

New lower bounds for Polynomial Calculus over non-Boolean bases

Yogesh Dahiya ✉ 

The Institute of Mathematical Sciences (A CI of Homi Bhabha National Institute), Chennai, India

Meena Mahajan ✉ 

The Institute of Mathematical Sciences (A CI of Homi Bhabha National Institute), Chennai, India

Sasank Mouli ✉

Hyderabad

Abstract

In this paper, we obtain new size lower bounds for proofs in the Polynomial Calculus (PC) proof system, in two different settings.

- When the Boolean variables are encoded using ± 1 (as opposed to $0, 1$): We establish a lifting theorem using an asymmetric gadget G , showing that for an unsatisfiable formula F , the lifted formula $F \circ G$ requires PC size $2^{\Omega(d)}$, where d is the degree required to refute F . Our lower bound does not depend on the number of variables n , and holds over every field. The only previously known size lower bounds in this setting were established quite recently in [Sokolov, STOC 2020] using lifting with another (symmetric) gadget. The size lower bound there is $2^{\Omega((d-d_0)^2/n)}$ (where d_0 is the degree of the initial equations arising from the formula), and is shown to hold only over the reals.
- When the PC refutation proceeds over a finite field \mathbb{F}_p and is allowed to use extension variables: We show that there is an unsatisfiable formula with N variables for which any PC refutation using $N^{1+\epsilon(1-\delta)}$ extension variables, each of arity at most $N^{1-\epsilon}$ and size at most N^ϵ , must have size $\exp(\Omega(N^{\epsilon\delta}/\text{poly log } N))$. Our proof achieves these bounds by an XOR-ification of the generalised PHP $_n^{m,r}$ formulas from [Razborov, CC 1998].
The only previously known lower bounds for PC in this setting are those obtained in [Impagliazzo-MouliPitassi, CCC 2023]; in those bounds the number of extension variables is required to be sub-quadratic, and their arity is restricted to logarithmic in the number of original variables. Our result generalises these, and demonstrates a tradeoff between the number and the arity of extension variables.

2012 ACM Subject Classification Theory of computation \rightarrow Proof complexity

Keywords and phrases Proof Complexity; Polynomial Calculus; degree; Fourier basis; extension variables

1 Introduction

Propositional proof complexity is the field of study of the complexity of proofs for tautological Boolean formulae. Cook and Reckhow [5] introduced this area in their seminal work with the ultimate goal of resolving the question of NP versus coNP using upper/lower bounds for stronger and stronger proof systems. Polynomial Calculus (PC) is one such propositional proof system that has received wide attention since its introduction by Clegg, Edmonds and Impagliazzo [4]. Degree lower bounds for PC and its variant PCR (PC with Resolution) have been proved starting with the work of Razborov [14], followed by a long series of works [10], [3], [1], [12], [6]. These translated to size lower bounds through a size-degree connection established in [4], [10]. Despite these works showing that we have a reasonable understanding of PC, there has been little progress towards lower bounds for the stronger

system $\text{AC}^0[p]$ -Frege, which was one of the main motives for defining PC. Therefore, this indicates that we have to look at systems stronger than PC in order to get new insights.

Grigoriev and Hirsch [7] introduced one such system called constant-depth PC, where in addition to the rules of PC we allow extension variables of constant depth to be introduced and used as new variables. They showed that this system simulates $\text{AC}^0[p]$ -Frege (at a proportional depth), making it a suitable target for lower bounds. However, Raz and Tzameret [13] showed that this system is already powerful enough to simulate the proof system Cutting Planes (which deals with integer linear inequalities) with polynomially bounded coefficients. Finally, Impagliazzo, Mouli and Pitassi [8] showed that this system can simulate at a fixed constant depth $\text{AC}^0[q]$ -Frege for any prime modulus q , Cutting Planes and the semialgebraic proof system sum-of-squares SOS with unbounded coefficients, and can also simulate TC^0 -Frege at a proportional depth. This implies that general lower bounds for this system are much harder than lower bounds for $\text{AC}^0[p]$ -Frege.

The simplest subsystem of the above for which size lower bounds were unknown until recently is PC over ± 1 variables (instead of $\{0, 1\}$). Degree lower bounds hold here as well, since a linear transformation can effect the change between the bases ± 1 and $\{0, 1\}$. However, known techniques based on the size-degree connection fail over this basis since they rely on terms vanishing when variables are set to zero. Moreover, the Tseitin tautologies require large PC degree but have small PC size over ± 1 , precluding the existence of such a generic connection. Sokolov [15] recently showed how to go past this barrier and proved size lower bounds for PC over ± 1 variables. Impagliazzo, Mouli and Pitassi [9] improved and generalized these bounds to PC over finite fields \mathbb{F}_p with a sub-quadratic number of extension variables, where each extension variable depends on $O(\log n)$ original variables.

Our Results.

In this work, we extend the results of both these works [15, 9]. For our first result, we show a generic degree-to-size lifting result for PC lower bounds over ± 1 .

► **Theorem 1.1.** *Let F be an unsatisfiable formula over variables $x_1 \cdots x_n$, with a polynomial encoding of degree d_0 , which requires degree $d > d_0$ to refute in PC. Let Ind denote the one-bit indexing gadget. Let $F \circ \text{Ind}$ be the formula obtained by replacing each x_i by $\text{Ind}(w_{i0}, w_{i1}, w_{i2})$ for a fresh set of variables $w_{i0}, w_{i1}, w_{i2} \in \{\pm 1\}$. Then $F \circ \text{Ind}$ requires size $2^{\Omega(d)}$ to refute in PC (over ± 1).*

Sokolov showed such a lifting result for the SOS proof system, using a symmetric gadget with certain properties (e.g. majority). Since SOS simulates PC over the reals [2], he obtains a degree-to-size lift for PC over \mathbb{R} . However, we use a different (non-symmetric) gadget, and our result holds over any field \mathbb{F} . Moreover, Sokolov's lower bounds are of the form $2^{\Omega((d-d_0)^2/n)}$ (where n is the number of variables, and d_0 is the degree of the initial equations arising from the formula). This gives meaningful results only when $d = \omega(\sqrt{n})$. In contrast, our result works for any superconstant d . Lastly, we believe that our result is arguably simpler to prove. This result is motivated by the work of Krause and Pudlák [11].

In our next result, we strengthen the lower bounds of [9] over finite fields to handle a sub-quadratic number of extension variables, each of which is of polynomial size and depends on a polynomial fraction of the original variables. This is an exponential improvement over the result of [9] in the *arity* of the extension variables, i.e. the number of original variables they can depend on (subject to the condition that the extension axioms are small).

► **Theorem 1.2.** *For every $N > 0$ large enough, any $1 > \epsilon, \delta > 0$, constant $c > 0$, and prime p , there exists a tautology F over N variables such that any PC refutation of F over \mathbb{F}_p with*

$N^{1+\epsilon(1-\delta)}$ extension variables, each depending on $N^{1-\epsilon}$ variables of F and of size at most N^c , requires size $\exp(\Omega(N^{\epsilon\delta}/\text{poly log } N))$.

Our Techniques.

In both our results, the notion of quadratic degree introduced in [15] plays a crucial role. This is the maximum degree that can be obtained from a PC refutation by multiplying any two terms that appear in the same line; see Definition 3.2. Sokolov’s insight was that it is easier to reason about this measure for a refutation, and at the same time, it does carry information about the usual degree. In particular, a refutation with low quadratic degree can be transformed into one with low degree; Lemma 3.7. An adaptation of this measure and this method was subsequently used in [9] to similarly reason about refutations using extension variables. A non-trivial part in both these is establishing that if a variable appears in the refutation but not in any axiom (other than the Boolean axiom or an extension axiom of a specific type), then it can be removed from the proof; Lemmas 3.4 and 4.13. This seems self-evident but needs to be done with care.

To prove our first result Theorem 1.1, we follow Sokolov’s approach, but we employ a lift by the (asymmetric) one-bit Indexing gadget. The gadget, on variables w_0, w_1, w_2 , selects the value of w_1 if $w_0 = -1$ (in this case, w_2 is irrelevant) and that of w_2 otherwise (now w_1 is irrelevant). If there is a small enough proof, then the probabilistic method guarantees the existence of an assignment to all the selector variables so that under this restriction, every high-degree quadratic term contains an irrelevant variable. The restricted refutation is in fact a refutation of the unlifted formula, but may use irrelevant variables along the way. Such irrelevant variables can be removed from the restricted proof as discussed above, and then Sokolov’s transformation from low-quadratic-degree to low-degree can be employed, yielding a contradiction.

