

The Weak Rank Principle: Lower Bounds and Applications*

Michal Garlík[†]
Imperial College London

Svyatoslav Gryaznov[‡]
Imperial College London

Hanlin Ren[§]
Institute for Advanced Study

Iddo Tzameret[¶]
Imperial College London

Abstract

Given two symbolic matrices X and Y of dimensions $m \times n$ and $n \times m$, respectively, the *rank principle* states that when $m = n + 1$ and A is a scalar matrix of rank $n + 1$, the equation $XY = A$ is unsatisfiable. When m is arbitrarily larger than n and A has rank exceeding n , we obtain the *weak rank principle*. We study this principle as an algebraic generalisation of the weak pigeonhole principle (WPHP), asserting that m pigeons cannot be injected into n holes. As a strengthening of WPHP, it admits proof complexity lower bounds in settings where none are known for WPHP, while still supporting analogous applications. Using new generalised random restrictions applied to the weak rank principle, which may be of independent interest, we resolve several open problems in proof complexity: we construct proof complexity generators for Polynomial Calculus Resolution over the two-element field ($\text{PCR}_{\mathbb{F}_2}$), new generators for Sherali–Adams (SA), and hardness results for circuit lower bound statements against $\text{PCR}_{\mathbb{F}_2}$.

Generators for $\text{PCR}_{\mathbb{F}_2}$: We prove exponential size lower bounds for several encodings—both algebraic and CNF—of the weak rank principle in PCR over $\mathbb{F}_2$, where no such bounds are known for the WPHP in the regime with arbitrarily many pigeons. In particular, we obtain $2^{\Omega(n)}$ size lower bounds for both algebraic and standard CNF encodings, including the *bamboo-tree encoding*, which is the most relevant for applications and corresponds to a circuit encoding, as considered by Alekhovich, Ben-Sasson, Razborov, and Wigderson (*SIAM J. Comput.*, 2004) and Razborov (*Ann. Math.*, 2015). Our bounds hold for every matrix A in $XY = A$, implying that the rank principle forms a proof complexity generator with nearly quadratic stretch. Using a standard iteration technique we amplify the stretch to $2^{n^{\Omega(1)}}$, thereby obtaining a function generator. This resolves the open problem posed by Alekhovich *et al.* (*SIAM J. Comput.*, 2004) and Razborov (*Ann. Math.*, 2015) concerning the construction of proof complexity generators with good stretch for $\text{PCR}_{\mathbb{F}_2}$.

Generators for Sherali–Adams: Since in SA even the *strong* pigeonhole principle is easy, we develop a new size lower-bound technique showing that the weak rank principle, encoded as a bamboo-tree CNF, serves as a proof complexity generator for SA. Our method introduces a new relaxed notion of degree and a corresponding pseudoexpectation tailored specifically to the rank principle (and incompatible with the pigeonhole principle).

*Preliminary version of this work appears in *58th Annual ACM Symposium on Theory of Computing (STOC) 2026*, Salt Lake City, Utah, June 2026. 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, EPRICOT project). It was also supported by the Engineering and Physical Sciences Research Council (EPSRC) under grant EP/Z534158/1, Integrated Approach to Computational Complexity: Structure, Self-Reference and Lower Bounds.

[†]Email: michal.garlik@gmail.com

[‡]Email: svyatoslav.i.gryaznov@gmail.com

[§]Email: h4n1in.r3n@gmail.com

[¶]Email: iddo.tzameret@gmail.com

Circuit lower bound formulas: We show that $\text{PCR}_{\mathbb{F}_2}$ does not admit short proofs of lower-bound statements against Boolean circuits, nor against weak models of algebraic circuits. This settles the open problem raised by Razborov (*Ann. Math.*, 2015) concerning the provability of such lower bounds in $\text{PCR}_{\mathbb{F}_2}$.

Strength of the weak rank principle: Finally, we show that the weak rank principle is *necessary* for proving NC^2 circuit lower bounds and, for odd primes p , *sufficient* within the theory corresponding to $\text{AC}^0[p]$ for deriving $\text{AC}^0[p]$ lower bounds.

Contents

1	Introduction	3
2	Our Results	6
2.1	Lower Bounds and Generators for $\text{PCR}_{\mathbb{F}_2}$	6
2.2	Generators for Sherali–Adams	13
2.3	Hardness of Circuit Lower Bounds from the Weak Rank Principle	14
2.4	On the Strength of WRank	15
3	Conclusions and Open Problems	16
4	Preliminaries	17
5	Proof Complexity Generators for $\text{PCR}_{\mathbb{F}_2}$	24
5.1	Algebraic Rank Principle	24
5.2	Perfect Matching Encoding of the Rank Principle	27
5.3	Bamboo-Tree Encoding of the Rank Principle	37
5.4	Iterated Generators	47
5.4.1	Iterability of $\text{Rank}_n^m(A)$	48
5.4.2	Iterability of $\text{BTRank}_n^m(A)$	52
6	Proof Complexity Generators for SA	68
6.1	Row Degree Lower Bound	69
6.2	Size to Row Degree Reduction	71
7	Structural Results	75
7.1	Amplification of the Rank Principle	75
7.2	Low-Degree Reductions for WRank in Strong Proof Systems	78
8	Hardness of Proving Circuit Lower Bounds	80
8.1	Lower Bounds for Non-commutative Algebraic Branching Programs	81
8.2	Boolean Circuit Lower Bounds	83
8.3	Lower Bounds for $\text{PCR}_{\mathbb{F}_2}$	86
8.3.1	Lower Bounds for Non-commutative ABPs	86
8.3.2	Lower Bounds for Boolean Circuits	87
9	Proving Circuit Lower Bounds via the Weak Rank Principle	89
9.1	Smolensky’s Lower Bound	90
9.1.1	Correlation Bound Against Low-Degree Polynomials	90
9.1.2	Approximating $\text{AC}^0[p]$ Circuits by Low-Degree Polynomials	93
9.2	Formalisation in $\overline{\mathbf{V}}^0(p) + \text{WRank}_p$	96

9.2.1	The Theory $\overline{\mathbf{V}^0(p)} + \text{WRank}_p$	96
9.2.2	Smolensky’s Lower Bound in $\overline{\mathbf{V}^0(p)} + \text{WRank}_p$	97
9.2.3	Proof of Theorem 9.9	99
A Expander Properties		100
References		102

1 Introduction

Proof complexity provides both a concrete and conceptual framework for studying computational lower bounds. On the one hand, it seeks to develop combinatorial and algebraic techniques for proving unconditional lower bounds on the lengths of proofs, with relations to complexity class separations. On the other hand, it offers a setting for formulating and exploring metamathematical questions—for example, which proof systems can efficiently establish which lower bounds in computational complexity theory.

For concrete proof-size lower bound questions, it remains a long-standing open problem to establish superpolynomial size lower bounds against sufficiently strong proof systems, such as textbook propositional logic (i.e., Frege systems). Such lower bounds are unknown even against significantly weaker fragments, such as those operating with $\text{AC}^0[2]$ circuits. Nonetheless, the development of theoretical frameworks and techniques aimed at tackling such lower bounds remains a highly active and vibrant area of research.

The Weak Pigeonhole Principle. A prominent example illustrating the broad reach of proof complexity is the *pigeonhole principle* (PHP for short), which is the (suitably encoded) statement that $m = n + 1$ pigeons cannot be mapped into n holes if each hole can accommodate at most one pigeon. The *weak pigeonhole principle* (WPHP) refers to the case where the number of pigeons is possibly much larger than the number of holes; namely, it is the collection of such statements for all pairs $m > n$, where m may be arbitrarily larger than n . In this regard, WPHP is logically weaker than PHP where $m = n + 1$.

The pigeonhole principle is perhaps the most influential example in proof complexity, serving as a driving force behind many lower bound techniques and frameworks. These include bottleneck counting [36], pigeonhole switching lemmas and k -evaluations [2, 53, 66], pigeon dance [41, 72], pseudo-width [68, 74], etc. As for upper bounds, already TC^0 -Frege admits polynomial-size proofs of PHP (and hence WPHP); see [16, 20]. Despite this progress, a notable open problem posed by Razborov [76] remains unresolved: proving size lower bounds for WPHP against the polynomial calculus resolution (PCR) proof system (see [Footnote 1](#) and discussion therein).

Beyond concrete lower bounds, the pigeonhole principle plays a central role in the development of theories in bounded arithmetic, which study the computational complexity of the concepts required to prove various statements within formal systems. In bounded arithmetic, the weak pigeonhole principle serves as an axiom from which important results—particularly those related to randomness in computation—can be derived. As early as 1981, Woods [85] observed that explicit counting of the number of elements in a finite set can often be replaced by applications of the pigeonhole principle for bounded formulas. Building on this idea, Paris, Wilkie, and Woods [64] introduced the weak pigeonhole principle, and showed that it often serves as a more suitable substitute for PHP in such contexts and that it is provable in bounded arithmetic T_2 . Initial work by Wilkie (unpublished; see Krajíček [50, Theorem 7.3.7]), further developed by Thapen [82] and systematically pursued by Jeřábek [44–46], established WPHP as a useful axiom for reasoning about randomized computation.

WPHP and proof complexity generators. A concept closely related to the weak pigeonhole

principle is that of a proof complexity generator. In fact, it is more directly connected to the *dual* weak pigeonhole principle (denoted dWPHP), which states that when n pigeons are mapped to $2n$ holes, at least one hole must remain empty. The notion of proof complexity generators was introduced independently by Alekhovich, Ben-Sasson, Razborov, and Wigderson [4], and by Krajíček [51, 52]; see also the monograph [55] for a comprehensive treatment.

A *proof complexity generator* is any encoding of the statement that expresses that a point $b \in \{0, 1\}^\ell$ is not in the image of a given polynomial time mapping $G: \{0, 1\}^m \rightarrow \{0, 1\}^\ell$, with $m < \ell$. The proof complexity of such statements depends on G and is sensitive to the encodings; indeed, [4] suggested three different encodings (see also Sokolov [79] for a discussion). We say that a proof complexity generator is *hard* against a proof system when for *every* point b there is no short proof that b is not in the image of G . This means that the proof system cannot prove that G is non-surjective. (Notice that even a single b for which the proof complexity generator is easy would constitute a proof of non-surjectivity, hence we insist that no such b exists; this point is explained further in Alekhovich *et al.* [4].)

The hope is that for strong propositional proof systems, one can establish (at least conditionally) that there are no $\text{poly}(\ell)$ -size proofs that g is non-surjective, under the assumption that the mapping g is sufficiently pseudorandom (cf. [76]). Razborov in [76] proved the existence of hard proof complexity generators for systems such as PCR and $\text{Res}(k)$. However, as Razborov mentions, the important case of generators for PCR over $\mathbb{F}_2$ is completely open (the importance of this case is explained after stating Theorem 2.4).

WPHP and the provability of circuit lower bounds. WPHP and proof complexity generators also play a central role in the metamathematics of complexity theory—particularly in understanding which systems can efficiently prove circuit lower bounds. This connection was first observed by Razborov [72], who showed that WPHP reduces to circuit lower bound statements. Informally, this means that if WPHP is hard for some (“nice”) proof system $\mathcal{P}$, then $\mathcal{P}$ cannot prove any circuit lower bound efficiently. This idea was employed by Razborov [72] and later by Raz [68]. The former established degree lower bounds for WPHP in polynomial calculus, and the latter established resolution size lower bounds. Both work then leveraged the above reduction to show that the statement “ $\text{NP} \not\subseteq \text{P/poly}$ ”, when encoded as a family of CNF formulas, does not admit efficient refutations in those systems.

More generally, one can define the *truth-table generator*, which maps a description of a small Boolean circuit to the truth table of the function it computes. Establishing that a string b lies outside the image of this generator is equivalent to proving that no small circuit computes the function with truth table b . Hence, the truth-table generator constitutes a proof-complexity generator against a system $\mathcal{P}$ exactly when $\mathcal{P}$ cannot efficiently prove any circuit lower bounds.

Linear algebraic instances. Beyond simple counting arguments like the pigeonhole principle, linear algebra is a key tool in discrete mathematics. Accordingly, Cook and Rackoff suggested linear algebraic statements as potential hard instances for strong propositional proof systems. Bonnet, Buss, and Pitassi [15, Sec. 3.1.1] further explored this idea and identified candidates hard tautologies such as the *oddtown principle*, which asserts $XX^T \neq I_m$ for an $m \times n$ variable matrix X with $m > n$, where I_m is the $m \times m$ identity matrix. Soltys and Cook [80] systematically studied the *hard matrix identities* in the context of bounded arithmetic, focusing on statements such as the *inversion principle*, which asserts that $XY = I_n$ implies $YX = I_n$ for square $n \times n$ variable matrices X, Y , as well as other equivalent formulations of basic linear algebraic facts. Their goal was to identify minimal formal theories capable of proving such statements. Soltys and Urquhart [81] showed that the pigeonhole principle is reducible to (hence, not stronger than) these matrix identities. Later, it was demonstrated by Hrubeš and Tzameret over $\mathbb{F}_2$ [39] and by Tzameret and Cook over the integers [84] that these linear algebraic principles admit quasipolynomial-size propositional (i.e., Frege) proofs and are provable in relatively weak theories of arithmetic. The latter two works roughly show that linear algebra sits within “ NC^2 -reasoning” (formally, the theory VNC^2). The hard matrix identities were also considered in the

context of the Ideal Proof System (IPS) by Grochow and Pitassi [35], and Andrews and Forbes [7] who investigated a close unsatisfiable instance $\{\det(X) = 0, XY = I\}$ for two square variable matrices X, Y .

A related algebraic principle is the (*strong*) rank principle: Let $m = n + 1$. Given two variable matrices X, Y of dimension $m \times n$ and $n \times m$, respectively, the product $XY = A$ is unsatisfiable whenever A is an $m \times m$ scalar matrix of rank m . This formula was considered implicitly in Krajíček [54], as a way to algebraically express proof complexity generators based on PHP.

For the specific case where A is the identity matrix I , Soltys and Cook [80, Equation (V), page 287] showed that the strong rank principle is equivalent over their theories (the base one denoted LA; hence over Frege) to the hard matrix identities, and specifically the inversion principle. Galesi, Grochow, Pitassi, and She [31] formulated the rank principle explicitly, for the special case when A is the identity matrix I , and when m is not necessarily $n + 1$. Galesi *et al.* established PC degree lower bounds for the rank principle, using the same reduction from PHP degree lower bounds that was used in [81].

The weak rank principle. In this work we show the advantage in considering what we call the *weak rank principle*, *WRank* for short, as a natural algebraic analogue and augmentation of the weak pigeonhole principle. Given two variable matrices X and Y of dimensions $m \times n$ and $n \times m$, respectively, *WRank* states that

WRank. For an arbitrary scalar matrix A of rank at least $n + 1$, the equation $XY = A$ is unsatisfiable, where $m > n$; here we mean that m is *arbitrarily larger* than n .

Notice that WPHP is (usually) considered in the same regime, namely when m is arbitrarily larger than n [74]. Accordingly, all our lower bounds are expressed in terms of n and are *independent* of m .

Our results show that the weak rank principle can be applied in settings where the weak pigeonhole principle fails, and that it provides a framework for both lower and upper bounds in proof complexity. In particular, we obtain lower bounds and corresponding proof-complexity generators with a good stretch for $\text{PCR}_{\mathbb{F}_2}$ which can also be iterated to a function generator (where none were previously known) and introduce a new generator for the Sherali–Adams system, establish the hardness of circuit lower-bound statements for $\text{PCR}_{\mathbb{F}_2}$ (which was open), and identify upper bounds—namely, feasibly constructive proofs—that are facilitated by the weak rank principle itself.

The proof systems we shall consider in this work are polynomial calculus resolution PCR and Sherali–Adams SA. See [Section 4](#) for the precise definitions and references.

What is known about WRank. For the *strong* rank principle (where $m = n + 1$, and in fact up to $m \leq n^2$), size lower bounds against PCR could be shown over any field, by a reduction to the (strong) pigeonhole principle (with m pigeons; see [31, 81] for this reduction), and then using a size-degree tradeoff argument. However, for the *weak* rank principle (when m is arbitrary) only *degree* lower bounds are known in polynomial calculus resolution (PCR); again, this follows easily from known degree lower bounds for WPHP. However, the more challenging and meaningful regime is that of *size* lower bounds for PCR and stronger systems. Notably, no PCR size lower bounds are currently known for WPHP for arbitrary many pigeons, and in fact when the number of pigeons exceeds n^2 (see Mikša and Nordström [57] and de Rezende, Nordström, Risse and Sokolov [28])¹. This motivates algebraically extending WPHP to *WRank*, which enables new size lower bounds in settings where WPHP has so far proven insufficient (as the current work shows).

With respect to *upper bounds*, while WPHP admits polynomial-size proofs in TC^0 -Frege, the best known upper bound for *WRank* is a polynomial-size proof in NC^2 -Frege. This could be shown as

¹ Beyond n^2 pigeons, size lower bound techniques against PCR based on degree break down. Note that for *resolution*, using different techniques by Raz [68] and subsequently Razborov [73, 75] as well as de Rezende *et al.* [28], the case of WPHP lower bounds with arbitrary many pigeons m is resolved.

follows: for strong enough proof systems closed under low degree reductions WRank is implied by the special case of WRank when $A = I$ (see [Lemma 7.4](#)), which in turn is implied by the inversion principle by [80]; finally, the inversion principle admits polynomial-size proofs in NC^2 -Frege by the work of Hrubeš and Tzameret [39]. On the other hand, it is reasonable to conjecture that Frege does not admit polynomial-size proofs of WRank by the fact that linear algebraic statements are expected not to have polynomial-size Frege proofs as noted by Buss *et al.* [15].

Note that the relation between the weak rank principle and the strong one (and hence, by [80], its relation to the hard matrix identities) is not entirely understood, since the latter concerns square matrices, while our focus is on highly non-square matrices (hence, the “weak” regime).

2 Our Results

We present four types of results based on the weak rank principle: Lower bounds and generators for PCR over $\mathbb{F}_2$, hereafter denoted $\text{PCR}_{\mathbb{F}_2}$ ([Section 2.1](#)), lower bounds and generators for SA ([Section 2.2](#)), hardness of circuit lower bound statements for $\text{PCR}_{\mathbb{F}_2}$ ([Section 2.3](#)), and upper bounds (namely, feasible constructive proofs) facilitated by the weak rank principle ([Section 2.4](#)).

2.1 Lower Bounds and Generators for $\text{PCR}_{\mathbb{F}_2}$

We establish lower bounds and construct generators for PCR over $\mathbb{F}_2$, under the following encodings of increasing strength:

- **Algebraic encoding.** This setting provides the most direct algebraic formulation and serves as a foundation for the subsequent encodings.
- **Perfect matching (PM) encoding.** In the PM encoding, we obtain lower bounds for the weak rank principle, when $A = I$, against $\text{PCR}_{\mathbb{F}_2}$. This encoding is in a CNF. To get a generator in this encoding, we should prove the lower bound for every A . Although it is achievable by similar techniques used in the next item, we do not formally prove this case, since this encoding does not appear to be directly useful to establish hardness of circuit lower bound statements. The case $A = I$ is conceptually simpler than the general case of arbitrary A , yet serves as an instructive case for the forthcoming generators because those results expand on similar ideas. It requires a different notion of degree from the algebraic encoding, that is useful in the sequel (e.g., for SA lower bounds).
- **Bamboo-tree encoding.** The bamboo-tree encoding yields the strongest construction, achieving the same stretch as the generators above ($2mn$ to m^2). In addition, we are able to prove that this generator is *iterable* (in the sense of [52, 76]) under this encoding, hence we establish a *function* generator (i.e., a generator with a stretch of $2^{n^{\Omega(1)}}$). A notable consequence of this function generator is that $\text{PCR}_{\mathbb{F}_2}$ does not admit short proofs of circuit lower bound statements such as $\text{NP} \not\subseteq \text{P/poly}$.

We start by providing more technical background to our new generators.

Context and motivation for our generators. Recall that a proof complexity generator encodes the statement that a point $b \in \{0, 1\}^\ell$ is not in the image of a polynomial-time map $G: \{0, 1\}^m \rightarrow \{0, 1\}^\ell$, with $m < \ell$. When the weak rank principle is viewed in this framework, the matrix product XY plays the role of the map G , and the point outside its image is the matrix A .

The known generators for $\text{PCR}_{\mathbb{F}_2}$ in the literature are based on reductions to the pigeonhole principle. For the *strong* pigeonhole principle, Krajíček [54, Section 1] constructed a generator stretching n bits to $n + 1$ bits. The same construction can be combined with the best known lower bounds for the *weak*

pigeonhole principle over $\text{PCR}_{\mathbb{F}_2}$ [5, 57] to yield a slightly superlinear stretch—from n^3 to n^4 bits². (For comparison, our results achieve a nearly quadratic stretch of $2mn$ to m^2 since our lower bounds hold in the regime where m is arbitrarily larger than n .)

In the present work, we aim to base the hardness of such generators on the hardness of WRank instead, which is *stronger* than WPHP. For instance, WPHP is easy for the SoS proof system, but WRank is plausibly hard for it; likewise, $\text{Res}(\log n)$ admits quasipolynomial-size refutations of WPHP [56], yet statements based on linear algebra (and in particular WRank) are expected to be hard for it [15].

As mentioned above, for arbitrary m , the hardness of WPHP against PCR over any field remains open. If such bounds were known, they could yield strong generators through Krajíček’s construction and further applications. In particular, a sufficiently strong lower bound for the Onto-WPHP would, via Raz’s method [68], yield that certain circuit lower bound statements for unbounded fan-in circuits cannot be efficiently proven in PCR.

This led Alekhovich *et al.* [4] to study hard instances distinct from WPHP on which to base proof-complexity generators and their applications to circuit lower-bound statements. They proposed the so-called *Nisan generator*, a special case of the Nisan–Wigderson generator [61], which is based on systems of linear equations over $\mathbb{F}_2$ defined over expander graphs. The stretch of this construction alone is not sufficient to yield application to circuit lower bound statements.

An important research direction is to obtain proof complexity generators with *as large stretch as possible*. When the stretch is exponential (i.e., $2^{n^{\Omega(1)}}$ in the seed length n) the generator is called a *function generator*. To achieve such amplification, Krajíček [52] introduced the notion of *s-iterability*, designed to boost the stretch of a variant of the Nisan generator [61]. Improving [4], Razborov [76] relaxed the expander requirements and showed that linear systems as in [4] can be *iterated* to achieve a much larger stretch—about 2^{n^ϵ} . In particular, the hardness of the base generator transfers to the iterated version, since an iterated expander remains an expander. However, systems of linear equations over $\mathbb{F}_2$ are *easy* to refute in $\text{PCR}_{\mathbb{F}_2}$; therefore, this construction cannot yield generators hard for $\text{PCR}_{\mathbb{F}_2}$.

The limitations of existing principles motivate the use of WRank which subsumes both the weak pigeonhole principle and systems of linear equations. This allows us to obtain lower bounds for WRank (over $\mathbb{F}_2$) for arbitrary m , thereby producing a hard *base generator*. The hardness of WRank persists under iteration, which enables us to construct function generators that remain hard and leads to the hardness of corresponding circuit lower-bound statements.

Among other results known about generators are Khaniki’s Nisan–Wigderson generator against $\text{AC}^0[p]$ -Frege [48] encoded as a (non-CNF) propositional formula. In addition, Sokolov [79] improved the lower bounds for a particular generator encoding (functional encoding as in [4]).

Algebraic encoding. The algebraic encoding of the negation of WRank, denoted $\text{Rank}_n^m(A)$, consists of m^2 polynomial equations of degree 2,

$$\text{Rank}_n^m(A): \quad \sum_{k=1}^n x_{i,k} y_{k,j} = A_{i,j}, \quad \text{for all } i, j \in [m]. \quad (1)$$

Theorem 2.1 (Algebraic encoding lower bound; see Theorems 5.7 and 5.44). *For every $m > n$ and $A \in \mathbb{F}_2^{m \times m}$, every $\text{PCR}_{\mathbb{F}_2}$ refutation of $\text{Rank}_n^m(A)$ requires size $2^{\Omega(n)}$. Consequently, $\text{Rank}_n^m(A)$ is a proof complexity generator stretching $2nm$ bits into m^2 . Moreover, the iterated variant of $\text{Rank}_n^m(A)$ is a function generator with stretch $2^{n^{\Omega(1)}}$.*

²Given a formula PHP_n^m , Krajíček [54] builds a function stretching $nm+tn$ bits into tm , for any parameter t ; the ratio $tm/(nm+tn)$ is maximised at $t = \Theta(n^2)$. The best known lower bounds for PHP_n^m against PCR hold for $m = o(n^2)$. If we plug these values of t and m into $nm+tn$ and tm , respectively, we obtain a stretch of n^3 to n^4 .

Here, size is measured by the number of distinct monomials appearing in the refutation (see precise definitions in [Section 4](#)).

The algebraic encoding lower bounds are obtained using random restrictions, which yield a size-to-degree reduction and already suffice to derive the hardness of certain circuit lower-bound statements (for example, for non-commutative algebraic branching programs; see [Section 8.3.1](#)). The argument in this setting is technically simpler than the proofs for the CNF encodings introduced below, as the latter requires translating the algebraic reasoning into a combinatorial framework. The lower bound for the algebraic instance in [Theorem 2.1](#) thus forms the basis for the CNF lower bounds presented later and underlies the high-level random-restriction approach used throughout.

The algebraic encoding, however, is specific to algebraic proof systems and therefore not directly comparable with encodings suited for other systems. To address this, we encode WRank as a CNF formula based on the perfect matching principle, and later consider the more demanding *bamboo-tree* encoding (see [Section 5.3](#)). The choice to work with CNF encodings is natural, as CNF is the standard formulation in proof complexity. Moreover, to extend lower bounds from algebraic proof systems such as PCR and SoS to systems like AC^0 -Frege or Res(lin) (that is, resolution over linear equations [69]), CNF encodings are the appropriate framework, whereas the algebraic formulation—involving degree-2 polynomial equations—lies outside the expressive power of these systems.

Perfect matching encoding. The perfect matching CNF encoding of WRank, denoted $PMRank_n^m(A)$, encodes each equation $\sum_{k=1}^n x_{i,k} y_{k,j} = A_{i,j}$ of $Rank_n^m(A)$ by stating there exists a perfect matching on the satisfied monomials (i.e., those equal 1) of this equation. To encode such a matching we use extension variables for each pair of monomials in the equation (together with an extra point if $A_{i,j} = 1$). See [Section 5.2](#) for the definition. Similar encodings of parity computations were considered for example in Impagliazzo–Segerlind and Ken [42, 47].

As mentioned above, we shall consider only the case $A = I$ and prove lower bounds for this encoding against $PCR_{\mathbb{F}_2}$. The case for every A , yielding a generator could be achieved with techniques similar to the ones in the bamboo-tree encoding.

Theorem 2.2 (Perfect matching encoding lower bound against $PCR_{\mathbb{F}_2}$; see [Theorem 5.10](#)). *For every $m > n$, every $PCR_{\mathbb{F}_2}$ refutation of the perfect matching CNF encoding $PMRank_n^m(I_m)$ requires size $2^{\Omega(n)}$.*

The case of $PCR_{\mathbb{F}_2}$ is the most interesting because the perfect matching principle (stating that there is no perfect matching on an odd size set) is in fact easy for $PCR_{\mathbb{F}_2}$, hence the hardness of $PMRank_n^m(I_m)$ does not stem from the encoding per se. This contrasts many other cases where the hardness follows from the hardness of the perfect matching principle itself [9, 17, 19, 33, 41, 50], as we show for the case of $PCR_{\mathbb{F}}$ when $\text{char}(\mathbb{F}) \neq 2$ in the next theorem.

The next result demonstrates the sensitivity of the encoding of the generators. Although we can construct generators for several proof systems using the PM encoding, it is not known how to get hardness of circuit lower bound statements from these generators.

Theorem 2.3 (Perfect matching encoding lower bound over any field; see [Theorem 5.18](#)). *For every $m > n$ and $A \in \{0, 1\}^{m \times m}$, any $PCR_{\mathbb{F}}$ refutation (for a field $\mathbb{F}$ with $\text{char}(\mathbb{F}) \neq 2$) or SoS refutation of the perfect matching CNF encoding $PMRank_n^m(A)$ requires size $2^{\Omega(n)}$. Similarly, any AC^0 -Frege refutation requires size $2^{n^{\Omega(1)}}$, where the constant in the exponent depends on the circuit depth.*

Bamboo-tree encoding. The main challenge and open problems regarding proof complexity generators and their applications is predominantly about proving lower bounds for good encodings; where an encoding is good, if it has extension variables that correspond to gates of a circuit, hence could be applied to circuit statement lower bounds [4, 76].

The bamboo-tree CNF encoding of WRank, denoted $\text{BTRank}_n^m(A)$, encodes each equation $\sum_{k=1}^n x_{i,k} y_{k,j} = A_{ij}$ of $\text{Rank}_n^m(A)$ by using extension variables introduced to sequentially compute all inner products involved. The extension variables for each inner product are arranged in a totally unbalanced binary tree known as a “bamboo”. Each extension variable $u_{i,j,\ell}$ encodes the partial sum $\sum_{k=1}^\ell x_{i,k} y_{k,j}$ (for $\ell \in [n]$). Moreover, each monomial $x_{i,k} y_{k,j}$ has its own extension variable $z_{i,j,k}$.




Figure 1: Example of one axiom of $\text{BTRank}_n^m(A)$. Here, $u_{i,j,\ell}$ are the extension variables for the partial sum $\sum_{k=1}^\ell x_{i,k} y_{k,j} = \sum_{k=1}^\ell z_{i,j,k}$ (for $\ell \in [n]$). See [Definition 5.19](#).

This natural encoding corresponds to a circuit in the sense of [4, 76], and similar encodings have been used repeatedly for parity computations [4, 76, 81].

Theorem 2.4 (Bamboo-tree encoding lower bound; see [Theorems 5.20](#) and [5.51](#)). *For every $m > n$ and $A \in \mathbb{F}_2^{m \times m}$, every $\text{PCR}_{\mathbb{F}_2}$ refutation of $\text{BTRank}_n^m(A)$ requires size $2^{\Omega(n)}$. Consequently, $\text{BTRank}_n^m(A)$ is a proof complexity generator stretching $2nm$ bits into m^2 bits. Moreover, there is an iterated variant of $\text{BTRank}_n^m(A)$ that yields a function generator with stretch $2^{n^{\Omega(1)}}$.*

This result resolves an open problem raised in [4, 76] concerning the construction of proof-complexity generators for $\text{PCR}_{\mathbb{F}_2}$. Razborov’s [76] established generators against PCR based on the Nisan construction only over fields of characteristic different from 2. Lower bounds over $\mathbb{F}_2$ are particularly significant and technically more challenging than those over other fields, as we now explain.

In general, when considering the weak rank principle, there are two distinct field parameters to consider: the *ground field*, over which the proof system operates, and the *expressed field*, over which the matrix product expressed in the weak rank principle is computed. Although these two fields may differ, the case where both are $\mathbb{F}_2$ —that is, WRank over $\mathbb{F}_2$ and PCR over $\mathbb{F}_2$ —is the most natural and the most difficult.

For example, unsatisfiable linear systems over $\mathbb{F}_2$ are easy for $\text{PCR}_{\mathbb{F}_2}$, and more generally, linear systems over $\mathbb{F}_p$ are easy for $\text{PCR}_{\mathbb{F}_p}$. When $p \neq 2$, however, such linear systems are non-Boolean and require explicit encoding of non-Boolean field elements, which makes the lower-bound task formally easier but conceptually less meaningful, as the hardness would then usually stem from the encoding rather than from the inherent structure of the weak rank principle. In contrast, working with both ground and expressed fields equal to $\mathbb{F}_2$ keeps all variables Boolean and ensures that the difficulty arises from the combinatorial and algebraic content of the principle itself.

Furthermore, $\mathbb{F}_2$ is usually the most challenging case for algebraic and semi-algebraic proof systems. For instance, lower bounds in the Nullstellensatz system or its variants often require substantially more complex designs over $\mathbb{F}_2$ than over larger or characteristic-0 fields (cf. [12, Theorem 12], [49]). Likewise, proof systems such as $\text{Res}(\text{lin})$ [69] admit lower bounds over characteristic 0 [65], whereas corresponding results over finite fields—and in particular over $\mathbb{F}_2$ (cf. [43])—remain open. Thus, establishing hard generators over $\mathbb{F}_2$ advances our understanding precisely in the setting that is both the most natural and the least understood.

Lower bounds techniques. All our lower bounds follow the same general *random self-reduction scheme*: we first reduce proof size to suitable notions of degree, tailored to the formulas and results

under consideration. For example, although total degree is the standard measure for monomials, one may also work with refined degree measures that count only variables from a designated set. This means that we shall start, by way of contradiction, from a small refutation (having only a few distinct monomials) and apply a random restriction or a random substitution (changing variables to different polynomials in different variables) that simplifies terms in the refutation *with respect to the appropriate notion of degree*, reducing the formula to a smaller instance of the same principle. The second step is proving a lower bound on the relevant notion of degree, which will result in a contradiction (namely, no small size refutation exists).

The novelty and challenge is in finding the appropriate notions of degree and in devising the right random substitutions which will decrease the corresponding notion of degree. These notions of degree increase in level of complication along with the complexity of the encoding. Accordingly, the random substitutions increase in level of complication.

One idea that is different from the usual proof complexity lower bounds is that our generalised random restrictions are *random substitutions* of both 0-1 and polynomials in formulas. In PHP lower bounds, the usual argument hinges on applying 0-1 restrictions (representing partial matchings). But in our case partial 0-1 random restrictions are not always sufficient, hence we need to substitute the original variables by both 0-1 and other *variables or their negations*. Our goal will be to get self-reducibility: a smaller version of the principle. In the literature such random substitutions were used in works by, e.g., Pitassi, Rossman, Servedio and Tan [67] and subsequently Hastad [37] for AC^0 -Frege lower bounds.

Algebraic and bamboo-tree encoding of rank principle. The following explains the lower bounds for both the algebraic encoding $\text{Rank}_n^m(A)$ and the bamboo-tree encoding $\text{BTRank}_n^m(A)$. We outline some ideas behind the proof of [Theorem 2.4](#), which establishes a $2^{\Omega(n)}$ size lower bound for $\text{BTRank}_n^m(A)$ in $\text{PCR}_{\mathbb{F}_2}$. The key challenge is that this lower bound must hold independently of the parameter m , which precludes the use of standard expander-based techniques. These techniques would typically involve restricting one of the matrices, say X , to an expander, zeroing out most of the entries $x_{i,k}$, for $k \in [n]$, in each row $i \in [m]$. However, when the gap between m and n is too large, there are no bipartite expanders with suitable parameters for achieving degree lower bounds.

Our solution hence avoids using expanders. We relax the degree notion, aiming to balance two competing requirements for establishing size lower bounds: (1) the existence of random self-reductions of $\text{BTRank}_n^m(A)$ that effectively reduce the relaxed degree, and (2) the existence of a lower bound on the relaxed notion of degree for refutations of $\text{BTRank}_n^m(A)$. The random self-reduction we design to achieve (1) is highly sensitive to the structure of the extension variables (the z -variables and u -variables defined above) and its necessarily technical nature leads us to present it in two stages for clarity. In this outline, we focus solely on motivating the notion of the relaxed degree that the self-reduction aims to decrease, and we primarily discuss requirement (2).

Consider a term t in x -, y - and z -variables of $\text{BTRank}_n^m(A)$, that is, a product of some x -, y - and z -variables and their negations. For a variable v , we write v^1 to denote v and v^0 to denote its negation, $\bar{v}$. Let C be the set of indices k such that $x_{i,k}^b$, $y_{k,j}^b$ or $z_{i,j,k}^b$ appears in t for some $b \in \{0,1\}$ and $i, j \in [m]$. We refer to C as the set of *columns* X -, Y^T - and Z -mentioned in t . Now, let τ be the substitution of $X^T A$ for Y , i.e., $\tau(x_{i,k}) = x_{i,k}$, $\tau(y_{k,j}) = \sum_{i' \in [m]} x_{i',k} A_{i',j}$, and $\tau(z_{i,j,k}) = \tau(x_{i,k})\tau(y_{k,j})$ for all $i, j \in [m]$ and $k \in [n]$, with $\tau(\bar{v}) = 1 + \tau(v)$ for any variable v . We see that if we apply τ to t , the columns X -mentioned in each term of the polynomial $\tau(t)$ remain in C (note that under τ the term t becomes a *polynomial*, not necessarily a single monomial, in the X -variables only). It follows that each term t' of $\tau(t)$ either has degree at most $|C|$ (since we are working over the ring of multilinear polynomials; this happens when each column index $k \in C$ contributes a single variable $x_{i,k}$ to t'), or there is $k \in C$ and $i < i' \in [m]$ such that t' is a multiple of $x_{i,k}x_{i',k}$ (by the pigeonhole principle; this happens when the column $k \in C$ contributes more than a single variable $x_{i,k}$ to t').

Why is the last observation useful? The substitution τ maps the set of polynomial equations $XY = A$ to the set of polynomial equations $XX^T A = A$, and the latter has a straightforward degree-2 derivation from the oddtown equations $XX^T = I_m$, which in turn have an immediate degree-2 derivation from the algebraic formulation of the functional pigeonhole principle FPHP_n^m . For FPHP_n^m , Razborov [72] established a degree lower bound of $n/2$, which holds even for PCR proofs operating with multilinear polynomials modulo the ideal generated by the hole axioms ($x_{i,k}x_{i',k}$ for $i < i' \in [m]$ and $k \in [n]$) and the functionality axioms ($x_{i,k}x_{i,k'}$ for $i \in [m]$ and $k < k' \in [n]$). We conclude that there is no refutation of the *algebraic* formulation $\text{Rank}_n^m(A)$ each term of which has less than $n/2$ X - or Y^T -mentioned columns. If such a refutation existed, then τ , along with the aforementioned degree-2 derivations, would convert it into a refutation of FPHP_n^m , where every term has either degree less than $n/2$ or is a multiple of a hole axiom, contradicting the cited lower bound. This forms the column-degree lower bound part in the proof of [Theorem 2.1](#). Similarly, this argument rules out the existence of a refutation of a variant of $\text{Rank}_n^m(A)$ that involves the x -, y -, and z -variables (but not the u -variables), where each term in the refutation has less than $n/2$ X -, Y^T - or Z -mentioned columns.

While the argument in the above paragraph works for the algebraic formulation $\text{Rank}_n^m(A)$, it is not sufficient for the bamboo-tree encoding $\text{BTRank}_n^m(A)$. This is because τ must be consistently extended to the u -variables by $\tau(u_{i,j,k}) = \sum_{\ell \in [k]} \tau(z_{i,j,\ell})$, and the column index k of $u_{i,j,k}$ is not particularly relevant here: when τ is applied to a product of multiple u -variables mentioning only a small number of columns, it leads to a blow up in the X -mentioned columns in the resulting terms. Thus, our new notion of degree for the u -variables concerns only the first of the three indices of $u_{i,j,k}$. We say that i is *U-left-row-mentioned* in a term t if $u_{i,j,k}^b$ appears in t for some $j \in [m]$, $k \in [n]$, and $b \in \{0, 1\}$.

To appreciate this definition, let us first calculate $\tau(u_{i,j,k}^{1-b})$ modulo the hole axioms. We have

$$\begin{aligned} \tau(u_{i,j,k}^{1-b}) &= b + \tau(u_{i,j,k}) = b + \sum_{\ell \in [k]} \tau(z_{i,j,\ell}) = b + \sum_{\ell \in [k]} \tau(x_{i,\ell})\tau(y_{\ell,j}) = b + \sum_{\ell \in [k]} x_{i,\ell} \sum_{i' \in [m]} x_{i',\ell} A_{i',j} \\ &= b + \sum_{\ell \in [k]} \sum_{i' \in [m]} x_{i,\ell} x_{i',\ell} A_{i',j} = b + \sum_{\ell \in [k]} x_{i,\ell} x_{i,\ell} A_{i,j} = b + A_{i,j} \sum_{\ell \in [k]} x_{i,\ell}, \end{aligned}$$

where the penultimate equality is by the hole axioms and the last equality is by the Boolean axioms. Hence, we see that, modulo the hole axioms, the only remaining row index is i . Let t be a term consisting of u -variables and their negations and let R be the set of indices *U-left-row-mentioned* in t . Based on the above calculation, we observe that each term t' of the polynomial $\tau(t)$ either has degree at most $|R|$, or is a multiple of a hole axiom, or (by the pigeonhole principle) is a multiple of a functionality axiom of FPHP_n^m .

The preceding discussion motivates the definition of the relaxed degree of a general term t in the variables of $\text{BTRank}_n^m(A)$ as the sum of the cardinality of the set of indices *U-left-row-mentioned* in t and the cardinality of the set of X -, Y^T - or Z -mentioned columns in t . For a precise definition of random self-reductions and a proof of their effectiveness in reducing the relaxed degree, we refer the reader to the proof of [Theorem 5.20](#).

Iterated rank principles. In order to address the iterated variants of the rank principle, we need additional techniques besides the ones that were sufficient for [Theorem 2.1](#). We start by describing the ideas behind the iterated algebraic principle. We consider the classical Goldwasser, Goldreich and Micali [32] iteration protocol, which was also used by Razborov [76]. This protocol achieves a generator with exponential stretch (i.e., a function generator) by iterating the base generator (the way to achieve such a function generator from the weak rank principle is described after [Theorem 2.8](#) in this section).

We start with the algebraic encoding. Consider the polynomial system $\text{IRank}_n^m(\{A^{(\pi)} : \pi \in \{0, 1\}^{\leq \kappa}\}, \kappa) := \bigcup_{\pi \in \{0, 1\}^{\leq \kappa}} X^{(\pi)}Y = A^{(\pi)}$. Here, each $X^{(\pi)}Y = A^{(\pi)}$ is a distinct instance of the weak rank principle with parameters m and n and the x -variables indexed by the binary string π , except

that the matrices $A^{(\pi)}$ are not necessarily constant: we also allow them to contain variables from $X^{(\pi')}$ for some strings π' with $|\pi'| > |\pi|$. This instance is a generalisation of WRank: letting “ $()$ ” be the empty string, when $\kappa = 0$ we have that $\text{IRank}_n^m(\{A^{()}\}, 0) = \text{Rank}_n^m(A^{()})$. Thus, to specify the entries of different $A^{(\pi)}$ ’s, it is helpful to view the formula arranged as a binary tree shown in Figure 2.




Figure 2: The Goldwasser, Goldreich and Micali-style iteration structure of IRank.

In this tree, the entries of each matrix $A^{(\pi)}$ can depend only on the entries of $X^{(\pi')}$ from the corresponding subtree. This differs from the standard notion of s -iterability [52, 76], which only assumes a flat structure; that is, the iterated copies are arranged as a path and not as a tree. Our lower bound depends on the depth of the arrangement (κ in the above notation), and the tree structure will maximise our lower bound. The tree structure was also used by Razborov [76]. Although this construction turns out to be sufficient for applications to circuit lower bounds statements. Our goal would be to show that any refutation of IRank with parameters m and n requires size $2^{(n/\kappa)^{\Omega(1)}}$.

We consider two notions of degree for IRank:

- the *column degree*, which measures the number of different columns of $X^{(\pi)}$ or Y^T the term mentions, where the same columns from different matrices are counted *at most once*; and
- the *row degree*, which measures the number of different rows of $X^{(\pi)}$ or Y^T the term mentions, where the same rows from different matrices *are counted individually*.

For example, the term $x_{1,1}^{(0)}x_{1,1}^{(1)}y_{1,1}$ has column degree 1 and row-degree 3. We perform two random self-reductions ρ and σ . The former reduces the row degree of every term in a refutation to $\leq d$ while the latter reduces the column degree to $\leq d$. This implies the total degree does not exceed $2d^2$.

We show that refuting $\text{IRank}_n^m(A^{(\pi)}, \kappa)$ requires degree $n/2$ by reducing it to a single copy of the formula $\text{Rank}_n^m(I_m)$ using a reduction that substitutes each $X^{(\pi)}$ with $A^{(\pi)}X$ for a new variable X . Since the entries of $A^{(\pi)}$ are not necessarily constant, this reduction can increase the degree of a refutation. However, it can grow by at most a factor of $\kappa + 1$, because for each copy the degree grows by at most 1. This explains the dependency on κ we mentioned earlier.

Up to now we discussed the algebraic encoding of the iterated rank principle. We now turn to the CNF encoding of $\bigcup_{\pi \in \{0,1\}^{\leq \kappa}} X^{(\pi)}Y^{(\pi)} = A^{(\pi)}$ (note that here the matrix Y is different in each iteration, unlike before) and again use the iteration protocol from [32]. We encode $\bigcup_{\pi \in \{0,1\}^{\leq \kappa}} X^{(\pi)}Y^{(\pi)} = A^{(\pi)}$ using the extension variables that are similar to the ones from BTRank and additionally impose some structure on the matrices $Y^{(\pi)}$. Our goal is again to prove that any $\text{PCR}_{\mathbb{F}_2}$ refutation requires size $2^{(n/\kappa)^{\Omega(1)}}$.

The notion of column degree is not well suited for the extension variables. For instance, given an extension variable $u_{i,j,k}^{(\pi)}$ which semantically encodes the sum $\sum_{\ell \in [k]} x_{i,k}^{(\pi)} y_{k,j}^{(\pi)}$, it is clear that it depends not only on the k th column but also on the preceding columns $\{1, \dots, k-1\}$. To address this, we restrict our matrices $Y^{(\pi)}$ (and hence the extension variables) to a bipartite expander G with a small

left-degree Δ . This automatically restricts the column degree of all but x -variables to be at most Δ , thus for all extension variables we only need to reduce the row degree.

Overall, the proof strategy is executed as follows: we start by reducing the row-degree (and thus the total degree) in u -variables that compute partial inner products. We then reduce the column degree in the x - and z -variables (recall that the latter compute the binary ANDs: $z_{i,j,k}^{(\pi)} = x_{i,k}^{(\pi)} y_{k,j}^{(\pi)}$). This is performed in the same way as in the algebraic construction. Finally, we reduce the row degree in the x -, y -, and z -variables by using a random substitution similar to the one that reduces the row degree for lRank. To deal with the z -variables we need to further modify the refutation. This is because of the terms of the form $s = \prod_{j \in J} \bar{z}_{i_0,j,k}$ (or of the symmetric form $\prod_{I \in J} \bar{z}_{i,j_0,k}$), which have the shape of a *star* (when the x - and y -variables are considered as two sides of a bipartite graph with an edge between $x_{i,k}^{(\pi)}$ and $y_{k,j}^{(\pi)}$ present whenever $\bar{z}_{i,j,k}^{(\pi)}$ appears in s). Such stars have common centre i_0 , and our substitutions do not, in fact, set them to 0 w.h.p. However, we can show that each such star with sufficiently many rays—that is, with large enough set J —collapses to its centre $\bar{x}_{i_0,k}$ w.h.p. This property limits the degree of all such star terms. For the rest of the variables, we can greedily identify a large set of variables in a term that are independently set to 0 w.h.p., implying that such a term vanishes from the proof w.h.p.

2.2 Generators for Sherali–Adams

We address the problem of establishing proof complexity generators for the Sherali–Adams (SA) semi-algebraic refutation system [25].

First, note that the Nisan generator construction, denoted $\tau(A, b)$ in [76], can be shown to work also for SA and even SoS. Specifically, its hardness follows from the lower bound on the *Gaussian width* of $\tau(A, b)$ as introduced in [14], and whose lower bound against SoS was established in [33, 77]. Combined with the random restriction from Razborov [76], this gives an SA size lower bound for $\tau(A, b)$.

We provide a different generator for SA than Razborov’s by proving that our bamboo-tree encoding of WRank is such a generator. Note that WPHP is easy for SA in terms of size. Therefore, WPHP cannot be directly used as a generator. We thus come up with a new technique that does not use the reduction to PHP and instead defines a relaxed notion of degree (row degree, in our case). The lower bounds are thus based on a reduction from size to row degree, followed by a lower bound on row degree. To establish the row degree lower bound we come up with a pseudoexpectation based on a new family of distributions tailored to WRank (and not to WPHP).

Lemma 2.5 (Row degree lower bound for SA (informal); Lemma 6.4). *For every $m > n$ and every $A \in \{0, 1\}^{m \times m}$, every SA refutation of (a variant of) $\text{BTRank}_n^m(A)$ requires row degree $n - 1$.*

Theorem 2.6 (Size lower bound for SA (informal); see Theorem 6.11). *Assume $m > n$ and $A \in \{0, 1\}^{m \times m}$ is a Boolean matrix. Then any SA refutation of (a variant of) $\text{BTRank}_n^m(A)$ requires size $2^{\Omega(n)}$.*

This result fits in the generator approach to proof complexity as stated by Razborov [76, p. 417] and proposed earlier by Krajíček [51, 52] and Alekhovich *et al.* [4]. Specifically, the *generator approach* seeks to establish proof complexity generators for as many proof systems as possible, roughly based on the computational hardness of the generator.

To rule out low row-degree SA refutations of $\text{BTRank}_n^m(A)$, we define a family of distributions over partial assignments. Each distribution in this family is indexed by a pair (I, J) , where $I \subseteq [m]$ and $J \subseteq [m]$ are sets of row indices of X and Y^T , respectively, with $|I| + |J| \leq n - 2$. The distribution corresponding to (I, J) is the uniform distribution supported on certain assignments to all variables $x_{i,k}$ with $i \in I, k \in [n]$ and $y_{\ell,j}$ with $\ell \in [n], j \in J$.

Before specifying these assignments, it is helpful to recall the family of distributions used against the functional pigeonhole principle FPHP_n^m by [26]. Each distribution in that family is indexed by a set $I \subseteq [m]$ of at most $n - 2$ pigeons and consists of the uniform distribution over all matchings of the pigeons in I to holes. Thus, the distribution indexed by I is supported on all assignments to the variables $x_{i,k} : i \in I, k \in [n]$ that do not violate any axiom of FPHP_n^m . In other words, no condition whatsoever limits these assignments except the axioms involving the pigeons in I .

In contrast, for $\text{BTRank}_n^m(A)$, we cannot allow similar freedom in the supports of the distributions in our family. It is straightforward to verify that if the support consisted of all assignments to $x_{i,k} : i \in I, k \in [n]$ and $y_{\ell,j} : \ell \in [n], j \in J$ that do not violate any axiom of $\text{BTRank}_n^m(A)$, the marginal distribution condition—a key ingredient of SA lower bounds³—would fail. In other words, limiting the support only by the requirement that the axioms concerning $A_{I,J}$ (the submatrix of A determined by rows I and columns J) hold is insufficient. Our approach therefore imposes stronger restrictions on the support than those dictated by the axioms of $\text{BTRank}_n^m(A)$ alone.

Specifically, we require that the row vectors assigned to the rows of X indexed by I be linearly independent, that the column vectors assigned to the columns of Y indexed by J form a linearly independent set of vectors, and that the product of these rows and columns evaluates to $A_{I,J}$. Of these three requirements, only the last is directly dictated by the axioms of $\text{BTRank}_n^m(A)$. We show that the uniform distribution over such assignments satisfies the marginal distribution conditions for members of our family. Notably, these assignments violate the FPHP_n^m axioms even in the special case $A = I_m$ and $Y = X^T$ (the oddtown case $XX^T = I_m$), highlighting the bespoke nature of our distribution family.

It is plausible that our method could be extended to the case of SoS, though we have not pursued this direction.

2.3 Hardness of Circuit Lower Bounds from the Weak Rank Principle

Next, we connect the weak rank principle to the provability of circuit lower bounds. We show that the weak rank principle is *necessary* for proving certain circuit lower bound statements, while in the sequel (Section 2.4) we show that it is also *sufficient* for proving some known circuit lower bounds of interest. These results suggest a fundamental relationship between the weak rank principle and the provability of circuit lower bounds.

Our results are inspired by—and, to some extent, mirror—the connection between the weak pigeonhole principles and the provability of complexity lower bounds. Razborov [72, 75] and Raz [68] showed the unprovability of circuit lower bounds such as $\text{NP} \not\subseteq \text{P/poly}$ in weak proof systems via reductions from WPHP. Müller and Pich [59] formalised a wide range of circuit lower bounds in Jeřábek’s theory for approximate counting in bounded arithmetic [44–46], which includes the dual weak pigeonhole principle as an axiom. More recently, Chen, Li, and Oliveira [18] showed that certain lower bound statements are *equivalent* to variants of weak pigeonhole principles.

We extend the above line of research in two aspects. First, we present low-degree reductions from the weak rank principle to circuit lower bound sentences, generalising the aforementioned connections by Razborov and Raz [68, 72, 75]. Then, as a concrete example, we show that $\text{PCR}_{\mathbb{F}_2}$ does not have efficient proofs of circuit lower bound statements such as $\text{NP} \not\subseteq \text{P/poly}$, by *iterating* the WRank-based proof complexity generators against $\text{PCR}_{\mathbb{F}_2}$ (as in [76]). A variant of circuit lower bound statements was recently proved to be hard for SoS [10].

Necessity of the weak rank principle. First, we present low-degree reductions showing that WRank is necessary for proving circuit lower bounds:

³In our case, the marginal distribution condition says that marginalising the distribution indexed by $(I \cup \{i_0\}, J)$ to the variables $x_{i,k} : i \in I, k \in [n]$ and $y_{\ell,j} : \ell \in [n], j \in J$ coincides with the distribution indexed by (I, J) , and similarly for the distribution indexed by $(I, J \cup \{j_0\})$.

Theorem 2.7 (WRank is necessary for Boolean circuit lower bounds (informal); see [Corollary 8.4](#) and [Remark 8.5](#)). *Let $\mathcal{P}$ be an algebraic proof system closed under low-degree reductions. If $\mathcal{P}$ cannot prove the weak rank principle efficiently, then for every Boolean function f (represented as a truth table), $\mathcal{P}$ cannot prove circuit lower bounds for f efficiently.*

Interestingly, our argument goes through algebraic circuit complexity. Specifically, we show that WRank is necessary for proving lower bounds against non-commutative algebraic branching programs (non-commutative ABPs, or ncABPs for short), which is a fairly weak algebraic circuit model (making our *unprovability* results stronger). Our reduction relies on Nisan’s characterisation of ncABP complexity via matrix rank [60]. We then establish that lower bounds on ncABP complexity are themselves necessary for proving lower bounds on (Boolean) circuit complexity.

Concrete lower bounds. Inspired by the above connection, we use the generators for $\text{PCR}_{\mathbb{F}_2}$ demonstrated in [Section 2.1](#) to show that this system cannot efficiently prove circuit lower bounds⁴.

Theorem 2.8 (Informal; see [Theorem 8.9](#)). *Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$. For every $s > n^{\Omega(1)}$, any $\text{PCR}_{\mathbb{F}_2}$ refutation of the statement “ f cannot be computed by Boolean circuits of size s ” (encoded as a system of polynomial equations) requires size $2^{s^{\Omega(1)}/n^{O(1)}}$.*

As in the case of Razborov [76] our Boolean circuits are in the basis $\neg, \wedge, \vee, \oplus$ where the last three connectives are binary. The proof idea is similar to [76]. We work with the iterated rank principle $\bigcup_{\pi \in \{0,1\}^{\leq n}} X^{(\pi)} Y^{(\pi)} = A^{(\pi)}$ encoded as a CNF using the bamboo-tree encoding. Given $\pi \in \{0, 1\}^n$, we choose the matrices $A^{(\pi)}$ such that $A_{1,1}^{(\pi)} = f(\pi)$ the rest of the matrices (for $|\pi| < n$) are arranged as $A^{(\pi)} = \begin{bmatrix} X^{(\pi*0)} & X^{(\pi*1)} & Y^{(\pi*0)T} & Y^{(\pi*1)T} \end{bmatrix}$. That is, given a copy $X^{(\pi)} Y^{(\pi)} = A^{(\pi)}$, we can compute both copies $X^{(\pi*0)} Y^{(\pi*0)} = A^{(\pi*0)}$ and $X^{(\pi*1)} Y^{(\pi*1)} = A^{(\pi*1)}$ recursively. This way, starting with the empty string “()”, we can gradually compute $A^{(\pi)}$ for $\pi \in \{0, 1\}^n$, thus obtaining $f(\pi)$. This computation can be naturally expressed as a Boolean circuit, using the gates that correspond to the extension variables.

2.4 On the Strength of WRank

Finally, we initiate the study of the weak rank principle as an axiom in bounded arithmetic. Linear algebra is used extensively throughout combinatorics and complexity theory, resulting in a variety of breakthroughs such as the resolution of the finite field Kakeya conjecture [29], Cap Set problem [24, 30] and the Sensitivity Conjecture [40]. Indeed, there is a whole book dedicated to linear algebraic methods in combinatorics [11].

Soltys and Cook [80] initiated the study of formal theories that “capture linear algebra” from the complexity perspective. Specifically, they considered a few formal theories for linear algebraic reasoning and showed that $\forall\text{LAP}$, the strongest theory considered in [80] which incorporates matrix powering and certain induction schema, proves many linear algebraic identities including the (strong) rank principle. Subsequently, Tzameret and Cook [84] showed that the seemingly weaker theory VNC^2 suffices for linear algebraic reasoning and specifically proving properties of the determinant. We take a somewhat similar route, but instead of powerful concepts such as matrix powering (which is DET -complete [21] and hence intuitively captures “full” linear algebra), we only consider the *weak* rank principle, hence our theories appear to be somewhat weaker.

An important motivation for considering the *weak* rank principle is that they capture arguments in combinatorics and complexity theory that are “loose” (i.e., not tight). Intuitively, the *weak* rank principle should already suffice when one is satisfied with a bound that is only tight within multiplicative

⁴Unfortunately, when measured by *size*, $\text{PCR}_{\mathbb{F}_2}$ is not closed under low-degree reductions. Hence our concrete lower bounds are not direct corollaries of [Theorem 2.7](#) and require more work.

factors. This is the case when we are satisfied with, say, a bound of $c \cdot 2.756^n$ for the Cap Set problem or a bound of $c \cdot 2^{n^{1/2d}}$ for $\text{AC}^0[p]$ circuit complexity, where $c > 0$ is an unspecified constant. This is similar to the case of approximate counting where we only care about the sizes of sets *approximately*, hence the *weak* pigeonhole principle is already applicable [46].

In this work, we put forward a theory $\overline{\mathbf{V}^0(p)} + \text{WRank}_p$, which incorporates the two-sorted theory $\overline{\mathbf{V}^0(p)}$ in the style of Cook–Nguyen [20] with an extra axiom WRank_p expressing the weak rank principle over $\mathbb{F}_p$. We demonstrate the power of the weak rank principle and confirm the above intuition by showing that this theory can prove Smolensky’s circuit lower bound against $\text{AC}^0[p]$ [78]:

Theorem 2.9 (Informal; Smolensky’s lower bound is implied by the weak rank principle; see [Theorem 9.12](#)). *For every prime number $p > 2$, $\overline{\mathbf{V}^0(p)} + \text{WRank}_p$ proves that MOD_2 cannot be computed by $\text{AC}^0[p]$ circuits of depth d and size $2^{o(n^{1/(2d)})}$.*

Previously, Müller and Pich [59] formalised Smolensky’s $\text{AC}^0[p]$ lower bounds in Jeřábek’s theory APC_1 [44, 46]. We point out a few differences between our formalisation of Smolensky’s lower bounds and the formalisation in [59].

