Good Degree Bounds on Nullstellensatz Refutations of the Induction Principle

Samuel R. Buss[*]
Department of Mathematics
University of California, San Diego

Toniann Pitassi[†]
Department of Computer Science
University of Arizona

Abstract

This paper gives nearly optimal, logarithmic upper and lower bounds on the minimum degree of Nullstellensatz refutations (i.e., polynomials) of the propositional induction principle.

1 Introduction

A new propositional proof system based on Hilbert's Nullstellensatz was recently introduced in [2]. (See [9] for a subsequent, more general treatment of algebraic proof systems.) In this system, one begins with an initial set of polynomial equations and the goal is to prove that they are not simultaneously solvable over a field such as GF_2. A proof of unsolvability is simply a linear combination of the initial polynomial equations plus *propositional* equations $x^2 - x = 0$ for all variables x, such that the linear combination is the unsatisfiable equation 1=0. The coefficients of the linear combination may be arbitrary polynomials; the inclusion of the propositional equalities $x^2 - x = 0$ restricts the variables x to take on only propositional values 0 and 1. If such a linear combination exists, there is no assignment of $0/1$ values to the variables that satisfies all the initial equations: therefore, the linear combination is called a *Nullstellensatz refutation* of the initial equations. We can obtain in this way a proof system for propositional formulas in CNF form by translating each clause into a polynomial equation; for instance, the clause $x_1 \vee \overline{x}_2 \vee x_3$ becomes the polynomial equality $(1 - x_1)x_2(1 - x_3) = 0$. By the weak form of Hilbert's Nullstellensatz, it follows that a propositional formula is unsatisfiable if and only if it has a Nullstellensatz refutation. To be precise, the Hilbert Nullstellensatz Theorem implies that the Nullstellensatz refutation system is complete in that any propositionally unsatisfiable set of polynomials has a Nullstellensatz refutation.[2]

[*] Supported in part by NSF grant DMS-9503247, and Grant INT-9600919/ME-103 from NSF and MŠMT (Czech Republic)

[†] Supported in part by NSF grant CCR-9457782, US-Israel BSF Grant 95-00238, and Grant INT-9600919/ME-103 from NSF and MŠMT (Czech Republic)

[2] The full generality of the Hilbert Nullstellensatz is not needed to prove the completeness of the propositional Nullstellensatz proof system. This is because variables take on only 0/1 values and the equation $x^2 - x$ is present for each variable x. A constructive proof of the completeness of the propositional Nullstellensatz proof system is given in [4, 9].

To measure the complexity of a Nullstellensatz refutation, we define the *degree* of a refutation to be the maximum degree of a coefficient in the linear combination. A family of propositional formulas has *constant degree* Nullstellensatz refutations if there is a constant c so that each formula in the family has a Nullstellensatz refutation of degree at most c.

The Nullstellensatz proof system is important for several reasons. Perhaps most importantly, it is a powerful propositional proof system which has a very simple deterministic, polynomial-time procedure to find a constant degree proof, if one exists. Because the total number of monomials in a degree d refutation over $x_1, ..., x_n$ is $n^{O(d)}$, one can solve the linear equations which determine the coefficients of the monomials in polynomial time. Therefore, the proof system has potential applications in automatic theorem proving as well as for satisfiability testing. The second (and original) reason for being interested in the Nullstellensatz proof system is that it has close connections to constant-depth Frege proofs; and lower bounds on the degrees of Nullstellensatz proofs can sometimes give lower bounds on the size of constant-depth Frege proofs.

In this paper, we are interested in understanding the power of constant degree Nullstellensatz proofs. In particular, how powerful are constant degree Nullstellensatz proofs compared with resolution? Most propositional theorem provers and many algorithms for satisfiability testing are based on deterministic versions of resolution. Thus if one could obtain a deterministic proof system that can polynomially simulate resolution, this would be a major breakthrough. Because one can count modulo 2 in the Nullstellensatz system over GF_2, constant degree Nullstellensatz refutations cannot be efficiently simulated by resolution. (For example, the propositional modulo 2 principle has a degree one Nullstellensatz refutation over GF_2, but it follows from [7] that there are no polynomial-sized resolution refutations of the mod 2 principle.)

We prove in this paper that constant degree Nullstellensatz refutations cannot polynomially simulate resolution. In particular, we show that the propositional induction principle requires degree $\lfloor \log n \rfloor - 1$ Nullstellensatz proofs. Our bound is tight up to a constant: we also give $\lceil \log(n-1) \rceil$ degree Nullstellensatz refutations of the induction principle. (All logarithms in this paper are base 2.) The induction principle is of particular importance because it formalizes sequential steps in a resolution proof. In addition, our hard examples are formalized as CNF formulas with clause size 2, and they have simple linear-size resolution proofs. Therefore, our lower bound actually shows that Horn-clause resolution as well as bounded-clause-size resolution cannot be simulated by constant degree Nullstellensatz proofs.

Other degree lower bounds for Nullstellensatz were known prior to our result, but the hard examples did not separate resolution from Nullstellensatz. For example, in [1] it was shown that the pigeonhole principle from m pigeons to n holes requires degree $\sqrt{n}$ Nullstellensatz refutations, and in [4], it was shown that the mod p counting principle requires degree $n^{\Omega(1)}$ Nullstellensatz refutations over GF_q (p, q distinct primes). However, both of these principles also require exponential size resolution proofs. In another line of research, various upper and lower bounds for the degree of solutions to the Nullstellensatz in general algebraic settings have been obtained ([3, 5, 8]); however these lower bounds do not hold in our finite field setting. More recently, [6] have proved linear degree bounds for Nullstellensatz refutations of a different principle.

Our lower bound method is similar to previous lower bounds. We first show that there exists degree d Nullstellensatz refutations if and only if a particular system of linear equations can be solved. This is shown to be equivalent to the existence of a particular type of combinatorial design. The novel part of our argument is in constructing the combinatorial design. In previous degree lower bounds, the designs were always explicitly given; in contrast, our lower bound is unique in that we show the existence of a design without explicitly giving the design construction. (Of course one could construct the design from our proof, but it would be complicated.)

Our paper is organized as follows. In Section 2, we define the induction principle and the Nullstellensatz proof system; Section 3 presents some simple examples of Nullstellensatz refutations. In Section 4 we give upper bounds for the induction principle, and in Sections 5 and 6 we give our lower bound. We conclude in Section 7 with open problems and related results.

2 Definitions

The well-known induction principle (or least number principle) on n Boolean variables, $x_1, ..., x_n$, states that if $x_1 = 1$ and $x_n = 0$, then there must be some point i, $1 \leq i \leq n-1$, such that $x_i = 1$ and $x_{i+1} = 0$. We can express the negation of this principle by the following equations over an arbitrary field k:

(1) $1 - x_1 = 0$;

(2) $x_i(1 - x_{i+1}) = 0$ for all $1 \leq i \leq n-1$;

(3) $x_n = 0$; and

(4) $x_i^2 - x_i = 0$ for all $1 \leq i \leq n$.

