Electronic Colloquium on Computational Complexity, Report No. 188 (2025)

Strong ETH Holds for Bounded-Depth Resolution over Parities

Klim Efremenko*[1] and Dmitry Itsykson†[1,2]

[1]Ben-Gurion University of the Negev, Israel
[2]On leave from Steklov Institute of Mathematics at St. Petersburg

November 20, 2025

Abstract

Strong lower bounds of the form $2^{(1-\epsilon)n}$, where n is the number of variables and $\epsilon > 0$ is arbitrarily small (i.e., bounds consistent with the Strong ETH), are exceptionally rare in proof complexity. The seminal work of Beck and Impagliazzo (STOC 2013) achieved such a bound for regular resolution, and the strongest extension known prior to our work was proved for $O(\epsilon)$-regular resolution by Bonacina and Talebanfard (Algorithmica, 2017).

We establish similar lower bounds for a significantly stronger proof system — a fragment of resolution over parities (Res($\oplus$)). This fragment captures Depth-n Res($\oplus$), and thus our result implies SETH-type lower bounds for both tree-like and regular Res($\oplus$). The core of our approach is a *lossless* lifting achieved by assigning distinct, randomly chosen gadgets to each variable.

Our result also yields a SETH-type lower bound for Depth-n resolution — a result that was previously unknown. We additionally provide a direct and simplified proof for this special case, which may be of independent interest.

1 Introduction

The Strong Exponential Time Hypothesis (SETH) [19] states that for every $\epsilon > 0$, there exists a sufficiently large k such that k-SAT cannot be solved in time $2^{n(1-\epsilon)}$. This hypothesis has become a starting point of fine-grained complexity theory, as assuming it allows us to show that the fastest known algorithms for many fundamental problems in P are essentially optimal [1, 4, 11, 12, 24]. From the other side, Williams [32, 33] demonstrated that any improvement in the exponent for certain satisfiability problems would lead to new circuit lower bounds. Building on this line of work, Jahanjou, Miles, and Viola [23] further established that refuting SETH itself would also imply circuit lower bounds.

Carmosino, Gao, Impagliazzo, Mihajlin, Paturi, and Schneider [14] introduced the *Nondeterministic Strong Exponential Time Hypothesis* (NSETH), which asserts that for every $\epsilon > 0$, there exists a sufficiently large k such that the set of tautologies in k-DNF cannot be decided in nondeterministic time $2^{n(1-\epsilon)}$ for unbounded k. Equivalently, NSETH states that the set of Boolean tautologies does not admit propositional proof systems capable of producing proofs significantly shorter than those obtained by brute-force enumeration of all possible variable assignments. The authors further showed that assuming NSETH leads to new fine-grained lower bounds, whereas refuting it would imply circuit lower bounds. In contrast, Williams [31] demonstrated that the analogue of NSETH for proof systems whose proofs can be verified in randomized time (i.e., MA-proof systems) does not hold — in other words, MASETH is false.

One might hope that SETH-type lower bounds could be established unconditionally for certain restricted classes of algorithms and proof systems. However, even for such specific classes, only a few SETH-type

*e-mail: klimefrem@gmail.com. Supported by European Research Council Grant No. 949707.
†e-mail: dmitrits@gmail.com. Supported by European Research Council Grant No. 949707.

ISSN 1433-8092

lower bounds are currently known. Pudlák and Impagliazzo [27] established an SETH-type lower bound for tree-like resolution, and consequently for DPLL algorithms. Scheder, Tang, Chen, and Talebanfard [29] proved SETH-type lower bounds for PPSZ algorithms [26, 25]. A major breakthrough was achieved by Beck and Impagliazzo [5], who constructed formulas requiring resolution width at least $(1-\epsilon)$ times the number of variables and used this result to derive an SETH-type lower bound for regular resolution, where all resolved variables along any path must be distinct. Subsequently, Bonacina and Talebanfard [10] slightly strengthened this result by proving a SETH-type lower bound for δ-regular resolution — a relaxed version of regular resolution in which, along any path from an initial clause to a contradiction, at most a fraction δ of variables may be resolved more than once. It is worth noting that a lower bound of $2^{(1-\epsilon)n}$ can be established for δ-regular refutations with $\delta = O(\epsilon)$.

Our goal is to substantially broaden the class of proof systems for which SETH-type lower bounds can be established unconditionally. Specifically, we aim to extend such results from subsystems of resolution to subsystems of resolution over parities (Res($\oplus$)) [21, 22] — an extension of resolution that allows reasoning with linear equations over $\mathbb{F}_2$.

1.1 Resolution over parities

A linear clause is defined as a disjunction of linear equations over $\mathbb{F}_2$ (equivalently, as the negation of a system of $\mathbb{F}_2$-linear equations). A Res($\oplus$) refutation of an unsatisfiable CNF formula φ is a sequence of linear clauses $C_1, C_2, \ldots, C_s$ satisfying the following conditions: 1) C_s is the empty clause (i.e., identically false); 2) for every i, the linear clause C_i is either a clause of φ, or derived from earlier linear clauses C_j, C_k with $j, k < i$ by the resolution rule, or derived from a linear clause C_j with $j < i$ by the weakening rule. The resolution rule allows deriving a linear clause $C \vee D$ from $C \vee (f = 0)$ and $D \vee (f = 1)$, where C and D are linear clauses and f is a linear form. The weakening rule permits deriving D from C whenever C semantically implies D. The classical resolution system is recovered as a special case of Res($\oplus$) when all linear clauses are restricted to be disjunctions of literals.

For unrestricted Res($\oplus$), no superpolynomial lower bounds on proof size are currently known. The system Res($\oplus$) forms a particularly simple and natural subsystem of $\mathrm{AC}^0[2]$-Frege (the constant-depth Frege system with parity gates). Establishing superpolynomial lower bounds for Res($\oplus$) would therefore represent a major breakthrough, marking significant progress toward the long-standing goal of proving superpolynomial lower bounds for $\mathrm{AC}^0[2]$-Frege and, ultimately, for stronger propositional proof systems.

Superpolynomial lower bounds on proof size are currently known only for certain subsystems of Res($\oplus$), such as tree-like, regular, and bounded-depth refutations. In this paper, we do not attempt to broaden the class of subsystems for which superpolynomial lower bounds are known. Instead, we focus on strengthening these bounds to the level predicted by SETH. We believe that achieving this goal will also provide new techniques that may prove useful for establishing size lower bounds for unrestricted Res($\oplus$).

The most extensively studied fragment of Res($\oplus$) is that of tree-like refutations, in which the proof has a tree structure. Tree-like Res($\oplus$) refutations of a CNF formula φ are equivalent to *parity decision trees* that, given an assignment, search for a clause of φ falsified by that assignment [21].

Parity decision trees have also been extensively studied outside of proof complexity as a model for representing Boolean functions. Exponential lower bounds of the form $2^{n(1-\epsilon)}$, where n is the number of variables, are known for parity decision trees computing certain Boolean functions. In particular, consider an *affine disperser* [6] — a Boolean function $f : \{0,1\}^n \to \{0,1\}$ that is nonconstant on every affine subspace of dimension at least ϵn. It is then straightforward to observe that no branch of a parity decision tree computing f can terminate in fewer than $(1-\epsilon)n$ steps.

This approach, however, does not carry over directly to search problems, because a clause of width k can be falsified by an affine subspace of dimension k, and therefore, affine dispersers have no natural analog for search problems.

Current lower bounds for tree-like Res($\oplus$) remain far from reaching SETH. The strongest known lower bound follows from the lifting theorem of Chattopadhyay, Mande, Sanyal, and Sherif [15], which states that if a formula φ has resolution depth d and is lifted via a k-stiffling gadget g, then any tree-like Res($\oplus$) refutation of the lifted formula $\varphi \circ g$ has size at least 2^{dk}. Unfortunately, k-stiffling gadgets must have size

at least $2k+1$, which implies that the resulting lower bound in terms of the total number of variables N cannot exceed $2^{N/2}$. Consequently, obtaining a SETH-type lower bound for tree-like Res($\oplus$) via lifting would require a *lossless* lifting method, in which the complexity of the lifted formula scales as $2^{(1-\epsilon)N}$, where N is the number of variables in the lifted formula.

Efremenko, Garlik, and Itsykson [16] made the first progress beyond tree-like Res($\oplus$) by proving an exponential lower bound for regular Res($\oplus$) under a natural notion of regularity. Building on this, Alekseev and Itsykson [2] established an exponential lower bound for a stronger model, for Res($\oplus$) refutations of depth at most $n \log\log n$, where n denotes the number of variables, using a technique based on random walks with restarts. Subsequent works [17, 13, 7, 20] extended this approach. In particular, Bhattacharya and Chattopadhyay [7] achieved superpolynomial lower bounds for depths up to $o(n^2/\log n)$, although this method appears unlikely to extend beyond this threshold.

Our goal is to establish a SETH-type lower bound for Res($\oplus$) refutations of depth at most n. Why focus on this particular subsystem? It is a relatively powerful proof system: it simulates all forms of regular Res($\oplus$), and hence also tree-like Res($\oplus$). While lower bounds are known for larger depths, they come at the cost of significantly weaker size lower bounds, which are insufficient for deriving SETH-type bounds.

1.2 Our contribution

We say that a Res($\oplus$) refutation contains an (a,b)-path if its graph contains a node v such that there are paths from v to the empty clause and from an initial clause to v, where the linear forms resolved along the path from v to the empty clause span a space of dimension at least a, and those resolved along the path from the initial clause to v span a space of dimension at least b.

Theorem 1.1 (Corollary 8.4)**.** For any $\epsilon > 0$ for all large enough n there exists $m = O(\log n \cdot \mathrm{poly}(1/\epsilon))$ and $N = O(n \log n/\epsilon)$ and an unsatisfiable m-CNF formula φ over N variables and of size $n^{\mathrm{poly}(1/\epsilon)}$ such that

- any Res($\oplus$) refutation of φ has width at least $N(1-\epsilon)$;
- any Res($\oplus$) refutation of φ that is free of $((1-\epsilon)N, 2\epsilon N)$-paths has size at least $2^{\left(1-\frac{7}{2}\epsilon\right)N}$.

This theorem implies that SETH holds for Res($\oplus$) proofs whose depth is bounded by the number of variables, and consequently for regular Res($\oplus$) (under any reasonable notion of regularity) as well as for tree-like Res($\oplus$).

The lower bound proof relies on the random walk approach, similar to that used in [16] and subsequent works. We employ a modified version of the *random walk with restarts* technique by Alekseev and Itsykson [2]: after the first execution of a random walk, we either obtain a sufficiently strong size lower bound or can proceed slightly deeper into the proof graph. A crucial ingredient of our argument is the use of lossless lifting, which we achieve by assigning distinct, randomly chosen gadgets to different variables. Random gadgets were previously used by Bonacina and Talebanfard [9] to obtain strong lower bounds on resolution width, but to the best of our knowledge, this is the first setting where employing different gadgets for different variables is essential.

Theorem 1.1 also yields new results for resolution. We further simplify the approach and present a direct, self-contained proof of a slightly stronger statement for resolution. The notion of an (a,b)-path is defined consistently as follows: there exist a node v such that there is a path from v to the empty clause and a path from an initial clause to v, the set of variables resolved along the path from v to the empty clause has size at least a, and the set of variables resolved along the path from the initial clause to v has size at least b.

Theorem 1.2 (Corollary 2.10)**.** For all large enough ℓ for infinitely many N there exists an unsatisfiable formula φ in ℓ-CNF that depends on N variables such that every its resolution refutation that is free of $((1-\delta)N, 2\delta N)$-paths has size at least $2^{(1-\tilde{\Theta}(\ell^{-1/4}))N}$, where $\delta = \tilde{\Theta}(\ell^{-1/4})$.

This theorem implies that SETH holds for resolution proofs whose depth is bounded by the number of variables, as well as for $\delta/2$-regular resolution refutations. Hence, Theorem 1.2 can be regarded as an

extension of the result of Bonacina and Talebanfard [10]. Moreover, Theorem 1.2 relies on the same family of formulas as used in [10].

We present the proof of Theorem 1.2 separately in Section 2, in a self-contained form, as it may be of independent interest.

1.3 Proof outline and overview of techniques

We first present the outline of the proof of Theorem 1.1. The proof of Theorem 1.2 proceeds analogously, with Res($\oplus$) replaced by resolution.

1. We start with the family of CNF formulas ψ_n over n variables that require resolution width $w \geq (1-\epsilon)n$ [5, 10].

   To handle them, we use the game-theoretic characterization of resolution width due to Atserias and Dalmau [3]. In this game, two players Prover and Adversary play on an unsatisfiable formula φ. They maintain a partial assignment ρ, which is initially empty. In each move, Prover may either erase some values from ρ or select an unassigned variable and ask Adversary to assign it a value. The game ends as soon as ρ falsifies a clause of φ. It is known [3] that φ has no resolution refutation of width smaller than w if and only if Adversary has a winning strategy in the following sense: there exists a nonempty set of partial assignments (called winning positions) such that 1) none of them falsify φ, 2) each has size at most w, 3) the set of winning positions is closed under erasure, and 4) from any winning position of size less than w, Adversary can always respond in a way that ensures the next position is also winning.

2. We apply a lifting of size ℓ to the formula ψ_n. By lifting we mean the following transformation: each variable of ψ_n is replaced by a gadget — that is, a Boolean function $g : \{0,1\}^\ell \to \{0,1\}$ defined on ℓ fresh variables — and the resulting formula is then converted into CNF form. Typically, the same gadget is used for all variables, but in our construction we allow different gadgets for different variables. We denote the lifted formula by $\mathrm{Lift}(\psi_n)$; it depends on $N := n\ell$ variables.