1. Müller and Pich proved a lower bound against circuits of size $n^{\log n}$ in the (single-sorted) theory APC_1 with the assumption that $n^{\log^{9d} n} \in \text{Log}$. This allows randomised quasi-polynomial-time reasoning over the purported $\text{AC}^0[p]$ circuit. We prove a lower bound against size $2^{\Omega(n^{1/2d})}$ in the (two-sorted) theory $\overline{\mathbf{V}^0(p)} + \text{WRank}$ where the circuit is represented as an object of string sort. Except for the weak rank principle, this theory only allows (quasi-polynomial-size) $\text{AC}^0[p]$ -reasoning over the purported $\text{AC}^0[p]$ circuit.
2. One drawback in the Müller–Pich proof is the requirement that $n^{\log^{9d} n} \in \text{Log}$. In fact, this allows them to prove the required rank principle by *brute force*. Our formalisation explicitly separates the weak rank principles from other $\text{AC}^0[p]$ -manipulations.

We also note that Razborov mentioned in [71, Section E.3] that the metamathematics of Razborov–Smolensky lower bounds are left open. Our work addresses this gap by showing that Smolensky’s lower bounds can be proved in $\overline{\mathbf{V}^0(p)} + \text{WRank}_p$. It is also worth mentioning the already Galesi *et al.* [31] considered the rank principle with $A = I$ as an axiom added to the polynomial calculus.

3 Conclusions and Open Problems

Our work introduces the weak rank principle as a new object of study in proof complexity and applies it to advance the *generator approach* to proof complexity, introduced by Krajíček [51, 52] and Alekhovich *et al.* [4], and later emphasized by Razborov [76, p. 417]. The goal of this approach is to construct *proof-complexity generators* for as many proof systems as possible, where the hardness of the generator reflects the computational limitations of the corresponding system. In our case, the hardness originates from linear-algebraic reasoning: the weak rank principle expresses the unsatisfiability of certain linear systems and matrix identities, and its difficulty is directly tied to the inability of a proof system to efficiently formalise arguments involving matrix rank. Consequently, the weak rank principle is expected to be hard for proof systems that cannot efficiently reason about linear algebra.

Future directions. Several directions appear promising for further study:

- The weak rank principle provides a useful family of tautologies to study in proof systems for which the pigeonhole principle is easy or for which no proof-complexity generators are known. The basic instance, $XY = I$, already gives a concrete candidate for proving lower bounds. Obtaining a lower bound for this instance is the first step toward establishing lower bounds for the more general

case $XY = A$, where A ranges over matrices of rank greater than n . Such results would yield a proof-complexity generator based on WRank. Once such a base generator is obtained, it can be iterated to obtain a *function generator*, showing that the proof system in question cannot efficiently prove certain complexity-class separations such as $\text{NP} \not\subseteq \text{P/poly}$.

- Develop bounded arithmetic theories that use WRank as axioms, such as the theory $\overline{\mathbf{V}^0(p)} + \text{WRank}_p$ developed in this paper. We speculate that many interesting results in combinatorics and complexity theory can be formalised in such theories and that the *weak* rank principles suffice. Our formalisation of Smolensky's lower bounds [78] can be seen as the first step towards this goal.
- The weak rank principle provides a natural candidate for lower bounds against $\text{AC}^0[p]$ -Frege (for which no lower bounds are currently known) and SoS. This connection follows from the conjecture that NC^2 captures the computational complexity of linear algebra. According to the informal correspondence between circuit classes and propositional proof systems whose proof-lines come from those classes, one expects that proofs operating with NC^2 -circuits (that is, NC^2 -Frege) are exactly those capable of efficiently reasoning about rank-based arguments. In contrast, the weak pigeonhole principle admits quasipolynomial-size proofs already in $\text{Res}(\log n)$ and thus in AC^0 -Frege [56, 64], indicating that stronger algebraic principles are required to obtain meaningful lower bounds.

4 Preliminaries

We will use the following common notations throughout the paper.

- For all $k > 0$, I_k denotes the $k \times k$ identity matrix $(I_k)_{i,j} = \delta_{i,j}$ for all $i, j \in [k]$.
- For all $k, \ell > 0$, $J_{k,\ell}$ denotes the $k \times \ell$ all-ones matrix, i.e., $(J_{k,\ell})_{i,j} = 1$ for all $i \in [k], j \in [\ell]$.

4.1 Concentration Bounds

Many of our constructions use probabilistic arguments and rely on concentration bounds. We first recall a standard version of the Chernoff bound, which can be found in many textbooks; see, for example, [58, Corollary 4.6].

Theorem 4.1 (Chernoff bound). *Let $X_1, \dots, X_n$ be independent Bernoulli random variables with $\Pr[X_i = 1] = p_i$ for all $i \in [n]$. Let $X = \sum_{i \in [n]} X_i$ and $\mu = \mathbb{E} X$. Then, for every $0 < \delta < 1$,*

$$\Pr[X \geq (1 + \delta)\mu] \leq \exp(-\delta^2\mu/3)$$

and

$$\Pr[X \leq (1 - \delta)\mu] \leq \exp(-\delta^2\mu/2).$$

Consequently,

$$\Pr[|X - \mu| \geq \delta\mu] \leq 2 \exp(-\delta^2\mu/3).$$

A useful special case of the Chernoff bound is sampling with replacement. Given a positive integer N and a subset $K \subseteq [N]$, we uniformly draw n samples from $[N]$ with replacement. For every $i \in [n]$, let X_i be the indicator variable of the event that the i th sample is in K . Then $X = \sum_{i \in [n]} X_i$ counts the number of samples that are in K . Theorem 4.1 gives a concentration bound for X around its expectation $\mu = n \cdot |K|/N$.

We also use the analogous bound for sampling without replacement. The following result follows, for example, from Hoeffding's theorem for samples drawn without replacement; see [38, Theorem 4]. In this setting, the number of samples belonging to the set K has a hypergeometric distribution.

Theorem 4.2 (Hypergeometric Chernoff bound). *Let $K \subseteq [N]$ be fixed, and let S be an s -subset of $[N]$ chosen uniformly at random. Let $X = |K \cap S|$ and $\mu = \mathbb{E} X$. Then, for every $0 < \delta < 1$,*

$$\Pr[X \geq (1 + \delta)\mu] \leq \exp(-\delta^2 \mu / 3)$$

and

$$\Pr[X \leq (1 - \delta)\mu] \leq \exp(-\delta^2 \mu / 2).$$

Consequently,

$$\Pr[|X - \mu| \geq \delta\mu] \leq 2 \exp(-\delta^2 \mu / 3).$$

4.2 Proof Systems

Let **TAUT** denote the **coNP**-complete set of propositional tautologies. Cook and Reckhow [23] defined *propositional proof systems* as nondeterministic algorithms for **TAUT**. Equivalently, a propositional proof system $\mathcal{P}(\phi, \Pi)$ is a deterministic algorithm that takes as input a propositional formula ϕ and a purported proof Π , and satisfies the following properties:

- **Completeness:** If $\phi \in \mathbf{TAUT}$, then there exists a proof Π such that $\mathcal{P}(\phi, \Pi)$ accepts.
- **Soundness:** If $\phi \notin \mathbf{TAUT}$, then for every Π , $\mathcal{P}(\phi, \Pi)$ rejects.
- **Efficiency:** $\mathcal{P}$ runs in time polynomial in $|\phi| + |\Pi|$.

A propositional proof system $\mathcal{P}$ is *polynomially bounded* (or *p-bounded*) if there is a polynomial $p(\cdot)$ such that every tautology $\phi \in \mathbf{TAUT}$ admits a proof Π of length at most $p(|\phi|)$ for which $\mathcal{P}(\phi, \Pi)$ accepts. Cook and Reckhow [23] showed that p-bounded propositional proof systems exist if and only if $\mathbf{NP} = \mathbf{coNP}$.

We will sometimes equivalently consider *refutation systems*, which are nondeterministic algorithms for the language $\overline{\mathbf{SAT}}$. A DNF F is a tautology if and only if $\neg F$ (which can be expressed as a CNF) is unsatisfiable. Thus, a *proof* of F can be equivalently viewed as a *refutation* of $\neg F$.

While traditional proof systems reason about Boolean formulas, *algebraic proof systems* reason about the unsatisfiability of collections of *polynomial equations*. Fix an underlying field $\mathbb{F}$, and let $f_1, f_2, \dots, f_m \in \mathbb{F}[x_1, x_2, \dots, x_n]$ be a collection of polynomials, where each f_i is called an *axiom* and represents the equation $f_i(\vec{x}) = 0$. The goal is to prove that these polynomials do not have a common zero over $\mathbb{F}^n$. It is standard in the literature to add the *Boolean axioms* $x_i^2 - x_i$ to the collection to enforce that all variables take values in $\{0, 1\}$; however, we do *not* include these Boolean axioms in our definition of rank principles.

For a polynomial f , let $\|f\|$ denote the number of monomials appearing with a non-zero coefficient in f .

Definition 4.3. Let $F = (f_1, f_2, \dots, f_m)$ be an unsatisfiable collection of polynomial equations over a field $\mathbb{F}$.

- A *Nullstellensatz* ($\mathbf{NS}_{\mathbb{F}}$) [13] refutation of F is a sequence of polynomials $(g_1, g_2, \dots, g_m)$ such that the following polynomial identity holds:

$$\sum_{i=1}^m f_i \cdot g_i = 1.$$

The *degree* of the refutation is $\max_{i \in [m]} \{\deg(f_i \cdot g_i)\}$. The *size* of the refutation is $\sum_{i=1}^m \|f_i\| \cdot \|g_i\|$.

- *Polynomial Calculus* ($\text{PC}_{\mathbb{F}}$) [19] is a dynamic proof system where proof lines are polynomials over $\mathbb{F}$. Starting from $(f_1, f_2, \dots, f_m)$ as axioms, this system has two inference rules:

$$\frac{f}{\alpha f + \beta g} \text{ (for } \alpha, \beta \in \mathbb{F} \text{)} \quad \text{and} \quad \frac{f}{x_i \cdot f}.$$

A $\text{PC}_{\mathbb{F}}$ *refutation* of F is a derivation of 1 from F . The *degree* of the refutation is the maximum degree of any polynomial appearing in it. The *size* of the refutation is the total number of monomials appearing in it.

- Suppose $\mathbb{F} = \mathbb{R}$. A *Sum-of-Squares* (SoS) [34, 62] refutation of F is a sequence of polynomials $(g_1, g_2, \dots, g_m)$ and $(h_1, h_2, \dots, h_\ell)$ such that the following polynomial identity holds:

$$\sum_{i=1}^m f_i \cdot g_i + \sum_{j=1}^{\ell} h_j^2 = -1.$$

The *degree* of the refutation is the maximum of $\max_{i \in [m]} \{\deg(f_i \cdot g_i)\}$ and $2 \max_{j \in [\ell]} \{\deg(h_j)\}$. The *size* of the refutation is $\sum_{i \in [m]} \|f_i\| \cdot \|g_i\| + \sum_{j \in [\ell]} \|h_j\|$.

Definition 4.4. Let $F = (f_1, f_2, \dots, f_m)$ be an unsatisfiable collection of polynomial equations over $\mathbb{R}$. We enforce that the variables $x_1, x_2, \dots, x_n$ are Boolean by adding every Boolean axiom $x_i^2 - x_i$ to F . A *Sherali–Adams* (SA) [25] refutation of F is a sequence of polynomials $(g_1, g_2, \dots, g_m)$ and a conical junta $\mathcal{J}$ such that the following polynomial identity holds.

$$\sum_{i \in [m]} f_i \cdot g_i + \mathcal{J} = -1.$$

Here, $\mathcal{J}$ is a *conical junta* if it is a nonnegative linear combination of terms ($\mathcal{J} = \sum_{j=1}^{\ell} \alpha_j t_j(x)$ for $\alpha_j \geq 0$), and each term is a product of factors of the form x_i and $(1 - x_j)$. The *degree* of the refutation is the maximum of $\deg(\mathcal{J})$ and $\max_i \{\deg(f_i) + \deg(g_i)\}$. The *size* of the refutation is $\sum_{i \in [m]} \|f_i\| \cdot \|g_i\| + \ell$.

Definition 4.5. Let $F = (f_1, f_2, \dots, f_m)$ be an unsatisfiable collection of polynomial equations over a field $\mathbb{F}$ with underlying variables $x_1, x_2, \dots, x_n$. We enforce that the variables are Boolean by adding every Boolean axiom $x_i^2 - x_i$ to F . A *Polynomial Calculus Resolution* ($\text{PCR}_{\mathbb{F}}$) [3] refutation of F is a $\text{PC}_{\mathbb{F}}$ refutation of F' , where F' is F with the *twin variables* $\bar{x}_1, \bar{x}_2, \dots, \bar{x}_n$ and the axioms $1 - x_i - \bar{x}_i$ and $\bar{x}_i^2 - \bar{x}_i$ added.

Note that $\text{PC}_{\mathbb{F}}$ and $\text{PCR}_{\mathbb{F}}$ are equivalent w.r.t. the degree measure, but a collection of polynomial equations might have much more efficient $\text{PCR}_{\mathbb{F}}$ proofs than $\text{PC}_{\mathbb{F}}$ proofs in terms of size.

4.3 Substitutions

In the setting of proof complexity, substitutions are usually defined as follows.

Definition 4.6 (Restrictions). Given variables $(x_1, \dots, x_n)$ and their twins $(\bar{x}_1, \dots, \bar{x}_n)$, the *restriction* ρ is a partial mapping from $\{x_1, \dots, x_n\} \cup \{\bar{x}_1, \dots, \bar{x}_n\}$ to $\{0, 1\}$.

Here we consider a generalised notion of substitutions in which we can substitute variables with fresh variables, possibly negated; several variables may be mapped to the same fresh variable.

Definition 4.7 (Substitutions). Given variables $(x_1, \dots, x_n)$ and $(y_1, \dots, y_m)$ and their twins $(\bar{x}_1, \dots, \bar{x}_n)$ and $(\bar{y}_1, \dots, \bar{y}_m)$, the *substitution* ρ is a mapping from $\{x_1, \dots, x_n\} \cup \{\bar{x}_1, \dots, \bar{x}_n\}$ to $\{0, 1\} \cup \{y_1, \dots, y_m\} \cup \{\bar{y}_1, \dots, \bar{y}_m\}$, which is consistent w.r.t. the twin variables: if $\rho(v) = b$ for some $b \in \{0, 1\}$, then $\rho(\bar{v}) = 1 - b$; if $\rho(v) = u$ for some variable u , then $\rho(\bar{v}) = \bar{u}$.

It coincides with the standard definition when $(y_1, \dots, y_m) = (x_1, \dots, x_n)$ and for every v with $\rho(v) \notin \{0, 1\}$, $\rho(v) = v$.

Given a polynomial $f \in \mathbb{F}[x_1, \dots, x_n, \bar{x}_1, \dots, \bar{x}_n]$ and a substitution ρ , the application of ρ to f , denoted $f \upharpoonright \rho$, is the polynomial $g(y_1, \dots, y_m, \bar{y}_1, \dots, \bar{y}_m) = f(\rho(x_1), \dots, \rho(x_n), \rho(\bar{x}_1), \dots, \rho(\bar{x}_n))$.

Given a proof Π (a sequence of polynomials), the application of a substitution ρ to Π , denoted $\Pi \upharpoonright \rho$, is obtained by applying the ρ to each line p of Π and removing any line that becomes the zero polynomial. Applying a substitution to a $\text{PCR}_{\mathbb{F}}$, SA , or SoS proof results in a proof with at most a linear increase in size: this is because [Definition 4.7](#) substitutes variables and constants for variables; when we replace a variable with a polynomial, size is not necessarily preserved.

4.4 Reductions in Proof Complexity

Let F and G be two collections of polynomial equations. In this paper, we use *low-degree algebraic reductions* [17, 27] to formalise the idea that F is “easier” to prove than G . Intuitively, this happens when F is a substitution instance of G , namely there is a substitution R to G under which $G \circ R$ equals F . More precisely, we require that there is an efficient (e.g., low degree) derivation from F of the substitution instance $G \circ R$. Hence, a short proof of $G \circ R$ can be transformed into a short proof of F . For example, to show that *circuit lower bounds* are as hard to prove as the *weak rank principle*, we exhibit a low-degree reduction from the polynomial equations encoding the weak rank principle to those encoding circuit lower bounds.

Definition 4.8 (Algebraic reductions; [27, Def. 8.2]). Let $F = (f_1, \dots, f_m)$ be a collection of polynomials over variables $(x_1, \dots, x'_m)$, and $G = (g_1, \dots, g_n)$ be a collection of polynomials over variables $(y_1, \dots, y_{n'})$. A *degree- d algebraic reduction*, denoted $F \leq_d^{\text{alg}} G$, consists of the following:

- (Reduction)** a mapping $R: \mathbb{F}^{m'} \rightarrow \mathbb{F}^{n'}$ such that each output element $R_i \in \mathbb{F}[x_1, \dots, x_{m'}]$, $i \in [n']$, is a polynomial of degree at most d in $(x_1, \dots, x_{m'})$;
- (Axioms)** for every $i \in [n]$, the polynomial $g_i \circ R$ has an $\text{NS}_{\mathbb{F}}$ -proof from F of degree $d \cdot \deg(g_i)$. That is, there exist polynomials $(q_{i,1}, \dots, q_{i,m})$ such that

$$g_i \circ R = \sum_{k=1}^m f_k q_{i,k},$$

and for every $k \in [m]$, $\deg(f_k) + \deg(q_{i,k}) \leq d \cdot \deg(g_i)$.

It is a standard and useful fact that algebraic reductions are closed under composition.

Fact 4.9. Let $F = (f_1, \dots, f_m)$, $G = (g_1, \dots, g_n)$, and $H = (h_1, \dots, h_\ell)$ be systems of polynomial equations. If $F \leq_{d_1}^{\text{alg}} G$ and $G \leq_{d_2}^{\text{alg}} H$, then $F \leq_{d_1 d_2}^{\text{alg}} H$.

Proof. Suppose that F is over variables $(x_1, \dots, x_{m'})$, G is over variables $(y_1, \dots, y_{n'})$, and H is over variables $(z_1, \dots, z_{\ell'})$. Let $R^1: \mathbb{F}^{m'} \rightarrow \mathbb{F}^{n'}$ be a degree- d_1 reduction from F to G and $R^2: \mathbb{F}^{n'} \rightarrow \mathbb{F}^{\ell'}$ be a degree- d_2 reduction from G to H . We claim that $R^2 \circ R^1: \mathbb{F}^{m'} \rightarrow \mathbb{F}^{\ell'}$ is a degree- $d_1 d_2$ reduction from F to H . Clearly, every output element of $R^2 \circ R^1$ is a polynomial of degree at most $d_1 d_2$ in $(x_1, \dots, x_{m'})$.

Now consider an axiom h_i from H . We want to show that $h_i \circ R^2 \circ R^1$ has an NS -proof from F of degree $d_1 d_2 \cdot \deg(h_i)$. We have

$$h_i \circ R^2 = \sum_{k=1}^n g_k q_{i,k}$$

for some polynomials $(q_{i,1}, \dots, q_{i,n})$ such that for every $k \in [n]$, $\deg(g_k) + \deg(q_{i,k}) \leq d_2 \cdot \deg(h_i)$. For each $k \in [n]$, we also have

$$g_k \circ R^1 = \sum_{j=1}^m f_j p_{k,j}$$

for some polynomials $(p_{k,1}, \dots, p_{k,m})$ such that for every $j \in [m]$, $\deg(f_j) + \deg(p_{k,j}) \leq d_1 \cdot \deg(g_k)$. Hence

$$h_i \circ R^2 \circ R^1 = \sum_{k=1}^n (g_k \circ R^1)(q_{i,k} \circ R^1) = \sum_{j=1}^m f_j \cdot \underbrace{\left(\sum_{k=1}^n p_{k,j} \cdot (q_{i,k} \circ R^1) \right)}_{r_{i,j}}. \quad (2)$$

Note that $\deg(f_j) + \deg(r_{i,j}) \leq \max_k \{\deg(f_j) + \deg(p_{k,j}) + \deg(q_{i,k}) \cdot d_1\} \leq \max_k \{d_1 \cdot (\deg(g_k) + \deg(q_{i,k}))\} \leq d_1 d_2 \cdot \deg(h_i)$. Hence, (2) is an NS-proof of $h_i \circ R^2 \circ R^1$ from F that has degree $d_1 d_2 \cdot \deg(h_i)$.

In conclusion, $R^2 \circ R^1$ is a degree- $d_1 d_2$ algebraic reduction from F to H . $\square$

Algebraic proof systems such as Nullstellensatz and Polynomial Calculus are closed under low-degree algebraic reductions when the efficiency measure is degree [27, Lemma 8.3]. This paper exhibits many low-degree algebraic reductions; for example, Corollary 8.4 shows a low-degree algebraic reduction from the weak rank principles to circuit lower bound statements. Such a result can be interpreted as follows: for any algebraic proof system $\mathcal{P}$ closed under low-degree algebraic reductions, if $\mathcal{P}$ cannot prove the weak rank principles, then $\mathcal{P}$ cannot prove any circuit lower bounds.

4.5 The Pigeonhole Principle

We define the *algebraic* encoding of the pigeonhole principle, first considered by Razborov in [72]—this version is sometimes explicitly called *functional* because of the presence of the Functionality Axioms below. Fix positive integers m and n with $m > n$. FPHP_n^m is defined over mn variables arranged as a Boolean matrix $X = (x_{i,k})_{i \in [m], k \in [n]}$. It consists of the following polynomial equations:

Pigeon Axioms: For $i \in [m]$, the polynomial

$$1 - \sum_{k \in [n]} x_{i,k}. \quad (3)$$

Hole Axioms: For $i, j \in [m], i \neq j, k \in [n]$, the polynomial

$$x_{i,k} x_{j,k}. \quad (4)$$

Functionality Axioms: For $i \in [m], k, \ell \in [n], k \neq \ell$, the polynomial

$$x_{i,k} x_{i,\ell}. \quad (5)$$

Boolean Axioms: For $i \in [m], k \in [n]$, the polynomial

$$x_{i,k}^2 - x_{i,k}. \quad (6)$$

We comment that the standard (non-functional) PHP that maps m pigeons to n holes is defined the same as FPHP_n^m , only excluding the Functionality Axioms (5).

Razborov [72] proved a $\text{PC}_{\mathbb{F}}$ degree lower bound on FPHP_n^m for arbitrary $m > n$.

Theorem 4.10 ([72]). *For every field $\mathbb{F}$, every $\text{PC}_{\mathbb{F}}$ refutation of FPHP_n^m has a term that mentions at least $n/2 + 1$ pigeons and at least $n/2 + 1$ holes.⁵*

Although not stated explicitly, the lower bound in [72] holds when the refutation is considered modulo the ideal generated by the Functionality and Hole Axioms. This implies that every refutation of FPHP_n^m contains a term $\prod_{\nu \in [\ell]} x_{i_\nu, k_\nu}$, where $\ell \geq n/2 + 1$, all i_ν are distinct and all k_ν are distinct.

There is a similar degree lower bound for SA. It was originally stated for FPHP_n^{n+1} , but exactly the same proof applies to FPHP_n^m .

Theorem 4.11 ([26]). *Every SA refutation of FPHP_n^m contains a term that mentions at least $n - 1$ pigeons.*

4.6 The Perfect Matching Principle

The perfect matching principle, denoted Count_n^2 , states that it is possible to find a perfect matching on n elements when only when n is even ([1, 63]). Formally, Count_n^2 is defined as a CNF over the variables $y_{\{k, \ell\}} : \{k, \ell\} \in \binom{[n]}{2}$. These variables give a partition of the elements into pairs. Each variable $y_{\{k, \ell\}}$ is true if and only if the elements k and ℓ are matched together. We assert that each element is matched with some other element by the following axioms:

$$\bigvee_{\ell \in [n] \setminus \{k\}} y_{\{k, \ell\}}, \quad k \in [n].$$

We also assert that no element is matched with more than one other element by the following axioms:

$$\neg y_{\{k, \ell\}} \vee \neg y_{\{k, \ell'\}}, \quad k \in [n], \ell, \ell' \in [n] \setminus \{k\}, \ell \neq \ell'.$$

When n is odd, it is clear that this formula is unsatisfiable.

4.7 Lossless Expanders

Let G be an $m \times n$ bipartite graph with parts L and R .

Definition 4.12. Given $i \in L$, the *neighbourhood* of i in G , denoted $\mathcal{N}_G(i)$, is the set of all $k \in R$ that are connected to i in G .

Definition 4.13. Given a set $I \subseteq L$, the *boundary* of I , denoted $\partial_G(I)$, is the set of all unique neighbours of I :

$$\partial_G(I) = \{k \in R : \text{there is exactly one } i \in I \text{ satisfying } k \in \mathcal{N}_G(i)\}.$$

Definition 4.14. G is an (r, c) -*boundary expander* if, for every set $I \subseteq L$ with $|I| \leq r$,

$$|\partial_G(I)| \geq c|I|.$$

Definition 4.15. G is an (r, d) -*lossless expander* if, for every set $I \subseteq L$ with $|I| \leq r$,

$$\sum_{i \in I} |\mathcal{N}_G(i)| - |\partial_G(I)| \leq d|I|. \quad (7)$$

This is called a *lossless* expander because the expansion from the left side to the right side preserves (almost) all edges without significant loss of distinctness. The following fact immediately follows from this definition.

⁵In fact, this term encodes a matching of size $n/2 + 1$.

Fact 4.16. Let $\Delta_{\min}$ be the minimum left-degree of the vertices in G . If G is an (r, d) -lossless expander, then it is also an (r, c) -boundary expander with $c = \Delta_{\min} - d$.

Lemma 4.17. Let $C > 1$. There exist constants $\delta, d > 0$ such that for every large enough positive integers m, n with $n < m < Cn$, there is an $m \times n$ bipartite graph G satisfying the following:

1. the minimum left-degree of G is at least $\Delta_{\min} = 50 \log n$,
2. the maximum left-degree of G is at most $\Delta_{\max} = 150 \log n$,
3. G is an (n^δ, d) -lossless expander.

In particular, G is an $(n^\delta, \Delta_{\min} - d)$ -boundary expander.

Let G be a bipartite graph with parts L and R and edges E . Given $K \subseteq R$, we define $G \setminus K$ as the bipartite graph with parts L and $R \setminus K$ and edges $\{(i, k) \in E : i \in L, k \in R \setminus K\}$. Observe that for every $i \in L$, $\mathcal{N}_{G \setminus K}(i) = \mathcal{N}_G(i) \setminus K$.

Lemma 4.18. Let $m = O(n)$ and let G be an $m \times n$ (n^δ, d) -lossless expander with parts L and R , minimum left-degree at least $\Delta_{\min} = 50 \log n$ and maximum left-degree at most $\Delta_{\max} = 150 \log n$. Let K be a random subset of R such that every $j \in R$ is added to K independently with probability $2/3$. With probability $1 - o(1)$, the graph $G \setminus K$ is an (n^δ, d) -lossless expander with a minimum left-degree $\Delta_{\min}/6$. In particular, $G \setminus K$ is an (n^δ, c) -boundary expander, where $c = \Delta_{\min}/6 - d$.

The proofs of these lemmas can be found in [Appendix A](#).

4.8 The Rank Principle and Oddtown Principle

In this section, we introduce our main instance. Here we focus on their purely algebraic formulation, as it is the simplest and most transparent. Later, we will also consider its CNF encodings.

Fix a field $\mathbb{F}$ and positive integers m and n with $m > n$. We denote by $\text{Rank}_n^m(A)$ the system of degree-2 polynomial equations stating that an $m \times m$ matrix A can be factorised as the product of an $m \times n$ matrix and an $n \times m$ matrix over $\mathbb{F}$. More precisely, the *rank principle* $\text{Rank}_n^m(A)$ is defined over $2mn$ variables arranged into two matrices $X \in \mathbb{F}^{m \times n}$ and $Y \in \mathbb{F}^{n \times m}$. For every $i, j \in [m]$, there is an equation in $\text{Rank}_n^m(A)$ stating that the (i, j) th entry of the product XY is equal to $A_{i,j}$. That is:

$$A_{i,j} - \sum_{k \in [n]} x_{i,k} y_{k,j}, \quad i, j \in [m]. \quad (8)$$

Every equation in Rank_n^m has degree 2 in the variables of X and Y . It follows from basic linear algebra that when the rank of A exceeds n , $\text{Rank}_n^m(A)$ is unsatisfiable.

When $A = I_m$, the $m \times m$ identity matrix, we may omit the matrix from the notation and simply write Rank_n^m .

We also need the *oddtown principle* $\text{OT}_n^m(A)$, which is the special case of $\text{Rank}_n^m(A)$ in which $Y = X^T$. That is, given a symmetric matrix $A \in \mathbb{F}^{m \times m}$, $\text{OT}_n^m(A)$ is defined over mn variables arranged as a matrix $X \in \mathbb{F}^{m \times n}$, and for every $i, j \in [m]$, there is an equation in OT_n^m stating that the (i, j) th entry of XX^T is equal to $A_{i,j}$:

$$A_{i,j} - \sum_{k \in [n]} x_{i,k} x_{j,k}, \quad i, j \in [m], i \leq j. \quad (9)$$

5 Proof Complexity Generators for $\text{PCR}_{\mathbb{F}_2}$

In this section, we study the following generator and show that, under various encodings, it is a proof complexity generator hard for $\text{PCR}_{\mathbb{F}_2}$.

The Rank Generator g

Let $m > n$.

Input: Two matrices $X \in \mathbb{F}^{m \times n}$ and $Y \in \mathbb{F}^{n \times m}$.

Output: Their product $XY \in \mathbb{F}^{m \times m}$.

Note that $g: \mathbb{F}^{2mn} \rightarrow \mathbb{F}^{m^2}$, so this function is stretching whenever $n < \frac{m}{2}$. This function was considered by Krajíček [54] as a proof complexity generator for algebraic proof systems based on the pigeonhole principle. It was also considered as a candidate hitting-set generator in algebraic complexity and polynomial identity testing by Andrews and Forbes [6, 7].

Given an $m \times m$ matrix A over $\mathbb{F}_2$, we consider different encodings of the statement “ $A \notin \text{Range}(g)$ ”.

In Section 5.1, we show the hardness of the algebraic encoding $\text{Rank}_n^m(A)$ for $\text{PCR}_{\mathbb{F}_2}$.

In Sections 5.2 and 5.3, we consider two Boolean encodings of the statement “ $A \notin \text{Range}(g)$ ”. The first encoding, $\text{PMRank}_n^m(A)$, employs the perfect matching principle to encode parities. The second encoding, $\text{BTRank}_n^m(A)$, uses extension variables arranged in unbalanced binary trees (bamboos). For the perfect matching encoding, when working over $\mathbb{F}_2$ and specifically $\text{PCR}_{\mathbb{F}_2}$, we consider only the case $A = I$ and establish its hardness for $\text{PCR}_{\mathbb{F}_2}$. For $\text{PCR}_{\mathbb{F}}$ with $\text{char}(\mathbb{F}) \neq 2$, SoS, and AC^0 -Frege, we consider $\text{PMRank}_n^m(A)$ with any matrix A and reduce its complexity to the perfect matching principle Count_n^2 , which is hard for these proof systems. For the bamboo-tree encoding $\text{BTRank}_n^m(A)$, we establish hardness for $\text{PCR}_{\mathbb{F}_2}$ for arbitrary matrices A .

In Section 5.4, we show that certain iterated variants of Rank and BTRank are hard for $\text{PCR}_{\mathbb{F}_2}$. Iteration links multiple copies of the rank generator by allowing the outputs of one copy to be used as inputs to copies at later levels. This amplifies the stretch while preserving the structure of the original rank principle.

In this section, we consider PCR in which every proof line is a multilinear polynomial. Thus the multiplication rule consists of multiplying by a variable followed by multilinearisation. Equivalently, we work in the quotient ring $\mathbb{F}_2[\vec{x}]/(\vec{x}^2 - \vec{x})$ factorised over the Boolean axioms for all variables (including the twin variables). This representation increases neither the monomial-size nor the degree of PCR refutations; see [19].

5.1 Algebraic Rank Principle

In this section, we use a special notion of degree, called *column degree*, which counts the number of distinct *columns* mentioned by a term in a refutation of $\text{Rank}_n^m(A)$. This is formally defined as follows.

Definition 5.1 (Column-mentioned variables). Let t be a term in the variables of $\text{Rank}_n^m(A)$. We say that $k \in [n]$ is *X-column-mentioned in t* if there is $i \in [m]$ such that $x_{i,k}$ or $\bar{x}_{i,k}$ appears in t . We say that $k \in [n]$ is *Y^T -column-mentioned in t* if there is $j \in [m]$ such that $y_{k,j}$ or $\bar{y}_{k,j}$ appears in t .

Definition 5.2 (Column degree). Let t be a term in the variables of $\text{Rank}_n^m(A)$. Let K be the set of all $k \in [n]$ that are either *X-column-mentioned* or *Y^T -column-mentioned* in t . We say that t *mentions columns K* . The *column degree* of t is defined as $|K|$.

We start by proving a lower bound on the column degree of $\text{PCR}_{\mathbb{F}_2}$ refutations of $\text{Rank}_n^m(A)$.

Lemma 5.3. *Suppose that $m > n > 0$ and $A \in \{0, 1\}^{m \times m}$ is an $m \times m$ Boolean matrix. Then any $\text{PCR}_{\mathbb{F}_2}$ refutation of $\text{Rank}_n^m(A)$ requires column degree at least $n/2 + 1$.*

Proof. We reduce the functional pigeonhole principle FPHP_n^m to this instance (in the sense of Definition 4.8); namely, we show that there is a substitution τ under which $\text{Rank}_n^m(A) \upharpoonright \tau$ is efficiently derivable from the axioms of FPHP_n^m . The reduction τ is defined as follows.

Let $P = (p_{i,k})_{i \in [m], k \in [n]}$ be fresh variables. It is useful to think of P as an $m \times n$ variable matrix.

1. For each $i \in [m]$ and $k \in [n]$, set

$$\tau(x_{i,k}) = (AP)_{i,k} = \sum_{s \in [m]} A_{i,s} p_{s,k}.$$

2. For each $j \in [m]$ and $k \in [n]$, set

$$\tau(y_{k,j}) = p_{j,k}.$$

Observe that $\tau(x_{i,k})$ and $\tau(y_{k,j})$ only use the variables of the form $p_{s,k} : s \in [m]$.

Let $i, j \in [m]$. The (i, j) th axiom of $\text{Rank}_n^m(A)$ under τ transforms as follows (recall that $A_{i,j} = -A_{i,j}$ over $\mathbb{F}_2$):

$$\begin{aligned} \left(\sum_{k \in [n]} x_{i,k} y_{k,j} + A_{i,j} \right) \upharpoonright \tau &= \sum_{k \in [n]} \left(\sum_{s \in [m]} A_{i,s} p_{s,k} \right) p_{j,k} + A_{i,j} \\ &= \sum_{s \in [m]} A_{i,s} \left(\sum_{k \in [n]} p_{s,k} p_{j,k} \right) + A_{i,j} \\ &= \sum_{s \in [m] \setminus \{j\}} A_{i,s} \left(\sum_{k \in [n]} p_{s,k} p_{j,k} \right) + A_{i,j} \left(\sum_{k \in [n]} p_{j,k} + 1 \right), \end{aligned}$$

which is a degree-2 $\text{NS}_{\mathbb{F}_2}$ derivation from the Hole Axioms $p_{s,k} p_{j,k} : s \in [m] \setminus \{j\}$, and the Pigeon Axiom $\sum_{k \in [n]} p_{j,k} + 1$.

Let Π be a $\text{PCR}_{\mathbb{F}_2}$ refutation of $\text{Rank}_n^m(A)$ of column degree d and let $\Pi' = \Pi \upharpoonright \tau$.

Under the substitution

$$\tau(x_{i,k}) = \sum_{s \in [m]} A_{i,s} p_{s,k} \quad \text{and} \quad \tau(y_{k,j}) = p_{j,k},$$

the image of any variable that either X - or Y^T -column mentions some column $k \in [n]$ depends only on the pigeonhole variables $p_{s,k} : s \in [m]$ associated with hole k . Thus each column k in the rank principle corresponds exactly to hole k in the pigeonhole principle. Hence, if a monomial mentions a set $K \subseteq [n]$ of columns, its image under τ involves only variables corresponding to the holes in K . Consequently, every term in Π' mentions at most d holes.

The refutation Π' can be extended to a refutation Π'' of FPHP_n^m such that each term in Π'' still mentions at most d holes, in a straightforward manner (we omit the details). Combined with the bound from Theorem 4.10, we conclude that the column degree of Π must be at least $n/2 + 1$. $\square$

We continue by demonstrating a trade-off between the size and the *column degree* of refutations of $\text{Rank}_n^m(A)$. We construct a random substitution σ that w.h.p. turns any small refutation Π of $\text{Rank}_n^m(A)$ into a refutation $\text{Rank}_{\tilde{n}}^m(A)$ with $\tilde{n} = \Omega(n)$ while simultaneously reducing the column degree of Π . The proof of the resulting size lower bound is relatively straightforward, since there are no extension variables to handle. This substitution will later be used to prove lower bounds for CNF encodings.

Definition 5.4 (Substitution σ). Assume $m > n \geq 2$ and $A \in \{0, 1\}^{m \times m}$ is an $m \times m$ Boolean matrix. A random substitution σ is chosen as follows.

1. For each $k \in [n - 1]$, uniformly and independently sample $c_k \in \{*, 1, 0\}$. Then set $c_n = |\{k \in [n - 1] : c_k = 1\}| \bmod 2$. This guarantees that $|\{k \in [n] : c_k = 1\}|$ is even. We will call c_k the *column type* of k .
2. Let $\tilde{n} = |\{k \in [n] : c_k = *\}|$ and let $\tilde{X} = (\tilde{x}_{i,k})_{i \in [m], k \in [\tilde{n}]}$ and $\tilde{Y} = (\tilde{y}_{k,j})_{k \in [\tilde{n}], j \in [m]}$ be fresh variables.
3. Let ι be a bijection between $\{k \in [n] : c_k = *\}$ and $[\tilde{n}]$.
4. For each $i \in [m]$ and $k \in [n]$ set

$$\sigma(x_{i,k}) = \begin{cases} \tilde{x}_{i,\iota(k)} & \text{if } c_k = *, \\ c_k & \text{if } c_k \in \{0, 1\}. \end{cases}$$

5. For each $j \in [m]$ and $k \in [n]$ set

$$\sigma(y_{k,j}) = \begin{cases} \tilde{y}_{\iota(k),j} & \text{if } c_k = *, \\ c_k & \text{if } c_k \in \{0, 1\}. \end{cases}$$

6. For any variable θ mapped to 0 or 1 by σ , set $\sigma(\bar{\theta}) = 1 - \sigma(\theta)$.
7. For any variable θ not mapped to a Boolean constant by σ , set $\sigma(\bar{\theta}) = \overline{\sigma(\theta)}$ (namely, the negated variable of $\sigma(\theta)$).

Note that, independently for each $k \in [n - 1]$, the substitution σ either keeps column k as a fresh variable column, sets it to the all-one column, or sets it to the all-zero column, each with probability $1/3$. The final column is set to a constant in order to preserve the parity of the number of all-one columns. The last column is unconditionally set to a constant, and thus we do not need to consider it in the analysis of the column degree. The resulting matrices X and Y^T are thus shrunk under the restriction to fewer columns, where the original variables are replaced with the new $\tilde{x}$ - and $\tilde{y}$ -variables.

Lemma 5.5. Assume $m > n \geq 2$ and $A \in \{0, 1\}^{m \times m}$ is a Boolean matrix. Let t be a term that does not mention the last n th column and mentions columns $K \subseteq [n - 1]$. Then $\Pr[t \upharpoonright \sigma \neq 0] \leq (2/3)^{|K|}$. Furthermore, $\Pr[\tilde{n} < (n - 1)/6] \leq e^{-\frac{n-1}{24}}$.

Proof. For the first statement, choose for each column k from the set K a single variable γ_k in t that mentions column k . Let Γ_k be the event that $(\gamma_k \upharpoonright \sigma \neq 0)$. Then $\{\Gamma_k : k \in K\}$ is a set of $|K|$ mutually independent events and each of them happens with probability at most $2/3$. The term t is non-zero under σ only if all events Γ_k happen, and thus

$$\Pr[t \upharpoonright \sigma \neq 0] \leq (2/3)^{|K|}.$$

For the second statement in the lemma, define for each $k \in [n - 1]$ the random variable

$$z_k = \begin{cases} 1 & \text{if } c_k = *, \\ 0 & \text{otherwise.} \end{cases}$$

Then $\tilde{n} = \sum_{k \in [n-1]} z_k$. We have $z_k = 1$ with probability $1/3$, and thus $\mu = \mathbb{E} \tilde{n} = (n-1)/3$. Therefore, by [Theorem 4.1](#), applied with $\delta = 1/2$,

$$\Pr \left[\tilde{n} < \frac{n-1}{6} \right] \leq e^{-\frac{n-1}{24}}.$$

□

The substitution σ transforms any refutation of $\text{Rank}_n^m(A)$ into a refutation of $\text{Rank}_{\tilde{n}}^m(A)$.

Lemma 5.6. *Assume $m > n \geq 2$ and $A \in \{0, 1\}^{m \times m}$ is a Boolean matrix. Then for every $\text{PCR}_{\mathbb{F}_2}$ refutation Π of $\text{Rank}_n^m(A)$, $\Pi \upharpoonright \sigma$ is a refutation of $\text{Rank}_{\tilde{n}}^m(A)$.*

Proof. For every $i, j \in [m]$, the (i, j) th axiom of $\text{Rank}_n^m(A)$ transforms under σ as follows:

$$\left(A_{i,j} + \sum_{k=1}^n x_{i,k} y_{k,j} \right) \upharpoonright \sigma = A_{i,j} + \sum_{k \in [n], c_k = *} \tilde{x}_{i, \iota(k)} \tilde{y}_{\iota(k), j} + |\{k \in [n] : c_k = 1\}| = A_{i,j} + \sum_{\ell=1}^{\tilde{n}} \tilde{x}_{i, \ell} \tilde{y}_{\ell, j},$$

since $|\{k \in [n] : c_k = 1\}|$ is even and ι is a bijection between $\{k \in [n] : c_k = *\}$ and $[\tilde{n}]$. □

Theorem 5.7. *Suppose that $m > n \geq 2$ and $A \in \{0, 1\}^{m \times m}$ is an $m \times m$ Boolean matrix. Then any $\text{PCR}_{\mathbb{F}_2}$ refutation of $\text{Rank}_n^m(A)$ requires size $2^{\Omega(n)}$.*

Proof. Assume that $\text{Rank}_n^m(A)$ has a refutation Π of size less than $(3/2)^{\frac{n-1}{12}} (1 - e^{-\frac{n-1}{24}}) = 2^{\Theta(n)}$. Then, by the union bound and [Lemmas 5.5](#) and [5.6](#) there exists a restriction σ such that $\Pi \upharpoonright \sigma$ is a refutation of $\text{Rank}_{\tilde{n}}^m(A)$ with $\tilde{n} \geq \frac{n-1}{6}$ and column degree less than $\frac{n-1}{12} \leq \frac{\tilde{n}}{2}$. This would contradict [Lemma 5.3](#). □

5.2 Perfect Matching Encoding of the Rank Principle

To encode the rank formula $\text{Rank}_n^m(A)$ as a CNF, we consider an encoding that uses the *perfect matching principle* to encode the parities appearing in the formula. Similar encodings were used in [\[42, 47\]](#). For every $i, j \in [m]$, we encode the parity $\sum_{k \in [n]} x_{i,k} y_{k,j} = A_{i,j}$ by exhibiting a perfect matching on the terms $x_{i,k} y_{k,j}$ that evaluate to one (together with an extra point when $A_{i,j} = 1$, added so that the set being matched has even cardinality).

Definition 5.8 (CNF encoding $\text{PMRank}_n^m(A)$). Assume $m > n > 0$ and $A \in \{0, 1\}^{m \times m}$.

Variables: $X = (x_{i,k})_{i \in [m], k \in [n]}$, $Y = (y_{k,j})_{k \in [n], j \in [m]}$, for each $i, j \in [m]$, $\{z_{i,j,\{k,\ell\}} : \{k, \ell\} \in \binom{[n]}{2}\}$ if $A_{i,j} = 0$ and $\{w_{i,j,\{k,\ell\}} : \{k, \ell\} \in \binom{[n+1]}{2}\}$ if $A_{i,j} = 1$.

Even Axioms: For every $i, j \in [m]$ with $A_{i,j} = 0$, we use the following clauses:

$$\neg x_{i,k} \vee \neg y_{k,j} \vee \bigvee_{\ell \in [n] \setminus \{k\}} z_{i,j,\{k,\ell\}}, \quad k \in [n], \quad (10)$$

$$\neg x_{i,k} \vee \neg y_{k,j} \vee \neg z_{i,j,\{k,\ell\}} \vee x_{i,\ell}, \quad k \in [n], \ell \in [n] \setminus \{k\}, \quad (11)$$

$$\neg x_{i,k} \vee \neg y_{k,j} \vee \neg z_{i,j,\{k,\ell\}} \vee y_{\ell,j}, \quad k \in [n], \ell \in [n] \setminus \{k\}, \quad (12)$$

$$\neg x_{i,k} \vee \neg y_{k,j} \vee \neg z_{i,j,\{k,\ell\}} \vee \neg z_{i,j,\{k,\ell'\}}, \quad k \in [n], \ell, \ell' \in [n] \setminus \{k\}, \ell \neq \ell'. \quad (13)$$

Odd Axioms: For every $i, j \in [m]$ with $A_{i,j} = 1$, we use the following clauses:

$$\neg x_{i,k} \vee \neg y_{k,j} \vee \bigvee_{\ell \in [n+1] \setminus \{k\}} w_{i,j,\{k,\ell\}}, \quad k \in [n], \quad (14)$$

$$\bigvee_{\ell \in [n]} w_{i,j,\{n+1,\ell\}}, \quad (15)$$

$$\neg x_{i,k} \vee \neg y_{k,j} \vee \neg w_{i,j,\{k,\ell\}} \vee x_{i,\ell}, \quad k \in [n], \ell \in [n] \setminus \{k\}, \quad (16)$$

$$\neg x_{i,k} \vee \neg y_{k,j} \vee \neg w_{i,j,\{k,\ell\}} \vee y_{\ell,j}, \quad k \in [n], \ell \in [n] \setminus \{k\}, \quad (17)$$

$$\neg w_{i,j,\{n+1,\ell\}} \vee x_{i,\ell}, \quad \ell \in [n], \quad (18)$$

$$\neg w_{i,j,\{n+1,\ell\}} \vee y_{\ell,j}, \quad \ell \in [n], \quad (19)$$

$$\neg x_{i,k} \vee \neg y_{k,j} \vee \neg w_{i,j,\{k,\ell\}} \vee \neg w_{i,j,\{k,\ell'\}}, \quad k \in [n], \ell, \ell' \in [n+1] \setminus \{k\}, \ell \neq \ell', \quad (20)$$

$$\neg w_{i,j,\{n+1,\ell\}} \vee \neg w_{i,j,\{n+1,\ell'\}}, \quad \ell \in [n], \ell' \in [n] \setminus \{\ell\}. \quad (21)$$

In words, for each polynomial equation $\sum_{k \in [n]} x_{i,k} y_{k,j} = A_{i,j}$ in $\text{Rank}_n^m(A)$, the statement that the equation is satisfied is encoded by exhibiting a perfect matching on the satisfied monomials appearing in the equation (namely, those monomials that evaluate to 1 under a given assignment). In particular, clauses (10)–(13) express that each polynomial equation $\sum_{k \in [n]} x_{i,k} y_{k,j} = 0$ is satisfied by exhibiting a perfect matching on the set consisting of the satisfied monomials among $x_{i,k} y_{k,j}$, with $k \in [n]$: each monomial $x_{i,k} y_{k,j}$ with $x_{i,k} = y_{k,j} = 1$ is matched with some other monomial $x_{i,\ell} y_{\ell,j}$ (that also satisfies $x_{i,\ell} = y_{\ell,j} = 1$) by the variable $z_{i,j,\{k,\ell\}}$. Similarly, clauses (14)–(21) express $\sum_{k \in [n]} x_{i,k} y_{k,j} = 1$ using an extra point: each monomial $x_{i,k} y_{k,j}$ with $x_{i,k} = y_{k,j} = 1$ can be either matched to another such monomial or to the extra point $n+1$ (by setting $w_{i,j,\{k,n+1\}} = 1$). The axioms about the extra point ((15), (18), (19) and (21)) are stated separately since they have a simpler formulation than their counterparts (14), (16), (17) and (20).

To simplify the reasoning and highlight the main proof ideas, we focus on the case where A is the identity matrix I_m . The general case requires additional arguments, which we present in the next section (Section 5.3) for the *bamboo-tree* encoding, as it is more useful in the context of this paper. A similar approach to that in Section 5.3 can be applied to $\text{PMRank}_n^m(A)$, but this would require repeating many parts of the proofs from Section 5.3.

Our goal is to prove a lower bound for $\text{PMRank}_n^m(I_m)$. It is enough to prove such a lower bound after the substitution $Y = X^T$, that is, after replacing each variable $y_{k,j}$ by $x_{j,k}$. Indeed, applying this substitution to any PCR refutation of $\text{PMRank}_n^m(I_m)$ gives a PCR refutation of the restricted formula, with no increase in size. Thus, a size lower bound for the restricted formula implies the same lower bound for $\text{PMRank}_n^m(I_m)$.

After the substitution $Y = X^T$, the encoded equations become $\sum_{k \in [n]} x_{i,k} x_{j,k} = \delta_{i,j}$. Since the ones lie only on the diagonal we can replace each w -variable $w_{i,i,\{k,\ell\}}$ with $w_{i,\{k,\ell\}}$, where $i \in [m]$. For $i, j \in [m]$ with $i \neq j$, the axioms for $A_{i,j}$ and $A_{j,i}$ become identical, thus we unify the variables $z_{i,j,\{k,\ell\}}$ and $z_{j,i,\{k,\ell\}}$ into a single variable $z_{\{i,j\},\{k,\ell\}}$. This leads to the following simplified set of axioms:

Definition 5.9 (CNF for perfect matching encoding of Odd Town principle PMOT_n^m).

Even Axioms: For every $i, j \in [m]$ with $i \neq j$:

$$\neg x_{i,k} \vee \neg x_{j,k} \vee \bigvee_{\ell \in [n] \setminus \{k\}} z_{\{i,j\},\{k,\ell\}}, \quad k \in [n], \quad (22)$$

$$\neg x_{i,k} \vee \neg x_{j,k} \vee \neg z_{\{i,j\},\{k,\ell\}} \vee x_{i,\ell}, \quad k \in [n], \ell \in [n] \setminus \{k\}, \quad (23)$$

$$\neg x_{i,k} \vee \neg x_{j,k} \vee \neg z_{\{i,j\},\{k,\ell\}} \vee \neg z_{\{i,j\},\{k,\ell'\}}, \quad k \in [n], \ell, \ell' \in [n] \setminus \{k\}, \ell \neq \ell'. \quad (24)$$

Odd Axioms: For every $i \in [m]$:

$$\neg x_{i,k} \vee \bigvee_{\ell \in [n+1] \setminus \{k\}} w_{i,\{k,\ell\}}, \quad k \in [n], \quad (25)$$

$$\bigvee_{\ell \in [n]} w_{i, \{n+1, \ell\}}, \quad (26)$$

$$\neg x_{i,k} \vee \neg w_{i, \{k, \ell\}} \vee x_{i, \ell}, \quad k \in [n], \ell \in [n] \setminus \{k\}, \quad (27)$$

$$\neg w_{i, \{n+1, \ell\}} \vee x_{i, \ell}, \quad \ell \in [n], \quad (28)$$

$$\neg x_{i,k} \vee \neg w_{i, \{k, \ell\}} \vee \neg w_{i, \{k, \ell'\}}, \quad k \in [n], \ell, \ell' \in [n+1] \setminus \{k\}, \ell \neq \ell', \quad (29)$$

$$\neg w_{i, \{n+1, \ell\}} \vee \neg w_{i, \{n+1, \ell'\}}, \quad \ell \in [n], \ell' \in [n] \setminus \{\ell\}. \quad (30)$$

We call the formula consisting of the above axioms PMOT_n^m (since it encodes the negation of the *oddtown principle*, namely $XX^T = I$).

Theorem 5.10. *Let $m > n \geq 15$ with n divisible by 5 and $n/5$ odd. Then every $\text{PCR}_{\mathbb{F}_2}$ refutation of PMOT_n^m (and thus $\text{PMRank}_n^m(I_m)$) requires size at least $(10/9)^{n/20}$.*

The proof of this theorem follows a similar approach to the lower-bound proof for the algebraic encoding in [Section 5.1](#), but uses row degree instead of column degree. Furthermore, due to the presence of the extension variables (the z - and w -variables), our substitution needs to be more structured: a simple restriction like the one used in [Section 5.1](#) would not suffice.

Row degree measures the number of different rows that a term mentions. Every x - and w -variable mentions exactly one row, while every z -variable mentions two distinct rows. The latter complicates the analysis: previously we relied on the fact that every variable mentions exactly one column, which allowed us to find many independent events that a term is eliminated by the random restriction. In contrast, some terms in the z -variables may not yield independent events. In particular, star-shaped terms like $\prod_{j \in J} \bar{z}_{\{i, j\}, \{k_j, \ell_j\}}$ for some fixed i may have high row degree, but all variables in this term simultaneously depend on the same row i .

The proof of [Theorem 5.10](#) consists of three main components.

- We define a random restriction ρ that independently assigns each row one of five *types*, and then substitutes the x -variables according to the chosen type. Depending on the x -variables, the z - and w -variables are substituted in a way that is consistent with the perfect matching axioms according to fixed matchings, and set arbitrarily otherwise. The structure of the row types guarantees that the latter case happens often enough to eliminate terms of high row degree with high probability.
- We show that under ρ , any restricted formula $\text{PMOT}_n^m \upharpoonright \rho$ can be further transformed into a smaller formula $\text{PMOT}_{n/5-1}^{\tilde{m}}$ for some $m/5 \leq \tilde{m} \leq m$. This is achieved by a reduction σ that preserves the row degree.
- We show the hardness of $\text{PMOT}_{n/5-1}^{\tilde{m}}$ based on the reduction from the pigeonhole principle, similarly to [Lemma 5.3](#).

If there is a purported refutation of PMOT_n^m of small size, the union bound implies the existence of a restriction ρ that eliminates all terms of high row degree, and thus gives a refutation of $\text{PMOT}_{n/5-1}^{\tilde{m}}$ of small row degree, which is impossible.

We start with the definition of row degree.

Definition 5.11 (Row-mentioned variables for the perfect matching encoding). Let t be a term in the variables of PMOT_n^m . We say that $i \in [m]$ is *X-row-mentioned* in t if there is $k \in [n]$ such that $x_{i,k}$ or $\bar{x}_{i,k}$ appears in t . We say that $i \in [m]$ is *W-row-mentioned* in t if there is $\{k, \ell\} \in \binom{[n]}{2}$ such that $w_{i, \{k, \ell\}}$ or $\bar{w}_{i, \{k, \ell\}}$ appears in t . We say that $i \in [m]$ is *Z-row-mentioned* in t if there are $j \in [m]$ and $\{k, \ell\} \in \binom{[n]}{2}$ such that $z_{\{i, j\}, \{k, \ell\}}$ or $\bar{z}_{\{i, j\}, \{k, \ell\}}$ appears in t .

Definition 5.12 (Row degree). Let t be a term in the variables of PMOT_n^m . Let I be the set of all $i \in [m]$ that are either X -row-mentioned, W -row-mentioned or Z -row-mentioned in t . The *row degree* of t is defined as $|I|$. We also say that t *mentions* each $i \in I$.

As before, we construct a random substitution ρ such that, for every fixed term of high row degree, the probability that the term is not eliminated by ρ is exponentially small. The construction is based on the matrix $\mathcal{B}_n$, which we define below.

Recall that $J_{\ell, \ell'}$ denotes the $\ell \times \ell'$ all-ones matrix. We use $\otimes$ for the Kronecker (tensor) product. For two matrices P and Q , the matrix $P \otimes Q$ is obtained by replacing each entry $P_{i,j}$ by the block $P_{i,j}Q$. In particular, if $Q = J_{1,t}$ is the all-ones row vector of size t , then $P \otimes J_{1,t}$ is obtained by repeating each column of P into a block of t identical columns.

Let n be a positive integer divisible by 5 such that $\tilde{n} = n/5$ is odd. Partition the set of columns $[n]$ into five consecutive blocks of size $\tilde{n}$ each. Specifically, we have five column blocks A_1, A_2, A_3, A_4, A_5 , where $A_r = \{(r-1)\tilde{n} + 1, (r-1)\tilde{n} + 2, \dots, r\tilde{n}\}$ for $r \in [5]$. For each row $i \in [m]$, we independently choose $r_i \in [5]$, which is called its *row type* and determines the block configuration of the i th row of X .

Let $\mathcal{B}$ be the 5×5 matrix

$$\mathcal{B} = \begin{pmatrix} 1 & * & 0 & 0 & 0 \\ 0 & 1 & * & 0 & 0 \\ 0 & 0 & 1 & * & 0 \\ 0 & 0 & 0 & 1 & * \\ * & 0 & 0 & 0 & 1 \end{pmatrix}. \quad (31)$$

We require the following properties from $\mathcal{B}$:

- every row $\mathcal{B}_i$ of $\mathcal{B}$ has exactly one $*$ -entry;
- for every two distinct rows $\mathcal{B}_i$ and $\mathcal{B}_j$ there is at most one index $k \in [5]$ satisfying $(\mathcal{B}_{i,k}, \mathcal{B}_{j,k}) \in \{(*, 1), (1, *)\}$, and the remaining indices ℓ have either $\mathcal{B}_{i,\ell} = 0$ or $\mathcal{B}_{j,\ell} = 0$;
- every column of $\mathcal{B}$ contains at least one 0 and at least one 1;
- for every pair $\{k, \ell\} \in \binom{[5]}{2}$, there is at least one row $\mathcal{B}_i$ of $\mathcal{B}$ such that $\mathcal{B}_{i,k} = \mathcal{B}_{i,\ell} = 0$.

We will use the $5 \times n$ matrix $\mathcal{B}_n = \mathcal{B} \otimes J_{1, n/5}$. For example,

$$\mathcal{B}_{15} = \mathcal{B} \otimes J_{1,3} = \begin{pmatrix} 1 & 1 & 1 & * & * & * & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & * & * & * & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & * & * & * & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & * & * & * \\ * & * & * & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \end{pmatrix}.$$

The rows of $\mathcal{B}_n$ define five different row types. This matrix serves as a template for the matrix X in PMOT_n^m : the i th row of X is assigned according to the r_i th row of $\mathcal{B}_n$: the constants are assigned as is, while the $*$ entries remain unassigned. Given a row type $r \in [5]$ and $s \in \{*, 1, 0\}$, let

$$C_n^s(r) = \{k \in [n] : (\mathcal{B}_n)_{r,k} = s\} \quad (32)$$

be the set of all the column indices of the occurrences of s in the r th row of $\mathcal{B}_n$. For example, $C_n^1(r) = \{(r-1)n/5 + 1, (r-1)n/5 + 2, \dots, rn/5\}$ for all $r \in [5]$. The sets $C_n^s(r)$ determine the column configuration of the x -variables in the i th row of X according to the row type r_i . In particular, the variables from $C_n^*(r_i)$ remain unassigned, while the remaining x -variables are substituted with constants.

