not $\vee$. It must be the case that T_A represents $\bigvee_{i \in I} Disj(T_{A_i})$.

Let $\mathbf{T}$ be a k-evaluation of a set of formulas containing formula A. We say that A *k-evaluates to true (false) under* $\mathbf{T}$ if and only if $T_A \equiv 1$ (respectively $T_A \equiv 0$).

Let the *size* of an axiom/rule in a Frege system $\mathcal{F}$ be the maximum number of distinct subformulas in it.

Lemma 2: Let P be a proof in Frege system $\mathcal{F}$ whose rules have size at most s, augmented by $Count_p$ axiom schemas. Suppose that $sk \leq |R| \leq |D|$ and let $\mathbf{T}$ be a k-evaluation of the set of subformulas of P. If every $Count_p$ axiom in $\mathcal{P}$ k-evaluates to true under $\mathbf{T}$ then all formulas in $\mathcal{P}$ k-evaluate to true under $\mathbf{T}$.

Proof By induction on the number of Frege axioms and inferences in P.

Consider a Frege axiom/inference in P:

$$\frac{A_1(B_1/p_1, \cdots, B_m/p_m), \ldots, A_\ell(B_1/p_1, \cdots, B_m/p_m)}{A_0(B_1/p_1, \cdots, B_m/p_m)}$$

where the inference rule $\mathcal{R}$ is:

$$\frac{A_1(p_1, \cdots, p_m), \ldots, A_\ell(p_1, \cdots, p_m)}{A_0(p_1, \cdots, p_m)}$$

and assume that each $A_i(B_1/p_1, \cdots, B_m/p_m)$ for $1 \leq i \leq \ell$ k-evaluates to true under $\mathbf{T}$. We now show that this also holds for $A_0(B_1/p_1, \ldots, B_m/p_m)$:

Let $\mathcal{A}$ be the set of distinct subformulas of $\mathcal{R}$ and let Γ be $\mathcal{A}(B_1/p_1, \cdots, B_m/p_m)$. By assumption $|\Gamma| \leq s$, say $\Gamma = \{A_0, \ldots, A_\ell, \ldots, A_j\}$ for $j < s$.

Let $\pi_0 \in \mathrm{Br}(T_{A_0})$. Since $sk \leq n$ we can apply the Lemma 1 to inductively find $\pi_i \in \mathrm{Br}(T_{A_i})$ compatible with $\pi_0 \cdots \pi_{i-1}$ for $1 \leq i \leq j$. Therefore all the π_i are mutually compatible. Let $\pi = \pi_0 \pi_1 \cdots \pi_j \in M_n$.

Observe that for any $A_i \in \Gamma$, $Disj(T_{A_i}) \restriction_\pi$ is the constant 0 or 1 and define $V : \Gamma \to \{0, 1\}$ by $V(A) = Disj(T_{A_i}) \restriction_\pi$. By the definition of k-evaluations, V is a consistent truth evaluation and by assumption

$$V(A_1) = \cdots = V(A_\ell) = 1.$$

Since the rule $\mathcal{R}$ is sound it follows that $V(A_0) = 1$, i.e.

$$Disj(T_{A_0}) \restriction_\pi = 1.$$

Since π extends branch π_0 of T_{A_0}, the leaf label of π_0 must be 1 as required. $\square$

On the other hand we show that the tree associated with the goal formula of the proof cannot k-evaluate to true.

Lemma 3: If $k + 1 \leq |R| < |D|$ and $\mathbf{T}$ is a k-evaluation of a set of formulas closed under subformulas and containing *onto-PHP*$_R^D$ then *onto-PHP*$_R^D$ does not k-evaluate to true under $\mathbf{T}$.

Proof In fact we show that every leaf of $T_{\textit{onto-PHP}_R^D}$ has label 0. By definition of a k-evaluation it is necessary and sufficient to show that $\mathrm{Br}_1(T_A) = \emptyset$ for each disjunct A in PHP_R^D.

Case 1: $A = (P_{ij} \wedge P_{i'j}) = \neg(\neg P_{ij} \vee \neg P_{i'j}) = \neg B$

Let $\pi \in \mathrm{Br}(T_A)$. It is also in $\mathrm{Br}(T_B)$. Since T_B represents $Disj(T_{\neg P_{ij}}) \vee Disj(T_{\neg P_{i'j}})$, it suffices to show that π is compatible with some element in $\mathrm{Br}_1(T_{\neg P_{ij}})$ or in $\mathrm{Br}_1(T_{\neg P_{i'j}})$.

By definition $T_{\neg P_{ij}}$ has height 1 with root label i and all its leaves are labelled 1 except the one below the out-edge with label $\langle i, j \rangle$.

Since $k + 1 \leq n$, $T_{\neg P_{ij}} \restriction_\pi$ is well-defined and consists of contractions of all branches compatible with π. If π does not contain $\langle i, j \rangle$ then some branch of $T_{\neg P_{ij}}$ other than $\langle i, j \rangle$ remains and this has leaf label 1.

If π does contain $\langle i, j \rangle$ then it does not contain $\langle i', j \rangle$ and we apply the same argument to $T_{\neg P_{i'j}}$.

Case 2: $A = \neg \bigvee_{j \in R} P_{ij}$

Similar to the previous case. Here, we show that $\pi \in \mathrm{Br}(T_A)$ is compatible with some element of $\mathrm{Br}_1(T_{P_{ij}})$ for some $j \in R$.

If π contains $\langle i, j \rangle$ for some $j \in R$ then every branch in $T_{P_{ij}}$ compatible with π will be in $\mathrm{Br}_1(T_{P_{ij}})$.

If π does not contain $\langle i, j \rangle$ for any $j \in R$ then let $j' \in R$ be unmatched by π (such a j' must exist). Since π matches neither i nor j' and $k + 1 \leq |R| < |D|$, π is compatible with the $\langle i, j' \rangle$ branch of $T_{P_{ij'}}$ which is what we need.

Case 3: The other *onto-PHP*$_R^D$ disjuncts are handled exactly as in Case 1. $\square$

5 Building a k-evaluation

Given an Frege proof $\mathcal{P}$ of limited size and depth we wish to find a restriction ρ such that after ρ is applied we have a suitable k-evaluation for all the subformulas in $\mathcal{P}$. This is too hard to do in a single step. Instead, we inductively build restrictions and k-evaluations for all depth i subformulas in $\mathcal{P}$ for $i = 0, \ldots, d$. The following lemma permits us to build upon previous k-evaluations.

Lemma 4: If Γ is a set of formulas closed under subformulas and $\mathbf{T}$ is a k-evaluation of Γ over $D \times R$ and ρ is a restriction on S with $|\rho| + k \leq |R|$ then the map $\mathbf{T}^\rho$ given by

$$T_F^\rho = \begin{cases} T_F \restriction_\rho & \text{if } \mathrm{Br}_1(T_F) \neq \emptyset \\ T_0 & \text{if } \mathrm{Br}_1(T_F) = \emptyset \end{cases}$$

is a k-evaluation of $\Gamma \restriction_\rho$ over $(D \times R) \restriction_\rho$.

Proof Note that for any matching decision tree T and formula F, if T represents F over $D \cup R$ then $T \restriction_\rho$ represents $F \restriction_\rho$ over $(D \cup R) \restriction_\rho$. Also note that for any matching decision tree T,

$$Disj(T) \restriction_\rho = Disj(T \restriction_\rho).$$

From this the Lemma follows easily by induction. (The extra condition when $\mathrm{Br}_1(T_F) = \emptyset$ is to make sure that $T_{P_{ij}} \restriction_\rho = T_0$ when $P_{ij} \restriction_\rho = 0$.) $\square$

The construction of decision trees for the higher level formulas of the proof uses the probabilistic method. The following so-called 'Switching Lemma' is the basis for that construction. We prove it as Lemma 7 below.

Lemma 5: Let F be an r-disjunction over $D \times R$ with $|R| = n$ and $|D| = n + m$. If $s \geq 0$ and $10m \leq \ell \leq (n/r)^{1/2}/10$ then, for ρ chosen uniformly at random from $\mathcal{M}^\ell_{D \times R}$, the probability that there does not exist a decision tree T over $(D \times R) \restriction_\rho$ of height less than s representing $F \restriction_\rho$ is less than $(1.5\ell^2 \sqrt{r/n})^s$.

Lemma 6: Let $|R| = n$, $|D| = n + m$. Let $n_0 = n$, $n_{i+1} = (n_i / 9 \log_2 S)^{1/4}$ for $i \geq 0$ and suppose that $n_d \geq \max\{10m, \log_2 S\}$. For any Frege proof $\mathcal{P}$ of size at most S and depth at most d in the pigeonhole variables on $D \times R$ there exists a restriction $\rho \in \mathcal{M}^{n_d}_{D \times R}$ such that there is a $\log_2 S$-evaluation $\mathbf{T}$ of the set of subformulas of $\mathcal{P} \restriction_\rho$ over $(D \times R) \restriction_\rho$.

Proof Let $k = \log_2 S$. We construct a sequence of restrictions $\rho_0, \ldots, \rho_d = \rho$ and maps $\mathbf{T}^0, \ldots, \mathbf{T}^d = \mathbf{T}$ such that for each $i = 0, \ldots, d$, $|R \restriction_{\rho_i}| = n_i$ and $\mathbf{T}^i$ is a k-evaluation of the set of formulas in $\mathcal{P}_i \restriction_{\rho_i}$, where $\mathcal{P}_i$ is the set of subformulas of depth at most i in $\mathcal{P}$. We only specify trees for unnegated formulas at each depth since negations do not add to depth and if we have a tree T_F then we easily have a tree $T_{\neg F} = T_F^c$ of the same height.

