

Extremely Deep Proofs

Noah Fleming[†]
University of California,
San Diego
Memorial University

Toniann Pitassi[‡]
Columbia University
University of Toronto
IAS

Robert Robere[†]
McGill University

November 24, 2021

Abstract

We further the study of *supercritical* tradeoffs in proof and circuit complexity, which is a type of tradeoff between complexity parameters where restricting one complexity parameter forces another to exceed its worst-case upper bound. In particular, we prove a new family of supercritical tradeoffs between *depth* and *size* for Resolution, $\text{Res}(k)$, and Cutting Planes proofs. For each of these proof systems we construct, for each $c \leq n^{1-\varepsilon}$, a formula with $n^{O(c)}$ clauses and n variables that has a proof of size $n^{O(c)}$ but in which any proof of size no more than roughly exponential in $n^{1-\varepsilon}/c$ must necessarily have depth $\approx n^c$. By setting $c = o(n^{1-\varepsilon})$ we therefore obtain exponential lower bounds on proof depth; this far exceeds the trivial worst-case upper bound of n . In doing so we give a simplified proof of a supercritical depth/width tradeoff for tree-like Resolution from [31]. Finally, we outline several conjectures that would imply similar supercritical tradeoffs between size and depth in circuit complexity via lifting theorems.

[†]Research supported by NSERC, the IAS School of Mathematics and NSF Grant No. CCF-1900460.

[‡]Research supported by NSERC.

1 Introduction

Beame, Beck, and Impagliazzo recently initiated the study of *supercritical* tradeoffs in proof and circuit complexity [4]. In its simplest form, a supercritical tradeoff is a tradeoff between two complexity parameters in which a restriction on one parameter forces an increase in the second parameter that goes *above* the trivial worst-case upper bound on the second parameter. Beame et. al. established a supercritical tradeoff between *size* and *space* in the *Resolution* proof system, showing there is an unsatisfiable CNF formula on n variables which admits quasi-polynomial size refutations, but in which any quasi-polynomial size refutation must also use quasi-polynomial clause-space. Since every formula on n variables has a Resolution proof of space $n + O(1)$ and *exponential* size [16], their result shows that limiting the size of refutations causes the space to go significantly beyond the trivial worst-case upper bound on the space complexity. Recently, Razborov showed that in certain cases restrictions in parameters can lead to extreme results [31]. He exhibited a strong supercritical size/width for tree-like Resolution [31], showing that there are unsatisfiable formulas on n variables such that if the width is restricted, then the tree-like Resolution size must be *doubly exponential* in n .

Despite the intuitive appeal of supercritical tradeoffs, only a few have been observed and, with the exception of length versus space, most of these tradeoffs have been for somewhat artificial combinations of parameters. Furthermore, to our knowledge, the existence of supercritical tradeoffs in circuit complexity has yet to be explored.

In this work we study supercritical tradeoffs between *size* and *depth* in both proof complexity and circuit complexity. For both proof systems and circuit models, depth captures the degree to which they — and therefore the families of algorithms which they capture — can be *parallelized*. Furthermore, for proof systems such as Cutting Planes, depth is closely related to measures of polytope rank which have been extensively studied in combinatorial optimization [12, 34]. For many standard proof systems (and, indeed, for all of the proof systems studied in this paper) the worst-case upper bound on depth is always at most n — this is because all of these proof systems can efficiently simulate tree-like Resolution, which trivially has depth- n proofs for every formula. A supercritical tradeoff in this setting for a particular proof system P would therefore exhibit a family of CNF formulas over n variables that have P refutations of size $s(n)$, however there is a size bound $s' \gg s$ such that any P -refutation of size $s'' \ll s'$ requires depth that is superlinear in n . Similarly, since any n -bit boolean function can be computed by a fanin-2 circuit of depth at most n , a supercritical tradeoff in this setting would prove the existence of a function f_n with $s(n)$ -size circuits, but for which there is a size bound s' such that any circuit of size $s'' \ll s'$ requires superlinear depth.

1.1 Supercritical Size/Depth Tradeoffs in Proof Complexity.

For some sufficiently strong propositional proof systems (such as Frege systems), it is known that proofs can be *balanced*: given a Frege proof Π of size s , there exists another Frege proof of the same tautology of size $O(s)$ and depth $O(\log s)$. However, for weaker proof systems such as Resolution (Res), k -DNF Resolution (Res(k)), and Cutting Planes (CP), balancing is not always possible. For example, there are known families of CNF formulas $\{F_n\}$ with polynomial-size Res proofs, but such that *any* proof requires depth $\Omega(n/\log n)$ [1, 11]. However, none of these results have breached the supercritical threshold, and it was not previously known whether an upper bound of $O(n)$ could be assumed on the depth of short proofs.

Our main result establishes the following supercritical depth/size tradeoffs for Res, Res(k), and CP refutations, which resolves conjectures made in [18].

Theorem 1.1. For any constant $\varepsilon > 0$, positive integers k, n sufficiently large, $\mathcal{P} \in \{\text{Res}, \text{Res}(k), \text{CP}\}$, and any arbitrary real parameter $1 \leq c < n^{\frac{1-\varepsilon}{2+\varepsilon}}$, there is a CNF formula F on n variables and $n^{O(c)}$ clauses such that

- There is a \mathcal{P} -refutation of F of size $n^{O(c)}$.
- If Π is a \mathcal{P} -refutation of F with $\text{size}(\Pi) = 2^{o\left(n^{\frac{1-\varepsilon}{2+\varepsilon}}/c\right)}$ then

$$\text{depth}(\Pi) \log^2 \text{size}(\Pi) = \Omega\left(\frac{n^{c/(2+\varepsilon)}}{c \log n}\right).$$




Figure 1: A depiction of our size/depth tradeoffs as we range the compression parameter c . The *depth* value represents the depth required to refute the formula F , while the *gap* value represents by how much the size of the minimal \mathcal{P} -proof of F needs to be increased before the supercritical depth lower bound no longer holds.

Varying the “compression parameter” c between $O(1)$ and n^δ (for some small constant δ) allows us to obtain an interesting family of tradeoff results; this is depicted in [Figure 1](#). In one extreme, when $c = O(1)$ we obtain a formula F which has refutations of size $\text{poly}(n)$, however any proof of size $\ll 2^{n^{1-\varepsilon}}$ must have depth that is *polynomial* (and superlinear!) in n . In the other extreme, setting $c = n^\delta$ implies formulas with exponential-size (in n ; the refutations are polynomial in the size of the formula) and exponential-depth refutations, but any refutation *must* have exponential depth.

Our proof builds on Razborov’s *hardness condensation* method [31]; we briefly describe Razborov’s method next, and a detailed description is given in [Section 3](#). Let F be a CNF formula over N variables satisfying the following properties:

1. There are Resolution refutations of F with small size.
2. Any Resolution refutation of F requires depth $d(N)$.

The strategy is to find a way to *compress* the variables y_1, \dots, y_N of F to a much smaller set of variables x_1, \dots, x_n with $n \ll N$ in a way which preserves the depth lower bound. Indeed, we will argue that any small refutation of the compressed formula must have depth approximately $d(N)$. We stress that the

lower bound of $d(N)$ is in terms of the *original* number of variables N , and therefore if n is sufficiently small compared to N , we obtain a supercritical depth lower bound for small proofs. Razborov proved that such a compression argument was possible for *tree-like* Resolution, assuming that the *width* of the refutation is bounded [31]. As a secondary contribution, we give a significantly simplified proof of this tradeoff in Section 5 using the “top-down” language of Prover-Adversary games.

Once we have established our size/depth tradeoff for Resolution using the aforementioned strategy, the lower bounds for CP and $\text{Res}(k)$ are straightforward applications of lifting theorems. For CP, we utilize the depth-preserving lifting theorem of Garg et al. [20]. For $\text{Res}(k)$, we show how the lower bound of Segerlind et al. [35] immediately gives rise to a lifting theorem for $\text{Res}(k)$ which preserves depth.

Supercritical Tradeoffs for Tseitin and Matching. The Tseitin formulas have a storied history in proof complexity. They were one of the first examples shown to be hard to refute in Resolution, and were long conjectured to also be exponentially hard for Cutting Planes. Recently, [13] disproved this conjecture, obtaining quasipolynomial-size Cutting Planes refutations for the Tseitin formulas by translating an upper bound in the Stabbing Planes proof system [5] into Cutting Planes. Their upper bound was generalized in [18], which showed that *any* unsatisfiable system of equations over a prime finite field also have quasipolynomial-size CP refutations. Quite interestingly, in both cases the CP refutations not only have quasipolynomial size but also have quasipolynomial depth — that is, the depth of these proofs is supercritical. In [18], it was conjectured that there are unsatisfiable formulas such that any small CP refutation requires superlinear depth, and moreover, they conjectured that the Tseitin formulas are an example of such formulas. We restate this conjecture for Tseitin formulas next.

Conjecture 1.2. *There is an unsatisfiable system of mod2 linear equations such that any proof of quasipolynomial size requires depth that is superlinear in n (or even quasipolynomial).*

One potential interest in resolving this conjecture is the similarity between the Tseitin formulas and related counting principles and the *perfect matching function*. The seminal paper of Razborov [30] proved that monotone circuits for matching require quasipolynomial size, and it is also known that monotone circuits require linear depth [29]. Furthermore, it is known [19, 22, 27] that small Cutting Planes proofs imply small monotone circuits. This brings up the intriguing possibility that the size/depth of Cutting Planes proofs the Tseitin formulas could inform the size/depth monotone circuits for perfect matching.

1.2 Supercritical Tradeoffs for Boolean Circuits.

It is well-known that every monotone function has monotone circuits of depth at most n . We ask if every monotone function with polynomial monotone circuit complexity can also be computed by a monotone circuit of both polynomial size and linear depth. For monotone circuits, it is known that it is not possible to reduce the depth below $n/\text{polylog}n$ while maintaining polynomial size [14]. However, these tradeoffs are proven using communication complexity, and are therefore limited to size/depth tradeoffs where the depth is at most n .

Conjecture 1.3. *There exists a monotone function which admits polynomial size monotone circuits, but any circuit of depth at most $n^{1+\epsilon}$ for some $\epsilon > 0$ requires superpolynomial-size monotone circuits.*

One promising avenue for resolving this conjecture is to *lift* our supercritical tradeoffs for Resolution to supercritical tradeoffs for monotone circuits. Indeed, lifting has already been used to obtain monotone circuit lower bounds from Resolution lower bounds [20].

