

Automating Algebraic Proof Systems Is NP-Hard

Susanna F. de Rezende
Institute of Mathematics of the Czech
Academy of Sciences
Prague, Czechia
rezende@math.cas.cz

Mika Göös
EPFL
Lausanne, Switzerland
mika.goos@epfl.ch

Jakob Nordström
University of Copenhagen & Lund
University
Copenhagen, Denmark and Lund,
Sweden
jakob.nordstrom@cs.lth.se

Toniann Pitassi
University of Toronto and IAS
Toronto, Canada and Princeton, U.S.A
toni@cs.toronto.edu

Robert Robere
McGill University
Montréal, Canada
robere@cs.mcgill.ca

Dmitry Sokolov
St. Petersburg State University and
PDMI RAS
St. Petersburg, Russia
sokolov.dmt@gmail.com

ABSTRACT

We show that algebraic proofs are hard to find: Given an unsatisfiable CNF formula F , it is NP-hard to find a refutation of F in the Nullstellensatz, Polynomial Calculus, or Sherali–Adams proof systems in time polynomial in the size of the shortest such refutation. Our work extends, and gives a simplified proof of, the recent breakthrough of Atserias and Müller (JACM 2020) that established an analogous result for Resolution.

CCS CONCEPTS

• Theory of computation → Proof complexity.

KEYWORDS

proof complexity, automatability, pigeonhole principle, algebraic proof systems, lower bounds

ACM Reference Format:

Susanna F. de Rezende, Mika Göös, Jakob Nordström, Toniann Pitassi, Robert Robere, and Dmitry Sokolov. 2021. Automating Algebraic Proof Systems Is NP-Hard. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC '21)*, June 21–25, 2021, Virtual, Italy. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3406325.3451080>

1 INTRODUCTION

Automatability. A proof system S is (polynomial-time) *automatable* [18] if there is an algorithm that takes as input an unsatisfiable CNF formula F and outputs an S -refutation of F in time polynomial in the size of the shortest S -refutation of F (plus the size of F). Intuitively, automatability addresses the proof search problem: How hard is it to *find* a proof? Automatability (or lack thereof) for well-studied proof systems is a central question for automated theorem proving and SAT solving.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
STOC '21, June 21–25, 2021, Virtual, Italy

© 2021 Association for Computing Machinery.
ACM ISBN 978-1-4503-8053-9/21/06...\$15.00
<https://doi.org/10.1145/3406325.3451080>

For example, state-of-the-art SAT solvers using conflict-driven clause learning (CDCL) [9, 48, 50] are based on the most basic propositional proof system, Resolution (R for short). This means that running a CDCL solver (without preprocessing) on an unsatisfiable formula F produces a Resolution refutation of F [12]. Thus, non-automatability of Resolution (studied in a long line of work [2–4, 7, 39, 49]) implies that any SAT solver based on Resolution will require superpolynomial time even on formulas that are easy, that is, admit a polynomial-size refutation.

Algebraic Proof Systems. In this paper, we study the automatability of *algebraic* proof systems. We show that it is NP-hard to automate any of the following standard systems:

- (NS) Nullstellensatz [11],
- (PC) Polynomial Calculus [1, 22],
- (SA) Sherali–Adams [63].

An important proof system that is missing above, and for which we still leave open the question of its automatability, is

- (SoS) Sum-of-Squares [44, 55, 64].

1.1 Our Results

For the aforementioned proof systems (excluding SoS), our main result shows that it is NP-hard to approximate the minimum refutation size up to a factor of 2^{n^ϵ} for some constant $\epsilon > 0$. We defer the standard definitions of the algebraic proof systems to Section 8. Our result holds regardless of definitional details such as which underlying field (real numbers, finite fields) we choose, or whether we allow twin variables (separate formal variables for negated literals).

THEOREM 1.1 (MAIN RESULT). *There is a polynomial-time algorithm \mathcal{A} that on input an n -variate 3-CNF formula F outputs a CNF formula $\mathcal{A}(F)$ such that for any system $S = R, NS, PC, SA$:*

- If F is satisfiable, then $\mathcal{A}(F)$ admits an S -refutation of size at most $n^{O(1)}$.
- If F is unsatisfiable, then $\mathcal{A}(F)$ requires S -refutations of size at least $2^{n^{\Omega(1)}}$.

A direct corollary of Theorem 1.1 is that we can rule out automatability in polynomial, quasi-polynomial and subexponential time under corresponding hardness assumptions. To state this more precisely, let QP be the class of problems that can be solved in

time $\exp(\log^{O(1)} n)$ and SUBEXP those that can be solved in time $\exp(n^{o(1)})$.

Corollary 1.2. *For any system $S = R, NS, PC, SA$:*

- *S is not automatable in polynomial time unless $NP \subseteq P$.*
- *S is not automatable in quasi-polynomial time unless $NP \subseteq QP$.*
- *S is not automatable in subexponential time unless $NP \subseteq SUBEXP$.*

We emphasize that our theorem handles all of the proof systems simultaneously. That is, there is one common polynomial-time constructible formula $\mathcal{A}(F)$ that is either easy for all the proof systems, or hard for all of them. This means that proof search is hard for R and NS even if we are allowed to search for proofs in a stronger system like PC and SA .

Previously, Galesi and Lauria [27], building on [3], proved that NS and PC are not polynomial-time automatable unless the fixed parameter hierarchy collapses. Our result upgrades this to an optimal hardness assumption, namely $P \neq NP$. For SA , no previous non-automatability results were known. As for upper bounds, the fastest-known search algorithms for PC , SA , and SoS run in exponential time $\exp(\tilde{O}(\sqrt{n \log s}))$, where s is the proof size and the \tilde{O} -notation hides $\text{poly}(\log n)$ factors. All these algorithms are based on general size-degree tradeoffs [6, 22, 38, 56].

1.2 Techniques

Our proof builds on the recent breakthrough of Atserias and Müller [7] that showed that automating Resolution is NP-hard. Namely, they proved [Theorem 1.1](#) for $S = R$. We give a simpler proof of their theorem that generalizes better, handling more systems simultaneously. The key new ingredient in our approach is a reduction from the *pigeonhole principle* to prove the lower bound in case F is unsatisfiable. As a further simplification, we show how standard size-width tradeoffs can be used to eliminate the “relativization/lifting” step in the Atserias and Müller proof by tweaking their construction of $\mathcal{A}(F)$ slightly. See [Section 2](#) for a detailed overview of our techniques.

1.3 Other Related Work

Degree-Automatability. Many algebraic proof systems possess a (weaker) form of automatability known as *degree-automatability* (as opposed to *size-automatability*), which enables proofs of *low degree* to be found efficiently. For our four systems, proofs of degree d can be found in time $n^{O(d)}$ for n -variate formulas: for NS and SA this can be achieved by solving an LP; for PC see [22]; for SoS (under technical assumptions that cover the case of CNF formulas) see [51, 61].

Degree (or size) automatability yields a meta-approach for search problems. Namely, when the *existence* of a solution can be proven via a low-degree (or small size) proof then degree (or size) automatability can be applied to generate an efficient algorithm for finding a solution. This proofs-as-algorithms approach has led to many beautiful and sometimes surprising new approximation algorithms for a variety of optimization and average-case parameter estimation problems. Examples include dictionary learning [8], tensor decomposition [47], learning mixtures of Gaussians [42], and constraint

satisfaction problems [36, 52]. What makes these algebraic proof systems special is that they hit a sweet spot, possessing strong power but also being weak enough to admit nontrivial proof search. For example, SA (resp. SoS) gives a standard way of tightening LP (resp. SDP) relaxations of boolean LPs in order to improve performance. Another example of their power is that SA and SoS are able to prove many useful (anti-)concentration inequalities in constant degree [53]. For a comprehensive introduction to the interplay between algebraic proofs and algorithms, see the monograph [26].

Size-Degree Tradeoffs. Degree-automatability has an interesting consequence for the way non-automatability results are proved: The formula $\mathcal{A}(F)$ we construct admits a short refutation when F is satisfiable, but every such refutation must require large degree (otherwise degree-automatability would allow us to find them quickly). Such formulas—admitting short proofs but none of small degree—were known to exist for R [17] and PC [28] (and for NS this is implicit in [20]). No such CNF formulas have yet been exhibited for SoS , although progress towards this goal was recently made in [57].

Other Proof Systems. For standard textbook-style proof systems (Frege and Extended Frege) *weak automatability* [4]—that is, being polynomially simulated by an automatable proof system—is equivalent to possessing *feasible interpolation*. More specifically, for any proof system that is closed under restrictions, weak automatability implies feasible interpolation [19], and for sufficiently strong proof systems (that admit short proofs of their soundness), the converse holds [59]. Under cryptographic assumptions, Frege, Extended Frege, and bounded-depth Frege systems are known to not have feasible interpolation and therefore are not even weakly automatable [16, 18, 43].

By contrast, for weak systems that seemingly cannot reason about their own soundness (R, NS, PC, SA, SoS), deciding whether they are automatable has proven more challenging. Until the recent breakthrough by Atserias and Müller [7], even the automatability of Resolution was unresolved. In an earlier important paper, Alekhovich and Razborov [3] ruled out automatability of Resolution under the assumption that the fixed parameter hierarchy is proper. However, the best upper bound on the time complexity remained exponential, and it remained open for a long time whether or not this upper bound could be improved until this question was finally resolved in [7]. Following in the wake of Atserias and Müller, non-automatability results have also been shown for other weak proof systems such as regular and ordered Resolution [13] (building on a preliminary version of this paper), k -DNF Resolution [31], and cutting planes [34].

2 PROOF OVERVIEW

Let us now give an overview of how we modify [7] to construct our new proof. In this section:

- (§2.1) We recall the the definition of the Resolution proof system.
- (§2.2) We outline a simpler proof of the Atserias–Müller theorem ([Theorem 1.1](#) for Resolution).

(§2.3) We outline why our simplified proof generalizes, with some additional work, to the setting of algebraic proof systems.

Readers who only care about our simplified proof of Atserias–Müller are in luck: We have organized the paper so that the initial Sections 3–7 present the simplified proof in a self-contained fashion. In particular, no knowledge of algebraic proof systems is required there.

2.1 Resolution Basics

Fix an unsatisfiable CNF formula F over variables x_1, \dots, x_n . We call the clauses of F *axioms* and often think of them as sets of literals (x_i or \bar{x}_i , where bar denotes negation). A Resolution refutation \mathcal{P} of F , or R-refutation for short, is a sequence of clauses $\mathcal{P} = (C_1, \dots, C_s)$ ending in the empty clause $C_s = \emptyset$ such that each C_i is either (i) an axiom of F ; or (ii) derived from clauses $C_j, C_{j'}$, where $j, j' < i$, using one of the following rules:

- *Resolution rule*: $C_i = (C_j \setminus \{x_k\}) \cup (C_{j'} \setminus \{\bar{x}_k\})$ where $x_k \in C_j$ and $\bar{x}_k \in C_{j'}$.
- *Weakening rule*: $C_i \supseteq C_j$.

The *size* of the refutation is $\|\mathcal{P}\| := s$. The Resolution size complexity of F , denoted $R(F)$, is the size of a smallest Resolution refutation of F . Another important complexity measure of a refutation \mathcal{P} is its *width* $w(\mathcal{P})$ defined as the maximum width $|C|$ of any of its clauses $C \in \mathcal{P}$. Define also the *width complexity* $w_R(F)$ of a formula F as the least width of a Resolution refutation of F .

For visualization purposes, a refutation \mathcal{P} can be thought of as a *directed acyclic graph (dag)*, also called the *refutation dag*: Introduce a node v_i for every clause C_i , and include a directed edge (j, i) if C_j is used to derive C_i . The final clause C_s becomes a *root node* (no parent), while the axioms are *leaves* (no children). A refutation is *tree-like* if this graph is a tree (note that the same clause can label several different nodes), and otherwise it is *dag-like*.

2.2 A Simpler Proof for the Non-Automatability of Resolution

Suppose we are given an n -variate 3-CNF formula F as input. The algorithm \mathcal{A} that Atserias and Müller devised computes in two steps: In the first step, the algorithm constructs a “refutation formula” denoted by $\text{Ref}(F)$. In the second step, this formula is “lifted” to produce $\text{Lift}(\text{Ref}(F))$, which is then output by \mathcal{A} . We explain these two steps in detail.