Base Case: $i = 0$. Let ρ_0 be the empty restriction. The only nodes of depth 0 are inputs and their negations. For each literal P_{ij}, let $T^0_{P_{ij}}$ be a tree of height 1 that queries i and has its only leaf label 1 on the node reached by edge labelled $\langle i, j \rangle$. Let T^0_b be a single node labeled b for $b = 0, 1$.

INDUCTION STEP: Now suppose that after ρ_i is applied we have a k-evaluation $\mathbf{T}^i$ of $\mathcal{P}_i\restriction_{\rho_i}$. We wish to find a π such that $\rho_{i+1}=\rho_i\pi$ and extend T^i to a k-evaluation $\mathbf{T}^{i+1}$ of $\mathcal{P}_{i+1}\restriction_{\rho_{i+1}}$.

Now for any choice of $\pi\in\mathcal{M}^{n_{i+1}}_{(D\times R)\restriction_{\rho_i}}$ and any $A\in\mathcal{P}_i$, using Lemma 4 we can define $T^{i+1}_A=(T^i)^{\pi}_A$. Observe that for such $A\in\mathcal{P}_i$, $Disj(T^{i+1}_A)=Disj(T^i_A\restriction_\pi)$ which is a k-disjunction.

It remains to choose π and define T^{i+1}_A for $A\in\mathcal{P}_{i+1}\setminus\mathcal{P}_i$ of the form $A=\bigvee_j A_j$ where $A\in\mathcal{P}_{i+1}\setminus\mathcal{P}_i$. and each $A_j\in\mathcal{P}_i$. We consider π chosen at random from $\mathcal{M}^{n_{i+1}}_{(D\times R)\restriction_{\rho_i}}$. By Lemma 5, the probability that π does not admit a choice for T^{i+1}_A is

$$<1.5n^2_{i+1}\sqrt{(\log_2 S)/n_i})^{\log_2 S}=2^{-\log_2 S}=1/S.$$

Since $|\mathcal{P}_{i+1}\setminus\mathcal{P}_i|\leq S$ the probability that some choice of π works for all formulas in $\mathcal{P}_{i+1}\setminus\mathcal{P}_i$ is strictly less than 1. We choose this π, fix $\rho_{i+1}=\rho_i\pi$ and set T^{i+1}_A according to that π for all $A\in\mathcal{P}_{i+1}\setminus\mathcal{P}_i$. The conditions for $\mathbf{T}^{i+1}$ are clearly satisfied. $\square$

We assume that there is a total order on the elements of $D\cup R$ with all elements of D preceding those of R. Let $K\subseteq D\cup R$ and define

$$\mathcal{M}_{D\times R}(K)=\{\pi\mid K\subseteq v(\pi)\text{ and }\forall e\in\pi.v(e)\cap K\neq\emptyset\},$$

i.e., all minimal partial matchings over $D\times R$ which involve all of the elements of K.

We define the *complete matching tree* for $K\subseteq D\cup R$ over $D\times R$ as a matching decision tree over $D\times R$ with no leaf labels. It is the unique tree T such that $\mathrm{Br}(T)=\mathcal{M}_{D\times R}(K)$ and the query at each node v is the smallest element of K that is not an endpoint of the matching associated with the path from the root to v.

Given a disjunction F over D, assume that F has a total order on its terms and an order on the variables within each term. A restriction ρ is applied to F in order, so that $F\restriction_\rho$ is the disjunction whose terms consist of those terms of F that are not falsified by ρ, each shortened by removing any variables that are satisfied by ρ, and taken in the order of occurrence of the original terms on which they are based.

The *canonical decision tree for F over $D\times R$*, $T_{D\times R}(F)$ is defined inductively as follows:

1. If F is the constant function 0 or 1 (contains no terms or has an empty first term, respectively) then $T_{D\times R}(F)$ consists of a single leaf node labelled by the appropriate constant value.

2. If the first term C_1 of F is not empty then let F' be the remainder of F so that $F=C_1\vee F'$. Let $K=v(C_1)$. We start with the complete matching tree for K. The paths of this tree correspond exactly to elements of $\mathcal{M}_{D\times R}(K)$. Let v_σ be the leaf node corresponding to the path labelled by $\sigma\in\mathcal{M}_{D\times R}(K)$. To obtain $T_{D\times R}(F)$, for each σ we replace the leaf node, v_σ, by the subtree $T_{(D\times R)\restriction_\sigma}(F\restriction_\sigma)$. (Note that

for the unique element $\sigma \in \mathcal{M}_{D\times R}(K)$ which satisfies C_1 the leaf label of v_σ will be 1. For all other choices of σ, $T_{(D\times R)\restriction_\sigma}(F\restriction_\sigma) = T_{(D\times R)\restriction_\sigma}(F'\restriction_\sigma)$.)

$T_{D\times R}(F)$ clearly represents F over $D \times R$. We'll show that for appropriately chosen restriction ρ the height of $T_{D\times R}(F\restriction_\rho)$, $|T_{D\times R}(F\restriction_\rho)|$, is small with high probability. This lemma is a switching lemma in the spirit of [12] because it will allow us to obtain a disjunction that approximates the negation of F by representing F by a matching decision tree T and then taking $Disj(T^c)$.

Lemma 7: Let F be an r-disjunction over $D \times R$ with $|R| = n$ and $|D| = n+m$. If $s \geq 0$ and $10m \leq \ell \leq (n/r)^{1/2}/10$ then

$$\frac{|\{\rho \in \mathcal{M}^{\ell}_{D\times R} \ : \ |T_{(D\times R)\restriction_\rho}(F\restriction_\rho)| \geq s\}|}{|\mathcal{M}^{\ell}_{D\times R}|} \leq (1.5\ell^2\sqrt{r/n})^s.$$

Proof We only need to consider $s > 0$. Let $S \in \mathcal{M}^{\ell}_{D\times R}$ be the set of restrictions ρ such that $|T_{(D\times R)\restriction_\rho}(F\restriction_\rho)| \geq s$. As in [4] we obtain a bound on $|S|/|\mathcal{M}^{\ell}_{D\times R}|$ by defining a 1-1 map from S to a small set.

Let $stars(r,s)$ to be the set of all sequences $\beta = (\beta_1,\ldots,\beta_k)$ such that for each j, $\beta_j \in \{*,-\}^r \setminus \{-\}^r$ and such that the total number of *'s in all the β_j is s. We will define a 1-1 map

$$S \quad \rightarrow \bigcup_{s/2\leq j\leq s} \mathcal{M}^{\ell-j}_{D\times R} \times stars(r,j) \times [1,\ell+m]^s$$

The map Let $F = C_1 \vee C_2 \vee \ldots$. Suppose that $\rho \in S$ and let π be the partial matching labelling the lexicographically first path in $T_{(D\times R)\restriction_\rho}(F\restriction_\rho)$ that has length $\geq s$. Trim the last few edges of π along the path from the root so that $|\pi| = s$. We use the formula F and π to determine the image of ρ. Let C_{ν_1} be the first term of F that is not set to 0 by ρ. Then $C_{\nu_1}\restriction_\rho$ will be the first term in $F\restriction_\rho$. Since $|\pi| > 0$, such a term must exist and is not the empty term. Let $K = v(C_{\nu_1}\restriction_\rho)$ and let σ_1 be the unique partial matching in $\mathcal{M}_{(D\times R)\restriction_\rho}(K)$ that satisfies $C_{\nu_1}\restriction_\rho$. Let π_1 be the portion of π that touches K. We have two cases based on whether or not $\pi_1 = \pi$.

1: If $\pi_1 \neq \pi$ then by the construction of π, $\pi_1 \in \mathcal{M}_{(D\times R)\restriction_\rho}(K)$. Note also that $C_{\nu_1}\restriction_{\rho\sigma_1} = 1$ but since $\pi_1 \neq \pi$, $\pi_1 \neq \sigma_1$, and thus $C_{\nu_1}\restriction_{\rho\pi_1} = 0$.

2: If $\pi_1 = \pi$ then it is possible that $v(\pi)$ does not contain all of K. In this case we shorten σ_1 so that it is the unique element of $\mathcal{M}_{(D\times R)\restriction_\rho}(K')$ that does not falsify $C_{\nu_1}\restriction_\rho$ where $K' = v(\pi_1) \cap K$.

Note that in either case $|\pi_1| \leq 2|\sigma_1|$.

We define β_1 to be a vector of length r based on the fixed ordering of the variables in term f_{ν_1}. The j-th component of β_1 is $*$ if and only if the j-th variable in C_{ν_1} is in $v(\sigma_1)$. Note that since $C_{\nu_1} \restriction_\rho$ is not the empty term then there is at least one $*$ in β_1. From C_{ν_1} and β_1 we can reconstruct σ_1.

Now, by the definition of $T_{(D\times R)\restriction_\rho}(F \restriction_\rho)$, $\pi \setminus \pi_1$ labels a path in the canonical tree $T_{(D\times R)\restriction_{\rho\pi_1}}(F \restriction_{\rho\pi_1})$. If $\pi_1 \neq \pi$, we repeat the above argument, with $\pi \setminus \pi_1$ in place of π, $\rho\pi_1$ in place of ρ and find a term C_{ν_2} which is the first term of F not set to 0 by $\rho\pi_1$. Based on this we generate π_2, σ_2, β_2, as before. We repeat this process until the round k in which $\pi_1\pi_2...\pi_k = \pi$.