The above set of equations over $x_1, .., x_n$ will be called IND_n. The polynomials appearing in the lefthand sides of the IND_n equations will be denoted P_i, for $1 \leq i \leq 2n+1$. Equations (1)-(3) assert the negation of the induction principle, and equations (4) are the propositional equations. Note that the equations (4) can be satisfied only by the values 0 and 1; so any solution to IND_n is a Boolean truth assignment. Since equations (2) force x_{i+1} to equal 1 provided x_i equals 1, it follows by the induction principle that these equations have no solution in (the algebraic closure of) the field k. Thus, by Hilbert's Nullstellensatz, there exists a set of polynomials $Q_j, 1 \leq j \leq 2n+1$ (with coefficients from k) such that $\sum_j P_j Q_j = 1$ (in the polynomial ring $k[\vec{x}]$).

These polynomials Q_j are called a "Nullstellensatz refutation" of the set of equations IND_n. As such they serve as a proof of the induction principle, since they show that the negation of the induction principle is unsatisfiable. Using the same method, one can give Nullstellensatz refutations of any unsatisfiable Boolean formula (in conjunctive normal form, say), see [2, 4].

The *size* of a Nullstellensatz refutation can be measured in several ways; most notably, by the total number of terms in the refutation or by maximum degree of the polynomials Q_j in the refutation. In this paper, we define the degree of a Nullstellensatz refutation to be the maximum degree of its polynomials Q_j and we measure the size of Nullstellensatz refutations by their degree.

3 Examples

To help clarify the nature of propositional Nullstellensatz refutations, we consider two examples.

Example 1: Consider the modus ponens principle with $n = 3$, that the following four propositional formulas

$$x_1 \qquad x_1 \rightarrow x_2 \qquad x_2 \rightarrow x_3 \qquad \neg x_3$$

are not simultaneously satisfiable. We can translate these four formulas into polynomials as follows:

Wff	Polynomial
x_1	$P_1 \overset{df}{=} 1 - x_1$
$x_1 \rightarrow x_2$	$P_2 \overset{df}{=} x_1(1 - x_2) = x_1 - x_1 x_2$
$x_2 \rightarrow x_3$	$P_3 \overset{df}{=} x_2(1 - x_3) = x_2 - x_2 x_3$
$\neg x_3$	$P_4 \overset{df}{=} x_3$

(Note that variables x_i in the left column are propositional variables, whereas variables in the right column are algebraic variables.) In addition to polynomials P_1–P_4, there are polynomials $P_{i+4} \overset{df}{=} x_i^2 - x_i$ for $i = 1, 2, 3$. It is clear that finding a Boolean assigment which satisfies the four propositional formulas is equivalent to finding an assigment (of field values to the variables x_i) that simultaneously sets $P_1, \ldots, P_7$ to zero.

We shall now derive is a stepwise fashion a Nullstellensatz refutation to $\{P_1, \ldots, P_7\}$. The first step is to derive

$$R_1 \ \overset{df}{=} \ (1 - x_2)P_1 + P_2 \ = \ 1 - x_2.$$

For the second step, one derives

$$R_2 \ \overset{df}{=} \ (1 - x_3)R_1 + P_3 \ = \ 1 - x_3.$$

Thirdly, we have $P_4 + R_2 = 1$. Putting these together yields

$$(1 - x_2)(1 - x_3)P_1 + (1 - x_3)P_2 + P_3 + P_4 \ = \ 1.$$

This is a Nullstellensatz refutation of $\{P_1, \ldots, P_7\}$ since it provides polynomials Q_i so that $\sum Q_i P_i = 1$. The existence of the Nullstellensatz refutation clearly implies that the polynomials P_i cannot be simultaneously equal to zero; this in turn implies that the original propositional formulas cannot be simultaneously true.

Note that the above Nullstellensatz refutation works for polynomials over any field (indeed over any ring). Since the maximum degree of the Q_i's is two, this refutation has degree two.

One can readily extend the idea behind the above refutation to get degree $n - 1$ Nullstellensatz proofs of the induction principles. However, in section 4, we give better, logarithmic degree Nullstellensatz proofs.

Example 2: For a second example, consider the following set of three formulas,

$$x_1 \leftrightarrow \neg x_2 \qquad x_2 \leftrightarrow \neg x_3 \qquad x_3 \leftrightarrow \neg x_1,$$

which is not satisfiable. We translate these into polynomials as follows:

Wff	Polynomial
$x_1 \leftrightarrow \neg x_2$	$P_1 \overset{df}{=} x_1 + x_2 - 1$
$x_2 \leftrightarrow \neg x_3$	$P_2 \overset{df}{=} x_2 + x_3 - 1$
$x_3 \leftrightarrow \neg x_1$	$P_3 \overset{df}{=} x_1 + x_3 - 1$

In addition P_{i+3} is $x_i^2 - x_i$, for $i = 1, 2, 3$. Note that the standard translation for any CNF formula mentioned in the introduction would give rise to different polynomials. In particular, we would always get a polynomial that on 0/1 values to the variables, returns a 0/1 value. (For example, $x_1 \leftrightarrow \neg x_2$ would translate to the polynomial $1 - x_1 - x_2 + 2x_1x_2$.) Our translation above is still fine, since it still has the property that a boolean truth assignment to the variables satisfies the wff if and only if that assignment sets the polynomial to zero.

To obtain a Nullstellensatz refutation of this example, let us first suppose that we are working with polynomials over Z_2 (or any field of characteristic 2). Then we just sum the polynomials P_1, P_2 and P_3 to get

$$P_1 + P_2 + P_3 \;=\; 2x_1 + 2x_2 + 2x_3 - 3 \;\equiv 1 \pmod 2.$$

This is therefore a valid Nullstellensatz refutation of degree zero.

For fields of characteristic not equal to two, we need different Nullstellensatz refutations. Indeed it is easy to prove that there is no degree zero Nullstellensatz refutation over fields of characteristic other than 2. Instead, one can use the identity

$$(2x_1 - 1)P_1 - (2x_1 - 1)P_2 + (2x_1 - 1)P_3 - 4P_4 \;=\; 1$$

which forms a Nullstellensatz refutation of degree one (using polynomials over an arbitrary field).

4 Upper bounds

4.1 Logarithmic upper bounds

We first show that the above system IND_n of equations has a degree $\lceil \log(n-1) \rceil$ Nullstellensatz refutation. Our arguments work over an arbitrary field, and for the rest of this section we assume we are working with polynomials over an arbitrary fixed field.