Step 1: A Block-Width Lower Bound. The refutation formula $\text{Ref}(F)$ (defined precisely in Section 3.1) intuitively states

$\text{Ref}(F) \equiv$ “ F admits a short dag-like Resolution refutation.”

For now, it suffices to say that the variables of $\text{Ref}(F)$ come partitioned into some number of *blocks*. For a clause C over the variables of $\text{Ref}(F)$, we define its *block-width* $\text{bw}(C)$ as the number of distinct blocks that C *touches*, that is, from which it contains a variable. For a Resolution refutation \mathcal{P} (resp. formula F), we define its *block-width* $\text{bw}(\mathcal{P})$ (resp. $\text{bw}(F)$) as the maximum block-width of its clauses. Finally, for a formula F , we define its *block-width complexity* $w_R(F)$ as the minimum block-width of a Resolution refutation of F .

The key property of $\text{Ref}(F)$ is that its block-width complexity depends drastically on F ’s satisfiability.

Lemma 2.1 (Atserias–Müller). *There is a polynomial-time algorithm that on input an n -variate 3-CNF formula F outputs a block-width- $O(1)$ CNF formula $\text{Ref}(F)$ such that*

- If F is satisfiable, then $\text{Ref}(F)$ admits a size- $n^{O(1)}$ block-width- $O(1)$ resolution refutation.*
- If F is unsatisfiable, then $\text{Ref}(F)$ requires resolution refutations of block-width $n^{\Omega(1)}$.*

We present the upper bound (i) in Section 7 for completeness (and also because our definition of $\text{Ref}(F)$ differs slightly from that of Atserias and Müller). Our main simplification is for the block-width lower bound (ii).

Simplifying Part (ii) of Lemma 2.1. Atserias and Müller originally proved the lower bound (ii) by a direct ad-hoc adversary argument. This was the most involved step in their proof. Our proof is by a mere *reduction* from the usual *pigeonhole principle*. We define in Section 3.3 a convenient, somewhat non-standard encoding of the principle, sometimes called the *retraction weak pigeonhole principle* [40, 60]. This encoding, denoted rPHP_m , is an $O(\log m)$ -width CNF formula that claims there exists an efficiently invertible injection, encoded in binary, from $2m$ pigeons to m holes. Our reduction in Section 5 translates, with modest loss, width complexity lower bounds for rPHP_{n^2} into block-width complexity lower bounds for $\text{Ref}(F)$.

Lemma 2.2. *For any n -variate unsatisfiable formula F we have $\text{bw}_R(\text{Ref}(F)) \geq \tilde{\Omega}(w_R(\text{rPHP}_{n^2})/n)$.*

Our simplified proof of (ii) is then concluded by invoking known width lower bounds for pigeonhole principles. Indeed, standard techniques [60, Proposition 3.4] show that

$$w_R(\text{rPHP}_m) \geq \Omega(m).$$

This lower bound and Lemma 2.2 imply that $\text{bw}_R(\text{Ref}(F)) \geq \tilde{\Omega}(n)$, which proves (ii).

Step 2: From Block-Width to Size. The goal of the second step is to transform the block-width gap in Lemma 2.1 into a size gap. There are two alternative approaches to achieve this.

Lifting. This technique was used by Atserias and Müller, although they called it *relativization* after [23]; see also [30]. Lifting techniques have produced a plethora of applications in proof complexity; recent examples include [24, 25, 29, 33–35, 37]. The general strategy is this: We start with a formula F that is hard in some weak sense (for us, block-width). Then we *compose* (or *lift*) the formula with a carefully chosen *gadget*—usually, each variable of F is replaced with a copy of the gadget—to produce a formula $\text{Lift}(F)$, which we then show is hard in a strong sense (for us, Resolution size).

Tradeoff. The famous size–width tradeoff of Ben-Sasson and Wigderson [14] states that any n -variate low-width formula F that has high width complexity, namely $w_R(F) \gg \sqrt{n}$, also has exponentially large size complexity. Atserias and Müller’s original proof did not use the tradeoff result, as their encoding of $\text{Ref}(F)$ did not admit a high enough width complexity. We observe that by defining $\text{Ref}(F)$ in a succinct enough way (technically speaking, using *binary* rather than *unary* encoding to represent numbers), the width complexity $w_R(\text{Ref}(F))$

ends up in a regime where the tradeoff result applies, which gives us an exponential size lower bound without the need for any gadget composition.

We think both approaches have merits. Lifting is the more robust technique: it is more widely applicable than the tradeoff, as it applies even if the starting formula F has only a small polynomial (block-)width complexity. However, given that we have been able (somewhat unintentionally) to optimize the encoding of $\text{Ref}(F)$, the tradeoff approach can give us a shorter proof.

We will opt to focus on the lifting approach in this paper. We do, however, outline briefly how the alternative tradeoff approach can be carried out in [Section 6.4](#).

Block Lifting. We prove in [Section 6](#) a lifting lemma whose notable feature is that it is *block-aware*: the gadgets corresponding to a single block will *share* some input variables. This allows us to lift block-width (rather than width) to Resolution size. The lemma is simple to prove via random restrictions: a proof is implicit in Atserias–Müller, and an even stronger version (lifting to Cutting Planes size) was proved in [34]. We formulate the lemma here for completeness, and also in order to generalize it to algebraic systems later (see [Section 2.3](#)).

Lemma 2.3 (Block lifting). *There is a polynomial-time algorithm that on input a block-width- $O(1)$ CNF formula F outputs a CNF formula $\text{Lift}(F)$ such that*

$$2^{\Omega(\text{bw}_R(F))} \leq R(\text{Lift}(F)) \leq 2^{O(\text{bw}(\mathcal{P}))} \cdot \|\mathcal{P}\|,$$

where \mathcal{P} is any Resolution refutation of F .

The main theorem for Resolution follows immediately by combining [Lemma 2.1](#) and [Lemma 2.3](#). Namely, the algorithm that computes $\mathcal{A}(F) := \text{Lift}(\text{Ref}(F))$ satisfies [Theorem 1.1](#) for Resolution. This completes the overview of our simplified proof of the non-automatability of Resolution.

2.3 Generalization to Algebraic Systems

Generalizing the proof from the previous subsection (using lifting) to algebraic systems $S = \text{NS}, \text{PC}, \text{SA}$ is now a matter of generalizing the block-width-based [Lemma 2.1](#) and [2.3](#).

Terminology. The algebraic proof systems are defined formally in [Section 8](#). For the purpose of this overview, we only sketch some notation. The analogue of width in an algebraic system S is *degree*. The degree of a monomial r is denoted $\text{deg}(r)$; the maximum degree of a monomial in an S -refutation \mathcal{P} is denoted $\text{deg}(\mathcal{P})$; the minimum degree of an S -refutation of a formula F is denoted $\text{deg}_S(F)$. Moreover, we define the *block-degree* $\text{bdeg}(r)$ of a monomial r as the number of blocks that r touches; we extend this definition to refutations and formulas as before. For convenience, when talking about Resolution, we use (block-)degree to mean (block-)width. Finally, we use $S(F)$ to denote the least size $\|\mathcal{P}\|$ of an S -refutation \mathcal{P} of F , measured as the number of monomials in \mathcal{P} .

Improved Lemmas. We now formulate the improved versions of [Lemma 2.1](#) and [2.3](#). The statements are as expected, except we

replace the formula $\text{Ref}(F)$ with a tree-like variant $\text{TreeRef}(F)$, discussed shortly. Our main result ([Theorem 1.1](#)) follows by considering $\mathcal{A}(F) := \text{Lift}(\text{TreeRef}(F))$ and applying the improved lemmas. The remainder of this section discusses how to prove these lemmas.

Lemma 2.4 (Improved [Lemma 2.1](#)). *There is a polynomial-time algorithm that on input an n -variate 3-CNF formula F outputs a block-width- $O(1)$ CNF formula $\text{TreeRef}(F)$ such that for proof systems $S = \text{R}, \text{NS}, \text{PC}, \text{SA}$ the following holds:*

- (i) *If F is satisfiable, then $\text{TreeRef}(F)$ admits a size- $n^{O(1)}$ block-degree- $O(1)$ S -refutation.*
- (ii) *If F is unsatisfiable, then $\text{TreeRef}(F)$ requires S -refutations of block-degree $n^{\Omega(1)}$.*

Lemma 2.5 (Improved [Lemma 2.3](#)). *There is a polynomial-time algorithm that on input a block-width- $O(1)$ CNF formula F outputs a CNF formula $\text{Lift}(F)$ such that for proof systems $S = \text{R}, \text{NS}, \text{PC}, \text{SA}$ it holds that*

$$2^{\Omega(\text{bdeg}_S(F))} \leq S(\text{Lift}(F)) \leq 2^{O(\text{bdeg}(\mathcal{P}))} \cdot \|\mathcal{P}\|,$$

where \mathcal{P} is any S -refutation of F .

Upper Bound (i). The first challenge in generalizing the proof for Resolution is that we do not know whether $\text{Ref}(F)$ for a satisfiable F admits a small Nullstellensatz refutation (we suspect not). This is why we introduce in [Section 3.2](#) a new tree-like variant of the formula that intuitively says

$\text{TreeRef}(F) \equiv$ “ F admits a short tree-like Resolution refutation where weakening is only applied on axioms.”

This formula is a *strengthening* of $\text{Ref}(F)$, meaning that it is obtained from $\text{Ref}(F)$ by adding new variables and axioms. The addition of the tree structure allows us to show the upper bound for Nullstellensatz. The upper bound for Resolution is inherited from $\text{Ref}(F)$, and for other systems they follow by simulations. See [Section 11](#) for the proof of [Lemma 2.4\(i\)](#).

Lower Bound (ii). Our simplified proof established the block-width lower bound for $\text{Ref}(F)$ by a reduction from rPHP_{n^2} . In fact, the same reduction works even for $\text{TreeRef}(F)$ without modification. Moreover, it is known that pigeonhole formulas require large degree for PC [62] and SA [32]. We show, via low-degree reductions, that these degree lower bounds apply also to our rPHP_{n^2} encoding, and hence to $\text{TreeRef}(F)$. See [Section 9](#) for the proof of [Lemma 2.4\(ii\)](#).

Lifting Block-Degree. Algebraic proofs are equally amenable to analysis via random restrictions (the key technique behind the proof of [Lemma 2.3](#)) as is Resolution. Hence it is straightforward to strengthen [Lemma 2.3](#) to [Lemma 2.5](#). See [Section 10](#) for the proof.

3 FORMULAS

In this section we define formulas that will be used throughout the paper. In [§3.1](#) we introduce a variant of the refutation formula $\text{Ref}(F)$ of Atserias and Müller [7]; in [§3.2](#) we modify $\text{Ref}(F)$ to obtain our tree-like variant, $\text{TreeRef}(F)$; and finally in [§3.3](#) we define a convenient version of the usual pigeonhole principle.

3.1 $\text{Ref}(F)$ Formula

Fix a CNF formula F with variables x_1, \dots, x_n and $m = \text{poly}(n)$ clauses. We define another CNF formula $\text{Ref}(F)$ that states that “ F

admits a short dag-like Resolution refutation.” Our definition differs slightly from that of Atserias and Müller [7]; the differences are discussed below.

Variables. The variables of $\text{Ref}(F)$ come partitioned into n^3 blocks B_1, \dots, B_{n^3} . The intention is for a block of variables to *encode* or *represent* a single clause in a purported Resolution refutation of F of length at most n^3 . More precisely, each block B_i contains the following variables.

- *Literal set.* There are $2n$ many indicator variables y_ℓ for the literals $\ell \in \{x_1, \bar{x}_1, \dots, x_n, \bar{x}_n\}$ of F . A boolean assignment to the y_ℓ variables is intended to define the set of literals for the clause represented by B_i . As a minor detail (relevant in Section 11), we interpret $y_\ell = 0$ to mean that literal ℓ is included in the block.
- *Block type.* There are two boolean variables $\tau = (\tau_1, \tau_2) \in \{0, 1\}^2$ encoding the block’s *type* as one of three options: *axiom* ($\tau = 00$), *derived* ($\tau = 01$), or *disabled* ($\tau_1 = 1$). We say a block is *enabled* if its type is *axiom* or *derived*. Accordingly, one of the following groups of variables becomes relevant.
 - (1) *Axiom.* There are $\log m$ many variables that encode an *axiom index* $j \in [m]$. The intention is for an axiom block B_i to be a weakening of the j -th axiom of F .
 - (2) *Derived.* There are $O(\log n)$ many variables that encode a pair of *child pointers* $(j, j') \in [n^3] \times [n^3]$ and a *resolved-variable index* $k \in [n]$. The intention is for a derived block B_i to be obtained from B_j and $B_{j'}$ by first resolving on variable x_k and then weakening.
 - (3) *Disabled.* There are no additional relevant variables.