In more detail, using lifting techniques it is known that given an unsatisfiable formula F , Resolution size lower bounds for F imply monotone circuit lower bounds for a related monotone function, mCSP-SAT_F . Stated contrapositively, given a small circuit C for mCSP-SAT_F , one can extract a small Resolution refutation for F of roughly the same size. Moreover, the structure of C is nearly the same as the structure of the Resolution refutation — thus if C has small size and depth, then F also has a Resolution refutation of small size and depth. It follows that a strong enough supercritical size/depth tradeoff for Resolution would imply a supercritical size/depth tradeoff for monotone circuits as well!

Unfortunately, our current supercritical tradeoff for Resolution isn't strong enough to obtain interesting size/depth tradeoffs for monotone circuits for the following reason. The reduction from monotone circuits for mCSP-SAT_F to Resolution proofs of F incurs a seemingly necessary blowup in the parameters — the number of clauses m in F corresponds to the number of variables in mCSP-SAT_F , and thus in order to prove a supercritical tradeoff for monotone circuit complexity via this route, one needs to prove the following strong supercritical tradeoff:

Conjecture 1.4. *There are families of unsatisfiable formulas $\{F_n\}$ with $m = m(n)$ clauses such that F_n has Resolution refutations of size polynomial in m , but any polynomial-size (or subexponential-size) refutation of F_n requires depth that is superlinear in m .*

Under this conjecture, we obtain supercritical tradeoffs for monotone circuits, which we prove in [Section 7](#). We note that it is entirely possible that the above conjecture is false — that is, it seems possible that Resolution refutations could be balanced to depth $O(m)$ without incurring a significant increase in size. Indeed, a potentially relevant construction is the classic work of Ben-Sasson and Wigderson [7] which converts a size s Resolution refutation into a refutation of low width ($\sqrt{n} \log s$). The next problem asks whether their construction can be improved to keep low width and low depth.

Problem 1.5. Prove or disprove that for any k -CNF formula on m clauses and n variables, any size s Resolution refutation can be converted into a refutation of depth $O(m)$ and width $k + O(\sqrt{n \log s})$.

A positive resolution to the above problem would give a counterexample to our conjecture and would also give a very surprising upper bound on the depths of arbitrary Resolution proofs, while a negative answer would imply (conditional) supercritical tradeoffs for monotone circuits (see [Section 7](#)). Given this “win-win” situation, we believe that studying this problem is of fundamental interest for future work.

Thirdly, we ask about supercritical size/depth tradeoffs for *non-monotone* circuits.

Conjecture 1.6. *There is a function f with polynomial-size circuits but for which any polynomial-size circuit requires superlinear depth.*

Given that an unconditional supercritical tradeoff for general circuits is beyond the reach of current techniques, it is even interesting to prove such a tradeoff under a standard complexity assumption, such as those concerning the existence of one-way functions or other cryptographic primitives.

1.3 Related Work

Besides the work of Beame, Beck and Impagliazzo [4] and Razborov [31], several supercritical tradeoffs have been observed in proof complexity. Razborov [32] considered a notion of *width* for Cutting Planes, which measures the number of distinct variables with non-zero coefficients which appear in a Cutting Planes line. Using similar machinery as in [31], Razborov established a supercritical width/depth tradeoff for CP.

Much of the work on supercritical tradeoffs has focused on notions of proof *space*. Extending Beame et al. [4], Beck et al. [6] exhibited supercritical size/clause-space tradeoffs for the Polynomial Calculus. Berkholz and Nordström [9] exhibited a strong supercritical tradeoff between width and clause-space for Resolution, and Razborov [33] established a supercritical tradeoff between variable-space and size for Resolution. Finally, Papamakarios and Razborov [25] showed that separating clause- or monomial-space from proof size for tree-like Resolution is the equivalent to showing a supercritical clause-space/depth tradeoff.

Proof Depth and Algorithm Analysis. One of the primary motivations for studying tradeoffs in proof complexity is the existence of deep connections to the analysis of practical algorithms. Proof systems in a specific, formal sense correspond to families of efficient, provably correct algorithms. For example, Resolution captures the reasoning used in state-of-the-art (CDCL) algorithms for SAT, while Cutting Planes (CP) was introduced in order to model a broad family of optimization algorithms used in integer programming based on Chvátal-Gomory cuts. Therefore resource tradeoffs for these proof systems apply to the corresponding family of algorithms. The depth in proofs captures the degree to which the associated families of algorithms can be paralellized, and so supercritical size/depth tradeoffs imply that the associated families of algorithms are inherently sequential in the worst case.

As well, the depth of Cutting Planes proofs has been thoroughly studied in integer programming theory under the name of *Chvátal rank* [12]. A number of results in this area obtained nearly tight worst-case bounds on the Chvátal rank of a number of polytopes [10, 15, 34]. Our main result implies Chvátal rank lower bounds that are exponentially stronger than previously known, under the condition that the Chátal-Gomory Cutting Planes procedures is *efficient*.

2 Preliminaries

We recall some preliminaries from proof complexity. A *clause* C is a disjunction of boolean literals, and the *width* of a clause is the number of literals in C . If $\mathcal{C} = C_1 \wedge C_2 \wedge \dots \wedge C_m$ is a CNF formula then the *width* of \mathcal{C} , denoted $\text{width}(\mathcal{C})$, is the maximum width of any clause in \mathcal{C} .

Resolution. Fix an unsatisfiable CNF formula F over variables x_1, \dots, x_n . A *Resolution* (Res) refutation of F is a sequence of clauses $\{C_i\}_{i \in [s]}$ ending with the empty clause $C_s = \perp$ such that each C_i is either in F or is derived from earlier clauses C_j, C_k with $j, k < i$ by one of the following inference rules:

- *Resolution.* $C_i = (C_j \setminus \{\ell_k\}) \cup (C_k \setminus \{\bar{\ell}_k\})$ where $\ell_k \in C_j, \bar{\ell}_k \in C_k$ is a literal.
- *Weakening.* $C_i \supseteq C_j$.

The *size* of a Res proof Π of F is s , the number of clauses in Π , and we denote by $\text{size}_{\text{Res}}(F)$ the minimum size of any Res proof of F . The size of resolution proofs is intimately related to their width, $\text{width}(\Pi)$, which is the maximum number of literals occurring in any clause in the proof.

Any Res proof can be represented as a directed acyclic graph (DAG) in which the nodes represent clauses in the proof, each clause of F has in-degree 0 and, and any other clause has at most two incoming arcs from the at most two clauses that produced it. With this in mind, we can define the *depth* of a Res proof to be the length of the longest root-to-leaf path in the proof, and denote by $\text{depth}_{\text{Res}}(F)$ the minimum depth of any Res proof of F . It is not difficult to see that for any CNF formula F , $\text{size}_{\text{Res}}(F) \leq 2^n$ and $\text{depth}_{\text{Res}}(F) \leq n$.

One can obtain generalizations of Res by allowing the proof system to have more general types of boolean formulas as lines. In this paper, we will also be interested in the proof system $\text{Res}(k)$ which operates with k -DNF formulas. The details of particular rules of this proof system will not be necessary for us, however we introduce them for completeness.

DNF Resolution. Let F be an unsatisfiable CNF formula over variables x_1, \dots, x_n and let k be some fixed constant. A k -DNF Resolution ($\text{Res}(k)$) refutation of F is a sequence of k -DNF formulas $\{D_i\}_{i \in [s]}$ such that $D_s = \Lambda$ and each D_i is either a clause of F or is deduced from earlier k -DNFs by one of the following inference rules, where a literal ℓ_i is either x_i or $\neg x_i$:

- *Cut.* From k -DNFs $A \vee (\wedge_{i \in I} \ell_i)$ and $B \vee (\vee_{i \in I} \neg \ell_i)$ deduce $A \vee B$.
- *Weakening.* From a $(k-1)$ -DNF A deduce $A \vee \ell$ for any literal ℓ .
- \wedge -Introduction. From $\{A \vee \ell_i\}_{i \in I}$ deduce $A \vee (\wedge_{i \in I} \ell_i)$.
- \wedge -Elimination. From $A \vee (\wedge_{i \in I} \ell_i)$ deduce $A \vee \ell_i$ for any $i \in I$.

The *size* of a $\text{Res}(k)$ refutation is s , the number of DNFs in the resolution, and the *depth* of the refutation is the length of the longest path in the proof DAG. Since $\text{Res} = \text{Res}(1)$, there is a $\text{Res}(k)$ refutation of size 2^n and depth n of any CNF formula.

The final proof system which we will be interested in is of the algebraic variety, refuting the existence of integer solutions to systems of linear inequalities. We can encode a CNF formula F over variables x_1, \dots, x_n as an equisatisfiable system of linear inequalities in the following way. First, introduce inequalities $0 \leq x_i \leq 1$ for all $i \in [n]$. If $C_i = \bigvee_{i \in I} x_i \vee \bigvee_{j \in J} \neg x_j$ is a clause in F , add the inequality

$$\sum_{i \in I} x_i + \sum_{j \in J} (1 - x_j) \geq 1.$$

It is not difficult to see that the $\{0, 1\}$ -solutions to F are exactly the solutions to this system of inequalities.

Cutting Planes and Semantic Cutting Planes. Let P be a system of integer-linear inequalities which is infeasible over \mathbb{Z}^n . A *Cutting Planes* (CP) refutation of P is a sequence of integer-linear inequalities $\{a_i x \geq b_i\}_{i \in [s]}$ ending with the trivially false inequality $0 \geq 1$ such that each inequality is either belongs to P , or is deduced from earlier inequalities by one of the following inference rules:

- *Linear Combination.* From $ax \geq b$, $cx \geq d$ deduce any non-negative linear combination with integer coefficients.
- *Division.* From $ax \geq b$, if $d \in \mathbb{Z}$ with $d \geq 0$ divides a , deduce $(a/d)x \geq \lceil b/d \rceil$.

The rules of Cutting Planes preserve integer solutions, and therefore a refutation exists if and only if P is infeasible over \mathbb{Z}^n .

