Depth Lower Bounds in Stabbing Planes for Combinatorial Principles

Stefan Dantchev ✉
Department of Computer Science, Durham University, UK

Nicola Galesi ✉
Department of Computer Science, Sapienza University of Rome, Italy

Abdul Ghani ✉
Department of Computer Science, Durham University, UK

Barnaby Martin ✉
Department of Computer Science, Durham University, UK

Abstract

Stabbing Planes is a proof system introduced very recently which, informally speaking, extends the DPLL method by branching on integer linear inequalities instead of single variables. The techniques known so far to prove size and depth lower bounds for Stabbing Planes are generalizations of those used for the *Cutting Planes* proof system established via communication complexity arguments. Rank lower bounds for Cutting Planes are also obtained by geometric arguments called *protection lemmas*.

In this work we introduce two new geometric approaches to prove size/depth lower bounds in Stabbing Planes working for any formula: (1) the *antichain method*, relying on *Sperner's Theorem* and (2) the *covering method* which uses results on *essential coverings* of the boolean cube by linear polynomials, which in turn relies on *Alon's combinatorial Nullenstellensatz*.

We demonstrate their use on classes of combinatorial principles such as the *Pigeonhole principle*, the *Tseitin contradictions* and the *Linear Ordering Principle*. By the first method we prove almost linear size lower bounds and optimal logarithmic depth lower bounds for the Pigeonhole principle and analogous lower bounds for the Tseitin contradictions over the complete graph and for the Linear Ordering Principle. By the covering method we obtain a superlinear size lower bound and a logarithmic depth lower bound for Stabbing Planes proof of Tseitin contradictions over a grid graph.

2012 ACM Subject Classification Theory of computation → Computational complexity and cryptography; Theory of computation → Proof complexity

Keywords and phrases proof complexity, computational complexity, lower bounds, cutting planes, stabbing planes

Digital Object Identifier 10.4230/LIPIcs.STACS.2022.24

Related Version *Previous Version*: https://arxiv.org/abs/2102.07622

Acknowledgements While finishing the writing of this manuscript we learned about [12] from Noah Fleming. We would like to thank him for answering some questions on his paper [2], and sending us the manuscript [12] and for comments on a preliminary version of this work.

1 Introduction

Finding a satisfying assignment for a propositional formula (SAT) is a central component for many computationally hard problems. Despite being older than 50 years and exponential time in the worst-case, the DPLL algorithm [10, 11, 26] is the core of essentially all high performance modern SAT-solvers. DPLL is a recursive boolean method: at each call one variable x of the formula $\mathcal{F}$ is chosen and the search recursively branches into the two cases obtained by setting x respectively to 1 and 0 in $\mathcal{F}$. It is well-known that the execution trace of the DPLL algorithm running on an unsatisfiable formula $\mathcal{F}$ is nothing more than a treelike refutation of $\mathcal{F}$ in the proof system of *Resolution* [26] (Res).

Since SAT can be viewed as an optimization problem the question whether Integer Linear Programming (ILP) can be made feasible for satisfiability testing received a lot of attention and is considered among the most challenging problems in local search [27, 17]. One proof system capturing ILP approaches to SAT is *Cutting Planes*, a system whose main rule implements the *rounding* (or *Chvátal cut*) approach to ILP. Cutting planes works with integer linear inequalities of the form $\mathbf{a}\mathbf{x} \leq b$, with $\mathbf{a}, b$ integers, and, like resolution, is a sound and complete refutational proof system for CNF formulas: indeed a clause $C = (x_1 \vee \ldots \vee x_r \vee \neg y_1 \vee \ldots \vee \neg y_s)$ can be written as the integer inequality $y_1 + \cdots + y_s - x_1 - \cdots - x_r \leq s - 1$.

Beame et al. [2], extended the idea of DPLL to a more general proof strategy based on ILP. Instead of branching only on a variable as in DPLL, in this method one considers a pair $(\mathbf{a}, b)$, with $\mathbf{a} \in \mathbb{Z}^n$ and $b \in \mathbb{Z}$, and branches limiting the search to the two half-planes: $\mathbf{a}\mathbf{x} \leq b - 1$ and $\mathbf{a}\mathbf{x} \geq b$. A *path* terminates when the LP defined by the inequalities in $\mathcal{F}$ and those forming the path is infeasible. This method can be made into a refutational treelike proof system for UNSAT CNF's called *Stabbing Planes* (SP) ([2]) and it turned out that it is polynomially equivalent to the treelike version of Res(CP), a proof system introduced by Krajíček [19] where clauses are disjunction of linear inequalities.

In this work we consider the complexity of proofs in SP focusing on the *length*, i.e. the number of queries in the proof; the *depth* (called also *rank* in [2]), i.e. the length of the longest path in the proof tree; and the *size*, i.e. the bit size of all the coefficients appearing in the proof.

1.1 Previous works and motivations

After its introduction as a proof system in the work [2] by Beame, Fleming, Impagliazzo, Kolokolova, Pankratov, Pitassi and Robere, *Stabbing Planes* received great attention. The quasipolynomial upper bound for the size of refuting Tseitin contradictions in SP given in [2] was surprisingly extended to CP in the work of [9] of Dadush and Tiwari refuting a long-standing conjecture. Recently in [12], Fleming, Göös, Impagliazzo, Pitassi, Robere, Tan and Wigderson were further developing the initial results proved in [2] making important progress on the question whether all Stabbing Planes proofs can be somehow efficiently simulated by Cutting Planes showing quasipolynomial simulation of bounded weight SP by CP.

Significant lower bounds for size can be obtained in SP, but in a limited way, using modern developments of a technique for CP based on communication complexity of search problems introduced by Impagliazzo, Pitassi, Urquhart in [16]: in [2] it is proven that size S and depth D SP refutations imply treelike Res(CP) proofs of size $O(S)$ and width $O(D)$; Kojevnikov [18], improving the *interpolation method* introduced for Res(CP) by Krajíček [19], gave exponential lower bounds for treelike Res(CP) when the width of the clauses (i.e. the number of linear inequalities in a clause) is bounded by $o(n/\log n)$. Hence these lower bounds are applicable only to very specific classes of formulas (whose hardness comes from boolean circuit hardness) and only to SP refutations of low depth.

Nevertheless SP appears to be a strong proof system. Firstly notice that the condition terminating a path in a proof is not a trivial contradiction like in resolution, but is the infeasibility of an LP, which is only a polynomial time verifiable condition. Hence linear size SP proofs might be already a strong class of SP proofs, since they can hide a polynomial growth into one final node whence to run the verification of the terminating condition.

Rank and depth in CP and SP

It is known that, contrary to the case of other proof systems like Frege, neither CP nor SP proofs can be balanced (see [2]), in the sense that a size $2^{O(d)}$ proof can always be converted into a depth $O(d)$ proof. The depth of CP-proofs of a set of linear inequalities L is measured in two ways: (1) as the depth of the dag representing to the proof, and (2) by the *Chvátal rank* of the associated polytope P.[1] It is known that rank in CP and depth in SP are separated, in the sense that Tseitin formulas can be proved in depth $O(\log n)$ in SP [2], but are known to require rank $\Theta(n)$ to be refuted in CP [6]. In this paper we further develop the study of proof depth for SP.

Rank lower bound techniques for Cutting Planes are essentially of two types. The main method is by reducing to the *real communication complexity* of certain search problem [16]. As such this method mainly works for classes of formulas *lifted* by certain gadgets capturing specific boolean functions. A second class of methods have been developed for Cutting Planes, which lower bound the rank measures of a polytope. In this setting, lower bounds are typically proven using a geometric method called *protection lemmas* [6]. These methods were recently extended in [12] also to the case of Semantic Cutting Planes. In principle this geometric method can be applied to any formula and not only to the lifted ones, furthermore for many formulas (such as the Tseitin formulas) it is known how to achieve $\Omega(n)$ rank lower bounds in CP via protection lemmas, while proving even $\omega(\log n)$ lower bounds via real communication complexity is impossible, due to a known folklore upper bound.

Lower bounds for depth in Stabbing Planes, proved in [2], are instead obtained only as a consequence of the real communication complexity approach extended to Stabbing Planes. In this paper we introduce two geometric approaches to prove depth lower bounds in SP.

