Lower Bounds against the Ideal Proof System in Finite Fields^{*}

Tal Elbaz[†] Nashlen Govindasamy[‡] Jiaqi Lu[§] Iddo Tzameret[¶]

Imperial College London Department of Computing

Abstract

Lower bounds against strong algebraic proof systems and specifically fragments of the Ideal Proof System (IPS), have been obtained in an ongoing line of work. All of these bounds, however, are proved only over large (or characteristic 0) fields,¹ yet finite fields are the more natural setting for propositional proof complexity, especially for progress toward lower bounds for Frege systems such as $AC^0[p]$ -Frege. This work establishes lower bounds against fragments of IPS over fixed finite fields. Specifically, we show that a variant of the knapsack instance studied by Govindasamy, Hakoniemi, and Tzameret (FOCS'22) has no polynomial-size IPS refutation over finite fields when the refutation is multilinear and written as a constant-depth circuit. The key ingredient of our argument is the recent set-multilinearization result of Forbes (CCC'24), which extends the earlier result of Limaye, Srinivasan, and Tavenas (FOCS'21) to all fields. We also separate this proof system from the one studied by Govindasamy, Hakoniemi, and Tzameret.

In addition, we present new lower bounds for read-once algebraic branching program refutations, roABP-IPS, in finite fields, extending results of Forbes, Shpilka, Tzameret, and Wigderson (Theor. of Comput.'21) and Hakoniemi, Limaye, and Tzameret (STOC'24).

Finally, we show that any lower bound against any proof system at least as strong as (nonmultilinear) constant-depth IPS over finite fields for *any* instance, even a purely algebraic instance (i.e., not a translation of a Boolean formula or CNF), implies a hard *CNF formula* for the respective IPS fragment, and hence an $AC^0[p]$ -Frege lower bound by known simulations over finite fields (Grochow and Pitassi (J. ACM'18)).

Note on independent concurrent work: Independently and concurrently with our work, Behera, Limaye, Ramanathan, and Srinivasan [BLRS25] using different arguments, obtained related results for fragments of IPS over fields of positive characteristic. Both works establish a lower bound for *constant-depth multilinear* IPS but the field assumptions differ: [BLRS25] requires the size of the field to grow with the instance, whereas our lower bound holds for any field of constant-depth multilinear (or small) positive characteristic. Hence, in the constant positive characteristic setting our constant-depth multilinear IPS lower bound strictly subsumes theirs as it also holds over any fixed finite field. We provide a detailed comparison of these and, where relevant, other results of [BLRS25] in Section 1.4.

^{*}This project was supported by the Engineering and Physical Sciences Research Council (grant EP/Z534158/1) and the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant agreement 101002742; EPRICOT project).

[†]E-mail: t.elbaz22@imperial.ac.uk

[‡]E-mail: nashlen.govindasamy@gmail.com

[§]E-mail: jiaqi.lu23@imperial.ac.uk

[¶]E-mail: iddo.tzameret@gmail.com. https://www.doc.ic.ac.uk/~itzamere/

¹Except for the *placeholder* lower bound model, where the instance itself lacks small circuits [HLT24].

1 Introduction

This work investigates lower bounds against the Ideal Proof System (IPS) over finite fields, motivated by two main considerations. First, existing lower bounds for IPS have not adequately addressed the case of finite fields. Second, focusing on finite fields—rather than large fields—offers a more natural setting for tackling a central open problem in proof complexity: proving superpolynomial lower bounds against $AC^0[p]$ -Frege.

1.1 Algebraic Proof Complexity

Proof complexity studies the size of proofs that certify membership in languages such as UNSAT, the set of unsatisfiable Boolean formulas. In this context, a proof is a witness that can be verified efficiently, and for UNSAT, such a proof is typically called a refutation. A central objective of the field is to establish lower bounds against increasingly powerful proof systems, with the overarching goal of demonstrating that no proof system admits polynomial-size refutations for all unsatisfiable formulas. This approach is often referred to as *Cook's Programme*, following Stephen Cook's suggestion in the 1970s that proof complexity lower bounds could yield insights into fundamental questions in computational complexity, such as the P versus NP problem. In particular, showing that no proof system can efficiently refute all unsatisfiable formulas would separate NP from coNP and thereby separate P from NP.

An important thread of proof complexity is to investigate algebraic proof systems, which certify that a given set of multivariate polynomials over a field has no common Boolean solution. Some of the foundational proof systems in this line are the Polynomial Calculus (PC) [CEI96] and its 'static' variant, Nullstellensatz [BIKPP96]. In PC, proofs proceed by algebraic manipulation, adding and multiplying polynomials, until deriving the contradiction 1 = 0. Contrastingly, in Nullstellensatz, a proof of the unsatisfiability of a set of axioms, written as polynomial equations $\{f_i(\bar{x}) = 0\}$ over a field, is a single polynomial identity expressing 1 as a combination of the axioms, that is:

$$\sum_{i} g_i(\overline{x}) \cdot f_i(\overline{x}) = 1, \qquad (1)$$

for some polynomials $\{g_i(\bar{x})\}$. These systems measure proof size by sparsity, defined as the total number of monomials involved, which makes them comparatively weak. An alternative way to measure proof size is by algebraic circuit size. This was suggested initially by Pitassi [Pit97; Pit98], and further investigated in the work of Grigoriev and Hirsch [GH03] and subsequently Raz and Tzameret [RT08b; RT08a], eventually leading to the Ideal Proof System [GP18] described in what follows.

1.2 Ideal Proof System

The Ideal Proof System (IPS, for short; Definition 9), introduced by Grochow and Pitassi [GP18], loosely speaking is the Nullstellensatz proof system where the polynomials $g_i(\bar{x})$ in (1) are represented by algebraic circuits. Formally, Forbes, Shpilka, Tzameret and Wigderson [FSTW21] showed that IPS is equivalent to Nullstellensatz in which the polynomials g_i in Equation (1) are written as algebraic circuits. In other words, an IPS refutation of the set of axioms $\{f_i(\bar{x}) = 0\}_i$ can be defined similarly to Equation (1) (here we display explicitly the Boolean axioms $x_j^2 - x_j$):

$$\sum_{i} g_{i}(\bar{x}) \cdot f_{i}(\bar{x}) + \sum_{j} h_{j}(\bar{x}) \cdot \left(x_{j}^{2} - x_{j}\right) = 1,$$
(2)

for some polynomials $\{g_i(\bar{x})\}_i$, where we think of the polynomials g_i, h_j written as algebraic circuits (instead of e.g., counting the number of monomials they have towards the size of the refutation). Thus, the size of the IPS refutation in Equation (0.2) is $\sum_i \text{size} (g_i(\bar{x})) + \sum_j \text{size} (h_j(\bar{x}))$, where size(g) stands for the (minimal) size of an algebraic circuit computing the polynomial g.

When considering algebraic circuit classes weaker than general algebraic circuits, one has to be a bit careful with the definition of IPS. For technical reasons the formalization in (2) does not capture the precise definition of IPS restricted to the relevant circuit class, rather the fragment which is denoted by C-IPS_{LIN} ("LIN" here stands for the linearity of the axioms f_i and the Boolean axioms; that is, they appear with power 1). In this work, we focus on C-IPS_{LIN} and a similar stronger variant denoted C-IPS_{LIN}'. Throughout the introduction, refutations in the system C-IPS_{LIN} are defined as in Equation (2) where the polynomials g_i, h_j are written as circuits in the circuit class C.

Technically, our lower bounds are proved by lower bounding the algebraic circuit size of the g_i 's in (2), namely the products of the axioms f_i , and not the products of the Boolean axioms (that is, we ignore the circuit size of the h_i 's). For this reason, our lower bounds are slightly stronger than lower bounds on C-IPS_{LIN}, rather they are lower bounds on the system denoted C-IPS_{LIN}' (see Definition 9).

Lower bounds methods and known results. Forbes, Shpilka, Tzameret, and Wigderson [FSTW21] introduced two approaches for turning algebraic circuit lower bounds into lower bounds for IPS: the *functional lower bound* method and the *lower bounds for multiples* method. Of the two, the functional approach has proved more instrumental, proving several concrete proof complexity lower bounds against fragments of IPS. These include lower bounds for variants of the subsetsum instance against IPS refutations written as read-once (oblivious) algebraic branching programs (roABPs), depth-3 powering formulas, and multilinear formulas [FSTW21]. A similar method underpinned the *conditional* lower bound against general IPS established by Alekseev, Grigoriev, Hirsch, and Tzameret [AGHT20] (leading to [Ale21]). Govindasamy, Hakoniemi, and Tzameret [GHT22] combined the functional method with the constant-depth algebraic circuit lower bound result of Limaye, Srinivasan, and Tavenas [LST21], obtaining constant-depth multilinear IPS lower bounds.

By contrast, the multiples method has so far matched the functional method only within the weaker *placeholder* model of IPS, where the hard instances themselves do not have small circuits in the fragment under study [FSTW21; AF22]. Other approaches have emerged as well: the *meta-complexity* approach of Santhanam and Tzameret [ST25], which obtains a *conditional* IPS size lower bound on a self-referential statement; the *noncommutative* approach of Li, Tzameret, and Wang [LTW18] (building on [Tza11], which reduced Frege lower bounds to matrix-rank lower bounds but has yet to yield concrete lower bounds; and recent lower bounds against *PC with extension variables* over finite fields of Impagliazzo, Mouli, and Pitassi [IMP23] (building on [Sok20] and improved by [DMM24]) which can be considered as a fragment of IPS sitting between depth-2 and depth-3.

The functional lower bound method was further investigated by Hakoniemi, Limaye, and Tzameret [HLT24]. There, Nullstellensatz degree lower bounds for symmetric instances and vector invariant polynomials were established, which were then lifted to yield IPS size lower bounds for the roABP and multilinear formula fragments of IPS. With invariant polynomials, the bounds hold over *finite fields*, though within the *placeholder* model. Building on recent advances in constantdepth algebraic circuit lower bounds from [AGKST23], they extend [GHT22] to constant-depth IPS refutations computing polynomials with $O(\log \log n)$ individual degree. Finally, they observe a barrier in that the functional method cannot yield lower bounds for any Boolean instance against sufficiently strong proof systems like constant-depth IPS.

1.3 IPS over Finite Fields

IPS lower bounds have so far been obtained almost exclusively over fields of large characteristic. Finite fields, however, are a more natural setting for propositional proof complexity, particularly for the long-standing open problem of establishing super-polynomial lower bounds for $AC^0[p]$ -Frege. This proof system operates with constant-depth propositional formulas equipped with modulo pcounting gates, where p is a prime. Grochow and Pitassi [GP18] showed that constant-depth IPS refutations over \mathbb{F}_p simulate $AC^0[p]$ -Frege, thus obtaining lower bounds against IPS over finite fields provides a concrete approach to settle the problem of obtaining lower bounds against $AC^0[p]$ -Frege. Although lower bounds against $AC^0[p]$ -Frege are sometimes thought to be within reach of current techniques, especially given existing lower bounds against both AC^0 -Frege and $AC^0[p]$ circuits, this problem and the problem of obtaining lower bounds against constant-depth IPS over \mathbb{F}_p remain elusive.

IPS lower bounds over finite fields face additional challenges that are not present in the characteristic 0 setting. A recurring obstacle is provided by Fermat's little theorem: for a nonzero $a \in \mathbb{F}_p$, $a^{p-1} = 1$ in \mathbb{F}_p . More generally, if \mathbb{F} is a finite field of size q, then for a nonzero $a \in \mathbb{F}$, $a^{q-1} = 1$. Hence, if a polynomial $f \in \mathbb{F}[\overline{x}]$ admits no satisfying Boolean assignment, then $(f(\overline{x}))^{(q-2)}f(\overline{x}) = (f(\overline{x}))^{(q-1)} = 1$ over Boolean assignments. The functional lower bound method of [FSTW21] requires a lower bound on the size of circuits computing $g(\overline{x})$ such that $g(\overline{x})f(\overline{x}) = 1$ over Boolean assignments, hence requires a lower bound on the size of $(f(\overline{x}))^{(q-2)}$. Thus we must simultaneously ensure that the hard instance $f(\overline{x})$ is easily computed by the subsystem of IPS under consideration while $(f(\overline{x}))^{(q-2)}$ is hard in that same subsystem. While this is not possible for proof systems closed under constant multiplication of polynomials, including certain constant-depth IPS subsystems (for example the one studied in [HLT24] which considered log log n individual degree refutations), for the multilinear constant-depth IPS subsystem considered in [GHT22], it is indeed possible, even for constant q.

1.4 Our Results

1.4.1 Bounds for Constant-depth IPS over Finite Fields

Our first contribution establishes a super-polynomial lower bound for constant-depth $\mathsf{IPS}_{\text{LIN}'}$ refutations over finite fields. As mentioned above, $\mathsf{IPS}_{\text{LIN}'}$ is the Nullstellensatz proof system (2) whose refutations are algebraic circuits (see Definition 9). This result is the finite field analogue of [GHT22], which was proved over characteristic 0 fields. Our hard instance is the knapsack mod ppolynomial $\mathsf{ks}_{w,p}$, a variant of the knapsack polynomial ks_w used in their work.

Theorem 1 (Informal; see Theorem 24). Let $p \ge 5$ be a prime, and let \mathbb{F} be a field of characteristic p. Every constant-depth multilinear $\mathsf{IPS}_{\mathrm{LIN}'}$ refutation over \mathbb{F} of the knapsack mod p instance $\mathsf{ks}_{w,p}$ requires super-polynomial size.

The proof in [GHT22] combines two main ingredients: first, the methods used by Limaye, Srinivasan, and Tavenas [LST21] to establish super-polynomial lower bounds for constant-depth algebraic circuits; and second, the functional lower bound framework of [FSTW21] for size lower bounds on IPS proofs. We adopt the same overall strategy, showing how the finite field setting introduces additional obstacles, and how we circumvent them. Following [GHT22], we reduce the task of lower-bounding the size of a constant-depth algebraic circuit computing the multilinear polynomial that constitutes the IPS proof into the task of lowerbounding the size of a constant-depth *set-multilinear circuit* computing a related *set-multilinear polynomial*. We derive this set-multilinear polynomial from the original multilinear IPS proof (which is not necessarily set-multilinear by itself) via the same variant of the functional lower bound method used in [GHT22].

[GHT22] subsequently applies a reduction presented in [LST21], which converts constant-depth general circuits into constant-depth set-multilinear circuits. Because the reduction presented in [LST21] requires fields of sufficiently large characteristic, we rely on the recent extension of Forbes [For24], which shows that this set-multilinearization reduction holds over all fields thereby removing this obstacle. [GHT22] also invokes another reduction presented in [LST21] from a size lower bound of a set-multilinear formula into a rank lower bound of its coefficient matrix. This second reduction already holds over all fields, and we use the improved parameters obtained by Bhargav, Dutta, and Saxena [BDS24].

The problem therefore reduces to constructing an unsatisfiable instance whose refutations, after the preceding reductions, have full rank. [GHT22] achieves this by introducing the knapsack polynomial, an instance that embeds a family of subset-sum instances $(\sum_{i=1}^{n} x_i - \beta = 0, \text{ for } \beta > n)$. They then use the *full* degree lower bound established in [FSTW21] for refutations of subset-sum instances to obtain the required full rank lower bound. Because the knapsack polynomial is satisfiable over Boolean assignments in finite fields, our task is to design a hard instance that both admits no satisfying Boolean assignment in finite fields and embeds a family of subset-sum-type instances that require full degree to refute. We proceed in two steps: first, we extend the full degree bound of [FSTW21] to more general subset-sum-type instances; and second, we introduce a variant of the knapsack instance, knapsack mod p, that embeds a family of these more general subset-sum-type instances. Thus we obtain the result. Although the theorem is stated over finite fields, it also holds over characteristic 0 fields, thereby providing additional hard instances for the proof system studied on [GHT22].

As noted earlier, [BLRS25] proves a lower bound for the same proof system as Theorem 1, but under different field assumptions. [BLRS25] requires size of the field to grow with the instance, whereas our lower bound holds for any field of constant positive characteristic. Consequently, in the constant positive characteristic setting, our result subsumes that of [BLRS25] as it also holds over fixed finite fields. By contrast, the [BLRS25] lower bound also covers characteristic 2 and 3 fields, which our result does not.

Both our work and [BLRS25] also establish upper bounds for subsystems of constant-depth IPS over fields of positive characteristic. While [BLRS25] proves a general upper bound for systems stronger than constant-depth multilinear $IPS_{LIN'}$, the system for which we obtain a superpolynomial lower bound, our upper bound is for a single explicit instance within the same system. Moreover, our specific instance is hard to refute in constant depth multilinear $IPS_{LIN'}$ over characteristic 0 fields, where a corresponding lower bound was shown in [GHT22]. Hence we obtain the first separation between constant depth multilinear $IPS_{LIN'}$ over finite fields and the same system over characteristic 0 fields.

Our separating instance ks_{w,e_2} , the symmetric knapsack of degree 2, is another variant of the knapsack instance used in [GHT22]. Note that the subset-sum instance can be viewed as an elementary symmetric sum of degree 1. In the same spirit as the knapsack polynomial, ks_{w,e_2} is designed so that it embeds a family of elementary symmetric sum of degree 2.

Theorem 2 (Informal; see Theorem 36). Let $p \ge 3$ be a prime, and let \mathbb{F} be a field of characteristic p. Then, for the symmetric knapsack ks_{w,e_2} of degree 2:

- ks_{w,e_2} has no satisfying Boolean assignment over \mathbb{F} , and over any field of characteristic 0;
- there is a polynomial-size, constant-depth multilinear $\mathsf{IPS}_{\mathrm{LIN}'}$ refutation of ks_{w,e_2} over \mathbb{F} ;
- for every characteristic 0 field E, every constant-depth multilinear $\mathsf{IPS}_{\mathrm{LIN}'}$ refutation over E of ks_{w,e_2} requires super-polynomial size.

For our separation, we need an instance that has no Boolean satisfying assignment in either field, both the finite field and the characteristic 0 field, rather than one whose satisfiability depends on the characteristic. We establish this for ks_{w,e_2} by showing that elementary symmetric sums of certain degrees have no Boolean satisfying assignment in finite fields. As ks_{w,e_2} admits no satisfying Boolean assignment in the finite field, it likewise admits none in the characteristic 0 field.

The upper bound for refutations of ks_{w,e_2} in constant-depth multilinear $\mathsf{IPS}_{\mathsf{LIN}'}$ over finite fields follows from Fermat's little theorem. What remains is the lower bound for ks_{w,e_2} over characteristic 0 fields. We show that, in characteristic 0, every refutation of the elementary symmetric sum of degree 2 must have full degree. The corresponding IPS lower bound for ks_{w,e_2} then follows similarly to the argument in [GHT22].