To prove our second result, we closely follow the approach of [9], which we outline in Section 4. However, to handle extension variables of large arity and polynomial size, we use a similar idea as above; composing a hard tautology with a simple gadget. The gadget we use here is just the XOR₂ gadget (the parity of two variables). We begin by considering family of restrictions where, for each XOR₂ gadget, one of the variables is assigned a random bit while the other variable remains free. Such restrictions recover the hard formula (possibly with some variables negated). Using the probabilistic method once again, it is easy to guarantee the existence of one such restriction under which every extension variable axiom, despite having a large arity, reduces to logarithmic degree. Thus the problem reduces to proving a size lower bound for the original unlifted tautology, where extension axioms are bounded by logarithmic degree. Such extension variables can be handled using the approach of [9] by further carefully chosen restrictions of small size and even smaller Hamming weight; this size lower bound is shown in Theorem 5.5. The hard formula we choose is the generalised Pigeon-Hole-Principle formula PHP _{n} ^{m,r} , introduced by Razborov in [14] (see Proposition 5.2), and we show that applying such small restrictions preserves the degree hardness shown by Razborov.

Organisation of the Paper.

In Section 2, we include the basic relevant definitions of the proof systems. In Section 3 we describe the setting where the encoding is over ± 1 , and prove Theorem 1.1. In Section 4, we describe the setting where extension variables are used, in particular describing all the facts and results from [9] that we crucially use. In Section 5 we prove Theorem 5.5 and

Theorem 1.2.

2 Preliminaries

The Polynomial Calculus proof system is a line-based proof system. It shows that any set of polynomial expressions with no common zero contains 1 in its ideal, by building up an explicit derivation of 1. It can be used to show that a Boolean formula in conjunctive normal form (CNF) is unsatisfiable (equivalently that a disjunctive normal form expression is a tautology) by appropriately encoding the variables and clauses and including expressions that force the variables to take only Boolean values; then common zeroes are exactly satisfying assignments. While it can also be used with polynomials not arising from CNFs, in this work we are only concerned with CNF formulas. We follow the notation from [15, 9].

► **Definition 2.1** (Polynomial Calculus/Polynomial Calculus Resolution). *Let $\Gamma = \{P_1 \dots P_m\}$ be an unsolvable system of polynomials in variables $\{x_1 \dots x_n\}$ over \mathbb{F} . A PC (Polynomial Calculus) refutation of Γ is a sequence of polynomials $\{R_1 \dots R_s\}$ where $R_s = 1$, and for every $\ell \in [s]$, R_ℓ is either a polynomial from Γ , or is obtained from two previous polynomials R_j, R_k , $j, k < \ell$, by one of the following derivation rules:*

$$\begin{aligned} R_\ell &= \alpha R_j + \beta R_k \text{ for } \alpha, \beta \in \mathbb{F} \\ R_\ell &= x_i R_k \text{ for some } i \in [n] \end{aligned}$$

The size of the refutation is $\sum_{\ell=1}^s |R_\ell|$, where $|R_\ell|$ is the number of monomials in the polynomial R_ℓ . The degree of the refutation is $\max_\ell \deg(R_\ell)$.

A PCR (Polynomial Calculus Resolution) refutation is a PC refutation over the set of Boolean variables $\{x_1 \dots x_n, \bar{x}_1 \dots \bar{x}_n\}$ where $\{\bar{x}_1 \dots \bar{x}_n\}$ are twin variables of $\{x_1 \dots x_n\}$. That is, over the $\{0, 1\}$ encoding, the equations $x_i^2 - x_i = 0$, $\bar{x}_i^2 - \bar{x}_i = 0$ and $x_i + \bar{x}_i - 1 = 0$ are treated as axioms. Similarly, over the ± 1 encoding, the equations $x_i^2 - 1 = 0$, $\bar{x}_i^2 - 1 = 0$ and $x_i \bar{x}_i + 1 = 0$ are treated as axioms.

In the literature, the terms PC and PCR are often used interchangeably. The notion of degree is the same in both, but size in PC with the $\{0, 1\}$ encoding of Boolean variables can be much larger than in PCR. Throughout this paper, we say PC but really mean PCR. In particular, our size lower bounds are for PCR.

Note that the minimal degree required to refute a formula is independent of whether Boolean variables are encoded over $\{0, 1\}$ or over ± 1 . However, the minimal size crucially depends on the encoding. As is well known [3], suitable Tseitin formulas require degree n , hence they require size $\exp(\Omega(n))$ over the $\{0, 1\}$ basis using size-degree connection, but have linear-size refutations over ± 1 .

As is standard, throughout this work we assume that we are working in the ideal modulo the Boolean axioms, and hence the polynomials in all lines are multilinear in the original variables. Technically, on deriving a higher degree term, it has to be cancelled by using suitable multiples of the Boolean axiom; however, these steps do not significantly alter the size or degree of the refutation.

► **Definition 2.2** (PC plus Extension Axioms). *Let $\Gamma = \{P_1 \dots P_m\}$ be a set of polynomials in variables $\{x_1 \dots x_n\}$ over a field \mathbb{F} , with no common zero. The polynomials in Γ are referred to as the (initial) axioms. Let $\mathbf{z} = z_1 \dots z_M$ be new extension variables with corresponding extension axioms $z_j - Q_j(x_1 \dots x_n)$. A PC + Ext (PC plus extension) refutation of Γ with M extension axioms $\text{Ext} = \{z_j - Q_j(x_1 \dots x_n) \mid j \in [M]\}$ is a PC refutation of the set of polynomials $\Gamma' = \{P_1 \dots P_m, z_1 - Q_1 \dots z_M - Q_M\}$. The size of the refutation is the total*

size of all lines in the refutation, including the polynomials in Γ' (where the size of a line $P \in \Pi$ is the number of monomials in P). The degree of the refutation is the maximum degree of any line in the refutation or in Γ' .

Similar to [9], our notion of extension variables is not recursive in the sense that new extension variables cannot be defined as functions of existing ones. Our extension variables are only allowed to depend on the original variables of the tautology.

3 PC size lower bounds over ± 1 by lifting with one-bit indexing

In this section we prove that if a tautology F requires PC degree d , then the tautology F' obtained by lifting each variable in F with a one-bit indexing gadget (over a fresh set of variables) requires PC size $2^{\Omega(d)}$ over ± 1 .

► **Definition 3.1** (One-bit indexing gadget). *Let w_0, w_1, w_2 be variables taking values in $\{\pm 1\}$. The function $\text{Ind}(w_0, w_1, w_2)$ is defined as follows: $\text{Ind}(-1, w_1, w_2) = w_1$ and $\text{Ind}(1, w_1, w_2) = w_2$. We call w_0 the selector variable and w_1, w_2 the data variables of the gadget.*

Lifting a Boolean formula F by this gadget for each variable means introducing three fresh variables w_0, w_1, w_2 corresponding to each variable w in F , and replacing each occurrence of w in F with an expression equivalent to $(w_0 \wedge w_1) \vee (\neg w_0 \wedge w_2)$. (This is because in this basis, true is encoded as -1 and false as $+1$.) In a PC formulation, it corresponds to replacing each occurrence of w in an axiom (other than the Boolean axiom) by the expression $((1 - w_0)w_1 + (1 + w_0)w_2)/2$, and adding the Boolean axioms $w_j^2 = 1$ for $j = 0, 1, 2$.

