Electronic Colloquium on Computational Complexity, Report No. 7 (2026)

New Polynomial-Depth Res(+) Lower Bounds

Yaroslav Alekseev*†1 and Nikita Gaevoy‡1

1Technion – Israel Institute of Technology, Haifa, Israel

Abstract

Res($\oplus$) is the simplest fragment of AC^0[2]-Frege for which no super-polynomial lower bounds on the size of proofs are known. Bhattacharya and Chattopadhyay [BC25] recently proved lower bounds of the form $\exp(\tilde{\Omega}(N^{\varepsilon}))$ on the size of Res($\oplus$) proofs whose depth is upper bounded by $O(N^{2-\varepsilon})$, where N is the number of variables in the unsatisfiable CNF formula. Their proof employs the "random walk with restarts" technique, which is unlikely to be used to prove lower bounds for proofs of depth greater than N^2. The next natural step would be to prove a lower bound for proofs of depth polynomial in the number of variables. In this work, we address this issue by proposing a new method for proving bounded-depth lower bounds.

We introduce a natural extension of the Bit Pigeon Hole Principle called the Constrained Bit Pigeon Hole Principle (CBPHP, for short), for which we will prove the following lower bounds:

- Under some natural combinatorial assumption, for any constant k, there is a collection of instances of CBPHP such that any Res($\oplus$) proof of CBPHP of depth N^k requires size $\exp(N^{\Omega(1)})$.
- Unconditionally, for any constant k, there is a collection of instances of CBPHP such that any RevRes($\oplus$) proof of CBPHP of depth N^k requires size $\exp(N^{\Omega(1)})$, where RevRes($\oplus$) is the fragment of Res($\oplus$) defined similarly to RevRes [GHJ+24].
- Unconditionally, for any small enough $\varepsilon > 0$ there is a collection of instances of CBPHP such that any Res($\oplus$) proof of CBPHP of depth $N^{2-\varepsilon}$ requires size $\exp(N^{\Omega(\varepsilon)})$.

1 Introduction

Propositional proof systems are used to certify that given CNF formulas are unsatisfiable. Cook and Rekhow [CR79] showed that NP $\neq$ coNP implies that for every propositional proof system, there is a family of hard formulas that require superpolynomial proof sizes. However, we currently cannot prove superpolynomial proof-size lower bounds for many particular proof systems.

One of the most well-studied proof systems is the Resolution proof system. Given an unsatisfiable collection of Boolean disjunctions of literals $C_1, \ldots, C_k$, one can always derive $\bot$ (the empty disjunction) with the following derivation rule:

$$\frac{A \vee x \quad B \vee \neg x}{A \vee B}.$$

*Supported by ISF grant 507/24.

†e-mail: tolstreg@gmail.com

‡e-mail: nikgaevoy@gmail.com

ISSN 1433-8092

This particular rule is called the *Resolution* rule. Many lower bounds are known for the Resolution proof system (see, for example, [Urq87, BSW01, ABSRW04]). One way to generalize the Resolution proof system is to allow the proof system to work with bounded-depth formulas instead of clauses. This is known as the AC^0-Frege proof system. For this proof system, we also know how to prove lower bounds (see, for example, [Ajt88, Hå20, HR25, Hå23]). If we allow formulas to contain XOR gates, we get the $\mathrm{AC}^0[2]$-Frege proof system. Obtaining superpolynomial lower bounds for $\mathrm{AC}^0[2]$-Frege for any unsatisfiable formula in CNF is still wide open.

Resolution over parities ($\mathrm{Res}(\oplus)$, for short), introduced in [IS14], is a generalization of the Resolution proof system that operates with disjunctions of XORs instead of disjunctions of literals. It is easy to see that $\mathrm{Res}(\oplus)$ is the subsystem of $\mathrm{AC}^0[2]$-Frege. So, the superpolynomial $\mathrm{Res}(\oplus)$ lower bounds are a stepping stone towards the $\mathrm{AC}^0[2]$-Frege lower bounds. Unfortunately, we still do not have superpolynomial-size lower bounds, even for $\mathrm{Res}(\oplus)$. Nevertheless, recent progress has been made for certain fragments of $\mathrm{Res}(\oplus)$.

Tree-like lower bounds There are plenty of tree-like $\mathrm{Res}(\oplus)$ lower bounds for particular formulas obtained by different techniques: Prover-Delayer games [IS14, IS20, Gry19, GOR24], reductions from communication complexity [IS14, IS20, IR21, Kra18], reductions from polynomial calculus degree [GK18]. Chattopadhyay et al. [CMSS23] proved that resolution depth can be lifted to tree-like $\mathrm{Res}(\oplus)$ size using stifling gadgets. Independently, Beame and Koroth [BK23] obtained similar results.

Bounded-depth $\mathrm{Res}(\oplus)$ **lower bounds** Inspired by the lower bound for Regular $\mathrm{Res}(\oplus)$[1] from [EGI24], Alekseev and Itsykson [AI25] proved an exponential lower bound for the $\mathrm{Res}(\oplus)$ refutations of depth up to $O(N \log\log N)$. Their result was subsequently refined to depths up to $O(n \log n)$ (see [EI25]), $O(N^{3/2-\varepsilon})$ (see [BI25]), and $O(N^{2-\varepsilon})$ (see [BC25]). All of these results use the same key technique: random walk with restarts. The main issue with this technique is that it seems impossible to prove $\mathrm{Res}(\oplus)$ superpolynomial size lower bounds for depth greater than N^2, as we explain later. The central goal of this work is to overcome this issue by presenting a technique that can potentially prove lower bounds for arbitrarily large polynomial depth.

1.1 Our Results

In this paper, we consider a new formula $\mathrm{CBPHP}^{N,M}_{k,f}$ over $O(N \log N)$ variables for which we prove the following conditional polynomial-depth lower bound:

Theorem 1.1. *Suppose that Conjecture 1.4 holds for $q > 1$ and $r = k - 2$. Then any* $\mathrm{Res}(\oplus)$ *refutation of* $\mathrm{CBPHP}^{N,M}_{k,f}$ *of size at most $2^{(\log(N))^{q/2}}$ requires depth d such that*

$$d \geq \Omega\left(N^{(k-2)\cdot c}\right),$$

where $c > 0$ only depends on q.

Conjecture 1.4 will be described in Section 1.2. For more details, see Section 4.

[1] Regular $\mathrm{Res}(\oplus)$ is a fragment of depth-n $\mathrm{Res}(\oplus)$ similar to regular Resolution, which we are not going to define in this paper.

By restricting ourselves to RevRes(⊕), we can get an unconditional lower bound. RevRes(⊕) is the fragment of Res(⊕), that was first defined in [AG25] and is similar to Reversible Resolution (see [GHJ+24] for reference). Informally, in this proof system, we make all the derivations on a blackboard, and every time we apply a derivation rule, we erase the premises from the board and add the conclusion. Göös et al. [GHJ+24] proved that Reversible Resolution is strictly weaker than Resolution. Therefore, proving RevRes(⊕) lower bounds might be a potentially easier task than proving Res(⊕) lower bounds. On the other hand, obtaining lower bounds on this system is also interesting as a potential approach to proving lower bounds on Res(⊕). Alekseev and Gaevoy [AG25] proved that RevRes(⊕) is equivalent to Res(⊕) $\curlywedge$ CatRes(⊕), which means, for example, that superpolynomial RevRes(⊕) lower bounds for Pigeon Hole Principle imply superpolynomial lower bounds for Res(⊕).

In this paper, we give the first example of a lower bound for RevRes(⊕) which does not immediately follow from a Res(⊕) lower bound. So, this is the first evidence that proving RevRes(⊕) lower bounds might be technically easier than proving Res(⊕) lower bounds.

Theorem 1.2. *Any* RevRes(⊕) *refutation of* $\mathrm{CBPHP}^{N,M}_{k,f}$ *of size at most* $\exp\left(N^{1/(2k+3)}\right)$ *requires depth d such that*

$$d \geq \Omega(N^{k-2}).$$

Finally, as a byproduct of the construction of our formula, we can directly apply the random walk with restarts, without using any of the advanced techniques such as lifting [BCD24, AI25, BC25], amortized closure [EI25, BC25], or any non-trivial sampling [BC25, BI25], to prove an unconditional lower almost N^2-depth Res(⊕) refutations of $\mathrm{CBPHP}^{N,M}_{k,f}$.