For each i, π_i matches all elements of $v(\sigma_i)$, so the $\sigma_1, \ldots, \sigma_k$ are mutually compatible and thus $\sigma_1...\sigma_k = \sigma_1 \cup \cdots \cup \sigma_k$. The image of ρ under the 1-1 map we define is a triple, $\langle \rho\sigma_1...\sigma_k, (\beta_1, ..., \beta_k), \delta \rangle$ where δ is defined below. Let $\sigma = \sigma_1...\sigma_k$ and $j = |\sigma|$. Clearly $\rho\sigma = \rho\sigma_1...\sigma_k \in \mathcal{M}^{\ell-j}_{D\times R}$ and $(\beta_1, ..., \beta_k) \in stars(r, j)$.

We now define the information δ. This will encode the relationships between all the σ_i and π_i. The set $v(\pi_i)$ contains $v(\sigma_i)$ possibly together with some nodes unset by $\rho\sigma$, each of which must be connected by π_i to some element of $v(\sigma_i)$. We list the edges of π_i using the total order induced on $v(\sigma_i)$ by the order on the elements of $D \cup R$. For each vertex of dom(σ_i) in order, we list the name of the other node to which it is matched in pi_i. This endpoint can be one of $|\sigma_i| + \ell - j \leq \ell$ possibilities (all of which are known) so a number between 1 and ℓ is sufficient to encode this using the order induced on these vertices. After this, for each vertex of range(σ_i) in order, that is not matched so far, we can similarly give a number between 1 and $\ell - j + m \leq \ell + m$ to indicate its mate in D. The information δ is then simply the vector of these numbers, one per edge of π and is thus contained in $[1, \ell + m]^s$. Thus the image of the map is as required.

Inverting the map It remains to show that the map we have just defined is indeed 1-1. To do this we show how to recover ρ from its image. The reconstruction is iterative. In the general stage of the reconstruction we will have recovered $\pi_1, ..., \pi_{i-1}$, $\sigma_1, ..., \sigma_{i-1}$, and will have constructed $\rho\pi_1...\pi_{i-1}\sigma_i...\sigma_k$. Recall that for $i < k$, $C_{\nu_i} \restriction_{\rho\pi_1...\pi_{i-1}\sigma_i} = 1$ and $C_j \restriction_{\rho\pi_1...\pi_{i-1}\sigma_i} = 0$ for all $j < \nu_i$. This clearly also holds when we append $\sigma_{i+1}...\sigma_k$ to the restriction. When $i = k$, something similar occurs except the only guarantee is that $C_{nu_i} \restriction_{\rho\pi_1...\pi_{k-1}\sigma_k} \neq 0$. Thus we can recover ν_i as the index of the first term of F that is not set to 0 by $\rho\pi_1...\pi_{i-1}\sigma_i...\sigma_k$.

Now, based on C_{ν_i} and β_i we can determine σ_i. Since we know $\sigma_1, ..., \sigma_i$ we can examine the entries in the vector δ associated with each of the vertices in $v(\sigma_i)$. At this point, although $\sigma_{i+1}, ..., \sigma_k$ are still undetermined, π_i can still be determined since π_i does not touch any of the vertices these restrictions touch.

We can now change $\rho\pi_1...\pi_{i-1}\sigma_i...\sigma_k$ to $\rho\pi_1...\pi_{i-1}\pi_i\sigma_{i+1}...\sigma_k$ using the knowledge of π_i

and σ_i. Finally, given all the values of the π_i we can reconstruct ρ.

The numbers Now we compute the value $|S|/|\mathcal{M}^{\ell}_{D\times R}|$. We can describe an element of $\mathcal{M}^{\ell}_{D\times R}$ by choosing ℓ elements of R and then, for each of the $n-\ell$ remaining vertices in turn, choosing an element of D with which it is to be matched. Thus $|\mathcal{M}^{\ell}_{D\times R}| = \binom{n}{\ell}(n+m)^{(n-\ell)} = \frac{n^{(\ell)}(n+m)^{(n-\ell)}}{\ell!}$ and

$$
\begin{aligned}
\frac{|\mathcal{M}^{\ell-j}_{D\times R}|}{|\mathcal{M}^{\ell}_{D\times R}|} &= \frac{n^{(\ell)}(n+m)^{(n-\ell)}(\ell-j)!}{(n-j)^{(\ell-j)}(n+m)^{(n-\ell+j)}\ell!} \\
&= \frac{(\ell+m)^{(j)}\ell^{(j)}}{(n-\ell)^{(j)}} \\
&\le \left(\frac{(\ell+m)\ell}{n-\ell}\right)^j
\end{aligned}
$$

There is an easy bound of $|stars(r,s)| \le 2^{s-1}r^s$ but we can also prove:

CLAIM: $|stars(r,s)| < (r/\ln 2)^s$.

For convenience in the proof we shall include the empty string in $stars(r,0)$ which would otherwise be empty. We shall show by induction on s that $|stars(r,s)| \le \gamma^s$ for $(1+1/\gamma)^r = 2$; the statement of the lemma follows by using $1+x < e^x$ for $x \ne 0$.

The base case $s = 0$ follows trivially. Now suppose that $s > 0$. It is easy to see from the definition that for any $\beta \in stars(r,s)$, if β_1 has $i \le s$ *'s then $\beta = (\beta_1, \beta')$ where $\beta' \in stars(r, s-i)$. (For $i = s$ we have used our augmentation of $stars(r,0)$.) There are $\binom{r}{i}$ choices of β_1 so

$$
\begin{aligned}
|stars(r,s)| &= \sum_{i=1}^{\min(r,s)} \binom{r}{i} |stars(r,s-i)| \\
&\le \sum_{i=1}^{r} \binom{r}{i} \gamma^{s-i} \\
&= \gamma^s \sum_{i=1}^{r} \binom{r}{i} (1/\gamma)^i \\
&= \gamma^s[(1+1/\gamma)^r - 1] \\
&= \gamma^s
\end{aligned}
$$

by the inductive hypothesis and the definition of γ. Thus the claim is proved.

Now applying our bounds we obtain

$$\begin{aligned}\frac{|S|}{|\mathcal{M}^{\ell}_{D\times R}|} &\leq \sum_{s/2\leq j} \frac{|\mathcal{M}^{\ell-j}_{D\times R}|}{|\mathcal{M}^{\ell}_{D\times R}|} \cdot |stars(r,j)| \cdot (\ell+m)^s \\ &\leq (\ell+1)^s \sum_{j\geq s/2} \left(\frac{r\ell(\ell+m)}{(n-\ell)\ln 2}\right)^j\end{aligned}$$

Since $1/\ln 2 < 1.4427$ and $10m \leq \ell \leq \sqrt{n/r}/10$, we have $(\ell+m) \leq 1.1\ell$ and $n \leq 1.02(n-\ell)$. Thus the series is at most

$$(1.1\ell)^s \sum_{j\geq s/2} (1.7r\ell^2/n)^j.$$

This is a geometric series with ratio $< .02$. Therefore it is at most

$$\begin{aligned}&1.03(1.1\ell)^s(1.7r\ell^2/n)^{s/2} \\ \leq\ &(1.5\ell^2\sqrt{r/n})^s.\end{aligned}$$

□

6 Generic systems and Exceptional Forests

Suppose that we have a k-evaluation $\mathbf{T}$ of the subformulas in a Frege proof $\mathcal{P}$ of *onto-PHP*^D_R with $Count_p$ axiom schemas. By Lemmas 2 and 3, there is some instance F of a $Count_p$ axiom in $\mathcal{P}$ that does not evaluate to true under $\mathbf{T}$. Therefore there is some $\pi \in \mathrm{Br}_0(T_F)$.

By Lemma 4, the map, $\mathbf{T}'$, given by

$$T'_A = \begin{cases} T_A\restriction_\pi & \text{if } \mathrm{Br}_1(T_A) \neq \emptyset \\ T_0 & \text{otherwise}\end{cases}$$

is also a k-evaluation of the formulas in F over $(D\times R)\restriction_\pi$ and T'_F is "false".

Let F_e for $e \in M^{(p)}$, $|M| \not\equiv 0 \pmod{p}$ be the formulas that substitute for Y_e in F and let $T_e = T'_{F_e}$. Using the T_e, we will see that if they 'locally' appear to define something that is a p-partition of M then $T'_F \equiv 1$. Then we show that it is impossible for the T_e to describe something that locally does appear to be a partition of M into blocks of size p. This latter is achieved by a reduction to a problem over polynomials.

Lemma 8: Suppose that some $\pi \in \mathrm{Br}_0(T_F)$ exists and $3k \leq |R| < |D|$.

(a) If $e, e' \in M^{(p)}$ with $e \perp e'$ there are no compatible branches $\sigma_e \in \mathrm{Br}_1(T_e)$ and $\sigma_{e'} \in \mathrm{Br}_1(T_{e'})$.

(b) For any restriction τ such that $|\tau| + k \leq |R|$, τ is compatible with some element of $\bigcup_{e \in M^{(p)}} \mathrm{Br}_1(T_e)$.