Our idea is to consider refutations of F lifted by the indexing gadget, then apply a restriction to the selector variables that yields a refutation of F with low quadratic degree, and from there obtain a small degree refutation of F . The quadratic degree of a refutation is defined in [15] using the notion of lazy representations of polynomials, and is rephrased here as follows:

► **Definition 3.2** (Quadratic set, Quadratic degree, Quadratic terms over ± 1 ; taken from [15], Section 3.2). *Given a proof Π over ± 1 variables, the Quadratic set of Π , denoted $\mathcal{Q}(\Pi)$, is the set of pairs of terms $\mathcal{Q}(\Pi) = \{(t_1, t_2) \mid t_1, t_2 \in P \text{ for some line } P \in \Pi\}$.*

Denote by $\mathcal{QT}(\Pi)$ the set of quadratic terms $\{t_1 t_2 \mid (t_1, t_2) \in \mathcal{Q}(\Pi)\}$, where the product is modulo the axioms $x_i^2 = 1$.

The Quadratic degree of Π is the max degree of a term in $\mathcal{QT}(\Pi)$.

Informally, Quadratic degree is the max degree of the square of each line (before cancellations).

When we apply the chosen restriction to the selector variables in $F \circ \text{Ind}$, the irrelevant variables no longer appear in any of the axioms (except the axioms $x_i^2 = 1$; we work modulo those anyway). However, they may still appear in the refutation, and we need to eliminate them. For this, we use the Split operation introduced in [15].

► **Definition 3.3** (Split operation over x [15], Section 5.4). *Given a proof $\Pi = (P_1, P_2, \dots, P_t)$ and a variable $x \in \{\pm 1\}$, each line P_i of Π is of the form $P_{i,1}x + P_{i,0}$, where $P_{i,1}, P_{i,0}$ do not contain x . The Split operation at x , denoted by $\text{Split}_x(\Pi)$, is the sequence Π' with the lines $\{P_{1,1}, P_{1,0}, P_{2,1}, P_{2,0}, \dots, P_{t,1}, P_{t,0}\}$.*

The following lemma shows that Split of a refutation is a valid refutation whenever the variable we are splitting on does not appear in any axioms except $x^2 = 1$. (That is, x has no role in the tautology we are considering, but is possibly introduced along the way and then eliminated. The gadget variables rendered irrelevant by our chosen restriction are like this.)

This is in fact a special case of a more general statement shown in [15], and we only need this case. For ease of reading, we include here a proof of just this special case.

► **Lemma 3.4.** *Suppose that Π is a proof and x is a variable that does not appear in any axioms of Π except $x^2 = 1$. Then $\text{Split}_x(\Pi)$ outputs a valid proof of the axioms of Π , with no line containing x .*

Proof. Let Π be the sequence P_1, \dots, P_t . We show by induction on the line number j that both $P_{j,1}$ and $P_{j,0}$ are derivable and x -free.

If P_i is an axiom, then it is free of x . So the Split version is $P_{i,1} = 0$, $P_{i,0} = P_j$, and both these polynomials are derivable.

If $P_i = \alpha P_j + \beta P_k$ for some $j, k < i$, then $P_{i,b} = \alpha P_{j,b} + \beta P_{k,b}$ for $b = 0, 1$.

If $P_i = y P_j$ for some $j < i$ and some variable $y \neq x$, then $P_{i,b} = y P_{j,b}$ for $b = 0, 1$.

If $P_i = x P_j$ for some $j < i$, then since $x^2 = 1$ we obtain $P_{i,1} = P_{j,0}$ and $P_{i,0} = P_{j,1}$.

Thus all the lines $P_{j,b}$ are derivable and do not contain x .

Since the last line of the proof is $P_t = 1$, we have $P_{t,1} = 0$ and $P_{t,0} = P_t = 1$. Thus $\text{Split}_x(\Pi)$ derives 1 and is a valid proof from the axioms of Π . ◀

► **Remark 3.5.** It may help to visualise the Split_x process as follows. Consider the case where the derivation structure underlying Π is tree-like. The tree T is rooted at P_t , and is unary-binary: linear combination nodes have two children and variable-multiplication nodes have one child. The Split_x process makes two nodes P_0, P_1 for each node P of T , and ends up creating a forest with two trees T_0, T_1 . The desired refutation is T_0 , since $P_{t,0} = 1$ whereas $P_{t,1} = 0$. It can be seen that T_0 may also be obtained directly from T as follows: for each axiom node, if the number of edges along the path to the root labeled by multiplication with x is odd, replace the axiom by 0. (The construction above would have ended at a source node which is a P_1 copy of an axiom, and since axioms are x -free, a P_1 copy of an axiom is 0.) Then, replace each edge label $\times x$ by the label $\times 1$.

The lemma below shows that Split_x removes all quadratic terms containing x from the proof, without introducing any new quadratic terms.

► **Lemma 3.6.** *Let $\mathcal{Q}_x(\Pi)$ be the set of pairs $(t_1, t_2) \in \mathcal{Q}(\Pi)$ such that $x \in t_1 t_2$, and let $\mathcal{QT}_x(\Pi)$ be the corresponding set of quadratic terms.*

If $(t_1, t_2) \in \mathcal{Q}(\text{Split}_x(\Pi))$, then t_1 and t_2 are both x -free, and at least one of (t_1, t_2) , $(t_1 x, t_2 x)$, is in $\mathcal{Q}(\Pi)$. Thus $\mathcal{QT}(\text{Split}_x(\Pi)) \subseteq \mathcal{QT}(\Pi) \setminus \mathcal{QT}_x(\Pi)$.

Proof. Consider a pair $(t_1, t_2) \in \mathcal{Q}(\text{Split}_x(\Pi))$. That t_1, t_2 are x -free follows from Lemma 3.4. The pair (t_1, t_2) is contributed to $\mathcal{Q}(\text{Split}_x(\Pi))$ by P_b for some line $P = x P_1 + P_0$ of Π and some $b \in \{0, 1\}$. If P_0 contributes the pair, then P also contributes the pair to $\mathcal{Q}(\Pi)$. If P_1 contributes the pair, then P contributes the pair $(t_1 x, t_2 x)$ to $\mathcal{Q}(\Pi)$. ◀

Finally, we note below that a proof with low quadratic degree can be transformed into a proof of low (usual) degree. This lemma is proved in [15], Lemma 3.6 using the notion of lazy representation of polynomials. For completeness, we include in the appendix a very similar proof but without explicitly using this notion.

► **Lemma 3.7.** *[[15], Lemma 3.6] Let Π be a refutation of a set of axioms F of degree d_0 with Quadratic degree at most d . Then there exists a refutation Π' of F with (usual) degree at most $2 \max(d, d_0)$.*

We conclude this section with a proof of Theorem 1.1, which we restate here for convenience.

► **Theorem 1.1.** *Let F be an unsatisfiable formula over variables $x_1 \cdots x_n$, with a polynomial encoding of degree d_0 , which requires degree $d > d_0$ to refute in PC. Let Ind denote the one-bit indexing gadget. Let $F \circ Ind$ be the formula obtained by replacing each x_i by $Ind(w_{i0}, w_{i1}, w_{i2})$ for a fresh set of variables $w_{i0}, w_{i1}, w_{i2} \in \{\pm 1\}$. Then $F \circ Ind$ requires size $2^{\Omega(d)}$ to refute in PC (over ± 1).*

Proof. Let F' denote the formula $F \circ Ind$.

Towards a contradiction, let Π be a refutation of F' of size 2^{cd} for a small enough $c > 0$. An assignment ρ_i to the selector variable w_{i0} sets the gadget to one of the two data variables w_{ij} , $j \in \{1, 2\}$; we say that the other data variable is irrelevant. (If ρ_i sets $w_{i0} = -1$ then w_{i2} is irrelevant, else w_{i1} is irrelevant.) We construct an assignment ρ to the selector variables such that for every pair $(t_1, t_2) \in \mathcal{Q}(\Pi|_\rho)$ with $\deg(t_1 t_2) \geq d/2$, $t_1 t_2$ contains an irrelevant variable. The rest of the proof is simple: we apply this ρ to Π to obtain a refutation of a copy of F without irrelevant variables. However, the irrelevant variables may still appear in the proof. We then repeatedly apply Split over each irrelevant variable, to obtain a refutation Π' of F with no irrelevant variables anywhere. (By Lemma 3.4, the result of the Split is a valid refutation.) Since every high-degree pair contains an irrelevant variable, and by Lemma 3.6 all pairs where the product contains a irrelevant variable are removed from the proof, Π' does not contain any high-degree pair and hence is of Quadratic degree less than $d/2$. We use Lemma 3.7 to get a refutation of degree less than d of F , contradicting our assumption about F .