3. A similar game characterization of width is known for Res($\oplus$) [18, 2], in which winning positions are systems of $\mathbb{F}_2$-linear equations rather than partial assignments, and the Prover queries the values of linear forms. We show that for the formula $\mathrm{Lift}(\psi_n)$, Adversary has a winning strategy in the Res($\oplus$)-game with width parameter $W \geq (1-\epsilon)\ell w \geq (1-2\epsilon)N$.

4. Moreover, we show that Adversary's winning strategy allows a remarkable degree of flexibility, which we refer to as the *freedom property*: for most linear forms f, a winning position in the game can be extended to another winning position under both assignments $f = 0$ and $f = 1$. Formally, we prove that along any sequence of moves extending winning positions by one linear form at a time, only an ϵ-fraction of these extensions are forced—that is, in at most an ϵ fraction of cases Adversary cannot choose the opposite value of the linear form without leaving the set of winning positions.

5. We consider the random full assignment σ of the formula $\mathrm{Lift}(\psi_n)$ and traverse a refutation from the empty clause, each time we go from linear clause to a premise that is falsified by σ. We finish the path if we go through at least W linearly independent linear forms that are used in resolution rules or if we reach a clause of the initial formula.

6. By the freedom property of Adversary's winning strategy, with probability at least $2^{-\epsilon W}$, the random walk reaches a linear clause C whose negation corresponds to a winning position.

7. Assume that the width of C is less than $(1-5\epsilon)W$. Then we can construct a path starting from C that passes through the negations of winning positions. This path can be extended until it contains $5\epsilon W$ linearly independent linear forms. Hence, the refutation must contain a $(W, 5\epsilon W)$-path. Under the assumption that the refutation does not contain any $((1-2\epsilon)N, 4\epsilon N)$-paths, we conclude that the width of C is at least $(1-5\epsilon)W$.

8. The probability that a random assignment σ falsifies a clause C of width t is exactly 2^{-t}. Therefore, the refutation must contain at least $2^{-\epsilon W}/2^{-(1-5\epsilon)W} = 2^{(1-6\epsilon)W} \geq 2^{(1-8\epsilon)N}$ distinct linear clauses.

In Section 2, we apply the above plan to the resolution case and prove Theorem 1.2. Here, all variables are lifted using the same gadget—the ℓ-bit parity function $\oplus_\ell$ with $\ell = 1/\epsilon$. The proof closely follows the general scheme, with only minor adjustments required for the resolution setting.

In the case of Res($\oplus$), the proof is considerably more technical, with the key challenge being the lifting.

Lifting of strategies from the game characterizing resolution width to game characterizing Res($\oplus$) width was obtained by Alekseev and Itsykson [2]; this was achieved by 1-stifling gadgets. But in that result the Res($\oplus$)-width of the lifted formula is at least the resolution-width of the initial formula, but the number of variables increases in the size of the gadget. Thus, to fulfill item 3 of the above plan, we need to construct *lossless* lifting.

The lifting construction of Alekseev and Itsykson [2] is essentially based on the notion of *closure* introduced in [16]. We work in the standard lifting setting, where the initial variables are $Y = \{y_1, y_2, \ldots, y_n\}$ and the lifted variables are $X = \{x_{i,j} \mid i \in [n],\ j \in [\ell]\}$. The lifted variables are partitioned into n blocks, each block corresponding to a distinct index i of the original variable y_i.

In Section 4, for every integer $k \in [\ell]$, we define the notion of k-*closure*, which generalizes the concept of closure in the sense that the standard closure corresponds to the case $k = 1$.

The notion of closure is based on the concept of *safeness*. We introduce the notion of k-*safe* matrices, which extends the concept of safe matrices. A matrix over the lifted variables is said to be k-safe if there exists a basis among its columns such that each block contains at most k basis elements. For $k = 1$, this definition coincides with the notion of safe matrices from [16]. The set of solutions of a linear system with a k-safe matrix has the following structure: variables that do not belong to the basis are *free* variables, while the variables corresponding to the basis are *dependent* and are expressed as affine functions of the free variables.

The main useful property of safe matrices and 1-stifling gadgets is the following [8, 2]: for every assignment to the unlifted variables, and for every satisfiable linear system over the lifted variables with a safe matrix, there exists a solution of the system that induces the initially chosen values on the unlifted variables via the gadget.

Our goal is to extend this property to k-safe matrices. Specifically, we construct gadgets with a property we call k-*surjectivity*: for every satisfiable linear system whose coefficient matrix is k-safe, every assignment to the unlifted variables can be realized by an appropriate solution of the system. When $k < \ell/2$, we can employ majority gadgets: in this case, the majority can be fixed by free variables only, allowing us to fix the unlifted variables arbitrarily by an appropriate choice of the free variables.

However, to make our lifting lossless, we need $k = (1-\epsilon)\ell$. It is easy to see that when $k \geq \ell/2$, we must use different gadgets for different variables in the lifting. Indeed, consider two blocks and assume that the first $\ell - k$ variables in the first block are free, while the second $\ell - k$ variables in the second block are free. Suppose further that each dependent variable in the second block is equal to the variable in the first block with the same index. If both variables are dependent, this means that we assign the same affine function to them; if one variable is free and the other is dependent, then the dependent variable is set equal to that free variable (that is, to an affine function coinciding with the value of that free variable). Finally, assume that the dependent variables in the first block are equal to the corresponding variables in the second block. Then, if we use identical gadgets for these two blocks, their values will coincide, and consequently, not all combinations of unlifted variable assignments will be realizable.

The above argument also implies that the gadget's size ℓ cannot be a constant. Indeed, if we have n blocks, then n must be greater than the number of Boolean functions on ℓ variables; otherwise, there would exist two blocks using the same gadget. Therefore, we obtain $\ell \geq \log\log n$. In Section 5, we present a randomized construction of a family of gadgets with $\ell = O(\log n/\epsilon)$, which satisfy the property of k-surjectivity.

In Section 6, we lift resolution width to Res($\oplus$) width by lifting strategies in the corresponding games, using the ideas of Itsykson and Alekseev [2] adapted to k-closure and our lifting. This realizes item 3 of the above plan.

In Section 7, we establish the *freedom property* of the lifted strategy, thus realizing item 4 of the plan. The proof is technically very similar to the proof of the tree-like Res($\oplus$) lower bound for the binary pigeonhole principle from [16].

Finally, we conclude the proof in Section 8, thereby realizing items 5–8 of the plan.

1.4 Further research

We outline two potential directions for strengthening our results:

- **Increasing depth.** For instance, even establishing a SETH-type lower bound for Depth-$2n$ resolution appears to require new ideas.
- **Reducing width.** It would be interesting to prove a $2^{(1-\epsilon)n}$ lower bound for Depth-n Res($\oplus$), or at least for tree-like Res($\oplus$), on k-CNF formulas where k depends only on ϵ and is independent of n.

2 SETH Holds for Bounded-Depth Resolution

2.1 Resolution and decision DAG

Let φ be an unsatisfiable CNF formula. A resolution refutation of φ is a sequence of clauses $C_1, C_2, \ldots, C_s$ such that C_s is the empty clause (i.e., identically false) and for every $i \in [s]$ the clause C_i is either a clause of φ or is obtained from previous clauses by the *resolution rule* that allows us to derive a clause $C \vee D$ from clauses $C \vee x$ and $D \vee \neg x$.

The *size* of a resolution refutation is the number of clauses in it. The *depth* of a resolution refutation is the length of the longest path between the empty clause and the clause of the original formula. The *width* of a resolution refutation is the maximal size of a clause from the refutation. The resolution width of an unsatisfiable CNF formula φ is the minimal possible width over all resolution refutations of φ.

A *decision DAG* for an unsatisfiable CNF formula φ is a directed acyclic graph with a single source and several sinks, satisfying the following properties:

- Each node v of the DAG is labeled with a partial ρ_v assignment over the variables of φ.
- The source is labeled with the empty assignment.
- For every sink v, there exists a clause C of φ such that ρ_v falsifies C.
- Every non-sink node v is additionally labeled with a *splitting* variable x_v and has two children v_0 and v_1. The edge (v, v_0) is labeled by the assignment $x_v := 0$, and the edge (v, v_1) is labeled by $x_v := 1$. Moreover, for each $\alpha \in \{0, 1\}$, the assignment ρ_{v_α} is a subassignment of $\rho_v \cup \{x_v := \alpha\}$.

The size of a parity decision DAG is its number of vertices, the depth is the length of the longest source-to-sink path, and the width is the maximum size of the assignments ρ_v.

It is known [28] that every resolution refutation of a formula φ can be efficiently transformed into a decision DAG for φ without increasing its size, depth, or width. The transformation is straightforward: the decision DAG has the same underlying graph as the refutation, but with all edges reversed. The empty clause becomes the source, while the clauses of the original formula become the sinks. Each node is labeled with the negation of the corresponding clause, and the splitting variables are exactly those used in the resolution steps.

A resolution refutation is called *regular* if, along every path in the associated decision DAG, all splitting variables are distinct. A resolution refutation of an n-variable formula is called δ-*regular* [10] if, along every path from the source to a sink in the decision DAG, all but at most δn of the splitting variables occur at most once.

2.2 Game characterization of resolution width

Atserias and Dalmau proposed a game characterization of resolution width [3]. Let φ be an unsatisfiable CNF formula. A w-winning resolution strategy for φ is a non-empty family $\mathcal{H}$ of partial truth assignments, none of which falsify any clause of φ, such that:

- For every $\rho \in \mathcal{H}$, we have $|\rho| \leq w$.
- If $\rho \in \mathcal{H}$ and $\tau \subseteq \rho$, then $\tau \in \mathcal{H}$.
- If $\rho \in \mathcal{H}$, $|\rho| < w$ and x is a variable, then there exists $a \in \{0,1\}$ such that $\rho \cup \{x := a\} \in \mathcal{H}$.

The following lemma characterizes the connection between resolution width and winning resolution strategies:

Lemma 2.1 ([3])**.** Let φ be a CNF formula and W an integer. Then φ admits no resolution refutation of width W if and only if there exists a $(W+1)$-winning resolution strategy for φ.

2.3 Lifting by a parity

We denote by $\oplus_k$ the parity function $\{0,1\}^k \to \{0,1\}$ that maps $(a_1, a_2, \ldots, a_k)$ to $a_1 + a_2 + \cdots + a_k \bmod 2$.

For every CNF formula φ over the variables $Y = \{y_1, y_2, \ldots, y_m\}$ and we define a CNF formula $\Phi \circ \oplus_k$ with variables $X = \{x_{i,j} \mid i \in [m], j \in [k]\}$ representing in CNF $\phi(x_{1,1} \oplus x_{1,2} \oplus \cdots \oplus x_{1,k}, x_{2,1} \oplus x_{2,2} \oplus \cdots \oplus x_{2,k}, \ldots, x_{m,1} \oplus x_{m,2} \oplus \cdots \oplus x_{m,k})$ (i.e. we substitute to every variable of φ the function $\oplus_k$ applied to k fresh variables).

For every partial assignment ρ to the variables of the formula $\varphi \circ \oplus_k$ we define the partial assignment $\text{induced}_{\oplus}(\rho)$ to the variables of φ as follows:

- $\text{induced}_{\oplus}(\rho)$ is defined on y_i, if and only if ρ is defined on all $x_{i,1}, x_{i,2}, \ldots, x_{i,k}$;
- $\text{induced}_{\oplus}(\rho)(y_i) = \bigoplus_{j=1}^{k} \rho(x_{i,j})$.

Let $\mathcal{A}$ be a set of partial assignments for the variables of φ.

Based on the formula $\varphi \circ \oplus_k$ we define a set $\mathcal{A}^{\oplus_k}$ that consists of partial assignments ρ to variables of $\varphi \circ \oplus_k$ such that $\text{induced}_{\oplus}(\rho) \in \mathcal{A}$.

Lemma 2.2. If $\mathcal{A}$ is w-winning resolution strategy for φ, then $\mathcal{A}^{\oplus_k}$ is wk-winning resolution strategy for $\varphi \circ \oplus_k$.

Proof. Consider $\rho \in \mathcal{A}^{\oplus_k}$, every clause of $\varphi \circ \oplus_k$ is a clause of $C \circ \oplus_k$, where C is a clause of φ. Since $\text{induced}_{\oplus}(\rho)$ doesn't falsify C, there is a variable y_j of C such that $\text{induced}_{\oplus}(\rho)$ is not defined on y_j. Hence, there is $i \in [k]$ such that ρ is not defined on $x_{j,i}$, hence ρ doesn't falsify $C \circ \oplus_k$.

It follows by the construction of $\mathcal{A}^{\oplus_k}$ that all elements of $\mathcal{A}$ have size at most kw.

Consider $\rho \in \mathcal{A}^{\oplus_k}$ and let $\rho' \subseteq \rho$. By the definition, $\text{induced}_{\oplus}(\rho') \subseteq \text{induced}_{\oplus}(\rho)$. Since $\text{induced}_{\oplus}(\rho) \in \mathcal{A}$, then $\text{induced}_{\oplus}(\rho') \in \mathcal{A}$, then $\rho' \in \mathcal{A}^{\oplus_k}$.

Let $\rho \in \mathcal{A}^{\oplus_k}$ such that $|\rho| < wk$, let $x_{i,j}$ be a variable not from domain of ρ. If the domain of ρ doesn't contain all $x_{i,t}$ for $t \in [k] \setminus \{j\}$, then we can arbitrarily extend ρ on $x_{i,j}$, since in this case $\text{induced}_{\oplus}(\rho \cup \{x_{i,j} := a\}) = \text{induced}_{\oplus}(\rho)$ for all $a \in \{0,1\}$.

