Lower bounds for Polynomial Calculus with extension variables over finite fields

Russell Impagliazzo
University of California San Diego, California, United States

Sasank Mouli
Hyderabad, India

Toniann Pitassi
Columbia University, New York, United States

Abstract

For every prime $p > 0$, every $n > 0$ and $\kappa = O(\log n)$, we show the existence of an unsatisfiable system of polynomial equations over $O(n \log n)$ variables of degree $O(\log n)$ such that any Polynomial Calculus refutation over $\mathbb{F}_p$ with M extension variables, each depending on at most κ original variables requires size $\exp\left(\Omega(n^2)/10^\kappa(M + n\log n)\right)$

2012 ACM Subject Classification Theory of computation → Proof complexity

Keywords and phrases Proof complexity, Algebraic proof systems, Polynomial Calculus, Extension variables, $\mathsf{AC}^0[p]$-Frege

Digital Object Identifier 10.4230/LIPIcs.CCC.2023.7

Funding *Russell Impagliazzo*: Supported by the Simons Foundation and NSF grant CCF-1909634.
Sasank Mouli: Supported by the Simons Foundation, NSF grant CCF-1909634 and the Swiss National Science Foundation project n. 200021_207429 / 1 "Ideal Membership Problems and the Bit Complexity of Sum of Squares Proofs".
Toniann Pitassi: Supported by NSF grant CCF-1900460, and by the IAS School of Mathematics.

Acknowledgements The authors would like to thank Paul Beame and Dmitry Sokolov for helpful discussions.

© Russell Impagliazzo, Sasank Mouli and Toniann Pitassi;
licensed under Creative Commons License CC-BY 4.0

1 Introduction

A major goal of proof complexity is to show limits on the types of reasoning formalizable with concepts of small computational complexity, usually formalized as circuits from small circuit classes. This makes results in proof complexity analogous to (and often building on) results in circuit complexity. However, despite having strong lower bounds for the class $\mathsf{AC}^0[p]$ since the 1980's, ([18, 19]) it is still an open problem in proof complexity to establish superpolynomial (or even quadratic) lower bounds for the corresponding proof system $\mathsf{AC}^0[p]$-Frege.

Motivated by the lack of progress towards proving $AC^0[p]$-Frege lower bounds, [4] defined the Nullstellensatz (Nullsatz) proof system for refuting systems of unsolvable polynomial equations. Given a system of polynomial equations $\mathcal{P} = \{P_1 = 0, \ldots, P_m = 0\}$ in Boolean variables $x_1, \ldots, x_n$ (where we enforce the Boolean condition by adding the equations $x_i^2 - x_i = 0$ to $\mathcal{P}$), a Nullsatz refutation of $\mathcal{P}$ over a field $\mathbb{F}$ is a set of polynomials $\mathcal{Q} = \{Q_1, \ldots, Q_m\}$ such that $\sum_i P_i Q_i = 1$. The degree of the refutation is the maximum degree of the $P_i Q_i$'s, and the size is the sum of the sizes of the polynomials in $\mathcal{P}, \mathcal{Q}$. A dynamic version of Nullsatz, called the Polynomial Calculus (PC) was later defined in [10].

While these and later papers showed strong lower bounds for these proof systems, often these lower bounds were brittle in that the tautologies where lower bounds were proved also had small upper bounds under changes of variables. Our work is intended to address the issue of proving algebraic proof lower bounds that are more robust under changes of variables. This can be viewed as a small but significant step towards proving lower bounds for $\mathsf{AC}^0[p]$-Frege, since the latter can simulate such changes of variables.

One reason for the brittleness of many of the earlier lower bounds is that these lower bounds were highly sensitive to the initial encoding. The known PC lower bounds hold for unsatisfiable CNF formulas which are converted to a corresponding system of unsolvable polynomial equations. Previous works established exponential PC lower bounds assuming a Boolean encoding, where the variables are Boolean, enforced by the initial equations $x_i^2 - x_i = 0$. Another natural encoding is the "Fourier" encoding which represents the constraints by polynomials over $\{-1, 1\}$-valued variables (by applying the linear transformation $x_i = 1 - 2x_i$ to the Boolean encoding). However under this second encoding, the size lower bounds all break down. This is due to the proof method, where size lower bounds were obtained from *degree* lower bounds. Over $\{0, 1\}$-valued variables, this can be accomplished by applying known size-degree tradeoffs for PC or by a random restriction argument to kill off all large monomials. But over $\{-1, 1\}$-valued variables, these methods no longer work: a generic size-degree tradeoff no longer holds (there are polynomial sized proofs of the Tseitin tautologies, although they require linear degree [8]), and since the monomials now correspond to parity equations, they are resilient to random restrictions.

However, recently, Sokolov [20] broke this barrier, and managed to prove exponential size lower bounds for PC refutations over the $\{-1, 1\}$ encoding. We note that while this may seem like a minor improvement over the known lower bounds which held for the $\{0, 1\}$-encoding, Sokolov had to invent a new and ingenious technique for proving size lower bounds. In this work, we generalize the methods of Sokolov to prove exponential PC lower bounds with up to $M = N^{2-\epsilon}$ extension variables which can depend on up to $\kappa = O(\log N)$ *original* variables (where N is the number of variables in the tautology). This shows that the Sokolov method can be used to prove highly robust lower bounds, that are not sensitive to local changes of variables. We state our result more precisely for two different choices of parameters, one that maximizes the size lower bound, and the other that maximizes the number of allowable extension variables.

▶ **Theorem 1** (high-end). *For n sufficiently large, there is a family of CNF tautologies F^{SEL} on $O(n \log n)$ variables with $poly(n)$ clauses of width $O(\log n)$ such that for any $M = npolylog(n)$ and $\kappa = O(\log \log n)$, any* PC *refutation over $\mathbb{F}_p$ of F^{SEL}, together with M κ-local extension axioms, requires size $2^{\Omega(n/polylog(n))}$.*

▶ **Theorem 2** (low-end). *For the same family of tautologies as above, there are $0 < \alpha, \beta, \gamma < 1$ so that, for $M = n^{1+\alpha}, \kappa = \beta logn$, any* PC *refutation of F^{SEL} together with any M κ-local extensions over $\mathbb{F}_p$ requires size $2^{\Omega(n^{\gamma})}$.*

We remark that our extension variables are only allowed to depend on the original variables, and not on previously defined extension variables. (In the more general case where extension variables are defined recursively, the proof system corresponds to $AC^0[p]$-Frege, where the level of recursion corresponds to the $AC^0[p]$ circuit depth.) Thus our lower bound can be (roughly) seen as proving exponential lower bounds for the following restricted class of depth-2.5 PC refutations. First, the refutation is given a *new* set of M variables, $z_1, \ldots, z_M$, and is allowed to define a corresponding set of M κ-local polynomials $Q_1, \ldots, Q_M$ (where each Q_i can only depend on κ original variables). Lines in the refutation are polynomials over the original variables, plus the new *extension* variables (which are placeholders for the Q_i's). Substituting the Q_i's for the new variables gives a set of depth 2.5 algebraic circuits using a pre-specified set of κ-local functions at the bottom layer of the circuit.

1.1 Related Work

The work that inspired us and that is most related to our result is the recent paper by Sokolov [20], proving exponential lower bounds on the size of PC refutations of CNF formulas, where the variables take on values in $\{1, -1\}$. We generalize Sokolov's result to hold over any finite field, even with the addition of superlinear many extension variables, each depending arbitrarily on a small number of original variables. Thus our result can be alternatively viewed as making progress towards proving exponential lower bounds for depth-3 $AC^0[p]$-Frege, for a family of CNF formulas.

We note that for systems of polynomial equations over the rationals, a body of recent work establishes much stronger lower bounds. First, [13] proved lower bounds for subsystems of IPS over the rationals by restricted classes of circuits, including low-depth formulas, multilinear formulas and read-once oblivious branching programs. Secondly, Alekseev [2] proved exponential lower bounds on the bit complexity of PC proofs with an arbitrary number of extension variables of unbounded depth over the rationals. Andrews and Forbes [3] prove quasipolynomial lower bounds on the circuit size of constant-depth IPS proofs for a different family of polynomials over the rationals; however, their hard instances do not have small-size constant-depth circuits. Finally, [14] establish a similar lower bound as [3], but for hard instances that have small constant-depth circuits.

We remark that these lower bounds are incomparable to ours for several reasons. First, they do not hold for finite fields, and secondly, the choice of hard polynomials are inherently nonboolean: [13, 2, 14] use the subset sum principle which when translated to a propositional statement is no longer hard, and the hard polynomials in [3] have logarithmic depth. Thus on the one hand they establish superpolynomial lower bounds for *much* stronger subsystems of IPS, but on the other hand, they do not translate to lower bounds for propositional proofs in the sense of Cook-Reckhow [11]. In particular, they don't imply lower bounds for proof systems dealing with Boolean formulae.

1.2 Our Result: Proof Overview

The standard way of proving size lower bounds for PC for an unsatisfiable formula F for Boolean-valued variables dates back to the celebrated superpolynomial lower bounds for Resolution [15, 7], where the basic tool is to reduce size lower bounds to degree lower bounds (or in the case of Resolution, size to clause-width) by way of either a general size-depth tradeoff, or by a more general random restriction argument. At a high level, both methods iteratively select a variable that occurs in a lot of high-degree terms, set this variable to zero (to kill off all high-degree terms containing it), while also ensuring (possibly by setting additional variables) that F remains hard to refute after applying the partial restriction. After applying this size-to-degree reduction, the main technical part is to prove degree lower bounds for the restricted version of F.

As mentioned in the Introduction, over the $\{-1, 1\}$ basis, the size to degree reduction breaks down. In fact, no generic reduction to degree can exist since random XOR instances over this basis require linear degree but have polynomial size PC refutations. Moreover, we lacked *any* method for proving PC lower bounds for unsatisfiable CNFs over the basis $\{-1, 1\}$, and more generally over an arbitrary linear transformation of the variables. In [16], we highlighted this as an open problem, noting that it is a necessary step toward proving superpolynomial $AC^0[2]$-Frege lower bounds, a major open problem in proof complexity.

Recently, Sokolov [20] made significant progress by proving exponential lower bounds for PC (as well as for SOS) for random CNF formulas over the domain $\{-1, 1\}$, by developing new formula-specific techniques to reduce size to degree over this domain. As this is the starting point for our work, we begin by describing the main method in [20] for reducing size to degree for certain families of formulas over $\{-1, 1\}$.

Let Π be an alleged PC refutation of F of small size which includes the axioms $w^2 = 1$ for all variables w. The first step in Sokolov's argument is to show how to remove all high degree terms containing a particular variable w, provided that w is *irrelevant* – meaning that it does not occur in any of the initial polynomials other than the equation $w^2 = 1$. Intuitively, we want to show that if our unsatisfiable system of polynomial equations doesn't contain w, then we should be able to eliminate high degree terms containing w altogether from the refutation. To show this, Sokolov introduced a new operation termed *Split* where he writes each line q in the refutation as $q_0 + q_1 w$, and proves by induction that if we replace each line q by the pair of lines q_0, q_1, then it is still a valid refutation of F (and no longer contains w). While the Split operation removes w from the proof, it doesn't kill off high degree terms. The crucial insight is that although this doesn't directly kill off high degree terms, a slightly different measure of degree (called Quadratic degree) can be used instead, since removing w via the Split operation removes all high Quadratic degree terms that w contributed to, and secondly low Quadratic degree implies low ordinary degree. The second and easier step in Sokolov's argument uses specific expansion properties of F to show that for any variable w, there exists a small restriction ρ (to some of the other variables) such that w becomes irrelevant under ρ.

Our main theorem significantly generalizes Sokolov's lower bound by proving exponential lower bounds for an unsatisfiable CNF formulas F, even when we allow the axioms $\mathcal{P}$ to contain superlinear many extension axioms, provided that each extension axiom depends on a small number of original variables. Note that the variables of F are Boolean, but the extension variables are not restricted to being Boolean. In particular, it may be the case that zero is not in the support of an extension variable (i.e. the set of all possible values that can be assigned to it without violating any Boolean axioms), for example if extension variable z is defined by the equation $z = x - 2$, then z cannot be set to zero without falsifying the

Boolean axiom $x^2 - x = 0$ for x. Intuitively we will handle extension variables z that cannot be set to zero in a similar manner to Sokolov, by first isolating z, and then generalizing the Split operation in order to kill off all large Quadratic degree terms that contain z. However, dealing with a general set of extension axioms presents new technical challenges that we address next.