Specifically the results we know at present relating SP and CP are:

1. SP polynomially simulates CP (Theorem 4.5 in [2]) and CP polynomially simulates SP with bounded coefficients [12]. Hence in particular the PHP^m_n can be refuted in SP by a proof of size $O(n^2)$ ([8]). Furthermore it can be refuted by a $O(\log n)$ depth proof since polynomial size CP proofs, by Theorem 4.4 in [2], can be balanced in SP.[2]
2. Beame et al. in [2] proved the surprising result that the class of Tseitin contradictions $\mathsf{Ts}(G, \omega)$ over any graph G of maximum degree D, with an odd charging ω, can be refuted in SP in size quasipolynomial in $|G|$ and depth $O(\log^2 |G| + D)$.

Depth lower bounds for SP are proved in [2]:

1. A $\Omega(n/\log^2 n)$ lower bound for the formula $\mathsf{Ts}(G, w) \circ \mathsf{VER}^n$, composing $\mathsf{Ts}(G, \omega)$ (over an expander graph G) with the gadget function VER^n (see Theorem 5.7 in [2] for details).
2. A $\Omega(\sqrt{n \log n})$ lower bound for the formula $\mathsf{Peb}(G) \circ \mathsf{IND}^n_l$ over $n^5 + n \log n$ variables obtained by lifting a pebbling formula $\mathsf{Peb}(G)$ over a graph with high pebbling number, with a *pointer function* gadget IND^n_l (see Theorem 5.5. in [2] for details).
3. There are also sub-linear lower bounds on SP depth when the coefficients in the SP proof are of magnitude at most 2^{n^δ} for some constant δ for random $O(\log n)$-CNF formulas over n variables. This lower bound is hence obtained (by communication complexity arguments) for an unlifted class of formulas by combining the Cutting Planes size lower bounds for

[1] This is the minimal d such that $P^{(d)}$ is empty, where $P^{(0)}$ is the polytope associated to L and $P^{(i+1)}$ is the polytope defined by all inequalities which can be inferred from those in $P^{(i)}$ using one Chvátal cut.

[2] Another way of proving this result is using Theorem 4.8 in [2] stating that if there are length L and space S CP refutations of a set of linear integral inequalities, then there are depth $O(S \log L)$ SP refutations of the same set of linear integral inequalities; and then use the result in [14] (Theorem 5.1) that PHP^m_n has polynomial length and constant space CP refutations.

random $O(\log n)$-CNF formulas of [15, 13] with the quasipolynomial transformation of Stabbing Planes proofs into Cutting Planes for bounded coefficients (2^{n^δ}, see [12]). This gives a $\exp(n^\delta/\,\mathsf{polylog}(n))$ size lower bound, and thus a $\Omega(n^\delta/\,\mathsf{polylog}(n))$ depth lower bound for SP proofs.

Similar to size, these depth lower bounds are applicable only to very specific classes of formulas. In fact they are obtained by extending to SP the technique introduced in [16, 20] for CP of reducing shallow proofs of a formula $\mathcal{F}$ to efficient real communication protocols computing a related search problem and then proving that such efficient protocols cannot exist.

Despite the fact that SP is at least as strong as CP, in SP the known lower bound techniques are derived from those of CP. Hence finding other techniques to prove depth and size lower bounds for SP is important to understand its proof strength. For instance, unlike CP where we know tight $\Theta(\log n)$ rank bounds for the PHP^m_n [6, 25] and $\Omega(n)$ rank bounds for Tseitin contradictions [6], for SP no depth lower bound is at present known for purely combinatorial statements.

In this work we address such problems.

1.2 Contributions and techniques

The main motivation of this work was to study size and depth lower bounds in SP through new methods, possibly geometric. Differently from weaker systems like Resolution, except for the technique highlighted above and based on reducing to the communication complexity of search problems, we do not know of other methods to prove size and depth lower bounds in SP. In CP and Semantic CP instead geometrical methods based on protection lemmas were used to prove rank lower bounds in [6, 12].

Our first steps in this direction were to set up methods working for truly combinatorial statements, like $\mathsf{Ts}(G, w)$ or PHP^m_n, which we know to be efficiently provable in SP, but on which we cannot use methods reducing to the complexity of boolean functions, like the ones based on communication complexity.

We present two new methods for proving depth lower bounds in SP which in fact are the consequence of proving length lower bounds that do not depend on the bit-size of the coefficients.

As applications of our two methods we respectively prove:

1. An exponential separation between the rank in CP and the depth in SP, using a new counting principle which we introduce and that we call the *Simple Pigeonhole Principle* SPHP_n. We prove that SPHP_n has $O(1)$ rank in CP and requires $\Omega(\log n)$ depth in SP. Together with the results proving that Tseitin formulas requires $\Omega(n)$ rank lower bounds in CP ([6]) and $O(\log^2 n)$ upper bounds for the depth in SP ([2]), this proves an incomparability between the two measures.
2. An almost linear lower bound for the size of SP proofs of the PHP^m_n and for Tseitin contradictions $\mathsf{Ts}(G, \omega)$ over the complete graph. These lower bounds immediately give an optimal $\Omega(\log n)$ lower bound for the depth of SP proofs of the corresponding principles.
3. A superlinear lower bound for the size of SP proofs of $\mathsf{Ts}(G, \omega)$, when G is a $n \times n$ grid graph H_n. In turn this implies an $\Omega(\log n)$ lower bound for the depth of SP proofs of $\mathsf{Ts}(H_n, \omega)$. Proofs of depth $O(\log^2 n)$ for $\mathsf{Ts}(H_n, \omega)$ are given in [2].
4. Finally we prove a linear lower bound for the size and a $\Omega(\log n)$ lower bound for the depth for the the Linear Ordering Principle LOP_n.

Our results are derived from the following initial geometrical observation: let $\mathbb{S}$ be a space of *admissible points* in $\{0, 1, 1/2\}^n$ satisfying a given unsatisfiable system of integer linear inequalities $\mathcal{F}(x_1, \ldots, x_n)$. In a SP proof for $\mathcal{F}$, at each branch $Q = (\mathbf{a}, b)$ the set of points in the $\mathsf{slab}(Q) = \{\mathbf{s} \in \mathbb{S} : b - 1 < \mathbf{ax} < b\}$ does not survive in $\mathbb{S}$. At the end of the proof on the leaves, where we have infeasible LP's, no point in $\mathbb{S}$ can survive the proof. So it is sufficient to find conditions such that, under the assumption that a proof of $\mathcal{F}$ is "small", even one point of $\mathbb{S}$ survives the proof. In pursuing this approach we use two methods.

The *antichain method.* Here we use a well-known bound based on Sperner's Theorem [7, 29] to upper bound the number of points in the slabs where the set of non-zero coefficients is sufficiently large. Trading between the number of such slabs and the number of points ruled out from the space $\mathbb{S}$ of admissible points, we obtain the lower bound.

We initially present the method and the $\Omega(\log n)$ lower bound on a set of unsatisfiable integer linear inequalities – the *Simple Pigeonhole Principle* (SPHP) – capturing the core of the counting argument used to prove the PHP efficiently in CP. Since SPHP_n has rank 1 CP proofs, it entails a strong separation between CP rank and SP depth. We then apply the method to PHP^m_n and to $\mathsf{Ts}(K_n, \omega)$.

The *covering method.* The antichain method appears too weak to prove size and depth lower bounds on $\mathsf{Ts}(G, w)$, when G is for example a grid or a pyramid. To solve this case, we consider another approach that we call the *covering method*: we reduce the problem of proving that one point in $\mathbb{S}$ survives from all the $\mathsf{slab}(Q)$ in a small proof of $\mathcal{F}$, to the problem that a set of polynomials which *essentially covers* the boolean cube $\{0, 1\}^n$ requires at least $\sqrt{n}$ polynomials, which is a well-known problem faced by Alon and Füredi in [1] and by Linial and Radhakrishnan in [21]. For this reduction to work we have to find a high dimensional projection of $\mathbb{S}$ covering the boolean cube and defined on variables effectively appearing in the proof. We prove that cycles of distance at least 2 in G work properly to this aim on $\mathsf{Ts}(G, \omega)$. Since the grid H_n has many such cycles, we can obtain the lower bound on $\mathsf{Ts}(H_n, \omega)$. The use of Linial and Radhakrishnan's result is not new in proof complexity. Part and Tzameret in [23], independently of us, were using this result in a completely different way from us in the proof system $\mathsf{Res}(\oplus)$ handling clauses over parity equations, and not relying on integer linear inequalities and geometrical reasoning.