Theorem 1.3. *Any* Res(⊕) *refutation of* $\mathrm{CBPHP}^{N,M}_{k,f}$ *of size S requires depth d such that*

$$d \geq \Omega\left(N^{2\frac{k-2}{k+2}} / \log S\right).$$

So, by choosing k, we can get a lower bound of the form $\Omega\left(N^{2-\varepsilon}/\log S\right)$ for any $\varepsilon > 0$.

1.2 Our Techniques

We showcase the technique of random walk with restarts on the case of lower bounds for the bounded depth Res(⊕) proof system on the formula BPHP. Informally, the random walk with restarts works in two phases: random walk and restart. The random walk starts in some vertex of a proof DAG and, using a random full assignment to all variables, traverses the proof graph for a certain number of steps. In the case of BPHP, we can afford doing $O(\sqrt{N})$ steps with 0.99 probability of not detecting a collision. If we are given a clause with a width much larger than the width of the starting clause, the probability that the random walk stops in this clause is small. Therefore, we can either obtain the lower bound on the proof size or find a clause of small width that does not contain a collision. In the latter case, we increase the width of the clause by $O(\log S)$. Now, we can do a restart and repeat the process with the start in this clause. Naively, we can do it for $O(\sqrt{N}/\log S)$ iterations and obtain a linear lower bound on the depth of the proof. One can also improve this process by changing the sampling method after the restart and increasing the number of iterations to $O(N/\log S)$, obtaining $O(N^{1.5-\varepsilon})$ depth lower bound (see [BI25]).

As one can see, this method has a fundamental limitation. One random walk iteration can only run for at most N steps because we can simply query all variables. Additionally, one can hope

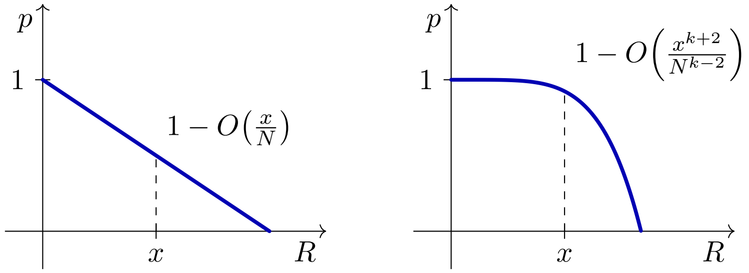

(a) BPHP_N^M (b) $\mathrm{CBPHP}_{k,f}^{N,M}$

Figure 1: Graphs of collision probability p for the new pigeon if R pigeons are already chosen without collisions.

for only a linear number of restarts. As a result, surpassing an N^2-depth lower bound with this technique appears highly unlikely. Bhattacharya and Chattopadhyay [BC25] achieved such a bound using a different formula, along with clever sampling strategies at each step and an amortized closure invariant. Using our formula $\mathrm{CBPHP}_{k,f}^{N,M}$, we can almost match this lower bound using the most naive approach that gives only a linear depth lower bound for BPHP. The detailed explanation of this proof would be given in Section 5. In the following paragraph, we describe our formula and how it gives us an advantage over BPHP.

New formula The *Constrained Bit Pigeon Hole Principle* ($\mathrm{CBPHP}_{k,f}^{N,M}$ for short) is the generalization of BPHP. This formula also works with pigeons and holes, but utilizes them in a different manner.

To better explain our formula, let us first imagine the following scenario: suppose we assigned values to R pigeons out of N in BPHP_N^M without collisions. Now, if we sample the value of any other pigeon uniformly at random, the probability of the collision would be $1 - \frac{R}{N}$.

We construct $\mathrm{CBPHP}_{k,f}^{N,M}$ in a way that the probability of collision in the same scenario is $1 - O\left(\frac{R^{k+2}}{N^{k-2}}\right)$, where $k > 0$ is a constant of our choice (see Figure 1).

In a process where pigeons are assigned to holes one by one, it would require more pigeons to detect a collision with constant probability in the case of our formula compared to BPHP This allows us to do a longer random walk on a parity decision DAG (equivalently, a Res($\oplus$) proof) without getting a collision. As a result, we can apply the random walk with restarts technique and almost immediately get the following lower bound for Res($\oplus$).

To construct such a formula, we take M *greater* than N, and then, to achieve unsatisfiability, we pick a random function $f\colon [M]^k \to [M]$, and for any $k+1$ pigeons out of N with values $p_1, \ldots, p_k, p_{k+1}$ we add to our formula a constraint encoding the fact that

$$f(p_1, \ldots, p_k) \neq p_{k+1}.$$

Polynomial-depth lower bounds Our new formula $\mathrm{CBPHP}_{k,f}^{N,M}$ allows us to implement a new counting technique for proving lower bounds. The new technique is similar in spirit to the previously known technique of random walk with restarts, but does not share the common process of random traversal over the decision DAG. Both techniques count the number of "nice" assignments on each layer of the proof DAG. The random walk with restarts implements it using several iterations of a random walk. In contrast, our technique does it more directly, counting the number of nice assignments on each layer one by one. This allows us to overcome the limitation of N^2 on the maximal depth for which our technique could be applied. However, it also comes with a drawback: a single step is harder to perform.

We propose two methods for proving a lower bound on the total number of nice assignments on each step, where both methods have the same core idea: focus on the number of nice assignments for each clause of "small" width that remain nice in the next level, and then use either a structure of the proof (in the case of RevRes($\oplus$)), or the combinatorial conjecture (in the case of Res($\oplus$)), which we believe to be true. In these two methods, the notion of nice assignments differs slightly. In the case of RevRes($\oplus$), due to the structure of the proof DAG, we count each assignment the same number of times it appears in the clauses on this particular layer. But in the case of Res($\oplus$), we can not do the same, so we use the conjecture to circumvent this issue. We first describe the first method used for RevRes($\oplus$) and then proceed with the Res($\oplus$) case.

Reversible Resolution over parities Introduced by Alekseev and Gaevoy [AG25], Reversible resolution over parities (RevRes($\oplus$), for short) is the following generalization of Reversible Resolution (see [GHJ$^+$24]), which is also a fragment of Res($\oplus$):

- Each line of the refutation in RevRes($\oplus$) is the multiset of affine clauses $C_1, C_2, \ldots, C_m$, where an affine clause is a disjunction of affine equations over $\mathbb{F}_2$.
- During application of any rule, we *replace the clauses from the premises with the clauses from the conclusion.* All other clauses of the current multiset remain unchanged.
- All the derivation rules are reversible in the sense that if we can derive the collection of clauses $\{A_i\}_{i\in r}$ from $\{B_j\}_{j\in q}$, then we can do the inverse: derive $\{B_j\}_{j\in q}$ from $\{A_i\}_{i\in r}$.
- RevRes($\oplus$) uses the resolution rule and its inverse, reversible weakening:

  $$\frac{A}{A \vee (\ell = 0) \quad A \vee (\ell = 1)}$$

  Additionally, we can replace each affine clause A with the semantically equivalent one B.

To prove the lower bound on the number of nice assignments on each layer of RevRes($\oplus$), we prove a lower bound on the number of remaining nice assignments for each particular clause on the layer and then use the structure of our refutation, i.e., the fact that conclusions replace premises, to get a lower bound for the total number of those assignments. For more details, we refer to Section 4.2.