Assume for now that n is of the form 2^{m+1}. First, for all odd values of i, from the equations $x_i(1 - x_{i+1}) = 0$ and $x_{i+1}(1 - x_{i+2}) = 0$, derive $x_i(1 - x_{i+2}) = 0$ by forming the linear combination:

$$(1 - x_{i+2})\,[x_i(1 - x_{i+1})] \;+\; (x_i)\,[x_{i+1}(1 - x_{i+2})] \;=\; x_i(1 - x_{i+2}).$$

The degree of this step is 1. Next, for all i such that 4 divides $i-1$, from $x_i(1-x_{i+2})=0$ and $x_{i+2}(1-x_{i+4})=0$, derive $x_i(1-x_{i+4})=0$ in a similar manner. Again, the degree of this step is 1. Continue for at most $\lceil \log(n-1) \rceil$ iterations until we derive $x_1(1-x_n)=0$. Now we can derive $x_1(1-x_n)+x_1x_n+(1-x_1)=0$, which is equivalent to $1=0$. The total degree of this derivation is $\lceil \log(n-1) \rceil$. Now suppose that $2^m < n < 2^{m+1}$. (That is, n is not a power of two.) Create $2^{m+1}-n$ new, "dummy" initial equations $0=0$ in order to pad out the initial equations to a power of two, and then apply the algorithm described above.

Theorem 1 *IND_n has a Nullstellensatz refutation of degree $\lceil \log(n-1) \rceil$ (over an arbitrary field).*

4.2 An equivalent principle

We will convert the above propositional induction principle into a restricted form of the pigeonhole principle for intuitive convenience in our lower bound argument. The pigeonhole principle is viewed as a mapping from pigeons numbered 1 through n to holes numbered 1 through $n-1$. The following set of equations expresses that each pigeon i, $1 \leq i \leq n$, is mapped to either hole $i-1$ or to hole i, and no hole has more than one pigeon mapped to it, where we let the variable $P_{i,0}$ denote the condition that pigeon i is mapped to hole $i-1$ and the variable $P_{i,1}$ denote the condition that it is mapped to hole i.

(1) $P_{1,0}=0$;

(2) $P_{n,1}=0$;

(3) $P_{i,0}+P_{i,1}-1=0$ for all i, $1 \leq i \leq n$;

(4) $P_{i,1}P_{i+1,0}=0$ for all i, $1 \leq i \leq n$; and

(5) $P_{i,j}^2-P_{i,j}=0$ for all i, $1 \leq i \leq n$, and all $j \in \{0,1\}$.

The above principle will be called $NEARPHP_n$. It is not hard to see that $NEARPHP_n$ is roughly equivalent to IND_n. First we will show that from a degree d refutation of $NEARPHP_n$, we can construct a degree d refutation of IND_n. Suppose we have a degree d Nullstellensatz refutation of $NEARPHP_n$. For all i, replace all occurrences of $P_{i,0}$ in the refutation of $NEARPHP_n$ by $1-x_i$, and replace all occurrences of $P_{i,1}$ by x_i. The equation of type (1) becomes $1-x_1=0$, which is an initial equation of IND_n; the equation of type (2) becomes $x_n=0$ which is an initial equation of IND_n; the equations of type (3) become $0=0$, so they can be removed; the equations of type (4) become equations of the form $(x_i)(1-x_{i+1})=0$, which is also an initial equation of IND_n; and lastly, equations of type (5) become initial equations of the form $x_i^2-x_i=0$. Thus, from a degree d refutation of $NEARPHP_n$ we easily get a degree d refutation of IND_n.

In the other direction, suppose we have a degree d Nullstellensatz refutation of IND_n:

$$Q(1-x_1)+R(x_n)+\sum_i P_i(x_i(1-x_{i+1}))+\sum_i S_i(x_i^2-x_i) \quad = \quad 1,$$

with Q, R, P_i and S_i polynomials of degree $\leq d$. First, replace all occurrences of x_i by $P_{i,1}$ for all i, $1 \leq i \leq n$, in the above equation. This yields:

$$Q^*(1-P_{1,1})+R^*(P_{n,1})+\sum_i P_i^*(P_{i,1}(1-P_{i+1,1}))+\sum_i S_i^*(P_{i,1}^2-P_{i,1}) \;=\; 1,$$

where Q^*, R^*, P_i^* and S_i^* are still polynomials of degree $\leq d$, now in the underlying variables $P_{i,1}$. We can write $(1-P_{1,1})$ as a linear combination of initial equations: $P_{1,0}-(P_{1,0}+P_{1,1}-1)$; $P_{n,1}$ is already an initial equation; $P_{i,1}^2 - P_{i,1}$ is also already an initial equation; and lastly, we can write $P_{i,1}(1-P_{i+1,1})$ as a degree 1 combination of the initial equations, namely, as

$$P_{i,1}P_{i+1,0} \;-\; P_{i,1}(P_{i+1,0}+P_{i+1,1}-1)$$

Applying these substitutions, we end up with a degree $d+1$ Nullstellensatz refutation of $NEARPHP_n$. We have now proved:

Theorem 2 *Fix an arbitrary field.*

(a) *If IND_n has a degree d Nullstellensatz refutation, then $NEARPHP_n$ has a degree $d+1$ refutation.*

(b) *If $NEARPHP_n$ has a degree d Nullstellensatz refutation, then IND_n has a degree d refutation.*

Corollary 3 *$NEARPHP_n$ has a degree $\lceil \log n \rceil + 1$ degree refutation (over an arbitrary field).*

The rest of the paper establishes a closely matching lower bound on the degree of Nullstellensatz refutations of $NEARPHP_n$ over a field k. We consider the field k to be fixed and let p be its characteristic. In particular, GF_p is a subfield of k. For $p=0$, we let GF_0 denote $\mathcal{Q}$, the rationals.

The lower bound proved below (and the main result of this paper) is:

Theorem 4 *Fix an arbitrary field. $NEARPHP_n$ does not have a degree $\lfloor \log n \rfloor - 1$ Nullstellensatz refutation. Therefore, any Nullstellensatz refutation of IND_n has degree at least $\lfloor \log n \rfloor - 1$.*

5 The reduction to designs

In this section we define *designs* for the induction principle and reduce the problem of obtaining the lower bound on the degree of Nullstellensatz refutations of the induction principle to the problem of constructing designs of sufficiently high degree.

5.1 Partial matchings and designs

Definition A *partial matching* is a function π with domain contained in $\{1,\ldots,n\}$ and range contained in $\{0,1,\ldots,n\}$, such that, for all i, if $\pi(i)$ is defined, then it equals either $i-1$ or i. If a partial matching is one-to-one and has range contained in $\{1,\ldots,n-1\}$, then it is a *proper* partial matching; otherwise it is *improper*. We view partial matchings as sets of variables (or rather, we view appropriate sets of variables as partial matchings) by using the following conventions. A set of variables $\{P_{i_1,j_1},\ldots,P_{i_d,j_d}\}$, with the values of $i_1,\ldots,i_d$ all distinct, represents the partial matching π such that $\pi(i_k)=i_k+j_k-1$ for $k=1,\ldots,d$.

Frequently we want to consider partial matchings with the variables sorted by their first subscript (i.e., by the pigeon-numbers). For this purpose, we denote a partial matching as an ordered tuple of variables, in the form

$$\langle P_{i_1,j_1},P_{i_2,j_2},\ldots,P_{i_d,j_d}\rangle$$

where the $\langle\cdots\rangle$ notation indicates that $i_1<i_2<\cdots<i_d$.