We remark that while we were writing this version of the paper, Yehuda and Yehudayoff in [30] slightly improved the results of [21] with the consequence, noticed in their paper too, that our size lower bounds for $\mathsf{Ts}(G, \omega)$ over a grid graph is in fact superlinear.

The paper is organized as follows: We give the preliminary definitions in the next section and then we move to other sections: one on the lower bounds by the antichain method and the other on lower bounds by the covering method. The antichain method is presented on the formulas SPHP_n, PHP^m_n, Tseitin formulas for the complete graph and the LOP_n. The covering method is used for the Tseitin formulas for the grid graph. The lower bound for the Linear Ordering Principle, LOP_n, is deferred to the appendix.

2 Preliminaries

We use $[n]$ for the set $\{1, 2, \ldots, n\}$, $\mathbb{Z}/2$ for $\mathbb{Z} \cup (\mathbb{Z} + \frac{1}{2})$ and $\mathbb{Z}^+$ for $\{1, 2, \ldots\}$.

2.1 Proof systems

Here we recall the definition of the Stabbing Planes proof system from [2].

▶ **Definition 1.** *A* linear integer inequality *in the variables $x_1, \ldots, x_n$ is an expression of the form $\sum_{i=1}^{n} a_i x_i \geq b$, where each a_i and b are integral. A set of such inequalities is said to be* unsatisfiable *if there are no* $0/1$ *assignments to the x variables satisfying each inequality simultaneously.*

Note that we reserve the term infeasible, in contrast to unsatisfiable, for (real or rational) linear programs.

▶ **Definition 2.** *Fix some variables $x_1, \ldots, x_n$. A* Stabbing Planes (SP) *proof of a set of integer linear inequalities $\mathcal{F}$ is a binary tree $\mathcal{T}$, with each node labeled with a* query $(\mathbf{a}, b)$ *with $\mathbf{a} \in \mathbb{Z}^n, b \in \mathbb{Z}$. Out of each node we have an edge labeled with $\mathbf{a}x \geq b$ and the other labeled with its integer negation $\mathbf{a}x \leq b-1$. Each leaf ℓ is labeled with a* LP *system P_ℓ made by a nonnegative linear combination of inequalities from $\mathcal{F}$ and the inequalities labelling the edges on the path from the root of $\mathcal{T}$ to the leaf ℓ.*

If $\mathcal{F}$ is an unsatisfiable *set of integer linear inequalities, $\mathcal{T}$ is a* Stabbing Planes (SP) refutation *of $\mathcal{F}$ if all the* LP*'s P_ℓ on the leaves of $\mathcal{T}$ are infeasible.*

▶ **Definition 3.** *The* slab *corresponding to a query $Q = (\mathbf{a}, b)$ is the set* $\mathsf{slab}(Q) = \{\mathbf{x} \in \mathbb{R}^n : b-1 < \mathbf{ax} < b\}$ *satisfying neither of the associated inequalities.*

Since each leaf in a SP refutation is labelled by an infeasible LP, throughout this paper we will actually use the following geometric observation on SP proofs $\mathcal{T}$: the set of points in $\mathbb{R}^n$ must all be ruled out by a query somewhere in $\mathcal{T}$. In particular this will be true for those points in $\mathbb{R}^n$ which satisfy a set of integer linear inequalities $\mathcal{F}$ and which we call *feasible points* for $\mathcal{F}$.

▶ **Fact 4.** *The slabs associated with a* SP *refutation must cover the feasible points of $\mathcal{F}$. That is,*

$$\{\mathbf{y} \in \mathbb{R}^n : \mathbf{cy} \geq r \text{ for all } (\mathbf{c}, r) \in \mathcal{F}\} \subseteq \bigcup_{(\mathbf{a},b)\in\mathcal{T}} \{\mathbf{x} \in \mathbb{R}^n : b-1 < \mathbf{ax} < b\}$$

The *length* of a SP refutation is the number of queries in the proof tree. The *depth* of a SP refutation $\mathcal{T}$ is the longest root-to-leaf path in $\mathcal{T}$. The size (respectively depth) of refuting $\mathcal{F}$ in SP is the *minimum* size (respectively depth) over all SP refutations of $\mathcal{F}$. We call *bit-size* of a SP refutation $\mathcal{T}$ the total number of bits needed to represent every inequality in the refutation.

▶ **Definition 5** ([8])**.** *The* Cutting Planes (CP) *proof system is equipped with boolean axioms and two inference rules:*

Boolean Axioms	*Linear Combination*	*Rounding*
$\dfrac{}{x \geq 0}$ $\quad$ $\dfrac{}{-x \geq -1}$	$\dfrac{\mathbf{ax} \geq c \qquad \mathbf{bx} \geq d}{\alpha\mathbf{ax} + \beta\mathbf{bx} \geq \alpha c + \beta d}$	$\dfrac{\alpha\mathbf{ax} \geq b}{\mathbf{ax} \geq \lceil b/\alpha \rceil}$

where $\alpha, \beta, b \in \mathbb{Z}^+$ and $\mathbf{a}, \mathbf{b} \in \mathbb{Z}^n$. A CP refutation of some unsatisfiable set of integer linear inequalities is a derivation of $0 \geq 1$ *by the aforementioned inference rules from the inequalities in $\mathcal{F}$.*

A CP refutation is *treelike* if the directed acyclic graph underlying the proof is a tree. The *length* of a CP refutation is the number of inequalities in the sequence. The *depth* is the length of the longest path from the root to a leaf (sink) in the graph. The *rank* of a CP proof is the maximal number of rounding rules used in a path of the proof graph. The *size* of a CP refutation is the bit-size to represent all the inequalities in the proof.

2.2 Restrictions

Let $V = \{x_1, \ldots, x_n\}$ be a set of n variables and let $\mathbf{ax} \le b$ be a linear integer inequality. We say that a variable x_i *appears in*, or is *mentioned by* a query $Q = (\mathbf{a}, b)$ if $a_i \neq 0$ and *does not appear* otherwise.

A *restriction* ρ is a function $\rho : D \to \{0, 1\}$, $D \subseteq V$. A restriction acts on a half-plane $\mathbf{ax} \le b$ setting the x_i's according to ρ. Notice that the variables $x_i \in D$ do not appear in the restricted half-plane.

By $\mathcal{T}\restriction_\rho$ we mean to apply the restriction ρ to all the queries in a SP proof $\mathcal{T}$. The tree $\mathcal{T}\restriction_\rho$ defines a new SP proof: if some $Q\restriction_\rho$ reduces to $0 \le -b$, for some $b \ge 1$, then that node becomes a leaf in $\mathcal{T}\restriction_\rho$. Otherwise in $\mathcal{T}\restriction_\rho$ we simply branch on $Q\restriction_\rho$. Of course the solution space defined by the linear inequalities labelling a path in $\mathcal{T}\restriction_\rho$ is a subset of the solution space defined by the corresponding path in $\mathcal{T}$. Hence the leaves of $\mathcal{T}\restriction_\rho$ define an infeasible LP.

We work with linear integer inequalities which are a translation of families of CNFs $\mathcal{F}$. Hence when we write $\mathcal{F}\restriction_\rho$ we mean the applications of the restriction ρ to the set of linear integer inequalities defining $\mathcal{F}$.

3 The antichain method

This method is based on Sperner's theorem. Using it we can prove depth lower bounds in SP for PHP^m_n and for Tseitin contradictions $\mathsf{Ts}(K_n, \omega)$ over the complete graph. To motivate and explain the main definitions, we use as an example a simplification of the PHP^m_n, the *Simplified Pigeonhole principle* SPHP_n, which has some interest since (as we will show) it exponentially separates CP rank from SP depth.

3.1 Simplified Pigeonhole Principle

As mentioned in the Introduction, the SPHP_n intends to capture the core of the counting argument used to efficiently refute the PHP in CP.