The following transformation will allow us to talk about Cutting Planes refutations of CNF formulas. We can encode a CNF formula F over variables x_1, \dots, x_n as an equisatisfiable system of integer-linear inequalities in the following way. First, introduce inequalities $0 \leq x_i \leq 1$ for all $i \in [n]$. If $C_i = \bigvee_{i \in I} x_i \vee \bigvee_{j \in J} \neg x_j$ is a clause in F , add the inequality

$$\sum_{i \in I} x_i + \sum_{j \in J} (1 - x_j) \geq 1.$$

Observe that the satisfying assignments to F are exactly the integer solutions to the resulting system of linear inequalities. With this transformation we will say that a Cutting Planes refutation is a refutation of F if it is a refutation of the associated system of linear inequalities.

The *semantic Cutting Planes* (sCP) proof system is a strengthening of CP to allow *any* deduction which is sound over boolean points. Like CP, an sCP refutation of a CNF formula F is a sequence of integer-linear inequalities $\{a_i x \geq b_i\}_{i \in [s]}$, however we are now permitted to use the following extremely strong semantic deduction rule:

- *Semantic Deduction.* From previously derived inequalities $ax \geq b$, $a'x \geq b'$, deduce any inequality $cx \geq d$ such that every $\{0, 1\}$ assignment satisfying both $ax \geq b$ and $a'x \geq b'$ also satisfies $cx \geq d$.

The size of a CP or sCP refutation is the number of inequalities s appearing in it. As well, analogous to Resolution, the depth of a refutation is the longest root-to-leaf path in its representation as a DAG. It is not difficult to see that any Res proof can be transformed into CP (and therefore sCP) while preserving both the size and depth up to a constant factor, and therefore there is always a CP refutation of any CNF formula F which has size 2^n and linear depth.

Filmus et al. [17] showed that sCP can be significantly more expressive than CP: there are formulas with polynomial size sCP proofs requiring exponential size proofs in CP. In fact, inferences made with the semantic deduction rule are not even verifiable in polynomial time unless $P = NP$.

3 Proof Overview

Our proof will follow the general approach of Razborov [31], which established a supercritical width/size tradeoff for tree-like Resolution. The key machinery in his proof is a technique — which has been termed *hardness condensation* — that compresses the number of variables of the formula F in such a way that the depth of any *bounded width* tree-like resolution refutation of the compressed formula remains proportional to the tree-like resolution depth of refuting F .

The compression is done by composing the formula F with an XOR gadget. However, unlike standard lifting theorems, the XOR gadgets will be defined on overlapping sets of variables. This will allow us to reduce the total number of variables of the composed function.

XOR Substitution. Let $G = ([N] \cup [n], E)$ be a bipartite graph and $\{y_1, \dots, y_N\}$, $\{x_1, \dots, x_n\}$ be sets of propositional variables. For a clause C in the variables $\{y_1, \dots, y_N\}$ we will denote by $C \circ \text{XOR}_G$ the CNF obtained from C by the \mathbb{F}_2 -linear substitution

$$y_i \mapsto \bigoplus_{j:(i,j) \in E} x_j,$$

and then rewriting the formula in CNF (see Figure 2). For a CNF F , let $F \circ \text{XOR}_G$ be the CNF formula that results from this substitution. If the clauses of F have width at most k and G has left-degree at most ℓ , then $F \circ \text{XOR}_G$ is a CNF formula on n variables and at most $m \cdot 2^{k\ell-1}$ clauses of width at most ℓk .

Similarly, we write $F \circ \text{XOR}_2^N$ to represent the CNF formula obtained by the substitution $y_i \mapsto u_i \oplus v_i$ for each i , where u_i and v_i are new variables.




Figure 2: The bipartite graph G , together with the XOR constraints that each left-vertex (y -variable) is replaced with.

Now, if G disperses the variables of $[n]$ among the gadgets in such a way that learning the value of y_i does not reveal too much information about the values of any y_j , then the best-possible Res proof should essentially be one which simply simulates the proof for F . Indeed, Razborov proved that this is the case under the assumption that G has sufficiently strong *expansion* properties. For $S \subseteq [n]$, let

$$\Gamma(S) := \{y_i \in [N] : \exists x_j \in S, (y_i, x_j) \in E\}$$

be the *neighbourhood* of S in G . The following definition records the expansion properties of bipartite graphs that we will use.

Boundary Expansion. For a bipartite graph $G = (U \cup V, E)$ the *boundary* of a set $W \subseteq U$ is

$$\delta(W) := \{v \in V : |\Gamma(v) \cap W| = 1\}.$$

The *boundary expansion* of a set $W \subseteq U$ is $|\delta(W)|/|W|$. The graph G is a (r, s) -*boundary expander* if the boundary expansion of every set $W \subseteq U$ with $|W| \leq r$ has boundary expansion at least s .

That is, G is a boundary expander if for any small enough subset of left-vertices, the number of *unique neighbours* is large.

We are now ready to state Razborov's hardness condensation theorem. [31]. We note that Razborov originally proved his theorem for the case of tree-like Resolution, but (as we will see) it also holds for general Resolution.

Depth Condensation Theorem. *Let F be an unsatisfiable CNF formula on N variables and let $G = ([N] \cup [n], E)$ be an $(r, 2)$ -boundary expander. If Π is a Resolution refutation of $F \circ \text{XOR}_G$ with $\text{width}(\Pi) \leq r/4$, then*

$$\text{depth}(\Pi)\text{width}(\Pi) \geq \frac{\text{depth}_{\text{Res}}(F)}{2}.$$

In [Section 5](#), we give a simplified proof of [Depth Condensation Theorem](#). By combining this theorem with known reductions from Resolution size to Resolution width [[20, 35](#)], we are able to establish the following depth-to-size lifting theorems. When instantiated, these will allow us to prove our tradeoffs.

Theorem 3.1. *Let F be any CNF formula on N variables and let $G = ([N] \cup [n], E)$ be an $(r, 2)$ -boundary expander. If Π is a resolution refutation of $F \circ \text{XOR}_G \circ \text{XOR}_2^n$ such that $\log(\text{size}(\Pi) + 1) \leq r/12$, then*

$$\text{depth}(\Pi) \log(\text{size}(\Pi) + 1) \geq \frac{\text{depth}_{\text{Res}}(F)}{6}.$$

Theorem 3.2. *Let $k \geq 1$ be any constant, let F be any CNF formula on N variables, and let $G = ([N] \cup [n], E)$ be an $(r, 2)$ -boundary expander. There is a constant $\delta := \delta(k) > 0$ such that if Π is a $\text{Res}(k)$ refutation of $F \circ \text{XOR}_G \circ \text{XOR}_2^n$ with $\log(\text{size}(\Pi) + 1) \leq \delta \cdot r$, then*

$$\text{depth}(\Pi) \log^2(\text{size}(\Pi)) = \Omega(\text{depth}_{\text{Res}}(F))$$

To obtain the depth-to-size lifting theorem for semantic Cutting Planes, we will instead use the t -bit index function as our outer gadget. Recall that $\text{IND}_t : [t] \times \{0, 1\}^t \rightarrow \{0, 1\}$ maps (x, y) to y_x . For a CNF formula on variables z_1, \dots, z_n , let $F \circ \text{IND}_t^n$ be the CNF formula obtained from F by the substitution $z_i \mapsto \text{IND}_t(x_i, y_i)$ on new sets of variables x_i, y_i . Note that if F is an unsatisfiable k -CNF formula with m clauses, then $F \circ \text{IND}_t^n$ is an unsatisfiable CNF formula on $O(nt)$ variables and $O(tm^k + n)$ clauses.

Theorem 3.3. *Let $\varepsilon > 0$ be any constant, let F be any CNF formula on N variables, and let $G = ([N] \cup [n], E)$ be an $(r, 2)$ -boundary expander. There is a constant $\delta > 0$ such that if Π is a semantic CP refutation of $F \circ \text{XOR}_G \circ \text{IND}_{n^{1+\varepsilon}}^n$ with $\log(\text{size}(\Pi)) \leq \delta \cdot r \log n$, then*

$$\text{depth}(\Pi) \log^2(\text{size}(\Pi)) = \Omega(\text{depth}_{\text{Res}}(F) \log^2 n).$$

We delay the proofs of these theorems until [Section 6](#). Instead, we will first instantiate them to obtain our tradeoffs.

4 Parameterizing the Tradeoffs

From the previous theorems we can obtain a family of supercritical tradeoffs. To do so, we will need a formula which has small resolution refutations, but requires large depth. The canonical example of such formulas are the *pebbling formulas* of [[1](#)] on some hard-to-pebble graph H . It is known that for N -vertex graph H , Peb_H has resolution refutations of size $O(N)$ and width $O(1)$. However, resolution depth required to refute these formulas is equal to the reversible pebbling number of the graph [[11](#)]. Furthermore, there exist $O(1)$ -degree graphs with reversible pebbling number $\Omega(N/\log N)$ [[26](#)].

By combining the lower and upper bounds for the pebbling formulas with the previous lemmas we can obtain a family of supercritical tradeoffs by varying the underlying expander graph $G = ([N] \cup [n], E)$. The following lemma provides us with a sufficient family of expander graphs.

Lemma 4.1 (Razborov [[31](#)]). *Let n be any sufficiently large positive integer, and let N, r, ℓ be positive integers depending on n such that $\ell \geq 4$. If*

$$rN^{4/\ell} = o(n/\ell)$$

then an $(r, 2)$ boundary expander $G = ([N] \cup [n], E)$ exists with left-degree ℓ .

Our tradeoffs will be in terms of the number of variables n that we are “compressing” $[N]$ into. It will be convenient to set $N = n^c$, for some real parameter $c \geq 1$, which we will call our *compression parameter*. As well, let $\varepsilon > 0$ be some arbitrarily small real parameter. We will set

$$\begin{aligned} N &:= n^c, \\ r &:= n^{1-\varepsilon}/c, \\ \ell &:= 8c/\varepsilon. \end{aligned}$$

We can verify that

$$rN^{4/\ell} = \frac{1}{c}n^{1-\varepsilon}n^{4c/\ell} \leq \frac{1}{c}n^{1-\varepsilon}n^{\varepsilon/2} = \frac{n^{1-\varepsilon/2}}{c} = o(n/\ell).$$

Now, choosing different ranges of c allow us to obtain the following interesting tradeoff results.

For convenience, we record the following proposition stating how the parameters of F transform under composition.