Consider the second case, where the domain of ρ contains all $x_{i,t}$ for $t \in [k] \setminus \{j\}$. Since $|\rho| < wk$, $|\text{induced}_{\oplus}(\rho)| < w$, and y_i does not belong to the domain of $\text{induced}_{\oplus}(\rho)$. By the properties of w-winning strategy $\mathcal{A}$ there exists $a \in \{0,1\}$ such that $\text{induced}_{\oplus}(\rho) \cup \{y_i := a\}$ belongs $\mathcal{A}$. So we can choose the value $b := a \oplus \bigoplus_{t \in [k] \setminus \{j\}} \rho(x_{i,j})$ and extend ρ by $x_{i,j} := b$ and then $\text{induced}_{\oplus}(\rho \cup \{x_{i,j} := b\}) = \text{induced}_{\oplus}(\rho) \cup \{y_i := a\}$. □

Consider an unsatisfiable CNF formula φ and a set $\mathcal{A}$ of partial assignments over the variables of φ such that (i) every assignment in $\mathcal{A}$ does not falsify any clause of φ, and (ii) $\mathcal{A}$ is closed under restrictions. We refer to such sets of assignments as *proper*. We define the $(\varphi, \mathcal{A})$*-game between Prover and Delayer*. In this game, two players, Prover and Delayer, maintain a partial assignment ρ for variables of φ; initially, ρ is empty. On every move, Prover chooses a variable x, and Delayer has two options:

- Delayer can earn a *white* coin and reports $*$. Then, Prover chooses a Boolean value a of x.
- Delayer can earn a *white* coin and pay a *black* coin to choose a Boolean value a of x by himself.

The current assignment ρ is updated: $\rho := \rho \cup \{x := a\}$. The game ends when $\rho \notin \mathcal{A}$.

Lemma 2.3. Let $\mathcal{A}$ be w-winning resolution strategy for φ. Then in the $(\varphi \circ \oplus_k, \mathcal{A}^{\oplus_k})$-game there is a strategy of Delayer that guaranties him to earn at least wk white coins while paying at most w black coins.

Proof. The lemma is in fact established in the proof of Lemma 2.2. The only instance in which the Delayer is required to assign a variable himself arises when $x_{i,j}$ is the last unassigned variable. Throughout wk rounds of the game, the number of such instances does not exceed w. □

2.4 Lower bound

We say that a resolution refutation contains an (a, b)*-path* if, in the decision DAG associated with the refutation, there exists a node v such that the path from the source to v passes through at least a distinct variables, and the path from v to a sink passes through at least b distinct variables.

Lemma 2.4. Let φ be an unsatisfiable CNF formula and let w be the minimal width of resolution refutations of φ. Then any resolution refutation of $\varphi \circ \oplus_k$ has either size $2^{(k-3)w}$ or contains $(wk, 2w)$ path.

Proof. Consider a resolution refutation of φ and its associated decision DAG. Suppose that the refutation does not contain any $(wk, 2w)$-path. Let τ be a random full assignment to the variables of $\varphi \circ \oplus_k$. We define a path in the decision DAG as follows: starting from the source, at each step we follow the outgoing edge whose label is consistent with τ. The path terminates once wk distinct variables have been queried, or earlier if a sink is reached. The random variable of interest is the partial assignment labeling the node at which this path terminates. An equivalent way to generate this distribution is as follows. Starting from the source, if the queried variable is new, we follow a uniformly random outgoing edge; if it has been queried before, we deterministically follow the edge consistent with its previously chosen value. Thus, we obtain two equivalent views of the same random process: in the first, a random full assignment τ is sampled in advance; in the second, the values of variables are chosen randomly only when they are queried.

By Lemma 2.1, there exists a w-winning resolution strategy $\mathcal{A}$ for φ. By Lemma 2.3, in the $(\varphi \circ \oplus_k, \mathcal{A}_k^{\oplus})$-game Delayer has a strategy that guarantees him at least wk white coins while paying at most w black coins.

Under the second view of the random process, the existence of the Delayer's strategy implies that, with probability at least 2^{-w}, the partial assignment encountered along the path is consistent with the Delayer's strategy. Consequently, with probability at least 2^{-w}, the partial assignment ρ defined by the edges of the random walk belongs to $\mathcal{A}^{\oplus_k}$. Moreover, the assignment σ labeling the endpoint of the path is a subassignment of ρ; hence, in this case, we also have $\sigma \in \mathcal{A}^{\oplus_k}$.

Claim 2.5. Assume that a node v of the decision DAG is at distance wk from the source and is labeled with a partial assignment $\rho_v \in \mathcal{A}^{\oplus_k}$. Then the size of ρ_v is at least $w(k-2)$.

Proof. By Lemma 2.2, $\mathcal{A}^{\oplus_k}$ is a wk-winning resolution strategy for $\varphi \circ \oplus_k$. Consider a path in the decision DAG starting from a node v, where at each step we move from a vertex u along the edge labeled by $x_u := a$ such that $\rho_u \cup \{x_u := a\} \in \mathcal{A}^{\oplus_k}$. Let u' denote the endpoint of this edge. Then, by the properties of a winning strategy, $\rho_{u'} \in \mathcal{A}^{\oplus_k}$. Since $\mathcal{A}^{\oplus_k}$ is a wk-winning strategy, we can construct a path whose edges involve at least $wk - |\rho_v|$ distinct variables. As the decision DAG does not contain $(wk, 2w)$-paths, there

is no path from v that contains at least $2w$ distinct variables. It follows that $wk - |\rho_v| \leq 2w$, and, thus, $|\rho_v| \geq w(k-2)$.

□

Any partial assignment of size at least $w(k-2)$ is consistent with a random assignment τ with probability at most $2^{-w(k-2)}$. Since with probability at least 2^{-w} the endpoint of the random path satisfies the conditions of Claim 2.5 (we refer to this as a *lucky event*), the number of distinct nodes that can serve as endpoints of the random path in lucky events is at least $\frac{2^{-w}}{2^{-w(k-2)}} = 2^{(k-3)w}$. □

We say that a resolution refutation is (a,b)-*path-free* if it does not contain any (a,b)-paths.

Proposition 2.6. If depth of a resolution refutation is less than $a+b$, then it is (a,b)-path-free.

Proof. The proof is straightforward, □

Proposition 2.7. If a resolution refutation is δ-regular, then it is $((1-\delta')n, 2\delta' n)$-path-free for every $\delta' > \delta$.

Proof. Assume that there exists a δ-regular resolution refutation containing a $((1-\delta')n, 2\delta' n)$-path for some $\delta' > \delta$. Then there exist a node v, a sink s of the decision DAG, a path from the source to v that passes through at least $(1-\delta')n$ distinct variables, and a path from v to s that passes through at least $2\delta' n$ distinct variables. Consequently, along the concatenated path source–v–s, at least $\delta' n$ variables appear more than once, contradicting δ-regularity of the refutation. □

Theorem 2.8. Let φ be an unsatisfiable CNF formula with n variables such that φ requires resolution width at least $(1-\delta)n$. Let k be integer such that $\delta(k+1) < 1$. Then any $((1-\delta)nk, 2\delta nk)$-path-free resolution refutation of $\varphi \circ \oplus_k$ has size at least $2^{(1-\delta-3/k)nk}$.

Proof. Notice that if $\delta(k+1) < 1$, then $\delta k < (1-\delta)$. By Lemma 2.4, any resolution refutation of $\varphi \circ \oplus_k$ has either size at least $2^{(1-\delta)n(k-3)} \geq 2^{(1-\delta-3/k)nk}$ or contains $((1-\delta)nk, 2(1-\delta)n)$-path, the second option is impossible since we consider only $((1-\delta)nk, 2\delta nk)$-paths free refutations. □

We use the standard tilde notation that suppresses logarithmic factors. Specifically, $f(n) = \tilde{O}(g(n))$ means that there is a constant c such that $f(n) = O\left(g(n) \cdot \log^c n\right)$, and $f(n) = \tilde{\Theta}(g(n))$ means that both $f(n) = \tilde{O}(g(n))$ and $g(n) = \tilde{O}(f(n))$ hold.

Lemma 2.9 ([9])**.** For large enough m and n there exists an unsatisfiable formula ψ in m-CNF such that resolution width of ψ is at least $(1-\delta)n$, where $\delta = \tilde{\Theta}(m^{-1/3})$.

Corollary 2.10. For all large enough ℓ for infinitely many N there exists an unsatisfiable formula φ in ℓ-CNF that depends on N variables such that every its $((1-\delta)N, 2\delta N)$-path-free resolution refutation has size at least $2^{(1-\tilde{\Theta}(\ell^{-1/4}))N}$, where $\delta = \tilde{\Theta}(\ell^{-1/4}))$.

Proof. By Lemma 2.9, for large enough m and n there exists an unsatisfiable formula ψ in m-CNF with n variables such that resolution width of ψ is at least $(1-\delta)n$, where $\delta = \tilde{\Theta}(m^{-1/3})$. Let us choose k such that $\delta(k+1) < 1$ and $k = \tilde{\Theta}\left(m^{1/3}\right)$. We define $\varphi := \psi \circ \oplus_k$. Then φ is mk-CNF formula, where $k = \tilde{\Theta}\left(m^{1/3}\right)$. Let us denote $N = nk$ and $\ell = mk$, then $\delta = \tilde{\Theta}\left(\ell^{-1/4}\right)$ and $k = \tilde{\Theta}\left(\ell^{1/4}\right)$. By Theorem 2.8 we get that any $((1-\delta)N, 2\delta N)$-path-free resolution refutation of φ has size at least $2^{(1-\delta-3/k)N} = 2^{(1-\tilde{\Theta}(\ell^{-1/4}))N}$. □

3 Resolution Over Parities

Here and after, all scalars are from the field $\mathbb{F}_2$. Let X be a set of variables taking values in $\mathbb{F}_2$. A linear form in variables from X is a homogeneous linear polynomial over $\mathbb{F}_2$ in variables from X or, in other words, a polynomial $\sum_i^n x_i a_i$, where $x_i \in X$ is a variable and $a_i \in \mathbb{F}_2$ for all $i \in [n]$. A linear equation is an equality $f = a$, where f is a linear form and $a \in \mathbb{F}_2$.

A *linear clause* is a disjunction of linear equations: $\bigvee_{i=1}^{t}(f_i = a_i)$. Note that over $\mathbb{F}_2$ a linear clause $\bigvee_{i=1}^{t}(f_i = a_i)$ may be represented as the negation of a linear system: $\neg\bigwedge_{i=1}^{t}(f_i = a_i + 1)$.

For a linear clause C we denote by $L(C)$ the set of linear forms that appear in C; i.e. $L\left(\bigvee_{i=1}^{t}(f_i = a_i)\right) = \{f_1, f_2, \ldots, f_t\}$. The same notation we use for linear systems: if Ψ is a $\mathbb{F}_2$-linear system, $L(\Psi)$ denotes the set of all linear forms from Ψ.

Now we define the proof system resolution over parities (Res($\oplus$)) [22].

Let φ be an unsatisfiable CNF formula. A Res($\oplus$) refutation of φ is a sequence of linear clauses C_1, C_2, $\ldots$, C_s such that C_s is the empty clause (i.e., identically false) and for every $i \in [s]$ the clause C_i is either a clause of φ or is obtained from previous clauses by one of the following inference rules:

- *Resolution rule* allows us to derive a linear clause $C \vee D$ from linear clauses $C \vee (f = a)$ and $D \vee (f = a + 1)$.
- *Weakening rule* allows us to derive from a linear clause C any linear clause D in the variables of φ that semantically follows from C (i.e., any assignment satisfying C also satisfies D).

The *size* of a Res($\oplus$) refutation is the number of linear clauses in it. The *depth* of a Res($\oplus$) refutation is the maximal number of resolution rules applied on a path between a clause of the initial formula and the empty clause. The *width* of a linear clause C is the rank of the linear system $\neg C$. The *width* of a Res($\oplus$) refutation is the maximal width of a linear clause in it.

A *parity decision DAG* (also known as an affine DAG) for an unsatisfiable CNF formula φ is a directed acyclic graph with a single source and several sinks, satisfying the following properties:

- Each node v of the DAG is labeled with an $\mathbb{F}_2$-linear system Φ_v over the variables of φ.
- The source is labeled with the empty system (i.e., identically true).
- For every sink v, there exists a clause C of φ such that Φ_v is inconsistent with C.
- Every non-sink node v is additionally labeled with a linear form f_v and has two children v_0 and v_1. The edge (v, v_0) is labeled by the equation $f_v = 0$, and the edge (v, v_1) is labeled by $f_v = 1$. Moreover, for each $\alpha \in \{0, 1\}$, the system $\Phi_v \wedge (f_v = \alpha)$ semantically implies Φ_{v_α}.

The size of a parity decision DAG is its number of vertices, the depth is the length of the longest source-to-sink path, and the width is the maximum rank of the systems Φ_v.

Similarly to the case of resolution, it is known [16] that every Res($\oplus$) refutation of φ can be efficiently transformed into a parity decision DAG for φ without increasing its size, depth, or width.

4 The k-Closure and Its Properties

In this section, we define the notion of k-closure, which generalizes the closure introduced by Efremenko, Garlik, and Itsykson [16]; in particular, 1-closure coincides with their definition. Since the proofs of all statements closely follow those in [16], we defer them to Appendix A.