The cardinality, d, of the partial matching is called its *degree*. The *domain* of the partial matching π above is $\{i_1,\ldots,i_d\}$ and is denoted $dom(\pi)$.

Definition A *design* D of degree d is a mapping from the partial matchings of degrees $\leq d$ to the range GF_p such that the conditions (1)-(3) below hold.

(1) $D(\emptyset)=1$, $\emptyset$ is the empty partial matching.

(2) If the degree of π is $<d$ and if $i\notin dom(\pi)$, then

$$D(\pi\cup P_{i,0})+D(\pi\cup P_{i,1})\equiv D(\pi)\pmod p$$

(3) For improper π, $D(\pi)=0$.

Note that the definition of a design depends in the characteristic p of the field k. In the special case where $p=2$ we can also consider D to be a subset of the proper partial matchings of degree $\leq d$ by identifying the subset with its characteristic function. In this case, the condition (2) can be restated as follows:

$(2')$ Let $\pi=\{P_{i_1,j_1},\ldots,P_{i_r,j_r}\}$ be a partial matching of degree $r<d$ and suppose $i_{r+1}\notin\{i_1,\ldots,i_r\}$. Then, if $\pi\in D$, there is exactly one $j_{r+1}\in\{0,1\}$ such that $\{P_{i_1,j_1},\ldots,P_{i_{r+1},j_{r+1}}\}$ is in D. And, if $\pi\notin D$, then there are either zero or two values of $j_{r+1}\in\{0,1\}$ such that $\{P_{i_1,j_1},\ldots,P_{i_{r+1},j_{r+1}}\}$ is in D.

The next theorem, which applies to any field k, reduces the problem of proving a lower bound on the degree of Nullstellensatz refutations to the problem of proving the existence of designs.

Theorem 5 *Suppose there is a design D of degree d. Then any Nullstellensatz refutation of $NEARPHP_n$ has degree $\geq d$.*

Proof Let $F_1,\ldots,F_{4n+2}$ be the polynomials on the lefthand sides of the $NEARPHP_n$ equations. Suppose, for sake of a contradiction, that there are polynomials Q_i of degree $<d$ such that

$$\sum_i Q_i\cdot F_i=1$$

in the polynomial ring $k[\vec{x}]$.

We define a *power product* to be an expression of the form $\mathbf{X}=\prod P_{i,j}^{a_{i,j}}$, i.e., a product of nonnegative powers of variables. A monomial is an expression of the form $\alpha\mathbf{X}$ with $\alpha\in k$

and $\mathbf{X}$ a power product. Since each Q_i is equal to a sum of monomials, the above equation implies that there are finitely many monomials $R_{i,j}$ so that

$$\sum_{i,j} R_{i,j} \cdot F_i \;=\; 1 \qquad (\text{in } k[\vec{x}]) \tag{1}$$

holds; where the summation is taken over all appropriate i's and j's.

We extend the degree d design D to be a mapping on polynomials as follows: Firstly, we define $D(\pi) = 0$ for any π of degree greater than d. Secondly, for any power product $\mathbf{X} = \prod P_{i,j}^{a_{i,j}}$, if $\pi_{\mathbf{X}} = \{P_{i,j} : a_{i,j} \neq 0\}$ is a partial one-to-one mapping, then $D(\mathbf{X})$ is defined to equal $D(\pi_{\mathbf{X}})$. However, if $\pi_{\mathbf{X}}$ is not a partial one-to-one mapping (by having two variables with the same first subscript), then $D(\pi_{\mathbf{X}}) = 0$. Thirdly, for any monomial $\alpha\mathbf{X}$, $D(\alpha\mathbf{X})$ is defined to equal $\alpha D(\mathbf{X})$. Finally, if g_i are monomials, $D(\sum_i g_i)$ is defined to equal $\sum_i D(g_i)$. In this way, D becomes a mapping from polynomials into k. Note that D respects addition, but not necessarily multiplication; i.e., $D(g + h) = D(g) + D(h)$ for all polynomials g, h.

Since equation (1) holds in $k[\vec{x}]$, we have,

$$\sum_{i,j} D(R_{i,j} \cdot F_i) \;=\; 1 \qquad (\text{in } k[\vec{x}]),$$

with each $R_{i,j}$ of degree less than d. But this immediately contradicts the next claim.

Claim For $\mathbf{X}$ any power product of degree less than d, and for F_i any $NEARPHP_n$ polynomial,

$$D(\mathbf{X} \cdot F_i) \;=\; 0.$$

Proof The proof of the claim is almost immediate from the definition of a design; one merely has to consider the possible cases for F_i. When F_i is $P_{1,0}$ or $P_{n,1}$, then $\mathbf{X}F_i$ is an improper monomial (i.e., $\pi_{\mathbf{X}F_i}$ is not a proper partial matching), and therefore $D(\mathbf{X}F_i) = 0$. When F_i is $P_{i,1}P_{i+1,0}$, then $\mathbf{X}F_i$ is improper (possibly of degree $d+1$) and hence $D(\mathbf{X}F_i) = 0$. When F_i is $P_{i,j}^2 - P_{i,j}$, then $D(\mathbf{X}P_{i,j}^2) = D(\mathbf{X}P_{i,j})$ by the definition of D as applied to power products. Finally, consider the case where F_i is $P_{i,0} + P_{i,1} - 1$. If $i \notin \mathit{dom}(\pi_{\mathbf{X}})$, then

$$D(\mathbf{X}F_i) \;=\; D(\mathbf{X}P_{i,0}) + D(\mathbf{X}P_{i,1}) - D(\mathbf{X}) \;=\; 0,$$

by condition (2) in the definition of designs since $\mathit{degree}(\mathbf{X}F_i) \le d$. If $i \in \mathit{dom}(\pi_{\mathbf{X}})$, then (at least) one of $\mathbf{X}P_{i,0}$ and $\mathbf{X}P_{i,1}$ is improper and again $D(\mathbf{X}F_i) = 0$.

That completes the proof of the claim and of Theorem 5. □

5.2 Block-respecting designs

In order to prove the lower bound for Theorem 4, it will suffice to build a design of degree $\lfloor \log n \rfloor$. It will be convenient to build a special, restricted kind of design called a "block respecting design". Recall that the n domain elements, the pigeons, are numbered from 1 to n. We group these pigeons into blocks; the blocks are arranged in a hierarchy of *levels.* At level number ℓ, there are 2^ℓ blocks which partition the pigeons into intervals of approximately equal size. More precisely, the blocks of level ℓ are denoted $\mathcal{B}_i^\ell$, for $0 \le i < 2^\ell$. The block $\mathcal{B}_j^\ell$ contains the pigeons numbered in the closed interval

$$\left[\left\lfloor \frac{jn}{2^\ell} \right\rfloor + 1, \left\lfloor \frac{(j+1)n}{2^\ell} \right\rfloor\right]$$

We call $\mathcal{B}_j^\ell$ the j-th block at level ℓ; note numbering of blocks starts with the zeroth block. Blocks are defined for levels $\ell = 0, \ldots, \lfloor \log n \rfloor$. We write B_{first}^ℓ and B_{last}^ℓ for B_0^ℓ and $B_{2^\ell - 1}^\ell$, respectively. Two blocks B_i^ℓ and B_{i+1}^ℓ are called *adjacent blocks.*

Lemma 6 *Block satisfy the following simple properties:*

(1) *The* 2^ℓ *many blocks at level* ℓ *are pairwise disjoint.*

(2) *For all* ℓ*,* $\bigcup_i \mathcal{B}_i^\ell = \{1, \ldots, n\}$.

