The Weak Rank Principle: Lower Bounds and Applications*

Michal Garlík
michal.garlik@gmail.com
Imperial College London
London, UK

Svyatoslav Gryaznov
svyatoslav.i.gryaznov@gmail.com
Imperial College London
London, UK

Hanlin Ren
h4n1in.r3n@gmail.com
Institute for Advanced Study
Princeton, NJ, USA

Iddo Tzameret
iddo.tzameret@gmail.com
Imperial College London
London, UK

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, extending its counting argument to an algebraic setting. As a strengthening of WPHP, it admits proof complexity lower bounds in settings where none are known for WPHP, yet we show that these still yield applications analogous to those of WPHP. In particular, using new generalised types of random restrictions, which may be interesting by themselves, this allows us to resolve a number of open problems in proof complexity, including the construction of proof complexity generators for Polynomial Calculus Resolution over the two-element field ($\mathrm{PCR}_{\mathbb{F}_2}$), new generators for Sherali–Adams (SA), and hardness results for circuit lower bound statements against $\mathrm{PCR}_{\mathbb{F}_2}$, as detailed below.

Generators for $\mathrm{PCR}_{\mathbb{F}_2}$. 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 useful and corresponds to a circuit encoding, as considered by Alekhnovich, 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 $2^{n^{\Omega(1)}}$, meaning we obtain a function generator. This resolves an open problem posed by Alekhnovich *et al.* (*SIAM J. Comput.*, 2004) and Razborov (*Ann. Math.*, 2015) concerning the construction of proof complexity generators with good stretch for $\mathrm{PCR}_{\mathbb{F}_2}$.

Generators for SA. 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 new corresponding pseudoexpectation tailored specifically to the rank principle (and incompatible with the pigeonhole principle).

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

Rank principle as an axiom. Finally, we demonstrate the centrality of the weak rank principle by showing that it is *necessary* for proving NC^2 circuit lower bounds and *sufficient* for proving $\mathrm{AC}^0[p]$ lower bounds.

CCS Concepts

• **Theory of computation** → **Proof complexity**.

Keywords

Proof Complexity, Lower bounds, Polynomial Calculus Resolution, Sherali–Adams, Bounded Arithmetic, Random Restrictions, Circuit Complexity

ACM Reference Format:
Michal Garlík, Svyatoslav Gryaznov, Hanlin Ren, and Iddo Tzameret. 2026. The Weak Rank Principle: Lower Bounds and Applications. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing (STOC '26), June 22–26, 2026, Salt Lake City, UT, USA.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3798129.3800735

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

*The full version of this paper is available at [24]. 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.*

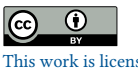

STOC '26, Salt Lake City, UT, USA

ACM ISBN 979-8-4007-2536-4/2026/06
https://doi.org/10.1145/3798129.3800735

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 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 [28], pigeonhole switching lemmas and k-evaluations [1, 44, 54], pigeon dance [32, 59], pseudo-width [56, 61], etc. As for upper bounds, already TC^0-Frege admits polynomial-size proofs of PHP (and hence WPHP); see [11, 15]. Despite this progress, a notable open problem posed by Razborov [63] remains unresolved: proving size lower bounds for WPHP against the polynomial calculus resolution (PCR) proof system (see Section 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 [71] 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 [52] 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 [40, Theorem 7.3.7]), further developed by Thapen [69] and systematically pursued by Jeřábek [35–37], 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 Alekhnovich, Ben-Sasson, Razborov, and Wigderson [2], and by Krajíček [41, 42]; see also the monograph [46] 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 polytime mapping $G\colon \{0,1\}^m \to \{0,1\}^\ell$, with $m < \ell$. The proof complexity of such statements depends on G and is sensitive to the encodings; indeed, [2] suggested three different encodings (see also Sokolov [66] 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 nonsurjective. (Notice that even a single b for which the proof complexity generator is easy would constitute a proof of nonsurjectivity, hence we insist that no such b exists; this point is explained further in Alekhnovich *et al.* [2].)

The hope is that for strong propositional proof systems, one can establish (at least conditionally) that there are no $\mathrm{poly}(\ell)$-size proofs that g is nonsurjective, under the assumption that the mapping g is sufficiently pseudorandom (cf. [63]). Razborov in [63] proved the existence of hard proof complexity generators for systems such as PCR and $\mathrm{Res}(k)$. 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 [59], 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 [59] and later by Raz [56]. 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 "$\mathsf{NP} \not\subseteq \mathsf{P/poly}$", 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. Bonet, Buss, and Pitassi [10, 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 [67] 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 [68] showed that the pigeonhole principle is reducible to (hence, not stronger than) these matrix identities. Later, it was demonstrated by Hrubeš and Tzameret over $\mathbb{F}_2$ [30] and by Tzameret and Cook over the integers [70] 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 [27], and Andrews and Forbes [4] 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 [45], 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 [67, 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 [23] 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 [68].

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 [61]. 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 $\mathrm{PCR}_{\mathbb{F}_2}$ 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 $\mathrm{PCR}_{\mathbb{F}_2}$ (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.

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 [23, 68] 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 [48] and de Rezende, Nordström, Risse and Sokolov [20])[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 follows: for strong enough proof systems closed under low degree reductions WRank is implied by the special case of WRank when $A = I$, which in turn is implied by the inversion principle by [67]; finally, the inversion principle admits polynomial-size proofs in NC^2-Frege by the work of Hrubeš and Tzameret [30]. 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 *el al.* [10].