Proof For part (a), suppose that there are compatible branches $\sigma_e \in \mathrm{Br}_1(T_e)$ and $\sigma_{e'} \in \mathrm{Br}_1(T_{e'})$. Let $\sigma = \sigma_e \cup \sigma_{e'}$ and apply it to all formulas in the k-evaluation $\mathbf{T}'$. We see that it will make $T'_{\neg F_e \vee \neg F_{e'}} \restriction_\sigma \equiv 0$. Thus $T'_{\neg(\neg F_e \vee \neg F_{e'})} \restriction_\sigma \equiv 1$ so $T'_F \restriction_\sigma = T'_F \equiv 1$ which is a contradiction.

For part (b), apply τ to all formulas in the k-evaluation $\mathbf{T}'$. If τ is incompatible with all elements of $\bigcup_{e \in M^{(p)}} \mathrm{Br}_1(T_e)$. then $T'_{\bigvee_{e \in M^{(p)}} F_e} \restriction_\tau \equiv 0$ so $T'_{\neg \bigvee_{e \in M^{(p)}} F_e} \restriction_\tau \equiv 1$ and thus $T'_F \restriction_\tau = T'_F \equiv 1$ which is a contradiction. $\Box$

Now for each $v \in M$ let $\mathcal{B}_v = \bigcup_{v \in e} \mathrm{Br}_1(T_e)$. Lemma 8 implies that any $\mathcal{B}_v$ consists of mutually incompatible elements and it contains an element compatible with any fixed τ with $|\tau| \leq |R| - k$. In the terminology of [18] this is a $(|R| - k)$-*basis of height* $\leq k$.

Definition 6.1: A (p, M)-*generic system of height* h *over* $D \times R$ is a collection of matching decision trees over $D \times R$: T_v, $v \in M$, with leaf labels that are p-subsets of M such that:
(1) each T_v has height at most h;
(2) each branch in T_v with leaf label e has $v \in e$;
(3) for all $e \in M^{(p)}$, for all $v, w \in e$, $\mathrm{Br}_e(T_v) = \mathrm{Br}_e(T_w)$.

Lemma 9: If F is an instance of a $Count_p^M$ axiom schema and there is some $\pi \in \mathrm{Br}_0(T_F)$ then if $n' = |R| < |D|$, $ph \leq N$, $N \leq \sqrt{(n'-k)/k}/10$, and $(1.5N^2\sqrt{k/(n'-k)})^h \leq 1/|M|$, there is restriction $\rho \in \mathcal{M}^N_{(D \times R)\restriction_\pi}$ such that there is a (p, M)-generic system over $(D \times R)\restriction_{\pi\rho}$ of height at most ph.

Proof For each $v \in M$, let $\mathcal{B}_v$ be as above and define

$$G_v = \bigvee_{\sigma \in \mathcal{B}_v} \bigwedge_{\langle i,j \rangle \in \sigma} P_{ij}.$$

By Lemma 7, for a ρ chosen uniformly at random from $\mathcal{M}^N_{(D \times R)\restriction_\pi}$, the probability that $G_v \restriction_\rho$ fails to have a canonical matching decision tree of height at most h is less than $1/|M|$. Therefore the probability that a ρ fails to do this for all $v \in M$ is less than one. Choose some ρ that achieves this for all $v \in M$ and let T_v^* be the tree associated with $G_v \restriction_\rho$.

By Lemma 8, if τ is a branch of T_v^* then $\rho\tau$ is compatible with some $\sigma \in \mathcal{B}_v$, i.e. some $\sigma \in \mathrm{Br}_1(T_e)$ for some e with $v \in e$, and thus the leaf label of τ must be 1. Therefore, since T_v^* is a canonical decision tree for $G_v \restriction_\rho$, σ must be contained in $\rho\tau$. Since the elements of $\mathcal{B}_v$ are mutually incompatible, the choice of σ must be unique. Therefore, each leaf of T_v^* is associated with a unique $e \in M^{(p)}$ with $v \in e$. Relabel the leaves of the T_v^* by their associated $e \in M^{(p)}$. Lemma 8 implies that for any $v, v' \in M$, if τ and τ' are compatible branches in T_v^* and $T_{v'}^*$ then their leaf labels e, e' are compatible.

In order to create the trees T_v, for each branch σ of T_v^* with leaf label e, extend σ in T_v^* by appending trees $T_w^* \upharpoonright_\sigma$ for each $w \in e$ in turn and labelling all leaves of the resulting branches by e. This at most multiplies the height of the trees by p. Observe that for $e = \{v_1, \ldots, v_p\}$ the branches with leaf label e in T_{v_j} are all elements of the form

$$\{\pi_1 \cdots \pi_p \mid \pi_i \in \mathrm{Br}_e(T_{v_i}^*) \text{ for } i = 1, \ldots, p\}$$

and are thus independent of the choice of j. $\square$

The above construction of a generic system is essentially from [6]. Alternatively, one could create the trees T_v^* by constructing the canonical decision trees for the G_v without applying any restriction. Using an argument like the one showing that any Boolean formula that has CNF clause size c and DNF term size d has Boolean decision tree of height $\leq cd$, one can show that the canonical decision tree constructed for G_v has height at most $k(k+1)$. This latter approach is very much like the one in [17].

The key property that we use about any (p, M)-generic system of height h over $D \times R$ is that it is a forest $\mathcal{T}$ of matching decision trees over $D \times R$ of height $\leq h$ such that each branch appears $0 \pmod p$ times in $\mathcal{T}$ and such that the total number of trees is $\not\equiv 0 \pmod p$. In the terminology of [16] this is a *p-exceptional forest of (D, R)-labelled trees.*

7 Nullstellensatz Proofs

DEFINITION 7.1: Given multivariate polynomials

$$Q_1(\vec{x}), \ldots, Q_m(\vec{x}) \in R[x_1, \ldots, x_n]$$

there is no solution to

$$\begin{aligned} Q_1(\vec{x}) &= 0 \\ \cdots &= \cdot \\ Q_m(\vec{x}) &= 0 \end{aligned}$$

over $\{0,1\}$ if $\exists P_1(\vec{x}), \ldots, P_m(\vec{x}) \in R[x_1, \ldots, x_n]$ such that $\sum_{i=1}^m P_i(\vec{x}) \cdot Q_i(\vec{x}) \equiv r \not\equiv 0 \pmod p$ in $R[x_1, \ldots, x_n]/(x_1^2 - x_1, \cdots, x_n^2 - x_n)$. We say that $P_1, \ldots, P_m$ are a *Nullstellensatz r-refutation* of $\{Q_1, \ldots, Q_m\}$. (We drop the r when $r = 1$. If p is a prime then the exact value of r is irrelevant. Also, if p is prime then a Nullstellensatz refutation is guaranteed to exist by Hilbert's Nullstellensatz whenever there is no $\{0,1\}$ solution.) The *degree* of the r-refutation is the maximum degree of the P_i

Definition 7.2: Let *onto-*$\mathcal{PHP}^D_R$ be the following system of polynomial equations in variables $x_{i,j}$ with $i \in D$, $j \in R$:

(1) $Q^D_i(\vec{x}) = (\sum_{j\in R} x_{i,j}) - 1 = 0$ one for each $i \in D$, and

(2) $Q^R_j(\vec{x}) = (\sum_{i\in D} x_{i,j}) - 1 = 0$ one for each $j \in R$, and

(3) $Q_{i,jk}(\vec{x}) = x_{i,j} \cdot x_{i,k} = 0$ one for each $i \in D$, $j, k \in R$, $j \neq k$, and

(4) $Q_{ij,k}(\vec{x}) = x_{i,k} \cdot x_{j,k} = 0$ one for each $i \neq j$, $i, j \in D$, $j \in R$.

Again we use *onto-*$\mathcal{PHP}^{|D|}_{|R|}$ to emphasize that the sizes of D and R are all that matter.

Definition 7.3: We relate monomials and sets of edges in $D \times R$ as follows: Given a set of edges $\pi \in D \times R$, define $X_\pi = \prod_{\langle i,j\rangle\in\pi} x_{i,j}$ and given a monomial $X = x^{e_1}_{i_1,j_1} \cdots x^{e_k}_{i_k,j_k}$ with $e_1, \ldots, e_k \geq 1$, define $\pi_X = \{\langle i_1, j_1\rangle, \ldots, \langle i_k, j_k\rangle\}$.

Lemma 10: If $|M| \equiv r \pmod{p}$ and a (p, M)-generic system of height h over $D \times R$ exists then there is a Nullstellensatz r-refutation of *onto-*$\mathcal{PHP}^D_R$ of degree at most $h - 1$ over $\mathbb{Z}_p$.

The basic idea is to consider the polynomial whose monomials are the products of the variables associated with each branch of the trees in the generic system. That is, with each tree T_v we get a polynomial

$$P_{T_v} = \sum_{\pi\in\mathrm{Br}(T_v)} X_\pi.$$

We first show that each P_{T_v} is $1 + L_v$ where L_v a linear combination of the Q polynomials of degree at most $h - 1$.

Lemma 11: Let T be a matching decision tree over $D \times R$. Then $P_T = \sum_{\pi\in\mathrm{Br}(T)} \prod_{e\in\pi} x_e$ is of the form $1 + L$ where L is a linear combination of the *onto-*$\mathcal{PHP}^D_R$ polynomials with coefficient polynomials of degree $\leq h - 1$.

Proof Proof by induction on the number of internal vertices of T.

Base Case: If T has no internal vertices then it has one branch of height 0, $P_T(\vec{x}) = 1$ and all coefficient polynomials are 0 which gives degree -1.