Proposition 4.2. *Let H be any graph on N vertices with indegree $O(1)$. Let $G = ([N] \cup [n], E)$ be a bipartite graph with left-degree at most ℓ . Then,*

- Peb_H has N variables, $N + 1$ clauses, and width $O(1)$.
- $\text{Peb}_H \circ \text{XOR}_G \circ \text{XOR}_2^n$ has $2n$ variables, $n2^{O(\ell)}$ clauses, and width $O(\ell)$.
- $\text{Peb}_H \circ \text{XOR}_G \circ \text{IND}_{n^{1+\varepsilon}}^n$ has $O(n^{2+\varepsilon})$ variables, $n^{O(\ell)}$ clauses, and width $O(\ell)$.

Proof. We obtain $\text{Peb}_H \circ \text{XOR}_G$ by replacing each variable y_i with an XOR of at most ℓ variables. After expanding, this yields a CNF formula with $n2^{(\ell-1)\text{width}(\text{Peb}_H)} + n = n2^{O(\ell)}$ clauses and width $O(\ell)$. Composing this with XOR_2 has the same effect. To handle composition with the index gadget, we use the encoding of [20] which, for any k -CNF formula F on n variables and m clauses, encodes $F \circ \text{IND}_{n^{1+\varepsilon}}^n$ as a $2k$ -CNF formula on $O(n^{2+\varepsilon})$ variables and $O(m \cdot n^{k(1+\varepsilon)})$ clauses. For our choice of parameters, including $\varepsilon > 0$ an arbitrarily small constant, this will be $n^{O(\ell)}$. \square

Now we can test different parameter regimes. In each of our regimes our tradeoffs are basically as follows: we have a trivial proof of size 2^n and depth n . However, if we demand that the proof has size $\ll 2^{n^{1-\varepsilon}}$, then the depth of the proof will explode to roughly n^c (which is *supercritical* in that it lies above the worst-case upper bound of n). Increasing c obviously increases the final depth lower bound, but since we must choose $\ell = O(c)$ it also increases the number of clauses proportionally. We first state a general tradeoff parameterized by c (a formal version of [Theorem 1.1](#)), and then instantiate c .

Theorem 4.3. *For all constants $\varepsilon > 0$, positive integers k , n sufficiently large, $\mathcal{P} \in \{\text{Res}, \text{Res}(k), \text{CP}\}$, and arbitrary real parameter $c \geq 1$, there is an unsatisfiable CNF formula $F_{\mathcal{P}}$ on n variables such that*

- If $\mathcal{P} = \text{Res}$, then $F_{\mathcal{P}}$ has a resolution refutation of size $n^c \cdot 2^{O(c)}$. However, any resolution refutation Π of $F_{\mathcal{P}}$ with $\text{size}(\Pi) = 2^{o(n^{1-\varepsilon}/c)}$ satisfies

$$\text{depth}(\Pi) \log \text{size}(\Pi) = \Omega\left(\frac{n^c}{c \log n}\right).$$

- If $\mathcal{P} = \text{Res}(k)$, then $F_{\mathcal{P}}$ has a $\text{Res}(k)$ refutation of size $n^c \cdot 2^{O(c)}$. However, any $\text{Res}(k)$ refutation Π of $F_{\mathcal{P}}$ with $\text{size}(\Pi) = 2^{o(n^{1-\varepsilon}/c)}$ satisfies

$$\text{depth}(\Pi) \log^2(\text{size}(\Pi)) = \Omega\left(\frac{n^c}{c \log n}\right)$$

- If $\mathcal{P} = \text{sCP}$, then $F_{\mathcal{P}}$ has a sCP refutation of size $n^{O(c)}$. However, any sCP refutation Π of $F_{\mathcal{P}}$ with $\text{size}(\Pi) = \exp(o(n^{(\frac{1-\varepsilon}{2+\varepsilon})}/c))$ satisfies

$$\text{depth}(\Pi) \log^2(\text{size}(\Pi)) = \Omega\left(\frac{n^{c/(2+\varepsilon)} \log n}{c}\right).$$

Proof. Let the parameters N, r, ℓ be set as above, and let $G = ([N] \cup [n], E)$ be an $(r, 2)$ -boundary expander whose existence is guaranteed by [Lemma 4.1](#). Let H be any directed graph on N vertices with indegree $O(1)$ whose reversible pebbling number is $\Omega(N/\log N)$. Then Peb_H requires resolution refutations of depth $\Omega(N/\log N)$, but has resolution refutations of size $O(N)$. For Res and $\text{Res}(k)$, let F be $\text{Peb}_H \circ \text{XOR}_G \circ \text{XOR}_2^n$. The lower bounds on F follow from [Theorem 3.1](#) and [Theorem 3.2](#). For the upper bound, we simulate the upper bound for Peb_H : every time the resolution proof for Peb_H would query a variable y_i , we query all of the variables that y_i was replaced with in F ; that is, we query all u_j, v_j such that (i, j) is an edge of G , in order to evaluate

$$\bigoplus_{j:(i,j) \in E} (u_j \oplus v_j),$$

which gives a value for y_j . Because the left-degree of G is at most ℓ , we query at most $2\ell = O(c)$ variables. This can be done in a subproof (a decision tree) of size $2^{O(c)}$. Altogether, this is a refutation of size $N \cdot 2^{O(c)} = n^c \cdot 2^{O(c)}$ of F .

For sCP , let F be $\text{Peb}_H \circ \text{XOR}_G \circ \text{IND}_{n^{1+\varepsilon}}^n$. By [Proposition 4.2](#), F has $O(n^{2+\varepsilon})$ variables. The lower bound follows from [Theorem 3.3](#) together with the lower bound on Peb_H . For the upper bound, note that $\text{IND}_{n^{1+\varepsilon}}$ can be evaluated in resolution by querying $(1+\varepsilon) \log n + 1$ variables (the $(1+\varepsilon) \log n$ ‘‘pointer variables’’ x together with the single bit y_x). Thus, by following the same strategy as before, we can simulate the resolution refutation of Peb_H by every time the Res refutation queries a variable y_i , evaluating the index gadgets of all $j \in [n]$ such that $(i, j) \in E$. As G has left-degree at most ℓ , evaluating y_i can be done by querying at most $\ell \cdot ((1+\varepsilon) \log n + 1)$ variables. This results in a Res (and therefore sCP) refutation of size $N \cdot 2^{\ell \cdot ((1+\varepsilon) \log n + 1)} = n^{O(c)}$. \square

In what follows we will explore different ranges of the compression parameter c . The next corollaries are somewhat lossy, as they are stated in order to hold simultaneously for resolution, $\text{Res}(k)$, and sCP . The first interesting regime is when $c = O(1)$. In this case, F is a polynomial size and constant width formula which has a trivial depth n and size 2^n proof, however any refutation of size $\ll 2^{n^{1-\varepsilon}}$, for some $\varepsilon > 0$, must have *polynomial* depth.

Corollary 4.4. *Let $c = O(1)$ be any constant, let $\varepsilon > 0$ be an arbitrarily small constant, and let $\Delta_{\text{sCP}} = 1 + \varepsilon$ and $\Delta_{\text{Res}} = \Delta_{\text{Res}(k)} = 0$. For any $\mathcal{P} \in \{\text{Res}, \text{Res}(k), \text{CP}\}$ there is a CNF formula $F_{\mathcal{P}}$ on n variables, such that*

- *There is a \mathcal{P} -refutation of $F_{\mathcal{P}}$ of size $\text{poly}(n)$.*

- Any \mathcal{P} -refutation Π of $F_{\mathcal{P}}$ with $\text{size}(\Pi) = \exp(o(n^{\frac{1-\varepsilon}{1+\Delta_{\mathcal{P}}}}))$ has $\text{depth}(\Pi) = \tilde{\Omega}(n^{c/(1+\Delta_{\mathcal{P}})})$.

The second interesting regime is when $c = \log^{O(1)} n$. In this case, we are compressing the number of variables quasipolynomially, and we obtain *quasipolynomial* depth lower bounds for small proofs.

Corollary 4.5. *Let $c = \log^{O(1)} n$ and let $\varepsilon > 0$ be an arbitrarily small constant, and let $\Delta_{\text{sCP}} = 1 + \varepsilon$ and $\Delta_{\text{Res}} = \Delta_{\text{Res}(k)} = 0$. For any $\mathcal{P} \in \{\text{Res}, \text{Res}(k), \text{CP}\}$ there is a CNF formula $F_{\mathcal{P}}$ on n variables, such that*

- There is a \mathcal{P} -refutation of $F_{\mathcal{P}}$ of size $2^{\log^{O(1)} n}$.
- Any \mathcal{P} -refutation Π of $F_{\mathcal{P}}$ has with $\text{size}(\Pi) = \exp(\tilde{o}(n^{\frac{1-\varepsilon}{1+\Delta_{\mathcal{P}}}}))$ has $\text{depth}(\Pi) = \tilde{\Omega}(n^{\log^{O(1)} n})$.

Finally, we would like to test how large we can set c . The best possible compression afforded by [Theorem 4.3](#) is $c = n^{\delta}$ for some small constant $\delta > 0$. Surprisingly, this implies an *exponential* blowup in the depth.

Corollary 4.6. *Let $\varepsilon > 0$ be an arbitrarily small constant, and let $\Delta_{\text{sCP}} = 1 + \varepsilon$ and $\Delta_{\text{Res}} = \Delta_{\text{Res}(k)} = 0$. For $\mathcal{P} \in \{\text{Res}, \text{Res}(k), \text{CP}\}$ and any $0 < \delta < (1 - \varepsilon)/(1 + \Delta)$. There is a CNF formula $F_{\mathcal{P}}$ on n variables such that*

- $F_{\mathcal{P}}$ has a \mathcal{P} -refutation of size $2^{O(n^{\delta}/\log n)}$.
- Any \mathcal{P} -refutation Π of $F_{\mathcal{P}}$ with $\text{size}(\Pi) = \exp(o(n^{\frac{1-\varepsilon}{1+\Delta_{\mathcal{P}}}-\delta}))$ has $\text{depth}(\Pi) = \exp(\Omega(n^{\delta}))$.

5 Proof of the Depth Condensation Theorem

In this section we prove the [Depth Condensation Theorem](#). To do so, it will be convenient to work with the following variant of the classic *Prover-Adversary* games of Pudlák [28] (this game was also considered in the work of Atserias and Dalmau [3]). Our variant characterizes the depth of bounded-width Resolution proofs.

