

Simple Hard Instances for Low-Depth Algebraic Proofs*

Nashlen Govindasamy[†] Tuomas Hakoniemi[‡] Iddo Tzameret[§]

Imperial College London
Department of Computing

Abstract

We prove super-polynomial lower bounds on the size of propositional proof systems operating with constant-depth algebraic circuits over fields of zero characteristic. Specifically, we show that the subset-sum variant $\sum_{i,j,k,\ell \in [n]} z_{ijkl} x_i x_j x_k x_\ell - \beta = 0$, for Boolean variables, does not have polynomial-size IPS refutations where the refutations are multilinear and written as constant-depth circuits.

Andrews and Forbes (STOC'22) established recently a constant-depth IPS lower bound, but their hard instance does not have itself small constant-depth circuits, while our instance is computable already with small depth-2 circuits.

Our argument relies on extending the recent breakthrough lower bounds against constant-depth algebraic circuits by Limaye, Srinivasan and Tavenas (FOCS'21) to the functional lower bound framework of Forbes, Shpilka, Tzameret and Wigderson (ToC'21), and may be of independent interest. Specifically, we construct a polynomial f computable with small-size constant-depth circuits, such that the multilinear polynomial computing $1/f$ over Boolean values and its appropriate set-multilinear projection are hard for constant-depth circuits.

1 Introduction

Proof complexity predominantly aims to establish lower bounds on proof size in different proof systems. From the perspective of complexity theory this can be viewed as the dual goal to circuit complexity. While circuit complexity aims to prove lower bounds on minimal circuit size required to decide membership in certain languages, e.g., SAT, proof complexity aims to establish lower bounds on the minimal size of proofs witnessing membership in a certain language, e.g., UNSAT, the language of unsatisfiable Boolean formulas. Here a proof is simply a witness that can be checked efficiently.

In circuit complexity it is usual to consider different or restricted types of circuits (e.g., constant-depth circuits or constant-depth circuits with counting gates modulo a prime). Similarly, in proof complexity it is standard to consider different or restricted types of proof systems, for example proofs using a prescribed set of inference rules that derive clauses from existing ones (i.e., resolution).

*This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 101002742).

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

[‡]E-mail: t.hakoniemi@imperial.ac.uk

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

The ideal goal of circuit complexity is to prove that the class NP is different from the class of languages decidable by polynomial-size circuits (and hence $\text{NP} \neq \text{P/poly}$). Similarly, the overarching view of proof complexity is that of an attempt to prove lower bounds on stronger and stronger proof systems in the hope to get as close as possible to ruling out the existence of any proof system that admits short proofs for all unsatisfiable formulas (namely, membership in UNSAT; and similarly for proving membership in other important languages). A language that admits no short efficiently verifiable proofs is by definition outside the class NP, from which we conclude, in the case of ruling out short proofs of UNSAT, that $\text{NP} \neq \text{coNP}$ (and hence $\text{NP} \neq \text{P}$). This view of proof complexity is usually called *The Cook’s Programme* of proof complexity.

1.1 Algebraic Proof Systems

One important strand of proof complexity deals with *algebraic proof systems* of increasing strength. Algebraic proof systems prove that a set of multivariate polynomials do not have a common 0-1 root over a field. The arguably canonical algebraic proof system is the relatively weak Polynomial Calculus (PC for short) [6] in which proofs start from a set of polynomial equations, and proceed to add and multiply existing polynomials until one reaches the unsatisfiable equation $1 = 0$ (proving that the initial polynomials do not have a common 0-1 solution). The “static” version of the polynomial calculus is called Nullstellensatz [4]. In Nullstellensatz a proof of the unsatisfiability of a set of axioms, given as polynomial equations $\{f_i(\bar{x}) = 0\}$ over a field, is simply a *single* polynomial combination of the axioms that equals 1 as a formal polynomial, namely:

$$\sum_i g_i(\bar{x}) \cdot f_i(\bar{x}) = 1, \tag{1}$$

for some polynomials $\{g_i(\bar{x})\}$ (it is said to be *static* because the proof is given as a single polynomial combination instead of deriving 1 “dynamically” step-by-step as in PC).

The size-complexity of proofs in both PC and Nullstellensatz is sparsity, namely the total number of monomials in all the polynomials appearing along the proof. The sparsity measure is what makes these proof systems weak (e.g., even a simple proof-line like $(x_1 - 1) \cdots (x_n - 1) = 0$ accounts for an exponential size because the number of monomials in it is 2^n).

While counting the total number of monomials in algebraic proofs towards their size-complexity yields comparatively weak proof systems, it is natural to think of stronger algebraic proofs by representing polynomials in a more compact manner than sparsity. In particular, one can consider writing polynomials using algebraic circuits. This idea has circulated in proof complexity starting from Pitassi [19, 20], and subsequently in Grigoriev and Hirsch [11], Raz and Tzameret [23, 22, 28], and finally in the introduction of the Ideal Proof System (IPS) by Grochow and Pitassi [10] which loosely speaking is the Nullstellensatz proof system in which proofs are written as algebraic circuits (indeed, [9] showed that IPS is equivalent to Nullstellensatz in which the polynomials g_i in Equation 1 are written as algebraic circuits).

Accordingly, it is natural to consider proof systems that sit between the weak Nullstellensatz on the one end and the strong IPS on the other end. This is done by writing polynomials in proofs with restricted kind of algebraic circuits, such as constant-depth circuits [11, 10, 13, 3], noncommutative formulas [28, 17, 28], algebraic branching programs [28, 9, 16], multilinear formulas [23, 22, 9] and very recently algebraic proofs with additional extension variables over large fields [1] or finite fields [14].

1.2 State of the Art in Algebraic Proof-Size Lower Bounds

For the *weaker* end of the algebraic proof systems' hierarchy many size lower bounds are known. Beginning in the works of Beame *et al.* [4] and Buss *et al.* [5] on Nullstellensatz, through the first Polynomial Calculus (PC) lower bound by Razborov [24], and the PC subset-sum lower bound by Impagliazzo, Pudlak and Sgall [15] (the simplest form of the subset sum principle, also called sometimes *Knapsack*, is the unsatisfiable over 0-1 values equation $\sum_{i=1}^n x_i - \beta = 0$, for $\beta > n$), as well as many other results.

Only recently, lower bounds against *stronger* algebraic proof systems were established. Forbes, Shpilka, Tzameret and Wigderson [9] considered subsystems of IPS using read-once oblivious algebraic programs (roABP) and multilinear formulas over large fields. However these subsystems are not necessarily comparable with constant-depth fragments of IPS which are the focus of the current work.

Alekseev [1] established lower bounds against the Polynomial Calculus with additional extension variables (i.e., variables that abbreviate polynomials with a single fresh variable) over large fields. This result is quite strong, since the proof system simulates strong propositional-logic systems like extended Frege¹. However, this proof system is (apparently) weaker than IPS. Furthermore, the complexity of proofs in this system is measured by bit-size (i.e., coefficients of monomials in each proof-line are written using binary notation; hence, even a short proof using a small number of steps can incur an exponential blow-up if it uses coefficients of super-exponential magnitude, as shown by Alekseev). Lastly, the hard instance in [1] uses coefficients of exponential magnitude, and this is crucial to the lower bound argument.

Very recently, Impagliazzo, Mouli and Pitassi [14] established lower bounds against PC with restricted number of extension variables over finite fields for CNF formulas. However this proof system is apparently weaker (or incomparable to) constant-depth IPS, and is weaker than PC with proof-lines written as constant-depth circuits, because of the restriction on the number of allowed extension variables.