Induction Step: Suppose that T has at least one internal vertex and has height h. Then it has one such vertex v all of whose children are leaves. Let T' be the matching decision tree obtained by removing all the children of v. Let π be the matching given along the path from the root to v.

If the query at v is $i \in D$, then

$$\begin{aligned} P_T(\vec{x}) &= P_{T'}(\vec{x}) + X_\pi - \sum_{j\in R\setminus \mathrm{range}(\pi)} X_\pi \cdot x_{i,j} \\ &= P_{T'}(\vec{x}) + X_\pi \cdot (1 - \sum_{j\in R\setminus \mathrm{range}(\pi)} x_{i,j}) \\ &= P_{T'}(\vec{x}) + X_\pi \cdot (1 - \sum_{j\in R} x_{i,j}) + X_\pi \cdot \sum_{k\in \mathrm{range}(\pi)} x_{i,k} \\ &= P_{T'}(\vec{x}) - X_\pi \cdot Q_i^D + X_\pi \cdot \sum_{k\in \mathrm{range}(\pi)} x_{i,k}. \end{aligned}$$

X_π has degree at most $h-1$, the last term is a degree $h-2$ combination of the $Q_{ij,k}$, and applying the induction hypothesis to $P_{T'}$ yields the desired result.

The case when the query is $j \in R$ is analogous. $\Box$

Proof of Lemma 10 Consider $\sum_{v\in M} P_{T_v}$ in $\mathbb{Z}_p$. On the one hand it is

$$\sum_{v\in M} (L_v + 1) = r + \sum_{v\in M} L_v.$$

On the other hand, every branch in the generic system appears some multiple of p times. Therefore over $\mathbb{Z}_p$,

$$\sum_{v\in M} P_{T_v} = 0.$$

We derive $r + \sum_{v\in M} L_v = 0$ and obtain the Nullstellensatz refutation by reversing signs. $\Box$

8 A Nullstellensatz degree lower bound for *onto*-$\mathcal{PHP}_N^{N+p^\ell}$

In this section we prove the following theorem which is of independent interest.

Theorem 12: Let $r \not\equiv 0(\bmod p)$. If $N \geq ((p+2)^\ell - p^\ell)/2$ then any Nullstellensatz r-refutation of *onto*-$\mathcal{PHP}_N^{N+p^\ell}$ over $\mathbb{Z}_p$ must have degree at least $2^\ell - 1$.

Definition 8.1: A *d-design for $D \times R$* is a mapping $\mathcal{D}$ from the partial matchings of size $\leq d$ on $D \times R$ into $\mathbb{Z}_p$ such that

(a) $\mathcal{D}(\emptyset) = 1$ for the empty matching $\emptyset$,

(b) For each partial matching π with $|\pi| < d$ and $i \in D \setminus \mathrm{dom}(\pi)$

$$\sum_{j \in R \setminus \mathrm{range}(\pi)} \mathcal{D}(\pi \cup \langle i, j \rangle) \equiv \mathcal{D}(\pi) \bmod p$$

(c) For each partial matching π with $|\pi| < d$ and $j \in R \setminus \mathrm{range}(\pi)$

$$\sum_{i \in D \setminus \mathrm{dom}(\pi)} \mathcal{D}(\pi \cup \langle i, j \rangle) \equiv \mathcal{D}(\pi) \bmod p$$

Lemma 13: Let $r \not\equiv 0(\bmod p)$. If there is a d-design for $D \times R$ then any r-refutation of *onto-*$\mathcal{PHP}_R^D$ over $\mathbb{Z}_p$ requires degree at least d.

Proof We extend the d-design $\mathcal{D}$ to be a function from the set of polynomials to $\mathbb{Z}_p$. For any monomial X in variables $x_{i,j}$ with $i \in D$ and $j \in R$ define

$$\mathcal{D}(X) = \begin{cases} \mathcal{D}(\pi_X) & \text{if } \pi_X \text{ is a matching with } |\pi_X| \le d \\ 0 & \text{otherwise} \end{cases}$$

and extend $\mathcal{D}$ linearly over $\mathbb{Z}_p$ to a map $\mathcal{D} : \mathbb{Z}_p[\vec{x}] \to \mathbb{Z}_p$ by setting $\mathcal{D}(P_1 + P_2) = \mathcal{D}(P_1) + \mathcal{D}(P_2)$ for $P_1, P_2 \in \mathbb{Z}_p[\vec{x}]$ and $\mathcal{D}(aP) = a\mathcal{D}(P)$ for $a \in \mathbb{Z}_p$ and $P \in \mathbb{Z}_p[\vec{x}]$.

Clearly $\mathcal{D}(1) = \mathcal{D}(\emptyset) = 1$ by part (a) of the design definition. We consider the polynomials in the definition of *onto-*$\mathcal{PHP}_R^D$ and show that for any $P \in \mathbb{Z}_p[\vec{x}]$ of degree $< d$,

$$\mathcal{D}(P \cdot Q_i^D) = \mathcal{D}(P \cdot Q_j^R) = \mathcal{D}(P \cdot Q_{i,jk}) = \mathcal{D}(P \cdot Q_{ij,k}) = \mathcal{D}(P \cdot (x_{i,j}^2 - x_{i,j})) = 0. \qquad (*)$$

We see that $(*)$ is sufficient by observing that it implies if $0 \neq r = \Sigma_i P_i Q_i$ is an r-refutation of *onto-*$\mathcal{PHP}_R^D$ over $\mathbb{Z}_p$ of degree $< d$ then $0 \neq r = \mathcal{D}(r) = \mathcal{D}(\Sigma_i P_i Q_i) = \Sigma_i \mathcal{D}(P_i Q_i) = 0$ which is a contradiction.

To prove $(*)$, by the linearity of $\mathcal{D}$ it clearly suffices to prove it when P is simply a monomial X of degree $< d$. Furthermore, if μ_X is not a partial matching then $\mathcal{D}(X) = 0$ so, by the linearity of $\mathcal{D}$, we can assume that μ_X is a partial matching.

Since $Q_{i,jk}$ and $Q_{ij,k}$ are monomials and both $\mu_{Q_{i,jk}}$ and $\mu_{Q_{ij,k}}$ are not partial matchings we immediately have $\mathcal{D}(X \cdot Q_{i,jk}) = \mathcal{D}(X \cdot Q_{ij,k}) = 0$ for any monomial X.

Also, since $\mu_{X \cdot x_{i,j}^2} = \mu_{X \cdot x_{i,j}}$, the linearity of $\mathcal{D}$ implies that $\mathcal{D}(X \cdot (x_{i,j}^2 - x_{i,j})) = 0$.

For $Q_i^D(\vec{x}) = \sum_{j \in R} x_{i,j} - 1 = 0$ we have two cases depending on whether or not $i \in \mathrm{dom}(\pi_X)$. If $i \notin \mathrm{dom}(\pi_X)$ then

$$\begin{aligned} \mathcal{D}(X \cdot Q_i^D(\vec{x})) &= \mathcal{D}(X \cdot (\sum_{j \in R} x_{i,j} - 1)) \\ &= \sum_{j \in R} \mathcal{D}(X \cdot x_{i,j}) - \mathcal{D}(X) \end{aligned}$$

$$
\begin{aligned}
&= \sum_{j\in R\setminus \mathrm{range}(\pi_X)} \mathcal{D}(X\cdot x_{i,j}) \ - \mathcal{D}(X) \\
&= \sum_{j\in R\setminus \mathrm{range}(\pi_X)} \mathcal{D}(\mu_X \cup \langle i,j\rangle) \ - \mathcal{D}(\mu_X) \\
&= 0
\end{aligned}
$$

over $\mathbb{Z}_p$ by part (b) of the definition of a d-design since $|\pi_X| < d$. If $i \in \mathrm{dom}(\pi_X)$ then let $\langle i, j^*\rangle \in \pi_X$. In this case

$$
\begin{aligned}
\mathcal{D}(X\cdot Q_i^D(\vec{x})) &= \mathcal{D}(X\cdot(\sum_{j\in R} x_{i,j} \ - 1)) \\
&= \sum_{j\in R} \mathcal{D}(X\cdot x_{i,j}) \ - \mathcal{D}(X) \\
&= \mathcal{D}(X\cdot x_{i,j^*}) \ - \mathcal{D}(X) \\
&= \mathcal{D}(\mu_X \cup \langle i,j^*\rangle) \ - \mathcal{D}(\mu_X) \\
&= 0
\end{aligned}
$$

since $\mu_{X\cdot x_{i,j}}$ is not a matching for $j \neq j^*$ and $\mu_X \cup \langle i,j^*\rangle = \mu_X$.

The result for $Q_j^R(\vec{x})$ follows similarly using part (c) of the definition of a d-design. $\square$

Lemma 14: If there is a d-design $\mathcal{D}$ for $D\times R$ over $\mathbb{Z}_p$ then there is a $2d+1$-design $\mathcal{D}'$ for $D'\times R'$ over $\mathbb{Z}_p$ where $|D'| = (p+1)|D| + |R|$ and $|R'| = |D| + (p+1)|R|$. (Observe that $|D'| - |R'| = p(|D| - |R|)$.)

Before we prove Lemma 14, we show how it implies Theorem 12.

Proof of Theorem 12 For $\ell \geq 0$ let $N_\ell = ((p+2)^\ell - p^\ell)/2$. We show by induction that there is a $2^\ell - 1$-design for $[1, N_\ell + p^\ell] \times [1, N_\ell]$ over $\mathbb{Z}_p$. The theorem then will follow by Lemma 13.