The first property of $\mathcal{B}$ guarantees that there is exactly one block that is left unassigned for each row type $r \in [5]$: for a block A we either have $C_n^*(r) = A$ or $A \cap C_n^*(r) = \emptyset$. Hence the column sets $C_n^*(r)$ are disjoint for different $r \in [5]$. The second property of $\mathcal{B}$ implies that for every two distinct row types r_i and r_j , there is at most one block A satisfying $C_n^*(r_i) \cap C_n^1(r_j) = A$ or $C_n^1(r_i) \cap C_n^*(r_j) = A$, and for all other blocks A' we have $A' \subseteq C_n^0(r_i)$ or $A' \subseteq C_n^0(r_j)$. The third property of $\mathcal{B}$ ensures that for every $k \in [n]$, each variable $x_{i,k}$ has a constant probability of being assigned to 0 or 1. Finally, the third and fourth properties of $\mathcal{B}$ guarantee that for every two distinct columns k and ℓ , there is at least one row type r_i such that both $x_{i,k}$ and $x_{i,\ell}$ are assigned to 0. This allows us to set the corresponding w - and z -variables arbitrarily, which is crucial for eliminating terms of high row degree.

To process the additional perfect matching variables, we fix a perfect matching $M(r)$ on the set $C_n^1(r) \cup \{n+1\}$ for each $r \in [5]$. This determines how the corresponding w -variables are substituted when the x -variables are assigned to 1 according to the row type r .

The interaction between two rows i and j depends on the relation between their row types r_i and r_j . In particular, if $r_i = r_j$, then the set of axioms for the inner product of the i th and j th rows is transformed into a perfect matching encoding on the surviving block with $\tilde{n}$ columns. The substitution σ flips these parities, yielding the usual oddtown constraints on one fewer column. If the two row types r_i and r_j are adjacent, then their axiom set is transformed into a perfect matching encoding on the inner product of the i th row with itself or the inner product of the j th row with itself, depending on the direction of adjacency. This does not add any new constraints. To restore the structure of this new axiom set, we fix an arbitrary index $k(r)$ in $C_n^1(r)$ for each $r \in [5]$ that plays the role of the new *extra point* and use another perfect matching $M'(r)$ on the set $C_n^1(r) \setminus \{k(r)\}$ that pairs all the remaining elements of $C_n^1(r)$ except for $k(r)$. Finally, in all other cases the corresponding axioms are satisfied.

Observe that the axioms of PMOT_n^m allow variables $w_{i,\{k,\ell\}}$ to be assigned arbitrarily when $\{k,\ell\} \subseteq C_n^0(r_i)$. These axioms are satisfied because we assign $x_{i,k} = x_{i,\ell} = 0$ for such k and ℓ . Similarly, the value of $z_{\{i,j\},\{k,\ell\}}$ can be chosen arbitrarily when $\{k,\ell\} \subseteq C_n^0(r_i) \cup C_n^0(r_j)$: the axioms it appears in are satisfied due to $x_{i,k}x_{j,k} = x_{i,\ell}x_{j,\ell} = 0$.

Definition 5.13 (Substitution ρ for PMOT_n^m). Assume that n is divisible by 5 and $\tilde{n} = n/5$ is odd. For each $r \in [5]$ fix an arbitrary perfect matching $M(r)$ on the set $C_n^1(r) \cup \{n+1\}$. For each $r \in [5]$ fix $k(r)$ to be the first (for definiteness) integer in $C_n^1(r)$ and fix an arbitrary perfect matching $M'(r)$ on the set $C_n^1(r) \setminus \{k(r)\}$. $k(r)$ will play the role of the new *extra point*. Define a random substitution ρ as a random partial map from the variables of PMOT_n^m and their twin variables to $\{0,1\}$ by the following random experiment:

1. Independently for each $i \in [m]$, choose uniformly at random an integer $r_i \in [5]$. Then, for each $i \in [m]$ and $k \in [n] \setminus C_n^*(r_i)$, set $\rho(x_{i,k}) = (\mathcal{B}_n)_{r_i,k}$. This assigns all the constant (0 or 1) entries of $(\mathcal{B}_n)_{r_i,k}$ in the i th row of X . The remaining variables $x_{i,k}$, where $i \in [m]$ and $k \in C_n^*(r_i)$, are left unassigned.
2. Independently for each $i \in [m]$ and each $\{k,\ell\} \subseteq C_n^0(r_i)$, ρ maps $w_{i,\{k,\ell\}}$ to 0 or 1 each with probability 1/2. These variables can be assigned to arbitrary values since all the axioms they appear in ((25), (27) and (29)) also contain $\neg x_{i,k}$ or $\neg x_{i,\ell}$, which we set to 1 in [Item 1](#), satisfying these axioms.

Next, set $\rho(w_{i,\{k,\ell\}}) = 1$ for every $\{k,\ell\} \in M(r_i)$, and $\rho(w_{i,\{k,\ell\}}) = 0$ for all variables $w_{i,\{k,\ell\}}$ that are not yet mapped, except for those with $\{k,\ell\} \subseteq C_n^*(r_i)$. This assignment sets all the remaining variables $\rho(w_{i,\{k,\ell\}})$ according to $M(r_i)$, leaving those with $\{k,\ell\} \subseteq C_n^*(r_i)$ unassigned.

Observe that, since $(n+1)$ is matched in $M(r_i)$, we always set all $w_{i,\{n+1,k\}}$, $k \in [n]$, to constants (exactly one of these variables to 1 and the rest to 0).

3. Independently for each $\{i, j\} \in \binom{[m]}{2}$ and each $\{k, \ell\} \subseteq C_n^0(r_i) \cup C_n^0(r_j)$, ρ maps $z_{\{i,j\},\{k,\ell\}}$ to 0 or 1 each with probability $1/2$. This is done similarly to [Item 2](#) since every axiom containing $z_{\{i,j\},\{k,\ell\}}$ ((22), (23) and (24)) also contains either $\neg x_{i,k} \vee \neg x_{j,k}$ or $\neg x_{i,\ell} \vee \neg x_{j,\ell}$ and both are satisfied under ρ .
4. For each $\{i, j\} \in \binom{[m]}{2}$ with $r_i = r_j$ and each $\{k, \ell\} \in M'(r_i)$, set $\rho(z_{\{i,j\},\{k,\ell\}}) = 1$. Set $\rho(z_{\{i,j\},\{k,\ell\}}) = 0$ for all variables with $i \neq j$ and $r_i = r_j$ that are not yet mapped, except for those with $\{k, \ell\} \subseteq C_n^*(r_i) \cup \{k(r_i)\}$.
5. For each $\{i, j\} \in \binom{[m]}{2}$ with $r_i \neq r_j$ and each $\{k, \ell\} \in \binom{[n]}{2}$ with $|\{k, \ell\} \cap (C_n^0(r_i) \cup C_n^0(r_j))| = 1$, set $\rho(z_{\{i,j\},\{k,\ell\}}) = 0$.
6. For any variable θ mapped to 0 or 1 by ρ in [Items 1 to 5](#), define $\rho(\bar{\theta}) = 1 - \rho(\theta)$.

Recall that PMOT_n^m semantically encodes the following statements:

$$\begin{aligned} \sum_{k=1}^n x_{i,k} &= 1, & i \in [m], \\ \sum_{k=1}^n x_{i,k} x_{j,k} &= 0, & \{i, j\} \in \binom{[m]}{2}. \end{aligned}$$

Under ρ , the former is transformed into

$$\sum_{k \in C_n^*(r_i)} x_{i,k} = 0, \quad i \in [m].$$

The latter depends on r_i and r_j . Let $\{i, j\} \in \binom{[m]}{2}$. Then it transforms into

$$\begin{aligned} \sum_{k \in C_n^*(r_i)} x_{i,k} x_{j,k} &= 1 & \text{if } r_i = r_j, \\ \sum_{k \in C_n^*(r_i)} x_{i,k} &= 0 & \text{if } r_j - r_i \equiv 1 \pmod{5}, \\ \sum_{k \in C_n^*(r_j)} x_{j,k} &= 0 & \text{if } r_j - r_i \equiv -1 \pmod{5}, \\ 0 &= 0 & \text{otherwise.} \end{aligned}$$

It follows from the above definition that for every variable θ of PMOT_n^m except the w -variables involving the extra point $n+1$, we have $\Pr[\theta \upharpoonright \rho = 0] \geq 1/10$ and $\Pr[\bar{\theta} \upharpoonright \rho = 0] \geq 1/10$. Moreover, these events depend only on the rows that θ mentions.

The following lemma shows that a term that mentions many rows is eliminated with high probability by ρ . We exclude the w -variables involving the extra point $n+1$ in the lemma because their assignments are correlated within each row; all of them are nevertheless assigned constants by ρ .

Lemma 5.14. *Suppose that n is divisible by 5 and $\tilde{n} = n/5$ is a positive odd integer. Let t be a term in the variables of PMOT_n^m and let t' be a subterm of t obtained by removing all variables $w_{i,\{k,n+1\}}$, $i \in [m]$, $k \in [n]$, and their twins. Let I be the set of all $i \in [m]$ that are X -, W - or Z -row mentioned by t' . Then $\Pr[t \upharpoonright \rho \neq 0] \leq (9/10)^{|I|/2}$.*

Proof. Given $q \in \{X, W, Z\}$, let $I(q)$ be the set of rows $i \in [m]$ that are q -row-mentioned in t' . For example, $I(X) = \{i \in [m] : \exists k \in [n], x_{i,k} \text{ or } \bar{x}_{i,k} \text{ occurs in } t'\}$.

There is a set $A \subseteq I(Z)$ of at least half of the elements $I(Z)$ and a function f from A to

$$\bigcup_{\substack{\{i,j\} \in \binom{[m]}{2}, \\ \{k,\ell\} \in \binom{[n]}{2}}} \{z_{\{i,j\},\{k,\ell\}}, \bar{z}_{\{i,j\},\{k,\ell\}}\}$$

such that f is injective⁶ and for each $i \in A$, $f(i)$ mentions i and $f(i)$ occurs in t' . To see this, consider a graph G with the vertex set $I(Z)$ and the edge set $\{\{i,j\} \in \binom{I(Z)}{2} : \text{for some } \{k,\ell\} \in \binom{[n]}{2}, z_{\{i,j\},\{k,\ell\}} \text{ or } \bar{z}_{\{i,j\},\{k,\ell\}} \text{ occurs in } t'\}$. For each connected component of G , fix a spanning tree and a root vertex. For each vertex i in the component other than the root, let j be the neighbour of i in the tree that is closer to the root (i.e., j is the parent of i in the rooted tree). We define $f(i)$ to be a z -variable (or its twin) that mentions both i and j and occurs in t' . This finishes the definition of A and f ; they clearly have the required properties. The size of A is at least the size of $I(Z)$ minus the number of connected components of G , which is at least $|I(Z)|/2$ since each component has at least two vertices.

We can further extend f to an injection f' defined on $A \cup I(X) \cup I(W)$ by setting $f'(i)$ for $i \in I(X) \setminus A$ to some x -variable (or its twin) that occurs in t and mentions i and by defining $f'(i)$ for $i \in I(W) \setminus (A \cup I(X))$ to be a w -variable (or its twin) that occurs in t and mentions i .

For every $i \in I(X) \setminus A$, define the event $\Gamma_i = (f'(i) = 0)$. By [Item 1](#) in [Definition 5.13](#), we have $\Pr[\Gamma_i] \geq 1/5$.

For every $i \in I(W) \setminus (A \cup I(X))$, we have $f'(i) = w_{i,\{k,\ell\}}$ or $f'(i) = \bar{w}_{i,\{k,\ell\}}$ and define the event $\Delta_i = (\{k,\ell\} \subseteq C_n^0(r_i) \wedge f'(i) = 0)$. It follows from the estimate $\Pr[\{k,\ell\} \subseteq C_n^0(r_i)] \geq 1/5$ (by [Item 1](#)) that $\Pr[\Delta_i] \geq 1/10$ (by [Item 2](#)).

For every $i \in A$, we have either $f'(i) = z_{\{i,j\},\{k,\ell\}}$ or $f'(i) = \bar{z}_{\{i,j\},\{k,\ell\}}$. Define the event $\Xi_i = (\{k,\ell\} \subseteq C_n^0(r_i) \wedge f'(i) = 0)$. It follows from the estimate $\Pr[\{k,\ell\} \subseteq C_n^0(r_i)] \geq 1/5$ (by [Item 1](#)) that $\Pr[\Xi_i] \geq 1/10$ (by [Item 3](#)). Note that Ξ_i depends only on the choice of r_i and the value of $f'(i)$.

Since all the events $\{\Gamma_i : i \in I(X) \setminus A\}$, $\{\Delta_i : i \in I(W) \setminus (I(X) \cup A)\}$ and $\{\Xi_i : i \in A\}$ are independent, we have

$$\Pr[t \upharpoonright \rho \neq 0] \leq (4/5)^{|I(X) \setminus A|} \cdot (9/10)^{|I(W) \setminus (A \cup I(X))|} \cdot (9/10)^{|I(Z)|/2} \leq (9/10)^{|I|/2}. \quad \square$$

Since the sets $C_n^*(r)$ are disjoint for distinct $r \in [5]$, the above formulas can be interpreted as a union of five disjoint copies of the same formula, up to the renaming of the variables. Within each row set $I(r)$, they form a copy of the same principle on the set $C_n^*(r)$, while the constraints between adjacent classes duplicate its diagonal constraints. Notice that the resulting formula asserts that the inner product of the row with itself is 0 while the inner product of two distinct rows of the same type is 1. This is the opposite of the original formula. This is why we need to further restrict the formula to invert these parities. Luckily, this can be done by a simple substitution that replaces the last column of each variable block $C_n^*(r)$ with the all-ones column that flips the signs, but reduces the number of columns only by one. This column matches the extra point in the even case and serves as a new extra point in the odd case. The following lemma formalizes this argument.

Lemma 5.15. *Let $m > n \geq 15$ with n divisible by 5 and $\tilde{n} = n/5$ odd. Let ρ be an element of the support of the distribution from [Definition 5.13](#). Then, there is an integer $\tilde{m}$ with $m/5 \leq \tilde{m} \leq m$ and a map σ from the variables of $\text{PMOT}_n^m \upharpoonright \rho$ and their twins to the variables of $\text{PMOT}_{\tilde{n}-1}^{\tilde{m}}$, their twins, and the constants 0, 1 such that $(\text{PMOT}_n^m \upharpoonright \rho) \upharpoonright \sigma = \text{PMOT}_{\tilde{n}-1}^{\tilde{m}}$. Moreover, there is a function $f : [m] \rightarrow [\tilde{m}]$ such that every variable of $\text{PMOT}_n^m \upharpoonright \rho$ that mentions the rows $i, j \in [m]$ (i may be equal to j) is mapped by σ to a variable that can only mention rows from $\{f(i), f(j)\}$.*

⁶We require f to injectively map A to the pairs of row indices $\{i, j\} \in \binom{[m]}{2}$.

Proof. Given $r \in [5]$, define

$$I(r) = \{i : r_i = r\}.$$

$I(r)$ gives a partition of $[m]$ into five disjoint sets. Let $\tilde{m} = \max\{|I(r)| : r \in [5]\}$. Then $m/5 \leq \tilde{m} \leq m$.

The surviving variables of $\text{PMOT}_n^m \upharpoonright \rho$ are the following variables.

- For each $i \in [m]$, the variables $x_{i,k}$ with $k \in C_n^*(r_i)$.
- For each $i \in [m]$, the variables $w_{i,\{k,\ell\}}$ with $\{k,\ell\} \subseteq C_n^*(r_i)$.
- For each $\{i,j\} \in \binom{[m]}{2}$ with $r_i = r_j$, the variables $z_{\{i,j\},\{k,\ell\}}$ with $\{k,\ell\} \subseteq C_n^*(r_i) \cup \{k(r_i)\}$.
- For each $\{i,j\} \in \binom{[m]}{2}$ with $r_j - r_i \equiv 1 \pmod{5}$, the variables $z_{\{i,j\},\{k,\ell\}}$ with $\{k,\ell\} \subseteq C_n^*(r_i)$.
- For each $\{i,j\} \in \binom{[m]}{2}$ with $r_i - r_j \equiv 1 \pmod{5}$, the variables $z_{\{i,j\},\{k,\ell\}}$ with $\{k,\ell\} \subseteq C_n^*(r_j)$.
- The twin variables of all the above variables.

Recall that for every $r \in [5]$, we have $r\tilde{n} \in C_n^*(r)$. Then for each $r \in [5]$, let $f_r : I(r) \rightarrow [\tilde{m}]$ be an arbitrary injection and $g_r : C_n^*(r) \rightarrow [\tilde{n}]$ be any bijection with $g_r(r\tilde{n}) = \tilde{n}$. The function f from the statement of the lemma is defined as $f(i) = f_r(i)$ for $i \in I(r)$.

For each $r \in [5]$, σ replaces the column $r\tilde{n}$ with the all-ones column, which flips the parities of the inner products. For the diagonal inner products, this column plays the role of the extra point, while for the off-diagonal inner products, it is matched with the existing extra point $k(r)$ to satisfy the corresponding axioms. The functions f_r and g_r are used to rename the surviving variables of $\text{PMOT}_n^m \upharpoonright \rho$ to the variables of $\text{PMOT}_{\tilde{n}-1}^{\tilde{m}}$. Formally, it is defined as follows.

1. For every $i \in [m]$ and $k \in C_n^*(r_i)$, set

$$\sigma(x_{i,k}) = \begin{cases} x_{f_{r_i}(i), g_{r_i}(k)}, & k \neq r_i\tilde{n}, \\ 1, & k = r_i\tilde{n}. \end{cases}$$

2. For every $i \in [m]$ and $\{k,\ell\} \subseteq C_n^*(r_i)$, set

$$\sigma(w_{i,\{k,\ell\}}) = w_{f_{r_i}(i), \{g_{r_i}(k), g_{r_i}(\ell)\}}.$$

3. For every $\{i,j\} \in \binom{[m]}{2}$ with $r_i = r_j$ and $\{k,\ell\} \subseteq C_n^*(r_i) \setminus \{r_i\tilde{n}\}$, set

$$\sigma(z_{\{i,j\},\{k,\ell\}}) = z_{\{f_{r_i}(i), f_{r_i}(j)\}, \{g_{r_i}(k), g_{r_i}(\ell)\}}.$$

4. For every $\{i,j\} \in \binom{[m]}{2}$ with $r_i = r_j$, set

$$\sigma(z_{\{i,j\},\{k(r_i), r_i\tilde{n}\}}) = 1.$$

5. For every $\{i,j\} \in \binom{[m]}{2}$ with $r_i = r_j$, $k \in C_n^*(r_i) \setminus \{r_i\tilde{n}\}$, and $\ell \in \{k(r_i), r_i\tilde{n}\}$, set

$$\sigma(z_{\{i,j\},\{k,\ell\}}) = 0.$$

6. For every $\{i,j\} \in \binom{[m]}{2}$ with $r_j - r_i \equiv 1 \pmod{5}$ and $\{k,\ell\} \subseteq C_n^*(r_i)$, set

$$\sigma(z_{\{i,j\},\{k,\ell\}}) = w_{f_{r_i}(i), \{g_{r_i}(k), g_{r_i}(\ell)\}}.$$

7. For every $\{i, j\} \in \binom{[m]}{2}$ with $r_i - r_j \equiv 1 \pmod{5}$ and $\{k, \ell\} \subseteq C_n^*(r_j)$, set

$$\sigma(z_{\{i,j\},\{k,\ell\}}) = w_{f_{r_j}(j),\{g_{r_j}(k),g_{r_j}(\ell)\}}.$$

8. For any variable θ mapped to 0 or 1 by σ , set $\sigma(\bar{\theta}) = 1 - \sigma(\theta)$.
9. For any variable θ not mapped to a Boolean constant by σ , set $\sigma(\bar{\theta}) = \overline{\sigma(\theta)}$.

By the description of the restricted axiom sets above, σ maps the diagonal, same-type off-diagonal, and adjacent-type axiom sets to the corresponding axiom sets of $\text{PMOT}_{\tilde{n}-1}^{\tilde{m}}$, while all remaining axiom sets are already satisfied. Moreover, for any $r \in [5]$ with $|I(r)| = \tilde{m}$, the map f_r is a bijection onto $[\tilde{m}]$, and hence all axiom sets of $\text{PMOT}_{\tilde{n}-1}^{\tilde{m}}$ are obtained.

Observe that our choice of f_r ensures that the row degree does not increase under σ . $\square$

Finally, we show a row degree lower bound on PMOT_n^m by constructing a reduction from FPHP_n^m to it.

Lemma 5.16. *Every $\text{PCR}_{\mathbb{F}_2}$ refutation of PMOT_n^m has a term t that mentions at least $n/2 + 1$ rows.*

Proof. We will reduce FPHP_n^m to PMOT_n^m . Let $p_{i,k} : i \in [m], k \in [n]$ be fresh variables.

We construct a reduction τ that maps the x -variables of PMOT_n^m to the p -variables of FPHP_n^m naturally: τ replaces $x_{i,k}$ with $p_{i,k}$. The hole axioms of FPHP ensure that all terms appearing in the inner product of two distinct rows are mapped to 0, so we can set all the z -variables to 0. The extra point is used to ensure that at least one variable $x_{i,k}$ is set to 1. Functional axioms ensure that at most one such variable is set to 1, so we can set all the w -variables involving the extra point to the corresponding p -variable – τ replaces $w_{i,\{n+1,k\}}$ with $p_{i,k}$ – and the remaining w -variables to 0.

Formally, we define τ as follows:

1. For every $i \in [m]$ and $k \in [n]$, set $\tau(x_{i,k}) = p_{i,k}$ and $\tau(\bar{x}_{i,k}) = \bar{p}_{i,k}$.
2. For every $i \neq j \in [m]$ and $\{k, \ell\} \in \binom{[n]}{2}$, set $\tau(z_{\{i,j\},\{k,\ell\}}) = 0$ and $\tau(\bar{z}_{\{i,j\},\{k,\ell\}}) = 1$.
3. For every $i \in [m]$ and $\{k, \ell\} \in \binom{[n]}{2}$, set $\tau(w_{i,\{k,\ell\}}) = 0$ and $\tau(\bar{w}_{i,\{k,\ell\}}) = 1$.
4. For every $i \in [m]$ and $k \in [n]$, set $\tau(w_{i,\{n+1,k\}}) = p_{i,k}$ and $\tau(\bar{w}_{i,\{n+1,k\}}) = \bar{p}_{i,k}$.

Let Π be a $\text{PCR}_{\mathbb{F}_2}$ refutation of PMOT_n^m such that every term in Π mentions at most d rows. Then every term in $\Pi' = \Pi \upharpoonright \tau$ mentions at most d pigeons. Now we show that each polynomial in $\text{PMOT}_n^m \upharpoonright \tau$ has a simple derivation from FPHP_n^m .

- The axioms (23) and (24) are satisfied since we set all z -variables to zero.
- The axiom (22) becomes $p_{i,k}p_{j,k}$, which is a Hole Axiom of FPHP_n^m (4).
- The axioms (25) and (28) become $p_{i,k}\bar{p}_{i,k}$ and thus trivially follow from the properties of twin variables.
- The axiom (26) becomes $\prod_{\ell \in [n]} \bar{p}_{i,\ell}$ that semantically follows from the Pigeon Axioms (3) and has a simple derivation from it, which mentions only pigeon i : we can multiply the Pigeon Axiom $\sum_{k \in [n]} p_{i,k} = 1$ by $\prod_{\ell \in [n]} \bar{p}_{i,\ell}$ and use the fact that $p_{i,k}\bar{p}_{i,k} = 0$ to obtain $\prod_{\ell \in [n]} \bar{p}_{i,\ell} = 0$.
- The axioms (27) and (29) are satisfied.
- The axiom (30) becomes $p_{i,\ell}p_{i,\ell'}$, which is a Functionality Axiom of FPHP_n^m (5).

Thus, Π' can be further extended to a refutation Π'' of FPHP_n^m such that each term in Π'' mentions at most d pigeons.

The lemma follows from the lower bound on the number of pigeons from [Theorem 4.10](#). $\square$

Combining [Lemmas 5.14](#) to [5.16](#), we proceed by proving [Theorem 5.10](#).

Proof of Theorem 5.10. Assume for the sake of contradiction that PMOT_n^m has a refutation Π of size less than $(10/9)^{n/20} = (10/9)^{\tilde{n}/4}$. We start by sampling a random substitution ρ from [Definition 5.13](#). It follows from [Lemma 5.14](#) that any term that mentions more than $\tilde{n}/2$ rows (excluding the w -variables involving the extra point $n+1$) is not set to 0 by ρ with probability at most $(9/10)^{\tilde{n}/4}$. By the union bound, there exists a substitution ρ such that every term in $\Pi \upharpoonright \rho$ mentions at most $\tilde{n}/2$ different rows.

By [Lemma 5.15](#) we can further restrict $\Pi \upharpoonright \rho$ and obtain a refutation $\Pi' = \Pi \upharpoonright \rho \upharpoonright \sigma$ of $\text{PMOT}_{\tilde{n}-1}^{\tilde{m}}$ with $m/5 \leq \tilde{m} \leq m$. Moreover, every term in Π' still mentions at most $\tilde{n}/2$ rows.

Applying [Lemma 5.16](#) to Π' , we obtain a refutation of $\text{FPHP}_{\tilde{n}-1}^{\tilde{m}}$ in which every term mentions at most $\tilde{n}/2$ pigeons. This contradicts the row degree bound from [Lemma 5.16](#). $\square$

The fact that we focused on the identity matrix I_m and not the general case of an arbitrary matrix A simplifies the argument since we only need one substitution ρ to reduce the degree of the terms. It is also helpful that we can use the natural notion of row degree in this case. We consider the general case for a different encoding – which encodes the parities using a balanced-tree structure – in [Section 5.3](#). There, we use a series of substitutions to reduce the appropriate notion of degree. Nevertheless, a similar approach can also be used to show the hardness of $\text{PMRank}_n^m(A)$, but it would require redoing many intermediate steps, which is why we do not present it here.

As we mentioned earlier, for proof systems in which the perfect matching principle is already hard, the hardness of $\text{PMRank}_n^m(A)$ follows from the hardness of Count_n^2 . We begin by describing this reduction. We describe it for the general case $\text{PMRank}_n^m(A)$, where the matrix A is arbitrary.

Lemma 5.17. *Let A be an arbitrary $m \times m$ Boolean matrix that is distinct from matrices of all ones or all zeros. Then there exists a reduction ρ such that $\text{PMRank}_n^m(A) \upharpoonright \rho$ coincides with $\text{Count}_n^2 \wedge \text{Count}_{n+1}^2$.*

Proof. Let $u_{\{k,\ell\}} : \{k,\ell\} \in \binom{[n]}{2}$ and $v_{\{k,\ell\}} : \{k,\ell\} \in \binom{[n+1]}{2}$ be fresh variables and consider the reduction ρ defined as follows:

1. For every $i \in [m]$ and $k \in [n]$ set $\rho(x_{i,k}) = 1$ and $\rho(\bar{x}_{i,k}) = 0$.
2. For every $j \in [m]$ and $k \in [n]$ set $\rho(y_{k,j}) = 1$ and $\rho(\bar{y}_{k,j}) = 0$.
3. For every $i, j \in [m]$ with $A_{i,j} = 0$ and every $\{k,\ell\} \in \binom{[n]}{2}$, set $\rho(z_{i,j,\{k,\ell\}}) = u_{\{k,\ell\}}$.
4. For every $i, j \in [m]$ with $A_{i,j} = 1$ and every $\{k,\ell\} \in \binom{[n+1]}{2}$, set $\rho(w_{i,j,\{k,\ell\}}) = v_{\{k,\ell\}}$.
5. For every variable θ , if $\rho(\theta)$ is not a Boolean constant, set $\rho(\bar{\theta}) = \overline{\rho(\theta)}$.

This restriction substitutes X with $J_{m,n}$, Y with $J_{n,m}$, and unifies all isomorphic instances of Count_n^2 and Count_{n+1}^2 .

We show that $\text{PMRank}_n^m(A) \upharpoonright \rho$ becomes a union of disjoint instances of Count_n^2 and Count_{n+1}^2 . Let $i, j \in [m]$. We consider the following cases.

- Suppose $A_{i,j} = 0$. In this case, we use the axioms (10)–(13), which under ρ become an instance of Count_n^2 in the variables $u_{\{k,\ell\}} : \{k,\ell\} \in \binom{[n]}{2}$.
- Suppose $A_{i,j} = 1$. In this case, we use the axioms (14)–(21), which under ρ become an instance of Count_{n+1}^2 in the variables $v_{\{k,\ell\}} : \{k,\ell\} \in \binom{[n+1]}{2}$.

Since A contains both ones and zeros, both formulas Count_n^2 and Count_{n+1}^2 appear in $\text{PMRank}_n^m(A) \upharpoonright \rho$. Clearly, exactly one of them is unsatisfiable. $\square$

This lemma immediately implies the hardness of $\text{PMRank}_n^m(A)$ for proof systems, in which Count_n^2 is hard.

Theorem 5.18. *Let $m > n > 0$ be positive integers. Let A be an arbitrary $m \times m$ Boolean matrix. Then any $\text{PCR}_{\mathbb{F}}$ refutation (for a field $\mathbb{F}$ with $\text{char}(\mathbb{F}) \neq 2$) or SoS refutation of the perfect matching CNF encoding $\text{PMRank}_n^m(A)$ requires size $2^{\Omega(n)}$. Similarly, any AC^0 -Frege refutation requires size $2^{n^{\Omega(1)}}$, where the constant in the exponent depends on the circuit depth.*

Proof. If all entries of A are all ones or all zeros, the formula is satisfiable and the theorem holds trivially. Otherwise, by applying Lemma 5.17, we conclude that there is a reduction from $\text{Count}_n^2 \wedge \text{Count}_{n+1}^2$ to $\text{PMRank}_n^m(A)$. Depending on the parity of n exactly one of these instances is unsatisfiable. They are defined on disjoint sets of variables, hence we can satisfy the other.

Count_k^2 requires size $2^{\Omega(k)}$ in $\text{PCR}_{\mathbb{F}}$ when $\text{char}(\mathbb{F}) \neq 2$, by combining the degree lower bounds from [17] with the size-degree trade-offs from [19, 41]. The analogous lower bound for SoS follows from [9, 33, 77]. Similarly, the lower bound for AC^0 -Frege follows from [50]. $\square$

5.3 Bamboo-Tree Encoding of the Rank Principle

To encode the rank formula $XY = A$ as a CNF, extension variables are introduced to sequentially compute all inner products involved. The extension variables for each inner product are arranged in a totally unbalanced binary tree known as a “bamboo”. This natural encoding corresponds to a circuit in the sense of [4, 76], and similar encodings have been used repeatedly for parity computations [4, 76, 81].

Definition 5.19 (The CNF encoding $\text{BTRank}_n^m(A)$). Let $m > n$ be positive integers and $A \in \{0, 1\}^{m \times m}$ an m by m Boolean matrix. The bamboo-tree CNF encoding, denoted $\text{BTRank}_n^m(A)$, uses input variables $x_{i,k}, y_{k,i} : i \in [m], k \in [n]$ together with extension variables $z_{i,j,k}, u_{i,j,k} : i, j \in [m], k \in [n]$, and consists of the following axioms.

Output Axioms: For every $i, j \in [m]$, the clause $u_{i,j,n}$ if $A_{i,j} = 1$ and the clause $\neg u_{i,j,n}$ if $A_{i,j} = 0$.

Binary AND Axioms: For every $i, j \in [m]$ and $k \in [n]$, $z_{i,j,k} = x_{i,k}y_{k,j}$, encoded by the clauses $x_{i,k} \vee \neg z_{i,j,k}, y_{k,j} \vee \neg z_{i,j,k}, z_{i,j,k} \vee \neg x_{i,k} \vee \neg y_{k,j}$.

Summation Base Axioms: For every $i, j \in [m]$, $u_{i,j,1} = z_{i,j,1}$, encoded by the clauses $u_{i,j,1} \vee \neg z_{i,j,1}, z_{i,j,1} \vee \neg u_{i,j,1}$.

Summation Axioms: For every $i, j \in [m]$, and $k \in \{2, \dots, n\}$, $u_{i,j,k} = u_{i,j,k-1} + z_{i,j,k}$, encoded by the following four clauses:

$$\begin{aligned} z_{i,j,k} \vee \neg u_{i,j,k-1} \vee u_{i,j,k}, \\ z_{i,j,k} \vee u_{i,j,k-1} \vee \neg u_{i,j,k}, \\ \neg z_{i,j,k} \vee \neg u_{i,j,k-1} \vee \neg u_{i,j,k}, \\ \neg z_{i,j,k} \vee u_{i,j,k-1} \vee u_{i,j,k}. \end{aligned}$$

We focus on the regime where m is arbitrarily larger than n . As with related formulas like PHP_n^m , obtaining size lower bounds in this regime is challenging across various proof systems, since standard measures such as resolution width or PC degree do not appear to be effective, at least when applied directly. Our technique applies to any matrix A and provides a $\text{PCR}_{\mathbb{F}_2}$ size lower bound of $2^{\Omega(n)}$. This makes $\text{BTRank}_n^m(A)$ a proof complexity generator, stretching $2mn$ input bits into m^2 output bits.

Theorem 5.20. *Suppose that $m > n \geq 16$ are integers, 8 divides n , and $A \in \{0, 1\}^{m \times m}$ is an m by m matrix over $\mathbb{F}_2$. Then any $\text{PCR}_{\mathbb{F}_2}$ refutation of $\text{BTRank}_n^m(A)$ has size at least $2^{\frac{\log e}{512} n - 2}$.*

Recall that the basic strategy for the simpler algebraic encoding of the rank principle is to use a random restriction that, with high probability, transforms a small-size refutation into a small-degree refutation. This then yields a contradiction with the PCR degree lower bounds for the pigeonhole principle. Accordingly, our goal here is to show that if $\text{BTRank}_n^m(A)$ admits a small $\text{PCR}_{\mathbb{F}_2}$ refutation, then it can be converted into one in which a certain *relaxed* form of degree is small. This relaxed notion of degree is carefully tailored to the structure of our formula, with distinct definitions for each variable sort. Moreover, it appears that if we require a random substitution to serve as a meaningful self-reduction (i.e., converting $\text{BTRank}_n^m(A)$ to a smaller instance of the same principle), then this relaxed degree measure is close to the limit of what can be effectively decreased under such a random transformation. While the definition can be made slightly stricter, any substantial move toward the standard notion of degree (i.e., the total degree of a term) does not seem to yield a measure that decreases adequately under random self-reductions. This limitation seems to stem from the distinct nature of the (“extension”) u -variables, whose behaviour differs from that of the other variables in the formula. Fortunately, our relaxed degree is still strong enough to yield, after an additional substitution and some proof manipulations, a $\text{PCR}_{\mathbb{F}_2}$ refutation of $\text{FPHP}_{n'}^m$ of degree less than $n'/2$, which is ruled out by a theorem of Razborov [72].

The initial conversion of a purported small refutation of $\text{BTRank}_n^m(A)$ begins by applying two distinct random substitutions, ρ and later σ . The random substitution ρ converts (with probability 1) a refutation of $\text{BTRank}_n^m(A)$ into a refutation of an instance with fewer columns $\text{BTRank}_{\tilde{n}}^m(A)$, for some $\tilde{n} < n$. Additionally, with high probability, this substitution simplifies the occurrences of u -variables in the terms of the refutation.

We now define the random substitution ρ (not to be confused with the substitution with the same notation in Section 5.2). This substitution randomly compresses the matrices X and Y from n columns to about $n/4$ columns, row by row, in a way that preserves the bamboo structure and makes different rows behave independently.

Definition 5.21 (Substitution ρ for $\text{BTRank}_n^m(A)$). Assume that $m > n \geq 8$ and n is divisible by 4. Let $\tilde{n} = (n - 4)/4$ and $\{\tilde{x}_{i,\ell} : i \in [m], \ell \in [\tilde{n}]\} \cup \{\tilde{y}_{\ell,i} : i \in [m], \ell \in [\tilde{n}]\} \cup \{\tilde{z}_{i,j,\ell} : i, j \in [m], \ell \in [\tilde{n}]\} \cup \{\tilde{u}_{i,j,\ell} : i, j \in [m], \ell \in [\tilde{n}]\}$ be a set of fresh variables. The random substitution ρ is defined as follows. Independently and uniformly at random, choose $4m$ bits $\{p_{i,0}, p_{i,1}, q_{i,0}, q_{i,1} : i \in [m]\} \in \{0, 1\}^{4m}$. These bits are used as follows: $p_{i,0}$ chooses how the boundary columns of row i in X are hard-fixed, $p_{i,1}$ chooses which half of the interior columns of row i carry $\tilde{x}$ -variables (the other half become 0), $q_{j,0}$ similarly fixes boundary entries for column j in Y , $q_{j,1}$ chooses which interior blocks of column j carry $\tilde{y}$ -variables (the rest become 0). See Figure 3 and Figure 4. We extend ρ to the u - and z -variables so that it agrees with what ρ already does to X and Y , and so that $\text{BTRank}_n^m(A)$ restricted by ρ becomes an instance of $\text{BTRank}_{\tilde{n}}^m(A)$ (with the same matrix A). Formally, this is done as follows.

1. For each $i \in [m]$ and $k \in [n]$, set

$$\rho(x_{i,k}) = \begin{cases} p_{i,0} & \text{if } k \in \{1, n-1\}, \\ 1 & \text{if } k \in \{2, n\}, \\ 0 & \text{if } k \in \{3, \dots, n-2\}, \lfloor \frac{k-3}{\tilde{n}} \rfloor \in \{0, 1\} \text{ and } p_{i,1} = 0, \\ \tilde{x}_{i,((k-3) \bmod \tilde{n})+1} & \text{if } k \in \{3, \dots, n-2\}, \lfloor \frac{k-3}{\tilde{n}} \rfloor \in \{2, 3\} \text{ and } p_{i,1} = 0, \\ \tilde{x}_{i,((k-3) \bmod \tilde{n})+1} & \text{if } k \in \{3, \dots, n-2\}, \lfloor \frac{k-3}{\tilde{n}} \rfloor \in \{0, 1\} \text{ and } p_{i,1} = 1, \\ 0 & \text{if } k \in \{3, \dots, n-2\}, \lfloor \frac{k-3}{\tilde{n}} \rfloor \in \{2, 3\} \text{ and } p_{i,1} = 1, \end{cases}$$

and similarly, for each $k \in [n]$ and $j \in [m]$, set

$$\rho(y_{k,j}) = \begin{cases} 1 & \text{if } k \in \{1, n-1\}, \\ q_{j,0} & \text{if } k \in \{2, n\}, \\ 0 & \text{if } k \in \{3, \dots, n-2\}, \lfloor \frac{k-3}{\tilde{n}} \rfloor \in \{0, 2\} \text{ and } q_{j,1} = 0, \\ \tilde{y}_{((k-3) \bmod \tilde{n})+1,j} & \text{if } k \in \{3, \dots, n-2\}, \lfloor \frac{k-3}{\tilde{n}} \rfloor \in \{1, 3\} \text{ and } q_{j,1} = 0, \\ \tilde{y}_{((k-3) \bmod \tilde{n})+1,j} & \text{if } k \in \{3, \dots, n-2\}, \lfloor \frac{k-3}{\tilde{n}} \rfloor \in \{0, 2\} \text{ and } q_{j,1} = 1, \\ 0 & \text{if } k \in \{3, \dots, n-2\}, \lfloor \frac{k-3}{\tilde{n}} \rfloor \in \{1, 3\} \text{ and } q_{j,1} = 1. \end{cases}$$

In other words, the first two variables $(x_{i,1}, x_{i,2})$ of the row $(x_{i,1}, \dots, x_{i,n})$ are set to $(0, 1)$ if $p_{i,0} = 0$, and to $(1, 1)$ if $p_{i,0} = 1$, and the same is true for the last two variables $(x_{i,n-1}, x_{i,n})$. The block of remaining variables $(x_{i,3}, \dots, x_{i,n-2})$ is set to $(0, 0, 1, 1) \otimes (\tilde{x}_{i,1}, \dots, \tilde{x}_{i,\tilde{n}})$ if $p_{i,1} = 0$, and to $(1, 1, 0, 0) \otimes (\tilde{x}_{i,1}, \dots, \tilde{x}_{i,\tilde{n}})$ if $p_{i,1} = 1$. Similarly, for $(y_{1,j}, \dots, y_{n,j})$, the first two (as well as the last two) variables are set to $(1, 0)$ or $(1, 1)$ depending on the value of $q_{j,0}$, and the remaining variables $(y_{3,j}, \dots, y_{n-2,j})$ are set to $(0, 1, 0, 1) \otimes (\tilde{y}_{1,j}, \dots, \tilde{y}_{\tilde{n},j})$ or $(1, 0, 1, 0) \otimes (\tilde{y}_{1,j}, \dots, \tilde{y}_{\tilde{n},j})$ depending on the value of $q_{j,1}$.

- For each $i, j \in [m]$ and $k \in [n]$, set $z_{i,j,k}$ to the correct value as determined by the values of $\rho(x_{i,k})$ and $\rho(y_{k,j})$ just defined and by Binary AND Axioms. That is, for $k \in \{1, n-1\}$ we define $\rho(z_{i,j,k}) = p_{i,0}$, for $k \in \{2, n\}$ we set $\rho(z_{i,j,k}) = q_{j,0}$, and for $k \in \{3, \dots, n-2\}$ we have four cases based on the values of $p_{i,1}, q_{j,1}$. First, let $r_{i,j} \in \{0, 1, 2, 3\}$ be defined by $r_{i,j} = 3 - 2p_{i,1} - q_{j,1}$. It follows from [Item 1](#) that the following definition is consistent with the Binary AND Axioms:

$$\rho(z_{i,j,k}) = \begin{cases} \tilde{z}_{i,j,((k-3) \bmod \tilde{n})+1} & \text{if } k \in \{r_{i,j}\tilde{n} + 3, \dots, (r_{i,j} + 1)\tilde{n} + 2\}, \\ 0 & \text{if } k \in \{3, \dots, n-2\} \setminus \{r_{i,j}\tilde{n} + 3, \dots, (r_{i,j} + 1)\tilde{n} + 2\}. \end{cases}$$

In other words, for $k \in \{3, \dots, n-2\}$ we define $\rho(z_{i,j,k}) = 0$ if either $\rho(x_{i,k}) = 0$ or $\rho(y_{k,j}) = 0$; otherwise, it follows from the previous item that $\rho(x_{i,k}) = \tilde{x}_{i,\ell}$ and $\rho(y_{k,j}) = \tilde{y}_{\ell,j}$, where $\ell = ((k-3) \bmod \tilde{n}) + 1$, and in this case we set $\rho(z_{i,j,k}) = \tilde{z}_{i,j,\ell}$.

- The u -variables are assigned by ρ consistently not only with Summation Base Axioms and Summation Axioms, but also with Output Axioms, taking into account that after applying ρ to $\text{BTRank}_n^m(A)$ we want Output Axioms to hold for the new variables $\tilde{u}_{i,j,\tilde{n}}$. Given $i \in [m]$ and $j \in [m]$, the definition of $\rho(u_{i,j,1}), \dots, \rho(u_{i,j,n})$ splits into cases based on the values $p_{i,0}, p_{i,1}, q_{j,0}, q_{j,1}$. It follows from [Item 1](#) that the values of the first two variables are constant, namely $\rho(u_{i,j,1}) = p_{i,0}$ and $\rho(u_{i,j,2}) = p_{i,0} \oplus q_{j,0}$. We now define the values $\rho(u_{i,j,3}), \dots, \rho(u_{i,j,n-2})$ with the help of the integer $r_{i,j}$ defined in [Item 2](#).

$$\rho(u_{i,j,k}) = \begin{cases} \tilde{u}_{i,j,((k-3) \bmod \tilde{n})+1} & \text{if } k \in \{r_{i,j}\tilde{n} + 3, \dots, (r_{i,j} + 1)\tilde{n} + 2\} \text{ and } p_{i,0} \oplus q_{j,0} = 0, \\ \bar{\tilde{u}}_{i,j,((k-3) \bmod \tilde{n})+1} & \text{if } k \in \{r_{i,j}\tilde{n} + 3, \dots, (r_{i,j} + 1)\tilde{n} + 2\} \text{ and } p_{i,0} \oplus q_{j,0} = 1, \\ p_{i,0} \oplus q_{j,0} & \text{if } k \in \{3, \dots, r_{i,j}\tilde{n} + 2\}, \\ p_{i,0} \oplus q_{j,0} \oplus A_{i,j} & \text{if } k \in \{(r_{i,j} + 1)\tilde{n} + 3, \dots, n-2\}. \end{cases}$$

To define ρ consistently at the last two variables, $u_{i,j,n-1}$ and $u_{i,j,n}$, we need to set $\rho(u_{i,j,n-1}) = (p_{i,0} \oplus q_{j,0} \oplus A_{i,j}) \oplus p_{i,0} = q_{j,0} \oplus A_{i,j}$ and $\rho(u_{i,j,n}) = (q_{j,0} \oplus A_{i,j}) \oplus q_{j,0} = A_{i,j}$.

- For any variable θ mapped to 0 or 1 by ρ , set $\rho(\bar{\theta}) = 1 - \rho(\theta)$.
- For any variable θ not mapped to a Boolean constant by ρ , set $\rho(\bar{\theta}) = \overline{\rho(\theta)}$.





Figure 3: Action of ρ on a fixed row of X (two cases: $p_{i,1} = 0$ and $p_{i,1} = 1$).




Figure 4: Action of ρ on a fixed column of Y (two cases: $q_{j,1} = 0$ and $q_{j,1} = 1$).

First, we demonstrate the desired effect of ρ on the u -variables occurring in terms.

Definition 5.22 (Row and column indices mentioned by u -variables inside a term). Let t be a term in the variables of $\text{BTRank}_n^m(A)$. We say that $i \in [m]$ is *U-left-row-mentioned* in t if there is $j \in [m], k \in [n]$, and $b \in \{0, 1\}$ such that $u_{i,j,k}^b$ appears in t . We say that $j \in [m]$ is *U-right-row-*

mentioned in t if there is $i \in [m], k \in [n]$, and $b \in \{0, 1\}$ such that $u_{i,j,k}^b$ appears in t . We say that $k \in [n]$ is *U-column-mentioned* in t if there are $i, j \in [m]$ and $b \in \{0, 1\}$ such that $u_{i,j,k}^b$ appears in t .

Note that row indices are those that originally range over $[m]$ while column indices originally range over $[n]$.

Lemma 5.23 (Random substitution ρ kills terms that mention many u -rows). *Suppose that $m > n \geq 8$ are integers, 4 divides n , and $A \in \{0, 1\}^{m \times m}$ is an m by m matrix over $\mathbb{F}_2$. Suppose that t is a term in the variables of $\text{BTRank}_n^m(A)$, $I \subseteq [m]$ is a set and t' is a subterm of t such that every $i \in I$ is *U-left-row-mentioned* in t' and no $k \in \{1, 2, n-1, n\}$ is *U-column-mentioned* in t' . Then $\Pr[t \upharpoonright \rho \neq 0] \leq (3/4)^{|I|}$.*

Proof. Let $r = |I|$. By the assumption of the lemma, t contains a subterm of the form $u_{i_1, j_1, k_1}^{b_1} u_{i_2, j_2, k_2}^{b_2} \cdots u_{i_r, j_r, k_r}^{b_r}$, where $1 \leq i_1 < i_2 < \cdots < i_r \leq m$ and for each $s \in [r]$ we have $b_s \in \{0, 1\}$, $j_s \in [m]$ and $k_s \in \{3, \dots, n-2\}$. It follows from the definition of ρ that for each row i that actually appears in the term (via some $u_{i,j,k}$ with $k \in \{3, \dots, n-2\}$), there is a $1/4$ chance that ρ sets that u -factor to the “opposite” constant (so the whole product becomes 0), and these events are independent across the r different rows; hence the term survives with probability at most $(3/4)^r$.

More formally, by the definition of ρ for each $(a_s)_{s \in [r]} \in \{0, 1\}^r$, for each $(a_{j,0}, a_{j,1})_{j \in \{j_1, \dots, j_r\}} \in \{0, 1\}^{2 \cdot |\{j_1, \dots, j_r\}|}$, and for each $h \in [r]$, we have

$$\Pr[\rho(u_{i_h, j_h, k_h}) = a_h \mid \text{for all } s \in [r], q_{j_s,0} = a_{j_s,0} \text{ and } q_{j_s,1} = a_{j_s,1}] \geq 1/4.$$

For each $h \in [r]$, at least one of the two possible values of a_h evaluates t to 0. Since all $i_1, \dots, i_r$ are distinct, these r conditional events are mutually independent, and hence

$$\Pr[t \upharpoonright \rho \neq 0 \mid \text{for all } s \in [r], q_{j_s,0} = a_{j_s,0} \text{ and } q_{j_s,1} = a_{j_s,1}] \leq (3/4)^r.$$

By averaging, the lemma follows. $\square$

Next, we verify that ρ transforms a refutation of $\text{BTRank}_n^m(A)$ into a refutation of $\text{BTRank}_{(n-4)/4}^m(A)$. Namely, we reduce the number of columns in X and, accordingly, the number of rows in Y . This is a structural consequence of the way the substitution is defined, and occurs with probability 1: the k -indices in the input variables $x_{i,k}, y_{k,i} : i \in [m], k \in [n]$ (excluding the four boundary ones) are partitioned into four consecutive blocks of equal size, and for each fixed row i the rule for ρ keeps exactly two of these blocks of $x_{i,k}$ as live variables while fixing the other two blocks to constants; similarly, for each fixed column j the rule for ρ keeps exactly two blocks of $y_{k,j}$ (see Figures 3 and 4). These two surviving blocks are deliberately misaligned, so that for every pair (i, j) their intersection is contained in a *single* block. Since $z_{i,j,k}$ can only remain potentially non-zero when both $x_{i,k}$ and $y_{k,j}$ survive (based on the definition of ρ , Item 2), it follows that after applying ρ the surviving z -variables (hence the relevant k -columns) for each i and j are confined to one block, i.e., to at most $(n-4)/4$ columns. Formally, we have the following.

Lemma 5.24. *Suppose that $m > n \geq 8$ are integers, 4 divides n , and $A \in \{0, 1\}^{m \times m}$ is an m by m matrix over $\mathbb{F}_2$. Let Π be a $\text{PCR}_{\mathbb{F}_2}$ refutation of $\text{BTRank}_n^m(A)$. Then $\Pi \upharpoonright \rho$ is a $\text{PCR}_{\mathbb{F}_2}$ refutation of $\text{BTRank}_{(n-4)/4}^m(A)$.*

Proof. We only need to check how ρ affects (the translations into polynomials of) the clauses of $\text{BTRank}_n^m(A)$. By Items 2 and 3, the clauses of each Output Axiom and each Summation Base Axiom are satisfied by ρ . Equivalently, ρ converts the translation of each such clause into the zero polynomial (which can be removed from $\Pi \upharpoonright \rho$). Further, by Item 2, each Binary AND Axiom is either satisfied

by ρ or turned into the corresponding axiom of $\text{BTRank}_n^m(A)$. Next, by [Items 2 and 3](#), Summation Axioms for $k \in \{2, n\}$ are satisfied by ρ .

Consider now a Summation Axiom for $k = n-1$, i.e., the CNF encoding of $u_{i,j,n-1} = u_{i,j,n-2} + z_{i,j,n-1}$. By the definition of ρ , we have $\rho(z_{i,j,n-1}) = p_{i,0}$, and $\rho(u_{i,j,n-1}) = q_{j,0} \oplus A_{i,j}$. If $r_{i,j} = 3$, we have $\rho(u_{i,j,n-2}) = \tilde{u}_{i,j,\tilde{n}}^{p_{i,0} \oplus q_{j,0} \oplus 1}$, which means that each of the four clauses encoding the axiom is either satisfied or is transformed into the corresponding Output Axiom clause $\tilde{u}_{i,j,\tilde{n}}^{A_{i,j}}$. If $r_{i,j} \neq 3$, then $\rho(u_{i,j,n-2}) = p_{i,0} \oplus q_{j,0} \oplus A_{i,j}$, and so all four clauses encoding the Summation Axiom are satisfied.

Finally, consider Summation Axioms for $k \in \{3, \dots, n-2\}$. For $k = r_{i,j}\tilde{n} + 3$, the axiom turns into the corresponding Summation Base Axiom (for either value of $p_{i,0} \oplus q_{j,0}$). For $k \in \{r_{i,j}\tilde{n} + 4, \dots, (r_{i,j} + 1)\tilde{n} + 2\}$, the axiom turns into a Summation Axiom (again, for either value of $p_{i,0} \oplus q_{j,0}$). For $k = (r_{i,j} + 1)\tilde{n} + 3$, when $r_{i,j} \neq 3$, we have $\rho(u_{i,j,(r_{i,j}+1)\tilde{n}+2}) = \tilde{u}_{i,j,\tilde{n}}^{p_{i,0} \oplus q_{j,0} \oplus A_{i,j}}$, $\rho(u_{i,j,(r_{i,j}+1)\tilde{n}+3}) = p_{i,0} \oplus q_{j,0} \oplus A_{i,j}$, and $\rho(z_{i,j,(r_{i,j}+1)\tilde{n}+3}) = 0$, so the axiom turns into the corresponding Output Axiom $\tilde{u}_{i,j,\tilde{n}}^{A_{i,j}}$. When $r_{i,j} = 3$, we have $k = n-1$, which is already covered above. For $k \in \{3, \dots, r_{i,j}\tilde{n} + 2\} \cup \{(r_{i,j} + 1)\tilde{n} + 4, \dots, n-2\}$, the Summation Axiom is satisfied by ρ . $\square$

The second random substitution, denoted by σ , is intended to simplify the x -, y - and z -variable occurrences in terms while, again, turning $\text{BTRank}_n^m(A)$ into a smaller instance. In contrast to the substitution ρ , which aims, with high probability, to reduce the number of *row* indices (of X) appearing through the u -variables in each term (as well as to reduce the instance from n to $\tilde{n}$ columns), the goal of σ is to reduce the number of *column* indices (of X and Y^T) involved via the x -, y - and z -variables.

Applying ρ reduces the number of columns of the matrices X and Y^T in the original instance from n to $(n-4)/4$ ([Lemma 5.24](#)). However, this does not substantially decrease the number of column indices mentioned by a given term: the term may still mention many columns within the surviving block via the x -, y -, and z -variables. The purpose of σ is to significantly reduce, for each term, the number of column indices involved through these variables.

Definition 5.25 (Substitution σ for $\text{BTRank}_n^m(A)$). Assume that $m > n \geq 3$ and $n-1$ is divisible by 2. A random substitution σ is chosen as follows.

1. Choose a subset $S \subseteq [n-1]$ uniformly at random from all $(n-1)/2$ -element subsets of $[n-1]$. For each $k \in [n-1] \setminus S$, independently choose $q_k \in \{0, 1\}$ at random with $\Pr[q_k = 0] = \Pr[q_k = 1] = 1/2$. For each $k \in S$, define $q_k = *$, and set $q_n = |\{\ell \in [n-1] : q_\ell = 1\}| \bmod 2$. Note that this choice of q_n makes the total number of ones even, i.e., if for $k \in [n]$ we let $o_k = |\{\ell \in [k] : q_\ell = 1\}| \bmod 2$, then we have $o_n = 0$. Denote by $s_k = |S \cap [k]|$ the number of stars among $q_1, \dots, q_k$.
2. For each $i, j \in [m]$ and $k \in [n]$, define

$$\sigma(x_{i,k}) = \begin{cases} q_k & \text{if } q_k \in \{0, 1\}, \\ \tilde{x}_{i,s_k} & \text{if } q_k = *, \end{cases} \quad \sigma(y_{k,j}) = \begin{cases} q_k & \text{if } q_k \in \{0, 1\}, \\ \tilde{y}_{s_k,j} & \text{if } q_k = *, \end{cases}$$

and extend σ consistently to the z - and u -variables by

$$\sigma(z_{i,j,k}) = \begin{cases} q_k & \text{if } q_k \in \{0, 1\}, \\ \tilde{z}_{i,j,s_k} & \text{if } q_k = *, \end{cases}$$

$$\sigma(u_{i,j,k}) = \begin{cases} o_k & \text{if } k \in \{1, \dots, \min(S) - 1\}, \\ \tilde{u}_{i,j,s_k}^{o_k \oplus 1} & \text{if } k \in \{\min(S), \dots, \max(S)\}, \\ A_{i,j} \oplus o_k & \text{if } k \in \{\max(S) + 1, \dots, n\}. \end{cases}$$

3. For any variable θ mapped to 0 or 1 by σ , set $\sigma(\bar{\theta}) = 1 - \sigma(\theta)$.

4. For any variable θ not mapped to a Boolean constant by σ , set $\sigma(\bar{\theta}) = \overline{\sigma(\theta)}$.

Example for Definition 5.25: choosing S and defining q_k, o_k, s_k (here $n = 9$)

k	1	2	3	4	5	6	7	8	9
q_k	1	*	0	*	1	*	0	*	0
o_k	1	1	1	1	0	0	0	0	0
s_k	0	1	1	2	2	3	3	4	4

Green = $k \in S$ (so $q_k = *$ survives), Gray = constant bits (fixed by σ).
Here $S = \{2, 4, 6, 8\}$ and q_9 is chosen so that $o_9 = 0$.

Figure 5: A concrete example of the random data in Definition 5.25: the set S , the values (q_k) , and the derived sequences (o_k) and $s_k = |S \cap [k]|$.

Example: action of σ on $(x_{i,1}, \dots, x_{i,9})$

k	1	2	3	4	5	6	7	8	9
$\sigma(x_{i,k})$	q_1	$\tilde{x}_{i,1}$	q_3	$\tilde{x}_{i,2}$	q_5	$\tilde{x}_{i,3}$	q_7	$\tilde{x}_{i,4}$	q_9

Rule: $\sigma(x_{i,k}) = q_k$ if $q_k \in \{0, 1\}$, and $\sigma(x_{i,k}) = \tilde{x}_{i,s_k}$ if $q_k = *$.

Figure 6: Example of σ on the x -variables.

Example: action of σ on $(y_{1,j}, \dots, y_{9,j})$

k	1	2	3	4	5	6	7	8	9
$\sigma(y_{k,j})$	q_1	$\tilde{y}_{1,j}$	q_3	$\tilde{y}_{2,j}$	q_5	$\tilde{y}_{3,j}$	q_7	$\tilde{y}_{4,j}$	q_9

Rule: $\sigma(y_{k,j}) = q_k$ if $q_k \in \{0, 1\}$, and $\sigma(y_{k,j}) = \tilde{y}_{s_k,j}$ if $q_k = *$.

Figure 7: Example of σ on the y -variables.

Example: action of σ on $(u_{i,j,1}, \dots, u_{i,j,9})$ (here $a = 2, b = 8$)

k	1	2	3	4	5	6	7	8	9
$\sigma(u_{i,j,k})$	o_1	$\tilde{u}_{i,j,s_2}^{o_2 \oplus 1}$	$\tilde{u}_{i,j,s_3}^{o_3 \oplus 1}$	$\tilde{u}_{i,j,s_4}^{o_4 \oplus 1}$	$\tilde{u}_{i,j,s_5}^{o_5 \oplus 1}$	$\tilde{u}_{i,j,s_6}^{o_6 \oplus 1}$	$\tilde{u}_{i,j,s_7}^{o_7 \oplus 1}$	$\tilde{u}_{i,j,s_8}^{o_8 \oplus 1}$	$A_{i,j} \oplus o_9$

Rule: $\sigma(u_{i,j,k}) = o_k$ for $k < a$; $\sigma(u_{i,j,k}) = \tilde{u}_{i,j,s_k}^{o_k \oplus 1}$ for $a \leq k \leq b$; $\sigma(u_{i,j,k}) = A_{i,j} \oplus o_k$ for $k > b$.

Figure 8: Example of σ on the u -chain (three regions determined by $a = \min(S)$ and $b = \max(S)$).