The following lower bounds form the frontiers of what is known about the complexity of strong algebraic proof systems most relevant to our work (i.e., IPS of increasing depth, beginning from Nullstellensatz, which is equivalent to depth-2 IPS, up to unbounded depth IPS):

- (i) *Conditional* lower bounds against IPS proofs for the Binary Value Principle $\sum_{i=1}^n 2^{i-1}x_i = -1$ (a subset sum instance with coefficients of exponential magnitude) by Alekseev, Hirsch, Grigoriev and Tzameret [2]. Apart from this result being conditional, the hard instances use coefficients of exponential magnitude, and this is crucial to the lower bound argument.
- (ii) Andrews and Forbes [3] very recently proved constant-depth IPS lower bounds. However, the hard instance itself cannot be computed by a polynomial-size constant-depth circuit, and this fact is crucial to the lower bound proof.

1.3 Our Results

We establish super-polynomial constant-depth IPS lower bounds for a subset sum instance with small coefficients (i.e., 0-1 coefficients) that is computable by an $O(n^5)$ -size *depth-2* circuits, and where the IPS proof is multilinear.

¹Though the hard instance is not a CNF, and so the lower bound in [1] does not imply Extended Frege lower bounds.

To understand better the proof system we work against, recall the proof shown in [Equation 1](#), in which the g_i 's are written as algebraic circuits—this is (equivalent) to the general IPS system. We shall work with the proof system multilinear $\text{IPS}_{\text{LIN}'}$, following the notation $\text{IPS}_{\text{LIN}'}$ in [9]. Proofs in multilinear $\text{IPS}_{\text{LIN}'}$ of the unsatisfiability of $\{f_i(\bar{x}) = 0\}$ are (roughly; the actual proof system is in fact stronger than this, see [Definition 2](#)) defined as the following polynomial identity

$$\sum_i g_i(\bar{x}) \cdot f_i(\bar{x}) + \sum_j h_j(\bar{x}) \cdot (x_j^2 - x_j) = 1$$

where $h_j(\bar{x})$ are some polynomials and the $g_i(\bar{x})$'s are *multilinear* polynomials, and the $g_i(\bar{x})$'s and $h_j(\bar{x})$'s are all written with constant-depth circuits (but *not* necessarily multilinear formulas, in contrast to multilinear-formula $\text{IPS}_{\text{LIN}'}$ as in [9]).

Theorem 1 (Informal; see [Theorem 3](#)). *Every constant-depth multilinear $\text{IPS}_{\text{LIN}'}$ refutation of the subset sum variant $\sum_{i,j,k,\ell \in [n]} z_{ijkl} x_i x_j x_k x_\ell - \beta$ (for $\beta \notin \{0, \dots, n^4\}$) requires super-polynomial (in n) size.*

Significance of the Results and Context. This is the first constant-depth IPS lower bound on an instance that is computable itself with small constant-depth circuits, when the polynomial that constitutes the IPS proof is multilinear (see [Definition 2](#)). Our hard instance has coefficients of small magnitude, and the lower bound is in the stronger unit cost model of algebraic circuits (i.e., in terms of the size of the circuits, not the size of the binary representation of the coefficients appearing in them). Thus, we rectify all the purported shortcomings of previous constant-depth algebraic proofs lower bounds (while paying by requiring that the IPS proofs are (partially) multilinear).

[Theorem 1](#) contributes to the tradition of showing that simple subset sum variants are hard for algebraic proofs. While Impagliazzo *et al.* [15] initially showed that the subset sum principle requires exponentially many monomials in PC refutations, and Forbes *et al.* [9] extended this to roABP and multilinear formulas, we show this hardness holds at least up to constant-depth IPS (when the proofs are multilinear).

Subset sum variants are not translations of CNF formulas or Boolean formulas more generally. Hence, algebraic-proofs lower bounds for them do not imply (immediately at least) propositional logic proof size lower bounds (i.e., Frege-style proofs). However, a major motivation behind investigating the complexity of algebraic proof systems is to understand the power of algebraic reasoning and proofs (and their algorithmic counterpart, e.g., Gröbner basis computations). For this purpose, it is enough to prove lower bounds on hard instances that are not necessarily translations of CNFs or Boolean formulas. Indeed many works on the complexity of algebraic proof systems are dedicated to establishing such lower bounds, most prominently the subset sum principle, and its variants (see also Razborov [24] non-CNF pigeonhole principle).

Furthermore, lower bounds on the size of algebraic proofs of subset-sum instances, and generally instances from the language of unsatisfiable 0-1 multivariate polynomials over a field, are as relevant to the Cook's programme mentioned above as much as Boolean formulas. The reason is that this language is a coNP-complete language, since we can efficiently check if a given 0-1 assignment satisfies all the polynomials in the system (assuming the polynomials and field elements are written in some standard way), and Boolean unsatisfiability is easily reducible to this language.

We explain in what follows our main technical contribution, which can be of independent interest.

1.4 Proof Technique

Our proof draws techniques from two sources. We use the methods presented in Limaye, Srinivasan and Tavenas [18] to prove superpolynomial lower bounds for constant-depth algebraic circuits, and combine these with the functional lower bound framework of Forbes, Shpilka, Tzameret and Wigderson [9] for size lower bounds on IPS proofs (see also [8] for the functional lower bound approach in algebraic circuit complexity in general).

In general, we prove [Theorem 1](#) by reducing the task of lower bounding the size of a constant-depth algebraic circuit computing the multilinear polynomial that constitutes the IPS proof into the following task: lower bound the size of a constant-depth *set-multilinear circuit* computing an associated *set-multilinear polynomial*. To get the new associated set-multilinear polynomial from the original multilinear IPS proof (which is not necessarily set-multilinear by itself) we use a variant of the functional lower bound approach with some additional arguments that we introduce to deal with the need to focus on set-multilinear monomials within a polynomial that is not set-multilinear.

Once, we have the associated set-multilinear polynomial we can use the reduction presented in [18] from constant-depth general circuits to constant-depth set-multilinear circuits. The reduction loses only a constant-factor in the depth, but pays quite heavily in the degree of the set-multilinear output polynomial. So in order to keep the size of the obtained set-multilinear circuit reasonable, we need to restrict the degree of the set-multilinear polynomial considerably.

Notice that unlike in Limaye *et al.* [18] (or in circuit complexity in general), we do not work from the get-go with a set-multilinear polynomial for which we need to prove a lower bound against constant-depth circuits computing it. In our case, we need to somehow show that this set-multilinear polynomial is “embedded” in some way in *any* multilinear IPS proof of our hard instance. This is the main technical challenge we face in this work.

Another point is that to show our simple degree-2 instance from [Theorem 1](#) is hard we use a substitution in this simple instance. Specifically, by assigning some of the z_{ijkl} and x_i variables in the hard instance, we show that one gets another variant of subset sum denoted ks_w (ks stands for Knapsack) that is defined with respect to some word $w \in \mathbb{Z}^d$ (in the sense of [18]). The definition of the subset sum over w is designed so that the multilinear IPS refutation of the simple hard instance from [Theorem 1](#) “embeds” P_w after applying the substitution to the z_{ijkl} (and x_i) variables, where P_w is the word polynomial from [18], which induces a full-rank coefficient matrix on a set of all set-multilinear monomials that arise from the given word w (a coefficient matrix of a polynomial is an associated matrix whose rank serves as a complexity measure for the polynomial’s circuit size, and in which each entry is a coefficient of a specific monomial in the polynomial).

The meaning of a polynomial “embedding” a set-multilinear polynomial refers to the set-multilinear polynomial being the set-multilinear *projection* of the original polynomial. Hence, we need to consider the *projection* to the space of all set-multilinear polynomials over a particular variable-partition of our original polynomial. We extend the evaluation dimension method from [9] to prove a rank lower bound for the coefficient matrix of this set-multilinear projection, which yields the set-multilinear circuit lower bound via a lemma from [18].

To add on the above, our proof diverges from its forbears in some essential ways. Firstly, as mentioned before, when [18] can work from the get-go with a low-degree set-multilinear polynomial, the multilinear refutations we consider are not of low-degree nor are they set-multilinear. Thus we need to find suitable set-multilinear polynomials within the refutations, and consider projections to the space of set-multilinear polynomials with respect to some variable-partition. Secondly, we use

the method based on partial assignments (or evaluations) from [9] to prove our rank lower bound. Our use of these partial assignments is however more subtle than the evaluation dimension method of [9].