For $\ell = 0$, letting $\mathcal{D}(\emptyset) = 1$ is sufficient to satisfy the conditions for a 0-design.

Suppose we have a $2^\ell - 1$-design $\mathcal{D}$ for $[1, N_\ell + p^\ell] \times [1, N_\ell]$. Observe that $N_{\ell+1} + p^{\ell+1} = (p+1)(N_\ell + p^\ell) + N_\ell$ and $N_{\ell+1} = (N_\ell + p^\ell) + (p+1)N^\ell$. Applying Lemma 14 we get a $2(2^\ell - 1) + 1 = 2^{\ell+1} - 1$-design $\mathcal{D}'$ for $[1, N_{\ell+1} + p^{\ell+1}] \times [1, N_{\ell+1}]$ over $\mathbb{Z}_p$ as required. $\square$

Proof of Lemma 14 Let $\mathcal{D}$ be a d-design for $D\times R$ over $\mathbb{Z}_p$. Let $D' = \{i_1, \ldots, i_{|D'|}\}$ and $R' = \{j_1, \ldots, j_{|R'|}\}$. Divide D' into $|R|$ individual points $i_1, \ldots, i_{|R|}$ and $|D|$ blocks $D_1, \ldots D_{|D|}$ each of size $p+1$ and divide R' into $|D|$ individual points $j_1, \ldots, j_{|D|}$ and $|R|$ blocks, $R_1, \ldots, R_{|R|}$ each of size $p+1$. Following [10], we also fix a cyclic ordering on the elements within each block, e.g. as a permutation $\sigma : D' \cup R' \to D' \cup R'$ which maps each of the individual points to itself and whose other orbits are the blocks of size $p+1$. We say that $\langle i,j\rangle$ is *parallel* to $\langle i',j'\rangle$ iff there is some r such that $i'=\sigma^r(i)$ and $j'=\sigma^r(j)$. Observe that this forms an equivalence relation on edges.

In matchings on $D'\times R'$, we say that an edge is a *cross edge* if it is in $D_a\times R_b$ for some a and b and is a *rung* if it is in $b\times R_b$ or $D_a\times a$ for some a or b, i.e. it joins some individual point to its corresponding block. Given $\pi\subseteq D'\times R'$, let $Im(\pi)=\{\langle a,b\rangle\mid \pi\cap D_a\times R_b\neq 0\}$, i.e. $Im(\pi)$ is the projection of the cross edges in π onto $D\times R$.

Definition 8.2: For each choice, V, of a set of $|D|+|R|$ representative elements, $u_i\in D_i$ for $i=1,\ldots,|D|$ and $v_j\in R_j$ for $j=1,\ldots,|R|$ and matching π on $D'\times R'$, we say that π *respects* V if

(A) the only edges of π are rungs or cross edges,

(B) each rung in π matches a representative element given by V, i.e. is of the form $\langle u_i,j_i\rangle$ for $i\leq|D|$ or $\langle i_j,v_j\rangle$ for $j\leq|R|$,

(C) for any a and b, each cross edge of π in $D_a\times R_b$ is parallel to $\langle u_a,v_b\rangle$ but not equal to it.

For each V as above, we can define a map $\mathcal{D}^V$ from the set of partial matchings of size $\leq d$ on $D'\times R'$ to $\mathbb{Z}_p$.

$$\mathcal{D}^V(\pi)=\begin{cases}\mathcal{D}(Im(\pi)) & \text{if } \pi \text{ respects } V \text{ and } Im(\pi) \text{ is a matching of size} \leq d\\ 0 & \text{otherwise}\end{cases}.$$

Finally, we define $\mathcal{D}'(\pi)=\sum_V\mathcal{D}^V(\pi)$.

Claim: $\mathcal{D}'$ is a $2d+1$-design for $D'\times R'$ over $\mathbb{Z}_p$.

Clearly $\mathcal{D}'(\emptyset)=\sum_V\mathcal{D}^V(\emptyset)=(p+1)^{|D|+|R|}=1$ over $\mathbb{Z}_p$ since there are exactly $(p+1)^{|D|+|R|}$ different choices of V and for each of these $\mathcal{D}^V(\emptyset)=1$. Thus condition (a) for a design is satisfied.

We now show that condition (b) for a $2d+1$-design is satisfied. The proof for condition (c) is analogous. Let $|\pi|\leq 2d$ be a matching on $D'\times R'$ and $i'\in D'\setminus\mathrm{dom}(\pi)$.

We can assume without loss of generality that $Im(\pi)$ is a matching of size $\leq d$ since otherwise $\mathcal{D}'(\pi)=0$ and $\mathcal{D}'(\pi\cup\langle i',j'\rangle)=0$ for all $j'\in R'$.

Now

$$\begin{aligned}\sum_{j'\in R'\setminus\mathrm{range}(\pi)}\mathcal{D}'(\pi\cup\langle i',j'\rangle) &= \sum_{j'\in R'\setminus\mathrm{range}(\pi)}\sum_V\mathcal{D}^V(\pi\cup\langle i',j'\rangle)\\ &= \sum_V\sum_{j'\in R'\setminus\mathrm{range}(\pi)}\mathcal{D}^V(\pi\cup\langle i',j'\rangle).\end{aligned}$$

We have several cases:

CASE 1: if $i' = i_j$ for $1 \leq j \leq |R|$ then $Im(\pi \cup \langle i', j' \rangle) = Im(\pi)$ and, if $\pi \cup \langle i', j' \rangle$ respects V, it must be the case that $j' = v_j$. Thus

$$\begin{aligned}
\sum_{j' \in R' \setminus \mathrm{range}(\pi)} \mathcal{D}'(\pi \cup \langle i', j' \rangle) &= \sum_V \sum_{j' \in R' \setminus \mathrm{range}(\pi)} \mathcal{D}^V(\pi \cup \langle i', j' \rangle) \\
&= \sum_V \mathcal{D}^V(\pi \cup \langle i_j, v_j \rangle) \\
&= \sum_V \mathcal{D}^V(\pi) \\
&= \mathcal{D}'(\pi)
\end{aligned}$$

as required.

CASE 2: $i' \in D_a$ for some a. We split this case into several subcases based on the structure of π. Let $\mathcal{V}_{i'}$ be the set of those V such that $u_a \neq i'$. For each subcase we first observe that we only need to consider those $V \in \mathcal{V}_{i'}$ such that π respects V. If π does not respect V then $\pi \cup \langle i', j' \rangle$ does not respect V, so $\mathcal{D}^V(\pi) = \mathcal{D}^V(\pi \cup \langle i', j' \rangle) = 0$. If V has $u_a = i'$ then $j' = a$ is the only value such that V respects $\pi \cup \langle i', j' \rangle$ and for this value, $Im(\pi \cup \langle i', j' \rangle) = Im(\pi)$. Thus for each $V \notin \mathcal{V}_{i'}$,

$$\sum_{j' \in R' \setminus \mathrm{range}(\pi)} \mathcal{D}^V(\pi \cup \langle i', j' \rangle) = \mathcal{D}^V(\pi).$$

SUBCASE (a): π has a cross edge $\langle i^*, j^* \rangle$ with $i^* \in D_a$ and $j^* \in R_b$ for some b.
If π respects $V \in \mathcal{V}_{i'}$ then there is exactly one j' such that $Im(\pi \cup \langle i' j' \rangle)$ is a matching and $\pi \cup \langle i', j' \rangle$ respects V. This is the unique $j' \in R_b$ such that such that $\langle i', j' \rangle$ is parallel to $\langle i^*, j^* \rangle$, i.e. $i' = \sigma^r(i^*)$ and $j' = \sigma^r(j^*)$ for some r. For this value of j', $Im(\pi \cup \langle i', j' \rangle) = Im(\pi)$, so for each $V \in \mathcal{V}_{i'}$,

$$\sum_{j' \in R' \setminus \mathrm{range}(\pi)} \mathcal{D}^V(\pi \cup \langle i', j' \rangle) = \mathcal{D}^V(\pi).$$

SUBCASE (b): π has no cross edges touching D_a and $|Im(\pi)| < d$.
In this case, given $V \in \mathcal{V}_{i'}$ such that π respects V, for each $b \in R \setminus \mathrm{range}(Im(\pi))$ there is one choice of $j' \in D' \setminus \mathrm{range}(\pi)$ such that $Im(\pi \cup \langle i', j' \rangle)$ is a matching and $\pi \cup \langle i', j' \rangle$ respects V. This j' is the unique member of R_b such that $\langle i', j' \rangle$ is parallel to $\langle u_a, v_b \rangle$ and for this value, $Im(\pi \cup \langle i', j' \rangle) = Im(\pi) \cup \langle a, b \rangle$. Thus

$$\sum_{j' \in R' \setminus \mathrm{range}(\pi)} \mathcal{D}^V(\pi \cup \langle i', j' \rangle) = \sum_{b \in R \setminus \mathrm{range}(Im(\pi))} \mathcal{D}(Im(\pi) \cup \langle a, b \rangle)$$

$$
\begin{aligned}
&= \mathcal{D}(Im(\pi)) \\
&= \mathcal{D}^V(\pi)
\end{aligned}
$$

using the fact that $\mathcal{D}$ satisfies condition (b) for a d-design over $D \times R$.