We now show the existence of ρ through a probabilistic argument. Let $t_1 t_2 \in \mathcal{QT}(\Pi)$ with degree in the data variables at least $d/2$. If, for some gadget, this product contains both data variables, then for any assignment ρ this product would contain an irrelevant variable. So without loss of generality we can assume that the product contains only one data variable from every copy of the gadget. Now, we pick a ρ which is a uniformly random assignment from $\{\pm 1\}^n$; i.e. pick the data variable at random in each gadget. For a data variable in $t_1 t_2$, the probability that it is picked is equal to $1/2$. Therefore, the probability that $t_1 t_2$ does not contain any irrelevant variable is at most $(1/2)^{d/2}$. Since there are only 2^{cd} terms and therefore 2^{2cd} pairs in the proof, we have by a union bound that there exists a restriction with the required property. ◀

We note here that a proof similar to above also works for the SOS proof system, where given a refutation $\sum_i p_i f_i + \sum_j q_j^2 = -1$ of axioms $f_i = 0$, the operation Split at an irrelevant variable x is defined (following [15]) as the refutation obtained by averaging the values of p_i and q_i^2 at $x = 1$ and $x = -1$ (this is a valid refutation since x is irrelevant). As a result, we also obtain a degree-to-size lifting theorem for SOS for the one-bit indexing gadget.

4 PC with extension variables over finite fields

We now consider the setting where the encoding is over $\{0, 1\}$, the arithmetic is over finite fields, and extension variables are allowed; this is the setting for our second main result Theorem 1.2. In this setting, a size lower bound was obtained in [9] provided the extension variables are subquadratic in number and at most logarithmic in arity. We follow that approach but improve the result substantially. In this section, we first outline the framework of [9], then describe at a high level the outline of our proof of Theorem 1.2, and then present the relevant definitions/lemmas from [9] that we need to use. The actual formal proof of Theorem 1.2 appears in the next section.

4.1 Proof outline

The approach in [9]

We first outline the framework of [9], whose lower bounds we improve. The proof of the lower bound in [9] proceeds as follows.

Given a small refutation of a well chosen tautology F in PC with extension variables, pick an extension variable z with extension axiom $z - Q$ that contributes to a lot of pairs of terms of high Quadratic degree (which is a notion similar to Quadratic degree for ± 1 variables as in Definition 3.2, but generalized to \mathbb{F}_p -valued variables; see Definition 4.5). Extension variables are not necessarily Boolean; z can take a subset of values in the underlying field (over all possible values to the Boolean variables in Q). If this subset includes zero, apply the partial assignment that sets $z = 0$ to the proof to remove all contributions of z to Quadratic degree.

If not, z appears in each line of the proof in the form $P_{\ell-1}z^{\ell-1} + \dots + P_1z + P_0$ where ℓ is the least value such that z^ℓ is a constant. The contributions of z to Quadratic degree therefore come from interactions of the polynomials $P_i z^i$ and $P_j z^j$, over all pairs (i, j) , $i \neq j$. Now pick a good pair (i, j) which contributes at least a $1/p^2$ fraction of the contributions of z to high Quadratic degree. The key step is to obtain a proof which separates the pair of polynomials P_i and P_j in each line into two different lines, using an operation called Split, see Definition 4.12. (Again, this is similar in spirit to the Split operation from Definition 3.3, but more nuanced.) Split essentially equates each line P to a polynomial of the form $R_1 z^i + R_0 z^j$, and solves for R_1 and R_0 in terms of P . In order for Split to output a valid proof, though, some preconditions need to be satisfied: the axioms need to be free of z except for the range axiom for z , and this range needs to be such that z^i and z^j are linearly independent, i.e. $z^i \neq cz^j$ or $z^{i-j} \neq c$. That is, z needs to take on at least two values a, b such that $a^{i-j} \neq b^{i-j}$. Therefore there are two tasks at hand: getting rid of the extension axiom $z - Q$, and doing it in such a way that z^{i-j} is not set to a constant. It is shown that a restriction to Q can be chosen that sets it to the form $(b - a)x + a$, with a, b satisfying the precondition for Split. Once this happens, Split is applied to reduce a fraction of high degree terms (after applying an additional restriction to make sure x does not occur in the axioms, and then setting $x = (b - a)^{-1}(z - a)$ in order to get rid of the extension axiom).

This process is repeated until the proof is of low Quadratic degree. Then an argument from [15] (adapting Lemma 3.7 to the extension-variables setting) is used to move to a low (usual) degree proof of the tautology $F|_{\rho}$, where ρ is the union of all restrictions ρ_i applied in this process. This contradicts the degree lower bound for $F|_{\rho}$. An important element of this proof is to ensure that none of the restrictions ρ_i s make the tautology easy. To ensure this, cleanup operations are performed at each iterations using additional restrictions, to restore to a sub-copy of F where hardness is preserved. These operations work correctly provided each extension variable depends on only $O(\log n)$ of the n variables of F . Also the number of iterations of this process has to be bounded as well, which gives an upper bound of $o(n^2)$ on the number of extension variables in order to get mildly exponential lower bounds.

Our proof outline

We largely follow the above approach for our lower bound, but show that sub-quadratically many extension variables with extension axioms that are polynomial sized and depend on a polynomial fraction of original variables can be handled.

For this we first reduce the above problem to handling low degree extension variables, by a simple parity lift. Let F be a tautology and let F' be obtained by replacing each input variable x of F by a two bit parity gadget, i.e. $x = w_1 \oplus w_2$. Suppose that there is a PC proof

of F' which uses extension variables of size bounded by N^c . In each copy of the gadget we select one variable at random and set the other variable to zero or one with equal probability, recovering a copy of F (possibly renaming some variables by their negations). It is easy to see using a probabilistic argument that there exists a restriction where each extension variable is of degree at most $O(\log N)$.

Thus the problem reduces to proving lower bounds for F where extension axioms $z - Q$ are degree bounded. We then show that when dealing with such extension variables, the aforementioned process that picks an extension variable z and either sets it to zero or restricts it to the form $(b - a)x + a$ can be performed with restrictions whose hamming weight is bounded by the degree of Q .

Finally, we observe that the tautology $\text{PHP}_n^{m,r}$ introduced by Razborov [14], which maps all r sized subsets of $[m]$ to n holes, is immune to such restrictions and can be cleaned up to restore hardness. This completes our proof.

For the rest of the article, we fix the finite field \mathbb{F}_p , $p > 2$ (since for the case of $p = 2$ lower bounds for PC with extension variables can be obtained through standard size-degree tradeoffs, see the discussion in [9], Section 3).

4.2 Relevant material from [9]: Support, Quadratic Degree, Split

Here we introduce the terminology of [9] and state some lemmas about Quadratic degree and Split from the same. The reader familiar with [9] can directly jump to Section 5.

As noted above, the notions of Quadratic degree and Split here are not the same as in Definitions 3.2 and 3.3. To define them appropriately in this setting, we first need some auxiliary notions.

► **Definition 4.1** (Support of a variable, Singular/Nonsingular variables. [9], Defs 10,11). *Let $z - Q(w_{i_1}, \dots, w_{i_\kappa}) = 0$ be an extension axiom associated with z . The set $\text{vars}(Q)$ is defined as $\text{vars}(Q) = \{w_{i_1}, \dots, w_{i_\kappa}\}$, and is sometimes also written as $\text{vars}(z)$, the set of variables that z depends on.*

The support of z , $\text{supp}(z) \subseteq [0, p-1]$, is equal to the set of all values that z can take under Boolean assignments to $\text{vars}(z)$. That is, $\text{supp}(z) = \{Q(\alpha) \mid \alpha \in \{0, 1\}^{|\text{vars}(Q)|}\} \subseteq [0, p-1]$. Sometimes this is also denoted by $\text{supp}(Q)$.

We say that z is a Singular variable if $0 \in \text{supp}(z)$, otherwise it is NonSingular.

For a Boolean variable w , $\text{supp}(w) = \{0, 1\}$ as enforced by the Boolean axiom $w^2 = w$.

As we apply restrictions to a proof (and hence to all the defining axioms), 0 may get removed from the support of a variable. Thus an extension variable can change from Singular to NonSingular, but not the other way around. However, Boolean variables that are not set by the restriction are always Singular.