Definition 5.26 (Column indices mentioned in a term by x -, y - and z -variables). Let t be a term in the variables of $\text{BTRank}_n^m(A)$ and let $k \in [n]$. We say that k is X -column-mentioned in t if there are $i \in [m]$ and $b \in \{0, 1\}$ such that $x_{i,k}^b$ appears in t . We say that k is Y^T -column-mentioned in t if there are $j \in [m]$ and $b \in \{0, 1\}$ such that $y_{k,j}^b$ appears in t . We say that k is Z -column-mentioned in t if there are $i, j \in [m]$ and $b \in \{0, 1\}$ such that $z_{i,j,k}^b$ appears in t .

Lemma 5.27. Suppose that $m > n \geq 3$ are integers, 2 divides $n - 1$, and $A \in \{0, 1\}^{m \times m}$ is an m by m matrix over $\mathbb{F}_2$. Suppose that t is a term in the variables of $\text{BTRank}_n^m(A)$ and $K \subseteq [n - 1]$ is a set such that every $k \in K$ is either X -column-mentioned or Y^T -column-mentioned or Z -column-mentioned in t . Then $\Pr[t \upharpoonright \sigma \neq 0] \leq \exp(-|K|/16) + 2^{-|K|/4}$.

Proof. Let $\bar{S} = [n - 1] \setminus S$ be the complement of S in $[n - 1]$. Since S is chosen uniformly at random from all $(n - 1)/2$ -element subsets of $[n - 1]$, the set $\bar{S}$ is also chosen uniformly at random from all $(n - 1)/2$ -element subsets of $[n - 1]$. Define the random variable X by

$$X = |K \setminus S| = |K \cap \bar{S}|$$

and $\mu = \mathbb{E} X = |K|/2$. Then [Theorem 4.2](#) (with $\delta = 1/2$) implies that

$$\Pr[X \leq \mu/2] \leq \exp(-\mu/8) = \exp(-|K|/16).$$

Now fix any choice of S such that $|K \setminus S| = X > |K|/4$. We pick for each $k \in K \setminus S$ a single variable γ_k from t witnessing that k is X -column-mentioned or Y^T -column-mentioned or Z -column-mentioned. Notice that σ evaluates these variables independently to 0 or 1, each with probability $1/2$. At most one of the $2^X \geq 2^{|K|/4}$ assignments to these variables does not zero out t , so $\Pr[t \upharpoonright \sigma \neq 0 \mid S \text{ with } X > |K|/4] \leq 2^{-|K|/4}$.

The union bound then gives $\Pr[t \upharpoonright \sigma \neq 0] \leq \exp(-|K|/16) + 2^{-|K|/4}$. $\square$

Lemma 5.28. Suppose that $m > n \geq 3$ are integers, 2 divides $n - 1$, and $A \in \{0, 1\}^{m \times m}$ is an m by m matrix over $\mathbb{F}_2$. Let Π be a $\text{PCR}_{\mathbb{F}_2}$ refutation of $\text{BTRank}_n^m(A)$. Then $\Pi \upharpoonright \sigma$ is a $\text{PCR}_{\mathbb{F}_2}$ refutation of $\text{BTRank}_{(n-1)/2}^m(A)$.

Proof. We inspect how σ acts on (the translations into polynomials of) the clauses of $\text{BTRank}_n^m(A)$. First, consider the Output Axioms. We know by [Item 1](#) that $o_n = 0$, and that $n \in \{\max(S) + 1, \dots, n\}$ for any choice of S . Hence by [Item 2](#), $\sigma(u_{i,j,n}) = A_{i,j}$, which means that Output Axioms are satisfied. Next, it is immediate by [Item 2](#) that the Binary AND Axiom for $k \in [n]$ is either satisfied by σ (if $q_k \in \{0, 1\}$) or turned into the Binary AND Axiom $\tilde{z}_{i,j,s_k} = \tilde{x}_{i,s_k} \tilde{y}_{s_k,j}$ of $\text{BTRank}_{(n-1)/2}^m(A)$. Further, the Summation Base Axiom is satisfied if $1 \notin S$, since $\sigma(u_{i,j,1}) = o_1 = q_1 = \sigma(z_{i,j,1})$; otherwise $s_1 = 1$ and it turns into a Summation Base Axiom of $\text{BTRank}_{(n-1)/2}^m(A)$.

Next, consider Summation Axioms. For $k \in \{2, \dots, \min(S) - 1\}$ the axiom is satisfied by σ . If $2 \leq k = \min(S)$, then each clause of the axiom is either satisfied or turned into a clause of the

Summation Base Axiom of $\text{BTRank}_{(n-1)/2}^m(A)$ (for either value of o_k). For $k \in \{\min(S), \dots, \max(S)\} \setminus S$, each clause of the axiom is either satisfied or turned into $u_{i,j,s_k} \vee \neg u_{i,j,s_k}$. For $k \in S \setminus \{\min(S)\}$, the axiom turns into the Summation Axiom for s_k of $\text{BTRank}_{(n-1)/2}^m(A)$ (for either value of o_k). For $k = \max(S) + 1$, each clause of the axiom is either satisfied or turned into an Output Axiom of $\text{BTRank}_{(n-1)/2}^m(A)$. Finally, for $k \in \{\max(S) + 2, \dots, n\}$, the Summation Axiom is satisfied by σ . $\square$

The next lemma combines [Lemmas 5.23, 5.24, 5.27](#) and [5.28](#).

Lemma 5.29. *Suppose that $m > n \geq 16$ are integers, 8 divides n , and $A \in \{0,1\}^{m \times m}$ is an m by m matrix over $\mathbb{F}_2$. If Π is a $\text{PCR}_{\mathbb{F}_2}$ refutation of $\text{BTRank}_n^m(A)$ of size s , then there exists a $\text{PCR}_{\mathbb{F}_2}$ refutation Π' of $\text{BTRank}_{(n-8)/8}^m(A)$ of size at most s such that, letting $d = \frac{16}{\log e}(\log s + 1)$, every term t in Π' satisfies the following:*

1. *the number of indices $i \in [m]$ that are U -left-row-mentioned in t is less than d , and*
2. *the number of indices $k \in [(n-8)/8]$ that are X -column-mentioned or Y^T -column-mentioned or Z -column-mentioned in t is less than d .*

Proof. Applying [Lemmas 5.24](#) and [5.28](#), let $\Pi_1 = \Pi \upharpoonright \rho$ be a refutation of $\text{BTRank}_{(n-4)/4}^m(A)$ and $\Pi_2 = \Pi_1 \upharpoonright \sigma$ be a refutation of $\text{BTRank}_{(n-8)/8}^m(A)$. A term t violating [Item 1](#) can only appear in Π_2 if $t = (\hat{t} \upharpoonright \rho) \upharpoonright \sigma$ for some term $\hat{t}$ in Π such that there exist a set $I \subseteq [m]$ and a subterm t' of $\hat{t}$ with the following properties: $|I| \geq d$, every $i \in I$ is U -left-row-mentioned in t' , and no $k \in \{1, 2, n-1, n\}$ is U -column-mentioned in t' . (The u -variables whose column indices belong to $\{1, 2, n-1, n\}$ can be excluded from t' since they are set to a constant by ρ with probability 1 and hence cannot contribute to a violation of [Item 1](#) by t .) However, for such a term $\hat{t}$ we have $\Pr[\hat{t} \upharpoonright \rho \neq 0] \leq (3/4)^{|I|} \leq (3/4)^d$ by [Lemma 5.23](#).

Similarly, a term t violating [Item 2](#) can only appear in Π_2 if $t = \tilde{t} \upharpoonright \sigma$ for some term $\tilde{t}$ in Π_1 such that there is a set $K \subseteq [(n-4)/4 - 1]$ with $|K| \geq d$ and every $k \in K$ is either X -column-mentioned or Y^T -column-mentioned or Z -column-mentioned in $\tilde{t}$. (The x -, y -, and z -variables that mention the last column index $\frac{n-4}{4}$ cannot contribute to a violation of [Item 2](#) by t as they are always evaluated to a constant by σ .) For such a term $\tilde{t}$, [Lemma 5.27](#) gives $\Pr[\tilde{t} \upharpoonright \sigma \neq 0] \leq \exp(-|K|/16) + 2^{-|K|/4} \leq \exp(-d/16) + 2^{-d/4}$.

Denote by T_r the set of all terms $\hat{t}$ in Π defined as above. In the event that all of these terms are zeroed out by ρ , there are at most $s - |T_r|$ terms $\tilde{t}$ in Π_1 defined as above. This is because $\Pi_1 = \Pi \upharpoonright \rho$ has no more terms than Π . By the union bound, the probability that there is a term t in Π_2 that violates [Item 1](#) or [Item 2](#) is bounded above by $|T_r| \cdot (3/4)^d + (s - |T_r|) \cdot (\exp(-d/16) + 2^{-d/4}) \leq s \cdot (\exp(-d/16) + 2^{-d/4}) < s \cdot 2^{-\frac{\log e}{16}d+1} = 1$ for $d = \frac{16}{\log e}(\log s + 1)$. Thus, we can fix choices of ρ and σ such that $\Pi' = (\Pi \upharpoonright \rho) \upharpoonright \sigma$ has the required properties. $\square$

Proof of Theorem 5.20. Assume, for contradiction, that there is a $\text{PCR}_{\mathbb{F}_2}$ refutation Π of size $s < 2^{\frac{\log e}{512}n-2}$ witnesses that the theorem fails. We let $\tilde{n} = (n-8)/8$ and construct a $\text{PCR}_{\mathbb{F}_2}$ refutation of $\text{FPHP}_{\tilde{n}}^m$ of degree less than $\tilde{n}/2$, thereby contradicting a theorem of Razborov [72]. By [Lemma 5.29](#), there is a $\text{PCR}_{\mathbb{F}_2}$ refutation Π_1 of $\text{BTRank}_{\tilde{n}}^m(A)$ such that [Items 1](#) and [2](#) hold with $d = \frac{16}{\log e}(\log s + 1) < \frac{n}{32} - \frac{16}{\log e} = \frac{8\tilde{n}+8}{32} - \frac{16}{\log e} < \frac{\tilde{n}}{4} - 10$.

Let τ be the following substitution on the variables of $\text{BTRank}_{\tilde{n}}^m(A)$: $\tau(x_{i,k}) = x_{i,k}$, $\tau(y_{k,j}) = \sum_{i \in [m]} x_{i,k} A_{i,j}$ (that is, τ substitutes $X^T A$ for Y), $\tau(z_{i,j,k}) = \tau(x_{i,k})\tau(y_{k,j})$, and $\tau(u_{i,j,k}) = \sum_{\ell \in [k]} \tau(z_{i,j,\ell})$. Extend τ to negations by $\tau(\bar{\theta}) = 1 + \tau(\theta)$ for any variable θ of $\text{BTRank}_{\tilde{n}}^m(A)$. Let $\Pi_2 = \Pi_1 \upharpoonright \tau$.

Consider a term t occurring in a proof line of Π_1 . We can write $t = t_1 t_2$, where t_1 consists of all x -, y - and z -variables and their twins that appear in t , and t_2 consists of all u -variables and their twins appearing in t .

Claim 5.30. *Every term of the polynomial $t_1 \upharpoonright \tau$ either has a degree less than d , or is a multiple of a Hole Axiom of $\text{FPHP}_{\tilde{n}}^m$.*

Proof. Let C denote the set of all indices $k \in [\tilde{n}]$ that are X -column-mentioned or Y^T -column-mentioned or Z -column-mentioned in t_1 . By the choice of Π_1 , we have $|C| < d$. Further, it is immediate from the definition of τ that the set $\{k \in [\tilde{n}] : k \text{ is } X\text{-column-mentioned in some term of the polynomial } t_1 \upharpoonright \tau\}$ is a subset of C . Thus, if a term t' in $t_1 \upharpoonright \tau$ has degree at least d , then, by the pigeonhole principle, t' must be a multiple of a Hole Axiom of $\text{FPHP}_{\tilde{n}}^m$. $\diamond$

Claim 5.31. *Every term of the polynomial $t_2 \upharpoonright \tau$ either has a degree less than d , or is a multiple of either a Hole Axiom or a Functionality Axiom of $\text{FPHP}_{\tilde{n}}^m$.*

Proof. Denote by R the set of all indices $i \in [m]$ that are U -left-row-mentioned in t_2 . By the choice of Π_1 , $|R| < d$. Assume that $u_{i,j,k}^{1-b}$ appears in t_2 for some $i \in R, j \in [m], k \in [\tilde{n}]$ and $b \in \{0, 1\}$. Then

$$\begin{aligned} \tau(u_{i,j,k}^{1-b}) &= b + \sum_{\ell \in [k]} \tau(z_{i,j,\ell}) = b + \sum_{\ell \in [k]} \tau(x_{i,\ell})\tau(y_{\ell,j}) = b + \sum_{\ell \in [k]} x_{i,\ell} \sum_{i' \in [m]} x_{i',\ell} A_{i',j} \\ &= b + \sum_{\ell \in [k]} \sum_{i' \in [m]} x_{i,\ell} x_{i',\ell} A_{i',j}. \end{aligned}$$

For every $\ell \in [k]$ and every $i' \in [m]$ such that $i' \neq i$ and $A_{i',j} \neq 0$, the summand $x_{i,\ell} x_{i',\ell} A_{i',j}$ in the last expression is a Hole Axiom. Since such summands only give rise to terms in $t_2 \upharpoonright \tau$ that are multiples of a Hole Axiom, we can continue the above calculation modulo Hole Axioms, obtaining $b + \sum_{\ell \in [k]} x_{i,\ell} x_{i,\ell} A_{i,j} = b + A_{i,j} \sum_{\ell \in [k]} x_{i,\ell}$. Thus, modulo Hole Axioms, $t_2 \upharpoonright \tau$ is either 0 or can be written as a product of degree-1 polynomials of the form

$$b + \sum_{\ell \in [k]} x_{i,\ell},$$

where $b \in \{0, 1\}, k \in [\tilde{n}]$, and $i \in R$. Whenever this product gives rise to a term of degree at least d , it follows from $|R| < d$ by the pigeonhole principle that such a term is a multiple of a Functionality Axiom of $\text{FPHP}_{\tilde{n}}^m$. $\diamond$

Combining the two claims, we see that every term of the polynomial $t \upharpoonright \tau = (t_1 \upharpoonright \tau)(t_2 \upharpoonright \tau)$ either has a degree at most $2d$, or is a multiple of either a Hole Axiom or a Functionality Axiom of $\text{FPHP}_{\tilde{n}}^m$. The same conclusion holds for every term appearing in $\Pi_2 = \Pi_1 \upharpoonright \tau$, since t was an arbitrary term occurring in Π_1 .

The next step is to transform the sequence of polynomials Π_2 into a $\text{PC}_{\mathbb{F}_2}$ refutation of $\text{FPHP}_{\tilde{n}}^m$.

First, we inspect the effect of τ on the axioms of Π_1 . By definition, τ satisfies all Complementarity Axioms. Additionally, τ satisfies the polynomial translation of every clause in the CNF encoding of Binary AND Axioms, Summation Base Axioms, and Summation Axioms. This can be verified by noting that each such clause is a semantic consequence of the algebraic version of the corresponding Binary AND Axiom (i.e., $z_{i,j,k} = x_{i,k}y_{k,j}$ for some $i, j \in [m]$ and $k \in [\tilde{n}]$), Summation Base Axiom (i.e., $u_{i,j,1} = z_{i,j,1}$ for some $i, j \in [m]$), or Summation Axiom (i.e., $u_{i,j,k} = u_{i,j,k-1} + z_{i,j,k}$ for some $i, j \in [m]$ and $k \in \{2, \dots, \tilde{n}\}$), and all of these are clearly satisfied by τ .

Each Output Axiom in Π_1 is converted by τ into one of the axioms algebraically encoding $XX^T A = A$. Each of the latter axioms is a sum of some axioms of $XX^T = I_m$, which is the algebraic oddtown formula $\text{OT}_{\tilde{n}}^m$. Each of the axioms of $\text{OT}_{\tilde{n}}^m$ is a sum of some axioms of $\text{FPHP}_{\tilde{n}}^m$. We append the degree-2 derivation just described to each axiom in Π_2 that is the result of an application of τ to an Output Axiom in Π_1 .

Next, we examine the non-axiom proof lines of Π_1 . The addition rule is not affected by τ , since if a polynomial p was inferred in Π_1 as a sum of q and r , then $p \upharpoonright \tau$ is a sum of $q \upharpoonright \tau$ and $r \upharpoonright \tau$. Regarding the multiplication rule, suppose that p and q are proof lines in Π_1 , and let p be inferred from q by multiplying q by θ , where θ is a variable or the negation of a variable of $\text{BTRank}_{\tilde{n}}^m(A)$. Then $p \upharpoonright \tau = (q \upharpoonright \tau)(\theta \upharpoonright \tau)$, and $\theta \upharpoonright \tau$ is a polynomial of degree at most 2. We add to Π_2 a derivation of $p \upharpoonright \tau$ from $q \upharpoonright \tau$ in which each term is a term of $q \upharpoonright \tau$ multiplied by at most two variables.

Let Π_3 denote the sequence of polynomials resulting from Π_2 by the modifications introduced in the last two paragraphs. It follows that Π_3 is a $\text{PC}_{\mathbb{F}_2}$ refutation of $\text{FPHP}_{\tilde{n}}^m$. Moreover, every term appearing in Π_3 either has a degree at most $2d + 2$, or is a multiple of either a Hole Axiom or a Functionality Axiom of $\text{FPHP}_{\tilde{n}}^m$.

Finally, using Hole Axioms and Functionality Axioms of $\text{FPHP}_{\tilde{n}}^m$, we transform Π_3 into a $\text{PC}_{\mathbb{F}_2}$ refutation of $\text{FPHP}_{\tilde{n}}^m$ in which each term has a degree at most $2d + 3 \leq \frac{\tilde{n}}{2} - 17 < \frac{\tilde{n}}{2}$, which is the desired contradiction. This can be done as follows. Delete from Π_3 each term that is a multiple of either a Hole Axiom or a Functionality Axiom. Then, for each line p that is an incorrect conclusion of the multiplication rule—which can only happen if p is missing some terms each of degree at most $2d + 3$ and having either exactly one row mentioned twice or exactly one column mentioned twice (both of which may occur simultaneously)—insert the correct conclusion p' of the rule before p and use Functionality Axioms and Hole Axioms to derive p from p' . $\square$

5.4 Iterated Generators

In the previous sections, we proved that the rank generator $g: \mathbb{F}_2^{2mn} \rightarrow \mathbb{F}_2^{m^2}$ is a hard proof complexity generator against $\text{PCR}_{\mathbb{F}_2}$. However, even when m is arbitrarily large compared to n , the stretch of this generator is at most quadratic. In this section, we *iterate* g (in the sense of [52, 76]) to obtain a proof complexity generator of arbitrarily large stretch that is still hard against $\text{PCR}_{\mathbb{F}_2}$.

The iteration process we consider is a weaker notion than s -iterability introduced by Krajíček [52] to boost the stretch of a variant of the Nisan generator [61]. Namely, we use a tree-based construction following Razborov [76], in which the copies of the rank generator are arranged in a q -ary tree. This notion of iterability is sufficient for our applications in Section 8.3.

At a high level, the iterated process is executed as follows.⁷ Let Σ be a finite alphabet of size q . We consider the complete q -ary tree of depth κ , whose nodes are indexed by strings π of length at most κ over the alphabet Σ . Each node π is associated with a copy of the principle

$$X^{(\pi)}Y^{(\pi)} = A^{(\pi)},$$

where $X^{(\pi)}$ and $Y^{(\pi)}$ are the input matrices.

We restrict the entries of $A^{(\pi)}$ to be either constants or variables from the input matrices of copies of g with indices π' with $|\pi'| > |\pi|$. For the applications in Section 8.3, we consider a concrete iteration procedure, where each $A^{(\pi)}$ is constant when $|\pi| = \kappa$ and is formed from the matrices $X^{(\pi*j)}$ and $Y^{(\pi*j)}$ for $j \in \Sigma$ when $|\pi| < \kappa$.

When $\kappa = 0$, the iterated formula is simply the original rank principle. Increasing κ increases the stretch of the generator, but also makes the resulting formulas easier to refute. Consequently, proving lower bounds for the iterated formulas is harder than proving them for the original weak rank principle.

⁷The actual constructions below differ slightly, but the underlying idea is the same.

5.4.1 Iterability of $\text{Rank}_n^m(A)$

Here we iterate the *algebraic* version of the rank principle. We will use a more limited notion of iterability than s -iterability as in [52, 76], which is sufficient for our purposes. We shall follow the general argument used by Razborov in [76].

We denote by $\Sigma^{\leq \kappa}$ the set of strings of length at most κ over the alphabet Σ . Similarly, $\Sigma^{< \kappa}$ denotes the strings of length less than κ over the alphabet Σ . We use $()$ to denote the empty string.

Definition 5.32 (The iterated formula $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma)$). Let $m > n > 0$ and $\kappa \geq 0$ be integers and Σ a finite alphabet. For every string $\pi \in \Sigma^{\leq \kappa}$, let $X^{(\pi)} = (x_{i,k}^{(\pi)})_{i \in [m], k \in [n]}$ be a variable matrix. Let $Y = (y_{k,i})_{k \in [n], i \in [m]}$ be a variable matrix. Let $\mathbf{A} = \{A^{(\pi)} : \pi \in \Sigma^{\leq \kappa}\}$ be a collection of $m \times m$ matrices, where, for every $\pi \in \Sigma^{\leq \kappa}$ and $i, j \in [m]$, the entry $A_{i,j}^{(\pi)}$ is either a Boolean constant or a variable from $X^{(\pi')}$, for some $\pi' \in \Sigma^{\leq \kappa}$ with $|\pi'| > |\pi|$. In particular, all entries of $A^{(\pi)}$, where $\pi \in \Sigma^\kappa$, are constant. The *iterated rank formula* $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma)$ is the polynomial system $\bigcup_{\pi \in \Sigma^{\leq \kappa}} X^{(\pi)} Y = A^{(\pi)}$.

Note that the formula $\text{IRank}_n^m(\mathbf{A}, 0, \Sigma)$, where $\mathbf{A} = \{A^{(0)}\}$, is precisely the (non-iterated) formula $\text{Rank}_n^m(A^{(0)})$. Our lower bound will depend on the number of iterations κ , and for the application, we need Σ with $|\Sigma| = n^{\Omega(1)}$.

We need the following measure, which describes how many rows are mentioned in a term.

Definition 5.33. Let t be a term in the variables of $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma)$. We say that $i \in [m]$ is $X^{(\pi)}$ -row-mentioned in t if there is $k \in [n]$ such that $x_{i,k}^{(\pi)}$ or $\bar{x}_{i,k}^{(\pi)}$ appears in t . We say that $j \in [m]$ is Y^T -row-mentioned in t if there is $k \in [n]$ such that $y_{k,j}$ or $\bar{y}_{k,j}$ appears in t .

Definition 5.34 (Row degree). Let t be a term in the variables of $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma)$. For every $\pi \in \Sigma^{\leq \kappa}$, let $I^{(\pi)}$ be the set of all $i \in [m]$ that are $X^{(\pi)}$ -row-mentioned in t . Similarly, let J be the set of all $j \in [m]$ that are Y^T -row-mentioned in t . In this case we shall denote the mentioned rows as a pair of row-sets, and say that t mentions rows $(\{I^{(\pi)} : \pi \in \Sigma^{\leq \kappa}\}, J)$. The *row degree* of t is defined as the sum $\sum_{\pi \in \Sigma^{\leq \kappa}} |I^{(\pi)}| + |J|$.

We will construct the restriction that w.h.p. sets to zero any term of large row degree. We need the following auxiliary matrices.

$$\mathcal{D} = \begin{pmatrix} 1 & * & 1 & 0 & 0 & * & * & * \\ 0 & 1 & * & 1 & * & * & * & 0 \\ * & * & 0 & * & 1 & 0 & 1 & * \\ * & 0 & * & * & * & 1 & 0 & 1 \end{pmatrix}, \quad \mathcal{E} = \begin{pmatrix} 1 & 0 & 0 & 1 & * & * & * & * \\ 0 & * & 1 & * & 1 & * & 0 & * \\ * & 1 & * & 0 & * & 0 & * & 1 \\ * & * & * & * & 0 & 1 & 1 & 0 \end{pmatrix}.$$

The properties of $\mathcal{D}$ and $\mathcal{E}$ that we will use are the following.

1. Every pair of rows $\mathcal{D}_i$ of $\mathcal{D}$ and $\mathcal{E}_j$ of $\mathcal{E}$ satisfies the following:
 - there is an odd number of indices $k \in [8]$ such that $\mathcal{D}_{i,k} = \mathcal{E}_{j,k} = *$;
 - there is an even number of indices $k \in [8]$ such that $\mathcal{D}_{i,k} = *$ and $\mathcal{E}_{j,k} = 1$;
 - there is an even number of indices $k \in [8]$ such that $\mathcal{D}_{i,k} = 1$ and $\mathcal{E}_{j,k} = *$.
2. Each column of $\mathcal{D}$ and $\mathcal{E}$ contains at least one 0 and at least one 1.

We will construct a random substitution ρ to reduce the row degree, and then follow with a substitution σ to reduce the total degree.

Intuitively, the substitution ρ partitions an $m \times n$ matrix into eight blocks. We randomly choose some row type $r_i^{(\pi)}$ defined by $\mathcal{D}$ for each row $i \in [m]$ of $X^{(\pi)}$. Then we apply this row type to the row

i : the value of the first entry of the row type $r_i^{(\pi)}$ is assigned to the first block, the value of the second entry to the second block, etc. The assignment for Y^T is similar. This results in matrices that have 0, 1 and $*$ entries. The $*$ entries are then replaced with fresh variables with the appropriate indices, forming $m \times \tilde{n}$ matrices.

Definition 5.35 (Substitution ρ). Assume $m > n \geq 16$ and n is divisible by 16. Let $\tilde{n} = n/8$ and then let $\tilde{X}^{(\pi)} = (\tilde{x}_{i,k}^{(\pi)})_{i \in [m], k \in [\tilde{n}]}$, where $\pi \in \Sigma^{\leq \kappa}$, and $\tilde{Y} = (\tilde{y}_{k,i})_{k \in [\tilde{n}], i \in [m]}$ be fresh variables. A random substitution ρ is chosen as follows.

1. For each $i \in [m]$ uniformly and independently sample $s_i \in [4]$ and $r_i^{(\pi)} \in [4]$, for all $\pi \in \Sigma^{\leq \kappa}$.
2. For each $\pi \in \Sigma^{\leq \kappa}$, $i \in [m]$ and $k \in [n]$, let $\nu = \lfloor \frac{k-1}{\tilde{n}} \rfloor + 1$ and set

$$\rho(x_{i,k}^{(\pi)}) = \begin{cases} \tilde{x}_{i,((k-1) \bmod \tilde{n})+1}^{(\pi)} & \text{if } \mathcal{D}_{r_i^{(\pi)}, \nu} = *, \\ \mathcal{D}_{r_i^{(\pi)}, \nu} & \text{if } \mathcal{D}_{r_i^{(\pi)}, \nu} \in \{0, 1\}. \end{cases}$$

3. For each $j \in [m]$ and $k \in [n]$, let $\nu = \lfloor \frac{k-1}{\tilde{n}} \rfloor + 1$ as before and set

$$\rho(y_{k,j}) = \begin{cases} \tilde{y}_{((k-1) \bmod \tilde{n})+1, j} & \text{if } \mathcal{E}_{s_j, \nu} = *, \\ \mathcal{E}_{s_j, \nu} & \text{if } \mathcal{E}_{s_j, \nu} \in \{0, 1\}. \end{cases}$$

4. For any variable θ mapped to 0 or 1 by ρ , set $\rho(\bar{\theta}) = 1 - \rho(\theta)$.
5. For any variable θ not mapped to a Boolean constant by ρ , set $\rho(\bar{\theta}) = \overline{\rho(\theta)}$.

Lemma 5.36. Assume $\kappa \geq 0$, $m > n \geq 16$, n divisible by 16 and $\mathbf{A} = \{A^{(\pi)} : \pi \in \Sigma^{\leq \kappa}\}$ is a collection of $m \times m$ matrices as in Definition 5.32. Let t be a term in the variables of $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma)$ that mentions rows $(\{I^{(\pi)} : \pi \in \Sigma^{\leq \kappa}\}, J)$. Then $\Pr[t \upharpoonright \rho \neq 0] \leq (3/4)^{\sum_{\pi \in \Sigma^{\leq \kappa}} |I^{(\pi)}| + |J|}$.

Proof. For each i from the set of rows $I^{(\pi)}$ choose a single variable $\gamma_i^{(\pi)}$ in t that mentions row i , and for each j from the set J choose a single variable δ_j in t . Let $\Gamma_i^{(\pi)}$ be the event that $(\gamma_i^{(\pi)} \upharpoonright \rho \neq 0)$ and Δ_j the event that $(\delta_j \upharpoonright \rho \neq 0)$. Then $\bigcup_{\pi \in \Sigma^{\leq \kappa}} \{\Gamma_i^{(\pi)} : i \in I^{(\pi)}\} \cup \{\Delta_j : j \in J\}$ is a set of mutually independent events and each of them happens with probability at most $3/4$. $\square$

The substitution ρ transforms any refutation of $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma)$ into a refutation of $\text{IRank}_{\tilde{n}}^m(\mathbf{A}_1, \kappa, \Sigma)$, where $\mathbf{A}_1 = \mathbf{A} \upharpoonright \rho$.

Lemma 5.37. Assume $\kappa \geq 0$, $m > n \geq 16$, n is divisible by 16 and $\mathbf{A} = \{A^{(\pi)} : \pi \in \Sigma^{\leq \kappa}\}$ is a collection of $m \times m$ matrices as in Definition 5.32. Let $\mathbf{A}_1$ be a collection of matrices defined as $\mathbf{A}_1 = \{A^{(\pi)} \upharpoonright \rho : A^{(\pi)} \in \mathbf{A}\}$. Then for every $\text{PCR}_{\mathbb{F}_2}$ refutation Π of $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma)$, $\Pi \upharpoonright \rho$ is a refutation of $\text{IRank}_{n/8}^m(\mathbf{A}_1, \kappa, \Sigma)$.

Proof. The proof is straightforward: each axiom $\sum_{k \in [n]} x_{i,k}^{(\pi)} y_{k,j} = A_{i,j}$ of $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma)$ transforms into an axiom $\sum_{k \in [n/8]} \tilde{x}_{i,k}^{(\pi)} \tilde{y}_{k,j} = \tilde{A}_{i,j}$ of $\text{IRank}_{n/8}^m(\mathbf{A}_1, \kappa, \Sigma)$. Consider the axiom

$$\sum_{k \in [n]} x_{i,k}^{(\pi)} y_{k,j} = A_{i,j}^{(\pi)}.$$

By our choice of $\mathcal{D}$ and $\mathcal{E}$, we have an odd number of $(*, *)$ -blocks, that are replaced with $\sum_{k \in [n/8]} \tilde{x}_{i,k}^{(\pi)} \tilde{y}_{k,j}$, and an even number of $(1, *)$ -blocks and $(*, 1)$ -blocks, which cancel out. There is no shift in the constant

term, since each $(1, 1)$ -block contributes $\sum_{k \in [n/8]} 1 = 0$ due to our choice of n . The remaining blocks contribute only zeroes. Thus, the axiom transforms into

$$\sum_{k \in [n/8]} \tilde{x}_{i,k}^{(\pi)} \tilde{y}_{k,j} = \rho(A_{i,j}^{(\pi)}).$$

□

We now construct the substitution σ that reduces the total degree of any term with small row degree. The construction is similar to [Definition 5.4](#).

Definition 5.38 (Substitution σ). Assume $\kappa \geq 0$, $m > n \geq 2$ and $\mathbf{A} = \{A^{(\pi)} : \pi \in \Sigma^{\leq \kappa}\}$ is a collection of $m \times m$ matrices as in [Definition 5.32](#). A random substitution σ is chosen as follows.

1. For each $k \in [n-1]$, uniformly and independently sample $c_k \in \{*, 1, 0\}$. Then set $c_n = |\{k \in [n-1] : c_k = 1\}| \bmod 2$. This guarantees that $|\{k \in [n] : c_k = 1\}|$ is even.
2. Let $\tilde{n} = |\{k \in [n] : c_k = *\}|$ and let $\tilde{X}^{(\pi)} = (\tilde{x}_{i,k}^{(\pi)})_{i \in [m], k \in [\tilde{n}]}$, where $\pi \in \Sigma^{\leq \kappa}$, and $\tilde{Y} = (\tilde{y}_{k,j})_{k \in [\tilde{n}], j \in [m]}$ be fresh variables.
3. Let ι be a bijection between $\{k \in [n] : c_k = *\}$ and $[\tilde{n}]$.
4. For each $\pi \in \Sigma^{\leq \kappa}$, $i \in [m]$ and $k \in [n]$ set

$$\sigma(x_{i,k}^{(\pi)}) = \begin{cases} \tilde{x}_{i,\iota(k)}^{(\pi)} & \text{if } c_k = *, \\ c_k & \text{if } c_k \in \{0, 1\}. \end{cases}$$

5. For each $j \in [m]$ and $k \in [n]$ set

$$\sigma(y_{k,j}) = \begin{cases} \tilde{y}_{\iota(k),j} & \text{if } c_k = *, \\ c_k & \text{if } c_k \in \{0, 1\}. \end{cases}$$

6. For any variable θ mapped to 0 or 1 by σ , set $\sigma(\bar{\theta}) = 1 - \sigma(\theta)$.
7. For any variable θ not mapped to a Boolean constant by σ , set $\sigma(\bar{\theta}) = \overline{\sigma(\theta)}$.

σ behaves similarly to the substitution from [Definition 5.4](#). We have the following analogue of [Lemma 5.5](#).

Lemma 5.39. Assume $\kappa \geq 0$, $m > n \geq 2$ and $\mathbf{A} = \{A^{(\pi)} : \pi \in \Sigma^{\leq \kappa}\}$ is a collection of $m \times m$ matrices as in [Definition 5.32](#). Let t be a term in the variables of $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma)$ of degree d that does not mention the last column and mentions rows $(\{I^{(\pi)} : \pi \in \Sigma^{\leq \kappa}\}, J)$. Then $\Pr[t \mid \sigma \neq 0] \leq (2/3)^{\frac{d}{\sum_{\pi \in \Sigma^{\leq \kappa}} |I^{(\pi)}| + |J|}}$. Furthermore, $\Pr[\tilde{n} < (n-1)/6] \leq e^{-\frac{n-1}{24}}$.

Proof. Let $K^{(\pi)}$ be the set of all $k \in [n]$ such that for some $i \in [m]$, either $x_{i,k}^{(\pi)}$ or $\bar{x}_{i,k}^{(\pi)}$ appears in t . Let L be the set of all $\ell \in [n]$ such that for some $j \in [m]$, either $y_{\ell,j}$ or $\bar{y}_{\ell,j}$ appears in t .

The degree of t is then upper-bounded by

$$\left(\sum_{\pi \in \Sigma^{\leq \kappa}} |I^{(\pi)}| \cdot |K^{(\pi)}| + |J| \cdot |L| \right) \leq \left(\sum_{\pi \in \Sigma^{\leq \kappa}} |I^{(\pi)}| + |J| \right) \cdot \max \left(\{|K^{(\pi)}| : \pi \in \Sigma^{\leq \kappa}\} \cup \{|L|\} \right).$$

By the pigeonhole principle, one of the sets $\{K^{(\pi)} : \pi \in \Sigma^{\leq \kappa}\} \cup \{L\}$ must have size at least $d' = \frac{d}{(\sum_{\pi \in \Sigma^{\leq \kappa}} |I^{(\pi)}| + |J|)}$. Let S be such a set and $\{\gamma_s : s \in S\}$ some variables in t corresponding to S . For every $s \in S$, let Γ_s be the event that $(\gamma_s \upharpoonright \sigma \neq 0)$. Then $\{\Gamma_s : s \in S\}$ is a set of at least d' mutually independent events, each of which happens with probability at most $2/3$.

The proof of the latter statement is the same as the proof of [Lemma 5.5](#). $\square$

The substitution σ transforms any refutation of $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma)$ into a refutation of $\text{IRank}_n^m(\mathbf{A}_2, \kappa, \Sigma)$, where $\mathbf{A}_2 = \mathbf{A} \upharpoonright \sigma$.

Lemma 5.40. *Assume $\kappa \geq 0$, $m > n \geq 2$ and $\mathbf{A} = \{A^{(\pi)} : \pi \in \Sigma^{\leq \kappa}\}$ is a collection of $m \times m$ matrices as in [Definition 5.32](#). Let $\mathbf{A}_2$ be a collection of matrices defined as $\mathbf{A}_2 = \{A^{(\pi)} \upharpoonright \sigma : \pi \in \Sigma^{\leq \kappa}\}$. Then for every $\text{PCR}_{\mathbb{F}_2}$ refutation Π of $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma)$, $\Pi \upharpoonright \sigma$ is a refutation of $\text{IRank}_n^m(\mathbf{A}_2, \kappa, \Sigma)$.*

Proof. The proof is the same as the proof of [Lemma 5.6](#). $\square$

We establish a lower bound on the degree of $\text{PCR}_{\mathbb{F}_2}$ refutations of $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma)$ by constructing a reduction to a single copy of $\text{Rank}_n^m(I_m)$.

Definition 5.41 (Reduction τ). Assume $\kappa \geq 0$, $m > n > 0$. Let $X = (x_{i,k})_{i \in [m], k \in [n]}$ be fresh variables. A substitution τ is defined inductively on $|\pi|$ in decreasing order.

1. For each $i \in [m]$ and $k \in [n]$ set

$$\begin{aligned} \tau(x_{i,k}^{(\pi)}) &= ((A^{(\pi)} \upharpoonright \tau)X)_{i,k} = \sum_{s \in [m]} \tau(A_{i,s}^{(\pi)})x_{s,k}, \\ \tau(y_{k,i}) &= y_{k,i}. \end{aligned}$$

2. For any variable θ , set $\tau(\bar{\theta}) = 1 - \tau(\theta)$.⁸

Note that $A^{(\pi)} \upharpoonright \tau = A^{(\pi)}$ for every $\pi \in \Sigma^{\leq \kappa}$ since the entries of $A^{(\pi)}$ can only be constant.

Lemma 5.42. *Assume $\kappa \geq 0$, $m > n > 0$ and $\mathbf{A} = \{A^{(\pi)} : \pi \in \Sigma^{\leq \kappa}\}$ is a collection of $m \times m$ matrices as in [Definition 5.32](#). Then for any $\text{PCR}_{\mathbb{F}_2}$ refutation Π of $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma)$ with degree d , $\Pi \upharpoonright \tau$ can be transformed into a refutation of $\text{Rank}_n^m(I_m)$ with degree at most $d(\kappa + 2)$.*

Proof. Let Π be a $\text{PCR}_{\mathbb{F}_2}$ refutation of $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma)$.

It follows by induction on $|\pi|$, that for every $\pi \in \Sigma^{\leq \kappa}$, $i \in [m]$, and $k \in [n]$, $\tau(x_{i,k}^{(\pi)})$ (and thus $\tau(x_{i,k}^{(\pi)})$) is a polynomial of degree at most $\kappa - |\pi| + 1$. In particular, the entries of $A^{(\pi)} \upharpoonright \tau$ are polynomials of degree at most $\kappa - |\pi|$.

It implies that all polynomials from $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma) \upharpoonright \tau$ can be easily derived from the axioms of $\text{Rank}_n^m(I_m)$.

Claim 5.43. *Every polynomial from $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma) \upharpoonright \tau$ has a degree- $(\kappa + 2)$ proof from $\text{Rank}_n^m(I_m)$.*

Proof. Fix $\pi \in \Sigma^{\leq \kappa}$. As shown above, all the entries of $A^{(\pi)} \upharpoonright \tau$ are polynomials of degree at most $\kappa - |\pi|$.

Let $i, j \in [m]$. The (i, j) th axiom of $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma)$ under τ transforms as follows

$$\left(\sum_{k \in [n]} x_{i,k}^{(\pi)} y_{k,j} + A_{i,j}^{(\pi)} \right) \upharpoonright \tau = \sum_{k \in [n]} \left(\sum_{s \in [m]} \tau(A_{i,s}^{(\pi)}) x_{s,k} \right) y_{k,j} + \tau(A_{i,j}^{(\pi)})$$

⁸The proof size may increase after applying this substitution; the degree increase is controlled in [Lemma 5.42](#). The resulting refutation will still be a PC refutation.

$$\begin{aligned}
&= \sum_{s \in [m]} \tau(A_{i,s}^{(\pi)}) \left(\sum_{k \in [n]} x_{s,k} y_{k,j} \right) + \tau(A_{i,j}^{(\pi)}) \\
&= \sum_{s \in [m] \setminus \{j\}} \tau(A_{i,s}^{(\pi)}) \left(\sum_{k \in [n]} x_{s,k} y_{k,j} \right) + \tau(A_{i,j}^{(\pi)}) \left(\sum_{k \in [n]} x_{j,k} y_{k,j} + 1 \right),
\end{aligned}$$

which is trivially derivable in degree at most $\kappa - |\pi| + 2$ from $\text{Rank}_n^m(I_m)$. $\diamond$

Thus, we can extend $\Pi \upharpoonright \tau$ to a refutation of $\text{Rank}_n^m(I_m)$ of degree at most $d(\kappa + 2)$. $\square$

We now have all the components for proving the size lower bound on the iterated formula $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma)$.

Theorem 5.44. *Assume $\kappa > 0$, $m > n \geq 16 + 192(\kappa + 2)$, $(n - 16)$ is divisible by $192(\kappa + 2)$ and $\mathbf{A} = \{A^{(\pi)} : \pi \in \Sigma^{\leq \kappa}\}$ is a collection of $m \times m$ matrices as in Definition 5.32. Then any $\text{PCR}_{\mathbb{F}_2}$ refutation of $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma)$ requires size $2^{\Omega(\sqrt{n/\kappa})}$.*

Proof. Let

$$\begin{aligned}
d_1 &= \frac{n - 16}{192(\kappa + 2)}, \\
d_2 &= \sqrt{d_1} = \sqrt{\frac{n - 16}{192(\kappa + 2)}}.
\end{aligned}$$

Assume, for the sake of contradiction, that $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma)$ has a refutation Π of size less than $(4/3)^{d_2} = 2^{\Theta(\sqrt{n/\kappa})}$.

By the union bound and Lemmas 5.36 and 5.37, there exists a substitution ρ such that $\Pi \upharpoonright \rho$ has row degree less than d_2 , and it is a refutation of $\text{IRank}_{n/8}^m(\mathbf{A} \upharpoonright \rho, \kappa, \Sigma)$. Further, extend ρ to ρ' by also fixing the last $(n/8)$ th column in each $X^{(\pi)}$ and Y^T to an all-zero column, thus obtaining $\text{IRank}_{n/8-1}^m(\mathbf{A} \upharpoonright \rho', \kappa, \Sigma)$.⁹

Since after applying ρ' every surviving term has row degree less than d_2 , Lemma 5.39 implies that any term of degree at least $2d_1$ survives σ with probability at most $(2/3)^{2d_1/d_2} = (2/3)^{2d_2}$. Hence, by the union bound, there is a choice of σ for which all terms of degree at least d_1 vanish and $\tilde{n} \geq \frac{n/8-2}{6} = \frac{n-16}{48}$. By Lemma 5.40, $\Pi \upharpoonright \rho' \upharpoonright \sigma$ is a refutation of $\text{IRank}_{\tilde{n}}^m(\mathbf{A} \upharpoonright \rho' \upharpoonright \sigma, \kappa, \Sigma)$.

By Lemma 5.42, we can extend $\Pi \upharpoonright \rho' \upharpoonright \sigma \upharpoonright \tau$ to a refutation of $\text{Rank}_{\tilde{n}}^m(I_m)$ with degree less than $2d_1(\kappa + 2) \leq \tilde{n}/2$, which contradicts Lemma 5.3. $\square$

5.4.2 Iterability of $\text{BTRank}_n^m(A)$

We start by defining all the formulas and constructions used in this section, followed by a high-level overview of the proof.

Notation. Given two sequences $(a_1, \dots, a_n)$ and $(b_1, \dots, b_m)$, we write $a \cup b$ to denote their concatenation $(a_1, \dots, a_n, b_1, \dots, b_m)$.

In this section, we make use of several specific variants of the rank principle in which the structure of the matrix Y is determined by a bipartite graph. When this graph has small left-degree (in particular, when it is an expander), each axiom $\sum_k x_{i,k} y_{k,j}$ only depends on a few variables and therefore its CNF encoding does not require too many extension variables. Unlike Section 5.4.1, the iteration structure is over the alphabet $\Sigma = \{0, 1\}$. We will use the following formulas.

⁹This step is needed because we require $16 \mid n$ in order to apply ρ .

Definition 5.45 (The iterated formula $\psi_G(\mathbf{A}, \kappa)$). Let $m, n > 0$ and $\kappa \geq 0$ be integers and G an $m \times n$ bipartite graph with parts $[m]$ and $[n]$. For convenience assume that $1 \in \mathcal{N}_G(i)$ for every $i \in [m]$. We define the formula $\psi_G(\mathbf{A}, \kappa)$ that encodes the statement of the form $\bigcup_{\pi \in \{0,1\}^{\leq \kappa}} X^{(\pi)} Y^{(\pi)} = A^{(\pi)}$, using extension variables that form a binary tree to compute partial parities. The structure of each matrix $Y^{(\pi)}$ is based on G .

The formula uses the following variables:

$$\bigcup_{\pi \in \{0,1\}^{\leq \kappa}} \left[\bigcup_{\mu \in [4]} \left(\{x_{i,k}^{(\pi,\mu)} : i \in [m], k \in [n]\} \right. \right. \\ \cup \{y_{k,j}^{(\pi,\mu)} : j \in [m], k \in \mathcal{N}_G(j)\} \\ \cup \{z_{i,j,k}^{(\pi,\mu)} : i, j \in [m], k \in \mathcal{N}_G(j)\} \\ \left. \left. \cup \{u_{i,j,\Sigma}^{(\pi,\mu)} : i, j \in [m], \Sigma \text{ is a non-empty initial segment of } \mathcal{N}_G(j)\} \right) \right. \\ \left. \cup \{w_{i,j,\ell}^{(\pi)} : i, j \in [m], \ell \in \{0,1,2\}\} \right].$$

Let $\mathbf{A} = \{A^{(\pi)} : \pi \in \{0,1\}^{\leq \kappa}\}$ be a collection of $m \times m$ matrices, where, for every $\pi \in \{0,1\}^{\leq \kappa}$ and $i, j \in [m]$, the entry $A_{i,j}^{(\pi)}$ is either a Boolean constant or a variable of the form $x_{i',k'}^{(\pi',\mu')}$ or $y_{k',j'}^{(\pi',\mu')}$ for some $i', j' \in [m]$, $k' \in [n]$, $\mu' \in [4]$ and $\pi' \in \{0,1\}^{\leq \kappa}$ with $|\pi'| > |\pi|$. In particular, all entries of $A^{(\pi)}$ are constant for $\pi \in \{0,1\}^{\leq \kappa}$.

The *iterated formula* $\psi_G(\mathbf{A}, \kappa)$ is a polynomial system consisting of the following axioms.

Binary AND Axiom: For every $\pi \in \{0,1\}^{\leq \kappa}$, $\mu \in [4]$, $i, j \in [m]$, $k \in \mathcal{N}_G(j)$, the polynomial $z_{i,j,k}^{(\pi,\mu)} = x_{i,k}^{(\pi,\mu)} y_{k,j}^{(\pi,\mu)}$.

Summation Base Axiom: For every $\pi \in \{0,1\}^{\leq \kappa}$, $\mu \in [4]$, $i, j \in [m]$, the polynomial $u_{i,j,(1)}^{(\pi,\mu)} = z_{i,j,1}^{(\pi,\mu)}$.

Summation Axiom: For every $\pi \in \{0,1\}^{\leq \kappa}$, $\mu \in [4]$, $i, j \in [m]$ and every initial segment $\Sigma \cup (k)$ of $\mathcal{N}_G(j)$ with $\Sigma \neq \emptyset$, the polynomial

$$u_{i,j,\Sigma \cup (k)}^{(\pi,\mu)} = u_{i,j,\Sigma}^{(\pi,\mu)} + z_{i,j,k}^{(\pi,\mu)}.$$

Graphically, the last two types of axioms can be represented as follows. Suppose that $\mathcal{N}_G(j) = (k_1, \dots, k_\Delta)$. For every $\ell \in [\Delta]$ define $\Sigma_\ell = (k_1, \dots, k_\ell)$. In particular, for each $\ell \in \{2, \dots, \Delta\}$, we have $\Sigma_\ell = \Sigma_{\ell-1} \cup (k_\ell)$. Then the u - and z -variables form a bamboo tree as shown in [Figure 9](#).




Figure 9: The bamboo-tree structure of the u - and z -variables.

Output Axioms: For every $\pi \in \{0,1\}^{\leq \kappa}$, $i, j \in [m]$, the following four axioms:

$$\begin{aligned} w_{i,j,1}^{(\pi)} &= u_{i,j,\mathcal{N}_G(j)}^{(\pi,1)} + u_{i,j,\mathcal{N}_G(j)}^{(\pi,2)}, \\ w_{i,j,2}^{(\pi)} &= u_{i,j,\mathcal{N}_G(j)}^{(\pi,3)} + u_{i,j,\mathcal{N}_G(j)}^{(\pi,4)}, \\ w_{i,j,0}^{(\pi)} &= w_{i,j,1}^{(\pi)} + w_{i,j,2}^{(\pi)}, \\ w_{i,j,0}^{(\pi)} &= A_{i,j}^{(\pi)}. \end{aligned}$$

Graphically, these axioms can be represented as shown in Figure 10.




Figure 10: The balanced tree structure of the w -variables.

Although $\psi_G(\mathbf{A}, \kappa)$ is defined as a polynomial system, it can be easily encoded as a CNF since each axiom of $\psi_G(\mathbf{A}, \kappa)$ depends only on a constant number of variables.

The following formula $\phi_G(\mathbf{A}, \kappa)$ has a simpler structure, and in particular does not involve the parameter μ or the w -variables. It can be viewed as a substitution instance of $\psi_G(\mathbf{A}, \kappa)$.

Definition 5.46 (Formula $\phi_G(\mathbf{A}, \kappa)$). Let $m, n > 0$ and $\kappa \geq 0$ be integers and let G be an $m \times n$ bipartite graph. Consider the following variables:

$$\bigcup_{\pi \in \{0,1\}^{\leq \kappa}} \left[\{x_{i,k}^{(\pi)} : i \in [m], k \in [n]\} \right]$$

$$\begin{aligned} & \cup \{y_{k,j}^{(\pi)} : j \in [m], k \in \mathcal{N}_G(j)\} \\ & \cup \{z_{i,j,k}^{(\pi)} : i, j \in [m], k \in \mathcal{N}_G(j)\} \\ & \cup \{u_{i,j,\Sigma}^{(\pi)} : i, j \in [m], \Sigma \text{ is a non-empty initial segment of } \mathcal{N}_G(j)\} \end{aligned} \Bigg].$$

Let $\mathbf{A} = \{A^{(\pi)} : \pi \in \{0,1\}^{\leq \kappa}\}$ be a collection of $m \times m$ matrices, where, for every $\pi \in \{0,1\}^{\leq \kappa}$ and $i, j \in [m]$, the entry $A_{i,j}^{(\pi)}$ is either a Boolean constant or a variable of the form $x_{i',k'}^{(\pi')}$ or $y_{k',j'}^{(\pi')}$ for some $i', j' \in [m]$, $k' \in [n]$ and $\pi' \in \{0,1\}^{\leq \kappa}$ with $|\pi'| > |\pi|$.

Binary AND Axiom: For every $\pi \in \{0,1\}^{\leq \kappa}$, $i, j \in [m]$, $k \in \mathcal{N}_G(j)$, the polynomial $z_{i,j,k}^{(\pi)} = x_{i,k}^{(\pi)} y_{k,j}^{(\pi)}$.

Summation Base Axiom: For every $\pi \in \{0,1\}^{\leq \kappa}$, $i, j \in [m]$, the polynomial $u_{i,j,(1)}^{(\pi)} = z_{i,j,1}^{(\pi)}$.

Summation Axiom: For every $\pi \in \{0,1\}^{\leq \kappa}$, $i, j \in [m]$ and every non-empty initial segment $\Sigma \cup (k)$ of $\mathcal{N}_G(j)$, the polynomial $u_{i,j,\Sigma \cup (k)}^{(\pi)} = u_{i,j,\Sigma}^{(\pi)} + z_{i,j,k}^{(\pi)}$.

Output Axiom: For every $\pi \in \{0,1\}^{\leq \kappa}$, $i, j \in [m]$, the polynomial $u_{i,j,\mathcal{N}_G(j)}^{(\pi)} = A_{i,j}^{(\pi)}$.

We also consider an even simpler formula $\tilde{\phi}_G(\mathbf{A}, \kappa)$, which is the algebraic analogue of $\phi_G(\mathbf{A}, \kappa)$.

Definition 5.47 (Formula $\tilde{\phi}_G(\mathbf{A}, \kappa)$). Let $m, n > 0$ and $\kappa \geq 0$ be integers and let G be an $m \times n$ bipartite graph. Consider the following variables:

$$\bigcup_{\pi \in \{0,1\}^{\leq \kappa}} \left[\{x_{i,k}^{(\pi)} : i \in [m], k \in [n]\} \cup \{y_{k,j}^{(\pi)} : j \in [m], k \in \mathcal{N}_G(j)\} \right].$$

Let $\mathbf{A}$ be defined as in [Definition 5.46](#). We define $\tilde{\phi}_G(\mathbf{A}, \kappa)$ as the polynomial system

$$\bigcup_{\pi \in \{0,1\}^{\leq \kappa}} \bigcup_{i,j \in [m]} \left\{ \sum_{k \in \mathcal{N}_G(j)} x_{i,k}^{(\pi)} y_{k,j}^{(\pi)} = A_{i,j}^{(\pi)} \right\}.$$

In what follows we use the graph gadgets $P(G)$ and $H(G)$.

Definition 5.48 (Graph $P(G)$). Let $m, n > 0$ and let G be an $m \times n$ bipartite graph. Given an adjacency matrix M_G of G , define $P(G)$ as the $m \times (n+2)$ bipartite graph with the adjacency matrix

$$M_{P(G)} = \begin{bmatrix} 1 & 1 \\ \vdots & \vdots & M_G \\ 1 & 1 \end{bmatrix}.$$

Definition 5.49 (Graph $H(G)$). Let $m, n > 0$ and let G be an $m \times n$ bipartite graph. Given an adjacency matrix M_G of G , define $H(G)$ as the $m \times (8n+8)$ bipartite graph with the adjacency matrix

$$M_{H(G)} = \underbrace{\begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ M_G & \vdots & M_G & \vdots & M_G & \vdots & M_G & \vdots & M_G & \vdots & M_G & \vdots \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix}}_8.$$

Finally, we will also need the following lower bound on the PC degree of $\text{FPHP}(G)$.

Proposition 5.50 ([57]). *Let G be an (r, c) -boundary expander with left-degree at most $\Delta_{\max}$. Then any $\text{PCR}_{\mathbb{F}}$ refutation of $\text{FPHP}(G)$ requires degree strictly larger than $\frac{cr}{2\Delta_{\max}}$.*

The main result of this subsection is a lower bound on the size of $\text{PCR}_{\mathbb{F}_2}$ refutations of $\psi_{P(H(G))}(\mathbf{A}, \kappa)$.

Theorem 5.51. *Let $n > 0$, $N = 8n + 8$, $m = 16(N + 2)$, $\kappa \leq n^\alpha$ for some $\alpha < 1$, and let G be an $m \times n$ (n^δ, d) -lossless expander with parts $[m]$ and $[n]$, the minimum left-degree at least $\Delta_{\min} = 50 \log n$, the maximum left-degree at most $\Delta_{\max} = 150 \log n$, and $\delta > 4\alpha$. Let $\mathbf{A} = \{A^{(\pi)} : \pi \in \Sigma^{\leq \kappa}\}$ be a collection of $m \times m$ matrices as in Definition 5.45. Then any $\text{PCR}_{\mathbb{F}_2}$ refutation of $\psi_{P(H(G))}(\mathbf{A}, \kappa)$ requires size $2^{\Omega((n^\delta/\kappa)^{1/3})}$.*

The proof of this theorem is executed as follows.

1. Assume, for the sake of contradiction, that $\psi_{P(H(G))}(\mathbf{A}, \kappa)$ has a refutation Π of size $2^{\epsilon(n^\delta/\kappa)^{1/3}}$. Let d_1 be a certain threshold parameter with $d_1 = \Theta(n^\delta/\kappa)$.
2. We construct a random substitution ρ_1 that turns Π into a refutation Π_1 of $\phi_{H(G)}(\mathbf{A}_1, \kappa)$, where $\mathbf{A}_1 = \mathbf{A} \upharpoonright \rho_1$. ρ_1 eliminates all terms t from Π that either have high degree in the u -variables or have high degree (exceeding d_1) in the *positive* occurrences of the z -variables. Moreover, it substitutes all the w -variables with either constants or x - or y -variables.
3. We proceed by constructing another random substitution ρ_2 that further eliminates all terms in Π_1 that mention many columns of G (specifically, exceeding $d_1^{1/3}$) using x -, y - or z -variables. The resulting refutation Π_2 is a refutation of $\phi_{H(G')}(\mathbf{A}_2, \kappa)$, where $\mathbf{A}_2 = \mathbf{A}_1 \upharpoonright \rho_2$ and G' is still a sufficiently good expander.
4. Then we construct a substitution ρ_3 that eliminates all terms of high degree (again, exceeding d_1) in x - and y -variables and negative occurrences in z -variables. The resulting refutation Π_3 is a refutation of $\phi_{G'}(\mathbf{A}_3, \kappa)$, where $\mathbf{A}_3 = \mathbf{A}_2 \upharpoonright \rho_3$.
5. Finally, we reduce $\text{FPHP}(G')$ to $\tilde{\phi}_{G'}(\mathbf{A}_3, \kappa)$, while increasing the degree of a refutation by at most $\Theta(\kappa)$. Combining these degree bounds, we obtain a refutation Π_4 of $\text{FPHP}(G')$ of degree $\Theta(d_1 \kappa)$, which contradicts Proposition 5.50 due to our choice of parameters.

Let d be a threshold parameter. We begin by defining a restriction ρ_1 that w.h.p. eliminates any term t that either $U^{(\pi)}$ -left-row-mentions or $U^{(\pi)}$ -right-row-mentions at least d rows, or positively- $Z^{(\pi)}$ -left-row-mentions or positively- $Z^{(\pi)}$ -right-row-mentions d rows. Moreover, ρ_1 is defined so that it substitutes all w -variables with either constants or x - or y -variables. The construction uses the auxiliary matrices $\mathcal{B}$ and $\mathcal{C}$.

$$\mathcal{B} = \begin{pmatrix} * & * & 0 & 0 \\ 0 & 0 & * & * \end{pmatrix}, \quad \mathcal{C} = \begin{pmatrix} * & 0 & * & 0 \\ 0 & * & 0 & * \end{pmatrix}.$$

Every row $\mathcal{B}_i$ of $\mathcal{B}$ and every row $\mathcal{C}_j$ of $\mathcal{C}$ satisfy the following properties:

- there is exactly one $k \in [4]$ satisfying $\mathcal{B}_{i,k} = \mathcal{C}_{j,k} = *$;
- the remaining indices $\ell \in [4] \setminus \{k\}$ satisfy $\mathcal{B}_{i,\ell} = 0$ or $\mathcal{C}_{j,\ell} = 0$.

Moreover, each column of $\mathcal{B}$ and $\mathcal{C}$ contains one 0-entry.