1.4.2 roABP-IPS Lower Bounds over Finite Fields

We also present new lower bounds for roABP-IPS_{LIN'} over finite fields, using two distinct techniques: the Functional Lower Bound method and the Lower Bound by Multiples method. In both cases, we obtain finite-field analogues of results from [HLT24] and [FSTW21] respectively, which originally required fields of large characteristic. Moreover, our proofs are significantly simpler. As a first step, we establish an exponential lower bound for roABP-IPS_{LIN'} in any variable order.

Theorem 3 (Informal; see Corollary 43). Let \mathbb{F}_q be a finite field with constant characteristic q. Then, there exists a polynomial $f \in \mathbb{F}[\overline{w}]$ such that any roABP-IPS_{LIN'} refutation (in any variable order) of f requires $2^{\Omega(n)}$ -size.

This proof employs the Functional Lower Bound method and closely follows the strategy of [HLT24]. As in their work, we first establish a lower bound in a fixed variable order, and then extend the result to any order. However, our hard instance differs from theirs—this not only simplifies the argument, but also allows us to prove the result over fields of constant characteristic. Additionally, we provide a lower bound for an unsatisfiable system of equations.

Theorem 4 (Informal; see Theorem 46). Let \mathbb{F}_q be a finite field of constant characteristic q. Then, there exist polynomials $f, g \in \mathbb{F}[x_1, \ldots, x_n]$ such that the system of equations $f, g, \overline{x}^2 - \overline{x}$ is unsatisfiable, and any roABP-IPS_{LIN'} refutation (in any order of the variables) requires size $\exp(\Omega(n))$.

In this case, we apply the Lower Bound by Multiples method from [FSTW21], and extend their result to finite fields. Our hard system of equations uses the same polynomial f as in their work, but a different choice of g, which allows us to avoid their reliance on large characteristic fields.

We emphasize that the lower bounds we obtain for roABP-IPS_{LIN'} are *placeholder* lower bounds—that is, the hard instances considered are not efficiently computable by roABPs. This makes the model strictly weaker than the non-placeholder setting. In fact, we show that it is *impossible* to obtain non-placeholder lower bounds for roABP-IPS_{LIN'} over finite fields using the functional lower bound method (see [HLT24, Theorem 1 in full version] for a precise definition of "the functional lower bound method").

Theorem 5 (Theorem 49). The functional lower bound method cannot establish non-placeholder lower bounds on the size of roABP-IPS_{LIN'} refutations when working in finite fields.

1.4.3 Towards Hard CNF Formulas

Several lower bounds are known for purely algebraic instances against subsystems of IPS. This raises an important question: Could we get lower bounds for CNF formulas against subsystems of IPS from those lower bounds?

Note that an instance consisting of a set of polynomials written as *circuit equations* $\{f_i(\bar{x}) = 0\}_i$, for $f_i(\bar{x}) \in \mathbb{F}[\bar{x}]$, does not necessarily correspond to a Boolean instance or a CNF formula. Specifically, we say that such an instance is *Boolean* whenever $f_i(\bar{x}) \in \{0,1\}$ for $\bar{x} \in \{0,1\}^{|\bar{x}|}$. For example, a CNF written as a set of (polynomials representing) clauses is a Boolean instance. Similarly, the standard arithmetization of propositional formulas leads to Boolean instances. On the other hand, the instances used in Theorem 1 as well as the standard subset sum $\sum_i x_i - \beta$ is non-Boolean, and thus said to be "purely algebraic": the image of the latter under $\{0, 1\}$ -assignments is $\{-\beta, 1-\beta, \ldots, n-\beta\}$, and thus cannot be considered a propositional or Boolean formula (formally, there is no known way to yield propositional proof lower bounds, say, in constant-depth Frege, from lower bounds for such purely algebraic instance even when such lower bounds are against proof systems that simulate constant-depth Frege).

We solve this problem, and show how to attain propositional proof lower bounds from purely algebraic instances lower bounds. This is done using efficient bit-arithmetic in finite fields: from a circuit we derive the statements that express its gate-by-gate bit-arithmetic description. we establish a *translation lemma*—that is, we show that CNF encoding can be efficiently derived from circuit equations and vice versa within these subsystems of IPS in finite fields. If a subsystem of IPS can efficiently derive the CNF encoding and then refute it, a lower bound for circuit equations implies a lower bound for CNF formulas.

In [ST25], Santhanam and Tzameret presented a translation lemma with *extension axioms* in IPS. In other words, given some additional axioms, IPS can efficiently derive the CNF encoding for circuit equations and vice versa. We eliminate the need to add additional extension axioms and extension variables altogether: we show that without those additional axioms, already bounded-depth IPS over a finite field can efficiently derive the CNF encoding for bounded-depth circuit equations. Following our translation lemma, every superpolynomial lower bound for bounded-depth circuit equations against bounded-depth IPS implies a superpolynomial lower bound for CNF formulas against bounded-depth IPS over a finite field, and hence an $AC^0[p]$ -Frege lower bound following standard simulation of $AC^0[p]$ -Frege by constant-depth IPS over \mathbb{F}_p .

We now explain our translation lemma. [ST25] used *unary* encoding to encode CNFs for circuit equations over finite fields. Each variable x over a finite field \mathbb{F}_q corresponds to q bits x_{q-1}, \ldots, x_0 where x_j equals 1 for $0 \le j \le q-1$ if and only if x = j; thus, these q bits can be viewed as "unary bits".

We use the Lagrange polynomial

$$\frac{\prod_{i=0, i\neq j}^{q-1} (x-i)}{\prod_{i=0, i\neq j}^{q-1} (j-i)}$$

to express each unary bit x_i with variable x, which we call UBIT

$$\mathsf{UBIT}_{j}(x) = \begin{cases} 1, & x = j, \\ 0, & \text{otherwise.} \end{cases}$$

We introduce a notation called semi-CNFs, which are CNFs where each Boolean variable is substituted by the corresponding UBIT. Hence, SCNFs are substitution instances of CNFs, which means a lower bound for SCNFs implies a lower bound for CNFs against sufficiently strong subsystems of IPS, including bounded-depth IPS. We show that the semi-CNF encoding of all the extension axioms in [ST25] can be efficiently proved in bounded-depth IPS over finite fields. Following the proof in [ST25], bounded-depth IPS can efficiently derive the semi-CNFs encoding of circuit equations. Hence, a lower bound for circuit equations implies a lower bound for CNFs.

Theorem 6 (Corollary 60). Let \mathbb{F}_q be a finite field, and let $\{C(\overline{x})\}$ be a set of circuits of depth at most Δ in the Boolean variable \overline{x} . Then, if a set of circuit equations $\{C(\overline{x}) = 0\}$ cannot be refuted in S-size, $O(\Delta')$ -depth IPS, then the CNF encoding of the set of circuit equations $\{\mathsf{CNF}(C(\overline{x}) = 0)\}$ cannot be refuted in $(S - \mathsf{poly}(|C|))$ -size, $O(\Delta' + \Delta)$ -depth IPS.

Notice that our lower bound Theorem 1 is against constant-depth IPS refutations which are *multilinear*. Since our algebraic-to-CNF translation lemma requires non-multilinear proofs it is unclear how to carry the translation lemma for our hard instance in constant-depth multilinear IPS. For this reason we cannot apply the translation lemma to our lower bound to obtain $AC^0[p]$ -Frege lower bounds.

This aligns with the barrier discovered in [HLT24], in which proof systems closed under ANDintroduction (i.e., from a set of formulas derive their conjunction), cannot use the Functional Lower Bound method (note that our lower bound in Theorem 1 employs this method).

Bit-arithmetic arguments are used in proof complexity in many works (beginning from [Bus12], and further in works such as [AGHT20; IMP20; Gro23], and as mentioned above [ST25]). However, in all prior works the bit-arithmetic of a given circuit was not efficiently derived *within the system* from the circuits themselves, rather it was used externally to argue about certain simulations. Thus, as far as we are aware of, our result is the first that shows how to efficiently derive internally within the proof system the bit-arithmetic from a circuit.

2 Preliminaries

2.1 Polynomials and Algebraic Circuits

For excellent treatises on algebraic circuits and their complexity see Shpilka and Yehudayoff [SY10] as well as Saptharishi [Sap22]. Let \mathbb{G} be a ring. Denote by $\mathbb{G}[X]$ the ring of (commutative) polynomials with coefficients from \mathbb{G} and variables $X := \{x_1, x_2, \ldots\}$. A *polynomial* is a formal linear combination of monomials, where a *monomial* is a product of variables. Two polynomials are *identical* if all their monomials have the same coefficients.

The (total) degree of a monomial is the sum of all the powers of variables in it. The (total) degree of a polynomial is the maximal total degree of a monomial in it. The degree of an *individual* variable in a monomial is its power. The *individual degree* of a monomial is the maximal individual degree of its variables. The individual degree of a polynomial is the maximal individual degree of its monomials. For a polynomial f in $\mathbb{G}[X, Y]$ with X, Y being pairwise disjoint sets of variables, the *individual* Y-degree of f is the maximal individual degree of a Y-variable only in f.

Algebraic circuits and formulas over the ring \mathbb{G} compute polynomials in $\mathbb{G}[X]$ via addition and multiplication gates, starting from the input variables and constants from the ring. More precisely, an *algebraic circuit* C is a finite directed acyclic graph (DAG) with *input nodes* (i.e., nodes of in-degree zero) and a single *output node* (i.e., a node of out-degree zero). Edges are labelled by ring \mathbb{G} elements. Input nodes are labelled with variables or scalars from the underlying ring. In this work (since we work with constant-depth circuits) all other nodes have unbounded *fan-in* (that is, unbounded in-degree) and are labelled by either an addition gate + or a product gate \times . Every node in an algebraic circuit C computes a polynomial in $\mathbb{G}[X]$ as follows: an input node computes the variable or scalar that labels it. A + gate computes the linear combination of all the polynomials computed by its incoming nodes, where the coefficients of the linear combination are determined by the corresponding incoming edge labels. A × gate computes the product of all the polynomials computed by its incoming nodes (so edge labels in this case are not needed). The polynomial computed by a node u in an algebraic circuit C is denoted \hat{u} . Given a circuit C, we denote by \hat{C} the polynomial computed by C, that is, the polynomial computed by the output node of C. The *size* of a circuit C is the number of nodes in it, denoted |C|, and the *depth* of a circuit is the length of the longest directed path in it (from an input node to the output node). The *product-depth* of the circuit is the maximal number of product gates in a directed path from an input node to the output node.

We say that a polynomial is *homogeneous* whenever every monomial in it has the same (total) degree. We say that a polynomial is *multilinear* whenever the individual degrees of each of its variables are at most 1.

Let $\overline{x} = \langle X_1, \ldots, X_d \rangle$ be a sequence of pairwise disjoint sets of variables, called a variablepartition. We call a monomial m in the variables $\bigcup_{i \in [d]} X_i$ set-multilinear over the variablepartition \overline{x} if it contains exactly one variable from each of the sets X_i , i.e. if there are $x_i \in X_i$ for all $i \in [d]$ such that $m = \prod_{i \in [d]} x_i$. A polynomial f is set-multilinear over \overline{x} if it is a linear combination of set-multilinear monomials over \overline{x} . For a sequence \overline{x} of sets of variables, we denote by $\mathbb{F}_{sml}[\overline{x}]$ the space of all polynomials that are set-multilinear over \overline{x} .

We say that an algebraic circuit C is set-multilinear over \overline{x} if C computes a polynomial that is set-multilinear over \overline{x} , and each internal node of C computes a polynomial that is set-multilinear over some sub-sequence of \overline{x} .

2.1.1 Oblivious Algebraic Branching Programs

An algebraic branching program (ABP) is a graph-based computational model for computing multivariate polynomials, providing a structured alternative to algebraic circuits. We state the formal definition below.

Definition 7 ([Nis91]; ABP). Let \mathbb{F} be a field. An algebraic branching program (ABP) of depth D and width $\leq r$ over variables x_1, \ldots, x_n is a directed acyclic graph (DAG) with the following properties:

- 1. The vertex set is partitioned into D + 1 layers V_0, V_1, \ldots, V_D , where V_0 contains a unique source node s and V_D contains a unique sink node t.
- 2. All edges are directed from layer V_{i-1} to V_i , for $1 \le i \le D$.
- 3. Each layer satisfies $|V_i| \leq r$ for all $0 \leq i \leq D$.
- 4. Each edge e is labeled by a polynomial $f_e \in \mathbb{F}[x_1, \ldots, x_n]$.

The (individual) degree of the ABP is the maximum individual degree of any polynomial label f_e . The size of the ABP is defined as $n \cdot r \cdot d \cdot D$, where d denotes the (individual) degree. Each s-t path computes a polynomial equal to the product of the edge labels along the path. The ABP as a whole computes the sum of these polynomials over all s-t paths.

We define the following restricted variants of ABPs:

• An ABP is called oblivious if, for every layer $1 \le \ell \le D$, all edge labels between $V_{\ell-1}$ and V_{ℓ} are univariate polynomials in a single variable $x_{i_{\ell}} \in \{x_1, \ldots, x_n\}$.

- An oblivious ABP is said to be a read-once oblivious ABP (roABP) if each variable x_i appears in the edge labels of exactly one layer. In this case, we have D = n, and the layers define a variable order, which we assume to be $x_1 < x_2 < \cdots < x_n$, unless otherwise stated.
- An oblivious ABP is said to be a read-k oblivious ABP if each variable x_i appears in the edge labels of exactly k layers, so that D = kn.

We have the following fact about roABPs.

Fact 8. roABPs are closed under the following operations:

- If $f(\overline{x}, \overline{y}) \in \mathbb{F}$ is computable by a width-r roABP in some variable order then the partial substitution $f(\overline{x}, \overline{\alpha})$, for $\alpha \in \mathbb{F}^{|}\overline{y}|$, is computable by a width-r roABP in the induced order on \overline{x} , where the degree of this roABP is bounded by the degree of the roABP for f.
- If $f(z_1, \ldots, z_n)$ is computable by a width-r roABP in variable order $z_1 < \ldots < z_n$, then $f(x_1y_1, \ldots, x_ny_n)$ is computable by a poly(r, ideg f)-width roABP in variable order $x_1 < y_1 < \ldots < x_n < y_n$.

2.2 Strong Algebraic Proof Systems

For a survey about algebraic proof systems and their relations to algebraic complexity see the survey [PT16]. Grochow and Pitassi [GP18] suggested the following algebraic proof system which is essentially a Nullstellensatz proof system [BIKPP96] written as an algebraic circuit. A proof in the Ideal Proof System is given as a *single* polynomial. We provide below the *Boolean* version of IPS (which includes the Boolean axioms), namely the version that establishes the unsatisfiability over 0-1 of a set of polynomial equations. In what follows we follow the notation in [FSTW21]:

Definition 9 (Ideal Proof System (IPS), Grochow-Pitassi [GP18]). Let $f_1(\overline{x}), \ldots, f_m(\overline{x}), p(\overline{x})$ be a collection of polynomials in $\mathbb{F}[x_1, \ldots, x_n]$ over the field \mathbb{F} . An IPS **proof of** $p(\overline{x}) = 0$ from **axioms** $\{f_j(\overline{x}) = 0\}_{j \in [m]}$, showing that $p(\overline{x}) = 0$ is semantically implied from the assumptions $\{f_j(\overline{x}) = 0\}_{j \in [m]}$ over 0-1 assignments, is an algebraic circuit $C(\overline{x}, \overline{y}, \overline{z}) \in \mathbb{F}[\overline{x}, y_1, \ldots, y_m, z_1, \ldots, z_n]$ such that (the equalities in what follows stand for formal polynomial identities²; recall the notation \widehat{C} for the polynomial computed by circuit C):

- 1. $\widehat{C}(\overline{x},\overline{0},\overline{0})=0;$
- 2. $\widehat{C}(\overline{x}, f_1(\overline{x}), \dots, f_m(\overline{x}), x_1^2 x_1, \dots, x_n^2 x_n) = p(\overline{x}).$

The size of the IPS proof is the size of the circuit C. An IPS proof $C(\overline{x}, \overline{y}, \overline{z})$ of 1 = 0 from $\{f_j(\overline{x}) = 0\}_{j \in [m]}$ is called an IPS refutation of $\{f_j(\overline{x}) = 0\}_{j \in [m]}$ (note that in this case it must hold that $\{f_j(\overline{x}) = 0\}_{j \in [m]}$ have no common solutions in $\{0, 1\}^n$). If \widehat{C} is of individual degree ≤ 1 in each y_j and z_i , then this is a linear IPS refutation (called Hilbert IPS by Grochow-Pitassi [GP18]), which we will abbreviate as IPS_{LIN}. If \widehat{C} is of individual degree ≤ 1 only in the y_j 's then we say this is an IPS_{LIN'} refutation (following [FSTW21]). If $\widehat{C}(\overline{x}, \overline{y}, \overline{0})$ is of individual degree ≤ 1 in each x_j and y_i variables, while $\widehat{C}(\overline{x}, \overline{0}, \overline{z})$ is not necessarily multilinear, then this is a multilinear IPS_{LIN'} refutation.

If C is of depth at most d, then this is called a depth-d IPS refutation, and further called a depth-d IPS_{LIN} refutation if \hat{C} is linear in $\overline{y}, \overline{z}$, and a depth-d IPS_{LIN} refutation if \hat{C} is linear in \overline{y} , and depth-d multilinear IPS_{LIN} refutation if $\hat{C}(\overline{x}, \overline{y}, \overline{0})$ is linear in $\overline{x}, \overline{y}$.

²That is, $C(\overline{x}, \overline{0}, \overline{0})$ computes the zero polynomial and $C(\overline{x}, f_1(\overline{x}), \dots, f_m(\overline{x}), x_1^2 - x_1, \dots, x_n^2 - x_n)$ computes the polynomial $p(\overline{x})$.

Notice that the definition above adds the equations $\{x_i^2 - x_i = 0\}_{i=1}^n$, called the **Boolean** axioms denoted $\overline{x}^2 - \overline{x}$, to the system $\{f_j(\overline{x}) = 0\}_{j=1}^m$. This allows to refute over $\{0, 1\}^n$ unsatisfiable systems of equations. The variables $\overline{y}, \overline{z}$ are called the *placeholder variables* since they are used as placeholders for the axioms. Also, note that the first equality in the definition of IPS means that the polynomial computed by C is in the ideal generated by $\overline{y}, \overline{z}$, which in turn, following the second equality, means that C witnesses the fact that 1 is in the ideal generated by $f_1(\overline{x}), \ldots, f_m(\overline{x}), x_1^2 - x_1, \ldots, x_n^2 - x_n$ (the existence of this witness, for unsatisfiable set of polynomials, stems from the Nullstellensatz [BIKPP96]).

In this work we focus on multilinear $\mathsf{IPS}_{\mathrm{LIN}'}$ refutations. This proof system is complete because its *weaker* subsystem multilinear-formula $\mathsf{IPS}_{\mathrm{LIN}'}$ was shown in [FSTW21, Corollary 4.12] to be complete (and to simulate Nullstellensatz with respect to sparsity by already depth-2 multilinear $\mathsf{IPS}_{\mathrm{LIN}'}$ proofs).