Forbes *et al.* [9] showed a rank lower bound against a coefficient matrix of a polynomial by reducing it to *dimension* lower bound for the space of all multilinearizations of the polynomials obtained after appropriate partial assignments (to the \bar{y} -variables of a polynomial in both \bar{x} and \bar{y} variables; see [9]). The dimension lower bound is then proved by showing that enough linearly independent monomials appear in the space as leading monomials. Such an argument is not enough in our case, because we need to prove the existence of specific set-multilinear monomials (with respect to a words w) in the original polynomial (namely, the IPS proof polynomial). The root of this difference is that the original argument defines partial assignments over *all* the monomials in the hard polynomial, while we try to argue for a rank lower bound for the coefficient matrix of only a *part* of the monomials, namely only the set-multilinear monomials within a non-set-multilinear polynomial—and this corresponds to only a *submatrix* of the full coefficient matrix. The drawback here is that this also forces us to consider only multilinear refutations (because multilinearization as used in [9] does not increase rank but preserves the high rank of the coefficient matrix of the original polynomial [before multilinearization], while multilinearization of a polynomial *can* increase the rank of the set-multilinear coefficient *sub*-matrix; e.g., multilinearizing the non-set-multilinear monomial $x_1^2 y_2^3$ produces $x_1 y_2$ which is set-multilinear, assuming the variables are partitioned into \bar{x} and \bar{y} variables).

Note also that our proof technique is completely different from Andrews and Forbes [3] who used the *hard multiples* framework from [9].

2 Preliminaries

2.1 Polynomials and Algebraic Circuits

For excellent treatises on algebraic circuits and their complexity see Shpilka and Yehudayoff [27] as well as Saptharishi [25]. Let \mathbb{G} be a ring. Denote by $\mathbb{G}[X]$ the ring of (commutative) polynomials with coefficients from \mathbb{G} and variables $X := \{x_1, x_2, \dots\}$. 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 \times 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 $\bar{X} = \langle X_1, \dots, X_d \rangle$ be a sequence of pairwise disjoint sets of variables, called *variable-partition*. We call a monomial m in the variables $\bigcup_{i \in [d]} X_i$ *set-multilinear* over the variable-partition \bar{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 \bar{X} if it is a linear combination of set-multilinear monomials over \bar{X} . For a sequence \bar{X} of sets of variables, we denote by $\mathbb{F}_{\text{sm}}[\bar{X}]$ the space of all polynomials that are set-multilinear over \bar{X} .

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

2.2 Strong Algebraic Proof Systems

For a survey about algebraic proof systems and their relations to algebraic complexity see the survey [21]. Grochow and Pitassi [12] suggested the following algebraic proof system which is essentially a Nullstellensatz proof system ([4]) 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 [9]:

Definition 2 (Ideal Proof System (IPS), Grochow-Pitassi [12]). *Let $f_1(\bar{x}), \dots, f_m(\bar{x}), p(\bar{x})$ be a collection of polynomials in $\mathbb{F}[x_1, \dots, x_n]$ over the field \mathbb{F} . An **IPS proof of $p(\bar{x}) = 0$ from axioms $\{f_j(\bar{x}) = 0\}_{j=1}^m$** , showing that $p(\bar{x}) = 0$ is semantically implied from the assumptions $\{f_j(\bar{x}) = 0\}_{j=1}^m$ over 0-1 assignments, is an algebraic circuit $C(\bar{x}, \bar{y}, \bar{z}) \in \mathbb{F}[\bar{x}, y_1, \dots, y_m, z_1, \dots, z_n]$ such that (the equalities in what follows stand for formal polynomial identities²):*

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

*The **size of the IPS proof** is the size of the circuit C . An IPS proof $C(\bar{x}, \bar{y}, \bar{z})$ of $1 = 0$ from $\{f_j(\bar{x}) = 0\}_{j \in [m]}$ is called an **IPS refutation** of $\{f_j(\bar{x}) = 0\}_{j \in [m]}$ (note that in this case it must*

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

hold that $\{f_j(\bar{x}) = 0\}_{j=1}^m$ have no common solutions in $\{0, 1\}^n$. If \widehat{C} (the polynomial computed by 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 [10]), 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 $\text{IPS}_{\text{LIN}'}$ refutation (following [9]). If $\widehat{C}(\bar{x}, \bar{y}, \bar{0})$ is of individual degree ≤ 1 in each x_j and y_i variables, while $\widehat{C}(\bar{x}, \bar{0}, \bar{z})$ is not necessarily multilinear, then this is a **multilinear** $\text{IPS}_{\text{LIN}'}$ 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 \widehat{C} is linear in \bar{y}, \bar{z} , and a depth- d $\text{IPS}_{\text{LIN}'}$ refutation if \widehat{C} is linear in \bar{y} , and depth- d multilinear $\text{IPS}_{\text{LIN}'}$ refutation if $\widehat{C}(\bar{x}, \bar{y}, \bar{0})$ is linear in \bar{x}, \bar{y} .

Notice that the definition above adds the equations $\{x_i^2 - x_i = 0\}_{i=1}^n$, called the **Boolean axioms** denoted $\bar{x}^2 - \bar{x}$, to the system $\{f_j(\bar{x}) = 0\}_{j=1}^m$. This allows to refute over $\{0, 1\}^n$ unsatisfiable systems of equations. The variables \bar{y}, \bar{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 \bar{y}, \bar{z} , which in turn, following the second equality, means that C witnesses the fact that 1 is in the ideal generated by $f_1(\bar{x}), \dots, f_m(\bar{x}), x_1^2 - x_1, \dots, x_n^2 - x_n$ (the existence of this witness, for unsatisfiable set of polynomials, stems from the Nullstellensatz theorem [4]).

In this work we focus on multilinear $\text{IPS}_{\text{LIN}'}$ refutations. This proof system is complete because its *weaker* subsystem multilinear-formula $\text{IPS}_{\text{LIN}'}$ was shown in [9, Corollary 4.12] to be complete (and to simulate Nullstellensatz with respect to sparsity by already depth-2 multilinear $\text{IPS}_{\text{LIN}'}$ proofs).

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

$$C(\bar{x}, \bar{y}, \bar{z}) = \sum_i g_i(\bar{x}) \cdot y_i + C'(\bar{x}, \bar{z}),$$

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

Important remark: Unlike the multilinear-formula $\text{IPS}_{\text{LIN}'}$ in [9], in multilinear $\text{IPS}_{\text{LIN}'}$ refutations $C(\bar{x}, \bar{y}, \bar{z})$ we do *not* require that the refutations are written as multilinear *formulas* or multilinear *circuits*, only that the *polynomial computed* by $C(\bar{x}, \bar{y}, \bar{0})$ is multilinear, hence the latter proof system easily simulates the former.

2.3 Set-Multilinear Monomials over a Word

We recall some notation from [18]. 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

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

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

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

³The P_w here is not to be confused with the canonical full-rank set-multilinear polynomial in [18] denoted as well by P_w mentioned in the introduction.

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

Given a word w , we associate with it a sequence $\overline{X}(w) = \langle X(w_1), \dots, X(w_d) \rangle$ 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 w if it is set-multilinear over the sequence $\overline{X}(w)$.

For a word w , let Π_w denote the projection onto the space $\mathbb{F}_{\text{sml}}[\overline{X}(w)]$ that maps the set-multilinear monomials over w identically to themselves and all other monomials to 0.