Axioms. It is now straightforward to write down a list of axioms expressing that a truth assignment to the above variables encodes a valid dag-like Resolution refutation of F . Namely, consider a list of constraints defined as follows, where each constraint involves $O(\log n)$ variables.

- *Root.* We require that the last block B_{n^3} (root of the dag) is enabled and that it represents the empty clause, that is, all literal indicator variables are set to 1. (This defines a list of $2n + 1$ constraints, each involving at most two variables.)
- *Derived.* For every derived block B_i with an associated triple $(j, j', k) \in [n^3] \times [n^3] \times [n]$ we require that $j, j' < i$; and that B_j (resp. $B_{j'}$) is enabled and contains literal x_k (resp. \bar{x}_k); and that every other literal in B_j (except x_k) or $B_{j'}$ (except \bar{x}_k) also appears in B_i .
- *Axiom.* For every axiom block B_i with an associated axiom index $j \in [m]$ we require that every literal appearing in the j -th axiom of F also appears in B_i .
- *Disabled.* We impose no constraints on disabled blocks.

Each of these constraints can be written as a CNF formula over $O(\log n)$ variables. While there are many ways of writing a given constraint in CNF, any choice of encoding will do. (In fact, any two encodings of a $O(\log n)$ -variate constraint can be proved equivalent by Resolution, or any other of the proof systems we are interested in, in size $\exp(O(\log n)) = \text{poly}(n)$.) We define $\text{Ref}(F)$ as the conjunction over all these constraints. In conclusion, $\text{Ref}(F)$ is an $O(\log n)$ -CNF formula with $\text{poly}(n)$ clauses of block-width ≤ 3 (the worst case is a clause involving a derived block and both children).

Comparison with Atserias–Müller. Our definition of $\text{Ref}(F)$ differs from that of Atserias and Müller in two ways. Firstly, we encode all pointers (and indices) in *binary* instead of *unary*. For the lifting-based proof, this difference is inconsequential and done for convenience as it yields a formula of low width $w(\text{Ref}(F)) \leq O(\log n)$, which is nice to work with. For the tradeoff-based proof, in contrast, binary encoding is crucial in order for the lower bound on $w_R(\text{Ref}(F))$ to be large enough (in case F is unsatisfiable), so that the lower bound on size in terms of width can be applied.

Secondly, we allow a block to be *disabled*, whereas Atserias and Müller only introduced this option in the *relativized* version of $\text{Ref}(F)$. In our simplified proof for Resolution, this difference is inconsequential: even if we cannot explicitly disable a block by setting its type to *disabled*, we can “manually” achieve the same effect by making the block represent an isolated axiom clause in the refutation dag. More interestingly, the option to disable blocks will be needed in extending our proof to the algebraic setting.

3.2 TreeRef(F) Formula

Next, we define a tree-like version of $\text{Ref}(F)$ that states “ F admits a short tree-like Resolution refutation where the weakening rule is only applied on axiom clauses.” Indeed, $\text{TreeRef}(F)$ is obtained by starting from $\text{Ref}(F)$ and adding some new variables and axioms.

- *New variables.* We add to each block $O(\log n)$ many new variables that encode a *parent pointer* $p \in [n^3]$. The intention is for p to point to the unique parent in a tree-like refutation.
- *New axioms (tree-likeness).* For a derived block B_i , we require that both of its children have their parent pointers set to i . In the other direction, for a non-root enabled block B_i , we require that its parent B_p is an enabled derived block having B_i as one of its children.
- *New axioms (no weakening).* For a derived block B_i , we require that every literal in B_i appears in *both* of its children. This new axiom implies (together with the old axioms) that if a derived block B_i (obtained by resolving on x_k) has literal set C , then its children have sets $\{x_k\} \cup C$ and $\{\bar{x}_k\} \cup C$. (Note that we still allow an axiom block to be a weakening of an axiom of F .)

3.3 rPHP Formula

Finally, we formulate the *retraction weak pigeonhole principle* rPHP_n [40, 60]. This variant features $2n$ pigeons and n holes. It uses a *binary encoding* of the pigeon-mapping, and provides an efficient way to *invert* the mapping. Specifically, the variables of rPHP_n describe two functions, $f: [2n] \rightarrow [n]$ and $g: [n] \rightarrow [2n]$, encoded as follows.

- *Pigeon map.* For every pigeon $i \in [2n]$ there are variables f_{ik} , $k \in [n]$. These variables encode in binary a hole $f(i) \in [n]$ that is expected to house pigeon i .
- *Hole map.* For every hole $j \in [n]$ there are variables $g_{j\ell}$, $\ell \in [2n]$. These variables encode in binary a pigeon $g(j) \in [2n]$ that is expected to occupy hole j .

The axioms of rPHP_n state that for every $i \in [2n]$ and $j \in [n]$,

$$f(i) = j \implies g(j) = i. \quad (1)$$

In other words, g is a left-inverse of f (meaning $g(f(i)) = i$). Note that we do not require g to be a right-inverse (meaning $f(g(j)) = j$), that is, the mapping f need not be surjective. In conclusion, rPHP_n can be written as a $O(\log n)$ -width CNF formula in the variables $(f, g) = (f_{ik}, g_{j\ell})$.

4 DECISION TREE REDUCTIONS

In this section, we define *decision tree reductions*, which will be used in Section 5 to prove a lower bound on the block-width $\text{bw}_R(\text{Ref}(F))$ of refuting the formula $\text{Ref}(F)$ in Resolution. We assume the reader is familiar with the standard notion of a decision tree computing a boolean function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ (see, e.g., the textbook [41, §14]). In particular, a depth- d decision tree \mathcal{T} computing f naturally gives rise to both a d -DNF and a d -CNF representation for f . Namely, the associated d -DNF is given by $\bigvee_{\ell} C_{\ell}$ where ℓ ranges over the leaves of \mathcal{T} that output 1, and C_{ℓ} is the conjunction of literals (query outcomes) on the path from root to leaf ℓ . The d -CNF is obtained by negating the d -DNF associated with the negated decision tree $\neg\mathcal{T}$ (that is, \mathcal{T} but with its output values flipped) computing $\neg f$.

4.1 What is a Reduction?

A decision tree reduction between formulas F and G is a reduction relating the variables of G to the variables of F via shallow decision trees, and moreover, showing that the axioms of F imply those of G . We formalize this as described next.

Definition 4.1 (Reduction). Let $F(x)$ and $G(y)$ be CNF formulas over variables $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_m)$. A *depth- d reduction*, denoted $F \leq_d^{\text{dt}} G$, consists of the following.

- **Variables.** The reduction is defined by a function $f: \{0, 1\}^n \rightarrow \{0, 1\}^m$ such that each output bit $f_i: \{0, 1\}^n \rightarrow \{0, 1\}$ (thought of as the value given to y_i) for $i \in [m]$ is computed by a depth- d decision tree.
- **Axioms.** Let $C(y)$ be a clause and view it as a function $C: \{0, 1\}^m \rightarrow \{0, 1\}$. Consider the composed function $C \circ f$. It can be computed by a depth- $(d \cdot |C|)$ decision tree, and hence we may naturally write it as a $(d \cdot |C|)$ -CNF formula. We require that for every axiom $C \in G$, every clause of $C \circ f$ is a weakening of an axiom of F .

The key property of a reduction is that it translates width complexity bounds.

Lemma 4.2. *If $F \leq_d^{\text{dt}} G$, then $\text{w}_R(F) \leq d \cdot \text{w}_R(G)$.*

This lemma is most elegantly proven using the standard *game semantics* (or *top-down*) characterization of $\text{w}_R(F)$ [5, 58]. We recall the details of this characterization but, due to space limitation, omit the proof of the lemma.

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

- (1) *Query a variable.* Prover chooses an $i \in [n]$ with $\rho_i = *$, after which Adversary chooses $b \in \{0, 1\}$. The state is updated by $\rho_i \leftarrow b$.
- (2) *Forget variables.* Prover chooses a (possibly empty) subset $I \subseteq [n]$. The state is updated by $\rho_i \leftarrow *$ for all $i \in I$.

An important detail is that if Prover queries the i -th variable, forgets it, and then queries it again, Adversary is free to respond with any value regardless of the answer given previously. The game ends when ρ falsifies an axiom of F . The width complexity $\text{w}_R(F)$ of F is characterized by the least w such that there is a Prover strategy of *width* w (maximum number of non- $*$ coordinates in the game state at the end of a round) to end the game no matter how Adversary plays.

4.2 Block-Aware Reductions

We also introduce a more fine-grained type of reduction, suitable for studying block-width.

Definition 4.3 (Block-aware reduction). Let $F(x) \leq_d^{\text{dt}} G(y)$ via $f: \{0, 1\}^n \rightarrow \{0, 1\}^m$ as in Definition 4.1. Suppose further that the variables $y = (y_1, \dots, y_m)$ of G are partitioned into blocks. We say that the reduction $F \leq_d^{\text{dt}} G$ is *block-aware* if for each block $B \subseteq [m]$ there is a depth- d decision tree that computes all the values $f_B(x) := (f_i(x) : i \in B) \in \{0, 1\}^B$ simultaneously.

Lemma 4.4. *If $F \leq_d^{\text{dt}} G$ via a block-aware reduction, then $\text{w}_R(F) \leq d \cdot \text{bw}_R(G)$.*

This lemma can also be proven by considering Prover–Adversary games, since these can equally well characterize block-width (defined naturally for a game state as the number of blocks that the state records values from).

5 BLOCK-WIDTH LOWER BOUND FOR $\text{Ref}(F)$

We have now collected the tools we need to prove Lemma 2.2 stating that $\text{bw}_R(\text{Ref}(F)) \geq \tilde{\Omega}(\text{w}_R(\text{rPHP}_{n^2})/n)$ holds, where F is any unsatisfiable n -variate CNF formula, and $\text{Ref}(F)$ and rPHP_m are as defined in Sections 3.1 and 3.3, respectively. Our goal is to describe a block-aware reduction

$$\text{rPHP}_{n^2} \leq_{\tilde{O}(n)}^{\text{dt}} \text{Ref}(F). \quad (2)$$

This reduction, together with Lemma 4.4, would complete the proof of Lemma 2.2.

5.1 Overview of Reduction

As in the original proof of Atserias and Müller [7], our reduction is guided by the *full tree-like* Resolution refutation \mathcal{T} of the unsatisfiable formula F . More specifically, \mathcal{T} viewed as a refutation dag is a binary tree of height n , it has the empty clause at its root, and at depth $i \in [n]$ the i -th variable is resolved. Thus, \mathcal{T} has 2^n leaves corresponding to all possible width- n clauses; each such leaf clause is a weakening of some axiom of F .

For any truth assignment to rPHP_{n^2} , our reduction is going to produce an assignment to $\text{Ref}(F)$ that represents a purported refutation of F isomorphic to a *subtree* \mathcal{T}' of the full tree \mathcal{T} . We refer to the set of nodes of \mathcal{T}' that have smaller degree in \mathcal{T}' than in \mathcal{T} as the *boundary* of the embedding $\mathcal{T}' \subseteq \mathcal{T}$. We note




Figure 1: Reduction from rPHP_{n^2} to $\text{Ref}(F)$. An assignment to the variables of rPHP_{n^2} defines a *partial matching* $h: [2n^2] \rightarrow [n^2]$ (drawn in blue). Using query access to h we construct an assignment to the variables of $\text{Ref}(F)$ that describes a purported refutation of F . The refutation consists of some n^3 blocks arranged in $n + 1$ layers. Each block has a type: either *derived* (yellow), *axiom* (purple), or *disabled* (gray). In the refutation dag (as defined in Section 2.1), we draw directed edges from children to parent (this is the reverse direction of the child pointers). The top-most $2 \log n$ layers are hardcoded with a tree topology, and between any two remaining layers we insert the partial matching h . The literal set (and other local structure) for each block is computed by locating its natural embedding in the full tree-like refutation \mathcal{T} .

that \mathcal{T}' will not be a valid refutation of F , because the nodes on the boundary are missing at least one child. However, the interior “local neighborhoods” of \mathcal{T}' will be indistinguishable from the corresponding neighborhoods of \mathcal{T} , and those parts do not violate any axioms of $\text{Ref}(F)$. The only axiom violations of $\text{Ref}(F)$ result from the boundary nodes.