► **Definition 4.2** ([9], Definition 12). *Let $A \subseteq [1, \dots, p-1]$, $A \neq \emptyset$. Define $\ell(A)$ to be the least $\ell \in [1, p-1]$ such that the set $\{a^\ell \mid a \in A\}$ is singleton. For a Nonsingular z , define $\ell(z) = \ell(\text{supp}(z))$.*

The following lemma from [9] is stated without proof.

► **Lemma 4.3** ([9], Lemma 13). *Let z be a Nonsingular extension variable with extension axiom $z - Q = 0$. Then the following polynomial equations are implied by (and therefore derivable from) the extension axiom for z plus the Boolean axioms for all variables in $\text{vars}(Q)$, in degree at most $|\text{vars}(Q)|$.*

1. $z - Q' = 0$, where Q' is the multilinear version of Q ;

2. For any $A' \subseteq [0, p-1]$ such that $\text{supp}(z) \subseteq A'$, $\Pi_{a \in A'}(z - a) = 0$;
3. $z^{\ell(z)} - c = 0$ for some $c \in \mathbb{F}_p^*$.

In particular, if z is Nonsingular, then the polynomial equation $z^{p-1} - 1 = 0$ is implied by $z - Q = 0$ together with the Boolean axioms for $\text{vars}(Q)$.

► **Definition 4.4** ([9], Definition 14). For a term t and a variable w , $\deg(t, w)$ is equal to the degree of w in t . Note that since we are working over \mathbb{F}_p , $\deg(t, w) < p$ for any variable w . For a term t , the degree of t , denoted $\deg(t)$, equals $\sum_{w \in \text{vars}(t)} \deg(t, w)$.

Quadratic degree

The following definition of Quadratic degree is taken from [9].

► **Definition 4.5** (Quadratic degree [9], Definition 10). Let V be a set of variables and let S be a subset of V . For a pair of terms t_1, t_2 over V , and a variable $w \in V$, we define $Qdeg^S(t_1, t_2, w)$ as follows. If $w \in S$, then $Qdeg^S(t_1, t_2, w) = 1$ if w occurs in at least one of t_1 or t_2 ; if $w \notin S$, then $Qdeg^S(t_1, t_2, w) = 1$ if and only if $\deg(t_1, w) \neq \deg(t_2, w)$. The overall quadratic degree of the pair t_1, t_2 , $Qdeg^S(t_1, t_2)$, is equal to $\sum_{w \in V} Qdeg^S(t_1, t_2, w)$. The quadratic degree of a polynomial P is equal to the maximum quadratic degree over all pairs (t_1, t_2) such that $t_1, t_2 \in P$. For a proof Π , the quadratic degree of Π is the maximum quadratic degree over all polynomials $P \in \Pi$.

► **Remark 4.6.** The above definition of Quadratic degree treats the variables in the set S differently from the rest of the variables. Typically S will not be explicitly specified, but will be assumed to be the set of Singular variables. This means that the notion of Quadratic degree depends on knowing which variables have zero in their support. For instance, for the pair $(z_1, z_1 z_2)$, the Quadratic degree is two if z_1 can take the value zero and one if z_1 does not take the value zero. We go on to observe that Quadratic degree always decreases when a variable changes from being Singular to Nonsingular, and make sure that this is the case when we prove our lower bound. This follows the approach of [9] whose lemmas we state below.

► **Lemma 4.7** ([9], Lemma 16). Let V be a set of variables and let S and T be subsets of V such that $T \subseteq S$. Then for any two terms t_1, t_2 over V , $Qdeg^T(t_1, t_2) \leq Qdeg^S(t_1, t_2)$.

Since applying a restriction cannot make NonSingular variables Singular, Lemma 4.7 implies that the Quadratic degree of any two terms t_1, t_2 , $Qdeg(t_1, t_2)$ with respect to the currently Singular variables cannot increase after applying the restriction. This is stated as Corollary 17 in [9].

► **Lemma 4.8** ([9], Lemma 20). Let Π be a PC + Ext refutation of F and let z be a Nonsingular variable. Let Π' be the proof obtained from Π by reducing each line of Π by $z^{\ell(z)} - c = 0$ for some $c \in \mathbb{F}_p^*$. Then for any $d \geq 0$, the number of pairs of terms of Quadratic degree at least d in Π' is at most that of Π .

We will use the following lemma from [9], which is a generalization of the argument from [15] that shows how to convert a proof with low Quadratic degree to one with low degree.

► **Lemma 4.9** ([9], Lemma 21). Let F be a set of unsatisfiable polynomials of degree d_0 with a PC refutation of Quadratic degree at most $d \geq d_0$ over \mathbb{F}_p . Then F has a PC refutation of degree at most $3pd$.

The Split operation

In this section, we define the operation Split and state its properties. We will only need to handle variables whose only axiom is $(z - a)(z - b) = 0$ for $a, b \in \mathbb{F}_p^*$, as we will apply an assignment to any general extension variable to reduce to this case. Below we state the relevant lemmas from [9].

► **Lemma 4.10** ([9], Lemma 23). *Let z be an extension variable such that $\text{supp}(z) = \{a, b\}$, where $a \neq b$ and $a, b \in \mathbb{F}_p^*$ and let P be any polynomial. Then, for any two distinct numbers i, j where $i < j$ and $a^{j-i} \neq b^{j-i}$, there exists a unique polynomial $R = R_0z^i + R_1z^j$ such that $R = P \pmod{(z - a)(z - b)}$.*

► **Remark 4.11.** It can be checked that for polynomial $P = \sum_{l < \ell(z)} P_l z^l$, the polynomials R_0, R_1 have the following form:

$$R_0 = P_i + \sum_{l < \ell(z), l \neq i, j} c_{0l} P_l$$

$$R_1 = P_j + \sum_{l < \ell(z), l \neq i, j} c_{1l} P_l$$

for some constants $c_{1i}, c_{0i} \in \mathbb{F}_p$. Note that any pair of terms (t_1, t_2) occurring in either R_1 or R_0 also occurs in P as $(t_1 z^{i'}, t_2 z^{j'})$ with $(i', j') \neq (i, j)$. That is, the contribution to Quadratic degree of P by the interaction of z^i and z^j is removed.

► **Definition 4.12** (Split [9], Definition 24). [9] *Let z be an extension variable with extension axiom $z - Q = 0$ such that $\text{supp}(z) = \{a, b\} \subseteq [1, \dots, p - 1]$. For any polynomial P and for every $i < j$ such that $a^{j-i} \neq b^{j-i}$, let $R = R_0z^i + R_1z^j$ be the unique polynomial given by Lemma 4.10 such that $R = P \pmod{(z - a)(z - b)}$. Then $\text{Split}_{z, i, j}(P)$ is defined to be the pair of polynomials $\{R_0, R_1\}$. For a proof Π , and an extension variable z such that $\text{supp}(z) = \{a, b\}$, we define $\text{Split}_{z, i, j}(\Pi)$ to be the sequence of lines $\text{Split}_{z, i, j}(P)$, over all $P \in \Pi$.*

► **Lemma 4.13** ([9], Lemma 25). *Let Π be a refutation of a set of unsatisfiable polynomials F . Let z be a variable that occurs in Π such that the polynomials in F do not contain z except for the axiom $(z - a)(z - b) = 0$ for some $a, b \in \mathbb{F}_p^*$. Then for any i, j such that $i < j$ and $a^{j-i} \neq b^{j-i}$, $\Pi' = \text{Split}_{z, i, j}(\Pi)$ forms a valid refutation of F modulo $(z - a)(z - b)$.*

5 Proof of the lower bound from Theorem 1.2

5.1 The tautology

We use the $\text{PHP}_n^{m, r}$ tautology defined in [14]; it is a variant of the Pigeonhole principle. In this variant, there are m “fractional” pigeons, r fractional parts add up to a whole “pigeon”, and there are n holes. The Boolean variables determine which part goes into which hole. A fractional part can participate in multiple pigeons, and can be assigned to multiple holes. The constraints enforce that no two r -sized subsets are mapped to the same hole, and no r -sized subset is mapped to more than one hole. When $\binom{m}{r} > n$, this is unsatisfiable. We describe the formula formally below.