Note that the relation between the weak rank principle and the strong one (and hence, by [67], its relation to the hard matrix identities) is not entirely understood, since the latter concerns square matrices, while our focus is on highly nonsquare 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 $\mathrm{PCR}_{\mathbb{F}_2}$ (Section 2.1), lower bounds and generators for SA (Section 2.2), hardness of circuit lower bound statements for $\mathrm{PCR}_{\mathbb{F}_2}$ (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 $\mathrm{PCR}_{\mathbb{F}_2}$

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 $A = I$,

[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 [56] and subsequently Razborov [60, 62] as well as de Rezende *et al.* [20], the case of WPHP lower bounds with arbitrary many pigeons m is resolved.

against $\mathrm{PCR}_{\mathbb{F}_2}$. 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 [42, 63]) 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 $\mathrm{PCR}_{\mathbb{F}_2}$ does not admit short proofs of circuit lower bound statements such as $\mathrm{NP} \not\subseteq \mathrm{P/poly}$.

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 $G\colon \{0,1\}^m \to \{0,1\}^\ell$, 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 $\mathrm{PCR}_{\mathbb{F}_2}$ in the literature are based on reductions to the pigeonhole principle. For the *strong* pigeonhole principle, Krajíček [45, 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 $\mathrm{PCR}_{\mathbb{F}_2}$ [3, 48] to yield a slightly superlinear stretch—from n^3 to n^4 bits[2]. (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, $\mathrm{Res}(\log n)$ admits refutations of quasipolynomial size of WPHP [47], yet statements based on linear algebra (and in particular WRank) are expected to be hard for it [10].

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 [56], yield that certain circuit lower bound statements for unbounded fan-in circuits cannot be efficiently proven in PCR.

[2]Given a formula PHP^m_n, Krajíček [45] 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^m_n 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.

This led Alekhnovich *et al.* [2] 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 [51], 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 [42] introduced the notion of *s-iterability*, designed to boost the stretch of a variant of the Nisan generator [51]. Improving [2], Razborov [63] relaxed the expander requirements and showed that linear systems as in [2] 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 $\mathrm{PCR}_{\mathbb{F}_2}$; therefore, this construction cannot yield generators hard for $\mathrm{PCR}_{\mathbb{F}_2}$.

The limitations of existing principles motivate the use of the 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 $\mathrm{AC}^0[p]$-Frege [39] encoded as a (non-CNF) propositional formula. In addition, Sokolov [66] improved the lower bounds for a particular generator encoding (functional encoding as in [2]).

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

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

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

Here, size is measured by the number of distinct monomials appearing in the refutation.

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 noncommutative algebraic branching programs). 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 4). 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 [57]), 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 $\mathsf{PMRank}_n^m(A)$, encodes each equation $\sum_{k=1}^n x_{i,k}y_{k,j} = A_{i,j}$ of $\mathsf{Rank}_n^m(A)$ 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 $A_{i,j} = 1$). Similar encodings of parity computations were considered in Impagliazzo–Segerlind and Ken [33, 38].

As mentioned above, we shall consider only the case $A = I$ and prove lower bounds for this encoding against $\mathsf{PCR}_{\mathbb{F}_2}$. 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 $\mathsf{PCR}_{\mathbb{F}_2}$). *For every $m > n$, every $\mathsf{PCR}_{\mathbb{F}_2}$ refutation of the perfect matching CNF encoding $\mathsf{PMRank}_n^m(I_m)$ requires size $2^{\Omega(n)}$.*

The case of $\mathsf{PCR}_{\mathbb{F}_2}$ 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 $\mathsf{PCR}_{\mathbb{F}_2}$, hence the hardness of $\mathsf{PMRank}_n^m(I_m)$ 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 [5, 12, 14, 26, 32, 40], as we show for the case of $\mathsf{PCR}_{\mathbb{F}}$ when $\mathrm{char}(\mathbb{F}) \neq 2$ 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). *For every $m > n$ and $A \in \{0,1\}^{m\times m}$, any $\mathsf{PCR}_{\mathbb{F}}$ refutation (for a field $\mathbb{F}$ with $\mathrm{char}(\mathbb{F}) \neq 2$) or SoS refutation of the perfect matching CNF encoding $\mathsf{PMRank}_n^m(A)$ 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 [2, 63]. The bamboo-tree CNF encoding of WRank, denoted $\mathsf{BTRank}_n^m(A)$, encodes each equation $\sum_{k=1}^n x_{i,k}y_{k,j} = A_{ij}$ from the algebraic encoding $\mathsf{Rank}_n^m(A)$ 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 $z_{i,j,k}$.

This natural encoding corresponds to a circuit in the sense of [2, 63], and similar encodings have been used repeatedly for parity computations [2, 63, 68].

Theorem 2.4 (Bamboo-tree encoding lower bound; see Theorem 4.2). *For every $m > n$ and $A \in \mathbb{F}_2^{m\times m}$, every $\mathsf{PCR}_{\mathbb{F}_2}$ refutation of $\mathsf{BTRank}_n^m(A)$ requires size $2^{\Omega(n)}$. Consequently, $\mathsf{BTRank}_n^m(A)$ is a proof complexity generator stretching $2nm$ bits into m^2 bits. Moreover, there is an iterated variant of $\mathsf{BTRank}_n^m(A)$ that yields a function generator with stretch $2^{n^{\Omega(1)}}$.*

This result resolves an open problem raised in [2, 63] concerning the construction of proof-complexity generators for $\mathsf{PCR}_{\mathbb{F}_2}$. Razborov's [63] established generators against PCR based on the Nisan construction only over fields of characteristic different from 2. Our result closes this gap by exhibiting hard generators for $\mathsf{PCR}_{\mathbb{F}_2}$. Lower bounds over $\mathbb{F}_2$ are particularly significant and technically more challenging than those over other fields, as we now explain.