SUBCASE (c): π has no cross edges or rung edges touching D_a and $|Im(\pi)| = d$. In this case, $|Im(\pi \cup \langle i', j' \rangle)| > d$ for any j' so $\mathcal{D}^V(Im(\pi \cup \langle i', j' \rangle)) = 0$ for all V. We show that the sum of $\mathcal{D}^V(\pi)$ for all $V \in \mathcal{V}_{i'}$ is also 0. We can group such V that π respects into equivalence classes based on their choices other than u_a. Observe that the choice of $u_a \neq i'$ does not affect the value of $\mathcal{D}^V(\pi)$ since π has no cross edges touching D_a. Within each equivalence class there are exactly p choices of $u_a \neq i'$, so for each such class C, $\sum_{V \in C} \mathcal{D}^V(\pi)$ is a multiple of p and thus equal to 0 in $\mathbb{Z}_p$. Therefore, the sum of $\mathcal{D}^V(\pi)$ for all $V \in \mathcal{V}_{i'}$ is 0, and thus $\sum_{V \in \mathcal{V}_{i'}} \sum_{j' \in R' \setminus \text{range}(\pi)} \mathcal{D}^V(\pi \cup \langle i', j' \rangle) = \sum_{V \in \mathcal{V}_{i'}} \mathcal{D}^V(\pi)$.

SUBCASE (d): π has a rung edge but no cross edges touching D_a and $|Im(\pi)| = d$. As in the previous case, $|Im(\pi \cup \langle i', j' \rangle)| > d$ for any j' such that $\pi \cup \langle i', j' \rangle$ is a matching so $\mathcal{D}^V(Im(\pi \cup \langle i', j' \rangle)) = 0$ for all V. Again we show that the sum of $\mathcal{D}^V(\pi)$ for all $V \in \mathcal{V}_{i'}$ is also 0. In this case, π has at least d cross edges and at most $2d$ total edges, one of which is a rung edge that does not touch the same block as any cross edge. Thus there is some cross edge of π, $\langle i^*, j^* \rangle \in D_e \times R_f$ for some e and f, such that no other edges of π touch D_e or R_f. We group all $V \in \mathcal{V}_{i'}$ that π respects into equivalence classes based on their choices of points other than u_e and v_f. Since π has no other edges touching D_e or R_f, the value of $\mathcal{D}^V(\pi)$ is the same for all V in each equivalence class. Within each equivalence class there are exactly p choices of V since π respects each V and there are exactly p choices of u_e and v_f such that $\langle u_e, v_f \rangle$ is parallel to $\langle i^*, j^* \rangle$ but $u_e \neq i^*$. Therefore for each class C', $\sum_{V \in C'} \mathcal{D}^V(\pi) = 0$ over $\mathbb{Z}_p$. It follows that the sum of $\mathcal{D}^V(\pi)$ for all such V is 0, and thus $\sum_{V \in \mathcal{V}_{i'}} \sum_{j' \in R' \setminus \text{range}(\pi)} \mathcal{D}^V(\pi \cup \langle i', j' \rangle) = \sum_{V \in \mathcal{V}_{i'}} \mathcal{D}^V(\pi)$.

Summarizing Case 2, we have $\sum_V \sum_{j' \in R' \setminus \text{range}(\pi)} \mathcal{D}^V(\pi \cup \langle i', j' \rangle) = \sum_V \mathcal{D}^V(\pi)$, i.e.

$$\sum_{j' \in R' \setminus \text{range}(\pi)} \mathcal{D}'(\pi \cup \langle i', j' \rangle) = \mathcal{D}'(\pi)$$

as required. □

When p is prime, using a construction from [16], we can see that the degree of a Nullstellensatz refutation is not too much larger than the above lower bound. To see this we first need the following:

Proposition 15: If p is prime and $m - p^\ell < p^\ell a \le m$ then $\binom{m}{p^\ell} \equiv a \pmod{p}$.

Proof For each r, $1 \leq r \leq p^\ell$, let $m(r)$ be the unique integer between $m - p^\ell + 1$ and m that is congruent to r modulo p^ℓ. Observe that

$$\binom{m}{p^\ell} \equiv \prod_{r=1}^{p^\ell} \frac{m(r)}{r} \; (\bmod\, p).$$

Since $m(p^\ell) = p^\ell a$, by assumption, $\frac{m(p^\ell)}{p^\ell} \equiv a \; (\bmod\, p)$.

For $1 \leq r < p^\ell$, write $r = p^k r'$ where $\gcd(p, r') = 1$ and $k < \ell$. Since $m(r) \equiv r \; (\bmod\, p^\ell)$, there is some m' such that $m(r) = p^\ell m' + r = p^k(p^{k-\ell} + r')$. Therefore $\frac{m(r}{r} = \frac{p^{k-\ell}+r'}{r'}$. Since $p^{k-\ell} + r' \equiv r' \; (\bmod\, p)$, we derive that $\frac{m(r)}{r} \equiv 1 \; (\bmod\, p)$ from which the proposition follows. $\square$

Lemma 16: If p is prime and $p^\ell \leq N$, there is a Nullstellensatz refutation of *onto-*$\mathcal{PHP}_N^{N+p^\ell}$ of degree $p^\ell - 1$.

Proof Let $D = [1, N + p^\ell]$ and $R = [1, N]$. Consider the polynomial

$$\sum_{A \subset D,\ |A| = p^\ell} \; \sum_{\pi,\ \mathrm{dom}(\pi) = A} X_\pi - \sum_{B \subset R,\ |B| = p^\ell} \; \sum_{\pi,\ \mathrm{range}(\pi) = B} X_\pi.$$

Since each X_π with $|\pi| = p^\ell$ appears exactly once in each sum, the value of the polynomial is 0. On the other hand, notice that $\sum_{\pi,\ \mathrm{dom}(\pi)=A} X_\pi$ is the polynomial P_T for a matching decision tree of height h that queries each element of A along each path. Therefore by Lemma 11, $\sum_{\pi,\ \mathrm{dom}(\pi)=A} X_\pi = 1 + L_A$ where L_A is a linear combination of the *onto-*$\mathcal{PHP}_R^D$ polynomials of degree $\leq p^\ell - 1$ over $\mathbb{Z}_p$. Similarly, $\sum_{\pi,\ \mathrm{range}(\pi)=B} X_\pi = 1 + L_B$ where L_B is a combination of degree $\leq p^\ell - 1$.

Therefore

$$\begin{aligned} 0 &= \sum_{A \subset D,\ |A| = p^\ell} \; \sum_{\pi,\ \mathrm{dom}(\pi) = A} X_\pi - \sum_{B \subset R,\ |B| = p^\ell} \; \sum_{\pi,\ \mathrm{range}(\pi) = B} X_\pi \\ &= \binom{N + p^\ell}{p^\ell} - \binom{N}{p^\ell} + L \end{aligned}$$

where L is a combination of the *onto-*$\mathcal{PHP}_R^D$ polynomials of degree at most $p^\ell - 1$. Since $\binom{N+p^\ell}{p^\ell} - \binom{N}{p^\ell} \equiv 1 \; (\bmod\, p)$ by Proposition 15, we obtain a Nullstellensatz refutation of *onto-*$\mathcal{PHP}_R^D$ of degree at most $p^\ell - 1$. $\square$

9 Putting it all Together

Theorem 17: For $\ell \leq \epsilon \log_2 n$ with $1/\epsilon = 3 \cdot 4^{d+1}(\frac{1}{2} + \log_2(p+2))$, any depth d proof of *onto-PHP*$_n^{n+p^\ell}$ in a Frege system augmented by $Count_p$ axiom schemas requires size at least $n^{2^\ell/(4^{d+1}p)}$.

Proof Suppose that ℓ satisfies the conditions of the statement and that $\mathcal{P}$ is a depth d Frege proof with $Count_p$ axioms of *onto-PHP*$_n^{n+p^\ell}$ of size $S < n^{2^\ell/(4^{d+1}p)}$.

Let $k = \log_2 S$, and $N = n^{1/(2\cdot 4^{d+1})}/\sqrt{k}$.

Since $\ell \leq \epsilon \log_2 n$,

$$\begin{aligned} 10(p+2)^\ell \sqrt{k} &\leq 10(p+2)^\ell 2^{\ell/2} \sqrt{\log_2 n} \\ &\leq 10 n^{\epsilon(\frac{1}{2}+\log_2(p+2))} \sqrt{\log n} \\ &\leq 10 n^{1/(3\cdot 4^{d+1})} \sqrt{\log_2 n} \\ &< n^{1/(2\cdot 4^{d+1})}. \end{aligned}$$

for n sufficiently large relative to d. Therefore $10(p+2)^\ell < N$.

Define $n_0, \ldots n_d$ as in the statement of Lemma 6. Then $n_d = n^{1/4^d}/(9k)^{\delta_d}$ where $\delta_d = \sum_{i=1}^d 4^{-i} < 1/3$ and thus $n_d \geq n^{1/4^d}/(9k)^{1/3} > N$. It follows that $n_d \geq k = \log_2 S$ and $n_d \geq 10p^\ell$.

Therefore by Lemma 6, there is a restriction $\rho \in \mathcal{M}_{D\times R}^{n_d}$ and a k-evaluation $\mathbf{T}$ of the set of subformulas of $\mathcal{P}\restriction_\rho$ over $(D \times R)\restriction_\rho = D' \times R'$ where $|D'| = |R'| + p^\ell = n_d + p^\ell$.