► **Definition 5.1** ($\text{PHP}_n^{m, r}$; Def 4.1 in [14]). *Let $m, n, r > 0$ be such that $\binom{m}{r} > n$. Let x_{ij} , for $i \in [m]$, $j \in [n]$, be variables that indicate the mapping of elements of $[m]$ to holes in $[n]$.*

For a subset I of $[m]$, abbreviate the term $\prod_{i \in I} x_{ij}$ to t_{Ij} ; note that t_{Ij} is only shorthand and not a variable in the formula. Then $\text{PHP}_n^{m,r}$ is the following set of equations.

$$\begin{aligned} t_{I_1} + t_{I_2} + \cdots + t_{I_n} &= 1 & \forall I \subset [m], |I| = r \\ t_{I_j} &= 0 & \forall I \subset [m], |I| = r + 1; \forall j \in [n] \\ t_{I_{j_1}} t_{I_{j_2}} &= 0 & \forall I \subset [m], |I| = r; \forall j_1, j_2 \in [n], j_1 \neq j_2. \end{aligned}$$

(Note that the last set of constraints is already implied by the first two constraint sets. It is nonetheless included, in [14], where a degree lower bound is shown even when these constraints are explicitly given and do not have to be derived.)

We state the lower bound from [14] on degree of PC proofs for $\text{PHP}_n^{m,r}$.

► **Proposition 5.2** ([14], Theorem 4.2). *For any ground field \mathbb{F} and any $m, r, n > 0$ such that $\binom{m}{r} > n$, $\text{PHP}_n^{m,r}$ requires proofs of degree $n/2 + 1$ to refute in PC over \mathbb{F} .*

5.2 The lower bound

We begin by showing in Theorem 5.5 a weak size lower bound for the $\text{PHP}_n^{m,r}$ formulas in PC when extension variables are allowed, provided the degree of the extension axioms is bounded. To establish the strong lower bound in Theorem 1.2, as discussed in Section 4.1, we show that a lift with the parity gadget (an XOR-ification of the formula), followed by a well-chosen restriction, achieves a degree-reduction of the extension axioms, and use Theorem 5.5.

The weak size lower bound of Theorem 5.5 also uses degree reduction, but it reduces the quadratic degree of the proof. A crucial ingredient in the quadratic-degree-reduction step is finding low-Hamming-weight assignments with certain nice properties. We first prove the existence of these assignments, and then establish the weak lower bound.

► **Lemma 5.3.** *Let z be an extension variable with the extension axiom $z - Q$ and let $l < p$ be a constant such that Q^l is not a constant, i.e. $\text{supp}(Q^l)$ is not singleton. Then there exists a partial assignment σ of Hamming weight at most $l \deg(Q)$, such that for some $x \in \text{vars}(z)$, $Q|_\sigma = (b - a)x + a$ for some a, b with $a^l \neq b^l$.*

Proof. Let X be the set of variables of Q and let x be a variable that appears in Q^l . Since Q^l is not a constant, such a variable always exists. Fix a total ordering over monomials in X that respects degree. Let \mathcal{M} be the least monomial in Q^l according to this ordering that contains x . Let σ be an assignment to $X \setminus \{x\}$ obtained as follows. We set every variable in \mathcal{M} other than x to one, and every other variable in $X \setminus \{x\}$ to zero. Note that this sets every monomial lesser than \mathcal{M} to a constant since it does not contain x . The same is true for monomials greater than \mathcal{M} that do not contain x . Since the ordering respects degree and \mathcal{M} is minimal according to it, any monomial greater than \mathcal{M} that contains x also contains at least one variable that is not in \mathcal{M} and hence is set to zero by σ . Therefore, \mathcal{M} is the only monomial containing x in Q^l that survives the restriction under σ ; thus $(Q^l)|_\sigma = \alpha x + \beta$ for some $\alpha \neq 0$. Since $(Q^l)|_\sigma = (Q|_\sigma)^l$, it follows that $Q|_\sigma$ has the form $(b - a)x + a$, where $a^l = \beta \neq \alpha + \beta = b^l$. By our choice of σ , it has Hamming weight at most the degree of the monomial \mathcal{M} , which is bounded above by $\deg(Q^l) \leq l \deg(Q)$. ◀

This lemma will enable us to satisfy the necessary precondition for applying a Split operation on the variable z , when required.

► **Corollary 5.4.** *Suppose that z is an extension variable with the extension axiom $z - Q$, where $0 \in \text{supp}(Q)$. Then there exists an assignment of Hamming weight at most $p \deg(Q)$ which sets Q to zero.*

Proof. If $\text{supp}(Q) = \{0\}$, then any assignment to $\text{vars}(Q)$ will do. If we are in the nontrivial case where $\text{supp}(Q)$ is not a singleton, then for every $l \in [p-1]$, Q^l is not a constant, and thus we can choose $l = p-1$ in Lemma 5.3. We thus obtain a partial assignment σ of Hamming weight $(p-1)\deg(Q)$ such that $Q|_\sigma = (b-a)x + a$, with $(Q|_\sigma)^{p-1}$ not a constant; i.e. $a^{p-1} \neq b^{p-1}$. This means that exactly one of a, b is zero. Setting the value σ_x to be 0 if $a = 0$ and 1 otherwise, we see that $Q|_{\sigma \cup \sigma_x} = 0$. The Hamming weight of $\sigma \cup \sigma_x$ is at most $(p-1)\deg(Q) + 1 \leq p\deg(Q)$. \blacktriangleleft

We now state and prove our main theorem.

▶ Theorem 5.5. *Let Π be a PC refutation of $\text{PHP}_n^{m,r}$ with M extension variables, each of degree $\leq k$ and depending on $\leq \kappa$ variables of $\text{PHP}_n^{m,r}$, such that $r > 2pk$. Then the size of Π is at least $\exp(\Omega(n^2/(M+mn)\kappa k))$.*

Proof. Let s be the size of the given refutation Π .

For a threshold d that we will choose later and depends on s , we will first show how to reduce the Quadratic degree to at most d . This will be achieved by finding a suitable restriction, in stages, that kills all quadratic terms of quadratic degree more than d . In the process, the restricted formula will become $\text{PHP}_{n'}^{m,r}$ for some $n' \in \Theta(n)$. Using Lemma 4.9 we will convert this to a PC proof of $\text{PHP}_{n'}^{m,r}$ of degree at most $3pd$ but with extension variables, and then by directly substituting the extension axioms, to a PC proof of degree at most $3pkd$ without extension variables. Finally, using the degree lower bound from Proposition 5.2, we will obtain the desired lower bound on s .

Let H be the set of all pairs of terms in Π of Quadratic degree more than d . We know that $|H| \leq s^2$. In each iteration, we will find a restriction that removes a fraction α of the pairs from H , for $\alpha = \frac{d}{4p^2(M+mn)}$, and removes no more than κ holes from the formula. Thus for a t satisfying $(1-\alpha)^t |H| \leq (1-\alpha)^t s^2 < 1$, after t iterations, no high-degree quadratic terms survive, and the number of remaining holes is $n' \geq n - t\kappa = n(1 - t\kappa/n)$. Since $1 - \alpha \leq e^{-\alpha}$, note that t is roughly $2 \log s/\alpha$. At this point, the choice for d is clearer; we choose d so that $t\kappa/n$ is a small enough constant; say $t\kappa/n \leq 1/2$. Choosing d so that $\alpha = \frac{4\kappa \log s}{n}$ does the trick; in particular, $d = \frac{16p^2(M+mn)\kappa \log s}{n}$. With this choice of d , continuing with the outline above, we obtain a PC proof of degree $3pkd$ without extension variables for $\text{PHP}_{n'}^{m,r}$ with $n' \geq n/2$. From Proposition 5.2, we conclude that $3pkd \geq n'/2 \geq n/4$, and plugging in the chosen value of d , we see that

$$\log s \geq \Omega\left(\frac{n^2}{kp^3(M+mn)\kappa}\right).$$

Now we come to the main part of the proof, namely showing how to obtain the desired restriction in each iteration.