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 $f = x_{i_1,j_1} + x_{i_2,j_2} + \cdots + x_{i_s,j_s}$ is the set $\{i_1, i_2, \ldots, i_s\}$ of blocks of variables that appear in f with non-zero coefficients. We denote the support by $\mathrm{supp}(f)$. 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 $\mathrm{supp}(F)$. Let k be a natural number. We say that a linearly independent set of linear forms F is *k-dangerous* if $|F| > k|\mathrm{supp}(F)|$. We say that a set of linear forms F is *k-safe* if $\langle F \rangle$ does not contain a k-dangerous set. If F is linearly dependent but $\langle F \rangle$ contains a k-dangerous set, instead of saying that F is k-dangerous, we say it is not k-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 $f_1, f_2, \ldots, f_s$, one may consider their coefficient matrix of size $s \times |X|$ in which the i-th row coincides with the coefficient vector of f_i.

Theorem 4.1. Let $f_1, f_2, \ldots, f_s$ be linearly independent linear forms and let M be their coefficient matrix. Then the following conditions are equivalent.

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

(2) For every set $T \subseteq [m]$, the dimension of the span of the set of columns of M corresponding to the variables with support in T is at least $k|T| - (km - s)$.

(3) One can choose s variables such that for every block at most k variables are chosen from this block and the columns of M corresponding to the s chosen variables are linearly independent. Since the rank of M is s, the chosen columns form the basis of the span of M's columns.

Let $S \subseteq [m]$ be a set of blocks; for a linear form f we denote by $f[\backslash S]$ a linear form obtained from f by substituting 0 for all variables with support in S. In other words, $f[\backslash S]$ is the projection of f to the linear space of all forms with support in $[m] \setminus S$. Being a projection, $[\backslash S]$ is a linear operator for every $S \subseteq [m]$.

For a set of linear forms F we will use the notation $F[\backslash S] = \{f[\backslash S] \mid f \in F\}$.

Lemma 4.2 ([16])**.** Let F be a set of linear forms and T be a subset of $[m]$. Then

$$\dim\langle F\rangle = \dim\langle F[\backslash T]\rangle + \dim\langle\{f \in \langle F\rangle \mid \mathrm{supp}(f) \subseteq T\}\rangle.$$

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

Lemma 4.3. (Uniqueness) For any F its k-closure is unique.

Let us denote the unique k-closure of F by $\mathrm{Cl}^{(k)}(F)$.

Lemma 4.4. (Monotonicity) If $F_1 \subseteq F_2$, then $\mathrm{Cl}^{(k)}(F_1) \subseteq \mathrm{Cl}^{(k)}(F_2)$.

Proof. $F_1[\backslash \mathrm{Cl}^{(k)}(F_2)] \subseteq F_2[\backslash \mathrm{Cl}^{(k)}(F_2)]$, hence $F_1[\backslash \mathrm{Cl}^{(k)}(F_2)]$ is k-safe. Consider an inclusion minimal set $S \subseteq \mathrm{Cl}^{(k)}(F_2)$ such that $F_1[\backslash S]$ is k-safe. Then S is a k-closure of F_1 and, by the uniqueness, $\mathrm{Cl}^{(k)}(F_1) = S \subseteq \mathrm{Cl}^{(k)}(F_2)$. □

Lemma 4.5. (Span invariance) $\mathrm{Cl}^{(k)}(F) = \mathrm{Cl}^{(k)}(\langle F\rangle)$.

Proof. Since $[\backslash S]$ is a linear operator, $\langle F\rangle[\backslash S] = \langle F[\backslash S]\rangle$. Hence for every S, the set $F[\backslash S]$ is k-safe iff $\langle F\rangle[\backslash S]$ is k-safe, and so $\mathrm{Cl}^{(k)}(F) = \mathrm{Cl}^{(k)}(\langle F\rangle)$. □

Lemma 4.6. (Size bound) $k|\mathrm{Cl}^{(k)}(F)| + \dim\langle F[\backslash \mathrm{Cl}^{(k)}(F)]\rangle \le \dim\langle F\rangle$, and hence $|\mathrm{Cl}^{(k)}(F)| \le \frac{1}{k}\dim\langle F\rangle$.

Lemma 4.7. (Increment) Let F be a set of linear forms and f be a linear form such that $\mathrm{Cl}^{(k)}(F \cup \{f\}) \neq \mathrm{Cl}^{(k)}(F)$. Then

$$\dim\langle F[\backslash \mathrm{Cl}^{(k)}(F)]\rangle - \dim\langle (F \cup \{f\})[\backslash \mathrm{Cl}^{(k)}(F \cup \{f\})]\rangle = k(|\mathrm{Cl}^{(k)}(F \cup \{f\})| - |\mathrm{Cl}^{(k)}(F)|).$$

5 Lifting by k-Surjective Collection of Gadgets

Let $\varphi(y_1, y_2, \ldots, y_m)$ be a CNF formula, and let $g_1, g_2, \ldots, g_m : \{0,1\}^\ell \to \{0,1\}$ be Boolean functions. We define the *lifting* of φ with respect to $(g_1, g_2, \ldots, g_m)$, denoted by $\mathrm{Lift}_{g_1,g_2,\ldots,g_m}(\varphi)$ as the CNF formula obtained from

$$\varphi(g_1(x_{1,1}, x_{1,2}, \ldots, x_{1,\ell}), g_2(x_{2,1}, x_{2,2}, \ldots, x_{2,\ell}), \ldots, g_m(x_{m,1}, x_{m,2}, \ldots, x_{m,\ell}))$$

by converting the resulting expression into CNF in the following way: for each clause of φ, we substitute the corresponding functions g_i, transform the resulting formula into CNF in any fixed manner, and finally take the conjunction of all such CNFs over all clauses of φ. We refer to the set $Y = y_1, y_2, \ldots, y_m$ as *the unlifted variables*, and to the set $X = \{x_{i,j} \mid i \in [m], j \in [\ell]\}$ as *the lifted variables*.

5.1 Definitions of k-special affine spaces and k-surjective collection of gadgets

Observe that every affine subspace of $\mathbb{F}^n$ of dimension k can be represented as follows: one selects k coordinates, called *free*, while the remaining coordinates are *dependent*, meaning they are determined by affine functions of the free coordinates.

Definition 5.1. Consider the vector space $\mathbb{F}_2^{m\ell}$, whose coordinates are partitioned into m blocks of size ℓ. The coordinates of the i-th block are denoted by $x_{i,1}, x_{i,2}, \ldots, x_{i,\ell}$. We say that a subspace $L \subseteq \mathbb{F}_2^{m\ell}$ is *k-special* if there exists a representation of L in which each block contains at least $\ell - k$ free coordinates.

The following proposition establishes that linear systems containing k-safe linear forms define k-special affine spaces.

Proposition 5.2. Assume that a set of linear forms F over variables X is k-safe. Then, for every satisfiable linear system Φ whose set of linear forms is F, the solution space of Φ is k-special.

Proof. By Theorem 4.1, the matrix of Φ contains a basis among its columns such that each block contains at most k basis elements. The general solution of Φ can then be described as follows: the variables corresponding to non-basis elements may be assigned arbitrarily (these are the *free* variables), while the values of the basis variables are uniquely determined as affine functions of the free variables. □

We call a collection of gadgets $g_1, g_2, \ldots, g_m : \{0,1\}^\ell \to \{0,1\}$ *k-surjective* if, for every k-special affine subspace $L \subseteq \mathbb{F}_2^{m\ell}$ and every tuple $(b_1, b_2, \ldots, b_m) \in \{0,1\}^m$, there exists an element $x \in L$ such that for all $i \in [m]$,

$$g_i(x_{i,1}, x_{i,2}, \ldots x_{i,\ell}) = b_i.$$

Let $g_1, g_2, \ldots, g_m : \{0,1\}^\ell \to \{0,1\}$ be a collection of gadgets. For every full assignment τ to the lifted variables X, we define the induced assignment $\mathrm{induced}_{g_1,g_2,\ldots,g_m}(\tau)$ as the assignment σ to the unlifted variables such that, for each $i \in [m]$, $\sigma(y_i) = g_i(\tau(x_{i,1}), \tau(x_{i,2}), \ldots, \tau(x_{i,\ell}))$.

5.2 Lifting and k-closure

In this subsection, we present the main technical tools that enable combining k-closure with lifting. The following lemmas are analogous to those in [8, 2], except that we employ lifting by a k-surjective collection of gadgets instead of 1-stifling gadgets and use the notion of k-closure in place of closure.

The next lemma states that for a satisfiable linear system whose set of linear forms is k-safe, every assignment to the unlifted variables can be induced by some solution of the system.

Lemma 5.3. Let Ψ be a satisfiable linear system in the lifted variables X and $L(\Psi)$ be k-safe. Let $g_1, g_2, \ldots, g_m : \{0,1\}^\ell \to \{0,1\}$ be an k-surjective collection of gadgets. Then for any full assignment σ to the unlifted variables Y there exists a full assignment τ to the lifted variables X such that τ satisfies Ψ and $\mathrm{induced}_{g_1,g_2,\ldots,g_m}(\tau) = \sigma$.

Proof. By Proposition 5.2, the affine space of solutions of Ψ is k-special. Then the lemma follows by the definition of a k-surjective collection of gadgets. □

The next lemma shows that for a satisfiable linear system, by varying its solution, one can arbitrarily change the values of unlifted variables outside the k-closure.

Lemma 5.4. Let Ψ be a satisfiable linear system in the lifted variables X. Let $g_1, g_2, \ldots, g_m : \{0,1\}^\ell \to \{0,1\}$ be k-surjective collection of gadgets. Suppose

- σ is a full assignment to lifted variables X satisfying Ψ.
- π is a full assignment to unlifted variables Y such that $\pi|_{\mathrm{Cl}^{(k)}(L(\Psi))} = \mathrm{induced}_{g_1,g_2,\ldots,g_m}(\sigma)|_{\mathrm{Cl}^{(k)}(L(\Psi))}$.

Then there exists a full assignment τ to the lifted variables X such that τ satisfies Ψ and $\mathrm{induced}_{g_1,g_2,\ldots,g_m}(\tau) = \pi$.

Proof. Let T be the set of all lifted variables with support in $\mathrm{Cl}^{(k)}(L(\Psi))$. Let σ_0 be the restriction of σ to T. The linear system $(\Psi)|_{\sigma_0}$ is satisfiable, and its set of linear forms is k-safe by the definition of closure. By Lemma 5.3, there exists an assignment γ to the lifted variables $\mathrm{Vars}(\Psi) \setminus T$ that satisfies $(\Psi)|_{\sigma_0}$ and such that $\mathrm{induced}_{g_1,g_2,\ldots,g_m}(\sigma_0 \cup \gamma) = \pi$. Thus, we can take $\tau = \sigma_0 \cup \gamma$. □

In most applications, we will only need the following consequence of Lemma 5.4.

Lemma 5.5. Let Ψ be a satisfiable linear system in the lifted variables X. Let $g_1, g_2, \ldots, g_m : \{0,1\}^\ell \to \{0,1\}$ be k-surjective set of gadgets. Suppose there exists a full assignment σ to lifted variables X satisfying Ψ such that $\mathrm{induced}_{g_1,g_2,\ldots,g_m}(\sigma)|_{\mathrm{Cl}^{(k)}(L(\Psi))}$ does not falsify any clause of φ. Then, Ψ does not contradict any clause of $\mathrm{Lift}_{g_1,g_2,\ldots,g_m}(\varphi)$.

Proof. Consider a clause C' from $\mathrm{Lift}_{g_1,g_2,\ldots,g_m}(\varphi)$ and a clause C from φ such that C' is a clause from $\mathrm{Lift}_{g_1,g_2,\ldots,g_m}(C)$.

Since $\mathrm{induced}_{g_1,g_2,\ldots,g_m}(\sigma)|_{\mathrm{Cl}^{(k)}(L(\Psi))}$ does not falsify any clause of φ, there exists a full assignment of unlifted variables π that extends $\mathrm{induced}_{g_1,g_2,g_m}(\sigma)|_{\mathrm{Cl}^{(k)}(L(\Psi))}$ and satisfies C. By Lemma 5.4, there exists an assignment τ of the lifted variables that satisfies Ψ and such that $\mathrm{induced}_{g_1,g_2,g_m}(\tau) = \pi$. Hence, τ satisfies $\mathrm{Lift}_{g_1,g_2,\ldots,g_m}(C)$ and, thus, τ satisfies C'. □

5.3 Existence of k-surjective gadget collections when $k \approx \ell/2$

Let $k < \ell$. A gadget $g : \{0,1\}^\ell \to \{0,1\}$ is called *k-stifling* [15] if for every $A \subset [\ell]$ of size k for every $c \in \{0,1\}$ there exists $a \in \{0,1\}^\ell$ such that for every $b \in \{0,1\}^\ell$ if a and b agree on set of indices $[\ell] \setminus A$, then $g(b) = c$.

It is easy to see that the majority function $Maj_{2k+1} : \{0,1\}^{2k+1} \to \{0,1\}$ is k-stifling.

Remark 5.6. If all $g_1, g_2, \ldots g_m : \{0,1\}^\ell \to \{0,1\}$ are k-stifling, then $g_1, g_2, \ldots g_m$ are k-surjective.

Proof. Since all gadgets are k-stifling, any block can be fixed arbitrarily using only free variables. □

Notice that if $g : \{0,1\}^\ell \to \{0,1\}$ is k-stifling, then necessarily $2k+1 \leq \ell$. In what follows, we prove the existence of a k-surjective collection of gadgets for the regime $k \geq (1-\epsilon)\ell$, where $\epsilon \leq \frac{1}{2}$.