Conditional lower bound To prove a lower bound on the total number of nice assignments in Res($\oplus$), we need to make sure that after we delete a small fraction of nice assignments for each clause, the total number of nice assignments will not decrease by much. This fact is captured in the following conjecture:

Conjecture 1.4. *Let r and q be some constants greater than 0. Let $\Phi_1, \ldots, \Phi_m$ be collections of affine subspaces of $\mathbb{F}_2^n$ of codimention at most $(\log n)^q$ such that*

$$\left| \bigcup_{j\in[m]} \Phi_j \right| \geq 2^{n-1}.$$

For each $i \in [m]$ let Φ'_i be any subset of Φ_i such that

$$|\Phi'_i| \geq \left(1 - \frac{1}{n^r}\right) |\Phi_i|.$$

Then there is a constant $c > 0$, depending only on q, such that

$$\left| \bigcup_{j \in [m]} \Phi'_j \right| \geq \left(1 - \frac{1}{n^{r \cdot c}}\right) \left| \bigcup_{j \in [m]} \Phi_j \right|.$$

The reason why we believe this conjecture is true is the following: any naive counterexample has the structure of a sunflower in the sense that all Φ'_j belong to some core, and $\Phi_j \setminus \Phi'_j$ are petals. Intuitively, it seems like it is impossible to construct such a sunflower using affine subspaces.

Assuming this conjecture, we can prove the superpolynomial lower bound on the size of polynomial depth Res($\oplus$) refutations for any fixed polynomials. For more details, see Section 4.3. Also, note that if we assume the conjectures with better parameters, we can achieve an exponential lower bound for any polynomial depth Res($\oplus$) refutations.

1.3 Organization of the Paper

In Section 2, we define the framework we work with and provide formal definitions of the proof systems (via DAGs). In Section 3, we define the formula we work with and prove that we can sample such a formula. In Section 4, we prove both polynomial-depth lower bounds. Finally, we give a proof of the almost-quadratic lower bound in Section 5.

2 Preliminaries

Definition 2.1 (Falsified clause search problem)**.** Consider an unsatisfiable CNF-formula $\varphi = C_1 \wedge C_2 \wedge \cdots \wedge C_m$. *The falsified clause search problem* is a problem of finding for a given assignment σ find a clause C_i that is falsified by σ.

One of the models for solving falsified clause search problem is parity decision DAG, which is a generalization of the standard model of parity decision tree. This model is also known under the name of linear branching program (see [GPT22, EGI24]). We define it next.

Definition 2.2 (Parity decision DAG)**.** *A parity decision DAG* solving falsified clause search problem for an unsatisfiable formula φ on variables $x_1, x_2, \ldots, x_n$ is a DAG on set of vertices V and set of edges $E = E_d \sqcup E_w$ with the following structure.

- Each vertex $v \in V$ is associated with a linear system L_v over $\mathbb{F}_2$ on variables x_i. We do not distinguish between different linear systems defining the same linear subspace.
- The DAG has exactly one vertex with in-degree equal to 0, which we call *the source*. The linear system corresponding to the source is empty.
- There are two types of edges: *decision* edges E_d and *weakening* edges E_w. All decision edges are mapped into linear equations (which we denote $\ell_e = b_e$, where $b_e \in \mathbb{F}_2$), and the following conditions hold.

- For each directed decision edge $v \to u$, $L_u = L_v \cup \{\ell_e = b_e\}$.
- For each directed weakening edge $v \to u$, L_u is a subsystem of L_v.

- For every vertex v, out-degree of v is at most 2. Moreover,
  - If out-degree is 0, then L_v encodes the negation of some clause C_i of φ. We call such vertex *a sink*.
  - If out-degree is 1, then the only outgoing edge is a weakening edge.
  - If out-degree is 2, then both the outgoing edges e_1 and e_2 are decision edges, and their corresponding equations are $\ell = b$ and $\ell = 1 - b$ for some linear form ℓ.
- Except for the source, each vertex may have either exactly one incoming decision edge or any positive number of incoming weakening edges.

Solving falsified clause search problem with parity decision DAG Using a parity decision DAG for a formula φ we can solve the falsified clause search problem for φ. To do this, we consider the given assignment σ and traverse the DAG with the following procedure starting in the source.

- If the out-degree of the current vertex is 1, we need to traverse the only outgoing edge.
- If the out-degree of the current vertex is 2, we take the linear form ℓ associated with two outgoing decision edges, compute it on our assignment σ and choose the edge, which equation we satisfy.
- Once we reach a sink (i.e. a vertex of out-degree 0), we stop.

Our process maintains the invariant that σ always satisfies the linear system associated with the current vertex. Therefore, when our process inevitably ends in a sink, we will find a clause falsified by σ, which will be our answer.

Equivalence with Res($\oplus$) Our general goal is to prove lower bounds on the number of vertices of parity decision DAG with a small diameter (longest path from the source to a sink) for a particular family of formulas. This problem is equivalent to the problem of finding the lower bounds for bounded-depth Res($\oplus$) proof system (see [EGI24], for example). Each Res($\oplus$) proof can be equivalently rewritten into a parity decision DAG preserving the following metrics.

- The length of Res($\oplus$) proof is equivalent to the number of vertices of the parity decision DAG.
- The depth of the proof is equivalent to the diameter (which we also call depth) of the DAG.
- The width of the proof is equivalent to the maximal size of all linear systems L_v associated with vertices.

Also, we would like to prove lower bounds for bounded depth RevRes($\oplus$) refutations. Those refutations could also be effectively associated with the following DAG.

Definition 2.3 (Reversible parity decision DAG)**.** *A reversible parity decision DAG* for formula φ is a parity decision DAG with the following two additional properties.

- The in-degree of each vertex is at most 2.
- If the in-degree of vertex v is exactly 2, then the pair w_1 and w_2 of parents of v satisfy the following property: there exist some linear form ℓ such that $L_{w_1} = L_v \cup \{\ell = b\}$ and $L_{w_2} = L_v \cup \{\ell = 1 - b\}$ for some constant $b \in \mathbb{F}_2$.

The reversible parity decision DAG has the same relation to the RevRes($\oplus$) proof system as the parity decision DAG relates to Res($\oplus$), maintaining the same complexity measures. The formal definition of RevRes($\oplus$) can be found in [AG25]. The proof of equivalence between RevRes($\oplus$) and reversible parity decision DAG is similar to the proof of equivalence between Res($\oplus$) proof system and parity decision DAG [EGI24]. We provide the formal proof in the full version of the paper.

Remark 2.4. If we add a requirement for each vertex to have the in-degree at most 1, we obtain the parity decision tree, which is equivalent to the tree-like Res($\oplus$) proof system.

Following the definitions from [EGI24], we define *safe* linear systems and *closure* of the linear systems.

2.1 Safe and Dangerous Sets of Linear Forms

For a set of vectors U from a vector space V we denote by $\langle U \rangle$ the linear span of U. We consider the set of propositional variables $X = \{x_{i,j} \mid i \in [m], j \in [\ell]\}$. The variables from X are divided into m blocks by the value of the first index. The variables $x_{i,1}, x_{i,2}, \ldots, x_{i,\ell}$ form the ith block, for $i \in [m]$.

Consider sets of linear forms using variables from X over the field $\mathbb{F}_2$. The *support* of a linear form $f = x_{i_1,j_1} + x_{i_2,j_2} + \cdots + x_{i_k,j_k}$ is the set $\{i_1, i_2, \ldots, i_k\}$ of blocks of variables that appear in f with non-zero coefficients. We denote the support by $\operatorname{supp}(f)$. The support of a set of linear forms F is the union of the supports of all linear forms in this set. We denote it by $\operatorname{supp}(F)$. We say that a linearly independent set of linear forms F is *dangerous* if $|F| > |\operatorname{supp}(F)|$. We say that a set of linear forms F is *safe* if $\langle F \rangle$ does not contain a dangerous set. If F is linearly dependent but $\langle F \rangle$ contains a dangerous set, instead of saying that F is dangerous, we say it is not safe.

Every linear form corresponds to a vector of its coefficients indexed by the variables from the set X. Given a list of linear forms $f_1, f_2, \ldots, f_k$, one may consider their coefficient matrix of size $k \times |X|$ in which the i-th row coincides with the coefficient vector of f_i.

Theorem 2.5 ([EGI24])**.** *Let $f_1, f_2, \ldots, f_k$ be linearly independent linear forms and let M be their coefficient matrix. Then, the following conditions are equivalent.*

(1) The set of linear forms $f_1, f_2, \ldots, f_k$ is safe.