We will define ρ_1 only for graphs $P(G)$, i.e., those that have two all-ones columns in their adjacency matrix. Let G be a bipartite graph with parts $[m]$ and $[n]$.

Definition 5.52. Let t be a term in the variables of $\psi_{P(G)}(\mathbf{A}, \kappa)$. We say that $i \in [m]$ is $U^{(\pi)}$ -left-row-mentioned in t if there are $\mu \in [4]$, $j \in [m]$ and a non-empty initial segment Σ of $\mathcal{N}_{P(G)}(j)$ with $\{1, 2\} \subsetneq \Sigma$ such that $u_{i,j,\Sigma}^{(\pi,\mu)}$ or $\bar{u}_{i,j,\Sigma}^{(\pi,\mu)}$ appears in t . We say that $j \in [m]$ is $U^{(\pi)}$ -right-row-mentioned in t if there are $\mu \in [4]$, $i \in [m]$ and a non-empty initial segment Σ of $\mathcal{N}_{P(G)}(j)$ with $\{1, 2\} \subsetneq \Sigma$ such that $u_{i,j,\Sigma}^{(\pi,\mu)}$ or $\bar{u}_{i,j,\Sigma}^{(\pi,\mu)}$ appears in t . We say that $i \in [m]$ is *positively- $Z^{(\pi)}$ -left-row-mentioned* in t if there are $\mu \in [4]$, $j \in [m]$ and $k \in \{3, \dots, n+2\}$ such that $z_{i,j,k}^{(\pi,\mu)}$ appears in t . We say that $j \in [m]$ is *positively- $Z^{(\pi)}$ -right-row-mentioned* in t if there are $\mu \in [4]$, $i \in [m]$ and $k \in \{3, \dots, n+2\}$ such that $z_{i,j,k}^{(\pi,\mu)}$ appears in t .

We substitute the first two columns with constants; therefore, we do not consider them in the above definition.

Definition 5.53 (Substitution ρ_1 for $\psi_{P(G)}(\mathbf{A}, \kappa)$). Let $\kappa \geq 0$, $m, n > 0$ and let G be an $m \times n$ bipartite graph with parts $[m]$ and $[n]$.

For every $\pi \in \{0, 1\}^{\leq \kappa}$, let $\{\tilde{x}_{i,k}^{(\pi)} : i \in [m], k \in [n]\}$, $\{\tilde{y}_{k,j}^{(\pi)} : j \in [m], k \in \mathcal{N}_G(j)\}$, $\{\tilde{z}_{i,j,k}^{(\pi)} : i, j \in [m], k \in \mathcal{N}_G(j)\}$ and $\{\tilde{u}_{i,j,\Sigma}^{(\pi)} : i, j \in [m], \Sigma \text{ is a non-empty initial segment of } \mathcal{N}_G(j)\}$ be fresh variables. The random substitution ρ_1 for $\psi_{P(G)}(\mathbf{A}, \kappa)$ is defined as follows. For every $\pi \in \{0, 1\}^{\leq \kappa}$:

1. For each $i \in [m]$, independently and uniformly sample $b_i^{(\pi)}, c_i^{(\pi)} \in [2]$ and $\alpha_i^{(\pi)}, \beta_i^{(\pi)} \in \{0, 1\}$.
2. For each $\mu \in [4]$, $i \in [m]$, $k \in [n+2]$, set

$$\rho_1(x_{i,k}^{(\pi,\mu)}) = \begin{cases} \alpha_i^{(\pi)} & \text{if } k = 1, \\ 1 & \text{if } k = 2, \\ 0 & \text{if } k \geq 3 \text{ and } \mathcal{B}_{b_i^{(\pi)}, \mu} = 0, \\ \tilde{x}_{i,k-2}^{(\pi)} & \text{if } k \geq 3 \text{ and } \mathcal{B}_{b_i^{(\pi)}, \mu} = *. \end{cases}$$

3. For each $\mu \in [4]$, $j \in [m]$, $k \in \mathcal{N}_{P(G)}(j)$, set

$$\rho_1(y_{k,j}^{(\pi,\mu)}) = \begin{cases} 1 & \text{if } k = 1, \\ \beta_j^{(\pi)} & \text{if } k = 2, \\ 0 & \text{if } k \geq 3 \text{ and } \mathcal{C}_{c_j^{(\pi)}, \mu} = 0, \\ \tilde{y}_{k-2,j}^{(\pi)} & \text{if } k \geq 3 \text{ and } \mathcal{C}_{c_j^{(\pi)}, \mu} = *. \end{cases}$$

4. For each $\mu \in [4]$, $i, j \in [m]$, $k \in \mathcal{N}_{P(G)}(j)$, set

$$\rho_1(z_{i,j,k}^{(\pi,\mu)}) = \begin{cases} \alpha_i^{(\pi)} & \text{if } k = 1, \\ \beta_j^{(\pi)} & \text{if } k = 2, \\ 0 & \text{if } k \geq 3 \text{ and either } \mathcal{B}_{b_i^{(\pi)}, \mu} = 0 \text{ or } \mathcal{C}_{c_j^{(\pi)}, \mu} = 0, \\ \tilde{z}_{i,j,k-2}^{(\pi)} & \text{if } k \geq 3 \text{ and } \mathcal{B}_{b_i^{(\pi)}, \mu} = \mathcal{C}_{c_j^{(\pi)}, \mu} = *. \end{cases}$$

5. For each $\mu \in [4]$, $i, j \in [m]$ and every non-empty initial segment Σ of $\mathcal{N}_{P(G)}(j)$, let $\Sigma' = (k-2 :$

$k \in \Sigma, k > 2$) and set

$$\rho_1(u_{i,j,\Sigma}^{(\pi,\mu)}) = \begin{cases} \alpha_i^{(\pi)} & \text{if } \Sigma = (1), \\ \alpha_i^{(\pi)} + \beta_j^{(\pi)} & \text{if } \Sigma = (1, 2), \\ \tilde{u}_{i,j,\Sigma'}^{(\pi)} & \text{if } \Sigma \setminus \{1, 2\} \neq \emptyset, \mathcal{B}_{b_i^{(\pi)}, \mu} = \mathcal{C}_{c_j^{(\pi)}, \mu} = * \text{ and } \alpha_i^{(\pi)} = \beta_j^{(\pi)}, \\ \bar{u}_{i,j,\Sigma'}^{(\pi)} & \text{if } \Sigma \setminus \{1, 2\} \neq \emptyset, \mathcal{B}_{b_i^{(\pi)}, \mu} = \mathcal{C}_{c_j^{(\pi)}, \mu} = * \text{ and } \alpha_i^{(\pi)} \neq \beta_j^{(\pi)}, \\ \alpha_i^{(\pi)} + \beta_j^{(\pi)} & \text{otherwise.} \end{cases}$$

6. For each $\mu \in [4]$, $i, j \in [m]$, set

$$\begin{aligned} \rho_1(w_{i,j,0}^{(\pi)}) &= \rho_1(A_{i,j}^{(\pi)}), \\ \rho_1(w_{i,j,1}^{(\pi)}) &= \begin{cases} \rho_1(A_{i,j}^{(\pi)}) & \text{if } b_i^{(\pi)} = 1, \\ 0 & \text{if } b_i^{(\pi)} = 2, \end{cases} \\ \rho_1(w_{i,j,2}^{(\pi)}) &= \begin{cases} 0 & \text{if } b_i^{(\pi)} = 1, \\ \rho_1(A_{i,j}^{(\pi)}) & \text{if } b_i^{(\pi)} = 2. \end{cases} \end{aligned}$$

7. For any variable θ mapped to 0 or 1 by ρ_1 , set $\rho_1(\bar{\theta}) = 1 - \rho_1(\theta)$.

8. For any variable θ not mapped to a Boolean constant by ρ_1 , set $\rho_1(\bar{\theta}) = \overline{\rho_1(\theta)}$.

Lemma 5.54. *Let $\kappa \geq 0$, $m, n > 0$ and let G be an $m \times n$ bipartite graph with parts $[m]$ and $[n]$ and the maximum left-degree at most $\Delta_{\max}$. Let $\mathbf{A} = \{A^{(\pi)} : \pi \in \Sigma^{\leq \kappa}\}$ be a collection of $m \times m$ matrices as in Definition 5.45. Let t be a term in the variables of $\psi_{P(G)}(\mathbf{A}, \kappa)$ and t_u be the subterm of t that consists of all the u -variables of the form $u_{i,j,\Sigma}^{(\pi,\mu)}$ or $\bar{u}_{i,j,\Sigma}^{(\pi,\mu)}$ with $\{1, 2\} \subsetneq \Sigma$ (that is, all variables that either $U^{(\pi)}$ -left-row-mention or $U^{(\pi)}$ -right-row-mention some indices from $[m]$). Similarly, let t_{+z} be the subterm of t consisting of all the positive z -variables of the form $z_{i,j,k}^{(\pi,\mu)}$ with $k > 2$ (such variables positively- $Z^{(\pi)}$ -left-row-mention or positively- $Z^{(\pi)}$ -right-row-mention some indices from $[m]$). Assume that either $\deg(t_u)$ or $\deg(t_{+z})$ is at least d . Then*

$$\Pr[t \upharpoonright \rho_1 \neq 0] \leq \left(\sqrt{3/4}\right)^{\sqrt{\frac{d}{4\Delta_{\max}}}}.$$

Proof. Assume that $\deg(t_u) \geq d$. For each $\pi \in \{0, 1\}^{\leq \kappa}$, let $I^{(\pi)}$ be the set of indices $i \in [m]$ that are $U^{(\pi)}$ -left-row-mentioned in t_u , and let $J^{(\pi)}$ be the set of indices $j \in [m]$ that are $U^{(\pi)}$ -right-row-mentioned in t_u . Let $D = \sum_{\pi \in \{0,1\}^{\leq \kappa}} (|I^{(\pi)}| + |J^{(\pi)}|)$. Then $d \leq \deg(t_u) \leq 4D^2\Delta_{\max}$, which implies that $D \geq \sqrt{\frac{d}{4\Delta_{\max}}}$.

Let t' be a minimal subterm of t_u that for each $\pi \in \{0, 1\}^{\leq \kappa}$, $U^{(\pi)}$ -left-row-mentions $I^{(\pi)}$ and $U^{(\pi)}$ -right-row-mentions $J^{(\pi)}$. If we think of the variables $u_{i,j,\Sigma}^{(\pi,\mu)}$ (or their negations) appearing in t' as edges $((\pi, i), (\pi, j))$ in a bipartite graph, then t' corresponds to a minimal edge cover of such a graph. This cover cannot contain a path of length three, thus its connected components are stars of the form $\prod_{(j,\mu,\Sigma) \in P} u_{i_0,j,\Sigma}^{(\pi_0,\mu)} \prod_{(j,\mu,\Sigma) \in N} \bar{u}_{i_0,j,\Sigma}^{(\pi_0,\mu)}$ or $\prod_{(i,\mu,\Sigma) \in P} u_{i,j_0,\Sigma}^{(\pi_0,\mu)} \prod_{(i,\mu,\Sigma) \in N} \bar{u}_{i,j_0,\Sigma}^{(\pi_0,\mu)}$, and each of these stars mentions its own disjoint set of indices.

Assume that we have a star with r rays, say $s = \prod_{(j,\mu,\Sigma) \in P} u_{i_0,j,\Sigma}^{(\pi_0,\mu)} \prod_{(j,\mu,\Sigma) \in N} \bar{u}_{i_0,j,\Sigma}^{(\pi_0,\mu)}$. For every

$b \in [2]$ and $\alpha \in \{0, 1\}$, the events

$$\left(u_{i_0, j, \Sigma}^{(\pi_0, \mu)} \upharpoonright \rho_1 \neq 0 \mid b_{i_0}^{(\pi_0)} = b \text{ and } \alpha_{i_0}^{(\pi_0)} = \alpha \right), \text{ where } (j, \mu, \Sigma) \in P,$$

and

$$\left(\bar{u}_{i_0, j, \Sigma}^{(\pi_0, \mu)} \upharpoonright \rho_1 \neq 0 \mid b_{i_0}^{(\pi_0)} = b \text{ and } \alpha_{i_0}^{(\pi_0)} = \alpha \right), \text{ where } (j, \mu, \Sigma) \in N,$$

are independent and each occurs with probability at most $3/4$. This is because when we fix the value of $b_{i_0}^{(\pi_0)}$ and $\alpha_{i_0}^{(\pi_0)}$, which determine the centre of the star (π_0, i_0) , the value of $u_{i_0, j, \Sigma}^{(\pi_0, \mu)}$ depends only on $\beta_j^{(\pi_0)}$ and $c_j^{(\pi_0)}$, which are independent for different j . Thus, all of them simultaneously occur with probability at most $(3/4)^r \leq (\sqrt{3/4})^{r+1}$. This implies that

$$\Pr[t \upharpoonright \rho_1 \neq 0] \leq \left(\sqrt{3/4} \right)^D \leq \left(\sqrt{3/4} \right)^{\sqrt{\frac{d}{4\Delta_{max}}}}.$$

The case when $\deg(t_{+z}) \geq d$ is proved analogously: each $z_{i, j, k}^{(\pi)}$ occurring in t_{+z} can be independently set to 0 by choosing the appropriate values for $b_i^{(\pi)}$ (assigning $\rho_1(x_{i, k}^{(\pi)}) = 0$) or for $c_j^{(\pi)}$ (assigning $\rho_1(y_{k, j}^{(\pi)}) = 0$). Note that we cannot apply the same argument to the negative occurrences of z -variables, since they depend on both x - and y -variables simultaneously. $\square$

The substitution ρ_1 transforms any refutation of $\psi_{P(G)}(\mathbf{A}, \kappa)$ into a refutation of $\phi_G(\mathbf{A}_1, \kappa)$, where $\mathbf{A}_1 = \mathbf{A} \upharpoonright \rho_1$.

Lemma 5.55. *Let $\kappa \geq 0$, $m, n > 0$ and let G be an $m \times n$ bipartite graph. Let $\mathbf{A} = \{A^{(\pi)} : \pi \in \Sigma^{\leq \kappa}\}$ be a collection of $m \times m$ matrices as in [Definition 5.45](#). Let $\mathbf{A}_1 = \mathbf{A} \upharpoonright \rho_1 = \{A^{(\pi)} \upharpoonright \rho_1 : A^{(\pi)} \in \mathbf{A}\}$. Then for every $\text{PCR}_{\mathbb{F}_2}$ refutation Π of $\psi_{P(G)}(\mathbf{A}, \kappa)$, $\Pi \upharpoonright \rho_1$ is a refutation of $\phi_G(\mathbf{A}_1, \kappa)$ (modulo the axioms for twin variables $\theta + \bar{\theta} + 1$).*

Proof. Follows from the definition of ρ_1 . Under ρ_1 , the axioms of $\psi_{P(G)}(\mathbf{A}, \kappa)$ are either satisfied or become axioms of the same type for $\phi_G(\mathbf{A}_1, \kappa)$. Since ρ_1 sometimes substitutes variables with negated variables, we need to apply the twin variable axioms $\theta + \bar{\theta} + 1$ to replace such occurrences of $\bar{\theta}$ with $\theta + 1$. This can lead to an insignificant increase in the refutation size, which we can safely ignore. $\square$

Next, we define a random substitution ρ_2 that w.h.p. turns any term t that mentions many columns into 0. As with ρ_1 , we only apply ρ_2 to graphs that have a specific structure, namely $H(G)$.

Definition 5.56. Let t be a term in the variables of $\phi_{H(G)}(\mathbf{A}, \kappa)$. Let $N = 8n + 8$. Given $k \in [N]$, let $\eta(k) = ((k - 1) \bmod (n + 1)) + 1$ and $\nu(k) = \lfloor \frac{k-1}{n+1} \rfloor + 1$.

We say that $k \in [n]$ is $X^{(\pi)}$ -column-mentioned in t if there are $i \in [m]$ and $\ell \in [N]$ with $\eta(\ell) = k$ such that $x_{i, \ell}^{(\pi)}$ or $\bar{x}_{i, \ell}^{(\pi)}$ appears in t . We say that $k \in [n]$ is $(Y^{(\pi)})^T$ -column-mentioned in t if there are $j \in [m]$ and $\ell \in [N]$ with $\eta(\ell) = k$ such that $y_{\ell, j}^{(\pi)}$ or $\bar{y}_{\ell, j}^{(\pi)}$ appears in t . We say that $k \in [n]$ is $Z^{(\pi)}$ -column-mentioned in t if there are $i, j \in [m]$ and $\ell \in [N]$ with $\eta(\ell) = k$ such that $z_{i, j, \ell}^{(\pi)}$ or $\bar{z}_{i, j, \ell}^{(\pi)}$ appears in t .

Definition 5.57 (Substitution ρ_2). Let $\kappa \geq 0$, $m, n > 0$, $N = 8n + 8$ and let G be an $m \times n$ bipartite graph with parts $[m]$ and $[n]$.

The random substitution ρ_2 for $\phi_{H(G)}(\mathbf{A}, \kappa)$ is defined as follows.

1. For each $k \in [n]$, independently and uniformly sample $c_k \in \{*, 1, 0\}$. Assign $c_{n+1} = *$.

2. Let $\tilde{K} = \{k \in [N] : c_{\eta(k)} \in \{0, 1\}\}$ and let $\tilde{G} = G \setminus \{k \in [n] : c_k \in \{0, 1\}\}$. Note that $\mathcal{N}_{H(\tilde{G})}(j) = \mathcal{N}_{H(G)}(j) \setminus \tilde{K}$ for all $j \in [m]$.

3. For every $\pi \in \{0, 1\}^{\leq \kappa}$:

- (a) Let $\{\tilde{u}_{i,j,\Sigma}^{(\pi)} : i, j \in [m], \Sigma \text{ is a non-empty initial segment of } \mathcal{N}_{H(\tilde{G})}(j)\}$ be fresh extension variables.
- (b) For every $i \in [m]$, and $k \in [N]$, set

$$\rho_2(x_{i,k}^{(\pi)}) = \begin{cases} c_{\eta(k)} & \text{if } c_{\eta(k)} \in \{0, 1\}, \\ x_{i,k}^{(\pi)} & \text{otherwise.} \end{cases}$$

- (c) For every $j \in [m]$, and $k \in \mathcal{N}_{H(G)}(j)$, set

$$\rho_2(y_{k,j}^{(\pi)}) = \begin{cases} c_{\eta(k)} & \text{if } c_{\eta(k)} \in \{0, 1\}, \\ y_{k,j}^{(\pi)} & \text{otherwise.} \end{cases}$$

- (d) For every $i, j \in [m]$, and $k \in \mathcal{N}_{H(G)}(j)$, set

$$\rho_2(z_{i,j,k}^{(\pi)}) = \begin{cases} c_{\eta(k)} & \text{if } c_{\eta(k)} \in \{0, 1\}, \\ z_{i,j,k}^{(\pi)} & \text{otherwise.} \end{cases}$$

- (e) For every $i, j \in [m]$ and every non-empty initial segment Σ of $\mathcal{N}_{H(G)}(j)$, set

$$\rho_2(u_{i,j,\Sigma}^{(\pi)}) = \begin{cases} 0 & \text{if } \Sigma \setminus \tilde{K} = \emptyset \text{ and } |\{k \in \Sigma : c_{\eta(k)} = 1\}| \text{ is even,} \\ 1 & \text{if } \Sigma \setminus \tilde{K} = \emptyset \text{ and } |\{k \in \Sigma : c_{\eta(k)} = 1\}| \text{ is odd,} \\ \tilde{u}_{i,j,(\Sigma \setminus \tilde{K})}^{(\pi)} & \text{if } \Sigma \setminus \tilde{K} \neq \emptyset \text{ and } |\{k \in \Sigma : c_{\eta(k)} = 1\}| \text{ is even,} \\ \bar{\tilde{u}}_{i,j,(\Sigma \setminus \tilde{K})}^{(\pi)} & \text{if } \Sigma \setminus \tilde{K} \neq \emptyset \text{ and } |\{k \in \Sigma : c_{\eta(k)} = 1\}| \text{ is odd.} \end{cases}$$

- (f) For any variable θ mapped to 0 or 1 by ρ_2 , set $\rho_2(\bar{\theta}) = 1 - \rho_2(\theta)$.
- (g) For any variable θ not mapped to a Boolean constant by ρ_2 , set $\rho_2(\bar{\theta}) = \overline{\rho_2(\theta)}$.

We show that under ρ_2 every term that mentions many columns becomes zero w.h.p.

Lemma 5.58. *Let $\kappa \geq 0$, $m, n > 0$ and let G be an $m \times n$ bipartite graph with parts $[m]$ and $[n]$. Let t be a term in the variables of $\phi_{H(G)}(\mathbf{A}, \kappa)$. Let K be the set of all $k \in [n]$ that are $X^{(\pi)}$ -column-mentioned, $(Y^{(\pi)})^T$ -column-mentioned or $Z^{(\pi)}$ -column-mentioned in t , for some $\pi \in \{0, 1\}^{\leq \kappa}$. Then $\Pr[t \upharpoonright \rho_2 \neq 0] \leq (2/3)^{|K|}$.*

Proof. For every x -, y - or z -variable θ appearing in t , $\Pr[\rho_2(\theta) \neq 0] \leq 2/3$, which occurs independently for distinct columns. Thus, $(t \upharpoonright \rho_2 \neq 0)$ occurs with probability at most $(2/3)^{|K|}$. $\square$

When G is an expander with sufficiently large minimum left-degree, the graph $\tilde{G}$ defined in Definition 5.57 remains a good expander due to Lemma 4.18.

Proposition 5.59. *Let $\kappa \geq 0$, $n > 0$, $N = 8n + 8$, $m = 16(N + 2)$ and let G be an $m \times n$ (n^δ, d) -lossless expander with parts $[m]$ and $[n]$, the minimum left-degree at least $\Delta_{\min} = 50 \log n$ and the maximum left-degree at most $\Delta_{\max} = 150 \log n$. Then with probability at least $1/2$, the graph $\tilde{G}$ from Definition 5.57 is an (n^δ, c) -boundary expander, where $c = \Delta_{\min}/6 - d$.*

We show that ρ_2 transforms any refutation of $\phi_{H(G)}(\mathbf{A}, \kappa)$ into a refutation of $\phi_{H(\tilde{G})}(\mathbf{A}_2, \kappa)$, where $\mathbf{A}_2 = \mathbf{A} \upharpoonright \rho_2$.

Lemma 5.60. *Let $\kappa \geq 0$, $m, n > 0$ and let G be an $m \times n$ bipartite graph with parts $[m]$ and $[n]$. Let $\mathbf{A} = \{A^{(\pi)} : \pi \in \Sigma^{\leq \kappa}\}$ be a collection of $m \times m$ matrices as in [Definition 5.46](#). Let $\mathbf{A}_2 = \mathbf{A} \upharpoonright \rho_2 = \{A^{(\pi)} \upharpoonright \rho_2 : A^{(\pi)} \in \mathbf{A}\}$. Then for every $\text{PCR}_{\mathbb{F}_2}$ refutation Π of $\phi_{H(G)}(\mathbf{A}, \kappa)$, $\Pi \upharpoonright \rho_2$ is a refutation of $\phi_{H(\tilde{G})}(\mathbf{A}_2, \kappa)$.*

Proof. Can be verified straightforwardly from the definition of ρ_2 , similarly to [Lemma 5.40](#). $\square$

Now we define a random restriction ρ_3 that reduces the total *degree* of the terms in any refutation of $\phi_{H(G)}(\mathbf{A}, \kappa)$ provided such a refutation has certain restrictions on its u -degree and column degree. We need the following matrices, which we used in [Section 5.4.1](#).

$$\mathcal{D} = \begin{pmatrix} 1 & * & 1 & 0 & 0 & * & * & * \\ 0 & 1 & * & 1 & * & * & * & 0 \\ * & * & 0 & * & 1 & 0 & 1 & * \\ * & 0 & * & * & * & 1 & 0 & 1 \end{pmatrix}, \quad \mathcal{E} = \begin{pmatrix} 1 & 0 & 0 & 1 & * & * & * & * \\ 0 & * & 1 & * & 1 & * & 0 & * \\ * & 1 & * & 0 & * & 0 & * & 1 \\ * & * & * & * & 0 & 1 & 1 & 0 \end{pmatrix}.$$

Definition 5.61 (Substitution ρ_3). Let $\kappa \geq 0$, $m, n > 0$, $N = 8n + 8$, and let G be an $m \times n$ bipartite graph with parts $[m]$ and $[n]$. Given $k \in [N]$, let $\eta(k) = ((k-1) \bmod (n+1)) + 1$ and $\nu(k) = \lfloor \frac{k-1}{n+1} \rfloor + 1$.

The substitution ρ_3 for $\phi_{H(G)}(\mathbf{A}, \kappa)$ is defined as follows. For every $\pi \in \{0, 1\}^{\leq \kappa}$:

1. For every $i \in [m]$ uniformly and independently sample $s_i^{(\pi)} \in [4]$ and $r_i^{(\pi)} \in [4]$.
2. For each $i \in [m]$ and $k \in [N]$, set

$$\rho_3(x_{i,k}^{(\pi)}) = \begin{cases} 1 & \text{if } \eta(k) = n+1 \text{ and } \mathcal{D}_{s_i^{(\pi)}, \nu(k)} = 1, \\ 0 & \text{if } \eta(k) = n+1 \text{ and } \mathcal{D}_{s_i^{(\pi)}, \nu(k)} \neq 1, \\ x_{i, \eta(k)}^{(\pi)} & \text{if } \eta(k) \neq n+1 \text{ and } \mathcal{D}_{s_i^{(\pi)}, \nu(k)} = *, \\ \mathcal{D}_{s_i^{(\pi)}, \nu(k)} & \text{if } \eta(k) \neq n+1 \text{ and } \mathcal{D}_{s_i^{(\pi)}, \nu(k)} \in \{0, 1\}. \end{cases}$$

3. For each $j \in [m]$ and $k \in \mathcal{N}_{H(G)}(j)$, set

$$\rho_3(y_{k,j}^{(\pi)}) = \begin{cases} |\mathcal{N}_G(j)| \bmod 2 & \text{if } \eta(k) = n+1 \text{ and } \mathcal{E}_{r_j^{(\pi)}, \nu(k)} = 1, \\ 0 & \text{if } \eta(k) = n+1 \text{ and } \mathcal{E}_{r_j^{(\pi)}, \nu(k)} \neq 1, \\ y_{\eta(k), j}^{(\pi)} & \text{if } \eta(k) \neq n+1 \text{ and } \mathcal{E}_{r_j^{(\pi)}, \nu(k)} = *, \\ \mathcal{E}_{r_j^{(\pi)}, \nu(k)} & \text{if } \eta(k) \neq n+1 \text{ and } \mathcal{E}_{r_j^{(\pi)}, \nu(k)} \in \{0, 1\}. \end{cases}$$

4. For each $i, j \in [m]$ and $k \in \mathcal{N}_{H(G)}(j)$ set

$$\rho_3(z_{i,j,k}^{(\pi)}) = \rho_3(x_{i,k}^{(\pi)}) \rho_3(y_{k,j}^{(\pi)}).$$

5. For each $i, j \in [m]$ and every non-empty initial segment Σ of $\mathcal{N}_{H(G)}(j)$, set

$$\rho_3(u_{i,j,\Sigma}^{(\pi)}) = \sum_{k \in \Sigma} \rho_3(z_{i,j,k}^{(\pi)}).$$

6. For every variable θ assigned by ρ_3 , set $\rho_3(\bar{\theta}) = 1 - \rho_3(\theta)$.

Note that ρ_3 does not substitute any negated variables and expands the extension variables as polynomials. It can be trivially extended to a PC refutation. The *size* of the resulting refutation can grow exponentially, but we will prove that its *degree* becomes small w.h.p. For every refutation Π of $\phi_{H(G)}(\mathbf{A}, \kappa)$, $\Pi \upharpoonright \rho_3$ can be transformed into a refutation of $\tilde{\phi}_G(\mathbf{A}_3, \kappa)$, where $\mathbf{A}_3 = \mathbf{A} \upharpoonright \rho_3$.

In order to handle the negative occurrences of the z -variables, we need to define a transformation R_d . It is a linear mapping defined on the terms in the variables of $\phi_{H(G)}(\mathbf{A}, \kappa)$. In the definitions below, when working with $\phi_{H(G)}$, we ignore all variables whose column k satisfies $\eta(k) = n + 1$. These variables are mapped to constants by ρ_3 , and hence do not contribute to the degree after the restriction. Accordingly, all z -stars considered by R_d are formed only from variables $\bar{z}_{i,j,k}^{(\pi)}$ with $\eta(k) \neq n + 1$.

Definition 5.62. Let t be a term in the variables of $\phi_{H(G)}(\mathbf{A}, \kappa)$. We define $R_d(t)$ by *contracting* every subterm of the form $\prod_{j \in J} \bar{z}_{i_0,j,k}^{(\pi)}$ or $\prod_{i \in I} \bar{z}_{i,j_0,k}^{(\pi)}$, where $|I|, |J| \geq d$, using the following procedure. Let $t = t_1 t_2$, where t_1 contains only negated z -variables that do not mention columns k with $\eta(k) = n + 1$ and t_2 contains the rest. Let

$$\begin{aligned} S_{left} &= \{(\pi, i_0, \{j \in [m] : \bar{z}_{i_0,j,k}^{(\pi)} \in \text{vars}(t_1)\}, k) : \pi \in \{0, 1\}^{\leq \kappa}, i_0 \in [m], \\ &\quad k \in [N], \eta(k) \neq n + 1\}, \\ S_{right} &= \{(\pi, \{i \in [m] : \bar{z}_{i,j_0,k}^{(\pi)} \in \text{vars}(t_1)\}, j_0, k) : \pi \in \{0, 1\}^{\leq \kappa}, j_0 \in [m], \\ &\quad k \in [N], \eta(k) \neq n + 1\}. \end{aligned}$$

Given a star $(\pi, i_0, J, k) \in S_{left}$ we can think of the set J as the set of *rays* of the star. Similarly, given a star $(\pi, I, j_0, k) \in S_{right}$ we can think of the set I as the set of rays of the star. Let C be the set of quadruples (π, i, j, k) with $\bar{z}_{i,j,k}^{(\pi)} \in \text{vars}(t_1)$ that do not belong to any star in $S_{left} \cup S_{right}$ with at least d rays. We then contract all stars with at least d rays to a single variable (its centre). Formally, we define

$$R_d(t) = \left(\prod_{\substack{(\pi, i_0, J, k) \in S_{left} \\ |J| \geq d}} \bar{x}_{i_0,k}^{(\pi)} \right) \cdot \left(\prod_{\substack{(\pi, I, j_0, k) \in S_{right} \\ |I| \geq d}} \bar{y}_{k,j_0}^{(\pi)} \right) \cdot \prod_{(\pi, i, j, k) \in C} \bar{z}_{i,j,k}^{(\pi)} \cdot t_2. \quad^{10}$$

We proceed by extending R_d to polynomials linearly.

When $d \geq 2$,¹¹ the mapping R_d does not increase the degree of any term. Indeed, suppose that R_d inserted q centres and removed r rays (negated z -variables appearing in large stars). The sum of sizes of all contracted stars is at least dq , while every removed ray belongs to at most one contracted left-star and one contracted right-star. Therefore, we remove at least $dq/2$ rays, and hence $r \geq dq/2 \geq q$, so the degree does not increase.

The intuition behind R_d is that every star with at least d rays is contracted to a single variable, which corresponds to its centre. The following lemma shows that ρ_3 reduces the degree of any term t with many x - and y -variables or many negative occurrences of z -variables after applying R_d .

Lemma 5.63. Let $\kappa \geq 0$, $m, n > 0$, $d, \tilde{d} > 0$ and let G be an $m \times n$ bipartite graph with parts $[m]$ and $[n]$. Let t be a term in the variables of $\phi_{H(G)}(\mathbf{A}, \kappa)$. Let K be the set of all $k \in [n]$ that are $X^{(\pi)}$ -column-mentioned, $(Y^{(\pi)})^T$ -column-mentioned or $Z^{(\pi)}$ -column-mentioned in t , for some $\pi \in \{0, 1\}^{\leq \kappa}$. Let $t_{x,y}$ be the subterm of $R_{\tilde{d}}(t)$ consisting of all x - and y -variables appearing in $R_{\tilde{d}}(t)$, excluding those that mention columns k with $\eta(k) = n + 1$. Similarly, let t_{-z} be the subterm of $R_{\tilde{d}}(t)$ consisting of all

¹⁰As usual, products are understood modulo the Boolean axioms, so the resulting terms are multilinear.

¹¹In our application, we choose d to be even larger.

negative occurrences of z -variables in $R_{\tilde{d}}(t)$, excluding those that mention columns k with $\eta(k) = n + 1$. Assume that $\deg(t_{x,y}) \geq d$ or $\deg(t_{-z}) \geq d$. Then

$$\Pr[R_{\tilde{d}}(t) \upharpoonright \rho_3 \neq 0] \leq (15/16)^{\frac{d}{16\tilde{d}|K|}}.$$

Proof. Suppose that $\deg(t_{x,y}) \geq d$. For every $\pi \in \{0,1\}^{\leq \kappa}$, let

$$\begin{aligned} I_x^{(\pi)} &= \{i \in [m] : \exists k \in [N] \text{ with } \eta(k) \neq n+1 \text{ s.t. } x_{i,k}^{(\pi)} \text{ or } \bar{x}_{i,k}^{(\pi)} \text{ appears in } t_{x,y}\}, \\ J_y^{(\pi)} &= \{j \in [m] : \exists k \in [N] \text{ with } \eta(k) \neq n+1 \text{ s.t. } y_{k,j}^{(\pi)} \text{ or } \bar{y}_{k,j}^{(\pi)} \text{ appears in } t_{x,y}\}. \end{aligned}$$

Then

$$d \leq \deg(t_{x,y}) \leq 8|K| \sum_{\pi \in \{0,1\}^{\leq \kappa}} (|I_x^{(\pi)}| + |J_y^{(\pi)}|).$$

For every $i \in I_x^{(\pi)}$ there is a variable in $t_{x,y}$ that is independently set to 0 with probability $1/4$ by ρ_3 . Similarly, for every $j \in J_y^{(\pi)}$, there is a variable in $t_{x,y}$ that is also set to 0 by ρ_3 with the same probability. This implies that

$$\Pr[t_{x,y} \upharpoonright \rho_3 \neq 0] \leq (3/4)^{\frac{d}{8|K|}}.$$

Now suppose instead that $\deg(t_{-z}) \geq d$. We greedily identify a large set of variables in t_{-z} that are independently set to 0 by ρ_3 w.h.p. Initially set $S = \text{vars}(t_{-z})$ and $T = \emptyset$.

- Pick any variable $\bar{z}_{i,j,k}^{(\pi)} \in S$ and add it to T .
- Remove all variables $\{\bar{z}_{i,j',k'}^{(\pi)} : j' \in [m], \eta(k') \in K\} \cup \{\bar{z}_{i',j,k'}^{(\pi)} : i' \in [m], \eta(k') \in K\}$ from S .
- Repeat until S is empty.

Observe that each variable in T mentions its own indices: for every $\pi \in \{0,1\}^{\leq \kappa}$ and $\ell \in [m]$, there is at most one variable of the form $\bar{z}_{i,\ell,k}^{(\pi)}$ or $\bar{z}_{\ell,j,k}^{(\pi)}$. Thus, every variable from T independently becomes zero under ρ_3 with probability at least $1/16$. It remains to show that the size of T is sufficiently large. At each step we remove at most $16\tilde{d}|K|$ z -variables from S due to the properties of $R_{\tilde{d}}$, therefore the size of T is at least $\frac{d}{16\tilde{d}|K|}$. This implies that

$$\Pr[t_{-z} \upharpoonright \rho_3 \neq 0] \leq (15/16)^{\frac{d}{16\tilde{d}|K|}}.$$

□

We proceed by showing that for every $\text{PCR}_{\mathbb{F}_2}$ refutation Π of $\phi_{H(G)}(\mathbf{A}, \kappa)$ of degree d , $R_{\tilde{d}}(\Pi) \upharpoonright \rho_3$ can be extended to a refutation of $\tilde{\phi}_G(\mathbf{A} \upharpoonright \rho_3, \kappa)$.

Lemma 5.64. *Let $\kappa \geq 0$, $m, n > 0$, $N = 8n + 8$, $\tilde{d} \geq 2$ and let G be an $m \times n$ bipartite graph with parts $[m]$ and $[n]$. Let Π be a $\text{PCR}_{\mathbb{F}_2}$ refutation of $\phi_{H(G)}(\mathbf{A}, \kappa)$ of size at most $\frac{(4/3)^{\tilde{d}}}{8mN^{2\kappa}}$.*

With probability at least $1/2$ over the choice of ρ_3 , the following holds: if $R_{\tilde{d}}(\Pi) \upharpoonright \rho_3$ has degree at most d for some $d \geq 0$, then it can be extended to a valid PC refutation $\tilde{\Pi}$ of $\tilde{\phi}_G(\mathbf{A} \upharpoonright \rho_3, \kappa)$ with degree at most $d + O(1)$.

Proof. We restrict our attention only to the terms that do not mention columns k with $\eta(k) = n + 1$, as these are mapped to constants by ρ_3 .

We begin by demonstrating the following property of ρ_3 .

Claim 5.65. Let $\pi \in \{0, 1\}^{\leq \kappa}$ and $k \in [N]$ with $\eta(k) \neq n + 1$. For every $i \in [m]$ and $J \subseteq \{\ell \in [m] : k \in \mathcal{N}_{H(G)}(\ell)\}$,

$$\Pr[\forall \ell \in J : \rho_3(y_{k,\ell}^{(\pi)}) \neq 1] \leq (3/4)^{|J|}.$$

Observe that if $\rho_3(y_{k,\ell}^{(\pi)}) = 1$, then $\rho_3(\bar{z}_{i,\ell,k}^{(\pi)}) = \rho_3(\bar{x}_{i,k}^{(\pi)})$ by the definition of ρ_3 . Similarly, for every $j \in [m]$ with $k \in \mathcal{N}_{H(G)}(j)$ and $I \subseteq [m]$,

$$\Pr[\forall \ell \in I : \rho_3(x_{\ell,k}^{(\pi)}) \neq 1] \leq (3/4)^{|I|}.$$

As before, if $\rho_3(x_{\ell,k}^{(\pi)}) = 1$, then $\rho_3(\bar{z}_{\ell,j,k}^{(\pi)}) = \rho_3(\bar{y}_{k,j}^{(\pi)})$.

Proof. We will show the first statement, and the second is symmetric. The events $\rho_3(y_{k,\ell}^{(\pi)}) = 1$ for $\ell \in J$ are independent, since they depend on the independent random choices of $r_\ell^{(\pi)}$ for $\ell \in J$. Each event occurs with probability at least $1/4$, since $\mathcal{E}$ contains at least one 1 in every column. Now fix $s_i^{(\pi)}$ and thus $\rho_3(x_{i,k}^{(\pi)})$. Then $\rho_3(\bar{z}_{i,\ell,k}^{(\pi)}) = \rho_3(\bar{x}_{i,k}^{(\pi)})$ when $\rho_3(y_{k,\ell}^{(\pi)}) = 1$. $\diamond$

There are at most $4mN2^\kappa|\Pi|$ distinct maximal stars in Π : for every term t choose a triple (π, ℓ, k) and consider two induced stars $\{\bar{z}_{i,\ell,k}^{(\pi)} : i \in [m]\} \cap \text{vars}(t)$ and $\{\bar{z}_{\ell,j,k}^{(\pi)} : j \in [m]\} \cap \text{vars}(t)$. Together with the union bound, the above claim shows that with probability at least $1/2$, the following holds for every term t appearing in Π :

- For every $\pi \in \{0, 1\}^{\leq \kappa}$, $k \in [N]$ with $\eta(k) \neq n + 1$, $i \in [m]$, if the star $\{\bar{z}_{i,j,k}^{(\pi)} : j \in [m]\} \cap \text{vars}(t)$ has at least $\tilde{d}$ rays, then there exists $\tilde{j} \in [m]$ such that $\bar{z}_{i,\tilde{j},k}^{(\pi)}$ appears in t and

$$\rho_3(y_{k,\tilde{j}}^{(\pi)}) = 1.$$

- For every $\pi \in \{0, 1\}^{\leq \kappa}$, $k \in [N]$ with $\eta(k) \neq n + 1$, $j \in [m]$, if the star $\{\bar{z}_{i,j,k}^{(\pi)} : i \in [m]\} \cap \text{vars}(t)$ has at least $\tilde{d}$ rays, then there exists $\tilde{i} \in [m]$ such that $\bar{z}_{\tilde{i},j,k}^{(\pi)}$ appears in t and

$$\rho_3(x_{\tilde{i},k}^{(\pi)}) = 1.$$

For the rest of the proof, fix such ρ_3 .

In what follows we omit the subscript $\tilde{d}$ and simply write R to denote $R_{\tilde{d}}$.

Claim 5.66. For every term t appearing in Π , we have

$$R(t) \upharpoonright \rho_3 = t \upharpoonright \rho_3.$$

Proof. We use the fact the conditions from Claim 5.65 hold for every term t in Π due to the choice of ρ_3 . Let t be a term in Π . R leaves all variables that are not negated z -variables unchanged, so we only need to consider the negative occurrences of z -variables in t . Moreover, for a fixed $\pi \in \{0, 1\}^{\leq \kappa}$ and $k \in [N]$ with $\eta(k) \neq n + 1$, R acts independently on the variable sets $\{\bar{z}_{i,j,k}^{(\pi)} : i, j \in [m]\}$.

Fix $\pi \in \{0, 1\}^{\leq \kappa}$ and $k \in [N]$ with $\eta(k) \neq n + 1$. Let t' be the subterm of t consisting of all negative occurrences of z -variables of the form $\bar{z}_{i,j,k}^{(\pi)}$ for some $i, j \in [m]$. It suffices to show that $R(t') \upharpoonright \rho_3 = t' \upharpoonright \rho_3$. Let $E = \{(i, j) : \bar{z}_{i,j,k}^{(\pi)} \text{ appears in } t'\}$.

Let C_{left} be the set of all centres $i \in [m]$ such that the left-star $(i, \{j \in [m] : (i, j) \in E\})$ has at least $\tilde{d}$ rays, and let C_{right} be the set of all centres $j \in [m]$ such that the right-star $(\{i \in [m] : (i, j) \in E\}, j)$

has at least $\tilde{d}$ rays. By the choice of ρ_3 , for every $i \in C_{\text{left}}$ there exists $j_i \in [m]$ such that $\bar{z}_{i,j_i,k}^{(\pi)}$ appears in t' and $\rho_3(y_{k,j_i}^{(\pi)}) = 1$. Similarly, for every $j \in C_{\text{right}}$ there exists $i_j \in [m]$ such that $\bar{z}_{i_j,j,k}^{(\pi)}$ appears in t' and $\rho_3(x_{i_j,k}^{(\pi)}) = 1$.

Hence we have

$$t' = \prod_{(i,j) \in E} \bar{z}_{i,j,k}^{(\pi)}$$

and

$$R(t') = \prod_{i \in C_{\text{left}}} \bar{x}_{i,k}^{(\pi)} \cdot \prod_{j \in C_{\text{right}}} \bar{y}_{k,j}^{(\pi)} \cdot \prod_{\substack{(i,j) \in E \\ i \notin C_{\text{left}}, j \notin C_{\text{right}}}} \bar{z}_{i,j,k}^{(\pi)}. \quad (33)$$

If for some $i \in C_{\text{left}}$ we have $\rho_3(x_{i,k}^{(\pi)}) = 1$, then $R(t') \upharpoonright \rho_3 = 0$. Our choice of j_i ensures that $\rho_3(\bar{z}_{i,j_i,k}^{(\pi)}) = 1 - \rho_3(x_{i,k}^{(\pi)}) \cdot \rho_3(y_{k,j_i}^{(\pi)}) = 0$, and hence $t' \upharpoonright \rho_3 = 0$ as well. The same argument applies to every $j \in C_{\text{right}}$ with $\rho_3(y_{k,j}^{(\pi)}) = 1$. The claim is proved if either of these cases occurs. Otherwise, we have $j_i \notin C_{\text{right}}$ for every $i \in C_{\text{left}}$ and $i_j \notin C_{\text{left}}$ for every $j \in C_{\text{right}}$. We can equivalently rewrite t' as

$$t' = \prod_{i \in C_{\text{left}}} \prod_{j: (i,j) \in E} \bar{z}_{i,j,k}^{(\pi)} \cdot \prod_{j \in C_{\text{right}}} \prod_{\substack{i: (i,j) \in E \\ i \notin C_{\text{left}}}} \bar{z}_{i,j,k}^{(\pi)} \cdot \prod_{\substack{(i,j) \in E \\ i \notin C_{\text{left}}, j \notin C_{\text{right}}}} \bar{z}_{i,j,k}^{(\pi)}. \quad (34)$$

Fix $i \in C_{\text{left}}$. By definition of ρ_3 , we have $\rho_3(\bar{z}_{i,j_i,k}^{(\pi)}) = \rho_3(\bar{x}_{i,k}^{(\pi)})$. Moreover, for every $j \in [m]$, we have

$$\begin{aligned} \rho_3(\bar{z}_{i,j,k}^{(\pi)}) \cdot \rho_3(\bar{x}_{i,k}^{(\pi)}) &= (1 - \rho_3(x_{i,k}^{(\pi)}) \cdot \rho_3(y_{k,j}^{(\pi)})) \cdot (1 - \rho_3(x_{i,k}^{(\pi)})) \\ &= 1 - \rho_3(x_{i,k}^{(\pi)}) \cdot \rho_3(y_{k,j}^{(\pi)}) - \rho_3(x_{i,k}^{(\pi)}) + \rho_3^2(x_{i,k}^{(\pi)}) \cdot \rho_3(y_{k,j}^{(\pi)}) \\ &= 1 - \rho_3(x_{i,k}^{(\pi)}) = \rho_3(\bar{x}_{i,k}^{(\pi)}) \end{aligned}$$

modulo the Boolean axioms. Thus,

$$\prod_{j: (i,j) \in E} \rho_3(\bar{z}_{i,j,k}^{(\pi)}) = \rho_3(\bar{z}_{i,j_i,k}^{(\pi)}) \cdot \prod_{\substack{j: (i,j) \in E \\ j \neq j_i}} \rho_3(\bar{z}_{i,j,k}^{(\pi)}) = \rho_3(\bar{x}_{i,k}^{(\pi)}) \cdot \prod_{\substack{j: (i,j) \in E \\ j \neq j_i}} \rho_3(\bar{z}_{i,j,k}^{(\pi)}) = \rho_3(\bar{x}_{i,k}^{(\pi)}).$$

Now fix $j \in C_{\text{right}}$. By definition of ρ_3 , we have $\rho_3(\bar{z}_{i_j,j,k}^{(\pi)}) = \rho_3(\bar{y}_{k,j}^{(\pi)})$. Moreover, for every $i \in [m]$, we have $\rho_3(\bar{z}_{i,j,k}^{(\pi)}) \cdot \rho_3(\bar{y}_{k,j}^{(\pi)}) = \rho_3(\bar{y}_{k,j}^{(\pi)})$. Thus,

$$\prod_{\substack{i: (i,j) \in E \\ i \notin C_{\text{left}}}} \rho_3(\bar{z}_{i,j,k}^{(\pi)}) = \rho_3(\bar{y}_{k,j}^{(\pi)}).$$

It implies that equations (33) and (34) are the same under ρ_3 modulo the Boolean axioms, and the claim follows. $\diamond$

Now we prove the lemma by induction on the proof lines in Π . Let f be a proof line in Π .

- If f is an axiom, then $R(f) = f$. After applying ρ_3 , this axiom is either an identity, or a simple consequence of the axioms of $\phi_G(\mathbf{A} \upharpoonright \rho_3, \kappa)$, by the definition of ρ_3 .
- If f is obtained from polynomials g and h as $f = \alpha g + \beta h$, then clearly $R(f) = \alpha R(g) + \beta R(h)$ by linearity of R .

- The remaining case is multiplication by a variable θ : $f = g \cdot \theta$. Applying [Claim 5.66](#) term-wise to f and g , we obtain

$$R(g) \upharpoonright \rho_3 = g \upharpoonright \rho_3$$

and

$$R(f) \upharpoonright \rho_3 = f \upharpoonright \rho_3 = (g \upharpoonright \rho_3) \cdot (\theta \upharpoonright \rho_3) = (R(g) \upharpoonright \rho_3) \cdot (\theta \upharpoonright \rho_3).$$

Hence $R(f) \upharpoonright \rho_3$ can be derived from $R(g) \upharpoonright \rho_3$ by multiplying by $\theta \upharpoonright \rho_3$, which is polynomial of degree at most 2. This can be done in PC by multiplying by the monomials of $\theta \upharpoonright \rho_3$ and taking linear combinations and can result in an increase of degree by at most a constant.

The absorption identities used in the multiplication step above are consequences of the Boolean axioms and can be derived in PC in constant degree. Hence the resulting derivation has degree at most $d + O(1)$. $\square$

Finally, we show how $\tilde{\phi}_G(\mathbf{A}, \kappa)$ can be reduced to FPHP(G).

Definition 5.67 (Reduction τ). Let $\kappa \geq 0$, $m, n > 0$ and let G be an $m \times n$ bipartite graph with parts $[m]$ and $[n]$. Let $P = (p_{i,k})_{i \in [m], k \in \mathcal{N}_G(i)}$ be fresh variables.¹² The reduction τ for $\tilde{\phi}_G(\mathbf{A}, \kappa)$ is defined as follows. For every $\pi \in \{0, 1\}^{\leq \kappa}$ in decreasing order:

1. For each $i \in [m]$ and $k \in [n]$ set

$$\tau(x_{i,k}^{(\pi)}) = \sum_{\substack{s \in [m] \\ \text{such that } k \in \mathcal{N}_G(s)}} \tau(A_{i,s}^{(\pi)}) p_{s,k}.$$

2. For each $j \in [m]$ and $k \in \mathcal{N}_G(j)$ set

$$\tau(y_{k,j}^{(\pi)}) = p_{j,k}.$$

This definition is well-defined: $A^{(\pi)} \upharpoonright \tau = A^{(\pi)}$ for every $\pi \in \{0, 1\}^\kappa$, because $A^{(\pi)}$ contains only constants in these cases.

Lemma 5.68. Let $\kappa \geq 0$, $m, n > 0$ and let G be an $m \times n$ bipartite graph with parts $[m]$ and $[n]$ and the maximum left-degree $\Delta_{\max}$. Let $\mathbf{A} = \{A^{(\pi)} : \pi \in \Sigma^{\leq \kappa}\}$ be a collection of $m \times m$ matrices as in [Definition 5.47](#). Then for any PCR $_{\mathbb{F}_2}$ refutation Π of $\tilde{\phi}_G(\mathbf{A}, \kappa)$ with degree d , $\Pi \upharpoonright \tau$ can be transformed into a refutation of FPHP(G) with degree at most $\max(\Delta_{\max}, d(\kappa + 1))$.

Proof. It follows by induction on $|\pi|$ that for every $\pi \in \{0, 1\}^{\leq \kappa}$, $i \in [m]$, and $k \in [n]$, $\tau(x_{i,k}^{(\pi)})$ (and thus $\tau(\overline{x_{i,k}^{(\pi)}})$) is a polynomial of degree at most $\kappa - |\pi| + 1$.

This implies that all polynomials from $\tilde{\phi}_G(\mathbf{A}, \kappa) \upharpoonright \tau$ can be easily derived from the axioms of FPHP(G).

Claim 5.69. Every polynomial from $\tilde{\phi}_G(\mathbf{A}, \kappa) \upharpoonright \tau$ has a degree- $(\kappa + 2)$ proof from FPHP(G).

Proof. As shown above, all the entries of $A^{(\pi)} \upharpoonright \tau$ are polynomials of degree at most $\kappa - |\pi|$.

Let $\pi \in \Sigma^{\leq \kappa}$ and $i, j \in [m]$. The (π, i, j) th axiom of $\tilde{\phi}_G(\mathbf{A}, \kappa)$ under τ transforms as follows:

$$\left(\sum_{k \in \mathcal{N}_G(j)} x_{i,k}^{(\pi)} y_{k,j}^{(\pi)} + A_{i,j}^{(\pi)} \right) \upharpoonright \tau = \sum_{k \in \mathcal{N}_G(j)} \left(\sum_{\substack{s \in [m] \\ \text{such that } k \in \mathcal{N}_G(s)}} \tau(A_{i,s}^{(\pi)}) p_{s,k} \right) p_{j,k} + \tau(A_{i,j}^{(\pi)})$$

¹²It is useful to think of P as an $m \times n$ variable matrix, where $P_{i,k} = 0$ for all $i \in [m]$ and $k \notin \mathcal{N}_G(i)$. Then $\tau(x_{i,k}^{(\pi)}) = ((A^{(\pi)} \upharpoonright \tau)P)_{i,k}$.

$$\begin{aligned}
&= \sum_{s \in [m]} \tau(A_{i,s}^{(\pi)}) \left(\sum_{k \in \mathcal{N}_G(j) \cap N_G(s)} p_{s,k} p_{j,k} \right) + \tau(A_{i,j}^{(\pi)}) \\
&= \sum_{s \in [m] \setminus \{j\}} \tau(A_{i,s}^{(\pi)}) \left(\sum_{k \in \mathcal{N}_G(j) \cap N_G(s)} p_{s,k} p_{j,k} \right) \\
&\quad + \tau(A_{i,j}^{(\pi)}) \left(\sum_{k \in \mathcal{N}_G(j)} p_{j,k} + 1 \right),
\end{aligned}$$

which is trivially derivable in degree $\kappa - |\pi| + 2$ from the hole axioms $p_{s,k} p_{j,k}$ and the polynomial

$$\sum_{k \in \mathcal{N}_G(j)} p_{j,k} + 1.$$

Note that the latter is an immediate consequence of the pigeon axiom for pigeon j together with its functional axioms, and thus has a derivation of degree at most Δ_{max} from $\text{FPHP}(G)$.

Thus, we can extend $\Pi \upharpoonright \tau$ to a refutation of $\text{FPHP}(G)$ of degree at most $\max(\Delta_{max}, d(\kappa + 1))$. $\diamond$

□

Finally, we are ready to prove [Theorem 5.51](#).

Proof of Theorem 5.51. Let

$$\begin{aligned}
c &= \Delta_{min}/6 - d, \\
d_1 &= \frac{cn^\delta}{32\Delta_{max}(\kappa + 1)}, \\
d_2 &= d_1^{1/3} = \left(\frac{cn^\delta}{32\Delta_{max}(\kappa + 1)} \right)^{1/3}.
\end{aligned}$$

Note that $d_1 = \Theta(n^\delta/\kappa)$ and $d_2 = \Theta((n^\delta/\kappa)^{1/3})$ because of our choice of Δ_{min} and Δ_{max} .

Let Π be a $\text{PCR}_{\mathbb{F}_2}$ refutation of $\psi_{P(H(G))}(\mathbf{A}, \kappa)$. Suppose, for the sake of contradiction, that $|\Pi| \leq 2^{\epsilon d_2}$ for some small $\epsilon > 0$. We will fix the value of ϵ later.

We start by applying the restriction ρ_1 from [Definition 5.53](#) to Π . Let $\Pi_1 = \Pi \upharpoonright \rho_1$. By [Lemma 5.54](#) and the union bound, there exists ρ_1 such that every term t in Π_1 contains at most d_1 u -variables and at most d_1 positively occurring z -variables. This holds since the maximum left-degree of $H(G)$ is

$8\Delta_{max} + 8$ and the size of Π is less than $(4/3) \sqrt{\frac{d_1}{4(8\Delta_{max}+8)}} = 2^{\Theta((\frac{d_1}{\log n})^{1/2})}$ for a sufficiently small ϵ . By [Lemma 5.55](#), Π_1 is a refutation of $\phi_{H(G)}(\mathbf{A} \upharpoonright \rho_1, \kappa)$. Since ρ_1 substitutes only constants and literals, the size and degree of Π_1 are at most the size and degree of Π , respectively.

We proceed by applying the restriction ρ_2 from [Definition 5.57](#) to Π_1 . Let $\Pi_2 = \Pi_1 \upharpoonright \rho_2$. By [Lemma 5.58](#) and the union bound, with probability greater than $1/2$, for every term t in Π_2 it holds that t mentions at most d_2 columns using x -, y - or z -variables; that is, there is a set $K \subseteq [n]$ of size at most d_2 such that no $k \notin K$ is $X^{(\pi)}$ -column-mentioned, $(Y^{(\pi)})^T$ -column-mentioned or $Z^{(\pi)}$ -column-mentioned in t . This holds since the size of Π_1 is less than $\frac{1}{2}(3/2)^{d_2} = 2^{\Theta(d_2)}$ for a sufficiently small ϵ . By [Proposition 5.59](#), the graph G' is an (n^δ, c) -boundary expander with probability at least $1/2$. Thus, by the union bound, there is ρ_2 for which both of these events occur. By [Lemma 5.60](#), Π_2 is a refutation of $\phi_{H(G')}(\mathbf{A} \upharpoonright \rho_1 \upharpoonright \rho_2, \kappa)$. Since ρ_2 substitutes only constants and literals, the size and degree of Π_2 are at most the size and degree of Π_1 (and thus also of Π), respectively.

Next, we apply the restriction ρ_3 from [Definition 5.61](#) to Π_2 . Let $\tilde{d} = d_2$ and $\Pi' = R_{d_2}(\Pi_2) \upharpoonright \rho_3$. By [Lemma 5.63](#) and the union bound, with probability larger than $1/2$, every term t of $R_{d_2}(\Pi_2)$ satisfying $t \upharpoonright \rho_3 \neq 0$ contains at most d_1 x - or y -variables and at most d_1 negative z -variables. Together with the bounds obtained after ρ_1 , this implies that every term of Π' has degree at most $8d_1$. This is because $|K|$ is less than d_2 due to ρ_2 , and the size of Π_2 is less than $\frac{1}{2}(16/15)^{\frac{d_1}{16d_2^2}} = 2^{\Theta(d_2)}$ for a sufficiently small ϵ .

Furthermore, since $\delta > 4\alpha$, we have $d_2 = \omega(\kappa + \log(mN))$. Hence, by choosing ϵ sufficiently small, the size assumption

$$|\Pi_2| \leq \frac{(4/3)^{d_2}}{8mN2^\kappa}$$

of [Lemma 5.64](#) is satisfied. Thus, Π' can be extended to a refutation Π_3 of $\tilde{\phi}_{G'}(\mathbf{A} \upharpoonright \rho_1 \upharpoonright \rho_2 \upharpoonright \rho_3, \kappa)$ of degree at most $8d_1 + O(1)$. Thus, by the union bound, there is ρ_3 for which both of the above events happen and Π_3 is a valid refutation of $\tilde{\phi}_{G'}(\mathbf{A} \upharpoonright \rho_1 \upharpoonright \rho_2 \upharpoonright \rho_3, \kappa)$ of degree at most $8d_1 + O(1)$.

Finally, we apply the reduction τ from [Definition 5.67](#) to Π_3 . By [Lemma 5.68](#), $\Pi_3 \upharpoonright \tau$ can be extended to a refutation Π_4 of $\text{FPHP}(G')$. The degree of Π_4 is at most $8(d_1 + O(1))(\kappa + 1) = \frac{cn^\delta}{4\Delta_{\max}} + O(\kappa) \leq \frac{cn^\delta}{2\Delta_{\max}}$. This contradicts the degree bound from [Proposition 5.50](#). $\square$

6 Proof Complexity Generators for SA

In this section, we will state our results for a stronger formula $\text{BT}^\bullet\text{Rank}_n^m(A)$, which is similar to $\text{BTRank}_n^m(A)$ except that it does not use z -variables. Formally, it is defined as follows.