5.4 Existence of k-surjective gadget collections when k is close to ℓ

Proposition 5.7. The total number of affine subspaces of $\mathbb{F}_2^n$ is at most 2^{n^2+n}.

Proof. Every affine subspace of $\mathbb{F}_2^n$ can be described as the solution set of a linear system $Ax = b$, where $A \in \mathbb{F}_2^{n \times n}$ and $b \in \mathbb{F}_2^n$. Since there are exactly 2^{n^2} choices for A and 2^n choices for b, the total number of such systems is 2^{n^2+n}. □

Theorem 5.8. Assume that $2^{\epsilon\ell} > (\ell m)^2 + \ell m + m$ and $k \le (1-\epsilon)\ell$. Then there exists a k-surjective collection of gadgets $g_1, g_2, \dots g_m : \{0,1\}^\ell \to \{0,1\}$.

Proof. We prove by induction on n that for every $n \in [m]$ there exists a k-surjective collection of gadgets $g_1, g_2, \dots g_n : \{0,1\}^\ell \to \{0,1\}$. The base of induction is $n = 1$. The number of k-special subspaces of $\mathbb{F}_2^\ell$ can be estimated from above as the total number of affine subspaces that is at most $2^{\ell^2+\ell}$ by Proposition 5.7. Every k-special subspace of $\mathbb{F}_2^\ell$ contains one block with at least $\ell - k$ free variables, hence its size is at least $2^{\ell-k}$. Let us choose g_1 uniformly at random from the set of all functions $\{0,1\}^\ell \to \{0,1\}$. For every k-special subspace, the probability that g_1 is constant on this subspace is at most $2^{-2^{\ell-k}+1}$. Since $2^{\ell^2+\ell} 2^{-2^{\ell-k}+1} \le 2^{\ell^2+\ell+1-2^{\epsilon\ell}} < 1$, there exists g_1 that is not constant on any k-special subspace of $\mathbb{F}_2^\ell$.

Induction step. Assume that there exist a k-surjective collection of gadgets $g_1, g_2, \dots, g_n$, where $n < m$. Let us show that we can extend this set by g_{n+1}.

Let L be a k-special affine subspace of $\mathbb{F}_2^{(n+1)\ell}$. Suppose that the last block of L contains t free variables; by definition, we have $t \ge \ell - k$. There are 2^t ways to assign values to these t variables. Each such assignment induces a k-special affine subspace on the first $n\ell$ coordinates, which we index by a string $\alpha \in \{0,1\}^t$ and denote by L_α.

We take g_{n+1} uniformly at random from the set of all functions $\{0,1\}^\ell \to \{0,1\}$.

Consider arbitrary bits $b_1, b_2, \dots, b_n \in \{0,1\}$. By the induction hypothesis, for every $\alpha \in \{0,1\}^t$, there exists $x^{(\alpha)} \in L_\alpha$ such that

$$g_i\left(x_{i,1}^{(\alpha)}, x_{i,2}^{(\alpha)}, \dots, x_{i,\ell}^{(\alpha)}\right) = b_i \text{ for all } i \in [n].$$

Let $y^{(\alpha)} \in L$ be the element whose projection onto the first n blocks coincides with $x^{(\alpha)}$ and whose t free coordinates in the last block are exactly α. Notice that for distinct α, the strings $y_{n+1,1}^{(\alpha)} y_{n+1,2}^{(\alpha)} \dots y_{n+1,\ell}^{(\alpha)}$ are different. Therefore, the probability that g_1 is constant on all these strings is at most $2^{-2^t+1} \le 2^{-2^{\epsilon\ell}+1}$.

Hence, the probability that $g_1, g_2, \dots, g_n, g_{n+1}$ fail to form a k-surjective collection of gadgets is at most

$$2^{-2^{\epsilon\ell}+1} \times (\text{number of affine subspaces } L) \times (\text{number of choices for } b_1, \dots, b_n).$$

Using Proposition 5.7, this probability is at most $2^{-2^{\epsilon\ell}+1} \cdot 2^{(\ell(n+1))^2+(n+1)\ell+n} \le 2^{-2^{\epsilon\ell}+1} \cdot 2^{(\ell m)^2+m\ell+m-1} < 1$. □

6 Lossless lifting from resolution width to Res(⊕) width

In this section, we generalize the lifting result of Itsykson and Alekseev [2] from resolution width to Res(⊕) width. Our approach replaces 1-stifling gadgets with a k-surjective collection of gadgets. Consequently, the width of the lifted formula increases by a factor of k relative to the original formula. Using the construction from Theorem 5.8, where k is nearly equal to the gadget size, we ensure that the width-to-variable ratio of the lifted formula remains nearly the same as in the original formula.

In Section 2 we presented the game-theoretic characterization of resolution width due to Atserias and Dalmau [3]. A similar width characterization is known for Res(⊕) [18, 2]. We follow the construction from [2], which is slightly more convenient for our purposes.

Let φ be an unsatisfiable CNF formula. A w-winning Res(⊕) strategy for φ is a non-empty family $\mathcal{G}$ of linear systems over the variables of φ such that:

- For every $\Phi \in \mathcal{G}$ and every clause C in φ, there exists a solution of Φ that satisfies C.
- For every $\Phi \in \mathcal{G}$, we have $\mathrm{rk}(\Phi) \le w$.
- If $\Phi \in \mathcal{G}$ and Φ semantically implies Ψ, then $\Psi \in \mathcal{G}$.
- If $\Phi \in \mathcal{G}$ and $\mathrm{rk}(\Phi) \le w - 1$ and f is a linear form, then there exists $a \in \mathbb{F}_2$ such that $\Phi \wedge \{f = a\} \in \mathcal{G}$.

Analogously to Lemma 2.1, the following holds:

Lemma 6.1 ([2])**.** Let φ be an unsatisfiable CNF formula. If there exists a $(W+1)$-winning Res($\oplus$) strategy for φ, then φ has no Res($\oplus$) refutation of width at most W.

Definition 6.2. Let $\mathcal{A}$ be a set of partial assignments to the variables Y, and let $g_1, g_2, \ldots, g_m : \{0,1\}^\ell \to \{0,1\}$ be a collection of gadgets. We define $\mathrm{S}(\mathcal{A}; g_1, g_2, \ldots, g_m)$ to be the set of linear systems Φ over the lifted variables X such that there exists a solution τ to Φ with the property that the restriction of $\mathrm{induced}_{g_1,g_2,\ldots,g_m}(\tau)$ to $\mathrm{Cl}^{(k)}(L(\Phi))$ belongs to $\mathcal{A}$.

The following lemma extends Theorem 3.1 of [2] to the setting of lifting with a k-surjective collection of gadgets.

Lemma 6.3. Let $\mathcal{H}$ be a w-winning resolution strategy for an unsatisfiable CNF formula $\varphi(y_1, y_2, \ldots, y_m)$. Let $g_1, g_2, \ldots, g_m$ be a k-surjective collection of gadgets. Define $\mathcal{G} := \{\Phi \in \mathrm{S}(\mathcal{H}; g_1, g_2, \ldots, g_m) \mid \mathrm{rk}(\Phi) \le wk\}$. Then $\mathcal{G}$ is a wk-winning Res($\oplus$) strategy for $\mathrm{Lift}_{g_1,g_2,\ldots,g_m}(\varphi)$.

Proof. Let us verify that $\mathcal{G}$ satisfies all the properties of wk-winning Res($\oplus$) strategy:

- By definition, for every $\Phi \in \mathcal{G}$, $\mathrm{rk}(\Phi) \le wk$.
- By Lemma 5.5, for every $\Phi \in \mathcal{G}$ and every clause C' from $\mathrm{Lift}_{g_1,g_2,\ldots,g_m}(\varphi)$, there exists a solution of Φ that satisfies C'.
- Let us show that if Ψ is a linear system and for some $\Phi \in \mathcal{G}$, Φ semantically implies Ψ, then $\Psi \in \mathcal{G}$. Indeed, since Φ semantically implies Ψ, $L(\Psi) \subseteq \langle L(\Phi)\rangle$. Then by Lemmas 4.4 and 4.5, $\mathrm{Cl}^{(k)}(L(\Psi)) \subseteq \mathrm{Cl}^{(k)}(L(\Phi))$. Clear that $\mathrm{rk}(\Psi) \le \mathrm{rk}(\Phi) \le wk$. Since, $\Phi \in \mathcal{G}$ there exist $\tau \in \mathcal{H}$ and there is a solution σ of Φ such that $\mathrm{induced}_{g_1,g_2,\ldots,g_m}(\sigma)$ coincides with τ on $\mathrm{Cl}^{(k)}(L(\Phi))$. Notice that σ is also a solution of Ψ and $\mathrm{induced}_{g_1,g_2,\ldots,g_m}(\sigma)$ coincides with h on $\mathrm{Cl}^{(k)}(L(\Psi))$. Hence Ψ is in $\mathcal{G}$.
- Finally, we need to show that for any $\Phi \in \mathcal{G}$ with $\mathrm{rk}(\Phi) < wk$ and for every linear form f, there exists a constant $a \in \mathbb{F}_2$ such that $\Phi \wedge (f = a) \in \mathcal{G}$.

  There exist $\tau \in \mathcal{H}$ and a solution σ of Φ such that $\mathrm{induced}_{g_1,g_2,\ldots,g_m}(\sigma)$ coincides with τ on $\mathrm{Cl}(L(\Phi))$. W.l.o.g., assume that the domain of τ is precisely $\mathrm{Cl}(L(\Phi))$. By Lemma 4.6, $|(\mathrm{Cl}^{(k)}(L(\Phi) \cup \{f\}))| \le (\mathrm{rk}(\Phi)+1)/k \le w$. By the properties of $\mathcal{H}$ there is $\pi \in \mathcal{H}$ such that $\tau \subseteq \pi$ and π is defined on $\mathrm{Cl}^{(k)}(L(\Phi) \cup \{f\})$; indeed, we can extend τ for all variables from $\mathrm{Cl}^{(k)}(L(\Phi) \cup \{f\}) \setminus \mathrm{Cl}(L(\Phi))$ one by one. Using Lemma 5.4, we can find a solution θ of Φ such that $\mathrm{induced}_{g_1,g_2,\ldots,g_m}(\theta)$ coincides with π on $\mathrm{Cl}(L(\Phi) \cup \{f\}))$. Let a be a value of linear form f on the solution θ. Then θ clearly satisfies $\Phi \wedge \{f = a\}$. On the other hand, $\mathrm{induced}_{g_1,g_2,\ldots,g_m}(\theta)$ coincides with π on $\mathrm{Cl}(L(\Phi) \cup \{f\}))$. Thus $\Phi \wedge \{f = a\} \in \mathcal{G}$.

□

7 Freedom in the lifted world

In this section, we demonstrate that the wk-winning Res($\oplus$)-strategy obtained through the lifting in Lemma 6.3 enjoys a high degree of freedom. Consider a sequence of winning positions $\Phi_1, \Phi_1, \ldots, \Phi_{wk+1}$, where Φ_0 is the empty position and $\Phi_{i+1} = \Phi_i \wedge (f_{i+1} = a_{i+1})$. We show that in the regime $k \approx \ell$, for almost all indices $i \in [wk]$, the current position Φ_i can be extended to a winning position for either assignment of the next linear form f_{i+1}.

Let $\mathcal{A}$ be a set of partial assignments for the variables of φ. We assume that $\mathcal{A}$ is *proper* for φ, that is, it satisfies the following properties:

- for all $\rho \in \mathcal{A}$ and $\sigma \subseteq \rho$, $\sigma \in \mathcal{A}$ (closure under restrictions);
- for all $\sigma \in \mathcal{A}$, σ does not falsify any clause of φ.

We introduce the $(\varphi, \mathcal{A})$-*game of Prover and Adversary*. The game proceeds as follows: the players maintain a partial assignment ρ to the variables of φ, starting from the empty assignment. In each move, Prover chooses a variable x, after which Adversary earns one coin and selects a Boolean value a for x. The assignment is updated to $\rho := \rho \cup \{x := a\}$. The game ends once ρ leaves the set $\mathcal{A}$. The Adversary's goal is to collect as many coins as possible.

Proposition 7.1. Let $\mathcal{A}$ be a w-winning resolution strategy for φ. Then $\mathcal{A}$ is proper, and in the $(\varphi, \mathcal{A})$-game Adversary has a strategy that guarantees him at least w coins.

Proof. The proof is straightforward. □

Let φ be an unsatisfiable CNF formula and let $\mathcal{G}$ be a set of linear systems over the variables of φ. By analogy, we say that $\mathcal{G}$ is *proper* if it satisfies the following conditions:

- Every linear system in $\mathcal{G}$ is consistent with all clauses of φ;
- For every $\Phi \in \mathcal{G}$, any linear system Ψ such that $\Phi \models \Psi$ also belongs to $\mathcal{G}$.

Similar to Prover-Delayer games from Section 2, we define a $(\varphi, \mathcal{G})$-$\oplus$-game of Prover and Delayer. In this game, two players, Prover and Delayer, maintain a linear system Φ in variables of φ that initially is the empty linear system (i.e., constant true). On every move, Prover chooses a linear form f that does not belong to $\langle L(\Phi)\rangle$, and Delayer has two options:

- Delayer can earn a *white* coin and reports $*$. Then, Prover chooses a Boolean value a of f.
- Delayer can earn a *white* coin and pay a *black* coin to choose a Boolean value a of f by himself.