Our first idea is to design the unsatisfiable formula F carefully so that we can force variables to be irrelevant in a more modular way. Specifically, let $F(x_1, \ldots, x_n)$ be an expanding unsatisfiable k-CSP formula with $m = O(n)$ constraints, such that any subset of $m' = \epsilon m$ constraints is unsatisfiable and requires proofs of large PC degree. We define an unsatisfiable formula F^{SEL} (based on F) that intuitively states that there is a subset S of $m' = \epsilon m$ constraints of F (as chosen by new selector variables $\mathbf{y}$) that is satisfiable. We will prove lower bounds on the set of constraints F^{SEL} even with the addition of an arbitrary set of extension axioms satisfying the conditions mentioned earlier. In order to make a variable of F^{SEL} irrelevant, we will simply make sure that our eventual assignment to the selector variables ($\mathbf{y}$) avoids constraints of F that contain this variable (we can also make a selector variable irrelevant in a slightly more complicated way, details are left to the relevant section).

A second challenge that we face (that doesn't come up in Sokolov's proof) is that extension variables may be defined so that originally they can be consistently set to zero, but can change status after applying a restriction. For example, suppose the proof uses the extension axiom $z = x_1 x_2 + x_1$. Then zero is in the support of z (since we can set $x_1 = x_2 = 0$), but if we set $x_1 = 1$, then zero is no longer in the support of z. In order to deal with this dynamically changing status of variables, our notion of Quadratic degree must pay attention to which category each of the extension variables is in at any particular time, and make sure that we do not lose progress that was made earlier due to variables changing from initially containing zero to disallowing zero in their support. Fortunately we observe that variables can only change unidirectionally, (since the support of a variable cannot increase under a restriction) and this is crucial for arguing that our measure of Quadratic degree always decreases so that we continually make progress.

Finally, we also have to generalize Sokolov's Split operation, which was previously defined only for $\{-1, 1\}$ variables. We give a generalization of how to do the Split for arbitrary valued variables.

2 Preliminaries

▶ **Definition 1** (Polynomial Calculus/Polynomial Calculus Resolution). *Let $\Gamma = \{P_1 \ldots P_m\}$ be an unsolvable system of polynomials in variables $\{x_1 \ldots x_n\}$ over $\mathbb{F}$. A* PC *(Polynomial Calculus) refutation of Γ is a sequence of polynomials $\{R_1 \ldots R_s\}$ such that $R_s = 1$ and for every $\ell \in [s]$, $R_\ell \in \Gamma$, 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:*

$R_\ell = \alpha R_j + \beta R_k$ for α, $\beta \in \mathbb{F}$

$R_\ell = x_i R_k$ for some $i \in [n]\}$

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 $\{x_1 \ldots x_n, \bar{x}_1 \ldots \bar{x}_n\}$ where $\{\bar{x}_1 \ldots \bar{x}_n\}$ are twin variables of $\{x_1 \ldots x_n\}$ i.e. the equations $x_i^2 - x_i = 0$, $\bar{x}_i{}^2 - \bar{x}_i = 0$ and $x_i \bar{x}_i = 0$ are treated as axioms.

▶ **Definition 2** (PC plus Extension Axioms). *Let $\Gamma = \{P_1 \ldots P_m\}$ be a set of polynomials in variables $\{x_1 \ldots x_n\}$ over a field $\mathbb{F}$. We will refer to the polynomials in Γ as (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* $\mathsf{PC} + \mathsf{Ext}$ *(*PC *plus extension) refutation of* Γ *with* M *extension axioms* $Ext = \{z_j - Q_j(x_1 \dots x_n)\}$ *is a* PC *refutation of the set of polynomials* $\Gamma' = \{P_1 \dots P_m, z_1 - Q_1 \dots z_M - Q_M\}$*. An extension axiom* $z_j = Q_j(x_1 \dots x_n)$ *is* κ*-local if* Q_j *is a* κ*-junta; that is, if the polynomial* Q_j *defining* z_j *involves at most* κ *of the* $\mathbf{x}$*-variables. We say that* Π *is a* $(M, \kappa) - \mathsf{PC} + \mathsf{Ext}$ *refutation of* Γ *if it is a* $\mathsf{PC} + \mathsf{Ext}$ *refutation of* Γ *with* M *extension axioms, each of which are* κ*-local. The size of the refutation is total size of all lines in the refutation, including the polynomials in* Γ *plus the extension axioms (where the size of a line* $P \in \Pi$ *is the number of monomials in* P*).*

We note that our definition of extension axioms is more limited than the general notion of extension axioms. Here we only allow the extension variables to depend on the *original* variables from Γ; the more general definition allows the extension variables to depend on the original $\mathbf{x}$-variables, and also on other extension variables.

▶ **Definition 3** (k-local CSPs)**.** *A constraint* C_i *over Boolean variables* $\{x_1, \dots, x_n\}$ *is simply a Boolean formula over these variables.* C_i *is a* k*-local constraint if* C_i *depends on at most* k *variables. A* k*-CSP* $\mathcal{C} = C_1 \wedge \dots \wedge C_m$ *over* $\{x_1, \dots, x_N\}$ *is the conjunction of a set of* k*-local constraints.*

We translate a k-CSP formula into a system of polynomial equations using the standard PCR translation which we define next.

▶ **Definition 4** (Converting k-CSPs into Polynomial Equations)**.** *Let* C *be a* k*-local constraint over variables* $x_{i_1}, \dots, x_{i_k}$*. We convert* C *to a polynomial equation,* $p(C)$*, using the translations* $p(x_{i_j}) = 1 - x_{i_j}$*,* $p(\neg A) = 1 - p(A)$*,* $p(A \vee B) = p(A) \cdot p(B)$*. It is easy to check that for any Boolean assignment* α *to the underlying variables,* $C(\alpha) = 1 \leftrightarrow p(\alpha) = 0$*, and* $C(\alpha) = 0 \leftrightarrow p(\alpha) = 1$*.*

A k*-CSP* $\mathcal{C} = C_1 \wedge \dots \wedge C_m$ *over* $\{x_1, \dots, x_n\}$ *converts to a set of polynomial equations* $\{E_j \mid j \in [m]\} \cup \{B_i \mid i \in [n]\}$ *over* $\{x_1, \dots, x_n\} \cup \{\bar{x}_1, \dots, \bar{x}_n\}$ *where* E_j *is the polynomial equation* $p(C_j)$ *In addition, we add the Boolean axioms* $\{B_i \mid i \in [n]\}$*, where* $B_i = \{x_i^2 - x_i = 0,\ \bar{x_i}^2 - \bar{x_i} = 0,\ x_i\bar{x_i} = 0\}$ *which force* $x_i, \bar{x_i}$ *to be zero-one valued, and force exactly one of* $x_i, \bar{x_i}$ *to be one.*

3 The Hard Formulas

We distinguish between the case $p = 2$ and the case $p > 2$, and concentrate on the latter. This is because the case $p = 2$ does not require any new technical ideas, and we can pick from a large number of known hard tautologies for this case, such as random CNF's. Over $\mathbb{F}_2$, every extension variable is zero-one valued, and so standard size-degree tradeoffs pertain even with respect to extension variables. Also, κ-local extension variables can change the degree by at most a factor of κ, therefore a degree lower bound of $\Omega(n)$ for the original tautology over n variables implies a degree lower bound of $\Omega(n/\kappa)$ after adding κ-local extension variables. Known size-degree tradeoffs imply that the degree must be at least square root of the number of variables in order to obtain exponential size lower bounds, this immediately gives a lower bound tolerating close to n^2/κ^2 many κ-local extension variables [10, 6, 17].

Over any field, there are unsatisfiable families of k-CNF formulas (e.g. the Tseitin tautologies as well as random parity equations) that require linear degree but have polynomial sized proofs with a linear number of extension variables [8, 6]. Therefore formulas that require high PC degree are not sufficient. Instead we will create our hard examples by taking a hard instance and then using selector variables to pick out a subset of the constraints. Similar ideas were used earlier (e.g., [12]). In more detail, our underlying hard unsatisfiable formulas, $\{F^{SEL}_{n,k}\}$, will be constructed from a family of k-CSP formulas, $F_{n,k}$, that have the property that any sufficiently large subset of the constraints of $F_{n,k}$ is unsatisfiable and still requires large PC degree.

▶ **Definition 5.** *Let $F_{n,k} = \{E_j \mid j \in [m]\} \cup \{B_i \mid i \in [n]\}$ be the system of degree-k polynomial equations over $\mathbf{x} = \{x_i, \bar{x_i} \mid i \in [n]\}$, obtained by converting a size-m k-CSP as given by Definition 4. For convenience, we will index the polynomial equations E_j in binary notation, so for example if $b_1 \dots b_{\log m} \in \{0,1\}^{\log m}$ is the binary notation for $j \in [m]$, we will write E_j as $E_{b_1 \dots b_{\log m}}$. We define a new set of polynomial equations $F^{SEL}_{n,k}$ with parameters m, m' as follows. The variables are $\mathbf{x} \cup \mathbf{y}$, where $\mathbf{x}$ are the original variables of $F_{n,k}$ and $\mathbf{y} = \{y_{i,j},\ \overline{y_{i,j}} \mid i \in [m'],\ j \in [\log m]\}$ are new "pigeon" variables. Let E^{SEL} be the following set of equations, where $y_i \neq b_1 \dots b_{\log m}$ abbreviates the monomial $\prod_{b_j=1} y_{i,j} \prod_{b_j=0} \overline{y_{i,j}}$:*

(i) $\forall i \in [m'], \forall b_1 \dots b_{\log m} \in \{0,1\}^{\log m}$, $(y_i \neq b_1 \dots b_{\log m}) \cdot E_{b_1 \dots b_{\log m}} = 0$;

(ii) $\forall i, i' \in [m']$, $i \neq i'$, $\forall b_1 \dots b_{\log m} \in \{0,1\}^{\log m}$, $(y_i \neq b_1 \dots b_{\log m}) \cdot (y_{i'} \neq b_1 \dots b_{\log m}) = 0$.

$F^{SEL}_{n,k}$ consists of the polynomial equations E^{SEL} together with the Boolean axioms $B_{i,j} = \{y^2_{i,j} - y_{i,j} = 0,\ \overline{y_{i,j}}^2 - \overline{y_{i,j}} = 0,\ y_{i,j}\overline{y_{i,j}} = 0\}$ for all $i \in [m'], j \in [m]$.

Intuitively we think of the $\mathbf{y}$ variables as a mapping from m' pigeons to m holes, where the holes correspond to the m axioms/constraints from E. For $i \in [m']$, the i^{th} "pigeon" y_i selects a hole (an equation from E).

The first set of polynomial equations in E^{SEL} states that if pigeon y_i selects the equation $E_{b_1 \dots b_{\log m}}$, then this equation must be satisfied; the second set of equations in E^{SEL} states that the mapping is one-to-one and thus altogether the $\mathbf{y}$ selector variables choose a subset E' of exactly m' equations from E. Thus $F^{SEL}_{n,k}$ asserts that there exists a subset of m' constraints of $F_{n,k}$ (chosen by the $\mathbf{y}$-variables) that are satisfiable.

Throughout this paper, the $\mathbf{x}$-variables are the variables that underly $F_{n,k}$; the $\mathbf{y}$-variables are the selector/pigeon variables described above that choose a subset of m' constraints from $F_{n,k}$, and the extension variables used in the PC + Ext refutation will be the $\mathbf{z}$-variables.

Our hard instances will be $F^{SEL}_{n,k}$, with $m = 10n$, $m' = (1-\epsilon)m$, where $F_{n,k}$ is (the polynomial translation of) an unsatisfiable k-CSP formula with $m = 10n$ k-local constraints over variables $\mathbf{x} = x_1 \dots x_n$, satisfying the follow property:

Property 1: Every subset of $(1-\epsilon)m'$ constraints is unsatisfiable and requires linear PC degree

The following Theorem shows that for sufficiently large n, such formulas exist. Similar proofs have appeared in several papers (e.g., [5]) but we give a proof in the Appendix for completeness.