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

$$\text{relrk}_w(f) = \frac{\text{rank}(M_w(f))}{\sqrt{|M_w^P| \cdot |M_w^N|}}.$$

2.5 Monomial Orders

Finally we recall some basic notions related to monomial orders. For an in-depth introduction see [7]. A monomial order (in a polynomial ring $\mathbb{F}[X]$) is a well-order \leq on the set of all monomials that respects multiplication:

$$\text{if } m_1 \leq m_2, \text{ then } m_1 m_3 \leq m_2 m_3 \text{ for any } m_3.$$

It is not hard to see that any monomial order extends the submonomial relation: if $m_1 m_2 = m_3$ 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 $\text{LM}(f)$, is the highest monomial with respect to \leq that appears in f with a non-zero coefficient.

3 The Lower Bound

Our main theorem is as follows:

Theorem 3 (Main). *Let $n, \Delta \in \mathbb{N}_+$ with $\Delta \leq 1/4 \log \log \log n$, and assume that $\text{char}(\mathbb{F}) = 0$. Then any product-depth at most Δ multilinear IPS_{LIN} refutation of the subset sum variant $\sum_{i,j,k,\ell \in [n]} z_{ijkl} x_i x_j x_k x_\ell - \beta$ (for $\beta \notin \{0, \dots, n^4\}$) is of size at least $n^{(\log n)^{\exp(-O(\Delta))}}$.*

Note that, when $\Delta \geq \omega(\log \log \log n)$, the lower bound above becomes trivial. To prove [Theorem 3](#) we need the following theorem, which is proved in the sequel:

Theorem 4. *Let $n, \Delta \in \mathbb{N}_+$ with $\Delta \leq 1/4 \log \log \log n$, and assume that $\text{char}(\mathbb{F}) = 0$. Let f be the multilinear polynomial such that*

$$f = \frac{1}{\sum_{i,j,k,\ell \in [n]} z_{ijkl} x_i x_j x_k x_\ell - \beta} \quad \text{over Boolean assignments.}$$

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

$$n^{(\log n)^{\exp(-O(\Delta))}}.$$

Proof of Theorem 3 from Theorem 4. Let $C(\bar{x}, \bar{y}, \bar{w})$ be a multilinear $\text{IPS}_{\text{LIN}'}$ refutation of $\sum_{i,j,k,\ell \in [n]} z_{ijkl} x_i x_j x_k x_\ell - \beta = 0$ (where here for notational clarity we renamed the Boolean axioms placeholder variables from \bar{z} to \bar{w}). Since there is only one non-Boolean axiom, C has in fact only a single \bar{y} variable denoted y (i.e., $\bar{y} = \{y\}$). Note that $\widehat{C}(\bar{x}, y, \bar{0}) = g(\bar{x}) \cdot y$, for some polynomial $g(\bar{x}) \in \mathbb{F}[\bar{x}]$, because by assumption \widehat{C} is linear in the \bar{y} variables (and since $\widehat{C}(\bar{x}, 0, \bar{0}) = 0$). Therefore, $C(\bar{x}, 1, \bar{0})$ computes the polynomial $g(\bar{x})$. Thus, the minimal product-depth- Δ circuit-size of $g(\bar{x})$ lower bounds the minimal product-depth- Δ circuit-size of $C(\bar{x}, y, \bar{w})$. It remains to lower bound the size of product depth at most Δ circuits computing $g(\bar{x})$.

Notice first that

$$\widehat{C}(\bar{x}, y, \bar{w}) = \widehat{C}(\bar{x}, y, \bar{0}) + \sum_i h_i \cdot w_i,$$

for some polynomials h_i in \bar{x}, y, \bar{w} . Thus, $\widehat{C}(\bar{x}, y, \bar{w}) = g(\bar{x}) \cdot y + \sum_i h_i \cdot w_i$.

By definition of an IPS refutation

$$\widehat{C}\left(\bar{x}, \sum_{i,j,k,\ell \in [n]} z_{ijkl} x_i x_j x_k x_\ell - \beta, \bar{x}^2 - \bar{x}\right) = 1$$

and so

$$g(\bar{x}) \cdot \left(\sum_{i,j,k,\ell \in [n]} z_{ijkl} x_i x_j x_k x_\ell - \beta \right) + \sum_i (h_i \cdot (x_i^2 - x_i)) = 1.$$

From this we get that over the Boolean cube $g(\bar{x}) \cdot \sum_{i,j,k,\ell \in [n]} z_{ijkl} x_i x_j x_k x_\ell - \beta \equiv 1$ (as a function, not necessarily as a polynomial identity), and hence $g(\bar{x}) = \frac{1}{\sum_{i,j,k,\ell \in [n]} z_{ijkl} x_i x_j x_k x_\ell - \beta}$ over the Boolean cube. This shows that the size of $C(\bar{x}, y, \bar{w})$ must be at least $n^{(\log n)^{\exp(-O(\Delta))}}$ by Theorem 4. \square

The rest of section 3 is devoted to proving Theorem 4.

We shall use the tight degree lower bound proved in [9] for functions defined by $\hat{f}(\bar{x}) = 1/f(\bar{x})$ for simple polynomials $f(\bar{x})$. Specifically, we use the fact that any multilinear polynomial agreeing with $1/f(\bar{x})$, where $f(\bar{x})$ is the subset-sum axiom $\sum_{i=1}^n x_i - \beta$ (where β is such that the axiom has no Boolean roots) must have degree n . We note that a degree lower bound of $\lceil n/2 \rceil + 1$ was established by Impagliazzo, Pudlák, and Sgall [15]. However, similar to [9], we need the tight bound of n here as it will be used crucially in the proof of Lemma 6 to obtain the rank lower bound, which is a stronger notion than degree lower bound. Recall the notation $\bar{x}^2 - \bar{x}$ for the Boolean axioms of the \bar{x} -variables.

Lemma 5 (Proposition 5.3 in [9]). *Let $n \geq 1$ and \mathbb{F} be a field with $\text{char}(\mathbb{F}) > n$ (or $\text{char}(\mathbb{F}) = 0$). Suppose that $\beta \in \mathbb{F} \setminus \{0, \dots, n\}$. Let $f \in \mathbb{F}[x_1, \dots, x_n]$ be a multilinear polynomial such that*

$$f(\bar{x}) \left(\sum_i x_i - \beta \right) = 1 \pmod{\bar{x}^2 - \bar{x}}.$$

Then $\deg f = n$.

3.1 Subset-Sum Based on a Word w

In this section we define an auxiliary polynomial we use to prove [Theorem 4](#). It is a variant of the subset-sum that is defined from a given word w . Let $w \in \mathbb{Z}^d$ be arbitrary word, and consider the sequence $\overline{X}(w) = \langle X(w_1), \dots, X(w_d) \rangle$ of sets of variables. We fix now a 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 $\{0, 1\}$):

$$A_w^{(i)} := \left[\sum_{\substack{i' \in P_w \\ i' < i}} w_{i'} + 1, \sum_{\substack{i' \in P_w \\ i' \leq i}} w_{i'} \right].$$

Hence, the size of $A_w^{(i)}$ is precisely w_i and binary strings on the interval $A_w^{(i)}$ (i.e., $\{0, 1\}^{A_w^{(i)}}$), allows $2^{|A_w^{(i)}|} = 2^{w_i}$ possible strings, each corresponds to a different variable in $X(w_i)$.

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

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

We call the variables in $x_\sigma^{(i)}$ the *positive variables*, or simply \bar{x} -variables, and the variables $y_\sigma^{(j)}$ the *negative variables*, or simply \bar{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 $\bigcup_{j \in T} B_w^{(j)}$ 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 , and any 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 , and for any binary string σ indexed by A_w^S we denote by $m(\sigma)$ the monomial it defines, and similarly for strings and monomials on the negative variables. Thus notice that for a (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 is of degree d with each \bar{x} -variable picked uniquely from the $X(w_i)$ -variables, for $i \in P_w$ the positive indices in w , and each \bar{y} -variable is picked uniquely from the $X(w_j)$ -variables, for $j \in N_w$ the negative indices in w , and moreover the monomial corresponds to a binary string of length $\sum_{i=1}^d |w_i|$.