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 $\mathsf{PCR}_{\mathbb{F}_2}$, and more generally, linear systems over $\mathbb{F}_p$ are easy for $\mathsf{PCR}_{\mathbb{F}_p}$. When $p \neq 2$, however, such 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 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. [8, Theorem 12], [43]). Likewise, proof systems such as Res(lin) admit lower bounds over characteristic 0 [53], whereas corresponding results over finite fields—and in particular over $\mathbb{F}_2$ (cf. [34])—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 share the following general *random self-reduction* structure: Use a reduction from size to appropriate, carefully defined notions of degree that vary from formula to formula and across our results. This means that we shall start, by way of contradiction, from a small refutation (having a small number of distinct monomials) and apply a random restriction or a generalised random restriction (changing variables

to different ones) 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 variables 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*. The goal is 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 [55] and subsequently Hastad [29] for AC^0-Frege lower bounds.

Bamboo-tree encoding of rank principle. The following explains the lower bounds for both the algebraic $\mathsf{Rank}_n^m(A)$ and the bamboo tree $\mathsf{BTRank}_n^m(A)$ encoding. We outline some ideas behind the proof of Theorem 2.4, which establishes a $2^{\Omega(n)}$ size lower bound for $\mathsf{BTRank}_n^m(A)$ in $\mathsf{PCR}_{\mathbb{F}_2}$. 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.

To overcome this, we relax the degree notion, aiming to balance two competing requirements for establishing size lower bounds: (1) the existence of random self-reductions of $\mathsf{BTRank}_n^m(A)$ that effectively reduce the relaxed degree, and (2) the existence of a lower bound on the relaxed degree for refutations of $\mathsf{BTRank}_n^m(A)$. 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 $\mathsf{BTRank}_n^m(A)$, 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, $\overline{v}$. 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^TA 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 $\tau(\overline{v}) = 1 + \tau(v)$ 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. 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), 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).

Why is the last observation useful? The substitution τ maps the set of polynomial equations $XY = A$ to the set of polynomial equations $XX^TA = A$, 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 [59] 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 $\mathsf{Rank}_n^m(A)$ 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 $\mathsf{Rank}_n^m(A)$ 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.

It is clear that the argument from the above paragraph alone is not sufficient for $\mathsf{BTRank}_n^m(A)$. 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 a small number of columns, it leads to a blow up in the X-mentioned columns in the resulting terms. Our notion of degree for the u-variables concerns 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

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

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 $\mathsf{BTRank}_n^m(A)$ 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 4.2.

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 [25] iteration protocol, which was also used in [63]. This protocol achieves a generator with exponential stretch (i.e., a function generator) by iterating the base generators (the way to achieve such a function generator from the weak rank principle is described after Theorem 2.8 in this section).

Let π be a binary string of length at most κ and consider the polynomial system $\mathsf{IRank}_n^m(\{A^{(\pi)}\},\kappa) = \bigcup_{\pi\in\{0,1\}^{\le\kappa}} X^{(\pi)}Y = A^{(\pi)}$. Here, each $X^{(\pi)}Y = A^{(\pi)}$ is an instance of the weak rank principle, except that the matrices $A^{(\pi)}$ are not necessarily constant: we also allow them to contain the variables from $X^{(\pi')}$ for some $|\pi'| > |\pi|$. These copies of Rank can be naturally arranged in a binary tree, where each node corresponding to $X^{(\pi)}Y = A^{(\pi)}$ with $|\pi| < \kappa$ has two children $X^{(\pi*0)}Y = A^{(\pi*0)}$ and $X^{(\pi*1)}Y = A^{(\pi*1)}$

In this tree, the entries of each matrix $A^{(\pi)}$ can depend only on the entries of $X^{(\pi')}$ from the layers below. This differs from the standard notion of s-iterability, which only assumes a flat structure; that is, the iterated copies are arranged as a path and not as a tree. We need to use a tree structure to maximise the stretch versus the iteration depth ratio, since our lower bound depends on the depth of the tree (that is, on κ in the above notation). The tree structure was also used by Razborov [63]. Nevertheless, 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 $x^{(0)}_{1,1}x^{(1)}_{1,1}y_{1,1}$ 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 $\le d$ while the latter reduces the column degree to $\le d$. This implies the total degree does not exceed $2d^2$. We show that $\mathsf{IRank}_n^m(A^{(\pi)},\kappa)$ requires degree $n/2$ by reducing it to one copy of the formula $\mathsf{Rank}_n^m(I_m)$ 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.

We now turn to the CNF encoding of $\bigcup_{\pi\in\{0,1\}^{\le\kappa}} X^{(\pi)}Y^{(\pi)} = A^{(\pi)}$ (note that here the matrix Y is different in each iteration, unlike before) and again use the iterated protocol from [25]. We encode $\bigcup_{\pi\in\{0,1\}^{\le\kappa}} X^{(\pi)}Y^{(\pi)} = A^{(\pi)}$ 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 $\mathrm{PCR}_{\mathbb{F}_2}$ 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 $u^{(\pi)}_{i,j,k}$ which semantically encodes the sum $\sum_{\ell\in[k]} x^{(\pi)}_{i,k}y^{(\pi)}_{k,j}$, it is clear that it depends not only on the kth 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: $z^{(\pi)}_{i,j,k} = x^{(\pi)}_{i,k}y^{(\pi)}_{k,j}$). 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 IRank. To deal with the z-variables we need to further modify the refutation. This is because of the terms of the form $s = \prod_{j\in J}\overline{z}_{i_0,j,k}$ (or of the symmetric form $\prod_{I\in J}\overline{z}_{i,j_0,k}$), 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 $x^{(\pi)}_{i,k}$ and $y^{(\pi)}_{k,j}$ present whenever $\overline{z}^{(\pi)}_{i,j,k}$ 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 a star with sufficiently many rays—that is, with large enough set J—collapses to its centre $\overline{x}_{i_0,k}$ 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 [18].