(3) $\mathcal{B}_j^\ell = \mathcal{B}_{2j}^{\ell+1} \cup \mathcal{B}_{2j+1}^{\ell+1}$.

(4) *For all* $\ell = 0, \ldots, \lfloor \log n \rfloor$ *and* $0 \le j < 2^\ell$*, the block* $\mathcal{B}_j^\ell$ *is nonempty.*

For $1 \le i \le n$ and $j = 0, 1$, we use $P_{i,j}$ as the variable which has truth value *True* when there is an edge from pigeon i to hole $i + j - 1$. For notational convenience, we allow also $P_{1,0}$ and $P_{n,1}$ as propositional variables; these will appear only in improper partial matchings, and one should think of them as variables which always take truth value *False.*

Definition We say that $P_{i,j}$ belongs to block B, written $P_{i,j} \in B$ provided that $i \in B$. We write $i \underset{\ell}{\sim} j$ if i and j are in the same level ℓ block.

For partial matchings $\pi = \langle P_{i_1,j_1} \ldots P_{i_d,j_d} \rangle$ and $\pi' = \langle P_{i'_1,j'_1} \ldots P_{i'_d,j'_d} \rangle$ of the same degree, we write $dom(\pi) \underset{d}{\sim} dom(\pi')$ to denote the condition that $i_k \underset{d}{\sim} i'_k$ for all $1 \le k \le d$.

Definition A degree d design D is *block respecting* provided the following conditions hold:

(α) *"Block Respecting Property".* Suppose $\pi = \langle P_{i_1,j_1}, \ldots P_{i_r,j_r} \rangle$ and $\pi' = \langle P_{i'_1,j'_1}, \ldots P_{i'_r,j'_r} \rangle$ are partial matchings of degree $r \le d$. Also suppose $dom(\pi_1) \underset{r}{\sim} dom(\pi')$ and that $j_k = j'_k$ for $1 \le k \le r$, then

$$D(\pi) \; = \; D(\pi').$$

(β) *"No Improper Matchings Property".* If $\pi = \langle P_{i_1,j_1}, \ldots P_{i_r,j_r} \rangle$ and if $D(\pi) \ne 0$, then the following hold:

(*i*) Let $1 \le k < r$. If $i_k \underset{r}{\sim} i_{k+1}$ or if i_k and i_{k+1} are in adjacent blocks at level r, then $j_k = 0$ or $j_{k+1} = 1$ (or both).

(*ii*) If $i_1 \in B_{\text{first}}^r$, then $j_1 = 1$.

(*iii*) If $i_r \in B_{\text{last}}^r$, then $j_r = 0$.

The property (α) is the crucial property that makes a design block respecting. Then property (β) is forced by (α) and the fact that all improper partial matchings are mapped to zero by a design.

Definition Let $\mathcal{I} = \langle i_1, \ldots, i_r \rangle$, with $i_1 < i_2 < \cdots < i_r$. The $\mathcal{I}$-hypercube, $H_{\mathcal{I}}$, is a r-dimensional hypercube containing the 2^r vertices

$$\langle j_1, \ldots, j_r \rangle, \text{ with } j_1, \ldots, j_r \in \{0, 1\}.$$

Each vertex is associated with the corresponding partial matchings (not necessarily proper):

$$\langle P_{i_1,j_1}, \ldots, P_{i_r,j_r} \rangle.$$

As usual, two vertices of the hypercube are *adjacent* if and only if they differ in only one coordinate: we use notations such as $\langle j_1, \ldots, j_s, \ldots, j_r \rangle$ and $\langle j_1, \ldots, j_s', \ldots, j_r \rangle$ to denote two adjacent vertices that differ only in their s-th component. The edge between these two adjacent vertices has an associated degree $r-1$ partial matching, namely, it is associated with

$$\langle P_{i_1,j_1} \ldots P_{i_{s-1},j_{s-1}}, P_{i_{s+1},j_{s+1}}, \ldots P_{i_r,j_r} \rangle.$$

The condition (2) for D to be a design amounts to saying that if π_1 and π_2 are associated with adjacent vertices in the $\mathcal{I}$-hypercube $H_{\mathcal{I}}$ and that if π_0 is the partial matching associated with the edge joining π_1 and π_2 in the hypercube, then

$$D(\pi_0) \equiv D(\pi_1) + D(\pi_2) \pmod{p}. \tag{2}$$

A *path* (resp., a *cycle*) in the hypercube $H_{\mathcal{I}}$ is defined to be a path (resp, a cycle), in the usual sense, in the hypercube when viewed as a graph. Suppose ρ is a path beginning at vertex x and ending at vertex y, and containing the edges $e_1, \ldots, e_k$. Let π_x and π_y be the degree r partial matchings associated with x and y, and let π_{e_i} be the degree $r-1$ partial matching associated with edge e_i. By applying equation (2) k times, we get that

$$D(\pi_x) + (-1)^{k-1} D(\pi_y) \equiv D(\pi_{e_1}) - D(\pi_{e_2}) + D(\pi_{e_3}) - \cdots + (-1)^{k-1} D(\pi_{e_k}) \pmod{p}.$$

When ρ is a cycle, and therefore $x = y$ and k is even, this becomes

$$0 \equiv D(\pi_{e_1}) - D(\pi_{e_2}) + D(\pi_{e_3}) - \cdots + D(\pi_{e_{k-1}}) - D(\pi_{e_k}) \pmod{p}. \tag{3}$$

These observations prompt the following definition.

Definition Let ρ be a path containing (in path order) edges $e_1, \ldots, e_k$. Then the *weight* of ρ is defined to be

$$\sum_{i=1}^{k} (-1)^{i-1} D(\pi_{e_i}) \bmod p.$$

6 The design construction

We now establish Theorem 7 which, together with Theorems 2 and 5, implies our main result Theorem 4.

Theorem 7 *For all $d \leq \lfloor \log n \rfloor$, there exists a block-respecting design of degree d (for any characteristic p).*

Proof. We will prove the above theorem by induction on d.

When $d = 0$, there is only one block, B_0^0, and the only matching of size 0 is $\emptyset$, the empty partial matching. Therefore there is only one degree zero design, denoted $\mathcal{D}_0$, and $\mathcal{D}_0(\emptyset) = 1$. Note that conditions (α) and (β) are vacuously satisfied for $\mathcal{D}_0$.