▶ **Definition 6.** *The* SPHP_n *is the following unsatisfiable family of inequalities:*

$$\sum_{i=1}^n x_i \ge 2$$
$$x_i + x_j \le 1 \qquad \textit{for all } i \neq j \in [n]$$
$$0 \le x_i \le 1 \qquad \textit{for all } i \in [n].$$

▶ **Lemma 7.** SPHP_n *has a rank* 1 CP *refutation, for* $n \ge 3$.

Proof. Let $S := \sum_{i=1}^n x_i$ (so we have $S \ge 2$). We fix some $i \in [n]$ and sum $x_i + x_j \le 1$ over all $j \in [n] \setminus \{i\}$ to find $S + (n-2)x_i \le n - 1$. We add this to $-S \le -2$ to get

$$x_i \le \frac{n-3}{n-2}$$

which becomes $x_i \le 0$ after a single cut. We do this for every i and find $S \le 0$ - a contradiction when combined with the axiom $S \ge 2$. ◀

It is easy to see that SPHP_n has depth $O(\log n)$, length $O(n)$ proofs in SP, either by a direct proof or appealing to the polynomial size proofs in CP of the PHP^m_n ([8]) and then using the Theorem 4.4 in [2] informally stating that "CP proofs can be balanced in SP".

▶ **Theorem 8.** *The* SPHP_n *has a* SP *refutation of size* $O(n)$ *and depth* $O(\log(n))$.

Proof. Note that no admissible point for the SPHP_n has any x_i set to 1. The SP refutation just performs a binary search looking for an x_i set to 1 – if it cannot find such an x_i, we contradict the axiom $\sum_{i=1}^n x_i \geq 2$,

In more detail, the root asks if $\sum_{i=1}^n x_i$ is at least 1 or at most 0. The at most 0 branch directly contradicts the axiom $\sum_{i=1}^n x_i \geq 2$ (and so terminates). The at least 1 branch asks if $\sum_{i=1}^{\lfloor n/2 \rfloor} x_i$ is again at least 1 or at most 0. If this is at most 0, we must have that $\sum_{i=\lfloor n/2 \rfloor+1}^{n} x_i \geq 1$ - in either case we have halved the range of the index of summation. ◀

We will prove that this depth bound is tight.

3.2 Sperner's Theorem

Let $\mathbf{a} \in \mathbb{R}^n$. The *width* $w(\mathbf{a})$ of $\mathbf{a}$ is the number of non-zero coordinates in $\mathbf{a}$. The width of a query $(\mathbf{a}, b)$ is $w(\mathbf{a})$, and the width of a SP refutation is the minimum width of its queries.

Let $n \in \mathbb{N}$. Fix $W \subseteq [0,1] \cap \mathbb{Q}^+$ of finite size $k \geq 2$ and insist that $0 \in W$. The W's we work with in this paper are $\{0, 1/2\}$ and $\{0, 1/2, 1\}$.

▶ **Definition 9.** *A (n, W)-word is an element in W^n.*

For two vectors $x, y \in \mathbb{R}^d$, say that $x \leq y$ *in the pointwise ordering* if $x_i \leq y_i$ for all $1 \leq i \leq d$. We consider the following extension of Sperner's theorem.

▶ **Theorem 10** ([22, 7])**.** *Fix any $t \geq 2, t \in \mathbb{N}$. For all $f \in \mathbb{N}$, with the pointwise ordering of $[t]^f$, any antichain has size at most $t^f \sqrt{\frac{6}{\pi(t^2-1)f}}(1 + o(1))$.*

We will use the simplified bound that any antichain $\mathcal{A}$ has size $|\mathcal{A}| \leq \frac{t^f}{\sqrt{f}}$.

▶ **Lemma 11.** *Let $\mathbf{a} \in \mathbb{Z}^n$ and $|W| = k \geq 2$. The number of (n, W)-words $\mathbf{s}$ such that $\mathbf{as} = b$, where $b \in \mathbb{Q}$, is at most $\frac{k^n}{\sqrt{w(\mathbf{a})}}$.*

Proof. Define $I_{\mathbf{a}} = \{i \in [n] : a_i \neq 0\}$. Let $\preceq$ be the partial order over $W^{I_{\mathbf{a}}}$ where $\mathbf{x} \preceq \mathbf{y}$ if $x_i \leq y_i$ for all i with $a_i > 0$ and $x_i \geq y_i$ for the remaining i with $a_i < 0$. Clearly the set of solutions (restricted to indices in $I_{\mathbf{a}}$) to $\mathbf{as} = b$ forms an antichain under $\preceq$. Noting that $\preceq$ is isomorphic to the typical pointwise ordering on $W^{I_{\mathbf{a}}}$, we appeal to Theorem 10 to upper bound the number of solutions in $W^{I_{\mathbf{a}}}$ by $\frac{k^{w(\mathbf{a})}}{\sqrt{w(\mathbf{a})}}$, each of which corresponds to at most $k^{n-w(\mathbf{a})}$ vectors in W^n. ◀

3.3 Large admissibility

A (n, W)-word s is *admissible* for an unsatisfiable set of integer linear inequalities $\mathcal{F}$ over n variables if s satisfies all constraints of $\mathcal{F}$. A set of (n, W)-words is admissible for $\mathcal{F}$ if all its elements are admissible. $\mathcal{A}(\mathcal{F}, W)$ is the set of all admissible (n, W)-words for $\mathcal{F}$.

The interesting sets W for an unsatisfiable set of integer linear inequalities $\mathcal{F}$ are those such that almost all (n, W)-words are admissible for $\mathcal{F}$. We will apply our method on sets of integer linear inequalities which are a translation of unsatisfiable CNF's generated over a given domain. Typically these formulas on a size n domain have a number of variables which is not exactly n but a function of n, $\nu(n) \geq n$. (For example, the PHP_n^{n+1} has $\nu(n) = n^2 + n$ variables.) Hence for the rest of this section we consider $\mathscr{F} := \{\mathcal{F}_n\}_{n \in \mathbb{N}}$ as a family of sets of unsatisfiable integer linear inequalities, where $\mathcal{F}_n$ has $\nu(n) \geq n$ variables. We call $\mathscr{F}$ an *unsatisfiable family.*

Consider then the following definition (recalling that we denote $k = |W|$):

▶ **Definition 12.** *$\mathscr{F}$ is* almost full *if* $|\mathcal{A}(\mathcal{F}_n, W)| \geq k^{\nu(n)} - o(k^{\nu(n)})$.

Notice that, because of the o notation, Definition 12 might be not necessarily true for all $n \in \mathbb{N}$, but only starting from some $n_{\mathscr{F}}$. Note also that being almost full is defined relative to some W.

▶ **Definition 13.** *Given some almost full family $\mathscr{F}$ (over $\nu(n)$ variables) we let $n_{\mathscr{F}}$ be the natural number with*

$$\frac{k^{\nu(n)}}{|\mathcal{A}(\mathcal{F}_n, W)|} \leq 2 \quad \textit{for all} \quad n \geq n_{\mathscr{F}}.$$

As an example we prove SPHP is almost full (notice that in the case of SPHP_n, $\nu(n) = n$).

▶ **Lemma 14.** SPHP_n *is almost full when* $W = \{0, 1/2\}$.

Proof. Let U be the set of all (n, W)-words with at least four coordinates set to $1/2$. U is admissible for SPHP_n since inequalities $x_i + x_j \leq 1$ are always satisfied for any value in W and inequalities $x_1 + \ldots + x_n \geq 2$ are satisfied by all points in U which contain at least four 1/2s. By a simple counting argument, in U there are at least $2^n - 4n^3 = 2^n - o(2^n)$ admissible (n, W)-words. Hence the claim. ◀

▶ **Lemma 15.** *Let $\mathscr{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$ be an almost full unsatisfiable family, where $\mathcal{F}_n$ has $\nu(n)$ variables. Further let $\mathcal{T}$ be a* SP *refutation of $\mathcal{F}$ of width w. If $n \geq n_{\mathscr{F}}$ then $|\mathcal{T}| = \Omega(\sqrt{w})$.*

Proof. We estimate at what rate the slab of the queries in $\mathcal{T}$ rule out admissible points in U. Let ℓ be the least common multiple of the denominators in W. Every (n, W)-word x falling in the slab of some query $(\mathbf{a}, b)$ satisfies one of ℓ equations $\mathbf{a}x = b + i/\ell, 1 \leq i < \ell$ (as $\mathbf{a}$ is integral). Note that as $|W|$ is a constant independent of n, so is ℓ.