First, note that the Nisan generator construction, denoted $\tau(A,b)$ in [63], 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 [9], and whose lower bound against SoS was established in [26, 64]. Combined with the random restriction from Razborov [63], 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)). *For every $m > n$ and every $A \in \{0,1\}^{m\times m}$, every* SA *refutation of (a variant of)* $\mathsf{BTRank}_n^m(A)$ *requires row degree $n-1$.*

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

This result fits in the generator approach to proof complexity as stated by Razborov [63, p. 417] and proposed earlier by Krajíček [41, 42] and Alekhnovich *et al.* [2]. 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 $\mathsf{BTRank}_n^m(A)$, 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 [19]. 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 $\mathsf{BTRank}_n^m(A)$, 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 $\mathsf{BTRank}_n^m(A)$, the marginal distribution condition—a key ingredient of SA lower bounds[3]—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 $\mathsf{BTRank}_n^m(A)$ 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 $\mathsf{BTRank}_n^m(A)$. 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.

[3]In 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 $(I, J \cup \{j_0\})$.

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 [59, 62] and Raz [56] showed the unprovability of circuit lower bounds such as $\mathsf{NP} \not\subseteq \mathsf{P/poly}$ in weak proof systems via reductions from WPHP. Müller and Pich [49] formalised a wide range of circuit lower bounds in Jeřábek's theory for approximate counting in bounded arithmetic [35–37], which includes the dual weak pigeonhole principle as an axiom. More recently, Chen, Li, and Oliveira [13] 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 [56, 59, 62]. Then, as a concrete example, we show that $\mathsf{PCR}_{\mathbb{F}_2}$ does not have efficient proofs of circuit lower bound statements such as $\mathsf{NP} \not\subseteq \mathsf{P/poly}$, by *iterating* the WRank-based proof complexity generators against $\mathsf{PCR}_{\mathbb{F}_2}$ (as in [63]). A variant of circuit lower bound statements was recently proved to be hard for SoS [6].

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

Theorem 2.7 (WRank is necessary for Boolean circuit lower bounds (informal)). *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 noncommutative algebraic branching programs (noncommutative 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 [50]. 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 $\mathsf{PCR}_{\mathbb{F}_2}$ demonstrated in Section 2.1 to show that this system cannot efficiently prove circuit lower bounds[4].

Theorem 2.8 (Informal). *Let $f\colon \{0,1\}^n \to \{0,1\}$. There exists a constant $c \geq 1$ such that for every $s > n^c$, any* $\mathsf{PCR}_{\mathbb{F}_2}$ *refutation*

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

of the statement "f cannot be computed by Boolean circuits of size s" (encoded as a system of polynomial equations) requires size $2^{\mathrm{poly}(s,1/n)}$, which is superpolynomial for $s = n^{\Omega(1)}$.

As in the case of Razborov [63] our Boolean circuits are in the basis $\neg, \wedge, \vee, \oplus$ where the last three connectives are binary. The proof idea is similar to [63]. 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 $A^{(\pi)}_{1,1} = f(\pi)$ 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 [21], Cap Set problem [17, 22] and the Sensitivity Conjecture [31]. Indeed, there is a whole book dedicated to linear algebraic methods in combinatorics [7].

Soltys and Cook [67] 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 ∀LAP, the strongest theory considered in [67] which incorporates matrix powering and certain induction schema, proves many linear algebraic identities including the (strong) rank principle. Subsequently, Tzameret and Cook [70] 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 [16] 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 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 $\mathsf{AC}^0[p]$ 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 [37].

In this work, we put forward a theory $\overline{\mathsf{V}^0(p)} + \mathrm{WRank}_p$, which incorporates the two-sorted theory $\overline{\mathsf{V}^0(p)}$ in the style of Cook–Nguyen [15] 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 $\mathsf{AC}^0[p]$ [65]:

Theorem 2.9 (Informal). *For every prime number $p > 2$, $\overline{\mathsf{V}^0(p)} + \mathrm{WRank}_p$ proves that MOD_2 cannot be computed by $\mathsf{AC}^0[p]$ circuits of depth d and size $2^{o(n^{1/(2d)})}$.*

Previously, Müller and Pich [49] formalised Smolensky's $\mathsf{AC}^0[p]$ lower bounds in Jeřábek's theory $\mathbf{APC}_1$ [35, 37]. We will discuss the differences between our formalisation and theirs. We also note that Razborov mentioned in [58, 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 $\overline{\mathsf{V}^0(p)} + \mathrm{WRank}_p$. In addition, Galesi *et al.* [23] 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 [41, 42] and Alekhnovich *et al.* [2], and later emphasised by Razborov [63, 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 $\mathsf{NP} \not\subseteq \mathsf{P/poly}$.
- Develop bounded arithmetic theories that use WRank as axioms, such as the theory $\overline{\mathsf{V}^0(p)} + \mathrm{WRank}_p$ 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 [65] can be seen as the first step towards this goal.
- The weak rank principle provides a natural candidate for lower bounds against $\mathsf{AC}^0[p]$-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 $\mathrm{Res}(\log n)$ and thus in AC^0-Frege [47, 52], indicating that stronger algebraic principles are required to obtain meaningful lower bounds.

4 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 [2, 63], and similar encodings have been used repeatedly for parity computations [2, 63, 68].

Definition 4.1 (The CNF encoding $\mathsf{BTRank}_n^m(A)$*).* Assume $m > n$ are positive integers and $A \in \mathbb{F}_2^{m\times m}$ is an m by m matrix over the two-element field. The CNF encoding the bamboo tree rank formula, denoted $\mathsf{BTRank}_n^m(A)$, 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 $k \in \{2, \ldots, n\}$, $u_{i,j,k} = u_{i,j,k-1} + z_{i,j,k}$, encoded by the following four clauses:

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

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 $\mathsf{PCR}_{\mathbb{F}_2}$ size lower bound of $2^{\Omega(n)}$. This makes $\mathsf{BTRank}_n^m(A)$ a proof complexity generator, stretching $2mn$ input bits into m^2 output bits.