▶ **Theorem 3.** *Let $m = 10n$. Then there exists constants $k > 0$, $0 < \epsilon < 1$ such that for sufficiently large n, there exists k-CSP formulas $\{F_{n,k}\}$ with m constraints such that Property 1 holds with $m' = (1-\epsilon)m$.*

4 The Lower Bound

4.1 Technical Proof Overview.

Conventionally, proof size lower bounds are reduced to degree lower bounds, a single step of which involves finding a *heavy* variable that occurs in a large fraction of high degree terms of the proof and setting it to zero. In our setting, if the heavy variable turns out to be an extension variable, z with extension axiom $z = Q(\mathbf{x}, \mathbf{y})$, it may be *Nonsingular* meaning that we cannot set $z = 0$ (without falsifying the extension axiom or a Boolean axiom), as opposed to *Singular* variables which can be set to zero in a consistent way[1]. In this case, we cannot simply eliminate the high degree terms containing z by setting $z = 0$. Sokolov [20] focused on the case where variables are over the ± 1 basis instead of the usual Boolean one, which is the simplest case where all variables are Nonsingular. Sokolov introduced *Quadratic degree* as a measure to be used instead of degree. Quadratic degree essentially measures the maximal degree of the *square* of each polynomial P occurring in the proof. For a ± 1 variable z, $z^2 = 1$, so squaring a polynomial P on ± 1 variables removes the contribution of a term $t \in P$ as it gets squared out, and what remain are the terms $t_1 t_2$ for t_1, $t_2 \in P$. Since any variable that appears in both terms gets squared out, the degree of these terms measures the symmetric difference between such terms, and this turns out to be a key complexity measure while dealing with Nonsingular variables. Sokolov showed that a refutation of low Quadratic degree can be turned into one of low degree. Thus the presence of Nonsingular variables is not necessarily a problem as long as the Quadratic degree of each line is low. Sokolov also introduced an operation *Split* that acts on a proof line by line in order to remove the contribution to Quadratic degree of any particularly *heavy* Nonsingular variable z, in the special case where they always take on values in ± 1, by replacing a line $P = P_1 z + P_0$ in the refutation with the lines P_1 and P_0. Sokolov managed to show that for some well chosen tautologies, the new Split lines still form a valid refutation of a hard subset of axioms. The crucial observation here is that this splitting of lines has eliminated from the square of the proof all pairs of terms whose product contained z. Thus, repeated application of Split would lead to contradiction of known degree lower bounds.

The first step for us was to generalize the notions of Quadratic degree and *Split* to any finite field. Motivated by the above definition of Quadratic degree, we generalize it as follows. Given two terms t_1 and t_2, a Nonsingular variable z contributes to the Quadratic degree between t_1 and t_2 if and only if it appears with different exponents in them, i.e. $z^i \in t_1$ and $z^j \in t_2$, for $i \neq j$. A Singular variable z contributes if and only if it appears in one of the terms with a nonzero exponent. The Quadratic degree of t_1 and t_2 is the total number of such variables z that contribute. Generalizing the *Split* operation proved a bit more difficult. We first focus on the case over $\mathbb{F}_p$ analogous to Sokolov's, where we have a variable z such that the identity $(z - a)(z - b) = 0$ holds for some constants $a, b \neq 0$ in the field. Note that a line $P(z)$ of the proof is of the form $P_{p-2} z^{p-2} + \cdots + P_1 z + P_0$. In the case of ± 1 variables, $p = 3$ and thus the contribution by z to Quadratic degree comes just from the interaction between two polynomials P_1 and P_0. Therefore separating P_1 and P_0 into different lines removes this contribution entirely. In the general case, however, the contribution by z to Quadratic degree is the sum total of interactions between polynomials P_i and P_j for every pair $i, j < p - 1$ such that i and j are distinct. We show how to separate P into two lines R_1, R_0 such that the interaction between P_i and P_j is completely removed, for any i, j satisfying $a^{i-j} \neq b^{i-j}$,

[1] This terminology is taken from singular and nonsingular matrices, since the key property we use is that a variable z is Nonsingular if and only if z^{p-2} is a "multiplicative inverse" of z, i.e. $z^{p-1} = 1$

or in other words, z^i and z^j are linearly independent over the two values that z takes. Let $R(z) = R_1 z^i + R_0 z^j$ be a polynomial such that R agrees with P for each possible value of z, i.e. $R(a) = P(a)$ and $R(b) = P(b)$. Since z^i and z^j are linearly independent over values $\{a, b\}$, these two equations can be solved for their coefficients R_1, R_0, expressed in terms of $P_{p-2} \dots P_0$. On closer observation, we find that P_i does not occur in the expression for R_1 and similarly P_j does not occur in R_0, and therefore we have successfully broken P into lines R_1 and R_0 while separating P_i and P_j. It is straightforward to show that this new set of lines forms a valid refutation, but an essential assumption we make here is that the initial axioms are free of z, except for $(z - a)(z - b) = 0$.

We now move to dealing with the case of a more general extension variable z with the extension axiom $z - Q$, where $Q(\mathbf{x}, \mathbf{y})$ is a polynomial that can depend on at most κ variables. Let H be the set of all pairs of terms (t_1, t_2) in a line of a given refutation that have high Quadratic degree between them. We would like to emulate Sokolov's strategy of eliminating this set of pairs from the refutation to drop its Quadratic degree. If an extension variable z which is Singular appears heavily in H, we apply the restriction that sets it to zero (which exists by the definition of Singular). In the case that z is Nonsingular, our goal is to reduce it to the above case in order to apply Split. But first, we will have to choose a "good" pair of indices ℓ_1, ℓ_0 such that Splitting them is effective in reducing H. We observe that for any pair of indices i, j, the set of pairs (t_1, t_2) in H such $z^i \in t_1$ and $z^j \in t_2$ is disjoint from the similar set defined for a distinct pair i', j'. Therefore by averaging we can pick a good pair ℓ_1, ℓ_0 that covers at least a $1/p^2$ fraction of z's appearances in H. We now have to reduce z to take on two distinct values a, b in order to apply Split, but these values need to be such that $a^{\ell_1 - \ell_0} \neq b^{\ell_1 - \ell_0}$. We show that there is a decision tree process (Lemma 10) that queries the variables underlying Q such that it is always possible to reduce z to the form $(b - a)w^* + a$, where a, b are useful to separate the indices ℓ_1, ℓ_0. It is fairly easy to see as a result of the discussion so far that if we are able to apply Split on z with indices ℓ_1, ℓ_0 at this stage, it causes a sizable reduction in H.

We are now almost ready to apply Split, but we still have to meet the requirement that the axioms are free of z. Since z is an extension variable it appears only in the extension axiom which has now been reduced to the form $(b - a)w^* + a$, and so the only way to remove this axiom is to make a substitution for $w^* = (z - a)/(b - a)$ in terms of z. This would get rid of this extension axiom and take the Boolean axiom for w^* to $(z - a)(z - b) = 0$ just like we need, but if w^* appears in any of the other axioms this substitution just creates new copies of z. Therefore we need to remove w^* from all the other axioms before we try to make this substitution. Here is where we make use of the structure of our tautology $F^{SEL}_{n,k}$ by defining an operation **Cleanup** which can remove any Boolean variable w^* from the axioms without actually setting it to a constant value. **Cleanup** also restores the structure of our tautology so that we are always working with a subset of equations and pigeons from $F^{SEL}_{n,k}$ that are untouched by previous restrictions. We describe this operation in detail in Section 4.5.1.

Once we perform the above cleanup operations we are ready to make the substitution for $w^* = (z - a)/(b - a)$ in terms of z to satisfy the requirements for Split. We are met with a final hurdle here: this substitution can potentially increase the number of pairs of terms in H. Fortunately it can be resolved by a simple case analysis: if the blowup is too large it must have been the case that w^* appeared frequently in H, and so setting it to zero will reduce H without the need for Split. Otherwise, Split is able to offset this blowup.

Therefore we have demonstrated above how to reduce the size of the high Quadratic degree set H by a constant fraction. Performing this for sufficiently many iterations would remove

H entirely and lower the Quadratic degree of any refutation. We then use a generalized version of Sokolov's argument that low Quadratic degree implies low degree in order to switch to a low degree refutation. For a small sized refutation, the number of iterations needed is bounded and thus we are able to keep most of the pigeons and equations alive at the end. We then select a hard subset of equations by assigning all remaining pigeons, and expand any remaining extension variables in order to obtain a low degree refutation of these equations, towards a contradiction.

4.2 Singular and Nonsingular variables

Let us fix the finite field $\mathbb{F}_p$, $p > 2$ for the rest of the article. We also fix a set of unsatisfiable polynomials F over Boolean variables $\mathbf{x} \cup \mathbf{y}$, and a set of extension axioms Ext of the form $z - Q$ over variables $\mathbf{z}$. Whenever we refer to a refutation Π, we assume that it is a PC + Ext refutation of $F \cup Ext$.

▶ **Definition 6** (Support of a variable). *Let $z - Q(w_{i_1}, \ldots, w_{i_\kappa}) = 0$ be a κ-local extension axiom associated with z. We define the set $vars(Q) = \{w_{i_1}, \ldots, w_{i_\kappa}\}$ and sometimes write $vars(z)$ to denote $vars(Q)$, the set of variables that z depends on. The support of z, $supp(z) \subseteq [0, p-1]$, is equal to the set of all values $a \in [0, p-1]$ such that there exists a Boolean assignment α to the variables of Q such that $Q(\alpha) = a$. Sometimes we also indicate this by $supp(Q)$.*

We extend the definition of support also to Boolean variables. For a Boolean variable w, $supp(w) = \{0, 1\}$ as enforced by the Boolean axiom $w^2 = w$.

▶ **Definition 7** (Singular and Nonsingular variables w.r.t. Ext). *Let Ext be a set of extension axioms and let z be an extension variable with an axiom in Ext. We say that z is Singular w.r.t. Ext iff $0 \in supp(z)$; otherwise we say that z is Nonsingular w.r.t. Ext. Any Boolean variable is considered Singular by default, independent of the set Ext, since zero always belongs to its support. For a term t, let $sing(t)$ be the subterm of t containing the Singular variables in t, and let $nsing(t)$ be the subterm of t containing the Nonsingular variables.*

Note that for a Singular extension variable z, it is possible to set z to zero, However, we note that this may falsify other polynomial equations in F. For example, if $xy = 0$ is a polynomial equation in F, and the extension axiom for z is $z - 1 + xy = 0$, then setting $x = y = 1$ forces $z = 0$, but this falsifies $xy = 0$.

▶ **Definition 8.** *Let $A \subseteq [1, \ldots, p-1]$, $A \neq \emptyset$. 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 $\ell(z) = \ell(supp(z))$.*

▶ **Lemma 4.** *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 $vars(Q)$, in degree at most $|vars(Q)|$.*

1. *$z - Q' = 0$, where Q' is the multilinear version of Q;*
2. *For any $A' \subseteq [0, p-1]$ such that $supp(z) \subseteq A'$, $\Pi_{a \in A'}(z - a) = 0$;*
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 $vars(Q)$.