To get the proof by induction started, we also need to construct a block-respecting design $\mathcal{D}_1$ of degree 1. There are two blocks, B_0^1 and B_1^1, which must be considered to ensure that $\mathcal{D}_1$ satisfies the properties (α) and (β). Condition $(\beta.ii)$ implies that $\mathcal{D}_1(\langle P_{1,1} \rangle) = 1$. Then condition (α) implies that $\mathcal{D}_1(\langle P_{i,1} \rangle) = 1$ for all $i \in B_0^1$. Similarly, condition $(\beta.iii)$ implies that $\mathcal{D}_1(\langle P_{n,0} \rangle) = 1$ and then condition (α) implies that $\mathcal{D}_1(\langle P_{j,0} \rangle) = 1$ for all $j \in B_1^1$. For all other matchings π of size 1, $\mathcal{D}_1(\pi) = 0$.

Thus, the design $\mathcal{D}_1$ (and by the same reasoning, any block-respecting design) is completely forced at levels $d = 0$ and $d = 1$. At level 1, all pigeons in block 0 (B_0^1) are mapped "up" and all pigeons in block 1 (B_1^1) are mapped "down". At higher levels, the design conditions will not uniquely force a particular design — instead we will show that there is at least one design extending the current one for all $d \leq \lfloor \log n \rfloor$.

The remainder of this section is dedicated to proving the inductive step. Let $1 < r \leq \lfloor \log n \rfloor$, and assume that there exists a design, $\mathcal{D}_{r-1}$, of degree $r - 1$. We want to show that there exists a design, $\mathcal{D}_r$, of degree r extending $\mathcal{D}_{r-1}$.

Fix a particular set of pigeons $\mathcal{I} = \langle i_1, \ldots, i_r \rangle$, with $i_1 < i_2 < \cdots < i_r$. We want to show that it is possible to extend $\mathcal{D}_{r-1}$ to assign values to the r-matchings over $\mathcal{I}$ so that all of the design conditions are satisfied for these r-matchings.

Recall the $\mathcal{I}$-hypercube, $H_{\mathcal{I}}$, where the 2^r vertices include all possible matchings on $\mathcal{I}$. As a first step to defining $\mathcal{D}_r$, we shall show that it is possible to consistently pick values for $\mathcal{D}_r(\pi)$ for all partial matchings π which are associated with vertices of $H_{\mathcal{I}}$. For this purpose, we define the edge labeling of $H_{\mathcal{I}}$ *induced by* $\mathcal{D}_{r-1}$ as follows: each edge e of $H_{\mathcal{I}}$ is labeled with the value $\mathcal{D}_{r-1}(\pi_e)$ This labeling imposes conditions on the allowable matchings over $\mathcal{I}$ in $\mathcal{D}_r$; namely the condition given by equation (2). As a first step, we want to show that it is possible to find values for $\mathcal{D}_r(\pi)$, for all π with $dom(\pi) = \mathcal{I}$, such that equation (2) holds for all π_1, π_2 associated with vertices in $H_{\mathcal{I}}$ connected by an edge labeled with π_0.

Definition An edge labeling of $H_{\mathcal{I}}$ is *consistent* if for all cycles ρ in $H_{\mathcal{I}}$, $weight(\rho) = 0$. Otherwise, the edge labeling is *inconsistent.*

Definition Let l_E be an edge labeling of $H_{\mathcal{I}}$ and let l_V be a vertex labeling of $H_{\mathcal{I}}$. Then (l_E, l_V) is *consistent* if for all edges e connecting vertices x, y in $H_{\mathcal{I}}$,

$$l_V(x) + l_V(y) \equiv l_E(e) \pmod{p}.$$

In other words, an edge labeling is consistent provided equation (3) holds for all cycles in the hypercube. Likewise, an edge and vertex labeling is consistent if all instances of equation (2) hold.

Lemma 8 *If the edge labeling l_E of $H_{\mathcal{I}}$ induced by $\mathcal{D}_{r-1}$ is consistent, then there are p vertex labelings l_V such that (l_E, l_V) is consistent. (When $p = 0$, there are infinitely many such vertex labelings.*

Proof Fix l_E, a consistent edge labeling, and fix some vertex $v \in H_{\mathcal{I}}$. Pick any $\alpha \in GF_p$. We define a vertex labeling l_V by setting $l_V(v) = \alpha$ and then extending l_V to the rest of the vertices. Since $H_{\mathcal{I}}$ is connected, there is at most one way to extend l_V to make (l_E, l_V) consistent; and because l_E is consistent, this extension is consistent. Thus there is exactly one way to extend the vertex labeling of v to all vertices so that (l_E, l_V) is consistent. Each $\alpha \in GF_p$ yields a different such vertex labeling l_V consistent with l_E. □

Note that if l_E is inconsistent, then there is no vertex labeling l_V such that (l_E, l_V) is consistent.

Lemma 9 *Let l_E be the edge labeling induced by $\mathcal{D}_{r-1}$. Then l_E is consistent.*

Proof First, we show that if all cycles of length 4 in $H_{\mathcal{I}}$ are consistent, then all cycles in $H_{\mathcal{I}}$ are consistent. This is essentially a finite analogue of Stokes' theorem; the point is that any cycle can be written as the sum of cycles of length four. For a formal proof of this, assume that all cycles of length four have weight zero, and let $\rho = \langle e_1, ..., e_m \rangle$ be a cycle in $H_{\mathcal{I}}$ with m edges. We must show that ρ has weight zero. The cycle ρ can be characterized by (a) its initial vertex and (b) by its *dimension sequence* $d_1, \ldots, d_m$ such that the edge e_k is pointed in the direction of the d_k-th dimension. (The latter condition means that the endpoints of e_k differ in the value of their d_k-th components.) Now if $d_k = d_{k+1}$ for some value of k, then the cycle can be shortened by removing the two edges $e_k = e_{k+1}$. If we fix a value for k and define ρ' to be the cycle with the same initial vertex and having dimension sequence equal that of ρ except with the values d_k and d_{k+1} interchanged, then it is immediate that ρ and ρ' differ by a four cycle, and the difference in their weights, $weight(\rho') - weight(\rho)$, is equal to the weight of that four cycle, i.e., their weights are equal. The upshot is that the dimension sequence can be permuted arbitrarily without affecting the weight of the cycle, and since each dimension value appears an even number of times in the cycle, they can all be cancelled out. Therefore every cycle has weight equal to the weight of the empty cycle, i.e., has weight zero.