Theorem 4.2. *Suppose that $m > n \geq 16$ are integers, 8 divides n, and $A \in \mathbb{F}_2^{m\times m}$ is an m by m matrix over $\mathbb{F}_2$. Then any $\mathsf{PCR}_{\mathbb{F}_2}$ refutation of $\mathsf{BTRank}_n^m(A)$ 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 restricting that reduces with high probability a small-size refutation into a small-degree refutation, which then leads to a contradiction with the pigeonhole principle PCR degree lower bounds. Accordingly, our goal here is to show that if $\mathsf{BTRank}_n^m(A)$ admits a small $\mathsf{PCR}_{\mathbb{F}_2}$ 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 $\mathsf{BTRank}_n^m(A)$ 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 $\mathsf{PCR}_{\mathbb{F}_2}$ refutation of $\mathsf{FPHP}_{n'}^m$ of degree less than $n'/2$, which is ruled out by a theorem of Razborov [59].

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

We define the random substitution ρ that randomly compresses the matrices X and Y^T 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. First, we demonstrate the desired effect of ρ on the u-variables occurring in terms.

Definition 4.3 (Row and column indices mentioned by u-variables in a term). Let t be a term in the variables of $\mathsf{BTRank}_n^m(A)$. 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 range originally over $[m]$ while column indices range originally over $[n]$.

Lemma 4.4 (Random substitution ρ kills terms that mention many u-rows). *Suppose that $m > n \geq 8$ are integers, 4 divides n, and $A \in \mathbb{F}_2^{m\times m}$ is an m by m matrix over $\mathbb{F}_2$. Suppose that t is a term in the variables of $\mathsf{BTRank}_n^m(A)$, $I \subseteq [m]$ 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, $\Pr[t \restriction \rho \neq 0] \leq (3/4)^{|I|}$.*

Next, we verify that ρ transforms a refutation of $\mathsf{BTRank}_n^m(A)$ into a refutation of $\mathsf{BTRank}_{\tilde{n}}^m(A)$. Namely, we reduce the number of columns in X (and accordingly, number of rows in Y). This is a

structural consequence of the way the substitution is defined, and happens with probability 1.

Lemma 4.5. *Suppose that $m > n \geq 8$ are integers, 4 divides n, and $A \in \mathbb{F}_2^{m \times m}$ is an m by m matrix over $\mathbb{F}_2$. Let Π be a $\mathsf{PCR}_{\mathbb{F}_2}$ refutation of $\mathsf{BTRank}_n^m(A)$. Then $\Pi \upharpoonright \rho$ is a $\mathsf{PCR}_{\mathbb{F}_2}$ refutation of $\mathsf{BTRank}_{(n-4)/4}^m(A)$.*

The second random substitution, denoted by σ, is intended to simplify the x-, y- and z-variable occurrences in terms while, again, turning $\mathsf{BTRank}_n^m(A)$ into a smaller instance. In contrast to the substitution ρ, which aims to reduce in each term with high probability the number of *row* indices (of X) appearing through the *u-variables* (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 of the original instance from n to $(n-4)/4$ (Lemma 4.5). However, this does not substantially decrease the number of column indices involved by a given term: the term may still mention many different 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 4.6. [Column indices mentioned in a term by x-,y- and z-variables] Let t be a term in the variables of $\mathsf{BTRank}_n^m(A)$ and let $k \in [n]$. We say that k is *X-column-mentioned in t* if there is $i \in [m]$ 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 is $j \in [m]$ 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 4.7. *Suppose that $m > n \geq 3$ are integers, 2 divides $n-1$, and $A \in \mathbb{F}_2^{m \times m}$ is an m by m matrix over $\mathbb{F}_2$. Suppose that t is a term in the variables of $\mathsf{BTRank}_n^m(A)$ 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 $\Pr[t \upharpoonright \sigma \neq 0] \leq e^{-|K|/16} + 2^{-|K|/4}$.*

Lemma 4.8. *Suppose that $m > n \geq 3$ are integers, 2 divides $n-1$, and $A \in \mathbb{F}_2^{m \times m}$ is an m by m matrix over $\mathbb{F}_2$. Let Π be a $\mathsf{PCR}_{\mathbb{F}_2}$ refutation of $\mathsf{BTRank}_n^m(A)$. Then $\Pi \upharpoonright \sigma$ is a $\mathsf{PCR}_{\mathbb{F}_2}$ refutation of $\mathsf{BTRank}_{(n-1)/2}^m(A)$.*

The next lemma combines Lemmas 4.4, 4.5, 4.7 and 4.8.

Lemma 4.9. *Suppose that $m > n \geq 16$ are integers, 8 divides n, and $A \in \mathbb{F}_2^{m \times m}$ is an m by m matrix over $\mathbb{F}_2$. If Π is a $\mathsf{PCR}_{\mathbb{F}_2}$ refutation of $\mathsf{BTRank}_n^m(A)$ of size s, then there exists a $\mathsf{PCR}_{\mathbb{F}_2}$ refutation Π' of $\mathsf{BTRank}_{(n-8)/8}^m(A)$ 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 $k \in [\frac{n-8}{8}]$ that are X-column-mentioned or Y^T-column-mentioned or Z-column-mentioned in t is less than d.*

The proof of Theorem 4.2 follows from the application of Lemma 4.9 and the subsequent reduction to $\mathsf{FPHP}_{\tilde{n}}^m$, where $\tilde{n} = (n-8)/8$. Assuming the size of the original refutation was less than $2^{\frac{\log e}{512} n - 2}$, we obtain a refutation of $\mathsf{FPHP}_{\tilde{n}}^m$ of degree less than $\tilde{n}/2$, contradicting the lower bound of Razborov [59].