To build an intuition for multilinear $\mathsf{IPS}_{\mathrm{LIN}'}$ it is useful to consider a subsystem of it in which refutations are written as

$$C(\overline{x}, \overline{y}, \overline{z}) = \sum_{i} g_i(\overline{x}) \cdot y_i + C'(\overline{x}, \overline{z}),$$

where $\widehat{C}'(\overline{x},\overline{0}) = 0$ and the g_i 's are multilinear. Note indeed that $C(\overline{x},\overline{0},\overline{0}) = 0$ so that the first condition of IPS proofs holds, and that $C(\overline{x},\overline{y},\overline{0})$ is indeed multilinear in $\overline{x},\overline{y}$.

Important remark: Unlike the multilinear-formula $\mathsf{IPS}_{\mathsf{LIN}'}$ in [FSTW21], in multilinear $\mathsf{IPS}_{\mathsf{LIN}'}$ refutations $C(\overline{x}, \overline{y}, \overline{z})$ we do not require that the refutations are written as multilinear formulas or multilinear circuits, only that the polynomial computed by $C(\overline{x}, \overline{y}, \overline{0})$ is multilinear, hence the latter proof system easily simulates the former.

We now formally state how we prove a functional lower bound for C-IPS systems.

Theorem 10 (Functional Lower Bound Method; Lemma 5.2 in [FSTW21]). Let $C \subseteq \mathbb{F}[\overline{x}]$ be a circuit class, and let $f(\overline{x}) \in C$ be a polynomial, which has no boolean roots. A functional lower bound against C-IPS_{LIN'} for $f(\overline{x})$ and $\overline{x}^2 - \overline{x}$ is a lower bound argument using the following circuit lower bound for $\frac{1}{f(\overline{x})}$: Suppose that $g \notin C$ for all $g \in \mathbb{F}[\overline{x}]$ with

$$g(\overline{x}) = \frac{1}{f(\overline{x})}, \qquad \forall \overline{x} \in \{0, 1\}^n.$$
(1.1)

Then, $f(\overline{x})$ and $\overline{x}^2 - \overline{x}$ do not have C-IPS_{LIN'} refutations. Moreover, if C is a set of multilinear polynomials, then $f(\overline{x})$ and $\overline{x}^2 - \overline{x}$ do not have C-IPS refutations.

2.3 Coefficient Matrix and Dimension

We define notions and measures used in this paper. Consider a polynomial $f \in \mathbb{F}[\overline{x}, \overline{y}]$. We can construct this polynomial by organizing the coefficients of f into a matrix format: the rows are indexed by monomials $\overline{x}^{\overline{a}}$ in the \overline{x} -variables, the columns are indexed by monomials $\overline{y}^{\overline{b}}$ in the \overline{y} -variables, and the entry at position $(\overline{x}^{\overline{a}}, \overline{y}^{\overline{b}})$ is the coefficient of the monomial $\overline{x}^{\overline{a}}\overline{y}^{\overline{b}}$ in f.

Definition 11 (Coefficient Matrix). Let $f \in \mathbb{F}[\overline{x}, \overline{y}]$ be a polynomial, where $\overline{x} = \{x_1, \ldots, x_n\}$ and $\overline{y} = \{y_1, \ldots, y_m\}$. Let $\operatorname{coeff}_{\overline{x}^{\overline{a}}\overline{y}^{\overline{b}}}(f)$ denote the coefficient of the monomial $\overline{x}^{\overline{a}}\overline{y}^{\overline{b}}$ in f. The coefficient matrix of f is the matrix C_f with entries

$$(C_f)_{\overline{a},\overline{b}} := \operatorname{coeff}_{\overline{x}^{\overline{a}}\overline{y}^{\overline{b}}}(f),$$

such that $\sum_{i=1}^{n} a_i + \sum_{j=1}^{m} b_j \leq \deg(f)$.

For our purposes, we care about the dimension of this matrix.

Definition 12 (Coefficient space). Let $\operatorname{coeff}_{\overline{x}|\overline{y}} : \mathbb{F}[\overline{x},\overline{y}] \to 2^{\mathbb{F}[\overline{x}]}$ be the space of $\mathbb{F}[\overline{x}][\overline{y}]$ coefficients, defined by

$$\operatorname{\mathbf{Coeff}}_{\overline{x}|\overline{y}}(f) := \left\{\operatorname{coeff}_{\overline{x}|\overline{y}^{\overline{b}}}(f)\right\}_{\overline{b} \in \mathbb{N}^n},$$

where the coefficients of f are in $\mathbb{F}[\overline{x}][\overline{y}]$. Similarly we have $\operatorname{Coeff}_{\overline{y}|\overline{x}}(f)$ by taking coefficients in $\mathbb{F}[\overline{y}][\overline{x}]$

That is, we use the above in the context of *coefficient dimension*, where we look at the dimension of the coefficient space of f, denoted dim $\mathbf{Coeff}_{\overline{x}|\overline{y}}(f)$. We state a result that connects this to the rank of the matrix.

Lemma 13 (Coefficient matrix rank equals dimension of polynomial space; Nisan [Nis91]). Consider $f \in \mathbb{F}[\overline{x}, \overline{y}]$, and let C_f denote the coefficient matrix of f (Definition 11). Then, the following holds:

rank
$$C_f = \dim \operatorname{Coeff}_{\overline{x}|\overline{y}}(f) = \dim \operatorname{Coeff}_{\overline{y}|\overline{x}}(f).$$

We now show that the coefficient dimension in fact characterizes the width of roABPs.

Lemma 14 (roABP width equals coefficient dimension). Let $f \in \mathbb{F}[\overline{x}]$ be a polynomial. If f is computed by a roABP of width r, then

$$r \geq \max_{i} \dim \operatorname{\mathbf{Coeff}}_{\overline{x}_{\leq i} | \overline{x}_{>i}}(f).$$

Conversely, f can be computed by a roABP of width $\max_i \dim \operatorname{Coeff}_{\overline{x}_{\leq i}|\overline{x}_{>i}}(f)$.

The coefficient dimension of a polynomial $f(\overline{x}, \overline{y})$ measures its complexity by considering the span of all coefficient vectors with respect to the \overline{y} -monomials. In a similar vein, we also consider the *evaluation dimension*, introduced by Saptharishi [Sap12]. Specifically, it captures the dimension of the span of all evaluations of $f(\overline{x}, \cdot)$ at points $\overline{y} \in \mathbb{F}^m$.

Definition 15 (Evaluation dimension). Let $S \subseteq \mathbb{F}$. Let $\mathbf{Eval}_{\overline{x}|\overline{y},S} : \mathbb{F}[\overline{x},\overline{y}] \to 2^{\mathbb{F}[\overline{x}]}$ be the space of $\mathbb{F}[\overline{x},\overline{y}]$ evaluations over S, defined by

$$\mathbf{Eval}_{\overline{x}|\overline{y},S}(f(\overline{x},\overline{y})) := \{f(\overline{x},\beta)\}_{\overline{\beta}\in S^{|\overline{y}|}}$$

The evaluation dimension is therefore the dimension of the above space, denoted dim $\mathbf{Eval}_{\overline{x}|\overline{u},S}(f)$.

That is, we consider the span of functions f over all assignments to the \overline{y} variables. This measure is particularly useful for our applications, as it is directly related to the coefficient dimension.

Lemma 16 (Evaluation dimension bounds coefficient dimension; Forbes-Shpilka [FS13]). Let $f \in \mathbb{F}[\overline{x}, \overline{y}]$, and let $S \subseteq \mathbb{F}$. Then,

$$\mathbf{Eval}_{\overline{x}|\overline{y},S}(f) \subseteq \operatorname{span} \mathbf{Coeff}_{\overline{x}|\overline{y}}(f),$$

and hence,

 $\dim \mathbf{Eval}_{\overline{x}|\overline{y},S}(f) \leq \dim \mathbf{Coeff}_{\overline{x}|\overline{y}}(f).$

Moreover, if |S| > ideg(f), then equality holds:

$$\dim \mathbf{Eval}_{\overline{x}|\overline{y},S}(f) = \dim \mathbf{Coeff}_{\overline{x}|\overline{y}}(f).$$

2.4 Set-Multilinear Monomials over a Word

We recall some notation from [LST21]. Let $w \in \mathbb{Z}^d$ be a word. For a subset $S \subseteq [d]$ denote by w_S the sum $\sum_{i \in S} w_i$, and by $w|_S$ the **subword** of w indexed by the set S. Let³

$$P_w := \{ i \in [d] : w_i \ge 0 \}$$

be the set of **positive indices** of w and let

$$N_w := \{i \in [d] : w_i < 0\}$$

be the set of **negative indices** of w.

Given a word w, we associate with it a sequence $\overline{X}(w) = \langle X(w_1), \ldots, X(w_d) \rangle$ of sets of variables, where for each $i \in [d]$ the size of $X(w_i)$ is $2^{|w_i|}$. We call a monomial set-multilinear over a word wif it is set-multilinear over the sequence $\overline{x}(w)$.

For a word w, let Π_w denote the projection onto the space $\mathbb{F}_{sml}[\overline{x}(w)]$, which maps set-multilinear monomials over w identically to themselves and all other monomials to 0. When the underlying variable partition is clear from context, we simply write Π_{sml} to denote the set-multilinear projection.

2.5 Relative Rank

Let M_w^P and M_w^N denote the set-multilinear monomials over $w|_{P_w}$ and $w|_{N_w}$, respectively. Let $f \in \mathbb{F}_{sml}[\overline{x}(w)]$ and denote by $M_w(f)$ the matrix with rows indexed by M_w^P and columns indexed by M_w^N , whose (m, m')-th entry is the coefficient of the monomial mm' in f.

For any $f \in \mathbb{F}_{sml}[\overline{x}(w)]$ define the **relative rank** with respect to w as follows

rel-rank_w(f) =
$$\frac{\operatorname{rank}(M_w(f))}{\sqrt{|M_w^P| \cdot |M_w^N|}}$$

2.6 Monomial Orders

Finally we recall some basic notions related to monomial orders. For an in-depth introduction see [CLO15]. A monomial order (in a polynomial ring $\mathbb{F}[X]$) is a well-order \leq on the set of all monomials that respects multiplication:

if $m_1 \leq m_2$, then $m_1m_3 \leq m_2m_3$ for any m_3 .

It is not hard to see that any monomial order extends the submonomial relation: if $m_1m_2 = m_3$ for some monomials m_1, m_2 and m_3 , then $m_1 \leq m_3$. This is essentially the only property we need of monomial orderings, and thus our results work for any monomial ordering. Given a polynomial $f \in \mathbb{F}[X]$, the leading monomial of f, denoted LM(f), is the highest monomial with respect to \leq that appears in f with a non-zero coefficient. We conclude this section with the following known fact.

Lemma 17. For any set of polynomials $S \subseteq \mathbb{F}[\overline{x}]$, the dimension of their span in $\mathbb{F}[\overline{x}]$ is equal to the number of unique distinct leading or trailing monomials in their span:

$$\dim \operatorname{span} S = |LM(\operatorname{span} S)| = |TM(\operatorname{span} S)|.$$

where LM and TM stand for leading and trailing monomials respectively. In particular, we have

 $\dim \operatorname{span} S \ge |LM(S)|, |TM(S)|.$

³The P_w here is not to be confused with the canonical full-rank set-multilinear polynomial in [LST21] denoted as well by P_w mentioned in the introduction.

3 Lower Bounds for Constant-depth Multilinear IPS

3.1 Notation for Knapsack

Before defining our hard instance, we introduce some notation. Our construction is based on the instance ks_w from [GHT22], and we adopt parts of their notation.

Let $w \in \mathbb{Z}^d$ be an arbitrary word. Consider the sequence $\overline{X}(w) = \langle X(w_1), \ldots, X(w_d) \rangle$ of sets of variables and the following useful representation of the variables in $\overline{X}(w)$. For any $i \in P_w$, we write the variables of $X(w_i)$ in the form $x_{\sigma}^{(i)}$, where σ is a binary string indexed by the set (formally, a binary string *indexed* by a set A is a function from A to $\{0, 1\}$):

$$A_{w}^{(i)} := \left[\sum_{\substack{i' \in P_{w} \\ i' < i}} w_{i'} + 1, \sum_{\substack{i' \in P_{w} \\ i' \le i}} w_{i'}\right].$$

Hence, the size of $A_w^{(i)}$ is precisely w_i , which implies that there are $2^{|A_w^{(i)}|} = 2^{w_i}$ possible strings indexed by $A_w^{(i)}$, each corresponding to a distinct variable in $X(w_i)$.

Similarly, for any $j \in N_w$, we write the variables of $X(w_i)$ in the form $y_{\sigma}^{(j)}$, where σ is a binary string indexed by the set

$$B_w^{(j)} := \left[\sum_{\substack{j' \in N_w \\ j' < j}} |w_{j'}| + 1, \sum_{\substack{j' \in N_w \\ j' \le j}} |w_{j'}| \right].$$

We call the variables in $x_{\sigma}^{(i)}$ the *positive variables*, or simply \overline{x} -variables, and the variables $y_{\sigma}^{(j)}$ the *negative variables*, or simply \overline{y} -variables. We write A_w^S for the set $\bigcup_{i \in S} A_w^{(i)}$ for any $S \subseteq P_w$, and B_w^T for the set $\bigcup_{i \in T} B_w^{(j)}$ for any $T \subseteq N_w$.

Each monomial that is set-multilinear on $w|_S$ for some $S \subseteq P_w$ corresponds to a binary string indexed by the set A_w^S . Similarly, each monomial that is set-multilinear on $w|_T$ for some $T \subseteq N_w$ corresponds to a binary string indexed by the set B_w^T . For any set-multilinear monomial m on some $w|_S$ with $S \subseteq P_w$, we denote by $\sigma(m)$ the corresponding binary string indexed by A_w^S . Conversely, for any binary string σ indexed by A_w^S , we denote by $m(\sigma)$ the monomial it defines. The same correspondence holds for strings and monomials on the negative variables. Thus, observe that for any (negative or positive) monomial m, we have $m(\sigma(m)) = m$. Moreover, if m is a negative monomial and $S \subseteq P_w$, we write $m(\sigma(m)|_{A_w^S})$ to denote the *positive* monomial determined by the string $\sigma(m)|_{A_w^S}$, which is a substring of $\sigma(m)$ restricted to A_w^S .

Therefore, every set-multilinear monomial on w has degree d, with each \overline{x} -variable picked uniquely from the $X(w_i)$ -variables for $i \in P_w$ (the positive indices in w), and each \overline{y} -variable picked uniquely from the $X(w_j)$ -variables for $j \in N_w$ (the negative indices). Moreover, such a set-multilinear monomial on w corresponds to a binary string of length $\sum_{i=1}^{d} |w_i|$.

We define the **overlap graph** G of the word w as the bipartite graph (P_w, N_w, E) , with an edge between $i \in P_w$ and $j \in N_w$ if $A_w^{(i)}$ and $B_w^{(j)}$ overlap, that is

$$E = \{ (i,j) \mid A_w^{(i)} \cap B_w^{(j)} \neq \emptyset \}.$$

We say that the word w is **balanced** if for every $i \in P_w \cup N_w$, the neighbourhood $N_G(i)$ is non-empty (see Figure 1). In what follows, we suppose that $|w_{N_w}| \ge |w_{P_w}|$, so that the negative monomials are determined by longer binary strings than the positive ones. Otherwise, we flip the roles of



Figure 1: From [GHT22]. Illustration of a word w. Each index w_i of w is shown as a box with w_i slots, so every variable $x_{\sigma}^{(i)}$ in $X(w_i)$ appears as the string σ written inside its corresponding box. The word w shown is balanced.

the positive and negative variables in the definition below. We define the **positive overlap** of w, denoted $\Delta_G(P_w)$, as the maximum degree of a vertex in P_w . Similarly, the **negative overlap**, denoted $\Delta_G(N_w)$, is the maximum degree of a vertex in N_w . A partition of the set of positive indices $P_w = P_w^{(1)} \sqcup \cdots \sqcup P_w^{(r)}$ is called **scattered** if for every part, the neighbourhoods of positive indices in that part are pairwise disjoint, that is

$$N_G(i_1) \cap N_G(i_2) = \emptyset \quad (\forall j \in [r], i_1, i_2 \in P_w^{(j)}, i_1 \neq i_2).$$

3.2 Hard Instance: Knapsack mod p

We now construct our hard instance $\mathsf{ks}_{w,p}$, knapsack mod p. Let $w \in \mathbb{Z}^d$ be a word with $|w_i| \leq b$ for every i. For $i \in P_w$ and $\sigma \in \{0,1\}^{A_w^{(i)}}$ let

$$f_{\sigma}^{(i)} := \prod_{\substack{j \in N_w \\ A_w^{(i)} \cap B_w^{(j)} \neq \emptyset}} f_{\sigma}^{(i,j)},$$

where

$$f_{\sigma}^{(i,j)} := 1 - \prod_{\sigma_j \in \{0,1\}^{B_w^{(j)}}} \left(1 - y_{\sigma_j}^{(j)} \right), \tag{3}$$

where the product in (3) ranges over all those σ_j that agree with σ on $A_w^{(i)} \cap B_w^{(j)}$ (see Figure 2). Let $P_w = P_w^{(1)} \sqcup \cdots \sqcup P_w^{(r)}$ be a scattered partition of the set of positive indices such that r < p. We define our hard instance

$$\mathsf{ks}_{w,p} := \sum_{j \in [r]} \prod_{i \in P_w^{(j)}} \mathsf{ks}_{w,p}^{(i)} - \beta,$$

where

$$\mathsf{ks}_{w,p}^{(i)} := 1 - \mathrm{ml}\left(\sum_{\sigma \in \{0,1\}^{A_w^{(i)}}} x_{\sigma}^{(i)} f_{\sigma}^{(i)}\right)^{p-1},$$

and $\beta \in \mathbb{F}$ is chosen such that $\mathsf{ks}_{w,p}$ is unsatisfiable over Boolean assignments.

Comment (the existence of β): We observe that each $f_{\sigma}^{(i)}$ is a Boolean function. Hence, by Fermat's little theorem, each $ks_{w,p}^{(i)}$ is a Boolean function. It follows that

$$\sum_{j \in [r]} \prod_{i \in P_w^{(j)}} \mathsf{ks}_{w,p}^{(i)} \in \{0, 1, \dots, r\}.$$

Since r < p, we can always choose $\beta \in \mathbb{F}$ such that $ks_{w,p}$ is unsatisfiable over Boolean assignments.



Figure 2: From [GHT22]. Here * represents either 0 or 1. In the construction of the polynomial $k_{s_{w,p}}$, for i = 2 and $\sigma = 011001$, we see that $f_{011001}^{(2)} = y_{011}^{(1)} \cdot y_{00}^{(3)} \cdot (1 - (1 - y_{1000}^{(4)})(1 - y_{1001}^{(4)}) \cdots (1 - y_{1111}^{(4)}))$. While our construction of $f_{\sigma}^{(i)}$ differs from [GHT22], it still functions as an indicator for the variable $x_{\sigma}^{(i)}$.