Proof. Let $z - Q(w_{i_1}, \ldots, w_{i_\kappa}) = 0$ be the extension axiom for z, and let $supp(z) = A \subseteq A' \subseteq [1, p-1]$. First, we can derive the multilinear version of Q, Q', from Q together with the Boolean axoms $w^2 - w = 0$ for all $w \in vars(Q)$. Secondly, by definition, $supp(z) = A$ means that the allowable values for z over Boolean assignments to $vars(Q)$ are the values in A. Therefore, $z-Q=0$ together with the Boolean axioms $w^2-w=0$ for all $w \in vars(Q)$ implies $\Pi_{a\in A}(z-a)=0$. Furthermore, this polynomial has a PC derivation, by the derivational completeness of PC. Since $A \subseteq A'$, $\Pi_{a\in A'}(z-a)=0$ is a weakening of $\Pi_{a\in A}(z-a)=0$ and is therefore derivable from $\Pi_{a\in A}(z=a)=0$. Lastly, we will argue that there exists some constant $c \in \mathbb{F}_p^*$ such that $z^{\ell(A)}-c=0$ is semantically implied by $z-Q=0$ plus the Boolean axioms for $vars(z)$ and therefore is derivable from these axioms. Since the only allowable values for z under the Boolean axioms are the values in A, and since by definition of $\ell(A)$, for every $a \in A$, $a^{\ell(A)}=c$ for some $c \in \mathbb{F}_p^*$, it follows that $z^{\ell(A)}-c=0$. ◀

▶ **Definition 9.** *For a term t and a variable w, $\deg(t,w)$ is equal to the degree of w in t. If w is $Nonsingular$, then $w^{p-1}=1 \mod p$, so $\deg(t,w)<p-1$. On the other hand if w is Singular then we have $w^p=w \mod p$ and therefore $\deg(t,w)<p$. For a term t the degree of t, $\deg(t)$, equals $\sum_{w\in vars(t)}\deg(t,w)$.*

4.3 Quadratic degree

The next definition is a generalization/modification of Sokolov's definition of Quadratic degree for the more general scenario where the proof contains extension variables that are $Singular$ as well as ones that are $Nonsingular$.

▶ **Definition 10** (Quadratic degree)**.** *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 $P\in\Pi$.*

We usually instantiate the above definition with $V=\mathbf{x}\cup\mathbf{y}\cup\mathbf{z}$ and with S being the set of Singular variables as defined by the extension axioms corresponding to $\mathbf{z}$. However since $Qdeg^S$ is a different measure for every S, and our set of Singular variables can change under the application of a restriction ρ to the variables in V, we must make sure that our measure of Quadratic degree does not change significantly under a restriction[2]. Fortunately, we can show that for any two sets S and T such that $T\subseteq S$, $Qdeg^T\leq Qdeg^S$. Along with the simple observation that the set of Singular extension variables can only shrink under a restriction, this implies that our measure of Quadratic degree can only decrease under a restriction. We make this formal below.

▶ **Lemma 5.** *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, $Qdeg^T(t_1,t_2)\leq Qdeg^S(t_1,t_2)$.*

Proof. Note that for a variable $w\in S-T$, $Qdeg^S(t_1,t_2,w)=1$ when w has a nonzero exponent in one of t_1 or t_2, otherwise zero. However, $Qdeg^T(t_1,t_2,w)=1$ if and only if the

[2] If the set S does not change under a restriction, $Qdeg^S$ can still change under the restriction as terms can shrink or disappear when variables are set by the restriction. However, this is no different from how the usual notion of degree changes under a restriction, and it is trivial to show that $Qdeg^S$ always decreases. Therefore we ignore this for the sake of simplicity.

previous condition is satisfied and the exponents of w in t_1 and t_2 are not equal. Thus the claim follows. ◀

Henceforth, when we refer to Quadratic degree, we always fix the set S to be the set of Singular variables w.r.t. the underlying extension axioms. We have the following important corollary that this measure always decreases under a restriction to the underlying variables.

▶ **Corollary 6.** *Let F be a set of unsatisfiable polynomials over variables $\mathbf{x} \cup \mathbf{y}$ and let Ext be a set of extension axioms of the form $z - Q(w_{i_1}, \ldots, w_{i_\kappa})$ for variables $z \in \mathbf{z}$ and $w_{i_1}, \ldots, w_{i_\kappa} \in \mathbf{x} \cup \mathbf{y}$. Let ρ be a restriction to $\mathbf{x} \cup \mathbf{y}$ and let $Ext|_\rho$ be the axioms given by $z - Q|_\rho$ for each axiom $z - Q \in Ext$. The Quadratic degree w.r.t. $Ext|_\rho$ is at most the Quadratic degree w.r.t. Ext.*

Proof. Since $supp(Q|_\rho) \subseteq supp(Q)$ for any polynomial Q, we have that the set of Singular variables under $Ext|_\rho$ is a subset of those under Ext. Therefore our claim follows from the previous lemma. ◀

▶ **Lemma 7** (Quadratic degree upper bounds degree of Singular variables)**.** *For any term t,* $\deg(sing(t)) \leq pQdeg(t, t)$

Proof. For any Singular variable w, $Qdeg(t, t, w) = 1$ if and only if w occurs in t. Since w can occur in t with degree at most $p - 1$, the claim follows. ◀

▶ **Definition 11** (High quadratic degree terms)**.** *For a proof Π and $d \geq 0$, let $\mathcal{H}_d(\Pi)$ denote the set of unordered pairs (t_1, t_2) of quadratic degree at least d. That is, $\mathcal{H}_d(\Pi)$ is the set of unordered pairs of terms (t_1, t_2) such that t_1, t_2 both occur in P for some polynomial $P \in \Pi$, and $Qdeg(t_1, t_2) \geq d$.*

▶ **Lemma 8.** *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 $|\mathcal{H}_d(\Pi')| \leq |\mathcal{H}_d(\Pi)|$ for any $d \geq 0$.*

Proof. Consider a polynomial $P \in \Pi$ and a pair of terms (t_1, t_2) that occur in P. For any variable w distinct from z, $Qdeg(t_1, t_2, w)$ is unaltered when P is reduced by $z^{\ell(z)} = c$. On the other hand, if z does not contribute to the Quadratic degree of (t_1, t_2) i.e. $Qdeg(t_1, t_2, z) = 0$, then it will still be 0 after reducing by $z^{\ell(z)} = c$. Therefore $Qdeg(t_1, t_2)$ never increases for any pair (t_1, t_2) and thus $|\mathcal{H}_d(\Pi')| \leq |\mathcal{H}_d(\Pi)|$. ◀

The following is a generalized version of the argument from [20] that shows how to convert a proof with low Quadratic degree to one with low degree.

▶ **Lemma 9.** *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 $3pd$.*

Proof. The proof of this lemma is largely based on (a slightly cleaner version of) Sokolov's argument ([20], Lemma 3.6) that low Quadratic degree over $\{\pm 1\}$ variables implies low degree. Our first observation is that Sokolov's argument can be applied to any refutation of low Quadratic degree over $\mathbb{F}_p$ such that every term contains only Nonsingular variables. In particular if $\{P_j\}$ is a refutation that only contains Nonsingular terms, then we can use his argument to show that $\{t_j^{p-2} P_j\}$ is also a valid refutation for some carefully chosen term $t_j \in P_j$. Moreover, the degree of the latter refutation is bounded by a constant times the Quadratic degree of the former one. To see this, first note that for two Nonsingular terms t_1 and t_2, we have that $deg(t_1 t_2^{p-2}) \leq (p - 1) \cdot Qdeg(t_1, t_2)$, because of the following. For a variable z that is Nonsingular such that z occurs in t_1 and t_2 with $deg(t_1, z) = deg(t_2, z)$, we have $deg(t_1 t_2^{p-2}, z) = Qdeg(t_1, t_2, z) = 0$ since it would appear in $t_1 t_2^{p-2}$ with an exponent that is a multiple of $p-1$, and $z^{p-1} = 1$ holds for Nonsingular variables. Any other Nonsingular z that occurs in at least one of t_1 and t_2 has $deg(t_1 t_2^{p-2}, z) < p-1$ and $Qdeg(t_1, t_2, z) = 1$. Therefore the degree of $t_1 t_2^{p-2}$ is at most $p \cdot Qdeg(t_1, t_2)$ when t_1 and t_2 contain only Nonsingular variables. This implies that the lines in the new refutation $\{t_j^{p-2} P_j\}$ have degree at most p times the Quadratic degree of the original refutation $\{P_j\}$. Sokolov additionally showed that each line in the new refutation can be derived from previous lines without exceeding degree equal to $2p$ times the Quadratic degree of the original refutation, completing the argument.

In our case we deal with terms containing both Singular and Nonsingular variables. The above argument cannot be applied directly to our case, since it crucially depends on the fact that Nonsingular variables can be raised to the power $p-1$ to make them vanish. Fortunately by Lemma 7, the degree of Singular variables in any term is at most p times the Quadratic degree with itself. Given this bound, we can ignore for each term the part that contains Singular variables, and apply the above argument only with respect to the Nonsingular part of each term, to reduce the degree of Nonsingular variables in each term of the refutation. Since we now have a bound on the degree of both Singular and Nonsingular variables in each term, we have bounded its degree. We describe this in full technical detail below.

Let $\{P_j\}$ be a refutation of F with Quadratic degree bounded by d. For any term t recall that $nsing(t)$ denotes the subterm of t containing only Nonsingular variables. Note that $nsing(t)^{p-1} = 1$ for any t. For every line P_j in the refutation, we pick a term $t_j \in P_j$ and define $P'_j = nsing(t_j)^{p-2} P_j$. Note that by the arguments outlined above, for any two terms t_1 and t_2 in P_j, we have $\deg(nsing(t_1)^{p-2} nsing(t_2)) \leq pd$ and thus the degree of Nonsingular variables in any term of P'_j is bounded by pd. Since the Singular variables in any term remain unchanged under multiplication by $nsing(t_j)^{p-2}$, the Singular degree of P'_j the same as that of P_j and is bounded by pd (Lemma 7) and therefore $\deg(P'_j) \leq pd + pd = 2pd$. We now show that the set $\{P'_j\}$ forms a valid refutation of F and each P'_j can be derived from previous lines in degree $3pd$. If P_j is one of the axioms, we multiply by $nsing(t_j)^{p-2}$ to get P'_j for an arbitrary $t_j \in P_j$, and this takes degree $pd_0 \leq pd$. If $P_j = wP_{j_1}$ for $j_1 < j$ and some variable w, we choose $t_j \in P_j$ such that $t_j = wt_{j_1}$ where $t_{j_1} \in P_{j_1}$ was chosen earlier. If w is Singular, we have $nsing(t_j) = nsing(t_{j_1})$ and therefore $P'_j = nsing(t_j)^{p-2} P_j = w \cdot nsing(t_{j_1})^{p-2} P_{j_1} = wP'_{j_1}$. On the other hand, if w is Nonsingular, we have $nsing(t_j) = w \cdot nsing(t_{j_1})$ and therefore $P'_j = nsing(t_j)^{p-2} P_j = w^{p-1} \cdot nsing(t_{j_1})^{p-2} P_{j_1} = P'_{j_1}$. Finally, let $P_j = P_{j_1} + P_{j_2}$ for $j_1, j_2 < j$. We pick an arbitrary term $t_j \in P_j$. Note that since $nsing(t)^{p-1} = 1$ for any term t, $P_{j_1} = nsing(t_{j_1}) P'_{j_1}$ and $P_{j_2} = nsing(t_{j_2}) P'_{j_2}$ and thus we have $P'_j = nsing(t_j)^{p-2} nsing(t_{j_1}) P'_{j_1} + nsing(t_j)^{p-2} nsing(t_{j_2}) P'_{j_2}$ for $t_{j_1} \in P_{j_1}$ and $t_{j_2} \in P_{j_2}$ chosen earlier. We now show that $\deg(nsing(t_j)^{p-2} nsing(t_{j_1})) \leq pd$ and $\deg(nsing(t_j)^{p-2} nsing(t_{j_2})) \leq pd$ to conclude the proof. Since every term in P_j appears in one of P_{j_1}, P_{j_2}, let $t_j \in P_{j_1}$ without loss of generality. Then we have that t_j, t_{j_1} both appear in P_{j_1} and thus $\deg(nsing(t_j)^{p-2} nsing(t_{j_1})) \leq pd$. If $t_{j_2} \in P_j$ i.e. it is not cancelled in the sum $P_{j_1} + P_{j_2}$, then we have t_j, t_{j_2} both appear in P_j and hence $\deg(nsing(t_j)^{p-2} nsing(t_{j_2})) \leq pd$. If $t_{j_2} \notin P_j$, this implies that it was cancelled in the sum $P_{j_1} + P_{j_2}$ and therefore $t_{j_2} \in P_{j_1}$ and $\deg(nsing(t_j)^{p-2} nsing(t_{j_2})) \leq pd$.