(2) One can choose k blocks and one variable from each of these blocks such that the columns of M corresponding to the k chosen variables are linearly independent.

2.2 Closure

Let $S \subseteq [m]$ be a set of blocks; for a linear form f we denote by $f[\setminus S]$ a linear form obtained from f by substituting 0 for all variables with support in S. For a set of linear forms F we will use the notation $F[\setminus S] = \{f[\setminus S] \mid f \in F\}$.

A *closure* of a set of linear forms F is any inclusion-wise minimal set $S \subseteq [m]$ such that $F[\setminus S]$ is safe.

Lemma 2.6 (Uniqueness [EGI24])**.** *For any F, its closure is unique.*

We denote the closure of F by $\mathrm{Cl}(F)$.

Lemma 2.7 (Monotonicity [EGI24])**.** *If $F_1 \subseteq F_2$, then $\mathrm{Cl}(F_1) \subseteq \mathrm{Cl}(F_2)$.*

Lemma 2.8 (Span invariance [EGI24])**.** $\mathrm{Cl}(F) = \mathrm{Cl}(\langle F \rangle)$.

Lemma 2.9 (Size bound [EGI24])**.** $|\mathrm{Cl}(F)| + \dim\langle F[\setminus \mathrm{Cl}(F)]\rangle \leq \dim\langle F\rangle$*, and hence* $|\mathrm{Cl}(F)| \leq \dim\langle F\rangle$*.*

We would abuse our notation in the sense that for any system of affine equations L with linear part F, we define $\mathrm{rk}(L) \coloneqq \dim\langle F\rangle$ and $\mathrm{Cl}(L) \coloneqq \mathrm{Cl}(F)$.

3 Constrained Bit Prigeon Hole Principle

We are going to consider the following generalization of BPHP, which we will denote as $\mathrm{CBPHP}_{k,f}^{N,M}$:

- Let M be the number of holes, N be the number of pigeons. Note that M is not necessarily less than N in our case. In fact, N would be equal to $M^{1/(k-1)+\varepsilon}$, where k is the constant we would choose later and $\varepsilon > 0$ is an arbitrary small constant of our choice.
- We want to consider a function $f\colon [M]^k \to [M]$ such that for any $S \subseteq [M], |S| = N - 1$, the size of the image of S^k under f is equal to M (i.e. $f(S^k) = [M]$). We will construct this function at random.
- The formula then would be the following: let $M = 2^m$ for some m. The variables would be $x_{i,j}$ where $i \in [N]$ and $j \in [m]$. For any $p_1, \ldots, p_{k+1} \in [N]^k$ such that $p_i \neq p_j$ for $i \neq j$ we write the following statement as a CNF:

  $$f(x_{p_1}, x_{p_2}, \ldots, x_{p_k}) \neq x_{p_{k+1}}$$

  and add it to our formula via conjunction. Also, for any $i \neq j \in [N]$, we add the following predicate via disjunction, expressed in the CNF:

  $$x_i \neq x_j.$$

  Now, we will explain how to write down the statement $f(x_{p_1}, x_{p_2}, \ldots, x_{p_k}) \neq x_{p_{k+1}}$ as CNF ($x_i \neq x_j$ is expressed similarly). Observe that

  $$[f(x_{p_1}, x_{p_2}, \ldots, x_{p_k}) \neq x_{p_{k+1}}] = \bigwedge_{\substack{h_1, \ldots, h_k \in [M]^k \\ f(h_1, \ldots, h_k) = h_{k+1}}} (x_{p_1} \neq h_1 \vee x_{p_2} \neq h_2 \vee \cdots \vee x_{p_{k+1}} \neq h_{k+1}),$$

  where each $x_{p_i} \neq h_i$ can be expressed bitwise as

  $$x_{p_i,1} \neq h_{i,1} \vee x_{p_i,2} \neq h_{i,2} \vee \cdots \vee x_{p_i,m} \neq h_{i,m},$$

  where $h_{i,1}, h_{i,2}, \ldots, h_{i,m}$ is a bit representation of h_i. Note that each expression $x_{p_i,j} \neq h_{i,j}$ is represented by a single literal.

- As we can see, the size of the formula is polynomial in M and N if k is a constant. This CNF is unsatisfiable for the following reason: let us consider the first $N-1$ pigeons. Suppose we assigned them different values from the set S of size $N-1$ (if two values coincide, then we get a contradiction immediately). Then, no matter which value we assign to the last pigeon, it clearly belongs to $f(S^k) = [M]$ and thus produces a contradiction to the formula.

3.1 Sampling the Function f

Let $\varepsilon > 0$ be a constant and $N \geq M^{1/(k-1)+\varepsilon}$. In this section, we show that each of the following two properties holds with probability greater than $1/2$ for a uniformly random f while M is large enough:

(i) The formula $\mathrm{CBPHP}_{k,f}^{N,M}$ is unsatisfiable.

(ii) Let us consider the smallest relation Sym_f such that

$$\forall \pi \in S_{k+1} \forall (h_1, h_2, \ldots, h_k) \in [M]^k : \pi(h_1, h_2, \ldots, h_k, f(h_1, h_2, \ldots, h_k)) \in \mathrm{Sym}_f.$$

For any collection of values $\gamma_1, \gamma_2, \ldots, \gamma_k \in [M]$ holds

$$|y : (h_1, h_2, \ldots, h_k, y) \in \mathrm{Sym}_f| \leq (\log M)^k,$$

Unsatisfiability probability lower bound First, let us denote $N_0 \coloneqq N-1$ for simplicity. Consider any fixed $S \subseteq [M], |S| = N_0$. We want to analyze the probability that $|f(S^k)| = M$ and use the union bound over all possible S.

This probability is the same as the following one: let $X_1, X_2, \ldots, X_{N_0^k}$ be i.i.d. variables distributed uniformly over $[M]$. We want to prove an upper bound on the probability that $|\{X_1, X_2, \ldots, X_{N_0^k}\}| \leq M-1$. We know that

$$\mathbb{E}\left[\left|\{X_1, X_2, \ldots, X_{N_0^k}\}\right|\right] = M \cdot \left(1 - (1-1/M)^{N_0^k}\right) \geq M \cdot \left(1 - e^{-N_0^k/M}\right).$$

Let us denote $N_0 \coloneqq M^\alpha$. By Markov's inequality, we get that

$$\Pr\left[\left|\{X_1, X_2, \ldots, X_{N_0^k}\}\right| \leq M-1\right] \leq M \cdot \exp\left(-M^{k\cdot\alpha-1}\right).$$

Now, the number of different S is equal to $\binom{M}{M^\alpha} \leq e^{\log M \cdot M^\alpha}$. So, altogether we want to find α such that

$$M \cdot \exp\left(-M^{k\cdot\alpha-1}\right) \cdot \exp(\log M \cdot M^\alpha) \leq \exp\left(\log M + \log M \cdot M^\alpha - M^{k\cdot\alpha-1}\right) < 1/2.$$

So, if we take $\alpha > 1/(k-1)$, we would get the desired bound.

Symmetrization of f and size of the image Sym_f consists of permutations of all the possible inputs together with the output of the function on these inputs. We want to show that with high probability for the random function $f : [M]^k \to [M]$ and for all $x_1, x_2, \ldots, x_k \in [M]^k$ we can prove an uniform upper bound on number of y's such that $(x_1, x_2, \ldots, x_k, y) \in \mathrm{Sym}_f$. This again can be done with a union bound. Let us fix $x_1, x_2, \dots, x_k \in [M]^k$ and the position of the variable y. Let us estimate the probability that for at least t different y's holds

$$f(x_1, \dots, x_i, y, x_{i+1}, \dots, x_{k-1}) = x_k.$$

This probability is equal to the probability that the sum of i.i.d. Bernoulli distributions $X_1, \dots, X_M$ is at most t, where $X_i = 0$ with probability $1 - 1/M$ and $X_i = 1$ with probability $1/M$. This probability is at most e^{-t} by Multiplicative Chernoff's inequality (for $t > 7$). So, if we take $t = (\log M)^k$, then probability that for all $x_1, x_2, \dots, x_k \in [M]^k$ the number of y's such that $(x_1, x_2, \dots, x_k, y) \in \mathrm{Sym}_f$ is at most t can by lower bounded by

$$1 - M^k \cdot k! \cdot e^{-t/k!} > 1/2.$$

This means that we can find a function $f : [M]^k \to [M]$ for which our formula is unsatisfiable, and at the same time for all $x_1, x_2, \dots, x_k \in [M]^k$ the number of y's such that $(x_1, x_2, \dots, x_k, y) \in \mathrm{Sym}_f$ is at most $(\log M)^k$.