Definition 6.1 (The CNF encoding $\text{BT}^\bullet\text{Rank}_n^m(A)$). Assume $m > n$ are positive integers and $A \in \{0, 1\}^{m \times m}$ is an m by m Boolean matrix. The (simple) bamboo-tree CNF encoding, denoted $\text{BT}^\bullet\text{Rank}_n^m(A)$, uses input variables $x_{i,k}, y_{k,i} : i \in [m], k \in [n]$ together with extension variables $u_{i,j,k} : i, j \in [m], k \in [n]$, and consists of the following axioms.

Output Axioms: For every $i, j \in [m]$, the clause $u_{i,j,n}$ if $A_{i,j} = 1$ and the clause $\neg u_{i,j,n}$ if $A_{i,j} = 0$.

Summation Base Axioms: For every $i, j \in [m]$, $u_{i,j,1} = x_{i,1}y_{1,j}$, encoded by the clauses $u_{i,j,1} \vee \neg x_{i,1} \vee \neg y_{1,j}$, $x_{i,1} \vee \neg u_{i,j,1}$ and $y_{1,j} \vee \neg u_{i,j,1}$.

Summation Axioms: For every $i, j \in [m]$, and $k \in \{2, \dots, n\}$, $u_{i,j,k} = u_{i,j,k-1} + x_{i,k}y_{k,j}$, encoded by the following six clauses:

$$\begin{aligned} & x_{i,k} \vee \neg u_{i,j,k-1} \vee u_{i,j,k}, \\ & x_{i,k} \vee u_{i,j,k-1} \vee \neg u_{i,j,k}, \\ & y_{k,j} \vee \neg u_{i,j,k-1} \vee u_{i,j,k}, \\ & y_{k,j} \vee u_{i,j,k-1} \vee \neg u_{i,j,k}, \\ & \neg x_{i,k} \vee \neg y_{k,j} \vee \neg u_{i,j,k-1} \vee \neg u_{i,j,k}, \\ & \neg x_{i,k} \vee \neg y_{k,j} \vee u_{i,j,k-1} \vee u_{i,j,k}. \end{aligned}$$

In what follows, it will be useful to additionally enforce that each row of X and each column of Y have odd weight. Let $m > n$ be positive integers and consider a matrix $A \in \{0, 1\}^{m \times m}$. We extend A to an $(m+1) \times (m+1)$ matrix $\tilde{A}$ by adding a new column and row. For every $i, j \in [m+1]$ we define

$$\tilde{A}_{i,j} = \begin{cases} A_{i,j} & i \in [m] \text{ and } j \in [m], \\ n \bmod 2 & \text{if } i = m+1 \text{ and } j = m+1, \\ 1 & \text{if } i = m+1 \text{ and } j \in [m], \\ 1 & \text{if } i \in [m] \text{ and } j = m+1. \end{cases} \quad (35)$$

Next, we construct a partial restriction σ , that we apply to $\text{BT}^\bullet\text{Rank}_n^{m+1}(\tilde{A})$ that fixes the last $((m+1)\text{th})$ row of X and the last $((m+1)\text{th})$ column of Y to the all-ones vector. Formally, for every $k \in [n]$, we have $\sigma(x_{m+1,k}) = \sigma(y_{k,m+1}) = 1$, $\sigma(\bar{x}_{m+1,k}) = \sigma(\bar{y}_{k,m+1}) = 0$, $\sigma(u_{m+1,m+1,k}) = (k \bmod 2)$ and $\sigma(\bar{u}_{m+1,m+1,k}) = 1 - (k \bmod 2)$. The variables of the form $u_{m+1,j,k}$, where $j \in [m]$ and $k \in [n]$, and $u_{i,m+1,k}$, where $i \in [m]$ and $k \in [n]$, (and their negated variants) are left unassigned. We denote by $\text{BT}^\bullet\text{Rank}_n^m(A)$ the restricted formula $\text{BT}^\bullet\text{Rank}_n^{m+1}(\tilde{A}) \upharpoonright \sigma$. Observe that $\text{BT}^\bullet\text{Rank}_n^m(A)$ is weaker than $\text{BT}^\bullet\text{Rank}_n^m(A)$ since it contains additional restrictions encoding $\sum_{k \in [n]} x_{i,k} = 1$, for all $i \in [m]$, and $\sum_{k \in [n]} y_{k,j} = 1$, for all $j \in [m]$, using the extra u -variables $u_{i,j,k}$ with either $i = m+1$ or $j = m+1$, and (augmented) Output, Summation Base and Summation Axioms.

Our plan for proving a lower bound for $\text{BT}^\bullet\text{Rank}_n^m(A)$ is to reduce it to $\text{BT}^\bullet\text{Rank}_{\tilde{n}}^m(A)$ for some $\tilde{n} = \Theta(n)$, in such a way that any small SA refutation is transformed into one of a small row degree (see Definition 6.3), and then to establish an explicit row degree lower bound for the latter formula. We start with the latter.

6.1 Row Degree Lower Bound

We start by proving a row degree bound on $\text{BT}^\bullet\text{Rank}_n^m(A)$.

Definition 6.2. Let t be a term in the variables of $\text{BT}^\bullet\text{Rank}_n^m(A)$. We say that $i \in [m]$ is *X-row-mentioned in t* if there is $k \in [n]$ such that $x_{i,k}$ or $\bar{x}_{i,k}$ appears in t . We say that $j \in [m]$ is *Y^T-row-mentioned in t* if there is $k \in [n]$ such that $y_{k,j}$ or $\bar{y}_{k,j}$ appears in t . We say that $i \in [m]$ is *U-left-row-mentioned in t* if there are $j \in [m+1]$ and $k \in [n]$ such that $u_{i,j,k}$ or $\bar{u}_{i,j,k}$ appears in t . We say that $j \in [m]$ is *U-right-row-mentioned in t* if there are $i \in [m+1]$ and $k \in [n]$ such that $u_{i,j,k}$ or $\bar{u}_{i,j,k}$ appears in t .

Definition 6.3. [Row degree] Let t be a term in the variables of $\text{BT}^\bullet\text{Rank}_n^m(A)$. Let I be the set of all $i \in [m]$ that are either X-row-mentioned or U-left-row-mentioned in t . Similarly, let J be the set of all $j \in [m]$ that are either Y^T-row-mentioned or U-right-row-mentioned in t . We say that t *mentions rows (I, J)* . The *row degree* of t is defined as the sum $|I| + |J|$.

Lower bounds on the SA degree are typically established by constructing a *pseudoexpectation*, which is a linear mapping that satisfies the following properties: it maps all low-degree multilinear polynomials to $\mathbb{R}$, vanishes on multiples of the axioms, and is non-zero on the constant-1 polynomial. We present such a mapping for multilinear terms with small row degree.

Lemma 6.4. For every $m > n$ and every $A \in \{0,1\}^{m \times m}$, every SA refutation of $\text{BT}^\bullet\text{Rank}_n^m(A)$ requires row degree $n - 1$.

Before proving the lemma, we construct a distribution of the pairs of matrices which will be used in the proof.

Definition 6.5 (The random family $\mathcal{D}_{I,J}$ for $\text{BT}^\bullet\text{Rank}_n^m(A)$). Let $I, J \subseteq [m]$, $d = |I| + |J| \leq n - 2$ and recall the definition of the matrix $\tilde{A}$ from (35). We define $\mathcal{D}_{I,J}$ as the uniform distribution on the pairs of full-rank matrices $M \in \{0,1\}^{(I \cup \{m+1\}) \times [n]}$ and $N \in \{0,1\}^{[n] \times (J \cup \{m+1\})}$ that satisfy $MN = \tilde{A}_{I \cup \{m+1\}, J \cup \{m+1\}}$ and for every $k \in [n]$, $M_{m+1,k} = N_{k,m+1} = 1$.

Based on $\mathcal{D}_{I,J}$ we construct a random assignment $\rho_{I,J}$.

Definition 6.6 (The random assignment $\rho_{I,J}$ for $\text{BT}^\bullet\text{Rank}_n^m(A)$). Let $I, J \subseteq [m]$ and $d = |I| + |J| \leq n - 2$. We define a random partial assignment $\rho_{I,J}$ as follows.

1. Sample a pair of matrices (M, N) according to $\mathcal{D}_{I,J}$.

2. Assign $\rho_{I,J}(X_{I,[n]}) = M_{I,[n]}$, meaning that for every $i \in I$ and $k \in [n]$, set $\rho_{I,J}(x_{i,k}) = M_{i,k}$. Similarly, assign $\rho_{I,J}(Y_{[n],J}) = N_{[n],J}$: for every $j \in J$ and $k \in [n]$, set $\rho_{I,J}(y_{k,j}) = N_{k,j}$.
3. For every $(i, j) \in ((I \cup \{m+1\}) \times (J \cup \{m+1\}) \setminus \{(m+1, m+1)\})$, and $k \in [n]$, set

$$\rho_{I,J}(u_{i,j,k}) = \left(\sum_{\ell=1}^k M_{i,\ell} N_{\ell,j} \bmod 2 \right).$$

4. For every variable θ assigned by $\rho_{I,J}$, set $\rho_{I,J}(\bar{\theta}) = 1 - \rho_{I,J}(\theta)$.

A key property of $\rho_{I,J}$ is that it remains consistent when restricted to smaller sets of rows. Let $I, J \subseteq [m]$ with $|I| + |J| \leq n - 2$ and consider any two subsets $I' \subseteq I$ and $J' \subseteq J$. Then, for every term t that mentions rows (I', J') , we have

$$\Pr[\rho_{I',J'}(t) = 1] = \Pr[\rho_{I,J}(t) = 1]. \quad (36)$$

We show this by induction on $|I \setminus I'| + |J \setminus J'|$. The base case is trivial and the inductive step follows from the following lemma.

Lemma 6.7. *Let $m > n$ be positive integers and let $A \in \{0, 1\}^{m \times m}$ be a Boolean matrix. Let $I, J \subseteq [m]$ with $|I| + |J| < n - 2$, $i_0 \in [m] \setminus I$ and $j_0 \in [m] \setminus J$. Then for every term t that mentions rows (I, J) we have*

$$\Pr[\rho_{I,J}(t) = 1] = \Pr[\rho_{I \cup \{i_0\}, J}(t) = 1] = \Pr[\rho_{I, J \cup \{j_0\}}(t) = 1].$$

Proof. By symmetry, it is sufficient to show

$$\Pr[\rho_{I,J}(t) = 1] = \Pr[\rho_{I, J \cup \{j_0\}}(t) = 1].$$

Let $\mathcal{D}'$ be the following distribution.

1. Sample a pair of matrices (M, N) according to $\mathcal{D}_{I,J}$.
2. Choose uniformly at random a vector $v \in \{0, 1\}^n$ such that $Mv = \tilde{A}_{I \cup \{m+1\}, j_0}$ and the matrix $N' = \begin{bmatrix} N & v \end{bmatrix}$ has full rank (that is, $\text{rank}(N') = |J| + 2$).
3. Output the pair (M, N') .

Showing that $\mathcal{D}'$ is the same as $\mathcal{D}_{I, J \cup \{j_0\}}$ concludes the proof.

Fix a pair of matrices (M, N) resulting from sampling from $\mathcal{D}_{I,J}$ in [Item 1](#). The number of choices of a vector $v \in \{0, 1\}^n$ that satisfy $Mv = \tilde{A}_{I \cup \{m+1\}, j_0}$ is $2^{n-|I|-1}$. This is because M has full rank. Given such v , the rank of the matrix $N' = \begin{bmatrix} N & v \end{bmatrix}$ is either $|J| + 1$ or $|J| + 2$. The former happens if and only if v is linearly dependent with the columns of N . That is, there exists a vector $\mu \in \{0, 1\}^{J \cup \{m+1\}}$ with $N\mu = v$. Moreover, each such μ corresponds uniquely to v since the rank of N is full.

Substituting $N\mu = v$ into $Mv = \tilde{A}_{I \cup \{m+1\}, j_0}$, we conclude that μ must satisfy $MN\mu = \tilde{A}_{I \cup \{m+1\}, j_0}$, which is equivalent to $\tilde{A}_{I \cup \{m+1\}, J \cup \{m+1\}}\mu = \tilde{A}_{I \cup \{m+1\}, j_0}$ due to our choice of (M, N) . Let ν be the number of vectors μ satisfying this last equality. Observe that $\nu \leq 2^{|J|+1}$ because of the dimension of μ .

Thus, there are $\nu \leq 2^{|J|+1}$ vectors v satisfying $Mv = \tilde{A}_{I \cup \{m+1\}, j_0}$ and $\text{rank}(N') = |J| + 1$. For the remaining $2^{n-|I|-1} - \nu > 0$ vectors, we have $\text{rank}(N') = |J| + 2$. This implies that the definition of $\mathcal{D}'$ is correct (that is, the vectors v from [Item 2](#) exist). Moreover, the number of choices of such v – that is, $2^{n-|I|-1} - \nu$ – is independent of the initial choice of (M, N) and is determined by A, I, J , and j_0 .

Every pair of matrices in $\mathcal{D}_{I,J \cup \{j_0\}}$ has a unique restriction to $\mathcal{D}_{I,J}$, and every pair of matrices in $\mathcal{D}_{I,J}$ has the same number $2^{n-|I|-1} - \nu$ of admissible extensions. Therefore, $\mathcal{D}'$ is the uniform distribution $\mathcal{D}_{I,J \cup \{j_0\}}$.

Observe that this also implies that each distribution $\mathcal{D}_{I,J}$ is non-empty: clearly $\mathcal{D}_{\emptyset,\emptyset}$ is non-empty and the above induction step shows that if $\mathcal{D}_{I,J}$ is non-empty, then so are $\mathcal{D}_{I \cup \{i_0\},J}$ and $\mathcal{D}_{I,J \cup \{j_0\}}$. $\square$

Proof of Lemma 6.4. Let S_{n-2} be the vector space of all multilinear polynomials in the variables of $\text{BT}^\bullet\text{Rank}_n^m(A)$ with row degree at most $n-2$. That is, $S_{n-2} \subseteq \mathbb{R}[\vec{x}, \vec{y}, \vec{u}] / (\vec{x}^2 - \vec{x}, \vec{y}^2 - \vec{y}, \vec{u}^2 - \vec{u})$. We construct a linear mapping

$$R: S_{n-2} \rightarrow \mathbb{R}$$

that satisfies the following properties.

- $R(1) = 1$.
- $R(t) \geq 0$ for every term t with row degree at most $n-2$.
- For every axiom axiom and every term t such that $\text{axiom} \cdot t$ has row degree at most $n-2$, $R(\text{axiom} \cdot t) = 0$.

This mapping rules out the existence of a refutation of $\text{BT}^\bullet\text{Rank}_n^m(A)$ with row degree at most $n-2$. Indeed, in any SA refutation of $\text{BT}^\bullet\text{Rank}_n^m(A)$, the left-hand side becomes non-negative under R , while the right-hand side evaluates to $R(-1) = -1$, leading to a contradiction.

Let t be a term in the variables of $\text{BT}^\bullet\text{Rank}_n^m(A)$ with row degree at most $n-2$; namely, it mentions rows (I, J) with $|I| + |J| \leq n-2$. We define $R(t) = \Pr[\rho_{I,J}(t) = 1]$. Note that $\rho_{I,J}$ assigns values to all the variables in t , which implies that $\rho_{I,J}(t) \in \{0, 1\}$. We then extend R to all multilinear polynomials of row degree at most $n-2$ linearly.

It immediately follows from the definition that $R(1) = 1$ and $R(t) \geq 0$ on all terms of row degree at most $n-2$.

(36) simplifies verifying the last properties of R . Let t be a term and axiom an axiom of $\text{BT}^\bullet\text{Rank}_n^m(A)$. Suppose the polynomial $t \cdot \text{axiom}$ mentions rows (I, J) with $|I| + |J| \leq n-2$. Expanding $t \cdot \text{axiom}$ as a sum of monomials $\sum_{\ell=1}^q \alpha_\ell t_\ell$, where each monomial t_ℓ mentions rows (I_ℓ, J_ℓ) with $I_\ell \subseteq I$ and $J_\ell \subseteq J$, we have:

$$\begin{aligned} R(t \cdot \text{axiom}) &= R\left(\sum_{\ell=1}^q \alpha_\ell t_\ell\right) = \sum_{\ell=1}^q \alpha_\ell R(t_\ell) = \sum_{\ell=1}^q \alpha_\ell \Pr[\rho_{I_\ell, J_\ell}(t_\ell) = 1] = \sum_{\ell=1}^q \alpha_\ell \Pr[\rho_{I, J}(t_\ell) = 1] \\ &= \sum_{\ell=1}^q \alpha_\ell \mathbb{E}[\rho_{I, J}(t_\ell)] = \mathbb{E}\left[\sum_{\ell=1}^q \alpha_\ell \rho_{I, J}(t_\ell)\right] = \mathbb{E}[\rho_{I, J}(t \cdot \text{axiom})] = 0. \end{aligned}$$

The last equality holds because $t \cdot \text{axiom}$ evaluates to zero under *every* choice of matrices M and N in $\rho_{I, J}$: the x - and y -variables are assigned values from M and N , respectively, and the u -variables are assigned the corresponding partial inner products. Since $MN = \tilde{A}_{I \cup \{m+1\}, J \cup \{m+1\}}$, every relevant axiom of $\text{BT}^\bullet\text{Rank}_n^m(A)$ is satisfied, and thus $t \cdot \text{axiom}$ evaluates to zero. $\square$

6.2 Size to Row Degree Reduction

We define a random substitution that transforms a refutation of $\text{BT}^\bullet\text{Rank}_n^m(A)$ into a refutation of $\text{BT}^\bullet\text{Rank}_{\tilde{n}}^m(A)$ (for some $\tilde{n} = \Theta(n)$) while decreasing the row degree of the terms appearing in the

refutation. Recall the matrices $\mathcal{D}$ and $\mathcal{E}$ (they were already defined at the beginning of [Section 5.4.1](#)).

$$\mathcal{D} = \begin{pmatrix} 1 & * & 1 & 0 & 0 & * & * & * \\ 0 & 1 & * & 1 & * & * & * & 0 \\ * & * & 0 & * & 1 & 0 & 1 & * \\ * & 0 & * & * & * & 1 & 0 & 1 \end{pmatrix}, \quad \mathcal{E} = \begin{pmatrix} 1 & 0 & 0 & 1 & * & * & * & * \\ 0 & * & 1 & * & 1 & * & 0 & * \\ * & 1 & * & 0 & * & 0 & * & 1 \\ * & * & * & * & 0 & 1 & 1 & 0 \end{pmatrix}.$$

The following definition is similar to [Definition 5.21](#).

Definition 6.8 (Random substitution for $\text{BT}^\bullet\text{Rank}_n^m(A)$). Assume $m > n \geq 20$ and $(n-4)$ is divisible by 16 and let $\tilde{n} = (n-4)/8$. Let $A \in \{0,1\}^{m \times m}$ be a Boolean matrix. Consider fresh variables $\{\tilde{x}_{i,k} : i \in [m], k \in [\tilde{n}]\}$, $\{\tilde{y}_{k,j} : j \in [m], k \in [\tilde{n}]\}$ and $\{\tilde{u}_{i,j,k} : (i,j) \in ([m+1] \times [m+1]) \setminus \{(m+1, m+1)\}, k \in [\tilde{n}]\}$, which serve as the variables of $\text{BT}^\bullet\text{Rank}_{\tilde{n}}^m$. The random substitution ρ is defined as follows.

For each $i \in [m]$, we sample its *row type* $s_i \in [4]$ (which corresponds to a row of $\mathcal{D}$) and a random bit $\alpha_i \in \{0,1\}$. Then we set the first and $(n-1)$ th bits of X_i (the i th row of the matrix X from $\text{BT}^\bullet\text{Rank}_n^m(A)$) to α_i and the second and n th bits of X_i to 1. The remaining $(n-4)$ bits are partitioned into eight equal consecutive parts (corresponding to the columns of $\mathcal{D}$). That is, given $k \in \{3, \dots, n-2\}$, let $\eta(k) = ((k-3) \bmod \tilde{n}) + 1$ and $\nu(k) = \lfloor \frac{k-3}{\tilde{n}} \rfloor + 1$. Then $\nu(k)$ gives the index of the part k belongs to and $\eta(k)$ gives its index within this part. Each variable $x_{i,k}$ will be replaced with either a constant or $\tilde{x}_{i,\eta(k)}$ according to the value of $\mathcal{D}_{s_i,\nu(k)}$.

Similarly, for each $j \in [m]$, we choose its row type $r_j \in [4]$ (corresponding to a row of $\mathcal{E}$) and $\beta_j \in \{0,1\}$. Then the first and $(n-1)$ th bits of Y_j (the j th column of Y) are set to 1, the second and n th bits to β_j and the rest are set to a constant or $\tilde{y}_{\eta(k),j}$ according to the value of $\mathcal{E}_{r_j,\nu(k)}$.

Then the value of the u -variables is defined consistently with the above assignments.

Formally, this is defined by the following random process.

1. For each $i \in [m]$, uniformly and independently sample $s_i, r_i \in [4]$ and $\alpha_i, \beta_i \in \{0,1\}$.
2. For each $i \in [m]$ and $k \in [n]$, set

$$\rho(x_{i,k}) = \begin{cases} \alpha_i & \text{if } k \in \{1, n-1\}, \\ 1 & \text{if } k \in \{2, n\}, \\ \tilde{x}_{i,\eta(k)} & \text{if } k \in \{3, \dots, n-2\} \text{ and } \mathcal{D}_{s_i,\nu(k)} = *, \\ \mathcal{D}_{s_i,\nu(k)} & \text{if } k \in \{3, \dots, n-2\} \text{ and } \mathcal{D}_{s_i,\nu(k)} \in \{0,1\}. \end{cases}$$

3. For each $j \in [m]$ and $k \in [n]$, set

$$\rho(y_{k,j}) = \begin{cases} 1 & \text{if } k \in \{1, n-1\}, \\ \beta_j & \text{if } k \in \{2, n\}, \\ \tilde{y}_{\eta(k),j} & \text{if } k \in \{3, \dots, n-2\} \text{ and } \mathcal{E}_{r_j,\nu(k)} = *, \\ \mathcal{E}_{r_j,\nu(k)} & \text{if } k \in \{3, \dots, n-2\} \text{ and } \mathcal{E}_{r_j,\nu(k)} \in \{0,1\}. \end{cases}$$

4. For each $i, j \in [m]$ and $k \in [n]$, we define $\rho(u_{i,j,k})$ according to $\sum_{\ell=1}^k \rho(x_{i,\ell})\rho(y_{\ell,j}) \bmod 2$.

For every $b \in [8]$, define

$$\gamma_{i,j,b} = \begin{cases} A_{i,j} & \text{if } (\mathcal{D}_{s_i,b}, \mathcal{E}_{r_j,b}) = (*, *), \\ 1 & \text{if } (\mathcal{D}_{s_i,b}, \mathcal{E}_{r_j,b}) = (*, 1) \text{ or } (1, *), \\ 0 & \text{otherwise,} \end{cases}$$

and

$$p_{i,j,b} = \alpha_i + \beta_j + \sum_{b'=1}^{b-1} \gamma_{i,j,b'} \pmod{2}.$$

Then for each $i, j \in [m]$ and $k \in [n]$, set

$$\rho(u_{i,j,k}) = \begin{cases} \alpha_i & \text{if } k = 1, \\ \alpha_i + \beta_j & \text{if } k = 2, \\ A_{i,j} + \beta_j & \text{if } k = n - 1, \\ A_{i,j} & \text{if } k = n, \\ \tilde{u}_{i,j,\eta(k)}^{1-p_{i,j,\nu(k)}} & \text{if } k \in \{3, \dots, n-2\}, \mathcal{D}_{s_i,\nu(k)} = \mathcal{E}_{r_j,\nu(k)} = *, \\ \tilde{u}_{i,m+1,\eta(k)}^{1-p_{i,j,\nu(k)}} & \text{if } k \in \{3, \dots, n-2\}, \mathcal{D}_{s_i,\nu(k)} = *, \mathcal{E}_{r_j,\nu(k)} = 1, \\ \tilde{u}_{m+1,j,\eta(k)}^{1-p_{i,j,\nu(k)}} & \text{if } k \in \{3, \dots, n-2\}, \mathcal{D}_{s_i,\nu(k)} = 1, \mathcal{E}_{r_j,\nu(k)} = *, \\ p_{i,j,\nu(k)} + (\eta(k) \pmod{2}) & \text{if } k \in \{3, \dots, n-2\}, \mathcal{D}_{s_i,\nu(k)} = \mathcal{E}_{r_j,\nu(k)} = 1, \\ p_{i,j,\nu(k)} & \text{if } k \in \{3, \dots, n-2\}, \mathcal{D}_{s_i,\nu(k)} = 0 \text{ or } \mathcal{E}_{r_j,\nu(k)} = 0. \end{cases}$$

We proceed by showing that ρ eliminates every term in the variables of $\text{BT}^\bullet \text{Rank}_n^m(A)$ of high row degree.

Lemma 6.9. *Assume $m > n \geq 20$, $(n-4)$ is divisible by 16 and $A \in \{0,1\}^{m \times m}$ is a Boolean matrix. Let t be a term in the variables of $\text{BT}^\bullet \text{Rank}_n^m(A)$ and t' be any subterm of t such that t' does not mention columns $\{1, 2, n-1, n\}$. Let $(I, J) \subseteq [m] \times [m]$ be the rows that t' mentions. Then $\Pr[t \upharpoonright \rho \neq 0] \leq (\sqrt{7/8})^{|I|+|J|}$.*

Proof. Consider the bipartite graph G with parts I and J and edges

$$\{(i, j) : \exists k \in [n] \text{ such that } u_{i,j,k} \text{ or } \bar{u}_{i,j,k} \text{ appears in } t'\}.$$

Let t'' be a minimal subterm of t' that mentions rows (I, J) . In terms of G , the term t'' corresponds to one of its minimal edge covers. Such a cover does not contain any paths of length three, thus its connected components are

- isolated vertices $x_{i,k}$, $\bar{x}_{i,k}$, $y_{k,j}$ or $\bar{y}_{k,j}$,
- stars $\prod_{(j,k) \in P} u_{i_0,j,k} \prod_{(j,k) \in N} \bar{u}_{i_0,j,k}$ and $\prod_{(i,k) \in P} u_{i,j_0,k} \prod_{(i,k) \in N} \bar{u}_{i,j_0,k}$.

The minimality of t'' ensures that each of these components mentions its own disjoint set of indices.

In the first case, we have $\Pr[\rho(x_{i,k}) \neq 0], \Pr[\rho(\bar{x}_{i,k}) \neq 0], \Pr[\rho(y_{k,j}) \neq 0], \Pr[\rho(\bar{y}_{k,j}) \neq 0] \leq 3/4$ (because each column of $\mathcal{D}$ and $\mathcal{E}$ contains one 0 and one 1).

In the second case, assume that we have a star with r rays, say $S = \prod_{(j,k) \in P} u_{i_0,j,k} \prod_{(j,k) \in N} \bar{u}_{i_0,j,k}$ where $|P| + |N| = r$. This star U -left-row-mentions index $i_0 \in I$ and U -right-row-mentions r indices $\{j : \exists k \text{ such that } (j, k) \in P \sqcup N\} \subseteq J$. For every $s \in [4]$ and $\alpha \in \{0, 1\}$, the events

$$(u_{i_0,j,k} \upharpoonright \rho \neq 0 \mid s_{i_0} = s, \alpha_{i_0} = \alpha), \text{ where } (j, k) \in P,$$

and

$$(\bar{u}_{i_0,j,k} \upharpoonright \rho \neq 0 \mid s_{i_0} = s, \alpha_{i_0} = \alpha), \text{ where } (j, k) \in N,$$

are independent and each occur with probability at most $7/8$. Thus, all of them simultaneously occur with probability at most $(7/8)^r \leq (\sqrt{7/8})^{r+1}$. It implies that

$$\Pr[t' \upharpoonright \rho \neq 0] \leq \Pr[t'' \upharpoonright \rho \neq 0] \leq \left(\sqrt{7/8}\right)^{|I|+|J|}. \quad \square$$

Next, we show that the random substitution ρ transforms any SA refutation of $\text{BT}^\bullet\text{Rank}_n^m(A)$ into a refutation of $\text{BT}^\bullet\text{Rank}_{(n-4)/8}^m(A)$.

Lemma 6.10. *Assume $m > n \geq 20$, $(n-4)$ is divisible by 16 and $A \in \{0,1\}^{m \times m}$ is a Boolean matrix. Then for every SA refutation Π of $\text{BT}^\bullet\text{Rank}_n^m(A)$, $\Pi \upharpoonright \rho$ can be converted into a SA refutation of $\text{BT}^\bullet\text{Rank}_{(n-4)/8}^m(A)$ without increasing its row degree.*

Proof. Observe that due to our choice of n , $\tilde{n} = (n-4)/8$ is even. Recall the properties of $\mathcal{D}$ and $\mathcal{E}$:

- there is an odd number of indices $k \in [8]$ such that $\mathcal{D}_{i,k} = \mathcal{E}_{j,k} = *$;
- there is an even number of indices $k \in [8]$ such that $\mathcal{D}_{i,k} = *$ and $\mathcal{E}_{j,k} = 1$;
- there is an even number of indices $k \in [8]$ such that $\mathcal{D}_{i,k} = 1$ and $\mathcal{E}_{j,k} = *$.

In particular, these properties ensure that the auxiliary values $p_{i,j,b}$ and $\gamma_{i,j,b}$ are defined consistently. In particular, we have $\sum_{b=1}^8 \gamma_{i,j,b} = A_{i,j} \bmod 2$. This implies that $p_{i,j,8} = \alpha_i + \beta_j + A_{i,j} + \gamma_{i,j,8} \bmod 2$.

Each Output Axiom is satisfied by ρ since it always assigns $\rho(u_{i,j,n}) = A_{i,j}$.

Similarly, every Summation Base Axiom is also satisfied by the definition of ρ .

Now consider a Summation Axiom. For readability, we reason about the parity relation $u_{i,j,k} = u_{i,j,k-1} + x_{i,k}y_{k,j} \bmod 2$ semantically instead of the CNF that encodes it. Each such equation involves only constantly many variables, thus the choice of encoding does not affect the argument: all the derivations below can be stated as constant-size SA derivations in the CNF encoding without mentioning any additional rows. For brevity, we omit the $\bmod 2$ in the analysis below, but it is understood that all equalities are modulo 2. When $k \in \{2, n-1, n\}$, we have the following cases.

- For $k = 2$, the axiom is satisfied since it becomes $\alpha_i + \beta_j = \alpha_i + 1 \cdot \beta_j$.
- For $k = n$, the axiom is satisfied since it becomes $A_{i,j} = (A_{i,j} + \beta_j) + 1 \cdot \beta_j$.
- For $k = n-1$, we have $A_{i,j} + \beta_j = \rho(u_{i,j,n-2}) + \alpha_i \cdot 1$. If $(\mathcal{D}_{s_i,8}, \mathcal{E}_{r_j,8}) = (*, *)$, then $\gamma_{i,j,8} = A_{i,j}$ and $p_{i,j,8} = \alpha_i + \beta_j \bmod 2$. In this case, $\rho(u_{i,j,n-2}) = \tilde{u}_{i,j,\tilde{n}}^{1-p_{i,j,8}}$. Thus we have

$$A_{i,j} + \beta_j = \tilde{u}_{i,j,\tilde{n}}^{1-(\alpha_i + \beta_j \bmod 2)} + \alpha_i \cdot 1,$$

which is a consequence of the Output Axiom $\tilde{u}_{i,j,\tilde{n}} = A_{i,j}$. Similarly, if $(\mathcal{D}_{s_i,8}, \mathcal{E}_{r_j,8}) \in \{(*, 1), (1, *)\}$, we have a consequence of the Output Axiom $\tilde{u}_{i,m+1,\tilde{n}} = 1$ or $\tilde{u}_{m+1,j,\tilde{n}} = 1$, respectively. Otherwise, $\rho(u_{i,j,n-2})$ is a constant and the axiom is satisfied.

The remaining case is $k \in \{3, \dots, n-2\}$. We have two options depending on the value of $\eta(k)$.

Suppose that $\eta(k) = 1$.

- If $k = 3$, then either the axiom is satisfied or it becomes a Summation Base Axiom for rows (i, j) , $(i, m+1)$ or $(m+1, j)$ (depending on the values of $\mathcal{D}_{s_i,\nu(k)}$ and $\mathcal{E}_{r_j,\nu(k)}$).
- If $k > 3$, we have $\nu(k) > 1$ and $\eta(k-1) = \tilde{n}$. Then either the axiom is satisfied or it becomes a simple consequence of the Output Axiom involving $\rho(u_{i,j,k-1})$ and the Summation Base Axiom for rows (i, j) , $(i, m+1)$ or $(m+1, j)$ (depending on the values of $\mathcal{D}_{s_i,\nu(k)}$ and $\mathcal{E}_{r_j,\nu(k)}$).

Now assume that $\eta(k) > 1$. Then $\nu(k) = \nu(k-1)$ and $\eta(k-1) = \eta(k) - 1$. Then the axiom is either satisfied, or it becomes a Summation Axiom involving rows (i, j) , $(i, m+1)$ or $(m+1, j)$ (depending on the values of $\mathcal{D}_{s_i, \nu(k)}$ and $\mathcal{E}_{r_j, \nu(k)}$).

Moreover, applying ρ to the conical junta in the original refutation yields another conical junta, since ρ maps variables to variables or constants.

It follows from this case analysis that the image of every axiom of $\text{BT}^\bullet\text{Rank}_n^m(A)$ is either satisfied, an axiom of $\text{BT}^\bullet\text{Rank}_{\tilde{n}}^m(A)$, or has a short SA derivation from the relevant Summation Base Axiom and Output Axiom. These derivations involve only the rows mentioned by the original axiom (and possibly $m+1$, which does not contribute to the row degree) and therefore do not increase row degree. Replacing every restricted source axiom by the corresponding derivation converts $\Pi \upharpoonright \rho$ into a SA refutation of $\text{BT}^\bullet\text{Rank}_{\tilde{n}}^m(A)$. $\square$

Theorem 6.11. *Assume $m > n \geq 20$, $(n-4)$ is divisible by 16 and $A \in \{0,1\}^{m \times m}$ is a Boolean matrix. Then any SA refutation of $\text{BT}^\bullet\text{Rank}_n^m(A)$ requires size $(\sqrt{8/7})^{(n-4)/8-1}$.*

Proof. Let $\tilde{n} = (n-4)/8$ and suppose, for the sake of contradiction, that there exists a refutation Π of $\text{BT}^\bullet\text{Rank}_n^m(A)$ with size $|\Pi| < (\sqrt{8/7})^{\tilde{n}-1}$. By the union bound and Lemma 6.9, there exists ρ such that the row degree of any term in $\Pi \upharpoonright \rho$ is at most $\tilde{n} - 2$. By Lemma 6.10, $\Pi \upharpoonright \rho$ can be converted into a refutation Π' of $\text{BT}^\bullet\text{Rank}_{\tilde{n}}^m(A)$ of row degree at most $\tilde{n} - 2$. This contradicts the row degree lower bound $\tilde{n} - 1$ by Lemma 6.4. $\square$

7 Structural Results

7.1 Amplification of the Rank Principle

In this section, we prove an amplification theorem for the weak rank principle, namely that the proof complexity of Rank_n^m is robust for different choices of parameters (m, n) , under low-degree algebraic reductions and assuming that m and n are not too close. The proof is an adaptation of known amplification theorems for weak pigeonhole principles [44, 64, 83]. This result should be regarded as a basic property of the weak rank principle that essentially follows from previous constructions. Nonetheless, it plays a foundational role in later sections, and we include it here for completeness. Moreover, this result provides another aspect in which the weak rank principle behaves similarly to the weak pigeonhole principles.

Theorem 7.1. *Let $m > 2n$ be positive integers, $\epsilon > 0$, and $d = O(\epsilon^{-1} \log(m/n))$, then there is a degree- d algebraic reduction*

$$\text{Rank}_n^{(1+\epsilon)n} \leq_d^{\text{alg}} \text{Rank}_n^m.$$

We prove Theorem 7.1 in two steps. First, we reduce Rank_n^{2n} to Rank_n^m for arbitrarily large m .

Theorem 7.2. *Let $m > 2n$ be positive integers and $d = \lceil \log(m/n) \rceil$. Then there is a degree- d algebraic reduction*

$$\text{Rank}_n^{2n} \leq_d^{\text{alg}} \text{Rank}_n^m.$$

Proof. Let $X \in \mathbb{F}^{2n \times n}$ and $Y \in \mathbb{F}^{n \times 2n}$ be the variables of Rank_n^{2n} . For every positive integer i , we will build a reduction from Rank_n^{2n} to $\text{Rank}_n^{2^i n}$ inductively. Let $X^{(i)} \in \mathbb{F}^{2^i n \times n}$ and $Y^{(i)} \in \mathbb{F}^{n \times 2^i n}$ be the variables of $\text{Rank}_n^{2^i n}$. Abusing notation, we will also specify our reduction $R_i : \mathbb{F}^{4n^2} \rightarrow \mathbb{F}^{2^{i+1}n^2}$ by defining the following two matrices $X^{(i)} \in \mathbb{F}^{2^i n \times n}$ and $Y^{(i)} \in \mathbb{F}^{n \times 2^i n}$ whose entries are polynomials over (X, Y) .

This is done inductively. First, $X^{(1)} = X$ and $Y^{(1)} = Y$. For each $i > 1$:

$$X^{(i+1)} = \underbrace{\begin{pmatrix} X^{(i)} & 0 \\ 0 & X^{(i)} \end{pmatrix}}_{2^{i+1}n \times 2n} \cdot \underbrace{X}_{2n \times n}, \quad Y^{(i+1)} = \underbrace{Y}_{n \times 2n} \cdot \underbrace{\begin{pmatrix} Y^{(i)} & 0 \\ 0 & Y^{(i)} \end{pmatrix}}_{2n \times 2^{i+1}n}.$$

The idea underlying our reduction is that if $XY = I_{2n}$, then $X^{(i)}Y^{(i)} = I_{2^{i+1}n}$ for every $i \geq 1$. This is easily seen by an induction on i :

$$\begin{aligned} X^{(i+1)}Y^{(i+1)} &= \begin{pmatrix} X^{(i)} & 0 \\ 0 & X^{(i)} \end{pmatrix} XY \begin{pmatrix} Y^{(i)} & 0 \\ 0 & Y^{(i)} \end{pmatrix} \\ &= \begin{pmatrix} X^{(i)} & 0 \\ 0 & X^{(i)} \end{pmatrix} \begin{pmatrix} Y^{(i)} & 0 \\ 0 & Y^{(i)} \end{pmatrix} \\ &= I_{2^{i+1}n}. \end{aligned}$$

It is easy to see that each entry of $X^{(i)}$ ($Y^{(i)}$ respectively) is a degree- i polynomial over X (Y respectively). Now we prove that each axiom of $\text{Rank}_n^{2^{i+1}n}$, when composed with our reduction R , admits a degree- $2i$ NS-proof from the axioms of $\text{Rank}_n^{2^n}$. Again, this is via an induction on i . The case when $i = 1$ is trivial; assuming this is true for $\text{Rank}_n^{2^i n}$, we prove this for $\text{Rank}_n^{2^{i+1}n}$:

$$\begin{aligned} &X^{(i+1)}Y^{(i+1)} - I_{2^{i+1}n} \\ &= \begin{pmatrix} X^{(i)} & 0 \\ 0 & X^{(i)} \end{pmatrix} XY \begin{pmatrix} Y^{(i)} & 0 \\ 0 & Y^{(i)} \end{pmatrix} - I_{2^{i+1}n} \\ &= \begin{pmatrix} X^{(i)} & 0 \\ 0 & X^{(i)} \end{pmatrix} \begin{pmatrix} Y^{(i)} & 0 \\ 0 & Y^{(i)} \end{pmatrix} + \begin{pmatrix} X^{(i)} & 0 \\ 0 & X^{(i)} \end{pmatrix} (XY - I_{2n}) \begin{pmatrix} Y^{(i)} & 0 \\ 0 & Y^{(i)} \end{pmatrix} - I_{2^{i+1}n} \\ &= \underbrace{\begin{pmatrix} X^{(i)}Y^{(i)} - I_{2^{i+1}n} & 0 \\ 0 & X^{(i)}Y^{(i)} - I_{2^{i+1}n} \end{pmatrix}}_{\text{(Part I)}} + \underbrace{\begin{pmatrix} X^{(i)} & 0 \\ 0 & X^{(i)} \end{pmatrix} (XY - I_{2n}) \begin{pmatrix} Y^{(i)} & 0 \\ 0 & Y^{(i)} \end{pmatrix}}_{\text{(Part II)}}. \end{aligned}$$

By our induction hypothesis, each entry of (Part I) admits a degree- $2i$ NS-proof from the axioms of $\text{Rank}_n^{2^n}$. Since the axioms of $\text{Rank}_n^{2^n}$ only express $XY = I_{2n}$, and each $X^{(i)}$ and $Y^{(i)}$ are degree- i polynomials over (X, Y) , it follows that each entry of (Part II) also admits a degree- $2i$ NS-proof from the axioms of $\text{Rank}_n^{2^n}$. This proves our claim for $i + 1$.

Now we wrap up the proof. Let $k = \lceil \log(m/n) \rceil$ be the smallest integer such that $2^k n \geq m$. Let $X' \in \mathbb{F}^{m \times n}$ and $Y' \in \mathbb{F}^{n \times m}$ be the variables of Rank_n^m . Given (X, Y) , our reduction sets X' to be the first m rows of $X^{(k)}$ and sets Y' to be the first m columns of $Y^{(k)}$. It is easy to see that the reduction can be computed in degree k and that each axiom of Rank_n^m admits a degree- $2k$ NS-proof from the axioms of $\text{Rank}_n^{2^n}$. $\square$

Now we reduce $\text{Rank}_n^{(1+\epsilon)n}$ to $\text{Rank}_n^{2^n}$. The proof idea is very similar to that of [Theorem 7.2](#) but we include the full proof here for completeness.

Theorem 7.3. *Let n be a positive integer, $\epsilon > 0$ such that ϵn is an integer, and $d = \lceil 1/\epsilon \rceil$. Then there is a degree- d algebraic reduction*

$$\text{Rank}_n^{(1+\epsilon)n} \leq_d^{\text{alg}} \text{Rank}_n^{2^n}.$$

Proof. Let $X \in \mathbb{F}^{(1+\epsilon)n \times n}$ and $Y \in \mathbb{F}^{n \times (1+\epsilon)n}$ be the variables of $\text{Rank}_n^{(1+\epsilon)n}$. For every positive integer i , we will build a reduction from $\text{Rank}_n^{(1+\epsilon)n}$ to $\text{Rank}_n^{(1+i\epsilon)n}$ inductively. Let $X^{(i)} \in \mathbb{F}^{(1+i\epsilon)n \times n}$ and

$Y^{(i)} \in \mathbb{F}^{n \times (1+i\epsilon)n}$ be the variables of $\text{Rank}_n^{(1+i\epsilon)n}$. Abusing notation, we will also specify our reduction $R : \mathbb{F}^{2(1+\epsilon)n^2} \rightarrow \mathbb{F}^{2(1+i\epsilon)n^2}$ by defining the following two matrices $X^{(i)}$ and $Y^{(i)}$ whose entries are polynomials over (X, Y) .

This is done inductively. First, $X^{(1)} = X$ and $Y^{(1)} = Y$. For each $i > 1$:

$$X^{(i+1)} = \underbrace{\begin{pmatrix} X^{(i)} & 0 \\ 0 & I_{\epsilon n} \end{pmatrix}}_{(1+(i+1)\epsilon)n \times (1+\epsilon)n} \cdot \underbrace{X}_{(1+\epsilon)n \times n}, \quad Y^{(i+1)} = \underbrace{Y}_{n \times (1+\epsilon)n} \cdot \underbrace{\begin{pmatrix} Y^{(i)} & 0 \\ 0 & I_{\epsilon n} \end{pmatrix}}_{(1+\epsilon)n \times (1+(i+1)\epsilon)n}.$$

For the sake of intuition, let us see that if $XY = I_{(1+\epsilon)n}$, then $X^{(i)}Y^{(i)} = I_{(1+i\epsilon)n}$ for every $i \geq 1$. This is via an induction over i :

$$\begin{aligned} X^{(i+1)}Y^{(i+1)} &= \begin{pmatrix} X^{(i)} & 0 \\ 0 & I_{\epsilon n} \end{pmatrix} XY \begin{pmatrix} Y^{(i)} & 0 \\ 0 & I_{\epsilon n} \end{pmatrix} \\ &= \begin{pmatrix} X^{(i)} & 0 \\ 0 & I_{\epsilon n} \end{pmatrix} \begin{pmatrix} Y^{(i)} & 0 \\ 0 & I_{\epsilon n} \end{pmatrix} \\ &= \begin{pmatrix} I_{(1+i\epsilon)n} & 0 \\ 0 & I_{\epsilon n} \end{pmatrix} \\ &= I_{(1+(i+1)\epsilon)n}. \end{aligned}$$

It is easy to see that each entry of $X^{(i)}$ ($Y^{(i)}$ respectively) is a degree- i polynomial over X (Y respectively). Now we prove that each axiom of $\text{Rank}_n^{(1+i\epsilon)n}$, when composed with our reduction R , admits a degree- $2i$ NS-proof from the axioms of $\text{Rank}_n^{(1+\epsilon)n}$. Again, this is via an induction on i . The case that $i = 1$ is trivial; assuming this is true for $\text{Rank}_n^{(1+i\epsilon)n}$, we prove this for $\text{Rank}_n^{(1+(i+1)\epsilon)n}$:

$$\begin{aligned} &X^{(i+1)}Y^{(i+1)} - I_{(1+(i+1)\epsilon)n} \\ &= \begin{pmatrix} X^{(i)} & 0 \\ 0 & I_{\epsilon n} \end{pmatrix} XY \begin{pmatrix} Y^{(i)} & 0 \\ 0 & I_{\epsilon n} \end{pmatrix} - I_{(1+(i+1)\epsilon)n} \\ &= \begin{pmatrix} X^{(i)} & 0 \\ 0 & I_{\epsilon n} \end{pmatrix} (XY - I_{(1+\epsilon)n}) \begin{pmatrix} Y^{(i)} & 0 \\ 0 & I_{\epsilon n} \end{pmatrix} + \begin{pmatrix} X^{(i)} & 0 \\ 0 & I_{\epsilon n} \end{pmatrix} \begin{pmatrix} Y^{(i)} & 0 \\ 0 & I_{\epsilon n} \end{pmatrix} - I_{(1+(i+1)\epsilon)n} \\ &= \underbrace{\begin{pmatrix} X^{(i)} & 0 \\ 0 & I_{\epsilon n} \end{pmatrix} (XY - I_{(1+\epsilon)n}) \begin{pmatrix} Y^{(i)} & 0 \\ 0 & I_{\epsilon n} \end{pmatrix}}_{\text{(Part I)}} + \underbrace{\begin{pmatrix} X^{(i)}Y^{(i)} - I_{(1+i\epsilon)n} & 0 \\ 0 & 0 \end{pmatrix}}_{\text{(Part II)}} \end{aligned}$$

By our induction hypothesis, each entry of (Part II) admits a degree- $2i$ NS-proof from the axioms of $\text{Rank}_n^{(1+\epsilon)n}$. Since the axioms of $\text{Rank}_n^{(1+\epsilon)n}$ only express $XY = I_{(1+\epsilon)n}$, and each $X^{(i)}$ and $Y^{(i)}$ are degree- i polynomials over (X, Y) , it follows that each entry of (Part I) also admits a degree- $2i$ NS-proof from the axioms of $\text{Rank}_n^{(1+\epsilon)n}$. This proves our claim for $i + 1$.

Now we wrap up the proof. Let $k = \lceil 1/\epsilon \rceil$ be the smallest integer such that $(1 + k\epsilon)n \geq 2n$. Let $X' \in \mathbb{F}^{2n \times n}$ and $Y' \in \mathbb{F}^{n \times 2n}$ be the variables of Rank_n^{2n} . Given (X, Y) , our reduction sets X' to be the first $2n$ rows of $X^{(k)}$ and sets Y' to be the first $2n$ columns of $Y^{(k)}$. It is easy to see that the reduction can be computed in degree k and that each axiom of Rank_n^{2n} admits a degree- $2k$ proof from the axioms of $\text{Rank}_n^{(1+\epsilon)n}$. $\square$

The takeaway message from the above two reductions is that the proof complexity of weak rank

formulas Rank_n^m are robust (in the regime when m and n are not *too close*), in the sense that they can be reduced to each other via low-degree reductions.

Theorem 7.1. *Let $m > 2n$ be positive integers, $\epsilon > 0$, and $d = O(\epsilon^{-1} \log(m/n))$, then there is a degree- d algebraic reduction*

$$\text{Rank}_n^{(1+\epsilon)n} \leq_d^{\text{alg}} \text{Rank}_n^m.$$

Proof Sketch. This follows from [Theorem 7.2](#), [Theorem 7.3](#), and [Fact 4.9](#). $\square$

7.2 Low-Degree Reductions for WRank in Strong Proof Systems

For sufficiently strong proof systems closed under low-degree reductions, we show that it is enough to prove the hardness of WRank with $A = I$ to get a generator (namely, hardness of WRank for every A). We demonstrate this simple property for the bamboo-tree encoding $\text{BT}^\bullet \text{Rank}_n^m(A)$.

We show that the formula $\text{BT}^\bullet \text{Rank}_n^m(A)$ can be reduced to $\text{BT}^\bullet \text{Rank}_n^r(I_r)$, where $r > n$, by a degree-1 reduction. The main idea consists of substitution $X' = AX$ and $Y' = Y$ into the statement $X'Y' = A$, thus obtaining $AXY = A$, which is derivable from $XY = I_r$.

Lemma 7.4. *Let $A \in \{0, 1\}^{m \times m}$ be an arbitrary matrix of rank r with $r > n$. Then there is a degree-1 reduction from $\text{BT}^\bullet \text{Rank}_n^r(I_r)$ to $\text{BT}^\bullet \text{Rank}_n^m(A)$.*

Proof. Let the variables of $\text{BT}^\bullet \text{Rank}_n^r(I_r)$ be $X = (x_{i,k})_{i \in [r], k \in [n]}$, $Y = (y_{k,j})_{k \in [n], j \in [r]}$, and $U = (u_{i,j,k})_{i,j \in [r], k \in [n]}$, and the variables of $\text{BT}^\bullet \text{Rank}_n^m(A)$ be $X' = (x'_{i,k})_{i \in [m], k \in [n]}$, $Y' = (y'_{k,j})_{k \in [n], j \in [m]}$, and $U' = (u'_{i,j,k})_{i,j \in [m], k \in [n]}$. Here we are going to slightly abuse the notation of degree-1 reductions: we are going to construct the reduction R that substitutes (X', Y', U') with degree-1 polynomials in (X, Y, U) , but the axioms of $\text{BT}^\bullet \text{Rank}_n^m(A) \circ R$ will have *constant degree* NS-proofs from the axioms of $\text{BT}^\bullet \text{Rank}_n^r(I_r)$. Since the axioms of $\text{BT}^\bullet \text{Rank}$ are themselves have constant degrees, this does not pose any problems and such a reduction is sufficient for our applications.

We can consider the following degree-1 reduction R that computes (X', Y', U') from (X, Y, U) :

- for every $i \in [m], k \in [n]$, $x'_{i,k} = x_{i,k}$;
- for every $k \in [n], j \in [m]$, $y'_{k,j} = \sum_{b \in [r]} y_{k,b} A_{b,j}$;
- for every $i, j \in [m]$ and $k \in [n]$, $u'_{i,j,k} = \sum_{b \in [r]} A_{b,j} u_{i,b,k}$.

We show that $\text{BT}^\bullet \text{Rank}_n^r(I_r)$ reduces to $\text{BT}^\bullet \text{Rank}_n^m(A)$ under R .

Recall that in algebraic proof systems, we encode a clause $C = \bigvee_{i \in P} v_i \vee \bigvee_{i \in N} \neg v_i$ as the polynomial $\prod_{i \in P} (1 + v_i) \prod_{i \in N} v_i$. We give the following names to the axioms in $\text{BT}^\bullet \text{Rank}_n^r(I_r)$:

$$\begin{aligned} \text{OutAx}_{i,j} &= \delta_{i,j} + u_{i,j,n}, & i, j &\in [r], \\ \text{MultBaseAx}_{i,j} &= u_{i,j,0}, & i, j &\in [r], \\ \text{MultAx}_{i,j,k}^1 &= (1 + x_{i,k}) u_{i,j,k-1} (1 + u_{i,j,k}), & i, j &\in [r], k \in [n], \\ \text{MultAx}_{i,j,k}^2 &= (1 + x_{i,k}) (1 + u_{i,j,k-1}) u_{i,j,k}, & i, j &\in [r], k \in [n], \\ \text{MultAx}_{i,j,k}^3 &= (1 + y_{k,j}) u_{i,j,k-1} (1 + u_{i,j,k}), & i, j &\in [r], k \in [n], \\ \text{MultAx}_{i,j,k}^4 &= (1 + y_{k,j}) (1 + u_{i,j,k-1}) u_{i,j,k}, & i, j &\in [r], k \in [n], \\ \text{MultAx}_{i,j,k}^5 &= x_{i,k} y_{k,j} u_{i,j,k-1} u_{i,j,k}, & i, j &\in [r], k \in [n], \\ \text{MultAx}_{i,j,k}^6 &= x_{i,k} y_{k,j} (1 + u_{i,j,k-1}) (1 + u_{i,j,k}), & i, j &\in [r], k \in [n]. \end{aligned}$$

To simplify the proofs, we also consider the following three types of auxiliary polynomials:

$$\begin{aligned}\text{MultAx}_{i,j,k}^7 &= (1 + x_{i,k})(u_{i,j,k-1} + u_{i,j,k}), & i, j \in [r], k \in [n], \\ \text{MultAx}_{i,j,k}^8 &= (1 + y_{k,j})(u_{i,j,k-1} + u_{i,j,k}), & i, j \in [r], k \in [n], \\ \text{MultAx}_{i,j,k}^9 &= x_{i,k}y_{k,j}(1 + u_{i,j,k-1} + u_{i,j,k}), & i, j \in [r], k \in [n].\end{aligned}$$

It is easy to see that for all $i, j \in [r]$ and $k \in [n]$ (over characteristic 2):

$$\begin{aligned}\text{MultAx}_{i,j,k}^7 &= \text{MultAx}_{i,j,k}^1 + \text{MultAx}_{i,j,k}^2, \\ \text{MultAx}_{i,j,k}^8 &= \text{MultAx}_{i,j,k}^3 + \text{MultAx}_{i,j,k}^4, \\ \text{MultAx}_{i,j,k}^9 &= \text{MultAx}_{i,j,k}^5 + \text{MultAx}_{i,j,k}^6.\end{aligned}$$

Let axiom be any polynomial from $\text{BT}^\bullet\text{Rank}_n^m(A)$. There are the following cases:

- **Output Axioms:** Suppose axiom = $A_{i,j} + u'_{i,j,n}$, where $i, j \in [m]$. Then

$$\begin{aligned}\text{axiom} \circ R &= A_{i,j} + \sum_{b \in [r]} A_{b,j} u_{i,b,n} \\ &= \sum_{b \in [r]} C_{b,j} ((\delta_{i,b} + u_{i,b,n}) + \delta_{i,b}) + A_{i,j} \\ &= \sum_{b \in [r]} A_{b,j} \text{OutAx}_{i,b},\end{aligned}$$

which is a linear combination of the axioms of $\text{BT}^\bullet\text{Rank}_n^r(I_r)$.

- **Multiplication Base Axioms:** Suppose axiom = $u'_{i,j,0}$, where $i, j \in [m]$. Then

$$\text{axiom} \circ R = \sum_{b \in [r]} A_{b,j} u_{i,b,0} = \sum_{b \in [r]} A_{b,j} \text{MultBaseAx}_{i,b},$$

which is a linear combination of the axioms of $\text{BT}^\bullet\text{Rank}_n^r(I_r)$.

- **Multiplication Axioms:**

1. Suppose axiom = $(1 + x'_{i,k})u'_{i,j,k-1}(1 + u'_{i,j,k})$, where $i, j \in [m]$, and $k \in [n]$. Then

$$\begin{aligned}\text{axiom} \circ R &= (1 + x_{i,k}) \left(\sum_{b_1 \in [r]} A_{b_1,j} u_{i,b_1,k-1} \right) (1 + \sum_{b_2 \in [r]} A_{b_2,j} u_{i,b_2,k}) \\ &= \sum_{b_1, b_2 \in [r]} A_{b_1,j} A_{b_2,j} u_{i,b_1,k-1} \text{MultAx}_{i,b_2,k}^7,\end{aligned}$$

which is a degree-4 NS proof of $\text{axiom} \circ R$ from $\text{BT}^\bullet\text{Rank}_n^r(I_r)$. This polynomial identity can be verified by induction on r .

2. Suppose axiom = $(1 + x'_{i,k})(1 + u'_{i,j,k-1})u'_{i,j,k}$, where $i, j \in [m]$, and $k \in [n]$. Then $\text{axiom} \circ R$ has a degree-4 NS-proof from $\text{BT}^\bullet\text{Rank}_n^r(I_r)$ due to symmetry: simply exchange $u_{a,b,k-1}$ and $u_{a,b,k}$ in [Item 1](#).
3. Suppose axiom = $(1 + y'_{k,j})u'_{i,j,k-1}(1 + u'_{i,j,k})$, where $i, j \in [m]$ and $k \in [n]$. Then

$$\text{axiom} \circ R^2 = (1 + \sum_{b_3 \in [r]} A_{b_3,j} y_{k,b_3}) \left(\sum_{b_1 \in [r]} A_{b_1,j} u_{i,b_1,k-1} \right) (1 + \sum_{b_2 \in [r]} A_{b_2,j} e_{i,b_2,k})$$

$$\begin{aligned}
&= \sum_{b_1, b_2 \in [r]} A_{b_1, j} A_{b_2, j} u_{i, b_1, k-1} \text{MultAx}_{i, b_2, k}^8 \\
&\quad + \sum_{\substack{b_1, b_2, b_3 \in [r] \\ b_2 \neq b_3}} A_{b_1, j} A_{b_2, j} A_{b_3, j} y_{k, b_3} u_{i, b_1, k-1} (\text{MultAx}_{i, b_2, k}^7 \\
&\quad \quad + \text{MultAx}_{i, b_2, k}^9 + x_{i, k} \text{MultAx}_{i, b_2, k}^8),
\end{aligned}$$

which is a degree-6 NS proof of $\text{axiom} \circ R$ from $\text{BT}^\bullet \text{Rank}_n^m(I)$.

4. Suppose $\text{axiom} = (1 + y'_{k, j})(1 + u'_{i, j, k-1})u'_{i, j, k}$, where $i, j \in [m]$, and $k \in [n]$. Then $\text{axiom} \circ R$ has a degree-6 NS-proof from $\text{BT}^\bullet \text{Rank}_n^r(I_r)$ due to symmetry: simply exchange $u_{a, b, k-1}$ and $u_{a, b, k}$ in [Item 3](#).
5. Suppose $\text{axiom} = x'_{i, k} y'_{k, j} u'_{i, j, k-1} u'_{i, j, k}$, where $i, j \in [m]$ and $k \in [n]$. Then

$$\begin{aligned}
\text{axiom} \circ R &= x_{i, k} \left(\sum_{b_3 \in [r]} A_{b_3, j} y_{k, b_3} \right) \left(\sum_{b_1 \in [r]} A_{b_1, j} u_{i, b_1, k-1} \right) \left(\sum_{b_2 \in [m]} A_{b_2, j} u_{i, b_2, k} \right) \\
&= \sum_{b_1, b_2 \in [r]} A_{b_1, j} A_{b_2, j} u_{i, b_1, k-1} \text{MultAx}_{i, b_2, k}^9 \\
&\quad + \sum_{\substack{b_1, b_2, b_3 \in [r] \\ b_2 \neq b_3}} A_{b_1, j} A_{b_2, j} A_{b_3, j} y_{k, b_3} u_{i, b_1, k-1} (\text{MultAx}_{i, b_2, k}^9 \\
&\quad \quad + x_{i, k} \text{MultAx}_{i, b_2, k}^8),
\end{aligned}$$

which is a degree-6 NS proof of $\text{axiom} \circ R$ from $\text{BT}^\bullet \text{Rank}_n^r(I_r)$.