Acknowledgments

We are indebted to Robert Andrews for very helpful discussions throughout this project.

References

[1] Miklós Ajtai. 1994. The Complexity of the Pigeonhole Principle. *Comb.* 14, 4 (1994), 417–433. doi:10.1007/BF01302964

[2] Michael Alekhnovich, Eli Ben-Sasson, Alexander A. Razborov, and Avi Wigderson. 2004. Pseudorandom Generators in Propositional Proof Complexity. *SIAM J. Comput.* 34, 1 (2004), 67–88. doi:10.1137/S0097539701389944

[3] Michael Alekhnovich and Alexander A. Razborov. 2001. 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*. IEEE Computer Society, 190–199. doi:10.1109/SFCS.2001.959893

[4] Robert Andrews and Michael A. Forbes. 2022. Ideals, determinants, and straightening: proving and using lower bounds for polynomial ideals. In *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, Stefano Leonardi and Anupam Gupta (Eds.). ACM, 389–402. doi:10.1145/3519935.3520025

[5] Albert Atserias and Tuomas Hakoniemi. 2019. 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*. 24:1–24:20. doi:10.4230/LIPIcs.CCC.2019.24

[6] Per Austrin and Kilian Risse. 2023. Sum-Of-Squares Lower Bounds for the Minimum Circuit Size Problem. In *38th Computational Complexity Conference, CCC 2023, July 17-20, 2023, Warwick, UK (LIPIcs, Vol. 264)*, Amnon Ta-Shma (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 31:1–31:21. doi:10.4230/LIPICS.CCC.2023.31

[7] László Babai and Péter Frankl. 1992. *Linear algebra methods in combinatorics*. Department of Computer Science, Univ. of Chicago.

[8] Paul Beame, Stephen A. Cook, Jeff Edmonds, Russell Impagliazzo, and Toniann Pitassi. 1998. The Relative Complexity of NP Search Problems. *J. Comput. Syst. Sci.* 57, 1 (1998), 3–19. doi:10.1006/JCSS.1998.1575

[9] Eli Ben-Sasson and Russell Impagliazzo. 2010. Random Cnf's are Hard for the Polynomial Calculus. *Comput. Complex.* 19, 4 (2010), 501–519. doi:10.1007/S00037-010-0293-1

[10] Maria Luisa Bonet, Samuel R. Buss, and Toniann Pitassi. 1995. Are there hard examples for Frege systems? In *Feasible mathematics, II (Ithaca, NY, 1992)*. Progr. Comput. Sci. Appl. Logic, Vol. 13. Birkhäuser Boston, Boston, MA, 30–56. doi:10.1007/978-1-4612-2566-9_3

[11] Samuel R. Buss. 1987. Polynomial Size Proofs of the Propositional Pigeonhole Principle. *J. Symb. Log.* 52, 4 (1987), 916–927. doi:10.2307/2273826

[12] Samuel R. Buss, Dima Grigoriev, Russell Impagliazzo, and Toniann Pitassi. 2001. Linear Gaps between Degrees for the Polynomial Calculus Modulo Distinct Primes. *J. Comput. Syst. Sci.* 62, 2 (2001), 267–289. doi:10.1006/JCSS.2000.1726

[13] Lijie Chen, Jiatu Li, and Igor C. Oliveira. 2024. 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*. IEEE, 505–527. doi:10.1109/FOCS61266.2024.00040

[14] Matthew Clegg, Jeff Edmonds, and Russell Impagliazzo. 1996. Using the Groebner Basis Algorithm to Find Proofs of Unsatisfiability. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, Gary L. Miller (Ed.). ACM, 174–183. doi:10.1145/237814.237860

[15] Stephen Cook and Phuong Nguyen. 2010. *Logical Foundations of Proof Complexity*. Cambridge University Press. doi:10.1017/CBO9780511676277

[16] Stephen A. Cook. 1985. A Taxonomy of Problems with Fast Parallel Algorithms. *Inf. Control.* 64, 1-3 (1985), 2–21. doi:10.1016/S0019-9958(85)80041-3

[17] Ernie Croot, Vsevolod F Lev, and Péter Pál Pach. 2017. Progression-free sets in $\mathbb{Z}_4^n$ are exponentially small. *Annals of Mathematics* (2017), 331–337.

[18] Stefan S. Dantchev and Barnaby Martin. 2013. Rank complexity gap for Lovász-Schrijver and Sherali-Adams proof systems. *Comput. Complex.* 22, 1 (2013), 191–213. doi:10.1007/S00037-012-0049-1

[19] Stefan S. Dantchev, Barnaby Martin, and Mark Nicholas Charles Rhodes. 2009. Tight rank lower bounds for the Sherali-Adams proof system. *Theor. Comput. Sci.* 410, 21-23 (2009), 2054–2063. doi:10.1016/J.TCS.2009.01.002

[20] Susanna F. de Rezende, Jakob Nordström, Kilian Risse, and Dmitry Sokolov. 2025. Exponential Resolution Lower Bounds for Weak Pigeonhole Principle and Perfect Matching Formulas over Sparse Graphs. *TheoretiCS* 4 (2025). doi:10.46298/THEORETICS.25.9

[21] Zeev Dvir. 2009. On the size of Kakeya sets in finite fields. *Journal of the American Mathematical Society* 22, 4 (2009), 1093–1097.

[22] Jordan S Ellenberg and Dion Gijswijt. 2017. On large subsets of $\mathbb{F}_q^n$ with no three-term arithmetic progression. *Annals of Mathematics* (2017), 339–343.