4 $\mathrm{poly}(n)$-Depth Lower Bounds for $\mathrm{CBPHP}_{k,f}^{N,M}$.

For the rest of this section, we will fix the parameter k. For each N we will take $M = N^{k-1-\varepsilon}$, where $\varepsilon > 0$ is a constant of our choice (we can take an arbitrarily small one). We consider a large enough N such that there exists a function $f\colon [M]^k \to [M]$ satisfying properties (i) and (ii) from Section 3.1. We want to prove the following theorem first, and then explain how to modify it to get a conditional polynomial-depth $\mathrm{Res}(\oplus)$ lower bound.

Theorem 4.1. *Any $\mathrm{RevRes}(\oplus)$ refutation of $\mathrm{CBPHP}_{k,f}^{N,M}$ of size S requires depth d such that*

$$d \geq \Omega\left(\frac{N^{k-1-\varepsilon}}{\left(\log S + (\log N)^{O(1)}\right)^{2k+2}}\right).$$

In particular, if $S \leq \exp\left(N^{1/(2k+3)}\right)$, then $d \geq \Omega(N^{k-2})$.

Moreover, we propose the following combinatorial conjecture, which implies the polynomial-depth $\mathrm{Res}(\oplus)$ lower bound.

Conjecture 4.2. *Let r and q be some constants greater than 0. Let $\Phi_1, \dots, \Phi_m$ be the collection of affine subspaces of $\mathbb{F}_2^n$ of codimention at most $(\log n)^q$ such that*

$$\left|\bigcup_{j\in[m]} \Phi_j\right| \geq 2^{n-1}.$$

For each $i \in [m]$ let Φ_i' be any subset of Φ_i such that

$$|\Phi_i'| \geq \left(1 - \frac{1}{n^r}\right)|\Phi_i|.$$

Then there is a constant $c(q) > 0$ depending only on q such that

$$\left|\bigcup_{j\in[m]} \Phi_j'\right| \geq \left(1 - \frac{1}{n^{r\cdot c}}\right)\left|\bigcup_{j\in[m]} \Phi_j\right|.$$

Theorem 4.3. *Suppose that Conjecture 4.2 holds for $q > 1$ and $r = k - 2$ with some parameter $c(q) > 0$. Then any* $\mathrm{Res}(\oplus)$ *refutation of* $\mathrm{CBPHP}^{N,M}_{k,f}$ *of size at most $2^{(\log(N))^{q/2}}$ requires depth d such that*

$$d \geq \Omega\left(\left(\frac{N^{k-1-\varepsilon}}{(\log N)^{2kq+O(1)}}\right)^c\right).$$

where c only depends on q.

Note 4.4. If we assume the conjecture for any constant q, it will imply a tradeoff between superpolynomial size and polynomial depth.

Before proving those theorems, we want to define the key concept we will be working with and explain the high-level idea behind the proofs.

4.1 Subsets of "Nice" Assignments and Main Tool

Definition 4.5 (Nice substitutions)**.** Let L be a system of linear equations over the variables of $\mathrm{CBPHP}^{N,M}_{k,f}$. Let $\gamma \in \{0,1\}^{N \times \log M}$ be a substitution into those variables that satisfies L. We abuse the notation and denote by $\gamma_i \in [M]$ the value of the pigeon with index $i \in [N]$.

We would say that γ is *nice* if it does not have any collisions on the pigeons from $\mathrm{Cl}(L)$. Formally, this means that $\gamma_{p_1} \neq \gamma_{p_2}$ for any pigeons $p_1 \neq p_2 \in \mathrm{Cl}(L)$ and $f(\gamma_{p_1}, \ldots, \gamma_{p_k}) \neq \gamma_{p_{k+1}}$ for any distinct $p_1, \ldots, p_{k+1} \in \mathrm{Cl}(L)$.

High-level proofs' idea We divide the parity decision DAG into layers. To do it, we compute for each vertex the length of the longest path with a start in the source that ends in this vertex. As the contents of the layer t, we define all vertices that are sinks from previous levels or have the computed distance of at least t and a parent with the computed distance of at most t. In other words, we split all vertices into levels by the maximal distance from the source, move all sinks to the last level, and then define layers as the contents of a level and the ends of all edges that pass through this level.

- Consider each layer of the decision DAG for $\mathrm{RevRes}(\oplus)$ (or $\mathrm{Res}(\oplus)$). On the top layer, we have a system $L_0 = \emptyset$. Let K_1 be the number of nice assignments for L_0. Since any substitution is nice for L_0, this number is M^N.
- We define K_i differently depending on the lower bound we are currently proving
  - In the case of $\mathrm{RevRes}(\oplus)$, let K_i be the sum of the numbers of nice assignments for each of the systems on layer i.
  - In the case of $\mathrm{Res}(\oplus)$, let K_i be the total number of assignments that are nice for at least one system L on layer i.

  The difference between the two metrics is motivated by the fact that the number of vertices associated with equivalent linear systems plays a crucial role for reversible parity decision DAG (similarly, using duplicate clauses is crucial for $\mathrm{RevRes}(\oplus)$ proofs). In contrast, in the non-reversible case, all such vertices can be combined into one in a single step using weakening edges.

We show that we can choose the parameter $\delta(k)$ growing to the ∞ with growth of k, if the size of the RevRes($\oplus$) (or Res($\oplus$)) refutation is bounded by some S, then

$$K_{i+1} \geq K_i \cdot \left(1 - \left(\frac{\operatorname{poly}\log S}{N}\right)^{\delta}\right)$$

if $K_i > 1/2 \cdot M^N$. In the case of Res($\oplus$), this lower bound would be conditional.

- Since for the bottom layer d we know that $K_d = 0$, we get that

$$d \geq \Omega\left(\frac{N}{\operatorname{poly}\log S}\right)^{\delta}.$$

Before proving the main result, we would need the following lemma, which we will use in both RevRes($\oplus$) and Res($\oplus$) lower bounds.

Lemma 4.6. *Let w be a parameter less than N. Let L be a linear form of rank at most w and let Φ be the collection of nice assignments in L. Let P be any set of pigeons such that $|P| \leq w$ and $\mathrm{Cl}(L) \subseteq P$. Let Φ' be the subset of the assignments Φ that do not have any collisions on P. Then*

$$|\Phi'| \geq |\Phi| \cdot \left(1 - \frac{(3w)^{k+1} \cdot (\log N)^{O(1)}}{N^{k-1-\varepsilon}}\right).$$

Proof of Lemma. Consider any partial assignment ρ on $\mathrm{Cl}(L)$ such that ρ can be extended to a satisfying assignment of L. Let Φ'_ρ and Φ_ρ be the full assignments from Φ' and Φ, which are consistent with ρ. We will show that

$$|\Phi'_\rho| \geq |\Phi_\rho| \cdot \left(1 - \frac{(3w)^{k+1} \cdot (\log N)^{O(1)}}{N^{k-1-\varepsilon}}\right).$$

This inequality implies the initial one, since the sets Φ_ρ do not intersect for different values of ρ.

Let L_ρ be a substitution of ρ into L. We know that L_ρ is *safe and satisfiable*. Let $w_0 = \mathrm{rk}(L_\rho) \leq w$. Consider a uniform distribution over Φ_ρ. Not that Φ_ρ is simply the set of all solutions of L_ρ. We can uniformly sample Φ_ρ with the following procedure:

- Choose at most one variable in a block such that our system L_ρ can be represented in the following way:

$$x_{p_1,j_1} = \ell_1, x_{p_2,j_2} = \ell_2, \ldots, x_{p_{w_0},j_{w_0}} = \ell_{w_0},$$

  where each of ℓ_r does not contain variables from $\{x_{p_z,j_z}\}_{z\in[w_0]}$ and all p_i are distinct. This is possible since L_ρ is safe.