Comment (computing $\mathsf{ks}_{w,p}$ by a $\mathsf{poly}(d, 2^{bp})$ -size, product-depth 3, multilinear formula of degree $O(pdb2^b)$): Fix $i \in P_w$ and consider computing $\mathsf{ks}_{w,p}^{(i)}$. Let $\sigma_1, \sigma_2 \in \{0, 1\}^{A_w^{(i)}}$ be distinct strings. Suppose there exists $j \in N_w$ such that $A_w^{(i)} \cap B_w^{(j)} \neq \emptyset$ and the polynomials $f_{\sigma_1}^{(i,j)}$ and $f_{\sigma_2}^{(i,j)}$ share \overline{y} -variables. Then, by construction, $f_{\sigma_1}^{(i,j)} = f_{\sigma_2}^{(i,j)}$. It follows that $\mathsf{ml}(f_{\sigma_1}^{(i)}f_{\sigma_2}^{(i)})$ can be computed in the same way as $f_{\sigma_1}^{(i)}f_{\sigma_2}^{(i)}$, but with the shared $f_{\sigma_2}^{(i,j)}$ terms excluded from the construction of $f_{\sigma_2}^{(i)}$. Hence, $\mathsf{ks}_{w,p}^{(i)}$ can be computed by a product-depth 2, multilinear formula of size $\mathsf{poly}(2^{bp})$. Since $P_w = P_w^{(1)} \sqcup \cdots \sqcup P_w^{(r)}$ is a scattered partition of the positive indices, the variables in each $\mathsf{ks}_{w,p}^{(i)}$ are disjoint across distinct *i*. Therefore, $\mathsf{ks}_{w,p}$ can be computed by a product-depth 3, multilinear formula of size $\mathsf{poly}(d, 2^{bp})$. Moreover, each $f_{\sigma}^{(i)}$ has degree at most $O(b2^b)$, so $\mathsf{ks}_{w,p}^{(i)}$ has degree at most $O(pb2^b)$.

3.3 Degree Lower Bound

We now state and prove the degree lower bound that we use in the rank lower bound. We begin with the bound that was used in [GHT22].

Lemma 18 ([FSTW21] Proposition 5.3). Let $n \ge 1$, char(\mathbb{F}) = 0 or char(\mathbb{F}) > n, and $\beta \in \mathbb{F} \setminus \{0, 1, ..., n\}$. If $f \in \mathbb{F}[x_1, ..., x_n]$ is the multilinear polynomial such that

$$f(\overline{x})\left(\sum_{i\in[n]}x_i-\beta\right)=1\pmod{\overline{x}^2-\overline{x}},$$

then $\deg f = n$.

Lemma 19. Let $\overline{x} = \bigsqcup_{i \in I} \overline{x}_i$ be a partition of the variables $\overline{x} = \{x_1, \ldots, x_n\}$, char(\mathbb{F}) = 0 or char(\mathbb{F}) = p > |I|, and $\beta \in \mathbb{F} \setminus \{0, 1, \ldots, |I|\}$. For $i \in I$, let $\psi_i \in \mathbb{F}[\overline{x}_i]$ be a polynomial over the \overline{x}_i -variables that is multilinear, full degree (that is, deg $\psi_i = |\overline{x}_i|$) and a Boolean function. If $f \in \mathbb{F}[\overline{x}]$ is the multilinear polynomial such that

$$f(\overline{x})\left(\sum_{i\in I}\psi_i-\beta\right)=1\pmod{\overline{x}^2-\overline{x}},\tag{4}$$

then $\deg f = n$.

Proof. Let $\{w_i\}_{i\in I}$ be Boolean variables and $f_w \in \mathbb{F}[\overline{w}]$ be the multilinear polynomial such that

$$f_w(w_1, \dots, w_{|I|})\left(\sum_{i \in I} w_i - \beta\right) = 1 \pmod{\overline{w}^2 - \overline{w}}.$$
(5)

By Lemma 18, we have deg $f_w = |I|$. We show that the following polynomial identity over the \overline{x} -variables holds:

$$f_w(\psi_1, \dots, \psi_{|I|})\left(\sum_{i \in I} \psi_i - \beta\right) = 1 \pmod{\overline{x}^2 - \overline{x}}.$$
(6)

We note that f_w is a polynomial over the \overline{w} -variables that is multilinear and for every $i \in I$, ψ_i is a polynomial over the \overline{x}_i -variables that is multilinear. Therefore, as $\overline{x} = \bigsqcup_{i \in I} \overline{x}_i$ is a partition of the \overline{x} -variables, $f_w(\psi_1, \ldots, \psi_{|I|})$ is a polynomial over the \overline{x} -variables that is multilinear. Thus, to show (6), it suffices to show that

$$f_w(\psi_1, \dots, \psi_{|I|})\left(\sum_{i \in I} \psi_i - \beta\right) = 1$$
(7)

holds for all $\overline{x} \in \{0,1\}^n$. Let $\alpha \in \{0,1\}^n$ be a Boolean assignment for the \overline{x} -variables and, for all $i \in I$, let $w_i = \psi_i(\alpha|_{\overline{x}_i})$. Since, for all $i \in I$, ψ_i is a Boolean function, the assignments on the \overline{w} -variables are all Boolean assignments. Therefore, for these Boolean assignments on the \overline{w} -variables, by (5),

$$f_w(w_1,\ldots,w_{|I|})\left(\sum_{i\in I}w_i-\beta\right)=1.$$

Since $w_i = \psi_i(\alpha|_{\overline{x}_i})$ for all $i \in I$, we see that (7) holds for the Boolean assignment α on the \overline{x} -variables. We therefore see that (6) holds. Thus, $f = f_w(\psi_1, \ldots, \psi_{|I|})$. Finally, as f_w has full degree and for all $i \in I$, ψ_i has full degree, we see that $f_w(\psi_1, \ldots, \psi_{|I|})$ has full degree. Therefore, deg $f = \deg f_w(\psi_1, \ldots, \psi_{|I|}) = n$.

Corollary 20. Let $\overline{x} = \bigsqcup_{i \in I} \overline{x}_i$ be a partition of the variables $\overline{x} = \{x_1, \ldots, x_n\}$, char(\mathbb{F}) = 0 or char(\mathbb{F}) = p > |I|, and $\beta \in \mathbb{F} \setminus \{0, 1, \ldots, |I|\}$. If $f \in \mathbb{F}[\overline{x}]$ is the multilinear polynomial such that

$$f(\overline{x})\left(\sum_{i\in I}\prod_{x\in\overline{x}_i}(1-x)-\beta\right) = 1 \pmod{\overline{x}^2 - \overline{x}},\tag{8}$$

then deg(f) = n.

Proof. This follows from Lemma 19, taking for all $i \in I$,

$$\psi_i(\overline{x}_i) = \prod_{x \in \overline{x}_i} (1-x)$$

and noting that ψ_i is a multilinear, full degree polynomial and a Boolean function.

3.4 Rank Lower Bound

Lemma 21. Let \mathbb{F} be a field with characteristic p and $w \in \mathbb{Z}^d$ be a balanced word. If f is the multilinear polynomial such that

$$f = \frac{1}{\mathsf{ks}_{w,p}}$$
 over Boolean assignments,

then $M_w(f)$ has full rank.

Proof. We recall the assumption that $|w_{N_w}| \ge |w_{P_w}|$ from the construction of ks_{w,p}. Now write

$$f = \sum_{m} g_m(\overline{x})m,\tag{9}$$

where the sum ranges over all multilinear monomials m in the \overline{y} -variables and $g_m(\overline{x})$ is some multilinear polynomial in the \overline{x} -variables.

Claim 22. For any monomial m that is set-multilinear on some $w|_T$, where $T \subseteq N_w$, the leading monomial of $g_m(\overline{x})$ is less than or equal to

$$m(\sigma(m)|_{A_m^S}),$$

where S is the maximal subset of P_w such that $A_w^S \subseteq B_w^T$. Moreover, if m is set-multilinear on $w|_{N_w}$, then the leading monomial of $g_m(\overline{x})$ equals

$$m(\sigma(m)|_{A^{Pw}_w}).$$

Proof. We prove this claim by induction on the size of T.

Base case: If $T = \emptyset$, consider the partial assignment τ_1 that maps all the \overline{y} -variables to 0. We have $\tau_1(f) = g_1(\overline{x})$, where $g_1(\overline{x})$ is the coefficient of the empty monomial 1. On the other hand, $\tau_1(\mathsf{ks}_{w,p}^{(i)}) = 1$ for all *i*. Since

$$f = \frac{1}{\mathsf{ks}_{w,p}}$$
 over Boolean assignments,

we see that $\tau_1(f) = 1/(r-\beta)$ over Boolean assignments. As $g_1(\overline{x})$ is multilinear, $g_1(\overline{x}) = 1/(r-\beta)$ as a polynomial identity, so the the leading monomial of $g_1(\overline{x})$ is the empty monomial 1.

Inductive step: Suppose that T is non-empty, and let m be a set-multilinear monomial over $w|_T$. Consider the partial assignment τ_m that maps any \overline{y} -variable in m to 1 and any other \overline{y} -variable to 0. By (9)

$$\tau_m(f) = \sum_{m'} g_{m'}(\overline{x}),\tag{10}$$

where m' ranges over all submonomials of m. On the other hand,

$$\tau_m(\mathsf{ks}_{w,p}) = \sum_{j \in [r]} \prod_{i \in P_w^{(j)}} \tau_m(\mathsf{ks}_{w,p}^{(i)}) - \beta.$$

For $i \in P_w$, if $A_w^{(i)} \not\subseteq B_w^T$, then $\tau_m(\mathsf{ks}_{w,p}^{(i)}) = 1$; however, if $A_w^{(i)} \subseteq B_w^T$, then $\tau_m(\mathsf{ks}_{w,p}^{(i)}) = 1 - x_{\sigma_i}^{(i)}$, where σ_i is the binary string indexed by $A_w^{(i)}$ that agrees with $\sigma(m)$ on $A_w^{(i)}$. Therefore

$$\tau_m(f)\left(\sum_{j\in[r]}\prod_{i\in P_w^{(j)}}(1-x_{\sigma_i}^{(i)})-\beta\right)=1 \text{ over Boolean assignments},$$

where the product ranges over $i \in P_w^{(j)}$ such that $A_w^{(i)} \subseteq B_w^T$. From Corollary 20, it follows that the leading monomial of $\tau_m(f)$ is the product of all the $x_{\sigma_i}^{(i)}$ appearing above, and thus the leading monomial is

$$m(\sigma(m)|_{A_w^S}),\tag{11}$$

	$A_w^{(2)}$				$A_w^{(5)}$			$A_w^{(6)}$						
					0	0	1	0	1	1	0	1		
1	0	0		1	0	0	1	0	1	1	0	1	1	
-	$B_w^{(1)}$		$B_w^{(3)}$		$B_w^{(4)}$				$B_{w}^{(7)}$				$B_{w}^{(8)}$	

Figure 3: From [GHT22]. In this example, $T = \{1, 4, 7, 8\} \subseteq N_w$ and $m = y_{100}^{(1)} \cdot y_{101}^{(4)} \cdot y_{0110}^{(7)} \cdot y_{11}^{(8)}$ is a set-multilinear monomial over $w|_T$. Like [GHT22], since $S = \{5, 6\}$ is the maximal subset of P_w with $A_w^S \subseteq B_w^T$, we have that the leading monomial of $g_m(\bar{x})$ is less than or equal to $x_{00}^{(5)} \cdot x_{101101}^{(6)}$. However, in contrast to [GHT22], in our polynomial $ks_{w,p}$, the partial assignment setting the \bar{y} -variables in m to 1 and the remaining \bar{y} -variables to 0 results in the polynomial $(1 - x_{00}^{(5)}) + (1 - x_{101101}^{(6)}) - \beta$.

where S is the maximal subset of P_w such that $A_w^S \subseteq B_w^T$ (see Figure 3). If the leading monomial of $g_m(\overline{x})$ were greater than $m(\sigma(m)|_{A_w^S})$, then it must be cancelled by some monomial of $g_{m'}(\overline{x})$ in (10) for some proper submonomial of m; however, by the inductive hypothesis, for all such proper submonomials m', the leading monomial of $g_{m'}(\overline{x})$ is less than or equal to $m(\sigma(m')|_{A_w^S})$. Therefore, the leading monomial of $g_m(\overline{x})$ must be less than or equal to (11), concluding the induction.

It remains to show that the leading monomial of $g_m(\overline{x})$ equals $m(\sigma(m)|_{A_w^{P_w}})$ whenever m is set-multilinear on $w|_{N_w}$. Let m' be a proper submonomial of m that is set-multilinear over $w|_T$ for some $T \subsetneq N_w$. As w is a balanced word, there is some $i \in P_w$ such that $A_w^{(i)} \not\subseteq B_w^T$, and thus the leading monomial of $g_{m'}(\overline{x})$ is strictly smaller than $m(\sigma(m)|_{A_w^{P_w}})$. From (11), it follows that the leading monomial of $g_m(\overline{x})$ must equal $m(\sigma(m)|_{A_w^{P_w}})$.

For each monomial m_P that is set-multilinear over $w|_{P_w}$, there exists a monomial m_N , setmultilinear over $w|_{N_w}$, such that the leading monomial of $g_{m_N}(\overline{x})$ is exactly m_P . Consequently, the (m_P, m_N) entry of $M_w(f)$ is non-zero in \mathbb{F} , while for every monomial $m'_P \neq m_P$, also set-multilinear over $w|_{P_w}$ and satisfying $m_P \leq m'_P$, the (m'_P, m_N) entry is zero. For $M_w(f)$, it follows that the dimension of the column space equals the number of rows, so $M_w(f)$ has full rank.

Corollary 23. Let \mathbb{F} be a field with characteristic p, and let $w \in \mathbb{Z}^d$ be a balanced word with $|w_i| \leq b$ for all $i \in [d]$. If f is the multilinear polynomial such that

$$f = \frac{1}{\mathsf{ks}_{w,p}}$$
 over Boolean assignments,

then rel-rank $(f) \ge 2^{-b/2}$.

Proof. Recall that, by the construction of $\mathsf{ks}_{w,p}$, we assume $|w_{N_w}| \ge |w_{P_w}|$. Since w is balanced and satisfies $|w_i| \le b$ for all $i \in [d]$, it follows that $|w_{P_w}| - |w_{N_w}| \ge -b$. By Lemma 21, $M_w(f)$ has rank $|M_w^P|$. Therefore

rel-rank_w(f) =
$$\sqrt{\frac{|M_w^P|}{|M_w^N|}} = \sqrt{2^{|w_{P_w}| - |w_{N_w}|}} \ge 2^{-b/2}.$$

3.5 IPS Lower Bound

We now state and prove our lower bound for constant-depth IPS over finite fields. We begin by recalling notation from [BDS24]. Let F(n) denote the *n*-th Fibonacci number, defined by F(0) = 1, F(1) = 2 and F(i) = F(i-1) + F(i-2) for $i \ge 2$; let G(i) = F(i) - 1 for all *i*. Fix the product-depth $\Delta \le \log \log \log n/4$, and let $d = \lfloor \log n/4 \rfloor$ and $\lambda = \lfloor d^{1/G(\Delta)} \rfloor$.

Theorem 24 ([GHT22] over Finite Fields). Let $p \ge 5$ be a prime, and let \mathbb{F} be a field of characteristic p. Let $n, \Delta \in \mathbb{N}_+$ with $\Delta \le \log \log \log n/4$. Then any product-depth at most Δ multilinear $\mathsf{IPS}_{\text{LIN'}}$ refutation over \mathbb{F} of $\mathsf{ks}_{w,p}$ has size at least

$$n^{\Omega(\lambda/\Delta)}$$

The proof of Theorem 24 relies on the following result:

Theorem 25. Let $p \ge 5$ be a prime, and let \mathbb{F} be a field of characteristic p. Let Δ be as above. If f is the multilinear polynomial that equals

$$\frac{1}{\mathsf{ks}_{w,p}}$$
 over Boolean assignments,

then any circuit of product-depth at most Δ computing f has size at least

$$n^{\Omega(\lambda/\Delta)}$$
.

Proof of Theorem 24 from Theorem 25. Let $C(\overline{x}, \overline{y}, \overline{z})$ be a multilinear $\mathsf{IPS}_{\mathrm{LIN}'}$ refutation of $\mathsf{ks}_{w,p}^4$. As there is only one non-Boolean axiom, C has a single \overline{y} -variable, which we denote by y. Since $\widehat{C}(\overline{x}, y, \overline{0})$ is linear in the y-variable and satisfies $\widehat{C}(\overline{x}, 0, \overline{0}) = 0$, it follows that

$$\widehat{C}(\overline{x}, y, \overline{0}) = g(\overline{x}) \cdot y$$

for some polynomial $g(\overline{x}) \in \mathbb{F}[\overline{x}]$. This polynomial $g(\overline{x})$ is computed by the circuit $C(\overline{x}, 1, \overline{0})$, so the minimal product-depth- Δ circuit size of $g(\overline{x})$ lower bounds that of $C(\overline{x}, \overline{y}, \overline{z})$. Therefore, it suffices to lower bound the size of product-depth at most Δ circuits computing $g(\overline{x})$.

We have

$$\widehat{C}(\overline{x}, y, \overline{z}) = \widehat{C}(\overline{x}, y, \overline{0}) + \sum_{i} h_{i} \cdot z_{i}$$

for some polynomials h_i in $\overline{x}, y, \overline{z}$, hence $\widehat{C}(\overline{x}, y, \overline{z}) = g(\overline{x}) \cdot y + \sum_i h_i \cdot z_i$. Since

$$\widehat{C}\left(\overline{x},\mathsf{ks}_{w,p},\overline{x}^2-\overline{x}\right)=1,$$

we see that

$$g(\overline{x}) \cdot (\mathsf{ks}_{w,p}) + \sum_{i} (h_i \cdot (x_i^2 - x_i)) = 1.$$

Therefore, over Boolean assignments, $g(\overline{x}) \cdot \mathsf{ks}_{w,p} \equiv 1$. The result now follows from Theorem 25. \Box

 $^{{}^{4}\}mathsf{ks}_{w,p}$ involves both \overline{x} - and \overline{y} -variables. As this distinction will not play a role in the proof, we treat all variables in $\mathsf{ks}_{w,p}$ as \overline{x} -variables. We therefore use the standard notation $C(\overline{x}, \overline{y}, \overline{z})$ for an IPS refutation, where \overline{x} are variables of the axioms, and $\overline{y}, \overline{z}$ serve as placeholder variables for the axioms.