Width-Bounded Prover-Adversary Game. The *Prover-Adversary* game associated with an n -variate formula F is played between two competing players, Prover and Adversary. The game proceeds in rounds, where in each round the state of the game is recorded by a partial assignment $\rho \in \{0, 1, *\}^n$ to the variables of F . Initially the state is the empty assignment $\rho = *^n$. Then, in each round, the Prover performs the following:

- *Query.* The Prover chooses an $i \in [n]$ with $\rho_i = *$, and the Adversary chooses $b \in \{0, 1\}$. The state is updated by $\rho_i \leftarrow b$ and play continues.
- *Forget.* The prover chooses a (possibly empty) set $S \subseteq [n]$ with $\rho_i \neq *$ for all $i \in S$. The state is updated by $\rho_i \leftarrow *$ for all $i \in S$.

The game ends when the state ρ falsifies an axiom of F .

Let $w > 0$ be an integer. A Prover-Adversary game is a *w-bounded Game* if at every step in game, the Prover's memory ρ remembers assignments to at most w variables; i.e., $|\rho^{-1}(*)| \leq w$.

We will say that a game is *non-bounded* if it is not necessarily w -bounded. The next lemma is precisely Lemma 6 in [8], and can be proved using the methods of Pudlák [28].

Lemma 5.1. *For any unsatisfiable k -CNF formula F , there is a depth d , width $w \geq k$ resolution refutation of F if and only if there the Prover has a strategy that ends the $(w + 1)$ -bounded game in d rounds, regardless of the strategy for the Adversary.*

Proof. Let Π be a resolution proof of width w and depth d . We extract a strategy for the Prover as follows: the Prover will take a root-to-leaf walk down the proof. At each step, corresponding to some clause C in the resolution proof, she will maintain that she is remembering exactly the unique falsifying assignment $\neg C$ to the clause C . If C was obtained by resolving $C_1 \vee x$ and $C_2 \vee \neg x$ then she will query the variable x . If the Adversary responds with $x = 0$ then she will move to $C_1 \vee x$ and forget all assignments except for $\neg C_1 \wedge \neg x$. Proceeding in this way, we arrive at a leaf C of Π in at most d steps. By our invariant, the Prover remembers at most $w + 1$ variables at any point.

As the converse direction will not be used in our proofs, we only provide a sketch of the argument. We can view the Prover's strategy for a $(w + 1)$ -bounded game as a dag, where every node is labelled with the memory of the Prover at that step in the strategy, along with the variable that the Prover queries, and there are two outgoing edges labelled 0 and 1 respectively, corresponding to possible answers of the Adversary. For each node labelled with some memory ρ and variable x , we will relabel it with the clause formed by the negation of the literals fixed by ρ , which will be obtained from its children by resolving on x . \square

Let us set up some notation. Let $G = ([N] \cup [n], E)$ be a bipartite graph; we will think of the left-vertices as the variables of F and the right-vertices as the variables of $F \circ \text{XOR}_G$. For a set $S \subseteq [n]$ of right-vertices of G , denote by

$$\text{Fixed}(S) := \{i \in [N] : \forall (i, j) \in E, j \in S\}$$

the set of left-vertices which become isolated after removing the set S of right vertices. Similarly, for a partial assignment $\rho \in \{0, 1, *\}^n$ let $\text{Fixed}(\rho) := \text{Fixed}(S)$, where $S = [n] \setminus \rho^{-1}(*)$, be the set of variables of F which are determined by ρ . Finally, let $G \upharpoonright \rho$ denote the graph obtained by removing all left-vertices that have been set by ρ (i.e., those in $[N] \setminus \rho^{-1}(*)$), and removing all isolated vertices.

Proof Overview. First, we give a high-level sketch of the proof. Let F be a CNF formula which requires depth d to refute in resolution. This gives us a strategy for the Adversary (for the non-bounded game) which ensures that it will proceed for at least d rounds. We will use this strategy to construct a strategy for the Adversary in the w -bounded-game for $F \circ \text{XOR}_G$. Ideally, we would like to proceed as follows: if the Prover (in the w -bounded game for $F \circ \text{XOR}_G$) queries a variable, we would like to set it according to the Adversary strategy of the non-bounded game for F if it would determine the value of some variable y_i of F , that is, i would be added to $\text{Fixed}(\rho)$, and set it arbitrarily otherwise. However, in this setting the XOR gadgets may share variables and so variables may be correlated. To circumvent this, we will exploit *expansion*. Indeed, if G is a good enough boundary expander, then we can always set the constraints to whatever value we like. That is, for any subset $I \subseteq [N]$ of XOR-constraints, we can always find a *strong system of distinct representative variables*.

Strong SDR. If $G = ([N] \cup [n], E)$ is a bipartite graph and $I = \{I_1, \dots, I_t\} \subseteq [N]$ then a *system of distinct representatives* (SDR) for I is a set $J = \{J_1, \dots, J_t\} \subseteq [n]$ such that I and J form a matching where $(I_i, J_i) \in E$ for all $i \in [t]$. The SDR of I is *strong* if, furthermore, I_i is not adjacent to J_j for all $j > i$.

The following lemma can be viewed as a strengthening of the claim that expanders have matchings on small sets.

Lemma 5.2. *If $G = ([N] \cup [n], E)$ is an $(r/2, 1/2)$ -boundary expander, then any $I \subseteq [N]$ with $|I| \leq r/2$ has a strong SDR.*

Proof. For $i = 1 \dots t$ perform the following. Because I has boundary at least $|I|/2$ within G , by the pigeonhole principle there exists $\ell \in I$ and a column $j \in [n]$ such that $(\ell, j) \in E$ and $(\ell', j) \notin E$ for every $\ell' \in I$ with $\ell' \neq \ell$; fix $I_i := \ell$ and $J_i := j$. Set G to be the graph obtained by removing vertices I_i and J_i and any edge incident to either of them, and update $I \leftarrow I \setminus I_i$. Because J_i was not adjacent to any vertex besides I_i , removing J_i does not decrease the expansion of I in G and we can recurse. \square

If a partial restriction $\rho \in \{0, 1, *\}^n$ (thought of as the Prover's memory) sets some variables, this may decrease the boundary expansion of the current graph $G \upharpoonright \rho$. Therefore, at each step of the simulation, the Adversary will track a *closure* of ρ which will set some additional variables, but will ensure that the residual graph is a good boundary expander. The following operator will allow us to restore the expansion of G after removing a subset of the vertices.

Closure Operator. For a $J \subseteq [n]$, denote by $G \setminus J$ the graph obtained by taking the subgraph induced by the vertex set $[N] \cup ([n] \setminus J)$ and removing any isolated vertices (i.e. y_i for which $i \in \text{Fixed}(J)$) from $[N]$. The following lemma states that for any small J there is a *closure* $\text{Cl}(J) \supseteq J$ such that $G \setminus \text{Cl}(J)$ is still expanding; a proof can be found in [31] (Lemma 2.3), building on ideas in [2, 35].

Lemma 5.3. *Let $G = ([N] \cup [n], E)$ be an $(r, 2)$ -boundary expander. For every $J \subseteq [n]$ with $|J| \leq r/4$ there exists $\text{Cl}(J) \supseteq J$ such that $|\text{Fixed}(\text{Cl}(J))| \leq 2|J|$ and $G \setminus \text{Cl}(J)$ is an $(r/2, 3/2)$ -boundary expander.*

We are now ready to prove the main theorem of this section.

Proof of the Depth Condensation Theorem. Fix an optimal strategy D for the Adversary (in the unbounded game) which delays the game for at least d rounds on F . We will construct a Adversary strategy for the w -bounded game which delays the game for at least $d/2w$ rounds on the composed formula $F \circ \text{XOR}_G$ for any $w \leq r/4$. We denote Prover's memory at each step in the game by $\rho \in \{0, 1, *\}^n$, and let $\text{set}(\rho)$ be the collection of coordinates $i \in [n]$ such that $\rho_i \in \{0, 1\}$. We also track a partial assignment $\rho^* \supseteq \rho$ that satisfies following invariants:

- *Expansion.* $G \setminus \text{set}(\rho^*)$ is an $(r/2, 3/2)$ -boundary expander.
- *Satisfying.* ρ^* does not falsify any of the constraints of $F \circ \text{XOR}_G$.

Initially $\rho = \rho^* = *^n$ and the invariants are satisfied.

In each round we will query the Adversary strategy D at most $2w$ times. Suppose that we (as the Adversary) have played for less than $d/2w$ rounds such the invariants are satisfied, then we can claim we can continue for another round and restore the invariants.

So, suppose in the current round that the Prover queries the variable x_i . If $i \in \text{set}(\rho)$ then we assign $\rho_i = \rho_i^*$; otherwise, we simply respond with an arbitrary bit. The Prover then chooses an arbitrary set of indices $I \subseteq \text{set}(\rho)$ and forgets all assignments in them. Let $\mu \subseteq \rho \cup \{x_i = b\}$ denote the partial assignment after the Adversary responds and the Prover forgets the chosen subset.

First, we note that by the *Expansion* invariant, each constraint not in $\text{Fixed}(\text{set}(\rho^*))$ had at least two free variables before x_i was set, and therefore setting x_i could not have fixed any of the XORs. So, we only need to describe how to restore the invariants. If $G \setminus \text{set}(\mu)$ is a $(r/2, 3/2)$ -boundary expander then we are done: we simply update $\rho = \rho^* = \mu$ and continue to the next stage. The interesting case is when $G \setminus \text{set}(\mu)$ is not an $(r/2, 3/2)$ -boundary expander.

Because $G \setminus \text{set}(\rho^*)$ was an $(r/2, 3/2)$ -boundary before setting x_i , it follows that after setting it, $G \setminus (\text{set}(\rho^*) \cup \{i\})$ is, at worst, an $(r/2, 1/2)$ -boundary expander. Applying [Lemma 5.3](#) to $\text{set}(\mu)$ we get a new collection of coordinates $\text{Cl}(\text{set}(\mu)) \supseteq \text{set}(\mu)$ such that $|\text{Fixed}(\text{Cl}(\text{set}(\mu)))| \leq 2|\text{set}(\mu)|$ and $G \setminus \text{Cl}(\text{set}(\mu))$ is an $(r/2, 3/2)$ boundary expander. We will now extend μ to μ^* such that $\text{set}(\mu^*) = \text{Cl}(\text{set}(\mu))$ and μ^* satisfies the invariants.