6. Suppose $\text{axiom} = x'_{i, k} y'_{k, j} (1 + u'_{i, j, k-1})(1 + u'_{i, j, k})$, where $i, j \in [m]$ and $k \in [n]$. Then

$$\begin{aligned}
\text{axiom} \circ R &= x_{i, k} \left(\sum_{b_3 \in [r]} A_{b_3, j} y_{k, b_3} \right) \left(\sum_{b_1 \in [r]} A_{b_1, j} u_{i, b_1, k-1} \right) \left(\sum_{b_2 \in [m]} A_{b_2, j} u_{i, b_2, k} \right) \\
&= \sum_{b_2 \in [r]} A_{b_2, j} \text{MultAx}_{i, b_2, k}^9 \\
&\quad + \sum_{\substack{b_2, b_3 \in [r] \\ b_2 \neq b_3}} A_{b_2, j} A_{b_3, j} y_{k, b_3} (\text{MultAx}_{i, b_2, k}^9 + x_{i, k} \text{MultAx}_{i, b_2, k}^8) \\
&\quad + (x'_{i, k} y'_{k, j} u'_{i, j, k-1} u'_{i, j, k}) \circ R,
\end{aligned}$$

which is a degree-6 NS proof of $\text{axiom} \circ R$ from $\text{BT}^\bullet \text{Rank}_n^r(I_r)$. $\square$

8 Hardness of Proving Circuit Lower Bounds

In this section, we show that there is a low-degree reduction from the weak rank principles to circuit lower bound statements. In other words, for every algebraic proof system $\mathcal{P}$ closed under low-degree reductions, if $\mathcal{P}$ cannot prove the weak rank principle, then $\mathcal{P}$ cannot prove circuit lower bounds for the truth table of any function. We also unconditionally show the hardness of such circuit lower bound statements for $\text{PCR}_{\mathbb{F}_2}$. Although $\text{PCR}_{\mathbb{F}_2}$, when measured by size, does not seem to be closed under low-degree reductions, we still manage to prove the hardness of circuit lower bound statements for $\text{PCR}_{\mathbb{F}_2}$ by reducing them to an appropriate iterated variant of the weak rank principle.

Remark 8.1 (Comparison with [72, 75]). Razborov showed that low-degree polynomial calculus proofs [72] and short resolution proofs [75] cannot prove any circuit lower bounds. Both proofs proceed by first

showing that (variants of) weak pigeonhole principles are hard for the underlying proof system and then reducing the weak pigeonhole principle to the circuit lower bound statements. Intuitively, these weak proof systems cannot prove circuit lower bounds because they cannot *count*.

Complementing Razborov's results above, in this section we show that there is a low-degree reduction from the rank principles to the circuit lower bound statements. Intuitively, weak proof systems that cannot reason about *linear algebra* cannot prove circuit lower bounds, either. Our result demonstrates yet another aspect in which the (weak) rank principles behave as an algebraic analogue of the (weak) pigeonhole principle.

Roughly speaking, our arguments uses *non-commutative algebraic branching programs* (ncABPs) as a proxy. Nisan [60] showed that the ncABP complexity of every non-commutative homogeneous polynomial f is *characterised* by the rank of certain matrix associated with f . Hence, if $\mathcal{P}$ cannot prove the weak rank principles, then $\mathcal{P}$ cannot prove any lower bound on the ncABP complexity of non-commutative homogeneous polynomials. In Section 8.1, we formalise the above intuition as a proof complexity reduction from the weak rank principles to polynomial equations expressing ncABP lower bound. Then, in Section 8.2, we show that the unprovability of ncABP lower bounds imply the unprovability of lower bounds for general Boolean circuits.

8.1 Lower Bounds for Non-commutative Algebraic Branching Programs

Fix an underlying field $\mathbb{F}$. We consider non-commutative polynomials over the variable set $X = \{x_1, x_2, \dots, x_n\}$. Denote

$$X^k = \{x_{i_1}x_{i_2}\dots x_{i_k} : i_1, i_2, \dots, i_k \in [n]\}$$

as the set of degree- k non-commutative monomials.

Formalisation of ncABP lower bounds. Let f be a non-commutative homogeneous polynomial of degree d over the variable set $X = \{x_1, x_2, \dots, x_n\}$. Let $\vec{r} = (r_0, r_1, \dots, r_d)$ where $r_0 = r_d = 1$, we now define $\text{lb}_{\text{ncABP}}(f, \vec{r})$, which is a set of polynomial equations expressing that f cannot be computed by a non-commutative ABP where the i -th layer has r_i nodes.

(Variables) For every layer $1 \leq i \leq d$ and every variable x_j ($1 \leq j \leq n$), there is a matrix $M_{i,x_j} \in \mathbb{F}^{r_{i-1} \times r_i}$ among the variables of $\text{lb}_{\text{ncABP}}(f, \vec{r})$. We will also denote M_i as the $r_{i-1} \times r_i$ matrix where the (u, v) -th entry is the (formal) linear form $\sum_{j \in [n]} M_{i,x_j}[u, v]x_j$. Note that here, $M_{i,x_j}[u, v]$ is a variable of the polynomial equation system $\text{lb}_{\text{ncABP}}(f, \vec{r})$ and x_j is a formal variable of f .

(Equations) For every degree- d non-commutative monomial $m \in X^d$, $\text{lb}_{\text{ncABP}}(f, \vec{r})$ contains a degree- d equation stating that the m -th coefficient of f is equal to the m -th coefficient of f' , where f' is the non-commutative polynomial computed by the ABP. Suppose that $m = x_{j_1}x_{j_2}\dots x_{j_d}$, then this equation can be written as

$$\prod_{i=1}^d M_{i,x_{j_i}} - \text{coeff}[f, m] = 0.$$

(Note that $\prod_{i=1}^d M_{i,x_{j_i}}$ is a 1×1 matrix over $\mathbb{F}$, i.e., a scalar in $\mathbb{F}$.)

Theorem 8.2. *Let f be any non-commutative homogeneous polynomial of degree d over the variable set $X = \{x_1, x_2, \dots, x_n\}$ and let $N = n^d$. Let $r \in \mathbb{N}$ and denote $\vec{r} = (1, \underbrace{r, r, \dots, r}_{(d-1) \text{ } r\text{'s}}, 1)$.*

Then there is a degree-2 reduction from Rank_r^N to $\text{lb}_{\text{ncABP}}(f, \vec{r})$.

Proof. Identify $[N]$ with X^d , i.e., the set of non-commutative degree- d monomials over X . Consider what we call the *brute-force ncABP* computing f : this is a ncABP consisting of $d + 1$ layers, with each internal layer of width N . The ncABP consists of N disjoint (with respect to nodes, except for the source and sink nodes) paths, where each such path from the root to the sink computes a single non-commutative monomial. More precisely, we can describe the brute-force ncABP in terms of matrices whose entries are linear forms over X as follows:

- For every $m \in X^d$, let $M_1^{\text{bf}}[1, m] = \text{coeff}[f, m]x_1$ where x_1 is the first variable in m .
- For every $2 \leq i < d$ and every $m \in X^d$, let $M_i^{\text{bf}}[m, m] = x_i$ where x_i is the i -th variable in m ; for every $m, m' \in X^d$ with $m \neq m'$, let $M_i^{\text{bf}}[m, m'] = 0$.
- For every $m \in X^d$, let $M_d^{\text{bf}}[m, 1] = x_d$ where x_d is the d -th variable in m .

Let $X \in \mathbb{F}^{N \times r}$ and $Y \in \mathbb{F}^{r \times N}$ be the variables of Rank_r^N and $(M_1, M_2, \dots, M_d)$ be the variables of $\text{lb}_{\text{ncABP}}(f, \vec{r})$. Consider the following degree-2 mapping R that maps (X, Y) into:

- $M_1 = M_1^{\text{bf}} \cdot X$, $M_d = Y \cdot M_d^{\text{bf}}$;
(What this actually means is that $M_{1,x_j} = M_{1,x_j}^{\text{bf}} \cdot X$ and $M_{d,x_j} = Y \cdot M_{d,x_j}^{\text{bf}}$ for every $x_j \in X$.)
- $M_i = Y \cdot M_i^{\text{bf}} \cdot X$ for every $2 \leq i < d$.
(Similarly, what this actually means is that $M_{i,x_j} = Y \cdot M_{i,x_j}^{\text{bf}} \cdot X$ for every $x_j \in X$.)

We prove this is a valid reduction from Rank_r^N to $\text{lb}_{\text{ncABP}}(f, \vec{r})$. Towards this end fix any monomial $m = x_{j_1}x_{j_2}\dots x_{j_d} \in X^d$ and consider the polynomial corresponding to m :

$$p_m = \prod_{i=1}^d M_{i,x_{j_i}} - \text{coeff}[f, m].$$

Then, $p_m \circ R$ is the following polynomial over (X, Y) :

$$p_m \circ R = M_{1,x_1}^{\text{bf}}(XY)M_{2,x_2}^{\text{bf}}(XY)M_{3,x_3}^{\text{bf}}\dots M_{d-1,x_{d-1}}^{\text{bf}}(XY)M_{d,x_d}^{\text{bf}} - \text{coeff}[f, m].$$

On the other hand, the following is immediate from the definitions of M_i^{bf} :

$$M_{1,x_1}^{\text{bf}}M_{2,x_2}^{\text{bf}}M_{3,x_3}^{\text{bf}}\dots M_{d-1,x_{d-1}}^{\text{bf}}M_{d,x_d}^{\text{bf}} - \text{coeff}[f, m] = 0.$$

Now we have the following telescoping sum:

$$\begin{aligned} p_m \circ R &= M_{1,x_1}^{\text{bf}}M_{2,x_2}^{\text{bf}}\dots M_{d,x_d}^{\text{bf}} - \text{coeff}[f, m] \\ &\quad + M_{1,x_1}^{\text{bf}}(XY - I)M_{2,x_2}^{\text{bf}}\dots M_{d,x_d}^{\text{bf}} \\ &\quad + M_{1,x_1}^{\text{bf}}(XY)M_{2,x_2}^{\text{bf}}(XY - I)M_{3,x_3}^{\text{bf}}\dots M_{d,x_d}^{\text{bf}} \\ &\quad + \dots \\ &\quad + M_{1,x_1}^{\text{bf}}(XY)M_{2,x_2}^{\text{bf}}(XY)M_{3,x_3}^{\text{bf}}\dots (XY - I)M_{d,x_d}^{\text{bf}} \\ &= \sum_{i=1}^{d-1} \text{Prefix}_i(XY - I)\text{Suffix}_{i+1}, \end{aligned} \tag{37}$$

where $\text{Prefix}_i = M_{1,x_1}^{\text{bf}}(XY)M_{2,x_2}^{\text{bf}}(XY)\dots (XY)M_{i,x_i}^{\text{bf}}$ and $\text{Suffix}_i = M_{i,x_i}^{\text{bf}}M_{i+1,x_{i+1}}^{\text{bf}}\dots M_{d,x_d}^{\text{bf}}$. It is easy to see that (37) is of the form $\sum_{u,v \in [N]} \text{axiom}_{u,v} p_{u,v}$, where $\text{axiom}_{u,v}$ is the (u, v) -th entry of $(XY - I)$, and each $p_{u,v}$ is a polynomial over (X, Y) of degree at most $2d - 2$. $\square$

As a corollary, there is a degree- $O(\log(n^d/r))$ reduction from Rank_r^{2r} to $\text{lb}_{\text{ncABP}}(f, \vec{r})$.

8.2 Boolean Circuit Lower Bounds

We show that there is a degree- $O(n)$ algebraic reduction from polynomial equations expressing ncABP lower bounds to polynomial equations expressing Boolean circuit lower bounds. This implies that for every proof system $\mathcal{P}$ closed under low-degree algebraic reductions, if $\mathcal{P}$ cannot prove the weak rank principle, then it also cannot prove circuit lower bounds for any Boolean function. (Note that in our formalisation, the circuit lower bound statement has size $2^{O(n)}$, hence a reduction of degree $O(n)$ can be seen as *low degree*.)

Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ denote a Boolean function (represented by its length- 2^n truth table), $s \in \mathbb{N}$, and $\text{lb}(f, s)$ denote the polynomial equation system expressing that the circuit complexity of f is greater than s .

Theorem 8.3. *For every $f : \{0, 1\}^n \rightarrow \{0, 1\}$, there is a non-commutative polynomial $p_f : \mathbb{F}^2 \rightarrow \mathbb{F}$ of degree n such that for every parameter r , if we let $\vec{r} = (1, \underbrace{r, r, \dots, r}_{(n-1) \text{ } r\text{'s}}, 1)$, then*

$$\text{lb}_{\text{ncABP}}(p_f, \vec{r}) \leq_{\text{alg}}^{O(n)} \text{lb}(f, s) \text{ where } s = O(nr^2).$$

Formalisation of $\text{lb}(f, s)$. We now recall the precise definition of $\text{lb}(f, s)$ from [72, 75, 76].

First, we list all variables of $\text{lb}(f, s)$ along with their intended meaning. (In what follows, $a \in \{0, 1\}^n$ represents the input of the circuit, $v, v' \in [s]$ represent gates of the circuit, $v' < v$, $b \in \{1, 2\}$, and $i \in [n]$ represents a variable.)

Name	Intended Meaning
y_{av}	The Boolean value computed at the computational node v on the input a
y_{abv}	The value on a brought to v by the b -th input to v
$\text{Fanin}(v)$	This is 0 if v is a NOT gate and 1 if v is an AND gate or an OR gate
$\text{Type}(v)$	When $\text{Fanin}(v) = 1$, this is 0 if v is an AND gate and 1 if v is an OR gate
$\text{InputType}_b(v)$	This is 0 if the b -th input to v is a constant or a variable and 1 if it is one of the previous computational gates
$\text{InputType}'_b(v)$	When $\text{InputType}_b(v) = 0$, this is 0 if the b -th input to v is a constant, and is 1 if it is a variable
$\text{InputType}''_b(v)$	When $\text{InputType}_b(v) = \text{InputType}'_b(v) = 0$, this equals the b -th input to v
$\text{InputVar}_b(v, i)$	When $\text{InputType}_b(v) = 0$ and $\text{InputType}'_b(v) = 1$, this is 1 if and only if the b -th input to v is the i -th variable
$\text{INPUTVAR}_b(v, i)$	Equals $\bigvee_{i' < i} \text{InputVar}_b(v, i')$, introduced to keep bottom fan-in bounded
$\text{InputNode}_b(v, v')$	When $\text{InputType}_b(v) = 1$, this is 1 if and only if the b -th input to v is the previous gate v'
$\text{INPUTNODE}_b(v, v')$	Analogous to $\text{INPUTVAR}_b(v, i)$ (i.e., equals $\bigvee_{v'' < v'} \text{InputNode}_b(v, v'')$).

Table 1: Variables of $\text{lb}(f, s)$

Then, $\text{lb}(f, s)$ consists of the following clauses (we can translate them into low-degree polynomials in the standard way):

(Group 1) If the b -th input of v is a constant, then this constant is equal to y_{abv} :

- $\neg \text{InputType}_b(v) \wedge \neg \text{InputType}'_b(v) \rightarrow (y_{abv} = \text{InputType}''_b(v));$

(Group 2) If the b -th input of v is a variable, then there is a unique i such that it is the i -th variable:

- $\neg \text{InputType}_b(v) \wedge \text{InputType}'_b(v) \rightarrow \neg(\text{InputVar}_b(v, i) \wedge \text{InputVar}_b(v, i'))$ where $i \neq i'$;
- $\neg \text{InputType}_b(v) \wedge \text{InputType}'_b(v) \rightarrow (\text{INPUTVAR}_b(v, i) = \text{INPUTVAR}_b(v, i-1) \vee \text{InputVar}_b(v, i))$ where we define $\text{INPUTVAR}_b(v, 0) = 0$;
- $\neg \text{InputType}_b(v) \wedge \text{InputType}'_b(v) \rightarrow \text{INPUTVAR}_b(v, n)$;
- $\neg \text{InputType}_b(v) \wedge \text{InputType}'_b(v) \wedge \text{InputVar}_b(v, i) \rightarrow y_{abv} = a_i$;

(Group 3) If the b -th input of v is a previous gate, then there is a unique $v' < v$ such that it is the v' -th gate:

- $\text{InputType}_b(v) \rightarrow \neg(\text{InputNode}_b(v, v') \wedge \text{InputNode}_b(v, v''))$ where $v' \neq v''$;
- $\text{InputType}_b(v) \rightarrow (\text{INPUTNODE}_b(v, v') = \text{INPUTNODE}_b(v, v' - 1) \vee \text{InputNode}_b(v, v'))$ where we define $\text{INPUTNODE}_b(v, 0) = 0$;
- $\text{InputType}_b(v) \rightarrow \text{INPUTNODE}_b(v, v - 1)$;
- $\text{InputType}_b(v) \wedge \text{InputNode}_b(v, v') \rightarrow y_{abv} = y_{av'}$;

(Group 4) Every gate operates as expected:

- $\neg \text{Fanin}(v) \rightarrow (y_{av} = \neg y_{a1v})$;
- $\text{Fanin}(v) \wedge \neg \text{Type}(v) \rightarrow (y_{av} = y_{a1v} \wedge y_{a2v})$;
- $\text{Fanin}(v) \wedge \text{Type}(v) \rightarrow (y_{av} = y_{a1v} \vee y_{a2v})$;

(Group 5) The circuit computes f :

- $y_{as} = f(a)$.

The reduction. We first provide an outline of the reduction. Given any Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, let $p_f : \mathbb{F}^2 \rightarrow \mathbb{F}$ be the degree- n non-commutative polynomial over the variables z_0, z_1 such that:

- For every $x \in \{0, 1\}^n$, the coefficient of the monomial $z_{x_1} z_{x_2} \dots z_{x_n}$ is equal to $f(x)$. (Here we treat $f(x) \in \{0, 1\}$ as an element in $\mathbb{F}$.)

Suppose that p_f can be computed by a small ncABP P . For every degree- n non-commutative monomial m , let $\text{Coeff}(P, m)$ denote the coefficient of m in the non-commutative polynomial computed by P (this non-commutative polynomial should be equal to p_f). Then Coeff can be computed by a small Boolean circuit. Plugging P into Coeff gives us a small circuit for f .

In more details, let $\vec{r} = (r_0, r_1, \dots, r_n)$ where $r_0 = r_d = 1$ and $r_i = r$ for every $1 \leq i < d$. Let $M_1, \dots, M_n$ denote the matrices representing the ncABP P , where each M_i is of dimension $r_{i-1} \times r_i$. Note that each entry of each M_i is a linear form over (the non-commutative variables) z_0 and z_1 , hence we can write $M_i = M_{i,z_0} z_0 + M_{i,z_1} z_1$ where each $M_{i,z_b} \in \mathbb{F}^{r_{i-1} \times r_i}$. Every $x \in \{0, 1\}^n$ corresponds to the non-commutative monomial $m_x = z_{x_1} z_{x_2} \dots z_{x_n}$, and the coefficient of m_x is

$$\prod_{i=1}^n M_{i,z_{x_i}}. \quad (38)$$

Let $\text{Coeff}(\{M_{i,z_b}\}, x)$ denote the circuit that given $\{M_{i,z_b}\}$ and x as inputs, outputs the value of (38). The size of Coeff is at most $s = O(nr^2)$. Moreover, it is possible to design a circuit implementation

of Coeff such that for every internal gate g and every $x \in \{0, 1\}^n$, there is an $\mathbb{F}$ -polynomial $q_{g,x}$ of degree $O(n)$ over the inputs $\{M_{i,z_b}\}$ that computes the value of g .

Now we are ready to specify the reduction from $\text{lb}_{\text{ncABP}}(p_f, \vec{r})$ to $\text{lb}(f, s)$, where $s = O(nr^2)$. Let $\{M_{i,z_b}\}$ be the inputs of $\text{lb}_{\text{ncABP}}(p_f, \vec{r})$. Consider the circuit $C(x) = \text{Coeff}(\{M_{i,z_b}\}, x)$ that takes $x \in \{0, 1\}^n$ as an input. The variables

$$\begin{aligned} &\text{Fanin}(v), \text{Type}(v), \text{InputType}_b(v), \text{InputType}'_b(v), \\ &\text{InputVar}_b(v, i), \text{INPUTVAR}_b(v, i), \text{InputNode}_b(v, v'), \text{INPUTNODE}_b(v, v') \end{aligned} \quad (39)$$

are fixed constants that reflect the wiring of C . Note that the circuit Coeff takes $\{M_{i,z_b}\}$ as inputs, but the circuit C treats $\{M_{i,z_b}\}$ as hardwired constants. Therefore, some variables of the form $\text{InputType}'_b(v)$ are projections over $\{M_{i,z_b}\}$. Finally, the variables y_{av} and y_{bv} are internal values on the computation $C(a)$; all of them can be computed by a degree- $O(n)$ polynomial over $\{M_{i,z_b}\}$ (i.e., $q_{g,a}$ for the corresponding gate g).

Next we verify that every clause in $\text{lb}(f, s)$ (composed with our reduction) admits a degree- $O(n)$ NS-proof from $\text{lb}_{\text{ncABP}}(p_f, \vec{r})$. In fact this essentially follows from definition:

- (Group 1) is true since we assign both y_{bv} and $\text{InputType}'_b(v)$ the same projection over $\{M_{i,z_b}\}$.
- (Group 2) and (Group 3) only contains variables in (39) (which are constants under our reduction). They assert that the circuit has a valid structure; since our circuit C indeed has a valid structure, they are trivially true.

The only two exceptions are

$$\begin{aligned} &\neg \text{InputType}_b(v) \wedge \text{InputType}'_b(v) \wedge \text{InputVar}_b(v, i) \rightarrow y_{bv} = a_i; \text{ and} \\ &\text{InputType}_b(v) \wedge \text{InputNode}_b(v, v') \rightarrow y_{bv} = y_{bv'}. \end{aligned}$$

as they contain not only variables in (39) but also variables of the form y_{bv} and $y_{bv'}$. But they are also trivially true by our definitions of y_{bv} and $y_{bv'}$.

- (Group 4) is true because the polynomials $q_{g,x}$ correctly captures the values being computed at the gate g on input x .
- (Group 5) asserts that $y_{as} = f(a)$. This is exactly the a -th axiom in $\text{lb}_{\text{ncABP}}(p_f, \vec{r})$ and hence trivially admits a degree- $O(n)$ NS-proof from $\text{lb}_{\text{ncABP}}(p_f, \vec{r})$.

Corollary 8.4. *For every $f : \{0, 1\}^n \rightarrow \{0, 1\}$ and every parameter r , if the circuit complexity of f is greater than $s = O(nr^2)$, then there is a degree- $O(n^2)$ algebraic reduction from Rank_r^{2r} to $\text{lb}(f, s)$.*

Proof. Let $\vec{r} = (1, \underbrace{r, r, \dots, r}_{(n-1) \text{ } r\text{'s}}, 1)$ and $p_f : \mathbb{F}^2 \rightarrow \mathbb{F}$ be the non-commutative polynomial guaranteed by

Theorem 8.3, then there is a degree- $O(n)$ reduction from $\text{lb}_{\text{ncABP}}(p_f, \vec{r})$ to $\text{lb}(f, s)$. By **Theorem 8.2**, there is a degree-2 reduction from Rank_r^{2n} to $\text{lb}_{\text{ncABP}}(p_f, \vec{r})$. By **Theorem 7.2**, there is a degree- $O(n)$ reduction from Rank_r^{2r} to Rank_r^{2n} . Composing these three reductions (via **Fact 4.9**) gives us a degree- $O(n^2)$ reduction from Rank_r^{2r} to $\text{lb}(f, s)$. $\square$

Remark 8.5. The circuit for Coeff can be computed in NC^2 (since it computes an iterated matrix multiplication over $\mathbb{F}$). Hence, the above argument also gives a reduction from the weak rank principle to sentences expressing NC^2 circuit lower bounds. In particular, let $\text{lb}_{\text{NC}^2}(f, s)$ denote the collection of polynomial equations encoding that f cannot be computed by a circuit of size s and depth $\log^2 s$ (we do not attempt to formally define $\text{lb}_{\text{NC}^2}(f, s)$ here). The same argument above shows that for every

Boolean function f and every parameter r , letting $s = \text{poly}(n, r)$, there is a degree- $O(n^2)$ algebraic reduction from Rank_r^{2r} to $\text{lb}_{\text{NC}^2}(f, s)$. Hence, if the rank principle is hard for some proof system $\mathcal{P}$ (that is closed under low-degree reductions), then $\mathcal{P}$ cannot prove any lower bound against NC^2 circuits as well.

8.3 Lower Bounds for $\text{PCR}_{\mathbb{F}_2}$

We apply the iterated rank principles from Section 5.4.1 to obtain $\text{PCR}_{\mathbb{F}_2}$ size lower bounds for $\text{lb}_{\text{ncABP}}(f, \vec{r})$ and $\text{lb}^\oplus(f, s)$. The latter is a variant of $\text{lb}(f, s)$ that allows binary $\oplus$ gates. The first lower bound uses the algebraic iterated rank principle, whereas the second uses its bamboo-tree encoding.

8.3.1 Lower Bounds for Non-commutative ABPs

We start by describing the formula that encodes the existence of small non-commutative algebraic branching programs for a given non-commutative polynomial f .

Definition 8.6 (Lower bound formula for ncABPs). Let $n \geq 2$, $d \geq 2$, and $r \geq 1$ be integers. Let $f(\vec{z}) = f(z_1, z_2, \dots, z_n)$ be a non-commutative homogeneous polynomial over $\mathbb{F}_2$ of degree d in n variables. Let $\vec{r} = (r_0, r_1, \dots, r_d)$ with $r_0 = r_d = 1$. For every $\Delta \in [d]$ and $k \in [n]$, let M_{Δ, z_k} be a variable matrix of dimension $r_{\Delta-1} \times r_\Delta$. Let $\mathcal{M}_\Delta$ be the set of all degree- Δ non-commutative monomials in $z_1, \dots, z_n$. For every $\Delta \in [d]$ and $m \in \mathcal{M}_\Delta$, let $P_{\Delta, m}$ be a $1 \times r_\Delta$ row vector of variables.

The lower bound formula $\text{ncABP}(f, \vec{r})$ is defined as the following polynomial system:

$$P_{1, z_k} = M_{1, z_k} \quad k \in [n], \quad (40)$$

$$P_{\Delta, m z_k} = P_{\Delta-1, m} M_{\Delta, z_k} \quad k \in [n], \Delta \in \{2, \dots, d\}, m \in \mathcal{M}_{\Delta-1}, \quad (41)$$

$$P_{d, m} = \text{coeff}(f, m) \quad m \in \mathcal{M}_d. \quad (42)$$

The above formula encodes the existence of a non-commutative ABP of depth $d+1$ and width $\vec{r}$. The Δ th layer, where $\Delta \in [d]$, computes the row vector of homogeneous non-commutative polynomials, represented by their coefficient vectors:

$$\sum_{m \in \mathcal{M}_\Delta} P_{\Delta, m} m.$$

Edge labels between consecutive layers of the program are encoded by the matrices M_{Δ, z_k} . The last set of equations (42) ensures that the program computes the given polynomial f .

We show that any satisfying assignment to the formula $\text{IRank}_n^m(\mathbf{A}, \kappa, \Sigma)$ (for a suitable choice of parameters) gives a construction of a polynomial-size ncABP circuit for any non-commutative polynomial.

Theorem 8.7. *Let $n \geq 2$, $d \geq 2$, and $r \geq 1$ be integers. Let $f(\vec{z}) = f(z_1, z_2, \dots, z_n)$ be a non-commutative homogeneous polynomial over $\mathbb{F}_2$ of degree d in n variables. Let $\vec{r} = (1, \underbrace{r, r, \dots, r}_{d-1}, 1)$.*

Then any $\text{PCR}_{\mathbb{F}_2}$ refutation of $\text{ncABP}(f, \vec{r})$ requires size $2^{\Omega(\sqrt{r/d})}$.

Proof. We will construct a reduction to $\text{ncABP}(f, \vec{r})$ from $\text{IRank}_r^{nr}(\mathbf{A}, d-1, [n])$ for some particular family of matrices $\mathbf{A}$. Let Π be a $\text{PCR}_{\mathbb{F}_2}$ refutation of $\text{ncABP}(f, \vec{r})$.

For every $\pi = (j_1, \dots, j_{d-1}) \in [n]^{d-1}$ define $A^{(\pi)}$ as any $nr \times nr$ Boolean constant matrix satisfying

$$A_{1, (j_{d-1})r+1}^{(\pi)} = \text{coeff}(f, z_{j_1} \cdots z_{j_{d-1}} \cdot z_{j_d})$$

for all $j_d \in [n]$. The remaining entries of $A^{(\pi)}$ can be filled arbitrarily.

For every $\pi \in [n]^{<d-1}$, define $A^{(\pi)}$ as an $nr \times nr$ block matrix

$$A^{(\pi)} = \begin{bmatrix} X^{(\pi*1)} & X^{(\pi*2)} & \dots & X^{(\pi*n)} \end{bmatrix}.$$

Let $\mathbf{A} = \{A^{(\pi)} : \pi \in [n]^{\leq d-1}\}$.

Also split the matrix Y into n consecutive blocks of dimension $r \times r$:

$$Y = \begin{bmatrix} Y^{(1)} & Y^{(2)} & \dots & Y^{(n)} \end{bmatrix}.$$

Given $\pi \in [n]^{\leq d-1}$, $X_1^{(\pi)}$ denotes the first *row* of the matrix $X^{(\pi)}$. Similarly, given $k \in [n]$, $Y_1^{(k)}$ denotes the first *column* of the matrix $Y^{(k)}$. Now define the reduction τ as follows.

$$\begin{aligned} \tau(P_{\Delta,m}) &= X_1^{(j_1, \dots, j_{\Delta})} & \Delta \in [d-1], m = z_{j_1} \dots z_{j_{\Delta}} \in \mathcal{M}_{\Delta}, \\ \tau(P_{d,m}) &= A_{1, (j_d-1)r+1}^{(j_1, \dots, j_{d-1})} & m = z_{j_1} \dots z_{j_d} \in \mathcal{M}_d, \\ \tau(M_{1,z_k}) &= X_1^{(k)} & k \in [n], \\ \tau(M_{\Delta,z_k}) &= Y^{(k)} & \Delta \in \{2, \dots, d-1\}, k \in [n], \\ \tau(M_{d,z_k}) &= Y_1^{(k)} & k \in [n]. \end{aligned}$$

For every $k \in [n]$, τ assigns P_{1,z_k} to $X_1^{(k)}$ and M_{1,z_k} to $X_1^{(k)}$, which satisfies (40).

For every $m = z_{j_1} \dots z_{j_d} \in \mathcal{M}_d$, we have

$$A_{1, (j_d-1)r+1}^{(j_1, \dots, j_{d-1})} = \text{coeff}(f, z_{j_1} \dots z_{j_d}),$$

which satisfies (42).

Now consider the axioms (41). For every $\Delta \in \{2, \dots, d-1\}$, $m = z_{j_1} \dots z_{j_{\Delta-1}} \in \mathcal{M}_{\Delta-1}$ and $j_{\Delta} \in [n]$, under τ the axiom becomes

$$X_1^{(j_1, \dots, j_{\Delta})} = X_1^{(j_1, \dots, j_{\Delta-1})} Y^{(j_{\Delta})}.$$

Similarly, for every $m = z_{j_1} \dots z_{j_{d-1}} \in \mathcal{M}_{d-1}$ and $j_d \in [n]$, we have

$$A_{1, (j_d-1)r+1}^{(j_1, \dots, j_{d-1})} = X_1^{(j_1, \dots, j_{d-1})} Y_1^{(j_d)}.$$

In both cases, these polynomial equations are among the axioms of $\text{IRank}_r^{nr}(\mathbf{A}, d-1, [n])$.

Hence every axiom of $\text{ncABP}(f, \vec{r}) \upharpoonright \tau$ is either satisfied or becomes an axiom of $\text{IRank}_r^{nr}(\mathbf{A}, d-1, [n])$, and $\Pi \upharpoonright \tau$ is a refutation of $\text{IRank}_r^{nr}(\mathbf{A}, d-1, [n])$. By Theorem 5.44, this implies that the size of Π is $2^{\Omega(\sqrt{r/d})}$. $\square$

Observe that the size of $\text{ncABP}(f, \vec{r})$ is polynomial in $n^d d r^2$, thus the bound becomes non-trivial when $r = \omega(d^3 \log^2 n)$.

8.3.2 Lower Bounds for Boolean Circuits

As in [76], we consider Boolean circuits in the basis $\neg, \wedge, \vee, \oplus$, where the last three connectives are binary.

Definition 8.8 (Formula $\text{lb}^\oplus(f, s)$). Let n and s be positive integers, and let $f: \{0, 1\}^n \rightarrow \{0, 1\}$ be a Boolean function. The formula $\text{lb}^\oplus(f, s)$ is defined similarly to $\text{lb}(f, s)$ except that it also allows binary $\oplus$ connectives.

We show that if, for a particular choice of parameters, the formula $\psi_{P(H(G))}(\mathbf{A}, \kappa)$ is satisfiable, we can construct a polynomial-size circuit for any Boolean function.

Theorem 8.9. *There exists $\alpha > 0$ such that for every sufficiently large integer n , every integer $s \geq n^5$, and every Boolean function $f: \{0, 1\}^n \rightarrow \{0, 1\}$, any $\text{PCR}_{\mathbb{F}_2}$ refutation of $\text{lb}^\oplus(f, s)$ requires size $2^{\Omega((s^\alpha/n)^{1/3})}$.*

Proof. Let $r = s^{1/5} \geq n$ and $q = (r - 10)/8$. Note that $q = \Theta(r) = \Theta(s^{1/5})$.

Choose $\delta, d > 0$ from [Lemma 4.17](#) such that there exists a $(16r) \times q$ (q^δ, d) -lossless expander G with the appropriate degree bounds and denote $G' = P(H(G))$. Then fix any $0 < \alpha < \delta/4$. If $q^\alpha < n$, the lower bound is trivial, so we can assume that $q^\alpha \geq n$. This allows us to apply [Theorem 5.51](#) later.

We will reduce $\psi_{P(H(G))}(\mathbf{A}, n)$ to $\text{lb}^\oplus(f, s)$ for some particular $\mathbf{A}$. Let Π be a $\text{PCR}_{\mathbb{F}_2}$ refutation of $\text{lb}^\oplus(f, s)$.

Define A_1 and A_0 as $16r \times 16r$ matrices as follows:

$$A_1 = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix} \text{ and } A_0 = \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix}.$$

The important property of these matrices is that $(A_1)_{1,1} = 1$ and $(A_0)_{1,1} = 0$. For every $\pi \in \{0, 1\}^n$, define $A^{(\pi)} = A_{f(\pi)}$.

For every $\pi \in \{0, 1\}^{\leq n}$, we define block matrices $X^{(\pi)}$ and $Y^{(\pi)}$ as follows:

$$X^{(\pi)} = \begin{bmatrix} X^{(\pi,1)} & X^{(\pi,2)} & X^{(\pi,3)} & X^{(\pi,4)} \end{bmatrix}, \text{ and } Y^{(\pi)} = \begin{bmatrix} Y^{(\pi,1)} \\ Y^{(\pi,2)} \\ Y^{(\pi,3)} \\ Y^{(\pi,4)} \end{bmatrix}.$$

Here, for every $\mu \in [4]$, $k \in [r]$, and $j \in [16r]$ with $k \notin \mathcal{N}_{G'}(j)$, the (k, j) th entry of $Y^{(\pi, \mu)}$ is interpreted as 0. Then for every $\pi \in \{0, 1\}^{\leq n}$, $A^{(\pi)}$ is a $16r \times 16r$ block matrix

$$A^{(\pi)} = \begin{bmatrix} X^{(\pi*0)} & X^{(\pi*1)} & Y^{(\pi*0)T} & Y^{(\pi*1)T} \end{bmatrix}.$$

Let

$$\mathbf{A} = \{A^{(\pi)} : \pi \in \{0, 1\}^{\leq n}\}.$$

We construct the circuit $\mathcal{C}(v_1, \dots, v_n)$ that computes f , provided that $\psi_{P(H(G))}(\mathbf{A}, n)$ is satisfied. The circuit uses gates that correspond to the extension variables of $\psi_{P(H(G))}(\mathbf{A}, n)$.

1. For every $\mu \in [4]$, $i \in [16r]$, $k \in [r]$:

$$g_x[0, \mu, i, k] = x_{i,k}^{((), \mu)}.$$

2. For every $\ell \in [n]$, $\mu \in [4]$, $i \in [16r]$, $k \in [r]$:

$$g_x[\ell, \mu, i, k] = (\neg v_\ell \wedge g_w[\ell - 1, i, (\mu - 1)r + k, 0]) \vee (v_\ell \wedge g_w[\ell - 1, i, (\mu - 1)r + k + 4r, 0]).$$

3. For every $\mu \in [4]$, $j \in [16r]$, $k \in \mathcal{N}_{G'}(j)$:

$$g_y[0, \mu, k, j] = y_{k,j}^{((), \mu)}.$$

4. For every $\ell \in [n]$, $\mu \in [4]$, $j \in [16r]$, $k \in \mathcal{N}_{G'}(j)$:

$$g_y[\ell, \mu, k, j] = (\neg v_\ell \wedge g_w[\ell - 1, j, (\mu - 1)r + k + 8r, 0]) \vee (v_\ell \wedge g_w[\ell - 1, j, (\mu - 1)r + k + 12r, 0]).$$

5. For every $\ell \in \{0, \dots, n\}$, $\mu \in [4]$, $i, j \in [16r]$, $k \in \mathcal{N}_{G'}(j)$:

$$g_z[\ell, \mu, i, j, k] = g_x[\ell, \mu, i, k] \wedge g_y[\ell, \mu, k, j].$$

6. For every $\ell \in \{0, \dots, n\}$, $\mu \in [4]$, $i, j \in [16r]$:

$$g_u[\ell, \mu, i, j, (1)] = g_z[\ell, \mu, i, j, 1].$$

7. For every $\ell \in \{0, \dots, n\}$, $\mu \in [4]$, $i, j \in [16r]$ and every initial segment $\Sigma \cup (k)$ of $\mathcal{N}_{G'}(j)$ with $\Sigma \neq \emptyset$:

$$g_u[\ell, \mu, i, j, \Sigma \cup (k)] = g_u[\ell, \mu, i, j, \Sigma] \oplus g_z[\ell, \mu, i, j, k].$$

8. For every $\ell \in \{0, \dots, n\}$, $i, j \in [16r]$:

$$\begin{aligned} g_w[\ell, i, j, 1] &= g_u[\ell, 1, i, j, \mathcal{N}_{G'}(j)] \oplus g_u[\ell, 2, i, j, \mathcal{N}_{G'}(j)], \\ g_w[\ell, i, j, 2] &= g_u[\ell, 3, i, j, \mathcal{N}_{G'}(j)] \oplus g_u[\ell, 4, i, j, \mathcal{N}_{G'}(j)], \\ g_w[\ell, i, j, 0] &= g_w[\ell, i, j, 1] \oplus g_w[\ell, i, j, 2]. \end{aligned}$$

The output of the circuit is $g_w[n, 1, 1, 0]$. The size of the circuit is $O(nr^3)$, which is $o(s)$ for large enough n . It follows from the construction that for every $\pi \in \{0, 1\}^n$, $\mathcal{C}(\pi) = A_{1,1}^{(\pi)} = f(\pi)$.

We plug $\mathcal{C}(v_1, \dots, v_n)$ into $\text{lb}^\oplus(f, s)$. Under the resulting substitution τ , every axiom encoding the circuit structure is satisfied. The remaining gate axioms become the corresponding Binary AND, Summation Base, Summation, or Output Axioms of $\psi_{G'}(\mathbf{A}, n)$. For every input $\pi \in \{0, 1\}^n$, the circuit-output axiom becomes

$$g_w[n, 1, 1, 0] = A_{1,1}^{(\pi)} = f(\pi).$$

Thus, every axiom of $\text{lb}^\oplus(f, s) \upharpoonright \tau$ is either an identity or an axiom of $\psi_{G'}(\mathbf{A}, n)$. Consequently, $\Pi \upharpoonright \tau$ is a refutation of $\psi_{G'}(\mathbf{A}, n)$. By [Theorem 5.51](#), this implies that $\text{lb}^\oplus(f, s)$ requires a refutation of size $2^{\Omega((q^\delta/n)^{1/3})} = 2^{\Omega((s^{\delta/5}/n)^{1/3})} \geq 2^{\Omega((s^{\alpha/5}/n)^{1/3})}$. $\square$

Observe that the size of $\text{lb}^\oplus(f, s)$ is $\text{poly}(2^n, s)$. Thus, the bound becomes non-trivial when $s = \omega(n^{20/\alpha})$.

9 Proving Circuit Lower Bounds via the Weak Rank Principle

We show that the weak rank principle, when added as an axiom to a weak formal theory corresponding to $\text{AC}^0[p]$, is capable of proving lower bounds against $\text{AC}^0[p]$ circuits by formalising Smolensky's lower bound [78].

Roadmap of Section 9. Smolensky’s proof consists of two parts: a correlation bound against low-degree polynomials (i.e., any low-degree polynomial over $\mathbb{F}_p$ has to make many errors when computing $\prod_{i=1}^n x_i$) and the “polynomial method” (i.e., every small $\text{AC}^0[p]$ circuit is well-approximated by a low-degree polynomial over $\mathbb{F}_p$). Hence we arrange this section as follows:

- [Section 9.1](#) provides an exposition of Smolensky’s proof. The exposition is equivalent to the original proof, but it will be presented in a way that only involves simple manipulations computable in $\text{AC}^0[p]$ and the weak rank principles.

In [Section 9.1.1](#) we prove the correlation bound against low-degree $\mathbb{F}_p$ -polynomials. This is the main step where the weak rank principle is involved. In [Section 9.1.2](#) we show that small $\text{AC}^0[p]$ circuits can be approximated by low-degree $\mathbb{F}_p$ polynomials, and put everything together to finish the circuit lower bound.

- In [Section 9.2](#) we formalise the lower bound in $\overline{\mathbf{V}^0(p)} + \text{WRank}_p$.

We start with a description of the relevant theories ($\overline{\mathbf{V}^0(p)}$ and $\overline{\mathbf{V}^0(p)} + \text{WRank}_p$) in [Section 9.2.1](#). The main formalisation will appear in [Section 9.2.2](#). Finally, [Section 9.2.3](#) contains a postponed proof of a technical theorem related to error-correcting codes.

It is recommended to read [Section 9.1](#) before delving into [Section 9.2](#) for the following reasons. First, to make the invocation of weak rank principles explicit, the proofs in [Section 9.1](#) are presented in a somewhat different way than it is usually presented in textbooks (e.g., [8]); it would be beneficial to be familiarised with the proof before formalising it. Second, [Section 9.1](#) contains important definitions (such as the matrices on which the weak rank principle is invoked) that the formalisation in [Section 9.2](#) needs to reason about.

9.1 Smolensky’s Lower Bound

9.1.1 Correlation Bound Against Low-Degree Polynomials

We need the following correlation bound against low-degree polynomials.

Theorem 9.1. *Let $P : \mathbb{F}_p^n \rightarrow \mathbb{F}_p$ be a polynomial of degree at most $0.1\sqrt{n}$. Let $S = \{x \in \{-1, 1\}^n : P(x) = \prod_{i=1}^n x_i\}$, then $|S| \leq 0.99 \cdot 2^n$.*

Proof Ideas. The idea is that every function $f : S \rightarrow \mathbb{F}_p$ is equivalent to a polynomial of degree at most $d = n/2 + 0.1\sqrt{n}$ over S . In fact, let $T \subseteq [n]$ and $m = \prod_{i \in T} x_i$ be a multilinear monomial. If $|T| = \deg(m) \leq n/2$, then m is trivially equivalent to a polynomial of degree at most d (i.e., m itself). If $|T| = \deg(m) > n/2$, then the equality

$$m \equiv \left(\prod_{i \in [n] \setminus T} x_i \right) \cdot P(x) \quad (43)$$

still holds over S , and the RHS of (43) is a polynomial of degree at most d . Hence every multilinear monomial m is equivalent to a polynomial of degree at most d over S . Since $S \subseteq \{-1, 1\}^n$, every function $f : S \rightarrow \mathbb{F}_p$ can be written as a linear combination of multilinear monomials, hence is also equivalent to a polynomial of degree at most d .

Given this, it follows that there are at most

$$p^{\sum_{i=0}^d \binom{n}{i}} \leq p^{0.99 \cdot 2^n}$$

different functions $f : S \rightarrow \mathbb{F}_p$, indicating that $|S| \leq 0.99 \cdot 2^n$. □

To aid our formalisations later, we need to explicitly construct the matrices over which the weak rank principle is used (instead of using concepts such as *dimensions of vector spaces*). We now rephrase the above proof so that the invocation of the weak rank principle becomes clear.

Theorem 9.2. *Let $S \subseteq \{-1, 1\}^n$, $|S| \geq 0.999 \cdot 2^n$, $P : S \rightarrow \mathbb{F}_p$ be any polynomial of degree at most $0.1\sqrt{n}$. Then there exists some $z \in S$ such that $P(z) \neq \prod_{i=1}^n z_i$.*

Proof. Given $S \subseteq \{-1, 1\}^n$, we construct the following matrices $A \in \mathbb{F}_p^{|S| \times r}$ and $B \in \mathbb{F}_p^{r \times |S|}$ where $r = 0.99 \cdot 2^n$. We identify $[r]$ (i.e., columns of A and rows of B) with monomials m of degree at most $d = n/2 + 0.1\sqrt{n}$. Now, for an input $x \in S$ and a monomial $m \in [r]$:

- Define $A[x, m] = m(x)$ to be the value of m on input x . Note that A can be seen as a linear mapping for *polynomial evaluation*: Let $\vec{e} \in \mathbb{F}_p^r$ and think of $\vec{e}$ as the coefficient list of a polynomial P with degree at most d . (That is, for every monomial m of degree at most d , $\vec{e}_m$ is the coefficient of m in P .) Then $A\vec{e} \in \mathbb{F}_p^S$ is exactly the list of values of P over S .
- Define $B[m, x]$ to be the sum of the following values:
 1. if $\deg(m) \leq n/2$, we add $(\frac{p+1}{2})^n m(x)$ to $B[m, x]$;
 2. for every monomial m' that appears in P with coefficient $\alpha_{m'}$, if the degree of $\tilde{m} = m' \cdot m \cdot \prod_{i \in [n]} x_i$ is greater than $n/2$, we add $(\frac{p+1}{2})^n \cdot \alpha_{m'} \cdot \tilde{m}(x)$ to $B[m, x]$.

The intended meaning of B is the following. It was shown in [Theorem 9.1](#) that every function $f : S \rightarrow \mathbb{F}_p$ is equivalent to a polynomial of degree at most $d = n/2 + 0.1\sqrt{n}$. Let $\vec{v}_f \in \mathbb{F}_p^S$ be the vector representing the value list of f , then $B\vec{v}_f \in \mathbb{F}_p^r$ is exactly the coefficient list of this polynomial.

To be more precise, we introduce the following notation. Let $\mathcal{M}$ denote the set of all multilinear monomials on variables $z_1, \dots, z_n$, $\mathcal{M}_{\leq n/2}$ and $\mathcal{M}_{> n/2}$ denote the subset of monomials with degree $\leq n/2$ and $> n/2$ respectively. Let $f : S \rightarrow \mathbb{F}_p$ be a function, then it can be expressed as the polynomial

$$f(z) = \sum_{m \in \mathcal{M}} \lambda_{f,m} m(z)$$

where $\lambda_{f,m} = (\frac{p+1}{2})^n \sum_{x \in S} f(x) m(x)$.

To see this, notice that

$$\begin{aligned} \sum_{m \in \mathcal{M}} \lambda_{f,m} m(z) &= \sum_{m \in \mathcal{M}} \left(\frac{p+1}{2} \right)^n \sum_{x \in S} f(x) m(x) m(z) \\ &= \left(\frac{p+1}{2} \right)^n \sum_{x \in S} f(x) \sum_{m \in \mathcal{M}} m(x) m(z) \\ &= \left(\frac{p+1}{2} \right)^n \sum_{x \in S} f(x) \cdot \delta_{x,z} \cdot 2^n \\ &= f(z). \end{aligned} \quad (\text{over } \mathbb{F}_p)$$

Given any function $f : S \rightarrow \mathbb{F}_p$, denote

$$f_{\text{low}}(z) = \sum_{m \in \mathcal{M}_{\leq n/2}} \lambda_{f,m} m(z) \quad \text{and} \quad f_{\text{high}}(z) = \sum_{m \in \mathcal{M}_{> n/2}} \lambda_{f,m} m(z).$$

The following claim expresses the “intended meaning of B ” mentioned above:

Claim 9.3. Let $f : S \rightarrow \mathbb{F}_p$ be any function, $\vec{v}_f \in \mathbb{F}_p^S$ be the list of values of f on S , and $\hat{f}$ denote the polynomial whose coefficient list is $B \cdot \vec{v}_f$. Then for every $z \in S$,

$$\hat{f}(z) = f_{\text{low}}(z) + f_{\text{high}}(z) \cdot P(z) \cdot \prod_{i=1}^n z_i.$$

Proof. We have that

$$\begin{aligned} & f_{\text{low}}(z) + f_{\text{high}}(z) \cdot P(z) \cdot \prod_{i=1}^n z_i \\ &= \sum_{m \in \mathcal{M}_{\leq n/2}} \lambda_{f,m} m(z) + \sum_{\tilde{m} \in \mathcal{M}_{> n/2}} \lambda_{f,\tilde{m}} \tilde{m}(z) P(z) \prod_{i=1}^n z_i \\ &= \left(\frac{p+1}{2} \right)^n \sum_{x \in S} \left(\underbrace{\sum_{m \in \mathcal{M}_{\leq n/2}} f(x) m(x) m(z)}_{\text{(Part I)}} + \underbrace{\sum_{\tilde{m} \in \mathcal{M}_{> n/2}} f(x) \tilde{m}(x) \tilde{m}(z) P(z) \prod_{i=1}^n z_i}_{\text{(Part II)}} \right). \end{aligned} \quad (44)$$

Let $x \in S$, $m \in \mathcal{M}$, consider the coefficient of $f(x)m(z)$ in (44). If $\deg(m) \leq n/2$ then the contribution of (Part I) is $(\frac{p+1}{2})^n m(x)$; otherwise it is 0. For every monomial m' of P with coefficient $\alpha_{m'}$, if $\tilde{m} = m \cdot m' \cdot \prod_{i=1}^n z_i$ has degree $> n/2$, then it contributes to (Part II) a coefficient of $(\frac{p+1}{2})^n \tilde{m}(x) \alpha_{m'}$; otherwise it contributes nothing. It follows from the definition of B that

$$(44) = \sum_{x \in S} \sum_{m \in \mathcal{M}} B[m, x] f(x) m(z),$$

which by definition is equal to $\hat{f}(z)$. $\diamond$

We also claim that any witness that $A \cdot B$ is not the identity matrix gives rise to a bad input on which P and $\prod_{i=1}^n z_i$ differ:

Claim 9.4. If $z, x \in S$ and $(AB)[z, x] \neq \delta_{z,x}$, then $P(z) \neq \prod_{i=1}^n z_i$.

Proof. Let $f : \{-1, 1\}^n \rightarrow \mathbb{F}_p$ be the function that evaluates to 1 on input x and evaluates to 0 on every input in $\{-1, 1\}^n \setminus \{x\}$; note that $\vec{1}_x$ is exactly the “value list” vector of f . By Claim 9.3, $B \cdot \vec{1}_x \in \mathbb{F}_p^r$ is the coefficient list of the polynomial $f'(z) = f_{\text{low}}(z) + f_{\text{high}}(z) \cdot P(z) \cdot \prod_{i=1}^n z_i$. Note that $f(z) = f_{\text{low}}(z) + f_{\text{high}}(z) = \delta_{z,x}$. However, $(AB)[z, x] = f'(z) \neq \delta_{z,x}$. This means that $P(z) \neq \prod_{i=1}^n z_i$. $\diamond$

Since $|S| \geq 0.999 \cdot 2^n \gg r$, it follows from the weak rank principle that there exists $z, x \in S$ such that $(AB)[z, x] \neq \delta_{z,x}$. This implies that $P(z) \neq \prod_{i=1}^n z_i$ as shown by Claim 9.4. $\square$

The matrices A and B in the above proof will be crucial for our formalisation, hence we reiterate their definitions here. (Note that A does not depend on the polynomial P but B does.)

Definition 9.5. Let $N = 2^n$, $r = 0.99 \cdot 2^n$, and $P : \{-1, 1\}^n \rightarrow \mathbb{F}_p$ be a polynomial of degree at most $0.1\sqrt{n}$. Identify $[N]$ with $\{-1, 1\}^n$ and $[r]$ with the set of monomials of degree at most $d = n/2 + 0.1\sqrt{n}$. Define:

- A to be the $N \times r$ matrix such that $A[x, m] = m(x)$ for every $x \in \{-1, 1\}^n$ and $m \in \mathcal{M}_{\leq d}$.

- $B(P)$ to be the $r \times N$ matrix such that for every $x \in \{-1, 1\}^n$ and $m \in \mathcal{M}_{\leq d}$, $B[m, x]$ is the sum of the following values. (1) If $\deg(m) \leq n/2$, we add $(\frac{p+1}{2})^n m(x)$ to $B[m, x]$; (2) for every monomial m' that appears in P with coefficient $\alpha_{m'}$, if the degree of $\tilde{m} = m' \cdot m \cdot \prod_{i \in [n]} x_i$ is greater than $n/2$, then we add $(\frac{p+1}{2})^n \cdot \alpha_{m'} \cdot \tilde{m}(x)$ to $B[m, x]$.

9.1.2 Approximating $\text{AC}^0[p]$ Circuits by Low-Degree Polynomials

The next ingredient in the proof is the *low-degree* approximation of $\text{AC}^0[p]$ circuits [70]: For every parameter Δ and every $\text{AC}^0[p]$ circuit C of depth d and size s , there is a polynomial $P : \mathbb{F}_p^n \rightarrow \mathbb{F}_p$ of degree at most $(p\Delta)^d$ that agrees with C on at least a $(1 - s \cdot 0.9^\Delta)$ fraction of inputs. Here we treat $\{0, 1\}$ as a subset of $\mathbb{F}_p$ and we guarantee that for every $x \in \{0, 1\}^n$, $P(x) \in \{0, 1\}$. We also assume that C only consists of NOT gates, MOD_p gates, and OR gates (AND gates can be simulated by NOT and OR gates using De Morgan's law), and the definition of MOD_p gates is

$$\text{MOD}_p(x_1, x_2, \dots, x_n) = \begin{cases} 0 & \text{if } \sum_{i=1}^n x_i \equiv 0 \pmod{p}; \\ 1 & \text{otherwise.} \end{cases}$$

Defining the low-degree polynomials. In fact, the Razborov–Smolensky proof provides us a stronger guarantee: There is a *distribution $\mathcal{P}$ of low-degree polynomials* such that for every input $x \in \{0, 1\}^n$, with high probability over $P \leftarrow \mathcal{P}$, we have $P(x) = C(x)$. Using error-correcting codes, we can make $\mathcal{P}$ the uniform distribution over a support of size $L = 2^{O(n)}$.¹³ That is, there is a list of L polynomials $P_{C,1}, P_{C,2}, \dots, P_{C,L}$ of degree at most $(p\Delta)^d$ such that for every $x \in \{0, 1\}^n$,

$$\Pr_{j \leftarrow [L]} [P_{C,j}(x) \neq C(x)] \leq s \cdot 0.9^\Delta.$$

In particular, let $\ell \leq \text{poly}(t)$, $\text{Enc} : \mathbb{F}_p^t \rightarrow \mathbb{F}_p^\ell$ be an $\mathbb{F}_p$ -linear error-correcting code such that for every $\vec{x} \in \mathbb{F}_p^t \setminus \{0^t\}$, at least a 0.1 fraction of indices $i \in [\ell]$ satisfy that $\text{Enc}(\vec{x})_i \neq 0$. Let $\vec{j} = (j_1, j_2, \dots, j_\Delta)$ where each $j_i \in [\ell]$. We will have a polynomial $P_{C,\vec{j}}$ for each such $\vec{j}$, hence there are $L = \ell^\Delta$ polynomials in total. We define $P_{C,\vec{j}}$ inductively over the structure of C :

1. If the top gate of C is a NOT gate, then $C = \neg C'$ for some circuit C' . We define $P_{C,\vec{j}} = 1 - P_{C',\vec{j}}$.
2. If the top gate of C is a MOD_p gate, then $C = \text{MOD}_p(C_1, C_2, \dots, C_t)$ for some $t \in \mathbb{N}$ and sub-circuits $C_1, C_2, \dots, C_t$. We define

$$P_{C,\vec{j}} = \left(\sum_{i=1}^t P_{C_i,\vec{j}} \right)^{p-1}.$$

The correctness of $P_{C,\vec{j}}$ follows from Fermat's little theorem.

3. If the top gate of C is an OR gate, then $C = \text{OR}(C_1, C_2, \dots, C_t)$ for some $t \in \mathbb{N}$ and sub-circuits $C_1, C_2, \dots, C_t$. The idea here is to approximate the OR function using the j_1 -th, $\dots$, j_Δ -th outputs of Enc : let $H_C(x) = (C_1(x), \dots, C_t(x))$, we “think” that $C(x)$ outputs 1 if there exists some $i \in [\Delta]$ such that $\text{Enc}(H_C(x))_{j_i} \neq 0 \pmod{p}$.

In particular, let $\text{OR} : \mathbb{F}_p^\Delta \rightarrow \mathbb{F}_p$ be any degree- Δ polynomial over $\mathbb{F}_p$ that computes the OR function over Δ Boolean inputs; for example, one could take $\text{OR}(x_1, \dots, x_\Delta) = 1 - \prod_{i=1}^\Delta (1 - x_i)$.

¹³As our formalisation in Section 9.2 will use the weak rank principles over matrices of dimension $\text{poly}(L, 2^n)$, it is crucial that L is not too large.

We define

$$P_{C,\vec{j}} = \text{OR}\left((\text{Enc}(H_C(x))_{j_1})^{p-1}, (\text{Enc}(H_C(x))_{j_2})^{p-1}, \dots, (\text{Enc}(H_C(x))_{j_\Delta})^{p-1}\right).$$

It can be seen from the above inductive proof that the degree of $P_{C,\vec{j}}$ is at most $(p\Delta)^d$.

Correctness of $P_{C,\vec{j}}$. Next we show that for a random $\vec{j}$, there are only a fraction of $0.9^\Delta \cdot s$ inputs $x \in \{0, 1\}^n$ such that $P_{C,\vec{j}}(x) \neq C(x)$. In fact, let **Bad** denote the set of $(x, \vec{j})$ such that $P_{C,\vec{j}}(x) \neq C(x)$, and let $K = 2^n L \cdot 0.9^\Delta \cdot s$, then $|\text{Bad}| \leq K$. Moreover, this inequality can be “feasibly witnessed” by a pair of functions

$$\phi : \text{Bad} \rightarrow [K] \quad \text{and} \quad \psi : [K] \rightarrow \text{Bad}$$

such that for every $(x, \vec{j}) \in \text{Bad}$, $\psi(\phi(x, \vec{j})) = (x, \vec{j})$. (Clearly, the existence of such functions (ϕ, ψ) implies that $|\text{Bad}| \leq K$. The complexity of computing ϕ and ψ will be crucial our formalisation and will be discussed in [Section 9.2.2](#).)