Lemma 26. Let $p \geq 5$ be a prime, and let \mathbb{F} be a field of characteristic p. Let Δ , d and λ be as above. There exist $\alpha \in \mathbb{Q}$ with $1/2 \leq \alpha < 1$, and $k \in \mathbb{N}_+$ with $k \in \lfloor \lfloor \log n \rfloor / 2, \lfloor \log n \rfloor \rfloor$ and $\alpha k \in \mathbb{Z}$, such that if $w \in \mathbb{Z}^d$ is a balanced word over the alphabet $\{\alpha k, -k\}$, and f is the multilinear polynomial which equals $1/k_{s_{w,p}}$ over Boolean assignments, then any set-multilinear circuit of product-depth Δ computing the set-multilinear projection $\Pi_w(f)$ has size at least

$$s \ge 2^{\frac{k(\lambda/256-1)}{2\Delta}}.$$

Proof. Let C be a set-multilinear circuit of size s and product-depth Δ computing $\Pi_w(f)$. By unwinding C into a formula, we obtain a set-multilinear formula F of size $s^{2\Delta}$ and product-depth Δ that also computes $\Pi_w(f)$. We now make use of the following claim from [BDS24]:

Claim 27 ([LST21],[BDS24] Lemma 4.3). Let $\delta \leq \Delta$ be an integer. There exist $\alpha \in \mathbb{Q}$ with $1/2 \leq \alpha < 1$, and $k \in \mathbb{N}_+$ with $k \in [\lfloor \log n \rfloor/2, \lfloor \log n \rfloor]$ and $\alpha k \in \mathbb{Z}$, such that if $w \in \mathbb{Z}^d$ is a word over the alphabet $\{\alpha k, -k\}$, and F is a set-multilinear formula of product-depth δ , degree at least $\lambda^{G(\delta)}/8$ and size at most s, then

$$\operatorname{rel-rank}_{w}(F) \leq s2^{-k\lambda/256}$$

As w is balanced, by Lemma 21, $M_w(f)$ has full rank and deg $F \ge d \ge \lambda^{G(\delta)}/8$. Thus, applying Corollary 23 and Claim 27, we obtain

$$2^{-k} \leq \operatorname{rel-rank}_w(\Pi_w(f)) \leq s^{2\Delta} 2^{-k\lambda/256}.$$

We therefore see that

$$s^{2\Delta} > 2^{k(\lambda/256-1)}.$$

from which the claim of the lemma follows.

Lemma 28 ([For24, Corollary 27]). Let \mathbb{F} be any field, and let the variables \overline{x} be partitioned into $\overline{x} = \overline{x}_1 \sqcup \cdots \sqcup \overline{x}_d$. Suppose $f \in \mathbb{F}[\overline{x}]$ can be computed by a size s, product-depth Δ algebraic circuit. Then the set-multilinear projection $\Pi_{sml}(f) \in \mathbb{F}[\overline{x}]$ can be computed by a size poly $(s, \Theta(\frac{d}{\ln d})^d)$, product-depth 2Δ set-multilinear circuit.

Proof of Theorem 24. Let C be a circuit of size $s \ge n$ and product-depth at most Δ computing f. Let $d = \lfloor \log n/4 \rfloor$ and $\lambda = \lfloor d^{1/G(\Delta)} \rfloor$ be as defined above, and let $1/2 \le \alpha < 1$ and $k \in \lfloor \log n \rfloor/2, \lfloor \log n \rfloor \rfloor$ be as constructed in Lemma 26. Construct, by induction, a balanced word $w \in \mathbb{Z}^d$ over the alphabet $\{\alpha k, -k\}$.

By Lemma 28, there exists a set-multilinear circuit C' of size $\mathsf{poly}(s, \Theta(\frac{d}{\ln d})^d)$ and product-depth 2Δ computing the set-multilinear projection $\Pi_w(f)$ of f.

Moreover, by Lemma 26, any set-multilinear circuit of product-depth 2Δ computing $\Pi_w(f)$ must have size at least

$$2^{\frac{k(\lambda/256-1)}{2\Delta}} \ge n^{\frac{\lambda/256-1}{8\Delta}},$$

where the inequality follows from the lower bound on k. Combining the two bounds above, we obtain

$$\mathsf{poly}(s,\Theta(\frac{d}{\ln d})^d) \geq n^{\frac{\lambda/256-1}{8\Delta}}$$

and therefore,

$$d^{O(d)}\mathsf{poly}(s) \ge n^{\Omega(\lambda/\Delta)}.$$

Since $\Delta \leq \log \log \log n/4 \leq \log \log d/2$, it follows that $\lambda \geq (\log d)^2$. Hence,

$$n^{\Omega(\lambda/\Delta)} \ge d^{\omega(d)}$$

from which the claim of the theorem follows.

1	2	3	1	. 2	3	1		
$A_w^{(2)}$	$A_w^{(4)}$	$A_w^{(5)}$	A	$\stackrel{(6)}{_{w}} A^{(10)}_{w}$	$A_w^{(11)}$	$A_w^{(12)}$		
					·			
$B_w^{(1)}$	$B_{w}^{(3)}$	$B_{w}^{(7)}$	$B_w^{(8)}$	$B_{w}^{(9)}$	E	$B_w^{(13)} = B_w^{(1)}$	$B_w^{(14)}$	
	•					•		

Figure 4: Illustration of the scattered partition induced by $\pi : P_w \to [\Delta_G(N_w)]$ with $\Delta_G(N_w) = 3$. The values of π appear above the positive boxes, while the neighbourhoods of the positive indices are shown below the negative boxes. Since this is a scattered partition, vertices in the same part have pairwise disjoint neighbourhoods.

Comment (constructing a scattered partition): Our instance $k_{w,p}$ and rank lower bound require a scattered partition $P_w = P_w^{(1)} \sqcup \cdots \sqcup P_w^{(r)}$ of the positive indices having fewer than p parts. Let $\xi := \Delta_G(N_w)$ be the negative overlap of the overlap graph. We construct a scattered partition with $r = \xi$ as follows. For each positive index $i \in P_w$, define $\pi'(i) := |\{i' \in P_w \mid i' \leq i\}|$, and let $\pi(i)$ be the least residue of $\pi'(i) \pmod{\xi}$; that is, $\pi(i) \in [\xi]$ with $\pi(i) \equiv \pi'(i) \pmod{\xi}$. The map π partitions P_w into ξ parts such that in each part, the neighbourhoods of the vertices are pairwise disjoint, hence π induces a scattered partition (see Figure 4).

Because our IPS lower bound assumes $p \geq 5$, it suffices to construct a scattered partition with r < 5. Since the word $w \in \mathbb{Z}^d$ is over the alphabet $\{\alpha k, -k\}$ with $1/2 \leq \alpha < 1$, it follows that $\Delta_G(N_w) \leq 3$. Therefore, π yields a scattered partition with $r \leq 3$.

4 Upper Bounds for Constant-depth Multilinear IPS

4.1 Elementary Symmetric Sums

Proposition 29. Over any field \mathbb{F} , for $|\overline{x}| = n \ge l \ge d \ge 0$,

$$e_l(\overline{x}) \cdot e_d(\overline{x}) = \sum_{i=k}^d \binom{l+d-i}{l} \binom{l}{i} e_{l+d-i}(\overline{x}) \pmod{\overline{x}^2 - \overline{x}},$$

where $k \ge 0$ is the smallest integer such that $l + d - k \le n$.

Proof. Since $ml(e_l(\overline{x}) \cdot e_d(\overline{x}))$ is symmetric in \overline{x} , we have

$$e_l(\overline{x}) \cdot e_d(\overline{x}) = \sum_{i=k}^d \gamma_i \cdot e_{l+d-i}(\overline{x}) \pmod{\overline{x}^2 - \overline{x}},$$

for $\gamma_i \in \mathbb{F}$. Let $S \subseteq [n]$ with |S| = l + d - i. The coefficient of the monomial x^S in $\operatorname{ml}(e_l(\overline{x}) \cdot e_d(\overline{x}))$ is $\binom{l+d-i}{l}\binom{l}{i}$. This is because each of the $\binom{l+d-i}{l}$ many sub-monomials x^A of x^S in $e_l(\overline{x})$ combines with $\binom{l}{i}$ many monomials x^B in $e_d(\overline{x})$, where $|A \cap B| = i$, to produce x^S in $\operatorname{ml}(e_l(\overline{x}) \cdot e_d(\overline{x}))$. Therefore, $\gamma_i = \binom{l+d-i}{l}\binom{l}{i}$.

We observe that elementary symmetric polynomials are unsatisfiable even over constant characteristic fields, when their degree meets a simple condition on their p-base expansion. **Lemma 30** ([Luc78] Lucas's Theorem). Let p be a prime and $m, n \in \mathbb{N}_+$. If $m = m_k p^k + \cdots + m_1 p + m_0$ and $n = n_k p^k + \cdots + n_1 p + n_0$ are the base p expansions of m and n respectively (where $0 \le m_i, n_i \le p - 1$ for $0 \le i \le k$), then

$$\binom{m}{n} \equiv \prod_{i=0}^{k} \binom{m_i}{n_i} \pmod{p}.$$

Lemma 31. Let \mathbb{F} be a field with characteristic p. If $d = d_k p^k + \cdots + d_1 p + d_0$ is the base p expansion of d, then

$$|\{e_d(\overline{x}) \mid x \in \{0,1\}^n\}| \le \prod_{\substack{i \in \{0,\dots,k\}\\d_i \neq 0}} (p-d_i) + 1.$$

In particular, if d has only one non-zero digit d_i in its base p expansion and $d_i \ge 2$, then there exists $\beta \in \mathbb{F}$ such that $e_d(\overline{x}) - \beta = 0$ is unsatisfiable over Boolean assignments.

Proof. If m is the Hamming weight of a Boolean assignment $\alpha \in \{0,1\}^n$, then $e_d(\alpha) = \binom{m}{d}$. By Lemma 30, we have

$$\binom{m}{d} \equiv \prod_{i \in \{0,\dots,k\}} \binom{m_i}{d_i} \equiv \prod_{\substack{i \in \{0,\dots,k\}\\d_i \neq 0}} \binom{m_i}{d_i} \pmod{p}.$$

We see that $e_d(\alpha) = 0$ in \mathbb{F} if and only if $m_i < d_i$ for some *i* with $d_i \neq 0$. Conversely, $e_d(\alpha)$ is non-zero in \mathbb{F} if and only if $d_i \leq m_i \leq p-1$ for all *i* with $d_i \neq 0$. Therefore, $e_d(\alpha)$ can attain at most

$$\prod_{\substack{i \in \{0,\dots,k\}\\d_i \neq 0}} ((p-1) - d_i + 1)$$

distinct non-zero values in \mathbb{F} . This completes the proof of the main claim of the lemma.

Now if $d_i \geq 2$ is the only non-zero digit in the base p expansion of d, then over Boolean assignments, $e_d(\overline{x})$ can attain at most $p - d_i + 1 \leq p - 1$ distinct values in \mathbb{F} . The existence of β follows.

Occasionally, instead of viewing \overline{x} as Boolean variables, we consider them more generally as Boolean functions. By a similar argument, it is straightforward to verify that Lemma 31 continues to hold in this more general setting.

4.2 Separation

Here, we separate the constant-depth IPS subsystem over finite fields, as considered in this work, from the constant-depth IPS subsystem over large fields studied in [GHT22].

Let $p \geq 3$ be a prime and let \mathbb{F} be a field of characteristic p. We construct the symmetric knapsack of degree 2, denoted as ks_{w,e_2} . Over Boolean assignments, ks_{w,e_2} is unsatisfiable in \mathbb{F} and in every field of characteristic 0. Moreover, constant-depth multilinear $\mathsf{IPS}_{\mathrm{LIN}'}$ over \mathbb{F} admits a polynomial-size refutation of ks_{w,e_2} , whereas constant-depth multilinear $\mathsf{IPS}_{\mathrm{LIN}'}$ over any characteristic 0 field does not.

Using the notation from Section 3, let $w \in \mathbb{Z}^d$ be a word with $|w_i| \leq b$ for every *i*. Our separating instance is defined as:

$$\mathsf{ks}_{w,e_2} := \mathrm{ml}\left(e_2\left(\left\{x_{\sigma}^{(i)}f_{\sigma}^{(i)}\right\}_{i\in P_w, \sigma\in\{0,1\}^{A_w^{(i)}}}\right)\right) - \beta.$$

where $\beta \in \mathbb{F}$ is chosen such that $\mathsf{ks}_{w,e_2} = 0$ admits no satisfying Boolean assignment in \mathbb{F} .

Comment (the existence of β): The existence of β follows from Lemma 31, specifically from the remark following the lemma concerning its application to Boolean functions rather than Boolean variables.

Comment (computing ks_{w,e_2} by a poly $(d, 2^b)$ -size, product-depth 2, multilinear formula): Since

$$\mathsf{ks}_{w,e_2} = \sum_{\substack{(i_1,\sigma_1), (i_2,\sigma_2) \in S\\(i_1,\sigma_1) \neq (i_2,\sigma_2)}} x_{\sigma_1}^{(i_1)} x_{\sigma_2}^{(i_2)} \mathrm{ml}(f_{\sigma_1}^{(i_1)} f_{\sigma_2}^{(i_2)}) - \beta$$

where $S = \{(i, \sigma) \mid i \in P_w, \sigma \in \{0, 1\}^{A_w^{(i)}}\}$, it suffices to verify that $\operatorname{ml}(f_{\sigma_1}^{(i_1)} f_{\sigma_2}^{(i_2)})$ can be computed by a suitable polynomial-size constant-depth multilinear formula. Since the positive overlap of w satisfies $\Delta_G(P_w) \leq b$, we see that each $f_{\sigma}^{(i)}$ can be written as $\sum \prod (1-y)$ where the fan-in of the sum gate is $O(2^b)$ and the product ranges over distinct \overline{y} -variables. Moreover, because $\operatorname{ml}((1-y)^2) = 1-y$, each $\operatorname{ml}(f_{\sigma_1}^{(i_1)} f_{\sigma_2}^{(i_2)})$ can likewise be written in this form. Altogether, ks_{w,e_2} can thus be written as a product-depth 2, multilinear formula of $\operatorname{poly}(d, 2^b)$ -size. Moreover, each $f_{\sigma}^{(i)}$ has degree at most $O(b2^b)$, so the overall degree of ks_{w,e_2} is therefore $O(b2^b)$.

Lemma 32. Over \mathbb{F} , there exists a product-depth 3 multilinear $\mathsf{IPS}_{\mathrm{LIN}'}$ refutation of ks_{w,e_2} of size $\mathsf{poly}(d^p, 2^{bp})$.

Proof. By Fermat's little theorem, we see that

$$\mathrm{ml}((\mathsf{ks}_{w,e_2})^{p-2}) \cdot \mathsf{ks}_{w,e_2} + \sum_{\psi \in \overline{x} \cup \overline{y}} h_{\psi}(\psi^2 - \psi) = 1$$
(12)

for some polynomials $h_{\psi} \in \mathbb{F}[\overline{x}, \overline{y}]$. From computing ks_{w,e_2} by a $\mathsf{poly}(d, 2^b)$ -size product-depth 2, multilinear formula, we see that $\mathsf{ml}((\mathsf{ks}_{w,e_2})^{p-2})$ can be computed by a product-depth 2, multilinear formula of size $\mathsf{poly}(d^p, 2^{bp})$. Moreover, we see that each h_{ψ} can be computed by a product-depth 2 formula of size $\mathsf{poly}(d^p, 2^{bp})$. Therefore, over \mathbb{F} , (12) is a product-depth 3 multilinear $\mathsf{IPS}_{\mathsf{LIN}'}$ refutation of ks_{w,e_2} .

Let E be a field of characteristic 0. Since ks_{w,e_2} admits no satisfying Boolean assignment in \mathbb{F} , it likewise admits none in E. Over E, we will prove a lower bound against constant-depth multilinear $\mathsf{IPS}_{\mathrm{LIN}'}$ for ks_{w,e_2} . We first prove a degree lower bound.

Lemma 33. Let $\operatorname{char}(\mathbb{F}) = 0$, n > 1 and $\beta \in \mathbb{Z}^+$ such that $e_2(x_1, \ldots, x_n) - \beta = 0$ is unsatisfiable over \mathbb{F} for $\overline{x} \in \{0, 1\}^n$. If $f \in \mathbb{F}[x_1, \ldots, x_n]$ is the multilinear polynomial such that

$$f(\overline{x}) (e_2(\overline{x}) - \beta) = 1 \pmod{\overline{x}^2 - \overline{x}},$$

then $\deg f = n$.

We note that a degree lower bound of deg $f \ge n-1$ follows from [HLT24] Corollary 1.2; however, we need the tight bound of deg $f \ge n$.

Proof of Lemma 33. As $f(\overline{x})$ is multilinear, we have

$$f(\overline{x}) = \sum_{T \subseteq [n]} f(\mathbb{1}_T) \prod_{i \in T} x_i \prod_{i \notin T} (1 - x_i),$$

where $\mathbb{1}_T \in \{0,1\}^n$ is the indicator vector of the set T. Therefore

$$f(\overline{x}) = \sum_{T \subseteq [n]} \frac{1}{\binom{|T|}{2} - \beta} \prod_{i \in T} x_i \prod_{i \notin T} (1 - x_i).$$

The coefficient of $\prod_{i \in [n]} x_i$ in $f(\overline{x})$ is thus

$$\sum_{T \subseteq [n]} \frac{1}{\binom{|T|}{2} - \beta} (-1)^{n-|T|} = \sum_{j=0}^{n} \binom{n}{j} \frac{1}{\binom{j}{2} - \beta} (-1)^{n-j}.$$
(13)

We show that (13) is nonzero. As (13) lies in the subfield of \mathbb{F} that is isomorphic to \mathbb{Q} , it suffices to show that it is nonzero over \mathbb{Q} . It therefore suffices to show that (13) is nonzero over \mathbb{R} . We have, over \mathbb{R} ,

$$\frac{1}{\binom{j}{2} - \beta} = \frac{2}{j^2 - j - 2\beta} = \frac{2}{(j - \gamma_1)(j - \gamma_2)} = 2\sqrt{1 + 8\beta} \left(\frac{1}{j - \gamma_2} - \frac{1}{j - \gamma_1}\right),$$

where $\gamma_1 = (1 - \sqrt{1 + 8\beta})/2$ and $\gamma_2 = (1 + \sqrt{1 + 8\beta})/2$. Hence

$$\sum_{j=0}^{n} \binom{n}{j} \frac{1}{\binom{j}{2} - \beta} (-1)^{n-j} = 2\sqrt{1+8\beta} \sum_{j=0}^{n} \binom{n}{j} \left(\frac{1}{j-\gamma_2} - \frac{1}{j-\gamma_1}\right) (-1)^{n-j}$$
$$= 2\sqrt{1+8\beta} \left(-\frac{n!}{\prod_{j=0}^{n}(\gamma_2-j)} + \frac{n!}{\prod_{j=0}^{n}(\gamma_1-j)}\right),$$

where the last equality follows from

Claim 34 ([FSTW21] Subclaim B.2).

$$\sum_{j=0}^{k} \binom{k}{j} \frac{1}{j-\beta} (-1)^{k-j} = -\frac{k!}{\prod_{j=0}^{k} (\beta-j)}.$$

Finally, to show that (13) is nonzero, it suffices to show that

$$\prod_{j=0}^{n} (\gamma_1 - j) \neq \prod_{j=0}^{n} (\gamma_2 - j).$$
(14)

We note that $\gamma_1(\gamma_1 - 1) = \gamma_2(\gamma_2 - 1) = 2\beta$; however, for k > 1, we have $|\gamma_1 - k| > |\gamma_2 - k|$. We therefore have

$$\left|\prod_{j=0}^{n} (\gamma_1 - j)\right| > \left|\prod_{j=0}^{n} (\gamma_2 - j)\right|,$$

hence (14) holds and (13) is nonzero.