◀

4.4 The Split Operation

In this section we will show how to apply a restriction and then use an operation *Split* (motivated by [20]) in order to eliminate high quadratic degree terms. Our main focus will be to handle the case where the variable to be set is an extension variable with extension axiom $z - Q = 0$ where z is $Nonsingular$, since in the other case we can potentially just set $z = 0$ to eliminate terms. We start by showing how to apply a small Boolean restriction ρ such that $Q|_\rho$ is a simple linear function of just one variable.

▶ **Lemma 10.** *Let z be an extension variable with extension axiom $z - Q(w_1, \ldots, w_k)$, for $k \leq \kappa$. Assume that z is Nonsingular (i.e. $supp(Q) \subseteq [1, \ldots, p-1]$) and $|supp(Q)| \geq 2$. Then for every $l \in [0, \ldots, \ell(supp(Q)) - 1]$, there exists a variable w^* in $vars(Q)$, and a restriction δ to $vars(Q) - w^*$ such that:*

(1) $Q|_\delta = (b-a)w^* + a$, where $b, a \in supp(Q)$. Thus $Q|_\delta$ is a linear function of w^* and $supp(Q|_\delta) = \{a, b\}$;

(2) $a^l \neq b^l$ (mod p)

Proof. We will create a decision tree that will query $vars(Q)$ one-by-one. Associated with the root r is the set of values $S_r = \{a^l \mid a \in supp(Q)\}$. That is, we label the root with the set of all possible values that z^l can take on. Since $l < \ell(supp(z))$, it follows that $|S_r| \geq 2$ (since otherwise we would have $l = \ell(supp(z))$). At the root we query the first variable w_1, with left edge labelled by $w_1 = 0$ and right edge labelled by $w_1 = 1$. Now we label the left vertex with the set $\{a^l \mid a \in supp(Q|_{w_1=0})\}$, of all values that z^l can take on under the restriction $w_1 = 0$. Similarly we label the right vertex with the set $\{a^l \mid a \in supp(Q|_{w_1=1})\}$. We continue recursively, querying the next variable at each vertex v of the decision tree, as long as the set of allowable values for z^l under the partial restriction ρ_v associated with v is greater than one. Now consider the longest path, ξ in T. The partial restriction ρ associated with ξ sets the first k' variables, where $k' \geq 1$ since initially z^l takes on at least two values. Also since ξ is a complete path, the associated set $\{a^l \mid a \in supp(Q|_\rho)\}$ contains exactly one element, call it q.

Now consider the twin path ξ' with associated restriction ρ', where ρ' is obtained from ρ by toggling the value of the last variable, $w_{k'}$, queried. Again since ξ' is a complete path, the associated set $\{a^l \mid a \in supp(Q|_{\rho'})\}$ contains exactly one element, call it q'. Note that q, q' must be distinct.

Let δ be the following assignment to $vars(Q) - w_{k'}$: for $1 \leq j < k'$, we set $\delta(w_j) = \rho(w_j) = \rho'(w_j)$, and for $k' < j \leq k$, we set $\delta(w_j) = 0$. Setting $w^* = w_{k'}$, $Q|_\delta$ is a linear equation of the form $(b-a)w^* + a$, where $b, a \in supp(Q)$. Finally, by construction, $a^l \neq b^l$ (since otherwise the two paths corresponding to ρ, ρ' would be the same).

◀

In the remainder of this subsection, we will be interested in the case where we want to eliminate some Nonsingular extension variable z from the refutation, and we have already applied the above Lemma so that the extension axiom for z is of the form $z-((b-a)w+a) = 0$, where w is some variable in $\mathbf{x} \cup \mathbf{y}$. Thus, $supp(z) = \{a, b\}$. The next two Lemmas generalizes a similar argument due to Sokolov, and show how to remove Quadratic degree pairs of the form $(t_1 z^i, t_2 z^j)$ for a carefully chosen pair i, j from the refutation via the Split operation.

▶ **Lemma 11.** *Let z be an extension variable such that $supp(z) = \{a, b\}$, where $a \neq b$ and $a, b \in \mathbb{F}_p^*$ and let P be any polynomial. Then, for any two distinct numbers ℓ_0, ℓ_1 where $\ell_0 < \ell_1$ and $a^{\ell_1 - \ell_0} \neq b^{\ell_1 - \ell_0}$, there exists a unique polynomial $R = R_0 z^{\ell_0} + R_1 z^{\ell_1}$ such that*

$R = P \mod (z-a)(z-b)$. *That is,* $R(a) = P(a)$ *and* $R(b) = P(b)$*, where* $P(a)$ *denotes the polynomial* P *under the substitution* $z = a$*.*

Proof. Let $z - Q = 0$ be the extension axiom for z, where $supp(z) = \{a, b\}$. Then by Lemma 4 the polynomial $(z-a)(z-b) = 0$ is implied by (and derivable from) the extension axiom for z plus the Boolean axioms. We can assume without loss of generality that P has the form $P_0 + zP_1 + \ldots + z^{p-2}P_{p-2}$.

Now we want to argue that there exists a polynomial $R = z^{\ell_0}R_0 + z^{\ell_1}R_1$, where R_0, R_1 are polynomials over $vars(P) - z$, and such that $R(a) = P(a)$, and $R(b) = P(b)$. We can find R_0 and R_1 by solving the following system of equations, where we view R_0, R_1 as the underlying variables, and treating $P(a), P(b)$ as constants:

$$a^{\ell_0}R_0 + a^{\ell_1}R_1 = P(a)$$

$$b^{\ell_0}R_0 + b^{\ell_1}R_1 = P(b)$$

This has a (unique) solution since the determinant of the associated matrix is $\begin{vmatrix} a^{\ell_0} & a^{\ell_1} \\ b^{\ell_0} & b^{\ell_1} \end{vmatrix} = a^{\ell_0}b^{\ell_0}(b^{\ell_1-\ell_0} - a^{\ell_1-\ell_0})$. By our assumption, this matrix is non-singular over $\mathbb{F}_p$ and therefore the above system of equations has a unique solution over $\mathbb{F}_p$, given by:

$$\begin{pmatrix} R_0 \\ R_1 \end{pmatrix} = \begin{pmatrix} a^{\ell_0} & a^{\ell_1} \\ b^{\ell_0} & b^{\ell_1} \end{pmatrix}^{-1} \begin{pmatrix} P(a) \\ P(b) \end{pmatrix}$$

Abbreviating $a^{\ell_0}, a^{\ell_1}, b^{\ell_0}, b^{\ell_1}$ by a_0, a_1, b_0, b_1 respectively, we have by definition of the inverse:

$$\begin{aligned} \begin{pmatrix} R_0 \\ R_1 \end{pmatrix} &= \begin{pmatrix} a_0 & a_1 \\ b_0 & b_1 \end{pmatrix}^{-1} \begin{pmatrix} P(a) \\ P(b) \end{pmatrix} \\ &= \frac{1}{a_0b_1 - a_1b_0} \begin{pmatrix} b_1 & -a_1 \\ -b_0 & a_0 \end{pmatrix} \begin{pmatrix} P(a) \\ P(b) \end{pmatrix} \end{aligned}$$

Solving for R_0 we have:

$$\begin{aligned} R_0 &= \frac{b_1}{a_0b_1 - a_1b_0}P(a) - \frac{a_1}{a_0b_1 - a_1b_0}P(b) \\ &= \frac{b_1}{a_0b_1 - a_1b_0}(a_0P_{\ell_0} + a_1P_{\ell_1} + \sum_{i \neq \ell_0, \ell_1} a^iP_i) - \frac{a_1}{a_0b_1 - a_1b_0}(b_0P_{\ell_0} + b_1P_{\ell_1} + \sum_{i \neq \ell_0, \ell_1} b^iP_i) \\ &= \frac{a_0b_1}{a_0b_1 - a_1b_0}P_{\ell_0} + \frac{a_1b_1}{a_0b_1 - a_1b_0}P_{\ell_1} + \frac{b_1}{a_0b_1 - a_1b_0}(\sum_{i \neq \ell_0, \ell_1} a^iP_i) \\ &\qquad - \frac{a_1b_0}{a_0b_1 - a_1b_0}P_{\ell_0} - \frac{a_1b_1}{a_0b_1 - a_1b_0}P_{\ell_1} - \frac{b_1}{a_0b_1 - a_1b_0}(\sum_{i \neq \ell_0, \ell_1} b^iP_i) \\ &= P_{\ell_0} + \sum_{i \neq \ell_0, \ell_1} c_{0i}P_i \end{aligned}$$

for some constants $c_{0i} \in \mathcal{F}_p$. And similarly solving for R_1, it has the following form:

$$R_1 = P_{\ell_1} + \sum_{i \neq \ell_0, \ell_1} c_{1i}P_i$$

for some constants $c_{1i} \in \mathcal{F}_p$.

◀

▶ **Definition 12** (*Split*). *Let z be an extension variable with extension axiom $z - Q = 0$ such that $supp(z) = \{a, b\} \subseteq [1, \ldots, p-1]$. For any polynomial P and for every $\ell_0 < \ell_1$ such that $a^{\ell_1 - \ell_0} \neq b^{\ell_1 - \ell_0}$, let $R = R_0 z^{\ell_0} + R_1 z^{\ell_1}$ be the unique polynomial given by Lemma 11 such that $R = P \bmod (z-a)(z-b)$. Then $Split_{z,\ell_1,\ell_0}(P)$ is defined to be the pair of polynomials $\{R_0, R_1\}$. For a proof Π, and an extension variable z such that $supp(z) = \{a, b\}$, we define $Split_{z,\ell_0,\ell_1}(\Pi)$ to be the sequence of lines $Split_{z,\ell_0,\ell_1}(P)$, over all $P \in \Pi$.*

▶ **Lemma 12.** *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 ℓ_0, ℓ_1 such that $\ell_0 < \ell_1$ and $a^{\ell_1 - \ell_0} \neq b^{\ell_1 - \ell_0}$, $\Pi' = Split_{z,\ell_0,\ell_1}(\Pi)$ forms a valid refutation of F modulo $(z-a)(z-b)$*

Proof. Fix an extension variable z in Π such that it does not occur in any axioms except $(z-a)(z-b) = 0$, and let ℓ_0, ℓ_1 be such that $\ell_0 < \ell_1$ and $a^{\ell_1 - \ell_0} \neq b^{\ell_1 - \ell_0}$. We will show by induction on the number of lines in Π that $Split_{z,\ell_0,\ell_1}(\Pi)$ is a valid derivation that meets the conditions of the lemma. For the base case, note that all of the axioms are either free of z or eliminated as a result of reducing by $(z-a)(z-b)$, and hence their Split versions are derivable. Now suppose that the Lemma holds for the first $j-1$ lines of Π; that is, $Split_{z,\ell_0,\ell_1}(\Pi_{j-1})$ is a valid derivation, where Π_{j-1} denotes the first $j-1$ lines of Π.

The first case is where P_j is a linear combination of two previously derived lines, so $P_j = \alpha P_{j_1} + \beta P_{j_2}$ for some j_1 and j_2 less than j and $\alpha, \beta \in \mathbb{F}_p$. Using the inductive hypothesis, we have:

$$\begin{aligned} P_j &= \alpha(z^{\ell_0} R_{j_1 0} + z^{\ell_1} R_{j_1 1}) + \beta(z^{\ell_0} R_{j_2 0} + z^{\ell_1} R_{j_2 1}) \mod (z-a)(z-b) \\ &= z^{\ell_0}(\alpha R_{j_1 0} + \beta R_{j_2 0}) + z^{\ell_1}(\alpha R_{j_1 1} + \beta R_{j_2 1}) \mod (z-a)(z-b) \end{aligned}$$

By the uniqueness of the polynomial $R_j = z^{\ell_0} R_{j0} + z^{\ell_1} R_{j1}$ that is equivalent to P_j mod $(z-a)(z-b)$ (by Lemma 11), this implies that $R_{j0} = \alpha R_{j_1 0} + \beta R_{j_2 0}$ and similarly $R_{j1} = \alpha R_{j_1 1} + \beta R_{j_2 1}$, and thus R_{j0} can be derived from a linear combination of $R_{j_1 0}$ and $R_{j_2 0}$ and similarly for R_{j1}.