We now describe the reduction in detail, relying heavily on the illustration in Figure 1.

5.2 Construction

We start by defining how the variables of $\text{Ref}(F)$ depend on the variables of rPHP_{n^2} . We think of the blocks of $\text{Ref}(F)$ as being arranged in $n + 1$ layers with layer $\ell \in \{0, 1, \dots, n\}$ containing

$\min\{2^\ell, n^2\}$ many blocks; see Figure 1. The top-most layer $\ell = 0$ contains just the root block B_{n^3} . The remaining layers host blocks in an arbitrary but fixed way that respects the block ordering: If block B_i is on a lower layer than block B_j , then $i < j$. A small detail is that so far we have not quite used up all the available n^3 blocks. Indeed, any such leftover blocks we define as *disabled*. From now on, we ignore them and do not draw them in Figure 1.

We proceed to define the child pointers—which determine the topology of the purported refutation—and then the literal sets (and other local structure).

Pointers. The pointers for the top-most $2 \log n$ layers we assign so as to build a full binary tree (which in particular matches the topology of \mathcal{T} on these top-most layers). We say this part of the pointer assignment is *hardcoded*, as it does not depend on the variables of rPHP_{n^2} .

Defining the topology for the remaining non-hardcoded layers is the crux of our reduction. Intuitively, we will *copy-and-paste* the pigeon-mapping described by the variables f_{ik} and $g_{j\ell}$ of rPHP_{n^2} (encoding the functions $f: [2n] \rightarrow [n]$ and $g: [n] \rightarrow [2n]$) between any two consecutive non-hardcoded layers. This results in several copies of the pigeon-mapping being used in defining the topology.

We first define a partial matching (partial injection) $h: [2n^2] \rightarrow [n^2] \cup \{*\}$ by

$$h(i) := \begin{cases} f(i) & \text{if } g(f(i)) = i, \\ * & \text{otherwise.} \end{cases} \quad (3)$$

Given a pigeon $i \in [2n^2]$, we can evaluate $h(i)$ by making $O(\log n)$ queries to the boolean variables defining f and g . Moreover, h is easy to invert with query access to f and g . Note that if $h(i) = *$, meaning $f(i) = j$ but $g(j) \neq i$, then this witnesses an axiom violation for rPHP_{n^2} associated with the pair (i, j) as per Equation (1). At the top of Figure 1, we illustrate one possible partial matching resulting from a particular assignment to rPHP_{n^2} .

Consider a layer $\ell \in \{2 \log n, \dots, n - 1\}$ that contains n^2 blocks. We think of the child pointers originating from layer ℓ as the $2n^2$ pigeons (each of the n^2 blocks names two children), and the blocks on the next layer $\ell + 1$ as the n^2 holes. More precisely, we define the left (resp. right) child of the i -th block on layer ℓ as the $h(2i - 1)$ -th (resp. $h(2i)$ -th) block on layer $\ell + 1$. If ever $h(i)$ is undefined (meaning an axiom of rPHP_{n^2} associated with i is violated), we define the corresponding pointer as *null* (say, by pointing to the root B_{n^3} , which results in an axiom violation for $\text{Ref}(F)$).

This completes the definition of the topology of the purported refutation described by the variables of $\text{Ref}(F)$. Note that the resulting topology (where we ignore null pointers) is a forest of binary trees: it is constructed by stitching together a binary tree at the top with a layered sequence of partial matchings where we have identified pairs of pigeons (each block couples two pigeons). The lower part of Figure 1 shows how the rPHP_{n^2} assignment at the top defines the structure of the refutation claimed to exist by $\text{Ref}(F)$.

Literal Sets. Recall that our goal is to make the purported proof isomorphic to a subtree $\mathcal{T}' \subseteq \mathcal{T}$ (plus some disabled blocks). But now that we have already defined the topology of our purported proof, the definitions of the literal sets (and other local structure) are already determined. To see this, consider the following algorithm

(implementable by a moderate-depth decision tree) for computing the literal set for a block B : Starting from B , walk up to its unique parent in the binary forest (this can be done with $O(\log n)$ queries by computing the inverse of h) and continue taking such upward steps until we reach a block without a parent. We have two cases depending on whether the walk terminates at the root block B_{n^3} .

(1) *Root is reached.* Consider the (reverse) path p (sequence of left/right turns) from B_{n^3} to B . This identifies a node v in the full tree \mathcal{T} , namely, the node obtained by following the path p starting at the root of \mathcal{T} . We simply copy all the local structure at v into B : We make the literal set of B equal that of v . If v is derived in \mathcal{T} by resolving the k -th variable, we make B a derived block and set its resolved-variable index to k . If v is a leaf of \mathcal{T} , that is, a weakening of some, say j -th, axiom of F , then we make B an axiom block and set its axiom index to j .

(2) *Root is not reached.* In this case we make B a *disabled* block.

This completes the definition of how the variables of $\text{Ref}(F)$ depend on the variables of rPHP_{n^2} . We finally note that the whole contents of a particular block can be computed by a single decision tree of depth $\tilde{O}(n)$. Indeed, the most expensive part is to perform the walk up the binary forest, which involves at most n (the depth of the purported proof) evaluations of the inverse of h .

5.3 Correctness

It remains to show that every axiom in the composed version of $\text{Ref}(F)$ is implied by some axiom of rPHP_{n^2} . We argue the contrapositive: any axiom violation for $\text{Ref}(F)$ implies an axiom violation for rPHP_{n^2} . Since our reduction, by construction, always produces a purported refutation isomorphic to a subtree $\mathcal{T}' \subseteq \mathcal{T}$ (plus some disabled blocks which do not violate axioms of $\text{Ref}(F)$), the only possible axiom violations are caused by a block on layer $\ell \in \{2 \log n, \dots, n-1\}$ containing a null pointer. Any null pointer is caused by the decision tree querying a pigeon i with $h(i) = *$. But this means the decision tree has witnessed a violation of (1), that is, an axiom violation for rPHP_{n^2} , by the discussion following (3). This completes the reduction (2).

5.4 Tree-Like Extension

To conclude this section, we observe for later use (in Section 9) that the reduction described above can be easily extended to a block-aware reduction

$$\text{rPHP}_{n^2} \stackrel{\text{dt}}{\leq_{\tilde{O}(n)}} \text{TreeRef}(F). \quad (4)$$

In order to do so, we simply define the *parent pointers* (which are the “new” variables) as the inverses (given by g outside the hardcoded region) of the child pointers defined by the original reduction. To see that the axioms of rPHP_{n^2} imply those of $\text{TreeRef}(F)$, we argue similarly as in Section 5.3: Since \mathcal{T} is a tree-like refutation that uses no weakening (except for the axioms), the output of our reduction (subtree \mathcal{T}' of \mathcal{T}) still has its axiom violations only at the boundary of the embedding $\mathcal{T}' \subseteq \mathcal{T}$.

6 LIFTING BLOCK-WIDTH TO SIZE

In this section, we prove Lemma 2.3, saying that for the lifted version $\text{Lift}(F)$ of a CNF formula F it holds that $2^{\Omega(\text{bw}_R(F))} \leq \text{R}(\text{Lift}(F)) \leq$

$2^{O(\text{bw}(\mathcal{P}))} \|\mathcal{P}\|$, where \mathcal{P} is any Resolution refutation of F . We start by describing how the formula $\text{Lift}(F)$ is constructed.

6.1 Lift(F) Formula

Fix a CNF formula F whose variables x_1, \dots, x_n are partitioned into m blocks. To construct the block-lifted formula $\text{Lift}(F)$, we replace each variable by a copy of a carefully chosen gadget, where gadgets corresponding to the same block partially share variables. Namely, we consider the 3-bit gadget $g: \{0, 1\}^3 \rightarrow \{0, 1\}$ defined by $g(x^0, x^1, s) := x^s$. Note that g is computed by a depth-2 decision tree. We now define $\text{Lift}(F)$ formally:

- *Variables.* For every variable x_i of F , the lifted formula will have two variables x_i^0 and x_i^1 . Moreover, for every block B of F , we introduce a *selector* variable s_B . Thus, altogether, $\text{Lift}(F)$ has $2n + m$ variables, called *lifted variables*.
- *Axioms.* Let $C \in F$ be a clause and view it as a function $C: \{0, 1\}^n \rightarrow \{0, 1\}$. We define a lifted constraint

$$\text{Lift}(C): \{0, 1\}^{2n+m} \rightarrow \{0, 1\}$$

over the lifted variables as the composition

$$\text{Lift}(C) := C(g(x_1^0, x_1^1, s_{B(x_1)}), \dots, g(x_n^0, x_n^1, s_{B(x_n)})),$$

where $B(x_i)$ denotes the unique block containing x_i . Note that $\text{Lift}(C)$ can be computed by composing a depth- $|C|$ decision tree for C with depth-2 decision trees for the gadgets. This results in a decision tree whose depth is only $d := |C| + \text{bw}(C)$ as the gadgets share selector variables. Hence we may write $\text{Lift}(C)$ naturally as a d -CNF formula (as discussed in Section 4). Finally, we define $\text{Lift}(F) := \bigwedge_{C \in F} \text{Lift}(C)$.

6.2 Upper Bound for Lift(F)

Let us prove the upper bound $\text{R}(\text{Lift}(F)) \leq 2^{O(\text{bw}(\mathcal{P}))} \|\mathcal{P}\|$. We again use the language of Prover–Adversary games from Section 4.1. Besides width, such games can also capture the refutation *size* [58]. Namely, size is characterized by *strategy size*: the total number of states that can ever arise in play (over any number of runs of the game). Thus, let \mathcal{P} be a Prover strategy for F of size $\|\mathcal{P}\|$ and block-width $\text{bw}(\mathcal{P})$. Our goal is to find a small-size strategy \mathcal{L} for $\text{Lift}(F)$.

We start by observing that $\text{Lift}(F) \stackrel{\text{dt}}{\leq} F$ via $f = (f_1, \dots, f_n)$ given by $f_i := g(x_i^0, x_i^1, s_{B(x_i)})$. The strategy \mathcal{L} is then constructed by simulating \mathcal{P} as in the proof of Lemma 4.2. We proceed to bound $\|\mathcal{L}\|$ by analyzing the simulation carefully. At the start of a simulation round, if \mathcal{P} is in state ρ , then \mathcal{L} is in one of $2^{\text{bw}(\rho)}$ many corresponding states; here the blow-up $2^{\text{bw}(\rho)}$ comes from having to record the values of $\text{bw}(\rho)$ many selector variables. During a simulation step, \mathcal{L} might have to evaluate an f_i , which gives rise to $O(1)$ intermediate states before the start of the next round. We conclude that there is a factor $O(2^{\text{bw}(\rho)})$ overhead in a single round of the simulation. Altogether, we get $\|\mathcal{L}\| \leq O(2^{\text{bw}(\mathcal{P}))} \|\mathcal{P}\|$, which proves the upper bound.

6.3 Lower Bound for Lift(F)

Finally, we establish the lower bound $2^{\Omega(\text{bw}_R(F))} \leq \text{R}(\text{Lift}(F))$. We show an equivalent claim, namely that $\text{bw}_R(F) \leq O(\log \|\mathcal{P}\|)$ holds

for any refutation \mathcal{P} of $\text{Lift}(F)$. Fix such a \mathcal{P} henceforth. We will proceed by a standard argument using random restrictions.

Recall that for a partial truth value assignment ρ and a clause C , the restricted clause $C \upharpoonright_\rho$ is defined to be the trivially true clause 1 if ρ satisfies some literal in C and otherwise the clause C with all literals falsified by ρ removed. This definition extends to sets/sequences of clauses \mathcal{A} in the natural way by restricting all clauses in \mathcal{A} , removing those which are satisfied. Given a Resolution refutation \mathcal{F} of a CNF formula F , it is a well-known fact that for any partial assignment ρ it holds that $\mathcal{F} \upharpoonright_\rho$ is a resolution refutation of the restricted formula $F \upharpoonright_\rho$ in at most the same size and width.