The current linear system Φ is updated: $\Phi := \Phi \wedge \{f = a\}$. The game ends when $\Phi \notin \mathcal{G}$.

Lemma 7.2. Let $\varphi(y_1, y_2, \ldots, y_m)$ be an unsatisfiable CNF formula, and let $\mathcal{A}$ be a proper set of partial assignments for φ. Suppose the Adversary has a strategy in the $(\varphi, \mathcal{A})$-game that guarantees him to earn at least w coins. Let $g_1, g_2, \ldots, g_m$ be a k-surjective collection of gadgets, and define $\mathcal{G} := \mathrm{S}(\mathcal{A}; g_1, g_2, \ldots, g_m)$ as in Definition 6.2. Then, in the Prover–Delayer Res($\oplus$)-game for $(\mathrm{Lift}_{g_1,g_2,\ldots,g_m}(\varphi), \mathcal{G})$, Delayer has a strategy that allows him to earn at least wk white coins while paying at most $(\ell - k)w$ black coins.

Proof. Let us describe the Delayer's strategy. In parallel, we will play the $(\varphi, \mathcal{A})$-game using the Adversary's strategy. Let Φ be the current linear system in the $(\mathrm{Lift}_{g_1,g_2,\ldots,g_m}(\varphi), \mathcal{G})$-game, and let ρ be the current partial assignment in the $(\varphi, \mathcal{A})$-game. Initially, Φ is the empty linear system and ρ is the empty assignment. Denote $F := L(\Phi)$. We maintain the following invariant: there exists a solution τ to Φ such that

$$\mathrm{induced}_{g_1,g_2,\ldots,g_m}(\tau)|_{\mathrm{Cl}^{(k)}(F)} = \rho.$$

The Delayer continues the game until $\dim\langle F\rangle \le wk$, which guarantees him wk white coins, since every request made by the Prover is linearly independent of all previous ones.

To describe Delayer's strategy, consider a step in which Prover requests the value of a linear form f.

- If $\mathrm{Cl}^{(k)}(F) = \mathrm{Cl}^{(k)}(F \cup \{f\})$ and $f[\backslash \mathrm{Cl}^{(k)}(F)]$ is in $\langle F[\backslash \mathrm{Cl}^{(k)}(F)]\rangle$, then Delayer chooses a value $f[\tau]$ and pays a black coin.
- If $\mathrm{Cl}^{(k)}(F) = \mathrm{Cl}^{(k)}(F \cup \{f\})$ and $f[\backslash \mathrm{Cl}^{(k)}(F)]$ is not in $\langle F[\backslash \mathrm{Cl}^{(k)}(F)]\rangle$, then Delayer answers $*$. In this case, $a \in \{0, 1\}$ is choosen by Prover. Let ρ denote the restriction of τ to variables with support $\mathrm{Cl}^{(k)}(F)$. The new linear system $\Phi \wedge (f = a)$ has the solution with the same values of variables from the domain of ρ as the solution τ since we just added a linearly independent equation to the satisfiable linear system $\Phi|_\rho$.

- In the last case, we have $\mathrm{Cl}^{(k)}(F\cup\{f\})\setminus\mathrm{Cl}^{(k)}(F) = T$ for some nonempty set T. If $|\mathrm{Cl}^{(k)}(F\cup\{f\})| \le w$, we extend ρ to a partial assignment ρ' on $\mathrm{Cl}^{(k)}(F\cup\{f\})$ according to the Adversary's strategy. By Lemma 5.4, there exists a solution τ' of Φ such that

$$\mathrm{induced}_{g_1,g_2,\ldots,g_m}(\tau')\big|_{\mathrm{Cl}^{(k)}(F\cup\{f\})} = \rho'.$$

  Delayer then selects a value $a = f[\tau']$, pays one black coin, and updates the state as follows:

$$\Phi := \Phi \wedge (f = a), \qquad \rho := \rho', \qquad \tau := \tau'.$$

  If instead $|\mathrm{Cl}^{(k)}(F\cup\{f\})| > w$, Delayer *gives up*, ensuring that $\mathrm{Cl}^{(k)}(F)$ always remains of size at most w.

We claim that at any time before Delayer gives up, the quantity $k|\mathrm{Cl}^{(k)}(F)| + \dim\langle F[\setminus\mathrm{Cl}^{(k)}(F)]\rangle$ records the number of answers $*$. We prove this by induction on the number of moves made. The base corresponds to the start of the game, and the statement is trivial.

If $\mathrm{Cl}^{(k)}(F) = \mathrm{Cl}^{(k)}(F\cup\{f\})$ and $f[\setminus\mathrm{Cl}^{(k)}(F)]$ is in $\langle F[\setminus\mathrm{Cl}^{(k)}(F)]\rangle$, then Delayer does not answer $*$, $\dim\langle F[\setminus\mathrm{Cl}^{(k)}(F)]\rangle$ and $\mathrm{Cl}^{(k)}(F)$ are not changed.

If $\mathrm{Cl}^{(k)}(F) = \mathrm{Cl}^{(k)}(F\cup\{f\})$ and $f[\setminus\mathrm{Cl}^{(k)}(F)]$ is not in $\langle F[\setminus\mathrm{Cl}^{(k)}(F)]\rangle$, then Delayer answers $*$, $\dim\langle F[\setminus\mathrm{Cl}^{(k)}(F)]\rangle$ increases by one, and $\mathrm{Cl}^{(k)}(F)$ does not change.

If $T = \mathrm{Cl}^{(k)}(F\cup\{f\}) \setminus \mathrm{Cl}^{(k)}(F) \neq \emptyset$, then Delayer does not earn a coin, $\mathrm{Cl}^{(k)}(F)$ increases by $|T|$ and, by Lemma 4.7, $\dim\langle F[\setminus \mathrm{Cl}(F)]\rangle$ decreases by $k|T|$. This finishes the inductive step.

At the moment when Delayer gives up, the number of black coins paid equals the number of answers different from $*$. This quantity is

$$\dim\langle F\rangle - k|\mathrm{Cl}^{(k)}(F)| - \dim\langle F[\setminus\mathrm{Cl}^{(k)}(F)]\rangle,$$

which, by Lemma 4.2, is equal to

$$\dim\{f \in \langle F\rangle \mid \mathrm{supp}(f) \subseteq \mathrm{Cl}^{(k)}(F)\} - k|\mathrm{Cl}^{(k)}(F)|.$$

Since this is at most $(\ell-k)|\mathrm{Cl}^{(k)}(F)| \le w(\ell-k)$, the claim follows.

□

The proof of Lemma 7.2 is very similar to the proof of the tree-like Res($\oplus$) lower bound for the binary pigeonhole principle from [16].

8 SETH-Type Lower Bound for Bounded-Depth Res($\oplus$)

In this section, we prove our main result.

We say that a Res($\oplus$) refutation contains an (a,b)-*path* if, in the parity decision DAG associated with the refutation, there exists a node v such that the path from the source to v passes through at least a linearly independent linear forms, and the path from v to a sink passes through at least b linearly independent linear forms.

We say that a Res($\oplus$) refutation is (a,b)-*path-free* if it does not contain any (a,b)-paths.

Proposition 8.1. If depth of a Res($\oplus$) refutation is less than $a+b$, then it is (a,b)-path-free.

Proof. The proof is straightforward. □

Theorem 8.2. Let $\varphi(y_1,y_2,\ldots,y_m)$ be an unsatisfiable CNF formula whose minimal resolution width w satisfies $w \ge (1-\delta)m$. Let $g_1,g_2,\ldots,g_m : \{0,1\}^\ell \to \{0,1\}$ form a k-surjective collection of gadgets, where $\ell - k \le \epsilon\ell$. Then

- the width of any Res($\oplus$) refutation of $\mathrm{Lift}_{g_1,g_2,\ldots,g_m}(\varphi)$ is at least $m\ell(1-\epsilon-\delta)$;
- every $((1-\epsilon-\delta)m\ell,\, 2(\epsilon+\delta)m\ell)$-path-free Res($\oplus$) refutation of $\mathrm{Lift}_{g_1,g_2,\ldots,g_m}(\varphi)$ has size at least $2^{m\ell(1-4\epsilon-3\delta)}$.

Proof. By Lemma 2.1, there exists a w-winning resolution strategy $\mathcal{A}$ for φ. Define

$$\mathcal{G} := \{\Phi \in \mathrm{S}(\mathcal{A}; g_1, g_2, \ldots, g_m) \mid \mathrm{rk}(\Phi) \leq wk\},$$

where S is as in Definition 6.2. Then, by Lemma 6.3, $\mathcal{G}$ forms a wk-winning Res($\oplus$) strategy for $\mathrm{Lift}_{g_1,g_2,\ldots,g_m}(\varphi)$. Consequently, by Lemma 6.1, the width of any Res($\oplus$) refutation of $\mathrm{Lift}_{g_1,g_2,\ldots,g_m}(\varphi)$ is at least $wk \geq (1-\epsilon-\delta)m\ell$.

Consider a Res($\oplus$) refutation of $\mathrm{Lift}_{g_1,g_2,\ldots,g_m}(\varphi)$ of depth at most $m\ell$, together with its associated parity decision DAG. Let τ be a random full assignment to the lifted variables X. We define a path in the parity decision DAG induced by τ as follows: we start at the source and, at each step, follow the outgoing edge whose labeling linear equation is satisfied by τ. We proceed the path until we meet wk linearly independent linear forms, if a sink is reached before, the path terminates there. The random variable of interest is the linear system labeling the node at the end of this path.

The same distribution can be generated by the following random process. We start at the source. Whenever the queried linear form is linearly independent of all previous queries along the current path, we follow a uniformly random outgoing edge. If instead the queried linear form is linearly dependent on the previous ones, we deterministically follow the edge dictated by the linear system defined by the labels of the edges already traversed.

Consider a $(\varphi, \mathcal{A})$-game of Prover and Adversary. By Proposition 7.1, in the $(\varphi, \mathcal{A})$-game Adversary has a strategy that guarantees him at least w coins. By Lemma 7.2, in the $(\mathrm{Lift}_{g_1,g_2,\ldots,g_m}(\varphi), \mathcal{G})$-game Delayer has a strategy that guarantees him at least wk white coins while paying at most $w(\ell-k)$ black coins.

Under the second view of the random process, the existence of the Delayer's strategy implies that, with probability at least $2^{-w(\ell-k)}$, the linear system encountered along the path is consistent with the Delayer's strategy. Consequently, with probability at least $2^{-w(\ell-k)}$, the linear system Ψ defined by the edges of the random walk belongs to $\mathcal{G}$. Moreover, the linear system Φ labeling the endpoint of the path is a semantic consequence of Ψ; hence, in this case, we also have $\Phi \in \mathcal{G}$.

Claim 8.3. Assume that a node v of the parity decision DAG is at distance wk from the source and is labeled with a linear system $\Phi_v \in \mathcal{G}$. Then the rank of Φ_v is at least $ml(1-3\epsilon-3\delta)$.

Proof. Consider a path in the parity decision DAG starting from a node v, where at each step we move from a vertex u along the edge labeled by $f_u := a$ such that $\Phi_u \cup \{f_u := a\} \in \mathcal{G}$. Let u' denote the endpoint of this edge. Since $\mathcal{G}$ is a winning Res($\oplus$) strategy and by the properties of the parity decision DAG, $\Phi_{u'} \in \mathcal{G}$. Since $\mathcal{G}$ is a wk-winning Res($\oplus$) strategy, we can construct a path that contains at least $wk - \mathrm{rk}(\Phi_v)$ linearly independent linear forms. As the parity decision DAG does not contain $((1-\epsilon-\delta)m\ell, 2(\epsilon+\delta)m\ell)$-path, $wk - \mathrm{rk}(\Phi_v) \leq 2(\epsilon+\delta)m\ell$, hence $\mathrm{rk}(\Phi_v) \geq wk - 2(\epsilon+\delta)m\ell \geq ml(1-3\epsilon-3\delta)$.

□

Note that any linear system of rank at least t is satisfied by a random assignment τ with probability at most 2^{-t}. Since, with probability at least $2^{-w(\ell-k)}$, the endpoint of a random path satisfies the conditions of Claim 8.3 (we refer to this as a *lucky event*), the number of distinct nodes that can appear as endpoints of the random path in lucky events is at least

$$\frac{2^{-w(\ell-k)}}{2^{-m\ell(1-3\epsilon-3\delta)}} \;\geq\; 2^{m\ell(1-4\epsilon-3\delta)}.$$

□

Corollary 8.4. For any $\epsilon > 0$ for all large enough n there exists $m = O(\log n \cdot \mathrm{poly}(1/\epsilon))$ and $N = O(\frac{n \log n}{\epsilon})$ and an unsatisfaible m-CNF formula φ over N variables and of size $n^{\mathrm{poly}(1/\epsilon)}$ such that

- any Res($\oplus$) refutation of φ has width at least $N(1-\epsilon)$;
- any Res($\oplus$) refutation of φ that is free of $((1-\epsilon)N, 2\epsilon N)$-paths has size at least $2^{(1-\frac{7}{2}\epsilon)N}$.

Proof. Let us choose $\delta = \epsilon/2$.

By Lemma 2.9, there exists $m_1 = \tilde{O}(1/\delta^3)$ such that for large enough n there exists an unsatisfiable formula ψ in m_1-CNF over n variables such that resolution width of ψ is at least $(1-\delta)n$.

Let us choose $\epsilon' = \epsilon/2$ and $\ell = O(\frac{\log n}{\epsilon})$ such that $2^{\epsilon'\ell/2} > (\ell n)^2 + \ell n + n$ and there exists integer number k such that $(1-\epsilon')\ell \leq k \leq (1-\epsilon'/2)\ell$. Then by Theorem 5.8, there exists a k-surjective collection of gadgets $g_1, g_2, \ldots g_n : \{0,1\}^\ell \to \{0,1\}$.