Lemma 35. Let E be a field of characteristic 0, and $n, \Delta \in \mathbb{N}_+$ with $\Delta \leq \log \log \log n/4$. Then any product-depth at most Δ multilinear $\mathsf{IPS}_{\mathrm{LIN}'}$ refutation of ks_{w,e_2} is of size at least

$$n^{\Omega(\lambda/\Delta)},$$

where $d = \lfloor \log n/4 \rfloor$ and $\lambda = \lfloor d^{1/G(\Delta)} \rfloor$.

Proof. The proof of this lemma is essentially the same as the proof of Theorem 24 (and [GHT22]) and is omitted here. We recall the general strategy of reducing an IPS lower bound to a rank lower bound, to a degree lower bound, which for this instance is Lemma 33).

Theorem 36 (Separation: [GHT22] over Finite Fields vs. [GHT22]). Let $p \ge 3$ be a prime, and let \mathbb{F} be a field of characteristic p. Let $n, \Delta \in \mathbb{N}_+$ with $\Delta \le \log \log \log n/4$. Then there exists a product-depth 2, multilinear formula $k_{\mathbf{s}_{w,e_2}}$ of size poly(n) such that:

- ks_{w,e_2} has no satisfying Boolean assignment over \mathbb{F} , and over any field of characteristic 0;
- there is a poly(n)-size, product-depth-3 multilinear $\mathsf{IPS}_{\mathrm{LIN}'}$ refutation of ks_{w,e_2} over \mathbb{F} ;
- for every field of characteristic 0, any product-depth at most Δ multilinear $\mathsf{IPS}_{\mathrm{LIN}'}$ refutation of ks_{w,e_2} requires size at least

 $n^{\Omega(\lambda/\Delta)}$.

where $d = \lfloor \log n/4 \rfloor$ and $\lambda = \lfloor d^{1/G(\Delta)} \rfloor$.

5 Lower Bounds for roABP-IPS

In this section, we work over field \mathbb{F}_q , where q is a constant greater than 2. In addition to proving a lower bound over finite fields, this work significantly simplifies [HLT24], though we note that we are still working in the placeholder model. Consider the following hard instance

$$f(\overline{x}) \coloneqq \prod_{i=1}^{n} (1 - x_i) - 2.$$

$$(15)$$

Clearly this function never evaluates to 0 over boolean assignments in \mathbb{F}_q (when q > 2). In contrast, the hard instance from [HLT24] is a subset-sum instance, therefore requiring large characteristic to be defined.

5.1 roABP-IPS Lower Bounds in Fixed Order

We begin by proving a lower bound where the roABPs are given a fixed order of the variables. By Corollary 20, we have the following lemma.

Lemma 37. Let \mathbb{F}_q be a finite field with q > 2, and let $f(\overline{x})$ be as in (15). If $g(\overline{x})$ is the multilinear polynomial such that

$$g(\overline{x}) \cdot f(\overline{x}) = 1 \mod \overline{x}^2 - \overline{x},$$

then $\deg(g) = n$.

For any $\overline{x}, \overline{y}$ variables, with $|\overline{x}| = |\overline{y}| = n$, we use $\overline{x} \circ \overline{y}$ to denote the entry-wise product (x_1y_1, \ldots, x_ny_n) . In other words, the *gadget* we use is the mapping

$$x_i \mapsto x_i y_i,$$

which substitutes the variables x_i by $x_i y_i$, for every *i*. We use $\mathbb{1}_S \in \{0,1\}^n$ to denote the indicator vector for a set *S*.

Theorem 38. Let $f(\overline{x})$ be as in (15). Let $g(\overline{x}, \overline{y}) \cdot f(\overline{x} \circ \overline{y}) = 1 \mod \overline{x}^2 - \overline{x}$. Then,

$$\left| \operatorname{LM}\left(\left\{ \operatorname{ml}(g(\overline{x}, \mathbb{1}_S)) : S \subseteq [n] \right\} \right) \right| = 2^n.$$
(16)

Proof. We first need the claim below.

Claim 39. Each $S \subseteq [n]$ induces a distinct leading monomial in $ml(g(\overline{x}, \mathbb{1}_S))$.

Proof of claim: Let $S \subseteq [n]$. By the assumption $g(\overline{x}, \overline{y}) \cdot f(\overline{x} \circ \overline{y}) = 1 \mod \overline{x}^2 - \overline{x}$, we also have

$$\mathrm{ml}(g(\overline{x}, \mathbb{1}_S)) \cdot f(\overline{x} \circ \mathbb{1}_S) = 1 \mod \overline{x}^2 - \overline{x},\tag{17}$$

since multilinearizing $g(\overline{x}, \mathbb{1}_S)$ does not affect the equality (as we work modulo $\overline{x}^2 - \overline{x}$). By the lifting defined above, $\operatorname{ml}(g(\overline{x}, \mathbb{1}_S))$ is a (multilinear symmetric) polynomial that depends on the variables x_i , for $i \in S$. Similarly, $f(\overline{x} \circ \mathbb{1}_S)$ is a polynomial of the same form as (15) that depends on the variables x_i , for $i \in S$. In addition, $f(\overline{x})$ has no Boolean roots, so neither does $f(\overline{x} \circ \overline{y})$. This together with (17) means the conditions of Lemma 37 are met, so we have

$$\deg(\operatorname{ml}(g(\overline{x}, \mathbb{1}_S))) = |S|$$

Since we assumed that our monomial ordering respects degree,

$$\deg(\operatorname{LM}(\operatorname{ml}(g(\overline{x}, \mathbb{1}_S)))) = |S|.$$
(18)

There is only one possible multilinear monomial of degree |S| on |S| variables; it follows that every S induces a unique leading monomial (consisting exactly of all variables in S).

This concludes the proof of Theorem 38.

Theorem 40. Let $f(\overline{x})$ be as in (15). Then, any roABP-IPS_{LIN'} refutation of $f(\overline{x} \circ \overline{y}) = 0$ is of size $2^{\Omega(n)}$, when the variables are ordered such that $\overline{x} < \overline{y}$ (i.e., \overline{x} -variables come before \overline{y} -variables).

Proof. Let $g(\overline{x}, \overline{y})$ be a polynomial such that $g(\overline{x}, \overline{y}) \cdot f(\overline{x} \circ \overline{y}) = 1$ over $\overline{x}, \overline{y} \in \{0, 1\}^n$. Hence,

$$g(\overline{x},\overline{y}) = \frac{1}{f(\overline{x}\circ\overline{y})}$$
 over $\overline{x},\overline{y} \in \{0,1\}^n$.

We show that dim $\operatorname{Coeff}_{\overline{x}|\overline{y}g} \geq 2^{\Omega(n)}$. This will conclude the proof by Lemma 14 which will give the roABP size (width) lower bound and by the functional lower bound in Theorem 10.

By lower bounding coefficient dimension by the evaluation dimension over the Boolean cube (Lemma 16),

$$\dim \operatorname{\mathbf{Coeff}}_{\overline{x}|\overline{y}}g \ge \dim \operatorname{\mathbf{Eval}}_{\overline{x}|\overline{y},\{0,1\}}g$$
$$= \dim\{g(\overline{x},\mathbb{1}_S): S \subseteq [n]\}$$
$$\ge \dim\{\operatorname{ml}(g(\overline{x},\mathbb{1}_S)): S \subseteq [n]\}.$$

Here we used that dimension is non-increasing under linear maps. For $S \subseteq [n]$, denoted by $\overline{x}_S := \{x_i : i \in S\}$ and note that for $\overline{x} \in \{0, 1\}^n$,

$$g(\overline{x}, \mathbb{1}_S) = \frac{1}{f(\overline{x}_S)},$$

and that $\operatorname{ml}(g(\overline{x}, \mathbb{1}_S))$ is a multilinear polynomial only depending on \overline{x}_S . By Theorem 38, we can lower bound the number of distinct leading monomials of $\operatorname{ml}(g(\overline{x}, \mathbb{1}_S))$, where S ranges over subsets of [n]:

$$\left| \operatorname{LM} \left(\{ \operatorname{ml}(g(\overline{x}, \mathbb{1}_S)) : S \subseteq [n] \} \right) \right| = 2^n.$$

Therefore, we can lower bound the dimension of the above space by the number of leading monomials (Lemma 17),

$$\dim \operatorname{\mathbf{Coeff}}_{\overline{x}|\overline{y}}g \ge \dim \{\operatorname{ml}(g(\overline{x}, \mathbb{1}_S)) : S \subseteq [n]\}$$
$$\ge \left| \operatorname{LM}\left(\{\operatorname{ml}(g(\overline{x}, \mathbb{1}_S)) : S \subseteq [n]\}\right) \right|$$
$$= 2^n.$$

5.2 roABP-IPS Lower Bounds in Any Order

We now extend this previous result to roABPs in any variable order. Consider a polynomial $f(\overline{w})$ over m variables, where $m = \binom{2n}{2}$ and $\overline{w} = \{w_{i,j}\}_{i < j \in [2n]}$. We apply the same gadget from [HLT24], defined by the mapping

$$w_{i,j} \mapsto z_{i,j} x_i x_j$$

which substitutes the *m* variables $w_{i,j}$ by m + 2n variables $\{z_{i,j}\}_{i < j \in [2n]}, x_1, \ldots, x_{2n}$ such that:

$$f^{\star}(\overline{z},\overline{x}) \coloneqq f(\overline{w})_{w_{i,j}\mapsto z_{i,j}x_ix_j},\tag{19}$$

where $f(\overline{w})_{w_{i,j}\mapsto z_{i,j}x_ix_j}$ means that we apply the lifting $w_{i,j}\mapsto z_{i,j}x_ix_j$ to the \overline{w} variables.

Let $f \in \mathbb{F}[\overline{x}, \overline{y}, \overline{z}]$. We denote by $f_{\overline{z}}$ the polynomial f considered as a polynomial in $\mathbb{F}[\overline{z}](\overline{x}, \overline{y})$, namely as a polynomial whose indeterminates are $\overline{x}, \overline{y}$ and whose scalars are from the ring $\mathbb{F}[\overline{z}]$. We will consider the dimension of a (coefficient) matrix when the entries are taken from the ring $\mathbb{F}[\overline{z}]$, and where the dimension is considered over the field of rational functions $\mathbb{F}(\overline{z})$. Note that for any $\overline{\alpha} \in \mathbb{F}^{\overline{z}}$, we have $f_{\overline{\alpha}}(\overline{x}, \overline{y}) = f(\overline{x}, \overline{y}, \overline{\alpha}) \in \mathbb{F}[\overline{x}, \overline{y}]$. We reference the following simple lemma.

Lemma 41 ([FSTW21]). Let $f \in \mathbb{F}[\overline{x}, \overline{y}, \overline{z}]$. Then for any $\overline{\alpha} \in \mathbb{F}^{|\overline{z}|}$

 $\dim_{\mathbb{F}(\overline{z})} \operatorname{\mathbf{Coeff}}_{\overline{x}|\overline{y}} f_{\overline{z}}(\overline{x},\overline{y}) \geq \dim_{\mathbb{F}} \operatorname{\mathbf{Coeff}}_{\overline{x}|\overline{y}} f_{\overline{\alpha}}(\overline{x},\overline{y}).$

We now prove the proposition below.

Proposition 42. Let $n \ge 1$, $m = \binom{2n}{2}$, and \mathbb{F}_q be a finite field of constant characteristic q. Let $f \in \mathbb{F}[\overline{w}]$ be as in (15), and $f^*(\overline{z}, \overline{x})$ be as in (19). Let $g \in \mathbb{F}[z_1, \ldots, z_m, x_1, \ldots, x_{2n}]$ be a polynomial such that

$$g(\overline{z},\overline{x}) = \frac{1}{f^{\star}(\overline{z},\overline{x})},$$

for $\overline{z} \in \{0,1\}^m$ and $\overline{x} \in \{0,1\}^{2n}$. Let $g_{\overline{z}}$ denote g as a polynomial in $\mathbb{F}[\overline{z}][\overline{x}]$. Then, for any partition $\overline{x} = (\overline{u}, \overline{v})$ with $|\overline{u}| = |\overline{v}| = n$,

$$\dim_{\mathbb{F}(\overline{z})} \operatorname{Coeff}_{\overline{u}|\overline{v}} g_{\overline{z}} \geq 2^{\Omega(n)}.$$

Proof. We embed $\frac{1}{f(\overline{u} \circ \overline{v})}$ in this instance via a restriction of \overline{z} . Define the \overline{z} -evaluation $\overline{\alpha} \in \{0,1\}^{\binom{2n}{2}}$ to restrict g to sum over those $x_i x_j$ in the natural matching between \overline{u} and \overline{v} , so that

$$\alpha_{ij} = \begin{cases} 1 & x_i = u_k, x_j = v_k, \\ 0 & \text{otherwise.} \end{cases}$$

It follows that $g(\overline{u}, \overline{v}, \overline{\alpha}) = \frac{1}{f(\overline{u} \circ \overline{v})}$ for $\overline{u}, \overline{v} \in \{0, 1\}^n$. Suppose for contradiction that there exists a partition $\overline{x} = (\overline{u}, \overline{v})$ with $|\overline{u}| = |\overline{v}| = n$, such that

$$\dim_{\mathbb{F}(\overline{z})} \operatorname{\mathbf{Coeff}}_{\overline{u}|\overline{v}} g_{\overline{z}}(\overline{u},\overline{v}) < 2^{\Omega(n)}$$

By Lemma 41, we get the relation between the coefficient dimensions of $g_{\overline{z}}$ and $g_{\overline{\alpha}}$

$$\dim_{\mathbb{F}} \operatorname{\mathbf{Coeff}}_{\overline{u}|\overline{v}} g_{\overline{\alpha}}(\overline{u}, \overline{v}) \leq \dim_{\mathbb{F}(\overline{z})} \operatorname{\mathbf{Coeff}}_{\overline{u}|\overline{v}} g_{\overline{z}}(\overline{u}, \overline{v}) \\ < 2^{\Omega(n)},$$

which contradicts our lower bound for a fixed partition (Theorem 40).

Corollary 43. Let $n \ge 1$, $m = \binom{2n}{2}$, and \mathbb{F}_q be a finite field with constant characteristics q. Let $f \in \mathbb{F}[\overline{w}]$ be as in (15), and $f^*(\overline{z}, \overline{x})$ be as in (19). Then, any roABP-IPS_{LIN'} refutation (in any variable order) of $f^*(\overline{z}, \overline{x})$ requires $2^{\Omega(n)}$ -size.

Proof. Consider the polynomial $g \in \mathbb{F}[z_1, \ldots, z_m, x_1, \ldots, x_{2n}]$ such that

$$g(\overline{z},\overline{x}) = \frac{1}{f^{\star}(\overline{z},\overline{x})}.$$

for $\overline{z} \in \{0, 1\}^m$ and $\overline{x} \in \{0, 1\}^{2n}$. We will show that any roABP computing g requires width $\geq 2^{\Omega(n)}$ in any variable order. The roABP-IPS_{LIN'} lower bound follows immediately from this functional lower bound on g along with the reduction (Theorem 10).

Suppose that $g(\overline{z}, \overline{x})$ is computable by a width-*r* roABP in *some* variable order. By pushing the \overline{z} variables into the fraction field, it follows that $f_{\overline{z}}$ (*f* as a polynomial in $\mathbb{F}[\overline{z}][\overline{x}]$) is also computable by a width-*r* roABP over $\mathbb{F}(\overline{z})$ in the induced variable order on \overline{x} (Fact 8). By splitting \overline{x} in half along its variable order and by the relation between the width of a roABP and its coefficient dimension (Lemma 14), we obtain

$$\dim_{\mathbb{F}(\overline{z})} \operatorname{Coeff}_{\overline{u}|\overline{v}} g_{\overline{z}} < 2^{\Omega(n)}$$

which contradicts the coefficient dimension lower bound of Proposition 42.

5.3 roABP-IPS Lower Bounds by Multiple

Here we present another roABP-IPS lower bound over finite fields, but this time using the lower bound for multiples method from [FSTW21]. We introduce the following two lemmas from their paper.

Lemma 44 (Corollary 6.23 in [FSTW21]). Let $f \in \mathbb{F}[x_1, \ldots, x_n]$ be defined by $f(\overline{x}) \coloneqq \prod_{i < j} (x_i + x_j + \alpha_{i,j})$ for $\alpha_{i,j} \in \mathbb{F}$. Then for any $0 \neq g \in \mathbb{F}[\overline{x}]$, $g \cdot f$ requires width- $2^{\Omega(n)}$ as a read-twice oblivious *ABP*.

Lemma 45 (Lemma 7.1 in [FSTW21]). Let $f, \overline{g}, \overline{x}^2 - \overline{x} \in \mathbb{F}[x_1, \ldots, x_n]$ be an unsatisfiable systems of equations, where $\overline{g}, \overline{x}^2 - \overline{x}$ is satisfiable. Let $C \in \mathbb{F}[\overline{x}, y, \overline{z}, \overline{w}]$ be an IPS refutation of $f, \overline{g}, \overline{x}^2 - \overline{x}$. Then

$$1 - C(\overline{x}, 0, \overline{g}, \overline{x}^2 - \overline{x})$$

is a nonzero multiple of f.

From here, consider the field \mathbb{F}_q for some constant q, and let the following two polynomials be our hard system of equations.

$$f \coloneqq \prod_{i < j} (x_i + x_j + 1), \quad g \coloneqq \prod_{i=1}^n (1 - x_i) - 1.$$
 (20)

Note that our f above is the same as in [FSTW21], but our g differs (they use $\sum_{i=1}^{n} x_i - n$, which is why they must work in fields of characteristic > n). We now state our lower bound.

Theorem 46. Let \mathbb{F}_q be a finite field of constant characteristic q. Let $f, g \in \mathbb{F}[x_1, \ldots, x_n]$, where $f := \prod_{i < j} (x_i + x_j + 1)$ and $g := \prod_{i=1}^n (1 - x_i) - 1$. Then, the system of equations $f, g, \overline{x}^2 - \overline{x}$ is unsatisfiable, and any roABP-IPS_{LIN'} refutation (in any order of the variables) requires size $\exp(\Omega(n))$.