We start by defining a random restriction ρ to a subset of the variables of $\text{Lift}(F)$ in two steps:

- (1) Let ρ_1 be a random restriction setting each selector variable s_B to a uniform random bit.
- (2) Define X_{ρ_1} as the set of variables that contains, for every variable x_i of F , the variable x_i^{1-s} where $s := s_B(x_i)$ is determined by ρ_1 . Let ρ_2 be a random restriction setting each variable in X_{ρ_1} to a uniform random bit. Let ρ be the concatenation of ρ_1 and ρ_2 .

Note that variables from different blocks are assigned independently. Moreover, each literal evaluates to true with probability at least $1/4$. Thus, the probability that a clause of block-width at least w is not satisfied by ρ is at most $(3/4)^w$. Consider the restricted refutation $\mathcal{P} \upharpoonright_\rho$. By a union bound, we see that

$$\Pr[\mathcal{P} \upharpoonright_\rho \text{ has a clause of block-width } \geq w] \leq \|\mathcal{P}\| \cdot (3/4)^w.$$

For $w := 3 \log \|\mathcal{P}\| > \log \|\mathcal{P}\| / \log(4/3)$ this probability is < 1 , and hence there exists some fixed ρ such that $\text{bw}(\mathcal{P} \upharpoonright_\rho) \leq w$. But $\mathcal{P} \upharpoonright_\rho$ is a refutation of the formula $\text{Lift}(F) \upharpoonright_\rho$, which is easily seen to be the same as F after renaming variables. Hence, $\text{bw}_R(F) \leq w = O(\log \|\mathcal{P}\|)$, which completes the proof of [Lemma 2.3](#).

6.4 Alternative Proof via Tradeoffs

We note that an alternative way of proving a Resolution size lower bound for $\text{Ref}(F)$ when F is unsatisfiable via the standard size-width tradeoff in [\[14\]](#) (without the need for gadget composition). We refer the reader to the full version of the paper for details.

7 RESOLUTION UPPER BOUND FOR $\text{Ref}(F)$

In order to establish [Lemma 2.1\(i\)](#), we must prove that if F is satisfiable, then $\text{Ref}(F)$ admits a polynomial-size Resolution refutation. As this was already shown by Atserias and Müller [\[7\]](#), our proof is modelled on theirs and only requires some minor changes due to the difference in formula encoding. We refer the reader to the full version of this paper for a formal proof.

8 ALGEBRAIC DEFINITIONS

In this section, we define: ([§8.1](#)) the algebraic proof systems Nullstellensatz (NS), Polynomial Calculus (PC), and Sherali-Adams (SA); and ([§8.2](#)) algebraic reductions.

8.1 Algebraic Proof Systems

All the algebraic proof systems are going to share the following basic setup. We work over the polynomial ring $\mathbb{F}[X]$ where \mathbb{F} is a fixed field and $X := \{x_1, x_2, \dots, x_n\}$ is a set of formal variables. We

define the *size* $\|p\|$ of a polynomial $p \in \mathbb{F}[X]$ as the number of its non-zero monomials (when expanded out as a linear combination of monomials). If the variables X are partitioned into blocks, we define the *block-degree* $\text{bdeg}(r)$ of a monomial r as the number of distinct blocks that r touches, and the block-degree of a polynomial as the largest block-degree of any of its monomials.

For a CNF formula F over variables X we use the standard translation of F into a set of polynomial equations F^* defined as follows. First, for each x_i we include in F^* the *boolean axiom* $x_i^2 - x_i = 0$ (enforcing $x_i \in \{0, 1\}$). Second, for each clause $\bigvee_{i \in I} x_i \vee \bigvee_{j \in J} \bar{x}_j$ of F we include in F^* the equation

$$\prod_{i \in I} (1 - x_i) \prod_{j \in J} x_j = 0. \quad (5)$$

This way, F and F^* have the same set of satisfying assignments. Henceforth, we will sometimes identify F and F^* . We are now ready to define our algebraic proof systems.

Nullstellensatz (NS). Nullstellensatz is a static algebraic proof system based on Hilbert's Nullstellensatz. An *NS-proof* of $f = 0$ from a set of polynomial equations $F = \{f_1 = 0, \dots, f_m = 0\}$ is a set of polynomials $\mathcal{P} = \{p_1, \dots, p_m\}$ such that, as formal polynomials,

$$\sum_{i \in [m]} p_i f_i = f.$$

The *size* of the proof is $\|\mathcal{P}\| := \sum_{i \in [m]} \|p_i\| \|f_i\|$, its *degree* is $\text{deg}(\mathcal{P}) := \max_{i \in [m]} (\text{deg}(p_i) + \text{deg}(f_i))$ and, if the variables X are partitioned into blocks, its *block-degree* is

$$\text{bdeg}(\mathcal{P}) := \max_{i \in [m]} (\text{bdeg}(p_i) + \text{bdeg}(f_i)).$$

An *NS-refutation* of F is an NS-proof of $1 = 0$ from F .

Polynomial Calculus (PC). Polynomial Calculus is a dynamic extension of Nullstellensatz. A *PC-proof* of $f = 0$ from a set of polynomial equations $F = \{f_1 = 0, \dots, f_m = 0\}$ is a sequence of polynomials $\mathcal{P} = (p_1, \dots, p_s)$ such that $p_s = f$ and for each $i \in [s]$ either (i) $p_i \in F$ or (ii) p_i is derived from polynomials earlier in the sequence using one of the following rules:

- *Linear combination*: From p_j and $p_{j'}$ derive $\alpha p_j + \beta p_{j'}$ for any $\alpha, \beta \in \mathbb{F}$.
- *Multiplication*: From p_j derive $x p_j$ for any $x \in X$.

The *size* of the proof is $\|\mathcal{P}\| := \sum_{i \in [s]} \|p_i\|$, its *degree* is $\text{deg}(\mathcal{P}) := \max_{i \in [s]} \text{deg}(p_i)$ and, if the variables X are partitioned into blocks, its *block-degree* is $\text{bdeg}(\mathcal{P}) := \max_{i \in [s]} \text{bdeg}(p_i)$. A *PC-refutation* of F is a PC-proof of $1 = 0$ from F .

Sherali-Adams (SA). Sherali-Adams is a static, (semi-)algebraic proof system that is based on the Sherali-Adams hierarchy of LP relaxations. The system is only defined over real numbers, so in this case we fix $\mathbb{F} = \mathbb{R}$. An *SA-proof* of the inequality $f \geq 0$ from a set of polynomial equations $F = \{f_1 = 0, \dots, f_m = 0\}$ is a set of polynomials $\mathcal{P} = \{p_1, \dots, p_m, q\}$ such that

$$\sum_{i \in [m]} p_i f_i + q = f,$$

and where q is a *conical junta*, that is, of the form

$$q = \sum_{I, J} \alpha_{I, J} \prod_{i \in I} x_i \prod_{j \in J} (1 - x_j)$$

for $\alpha_{I,J} \geq 0$ non-negative reals. The *size* of the proof is $\|\mathcal{P}\| := \|q\| + \sum_{i \in [m]} \|p_i\| \|f_i\|$, its *degree* is $\deg(\mathcal{P}) := \max\{\deg(q), \deg(p_i) + \deg(f_i) : i \in [m]\}$ and, if the variables X are partitioned into blocks, its *block-degree* is $\text{bdeg}(\mathcal{P}) := \max\{\text{bdeg}(q), \text{bdeg}(p_i) + \text{bdeg}(f_i) : i \in [m]\}$. An *SA-refutation* of F is an SA-proof of $-1 \geq 0$ from F .

Complexity Measures. We define complexity measures uniformly across $S = \text{NS}, \text{PC}, \text{SA}$.

- The *size complexity* $S(F)$ of a formula F is the minimum size of an S -refutation of F .
- The *degree complexity* $\text{deg}_S(F)$ is the minimum degree of an S -refutation of F .
- The *block degree complexity* $\text{bdeg}_S(F)$ is the minimum block-degree of an S -refutation of F .

Twin Variables. Every algebraic proof systems can be extended using so-called *twin variables*. This means that for every variable $x \in X$ we add another formal variable \bar{x} , and include the *complementary axiom* $x + \bar{x} - 1 = 0$. The translation of CNF formulas to polynomial equations can be made more concise by the use of twin variables. Polynomial Calculus with twin variables is often called *Polynomial Calculus Resolution* (PCR). Using twin variables does not affect the degree complexity in any of the proof systems, but their introduction could potentially reduce size quite drastically. Our main result ([Theorem 1.1](#)) holds in the best of all possible worlds: All upper bounds hold *without* twin variables, and the lower bounds hold *with* twin variables.

Relationships. It is well-known and easy to see that PC (and SA if the field is \mathbb{R}) can efficiently simulate NS. A surprising result of Berkholz [15] is that SoS efficiently simulates PC over \mathbb{R} . In this paper, we need only the easy simulations.

Fact 8.1 (Simulations). *Suppose a polynomial f admits an NS-proof from a set of n -variate polynomials F in size s and (block-)degree d . Then there is a PC-proof (and an SA-proof if the field is \mathbb{R}) of f from F in size $\text{poly}(s, n)$ and (block-)degree d . \square*

Multilinear Polynomials. The *multilinearization* of a polynomial p is defined as the polynomial obtained by replacing all terms in p of the form x^i , $i \geq 2$, with x ; that is, we work modulo the boolean axioms. It will be convenient to assume that all polynomials appearing in our algebraic manipulations are implicitly multilinearized. For example, the product pq of two multilinear polynomials p and q may not itself be multilinear, but pq can be efficiently proven equivalent to its multilinearization by an application of the boolean axioms. It is well known that this implicit multilinearization does not affect the degree complexity of a formula except by a constant factor, and the size complexity can increase at most polynomially. When we work in a multilinear setting we can equate the *syntactic* representation of a polynomial as an element of $\mathbb{F}[X]$ with its *semantic* representation as a boolean function $\{0, 1\}^n \rightarrow \{0, 1\}$, since each boolean function has a unique representation as a multilinear polynomial.

8.2 Algebraic Reductions

We now develop algebraic analogues of the decision tree reductions introduced in [Section 4](#). Notions similar to the next definition have occurred before in, for instance, [21, 45, 46]. As the proofs are straightforward we omit them due to space limitations.

Definition 8.2 (Algebraic reduction). Let F and G be two sets of polynomials encoding CNF formulas over a field \mathbb{F} , defined on variables $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_m)$, respectively. An *algebraic reduction*, denoted $F \leq^{\text{alg}} G$, of degree d consists of the following.

- **Variables.** The reduction is computed by a function

$$r: \{0, 1\}^n \rightarrow \{0, 1\}^m$$

such that each output bit $r_i: \{0, 1\}^n \rightarrow \{0, 1\}$ is computed by a degree- d polynomial.

- **Axioms.** For any $g \in G$, the *multilinearization* of the polynomial $g \circ r$ has an NS-proof from F (over any field) of degree $d \cdot \deg(g)$.

This definition allows us to transform algebraic refutations of G into refutations of F .

Lemma 8.3. *If $F \leq^{\text{alg}} G$ with degree d , then $\text{deg}_S(F) \leq d \cdot \text{deg}_S(G)$ for all $S = \text{NS}, \text{PC}, \text{SA}$.*

Next, we define the algebraic analogue of a *block-aware* reduction. Note that when the algebraic reduction is applied to a monomial $\prod_{i \in I} y_i$ this will produce the polynomial $\prod_{i \in I} r_i$ on the variables $x = (x_1, \dots, x_n)$, and so we will need to control the degree and block-degree of such polynomials.

Definition 8.4 (Algebraic block-aware reduction). Let F and G be two sets of polynomials encoding CNF formulas over a field \mathbb{F} , and suppose that $F \leq^{\text{alg}} G$ by a degree- d reduction $r: \{0, 1\}^n \rightarrow \{0, 1\}^m$ as in the previous definition. Suppose further that the variables of G are partitioned into blocks. The reduction r is a degree- d , size- s *block-aware reduction* if the following two conditions hold:

- **Blocks.** For each block B and each $T \subseteq B$ the polynomial

$$r_T := \text{multilinearization of } \prod_{i \in T} r_i$$

has degree at most d and size at most s .

- **Axioms.** For any $g \in G$, the *multilinearization* of the polynomial $g \circ r$ has an NS-proof from F (over any field) of degree $d \cdot \text{bdeg}(g)$ and size s .