[23] Nicola Galesi, Joshua A. Grochow, Toniann Pitassi, and Adrian She. 2023. On the Algebraic Proof Complexity of Tensor Isomorphism. In *38th Computational Complexity Conference, CCC 2023, Warwick, UK, July 17-20, 2023 (LIPIcs)*, Amnon Ta-Shma (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 4:1–4:40. doi:10.4230/LIPICS.CCC.2023.4

[24] Michal Garlík, Svyatoslav Gryaznov, Hanlin Ren, and Iddo Tzameret. 2026. The Weak Rank Principle: Lower Bounds and Applications. (2026). https://symplectic.imperial.ac.uk/viewobject.html?cid=1&id=1653030

[25] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. 1986. How to construct random functions. *J. ACM* 33, 4 (1986), 792–807. doi:10.1145/6490.6503

[26] Dima Grigoriev. 2001. Linear lower bound on degrees of Positivstellensatz calculus proofs for the parity. *Theor. Comput. Sci.* 259, 1-2 (2001), 613–622. doi:10.1016/S0304-3975(00)00157-2

[27] Joshua A. Grochow and Toniann Pitassi. 2018. Circuit Complexity, Proof Complexity, and Polynomial Identity Testing: The Ideal Proof System. *J. ACM* 65, 6 (2018), 37:1–37:59. doi:10.1145/3230742

[28] Armin Haken. 1985. The Intractability of Resolution. *Theor. Comput. Sci.* 39 (1985), 297–308. doi:10.1016/0304-3975(85)90144-6

[29] Johan Håstad. 2021. On Small-depth Frege Proofs for Tseitin for Grids. *J. ACM* 68, 1 (2021), 1:1–1:31. doi:10.1145/3425606

[30] Pavel Hrubes and Iddo Tzameret. 2015. Short Proofs for the Determinant Identities. *SIAM J. Comput.* 44, 2 (2015), 340–383. doi:10.1137/130917788

[31] Hao Huang. 2019. Induced subgraphs of hypercubes and a proof of the sensitivity conjecture. *Annals of Mathematics* 190, 3 (2019), 949–955.

[32] Russell Impagliazzo, Pavel Pudlák, and Jirí Sgall. 1999. Lower Bounds for the Polynomial Calculus and the Gröbner Basis Algorithm. *Comput. Complex.* 8, 2 (1999), 127–144. doi:10.1007/S000370050024

[33] Russell Impagliazzo and Nathan Segerlind. 2001. 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*. IEEE Computer Society, 200–209. doi:10.1109/SFCS.2001.959894

[34] Dmitry Itsykson and Dmitry Sokolov. 2020. Resolution over linear equations modulo two. *Ann. Pure Appl. Log.* 171, 1 (2020). doi:10.1016/j.apal.2019.102722 Extended abstract appeared initially in MFCS 2014.

[35] Emil Jerábek. 2004. Dual weak pigeonhole principle, Boolean complexity, and derandomization. *Ann. Pure Appl. Log.* 129, 1-3 (2004), 1–37. doi:10.1016/J.APAL.2003.12.003

[36] Emil Jeřábek. 2005. *Weak pigeonhole principle, and randomized computation.* Ph. D. Dissertation. Faculty of Mathematics and Physics, Charles University, Prague.

[37] Emil Jerábek. 2007. Approximate counting in bounded arithmetic. *J. Symb. Log.* 72, 3 (2007), 959–993. doi:10.2178/JSL/1191333850

[38] Eitetsu Ken. 2025. On some Σ_0^B-formulae generalizing counting principles over V^0. *Arch. Math. Log.* 64, 1-2 (2025), 117–158. doi:10.1007/S00153-024-00938-1

[39] Erfan Khaniki. 2022. Nisan-Wigderson Generators in Proof Complexity: New Lower Bounds. In *37th Computational Complexity Conference, CCC 2022, Philadelphia, PA, USA, July 20-23, 2022 (LIPIcs)*, Shachar Lovett (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 17:1–17:15. doi:10.4230/LIPICS.CCC.2022.17

[40] Jan Krajíček. 1995. *Bounded arithmetic, propositional logic, and complexity theory.* Encyclopedia of Mathematics and its Applications, Vol. 60. Cambridge University Press, Cambridge. xiv+343 pages. doi:10.1017/CBO9780511529948

[41] Jan Krajíček. 2001. On the weak pigeonhole principle. *Fund. Math.* 170, 1-2 (2001), 123–140. doi:10.4064/fm170-1-8

[42] Jan Krajícek. 2004. Dual weak pigeonhole principle, pseudo-surjective functions, and provability of circuit lower bounds. *J. Symb. Log.* 69, 1 (2004), 265–286. doi:10.2178/JSL/1080938841

[43] J. Krajíček. 2019. *Proof Complexity.* Cambridge University Press. https://books.google.co.uk/books?id=uOyKuQEACAAJ

[44] Jan Krajíček, Pavel Pudlák, and Alan Woods. 1995. An exponential lower bound to the size of bounded depth Frege proofs of the pigeonhole principle. *Random Structures and Algorithms* 7, 1 (1995), 15–39. doi:10.1002/rsa.3240070103

[45] Jan Krajíček. 2009. 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.