The second case is when P_j is derived from a previously derived line $P_{j'}$ by multiplying $P_{j'}$ by a variable w. That is, $P_j = w P_{j'}$ for some $j' < j$. If $w \neq z$, then we have that $R_{j1} = w R_{j'1}$ (similarly for R_{j0}). If $w = z$ then we have:

$$\begin{pmatrix} R_{j'1} \\ R_{j'0} \end{pmatrix} = \begin{pmatrix} a^{\ell_1} & a^{\ell_0} \\ b^{\ell_1} & b^{\ell_0} \end{pmatrix}^{-1} \begin{pmatrix} P_{j'}(a) \\ P_{j'}(b) \end{pmatrix}$$

from which we need to derive

$$\begin{aligned} \begin{pmatrix} R_{j1} \\ R_{j0} \end{pmatrix} &= \begin{pmatrix} a^{\ell_1} & a^{\ell_0} \\ b^{\ell_1} & b^{\ell_0} \end{pmatrix}^{-1} \begin{pmatrix} P_j(a) \\ P_j(b) \end{pmatrix} \\ &= \begin{pmatrix} a^{\ell_1} & a^{\ell_0} \\ b^{\ell_1} & b^{\ell_0} \end{pmatrix}^{-1} \begin{pmatrix} aP_{j'}(a) \\ bP_{j'}(b) \end{pmatrix} \\ &= \begin{pmatrix} a^{\ell_1} & a^{\ell_0} \\ b^{\ell_1} & b^{\ell_0} \end{pmatrix}^{-1} \begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix} \begin{pmatrix} a^{\ell_1} & a^{\ell_0} \\ b^{\ell_1} & b^{\ell_0} \end{pmatrix} \begin{pmatrix} R_{j'1} \\ R_{j'0} \end{pmatrix}. \end{aligned}$$

Thus, R_{j1} can be derived as a linear combination of $R_{j'1}$ and $R_{j'0}$, and similarly for R_{j0}.

◀

4.5 Proof of Main Theorem

The proof of our lower bound for the tautology $F^{SEL}_{n,k}$ with extension axioms Ext proceeds by choosing a variable in the given refutation Π that contributes to a lot of high quadratic degree pairs of terms in Π. If this variable is Singular, we apply the restriction that sets it to zero. On the other hand, if it is Nonsingular and therefore an extension variable z, we first reduce it to depend on a single variable w^* by applying a restriction chosen from Lemma 10, and then use a more complicated case analysis (see Lemma 15) in order to apply the Split operation from Lemmas 11 and 12 on z. In both of these cases we are able to remove a small fraction of high Quadratic degree terms, and thus after sufficiently many iterations we obtain a refutation of low Quadratic degree. We convert this to a refutation of low (usual) degree using Lemma 9, and then substitute for the pigeon variables $\mathbf{y}$ to select a subset of equations from $F_{n,k}$ that require high degree, obtaining a contradiction.

4.5.1 Cleanup operations

In order to get the contradiction at the end of the above argument, we need to ensure that our process above is always working with a subset of equations of $F_{n,k}$ that are untouched, i.e. unaffected by earlier restrictions to variables. We also need to eliminate any partially assigned pigeons so that we have full choice over the equations we are able to pick at the end. Additionally, a key requirement of the Split lemmas (Lemmas 11 and 12) is that the variable z we Split on must not appear in any axioms except for one of the form $(z-a)(z-b)=0$, which indicates that it takes two distinct values. In particular, we cannot set z or the underlying variable w^* described above in order to eliminate them from the refutation. This presents us with a unique requirement: for any choice of a variable $w^* \in \mathbf{x} \cup \mathbf{y}$, we need to be able to eliminate all axioms containing w^* without actually setting it. We show how to perform these operations by making use of the structure of our tautology $F^{SEL}_{n,k}$.

We first show how to "ban" an equation $E_{b_1 \ldots b_{\log m}}$ from $F_{n,k}$ by switching to a set of axioms that prevent any pigeon from being assigned to $b_1 \ldots b_{\log m}$.

▶ **Lemma 13.** *Let Π be a refutation of $F^{SEL}_{n,k}|_\rho$ for some restriction ρ and let $(y_i \neq b_1 \ldots b_{\log m}) \cdot E_{b_1 \ldots b_{\log m}} = 0$ be one of its axioms. Then there exists another valid refutation Π' with the latter axiom replaced by the axiom $(y_i \neq b_1 \ldots b_{\log m}) \equiv \prod_{b_j=1} y_{i,j} \prod_{b_j=0} \overline{y_{i,j}}$, such that the quadratic degree of Π' is at most that of Π.*

Proof. Note that the axiom $(y_i \neq b_1 \ldots b_{\log m}) \cdot E_{b_1 \ldots b_{\log m}} = 0$ can be derived from the axiom $(y_i \neq b_1 \ldots b_{\log m}) \equiv \prod_{b_j=1} y_{i,j} \prod_{b_j=0} \overline{y_{i,j}}$ by multiplying by the polynomial $E_{b_1 \ldots b_{\log m}}$. Since this derivation involves only singular variables, the degree can never drop and therefore the quadratic degree of this derivation is at most that of the final polynomial. We construct Π' as follows. We first derive the former axiom from the latter in Π'. Besides this derivation, Π' involves the same steps as Π. ◀

▶ **Definition 13.** *An equation $E_{b_1 \ldots b_{\log m}}$ is said to be banned when the previous lemma is applied repeatedly to eliminate all occurrences of it from the axioms.*

▶ **Definition 14.** *A clean version of $F^{SEL}_{n,k}$ is any subset of axioms of $F^{SEL}_{n,k}$ along with axioms that ban some subset of equations of the form $E_{b_1 \ldots b_{\log m}}$.*

4.5.1.1 Cleanup(ρ)

We now describe how to perform the cleanup operations, which we collectively call **Cleanup(ρ)**, that takes as input an "unclean" version of $F^{SEL}_{n,k}$ derived by applying a restriction ρ to a clean version, and outputs another clean version that is in some sense a subset of the input. Suppose that we are given a restriction ρ that has been applied to a clean version of $F_{n,k}^{SEL}$, with a variable $w^* \in \rho$ possibly set to $\star$, indicating that it must remain unset. To eliminate an axiom that has been affected by a $\mathbf{x}$ variable in ρ not set to $\star$, we simply obtain the refutation that bans the corresponding equation $E_{b_1 \dots b_{\log m}}$ as described in the above lemma. Note that since we are eliminating the axiom without setting any variables in it, we can also do this in case our variable $w^* \in \mathbf{x}$. Suppose that y_{ij} is a $\mathbf{y}$ variable in ρ not set to $\star$. We first note that any axiom that contained y_{ij} before the application of ρ contains all the variables $y_{i1} \dots y_{i \log m}$ corresponding to the i^{th} pigeon y_i. We first make sure that this i^{th} pigeon does not contain our variable w^* that must remain unset. If it doesn't, we proceed as follows. We set all the other variables in this pigeon to select some equation $E_{b_1 \dots b_{\log m}}$ that has not been banned. Such an equation exists provided that the number of banned equations so far is bounded, and the size of the restriction ρ is also bounded (we formalize this in the lemma below). We then apply an additional restriction to the $\mathbf{x}$ variables that satisfies this equation $E_{b_1 \dots b_{\log m}}$ picked above. We then ban all the equations affected by this additional restriction, like we did above for the part of ρ containing $\mathbf{x}$ variables. This eliminates the pigeon y_i. We are left with the case where our variable w^* belongs to some pigeon y_j. We set all the variables in the pigeon y_j except for w^*, such that neither of the two equations $E_{b_1 \dots b_{\log m}}$ and $E_{b'_1 \dots b'_{\log m}}$ that would be selected if w^* is set to zero or one are banned (again, these exist under the same conditions as above). We then proceed as before, i.e. apply an additional restriction to satisfy both these equations, and then ban any other equations that have been affected by this additional restriction. With this we have eliminated the axioms of pigeon y_j which select an equation, but we are still left with the axioms that prevent y_j from colliding with any other pigeon, which are now of the form $w^* \cdot (y_{j'} \neq b_1 \dots b_{\log m})$ and $\overline{w^*} \cdot (y_{j'} \neq b'_1 \dots b'_{\log m})$ indicating that any pigeon $y_{j'}$ distinct from y_j must not be mapped to the equations $E_{b_1 \dots b_{\log m}}$ and $E_{b'_1 \dots b'_{\log m}}$ if one of them is selected by setting w^* to zero or one. To remove the latter axioms we do something similar to the process of banning an equation, where we simply replace these axioms by the axioms $(y_{j'} \neq b_1 \dots b_{\log m})$ and $(y_{j'} \neq b'_1 \dots b'_{\log m})$, effectively banning the equations $E_{b_1 \dots b_{\log m}}$ and $E_{b'_1 \dots b'_{\log m}}$ for the remaining pigeons.

4.5.1.2 Correctness of Cleanup(ρ)

We note that the above cleanup operations over $\mathbf{y}$ variables terminate successfully only when there are enough equations that have not been banned by prior calls to cleanup, and also the size of the restriction ρ is bounded. We make this formal by the below lemma.

▶ **Lemma 14** (Correctness of **Cleanup(ρ)**). *Let ρ be a restriction of size κ. If the number of banned equations (from previous calls to **Cleanup**) is $\ll m/2^\kappa$, then **Cleanup(ρ)** terminates correctly. Moreover, it bans at most $O(\kappa)$ additional equations and removes at most $O(\kappa)$ pigeons in its run.*

Proof. In **Cleanup(ρ)**, note that we can remove the axioms that contain $\mathbf{x}$ variables unconditionally. When we remove a pigeon $y_i = y_{i1} \dots y_{i \log m}$, we rely on having an equation it can be set to that is not already banned. Since the size of ρ is bounded by κ, note that at most κ variables from $y_{i1} \dots y_{i \log m}$ can be set by ρ. Therefore there are at least $\log m - \kappa$ of them unset, corresponding to selecting $m/2^\kappa$ many equations. Since we assume that the number of banned equations is much less than this, we can always find one that is not banned to assign this pigeon to.

We now count the number of new equations banned and the number of pigeons removed by this call to **Cleanup(ρ)**. Since each $\mathbf{x}$ variable appears in a constant number of equations, the number of equations we ban while processing it is a constant. When we process a $\mathbf{y}$ variable, we pick and satisfy an equation, and ban all other equations affected in the process. Since every equation also contains a constant number of variables, satisfying it affects only a constant number of other equations. Therefore, for every variable we process we ban only a constant number of equations, and thus the total number of equations banned is $O(\kappa)$. We remove only those pigeons with a variable in ρ, so this is also bounded by $O(\kappa)$. ◀

4.5.2 The Main Theorem

We need first the following key lemma that shows how to apply the Split operation to reduce high quadratic degree terms.