Suppose in addition that the variables of F are also partitioned into blocks. Then the reduction r is *block-preserving* if $\text{bdeg}(r_T) = O(1)$ for every T contained in a block of G and if for every $g \in G$ the NS-proof of $g \circ r$ from F that satisfies the *Axioms* property above has block-degree $O(1)$.

We note that although the definition specifies both the degree and the size of the reduction, often only one of these measures will be relevant and hence mentioned.

Lemma 8.5. *If $F \leq^{\text{alg}} G$ via a degree- d block-aware reduction, then $\text{deg}_S(F) \leq d \cdot \text{bdeg}_S(G)$ for all $S = \text{NS}, \text{PC}, \text{SA}$.*

We end this section with the (intuitively clear) claim that block-aware decision-tree reductions are block-aware algebraic reductions.

Lemma 8.6. *If $F \leq_d^{\text{dt}} G$ via a block-aware reduction then $F \leq^{\text{alg}} G$ via a degree- d block-aware reduction.*

9 BLOCK-DEGREE LOWER BOUND FOR TreeRef(F)

In this section we sketch the proof of [Lemma 2.4\(ii\)](#), which says that $\text{bdeg}_S(\text{TreeRef}(F)) \geq n^{\Omega(1)}$, where F is unsatisfiable and $S = R, NS, PC, SA$. We already know that

$$\text{bdeg}_S(\text{TreeRef}(F)) \geq \tilde{\Omega}(\text{deg}_S(\text{rPHP}_{n^2})/n)$$

by the reduction of [Section 5.4](#) and [Lemmas 8.5](#) and [8.6](#). Hence it suffices to prove

$$\text{deg}_S(\text{rPHP}_n) \geq \tilde{\Omega}(n). \quad (6)$$

This turns out to follow immediately from known degree lower bounds due to Razborov (for PC over any field) [\[62\]](#) and Georgiou and Magen (for SA) [\[32\]](#). These papers use a different algebraic encoding of the pigeonhole principle, but we can prove [\(6\)](#) by a reduction to their lower bound. We omit the details of this reduction due to space limitations.

10 LIFTING BLOCK-DEGREE TO SIZE

In this section, we prove [Lemma 2.5](#) that states that $2^{\Omega(\text{bdeg}_S(F))} \leq S(\text{Lift}(F)) \leq 2^{O(\text{bdeg}(\mathcal{P}))} \|\mathcal{P}\|$ where \mathcal{P} is any S -refutation of F and $S = R, NS, PC, SA$. We use the same definition of the formula $\text{Lift}(F)$ as in [Section 6](#). For Resolution this is exactly [Lemma 2.3](#).

10.1 Upper Bound for Lift(F)

To prove the upper bound $S(\text{Lift}(F)) \leq 2^{O(\text{bdeg}(\mathcal{P}))} \|\mathcal{P}\|$ for the algebraic proof systems, we start by observing that $\text{Lift}(F) \leq^{\text{alg}} F$ via the degree-2 reduction $r = (r_1, \dots, r_n)$ given by $r_i := g(x_i^0, x_i^1, s_{B(x_i)}) = x_i^0(1 - s_{B(x_i)}) + x_i^1 s_{B(x_i)}$. Note that for any polynomial p over the variables of F ,

$$\|p \circ r\| \leq 3^{\text{bdeg}(p)} \cdot \|p\|.$$

We first prove the upper bound for Nullstellensatz by analyzing this reduction (the proof for Sherali–Adams is analogous). Let $F = \{f_1, \dots, f_m\}$ and let $\mathcal{P} = \{p_1, \dots, p_m\}$ be an NS-refutation of F . Recall that $\|\mathcal{P}\| = \sum_{i \in [m]} \|p_i\| \|f_i\|$. Consider the expression

$$\sum_{i \in [m]} (p_i f_i) \circ r = \sum_{i \in [m]} (p_i \circ r)(f_i \circ r) = 1,$$

which, as argued in the proof of [Lemma 8.3](#), is a refutation of $\text{Lift}(F)$. Note that the polynomial $p_i \circ r$ has at most $3^{\text{bdeg}(p_i)} \cdot \|p_i\| \leq 3^{\text{bdeg}(\mathcal{P})} \cdot \|p_i\|$ monomials and that $f_i \circ r$ is equal to the sum of the $2^{\text{bdeg}(f_i)} = O(1)$ axioms of $\text{Lift}(f_i)$, each of which has $3^{\text{bdeg}(f_i)} \|f_i\| = O(\|f_i\|)$ monomials. Therefore, we can conclude there is a NS-refutation of size $\sum_{i \in [m]} 3^{\text{bdeg}(\mathcal{P})} \cdot \|p_i\| \cdot O(\|f_i\|) \leq O(3^{\text{bdeg}(\mathcal{P})} \|\mathcal{P}\|)$.

We now prove the upper bound for PC. Let \mathcal{P} be a PC-refutation of F . We construct a PC-refutation \mathcal{P}' of $\text{Lift}(F)$ in the same way as done in the proof of [Lemma 8.3](#): whenever \mathcal{P} derives p , in \mathcal{P}' we will derive the polynomial $p \circ r$ (which has at most $3^{\text{bdeg}(p)} \|p\|$ monomials).

- *Axioms.* For any axiom $f \in F$, we noted already that the polynomial $f \circ r$ is equal to the sum of the $2^{\text{bdeg}(f)} = O(1)$ axioms of $\text{Lift}(f)$, each of which has $3^{\text{bdeg}(f)} \|f\| = O(\|f\|)$ monomials. Thus, $f \circ r$ can be derived in PC in size $O(\|f\|)$.
- *Linear Combination.* If the polynomial p_3 is derived from p_1 and p_2 using a linear combination, then we derive $p_3 \circ r$ from $p_1 \circ r$ and $p_2 \circ r$ using the same linear combination in \mathcal{P}' .

- *Multiplication.* If $y_i p$ is derived from p by the multiplication rule, then we can derive $(y_i p) \circ r = r_i(p \circ r)$ from $p \circ r$ in size $O(\|p \circ r\|)$.

Therefore, the PC-refutation has size $O(3^{\text{bdeg}(\mathcal{P})} \|\mathcal{P}\|)$.

10.2 Lower Bound for Lift(F)

The proof of the lower bound $2^{\Omega(\text{bdeg}_S(F))} \leq S(\text{Lift}(F))$ for $S = NS, SA, PC$ follows the random restriction argument used for Resolution exactly ([Section 6](#)), so due to space limitations we omit the argument and refer the reader to the full version.

11 ALGEBRAIC UPPER BOUND FOR TreeRef(F)

In this section, we prove [Lemma 2.4\(i\)](#) that states that $\text{TreeRef}(F)$, where F is satisfiable, admits a size- $n^{O(1)}$ block-degree- $O(1)$ S -refutation for $S = R, NS, PC, SA$. We prove this for NS, which implies the result for PC and SA by simulations ([Fact 8.1](#)). The result holds for R by the upper bound for $\text{Ref}(F)$ (in [Section 7](#)) and the fact that $\text{TreeRef}(F)$ was defined as a strengthening of $\text{Ref}(F)$. Therefore, the goal of this section is to prove the following lemma.

Lemma 11.1 (Algebraic upper bound). *Let F be a satisfiable n -variate formula. There is a size- $n^{O(1)}$ block-degree- $O(1)$ NS-refutation of $\text{TreeRef}(F)$ (over any field, without twin variables).*

The proof of this lemma essentially implements the algorithm refuting $\text{Ref}(F)$ as an algebraic reduction to the end-of-line formula EoL_n . We proceed in three steps:

(§11.1) First we define EoL_n , which is a size- $n^{O(1)}$ block-degree- $O(1)$ CNF formula.

(§11.2) Then we reduce $\text{TreeRef}(F)$ to EoL_n .

(§11.3) Finally, we recall from prior work [\[33\]](#) that EoL_n admits a small NS-refutation.

The last two steps are formalized in the following two claims. As we want our result to be as general as possible, our algebraic proofs will be implemented over the integers \mathbb{Z} (hence the computations are valid over any field), and assume no twin variables.

Claim 11.2 (Reduction to EoL). *Fix an n -variate satisfiable F . There is a block-aware, block-preserving algebraic reduction $\text{TreeRef}(F) \leq^{\text{alg}} \text{EoL}_n$ of size $n^{O(1)}$.*

Claim 11.3 (Upper bound for EoL). *EoL_n has a size- $n^{O(1)}$ block-degree- $O(1)$ NS-refutation over \mathbb{Z} .*

The algebraic upper bound ([Lemma 11.1](#)) follows by combining these two lemmas.

11.1 EoL Formula

The end-of-line formula EoL_n states that “there is an n -node digraph where every node has in/out-degree 1, except for one distinguished node that has in-degree 0 and out-degree 1.” The combinatorial principle underlying EoL_n is central in the theory of total NP search problems [\[10, 54\]](#).

The variables of EoL_n are intended to describe a digraph on vertices $[n]$ where $n \in [n]$ is thought of as a distinguished node. Namely, for each $i \in [n]$, there is a block of $2 \log n$ boolean variables $z_i := (\tilde{z}_i, \bar{z}_i)$ that encode, in binary, a predecessor pointer $\tilde{z}_i \in [n]$

and a *successor* pointer $\vec{z}_i \in [n]$. An assignment to the variables $z = (z_1, \dots, z_n)$ defines a digraph $G_z := ([n], E_z)$ where

$$(i, j) \in E_z \quad \text{iff} \quad \vec{z}_i = j \text{ and } \vec{z}_j = i.$$

A small detail is that we allow any node to be a *self-loop*, achieved by setting $\vec{z}_i = \vec{z}_i = i$.

The axioms of EoL_n are:

- *Distinguished.* The node $n \in [n]$ has $\text{indeg}_{G_z}(n) = 0$ and $\text{outdeg}_{G_z}(n) = 1$.
- *Non-distinguished.* Every node $i \in [n - 1]$ has $\text{indeg}_{G_z}(i) = \text{outdeg}_{G_z}(i) = 1$.

In particular, EoL_n can be written as an $O(\log n)$ -CNF formula of block-width 2. The reader familiar with pigeonhole principles can observe that our definition is equivalent to a variant of the bijective pigeonhole principle: EoL_n claims the edges of G_z define a bijection $[n] \rightarrow [n - 1]$.

11.2 Reduction to EoL

Next we prove [Claim 11.2](#): For an n -variate satisfiable F , we give a size- $n^{O(1)}$ block-aware, block-preserving algebraic reduction

$$\text{TreeRef}(F) \leq^{\text{alg}} \text{EoL}_{n^3}.$$

\wedge -Decision Trees. For ease of understanding, we describe the reduction as an \wedge -decision tree, that is, a decision tree that is allowed to query, in a single step, the logical-and $\bigwedge_{x \in A} x$ of any subset A of variables. Note that ordinary “singleton” queries are still supported by choosing A to contain a single variable. Such trees can be converted into polynomials by a standard method.

Fact 11.4. *If r is computed by a depth- d \wedge -decision tree, then r is computed by size- $2^{O(d)}$ polynomial.*

PROOF. For each leaf ℓ in the tree, let $r_\ell(x)$ denote the indicator function that is 1 iff the leaf ℓ is reached on input x . Every query $\bigwedge_{x \in A} x$ can be simulated by the monomial $x^A := \prod_{x \in A} x$. Hence we can compute r_ℓ by taking the product along the unique path from the root to ℓ of either x^A or $1 - x^A$ (depending on the outcome of the query on the path). Hence, as a multilinear polynomial, r_ℓ satisfies $\|r_\ell\| \leq 2^d$. Moreover, r can be written as $r = \sum_\ell r_\ell$ where the sum is over leaves ℓ that output 1. There are at most 2^d leaves, and thus $\|r\| \leq 2^{2d}$. \square

Reduction. We describe a family of \wedge -decision trees

$$\mathcal{T} = (\mathcal{T}_1, \dots, \mathcal{T}_{n^3})$$

where each \mathcal{T}_i has depth $O(\log n)$, queries at most 4 blocks of $\text{TreeRef}(F)$ and outputs values for the variables $z_i = (\vec{z}_i, \bar{z}_i)$. Our goal is to satisfy the following condition, which will imply (as we will argue below) the *Axiom* property of a block-aware reduction.