- Sample values for the Boolean variables corresponding to the pigeons from $[N] \setminus \mathrm{Cl}(L)$ excluding $\{x_{p_z,j_z}\}_{z\in[w_0]}$ independently and uniformly at random.

- Assign the values for the variables from $\{x_{p_z,j_z}\}_{z\in[w_0]}$ according to L_ρ.

Now, we show that with a high enough probability, even after the second stage of this sampling, we will get a partial assignment such that no matter how we assign the values for $\{x_{p_z,j_z}\}_{z\in[w_0]}$, we would not get a collision on pigeons from P. Indeed, on the second step, we uniformly independently sample the pairs of values for each pigeon from $P \setminus \mathrm{Cl}(L)$ (the pair of the values is connected via the value of x_{p_z,j_z}). Let us sample these pairs one by one $|P \setminus \mathrm{Cl}(L)|$ times and upper bound the probability that the next pair contains a value that can produce a collision.

Let Q be the values of pigeons from ρ, and H be the union of all chosen pairs so far. The probability that one of the elements from the next pair forms a collision with $Q \cup H$ is at most

$$\frac{(|Q|+|H|)^k \cdot (\log N)^{O(1)}}{M/2}.$$

Indeed, there are at most $\binom{|Q|+|H|}{k} \le (|Q|+|H|)^k$ choices of k values $h_1, \dots, h_k$ among $Q \cup H$ and by the choice of f we know that

$$|y : (h_1, h_2, \dots, h_k, y) \in \mathrm{Sym}_f| \le (\log N)^{O(1)},$$

where Sym_f is the relation constructed by symmetrization of f. Now, since we know that $|Q| \le w$ and $|H| \le 2w$ and since we repeat this process $|P \setminus \mathrm{Cl}(L)| \le w$ times, we can bound the collision probability as

$$\frac{(3w)^{k+1} \cdot (\log N)^{O(1)}}{M/2}.$$

So, since $M = N^{k-1-\varepsilon}$, we get the desired bound

$$|\Phi'_\rho| \ge |\Phi_\rho| \cdot \left(1 - \frac{(3w)^{k+1} \cdot (\log N)^{O(1)}}{N^{k-1-\varepsilon}}\right). \qquad \square$$

4.2 Proof of Theorem 4.1

Suppose that $K_i \ge 1/2$. We want to prove a lower bound on K_{i+1}, so we consider what happens with each system of linear equations L on the layer $i+1$.

- On the layer $i+1$, system L stays the same. Then the amount of nice assignments for this system does not change.
- On the layer $i+1$, system $L_0 = L \cup (\ell = 0)$ is merged with some system $L_1 = L \cup (\ell = 1)$ into system L. Then L contains all the nice assignments from both L_0 and L_1, and those sets do not intersect. This means that the number of nice assignments cannot decrease for this pair of systems.
- On the layer $i+1$, system L is splitted into $L_0 = L \cup (\ell = 0)$ and $L_1 = L \cup (\ell = 1)$. Let $P = \mathrm{Cl}(L_0) = \mathrm{Cl}(L_1)$ and Φ_L be nice assignments for L. Let Φ'_L be the subset of the assignments Φ_L such that all the assignments from Φ'_L do not have any collisions on P. We know that $|\Phi'_L|$ is equal to the number of nice substitutions in L_0 and L_1.

  Lemma 4.6 implies that

  $$|\Phi'_L| \ge |\Phi_L| \cdot \left(1 - \frac{(3w)^{k+1} \cdot (\log N)^{O(1)}}{N^{k-1-\varepsilon}}\right),$$

  where w is the rank of L. Now, we have two subcases here:

i If $w \leq \left(\log S + (\log N)^{O(1)}\right)^2$, then the decrease in the number of nice assignments is at most $\frac{(\log S + (\log N)^{O(1)})^{2k+2}}{N^{k-1-\varepsilon}}$ for this particular linear system.

ii If $w > \left(\log S + (\log N)^{O(1)}\right)^2$, then $|\Phi_L| < \frac{M^N}{2^w} \leq \frac{M^N}{(SN^k)^2}$. Since we have at most S different systems in our refutation and $|K_i|/M^N \geq 1/2$, all such clauses would only give us a $1/(SN^{2k})$ decrease in the total size.

Altogether, if $K_i > \frac{M^N}{2}$ we get that

$$K_{i+1} \geq K_i \cdot \left(1 - \frac{\left(\log S + (\log N)^{O(1)}\right)^{2k+2}}{N^{k-1-\varepsilon}}\right).$$

Since $K_1 = M^N$ and $K_d = 0$, this means that the total number of layers d should be at least

$$\Omega\left(\frac{N^{k-1-\varepsilon}}{\left(\log S + (\log N)^{O(1)}\right)^{2k+2}}\right).$$

In particular, if $S \leq \exp\left(N^{1/(2k+3)}\right)$, then $d \geq \Omega(N^{k-2})$.

4.3 Proof of Theorem 4.3

Suppose $S \leq 2^{(\log N)^{q/2}}$. Let $L_1, L_2, \ldots, L_g$ be the collection of linear systems on layer i. For each $j \in [g]$ let Φ_j be the collection of nice assignments for L_j. So, the $K_i = \left|\bigcup_{j \in [g]} \Phi_j\right|$. We assume that $K_i \geq \frac{M^N}{2}$. Note that Φ_j is not necessarily an affine subspace of $\mathbb{F}_2^{N \cdot \log M}$. However, for each $j \in [g]$, we can split Φ_j into the disjoint union of *affine subspaces* Φ_j^ρ of codimension at most $\mathrm{rk}(L_j)$ where each ρ is the assignment of the variables from $\mathrm{Cl}(L_j)$ that does not have any collisions and Φ_j^ρ is the subset of the assignments from Φ_j which are consistent with ρ.

First, observe that the union of the subspaces Φ_j corresponding to the systems L_j with $\mathrm{rk}(L_j) > (\log S + \log N)^2$ have size at most $S \cdot \frac{M^N}{2^{(\log S + \log N)^2}} \leq K_i/N^{2k}$. This means that the union of these subspaces cannot contain more than K_i/N^{2k} nice assignments.

Now, let $L_1, L_2, \ldots, L_g$ be the systems with rank at most $(\log S + \log N)^2$. We want to estimate the number of nice substitutions that are nice for one of the systems L_j, but are not nice in the layer $i+1$. Let $K_i' = |\bigcup_{j \in [g]} \Phi_j|$. For each L_j that is splited by the decision edges on the affine form ℓ_j and an assignment ρ on $\mathrm{Cl}(L)$, that does not contain collisions, let Ψ_j^ρ be the collection of full assignments from Φ_j^ρ which do not have any collisions on $L_j \cup \{\ell_j = \beta\}$. For any other L_j, we know that all the nice substitutions for L_j would also be nice on layer $i+1$ for some system. So, for those systems, we define $\Psi_j^\rho = \Phi_j^\rho$. Let $K_{i+1}' = \left|\bigcup_{j \in [g], \rho} \Psi_j^\rho\right|$.