We call the word w **balanced** if for every $i \in P_w$ there is some $j \in N_w$ such that $A_w^{(i)} \cap B_w^{(j)} \neq \emptyset$ and for every $j \in N_w$ there is some $i \in P_w$ such that $A_w^{(i)} \cap B_w^{(j)} \neq \emptyset$. This means that any positive variable as defined above has some overlap with some negative variable within the given indexing scheme, and vice versa. [Figure 1](#) and [Figure 2](#) give examples of unbalanced and balanced words, respectively.

Our notion of a balanced word is different from, but related to, the notion of a b -unbiased word in [\[18\]](#): a word $w \in \mathbb{Z}^d$ is *b-unbiased* if $|w_{[t]}| \leq b$ for every $t \leq d$. If a balanced word $w \in \mathbb{Z}^d$ has all




Figure 1: An example of an *unbalanced* word $w \in \mathbb{Z}^d$, and where each w_i is written as a box with its corresponding binary string index set inside it, such that $|A_w^{(i)}| = w_i$ for $i \in P_w$ and $|B_w^{(j)}| = w_j$ for $j \in N_w$. A word is balanced if for every positive box $i \in P_w$, there exists some overlapping negative box, i.e., a $j \in N_w$ such that $A_w^{(i)} \cap B_w^{(j)} \neq \emptyset$, and vice versa. Here however, $B_w^{(8)} \cap (A_w^{(2)} \cup A_w^{(5)} \cup A_w^{(6)}) = \emptyset$.




Figure 2: An example of a balanced word. The construction of the knapsack polynomial \mathbf{ks}_w makes use of a balanced word w .

its entries bounded by b in absolute value, i.e., $|w_i| \leq b$ for every $i \in [d]$, then the sum of all entries of w is bounded by b in the sense that $|w_{[d]}| \leq b$. The notion of balance is more relaxed than that of unbiased in the sense that we do not need the property that all initial segments are also bounded by b . On the other hand the construction and proof in [18] does not require a balanced word, when this property is essential for us in the proof of Lemma 6.

We define the polynomial \mathbf{ks}_w as follows. Below we suppose that $|w_{N_w}| \geq |w_{P_w}|$, so that the negative monomials are determined by a longer binary string than the positive ones. Otherwise we flip the roles of the negative and positive variables in the definition below.

For a positive index $i \in P_w$ and $\sigma \in \{0, 1\}^{A_w^{(i)}}$, define the polynomial

$$f_\sigma^{(i)} := \prod_{\substack{j \in N_w: \\ A_w^{(i)} \cap B_w^{(j)} \neq \emptyset}} \sum_{\sigma_j \in \{0, 1\}^{B_w^{(j)}}} y_{\sigma_j}^{(j)}, \quad (2)$$

where the sum ranges over all those σ_j that agree with σ on $A_w^{(i)} \cap B_w^{(j)}$. The degree of $f_\sigma^{(i)}$ equals the number of those $B_w^{(j)}$ that overlap with $A_w^{(i)}$. Note that the degree is always at least 1. Now define the polynomial

$$\mathbf{ks}_w := \sum_{i \in P_w} \sum_{\sigma \in \{0, 1\}^{A_w^{(i)}}} x_\sigma^{(i)} f_\sigma^{(i)} - \beta,$$

where β is any field \mathbb{F} element so that the polynomial \mathbf{ks}_w has no Boolean roots (here we use the fact that $\text{char}(\mathbb{F}) = 0$). Figure 3 gives an example of the construction.

The basic idea behind the above construction is simple. Given a monomial m that is set-multilinear over $w|_{N_w}$, consider the partial assignment τ_m to the negative variables that sends any variable in m to 1 and all other negative variables to 0. Now after applying τ_m to \mathbf{ks}_w we are left with a simple subset sum instance

$$\sum_{i \in P_w} x_{\sigma_i}^{(i)} - \beta,$$

where σ_i is the binary string indexed by $A_w^{(i)}$ that agrees with $\sigma(m)$ on $A_w^{(i)}$. Similarly for any monomial that is set-multilinear over $w|_T$ for some $T \subseteq N_w$, the polynomial $\tau_m(\mathbf{ks}_w)$ is the subset




Figure 3: Here $*$ represents either 0 or 1. In the construction of the polynomial ks_w , for $i = 2$ and $\sigma = 011001$, we have $f_{011001}^{(2)} = y_{011}^{(1)} \cdot y_{00}^{(3)} \cdot (y_{1000}^{(4)} + y_{1001}^{(4)} + \dots + y_{1111}^{(4)})$.

sum instance

$$\sum_{i \in S} x_{\sigma_i}^{(i)} - \beta,$$

where S is the maximal subset of P_w such that $A_w^S \subseteq B_w^T$. With [Lemma 5](#) we have a very good understanding of the multilinear IPS refutations of such subset sum instances, namely we know that the multilinear polynomial f that equals

$$\frac{1}{\sum_{i \in S} x_{\sigma_i}^{(i)} - \beta} \text{ over Boolean assignments}$$

has as its leading monomials the product $\prod_{i \in S} x_{\sigma_i}^{(i)}$. This observation allows us to prove our rank lower bound in the following section. [Figure 5](#) in [subsection 3.2](#) illustrates the way assignments to the negative variables give rise to simple subset sum instances.

3.2 Rank Lower Bound Lemma

In this section we prove the main technical lemma of this paper – a rank lower bound for subset sum over any balanced word. Let $w \in \mathbb{Z}^d$ be any word, and let f be a multilinear polynomial in the variables $\bar{X}(w)$. Denote by $M(f)$ the coefficient matrix of f with rows indexed by all multilinear monomials (of *any* degree) in the positive variables and columns indexed by all multilinear monomials (again, of any degree) in the negative variables, and denote by $M_w(f)$ its *submatrix* with rows indexed by all monomials that are set-multilinear over $w|_{P_w}$ and columns indexed by all monomials that are set-multilinear over $w|_{N_w}$.

Lemma 6. *Let $w \in \mathbb{Z}^d$ be a balanced word, and let f be the multi-linear polynomial so that*

$$f = \frac{1}{\text{ks}_w} \text{ over Boolean assignments.}$$

Then $M_w(f)$ is full-rank.

Proof. Without loss of generality we assume that $|w_{N_w}| \geq |w_{P_w}|$, so that our notation matches that of [Section 3.1](#). Write

$$f = \sum_m g_m(x)m, \tag{3}$$

where the sum ranges over all multilinear monomials m in the y -variables and $g_m(x)$ is some multilinear polynomial in the x -variables (note that here we include all multilinear monomials and

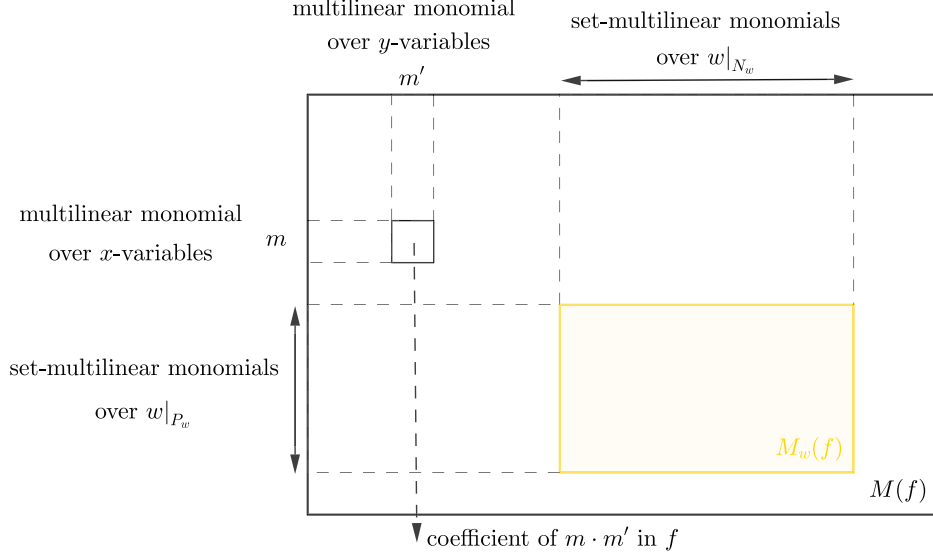


Figure 4: The coefficient matrix $M(f)$ with rows and columns indexed by the multilinear monomials over the x -variables and y -variables respectively; and its submatrix $M_w(f)$ with rows and columns indexed by the set-multilinear monomials over $w|_{P_w}$ and $w|_{N_w}$ respectively. If w is a balanced word and f is a multilinear polynomial agreeing with $1/k_{S_w}$ over boolean assignments, then $M_w(f)$ has full-rank

not only the set-multilinear ones). We show that for any m that is set-multilinear on $w|_{N_w}$ the leading monomial of $g_m(x)$ is the set-multilinear monomial $m(\sigma(m)|_{A_w^{P_w}})$. This is where we focus on the submatrix of set-multilinear monomials within the bigger matrix of all multilinear monomials, as depicted in Figure 4. To prove this we need the following claim.