(\dagger) *For each assignment y to $\text{TreeRef}(F)$, if the output $\mathcal{T}(y)$ violates an axiom of EoL_{n^3} involving node-blocks j and j' , then the execution of $\mathcal{T}_j(y)$ or $\mathcal{T}_{j'}(y)$ has witnessed by its **singleton queries** an axiom violation for $\text{TreeRef}(F)$.*

Henceforth, fix a satisfying assignment x^* of F . Given an assignment y to $\text{TreeRef}(F)$, we say a block B is *feasible* iff the clause encoded by B is falsified by x^* . Note that the feasibility of a given block can be decided by a single \wedge -query (involving n indicator

variables; here we use our convention that literal indicators are set to 1 iff the literal is *not* included in the block). The tree \mathcal{T}_i computes $z_i = (\vec{z}_i, \bar{z}_i)$ as follows. We start by checking whether B_i is feasible:

B_i is not feasible: Two cases depending on whether B_i is root (that is, $i = n^3$).

- *Non-root.* We make node i into a self-loop by outputting $\vec{z}_i = \bar{z}_i := i$.
- *Root.* We know that B_{n^3} contains some literal consistent with x^* . By binary search (using $O(\log n)$ many \wedge -queries) we can discover a specific literal indicator of B_{n^3} that is set to 0. This violates an axiom of $\text{TreeRef}(F)$. Hence by (\dagger), it is safe to output anything for (\vec{z}_i, \bar{z}_i) .

B_i is feasible: Query B_i 's type.

- *Disabled:* If B_i is non-root, we make node i into a self-loop. If B_i is root, then we have found an axiom violation for $\text{TreeRef}(F)$ (and by (\dagger) we can output anything).
- *Axiom:* Here we can find an axiom violation. Query B_i 's axiom index j . Since the j -th axiom of F is satisfied by x^* , it contains some literal ℓ consistent with x^* . But since B_i is feasible, B_i does not contain ℓ . Hence ℓ is a literal in the j -th axiom not in B_i , which is a violation.
- *Derived:* Query B_i 's child pointers (j, j') , the resolved-variable index k , and the parent pointer p . Query whether B_j and $B_{j'}$ are feasible, and query their type and parent pointers. If B_i is non-root, query also the type and child pointers of B_p .

We may assume the variables that are singleton-queried above cause no axiom violations for $\text{TreeRef}(F)$ (as otherwise we are free to output anything). We may also assume we are in the case where exactly one of B_j and $B_{j'}$ is feasible, say B_j (otherwise we may use binary search to find a violation related to a literal indicator), and both have their parent pointers set to i . We also assume that, if B_i is non-root, then it is a child of B_p . We output $(\vec{z}_i, \bar{z}_i) := (p, j)$.

We claim the condition (\dagger) is satisfied: If the decision trees $\mathcal{T}_{i'}$ for $i' = j, j', p$ do not find a violation either, then they will not produce an axiom violation involving node i . Namely, they output $\vec{z}_j := i$ and $\bar{z}_p := i$ (if B_i is non-root) and the node j' will be made a self-loop.

We need to prove that this reduction is an $n^{O(1)}$ -size block-aware, block-preserving algebraic reduction. First, we show that each polynomial r_T generated by the reduction has size $n^{O(1)}$ and block-degree $O(1)$. Since each \wedge -decision tree \mathcal{T}_i has depth $O(\log n)$, by [Fact 11.4](#) each output bit (or even the product polynomial r_T for a subset T of output bits) of \mathcal{T}_i can be converted to a polynomial of size $n^{O(1)}$. Furthermore, since \mathcal{T}_i queries variables from at most 4 blocks of $\text{TreeRef}(F)$ it follows that $\text{bdeg}(r_T) = O(1)$.

It remains to show that for each polynomial g encoding an axiom of EoL_{n^3} , the polynomial $g \circ r$ has an NS proof from the axioms of F in size $n^{O(1)}$ and block-degree $O(1)$. So, suppose that g is associated to node-blocks j, j' in EoL_{n^3} . There is a unique partial assignment α to the variables of g such that $g(\alpha) = 1$; this assignment falsifies the clause of EoL_{n^3} corresponding to g . Since the block-degree of g is at most two we can write $\alpha = \alpha_{T_j} \alpha_{T_{j'}}$, where $\alpha_{T_j}, \alpha_{T_{j'}}$ assign the variables in the two blocks of g . For $i = j, j'$ let \mathcal{L}_i denote the leaves in the tree \mathcal{T}_i that are consistent with the partial assignment α_{T_i} .

We can express

$$g \circ r = \left(\sum_{\ell \in \mathcal{L}_j} r_\ell \right) \left(\sum_{\ell \in \mathcal{L}_{j'}} r_\ell \right)$$

where each polynomial r_ℓ is an indicator function for the corresponding leaf ℓ in each \wedge -decision tree. Since $\mathcal{T}_j, \mathcal{T}_{j'}$ are \wedge -decision trees, if $(g \circ r)(y) = 1$ it follows that there are leaves $\ell_1 \in \mathcal{L}_j, \ell_2 \in \mathcal{L}_{j'}$ with indicators satisfying $r_{\ell_1}(y) = 1, r_{\ell_2}(y) = 1$, and all other leaf indicators in both trees are 0. By (\dagger) , one of these two leaf indicators r_{ℓ_1}, r_{ℓ_2} must witness an axiom violation for $\text{TreeRef}(F)$ in its singleton queries, and thus this leaf indicator is a weakening of an axiom of $\text{TreeRef}(F)$. Ranging over all y such that $(g \circ r)(y) = 1$, this implies that $g \circ r$ can be written as a sum of weakenings of axioms of $\text{TreeRef}(F)$. Since each \wedge -decision tree has depth $O(\log n)$ and queries $O(1)$ blocks from $\text{TreeRef}(F)$ we can prove $g \circ r$ from $\text{TreeRef}(F)$ in NS in size $n^{O(1)}$ and block-degree $O(1)$. This concludes the proof of [Claim 11.2](#).

11.3 Upper Bound for EoL

For the proof of [Claim 11.3](#), we refer the reader to [[33](#), Remark 4.2] or to the full version of this paper.

ACKNOWLEDGEMENTS

We thank Jan Pich, Jan Krajíček, and Pavel Pudlák for comments on an early version of this manuscript. We also thank Shuo Pang, Aaron Potechin, and Madhur Tulsiani for discussions. Finally, we thank the anonymous reviewers for many comments that helped us improve the presentation.

Susanna F. de Rezende was supported by Knut and Alice Wallenberg grant KAW 2018.0371. Jakob Nordström received funding from the Swedish Research Council grant 2016-00782 and the Independent Research Fund Denmark grant 9040-00389B. Toniann Pitassi did this work supported by NSERC and NSF Grant CCF-1900460. Robert Robere was supported by NSERC, the Charles Simonyi Endowment, and indirectly supported by the National Science Foundation Grant No. CCF-1900460. This material is based on work supported by NSERC and NSF Grant CCF-1900460. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- [1] Michael Alekhnovich, Eli Ben-Sasson, Alexander Razborov, and Avi Wigderson. 2002. Space Complexity in Propositional Calculus. *SIAM J. Comput.* 31, 4 (2002), 1184–1211. <https://doi.org/10.1137/S0097539700366735>
- [2] Michael Alekhnovich, Sam Buss, Shlomo Moran, and Toniann Pitassi. 2001. Minimum propositional proof length is NP-hard to linearly approximate. *Journal of Symbolic Logic* 66, 1 (2001), 171–191. <https://doi.org/10.2307/2694916>
- [3] Michael Alekhnovich and Alexander Razborov. 2008. Resolution Is Not Automatizable Unless W[P] Is Tractable. *SIAM J. Comput.* 38, 4 (2008), 1347–1363. <https://doi.org/10.1137/06066850X>
- [4] Albert Atserias and Maria Luisa Bonet. 2004. On the automatizability of resolution and related propositional proof systems. *Information and Computation* 189, 2 (2004), 182–201. <https://doi.org/10.1016/j.ic.2003.10.004>
- [5] Albert Atserias and Victor Dalmau. 2008. A combinatorial characterization of resolution width. *J. Comput. System Sci.* 74, 3 (2008), 323–334. <https://doi.org/10.1016/j.jcss.2007.06.025>
- [6] Albert Atserias and Tuomas Hakoniemi. 2019. Size-Degree Trade-Offs for Sums-of-Squares and Positivstellensatz Proofs. In *Proceedings of the 34th Computational Complexity Conference (CCC)*, Vol. 137. 24:1–24:20. <https://doi.org/10.4230/LIPIcs.CCC.2019.24>
- [7] Albert Atserias and Moritz Müller. 2020. Automating Resolution is NP-Hard. *J. ACM* 67, 5, Article 31 (2020). <https://doi.org/10.1145/3409472> Preliminary version in FOCS '19.
- [8] Boaz Barak, Jonathan Kelner, and David Steurer. 2015. Dictionary Learning and Tensor Decomposition via the Sum-of-Squares Method. In *Proceedings of the 47th Symposium on Theory of Computing (STOC)*, 143–151. <https://doi.org/10.1145/2746539.2746605>
- [9] Roberto J. Bayardo Jr. and Robert Schrag. 1997. Using CSP Look-Back Techniques to Solve Real-World SAT Instances. In *Proceedings of the 14th National Conference on Artificial Intelligence (AAAI)*, 203–208.
- [10] Paul Beame, Stephen Cook, Jeff Edmonds, Russell Impagliazzo, and Toniann Pitassi. 1998. The Relative Complexity of NP Search Problems. *J. Comput. System Sci.* 57, 1 (1998), 3–19. <https://doi.org/10.1006/jcss.1998.1575>
- [11] Paul Beame, Russell Impagliazzo, Jan Krajíček, Toniann Pitassi, and Pavel Pudlák. 1994. Lower bounds on Hilbert's Nullstellensatz and propositional proofs. In *Proceedings of the 35th Symposium on Foundations of Computer Science (FOCS)*, 794–806. <https://doi.org/10.1109/SFCS.1994.365714>
- [12] Paul Beame, Henry Kautz, and Ashish Sabharwal. 2004. Towards Understanding and Harnessing the Potential of Clause Learning. *Journal of Artificial Intelligence Research* 22 (2004), 319–351. <https://doi.org/10.1613/jair.1410>
- [13] Zoë Bell. 2020. *Automating Regular or Ordered Resolution is NP-Hard*. Technical Report TR20-105. Electronic Colloquium on Computational Complexity (ECCC). <https://eccc.weizmann.ac.il/report/2020/105/>
- [14] Eli Ben-Sasson and Avi Wigderson. 2001. Short Proofs Are Narrow—Resolution Made Simple. *J. ACM* 48, 2 (2001), 149–169. <https://doi.org/10.1145/375827.375835>
- [15] Christoph Berkholz. 2018. The Relation between Polynomial Calculus, Sherali-Adams, and Sum-of-Squares Proofs. In *Proceedings of the 35th Symposium on Theoretical Aspects of Computer Science (STACS)*, Vol. 96. 11:1–11:14. <https://doi.org/10.4230/LIPIcs.STACS.2018.11>
- [16] Maria Luisa Bonet, Carlos Domingo, Ricard Gavaldà, Alexis Maciel, and Toniann Pitassi. 2004. Non-Automatizability of Bounded-Depth Frege Proofs. *Computational Complexity* 13, 1-2 (2004), 47–68. <https://doi.org/10.1007/s00037-004-0183-5>
- [17] Maria Luisa Bonet and Nicola Galesi. 2001. Optimality of size-width tradeoffs for resolution. *Computational Complexity* 10, 4 (2001), 261–276. <https://doi.org/10.1007/s000370100000>
- [18] Maria Luisa Bonet, Toniann Pitassi, and Ran Raz. 1997. No Feasible Interpolation for TC^0 -Frege Proofs. In *Proceedings of the 38th Symposium on Foundations of Computer Science (FOCS)*, 254–263. <https://doi.org/10.1109/SFCS.1997.646114>
- [19] Maria Luisa Bonet, Toniann Pitassi, and Ran Raz. 2000. On Interpolation and Automatization for Frege Systems. *SIAM J. Comput.* 29, 6 (2000), 1939–1967. <https://doi.org/10.1137/S0097539798353230>
- [20] Joshua Buresh-Oppenheim, Matthew Clegg, Russell Impagliazzo, and Toniann Pitassi. 2002. Homogenization and the Polynomial Calculus. *Computational Complexity* 11, 3-4 (2002), 91–108. <https://doi.org/10.1007/s00037-002-0171-6> Preliminary version in ICALP '00.
- [21] Sam Buss, Dima Grigoriev, Russell Impagliazzo, and Toniann Pitassi. 2001. Linear Gaps between Degrees for the Polynomial Calculus Modulo Distinct Primes. *J. Comput. System Sci.* 62, 2 (2001), 267–289. <https://doi.org/10.1006/jcss.2000.1726>
- [22] Matthew Clegg, Jeff Edmonds, and Russell Impagliazzo. 1996. Using the Groebner Basis Algorithm to Find Proofs of Unsatisfiability. In *Proceedings of the 28th Symposium on Theory of Computing (STOC)*, 174–183. <https://doi.org/10.1145/237814.237860>
- [23] Stefan Dantchev and Søren Riis. 2003. On Relativisation and Complexity Gap for Resolution-Based Proof Systems. In *Computer Science Logic*. Springer, 142–154. https://doi.org/10.1007/978-3-540-45220-1_14
- [24] Susanna F. de Rezende, Or Meir, Jakob Nordström, Toniann Pitassi, Robert Robere, and Marc Vinyals. 2020. Lifting with Simple Gadgets and Applications to Circuit and Proof Complexity. In *Proceedings of the 61st Symposium on Foundations of Computer Science (FOCS)*, 24–30. <https://doi.org/10.1109/focs46700.2020.00011>
- [25] Susanna F. de Rezende, Jakob Nordström, and Marc Vinyals. 2016. How Limited Interaction Hinders Real Communication (and What It Means for Proof and Circuit Complexity). In *Proceedings of the 57th Symposium on Foundations of Computer Science (FOCS)*, 295–304. <https://doi.org/10.1109/FOCS.2016.40>
- [26] Noah Fleming, Pravesh Kothari, and Toniann Pitassi. 2019. Semialgebraic Proofs and Efficient Algorithm Design. *Foundations and Trends in Theoretical Computer*