Secondly, we will show that all cycles of length four in $H_{\mathcal{I}}$ are consistent. Assume for sake of contradiction that some cycle of length four is inconsistent — that is, it has an odd edge labeling. Without loss of generality, assume the vertices of the cycle are as follows: $\langle 0, 0, \bar{\jmath} \rangle$, $\langle 0, 1, \bar{\jmath} \rangle$, $\langle 1, 1, \bar{\jmath} \rangle$, $\langle 1, 0, \bar{\jmath} \rangle$, $\langle 0, 0, \bar{\jmath} \rangle$, where $\bar{\jmath}$ is some fixed $0 - 1$ setting to $j_3, .., j_r$. Denote the edges of this cycle as e_1, e_2, e_3, e_4 and let the four associated partial matchings of degree $r-1$ be $\pi_1, \pi_2, \pi_3, \pi_4$. Thus, π_1 and π_3 are partial matchings with domain $i_1, i_3, \ldots, i_r$ so that $\pi_1(i_1) = i_1 - 1$ and $\pi_3(i_1) = i_1$ and so that $\pi_1(i_k) = \pi_3(i_k)$ are set according to $\bar{\jmath}$ for $k \geq 3$ (i.e., $\pi_1(i_k) = \pi_3(i_k) = i_k + j_k - 1$, for $k \geq 3$). Likewise, π_2 and π_4 have domain $i_2, i_3, \ldots, i_r$ and $\pi_2(i_2) = i_2$ and $\pi_4(i_2) = i_2 - 1$. The weight of the cycle ρ is equal to

$$D_{r-1}(\pi_1) - D_{r-1}(\pi_2) + D_{r-1}(\pi_3) - D_{r-1}(\pi_4).$$

Let π_0 be the degree $r-2$ partial matching with domain $i_3, \ldots, i_r$ such that $\pi_0(i_k)$ is set according to $\bar{\jmath}$. Then we have

$$\begin{aligned} D_{r-1}(\pi_1) + D_{r-1}(\pi_3) &\equiv \mathcal{D}_{r-2}(\pi_0) \pmod{p} \\ D_{r-1}(\pi_2) + D_{r-1}(\pi_4) &\equiv \mathcal{D}_{r-2}(\pi_0) \pmod{p} \end{aligned}$$

since D_{r-1} is a design and satisfies equation (2). It follows immediately that the weight of ρ equals zero. □

Lemmas 8 and 9 show that there are potentially p possible choices of ways to set the values of $\mathcal{D}_r(\pi)$ for partial matchings π which are associated with vertices of $\mathcal{I}$; namely, for each fixed vertex labeling l_V such that (l_E, l_V) is consistent, it may be possible to define

$$\mathcal{D}_r(\langle P_{i_1,j_1}, \ldots, P_{i_r,j_r}\rangle) \ = \ l_V(\langle j_1, \ldots, j_r\rangle).$$

However, many of these p potential choices may be disallowed by the design condition that $D(\pi) = 0$ for all improper partial matchings π. So it still remains to show that there is a vertex labeling l_V such that all improper matchings are labeled with 0. Such a vertex labeling will give a truly valid way to extend $\mathcal{D}_{r-1}$ to r-matchings over $\mathcal{I}$. This motivates the following definition and final lemma.

Definition Let $\mathcal{I} = \langle i_1, .., i_r\rangle$ and let $\langle j_1, ..., j_r\rangle$ be a vertex of the r-dimensional hypercube, $H_{\mathcal{I}}$. Then vertex $\langle j_1, .., j_r\rangle$ of $H_{\mathcal{I}}$ is *improper* if any of the following conditions hold: (1) i_k and i_{k+1} are in the same block or are in adjacent blocks and $j_k = 1$ and $j_{k+1} = 0$, (2) $i_1 \in B^r_{\text{first}}$ and $j_1 = 0$, or (3) $i_r \in B^r_{\text{last}}$ and $j_r = 1$.

In other words, a vertex of $H_{\mathcal{I}}$ is defined to be *improper* if its associated degree r partial matching fails to satisfy conditions $(\beta.i)$, $(\beta.ii)$ and $(\beta.iii)$.

Lemma 10 *Let l_E be the consistent edge labeling induced by $\mathcal{D}_{r-1}$. Then there is a vertex labeling l_V which is consistent with l_E such that all improper vertices are labeled with zero.*

Proof Let x and y be two improper vertices with associated improper partial matchings π_x and π_y. It will suffice to show that there is a path in $H_{\mathcal{I}}$ connecting x and y such that every edge in the path has label zero in the edge labeling induced by $\mathcal{D}_{r-1}$.

Suppose x is improper according to case $(\beta.i)$ of the definition and choose a value k_x such that i_{k_x} and i_{k_x+1} are in the same block or in adjacent blocks and such that $\pi_x(i_{k_x}) = 1$ and $\pi_x(i_{k_x+1}) = 0$. Likewise, suppose y is improper according to case $(\beta.i)$, choose k_y similarly. We need to choose a path from x to y: the best way to view such a path is as a sequence of changes to single values of x which transform x into y. There are several cases to consider, depending on how the sets $\{k_x, k_x+1\}$ and $\{k_y, k_y+1\}$ intersect. First, if $k_x = k_y$, then one can choose a path from x to y by changing, one at a time, the values where the two partial matchings differ. Since they agree at $\{i_{k_x}, i_{k_x+1}\}$, then each edge in this path has an associated partial matching which also maps i_{k_x} and i_{k_x+1} to the same values; such degree $r-1$ matchings are improper and therefore the edges to which they are associated have label zero. Second, if $\{k_x, k_x+1\}$ and $\{k_y, k_y+1\}$ are disjoint, then the path from

x to y may be picked by first changing the values of the matching x at i_{k_y} and i_{k_y+1} and then changing the rest of the values of the matching. Again, each edge in this path is associated with an improper degree $r-1$ partial matching and thus has label zero. Third, suppose $k_x + 1 = k_y$ (the case where $k_y + 1 = k_x$ is the same). In this case, start by changing the values of x which differ from those of y except for the values at i_{k_x} and i_{k_y}. In this way, by traversing a path in which all edges are associated with improper matchings and so have label zero, a vertex u is reached, which differs from y in the value at i_{k_y} and possibly in the value at i_{k_x} (it is possible for u to equal x). From vertex u, change the value at i_{k_y} to reach a vertex v with $v(i_{k_y}) = 1$. Either $v = y$ or v is adjacent to y and they are connected by an edge which has an improper associated matching. It therefore, will suffice to show that the edge e connecting u and v is labeled with an improper matching. The edge e has associated degree $r-1$ matching π such that $\pi(i_{k_x}) = 1$ and $\pi(i_{k_y+1}) = 0$; furthermore, i_{k_x} and i_{k_y+1} are in adjacent or equal blocks at level $r-1$, since i_{k_x} and i_{k_y} are in adjacent or equal blocks at level r and i_{k_y} and i_{k_y+1} are in adjacent or equal blocks at level r. Therefore $\mathcal{D}_{r-1}(\pi_e) = 0$ and the edge e is labeled with zero.