[46] Jan Krajíček. 2025. *Proof Complexity Generators.* London Mathematical Society Lecture Note Series, Vol. 497. Cambridge University Press. doi:10.1017/9781009611664

[47] Alexis Maciel, Toniann Pitassi, and Alan R. Woods. 2002. A New Proof of the Weak Pigeonhole Principle. *J. Comput. Syst. Sci.* 64, 4 (2002), 843–872. doi:10.1006/JCSS.2002.1830

[48] Mladen Mikša and Jakob Nordström. 2024. A Generalized Method for Proving Polynomial Calculus Degree Lower Bounds. *J. ACM* 71, 6 (2024), 37:1–37:43. doi:10.1145/3675668

[49] Moritz Müller and Ján Pich. 2020. Feasibly constructive proofs of succinct weak circuit lower bounds. *Ann. Pure Appl. Log.* 171, 2 (2020). doi:10.1016/J.APAL.2019.102735

[50] Noam Nisan. 1991. Lower Bounds for Non-Commutative Computation (Extended Abstract). In *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing, May 5-8, 1991, New Orleans, Louisiana, USA*, Cris Koutsougeras and Jeffrey Scott Vitter (Eds.). ACM, 410–418. doi:10.1145/103418.103462

[51] Noam Nisan and Avi Wigderson. 1994. Hardness vs Randomness. *J. Comput. Syst. Sci.* 49, 2 (1994), 149–167. doi:10.1016/S0022-0000(05)80043-1

[52] Jeff B. Paris, A. J. Wilkie, and Alan R. Woods. 1988. Provability of the Pigeonhole Principle and the Existence of Infinitely Many Primes. *J. Symb. Log.* 53, 4 (1988), 1235–1244. doi:10.1017/S0022481200028061

[53] Fedor Part and Iddo Tzameret. 2021. Resolution with Counting: Dag-Like Lower Bounds and Different Moduli. *Comput. Complex.* 30, 1 (2021), 2. doi:10.1007/s00037-020-00202-x

[54] Toniann Pitassi, Paul Beame, and Russell Impagliazzo. 1993. Exponential Lower Bounds for the Pigeonhole Principle. *Comput. Complex.* 3 (1993), 97–140. doi:10.1007/BF01200117

[55] Toniann Pitassi, Benjamin Rossman, Rocco A. Servedio, and Li-Yang Tan. 2016. Poly-logarithmic Frege depth lower bounds via an expander switching lemma. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, Daniel Wichs and Yishay Mansour (Eds.). ACM, 644–657. doi:10.1145/2897518.2897637

[56] Ran Raz. 2004. Resolution lower bounds for the weak pigeonhole principle. *J. ACM* 51, 2 (2004), 115–138. doi:10.1145/972639.972640

[57] Ran Raz and Iddo Tzameret. 2008. Resolution over linear equations and multilinear proofs. *Ann. Pure Appl. Log.* 155, 3 (2008), 194–224. doi:10.1016/J.APAL.2008.04.001

[58] Alexander A. Razborov. 1995. Bounded Arithmetic and lower bounds in Boolean complexity. In *Feasible Mathematics II. Progress in Computer Science and Applied Logic*. Vol. 13. Birkhäuser, 344–86. doi:10.1007/978-1-4612-2566-9_12

[59] Alexander A. Razborov. 1998. Lower Bounds for the Polynomial Calculus. *Comput. Complex.* 7, 4 (1998), 291–324. doi:10.1007/S000370050013

[60] Alexander A. Razborov. 2001. Improved Resolution Lower Bounds for the Weak Pigeonhole Principle. *Electron. Colloquium Comput. Complex.* TR01, Article TR01-055 (2001). ECCC:TR01-055 https://eccc.weizmann.ac.il/eccc-reports/2001/TR01-055/index.html

[61] Alexander A. Razborov. 2001. Proof Complexity of Pigeonhole Principles. In *Developments in Language Theory, 5th International Conference, DLT 2001, Vienna, Austria, July 16-21, 2001, Revised Papers (Lecture Notes in Computer Science)*, Werner Kuich, Grzegorz Rozenberg, and Arto Salomaa (Eds.). Springer, 100–116. doi:10.1007/3-540-46011-X_8

[62] Alexander A. Razborov. 2004. Resolution lower bounds for perfect matching principles. *J. Comput. Syst. Sci.* 69, 1 (2004), 3–27. doi:10.1016/J.JCSS.2004.01.004

[63] Alexander A. Razborov. 2015. Pseudorandom Generators Hard for k-DNF Resolution and Polynomial Calculus Resolution. *Annals of Mathematics* 181 (2015), 415–472.

[64] Grant Schoenebeck. 2008. 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*. IEEE Computer Society, 593–602. doi:10.1109/FOCS.2008.74

[65] Roman Smolensky. 1987. Algebraic Methods in the Theory of Lower Bounds for Boolean Circuit Complexity. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, New York, New York, USA*, Alfred V. Aho (Ed.). ACM, 77–82. doi:10.1145/28395.28404

[66] Dmitry Sokolov. 2022. Pseudorandom Generators, Resolution and Heavy Width. In *37th Computational Complexity Conference, CCC 2022, July 20-23, 2022, Philadelphia, PA, USA (LIPIcs, Vol. 234)*, Shachar Lovett (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 15:1–15:22. doi:10.4230/LIPICS.CCC.2022.15

[67] Michael Soltys and Stephen A. Cook. 2004. The proof complexity of linear algebra. *Ann. Pure Appl. Log.* 130, 1-3 (2004), 277–323. doi:10.1016/J.APAL.2003.10.018

[68] Michael Soltys and Alasdair Urquhart. 2004. Matrix identities and the pigeonhole principle. *Arch. Math. Log.* 43, 3 (2004), 351–358. doi:10.1007/S00153-003-0205-Z

[69] Neil Thapen. 2002. A model-theoretic characterization of the weak pigeonhold principle. *Ann. Pure Appl. Log.* 118, 1-2 (2002), 175–195. doi:10.1016/S0168-0072(02)00038-6

[70] Iddo Tzameret and Stephen A. Cook. 2021. Uniform, Integral, and Feasible Proofs for the Determinant Identities. *J. ACM* 68, 2 (2021), 12:1–12:80. doi:10.1145/3431922

[71] A. Woods. 1981. *Some problems in logic and number theory, and their connections.* PhD thesis. University of Manchester. 14.

Received 4 November 2025; accepted 1 February 2026