Since all the queries in $\mathcal{T}$ have width at least w, according to Lemma 11, each query in $\mathcal{T}$ rules out at most $\ell \cdot \frac{k^{\nu(n)}}{\sqrt{w}}$ admissible points. By Fact 4 no point survives at the leaves, in particular the admissible points. Then it must be that

$$|\mathcal{T}|\ell \cdot \frac{k^{\nu(n)}}{\sqrt{w}} \geq |\mathcal{A}(\mathcal{F}_n, W)| \quad \text{which means} \quad |\mathcal{T}|\ell \cdot \frac{k^{\nu(n)}}{|\mathcal{A}(\mathcal{F}_n, W)|} \geq \sqrt{w}$$

We finish by noting that, by the assumption $n \geq n_{\mathscr{F}}$, and then by Definition 13, we have $2 \geq \frac{k^{\nu(n)}}{|\mathcal{A}(\mathcal{F}_n, W)|}$, so $|\mathcal{T}| \geq \sqrt{w}/(2\ell) \in \Omega(\sqrt{w})$. ◀

3.4 Main theorem

We focus on restrictions ρ that after applied to an unsatisfiable family $\mathscr{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$, reduce the set $\mathcal{F}$ to another set in the same family.

▶ **Definition 16.** *Let $\mathscr{F} = \{\mathcal{F}_n\}_{n \in \mathbb{N}}$ be an unsatisfiable family and c a positive constant. $\mathscr{F}$ is* c-self-reducible *if for any set V of variables, with $|V| = v < n/c$, there is a restriction ρ with domain $V' \supseteq V$, such that $\mathcal{F}_n\restriction_\rho = \mathcal{F}_{n-cv}$ (up to renaming of variables).*

Let us motivate the definition with an example.

▶ **Lemma 17.** SPHP_n *is* 1*-self-reducible.*

Proof. Whatever set of variables x_i, $i \in I \subset [n]$ we consider, it is sufficient to set x_i to 0 to fulfill Definition 16. ◀

▶ **Theorem 18.** *Let $\mathscr{F} := \{\mathcal{F}_n\}_{n\in\mathbb{N}}$ be a unsatisfiable set of integer linear inequalities which is almost full and c-self-reducible, for some constant c. If $\mathcal{F}_n$ defines a feasible* LP *whenever $n > n_{\mathscr{F}}$, then for n large enough, the shortest* SP *proof of $\mathcal{F}_n$ is of length $\Omega(\sqrt[4]{n})$.*

Proof. Take any SP proof $\mathcal{T}$ refuting $\mathcal{F}_n$ and fix $t = \sqrt[4]{n}$.

The proof proceeds by stages $i \geq 0$ where $\mathcal{T}_0 = \mathcal{T}$. The stages will go on while the invariant property (which at stage 0 is true since $n > n_{\mathscr{F}}$ and c a positive constant)

$$n - ict^3 > \max\{n_{\mathscr{F}}, n(1 - 1/c)\}$$

holds.

At the stage i we let $\Sigma_i = \{(\mathbf{a}, b) \in \mathcal{T}_i : w(\mathbf{a}) \leq t^2\}$ and $s_i = |\Sigma_i|$. If $s_i \geq t$ the claim is trivially proven. If $s_i = 0$, then all queries in $\mathcal{T}_i$ have width at least t^2 and by Lemma 15 (which can be applied since $n - ict^3 > n_{\mathscr{F}}$) the claim is proven (for n large enough).

So assume that $0 < s_i < t$. Each of the queries in Σ_i involves at most t^2 nonzero coefficients, hence in total they mention at most $s_i t^2 \leq t^3$ variables. Extend this set of variables to some V' in accordance with Definition 16 (which can be done since, by the invariant, $ict^3 < n/c$). Set all these variables according to self-reducibility of $\mathcal{F}$ in a restriction ρ_i and define $\mathcal{T}_{i+1} = \mathcal{T}_i\restriction_{\rho_i}$. Note that by Definition 16 and by that of restriction, $\mathcal{T}_{i+1}$ is a SP refutation of $\mathcal{F}_{n-ict^3}$ and we can go on with the next stage. (Also note that we do not hit an empty refutation this way, due to the assumption that $\mathcal{F}_n$ defines a feasible LP.)

Assume that the invariant does not hold. If this is because $n - ict^3 < n_{\mathscr{F}}$ then, as each iteration destroys at least one node,

$$|\mathcal{T}| \geq i > \frac{n - n_{\mathscr{F}}}{ct^3} \in \Omega(n^{1/4}).$$

If this is because $n - ict^3 < n - n/c$, then again for the same reason it holds that

$$|\mathcal{T}| \geq i > \frac{n}{c^2 n^{3/4}} \in \Omega(n^{1/4}).$$ ◀

Using Lemmas 14 and 17 and the previous Theorem we get:

▶ **Corollary 19.** *The length of any* SP *refutation of* SPHP_n *is $\Omega(\sqrt[4]{n})$. Hence the minimal depth is $\Omega(\log n)$.*

3.5 Lower bounds for the Pigeonhole principle

▶ **Definition 20.** *The* Pigeonhole principle PHP^m_n*, $m(n) > n$, is the family of unsatisfiable integer linear inequalities defined over the variables $\{P_{i,j} : i \in [m], j \in [n]\}$ consisting of the following inequalities:*

$$\begin{array}{lll} \sum_{j=1}^{n} P_{i,j} \geq 1 & \forall i \in [m] & \text{(every pigeon goes into some hole)} \\ P_{i,k} + P_{j,k} \leq 1 & \forall k \in [n], i \neq j \in [m] & \text{(at most one pigeon enters any given hole)} \end{array}$$

We present a lower bound for PHP^m_n closely following that for SPHP_n, in which we largely ignore the diversity of different pigeons (which makes the principle rather like SPHP_n).

In this subsection we fix $W = \{0, 1/2\}$, and for the sake of brevity refer to (n, W)-words as *biwords.*

In this section we fix m to be $n + d$, for any fixed $d \in \mathbb{N}$ at least one.

▶ **Lemma 21.** *The* PHP^{n+d}_n *is almost full.*

Proof. We show that there are at least 2^{mn-1} admissible biwords (for sufficiently large n). For each pigeon i, there are admissible valuations to holes so that, so long as at least two of these are set to $1/2$, the others may be set to anything in $\{0, 1/2\}$. This gives at least $2^n - (n+1)$ possibilities. Since the pigeons are independent, we obtain at least $(2^n - (n+1))^m$ biwords. Now this is $2^{mn}\left(1 - \frac{n+1}{2^n}\right)^m$ where $\left(1 - \frac{n+1}{2^n}\right)^m \sim e^{\frac{-(n+1)m}{2^n}}$ whence, $\left(1 - \frac{n+1}{2^n}\right)^m \geq e^{\frac{-(n+2)m}{2^n}}$ for sufficiently large n. It follows there is a constant c so that:

$$2^{mn}\left(1 - \frac{n+1}{2^n}\right)^m \geq 2^{mn - \frac{c(n+2)m}{2^n}} \geq 2^{mn-1}$$

for sufficiently large n. ◀

▶ **Lemma 22.** *The* PHP_n^{n+d} *is* 1*-self-reducible.*

Proof. We are given some set I of variables from PHP_n^{n+d}. These variables will mention some set of holes $H := \{j : P_{i,j} \in I \text{ for some i}\}$ and similarly a set of pigeons P. Each of P, H have size at most $|I|$ and we extend them both arbitrarily to have size exactly $|I|$. Our restriction matches P and H in any way and then sets any other variable mentioning a pigeon in P or a hole in H to 0. ◀

▶ **Theorem 23.** *The length of any* SP *refutation of* PHP_n^{n+d} *is* $\Omega(\sqrt[4]{n})$*.*

Proof. Note that the all 1/2 point is feasible for PHP_n^{n+d}. Then with Lemma 21 and Lemma 22 in hand we meet all the prerequisites for Theorem 18. ◀

By simply noting that a SP refutation is a binary tree, we get the following corollary.