- Science 14, 1-2 (2019), 1–221. <https://doi.org/10.1561/04000000086>
- [27] Nicola Galesi and Massimo Lauria. 2010. On the Automatizability of Polynomial Calculus. *Theory of Computing Systems* 47, 2 (2010), 491–506. <https://doi.org/10.1007/s00224-009-9195-5>
- [28] Nicola Galesi and Massimo Lauria. 2010. Optimality of Size-Degree Tradeoffs for Polynomial Calculus. *ACM Transactions on Computational Logic* 12, 1 (2010). <https://doi.org/10.1145/1838552.1838556>
- [29] Ankit Garg, Mika Göös, Prithish Kamath, and Dmitry Sokolov. 2018. Monotone Circuit Lower Bounds from Resolution. In *Proceedings of the 50th Symposium on Theory of Computing (STOC)*. 902–911. <https://doi.org/10.1145/3188745.3188838>
- [30] Michal Garlík. 2019. Resolution Lower Bounds for Refutation Statements. In *Proceedings of the 44th Mathematical Foundations of Computer Science (MFCS)*, Vol. 138. 37:1–37:13. <https://doi.org/10.4230/LIPIcs.MFCS.2019.37>
- [31] Michal Garlík. 2020. *Failure of Feasible Disjunction Property for k -DNF Resolution and NP-hardness of Automating It*. Technical Report. Electronic Colloquium on Computational Complexity (ECCC). <https://eccc.weizmann.ac.il/report/2020/037/>
- [32] Konstantinos Georgiou and Avner Magen. 2008. *Limitations of the Sherali-Adams lift and project system: Compromising local and global arguments*. Technical Report. University of Toronto. <http://www.cs.utoronto.ca/pub/reports/csr/587/CSRG-587.pdf>
- [33] Mika Göös, Prithish Kamath, Robert Robere, and Dmitry Sokolov. 2019. Adventures in Monotone Complexity and TFP. In *Proceedings of the 10th Innovations in Theoretical Computer Science Conference (ITCS)*. 38:1–38:19. <https://doi.org/10.4230/LIPIcs.ITCS.2019.38>
- [34] Mika Göös, Sajin Koroth, Ian Mertz, and Toniann Pitassi. 2020. Automating Cutting Planes is NP-Hard. In *Proceedings of the 52nd Symposium on Theory of Computing (STOC)*. 68–77. <https://doi.org/10.1145/3357713.3384248>
- [35] Mika Göös and Toniann Pitassi. 2018. Communication Lower Bounds via Critical Block Sensitivity. *SIAM J. Comput.* 47, 5 (2018), 1778–1806. <https://doi.org/10.1137/16M1082007> Preliminary version in *STOC '14*.
- [36] Samuel Hopkins, Pravesh Kothari, Aaron Potechin, Prasad Raghavendra, Tselil Schramm, and David Steurer. 2017. The Power of Sum-of-Squares for Detecting Hidden Structures. In *Proceedings of the 58th Symposium on Foundations of Computer Science (FOCS)*. 720–731. <https://doi.org/10.1109/FOCS.2017.72>
- [37] Trinh Huynh and Jakob Nordström. 2012. On the Virtue of Succinct Proofs: Amplifying Communication Complexity Hardness to Time-Space Trade-Offs in Proof Complexity. In *Proceedings of the 44th Symposium on Theory of Computing (STOC)* (New York, New York, USA). 233–248. <https://doi.org/10.1145/2213977.2214000>
- [38] Russell Impagliazzo, Pavel Pudlák, and Jiří Sgall. 1999. Lower Bounds for the Polynomial Calculus and the Gröbner Basis Algorithm. *Computational Complexity* 8, 2 (1999), 127–144. <https://doi.org/10.1007/s000370050024>
- [39] Kazuo Iwama. 1997. Complexity of finding short resolution proofs. In *Mathematical Foundations of Computer Science (MFCS)*. 309–318. <https://doi.org/10.1007/BFb0029974>
- [40] Emil Jeřábek. 2007. On Independence of Variants of the Weak Pigeonhole Principle. *Journal of Logic and Computation* 17, 3 (2007), 587–604. <https://doi.org/10.1093/logcom/exm017>
- [41] Stasys Jukna. 2012. *Boolean Function Complexity: Advances and Frontiers*. Algorithms and Combinatorics, Vol. 27. Springer.
- [42] Pravesh Kothari, Jacob Steinhardt, and David Steurer. 2018. Robust moment estimation and improved clustering via sum of squares. In *Proceedings of the 50th Symposium on Theory of Computing (STOC)*. 1035–1046. <https://doi.org/10.1145/3188745.3188970>
- [43] Jan Krajíček and Pavel Pudlák. 1998. Some Consequences of Cryptographical Conjectures for S_2^1 and EF. *Information and Computation* 140, 1 (1998), 82–94. <https://doi.org/10.1006/inco.1997.2674>
- [44] Jean Lasserre. 2001. An Explicit Exact SDP Relaxation for Nonlinear 0–1 Programs. In *Proceedings of the 8th International Conference on Integer Programming and Combinatorial Optimization (IPCO)*. 293–303. https://doi.org/10.1007/3-540-45535-3_23
- [45] Massimo Lauria and Jakob Nordström. 2017. Graph Colouring is Hard for Algorithms Based on Hilbert’s Nullstellensatz and Gröbner Bases. In *Proceedings of the 32nd Computational Complexity Conference (CCC)*. 2:1–2:20. <https://doi.org/10.4230/LIPIcs.CCC.2017.2>
- [46] Massimo Lauria and Jakob Nordström. 2017. Tight Size-Degree Bounds for Sums-of-Squares Proofs. *Computational Complexity* 26, 4 (2017), 911–948. <https://doi.org/10.1007/s00037-017-0152-4>
- [47] Tengyu Ma, Jonathan Shi, and David Steurer. 2016. Polynomial-Time Tensor Decompositions with Sum-of-Squares. In *Proceedings of the 57th Symposium on Foundations of Computer Science (FOCS)*. 438–446. <https://doi.org/10.1109/FOCS.2016.54>
- [48] João P. Marques-Silva and Karem A. Sakallah. 1999. GRASP: A Search Algorithm for Propositional Satisfiability. *IEEE Trans. Comput.* 48, 5 (May 1999), 506–521. <https://doi.org/10.1109/12.769433> Preliminary version in *ICCAD '96*.
- [49] Ian Mertz, Toniann Pitassi, and Yuanhao Wei. 2019. Short Proofs Are Hard to Find. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming (ICALP)*, Vol. 132. 84:1–84:16. <https://doi.org/10.4230/LIPIcs.ICALP.2019.84>
- [50] Matthew W. Moskewicz, Conor F. Madigan, Ying Zhao, Lintao Zhang, and Sharad Malik. 2001. Chaff: Engineering an Efficient SAT Solver. In *Proceedings of the 38th Design Automation Conference (DAC)*. 530–535. <https://doi.org/10.1145/378239.379017>
- [51] Ryan O’Donnell. 2017. SOS Is Not Obviously Automatizable, Even Approximately. In *Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS)*, Vol. 67. Schloss Dagstuhl, 59:1–59:10. <https://doi.org/10.4230/LIPIcs.ITCS.2017.59>
- [52] Ryan O’Donnell and Tselil Schramm. 2019. Sherali–Adams Strikes Back. In *Proceedings of the 34th Computational Complexity Conference (CCC)*. 8:1–8:30. <https://doi.org/10.4230/LIPIcs.CCC.2019.8>
- [53] Ryan O’Donnell and Yuan Zhou. 2013. Approximability and proof complexity. In *Proceedings of the 24th Symposium on Discrete Algorithms (SODA)*. 1537–1556. <https://doi.org/10.1137/1.9781611973105.111>
- [54] Christos Papadimitriou. 1994. On the Complexity of the Parity Argument and Other Inefficient Proofs of Existence. *J. Comput. System Sci.* 48, 3 (1994), 498–532. [https://doi.org/10.1016/S0022-0000\(05\)80063-7](https://doi.org/10.1016/S0022-0000(05)80063-7)
- [55] Pablo Parrilo. 2000. *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*. Ph.D. Dissertation. California Institute of Technology.
- [56] Toniann Pitassi and Nathan Segerlind. 2012. Exponential Lower Bounds and Integrality Gaps for Tree-Like Lovász–Schrijver Procedures. *SIAM J. Comput.* 41, 1 (2012), 128–159. <https://doi.org/10.1137/100816833>
- [57] Aaron Potechin. 2020. Sum of Squares Bounds for the Ordering Principle. In *Proceedings of the 35th Computational Complexity Conference (CCC)*, Vol. 169. 38:1–38:37. <https://doi.org/10.4230/LIPIcs.CCC.2020.38>
- [58] Pavel Pudlák. 2000. Proofs as Games. *The American Mathematical Monthly* 107, 6 (2000), 541–550. <https://doi.org/10.2307/2589349>
- [59] Pavel Pudlák. 2003. On reducibility and symmetry of disjoint NP pairs. *Theoretical Computer Science* 295 (2003), 323–339. [https://doi.org/10.1016/S0304-3975\(02\)00411-5](https://doi.org/10.1016/S0304-3975(02)00411-5)
- [60] Pavel Pudlák and Neil Thapen. 2019. Random resolution refutations. *Computational Complexity* 28, 2 (2019), 185–239. <https://doi.org/10.1007/s00037-019-00182-7>
- [61] Prasad Raghavendra and Benjamin Weitz. 2017. On the Bit Complexity of Sum-of-Squares Proofs. In *Proceedings of the 44th International Colloquium on Automata, Languages, and Programming (ICALP)*. 80:1–80:13. <https://doi.org/10.4230/LIPIcs.ICALP.2017.80>
- [62] Alexander Razborov. 1998. Lower bounds for the polynomial calculus. *Computational Complexity* 7 (1998), 291–324. <https://doi.org/10.1007/s000370050013>
- [63] Hanif Sherali and Warren Adams. 1994. A hierarchy of relaxations and convex hull characterizations for mixed-integer zero–one programming problems. *Discrete Applied Mathematics* 52, 1 (1994), 83–106. [https://doi.org/10.1016/0166-218X\(92\)00190-W](https://doi.org/10.1016/0166-218X(92)00190-W)
- [64] Naum Shor. 1987. An Approach to Obtaining Global Extremums in Polynomial Mathematical Programming Problems. *Cybernetics* 23, 5 (1987), 695–700. <https://doi.org/10.1007/BF01074929>