Claim 7. *For any monomial m set-multilinear on some $w|_T$, where $T \subseteq N_w$, the leading monomial of $g_m(x)$ is less or equal to*

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

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 $g_m(x)$ equals

$$m(\sigma(m)|_{A_w^{P_w}}).$$

Proof of Claim. Proof by induction on the size of T .

Base case: If $T = \emptyset$, the only monomial set-multilinear on $w|_\emptyset$ is the empty monomial 1. Now consider the partial assignment τ_1 that maps all the \bar{y} -variables to 0. Now $\tau_1(f) = g_1(x)$, where $g_1(x)$ is the coefficient of the empty monomial 1. On the other hand, since

$$f = \frac{1}{k_{S_w}} \quad \text{over Boolean assignments,}$$

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

Induction step: Suppose then that T is non-empty, and let m be a set-multilinear monomial on




Figure 5: In this example, $T = \{1, 4, 7, 8\} \subseteq N_w$ and $m = y_{100}^{(1)} \cdot y_{1001}^{(4)} \cdot y_{0110}^{(7)} \cdot y_{11}^{(8)}$ is a set-multilinear monomial over $w|_T$. As $S = \{5, 6\}$ is the maximal subset of P_w such that $A_w^S \subseteq B_w^T$, we therefore have that the leading monomial of $g_m(x)$ (from Equation 3) is less than or equal to $x_{00}^{(5)} \cdot x_{101101}^{(6)}$. Moreover, in the polynomial \mathbf{ks}_w , the partial assignment setting the y -variables in m to 1 and the remaining y -variables to 0 results in the polynomial $x_{00}^{(5)} + x_{101101}^{(6)} - \beta$.

$w|_T$. Now consider the partial assignment τ_m that maps any variable in m to 1 and any other \bar{y} -variable to 0. By Equation 3

$$\tau_m(f) = \sum_{m'} g_{m'}(x), \quad (4)$$

where m' ranges over all submonomials of m . On the other hand, by the construction of \mathbf{ks}_w

$$\tau_m(f) = \frac{1}{\sum x_\sigma^{(i)} - \beta} \quad \text{over Boolean assignments,}$$

where the sum in the denominator ranges over those i and σ such that $A_w^{(i)} \subseteq B_w^T$ and σ agrees with $\sigma(m)$ on the interval $A_w^{(i)}$. Note that for any $i \in P_w$ there is at most one σ such that $x_\sigma^{(i)}$ appears in the sum (because $A_w^{(i)} \subseteq B_w^T$, and by construction of \mathbf{ks}_w , $x_\sigma^{(i)}$ for a fixed σ is multiplied in \mathbf{ks}_w by all products $\prod_j y_{\rho_j}^{(j)}$ such that the concatenation of the ρ_j 's extends the string σ ; but since we assigned 0-1 to all the \bar{y} -variables there is a *single* such concatenation, induced by our assignment of 1's; see Figure 3).

It follows by Lemma 5 that the leading monomial of $\tau_m(f)$ is the product of all the $x_\sigma^{(i)}$ appearing in the sum in the denominator above, and thus the leading monomial equals

$$m\left(\sigma(m)|_{A_w^S}\right), \quad (5)$$

where S is the maximal subset of P_w such that $A_w^S \subseteq B_w^T$. By Equation 4, this means that either the leading monomial of $g_m(x)$ is less than or equal to (5), or otherwise is greater than (5) but is cancelled out in (4) by some monomial in $g_{m'}(x)$ for m' a proper submonomial of m . But by induction assumption, for a *proper* submonomial m' of m with $T' \subsetneq T$ and m' set-multilinear on $w|_{T'}$, the leading monomial of $g_{m'}(x)$ is less or equal to $m\left(\sigma(m')|_{A_w^{S'}}\right)$, where S' is the maximal subset of P_w such that $A_w^{S'} \subseteq B_w^{T'}$. Since, m' is a submonomial of m the monomial $m\left(\sigma(m')|_{A_w^{S'}}\right)$ is less than or equal to $m\left(\sigma(m)|_{A_w^S}\right)$, and so the above mentioned cancellation cannot occur, and we conclude that the leading monomial of $g_m(x)$ is less than or equal to (5).

It remains to show that the leading monomial of $g_m(x)$ is *precisely* (5), when m is a set-multilinear monomial on $w|_{N_w}$. Let m' be a proper submonomial of m that is set-multilinear over $w|_T$ for some $T \subsetneq N_w$. By the assumption that w is balanced there is some $i \in P_w$ such

that $A_w^{(i)} \not\subseteq B_w^T$, and thus the leading monomial of $g_{m'}(x)$ is properly smaller than $m(\sigma(m)|_{A_w^{P_w}})$. Hence, by [Equation 5](#) (and the sentence preceding it) the leading monomial of $g_m(x)$ must equal $m(\sigma(m)|_{A_w^{P_w}})$. \square

By the claim for any m that is set-multilinear on $w|_{N_w}$ the leading monomial of $g_m(x)$ is the monomial $m(\sigma(m)|_{A_w^{P_w}})$, and these include all the set-multilinear monomials on $w|_{P_w}$. Thus the column space of $M_w(f)$ spans the space of all set-multilinear polynomials on $w|_{P_w}$, and $M_w(f)$ is full-rank (where a column in $M_w(f)$ determines a polynomial that is a linear combination of the positive monomials in its rows). \square

Corollary 8. *Let $w \in \mathbb{Z}^d$ be a balanced word with $|w_i| \leq b$ for all $i \in [d]$, and let f be the multi-linear polynomial so that*

$$f = \frac{1}{\text{ks}_w} \text{ over Boolean assignments.}$$

Then $\text{relrk}_w(f) \geq 2^{-b/2}$.

Proof. Assume again without a loss of generality that $|w_{N_w}| \geq |w_{P_w}|$. By the ‘‘balanced-ness’’ and the assumption that $|w_i| \leq b$ for all $i \in [d]$, we know that $|w_{P_w}| - |w_{N_w}| \geq -b$. By [Lemma 6](#), $M_w(f)$ is of rank $|M_w^P|$, and so

$$\text{relrk}_w(f) = \sqrt{\frac{|M_w^P|}{|M_w^N|}} = \sqrt{2^{|w_{P_w}| - |w_{N_w}|}} \geq 2^{-b/2}.$$

\square

3.3 Lower Bound for Constant-Depth Set-Multilinear Circuits

In this section we prove the following lower bound on bounded-depth set-multilinear circuits.

Lemma 9. *Let $d, k, \Delta \in \mathbb{N}_+$ with $k \geq 10d$. Let $w \in \mathbb{Z}^d$ be a balanced word in the vocabulary $\{\lfloor \alpha k \rfloor, -k\}$, where $\alpha = 1/\sqrt{2}$, and let f be the multilinear polynomial which equals $1/\text{ks}_w$ over Boolean assignments. Then any set-multilinear circuit of product-depth Δ computing the set-multilinear projection $\Pi_w(f)$ has size at least*

$$2^k \binom{\frac{d^{1/(2^\Delta-1)} - 20}{40\Delta}}{k}$$

Proof. We prove the lemma by using the following claim from [\[18\]](#).