▶ **Corollary 24.** *The* SP *depth of the* PHP_n^{n+d} *is* $\Omega(\log n)$*.*

3.6 Lower bounds for Tseitin contradictions over the complete graph

▶ **Definition 25.** *For a graph* $G = (V, E)$ *along with a charging function* $\omega : V \to \{0, 1\}$ *satisfying* $\sum_{v \in V} \omega(v) = 1 \mod 2$*. The* Tseitin contradiction $\mathsf{Ts}(G, \omega)$ *is the set of linear inequalities which translate the CNF encoding of*

$$\sum_{\substack{e \in E \\ e \ni v}} x_e = \omega(v) \mod 2.$$

for every $v \in V$*, where the variables* x_e *range over the edges* $e \in E$*.*

In this subsection we consider $\mathsf{Ts}(K_n, \omega)$ and ω will always be an odd charging for K_n. We let $N := \binom{n}{2}$ and we fix $W = \{0, 1/2, 1\}$, $k = 3$ and for the sake of brevity refer to (n, W)-words as *triwords*. We will abuse slightly the notation of Section 3.3 and consider the family $\{\mathsf{Ts}(K_n, \omega)\}_{n \in \mathbb{N},\ \omega \text{ odd}}$ as a single parameter family in n. The reason we can do this is because the following proofs of almost fullness and self reducibility do not depend on ω at all (so long as it is odd, which we will always ensure).

▶ **Lemma 26.** $\mathsf{Ts}(K_n, \omega)$ *is almost full.*

Proof. We show that $\mathsf{Ts}(K_n, \omega)$ has at least $c3^N$ admissible triwords, for any constant $0 < c < 1$ and n large enough. We define the assignment ρ setting all edges (i.e. x_e) to a value in $W = \{0, 1, 1/2\}$ independently and uniformly at random, and inspecting the probability that some fixed constraint for a node v is violated by ρ.

Clearly if at least 2 edges incident to v are set to $1/2$ its constraint is satisfied. If none of its incident edges are set to $1/2$ then it is satisfied with probability $1/2$. Let $A(v)$ be the event "*no edge incident to v is set to $1/2$ by ρ*" and let $B(v)$ be the event that "*exactly one edge incident to v is set to $1/2$ by ρ*". Then:

$$\Pr[v \text{ is violated}] \leq \frac{1}{2}\Pr[A(v)] + \Pr[B(v)] = \frac{1}{2}\frac{2^{n-1}}{3^{n-1}} + \frac{(n-1)2^{n-2}}{3^{n-1}} = n\frac{2^{n-2}}{3^{n-1}}.$$

Therefore, by a union bound, the probability that there exists a node with violated parity is bounded above by $n^2\frac{2^{n-2}}{3^{n-1}}$, which approaches 0 as n goes to infinity. ◀

▶ **Lemma 27.** *$\mathsf{Ts}(K_n, \omega)$ is 2-self-reducible.*

Proof. We are given some set of variables I. Each variable mentions 2 nodes, so extend these mentioned nodes arbitrarily to a set S of size exactly $2|I|$, which we then hit with the following restriction: if S is evenly charged, pick any matching on the set $\{s \in S : w(s) = 1\}$, set those edges to 1, and set any other edges involving some vertex in S to 0. Otherwise (if S is oddly charged) pick any $l \in \{s \in S : w(s) = 1\}$ and $r \in [n] \setminus S$ and set x_{lr} to 1. $\{s \in S : w(s) = 1\} \setminus l$ is now even so we can pick a matching as before. And as before we set all other edges involving some vertex in S to 0. In the first case the graph induced by $[n] \setminus S$ must be oddly charged (as the original graph was). In the second case this induced graph was originally evenly charged, but we changed this when we set x_{lr} to 1. ◀

▶ **Lemma 28.** *For any oddly charged ω and n large enough, all SP refutations of $\mathsf{Ts}(K_n, \omega)$ have length $\Omega(\sqrt[4]{n})$.*

Proof. We have that the all $1/2$ point is feasible for $\mathsf{Ts}(K_n, \omega)$. Then we can simply apply Theorem 18. ◀

▶ **Corollary 29.** *The depth of any SP refutation of $\mathsf{Ts}(K_n, \omega)$ is $\Omega(\log n)$.*

4 The covering method

▶ **Definition 30.** *A set L of linear polynomials with real coefficients is said to be a* cover *of the cube $\{0,1\}^n$ if for each $v \in \{0,1\}^n$, there is a $p \in L$ such that $p(v) = 0$.*

In [21] Linial and Radhakrishnan considered the problem of the minimal number of hyperplanes needed to cover the cube $\{0,1\}^n$. Clearly every such cube can be covered by the zero polynomial, so to make the problem more meaningful they defined the notion of an *essential covering* of $\{0,1\}^n$.

▶ **Definition 31** ([21])**.** *A set L of linear polynomials with real coefficients is said to be an* essential cover *of the cube $\{0,1\}^n$ if*

(E1) *L is a cover of $\{0,1\}^n$,*

(E2) *no proper subset of L satisfies (E1), that is, for every $p \in L$, there is a $v \in \{0,1\}^n$ such that p alone takes the value 0 on v, and*

(E3) *every variable appears (in some monomial with non-zero coefficient) in some polynomial of L.*

They then proved that any essential cover E of the hypercube $\{0,1\}^n$ must satisfy $|E| \geq \sqrt{n}$. We will use the slightly strengthened lower bound given in [31]:

▶ **Theorem 32.** *Any essential cover L of the cube with n coordinates satisfies $|L| \in \Omega(n^{0.52})$.*

We will need an auxillary definition and lemma.

▶ **Definition 33.** *Let L be a cover of $\{0,1\}^I$ for some index set I. Some subset L' of L is an* essentialisation *of L if L' also covers $\{0,1\}^I$ but no proper subset of it does.*

▶ **Lemma 34.** *Let L be a cover of the cube $\{0,1\}^n$ and L' be any essentialisation of L. Let M' be the set of variables appearing with nonzero coefficient in L'. Then L' is an essential cover of $\{0,1\}^{M'}$.*

Proof.

(E1) Given any point $x \in \{0,1\}^{M'}$, we can extend it arbitrarily to a point $x' \in \{0,1\}^M$. Then there is some $p \in L'$ with $p(x') = 0$ – but $p(x') = p(x)$, as p doesn't mention any variable outside of M'.

(E2) Similarly to the previous point, this will follow from the fact that if some set $\mathcal{T}$ covers a hypercube $\{0,1\}^I$, it also covers $\{0,1\}^{I'}$ for any $I' \supseteq I$.
Suppose some proper subset $L'' \subset L'$ covers $\{0,1\}^{M'}$, then it covers $\{0,1\}^n$ – but we picked L' to be a minimal set with this property.

(E3) We defined M' to be the set of variables appearing with nonzero coefficient in L'. ◀

4.1 The covering method and Tseitin

Let H_n denote the $n \times n$ grid graph. Fix some ω with odd charge and a SP refutation $\mathcal{T}$ of $\mathsf{Ts}(H_n, \omega)$. Fact 4 tells us that for every point x admissible for $\mathsf{Ts}(H_n, \omega)$, there exists a query $(\mathbf{a}, b) \in \mathcal{T}$ such that $b < \mathbf{a}x < b + 1$. In this section we will only consider admissible points with entries in $\{0, 1/2, 1\}$, turning the slab of a query $(\mathbf{a}, b)$ into the solution set of the single linear equation $\mathbf{a} \cdot x = b + 1/2$. So we consider $\mathcal{T}$ as a set of such equations.

We say that an edge of H_n is *mentioned* in $\mathcal{T}$ if the variable x_e appears with non-zero coefficient in some query in $\mathcal{T}$. We can see H_n as a set of $(n-1)^2$ squares (4-cycles), and we can index them as if they were a Cartesian grid, starting from 1. Let S be the set of $\lfloor (n/3)^2 \rfloor$ squares in H_n gotten by picking squares with indices that become 2 (mod 3). This ensures that every two squares in S in the same row or column have at least two other squares between them, and that no selected square is on the perimeter.