By Lemmas 2 and 3, there must be some instance F of a $Count_p^M$ axiom schema in $\mathcal{P}\restriction_\rho$ and $\pi \in \mathrm{Br}_0(T_F)$. We now let $h = 4^{d+1}\log_n S$. Observe that by assumption about S, $h < 2^\ell/p$ and that

$$\begin{aligned} (1.5N^2\sqrt{k/(n_d-k)})^h &< (3N^2\sqrt{k}(9k)^{1/6}/n^{1/(2\cdot 4^d)})^h \\ &< (n^{1/4^{d+1}-1/(2\cdot 4^d)})^h \\ &= n^{h/4^{d+1}} \leq 1/S \leq 1/|M| \end{aligned}$$

and apply Lemma 9 to obtain a (p,M)-generic system of height $ph < 2^\ell$ over $D'' \times R''$ where $|D''| = |R''| + p^\ell = N + p^\ell$. Applying Lemma 10, we obtain a Nullstellensatz $|M|$-refutation of *onto*-$\mathcal{PHP}_N^{N+p^\ell}$ of degree less than $2^\ell - 1$ which contradicts Theorem 12. $\square$

Riis [16], by considering all possible domain and range subsets of size p^ℓ, as in Lemma 16, has shown that one can prove *onto*-$\mathcal{PHP}_n^{n+p^\ell}$ from $Count_p$ using a constant-depth proof of size $n^{O(p^\ell)}$ so the above bound is relatively tight.

Corollary 18: Any depth d Frege proof of PHP_n^{n+1} requires size $n^{\Omega(n^{1/(30\cdot 4^d)})}$ even if axiom schemas for *onto-*PHP_n^{n+1} are permitted.

Proof Apply Theorem 17 with $p = 2$, $\ell = (\log_2 n)/(30 \cdot 4^d) - 1$, and $n' = n + 1 - p^\ell$. (It is not hard to check that the conditions hold.) This implies that any depth d Frege proof of *onto-*$PHP_{n'}^{n'+p^\ell}$ using axiom schemas for $Count_2$ requires size $n^{\Omega(n^{1/(30\cdot 4^d)})}$. Now it is easy to see that *onto-*PHP_n^{n+1} is an immediate consequence of $Count_2^{2n+1}$ so the same lower bound applies to the size of the proofs with *onto-*PHP_n^{n+1} schemas instead of $Count_2$ axiom schemas. Finally, observe that *onto-*$PHP_{n'}^{n'+p^\ell}$ is an immediate consequence of $PHP_{n'+p^\ell-1}^{n'+p^\ell}$, i.e. of PHP_n^{n+1}. $\square$

Corollary 19: [10] If p and q are positive integers such that q contains a prime factor not dividing p then any depth d Frege proof of $Count_q$ requires size $2^{n^{\Omega(1/4^d)}}$ even if axiom schemas for $Count_p$ are permitted.

More generally:

Corollary 20: If p and $q_1, \ldots, q_k$ are positive integers such that each q_i contains a prime factor not dividing p then any depth d Frege proof of $\bigvee_{i=1}^k Count_{q_i}$ requires size $2^{n^{\Omega(1/4^d)}}$ even if axiom schemas for $Count_p$ are permitted.

Proof If q_i contains a prime factor not dividing p then there is an easy proof of *onto-*$PHP_n^{n+p^\ell}$ from $Count_{q_i}$ by counting the number of edges touching the domain and range, respectively, and observing that these must be different modulo q_i. The implementation of this as a proof of size $(2n + p^\ell)^{O(q_i)}$ is quite straightforward. The $\Omega(\cdot)$ in the lower bound depends on the sizes of p and the q_i but does not depend on n or d. The overall argument is easily handled by cases. $\square$

Following standard connections between bounded-depth Frege systems and bounded arithmetic (see [14]) the results above also have implications for the relativized system of bounded arithmetic $S_2(R)$, defined by Buss [11], in which R is an uninterpreted function symbol. In general, lower bounds for $S_2(R)$ follow from $2^{(\log n)^{\omega(1)}}$ size lower bounds. If we let $PHP_*^{*+1}(R)$ (respectively *onto-*$PHP_*^{*+p^\ell}(R)$, $Count_p(R)$, etc.) denote the first-order version of the pigeonhole principle (etc.) for the relation R then the following are immediate corollaries of the above results.

Corollary 21: (1) Let $\ell(n)$ be an integer function of n such that $\ell(n) = \omega(\log\log n)$ and $\ell(n) = o(\log n)$. There is no proof of $PHP_*^{*+p^{\ell(*)}}(R)$ in $S_2(R) + Count_p(R)$.

(2) There is no proof of $PHP_*^{*+1}(R)$ in $S_2(R) + onto\text{-}PHP_*^{*+1}(R)$.

(3) [10] If q contains a prime factor not dividing p then there is no proof of $Count_q(R)$ in $S_2(R) + Count_p(R)$.

10 Remarks

It is interesting to compare the degree lower bound for the Nullstellensatz refutations of $onto\text{-}\mathcal{PHP}_N^{N+p^\ell}$ with the degree lower bound for $\mathcal{PHP}_N^{N+s}$ using the quite different construction in [5]. If we take $p = 2$ and $N = (4^\ell - 2^\ell)/2$, then the degree lower bound from Theorem 12 is $d = 2^\ell - 1$ which satisfies $N = d(d+1)/2$, i.e. the same degree as in [5] despite the more stringent conditions required in Theorem 12. (For $p > 2$ Theorem 12 does not give as large a degree bound.) It would be interesting to improve the lower bound and close the gap between $p^\ell - 1$ and $2^\ell - 1$ or to reduce the size of N required to achieve it.

Acknowledgements

The authors would like to thank Toni Pitassi and Russell Impagliazzo for discussions that clarified many of the issues in this paper, and Alasdair Urquhart for his insight as to how best to formulate these arguments.

References

[1] M. Ajtai. The independence of the modulo p counting principles. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing*, Montréal, Québec, Canada, May 1994.

[2] Miklós Ajtai. The complexity of the pigeonhole principle. In *29th Annual Symposium on Foundations of Computer Science*, pages 346–355, White Plains, NY, October 1988. IEEE.

[3] Miklós Ajtai. Parity and the pigeonhole principle. In Samuel R. Buss and P. J. Scott, editors, *Feasible Mathematics*, pages 1–24, A Mathematical Sciences Institute Workshop, Ithaca, NY, 1990. Birkhäuser.

[4] Paul W. Beame. A switching lemma primer. Technical Report 95–07–01, Department of Computer Science and Engineering, University of Washington, November 1994.

[5] Paul W. Beame, Stephen A. Cook, Jeff Edmonds, and Russell Impagliazzo. The relative complexity of NP search problems. In *Proceedings of the Twenty-Seventh Annual ACM Symposium on Theory of Computing*, Las Vegas, NV, May 1995. To appear.

[6] Paul W. Beame, Russell Impagliazzo, Jan Krajíček, Toniann Pitassi, and Pavel Pudlák. Lower bounds on Hilbert's Nullstellensatz and propositional proofs. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pages 794–806, Santa Fe, NM, November 1994. IEEE.

[7] Paul W. Beame, Russell Impagliazzo, Jan Krajíček, Toniann Pitassi, Pavel Pudlák, and Alan Woods. Exponential lower bounds for the pigeonhole principle. In *Proceedings of the Twenty-Fourth Annual ACM Symposium on Theory of Computing*, pages 200–220, Victoria, B.C., Canada, May 1992.

[8] Paul W. Beame and Toniann Pitassi. An exponential separation between the matching principle and the pigeonhole principle. In *8th Annual IEEE Symposium on Logic in Computer Science*, Montreal, Quebec, June 1993.

[9] S. Bellantoni, T. Pitassi, and A. Urquhart. Approximation and small depth Frege proofs. In *Proceedings, Structure in Complexity Theory, Sixth Annual Conference*, pages 367–391, Chicago, IL, June 1991. IEEE.

[10] S. Buss, R. Impagliazzo, J. Krajíček, P. Pudlák, A. A. Razborov, and J. Sgall. Proof complexity in algebraic systems and bounded depth Frege systems with modular counting. Manuscript, 1996.

[11] Samuel R. Buss. *Bounded Arithmetic*. Bibliopolis, Napoli, 1986. Volume 3 of Studies in Proof Theory.

[12] Johan Håstad. *Computational Limitations of Small-Depth Circuits*. MIT Press, 1987. ACM Doctoral Dissertation Award Series (1986).

[13] 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(1), 1995.

[14] J. Paris and A. Wilkie. Counting problems in bounded arithmetic. In *Methods in Mathematical Logic: Proceedings of the 6th Latin American Symposium on Mathematical Logic 1983*, volume 1130 of *Lecture notes in Mathematics*, pages 317–340, Berlin, 1985. Springer-Verlag.

[15] Toniann Pitassi, Paul W. Beame, and Russell Impagliazzo. Exponential lower bounds for the pigeonhole principle. *Computational Complexity*, 3(2):97–140, 1993.

[16] Søren Riis. *Independence in Bounded Arithmetic.* PhD thesis, Oxford University, 1993.

[17] Søren Riis. $Count(q)$ does not imply $Count(p)$: revised version. Technical Report RS 94-21, BRICS, 1994.

[18] Søren Riis. $Count(q)$ versus the pigeonhole principle. Technical Report RS 94-25, BRICS, 1994.

[19] Søren Riis. $Count(q)$ versus the pigeonhole principle: revised version. Manuscript, 1996.