Claim 10 ([\[18\]](#) Claim 16). *Let $k \geq 10d$. Let w be any word of length d with entries in $\{\lfloor \alpha k \rfloor, -k\}$, where $\alpha = 1/\sqrt{2}$. Then for any $\Delta \geq 1$, any set-multilinear formula C of product-depth Δ of size at most s satisfies*

$$\text{relrk}_w(C) \leq s \cdot 2^{\frac{-kd^{1/(2^\Delta-1)}}{20}}$$

Let C be a set-multilinear circuit of size s and product-depth Δ computing $\Pi_w(f)$. We can transform C into a set-multilinear formula F of size $s^{2\Delta}$ and product-depth Δ computing $\Pi_w(f)$.

Now by [Lemma 6](#) and [Claim 10](#), we have that

$$2^{-k} \leq \text{relrk}_w(\Pi_w(f)) \leq s^{2\Delta} \cdot 2^{-\frac{k d^{1/(2^\Delta-1)}}{20}},$$

and thus

$$s^{2\Delta} \geq 2^k \left(\frac{d^{1/(2^\Delta-1)} - 20}{20} \right),$$

from which the claim follows. \square

3.4 Lower Bound for Constant-Depth Circuits

Finally in this section we prove [Theorem 4](#). To prove this theorem we reduce the task of computing f in [Theorem 4](#) to computing the set-multilinear projection $\Pi_w(f')$ of the multilinear f' that equals $1/\text{ks}_w$ over the Boolean values for a suitable word w . For this we require the following lemma which can be proved in a manner similar to [Proposition 9](#) in [[18](#)].

Lemma 11. *Let s, N, d and Δ be growing parameters with $s \geq Nd$. Assume that $\text{char}(\mathbb{F}) = 0$ or $\text{char}(\mathbb{F}) > d$. Let C be a circuit of size at most s and product-depth at most Δ computing a polynomial f . Let $\bar{X} = \langle X_1, \dots, X_d \rangle$ be a sequence of pairwise disjoint sets of variables, each of size at most N . Then there is a set-multilinear circuit \tilde{C} of size at most $d^{O(d)} \text{poly}(s)$ and product-depth at most 2Δ computing the set-multilinear projection $\Pi_{\bar{X}}(f)$ of f .*

With this lemma at hand, we are ready to prove [Theorem 4](#).

Proof of [Theorem 4](#). Let C be a circuit of size at most $s \geq n$ and product-depth Δ computing f . Now let $k = \lfloor \log n/2 \rfloor$ and $d = \lfloor \log n/25 \rfloor$, and note that $d2^k < n \leq s$ for large enough n . Note also that $k \geq 10d$ for large enough n . Let w be a balanced word of length d on the alphabet $\{\lfloor \alpha k \rfloor, -k\}$. One can easily construct such a word by induction on d .

Note that the polynomial ks_w is now of degree at most 4, as any $B_w^{(j)}$ overlaps at most 3 different $A_w^{(i)}$'s and any $A_w^{(i)}$ overlaps at most 2 different $B_w^{(j)}$'s. Also, by choice of parameters, ks_w involves less than n many variables. Hence there is some partial assignment τ_w to the variables $\{z_{ijkl}, x_i : i, j, k, \ell \in [n]\}$ that maps $\sum_{i,j,k,\ell \in [n]} z_{ijkl} x_i x_j x_k x_\ell - \beta$ to the polynomial ks_w (up to renaming of variables). By applying this partial mapping to C , we obtain a circuit C' of size at most s and product-depth Δ that computes the multilinear f' that equals $1/\text{ks}_w$ over Boolean assignments. Now, by [Lemma 11](#), there is a set-multilinear circuit C' of size $d^{O(d)} \text{poly}(s)$ and product-depth 2Δ computing the set-multilinear projection $\Pi_w(f')$ of f' .

By [Lemma 9](#) any set-multilinear circuit of product-depth 2Δ computing $\Pi_w(f')$ has size at least

$$2^k \left(\frac{d^{1/(2^{2\Delta}-1)} - 20}{80\Delta} \right) \geq n \frac{d^{1/(2^{2\Delta}-1)} - 20}{200\Delta}.$$

Putting everything together we have that

$$d^{O(d)} \text{poly}(s) \geq n \frac{d^{1/(2^{2\Delta}-1)} - 20}{200\Delta}.$$

Now, given that $\Delta \leq 1/4 \log \log \log n$, by the choice of d , we have that $2\Delta \leq 1/2 \log \log 30d$. Then for large enough n we have that

$$n^{\frac{d^{1/(2^{2\Delta}-1)}-20}{200\Delta}} \geq n^{\frac{2^{\sqrt{\log(d/30)}/2}}{50 \log \log 30d}},$$

and as

$$n^{\frac{2^{\sqrt{\log(d/30)}/2}}{50 \log \log 30d}} \geq d^{\omega(d)},$$

the lower bound follows. □

Comment: We remark that if we make sure that the word w above leans towards the negative monomials, meaning that $|w_{N_w}| \geq |w_{P_w}|$, we can actually prove the lower bound for the following degree-4 variant of subset-sum

$$\sum_{i,j,k \in [n]} z_{ijk} x_i x_j x_k - \beta.$$

We have however opted for the proof above for simplicity, as it works for any balanced word over the vocabulary $\{[\alpha k], -k\}$ and does not involve any set-up of a suitable word that could distract from the main idea.

4 Conclusions and Open Problems

The main goal of this work is to advance on the frontiers of strong propositional proofs lower bounds. We provide the first lower bounds against algebraic proof systems operating with constant-depth circuits, where the hard instance is computable itself with small constant-depth circuits. Our hard instances are combinatorial, simple, and have 0-1 coefficients, and the lower bounds work in the unit-cost model of algebraic circuits, namely, where the size does not depend on the magnitude of coefficients used in the polynomials appearing in the proof

Thus, our result brings us to the natural and standard setting of proof complexity lower bounds, while coming closer to CNF hard instances (since the magnitude of coefficients in the hard instances do not play a role in our lower bound proofs).

On the other hand, establishing lower bounds against CNF formulas in strong algebraic proof systems stays a remarkable open problem, since it necessitates a lower bound technique that is different from the functional lower bound approach we used or the bit-complexity/large-coefficients approach of [1, 2] (or at least a substantially modified technique than those two techniques). Note that such lower bounds would also imply constant-depth Frege with counting gates ($AC^0[p]$ -Frege) lower bounds, which is an important long-standing open problem in proof complexity. This leads us to the following set of *open problems*:

1. **CNF hard instances:** Can we establish lower bounds against strong algebraic proof systems, and specifically constant-depth IPS proofs for a family of CNF formulas? As mentioned above, this is a very challenging problem, or at least one with important consequences in proof complexity.