First, for every XOR constraint $I_j \in \text{Fixed}(\text{Cl}(\text{set}(\mu))) \cap \text{Fixed}(\text{set}(\rho^*))$, we have that $\text{Vars}(I_j)$ is contained in both $\text{Cl}(\text{set}(\mu))$ and $\text{set}(\rho^*)$, so for every such variable x_i we simply assign $\mu_i^* = \rho_i^*$. Let

$$I = \{I_1, \dots, I_t\} := \text{Fixed}(\text{Cl}(\text{set}(\mu))) \setminus \text{Fixed}(\text{set}(\rho^*))$$

be the set of indices of XOR-constraints whose variables must be updated. Because $G \setminus (\text{set}(\rho^*) \cup \{i\})$ is an $(r/2, 1/2)$ -boundary expander, by [Lemma 5.2](#) we can find a strong SDR $J = \{J_1, \dots, J_t\}$ for I . Therefore, for every variable x_j which is in $\text{Cl}(\text{set}(\mu)) \setminus J$ and which has not yet been assigned, we assign x_j arbitrarily in μ^* . This leaves just the variables in the strong SDR J , which we will assign according to the Adversary strategy D .

For $\ell = 1, \dots, t$ perform the following: Query the Adversary strategy D for the response $\beta \in \{0, 1\}$ when the current state is $\text{XOR}_G(\mu^*)$ and the Prover is querying the variable y_{I_ℓ} (of F). Observe that because we have already fixed all of the other variables in the neighbourhood of y_{I_ℓ} , the only free variable in the constraint corresponding to y_{I_ℓ} is x_{J_ℓ} . Therefore, we fix $\mu_{J_\ell}^*$ so that

$$\bigoplus_{j:(I_\ell, k) \in E} x_j = \beta.$$

Doing this for all ℓ satisfies all constraints by the correctness of the Adversary strategy and, as we have stated before, $G \setminus \text{Cl}(\text{set}(\mu)) = G \setminus \text{set}(\mu^*)$ is an $(r/2, 3/2)$ -expander. We can now update $\rho = \mu$ and $\rho^* = \mu^*$, restoring our invariants.

Finally, observe that since $|\text{set}(\mu)| \leq w$ by definition of a w -bounded game, it follows by [Lemma 5.3](#) that

$$|\text{Fixed}(\text{Cl}(\text{set}(\mu)))| \leq 2w,$$

and we can conclude that we query the Adversary strategy at most $2w$ times during this round. Because we have played for less than $d/2w - 1$ rounds, the Adversary has answered at most $d - 2w$ queries so far. This means that the Adversary can still provide answers to these $\leq 2w$ queries so that no constraint of F (and thus, $F \circ \text{XOR}_G$) is falsified.

6 Proofs of the Tradeoffs

6.1 Proof of the Resolution Tradeoff

We begin with [Theorem 3.1](#), which we restate next for convenience.

Theorem 3.1. *Let F be any CNF formula on N variables and let $G = ([N] \cup [n], E)$ be an $(r, 2)$ -boundary expander. If Π is a resolution refutation of $F \circ \text{XOR}_G \circ \text{XOR}_2^n$ such that $\log(\text{size}(\Pi) + 1) \leq r/12$, then*

$$\text{depth}(\Pi) \log(\text{size}(\Pi) + 1) \geq \frac{\text{depth}_{\text{Res}}(F)}{6}.$$

We require the following simple size-width lifting theorem for resolution.

Lemma 6.1. *Let F be any unsatisfiable CNF formula. For any resolution refutation Π^* of $F \circ \text{XOR}_2^n$ there is a resolution refutation Π of F such that*

$$\begin{aligned} \text{width}(\Pi) &\leq 3 \log(\text{size}(\Pi^*) + 1), \\ \text{depth}(\Pi) &\leq \text{depth}(\Pi^*). \end{aligned}$$

Proof. Let x_1, \dots, x_n be the variables of F and let $u_1, v_1, \dots, u_n, v_n$ be the variables of $F \circ \text{XOR}_2^n$. Let \mathcal{D} be the collection of partial restrictions $\rho \in \{0, 1, *\}$ that, for every $i \in [n]$, set exactly one of u_i or v_i to a value in $\{0, 1\}$ and leave the other unset. Denote by $\rho \sim \mathcal{D}$ sampling a restriction ρ uniformly at random from \mathcal{D} . It is easy to see that for any resolution refutation Π^* of $F \circ \text{XOR}_2^n$ and any $\rho^* \in \mathcal{D}$, $\Pi^* \upharpoonright \rho^*$ is a resolution refutation of F , and furthermore by closure under restrictions it follows that $\text{depth}(\Pi^* \upharpoonright \rho) \leq \text{depth}(\Pi^*)$.

Let t be a positive integer to be set later. For any clause C of $\text{width}(C) \geq t$ in Π^* , it follows that for $\rho \sim \mathcal{D}$, the probability $C \upharpoonright \rho$ is not satisfied is at most $(3/4)^t$. By a union bound, it follows that the probability that $\Pi^* \upharpoonright \rho$ has a clause of width $\geq t$ is at most $\text{size}(\Pi^*)(3/4)^t$, which is strictly less than 1 as long as $\text{size}(\Pi^*) \leq (4/3)^t$. Choosing $t = \log_{4/3}(\text{size}(\Pi^*) + 1) \leq 3 \log(\text{size}(\Pi^*) + 1)$ completes the proof. \square

By combining this lemma with the [Depth Condensation Theorem](#), we can prove [Theorem 3.1](#).

Proof of Theorem 3.1. Let Π^* be a resolution refutation of $F \circ \text{XOR}_G \circ \text{XOR}_2^n$. By [Lemma 6.1](#), there is a resolution refutation Π of $F \circ \text{XOR}_G$ with $\text{depth}(\Pi) \leq \text{depth}(\Pi^*)$ and

$$\text{width}(\Pi) \leq 3 \log(\text{size}(\Pi^*) + 1) \leq 3r/14 = r/4.$$

By the [Depth Condensation Theorem](#), it follows that

$$\text{depth}(\Pi) \log(\text{size}(\Pi) + 1) \geq \text{depth}_{\text{Res}}(F)/6.$$

\square

6.2 Proof of the k -DNF Resolution Tradeoff

Next, we establish [Theorem 3.2](#), which we restate next.

Theorem 3.2. *Let $k \geq 1$ be any constant, let F be any CNF formula on N variables, and let $G = ([N] \cup [n], E)$ be an $(r, 2)$ -boundary expander. There is a constant $\delta := \delta(k) > 0$ such that if Π is a $\text{Res}(k)$ refutation of $F \circ \text{XOR}_G \circ \text{XOR}_2^n$ with $\log(\text{size}(\Pi) + 1) \leq \delta \cdot r$, then*

$$\text{depth}(\Pi) \log^2(\text{size}(\Pi)) = \Omega(\text{depth}_{\text{Res}}(F))$$

To do so, we will prove a generic lifting theorem for $\text{Res}(k)$. For this, it will not be necessary to recall the specific rules of $\text{Res}(k)$, only that every line in a $\text{Res}(k)$ proof is a k -DNF formula.

Res(k) Lifting Theorem. *Let $k \geq 1$ be an integer and F be any CNF formula. For any $\text{Res}(k)$ refutation Π^* of $F \circ \text{XOR}_2^n$ there is a resolution refutation Π of F such that*

$$\begin{aligned} \text{width}(\Pi) &\leq k \left(\log \text{size}(\Pi^*) - \log(4k) \right) \left(\frac{4^{k+1}k}{\log e} \right)^k, \\ \text{depth}(\Pi) &\leq k \cdot \text{depth}(\Pi^*) \left(\log \text{size}(\Pi^*) - \log(4k) \right) \left(\frac{4^{k+1}k}{\log e} \right)^k. \end{aligned}$$

This theorem follows in a straightforward way from the *switching lemma* of Segerlind Buss and Impagliazzo [35], which shows that low-width DNFs can be converted into short decision trees under a random restriction.

Definition 6.2. A *decision tree* is a rooted binary tree in which every non-leaf node is labelled with a variable, the edges leaving a node are labelled with 0 and 1, and the leaves are labelled either 0 or 1. Every root-to-leaf path π in a decision tree T can be viewed as a partial assignment $\pi \in \{0, 1, *\}^n$, where, if the π takes the edge labelled $\alpha \in \{0, 1\}$ at node x_i then $\pi_i = \alpha$. We say that T computes a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ if for every $x \in \{0, 1\}^n$, the leaf of the unique root-to-leaf path π in T which agrees with x is labelled with $f(x)$. The *decision tree complexity* of computing f , $\text{DT}(f)$ is the minimum depth of any decision tree computing f .

For a DNF D over variables $\{x_1, \dots, x_n\}$, let $C(D)$ denote the *covering number* of D — the minimum size of a set $S \subseteq \{x_1, \dots, x_n\}$ such that for every term T of D , S contains at least one variable in T . The switching lemma is argued by showing that if the size of $C(D)$ is large, then many terms of D are independent and thus D is set to 1 with high probability, and if $C(D)$ is small then we can build a small decision tree computing D .

Lemma 6.3 ([35]). *Let s_0, \dots, s_{k-1} and p_1, \dots, p_k be positive numbers and let \mathcal{D} set of partial assignments such that for every $i \leq k$ and every i -DNF D' , if $C(D') > s_{i-1}$ then $\Pr_{\rho \sim \mathcal{D}}[D' \upharpoonright \rho \neq 1] \leq p_i$. Then, for every k -DNF D ,*

$$\Pr_{\rho \sim \mathcal{D}} \left[\text{DT}(D \upharpoonright \rho) > \sum_{i=0}^{k-1} s_i \right] \leq \sum_{i=1}^k p_i \cdot 2^{(\sum_{j=i}^{k-1} s_j)}.$$

In the same paper, they showed that $\text{Res}(k)$ refutations in which every line can be represented by a short decision tree can be transformed into a low-width resolution refutation.