Proof. The system $g(\overline{x}) = 0$ and $\overline{x}^2 - \overline{x} = 0$ is satisfiable and has the unique satisfying assignment $\overline{0}$. However, this single assignment does not satisfy f as $f(\overline{0}) = \prod_{i < j} (0 + 0 + 1) = 1 \neq 0$, so the entire system is unsatisfiable. Thus by Lemma 45, for any roABP-IPS_{LIN'} refutation $C(\overline{x}, y, z, \overline{w})$ of $f, g, \overline{x}^2 - \overline{x}, 1 - C(\overline{x}, 0, g, \overline{x}^2 - \overline{x})$ is a nonzero multiple of f.

Let s be the size of C as an roABP. We now argue that $1 - C(\overline{x}, 0, g, \overline{x}^2 - \overline{x})$ has a small read-*twice* oblivious ABP. First, note that we can expand $C(\overline{x}, 0, z, \overline{w})$ into powers of z:

$$C(\overline{x}, 0, z, \overline{w}) = C_0(\overline{x}, \overline{w}) + C_1(\overline{x}, \overline{w})z.$$

There are only two terms because $C(\overline{x}, y, z, \overline{w})$ is a roABP-IPS_{LIN'} refutation, implying the degree of z in $C(\overline{x}, y, z, \overline{w})$ is at most 1. Each $C_i(\overline{x}, \overline{w})$ has a poly(s)-size roABP (in the order of the variables of C where z is omitted), as we can compute C_i via interpolation over z (since each evaluation preserves roABP size by Fact 8). Furthermore, as g can also be computed by a poly(n)-size roABP, we see that

$$1 - C(\overline{x}, 0, g, \overline{w}) = 1 - C_0(\overline{x}, \overline{w}) - C_1(\overline{x}, \overline{w})g$$

has a poly(s, n)-size roABP in the order of variables that C induces on $\overline{x}, \overline{w}$. As each Boolean axiom $x_i^2 - x_i$ only refers to a single variable, substituting $\overline{w} \leftarrow \overline{x}^2 - \overline{x}$ for $1 - C(\overline{x}, 0, g, \overline{w})$ in the roABP will preserve obliviousness, but now each variable is read twice. Therefore, $1 - C(\overline{x}, 0, g, \overline{x}^2 - \overline{x})$ has a poly(s, n)-size read-twice oblivious ABP. Finally, using the fact that a nonzero multiple of f requires $\exp(\Omega(n))$ size to be computed as read-twice oblivious ABPs (Lemma 44), it follows that $\operatorname{poly}(s, n) \ge \exp(\Omega(n))$, implying $s \ge \exp(\Omega(n))$ as desired.

5.4 Limitations

The following discusses the limitations of the functional lower bound method for roABP-IPS. Namely, we show that it is impossible to get a non-placeholder functional lower bound against roABP-IPS over finite fields, even if the refutation is restricted to a multilinear polynomial. We first recall this fact about roABPs.

Fact 47. If $f, g \in \mathbb{F}[\overline{x}]$ are computable by width-r and width-s roABPs respectively, then

- f + g is computable by a width-(r + s) roABP.
- $f \cdot g$ is computable by a width-(rs) roABP.

Now, as discussed in Section 4, for a given unsatisfiable instance f in finite field \mathbb{F}_p , by Fermat's Little Theorem we have the following refutation:

$$f(\overline{x})^{p-2}f(\overline{x}) = 1 \mod \overline{x}^2 - \overline{x}.$$
(21)

Thus, if f is easy for roABPs then by Fact 47, so is $f(\overline{x})^{p-2}$ (as p is constant), so in this case it is impossible to acheive a lower bound on roABP-IPS refutations. Now, consider the case where refutations must be multilinear (that is, an analogue to the constant-depth multilinear IPS proof system from Section 3). In this proof system, the refutation in (21) cannot work, as it is not multilinear. However, it is shown in [FSTW21] that roABPs are closed under multilinearization. We restate their result for concreteness.

Proposition 48 (Proposition 4.5 from [FSTW21]). Let $f \in \mathbb{F}[\overline{x}]$ be computable by a width-r roABP, in order of the variables $x_1 < \cdots < x_n$, and with individual degrees at most d. Then, ml(f) has a poly(r, n, d)-explicit width-r roABP in order of the variables $x_1 < \cdots < x_n$.

Thus, we simply consider the multilinear polynomial $g = ml(f^{p-2})$ to be our refutation (as g agrees with f^{p-2} over the Boolean cube, implying (21) holds). By the above proposition, since f^{p-2} has a small roABP computing it, so does g. This leads to the following theorem.

Theorem 49. The functional lower bound method cannot establish non-placeholder lower bounds on the size of roABP-IPS refutations when working in finite fields.

6 Towards Lower Bounds for CNF Formulas

We now turn to the problem of establishing lower bounds for CNF formulas. In the previous sections, the lower bounds we presented were for algebraic instances. In contrast, we show that an IPS lower bound against an unsatisfiable set of one or more polynomial equations over finite fields implies the existence of a hard Boolean instance. However, this implication requires a subsystem of IPS which can reason with large degree, therefore our results do not meet this criteria. Accordingly, the existence of any hard instance for IPS over finite fields (even when the equations are given as algebraic circuits), allowing refutations of possibly exponential total degree, implies the existence of hard Extended Frege instances. Similarly, if the hard instance is only against IPS refutations of polynomial total degree, there are hard instances against Frege.

We work in finite field \mathbb{F}_q where q is a constant (independent of the size of the formulas and their number of variables). When we work with CNF formulas in IPS we assume that the CNF formulas are translated according to the following definition.

Definition 50 (Algebraic translation of CNF formulas). Given a CNF formula in the variables \overline{x} , every clause $\bigvee_{i \in P} x_i \lor \bigvee_{j \in N} \neg x_j$ is translated into $\prod_{i \in P} (1 - x_i) \cdot \prod_{j \in N} x_j = 0$. (Note that these expressions are represented as algebraic circuits, where the products are not expanded.)

Notice that a CNF formula is satisfiable by 0-1 assignment if and only if the assignment satisfies all the equations in the algebraic translation of the CNF. The following definitions are taken from [ST25], and we supply them here for completeness.

Definition 51 (Algebraic extension axioms and unary bits [ST25]). Given a circuit C and a node g in C, we call the equation

$$x_g = \sum_{i=0}^{q-1} i \cdot x_{g_i}$$

the algebraic extension axiom of g, with each variable x_{g_i} being the *i*th unary-bit of g.

Definition 52 (Plain CNF encoding of constant-depth algebraic circuit $\operatorname{cnf}(C(\overline{x}))$ [ST25]). Let $C(\overline{x})$ be a circuit in the variables \overline{x} . The plain CNF encoding of the circuit $C(\overline{x})$, denoted $\operatorname{cnf}(C(\overline{x}))$ consists of the following CNFs in the unary-bits variables of all the gates in C and extra extension variables (and only in the unary-bit variables):

1. If x_i is an input node in C, the plain CNF encoding of C uses the variables $x_{x_{i0}}, \dots, x_{x_{i(q-1)}}$ that are the unary-bits of x_i , and includes clauses ensuring that exactly one unary bit is 1 and all others are 0:

$$\bigvee_{j=0}^{q-1} x_{x_ij} \wedge \bigwedge_{j \neq l \in \{0, \cdots, q-1\}} (\neg x_{x_ij} \vee \neg x_{x_il})$$

- 2. If $\alpha \in \mathbb{F}_q$ is a scalar input node in C, the plain CNF encoding of C contains the $\{0, 1\}$ constants corresponding to the unary-bits of α . These constants are used when fed to (translation of) gates according to the wiring of C in item 4.
- 3. For every node g in $C(\overline{x})$ and every satisfying assignment $\overline{\alpha}$ to the plain CNF encoding, the corresponding unary-bit x_{g_i} evaluates to 1 if and only if the value of g equals $i \in 0, \ldots, q-1$ (when the algebraic inputs $\overline{x} \in (\mathbb{F}_q)^*$ to $C(\overline{x})$ take on the values corresponding to the Boolean assignment $\overline{\alpha}$; "*" here means the Kleene star). This is ensured with the following equations: if g is $a \circ \in \{+, \times\}$ node that has inputs u_1, \cdots, u_t . Then we consider the following equations:

$$u_{1} \circ u_{2} = v_{1}^{g}$$

$$u_{i+2} \circ v_{i}^{g} = v_{i+1}^{g}, \qquad 1 \le i \le t-3$$

$$u_{t} \circ v_{t-2}^{g} = g.$$

Then, for each equation above, for simplicity, we denote as $x \circ y = z$. For each x + y = zwe have a CNF ϕ in the unary-bits variables of x, y, z that is satisfied by assignment precisely when the output unary-bits of z get their correct values based on the (constant-size) truth table of \circ over \mathbb{F}_q and the input unary-bits of x and y (we ensure that if more than one unary-bit is assigned 1 in any of the unary-bits of x, y, z then the CNF is unsatisfiable).

4. For every unary-bit variable x_{g_i} , we have the Boolean axiom (recall we write these Boolean axioms explicitly since we are going to work with IPS^{alg}):

$$x_{g_i}^2 - x_{g_i} = 0$$

Therefore, we can see that the formula size of $\operatorname{cnf}(C(\overline{x}) = 0)$ is $\operatorname{poly}(q^2 \cdot |C|)$.

Definition 53 (Plain CNF encoding of a constant-depth circuit equation $\operatorname{cnf}(C(\overline{x}) = 0)$ [ST25]). Let $C(\overline{x})$ be a circuit in the variables \overline{x} . The plain CNF encoding of the circuit equation $C(\overline{x}) = 0$ denoted $\operatorname{cnf}(C(\overline{x}) = 0)$ consists of the following CNF encoding from Definition 52 in the unary-bits variables of all the gates in C (and only in the unary-bit variables), together with the equations:

$$x_{g_{out}0} = 1$$
 and $x_{g_{out}i} = 0$, for all $i = 1, \cdots, q - 1$,

which express that $g_{out} = 0$, where g_{out} is the output node of C.

Definition 54 (Extended CNF encoding of a circuit equation (circuit, resp.); $ecnf(C(\overline{x}) = 0)$ ($ecnf(C(\overline{x}))$, resp.) [ST25]). Let $C(\overline{x})$ be a circuit in the variables \overline{x} over the finite field \mathbb{F}_q . The extended CNF encoding of the circuit equation $C(\overline{x}) = 0$ (circuit $C(\overline{x})$, resp.), in symbols $ecnf(C(\overline{x}) = 0)$ ($ecnf(C(\overline{x}))$, resp.), is defined to be a set of algebraic equations over \mathbb{F}_q in the variables x_g and x_{g0}, \dots, x_{gq-1} which are the unary-bit variables corresponding to the node g in C, that consist of:

- 1. the plain CNF encoding of the circuit equation $C(\overline{x}) = 0$ (circuit $C(\overline{x})$, resp.), namely, $cnf(C(\overline{x}) = 0)$ ($cnf(C(\overline{x}))$), resp.); and
- 2. the algebraic extension axiom of g, for every gate g in C.

Since we work with extension variables, it is more convenient to express a circuit equation $C(\bar{x}) = 0$ as a set of equations encoding each gate of C, which we call the straight line program of $C(\bar{x})$ (and is equivalent in strength to algebraic circuits).

Definition 55 (Straight line program (SLP)). An SLP of a circuit $C(\overline{x})$, denoted by $SLP(C(\overline{x}))$, is a sequence of equations between variables such that the extension variable for the output node computes the value of the circuit assuming all equations hold. Formally, we choose any topological order $g_1, g_2, \dots, g_i, \dots, g_{|C|}$ on the nodes of the circuit C (that is, if g_j has a directed path to g_k in C then j < k) and define the following set of equations to be the SLP of $C(\overline{x})$:

$$g_i = g_{j1} \circ g_{j2} \circ \cdots \circ g_{jt}$$
 for $o \in \{+, \times\}$ iff g_i is a o node in C with t incoming edges from g_{j1}, \cdots, g_{jt} .

An SLP representation of a circuit equation $C(\bar{x}) = 0$ means that we add to the SLP above the equation $g_{|C|} = 0$, where $g_{|C|}$ is the output node of the circuit.

The following lemma, which we refer to as the *translation lemma* throughout this paper, shows that we can derive (with some additional axioms) the circuit equation $C(\overline{x}) = 0$ given the extended CNF encoding of this circuit equation $ecnf(C(\overline{x}) = 0)$, and vice versa.

Lemma 56 (Translating between extended CNFs and circuit equations in Fixed Finite Fields [ST25]). Let \mathbb{F}_q be a finite field, and let $C(\overline{x})$ be a circuit of depth Δ in the \overline{x} variables over \mathbb{F}_q . Then, both of the following hold:

$$\operatorname{ecnf}(C(\overline{x})=0) \left| \frac{*,O(\Delta)}{\operatorname{IPS^{alg}}} C(\overline{x}) = 0. \right.$$
(22)

$$\{x_g = \sum_{i=0}^{q-1} i \cdot x_{gi} : g \text{ is a node in } C\},\$$

$$\{x_{gi}^2 - x_{gi} = 0 : g \text{ is a node in } C, 0 \le i < q\},\$$

$$\{\sum_{i=0}^{q-1} x_{gi} = 1 : g \text{ is a node in } C\}, \quad \left|\frac{*,O(\Delta)}{\mathsf{IPS}^{\mathsf{alg}}}\right| = \mathsf{cenf}(C(\overline{x}) = 0).\$$

$$\mathsf{SLP}(C(\overline{x})),\$$

$$C(\overline{x}) = 0$$

$$(23)$$

Proposition 57 (Proposition 3.7 in [ST25]). Let $C(\overline{x}) = 0$ be a circuit equation over \mathbb{F}_q where q is any constant prime. Then, $C(\overline{x}) = 0$ is unsatisfiable over \mathbb{F}_q iff $cnf(C(\overline{x}) = 0)$ is an unsatisfiable CNF iff $cnf(C(\overline{x}) = 0)$ is an unsatisfiable set of equations over \mathbb{F}_q .

From here, we extend their result by eliminating these additional axioms in both directions. The only additional axioms we need are the *field axioms* $\{x^q - x = 0 : x \text{ is a variable in } C\}$, which can be easily derived from the Boolean axioms if the variables in the circuit are Boolean (as we are working in finite fields). We use $\mathsf{UBIT}_i(x)$ to denote the following Lagrange polynomial:

$$\frac{\prod_{i=0, i\neq j}^{q-1} (x-i)}{\prod_{i=0, i\neq j}^{q-1} (j-i)},$$

where x can be a single variable or an algebraic circuit. Hence, it is easy to observe that

$$\mathsf{UBIT}_{j}(x) = \begin{cases} 1, & x = j, \\ 0, & \text{otherwise.} \end{cases}$$

Also, suppose x has size |x| and depth Depth(x) (when x is a single variable, it has size 1 and depth 1). Then, $\text{UBIT}_j(x)$ can be computed by an algebraic circuit of size $O(|x|^{q-1})$ and depth Depth(x) + 2. In addition, we introduce a Semi-CNF SCNF, which is a substitution instance of a CNF.

Definition 58 (Semi-CNF SCNF encoding of a constant-depth circuit equation $SCNF(C(\overline{x}) = 0)$). Let $C(\overline{x})$ be a circuit in the variables \overline{x} . The semi-CNF encoding of the circuit equation $C(\overline{x}) = 0$ denoted $SCNF(C(\overline{x}))$ is a substitution instance of the plain CNF encoding in Definition 53 where each unary-bits x_{uj} of all the gates and extra extension variables u is substituted with $UBIT_j(C_u)$ where C_u is the constant-depth algebraic circuit computes u.⁵

We now demonstrate the connection between semi-CNFs and circuit equations.

Theorem 59 (Translate semi-CNFs from circuit equations in Fixed Finite Fields). Let \mathbb{F}_q be a finite field, and let $C(\overline{x})$ be a circuit of depth Δ in the \overline{x} variables over \mathbb{F}_q . Then, the following holds

$$\{x^q - x = 0 : x \text{ is a variable in } C\}, \ C(\overline{x}) = 0 \mid_{\mathsf{IPS}^{\mathsf{alg}}}^{*,O(\Delta)} \mathsf{SCNF}(C(\overline{x}) = 0).$$

Since the field equations $x^q - x = 0$ are efficiently derivable from the Boolean axioms, we get the following for IPS (which by default contains the Boolean axioms):

$$C(\overline{x}) = 0 \left| \frac{*, O(\Delta)}{\mathsf{IPS}^{\mathsf{alg}}} \operatorname{SCNF}(C(\overline{x}) = 0) \right|.$$

Proof. In Lemma 56, the given axioms include:

(i)
$$\{x_g = \sum_{i=0}^{q-1} i \cdot x_{gi} : g \text{ is a node in } u \circ v = w\}$$

(ii) $\{x_{gi}^2 - x_{gi} = 0 : g \text{ is a node in } u \circ v = w\}$,
(iii) $\{\sum_{i=0}^{q-1} x_{gi} = 1 : g \text{ is a node in } u \circ v = w\}$,
(iv) $\sum_{i=0}^{q-1} i \cdot x_{ui} \circ \sum_{i=0}^{q-1} i \cdot x_{vi} = \sum_{i=0}^{q-1} i \cdot x_{wi}$,

and there is a constant-depth constant-size IPS derivation of the plain CNF encoding of $u \circ v = w$. Thus, we must show that we can derive the above four axioms when we substitute x_{gi} with $\mathsf{UBIT}_i(C_g)$. Due to the standard property of Lagrange polynomials, the following circuit equation is a polynomial identity, which can be proved freely in IPS (in finite field \mathbb{F}_q):

$$x = \sum_{j=0}^{q-1} j \cdot \mathsf{UBIT}_j(x),$$

which is exactly the axiom in (i). Hence, we know that $C_u = \sum_{j=0}^{q-1} j \cdot \mathsf{UBIT}_j(C_u), C_v = \sum_{j=0}^{q-1} j \cdot \mathsf{UBIT}_j(C_v)$ and $C_u \circ C_v = C_w = \sum_{j=0}^{q-1} j \cdot \mathsf{UBIT}_j(C_w)$. These polynomial identities give us the substitution instance of the last equation (iv):

$$\sum_{j=0}^{q-1} j \cdot \mathsf{UBIT}_j(C_u) \circ \sum_{j=0}^{q-1} j \cdot \mathsf{UBIT}_j(C_v) = \sum_{j=0}^{q-1} j \cdot \mathsf{UBIT}_j(C_w).$$

The second set of equations (*ii*) is in the ideal of the field axioms for g. We show that in depth- $O(\Delta)$ and polynomial size, we can derive the field axioms $C_g^q - C_g = 0$ for all circuits that compute the nodes and extension variables (using the field axioms $x^q - x = 0$, for every input variable).