Algorithm 1 Algorithm for Lemma 15

Input: A refutation Π, and a nonsingular variable z with extension axiom $z - Q = 0$ satisfying the pre-conditions of Lemma 15
Output: A refutation Π' satisfying post-conditions of Lemma 15

Let $\ell_0 < \ell_1$ be such that $|\mathcal{H}_d(\Pi, z, \ell_0, \ell_1)| \geq |\mathcal{H}_d(\Pi, z)|/p^2$.
Apply Lemma 10 with $l = \ell_1 - \ell_0$ to obtain δ, w^*, a, b satisfying post-conditions of Lemma 10.
$\Pi = \Pi|_\delta$ (and in particular $z - Q|_\delta = z - (b-a)w^* - a$)
Cleanup($\delta \cup \{w^* = \star\}$) (Cleanup axioms affected by δ and remove w^* from all axioms other than $z - (b-a)w^* - a$ while keeping it alive.)
if w^* *contributes to* $\geq \epsilon/4p^2$ *fraction of pairs in* $\mathcal{H}_d(\Pi)$ **then**
  $\Pi = \Pi|_{w^*=0}$
end
else
  Apply the substitution $(z-a)/(b-a)$ for w^* in Π
  Let $\Pi' = Split_{z,l_0,l_1}(\Pi)$
end

▶ **Lemma 15.** *Let F be a system of unsatisfiable polynomials and let z be a nonsingular extension variable with the extension axiom $z - Q$. Let $\ell = \ell(supp(Q))$ so that $z^\ell = c$ holds for some $c \in \mathbb{F}_p$. Let Π be a refutation of $F \cup \{z - Q\}$ modulo $z^\ell = c$ such that for at least an ϵ fraction of pairs (t_1, t_2) in $\mathcal{H}_d(\Pi)$, $Qdeg(t_1, t_2, z) = 1$, for some $d \geq 0$. Then there exists a refutation Π' of F such that $|\mathcal{H}_d(\Pi')| \leq (1 - \epsilon/4p^2)|\mathcal{H}_d(\Pi)|$*

Proof. We will apply a procedure as described by Algorithm 1 in order to modify the proof to satisfy the post-conditions of the Lemma. Here we give a detailed description of the algorithm, together with its correctness. Let $\mathcal{H}_d(\Pi, z)$ be the set of all unordered pairs $(t_1, t_2) \in \mathcal{H}_d(\Pi)$ that z contributes to. That is, $\mathcal{H}_d(\Pi, z)$ is the set of all unordered pairs $(t_1, t_2) \in \mathcal{H}_d(\Pi)$ such that $Qdeg(t_1, t_2, z) = 1$. There are many different ways that z can contribute to $\mathcal{H}_d(\Pi, z)$: namely, for all i, j such that $i < j < \ell$, let $\mathcal{H}_d(\Pi, z, i, j)$ be the set of all unordered pairs $(t_1, t_2) \in \mathcal{H}_d(\Pi, z)$, such that the degree of z in t_1 is i and the degree of z in t_2 is j. Note that for any two pairs (i, j) and (i', j') such that $i \neq i'$ or $j \neq j'$, $\mathcal{H}_d(\Pi, z, i, j)$ and $\mathcal{H}_d(\Pi, z, i', j')$ are disjoint. Therefore, there exists a "good" pair $\ell_0 < \ell_1 < \ell$ such that removing $\mathcal{H}_d(\Pi, z, \ell_1, \ell_0)$ from $\mathcal{H}_d(\Pi, z)$ will remove at least a $1/p^2$ fraction of $\mathcal{H}_d(\Pi, z)$ and therefore a ϵ/p^2 fraction of pairs in $\mathcal{H}_d(\Pi)$, since $|\mathcal{H}_d(\Pi, z)| \geq \epsilon|\mathcal{H}_d(\Pi)|$.

We want to apply the Split operation $Split_{z,\ell_0,\ell_1}$ to remove all such pairs. But in order to do this we have to satisfy the preconditions of Lemmas 11 and 12: we need two values a, b such that $a^{\ell_1-\ell_0} \neq b^{\ell_1-\ell_0}$ and all the axioms should be free of z except for $(z-a)(z-b) = 0$. The first step (Line 2 of 1) is to apply Lemma 10 with $l = \ell_1 - \ell_0$. This gives us $w^* \in vars(Q)$, $a, b \in supp(Q)$ and a partial restriction δ to $vars(Q) - w^*$ such that $(z-Q)|_\delta = z - (b-a)w^* - a$, where $a^{\ell_1-\ell_0} \neq b^{\ell_1-\ell_0} \mod p$. Next, we apply the restriction δ to Π (Line 3).

Now we have a simpler *linear* extension axiom for z of the form $z - (b-a)w^* - a = 0$. Next we would like to make the substitution $w^* = (z-a)/(b-a)$ in Π in order to satisfy this extension axiom, towards the goal of eliminating z from the axioms so that we have the preconditions of Lemma 12 and therefore are able to apply $Split_{z,\ell_1,\ell_0}$. However, if w^* appears in any of the axioms in F, this would create additional occurrences of z and we would not make any progress. Therefore, we have to make sure that none of the axioms of F contain w^*. But we also cannot set w^* to zero or one in an attempt to get rid of it, since this would set z to either a or b through the above extension axiom, and Split requires that z take on two distinct values. We thus have to get rid of all axioms mentioning w^* either by setting other variables or by replacing these axioms with stronger versions, such that the former can be derived from the latter. This is what the subroutine **Cleanup** does, in addition to removing the axioms in F that were affected by our earlier restriction δ, so that we have a clean version of $F^{SEL}_{n,k}$ as defined in the previous section.

We are now ready to make the substitution $w^* = (z-a)/(b-a)$. Under this substitution, the Boolean axiom ${w^*}^2 - w = 0$ reduces to $(z-a)(z-b) = 0$, and the original extension axiom for z disappears (since under this substitution it becomes $0 = 0$.) Thus this substitution would satisfy all of the preconditions of Lemmas 11, 12. However, this substitution can create a new problem: it can cause a blow up in the size of $\mathcal{H}_d(\Pi)$ since for every pair of terms (t_1, t_2) such that one of them contains w^*, we could have up to four new terms after the substitution. In order to deal with this potential blow up we do a simple case analysis: If w^* contributes to at least an $\epsilon/4p^2$ fraction of pairs (t_1, t_2) in $\mathcal{H}_d(\Pi)$, then we set $w^* = 0$ (Lines 4-5). This gives us the required reduction in the size of $\mathcal{H}_d(\Pi)$ (z is also set to a constant by setting $w^* = 0$, but we don't care about that since we have obtained a reduction in high Quadratic degree terms without needing to use Split). Otherwise, the blowup caused by the substitution $w^* = (z-a)/(b-a)$ adds at most $3\epsilon/4p^2$ fraction of pairs to $\mathcal{H}_d(\Pi)$, and thus if we remove all pairs in $\mathcal{H}_d(\Pi, z, \ell_0, \ell_1)$ (after this blowup) then overall we will will have reduced the size of $\mathcal{H}_d(\Pi)$ to $(1 - \epsilon/4p^2)|\mathcal{H}_d(\Pi)|$. So in this latter case, we apply the substitution mentioned above (Line 8) which simultaneously removes w^* from all axioms, and replaces the linear axiom for z by $(z-a)(z-b) = 0$. Now all preconditions for Lemma 8 hold so we can apply $Split_{z,\ell_0,\ell_1}$ (Line 9) to get a valid refutation. It is left to argue that this indeed removes the set $\mathcal{H}_d(\Pi, z, \ell_1, \ell_0)$. More precisely, we argue that high Quadratic degree pairs of terms in the refutation obtained after applying Split have a one to one mapping to the set $\mathcal{H}_d(\Pi) - \mathcal{H}_d(\Pi, z, \ell_1, \ell_0)$. Fix a line $P \in \Pi$. Since we are working modulo $z^\ell = c$, we can assume that $P = P_0 + zP_1 + \ldots z^{\ell-1}P_{\ell-1}$. Let $R = z^{\ell_0}R_0 + z^{\ell_1}R_1$ be the unique polynomial equivalent to P mod $(z-a)(z-b)$. $Split_{z,\ell_0,\ell_1}(\Pi)$ is the refutation with lines R_1, R_0 for all $P \in \Pi$. By the proof of Lemma 11 R_0, R_1 have the form:

$$R_1 = P_{\ell_1} + \sum_{i<\ell, i\neq\ell_0} c_{1i}P_i$$

$$R_0 = P_{\ell_0} + \sum_{i<\ell, i\neq\ell_1} c_{0i}P_i$$

for some constants $c_{1i}, c_{0i} \in \mathbb{F}_p$.

For a pair of terms (t_i, t_j) in R_1 such that $t_i \in P_i$ and $t_j \in P_j$ and $Qdeg(t_i, t_j) \geq d$, we map it to the pair $(t_i z^i, t_j z^j) \in P$, and similarly for R_0. Clearly this is a one-one mapping, and since P_{ℓ_0} does not occur in R_1 and P_{ℓ_1} does not occur in R_0, it is a mapping to $\mathcal{H}_d(\Pi) - \mathcal{H}_d(\Pi, z, \ell_1, \ell_0)$. Therefore we have that for the refutation $\Pi' = Split_{z,\ell_0,\ell_1}(\Pi)$ whose lines are $\{R_1, R_0\}$, $|\mathcal{H}_d(\Pi')| \leq |\mathcal{H}_d(\Pi) - \mathcal{H}_d(\Pi, z, \ell_1, \ell_0)| \leq (1 - \epsilon/4p^2)|\mathcal{H}_d(\Pi)|$.

◀

▶ **Theorem 16.** *For n sufficiently large, any (M, κ)-PC + Ext refutation of $F^{SEL}_{n,k}$ has size* $\exp\left(\frac{\Omega(n^2)}{10^\kappa (M + n \log n)}\right)$.

Algorithm 2 Eliminating high Quadratic degree terms from the proof

Input: A refutation Π of $F^{SEL}_{n,k}$ with extension axioms Ext
Output: A refutation Π' with Quadratic degree less than d

$d \leftarrow \nu n/\kappa$, where ν is a sufficiently small constant.
$M' \leftarrow M + n\log(n)$. ($M'$ upper bounds $|\mathbf{x} \cup \mathbf{y} \cup \mathbf{z}|$, the total number of variables)
$S \leftarrow \mathbf{x} \cup \mathbf{y} \cup \mathbf{z}$ (the current set of singular variables: all Boolean variables are singular by default and we initialize all extension variables to also be singular. This could possibly reduce in each iteration.)
$H \leftarrow \{(t_1, t_2) \mid t_1, t_2 \in \Pi \text{ and } Qdeg^S(t_1, t_2) \geq d\}$ (the set of all pairs of terms of large Quadratic degree according to S)
while *H is non empty* **do**
  for *every extension axiom $z - Q \in Ext$* **do**
    if $0 \notin supp(Q)$ **then**
      $S \leftarrow S - \{z\}$
      Compute c such that $z^{\ell(z)} = c$ and reduce Π by the latter identity
    end
  end
  $H \leftarrow \{(t_1, t_2) \mid t_1, t_2 \in \Pi \text{ and } Qdeg^S(t_1, t_2) \geq d\}$ (update H to reflect changes due to the above **for** loop)
  Pick a variable w that, by an averaging argument, occurs in at least an ϵ fraction of terms in H, where we choose $\epsilon = d/M'$.
  if $w \in S$ **then**
    Let σ be a restriction on $\mathbf{x} \cup \mathbf{y}$ such that $w|_\sigma = 0$
    $\Pi \leftarrow \Pi|_\sigma$
    Cleanup(σ)
  end
  else
    Apply Algorithm 1, which by Lemma 15 satisfies the post-conditions of Lemma 15
  end
end

Proof. Let Π be an alleged (M, κ)-PC + Ext refutation of $F^{SEL}_{n,k}$ with logarithm of its size less than $\gamma n^2/(10^\kappa(M + n\log n))$, for a small enough constant γ. Given Π, Algorithm 2 (defined below) will apply a sequence of restrictions and cleanup steps in order to produce a refutation Π' of a clean version of $F^{SEL}_{n,k}$ (see Definition 14) with the property that the

Quadratic degree of Π' is at most $d = \nu n/\kappa$ for a small enough constant $\nu > 0$. The algorithm contains a while loop which iteratively removes all pairs of terms of high Quadratic degree. From Π', we will apply a further restriction to all of the remaining unset $\mathbf{y}$-variables (i.e. pigeons that select equations from $F_{n,k}$), to extract a refutation of a subset of m' equations from $F_{n,k}$ of low degree, contradicting the degree lower bound given in Lemma 19. Recall that $F_{n,k}$ is defined over variables $\mathbf{x}$ and we pick a subset of these equations by matching pigeons y_i to equations in $F_{n,k}$ through a complete bipartite graph.

The algorithm first initializes a few things. Set $d = \nu n/\kappa$ for a small enough constant $\nu > 0$. Let $M' = M + n \log n$, which upper bounds the total number of variables occurring in the refutation. Let S be the set of all variables that are Singular w.r.t. the current set of extension axioms. We initialize S to be the set of all variables $\mathbf{x} \cup \mathbf{y} \cup \mathbf{z}$ since this is the largest possible set we will be dealing with; this will be updated at every iteration of the while loop, although we note that it can only reduce as we apply restrictions. Henceforth when we refer to Quadratic degree, we mean $Qdeg^S$. Finally, we initialize H to be the set of all pairs of terms in Π with Quadratic degree greater than d.