We will assume WLOG that n is a multiple of 3, so $|S| = (n/3)^2$. Let $K = \bigcup_{t \in S} t$ be the set of edges mentioned by S, and for some $s \in S$, let $K_s := \bigcup_{t \in S, t \neq s} t$ be the set of edges mentioned in S by squares other than s.

▶ **Lemma 35.** *For every $s \in S$ we can find an admissible point $b_s \in \{0, 1/2, 1\}^{E(H_n)}$ such that*

1. *$b_s(x_e) = 0$ for all $e \in K_s$, and*
2. *b_s is fractional only on the edges in s.*

Proof. We use the following fact due to A. Urquhart in [28]

▶ **Fact 36.** *For each vertex v in H_n there is a totally binary assignment, called v-critical in [28], satisfying all parity axioms in $\mathsf{Ts}(H_n, \omega)$ except the parity axiom of node v.*

Pick any corner c of s. Let b_s be the result of taking any c-critical assignment of the variables of $\mathsf{Ts}(H_n, \omega)$ and setting the edges in s to $1/2$. b_s is admissible, as c is now adjacent to two variables set to $1/2$ (so its originally falsified parity axiom becomes satisfied) and every other vertex is either unaffected or also adjacent to two $1/2$s. While b_s sets some edge $e \in K_s$ to 1, flip all of the edges in the unique other square containing e. This other square always exists (as no square touches the perimeter) and also contains no other edge in K_s (as there are at least two squares between any two squares in S). Flipping the edges in a cycle preserves admissibility, as every vertex is adjacent to 0 or 2 flipped edges. ◀

▶ **Definition 37.** *Let $V_S := \{v_s : s \in S\}$ be a set of new variables. For $s \in S$ define the substitution h_s, taking the variables of $\mathsf{Ts}(H_n, \omega)$ to $V_S \cup \{0, 1/2, 1\}$, as*

$$h_s(x_e) := \begin{cases} b_s(e) & \text{if } e \text{ is not mentioned in } S\text{, or if } e \text{ is mentioned by } s, \\ v_t & \text{if } e \text{ is mentioned by some square } t \neq s \in S. \end{cases}$$

(where b_s is from Lemma 35).

▶ **Definition 38.** *Say that a linear polynomial $p = c + \sum_{e \in E(H_n)} \mu_e x_e$ with coefficients $\mu_e \in \mathbb{Z}$ and some constant part $c \in \mathbb{R}$ has* odd coefficient *in $X \subseteq E(H_n)$ if $\sum_{e \in X} \mu_e$ is an odd integer. Given some polynomial p in the variables x_e of Tseitin, and some square $s \in S$, let p_s be the polynomial in variables V_S gotten by applying the substitution $x_e \to h_s(x_e)$. Also, for any set of polynomials $\mathcal{T}$ in the variables x_e let $\mathcal{T}_s := \{p_s : p \in \mathcal{T}, p$ has odd coefficient in $s\}$.*

Given some assignment $\alpha \in \{0,1\}^{V_S \setminus \{v_s\}}$, and some h_s as in Definition 37, we let $\alpha(h_s)$ be the assignment to the variables of $\mathsf{Ts}(H_n, \omega)$ gotten by replacing the v_t in the definition of h_s by $\alpha(v_t)$.

▶ **Lemma 39.** *Let $s \in S$. For all $2^{|S|-1}$ settings α of the variables in $V_S \setminus \{v_s\}$, $\alpha(h_s)$ is admissible.*

Proof. When $\alpha(v_t)$ is all 0, $h_s = b_s$ is admissible (by Lemma 35). Toggling some v_t only has the effect of flipping every edge in a cycle, which preserves admissibility. ◀

▶ **Lemma 40.** *$\mathcal{T}_s$ covers $\{0,1\}^{V_S \setminus \{v_s\}}$.*

Proof. For every setting of $\alpha \in \{0,1\}^{V_S \setminus \{v_s\}}$, $\alpha(h_s)$ as defined above is admissible and therefore covered by some $p \in \mathcal{T}$, which has constant part $1/2 + b$ for some $b \in \mathbb{Z}$. Furthermore, as $\alpha(h_s)$ sets every edge in s to $1/2$, every such p must have odd coefficient in front of s - otherwise

$$p(\alpha(h_s)) = 1/2 + b + (1/2)\left(\sum_{e \in s} \mu_e\right) + \sum_{e \notin s} \mu_e \alpha(h_s)(x_e)$$

can never be zero, as the $1/2$ is the only non integral term in the summation. ◀

▶ **Theorem 41.** *Any SP refutation $\mathcal{T}$ of $\mathsf{Ts}(H_n, \omega)$ must have $|\mathcal{T}| \in \Omega(n^{1.04})$.*

Proof. We are going to find a set of pairs $(L_1, M_1), (L_2, M_2), \ldots, (L_q, M_q)$, where the L_i are pairwise disjoint nonempty subsets of $\mathcal{T}$, the M_i are subsets of V_S, and for every i there is some $s_i \in S \setminus \bigcup_{i=1}^q M_i$ such that $|(L_i)_{s_i}| \geq |M_i|^{0.52}$. These pairs will also satisfy the property that

$$\{s_i : 1 \leq i \leq q\} \cup \bigcup_{i=1}^{q} M_i = S. \tag{1}$$

As $|S| = (n/3)^2$ this would imply that $\sum_{i=1}^q |M_i| \geq (n/3)^2 - q$. If $q \geq (n/3)^2/2$, then (as the L_i are nonempty and pairwise disjoint) we have $|\mathcal{T}| \geq (n/3)^2/2 \in \Omega(n^{1.04})$. Otherwise $\sum_{i=1}^q |M_i| \geq (n/3)^2/2$, and as (by Theorem 32) each $|L_i| \geq |M_i|^{0.52}$,

$$|\mathcal{T}| \geq \sum_{i=1}^{q} |L_i| \geq \sum_{i=1}^{q} |M_i|^{0.52} \geq \left(\sum_{i=1}^{q} |M_i|\right)^{0.52} \geq \left((n/3)^2/2\right)^{0.52} \in \Omega(n^{1.04}). \tag{2}$$

We create the pairs by stages. Let $S_1 = S$ and start by picking any $s_1 \in S_1$. By Lemma 40 $\mathcal{T}_{s_1}$ covers $\{0,1\}^{V_{S_1} \setminus \{v_{s_1}\}}$ and has as an essentialisation E, which will be an essential cover of $\{0,1\}^{V'}$ for some $V' \subseteq V_{S_1} \setminus \{v_{s_1}\}$. We create the pair $(L_1, M_1) = (\{p : p_{s_1} \in E\}, V')$ and update $S_2 = S_1 \setminus (\{s : v_s \in V'\} \cup \{s_1\})$. (Note that V' could possibly be empty - for example, if the polynomial $x_e = 1/2$ appears in $\mathcal{T}$, where $e \in s_1$. In this case however we still have $|L_1| \geq |M_1|^{0.52}$. If V' is not empty we have the same bound due to Theorem 32.) If S_2 is nonempty we repeat with any $s_2 \in S_2$, and so on.

We now show that as promised the left hand sides of these pairs partition a subset of $\mathcal{T}$, which will give us the first inequality in Equation (2). Every polynomial p with $p \in L_i$ has every v_t mentioned by p_{s_i} removed from S_j for all $j \geq i$, so the only way p could reappear in some later L_j is if $p_{s_j} \in \mathcal{T}_{s_j}$, where v_{s_j} does *not* appear in p_{s_i}. Let $\mu_e, e \in s_j$ be the coefficients of p in front of the four edges of s_j. The coefficient in front of v_{s_j} in p_{s_i} is just $\sum_{e \in s_j} \mu_e$. As v_{s_j} failed to appear this sum is 0 and p does not have the odd coefficient sum it would need to appear in $\mathcal{T}_{s_j}$. ◀

5 Conclusions

The $\Omega(\log n)$ depth lower bound for $\mathsf{Ts}(H_n, \omega)$ is not optimal since [2] proved an $O(\log^2 n)$ upper bound for $\mathsf{Ts}(G, \omega)$, for any bounded-degree G. Even to apply the covering method to prove a depth $\Omega(\log^2 n)$ lower bound on $\mathsf{Ts}(K_n, \omega)$ (notice that it would imply a superpolynomial length lower bound), the polynomial covering of the boolean cube should be improved to work on general cubes. To this end the algebraic method used in [21] should be improved to work with generalizations of multilinear polynomials.