In each iteration, we first perform the following preprocessing steps. For each extension variable z with extension axiom $z - Q$, we compute its support and check whether zero is in it. If not, we compute $\ell(z)$ (Definition 4.2) and reduce the proof by $z^{\ell(z)} = c$. By Lemma 4.3 the latter is derivable from the extension axiom, and by Lemma 4.8, it does not raise the size of H . Moreover, our measure of Quadratic degree can only decrease when variables switch from Singular to Nonsingular; see comment after Lemma 4.7.

We then pick a variable y that by an averaging argument contributes to the quadratic degree of at least a $d/(M+mn)$ fraction of pairs in H . There are three cases to consider.

Case 1: y is an original Boolean variable, say x_{uv} for some $u \in [m]$, $v \in [n]$. We choose the restriction that sets all variables $x_{w,v}$ to 0, thus removing the hole v from the formula. Since x_{uv} is also set to 0 this way, $d/(M+mn)$ fraction of pairs in H are killed.

Case 2: y is an extension variable, say z , with the extension axiom $z - Q$, with $0 \in \text{supp}(Q)$, with $\deg(Q) \leq k$ and $|\text{vars}(Q)| \leq \kappa$.

By Corollary 5.4, we can find an assignment σ to $\text{vars}(Q)$ that has Hamming weight at most pk and sets Q to 0. We apply this assignment to the proof, additionally setting z to zero in the proof as well. We then look at how this assignment affects the tautology, and apply an additional assignment to restore to $\text{PHP}_{n'}^{m,r}$ where $n' \geq n - \kappa$.

We say that a hole v is affected if for some u , the variable x_{uv} is set by σ . Note that at most κ holes are affected since σ only sets variables in $\text{vars}(Q)$. We say that an assignment commits a pigeon $I \in \binom{[m]}{r}$ to a hole v if it sets the term t_{Iv} to 1. Now note that σ does not commit any pigeon I to any affected hole v , because each term t_{Iv} is the product of r variables, and the Hamming weight of σ is at most pk which is less than $r/2$. Thus we are free to remove an affected hole from the formula. We do so by setting to zero all unset variables x_{uv} for each affected hole v ; this makes $t_{Iv} = 0$ for all affected holes v and all pigeons I . The resulting formula is $\text{PHP}_{n'}^{m,r}$ where $n' \geq n - \kappa$, and applying the restriction to the current refutation gives a new refutation of this reduced formula with at most $(1 - d/(M + mn))|H|$ pairs of Quadratic degree d or more.

Case 3: y is an extension variable, say z , with the extension axiom $z - Q$, and $0 \notin \text{supp}(Q)$. This is the trickiest case.

We need to find a suitable assignment to the variables in $\text{vars}(Q)$ to kill many high-degree quadratic terms involving z . Recall from Definition 4.5 that, since z is non-singular, z contributes to the degree of a quadratic term pair through terms t_1, t_2 where the degrees of z in t_1 and t_2 are different. By averaging, we can pick indices $0 \leq i < j \leq p - 1$ such that pairs of terms of the form $(t_1 z^i, t_2 z^j)$ contribute at least a $d/p^2(M + mn)$ fraction of the contribution of z . Since we had preprocessed using Lemmas 4.3 and 4.8, we know that $i, j < \ell(z)$, so z^{j-i} is not a constant. Hence, using Lemma 5.3 with $l = j - i$, once again we obtain an assignment σ of Hamming weight $\leq lk$ such that $Q|_\sigma = (b - a)x_{uv} + a$, where $a^{j-i} \neq b^{j-i}$ and x_{uv} is some variable of $\text{PHP}_n^{m,r}$.

We would like to apply $\text{Split}_{z,i,j}$ to remove the contribution of these term pairs with z^i, z^j (see Section 4.2 and Remark 4.11) and reduce high Quadratic degree terms. But first we need to meet the preconditions for applying $\text{Split}_{z,i,j}$. In particular, we need to get rid of all axioms containing z , except for $(z - a)(z - b) = 0$; even the extension axiom $z = (b - a)x_{uv} + a$ must be eliminated. We also need to restore to a version $\text{PHP}_{n'}^{m,r}$. To this end, we apply σ and perform cleanup in a way similar to Case 2, before applying $\text{Split}_{z,i,j}$. The only difference is that here we need to get rid of all axioms containing x_{uv} , without actually setting the latter.

We handle holes v' other than v affected by σ exactly as in Case 2; all variables touching this hole ($x_{wv'}$) but unset by σ are now set to 0, eliminating hole v' .

For the hole v , all variables touching this hole (x_{wv}) but unset by σ are now set to 0, but x_{uv} is left unset. Nonetheless, we claim that all occurrences of x_{uv} in the axioms are now eliminated. This is because every such occurrence is in a term t_{Iv} , for some subset I of $[m]$ of size at least r . Since σ sets at most $lk < pk$ variables to one and $r > 2pk$, each such occurrence contains at least two variables unset by σ , and in particular contains an unset variable other than x_{uv} . Therefore, setting all variables of hole v (other than x_{uv}) which are unset by σ to zero gets rid of all such occurrences and thereby eliminates hole v .

Thus, we end up with a refutation of $\text{PHP}_{n'}^{m,r}$ for $n' \geq n - \kappa$ such that all affected holes are eliminated, and x_{uv} is still unset but does not appear in the axioms.

We now intend to substitute $x_{uv} = (b - a)^{-1}(z - a)$. Note that under this substitution, the extension axiom gets eliminated (becomes $0 = 0$) and the Boolean axiom $x_{uv}^2 - x_{uv} = 0$

reduces to $(z - a)(z - b) = 0$. This is possible by Lemma 4.3(2), and will enable us to satisfy all the preconditions to apply Split on z . However, there is a still a catch. The substitution might actually blow up the number of pairs in H , because it creates three additional pair of terms for every pair of terms (t_1, t_2) containing x . To handle this, we note that if the substitution blows up the number of high Quadratic degree pairs by more than a factor of $3d/4p^2(M + mn)$, then this implies that at least a $d/4p^2(M + mn)$ fraction of pairs of terms must have contained x_{uv} before this substitution. (The same argument is also used in [9].) In this case, we can just set $x_{uv} = 0$ instead of the above substitution; this will remove a $d/4p^2(M + mn)$ fraction of pairs of terms (and prevent the need to use Split). Otherwise, we apply the substitution, introducing at most $3d|H|/4p^2(M + mn)$ new pairs of terms, and then use $\text{Split}_{z,i,j}$ to obtain a valid refutation of the reduced formula (Lemma 4.13), removing at least the $d|H|/p^2(M + mn)$ pairs of terms which had quadratic degree with a contribution from z, i, j . Either way, the number of high-degree quadratic terms reduces by a fraction at least $d/4p^2(M + mn)$. Thus, we obtain a refutation of $\text{PHP}_n^{m,r}$ with at most $(1 - d/4p^2(M + mn))|H|$ pairs of terms of Quadratic degree at least d , with $n' \geq n - \kappa$.

This completes the description of how to extract a good restriction in each iteration. The fraction of high-degree Quadratic pairs eliminated is at least $d/(M + mn)$ in the first two cases and at least $d/4p^2(M + mn)$ in case 3. So in every case, at least $\alpha = d/4p^2(M + mn)$ fraction of the pairs is removed. With the analysis given in the beginning of this proof, the proof of Theorem 5.5 is now complete. \blacktriangleleft

Finally, applying a lift with the XOR_2 gadget and by choosing the parameters carefully, we obtain our claimed lower bound of Theorem 1.2. For convenience, we restate the theorem here.

▶ Theorem 1.2. *For every $N > 0$ large enough, any $1 > \epsilon, \delta > 0$, constant $c > 0$, and prime p , there exists a tautology F over N variables such that any PC refutation of F over \mathbb{F}_p with $N^{1+\epsilon(1-\delta)}$ extension variables, each depending on $N^{1-\epsilon}$ variables of F and of size at most N^c , requires size $\exp(\Omega(N^{\epsilon\delta}/\text{poly log } N))$.*

Proof. Pick an arbitrary n , and set $r = 100p(c + 2) \log n$ and $m = 2r$, so that $\binom{m}{r} > n$. Let G be the formula $\text{PHP}_n^{m,r}$. Let F be the formula obtained by composing G with the parity gadget on two variables; $F = G \circ \text{XOR}_2$. That is, replace each variable x in G by the XOR of two new variables x^1 and x^2 . The number of variables in the formula F is $N = 2mn = \Theta(n \log n)$.