Let $\varphi = \mathrm{Lift}_{g_1,g_2,\ldots,g_n}(\psi)$ and let $N := n\ell$ denote the number of variables of φ. By Theorem 8.2,

- any Res($\oplus$) refutation of φ has width at least $N(1-\epsilon_1-\delta) = N(1-\epsilon)$;
- any $((1-\epsilon)N, 2\epsilon N)$-path-free Res($\oplus$) refutation of φ has size at least $2^{N(1-\frac{7}{2}\epsilon)}$.

The number of variables of φ is $N = n\ell = O(n \log n/\epsilon)$. The CNF representation of $\mathrm{Lift}_{g_1,g_2,\ldots,g_n}(\psi)$ is constructed as follows. For each clause C of ψ, we translate $\mathrm{Lift}_{g_1,g_2,\ldots,g_n}(C)$ into CNF and then take the conjunction of all these translations over all clauses C of ψ.

The formula $\mathrm{Lift}_{g_1,g_2,\ldots,g_n}(C)$ involves $\ell m_1 = \log n \cdot \mathrm{poly}(1/\epsilon)$ variables. Hence, for any clause C of ψ, the formula $\mathrm{Lift}_{g_1,g_2,\ldots,g_n}(C)$ can be represented as an $O(\log n \cdot \mathrm{poly}(1/\epsilon))$-CNF formula of size $n^{\mathrm{poly}(1/\epsilon)}$. Since ψ is an m_1-CNF formula, it contains at most $n^{\mathrm{poly}(1/\epsilon)}$ clauses. Therefore, $\mathrm{Lift}_{g_1,g_2,\ldots,g_n}(\psi)$ can be represented as an $O(\log n \cdot \mathrm{poly}(1/\epsilon))$-formula of overall size $n^{\mathrm{poly}(1/\epsilon)}$.

□

A Proofs of facts about k-closure

A.1 Proof of Theorem 4.1

The following theorem generalizes the classical Hall's theorem.

Theorem A.1 ([30])**.** Suppose that L is a vector space, $V_1, V_2, \ldots, V_n$ are sets of vectors from L such that for every $A \subseteq [n]$ the dimension of $\langle \cup_{i\in A} V_i \rangle$ is at least $|A|$. Then there exist vectors $v_1 \in V_1, v_2 \in V_2, \ldots, v_n \in V_n$ such that $v_1, v_2, \ldots, v_n$ are linearly independent.

Corollary A.2. Suppose that L is a vector space, $V_1, V_2, \ldots, V_n$ are sets of vectors from L and $s \in [n]$ is such that for every $A \subseteq [n]$ the dimension of $\langle \cup_{i\in A} V_i \rangle$ is at least $k|A| - s$. Then there exist pairwise disjoint sets $U_1, U_2, \ldots, U_n$ such that for every $j \in [n]$, $U_j \subseteq V_j$, $|U_j| \leq k$, all vectors from the set $\bigcup_{i\in[n]} U_i$ are linearly independent, and $\sum_{i=1}^n |U_i| = kn - s$.

Proof. The proof proceeds by a reduction to Theorem A.1. First, choose a set S of s linearly independent vectors that is also independent of $\langle \bigcup_{i=1}^n V_i \rangle$. If L does not contain such s vectors, we embed L into a larger vector space and locate these s vectors there. Now construct nk sets as follows: for each $i \in [n]$, take k copies of $V_i \cup S$. By construction, these nk sets satisfy the assumptions of Theorem A.1, and therefore we can select linearly independent representatives, one from each set. Finally, discard those representatives that lie in S. Since at most s of them can come from S, the remaining representatives form the desired family. □

Theorem 4.1. Let $f_1, f_2, \ldots, f_s$ be linearly independent linear forms and let M be their coefficient matrix. Then the following conditions are equivalent.

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

(2) For every set $T \subseteq [m]$, the dimension of the span of the set of columns of M corresponding to the variables with support in T is at least $k|T| - (km - s)$.

(3) One can choose s variables such that for every block at most k variables are chosen from this block and the columns of M corresponding to the s chosen variables are linearly independent. Since the rank of M is s, the chosen columns form the basis of the span of M's columns.

Proof of Theorem 4.1. Let us prove the equivalence of the first two conditions. Consider an arbitrary set of blocks $T \subseteq [m]$. Consider a submatrix M_T of M that contains only the columns indexed by variables with support in T. Consider the vector space $V_T \subseteq \{0,1\}^s$ consisting of all vectors that have zero inner product with every column of M_T. The dimension of V_T equals $s - \operatorname{rk}(M_T)$. Consider the space $H_T = \langle \sum \alpha_i f_i \mid \alpha \in V_T \rangle$. Notice that $H_T = \{g \in \langle f_1, f_2, \ldots, f_s \rangle \mid \operatorname{supp}(g) \subseteq [m] \setminus T\}$. Since $f_1, f_2, \ldots, f_s$ are linearly independent, $\dim H_T = \dim V_T = s - \operatorname{rk}(M_T)$.

The set $f_1, f_2, \ldots, f_k$ is k-safe if and only if for every $T \subseteq [m]$, $\dim H_T \leq k(m - |T|)$ which is equivalent to $\operatorname{rk}(M_T) \geq s - k(m - |T|)$. Thus, items (1) and (2) are equivalent.

Now assume (2) and let us prove (3). Consider sets of vectors $V_1, V_2, \ldots, V_m$, where V_i consists of columns of M corresponding to the block i (i.e. to variables with support $\{i\}$). By Corollary A.2 applied to $V_1, V_2, \ldots, V_m$ and $(km - s)$, there exist pairwise disjoint sets $U_1, U_2, \ldots, U_m$ such that for every $j \in [m]$ $U_j \subseteq V_j$, $|U_j| \leq k$, all vectors from the set $\bigcup_{i \in [m]} U_i$ are linearly independent, and $\sum_{i=1}^{m} |U_i| = s$. Let us choose variables corresponding to the columns from $\cup_{i \in [m]} U_i$, the number of chosen variables is s, every block contains at most k chosen variables, and all columns corresponding to them are linearly independent. Thus, the third condition holds.

Finally, assume that the third condition holds and there are s chosen linearly independent columns of M such that each block contains at most k chosen columns. Let $T \subseteq [m]$. At most $km - k|T|$ of the chosen columns have their corresponding block in $[m] \setminus T$, hence there are at least $s - km + k|T|$ of the chosen columns with their corresponding block in T. Therefore, the dimension of the span of the set of columns of M corresponding to variables with support in T is at least $k|T| - (km - s)$. I.e., the second condition holds. □

A.2 Uniqueness of k-closure

A set of linear forms F is *minimally k-dangerous* if it is k-dangerous and $\langle F \rangle$ does not contain a k-dangerous set with strictly smaller support than the support of F. Recall that a k-dangerous set is necessarily linearly independent.

Lemma A.3. Let H be a minimally k-dangerous set and S be a strict subset of $\operatorname{supp}(H)$. Then $H[\setminus S]$ is not k-safe.

Proof. Because H is k-dangerous, $\dim\langle H \rangle = |H| > k|\operatorname{supp}(H)|$. Since H is minimally k-dangerous, $k|S| \geq \dim\langle \{h \in \langle H \rangle \mid \operatorname{supp}(h) \subseteq S\} \rangle$. By Lemma 4.2, $\dim\langle H[\setminus S] \rangle = \dim\langle H \rangle - \dim\langle \{h \in \langle H \rangle \mid \operatorname{supp}(h) \subseteq S\} \rangle > k|\operatorname{supp}(H)| - k|S|$. Hence a basis of $H[\setminus S]$ is k-dangerous. □

Lemma 4.3. (Uniqueness) For any F its k-closure is unique.

Proof. Let S_1 and S_2 be two different k-closures of F. Then $S_1 \cap S_2$ is not a closure. Hence $\langle F[\setminus (S_1 \cap S_2)] \rangle$ contains a k-dangerous set and hence it contains a minimally k-dangerous set H. Since $\operatorname{supp}(H) \subseteq [m] \setminus (S_1 \cap S_2)$, either S_1 or S_2 does not contain $\operatorname{supp}(H)$. W.l.o.g. assume that S_1 does not contain $\operatorname{supp}(H)$. Then by Lemma A.3, the set $H[\setminus S_1] = H[\setminus (S_1 \cap \operatorname{supp}(H))]$ is not k-safe. Since $H \subseteq \langle F[\setminus (S_1 \cap S_2)] \rangle$, we have $H[\setminus S_1] \subseteq \langle F[\setminus S_1] \rangle$. This is a contradiction since S_1 is a k-closure of F and so $\langle F[\setminus S_1] \rangle$ (and hence all its subsets) has to be k-safe. □

A.3 Size bound on k-closure

Lemma A.4. Let $S \subseteq \mathrm{Cl}^{(k)}(F)$ and let $\langle F[\setminus S] \rangle$ contain a minimally k-dangerous set H. Then $\operatorname{supp}(H) \subseteq \mathrm{Cl}^{(k)}(F)$.

Proof. Assume that $\mathrm{supp}(H) \not\subseteq \mathrm{Cl}^{(k)}(F)$, then $(\mathrm{Cl}^{(k)}(F) \cap \mathrm{supp}(H)) \subsetneq \mathrm{supp}(H)$. By Lemma A.3, $H[\setminus(\mathrm{Cl}^{(k)}(F) \cap \mathrm{supp}(H))] = H[\setminus \mathrm{Cl}^{(k)}(F)]$ is not k-safe. Since $H \subseteq \langle F[\setminus S]\rangle$, we have $H[\setminus \mathrm{Cl}^{(k)}(F)] \subseteq \langle F[\setminus \mathrm{Cl}^{(k)}(F)]\rangle$ and this is a contradiction, since $\langle F[\setminus \mathrm{Cl}^{(k)}(F)]\rangle$ and all its subsets have to be k-safe by the definition of the k-closure. □

Algorithm A.5. Input: a set of linear forms F.

1. $S \leftarrow \emptyset$;
2. While $\langle F[\setminus S]\rangle$ contains k-dangerous sets:
   - Find a minimally k-dangerous set in $\langle F[\setminus S]\rangle$. Let T be its support.
   - $S \leftarrow S \cup T$.
3. Return S.

Corollary A.6. Algorithm A.5 computes $\mathrm{Cl}^{(k)}(F)$.

Proof. Each iteration of the loop increases S. Since $S \subseteq [m]$, the algorithm stops in a finite number of steps. Let $S' \subseteq [m]$ be the output of the algorithm.

Let us prove by induction that $S \subseteq \mathrm{Cl}^{(k)}(F)$ at every moment during the execution of Algorithm A.5. Initially, $S := \emptyset$, so the assertion holds. The induction step follows by Lemma A.4.

It follows that $S' \subseteq \mathrm{Cl}^{(k)}(F)$. We also know that $F[\setminus S']$ is k-safe. Thus, $S' = \mathrm{Cl}^{(k)}(F)$. □

Lemma 4.6. (Size bound) $k|\mathrm{Cl}^{(k)}(F)| + \dim\langle F[\setminus \mathrm{Cl}^{(k)}(F)]\rangle \le \dim\langle F\rangle$, and hence $|\mathrm{Cl}^{(k)}(F)| \le \frac{1}{k}\dim\langle F\rangle$.

Proof. We prove by induction that during the execution of Algorithm A.5 the following inequality holds: $k|S| + \dim\langle F[\setminus S]\rangle \le \dim\langle F\rangle$. Since the algorithm outputs $S = \mathrm{Cl}^{(k)}(F)$, we get the required inequality.

At the start of the algorithm, the inequality holds. Let us show that it holds after each step. Suppose the algorithm has found in $\langle F[\setminus S]\rangle$ a minimally k-dangerous set H with support T. As $H \subseteq \{f \in \langle F[\setminus S]\rangle \mid \mathrm{supp}(f) \subseteq T\}$, we have $\dim\{f \in \langle F[\setminus S]\rangle \mid \mathrm{supp}(f) \subseteq T\} \ge \dim\langle H\rangle > k|T|$.

By Lemma 4.2, $\dim\{f \in \langle F[\setminus S]\rangle \mid \mathrm{supp}(f) \subseteq T\} = \dim\langle F[\setminus S]\rangle - \dim\langle F[\setminus (S \cup T)]\rangle$. Therefore, $\dim\langle F[\setminus (S \cup T)]\rangle < \dim\langle F[\setminus S]\rangle - k|T|$.

Finally, $k|S \cup T| + \dim\langle F[\setminus (S \cup T)]\rangle < k|S| + k|T| + \dim\langle F[\setminus S]\rangle - k|T| = k|S| + \dim\langle F[\setminus S]]\rangle \le \dim\langle F\rangle$. In the last inequality, we use the inductive hypothesis. □

A.4 Increment of k-closure

Lemma A.7. Let F be k-safe and $f_1, f_2, \ldots, f_{ks}$ be linearly independent elements of $\langle F\rangle$ such that $\mathrm{supp}(f_1, f_2, \ldots, f_{ks}) = T$ and $|T| = s$. Then the set $F[\setminus T]$ is k-safe.

Proof. We argue by contradiction. Let $g_1, g_2, \ldots, g_t$ be a linearly independent set from $\langle F[\setminus T]\rangle$ with support S and $k|S| \le t - 1$. Let $g'_1, g'_2, \ldots, g'_t$ be elements of $\langle F\rangle$ such that $g'_i[\setminus T] = g_i$.