Then, we know that

$$K_{i+1} - K_i \leq K_{i+1}' - K_i' + \frac{K_i}{N^{2k}}.$$

So, we want to estimate $K_{i+1}' - K_i'$. For $j \in [g]$, by Lemma 4.6 and $S \leq 2^{(\log N)^{q/2}}$ we get that

$$|\Psi_j^\rho| \geq |\Phi_j^\rho| \cdot \left(1 - \frac{(\log N)^{2kq+O(1)}}{N^{k-1-\varepsilon}}\right).$$

So, if we assume the Conjecture 4.2, we get that

$$\left|\bigcup_{j\in[g],\rho} \Psi_j^\rho\right| \geq \left|\bigcup_{j\in[g],\rho} \Phi_j^\rho\right| \cdot \left(1 - \left(\frac{(\log N)^{2kq+O(1)}}{N^{k-1-\varepsilon}}\right)^c\right)$$

for some parameter $c(q)$ from Conjecture 4.2 such that $0 < c(q) \leq 1$. So, in total, we get that

$$K_{i+1} - K_i \leq \left(\frac{(\log N)^{2kq+O(1)}}{N^{k-1-\varepsilon}}\right)^c \cdot K_i + \frac{K_i}{N^{2k}} \leq 2\left(\frac{(\log N)^{2kq+O(1)}}{N^{k-1-\varepsilon}}\right)^c \cdot K_i.$$

So, the refutation in Res($\oplus$) should contain at least d layers, where

$$d \geq \Omega\left(\left(\frac{N^{k-1-\varepsilon}}{(\log N)^{2kq+O(1)}}\right)^c\right).$$

5 Almost n^2-Depth Lower Bound

For the purposes of this section, we need to introduce one more metric for linear systems.

Definition 5.1 (Safe rank)**.** *The safe rank* of an affine system L ($\mathrm{srk}(L)$) is the sum of the size of the closure of A and the rank of the remaining safe system. Formally, if we pick any solution σ of A and consider partial assignment $\pi = \sigma|_{\mathrm{Cl}(L)}$, then

$$\mathrm{srk}(L) = |\mathrm{Cl}(L)| + \mathrm{rk}(L|_\pi).$$

By Lemma 2.9 we know that $\mathrm{srk}(L) \leq \mathrm{rk}(L)$. It is important to note that the safe rank of linear systems is not a monotone metric, in the sense that adding linear equations to a system may decrease its safe rank.

Random walk with restarts To prove an almost n^2-depth lower bound for Res($\oplus$), we want to do a random walk with restarts. The random walk with restarts works by repeating two phases. The first phase is the random walk. We start in a system L with a small (less than $O(N^{\frac{k-2}{k+2}})$) safe rank and traverse the DAG randomly and consistently with L and some assignment to $\mathrm{Cl}(L)$. We show that if we make $O(N^{\frac{k-2}{k+2}})$ steps of the random walk, we do not find a collision and thus do not end up in a sink with high probability.

Then, we can run the second phase: restart. Among the vertices reached by the random walk, we can find a vertex z associated with a linear system that could be strengthened to a system L' with $\mathrm{srk}(L') \leq \mathrm{srk}(L) + O(\log S)$. Now we can take the vertex z and consider the subgraph of the decision parity DAG that has a source at z, induced on the vertices reachable from z. Then we add some linear equations for all vertices in this subgraph to make L_z equal to L', which allows us to repeat our process from the first phase, starting at vertex z. Note that this operation does not increase the size or depth of the graph.

One iteration of our process increases the safe rank of the system by at most $O(\log S)$, which means that we can repeat it $\Omega\left(N^{\frac{k-2}{k+2}}/\log S\right)$ times. Hence, we obtain the lower bound on the depth of $\Omega\left(N^{2\frac{k-2}{k+2}}/\log S\right)$.

Formally, we prove the following theorem.

Theorem 5.2. *Any* $\mathrm{Res}(\oplus)$ *refutation of* $\mathrm{CBPHP}_{k,f}^{N,M}$ *of size* S *requires depth* d *such that*

$$d \geq \Omega\left(N^{2\frac{k-2}{k+2}} / \log S\right).$$

To prove this theorem, we need the following lemma.

Lemma 5.3. *Let* $w \leq N^{\frac{k-2}{k+2}}$. *Let* L *be a system of linear equations such that the following holds:*

- *There is a solution* π *of this system, such that partial assignment* $\pi_0 = \pi|_{\mathrm{Cl}(L)}$ *does not contain any collisions.*
- $\mathrm{srk}(L) = |\mathrm{Cl}(L)| + \mathrm{rk}(L|_{\pi_0}) \leq w$.

Consider any parity decision tree T *of depth* w. *Let* $\boldsymbol{\sigma}$ *be a uniformly random full assignment of variables of* $\mathrm{CBPHP}_{k,f}^{N,M}$, *consistent with* π_0 *and satisfying* L. *Let* $L_{\boldsymbol{\sigma}}$ *be a system in the leaf of* T *corresponding to the assignment* $\boldsymbol{\sigma}$. *Then for a large enough* N

$$\Pr\left[\textit{there is a collision at } \boldsymbol{\sigma}|_{\mathrm{Cl}(L \cup L_{\boldsymbol{\sigma}})}\right] \leq \delta$$

for some fixed constant $\delta > 0$.

Proof. We would need the following lemma from [AI25].

Lemma 5.4 ([AI25])**.** *Consider a binary tree with root* r *and a set of leaves* L. *We associate every node* v *except the leaves with a number* $p_v \neq 0$. *For every node* v *of the tree, there is a number* n_v *such that if* u *and* w *are children of* v, *then* $n_v p_v = (n_u + n_w)$. *Let for every leaf* l *the unique path from the root to* l *be denoted* $\pi_l = (s_1 = r, s_2, \ldots, s_t = l)$; *let us denote* $p(\pi_l) = \prod_{i=1}^{t-1} p_{s_i}$. *Then* $n_r = \sum_{l \in L} n_l \frac{1}{p(\pi_l)}$.

For each v, let n_v be the total number of full assignments σ which are nice for $L \cup L_v$ (see Section 4.1). Then, for each vertex of the tree v, the probability p_v is equal to the following:

$$p_v = \Pr[\boldsymbol{\sigma} \text{ do not have a collision on } \mathrm{Cl}(L \cup L_u) | \boldsymbol{\sigma} \text{ is nice for } L \cup L_v],$$

where u is any of the children of v. So, for each path π_ℓ we want to prove an lower bound on $\prod_{i=1}^{w-1} p_{s_i}$. This can be done by Lemma 4.6. More precisely, for each $i \in [w]$ let $c_i = |\mathrm{Cl}(L \cup L_{s_i})|$. Then, since $\mathrm{rk}(L \cup L_{s_i}) \leq w \cdot \log M$, by Lemma 4.6 we get that for large enough N

$$p_{s_i} \geq 1 - \frac{(3c_{i+1})^{k+1} \cdot (\log N)^{O(1)}}{N^{k-1-\varepsilon}} \geq 1 - \frac{(3N^{\frac{k-2}{k+2}})^{k+1} \cdot (\log N)^{O(1)}}{N^{k-1-\varepsilon}} \geq$$
$$1 - \frac{N^{k-2}}{N^{k-1-\varepsilon}} = 1 - N^{-(1-\varepsilon)}.$$

So, for $\prod_{i=1}^{t-1} p_{s_i}$ we get the following lower bound:

$$\prod_{i=1}^{w-1} p_{s_i} \geq \left(1 - N^{-(1-\varepsilon)}\right)^w \geq 1 - w \cdot N^{-(1-\varepsilon)} \geq 1 - N^{\frac{k-2}{k+2} + \varepsilon - 1} \geq 1/2$$

for large enough N while ε is a small enough constant.

□

Proof of Theorem 5.2. Consider an iteration of the random walk procedure. We start at a vertex associated with a system L with $\mathrm{srk}(L) \leq N^{\frac{k-2}{k+2}}$ which admits a nice partial solution π_0. Now we can choose the random assignment $\boldsymbol{\sigma}$ consistent with π_0 and L and traverse our DAG according to $\boldsymbol{\sigma}$ for $N^{\frac{k-2}{k+2}}$ steps. This random walk succeeds with at least constant probability due to Lemma 5.3 and the fact that we can transform a parity decision DAG into a parity decision tree by duplicating vertices. Among the vertices where our random walk may stop, we want to choose a vertex z such that

- L_z admits a nice solution σ, consistent with π_0 and L.
- $\mathrm{rk}(L_z|_{\pi_0}) \leq \mathrm{rk}(L|_{\pi_0}) + O(\log S)$.

Such a vertex z exists since we have at most S different systems in the leaves of our random walk, and for any system L_w in the leaf of our random walk, probability that $\boldsymbol{\sigma}$ end up in L_w is at most $2^{\mathrm{rk}(L|_{\pi_0})-\mathrm{rk}(L_w|_{\pi_0})}$.