Suppose we are given a PC refutation Π of F of size s , that uses no more than $N^{1+\epsilon(1-\delta)}$ extension variables, each of arity bounded by $N^{1-\epsilon}$ and size bounded by N^c . We will recover from Π a refutation of G , and then use Theorem 5.5 to obtain the stated lower bound on s .

Set $k = 10(c + 2) \log N$; then $r > 2pk$. We will find a restriction that reduces F to G , and reduces the degree of all extension axioms in Π to at most k . Note that the total size of all the extension axioms put together is at most $N^c \times N^{1+\epsilon(1-\delta)} < N^{c+2}$. Let ρ be a restriction that independently, for each variable x of G , picks one of x^1, x^2 uniformly at random, and sets it to 0 or 1 with equal probability. For any term t of degree at least k , the probability that t survives after applying ρ is at most $(3/4)^k$. By the union bound, the probability that some term in an extension axiom survives ρ is at most $N^{c+2}(3/4)^k$, which is strictly less than 1 for our choice of k . Hence there exists a restriction ρ that sets exactly one variable in each XOR gadget, and which reduces all extension axioms to degree at most k . A suitable renaming of the surviving variables (and interchanging with the negated literal if necessary) recovers G .

We thus have a PC refutation $\Pi' = \Pi|_{\rho}$ of G of size at most s . The number of extension variables in Π' is $M = N^{1+\epsilon(1-\delta)} = \tilde{O}(n^{1+\epsilon(1-\delta)})$, and each has arity at most $\kappa = \tilde{O}(n^{1-\epsilon})$ and degree at most $k = \Theta(\log n)$. Further, $M + mn = \theta(M)$. Also as already noted, $r > 2pk$. Hence by Theorem 5.5 we conclude that s is at least $\exp(\tilde{\Omega}(n^2/(M + mn)\kappa k)) = \exp(\tilde{\Omega}(n^{1+\epsilon}/n^{1+\epsilon(1-\delta)})) = \exp(\tilde{\Omega}(N^{\epsilon\delta}))$ (where the $\tilde{O}, \tilde{\Omega}$ notation hides poly log factors). This is the claimed bound. \blacktriangleleft

References

- 1 Michael Alekhnovich and Alexander A. Razborov. Lower bounds for Polynomial Calculus: Non-binomial case. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14-17 October 2001, Las Vegas, Nevada, USA*, pages 190–199. IEEE Computer Society, 2001.
- 2 Christoph Berkholz. The relation between Polynomial Calculus, Sherali-Adams, and Sum-of-Squares proofs. In *35th Symposium on Theoretical Aspects of Computer Science (STACS 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- 3 Sam Buss, Dima Grigoriev, Russell Impagliazzo, and Toniann Pitassi. Linear gaps between degrees for the Polynomial Calculus modulo distinct primes. *Journal of Computer and System Sciences*, 62(2):267–289, 2001.
- 4 Matthew Clegg, Jeffery Edmonds, and Russell Impagliazzo. Using the Groebner basis algorithm to find proofs of unsatisfiability. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 174–183, 1996.
- 5 Stephen A Cook and Robert A Reckhow. The relative efficiency of propositional proof systems. *The Journal of Symbolic Logic*, 44(1):36–50, 1979.
- 6 Nicola Galesi and Massimo Lauria. Optimality of size-degree tradeoffs for Polynomial Calculus. *ACM Transactions on Computational Logic (TOCL)*, 12(1):1–22, 2010.
- 7 Dima Grigoriev and Edward A Hirsch. Algebraic proof systems over formulas. *Theoretical Computer Science*, 303(1):83–102, 2003.
- 8 Russell Impagliazzo, Sasank Mouli, and Toniann Pitassi. The surprising power of constant depth algebraic proofs. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 591–603, 2020.
- 9 Russell Impagliazzo, Sasank Mouli, and Toniann Pitassi. Lower bounds for Polynomial Calculus with extension variables over finite fields. In *38th Computational Complexity Conference (CCC 2023)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2023.
- 10 Russell Impagliazzo, Pavel Pudlák, and Jiri Sgall. Lower bounds for the Polynomial Calculus and the Gröbner basis algorithm. *Computational Complexity*, 8:127–144, 1999.
- 11 Matthias Krause and Pavel Pudlák. On the computational power of depth 2 circuits with threshold and modulo gates. In *Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, pages 48–57, 1994.
- 12 Mladen Miksa and Jakob Nordström. A generalized method for proving Polynomial Calculus degree lower bounds. In *30th Conference on Computational Complexity, CCC 2015*, volume 33 of *LIPICs*, pages 467–487, 2015.
- 13 Ran Raz and Iddo Zameret. Resolution over linear equations and multilinear proofs. *Annals of Pure and Applied Logic*, 155(3):194–224, 2008.
- 14 Alexander A Razborov. Lower bounds for the Polynomial Calculus. *Computational Complexity*, 7:291–324, 1998.
- 15 Dmitry Sokolov. (Semi)Algebraic proofs over $\{\pm 1\}$ variables. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 78–90, 2020.

Appendix

We include here a self-contained proof of Lemma 3.7.

► **Lemma 3.7.** *[[15], Lemma 3.6] Let Π be a refutation of a set of axioms F of degree d_0 with Quadratic degree at most d . Then there exists a refutation Π' of F with (usual) degree at most $2\max(d, d_0)$.*

Proof. Let $\Pi = \{P_j\}_j$. Now, consider $\Pi' = \{P'_j\}_j$ where $P'_j = t_j P_j$, with each $t_j \in P_j$ carefully selected. Since the degree of $t_j P_j$ is bounded by the Quadratic degree of P_j , every line in Π' is of degree at most d . However, Π' is not an immediate valid refutation of F , but it can be transformed into one. We will show that each line of Π' can be derived from previous lines and axioms of F in degree at most $2\max(d, d_0)$, completing the proof. We proceed by induction on line number j .

If P_j is an axiom, then we set t_j to be an arbitrary term in P_j and derive $P'_j = t_j P_j$ in degree $2d_0$ starting from P_j .

(Note that in [15], it is claimed that this step can be derived in degree d_0 . But this is not always so. For instance, if $p = x_1 x_2 x_3 + x_2 x_3 x_4 + x_3 x_4 x_1 + x_4 x_1 x_2$ has degree $d_0 = 3$, and $d = 2$, then for any term $t \in p$, tp has degree 2 but needs degree $4 > \max\{d, d_0\}$ for the derivation.)

If $P_i = xP_j$ for some $j < i$, then we select $t_i = x t_j$, and consequently, $P'_i = t_i P_i = t_j P_j = P'_j$ is derived without raising the degree.

Finally, if $P_i = P_{j_1} + P_{j_2}$, we choose t_i to be an arbitrary term in P_i and derive $P'_i = t_i P_i = t_i t_{j_1} P'_{j_1} + t_i t_{j_2} P'_{j_2}$. We argue that the degree of both $t_i t_{j_1}$ and $t_i t_{j_2}$ is at most d , and as a result, P'_i can be derived from P'_{j_1} and P'_{j_2} in degree at most $2d$, which completes the proof. To justify this assertion, let $t_i \in P_{j_1}$ without loss of generality (every term in P_i appears in either P_{j_1} or P_{j_2}). Then degree of $t_i t_{j_1}$ is bounded by the Quadratic degree of P_{j_1} and hence by d . Additionally, if $t_{j_2} \in P_i$, then the degree of $t_i t_{j_2}$ is bounded by the quadratic degree of P_i and is also bounded by d . In the case where $t_{j_2} \notin P_i$, it means that it was cancelled in the sum and therefore $t_{j_2} \in P_{j_1}$ and so degree of $t_i t_{j_2}$ is bounded by the Quadratic degree of P_{j_1} and is again bounded by d .

Thus all lines in Π' can be derived from previous lines and axioms of F in degree at most $2\max(d, d_0)$. Since the last line of Π' is 1, we get that Π' can be successfully transformed into a valid proof of F of degree $2\max(d, d_0)$. ◀