Then $\mathrm{supp}(\{f_1, f_2, \ldots, f_{ks}, g'_1, \ldots, g'_t\}) \subseteq S \cup T$ and the size of $S \cup T$ is at most $s + \frac{t-1}{k}$. To get a contradiction, we verify that all these forms are linearly independent. Indeed, assume that $\sum_{i=1}^{ks} \alpha_i f_i + \sum_{j=1}^{t} \beta_i g'_i = 0$. By applying $[\setminus T]$ operator to this equation we get $\sum_{j=1}^{t} \beta_i g_t = 0$, hence $\beta_i = 0$ for $i \in [t]$. Since $f_1, f_2, \ldots, f_{ks}$ are linearly independent, we get that $\alpha_i = 0$ for $i \in [ks]$. □

Lemma 4.7. (Increment) Let F be a set of linear forms and f be a linear form such that $\mathrm{Cl}^{(k)}(F \cup \{f\}) \neq \mathrm{Cl}^{(k)}(F)$. Then

$$\dim\langle F[\setminus \mathrm{Cl}^{(k)}(F)]\rangle - \dim\langle (F \cup \{f\})[\setminus \mathrm{Cl}^{(k)}(F \cup \{f\})]\rangle = k(|\mathrm{Cl}^{(k)}(F \cup \{f\})| - |\mathrm{Cl}^{(k)}(F)|).$$

Proof. Since $\mathrm{Cl}^{(k)}(F \cup \{f\})$ is strictly greater than $\mathrm{Cl}^{(k)}(F)$, the set $(F \cup f)[\setminus \mathrm{Cl}^{(k)}(F)]$ is not k-safe. Consider an arbitrary minimally k-dangerous set of linear forms $h_1, h_2, \ldots, h_s$ in $\langle (F \cup f)[\setminus \mathrm{Cl}^{(k)}(F)] \rangle$. For every $i \in [k]$, either $h_i \in \langle F[\setminus \mathrm{Cl}^{(k)}(F)] \rangle$ or $h_i \in f[\setminus \mathrm{Cl}^{(k)}(F)] + \langle F[\setminus \mathrm{Cl}^{(k)}(F)] \rangle$. We can assume that $h_1, h_2, \ldots, h_s$ have been chosen such that $I := \{i \in [s] \mid h_i \in f[\setminus \mathrm{Cl}^{(k)}(F)] + \langle F[\setminus \mathrm{Cl}^{(k)}(F)] \rangle\}$ has the minimum cardinality. We know $|I| \geq 1$, otherwise $\mathrm{Cl}^{(k)}(F)$ is not the correct closure. Moreover, it is easy to see that $|I| = 1$. Indeed, if $i_1 \neq i_2 \in I$, then we can replace the form h_{i_1} in $h_1, h_2, \ldots, h_k$ with $h_{i_1} + h_{i_2}$; this alters neither the linear independence nor the support, but $h_{i_1} + h_{i_2} \in \langle F[\setminus \mathrm{Cl}^{(k)}(F)] \rangle$, a contradiction with the minimality of $|I|$. W.l.o.g. assume that $h_i \in \langle F[\setminus \mathrm{Cl}^{(k)}(F)] \rangle$ for $i \in [s-1]$ and $h_s \in f[\setminus \mathrm{Cl}^{(k)}(F)] + \langle F[\setminus \mathrm{Cl}^{(k)}(F)] \rangle$.

Let $T = \mathrm{supp}(h_1, h_2, \ldots, h_s)$, then $T \subseteq [m] \setminus \mathrm{Cl}^{(k)}(F)$. Note that $s = |T|k + 1$, since if the support of the set $h_1, h_2, \ldots, h_s$ were smaller, then the set $h_1, h_2, \ldots, h_{s-1}$ would be dangerous and in $\langle F[\setminus \mathrm{Cl}^{(k)}(F)] \rangle$.

Claim A.8. $f[\setminus (\mathrm{Cl}^{(k)}(F) \cup T)] \in \langle F[\setminus (\mathrm{Cl}^{(k)}(F) \cup T)] \rangle$.

Proof. Let us apply the linear operator $[\setminus T]$ to the statement $f[\setminus \mathrm{Cl}^{(k)}(F)] + h_s \in \langle F[\setminus \mathrm{Cl}^{(k)}(F)] \rangle$. Since $h_s[\setminus T] = 0$, we get $f[\setminus (\mathrm{Cl}^{(k)}(F) \cup T)] \in \langle F[\setminus (\mathrm{Cl}^{(k)}(F) \cup T)] \rangle$. □

Claim A.9. $T = \mathrm{Cl}^{(k)}(F \cup \{f\}) \setminus \mathrm{Cl}^{(k)}(F)$.

Proof. By monotonicity, $\mathrm{Cl}^{(k)}(F) \subseteq \mathrm{Cl}^{(k)}(F \cup \{f\})$. Since $h_1, h_2, \ldots, h_s$ is minimally k-dangerous, it follows by Lemma A.4 that $T \subseteq \mathrm{Cl}^{(k)}(F \cup \{f\})$. The set $h_1, h_2, \ldots, h_{s-1}$ is safe, hence $|\mathrm{supp}(\{h_1, h_2, \ldots, h_{s-1}\})| = \frac{s-1}{k}$, and so $\mathrm{supp}(\{h_1, h_2, \ldots, h_{s-1}\}) = T$. Consequently, Lemma A.7 applied to $F[\setminus \mathrm{Cl}^{(k)}(F)]$ and $h_1, h_2, \ldots, h_{s-1}$ shows that $F[\setminus (\mathrm{Cl}^{(k)}(F) \cup T)]$ is safe. By Claim A.8, $\langle F[\setminus (\mathrm{Cl}^{(k)}(F) \cup T)] \rangle = \langle (F \cup \{f\})[\setminus (\mathrm{Cl}^{(k)}(F) \cup T)] \rangle$, hence $(F \cup \{f\})[\setminus (\mathrm{Cl}^{(k)}(F) \cup T)]$ is also safe. Thus, $\mathrm{Cl}^{(k)}(F \cup \{f\}) = \mathrm{Cl}^{(k)}(F) \cup T$. As $T \subseteq [m] \setminus \mathrm{Cl}^{(k)}(F)$, the claim follows. □

Consider the space $\{g \in \langle F[\setminus \mathrm{Cl}^{(k)}(F)] \rangle \mid \mathrm{supp}(g) \subseteq T\}$; by the definition of closure its dimension is at most $k|T|$, but as it contains all $h_1, h_2, \ldots, h_{s-1}$, the dimension is exactly $k|T|$.

By Lemma 4.2, $\dim \langle F[\setminus \mathrm{Cl}^{(k)}(F)] \rangle - \dim \langle F[\setminus (\mathrm{Cl}^{(k)}(F) \cup T)] \rangle = \dim \{g \in \langle F[\setminus \mathrm{Cl}^{(k)}(F)] \rangle \mid \mathrm{supp}(g) \subseteq T\} = k|T|$. By Claim A.8, $\langle F[\setminus (\mathrm{Cl}^{(k)}(F) \cup T)] \rangle = \langle (F \cup \{f\})[\setminus (\mathrm{Cl}^{(k)}(F) \cup T)] \rangle = \langle (F \cup \{f\})[\setminus \mathrm{Cl}^{(k)}(F \cup \{f\})] \rangle$. Thus,

$$\dim \langle F[\setminus \mathrm{Cl}^{(k)}(F)] \rangle - \dim \langle (F \cup \{f\})\,[\setminus \mathrm{Cl}^{(k)}(F \cup \{f\})] \rangle = k|T| = k(|\mathrm{Cl}^{(k)}(F \cup \{f\})| - |\mathrm{Cl}^{(k)}(F)|).$$

□

Acknowledgements The authors are grateful to Jonathan Mosheiff and Sergey Komech for fruitful discussions on the existence of k-surjective gadget collections, and to Navid Talebanfard for pointing out earlier uses of random gadgets. The authors also thank Susanna de Rezende for her motivating question, which ultimately led to substantially strengthening the results of this paper.

References

[1] Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. Tight hardness results for LCS and other sequence similarity measures. In Venkatesan Guruswami, editor, *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 59–78. IEEE Computer Society, 2015.

[2] Yaroslav Alekseev and Dmitry Itsykson. Lifting to bounded-depth and regular resolutions over parities via games. In Michal Koucký and Nikhil Bansal, editors, *Proceedings of the 57th Annual ACM Symposium on Theory of Computing, STOC 2025, Prague, Czechia, June 23-27, 2025*, pages 584–595. ACM, 2025.

[3] Albert Atserias and Víctor Dalmau. A combinatorial characterization of resolution width. *Journal of Computer and System Sciences*, 74(3):323–334, 2008. Computational Complexity 2003.

[4] Arturs Backurs and Piotr Indyk. Edit distance cannot be computed in strongly subquadratic time (unless SETH is false). *SIAM J. Comput.*, 47(3):1087–1097, 2018.

[5] Christopher Beck and Russell Impagliazzo. Strong ETH holds for regular resolution. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 487–494. ACM, 2013.

[6] Eli Ben-Sasson and Swastik Kopparty. Affine dispersers from subspace polynomials. *SIAM J. Comput.*, 41(4):880–914, 2012.

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

[8] 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.

[9] Ilario Bonacina and Navid Talebanfard. Improving resolution width lower bounds for k-cnfs with applications to the strong exponential time hypothesis. *Inf. Process. Lett.*, 116(2):120–124, 2016.

[10] Ilario Bonacina and Navid Talebanfard. Strong ETH and resolution via games and the multiplicity of strategies. *Algorithmica*, 79(1):29–41, 2017.

[11] Michele Borassi, Pierluigi Crescenzi, and Michel Habib. Into the square: On the complexity of some quadratic-time solvable problems. In Pierluigi Crescenzi and Michele Loreti, editors, *Proceedings of the 16th Italian Conference on Theoretical Computer Science, ICTCS 2015, Firenze, Italy, September 9-11, 2015*, volume 322 of *Electronic Notes in Theoretical Computer Science*, pages 51–67. Elsevier, 2015.

[12] Karl Bringmann. Why walking the dog takes time: Frechet distance has no strongly subquadratic algorithms unless SETH fails. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 661–670. IEEE Computer Society, 2014.

[13] 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.

[14] Marco L. Carmosino, Jiawei Gao, Russell Impagliazzo, Ivan Mihajlin, Ramamohan Paturi, and Stefan Schneider. Nondeterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility. In Madhu Sudan, editor, *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, Cambridge, MA, USA, January 14-16, 2016*, pages 261–270. ACM, 2016.

[15] 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.

[16] Klim Efremenko, Michal Garlík, and Dmitry Itsykson. Lower bounds for regular resolution over parities. *SIAM J. Comput.*, 54(4):887–915, 2025. Preliminary version appeared in Proceedings of STOC 2024.

[17] Klim Efremenko and Dmitry Itsykson. Amortized closure and its applications in lifting for resolution over parities. In Srikanth Srinivasan, editor, *40th Computational Complexity Conference, CCC 2025, August 5-8, 2025, Toronto, Canada*, volume 339 of *LIPIcs*, pages 8:1–8:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025.

[18] 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.

[19] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-sat. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001.

[20] Dmitry Itsykson and Alexander Knop. Supercritical tradeoff between size and depth for resolution over parities. *Electron. Colloquium Comput. Complex.*, TR25-116, 2025. To appear in Proceedings of ITCS 2026.

[21] 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.

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

[23] Hamidreza Jahanjou, Eric Miles, and Emanuele Viola. Local reduction. *Inf. Comput.*, 261:281–295, 2018.

[24] Mihai Pătraşcu and Ryan Williams. On the possibility of faster SAT algorithms. In Moses Charikar, editor, *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*, pages 1065–1075. SIAM, 2010.

[25] Ramamohan Paturi, Pavel Pudlák, Michael E. Saks, and Francis Zane. An improved exponential-time algorithm for k-sat. *J. ACM*, 52(3):337–364, 2005.

[26] Ramamohan Paturi, Pavel Pudlák, and Francis Zane. Satisfiability coding lemma. *Chic. J. Theor. Comput. Sci.*, 1999, 1999.

[27] Pavel Pudlák and Russell Impagliazzo. A lower bound for DLL algorithms for k-sat (preliminary version). In David B. Shmoys, editor, *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms, January 9-11, 2000, San Francisco, CA, USA*, pages 128–136. ACM/SIAM, 2000.

[28] Pavel Pudlák. Proofs as games. *The American Mathematical Monthly*, 107(6):541–550, 2000.

[29] Dominik Scheder, Bangsheng Tang, Shiteng Chen, and Navid Talebanfard. Exponential lower bounds for the PPSZ k-sat algorithm. In Sanjeev Khanna, editor, *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 1253–1263. SIAM, 2013.

[30] D.J.A. Welsh. Generalized versions of Hall's theorem. *Journal of Combinatorial Theory, Series B*, 10(2):95–101, 1971.

[31] Richard Ryan Williams. Strong ETH breaks with merlin and arthur: Short non-interactive proofs of batch evaluation. In Ran Raz, editor, *31st Conference on Computational Complexity, CCC 2016, May 29 to June 1, 2016, Tokyo, Japan*, volume 50 of *LIPIcs*, pages 2:1–2:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.

[32] Ryan Williams. Improving exhaustive search implies superpolynomial lower bounds. *SIAM J. Comput.*, 42(3):1218–1244, 2013.

[33] Ryan Williams. Nonuniform ACC circuit lower bounds. *J. ACM*, 61(1):2:1–2:32, 2014.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il