In the while loop, we first update the set S by checking which of the extension variables z have zero in their support according to their current extension axioms, and deleting those that don't. For each extension variable z that we delete from S, we reduce the refutation Π by $z^{\ell(z)} = c$. Such an identity exists and is derivable by Lemma 4, and does not increase the size of H by Lemma 8. Once we have updated S, we recompute the set of high Quadratic degree pairs H with respect to the updated set S. This also does not increase the size of H, by Lemma 5. We then pick a variable w that contributes to the Quadratic degree of at least a d/M' fraction of pairs in H by averaging.

There are two cases depending on whether $w \in S$ or not. In the first case (lines 14-18), w is Singular so we apply the restriction σ such that $w|_\sigma = 0$ and call **Cleanup**(σ) to restore to a clean version of our tautology. This eliminates the contribution to high Quadratic degree from terms containing w, and hence obtains a $(1 - d/M')$-factor reduction in the size of H. In the second case (lines 19-34), w is Nonsingular so we apply Algorithm 1, which uses the Split operation non-trivially to reduce the size of H. Lemma 15 proves correctness of the algorithm, and thus upon termination of one call to Algorithm 1, we have obtained a $(1 - d/(4p^2M'))$-factor reduction in the number of high Quadratic degree terms.

Repeating the above for $-\log|H|/\log(1 - d/4p^2M') \approx 4p^2M' \log|H|/d \le O(\gamma)\kappa n/10^\kappa$ iterations, we eliminate all terms in H from the proof and thus obtain a refutation of Quadratic degree less than d. Since we call **Cleanup** once per iteration, and in each call it bans at most $O(\kappa)$ many equations and removes at most $O(\kappa)$ many pigeons (by Lemma 14), we have banned at most $O(\gamma)\kappa^2 n/10^\kappa$ equations and removed at most those many pigeons in total. Therefore, we always satisfy the invariant that the number of banned equations is much less than $m/2^\kappa$ (where $m = 10n$), satisfying the required conditions for correctness of **Cleanup** from Lemma 14.

Let Π' denote the modified proof upon termination of Algorithm 2. Note that out of the $m' = (1 - \epsilon)m$ pigeons, there are at least a $1 - O(\gamma)$ fraction of pigeons still alive (i.e. not removed by **Cleanup**) and a $1 - O(\gamma)$ fraction of the m equations not banned. We now substitute for the remaining pigeons $\mathbf{y}$ so that we select a subset of at least $(1 - 2\epsilon)m$ unsatisfiable equations from $F_{n,k}$ that are not banned and obtain a refutation of them of Quadratic degree at most d (assuming γ is small enough). By Lemma 9, we can obtain a refutation of these equations of degree at most $3pd$. Now, for all surviving extension variables we substitute them with their definitions in terms of the variables $\mathbf{x}$. Note that since each extension variable is a degree κ polynomial this raises the degree to at most $3\kappa pd$. Since

$d = \nu n/\kappa$, for sufficiently small ν we end up with a refutation of $(1-2\epsilon)$ equations from $F_{n,k}$ of degree less than $c_2 n$, contradicting Lemma 19.

◀

5 Appendix

We will prove Theorem 3, which we state again here for convenience.

▶ **Theorem 17** (Theorem 3). *Let $m = 10n$. Then there exists constants $k > 0$, $0 < \epsilon < 1$ such that for sufficiently large n, there exists k-CSP formulas $\{F_{n,k}\}$ with m k-local constraints such that for $m' = (1-\epsilon)m$, every subset of m' constraints is unsatisfiable and requires linear degree* PC *refutations.*

First we'll show that a random regular bipartite graph has good boundary expansion. This has been used implicitly in other works ([9], [5]), but for completeness we state and prove it here. Let $G = (L, R, E)$ be a bipartite graph, and let $A \subseteq R$. The *boundary* for A, $\partial(A)$, is the set of vertices x in L so that $|N(x) \cap A| = 1$, i.e., vertices with a unique neighbor in A. A bipartite graph is (d, k) regular if every vertex in L has degree d and every vertex in R has degree k. In this case, for $n = |L|, m = |R|$, we have $dn = km$.

▶ **Theorem 18.** *Let d, k, n, m be positive integers with $dn = km$, $k \geq 12$. Then with high probability for a random (d, k) regular bipartite graph with $|L| = n, |R| = m$, for all $A \subset R$, $|A| < n/(e^6 k^2)$, we have $\partial(A) \geq k|A|/2$.*

Proof. Let $N(A)$ be all the neighbors of A. Since the total degrees of vertices in A is $k|A|$, and each element of $N(A) - \partial(A)$ is contingent on two such edges, $k|A| \geq 2(|N(A)| - |\partial(A)|) + |\partial(A)$, or $\partial(A) \geq 2|N(A)| - k|A|$. We will show that with high probability for all such A, $|N(A)| > 3k|A|/4$, and hence $\partial(A) \geq k|A|/2$.

If not, there are sets $A \subset R$ and $B \subset L$ so that $N(A) \subseteq B$ and $|B| = 3k|A|/4$. We will bound the probability that this is true for fixed sets A, B and then take a union bound. We can view picking a random (d, k) bipartite graph as picking a random matching between d half-edges adjacent to each $x \in L$ and k such half-edges adjacent to each $y \in R$; if a half edge for x is matched to a half-edge for y, it forms an edge between x and y.

We can form this matching by going through the half edges for nodes in R and for each randomly selecting an unmatched half-edge for some node in L. We start with the edges for A in an arbitrary order. If we condition on all previous neighbors for A being in B, the number of half-edges left still available for B is less than $d|B|$, whereas the number for $\overline{B}$ stays at exactly $d(n - |B|)$. Thus, the conditional probability that the next edge formed is also in B is at most $|B|/n$, and we do this for each of $k|A|$ edges, meaning the probability that all neighbors are in B is at most $(|B|/n)^{k|A|}$.

Now, for a fixed $|A|$ and setting $|B| = 3k|A|/4$, we take the union bound over all subsets A and B. This gives a total probability of failure for some set A of size a as :

$$\binom{m}{a}\binom{n}{3ka/4}(3ka/4n)^{ka}$$

$$\leq (em/a)^a (4en/3ka)^{3ka/4} (3ka/4n)^{ka}$$

$$\leq (em/a)^a (e^3 ka/n)^{ka/4} = (ekn/da)^a (e^3 ka/n)^{ka/4} = (e^{3k/4+1} a^{k/4-1} k^{k/4+1}/dn^{k/4-1})^a$$

Since we are assuming $a < n/(e^6 k^2)$, the base in the above expression is at most

$$e^{3k/4+1}(n/e^6 k^2)^{k/4-1} k^{k/4+1}/dn^{k/4-1}$$

$$= e^{7-3k/4}k^{3-k/4}/d$$

which for $k \geq 12$ is bounded below e^{-2}, meaning the probability of such a bad set existing is exponentially small in a, and the probability of such a bad set existing for any a is less than $1/2$.

◀

▶ **Definition 15.** *For a Boolean vector $X = \{x_1, \ldots, x_n\}$, we define $\mathcal{L}_{n,m,k_1,k}(X)$ to be the distribution over k-CSP formulas over n variables $X = \{x_1, \ldots, x_n\}$ obtained by selecting m parity equations, where each parity is represented by a node on the right of a randomly chosen bipartite graph $G(L, R, E)$, with $|L| = n$, $|R| = m$, and with left degree bounded by k_1 and right degree bounded by k.*

▶ **Lemma 19.** *Let $F_{n,k}$ be a tautology given by the system of parity equations $AX = b$ over variables $X = \{x_1, \ldots, x_n\}$ drawn at random from $\mathcal{L}_{n,m,k_1,k}$ where $m = 10n$, for large enough constants $k_1, k > 0$, and b is chosen randomly. Then the following hold with high probability for a small enough $\epsilon > 0$:*

a) Any subset of a $(1 - \epsilon)$-fraction of the equations in $F_{n,k}$ is unsatisfiable
b) Any subset of a $(1 - \epsilon)$-fraction of the equations in $F_{n,k}$ requires PC *degree $c_2(n)$ to refute, for some $c_2 > 0$.*

Proof. a) The probability that a set of $(1 - \epsilon)10n$ random parities (i.e. for a random choice of b) is satisfiable is at most 2^{-9n} for a small enough ϵ. The probability that any such subset of $F_{n,k}$ is satisfiable is therefore at most $2^{(-n(9-10H(\epsilon)))}$, which is exponentially small for a small enough ϵ (where $H(\epsilon)$ is the binary entropy function).

b) This follows directly from [1], Theorem 3.8 and Theorem 4.4, since by Theorem 18 the bipartite graph underlying the system of parity equations A has good boundary expansion with high probability.

◀

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. doi:10.1109/SFCS.2001.959893.

2 Yaroslav Alekseev. A lower bound for polynomial calculus with extension rule. In Valentine Kabanets, editor, *36th Computational Complexity Conference, CCC 2021, July 20-23, 2021, Toronto, Ontario, Canada (Virtual Conference)*, volume 200 of *LIPIcs*, pages 21:1–21:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.CCC.2021.21.

3 Robert Andrews and Michael A. Forbes. Ideals, determinants, and straightening: Proving and using lower bounds for polynomial ideals. *CoRR*, abs/2112.00792, 2021. URL: https://arxiv.org/abs/2112.00792, arXiv:2112.00792.

4 Paul Beame, Russell Impagliazzo, Jan Krajíček, Toniann Pitassi, and Pavel Pudlák. Lower bounds on hilbert's nullstellensatz and propositional proofs. *Proceedings of the London Mathematical Society*, 3(1):1–26, 1996.

5 Paul Beame, Richard Karp, Toniann Pitassi, and Michael Saks. The efficiency of resolution and davis–putnam procedures. *SIAM Journal on Computing*, 31(4):1048–1075, 2002.

6 Eli Ben-Sasson and Russell Impagliazzo. Random cnfs are hard for the polynomial calculus. In *40th Annual Symposium on Foundations of Computer Science (Cat. No. 99CB37039)*, pages 415–421. IEEE, 1999.

7 Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow – resolution made simple. In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pages 517–526, 1999.

8 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.

9 Vašek Chvátal and Endre Szemerédi. Many hard examples for resolution. *Journal of the ACM (JACM)*, 35(4):759–768, 1988.

10 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.

11 Stephen A Cook and Robert A Reckhow. The relative efficiency of propositional proof systems. *The journal of symbolic logic*, 44(1):36–50, 1979.

12 Stefan S. Dantchev and Søren Riis. On relativisation and complexity gap. In Matthias Baaz and Johann A. Makowsky, editors, *Computer Science Logic, 17th International Workshop, CSL 2003, 12th Annual Conference of the EACSL, and 8th Kurt Gödel Colloquium, KGC 2003, Vienna, Austria, August 25-30, 2003, Proceedings*, volume 2803 of *Lecture Notes in Computer Science*, pages 142–154. Springer, 2003. doi:10.1007/978-3-540-45220-1_14.

13 Michael A. Forbes, Amir Shpilka, Iddo Tzameret, and Avi Wigderson. Proof complexity lower bounds from algebraic circuit complexity. *Theory Comput.*, 17:1–88, 2021. URL: https://theoryofcomputing.org/articles/v017a010/.

14 Nashlen Govindasamy, Tuomas Hakoniemi, and Iddo Tzameret. Simple hard instances for low-depth algebraic proofs. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022*, pages 188–199. IEEE, 2022.

15 Armin Haken. The intractability of resolution. *Theoretical computer science*, 39:297–308, 1985.

16 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.

17 Russell Impagliazzo, Pavel Pudlák, and Jirí Sgall. Lower bounds for the polynomial calculus and the gröbner basis algorithm. *Comput. Complex.*, 8(2):127–144, 1999. doi:10.1007/s000370050024.

18 Alexander A Razborov. Lower bounds for the size of circuits of bounded depth with basis fˆ; g. *Math. notes of the Academy of Sciences of the USSR*, 41(4):333–338, 1987.

19 Roman Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 77–82, 1987.

20 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.