Now, let π_1 be a restriction on $\mathrm{Cl}(L_z|_{\pi_0})$ of the nice solution σ of L_z. We want to prove that $L' = L_z \cup \pi_0 \cup \pi_1$[2] satisfies the following properties:

- $L'_{\pi_0 \cup \pi_1}$ is safe since $(L_z|_{\pi_0})|_{\pi_1}$ is safe.
- Since π_0 and π_1 are full assignments of the corresponding blocks, $\mathrm{Cl}(L') = \mathrm{Vars}(\pi_0 \cup \pi_1)$.

Lemma 2.9 allows us to write the following inequality:

$$\begin{aligned}
&\mathrm{srk}(L') = |\pi_0| + |\pi_1| + \mathrm{rk}(L_z|_{\pi_0 \cup \pi_1}) = \\
&|\mathrm{Cl}(L)| + |\mathrm{Cl}(L_z|_{\pi_0})| + \mathrm{rk}(L_z|_{\pi_0 \cup \pi_1}) \leq |\mathrm{Cl}(L)| + \mathrm{rk}(L_z|_{\pi_0}) \leq \\
&|\mathrm{Cl}(L)| + \mathrm{rk}(L|_{\pi_0}) + O(\log S).
\end{aligned}$$

So, we can do another iteration of our random walk, starting with a clause L'. We can repeat the process for $\Omega(N^{\frac{k-2}{k+2}} / \log S)$ times, after which reach a vertex on the depth

$$\Omega\left(N^{2\frac{k-2}{k+2}} / \log S\right). \qquad \square$$

Acknowledgements

We want to thank Yuval Filmus for the fruitful discussions.

References

[ABSRW04] Michael Alekhnovich, Eli Ben-Sasson, Alexander A. Razborov, and Avi Wigderson. Pseudorandom generators in propositional proof complexity. *SIAM Journal on Computing*, 34(1):67–88, 2004.

[AG25] Yaroslav Alekseev and Nikita Gaevoy. Intersection theorems: A potential approach to proof complexity lower bounds. *Electron. Colloquium Comput. Complex.*, TR25-160, 2025.

[2] Here we view π_i as a system of equations $x_{i,j} = \gamma_{i,j}$, where $\gamma_{i,j} \in \{0, 1\}$

[AI25] Yaroslav Alekseev and Dmitry Itsykson. Lifting to bounded-depth and regular resolutions over parities via games. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC '25, page 584–595, New York, NY, USA, 2025. Association for Computing Machinery.

[Ajt88] Miklós Ajtai. The complexity of the pigeonhole principle. *Combinatorica*, 14:417–433, 1988.

[BC25] Sreejata Bhattacharya and Arkadev Chattopadhyay. Exponential lower bounds on the size of reslin proofs of nearly quadratic depth. *Electron. Colloquium Comput. Complex.*, TR25-106, 2025.

[BCD24] Sreejata Kishor Bhattacharya, Arkadev Chattopadhyay, and Pavel Dvorák. Exponential separation between powers of regular and general resolution over parities. In Rahul Santhanam, editor, *39th Computational Complexity Conference, CCC 2024, July 22-25, 2024, Ann Arbor, MI, USA*, volume 300 of *LIPIcs*, pages 23:1–23:32. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.

[BI25] Farzan Byramji and Russell Impagliazzo. Lower bounds for the bit pigeonhole principle in bounded-depth resolution over parities. *Electron. Colloquium Comput. Complex.*, TR25-118, 2025.

[BK23] Paul Beame and Sajin Koroth. On Disperser/Lifting Properties of the Index and Inner-Product Functions. In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, volume 251 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 14:1–14:17, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[BSW01] Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow—resolution made simple. *J. ACM*, 48(2):149–169, March 2001.

[CMSS23] Arkadev Chattopadhyay, Nikhil S. Mande, Swagato Sanyal, and Suhail Sherif. Lifting to parity decision trees via stifling. In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference, ITCS 2023, January 10-13, 2023, MIT, Cambridge, Massachusetts, USA*, volume 251 of *LIPIcs*, pages 33:1–33:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.

[CR79] Stephen A. Cook and Robert A. Reckhow. The relative efficiency of propositional proof systems. *Journal of Symbolic Logic*, 44(1):36–50, 1979.

[EGI24] Klim Efremenko, Michal Garlík, and Dmitry Itsykson. Lower bounds for regular resolution over parities. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, June 24-28, 2024*, pages 640–651. ACM, 2024. The full version is available as ECCC technical report TR23-187.

[EI25] Klim Efremenko and Dmitry Itsykson. Amortized closure and its applications in lifting for resolution over parities. *Electron. Colloquium Comput. Complex.*, TR25-039, 2025.

[GHJ$^+$24] Mika Göös, Alexandros Hollender, Siddhartha Jain, Gilbert Maystre, William Pires, Robert Robere, and Ran Tao. Separations in proof complexity and TFNP. *J. ACM*, 71(4), August 2024.

[GK18] Michal Garlík and Leszek Aleksander Kolodziejczyk. Some subsystems of constant-depth frege with parity. *ACM Trans. Comput. Log.*, 19(4):29:1–29:34, 2018.

[GOR24] Svyatoslav Gryaznov, Sergei Ovcharov, and Artur Riazanov. Resolution over linear equations: Combinatorial games for tree-like size and space. *ACM Trans. Comput. Theory*, jul 2024. Just Accepted.

[GPT22] Svyatoslav Gryaznov, Pavel Pudlák, and Navid Talebanfard. Linear branching programs and directional affine extractors. In Shachar Lovett, editor, *37th Computational Complexity Conference, CCC 2022, July 20-23, 2022, Philadelphia, PA, USA*, volume 234 of *LIPIcs*, pages 4:1–4:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

[Gry19] Svyatoslav Gryaznov. Notes on resolution over linear equations. In René van Bevern and Gregory Kucherov, editors, *Computer Science - Theory and Applications - 14th International Computer Science Symposium in Russia, CSR 2019, Novosibirsk, Russia, July 1-5, 2019, Proceedings*, volume 11532 of *Lecture Notes in Computer Science*, pages 168–179. Springer, 2019.

[HR25] Johan Håstad and Kilian Risse. On bounded depth proofs for tseitin formulas on the grid; revisited. *SIAM Journal on Computing*, 54(5):FOCS22–288–FOCS22–339, 2025.

[Hå20] Johan Håstad. On small-depth frege proofs for tseitin for grids. *J. ACM*, 68(1), November 2020.

[Hå23] Johan Håstad. On small-depth Frege proofs for PHP . In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 37–49, Los Alamitos, CA, USA, November 2023. IEEE Computer Society.

[IR21] Dmitry Itsykson and Artur Riazanov. Proof complexity of natural formulas via communication arguments. In Valentine Kabanets, editor, *36th Computational Complexity Conference, CCC 2021, July 20-23, 2021, Toronto, Ontario, Canada (Virtual Conference)*, volume 200 of *LIPIcs*, pages 3:1–3:34. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.

[IS14] Dmitry Itsykson and Dmitry Sokolov. Lower bounds for splittings by linear combinations. In Erzsébet Csuhaj-Varjú, Martin Dietzfelbinger, and Zoltán Ésik, editors, *Mathematical Foundations of Computer Science 2014 - 39th International Symposium, MFCS 2014, Budapest, Hungary, August 25-29, 2014. Proceedings, Part II*, volume 8635 of *Lecture Notes in Computer Science*, pages 372–383. Springer, 2014.

[IS20] Dmitry Itsykson and Dmitry Sokolov. Resolution over linear equations modulo two. *Ann. Pure Appl. Log.*, 171(1), 2020.

[Kra18] Jan Krajíček. Randomized feasible interpolation and monotone circuits with a local oracle. *J. Math. Log.*, 18(2):1850012:1–1850012:27, 2018.

[Urq87] A. Urquhart. Hard examples for resolution. *JACM*, 34(1):209–219, 1987.