Before tackling this difficult open problem, we identify several possibly less challenging ones that seem to be prerequisites for solving open problem 1 (at least as far as taking the proof complexity lower bound approach in the current paper):

2. **Finite fields:** Can we establish lower bounds against strong algebraic proof systems, and specifically constant-depth IPS proofs over finite fields? Both the functional lower bound argument from Forbes *et al.* [9] and the Limaye *et al.* [18] technique, use the fact that the fields are sufficiently big, or have characteristic 0. For the former approach, characteristic 0 fields are essential: first, the subset sum instance is not necessarily unsatisfiable over finite fields. But more crucially, the whole functional lower bound approach for IPS hinges on lower bounding the size of a circuit computing the function $1/f$ over the Boolean cube, for some efficiently computable polynomial f . However, over finite fields $1/f$ is efficiently computed, when f is, over the Boolean cube. For the latter [18] approach, large fields do not seem to be as crucial to the argument (it is used only in the homogenization procedure to yield low-depth circuits, using polynomial interpolation, the latter requires a sufficiently large field; this homogenization is based on a generalization of Shpilka and Wigderson [26]). However, there may be a different way to homogenize constant-depth circuits without increasing too much the depth even over finite fields.
3. **No multilinear requirement:** Can we get rid of the requirement for multilinearity of the IPS refutations in our lower bounds? Namely, can we use a stronger proof system than multilinear IPS_{LIN} ? We discussed this requirement in the introduction. It comes from the requirement in the Limaye *et al.* [18] technique to consider set-multilinear polynomials, as well as the use of the functional lower bound approach from [9] which focuses on functions computed on the Boolean cube alone. A hard set-multilinear polynomial can compute the same function over the Boolean cube as a polynomial whose set-multilinear projection is in fact zero (and hence easy to compute), which breaks our argument. It is unclear at the moment how to overcome this obstacle. This leads us to the following interesting question in algebraic circuit complexity proper:
4. **Functional lower bounds for constant-depth circuits:** Can we prove a lower bound on the size of every constant-depth circuit computing a certain *function* (in contrast to a certain specific polynomial; this is the difference between “semantic” lower bounds, and the weaker notion of “syntactic” lower bounds in algebraic circuit complexity)?

References

- [1] Yaroslav Alekseev. A lower bound for polynomial calculus with extension rule. In Valentine Kabanets, editor, *36th Computational Complexity Conference, CCC 2021, July 20-23, 2021, Toronto, Ontario, Canada (Virtual Conference)*, volume 200 of *LIPICs*, pages 21:1–21:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. [1.1](#), [1.2](#), [1](#), [4](#)
- [2] 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*, pages 54–67. ACM, 2020. [1.2](#), [4](#)
- [3] 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. [1.1](#), [1.2](#), [1.4](#)

- [4] Paul Beame, Russell Impagliazzo, Jan Krajíček, Toniann Pitassi, and Pavel Pudlák. Lower bounds on Hilbert’s Nullstellensatz and propositional proofs. *Proc. London Math. Soc. (3)*, 73(1):1–26, 1996. [1.1](#), [1.2](#), [2.2](#), [2.2](#)
- [5] Samuel R. Buss, Russell Impagliazzo, Jan Krajíček, Pavel Pudlák, Alexander A. Razborov, and Jiří Sgall. Proof complexity in algebraic systems and bounded depth Frege systems with modular counting. *Computational Complexity*, 6(3):256–298, 1996. [1.2](#)
- [6] 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)*, pages 174–183, New York, 1996. ACM. [1.1](#)
- [7] David Cox, John Little, and Donal O’Shea. *Ideals, varieties, and algorithms*. Undergraduate Texts in Mathematics. Springer, New York, third edition, 2007. An introduction to computational algebraic geometry and commutative algebra. [2.5](#)
- [8] Michael A. Forbes, Mrinal Kumar, and Ramprasad Saptharishi. Functional lower bounds for arithmetic circuits and connections to boolean circuit complexity. In Ran Raz, editor, *31st Conference on Computational Complexity, CCC 2016, May 29 to June 1, 2016, Tokyo, Japan*, volume 50 of *LIPICs*, pages 33:1–33:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. [1.4](#)
- [9] Michael A. Forbes, Amir Shpilka, Iddo Zameret, and Avi Wigderson. Proof complexity lower bounds from algebraic circuit complexity. *Theory Comput.*, 17:1–88, 2021. [1.1](#), [1.2](#), [1.3](#), [1.3](#), [1.4](#), [2.2](#), [2](#), [2.2](#), [3](#), [5](#), [2](#), [3](#)
- [10] Mika Göös and Toniann Pitassi. Communication lower bounds via critical block sensitivity. *SIAM J. Comput.*, 47(5):1778–1806, 2018. [1.1](#), [2](#)
- [11] Dima Grigoriev and Edward A. Hirsch. Algebraic proof systems over formulas. *Theoret. Comput. Sci.*, 303(1):83–102, 2003. Logic and complexity in computer science (Créteil, 2001). [1.1](#)
- [12] Joshua A. Grochow and Toniann Pitassi. Circuit complexity, proof complexity, and polynomial identity testing: The ideal proof system. *J. ACM*, 65(6):37:1–37:59, 2018. [2.2](#), [2](#)
- [13] Russell Impagliazzo, Sasank Mouli, and Toniann Pitassi. The surprising power of constant depth algebraic proofs. In Holger Hermanns, Lijun Zhang, Naoki Kobayashi, and Dale Miller, editors, *LICS ’20: 35th Annual ACM/IEEE Symposium on Logic in Computer Science, Saarbrücken, Germany, July 8-11, 2020*, pages 591–603. ACM, 2020. [1.1](#)
- [14] Russell Impagliazzo, Sasank Mouli, and Toniann Pitassi. Lower bounds for polynomial calculus with extension variables over finite fields. *Electron. Colloquium Comput. Complex.*, 2022. [1.1](#), [1.2](#)
- [15] Russell Impagliazzo, Pavel Pudlák, and Jiří Sgall. Lower bounds for the polynomial calculus and the Gröbner basis algorithm. *Computational Complexity*, 8(2):127–144, 1999. [1.2](#), [1.3](#), [3](#)

- [16] Alexander Knop. Ips-like proof systems based on binary decision diagrams. *Electron. Colloquium Comput. Complex.*, page 179, 2017. [1.1](#)
- [17] Fu Li, Iddo Tzameret, and Zhengyu Wang. Characterizing propositional proofs as noncommutative formulas. In *SIAM Journal on Computing*, volume 47, pages 1424–1462, 2018. Full Version: <http://arxiv.org/abs/1412.8746>. [1.1](#)
- [18] 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*, pages 804–814. IEEE, 2021. [1.4](#), [2.3](#), [3](#), [3.1](#), [3.3](#), [10](#), [3.4](#), [2](#), [3](#)
- [19] Toniann Pitassi. Algebraic propositional proof systems. In *Descriptive complexity and finite models (Princeton, NJ, 1996)*, volume 31 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 215–244. Amer. Math. Soc., Providence, RI, 1997. [1.1](#)
- [20] Toniann Pitassi. Unsolvability systems of equations and proof complexity. In *Proceedings of the International Congress of Mathematicians, Vol. III (Berlin, 1998)*, number Vol. III, pages 451–460, 1998. [1.1](#)
- [21] Toniann Pitassi and Iddo Tzameret. Algebraic proof complexity: Progress, frontiers and challenges. *ACM SIGLOG News*, 3(3), 2016. [2.2](#)
- [22] Ran Raz and Iddo Tzameret. Resolution over linear equations and multilinear proofs. *Ann. Pure Appl. Logic*, 155(3):194–224, 2008. [1.1](#)
- [23] Ran Raz and Iddo Tzameret. The strength of multilinear proofs. *Computational Complexity*, 17(3):407–457, 2008. [1.1](#)
- [24] Alexander A. Razborov. Lower bounds for the polynomial calculus. *Comput. Complexity*, 7(4):291–324, 1998. [1.2](#), [1.3](#)
- [25] Ramprasad Saptharishi. A survey of lower bounds in arithmetic circuit complexity, 2016-2022. <https://github.com/dasarpmar/lowerbounds-survey/releases>. [2.1](#)
- [26] Amir Shpilka and Avi Wigderson. Depth-3 arithmetic circuits over fields of characteristic zero. *Comput. Complexity*, 10:1–27, 2001. [2](#)
- [27] Amir Shpilka and Amir Yehudayoff. Arithmetic circuits: A survey of recent results and open questions. *Foundations and Trends in Theoretical Computer Science*, 5(3-4):207–388, 2010. [2.1](#)
- [28] Iddo Tzameret. Algebraic proofs over noncommutative formulas. *Inf. Comput.*, 209(10):1269–1292, 2011. [1.1](#)

— Page left blank for ECCC stamp —