Lemma 6.4 ([35]). *Let F be any unsatisfiable CNF formula. If Π is a $\text{Res}(k)$ refutation of F such that for every line $D \in \Pi$, $\text{DT}(D) \leq t$ then there is a resolution refutation Π^* of F with*

$$\begin{aligned} \text{width}(\Pi^*) &\leq kt, \\ \text{depth}(\Pi^*) &\leq kt \cdot \text{depth}(\Pi). \end{aligned}$$

To prove the **Res(k) Lifting Theorem**, our strategy will be to show that for any small $\text{Res}(k)$ refutation of $F \circ \text{XOR}_2$ there is a restriction such that under this restriction every line in the proof can be computed by a short decision tree.

Proof of the Res(k) Lifting Theorem. Let $F \circ \text{XOR}_2^n$ be defined over variables $u_1, v_1, \dots, u_n, v_n$. Let \mathcal{D} be the set of restrictions $\rho \in \{0, 1, *\}^{2n}$ such that for every $i \in [n]$, ρ sets exactly one of u_i, v_i to some value in $\{0, 1\}$ and leaves the other unset. Note that for any $\rho \in \mathcal{D}$, $F \circ \text{XOR}_2^n \upharpoonright \rho = F$. Fix a $\text{Res}(k)$ refutation Π of $F \circ \text{XOR}_2^n$ satisfying the assumption of the lemma. It remains to argue that there exists a restriction $\rho \in \mathcal{D}$ such that every line in $\Pi \upharpoonright \rho$ can be computed by a short decision tree.

Fix any k -DNF D . By the pigeonhole principle, there is a set of at least $C(D)/k$ variable-disjoint terms $T_1, \dots, T_{C(D)/k} \in D$. Denote by $\rho \sim \mathcal{D}$ sampling a restriction ρ from \mathcal{D} uniformly at random, and observe that the probability that ρ satisfies any term T is at most $(1/4)^k$. Therefore,

$$\Pr_{\rho \sim \mathcal{D}} [D \upharpoonright \rho \neq 1] \leq (1 - (1/4)^k)^{C(D)/k} \leq \exp(-(1/4)^k C(D)/k).$$

Denote $w := (\log \text{size}(\Pi) - \log(2k))(4^{k+1}k/\log e)^k$ and let $s_i := (w/2)(\log e/4^{i+1})^k$ and $p_i := 2^{-4s_i}$. Observe that $s_{i-1}/4 \leq s_i$. Therefore,

$$\sum_{j=i}^{k-1} s_j \leq \sum_{j=i}^{k-1} s_i/4^{j-i} \leq 2s_i,$$

and in particular, $\sum_{i=0}^{k-1} s_i \leq 2s_0 \leq (w/2) \log e \leq w$. For any i -DNF D with $C(D) \geq s_{i-1}$, we have

$$\Pr_{\rho \sim \mathcal{D}}[D \upharpoonright \rho \neq 1] \leq \exp(-(1/4)^i C(D)/i) \leq \exp(-s_{i-1} p_i/i).$$

Therefore, for any k -DNF, it follows from the switching lemma ([Lemma 6.3](#)) that

$$\Pr_{\rho \sim \mathcal{D}}[\text{DT}(D) > w] \leq \Pr_{\rho \sim \mathcal{D}}\left[\text{DT}(D) > \sum_{i=0}^{k-1} s_i\right] \leq \sum_{i=1}^k p_i \cdot 2^{\left(\sum_{j=i}^{k-1} s_j\right)} \leq \sum_{i=1}^k k 2^{2s_i} (2^{-4s_i}) \leq k 2^{-2s_k}.$$

Finally, we can conclude the lemma by taking a union bound over all DNFs in Π ,

$$\Pr_{\rho \sim \mathcal{D}}[\exists D \in \Pi : \text{DT}(D \upharpoonright \rho) > w] \leq \text{size}(\Pi) k 2^{-2s_k} = \text{size}(\Pi) \cdot k 2^{-w \left(\frac{\log e}{4^{k+1}k}\right)^k} = 1/2,$$

where the final equality follows by our setting of w . Thus, there exists some restriction $\rho \in \mathcal{D}$ such that every $D \in \Pi \upharpoonright \rho$ has $\text{DT}(D) \leq w$. Applying [Lemma 6.4](#) we can conclude that there is a resolution refutation of width at most $k w$ and depth $d k w$. \square

With this lifting theorem in hand, we are ready to prove [Theorem 3.2](#).

Proof of Theorem 3.2. Set $\delta > 0$ such that $\delta = (4^{k+1}k/\log e)^{-k}/4k + \log(4k)/r$. Let Π^* be any $\text{Res}(k)$ refutation of $F \circ \text{XOR}_2^n$ with $\log \text{size}(\Pi^*) \leq \delta \cdot r$, and denote by $t := (\log \text{size}(\Pi^*) - \log(4k))(4^{k+1}k/\log e)^k$. By the [Res\(k\) Lifting Theorem](#), there exists a resolution refutation Π with $\text{depth}(\Pi) \leq k t \cdot \text{depth}(\Pi^*)$ and

$$\text{width}(\Pi) \leq k t \leq k(\delta r - \log(4k)) \left(\frac{4^{k+1}k}{\log e}\right)^k = r/4.$$

Applying the [Depth Condensation Theorem](#), we have that

$$\frac{\text{depth}_{\text{Res}}(F)}{2} \leq \text{depth}(\Pi) \text{size}(\Pi) \leq (k t)^2 \text{depth}(\Pi^*) = O\left(\log^2 \text{size}(\Pi^*) \text{depth}(\Pi)\right),$$

which completes the proof. \square

6.3 Proof of the Semantic Cutting Planes Tradeoff

Finally, we establish [Theorem 3.3](#).

Theorem 3.3. *Let $\varepsilon > 0$ be any constant, let F be any CNF formula on N variables, and let $G = ([N] \cup [n], E)$ be an $(r, 2)$ -boundary expander. There is a constant $\delta > 0$ such that if Π is a semantic CP refutation of $F \circ \text{XOR}_G \circ \text{IND}_{n^{1+\varepsilon}}^n$ with $\log(\text{size}(\Pi)) \leq \delta \cdot r \log n$, then*

$$\text{depth}(\Pi) \log^2(\text{size}(\Pi)) = \Omega(\text{depth}_{\text{Res}}(F) \log^2 n).$$

This theorem follows almost immediately by applying the dag-like lifting theorem of Garg et al. [20], with the improved parameters from [23], and observing that their proof also preserves depth. We state this theorem next, specialized to semantic Cutting Planes.

Theorem 6.5 ([20, 23]). *Let $\varepsilon > 0$ be any constant and let F be an unsatisfiable CNF formula on n variables. For any semantic CP Π of $F \circ \text{IND}_{n^{1+\varepsilon}}^n$ there is a resolution refutation Π^* of F satisfying*

$$\begin{aligned} \text{width}(\Pi^*) &= O\left(\frac{\log \text{size}(\Pi)}{\log n}\right), \\ \text{depth}(\Pi^*) &= O\left(\frac{\text{depth}(\Pi) \log \text{size}(\Pi)}{\log n}\right). \end{aligned}$$

By combining this lifting theorem with the **Depth Condensation Theorem**, we can prove **Theorem 3.3**.

Proof of Theorem 3.3. Let Π be a semantic CP refutation of $F \circ \text{XOR}_G \circ \text{IND}_{n^{1+\varepsilon}}^n$. By **Theorem 6.5** there is a semantic Cutting Planes refutation Π^* of $F \circ \text{XOR}_G$ with $\text{depth}(\Pi^*) = O(\log \text{size}(\Pi) / \log n)$ and $\text{width}(\Pi^*) = \alpha \cdot \log \text{size}(\Pi) / \log n$ for some constant $\alpha > 0$. Setting $\delta > 0$ so that $\alpha\delta = 1/4$,

$$\text{width}(\Pi^*) = \frac{\alpha \cdot \log \text{size}(\Pi)}{\log n} \leq \alpha\delta \cdot r = r/4.$$

By the **Depth Condensation Theorem**, it follows that

$$\text{depth}(\Pi) \log^2 \text{size}(\Pi) = \Omega(\text{depth}_{\text{Res}}(F) \log^2 n).$$

□

7 Conditional Supercritical Tradeoffs for Monotone Circuits

We end by recording a supercritical size/depth tradeoff for monotone circuits, assuming the following conjecture which asserts a (quantatively) stronger supercritical size/depth tradeoff for Resolution.

Conjecture 1.4. *There are families of unsatisfiable formulas $\{F_n\}$ with $m = m(n)$ clauses such that F_n has Resolution refutations of size polynomial in m , but any polynomial-size (or subexponential-size) refutation of F_n requires depth that is superlinear in m .*

Denote by $\text{Search}(F)$ the canonical CNF search problem associated with a CNF formula F . For any k -CNF formula F with m variables and n clauses, and any gadget $g : [D] \times \{0, 1\}^D \rightarrow \{0, 1\}$, it is known that the dag-like communication complexity of $\text{Search}(F) \circ g^n$ is equivalent to the monotone circuit complexity of computing an associated monotone function $\text{mCSP-SAT}_{F,g}$ on $N = O(mD^k)$ many variables [20, 21]. Garg et al. [20] proved a dag-like lifting theorem, showing that there exists a gadget g such that from a dag-like communication protocol for $\text{Search}(F) \circ g^n$, one can extract a Resolution proof of essentially the same size. We state this lifting theorem next, with the improved parameters due to Lovett et al. [24].

Theorem 7.1 ([20, 24]). *Let F be any unsatisfiable k -CNF formula on n variables and m clauses and let $\varepsilon > 0$ be any constant. There is a monotone boolean function f_F on $mn^{k(1+\varepsilon)}2^k$ variables such that any monotone circuit C computing f_F implies a resolution refutation Π^* of F satisfying*

$$\begin{aligned} \text{size}(\Pi^*) &= O(\text{size}(C)), \\ \text{depth}(\Pi^*) &= O\left(\frac{\text{depth}(C) \log \text{size}(C)}{\log n}\right). \end{aligned}$$

Note that Garg et al. [20] state the lifting theorem for Resolution *width*, rather than size; the statement for size simply follows by the simple observation that there are at most n^{w+1} distinct clauses of width at most w .