⁵This C_u can be constructed from SLPs easily.

We derive the field axioms for nodes and extension variables by induction on depth. When g is a $o \in \{+, \times\}$ node that has inputs u_1, \dots, u_t , the SLPs includes:

$$u_{1} \circ u_{2} = v_{1}^{g}$$

$$u_{i+2} \circ v_{i}^{g} = v_{i+1}^{g}, \qquad 1 \le i \le t-3$$

$$u_{t} \circ v_{t-2}^{g} = g.$$

For each $v_i^g = u_1 \circ \cdots \circ u_{i+2}$, $C_{v_i^g} = C_{u_1} \circ \cdots \circ C_{u_{i+2}}$ is a polynomial identity. By induction, we already have the field axioms for all C_{u_i} . We show that we can derive the field axioms for all $C_{v_i^g}$ and C_q simultaneously. Now suppose $\circ = +$, then the following equations hold over \mathbb{F}_q :

$$C_{v_i^g}^q \equiv (C_{u_1} + \dots + C_{u_{i+2}})^q$$
$$\equiv C_{u_1}^q + \dots + C_{u_{i+2}}^q$$
$$\equiv C_{u_1} + \dots + C_{u_{i+2}}$$
$$\equiv C_{v_i^g}.$$

The proof for the node g is the same. We can therefore conclude that if $\circ = +$, we can derive the field axioms for all $C_{v_i^g}$ and C_g simultaneously. Suppose $\circ = \times$, then the following equations hold over \mathbb{F}_q :

$$C_{v_i^q}^q \equiv (C_{u_1} \times \dots \times C_{u_{i+2}})^q$$
$$\equiv C_{u_1}^q \times \dots \times C_{u_{i+2}}^q$$
$$\equiv C_{u_1} \times \dots \times C_{u_{i+2}}$$
$$\equiv C_{v_i^g}.$$

Again, the proof for the node g is the same and thus we conclude that given the field axioms for the input variables, we can derive the field axioms for all circuits that compute the nodes and extension variables in depth $O(\Delta)$ and polynomial size. It remains to show that $\mathsf{UBIT}_j(x)^2 - \mathsf{UBIT}_j(x) = 0$ is in the ideal of the field axiom of x, for any x. The equation $x^q - x = 0$ is the unique monic polynomial of degree q that has all elements of \mathbb{F}_q as roots. Therefore, any polynomial $f(x) \in \mathbb{F}_q[x]$ that vanishes when evaluated to any $x \in \mathbb{F}_q$ must be divisible by $x^q - x$. It is easy to check that $\mathsf{UBIT}_j(x)^2 - \mathsf{UBIT}_j(x)$ vanishes at all points, implying it is in the ideal generated by $x^q - x$. Hence, there is a degree (of x) q - 2 polynomial Q(x) such that $Q(x) \cdot (x^q - x) = \mathsf{UBIT}_j(x)^2 - \mathsf{UBIT}_j(x)$, and as a result there is a depth- Δ polynomial-size proof for $\mathsf{UBIT}_j(x)^2 - \mathsf{UBIT}_j(x) = 0$ from $x^q - x$.

Finally, $\sum_{j=0}^{q-1} \mathsf{UBIT}_j(x) = 1$ is a polynomial identity, for every x. This follows from the fact that it is a single-variable polynomial with degree q-1, but has q many distinct roots. By the fundamental theorem of algebra, it must be a zero polynomial, and consequently we get axioms in *(iii)* for free in IPS. All together, we can conclude that

$$\{x^q - x = 0 : x \text{ is a variable in } C\}, C(\overline{x}) = 0 \left| \begin{smallmatrix} *, O(\Delta) \\ \mathsf{IPS}^{\mathsf{alg}} \end{smallmatrix} \right| \mathsf{SCNF}(C(\overline{x}) = 0)$$

by first deriving the substitution instance above, and then substituting the derivation for the CNF to get the derivation for the semi-CNF. $\hfill \Box$

Since SCNFs are substitution instances of CNFs, lower bounds for CNFs imply lower bounds for SCNFs, which gives the following corollary.

Corollary 60 (Lower bounds for circuit equations imply lower bounds for CNFs). Let \mathbb{F}_q be a finite field, and let $\{C(\overline{x})\}$ be a set of circuits of depth at most Δ in the Boolean variable \overline{x} . Then, if a set of circuit equations $\{C(\overline{x}) = 0\}$ cannot be refuted in S-size, $O(\Delta')$ -depth IPS, then the CNF encoding of the set of circuit equations $\{\mathsf{CNF}(C(\overline{x}) = 0)\}$ cannot be refuted in $(S - \mathsf{poly}(|C|))$ -size, $O(\Delta' + \Delta)$ -depth IPS.

Lemma 61 (Translate circuit equations from semi-CNFs in fixed finite fields). Let \mathbb{F}_q be a finite field, and let $C(\overline{x})$ be a circuit of depth Δ in the \overline{x} variables over \mathbb{F}_q . Then, the following holds:

$$\{x^q - x = 0 : x \text{ is a variable in } C\}, \ \mathsf{SCNF}(C(\overline{x}) = 0) \mid_{\mathsf{IPS^{alg}}}^{*,O(\Delta)} C(\overline{x}) = 0.$$

Proof. From the CNF encoding of each SLP axiom $u \circ v = w$ and the Boolean axioms for each unary bit, we have

$$\sum_{j=0}^{q-1} j \cdot x_{uj} \circ \sum_{j=0}^{q-1} j \cdot x_{vj} = \sum_{j=0}^{q-1} j \cdot x_{wj}$$

in constant-depth polynomial-size IPS. As we showed in the proof of Theorem 59, the field axioms for all circuits that compute nodes and extension variables can be derived from the field axioms of the input variables, in constant-depth polynomial-size IPS. We also showed that these derived field axioms can in turn also derive the Boolean axioms for all UBIT polynomials (of circuits that compute nodes and extension variables) in constant-depth polynomial-size IPS. As a result, substituting each x_{gj} above with $\text{UBIT}_j(C_g)$ for $g \in \{u, v, w\}$ and $0 \le j \le q-1$, we get a constantdepth polynomial-size IPS derivation of

$$\sum_{j=0}^{q-1} j \cdot \mathsf{UBIT}_j(C_u) \circ \sum_{j=0}^{q-1} j \cdot \mathsf{UBIT}_j(C_v) = \sum_{j=0}^{q-1} j \cdot \mathsf{UBIT}_j(C_w)$$

from the semi-CNF encoding of $u \circ v = w$ and the Boolean axioms for each UBIT. Lastly, as mentioned in the proof of Theorem 59, in finite field \mathbb{F}_q we get the following circuit equation for free in IPS (as it is a polynomial identity):

$$x = \sum_{j=0}^{q-1} j \cdot \mathsf{UBIT}_j(x).$$

Therefore, we get the full SLP for the circuit equation $C(\overline{x}) = 0$, and consequently the circuit equation can easily be obtained from this SLP.

Acknowledgment

We are thankful to Rahul Santhanam for helpful discussions on the translation lemmas (Section 6) and for Behera, Limaye, Ramanathan, and Srinivasan for letting us know about their independent work.

References

[AF22] Robert Andrews and Michael A. Forbes. "Ideals, Determinants, and Straightening: Proving and Using Lower Bounds for Polynomial Ideals". In: 54th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2022. 2022. arXiv: 2112.00792. URL: https://arxiv.org/ abs/2112.00792.

- [AGHT20] Yaroslav Alekseev, Dima Grigoriev, Edward A. Hirsch, and Iddo Tzameret. "Semi-algebraic proofs, IPS lower bounds, and the τ-conjecture: can a natural number be negative?" In: Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020. ACM, 2020, pp. 54–67.
- [AGKST23] Prashanth Amireddy, Ankit Garg, Neeraj Kayal, Chandan Saha, and Bhargav Thankey. "Low-Depth Arithmetic Circuit Lower Bounds: Bypassing Set-Multilinearization". In: 50th International Colloquium on Automata, Languages, and Programming (ICALP 2023). Ed. by Kousha Etessami, Uriel Feige, and Gabriele Puppis. Vol. 261. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl Leibniz-Zentrum für Informatik, 2023, 12:1–12:20. ISBN: 978-3-95977-278-5. DOI: 10.4230/LIPIcs.ICALP.2023.12. URL: https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ICALP.2023.12.
- [Ale21] Yaroslav Alekseev. "A Lower Bound for Polynomial Calculus with Extension Rule". In: 36th Computational Complexity Conference, CCC 2021, July 20-23, 2021, Toronto, Ontario, Canada (Virtual Conference). Ed. by Valentine Kabanets. Vol. 200. LIPIcs. Schloss Dagstuhl -Leibniz-Zentrum für Informatik, 2021, 21:1–21:18. DOI: 10.4230/LIPIcs.CCC.2021.21. URL: https://doi.org/10.4230/LIPIcs.CCC.2021.21.
- [BDS24] C.S. Bhargav, Sagnik Dutta, and Nitin Saxena. "Improved Lower Bound, and Proof Barrier, for Constant Depth Algebraic Circuits". In: ACM Trans. Comput. Theory 16.4 (Nov. 2024).
 ISSN: 1942-3454. DOI: 10.1145/3689957. URL: https://doi.org/10.1145/3689957.
- [BIKPP96] Paul Beame, Russell Impagliazzo, Jan Krajíček, Toniann Pitassi, and Pavel Pudlák. "Lower bounds on Hilbert's Nullstellensatz and propositional proofs". In: Proc. London Math. Soc. (3) 73.1 (1996), pp. 1–26. DOI: 10.1112/plms/s3-73.1.1.
- [BLRS25] Amik Raj Behera, Nutan Limaye, Varun Ramanathan, and Srikanth Srinivasan. "New Bounds for the Ideal Proof System in Positive Characteristic". In: 52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025). To appear. Aarhus, Denmark, July 2025.
- [Bus12] Samuel Buss. "Towards **NP-P** via Proof Complexity and Search". In: Annals of Pure and Applied Logic 163.7 (2012), pp. 906–917.
- [CEI96] Matthew Clegg, Jeffery Edmonds, and Russell Impagliazzo. "Using the Groebner basis algorithm to find proofs of unsatisfiability". In: Proceedings of the 28th Annual ACM Symposium on the Theory of Computing (Philadelphia, PA, 1996). New York: ACM, 1996, pp. 174–183.
- [CLO15] David Cox, John Little, and Donal O'Shea. Ideals, Varieties, and Algorithms: An Introduction to Computational Algebraic Geometry and Commutative Algebra. Fourth. Undergraduate Texts in Mathematics. Springer Cham, 2015, pp. XVI, 646. DOI: 10.1007/978-3-319-16721-3.
- [DMM24] Yogesh Dahiya, Meena Mahajan, and Sasank Mouli. "New Lower Bounds for Polynomial Calculus over Non-Boolean Bases". In: 27th International Conference on Theory and Applications of Satisfiability Testing (SAT 2024). Vol. 305. LIPIcs. Leibniz Int. Proc. Inform. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2024, Art. No. 10, 20. DOI: 10.4230/lipics.sat. 2024.10. URL: https://doi.org/10.4230/lipics.sat.2024.10.
- [For24] Michael A. Forbes. "Low-depth algebraic circuit lower bounds over any field". In: 39th Computational Complexity Conference. Vol. 300. LIPIcs. Leibniz Int. Proc. Inform. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2024, Art. No. 31, 16. DOI: 10.4230/lipics.ccc.2024.31. URL: https://doi.org/10.4230/lipics.ccc.2024.31.
- [FS13] Michael A. Forbes and Amir Shpilka. "Quasipolynomial-Time Identity Testing of Noncommutative and Read-Once Oblivious Algebraic Branching Programs". In: FOCS 2013. ArXiv 1209.2408. 2013, pp. 243–252. DOI: 10.1109/F0CS.2013.34.
- [FSTW21] Michael A. Forbes, Amir Shpilka, Iddo Tzameret, and Avi Wigderson. "Proof Complexity Lower Bounds from Algebraic Circuit Complexity". In: *Theory Comput.* 17 (2021), pp. 1–88. URL: https://theoryofcomputing.org/articles/v017a010/.

- [GH03] Dima Grigoriev and Edward A. Hirsch. "Algebraic proof systems over formulas". In: Theoret. Comput. Sci. 303.1 (2003). Logic and complexity in computer science (Créteil, 2001), pp. 83– 102. ISSN: 0304-3975.
- [GHT22] N. Govindasamy, T. Hakoniemi, and I. Tzameret. "Simple Hard Instances for Low-Depth Algebraic Proofs". In: 2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS). Los Alamitos, CA, USA: IEEE Computer Society, Nov. 2022, pp. 188–199. DOI: 10.1109/F0CS54457.2022.00025. URL: https://doi.ieeecomputersociety.org/10.1109/ F0CS54457.2022.00025.
- [GP18] Joshua A. Grochow and Toniann Pitassi. "Circuit Complexity, Proof Complexity, and Polynomial Identity Testing: The Ideal Proof System". In: J. ACM 65.6 (2018), 37:1–37:59. DOI: 10.1145/3230742. URL: https://doi.org/10.1145/3230742.
- [Gro23] Joshua A. Grochow. Polynomial Identity Testing and the Ideal Proof System: PIT is in NP if and only if IPS can be p-simulated by a Cook–Reckhow proof system. arXiv:2306.02184 [cs.CC]. 2023.
- [HLT24] Tuomas Hakoniemi, Nutan Limaye, and Iddo Tzameret. "Functional Lower Bounds in Algebraic Proofs: Symmetry, Lifting, and Barriers". In: Proceedings of the 56th Annual ACM Symposium on Theory of Computing. STOC 2024. Vancouver, BC, Canada: Association for Computing Machinery, 2024, pp. 1396–1404. ISBN: 9798400703836. DOI: 10.1145/3618260. 3649616. URL: https://doi.org/10.1145/3618260.3649616.
- [IMP20] Russell Impagliazzo, Sasank Mouli, and Toniann Pitassi. "The Surprising Power of Constant Depth Algebraic Proofs". In: LICS '20: 35th Annual ACM/IEEE Symposium on Logic in Computer Science, Saarbrücken, Germany, July 8-11, 2020. Ed. by Holger Hermanns, Lijun Zhang, Naoki Kobayashi, and Dale Miller. ACM, 2020, pp. 591–603. DOI: 10.1145/3373718. 3394754. URL: https://doi.org/10.1145/3373718.3394754.
- [IMP23] Russell Impagliazzo, Sasank Mouli, and Toniann Pitassi. "Lower Bounds for Polynomial Calculus with Extension Variables over Finite Fields". In: CCC '23. Warwick, United Kingdom: Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2023. ISBN: 9783959772822. DOI: 10.4230/ LIPIcs.CCC.2023.7. URL: https://doi.org/10.4230/LIPIcs.CCC.2023.7.
- [LST21] Nutan Limaye, Srikanth Srinivasan, and Sébastien Tavenas. "Superpolynomial Lower Bounds Against Low-Depth Algebraic Circuits". In: 62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022. IEEE, 2021, pp. 804– 814. DOI: 10.1109/F0CS52979.2021.00083. URL: https://doi.org/10.1109/F0CS52979. 2021.00083.
- [LTW18] Fu Li, Iddo Tzameret, and Zhengyu Wang. "Characterizing Propositional Proofs as Noncommutative Formulas". In: SIAM Journal on Computing. Vol. 47. 4. Full Version: http: //arxiv.org/abs/1412.8746. 2018, pp. 1424–1462.
- [Luc78] Edouard Lucas. "Theorie des Fonctions Numeriques Simplement Periodiques". In: Amer. J. Math. 1.2 (1878), pp. 184–196. ISSN: 0002-9327. DOI: 10.2307/2369308. URL: https://doi. org/10.2307/2369308.
- [Nis91] Noam Nisan. "Lower Bounds for Non-Commutative Computation". In: *STOC 1991*. 1991, pp. 410–418. DOI: 10.1145/103418.103462.
- [Pit97] Toniann Pitassi. "Algebraic propositional proof systems". In: Descriptive complexity and finite models (Princeton, NJ, 1996). Vol. 31. DIMACS Ser. Discrete Math. Theoret. Comput. Sci. Providence, RI: Amer. Math. Soc., 1997, pp. 215–244.
- [Pit98] Toniann Pitassi. "Unsolvable systems of equations and proof complexity". In: Proceedings of the International Congress of Mathematicians, Vol. III (Berlin, 1998). Vol. III. 1998, pp. 451– 460.
- [PT16] Tonnian Pitassi and Iddo Tzameret. "Algebraic Proof Complexity: Progress, Frontiers and Challenges". In: ACM SIGLOG News 3.3 (2016). Ed. by Andrzej Murawski.

- [RT08a] Ran Raz and Iddo Tzameret. "Resolution over linear equations and multilinear proofs". In: Ann. Pure Appl. Logic 155.3 (2008), pp. 194–224. DOI: 10.1016/j.apal.2008.04.001. URL: http://dx.doi.org/10.1016/j.apal.2008.04.001.
- [RT08b] Ran Raz and Iddo Tzameret. "The Strength of Multilinear Proofs". In: Computational Complexity 17.3 (2008), pp. 407–457. DOI: 10.1007/s00037-008-0246-0. URL: http://dx.doi. org/10.1007/s00037-008-0246-0.
- [Sap12] Ramprasad Saptharishi. Personal communication to Forbes-Shpilka [FS13]. 2012.
- [Sap22] Ramprasad Saptharishi. A survey of lower bounds in arithmetic circuit complexity. 2016-2022. URL: %5Curl%7Bhttps://github.com/dasarpmar/lowerbounds-survey/releases%7D.
- [Sok20] Dmitry Sokolov. "(Semi)Algebraic proofs over ±1 variables". In: Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing. STOC 2020. Chicago, IL, USA: Association for Computing Machinery, 2020, pp. 78–90. ISBN: 9781450369794. DOI: 10.1145/3357713.3384288. URL: https://doi.org/10.1145/3357713.3384288.
- [ST25] Rahul Santhanam and Iddo Tzameret. "Iterated Lower Bound Formulas: A Diagonalization-Based Approach to Proof Complexity". In: SIAM Journal on Computing 0.0 (2025), STOC21-313-STOC21-349. DOI: 10.1137/21M1447519. eprint: https://doi.org/10.1137/ 21M1447519. URL: https://doi.org/10.1137/21M1447519.
- [SY10] Amir Shpilka and Amir Yehudayoff. "Arithmetic Circuits: A survey of recent results and open questions". In: Foundations and Trends in Theoretical Computer Science 5.3-4 (2010), pp. 207– 388. DOI: 10.1561/0400000039.
- [Tza11] Iddo Tzameret. "Algebraic proofs over noncommutative formulas". In: *Inf. Comput.* 209.10 (2011), pp. 1269–1292. DOI: 10.1016/j.ic.2011.07.004. URL: http://dx.doi.org/10. 1016/j.ic.2011.07.004.

— Page left blank for ECCC stamp —