It remains to consider the cases where one or both of x and y are improper according to cases ($\beta.ii$) or ($\beta.iii$) of the definition. These cases can be handled by an analogous argument. Alternatively, one can allow $k = 0$ or $k = n+1$ to treat these as special cases of the previous case; then one adopts the convention that all partial matchings map 0 to 0 and $n+1$ to n and the above argument applies verbatim. □

Conclusion of proof of Theorem 7. Let $\mathcal{I} = \langle i_1, ..., i_r \rangle$ be a collection of pigeons, $i_1 < i_2 < ... < i_r$, and let $l_E^{\mathcal{I}}$ be the edge labeling of $H_{\mathcal{I}}$ induced by $\mathcal{D}_{r-1}$. By Lemmas 8-10 there is a vertex labeling $l_V^{\mathcal{I}}$ of $H_{\mathcal{I}}$ such that $(l_E^{\mathcal{I}}, l_V^{\mathcal{I}})$ is consistent and such that each improper vertex has label zero. In addition, we can choose the labelings $l_V^{\mathcal{I}}$ for all $\mathcal{I}$, so that $l_V^{\mathcal{I}} = l_V^{\mathcal{I}'}$ whenever $\mathcal{I} \underset{r}{\sim} \mathcal{I}'$; this is possible, firstly, since if $\mathcal{I} \underset{r}{\sim} \mathcal{I}'$ then, because $\mathcal{D}_{r-1}$ is block-respecting, the same edge labelings are induced on $H_{\mathcal{I}}$ and on $H_{\mathcal{I}'}$, and secondly, since the same vertices are improper in $H_{\mathcal{I}}$ and in $H_{\mathcal{I}'}$.

Define $\mathcal{D}_r$ by letting $\mathcal{D}_r(\pi) = \mathcal{D}_{r-1}(\pi)$ for all π of degree $< r$ and letting

$$\mathcal{D}_r(P_{i_1,j_1}, P_{i_2,j_2}, ..., P_{i_r,j_r}) \ = \ l_V^{\langle\vec{i}\rangle}(j_1, ..., j_r).$$

We claim that $\mathcal{D}_r$ is a block-respecting design of degree r. The induction hypothesis that D_{r-1} is a block-respecting $r-1$ design means that it is only necessary to check that equation (2) and conditions (α) and (β) are valid for the maximum degree partial matchings. Equation (2) is valid because of the consistency of $(l_E^{\mathcal{I}}, l_V^{\mathcal{I}})$ for all $\mathcal{I}$. The block respecting property (α) holds because of the fact that $l_V^{\mathcal{I}} = l_V^{\mathcal{I}'}$ whenever $\mathcal{I} \underset{r}{\sim} \mathcal{I}'$. The property ($\beta$) is valid since the choices for $l_V^{\mathcal{I}}$ satisfy the property of Lemma 10. □

7 Conclusion

The above lower bound shows that resolution proofs cannot be translated into constant degree Nullstellensatz refutations. Recently, [6] have shown that the separation between resolution and Nullstellensatz is even worse. They give another principle, based on strong induction, and prove that it has efficient resolution proofs, but requires linear degree Nullstellensatz refutations. However, their principle is not in 3CNF form, and when it is converted into 3CNF form, the lower bound is not linear but only $O(\log n)$. Thus, it is an open problem whether or not there are 3CNF formulas with polynomial size resolution proofs, but requiring linear degree Nullstellensatz refutations.

While these results are quite negative, there is a natural generalization of the Nullstellensatz proof system, the Gröbner proof system.[3] A proof in this system is still a linear combination of the initial polynomials; however, now in a degree d proof, one can derive intermediate polynomials of degree at most d, and iteratively obtain the constant 1 by using these intermediate polynomials. (The difference between Gröbner proofs and Nullstellensatz proofs is the degree; any degree d Gröbner proof can be easily converted into a Nullstellensatz proof, but the degree of the Nullstellensatz proof may be larger.) In [6], the Gröbner system is introduced, and a deterministic, polynomial-time procedure is given to find small degree Gröbner proofs; moreover this algorithm is used to show that tree-like resolution proofs can be efficiently simulated by small-degree Gröbner proofs. It is an open question whether or not small degree Gröbner proofs can polynomially simulate (unrestricted) resolution proofs. [4] also discusses properties of Gröbner proof systems.

We are also interested in the extent to which strong lower bounds for the Nullstellensatz proof system can be used to obtain lower bounds for Frege systems. In fact, the Nullstellensatz proof system was originally defined as a means to obtaining lower bounds for bounded-depth Frege proofs with parity gates ($AC^0[2]$-proofs.) It is not too hard to see that small degree Nullstellensatz proofs as well as small degree Gröbner proofs can be efficiently simulated by $AC^0[2]$ proofs. Thus, superpolynomial lower bounds for for $AC^0[2]$ proofs imply good lower bounds for small degree Gröbner proofs. In fact, [4] show that superpolynomial lower bounds for AC^0_2 proofs imply lower bounds for small degree Gröbner proofs with a constant number of levels of extension axioms. Therefore, the next obvious step is to obtain nonconstant degree lower bounds for Gröbner proofs.

Acknowledgements. The authors wish to thank Paul Beame for discussions during the course of preparing this paper.

References

[1] P. Beame, S. Cook, J. Edmonds, R. Impagliazzo, and T. Pitassi, *The relative complexity of NP search problems*, in Proceedings of the 27th ACM Symposium on Theory of Computing, 1995, pp. 303–314.

[2] P. Beame, R. Impagliazzo, J. Krajíček, T. Pitassi, and P. Pudlák, *Lower bounds on Hilbert's Nullstellensatz and propositional proofs*, in Thirty-fifth Annual Symposium on Foundations of Computer Science, IEEE Press, 1994, pp. 794–806. Revised version to appear in Proceedings of the London Mathematical Society.

[3] The Gröbner proof system is variously called the "equational calculus", the "polynomial calculus" or the"iterated ideal proof system".

[3] W. D. Brownawell, *Bounds for the degrees in the Nullstellensatz*, Annals of Mathematics (second series), 126 (1987), pp. 577–591.

[4] S. R. Buss, R. Impagliazzo, J. Krajíček, P. Pudlák, A. A. Razborov, and J. Sgall, *Proof complexity in algebraic systems and constant depth Frege systems with modular counting.* To appear in *Computational Complexity*, 199?

[5] L. Caniglia, A. Galligo, and J. Heintz, *Some new effectivity bounds in computational geometry*, in Applied Algebra, Algebraic Algorithms and Error Correcting Codes, Proceedings, Sixth International Conference, Rome 1988, Lecture Notes in Computer Science #357, T. Mora, ed., Berlin, 1989, Springer-Verlag, pp. 131–151.

[6] M. Clegg, J. Edmonds, and R. Impagliazzo, *Using the Groebner basis algorithm to find proofs of unsatisfiability*, in Proceedings of the Twenty-eighth Annual ACM Symposium on the Theory of Computing, Association for Computing Machinery, 1996, pp. 174–183.

[7] A. Haken, *The intractability of resolution*, Theoretical Computer Science, 39 (1985), pp. 297–308.

[8] J. Kollár, *Sharp effective Nullstellensatz*, Journal of the American Mathematical Society, 1 (1988), pp. 963–975.

[9] T. Pitassi, *Algebraic propositional proof systems*, in Descriptive Complexity and Finite Models, N. Immerman and P. Kolaitis, eds., DIMACS Series in Discrete Mathematics and Theoretical Computer Science #31, American Mathematics Society, 1996, pp. 215–244.