Now, suppose that there is an $O(1)$ -CNF formula F on m clauses such that any polynomial size Resolution refutation of F requires depth $\Omega(m^{4+\varepsilon})$ for an arbitrarily small constant $\varepsilon > 0$. Then, applying this lifting theorem we obtain a supercritical size/depth tradeoff for Resolution. We state this formally for k -CNF formulas next.

Theorem 7.2. *Let $\varepsilon > 0$ be an arbitrarily small constant and let F be an unsatisfiable 3-CNF formula with m clauses on n variables and define $N := 8mn^{3(1+\varepsilon)}$. If any polynomial size Resolution refutation of F requires depth $\omega(N)$, then there is a size/depth supercritical tradeoff for monotone circuits.*

Proof. Suppose that any polynomial size Resolution refutation of F requires depth at least d for $d = \omega(N)$. Applying Theorem 7.1 completes the proof. \square

Note that the assumption that F is a 3-CNF formula is essentially without loss of generality. Indeed, if F is a k -CNF formula on m clauses and n variables, then by introducing additional “extension” variables, we can transform it into an equivalent 3-CNF formula as follows: for each k -clause $C = \ell_1 \vee \dots \vee \ell_k$ with $k > 3$, introduce $k - 1$ new variables y_1, \dots, y_{k-1} , and replace C by the following k -clauses

$$y_1 \vee \ell_1, \neg y_1 \vee \ell_2 \vee y_2, \dots, \neg \ell_{k-1} \vee \ell_k.$$

The resulting CNF formula has at most $n + mk$ variables and at most mk clauses. Finally, observe that the original CNF formula can be derived from the extended formula in size at most $O(mk)$ and depth $\log(k)$ by simply resolving on the y -variables. Thus, for k -CNF formulas we obtain Theorem 7.2 with $N = O(mk(n + mk)^{3(1+\varepsilon)})$.

Finally, note that a lifting theorems for a constant-size gadget would allow us to obtain supercritical size/depth tradeoffs for monotone circuits from supercritical tradeoffs for Resolution that are *barely* superlinear in m (in contrast with the $\Omega(m^4)$ required by the current lifting theorem). We state this next.

Observation 7.3. Let F be any unsatisfiable k -CNF formula on n variables and m clauses and let $g : [D] \times \{0, 1\}^D \rightarrow \{0, 1\}$ be any function such that any dag-like communication protocol for $\text{Search}(F) \circ g^n$ implies a Resolution refutation Π^* of F satisfying

$$\begin{aligned} \text{size}(\Pi^*) &= \text{poly}(\text{size}(C)), \\ \text{depth}(\Pi^*) &= O\left(\frac{\text{depth}(C) \log \text{size}(C)}{\log n}\right). \end{aligned}$$

If any polynomial size Resolution refutation of F requires depth at least $\omega(mD^k)$ then there is a supercritical size/depth tradeoff for monotone circuits.

Proof. Recall from the beginning of this section that a dag-like communication protocol for solving $\text{Search}(F) \circ g^n$ is equivalent to a monotone circuit computing the monotone function $\text{mCSP-SAT}_{F,g}$ on $N = O(mD^k)$ many variables. By the assumption that any polynomial-size Resolution refutation of F requires depth $\omega(mD^k)$, and the assumed lifting theorem, we obtain a supercritical size/depth tradeoff for monotone circuits. \square

References

- [1] Karen Aardal, Robert E. Bixby, Cor A. J. Hurkens, Arjen K. Lenstra, and Job W. Smeltink. Market split and basis reduction: Towards a solution of the cornuéjols-dawande instances. *INFORMS J. Comput.*, 12(3):192–202, 2000.
- [2] Michael Alekhnovich. Lower bounds for k-DNF resolution on random 3-CNFs. In Harold N. Gabow and Ronald Fagin, editors, *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 251–256. ACM, 2005.
- [3] Albert Atserias and Víctor Dalmau. A combinatorial characterization of resolution width. *J. Comput. Syst. Sci.*, 74(3):323–334, 2008.
- [4] Paul Beame, Chris Beck, and Russell Impagliazzo. Time-space trade-offs in resolution: Super-polynomial lower bounds for superlinear space. *SIAM J. Comput.*, 45(4):1612–1645, 2016.
- [5] Paul Beame, Noah Fleming, Russell Impagliazzo, Antonina Kolokolova, Denis Pankratov, Toniann Pitassi, and Robert Robere. Stabbing planes. In *9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA, USA*, pages 10:1–10:20, 2018.
- [6] Chris Beck, Jakob Nordström, and Bangsheng Tang. Some trade-off results for polynomial calculus: extended abstract. 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 813–822. ACM, 2013.
- [7] Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow - resolution made simple. *J. ACM*, 48(2):149–169, 2001.
- [8] Christoph Berkholz. On the complexity of finding narrow proofs. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 351–360. IEEE Computer Society, 2012.
- [9] Christoph Berkholz and Jakob Nordström. Supercritical space-width trade-offs for resolution. *SIAM J. Comput.*, 49(1):98–118, 2020.
- [10] Joshua Buresh-Oppenheim, Nicola Galesi, Shlomo Hoory, Avner Magen, and Toniann Pitassi. Rank bounds and integrality gaps for cutting planes procedures. *Theory of Computing*, 2(4):65–90, 2006.
- [11] Siu Man Chan. Just a pebble game. In *Proceedings of the 28th Conference on Computational Complexity, CCC 2013, K.lo Alto, California, USA, 5-7 June, 2013*, pages 133–143. IEEE Computer Society, 2013.
- [12] Vašek Chvátal, William Cook, and Mark Hartmann. On cutting-plane proofs in combinatorial optimization. *Linear algebra and its applications*, 114:455–499, 1989.
- [13] Daniel Dadush and Samarth Tiwari. On the complexity of branching proofs. In Shubhangi Saraf, editor, *35th Computational Complexity Conference, CCC 2020, July 28-31, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 169 of *LIPICs*, pages 34:1–34:35. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

- [14] Susanna F. de Rezende, Or Meir, Jakob Nordström, Toniann Pitassi, Robert Robere, and Marc Vinyals. Lifting with simple gadgets and applications to circuit and proof complexity. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 24–30. IEEE, 2020.
- [15] Friedrich Eisenbrand and Andreas S. Schulz. Bounds on the chvátal rank of polytopes in the 0/1-cube. In Gérard Cornuéjols, Rainer E. Burkard, and Gerhard J. Woeginger, editors, *Integer Programming and Combinatorial Optimization, 7th International IPCO Conference, Graz, Austria, June 9-11, 1999, Proceedings*, volume 1610 of *Lecture Notes in Computer Science*, pages 137–150. Springer, 1999.
- [16] Juan Luis Esteban and Jacobo Torán. A combinatorial characterization of treelike resolution space. *Inf. Process. Lett.*, 87(6):295–300, 2003.
- [17] Yuval Filmus, Pavel Hrubeš, and Massimo Lauria. Semantic versus syntactic cutting planes. In Nicolas Ollinger and Heribert Vollmer, editors, *33rd Symposium on Theoretical Aspects of Computer Science, STACS 2016, February 17-20, 2016, Orléans, France*, volume 47 of *LIPICs*, pages 35:1–35:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.
- [18] Noah Fleming, Mika Göös, Russell Impagliazzo, Toniann Pitassi, Robert Robere, Li-Yang Tan, and Avi Wigderson. On the power and limitations of branch and cut. In Valentine Kabanets, editor, *36th Computational Complexity Conference, CCC 2021, July 20-23, 2021, Toronto, Ontario, Canada (Virtual Conference)*, volume 200 of *LIPICs*, pages 6:1–6:30. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [19] Noah Fleming, Denis Pankratov, Toniann Pitassi, and Robert Robere. Random $\Theta(\log n)$ -CNFs are hard for cutting planes. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 109–120, 2017.
- [20] Ankit Garg, Mika Göös, Prithish Kamath, and Dmitry Sokolov. Monotone circuit lower bounds from resolution. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 902–911. ACM, 2018.
- [21] Mika Göös and Toniann Pitassi. Communication lower bounds via critical block sensitivity. *SIAM J. Comput.*, 47(5):1778–1806, 2018.
- [22] Pavel Hrubeš and Pavel Pudlák. Random formulas, monotone circuits, and interpolation. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 121–131, 2017.
- [23] Shachar Lovett, Raghu Meka, Ian Mertz, Toniann Pitassi, and Jiapeng Zhang. Lifting with sunflowers. unpublished.
- [24] Shachar Lovett, Raghu Meka, Ian Mertz, Toniann Pitassi, and Jiapeng Zhang. Lifting with sunflowers. In *Electron. Colloquium Comput. Complex.*, page 111, 2020.
- [25] Theodoros Papamakarios and Alexander A. Razborov. Space characterizations of complexity measures and size-space trade-offs in propositional proof systems. *Electron. Colloquium Comput. Complex.*, 28:74, 2021.

- [26] Wolfgang J. Paul, Robert Endre Tarjan, and James R. Celoni. Space bounds for a game on graphs. *Math. Syst. Theory*, 10:239–251, 1977.
- [27] Pavel Pudlák. Lower bounds for resolution and cutting plane proofs and monotone computations. *J. Symb. Log.*, 62(3):981–998, 1997.
- [28] Pavel Pudlák. Proofs as games. *Am. Math. Mon.*, 107(6):541–550, 2000.
- [29] Ran Raz and Avi Wigderson. Monotone circuits for matching require linear depth. *J. ACM*, 39(3):736–744, 1992.
- [30] Alexander A. Razborov. Lower bounds on monotone complexity of the logical permanent. *Mathematical Notes of the Academy of Sciences of the USSR*, 37(6):485–493, 1985.
- [31] Alexander A. Razborov. A new kind of tradeoffs in propositional proof complexity. *J. ACM*, 63(2):16:1–16:14, 2016.
- [32] Alexander A. Razborov. On the width of semialgebraic proofs and algorithms. *Math. Oper. Res.*, 42(4):1106–1134, 2017.
- [33] Alexander A. Razborov. On space and depth in resolution. *Comput. Complex.*, 27(3):511–559, 2018.
- [34] Thomas Rothvoß and Laura Sanita. 0/1 polytopes with quadratic chvátal rank. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 349–361. Springer, 2013.
- [35] Nathan Segerlind, Samuel R. Buss, and Russell Impagliazzo. A switching lemma for small restrictions and lower bounds for k-DNF resolution. *SIAM J. Comput.*, 33(5):1171–1200, 2004.