We identify $[K]$ with $\{0, 1\}^n \times [s] \times [0.9^\Delta \cdot L]$. That is, each element in $[K]$ is treated as a tuple (x, g, c) where $x \in \{0, 1\}^n$ is an input, $g \in [s]$ is the index of some OR gate in C , and $c \in [0.9^\Delta \cdot L]$ is the “compressed” representation of some $\vec{j} \in [L] = [\ell]^\Delta$ that introduced an error on the g -th gate of C when computing $C(x)$.

Now, let $(x, \vec{j}) \in \text{Bad}$. Let g be the bottommost gate of C that introduced an error; that is, if we denote by C^g the sub-circuit computed by g , then g is the bottommost gate such that $C^g(x) \neq P_{C^g, \vec{j}}(x)$. Suppose that C^g is the OR of sub-circuits $C_1, C_2, \dots, C_t$. Then we have:

- $H_g(x) = (C_1(x), C_2(x), \dots, C_t(x))$ is not the all-zero vector, but for every $k \in [\Delta]$ we have $\text{Enc}(H_g(x))_{j_k} = 0$.

For every fixed x and g , the probability over a random $j \leftarrow [\ell]$ that $H_g(x) \neq \vec{0}$ but $\text{Enc}(H_g(x))_j = 0$ is at most 0.9; hence the probability over a random $\vec{j}$ that the above bullet happens is at most 0.9^Δ . It follows that given x and g , every $\vec{j}$ that makes the above bullet happens can be encoded by a number $c \in [0.9^\Delta \cdot L]$. We then define $\phi(x, \vec{j}) = (x, g, c)$. It is easy to define the inverse function ψ such that $\psi(x, g, c) = (x, \vec{j})$ so that for every $(x, \vec{j}) \in \text{Bad}$, we have that $\psi(\phi(x, \vec{j})) = (x, \vec{j})$.

Basis change. A minor but annoying issue is that [Theorem 9.2](#) deals with low-degree polynomials $P : \{-1, 1\}^n \rightarrow \{-1, 1\}$ but our $P_{C,\vec{j}}$ are polynomials from $\{0, 1\}^n$ to $\{0, 1\}$. Observe that the linear map $\tau(x) = 1 - 2x$ maps 0 to 1 and maps 1 to -1 , and its inverse map is $\tau^{-1}(x) = \frac{x-1}{2}$. We define

$$Q_{C,\vec{j}}(x_1, \dots, x_n) = \tau(P_{C,\vec{j}}(\tau^{-1}(x_1), \dots, \tau^{-1}(x_n))).$$

Clearly, the degree of $Q_{C,\vec{j}}$ is the same as the degree of $P_{C,\vec{j}}$. Moreover, for every $x \in \{0, 1\}^n$, we have that $P_{C,\vec{j}}(x) = \text{MOD}_2(x)$ if and only if $Q_{C,\vec{j}}(\tau(x)) = \prod_{i=1}^n \tau(x_i)$. (Hereafter, for $x \in \mathbb{F}_p^n$, $\tau(x)$ is a vector in $\mathbb{F}_p^n$ obtained by applying τ to each coordinate of x individually.)

Proof of the main lower bound. Now we are ready to prove the main lower bound against $\text{AC}^0[p]$ circuits. Let $C : \{0, 1\}^n \rightarrow \{0, 1\}$ be an $\text{AC}^0[p]$ circuit that purportedly computes MOD_2 . A naïve attempt is to find a good $\vec{j}$, construct the matrices $A \in \mathbb{F}_p^{2^n \times 0.99 \cdot 2^n}$ and $B = B(Q_{C,\vec{j}}) \in \mathbb{F}_p^{0.99 \cdot 2^n \times 2^n}$ (see [Definition 9.5](#)), use the weak rank principle to find inputs $z, x \in \{-1, 1\}^n$ such that $(AB)[z, x] \neq \delta_{z,x}$, and conclude that $C(\tau^{-1}(z)) \neq \text{MOD}_2(\tau^{-1}(z))$. However, it may be the case that $(\tau^{-1}(z), \vec{j}) \in \text{Bad}$, which does not give us the desired conclusion.

To remedy this, we put the matrices A and $B(Q_{C,\vec{j}})$ together for every possible $\vec{j}$ (there are L many possibilities) in a block-diagonal manner. This results in a giant matrix multiplication of dimension $(2^n L) \times (0.99 \cdot 2^n L) \times (2^n L)$. We still need to handle the entries in **Bad**. However, since $|\mathbf{Bad}| \ll 0.01 \cdot 2^n L$, we can afford to increase the middle dimension (i.e., $0.99 \cdot 2^n L$) by $|\mathbf{Bad}|$ to “correct” all the values related to entries in **Bad**.

Theorem 9.6. *Let C be an $\text{AC}^0[p]$ circuit of depth d and size $s < 2^{n^{1/2d}/(20p)}$, then there exists an input $x \in \{0, 1\}^n$ such that $C(x) \neq \text{MOD}_2(x)$.*

Proof. We set the following parameters:

Value	Intended meaning
$\Delta = \frac{n^{1/2d}}{10p}$	Number of “samples” taken in $\mathcal{S}$. Equiv. the degree of the “polynomial approximation” of OR
$\ell = \text{poly}(s)$	Output length of Enc
$L = \ell^\Delta \leq \exp(n)$	Number of polynomials in $\mathcal{P}$
$N = 2^n L$	Dimension of the weak rank principle
$K = 2^n L \cdot 0.6^\Delta \cdot s \leq 0.001N$	$ \mathbf{Bad} $
$R = 0.99 \cdot 2^n L + K \leq 0.991N$	Rank parameter of the weak rank principle

We construct the following two matrices $\hat{A}$ and $\hat{B}$ where $\hat{A} \in \mathbb{F}_p^{N \times R}$ and $\hat{B} \in \mathbb{F}_p^{R \times N}$. It would be convenient to identify $[N]$ with $\{-1, 1\}^n \times [L]$, so every element in $[N]$ can be parsed as $(x, \vec{j})$ where $x \in \{-1, 1\}^n$ and $\vec{j} \in [\ell]^\Delta$. Similarly, we identify $[R]$ with the disjoint union of $[0.99 \cdot 2^n] \times [\ell]^\Delta$ and $[K]$. Recall from the proof of [Theorem 9.2](#) that we also identified $[0.99 \cdot 2^n]$ with monomials of degree at most $d = n/2 + 0.1\sqrt{n}$.

Let $(z, \vec{j}) \in [N]$ and $a \in [R]$, we define $\hat{A}[(z, \vec{j}), a]$ as follows:

- Suppose $a = (m, \vec{j}') \in [0.99 \cdot 2^n] \times [\ell]^\Delta$. If $\vec{j} \neq \vec{j}'$ then $\hat{A}[(z, \vec{j}), a] = 0$; otherwise $\hat{A}[(z, \vec{j}), a] = m(z)$.
- Suppose $a \in [K]$. If $(\tau^{-1}(z), \vec{j}) \in \mathbf{Bad}$ and $a = \phi(\tau^{-1}(z), \vec{j})$ then $\hat{A}[(z, \vec{j}), a] = 1$; otherwise $\hat{A}[(z, \vec{j}), a] = 0$.

Similarly, let $a \in [R]$ and $(x, \vec{j}) \in [N]$, we define $\hat{B}[a, (x, \vec{j})]$ as follows:

- Suppose $a = (m, \vec{j}') \in [0.99 \cdot 2^n] \times [\ell]^\Delta$. If $\vec{j} = \vec{j}'$ then $\hat{B}[a, (x, \vec{j})] = B(Q_{C,\vec{j}})[a, x]$ (recall the definition of $B(Q_{C,\vec{j}})$ in [Definition 9.5](#)); otherwise $\hat{B}[a, (x, \vec{j})] = 0$.
- Suppose $a \in [K]$. Let $(\tau^{-1}(z), \vec{j}') = \psi(a)$, we define

$$\hat{B}[a, (x, \vec{j})] = \delta_{(z, \vec{j}'), (x, \vec{j})} - \sum_{m \in [0.99 \cdot 2^n]} \hat{A}[(z, \vec{j}'), (m, \vec{j}')] \cdot \hat{B}[(m, \vec{j}'), (x, \vec{j})].$$

(The intention is to set $\hat{B}[a, (x, \vec{j})]$ so that $(\hat{A} \cdot \hat{B})[(z, \vec{j}'), (x, \vec{j})] = \delta_{(z, \vec{j}'), (x, \vec{j})}$.)

From the above definition it is clear that for every $(\tau^{-1}(z), \vec{j}') \in \mathbf{Bad}$ and every $(x, \vec{j}) \in [N]$, we have $(\hat{A} \hat{B})[(z, \vec{j}'), (x, \vec{j})] = \delta_{(z, \vec{j}'), (x, \vec{j})}$. On the other hand, the weak rank principle implies that there exists $(z, \vec{j}'), (x, \vec{j}) \in [N]$ such that $(\hat{A} \hat{B})[(z, \vec{j}'), (x, \vec{j})] \neq \delta_{(z, \vec{j}'), (x, \vec{j})}$. It is easy to see that $\vec{j} = \vec{j}'$ as otherwise $(\hat{A} \hat{B})[(z, \vec{j}'), (x, \vec{j})] = 0 = \delta_{(z, \vec{j}'), (x, \vec{j})}$. Then we can see that $(A \cdot B(Q_{C,\vec{j}}))[z, x] \neq \delta_{z,x}$. It follows from [Claim 9.4](#) that $Q_{C,\vec{j}}(z) \neq \prod_{i=1}^n z_i$, hence $P_{C,\vec{j}}(\tau^{-1}(z)) \neq \text{MOD}_2(\tau^{-1}(z))$. However, since $(\tau^{-1}(z), \vec{j}') \notin \mathbf{Bad}$, we have that $P_{C,\vec{j}}(\tau^{-1}(z)) = C(\tau^{-1}(z))$. Hence $C(\tau^{-1}(z)) \neq \text{MOD}_2(\tau^{-1}(z))$ and C fails to compute the MOD_2 function. $\square$

9.2 Formalisation in $\overline{\mathbf{V}^0(p)} + \text{WRank}_p$

9.2.1 The Theory $\overline{\mathbf{V}^0(p)} + \text{WRank}_p$

We start with a brief introduction to the theories $\mathbf{V}^0(p)$, $\overline{\mathbf{V}^0(p)}$, and $\overline{\mathbf{V}^0(p)} + \text{WRank}_p$. A more comprehensive treatment for the theories $\mathbf{V}^0(p)$ and $\overline{\mathbf{V}^0(p)}$ can be found in Cook and Nguyen’s textbook [20]. The theory $\overline{\mathbf{V}^0(p)} + \text{WRank}_p$ is introduced and studied for the first time in this work.

We consider *two-sorted* logic: the first sort are natural numbers (i.e., indices) and the second sort are bit strings (or more precisely, finite sets of natural numbers whose characteristic vectors are bit strings). In this context, we usually use capital letters (X, Y, Z) to denote strings and lower case letters (x, y, z) to denote numbers; however our notation does not strictly follow this convention and we list the exceptions in [Remark 9.7](#).

The vocabulary for two-sorted logic is $\mathcal{L}_A^2$, which contains symbols

$$0, 1, +, \cdot, | \cdot |; =_1, =_2, \leq, \in.$$

Here, $=_1$ and $=_2$ denotes equality over numbers and strings respectively; it is usually clear in the context whether numbers or strings are being compared, so we often omit the subscripts 1 and 2 and simply write $=$ for convenience. We will also use $X(t)$ or X_t to denote “ $t \in X$ ”. These symbols are governed by a set of “basic axioms” called **2-BASIC** (see [20, Figure 2]).

Let Φ be a set of formulas, then we can define the following *axiom schema*:

- **Φ -COMP** (comprehension): $\exists X \leq y \forall z < y (X(z) \leftrightarrow \varphi(z))$, where $\varphi(z)$ is any formula in Φ and X does not occur free in $\varphi(z)$.
- **Φ -IND** (number induction): $(\varphi(0) \wedge \forall x (\varphi(x) \rightarrow \varphi(x+1))) \rightarrow \forall z \varphi(z)$, where $\varphi(z)$ is any formula in Φ .
- **Φ -MIN** (number minimisation): $\varphi(y) \implies (\exists x \leq y (\varphi(x) \wedge \neg \exists z < x \varphi(z)))$, where $\varphi(z)$ is any formula in Φ .

Let $\mathcal{L}$ be a two-sorted vocabulary that contains $\mathcal{L}_A^2$, then $\Sigma_0^B(\mathcal{L})$ is the set of $\mathcal{L}$ -formulas whose only quantifiers are bounded number quantifiers (there can be free string variables).

Let m be a constant,¹⁴ the theory $\mathbf{V}^0(m)$ has vocabulary $\mathcal{L}_A^2$ and is axiomatised by **2-BASIC**, $\Sigma_0^B(\mathcal{L}_A^2)$ -**COMP**, and an axiom defining the function MOD_m ; see [20, Section IX.4.6] for details.

We actually work with the theory $\overline{\mathbf{V}^0(m)}$ instead of $\mathbf{V}^0(m)$. This theory has vocabulary $\mathcal{L}_{\text{FAC}^0(m)}$ that contains a symbol for every function in $\text{FAC}^0(m)$ ($\text{FAC}^0(m)$ is the class of (possibly non-Boolean) *functions* computable in $\text{AC}^0(m)$). It proves $\Sigma_0^B(\mathcal{L}_{\text{FAC}^0(m)})$ -**COMP**, $\Sigma_0^B(\mathcal{L}_{\text{FAC}^0(m)})$ -**IND**, and $\Sigma_0^B(\mathcal{L}_{\text{FAC}^0(m)})$ -**MIN**: for example, for every $\text{FAC}^0(m)$ function f that maps a string X to the string $f(X)$, we can use the number minimisation axiom inside $\overline{\mathbf{V}^0(m)}$ to conclude that if $f(X)$ is not the all-zero string, then there exists the smallest i such that $f(X)_i = 1$. Again, we refer to [20, Section IX.4.6] for more details.

Although $\overline{\mathbf{V}^0(m)}$ “contains much more features” than $\mathbf{V}^0(m)$, it is a conservative extension of $\mathbf{V}^0(m)$, i.e., every formula in the vocabulary $\mathcal{L}_A^2$ that can be proved in $\overline{\mathbf{V}^0(m)}$ can be proved in $\mathbf{V}^0(m)$ as well. Hence it is without loss of generality to work with $\overline{\mathbf{V}^0(m)}$ instead of $\mathbf{V}^0(m)$.

Finally, let $m = p$ be a prime. The *weak rank principle* over $\mathbb{F}_p$, denoted as WRank_p , is the following axiom. Let n be a number, $A \in \mathbb{F}_p^{2n \times n}$ and $B \in \mathbb{F}_p^{n \times 2n}$ be matrices coded as strings, then there exists $x, y \in [2n]$ such that

$$\sum_{z \in [n]} a_{x,z} b_{z,y} \not\equiv \delta_{x,y} \pmod{p}.$$

¹⁴Here we mainly consider the case that $m = p$ is a prime, but it also makes perfect sense to consider the theory $\mathbf{V}^0(m)$ when m is a composite number.

(Note that $\sum_{z \in [n]} a_{x,z} b_{z,y} \bmod p$ can be expressed as a function in $\text{FAC}^0[p]$, hence WRank_p is indeed a universal sentence in the language $\mathcal{L}_{\text{FAC}^0(p)}$.) For every constant $\epsilon > 0$, one can use a similar argument as [Theorem 7.3](#) to show that $\overline{\mathbf{V}^0(p)} + \text{WRank}_p$ proves the following: Let n be a number, $A \in \mathbb{F}_p^{(1+\epsilon)n \times n}$ and $B \in \mathbb{F}_p^{n \times (1+\epsilon)n}$ be two matrices encoded as strings, then $AB \neq I_{(1+\epsilon)n}$. In other words, “ $\text{rank}(I_{2n}) > n$ ” and “ $\text{rank}(I_{(1+\epsilon)n}) > n$ ” are equivalent, and replacing the constant 2 in the definition of WRank_p with other constants greater than 1 does not affect the strength of the theory $\overline{\mathbf{V}^0(p)} + \text{WRank}_p$. In fact, our formalisation will use “ $\text{rank}(I_{(1+\epsilon)n}) > n$ ” for some small constant $\epsilon > 0$.

9.2.2 Smolensky’s Lower Bound in $\overline{\mathbf{V}^0(p)} + \text{WRank}_p$

Fix a constant $d \in \mathbb{N}$ (in the meta-theory). Let $n \in \text{Log}$ (that is, we assume 2^n is a “number”), $s = 2^{n^{1/2d}/(20p)}$, and C be a string that codes an $\text{AC}^0[p]$ circuit of depth d and size s . Let $\text{Eval}_d(C, x)$ denote the circuit evaluation function on depth- d $\text{AC}^0[p]$ circuits, where $x \in [2^n]$ is a number that denotes the input of C . Since Eval_d is in $\text{AC}^0[p]$, it is in the language of $\overline{\mathbf{V}^0(p)}$. Also, since x is a number, we can talk about $\text{MOD}_2(x)$ in $\overline{\mathbf{V}^0(p)}$. Our goal is to prove that

$$\exists x \text{ Eval}_d(C, x) \neq \text{MOD}_2(x).$$

Remark 9.7. It should be clear from the context whether a variable is of the number sort or the string sort. Most of the time we use capital letters (X, Y, Z) to denote strings and use lower-case letters (x, y, z) to denote numbers. However, one exception is that L, N, K, R (in the proof of [Theorem 9.6](#)) are numbers.

The matrix $B(P)$ ([Definition 9.5](#)). It is clear from the definition that there is an $\text{FAC}^0[p]$ function that given (m, x) and the list of coefficients in P , outputs the value of $B(P)[m, x]$.

Claim 9.8 ($\overline{\mathbf{V}^0(p)} \vdash \cdot$). *Let $z, x \in \{-1, 1\}^n$ such that $(A \cdot B(P))[z, x] \neq \delta_{z,x}$. Then $P(z) \neq \prod_{i=1}^n z_i$.*

Proof Sketch. The proofs of both [Claim 9.3](#) and [Claim 9.4](#) can be carried out in $\overline{\mathbf{V}^0(p)}$. $\square$

The polynomials $P_{C, \vec{j}}$ and $Q_{C, \vec{j}}$. For each gate g , let C^g denote the sub-circuit computed by g , we represent the polynomials $P_{C^g, \vec{j}}$ and $Q_{C^g, \vec{j}}$ by their list of coefficients. There is an $\text{FAC}^0[p]$ circuit that given the description of C , an index g , and a list of vectors $\vec{j}$, outputs the coefficient list of $P_{C^g, \vec{j}}$. (Note that since $\vec{j} \in [\ell]^\Delta$ and $\ell^\Delta = L$ is a number in the theory, $\vec{j}$ can be represented as a single number.) The coefficient list of $Q_{C, \vec{j}}(x) = P_{C, \vec{j}}(x - \vec{1}) + 1$ can be computed in $\text{FAC}^0[p]$ as well. If C^g is an $\text{AC}^0[p]$ circuit of depth d , then the degree of $P_{C^g, \vec{j}}$ is at most $(p\Delta)^d$ (and we only consider monomials up to this degree in our representation).

A good error-correcting code. Before formalising the low-degree approximations of $\text{AC}^0[p]$ circuits, we need the following $\mathbb{F}_p$ -linear error-correcting code whose correctness can be proved feasibly. In fact, we need to feasibly prove the following formalisation of “every non-zero codeword of Enc has constant relative Hamming weight”:

Theorem 9.9. *Suppose that $\log^2 s \in \text{Log}$. For some $\ell \leq \text{poly}(s)$, there is an $\mathbb{F}_p$ -linear code $\text{Enc} : \mathbb{F}_p^s \rightarrow \mathbb{F}_p^\ell$ and functions $\tilde{\phi}(X, i_1, i)$, $\tilde{\psi}(X, i_1, c)$ computable in $\text{FAC}^0[p]$ such that $\overline{\mathbf{V}^0(p)}$ proves the following. For every string $X \in \mathbb{F}_p^s$, every numbers $i_1 \in [s]$ and $i \in [\ell]$, if $X_{i_1} \neq 0$ but $\text{Enc}(X)_i = 0$, then:*

- $\tilde{\phi}(X, i_1, i)$ returns a number in $[0.9\ell]$, and
- $\tilde{\psi}(X, i_1, \tilde{\phi}(X, i_1, i)) = i$.

Note that the above statement implies that the relative Hamming weight of $\text{Enc}(X)$ is at least 0.1. Roughly speaking, the error-correcting code is the concatenation of Reed–Muller code and Hadamard code; the proof proceeds by formalising the local (unique-)decoding algorithm of the Reed–Muller code (see e.g., [8, Section 19.4.2]). We postpone the proof to [Section 9.2.3](#).

The functions ϕ, ψ . We need to show that ϕ and ψ are computable in $\text{FAC}^0[p]$. Given a string C representing our $\text{AC}^0[p]$ circuit and two numbers $x, \vec{j}$, we can compute $\phi(x, \vec{j})$ as follows:

1. First, find the bottommost gate g for which $C^g(x) \neq P_{C^g, \vec{j}}(x)$.

For every g , we can compute the indicator of whether $C^g(x) = P_{C^g, \vec{j}}(x)$ in $\text{AC}^0[p]$. Then we use an FAC^0 circuit to compute the number g , which is the index of the bottommost gate on which the indicator is False. (If there is no such gate then we set $g = \perp$.)

2. Let $H_g(x) = (C_1(x), \dots, C_t(x))$ where $C_1, \dots, C_t$ are input gates of g . For each $k \in [\Delta]$, we have that $\text{Enc}(H_g(x))_{j_k} = 0$. However, there exists an index $j' \in [t]$ such that $C_{j'}(x) = 1$, and the smallest such index can be computed in FAC^0 given $H_g(x)$. Hence we compute $c_k = \tilde{\phi}(H_g(x), j', j_k) \in [0.9\ell]$ using [Theorem 9.9](#). This gives us a vector $\vec{c} = (c_1, \dots, c_\Delta)$. (Note that $\ell^\Delta = L$ is a number in the theory, hence $\vec{c}$ can be represented as a single number.) We output $\phi(x, \vec{j}) = (x, g, \vec{c})$.

Similarly, one can compute ψ in $\text{FAC}^0[p]$ by invoking the function $\tilde{\psi}$ in [Theorem 9.9](#). Namely, given a string C representing our $\text{AC}^0[p]$ circuit and three numbers $x, g, \vec{c}$, we can compute $\psi(x, g, \vec{c})$ as follows. First, let $C_1, \dots, C_t$ be the input gates of g , we can compute $H_g(x) = (C_1(x), \dots, C_t(x))$ in $\text{FAC}^0[p]$. We can also compute the smallest index j' such that $C_{j'}(x) = 1$ in FAC^0 given $H_g(x)$. Then, for each $K \in [\Delta]$, we let $j_k = \tilde{\psi}(H_g(x), j', c_k)$. We return $\psi(x, g, \vec{c}) = (x, \vec{j})$.

Claim 9.10 ($\overline{\mathbf{V}^0(p)} \vdash \cdot$). *Suppose that $\log^2 s \in \text{Log}$. For every $x \in \{0, 1\}^n$ and $\vec{j} \in [\ell]^\Delta$, if $P_{C, \vec{j}}(x) \neq C(x)$, then $\psi(\phi(x, \vec{j})) = (x, \vec{j})$.*

Proof Sketch. Since $P_{C, \vec{j}}(x) \neq C(x)$, it follows from $\Sigma_0^B(\mathcal{L}_{\text{FAC}^0(p)})\text{-MIN}$ that there exists a bottommost gate g for which $C^g(x) \neq P_{C^g, \vec{j}}(x)$. It follows that $H_g(x)_{j'} \neq 0$ but $\text{Enc}(H_g(x))_{j_k} = 0$ for every $k \in [\Delta]$. Let $c_k = \tilde{\phi}(H_g(x), j', j_k)$, then $\phi(x, \vec{j}) = (x, g, \vec{c})$. Let $\psi(x, g, \vec{c}) = (x, \vec{j}')$, then for every $k \in [\Delta]$,

$$j'_k = \tilde{\psi}(H_g(x), j', c_k) = \tilde{\psi}(H_g(x), j', \phi(H_g(x), j', j_k)) = j_k$$

where the last equality follows from [Theorem 9.9](#). Hence $\psi(\phi(x, \vec{j})) = (x, \vec{j})$. $\square$

The matrices $\hat{A}, \hat{B}$ as defined in the proof of [Theorem 9.6](#). It is easy to see that these matrices can be computed in $\text{FAC}^0[p]$. That is, there is an $\text{FAC}^0[p]$ function that takes a string C and numbers $z, \vec{j}, a$ as inputs and outputs the value of $\hat{A}[(z, \vec{j}), a]$, and the analogous statement holds for $\hat{B}$.

Claim 9.11 ($\overline{\mathbf{V}^0(p)} \vdash \cdot$). *Let $z, x \in \{-1, 1\}^n$, $\vec{j}, \vec{j}' \in [\ell]^\Delta$, then:*

- (a) *If $\vec{j} \neq \vec{j}'$, then $(\hat{A}\hat{B})[(z, \vec{j}), (x, \vec{j}')] = 0$.*
- (b) *If $\psi(\phi(z - \vec{1}, \vec{j})) = (z - \vec{1}, \vec{j})$, then $(\hat{A}\hat{B})[(z, \vec{j}), (x, \vec{j})] = \delta_{(z, \vec{j}), (x, \vec{j})}$.*
- (c) *If $\psi(\phi(z - \vec{1}, \vec{j})) \neq (z - \vec{1}, \vec{j})$, then $(\hat{A}\hat{B})[(z, \vec{j}), (x, \vec{j})] = (A \cdot B(Q))[z, x]$.*

Proof Sketch. These follow easily from the definitions. $\square$

Putting it together. Fix a constant $d \in \mathbb{N}$ (in the meta-language), we have:

Theorem 9.12 ($\overline{\mathbf{V}^0(p)} + \text{WRank}_p \vdash \cdot$). *Let $n \in \text{Log}$, $s = 2^{n^{1/2d}/(20p)}$, and C be a string that codes an $\text{AC}^0[p]$ circuit of depth d and size s . Then there exists an input $x \in \{0, 1\}^n$ (coded as a number) such that $\text{Eval}_d(C, x) \neq \text{MOD}_2(x)$.*

Proof. Consider the matrices $\hat{A}$ and $\hat{B}$ as defined above. By WRank_p , there exists $(z, \vec{j}'), (x, \vec{j}) \in [N]$ such that $(\hat{A}\hat{B})[(z, \vec{j}'), (x, \vec{j})] \neq \delta_{(z, \vec{j}'), (x, \vec{j})}$. Then:

- By [Claim 9.11 \(item a\)](#), $\vec{j} = \vec{j}'$;
- By [Claim 9.11 \(item b\)](#), $\psi(\phi(z - \vec{1}, \vec{j})) \neq (z - \vec{1}, \vec{j})$;
- By [Claim 9.11 \(item c\)](#), $(AB)[(z, \vec{j}), (x, \vec{j})] = (A \cdot B(Q_{C, \vec{j}}))[z, x] \neq \delta_{z, x}$.
- By [Claim 9.8](#), $Q_{C, \vec{j}}(z) \neq \prod_{i=1}^n z_i$, hence $P_{C, \vec{j}}(z - \vec{1}) \neq \text{MOD}_2(z - \vec{1})$;
- By [Claim 9.10](#) (and note that it is indeed the case that $\log^2 s \in \text{Log}$), $P_{C, \vec{j}}(z - \vec{1}) = C(z - \vec{1})$.
Hence $C(z - \vec{1}) \neq \text{MOD}_2(z - \vec{1})$. $\square$

9.2.3 Proof of [Theorem 9.9](#)

Theorem 9.9. *Suppose that $\log^2 s \in \text{Log}$. For some $\ell \leq \text{poly}(s)$, there is an $\mathbb{F}_p$ -linear code $\text{Enc} : \mathbb{F}_p^s \rightarrow \mathbb{F}_p^\ell$ and functions $\tilde{\phi}(X, i_1, i), \tilde{\psi}(X, i_1, c)$ computable in $\text{FAC}^0[p]$ such that $\overline{\mathbf{V}^0(p)}$ proves the following. For every string $X \in \mathbb{F}_p^s$, every numbers $i_1 \in [s]$ and $i \in [\ell]$, if $X_{i_1} \neq 0$ but $\text{Enc}(X)_i = 0$, then:*

- $\tilde{\phi}(X, i_1, i)$ returns a number in $[0.9\ell]$, and
- $\tilde{\psi}(X, i_1, \tilde{\phi}(X, i_1, i)) = i$.

Proof. Let $n = O(\log s / \log \log s)$, $h = \lceil s^{1/n} \rceil$, q be the smallest power of p that is greater than $10nh$. Note that $q \leq \log^2 s \in \text{Log}$, hence $\overline{\mathbf{V}^0(p)}$ can reason about $\mathbb{F}_q$ by brute force. (For example, $\overline{\mathbf{V}^0(p)}$ can prove the existence of an irreducible degree- $\log_p q$ polynomial over $\mathbb{F}_p$. See e.g., [45, Section 4.3].) Let H be the set of lexicographically first h elements in $\mathbb{F}_q$. We interpret $X \in \mathbb{F}_p^s \subseteq \mathbb{F}_q^s$ as a function $X : H^n \rightarrow \mathbb{F}_q$. Then there is a unique polynomial $\tilde{X} : \mathbb{F}_q^n \rightarrow \mathbb{F}_q$ of individual degree at most h that is consistent with X on H^n . Moreover, one can compute $\tilde{X}(\vec{v})$ using an $\text{FAC}^0[p]$ circuit (in fact, an $\mathbb{F}_q$ -linear transformation) over the string input X and the number input $\vec{v}$.

We now concatenate the above Reed–Muller code with the Hadamard code. Treat $\mathbb{F}_q$ as the $\log_p q$ -dimensional vector space over $\mathbb{F}_p$, then each $\alpha \in \mathbb{F}_q$ is encoded as a vector $\text{Had}(\alpha) \in \mathbb{F}_p^q$ defined as follows. For each $u \in \mathbb{F}_q$, we interpret both α and u as vectors in $\mathbb{F}_p^{\log_p q}$, and the u -th element of $\text{Had}(\alpha)$ is equal to $\text{IP}(\alpha, u) = \sum_{j=1}^{\log_p q} \alpha_j u_j$. Again, since $q \in \text{Log}$, $\overline{\mathbf{V}^0(p)}$ can reason about the Hadamard code using brute force.

Define $\text{Enc}(X)$ be the concatenation of $\text{Had}(\tilde{X}(\vec{v}))$ for each $\vec{v} \in \mathbb{F}_q^n$. That is, the length of $\text{Enc}(X)$ is $\ell = q^{n+1} \leq \text{poly}(s)$. Given an index $i \in [\ell]$, we interpret i as a pair of vectors $(\vec{v} \in \mathbb{F}_q^n, u \in \mathbb{F}_q = \mathbb{F}_p^{\log_p q})$, and define $\text{Enc}(X)_i = \text{IP}(\tilde{X}(\vec{v}), u)$.

The local decoder of the Reed–Muller code requires the following equivalence class over “directions” in $\mathbb{F}_q^n$: two non-zero vectors $\vec{v}_1, \vec{v}_2 \in \mathbb{F}_q^n$ are equivalent if and only if there is a non-zero scalar $\alpha \in \mathbb{F}_q$ such that $\vec{v}_1 = \alpha \vec{v}_2$. Let $\vec{v}$ be a non-zero vector in $\mathbb{F}_q^n$ and define $[\vec{v}]$ as the following vector: let $i \in [n]$ be the first index such that $v_i \neq 0$, and for each $j \in [n]$, the j -th element of $[\vec{v}]$ is equal to $v_i^{-1} v_j$ (over $\mathbb{F}_q$). Again, $\overline{\mathbf{V}^0(p)}$ can prove (by brute force) that v_1 and v_2 are equivalent if and only if $[v_1] = [v_2]$, and that there are $q^{n-1} + q^{n-2} + \dots + q + 1 \leq 2q^{n-1}$ equivalence classes.

Let $X \in \mathbb{F}_p^s$ be a non-zero string, $i_1 \in [s]$ be an index such that $X_{i_1} = 1$, and $i \in [\ell]$ be an index such that $\text{Enc}(X)_i = 0$. We interpret i as a pair of vectors $(\vec{v} \in \mathbb{F}_q^n, u \in \mathbb{F}_q)$. To compute $\hat{\phi}(X, i_1, i)$, we consider the following two cases:

- If $\tilde{X}(\vec{v}) \neq 0$, then we consider the Hadamard code. In particular, u is among the vectors in $\mathbb{F}_q$ for which $\text{IP}(\tilde{X}(\vec{v}), u) = 0$. There are at most q/p such elements, hence we let $\tilde{\phi}(X, i_1, i) = (\text{Hadamard}, \vec{v}, k)$ where u is the k -th smallest element in $\mathbb{F}_q$ such that $\text{IP}(\tilde{X}(\vec{v}), u) = 0$.
- If $\tilde{X}(\vec{v}) = 0$, then we consider the Reed–Muller code. In particular, consider the directional vector $\vec{w} = \vec{v} - i_1$ (where we interpret $i_1 \in [s] \subseteq \mathbb{F}_q^n$ and the subtraction happens in $\mathbb{F}_q^n$). The unique line determined by $\vec{v}$ and i_1 is $\ell(t) = t \cdot [\vec{w}] + i_1$. Composing $\tilde{X}$ with ℓ gives us a univariate polynomial $f(t) = \tilde{X}(t \cdot [\vec{w}] + i_1)$ of degree at most nh . Since $f(0) \neq 0$, this polynomial has at most nh roots. We define $\tilde{\phi}(X, i_1, i) = (\text{Reed–Muller}, [\vec{w}], k, u)$ where $\vec{v}$ corresponds to the k -th root of f . (Recall that k can be found by brute force.)

Now we can see that $\tilde{\phi}(X, i_1, i)$ either returns $(\text{Hadamard}, \vec{v}, k)$ where $\vec{v} \in \mathbb{F}_q^n, k \in [q/p]$ or returns $(\text{Reed–Muller}, [\vec{w}], k, u)$ where $[\vec{w}] \in [2q^{n-1}]$, $k \in [q/10]$, $u \in \mathbb{F}_q$. It follows that the size of $\text{Range}(\tilde{\phi})$ is at most $q^{n+1}/p + q^{n+1}/5 \leq 0.9\ell$.

We can similarly define the function $\tilde{\psi}(X, i_1, c)$:

- If c is of the form $(\text{Hadamard}, \vec{v}, k)$, then we compute the k -th smallest $u \in \mathbb{F}_q$ such that $\text{IP}(\tilde{X}(\vec{v}), u) = 0$ and output $(\vec{v}, u)$.
- If c is of the form $(\text{Reed–Muller}, [\vec{w}], k, u)$, then we compute the k -th root $\vec{v}$ of the degree- nh polynomial $f(t) = \tilde{X}(t \cdot [\vec{w}] + i_1)$ and output $(\vec{v}, \vec{u})$.

It is easy to see that $\psi(X, i_1, \phi(X, i_1, i)) = i$ for every i such that $\text{Enc}(X)_i = 0$. □

A Expander Properties

The proofs here mostly mimic the corresponding proofs from [76].

Lemma 4.17. *Let $C > 1$. There exist constants $\delta, d > 0$ such that for every large enough positive integers m, n with $n < m < Cn$, there is an $m \times n$ bipartite graph G satisfying the following:*

1. *the minimum left-degree of G is at least $\Delta_{\min} = 50 \log n$,*
2. *the maximum left-degree of G is at most $\Delta_{\max} = 150 \log n$,*
3. *G is an (n^δ, d) -lossless expander.*

In particular, G is an $(n^\delta, \Delta_{\min} - d)$ -boundary expander.

Proof. The proof essentially follows [76] but with a slightly different choice of parameters. We construct a bipartite graph G with parts $[m]$ and $[n]$. For every $i \in [m]$ and $k \in [n]$, we add an edge between i and k with probability $\frac{100 \log n}{n}$. Let $M_G = (a_{i,j})_{i \in [m], k \in [n]}$ be the adjacency matrix of such graph.

For each $i \in [m]$, the left-degree of i is the sum of n independent variables $a_{i,k} : k \in [n]$, with expectation $\mu = 100 \log n$. By [Theorem 4.1](#), applied with $\delta = 1/2$,

$$\Pr \left[\left| \sum_{k \in [n]} a_{i,k} - 100 \log n \right| > 50 \log n \right] \leq 2e^{-100 \log n / 12} = o(1/n).$$

By the union bound, these events happen simultaneously for all $i \in [m]$ with probability $1 - o(1)$.

Now let $r = n^\delta$.

Claim A.1. *If property (7) is violated for some set $I \subseteq [m]$ of size $l \leq r$, then there exist a set $K \subseteq [n]$ of size $ld/2$ such that the rectangle $I \times K$ contains at least ld ones in M_G .*

Proof. Let $K = \bigcup_{i \in I} \mathcal{N}_G(i) \setminus \partial_G(I)$, i.e., it contains all non-unique neighbours of I . In particular, every column of the rectangle $I \times K$ contains at least two ones. Thus, if $|K| \geq ld/2$, we can take any $K' \subseteq K$ of size exactly $ld/2$. Now assume that $|K| < ld/2$. The number of ones in $I \times K$ is at least

$$\sum_{i \in I} |\mathcal{N}_G(i) \setminus \partial_G(I)| = \sum_{i \in I} |\mathcal{N}_G(i)| - |\partial_G(I)|,$$

since each $k \in \partial_G(I)$ is contained in exactly one $\mathcal{N}_G(i)$. Thus, since (7) is violated for I , this is further lower bounded by $d|I| = ld$. To conclude the proof, we can arbitrarily enlarge K to have size exactly $ld/2$. $\diamond$

Now we will bound the probability of the existence of such $I \times K$. For every $\ell \in [r]$, there is at most $(Cn)^\ell$ choices of I , $n^{\ell d/2}$ choices of K , and $(\ell d)^{2\ell d}$ choices of the positions of ones in $I \times K$. Thus, the probability that there exists such a rectangle $I \times K$ with $|I| = \ell \leq r = n^\delta$ is at most

$$(Cn)^\ell n^{\ell d/2} (\ell d)^{2\ell d} \left(\frac{100 \log n}{n} \right)^{\ell d} \leq n^{2\ell} \left(\frac{n^{4\delta} d^4 (100 \log n)^2}{n} \right)^{\ell d/2} \leq n^{-2\ell}$$

for small enough δ and large enough n and d . Observing that $\sum_{\ell=1}^r n^{-2\ell} = o(1)$, the proof follows. $\square$

Lemma 4.18. *Let $m = O(n)$ and let G be an $m \times n$ (n^δ, d) -lossless expander with parts L and R , minimum left-degree at least $\Delta_{\min} = 50 \log n$ and maximum left-degree at most $\Delta_{\max} = 150 \log n$. Let K be a random subset of R such that every $j \in R$ is added to K independently with probability $2/3$. With probability $1 - o(1)$, the graph $G \setminus K$ is an (n^δ, d) -lossless expander with a minimum left-degree $\Delta_{\min}/6$. In particular, $G \setminus K$ is an (n^δ, c) -boundary expander, where $c = \Delta_{\min}/6 - d$.*

Proof. Let $G' = G \setminus K$. For every $i \in L$, let $M_i \subseteq \mathcal{N}_G(i)$ be any fixed subset of size exactly $\Delta_{\min}$. Hence $|M_i \setminus K|$ is the sum of $\Delta_{\min}$ independent indicator variables, each of which is 1 with probability $1/3$. The expectation of this sum is $\mu = \mathbb{E}[|M_i \setminus K|] = \Delta_{\min}/3$. By Theorem 4.1, applied with $\delta = 1/2$,

$$\Pr[|M_i \setminus K| < \Delta_{\min}/6] \leq e^{-\Delta_{\min}/24} = o(1/n),$$

since $\Delta_{\min} = 50 \log n$.

By the union bound, the left-degree of G' is at least $\Delta_{\min}/6$ with probability $1 - o(1)$.

The remaining part of the proof repeats the argument from [76]. Let $I \subseteq L$ with $|I| \leq n^\delta$. For every $k \in R$, let $\gamma_k = |\{i \in I : k \in \mathcal{N}_G(i)\}|$. Using the bound from (7), it is easy to see that

$$\sum_{\substack{k \in R \\ \gamma_k \geq 2}} \gamma_k = \sum_{i \in I} |\mathcal{N}_G(i)| - |\partial_G(I)| \leq d|I|.$$

Now let $\gamma'_k = |\{i \in I : k \in \mathcal{N}_{G'}(i)\}|$. Clearly, $\gamma'_k \leq \gamma_k$ for all k , since $\gamma'_k = 0$ for all $k \in K$. This implies that

$$\sum_{i \in I} |\mathcal{N}_{G'}(i)| - |\partial_{G'}(I)| = \sum_{\substack{k \in R \\ \gamma'_k \geq 2}} \gamma'_k \leq \sum_{\substack{k \in R \\ \gamma_k \geq 2}} \gamma_k \leq d|I|. \quad \square$$

Acknowledgment

We thank the anonymous reviewers for helpful comments that improved the exposition of this work.

AI usage notice. AI tools played no role in conducting the research, developing or drafting the proofs, or producing the scientific content in this work. Their use was limited to basic copy-editing and assistance in preparing some figures.

References

- [1] Miklós Ajtai. Parity and the pigeonhole principle. In Samuel R. Buss and Philip J. Scott, editors, *Feasible Mathematics: A Mathematical Sciences Institute Workshop, Ithaca, New York, June 1989*, pages 1–24. Birkhäuser Boston, Boston, MA, 1990. [22](#)
- [2] Miklós Ajtai. The complexity of the pigeonhole principle. *Comb.*, 14(4):417–433, 1994. [3](#)
- [3] Michael Alekhnovich, Eli Ben-Sasson, Alexander A. Razborov, and Avi Wigderson. Space complexity in propositional calculus. *pppSIAM J. Comput.*, 31(4):1184–1211, 2002. [19](#)
- [4] Michael Alekhnovich, Eli Ben-Sasson, Alexander A. Razborov, and Avi Wigderson. Pseudorandom generators in propositional proof complexity. *SIAM J. Comput.*, 34(1):67–88, 2004. [4](#), [7](#), [8](#), [9](#), [13](#), [16](#), [37](#)
- [5] Michael Alekhnovich and Alexander A. Razborov. Lower bounds for polynomial calculus: Non-binomial case. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, Las Vegas, Nevada, USA, October 14-17, 2001*, pages 190–199. IEEE Computer Society, 2001. [7](#)
- [6] Robert Andrews. On matrix multiplication and polynomial identity testing. *SIAM J. Comput.*, 54(5):S22–1, 2025. [24](#)
- [7] Robert Andrews and Michael A. Forbes. Ideals, determinants, and straightening: proving and using lower bounds for polynomial ideals. In Stefano Leonardi and Anupam Gupta, editors, *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, pages 389–402. ACM, 2022. [5](#), [24](#)
- [8] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, New York, NY, USA, 2009. [90](#), [98](#)
- [9] Albert Atserias and Tuomas Hakoniemi. Size-degree trade-offs for sums-of-squares and Positivstellensatz proofs. In *34th Computational Complexity Conference, CCC 2019, July 18-20, 2019, New Brunswick, NJ, USA*, pages 24:1–24:20, 2019. [8](#), [37](#)
- [10] Per Austrin and Kilian Risse. Sum-of-squares lower bounds for the minimum circuit size problem. In Amnon Ta-Shma, editor, *38th Computational Complexity Conference, CCC 2023, July 17-20, 2023, Warwick, UK*, volume 264 of *LIPICs*, pages 31:1–31:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. [14](#)
- [11] László Babai and Péter Frankl. *Linear algebra methods in combinatorics*. Department of Computer Science, Univ. of Chicago, 1992. [15](#)
- [12] Paul Beame, Stephen A. Cook, Jeff Edmonds, Russell Impagliazzo, and Toniann Pitassi. The relative complexity of NP search problems. *J. Comput. Syst. Sci.*, 57(1):3–19, 1998. [9](#)
- [13] 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. [18](#)
- [14] Eli Ben-Sasson and Russell Impagliazzo. Random cnf’s are hard for the polynomial calculus. *Comput. Complex.*, 19(4):501–519, 2010. [13](#)
- [15] Maria Luisa Bonet, Samuel R. Buss, and Toniann Pitassi. Are there hard examples for Frege systems? In *Feasible mathematics, II (Ithaca, NY, 1992)*, volume 13 of *Progr. Comput. Sci. Appl. Logic*, pages 30–56. Birkhäuser Boston, Boston, MA, 1995. [4](#), [6](#), [7](#)
- [16] Samuel R. Buss. Polynomial size proofs of the propositional pigeonhole principle. *J. Symb. Log.*, 52(4):916–927, 1987. [3](#)
- [17] Samuel R. Buss, Dima Grigoriev, Russell Impagliazzo, and Toniann Pitassi. Linear gaps between degrees for the polynomial calculus modulo distinct primes. *J. Comput. Syst. Sci.*, 62(2):267–289, 2001. [8](#), [20](#), [37](#)

- [18] Lijie Chen, Jiatu Li, and Igor C. Oliveira. Reverse mathematics of complexity lower bounds. In *65th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2024, Chicago, IL, USA, October 27-30, 2024*, pages 505–527. IEEE, 2024. [14](#)
- [19] Matthew Clegg, Jeff Edmonds, and Russell Impagliazzo. Using the groebner basis algorithm to find proofs of unsatisfiability. In Gary L. Miller, editor, *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 174–183. ACM, 1996. [8](#), [19](#), [24](#), [37](#)
- [20] Stephen Cook and Phuong Nguyen. *Logical Foundations of Proof Complexity*. ASL Perspectives in Logic. Cambridge University Press, 2010. [3](#), [16](#), [96](#)
- [21] Stephen A. Cook. A taxonomy of problems with fast parallel algorithms. *Inf. Control.*, 64(1-3):2–21, 1985. [15](#)
- [22] Stephen A. Cook and Robert A. Reckhow. Corrections for “On the lengths of proofs in the propositional calculus (preliminary version)”. *SIGACT News*, 6(3):15–22, July 1974. [103](#)
- [23] Stephen A. Cook and Robert A. Reckhow. On the lengths of proofs in the propositional calculus (preliminary version). In *Proceedings of the Annual ACM Symposium on the Theory of Computing 1974*, pages 135–148, 1974. For corrections see Cook-Reckhow [22]. [18](#)
- [24] Ernie Croot, Vsevolod F Lev, and Péter Pál Pach. Progression-free sets in $\mathbb{Z}_4^n$ are exponentially small. *Annals of Mathematics*, pages 331–337, 2017. [15](#)
- [25] Stefan S. Dantchev and Barnaby Martin. Rank complexity gap for lovász-schrijver and sherali-adams proof systems. *Comput. Complex.*, 22(1):191–213, 2013. [13](#), [19](#)
- [26] Stefan S. Dantchev, Barnaby Martin, and Mark Nicholas Charles Rhodes. Tight rank lower bounds for the sherali-adams proof system. *Theor. Comput. Sci.*, 410(21-23):2054–2063, 2009. [14](#), [22](#)
- [27] Susanna F. de Rezende, Mika Göös, Jakob Nordström, Toniann Pitassi, Robert Robere, and Dmitry Sokolov. Automating algebraic proof systems is NP-hard. In *STOC*, pages 209–222. ACM, 2021. [20](#), [21](#)
- [28] Susanna F. de Rezende, Jakob Nordström, Kilian Risse, and Dmitry Sokolov. Exponential resolution lower bounds for weak pigeonhole principle and perfect matching formulas over sparse graphs. *TheoretCS*, 4, 2025. [5](#)
- [29] Zeev Dvir. On the size of Kakeya sets in finite fields. *Journal of the American Mathematical Society*, 22(4):1093–1097, 2009. [15](#)
- [30] Jordan S Ellenberg and Dion Gijswijt. On large subsets of $\mathbb{F}_q^n$ with no three-term arithmetic progression. *Annals of Mathematics*, pages 339–343, 2017. [15](#)
- [31] Nicola Galesi, Joshua A. Grochow, Toniann Pitassi, and Adrian She. On the algebraic proof complexity of tensor isomorphism. In Amnon Ta-Shma, editor, *38th Computational Complexity Conference, CCC 2023, Warwick, UK, July 17-20, 2023*, LIPIcs, pages 4:1–4:40. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. [5](#), [16](#)
- [32] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986. [11](#), [12](#)
- [33] Dima Grigoriev. Linear lower bound on degrees of positivstellensatz calculus proofs for the parity. *Theor. Comput. Sci.*, 259(1-2):613–622, 2001. [8](#), [13](#), [37](#)
- [34] Dima Grigoriev and Nicolai N. Vorobjov Jr. Complexity of null-and positivstellensatz proofs. *Ann. Pure Appl. Log.*, 113(1-3):153–160, 2001. [19](#)
- [35] 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. [5](#)
- [36] Armin Haken. The intractability of resolution. *Theor. Comput. Sci.*, 39:297–308, 1985. [3](#)
- [37] Johan Håstad. On small-depth frege proofs for tseitin for grids. *J. ACM*, 68(1):1:1–1:31, 2021. [10](#)

- [38] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. [17](#)
- [39] Pavel Hrubes and Iddo Tzameret. Short proofs for the determinant identities. *SIAM J. Comput.*, 44(2):340–383, 2015. [4](#), [6](#)
- [40] Hao Huang. Induced subgraphs of hypercubes and a proof of the sensitivity conjecture. *Annals of Mathematics*, 190(3):949–955, 2019. [15](#)
- [41] Russell Impagliazzo, Pavel Pudlák, and Jirí Sgall. Lower bounds for the polynomial calculus and the gröbner basis algorithm. *Comput. Complex.*, 8(2):127–144, 1999. [3](#), [8](#), [37](#)
- [42] Russell Impagliazzo and Nathan Segerlind. Counting axioms do not polynomially simulate counting gates. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14–17 October 2001, Las Vegas, Nevada, USA*, pages 200–209. IEEE Computer Society, 2001. [8](#), [27](#)
- [43] Dmitry Itsykson and Dmitry Sokolov. Resolution over linear equations modulo two. *Ann. Pure Appl. Log.*, 171(1), 2020. Extended abstract appeared initially in MFCS 2014. [9](#)
- [44] Emil Jerábek. Dual weak pigeonhole principle, boolean complexity, and derandomization. *Ann. Pure Appl. Log.*, 129(1-3):1–37, 2004. [3](#), [14](#), [16](#), [75](#)
- [45] Emil Jeřábek. *Weak pigeonhole principle, and randomized computation*. PhD thesis, Faculty of Mathematics and Physics, Charles University, Prague, 2005. [3](#), [14](#), [99](#)
- [46] Emil Jerábek. Approximate counting in bounded arithmetic. *J. Symb. Log.*, 72(3):959–993, 2007. [3](#), [14](#), [16](#)
- [47] Eitetsu Ken. On some Σ_0^B -formulae generalizing counting principles over V^0 . *Arch. Math. Log.*, 64(1-2):117–158, 2025. [8](#), [27](#)
- [48] Erfan Khaniki. Nisan-wigderson generators in proof complexity: New lower bounds. In Shachar Lovett, editor, *37th Computational Complexity Conference, CCC 2022, Philadelphia, PA, USA, July 20–23, 2022*, LIPIcs, pages 17:1–17:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. [7](#)
- [49] J. Krajíček. *Proof Complexity*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2019. [9](#)
- [50] Jan Krajíček. *Bounded arithmetic, propositional logic, and complexity theory*, volume 60 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 1995. [3](#), [8](#), [37](#)
- [51] Jan Krajíček. On the weak pigeonhole principle. *Fund. Math.*, 170(1-2):123–140, 2001. [4](#), [13](#), [16](#)
- [52] Jan Krajíček. Dual weak pigeonhole principle, pseudo-surjective functions, and provability of circuit lower bounds. *J. Symb. Log.*, 69(1):265–286, 2004. [4](#), [6](#), [7](#), [12](#), [13](#), [16](#), [47](#), [48](#)
- [53] Jan Krajíček, Pavel Pudlák, and Alan Woods. An exponential lower bound to the size of bounded depth frege proofs of the pigeonhole principle. *Random Structures and Algorithms*, 7(1):15–39, 1995. [3](#)
- [54] Jan Krajíček. A proof complexity generator. In *Proc. from the 13th International Congress of Logic, Methodology and Philosophy of Science (Beijing, August 2007)*, Studies in Logic and the Foundations of Mathematics. King’s College Publications, London, 2009. [5](#), [6](#), [7](#), [24](#)
- [55] Jan Krajíček. *Proof Complexity Generators*, volume 497 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, 2025. [4](#)
- [56] Alexis Maciel, Toniann Pitassi, and Alan R. Woods. A new proof of the weak pigeonhole principle. *J. Comput. Syst. Sci.*, 64(4):843–872, 2002. [7](#), [17](#)
- [57] Mladen Mikša and Jakob Nordström. A generalized method for proving polynomial calculus degree lower bounds. *J. ACM*, 71(6):37:1–37:43, 2024. [5](#), [7](#), [56](#)
- [58] Michael Mitzenmacher and Eli Upfal. *Probability and Computing – Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005. [17](#)
- [59] Moritz Müller and Ján Pich. Feasibly constructive proofs of succinct weak circuit lower bounds. *Ann. Pure Appl. Log.*, 171(2), 2020. [14](#), [16](#)

- [60] Noam Nisan. Lower bounds for non-commutative computation (extended abstract). In Cris Koutsougeras and Jeffrey Scott Vitter, editors, *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing, May 5-8, 1991, New Orleans, Louisiana, USA*, pages 410–418. ACM, 1991. [15](#), [81](#)
- [61] Noam Nisan and Avi Wigderson. Hardness vs randomness. *J. Comput. Syst. Sci.*, 49(2):149–167, 1994. [7](#), [47](#)
- [62] Ryan O’Donnell and Yuan Zhou. Approximability and proof complexity. In Sanjeev Khanna, editor, *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 1537–1556. SIAM, 2013. [19](#)
- [63] J. Paris and A. Wilkie. Counting problems in bounded arithmetic. In *Methods in mathematical logic (Caracas, 1983)*, volume 1130 of *Lecture Notes in Math.*, pages 317–340. Springer, Berlin, 1985. [22](#)
- [64] Jeff B. Paris, A. J. Wilkie, and Alan R. Woods. Provability of the pigeonhole principle and the existence of infinitely many primes. *J. Symb. Log.*, 53(4):1235–1244, 1988. [3](#), [17](#), [75](#)
- [65] Fedor Part and Iddo Tzameret. Resolution with counting: Dag-like lower bounds and different moduli. *Comput. Complex.*, 30(1):2, 2021. [9](#)
- [66] Toniann Pitassi, Paul Beame, and Russell Impagliazzo. Exponential lower bounds for the pigeonhole principle. *Comput. Complex.*, 3:97–140, 1993. [3](#)
- [67] Toniann Pitassi, Benjamin Rossman, Rocco A. Servedio, and Li-Yang Tan. Poly-logarithmic frege depth lower bounds via an expander switching lemma. In Daniel Wichs and Yishay Mansour, editors, *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 644–657. ACM, 2016. [10](#)
- [68] Ran Raz. Resolution lower bounds for the weak pigeonhole principle. *J. ACM*, 51(2):115–138, 2004. [3](#), [4](#), [5](#), [7](#), [14](#)
- [69] Ran Raz and Iddo Tzameret. Resolution over linear equations and multilinear proofs. *Ann. Pure Appl. Log.*, 155(3):194–224, 2008. [8](#), [9](#)
- [70] Alexander A. Razborov. Lower bounds on the size of bounded depth circuits over a complete basis with logical addition. *Mathematical Notes of the Academy of Sciences of the USSR*, 41(4):333–338, 1987. [93](#)
- [71] Alexander A. Razborov. Bounded arithmetic and lower bounds in Boolean complexity. In *Feasible Mathematics II. Progress in Computer Science and Applied Logic*, volume 13, pages 344–86. Birkhäuser, 1995. [16](#)
- [72] Alexander A. Razborov. Lower bounds for the polynomial calculus. *Comput. Complex.*, 7(4):291–324, 1998. [3](#), [4](#), [11](#), [14](#), [21](#), [22](#), [38](#), [45](#), [80](#), [83](#)
- [73] Alexander A. Razborov. Improved resolution lower bounds for the weak pigeonhole principle. *Electron. Colloquium Comput. Complex.*, TR01, 2001. [5](#)
- [74] Alexander A. Razborov. Proof complexity of pigeonhole principles. In Werner Kuich, Grzegorz Rozenberg, and Arto Salomaa, editors, *Developments in Language Theory, 5th International Conference, DLT 2001, Vienna, Austria, July 16-21, 2001, Revised Papers*, Lecture Notes in Computer Science, pages 100–116. Springer, 2001. [3](#), [5](#)
- [75] Alexander A. Razborov. Resolution lower bounds for perfect matching principles. *J. Comput. Syst. Sci.*, 69(1):3–27, 2004. [5](#), [14](#), [80](#), [83](#)
- [76] Alexander A. Razborov. Pseudorandom generators hard for k -DNF resolution and polynomial calculus resolution. *Annals of Mathematics*, 181:415–472, 2015. [3](#), [4](#), [6](#), [7](#), [8](#), [9](#), [11](#), [12](#), [13](#), [14](#), [15](#), [16](#), [37](#), [47](#), [48](#), [83](#), [87](#), [100](#), [101](#)
- [77] Grant Schoenebeck. Linear level Lasserre lower bounds for certain k -CSPs. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA*, pages 593–602. IEEE Computer Society, 2008. [13](#), [37](#)
- [78] Roman Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In Alfred V. Aho, editor, *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, New York, New York, USA*, pages 77–82. ACM, 1987. [16](#), [17](#), [89](#)

- [79] Dmitry Sokolov. Pseudorandom generators, resolution and heavy width. In Shachar Lovett, editor, *37th Computational Complexity Conference, CCC 2022, July 20-23, 2022, Philadelphia, PA, USA*, volume 234 of *LIPICs*, pages 15:1–15:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. [4](#), [7](#)
- [80] Michael Soltys and Stephen A. Cook. The proof complexity of linear algebra. *Ann. Pure Appl. Log.*, 130(1-3):277–323, 2004. [4](#), [5](#), [6](#), [15](#)
- [81] Michael Soltys and Alasdair Urquhart. Matrix identities and the pigeonhole principle. *Arch. Math. Log.*, 43(3):351–358, 2004. [4](#), [5](#), [9](#), [37](#)
- [82] Neil Thapen. A model-theoretic characterization of the weak pigeonhold principle. *Ann. Pure Appl. Log.*, 118(1-2):175–195, 2002. [3](#)
- [83] Neil Thapen. *The weak pigeonhole principle in models of bounded arithmetic*. PhD thesis, University of Oxford, 2002. [75](#)
- [84] Iddo Tzameret and Stephen A. Cook. Uniform, integral, and feasible proofs for the determinant identities. *J. ACM*, 68(2):12:1–12:80, 2021. [4](#), [15](#)
- [85] A. Woods. *Some problems in logic and number theory, and their connections*. PhD thesis, University of Manchester, 1981. 14. [3](#)

— Page left blank for ECCC stamp —