One weakness of the lower bound techniques presented in this work is that we consider coverings of polytopes with slabs, rather than *recursive* coverings. That is, if we branch on $ax \geq b$ then any further query that we do on the branch $ax \leq b - 1$ will not affect the points on the branch $ax \geq b$. Thus our method is overstating the number of points ruled out by each slab. In treelike proof systems where proofs can be balanced and depth lower bounds give size lower bounds [24, 3, 4, 23] such a recursive method can be approached through the Prover-Delayer game-theoretic tool [24] or generalizations of this game [4, 5]. Proving stronger direct lower bounds on Stabbing Planes by recursive methods is a direction for further research left open in this work.

References

1 Noga Alon and Zoltán Füredi. Covering the cube by affine hyperplanes. *Eur. J. Comb.*, 14(2):79–83, 1993. doi:10.1006/eujc.1993.1011.

2 Paul Beame, Noah Fleming, Russell Impagliazzo, Antonina Kolokolova, Denis Pankratov, Toniann Pitassi, and Robert Robere. Stabbing planes. In Anna R. Karlin, editor, *9th Innovations in Theoretical Computer Science Conference, ITCS 2018, January 11-14, 2018, Cambridge, MA, USA*, volume 94 of *LIPIcs*, pages 10:1–10:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.ITCS.2018.10.

3 Eli Ben-Sasson, Russell Impagliazzo, and Avi Wigderson. Near optimal separation of tree-like and general resolution. *Comb.*, 24(4):585–603, 2004. doi:10.1007/s00493-004-0036-5.

4 Olaf Beyersdorff, Nicola Galesi, and Massimo Lauria. A lower bound for the pigeonhole principle in tree-like resolution by asymmetric prover-delayer games. *Inf. Process. Lett.*, 110(23):1074–1077, 2010. doi:10.1016/j.ipl.2010.09.007.

5 Olaf Beyersdorff, Nicola Galesi, and Massimo Lauria. A characterization of tree-like resolution size. *Inf. Process. Lett.*, 113(18):666–671, 2013. doi:10.1016/j.ipl.2013.06.002.

6 Joshua Buresh-Oppenheim, Nicola Galesi, Shlomo Hoory, Avner Magen, and Toniann Pitassi. Rank bounds and integrality gaps for cutting planes procedures. *Theory of Computing*, 2(4):65–90, 2006.

7 Teena Carroll, Joshua Cooper, and Prasad Tetali. Counting antichains and linear extensions in generalizations of the boolean lattice, 2009.

8 W. Cook, C. R. Coullard, and G. Turán. On the complexity of cutting-plane proofs. *Discrete Appl. Math.*, 18(1):25–38, 1987. doi:10.1016/0166-218X(87)90039-4.

9 Daniel Dadush and Samarth Tiwari. On the complexity of branching proofs. In Shubhangi Saraf, editor, *35th Computational Complexity Conference, CCC 2020, July 28-31, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 169 of *LIPIcs*, pages 34:1–34:35. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.CCC.2020.34.

10 Martin Davis, George Logemann, and Donald W. Loveland. A machine program for theorem-proving. *Commun. ACM*, 5(7):394–397, 1962. doi:10.1145/368273.368557.

11 Martin Davis and Hilary Putnam. A computing procedure for quantification theory. *J. ACM*, 7(3):201–215, 1960. doi:10.1145/321033.321034.

12 Noah Fleming, Mika Göös, Russell Impagliazzo, Toniann Pitassi, Robert Robere, Li-Yang Tan, and Avi Wigderson. On the power and limitations of branch and cut. In Valentine Kabanets, editor, *36th Computational Complexity Conference, CCC 2021, July 20-23, 2021, Toronto, Ontario, Canada (Virtual Conference)*, volume 200 of *LIPIcs*, pages 6:1–6:30. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.CCC.2021.6.

13 Noah Fleming, Denis Pankratov, Toniann Pitassi, and Robert Robere. Random Θ(log n)-cnfs are hard for cutting planes. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 109–120. IEEE Computer Society, 2017. doi:10.1109/FOCS.2017.19.

14 Nicola Galesi, Pavel Pudlák, and Neil Thapen. The space complexity of cutting planes refutations. In David Zuckerman, editor, *30th Conference on Computational Complexity, CCC 2015, June 17-19, 2015, Portland, Oregon, USA*, volume 33 of *LIPIcs*, pages 433–447. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPIcs.CCC.2015.433.

15 Pavel Hrubes and Pavel Pudlák. Random formulas, monotone circuits, and interpolation. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 121–131. IEEE Computer Society, 2017. doi:10.1109/FOCS.2017.20.

16 Russell Impagliazzo, Toniann Pitassi, and Alasdair Urquhart. Upper and lower bounds for tree-like cutting planes proofs. In *Proceedings Ninth Annual IEEE Symposium on Logic in Computer Science*, pages 220–228. IEEE, 1994.

17 Henry A. Kautz and Bart Selman. Ten challenges redux: Recent progress in propositional reasoning and search. In Francesca Rossi, editor, *Principles and Practice of Constraint Programming - CP 2003, 9th International Conference, CP 2003, Kinsale, Ireland, September 29 - October 3, 2003, Proceedings*, volume 2833 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2003. doi:10.1007/978-3-540-45193-8_1.

18 Arist Kojevnikov. Improved lower bounds for tree-like resolution over linear inequalities. In João Marques-Silva and Karem A. Sakallah, editors, *Theory and Applications of Satisfiability Testing - SAT 2007, 10th International Conference, Lisbon, Portugal, May 28-31, 2007, Proceedings*, volume 4501 of *Lecture Notes in Computer Science*, pages 70–79. Springer, 2007. doi:10.1007/978-3-540-72788-0_10.

19 Jan Krajícek. Discretely ordered modules as a first-order extension of the cutting planes proof system. *J. Symb. Log.*, 63(4):1582–1596, 1998. doi:10.2307/2586668.

20 Jan Krajícek. Interpolation by a game. *Math. Log. Q.*, 44:450–458, 1998. doi:10.1002/malq.19980440403.

21 Nathan Linial and Jaikumar Radhakrishnan. Essential covers of the cube by hyperplanes. *Journal of Combinatorial Theory, Series A*, 109(2):331–338, 2005.

22 Lutz Mattner and Bero Roos. Maximal probabilities of convolution powers of discrete uniform distributions. *Statistics & probability letters*, 78(17):2992–2996, 2008.

23 Fedor Part and Iddo Tzameret. Resolution with counting: Dag-like lower bounds and different moduli. *Comput. Complex.*, 30(1):2, 2021. doi:10.1007/s00037-020-00202-x.

24 Pavel Pudlák and Russell Impagliazzo. A lower bound for DLL algorithms for k-sat (preliminary version). In David B. Shmoys, editor, *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms, January 9-11, 2000, San Francisco, CA, USA*, pages 128–136. ACM/SIAM, 2000. URL: `http://dl.acm.org/citation.cfm?id=338219.338244`.

25 Mark Nicholas Charles Rhodes. On the chvátal rank of the pigeonhole principle. *Theor. Comput. Sci.*, 410(27-29):2774–2778, 2009. `doi:10.1016/j.tcs.2009.03.035`.

26 John Alan Robinson. A machine-oriented logic based on the resolution principle. *J. ACM*, 12(1):23–41, 1965. `doi:10.1145/321250.321253`.

27 Bart Selman, Henry A. Kautz, and David A. McAllester. Ten challenges in propositional reasoning and search. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence, IJCAI 97, Nagoya, Japan, August 23-29, 1997, 2 Volumes*, pages 50–54. Morgan Kaufmann, 1997. URL: `http://ijcai.org/Proceedings/97-1/Papers/008.pdf`.

28 Alasdair Urquhart. Hard examples for resolution. *J. ACM*, 34(1):209–219, 1987. `doi:10.1145/7531.8928`.

29 Jacobus Hendricus van Lint and Richard Michael Wilson. *A Course in Combinatorics*. Cambridge University Press, Cambridge, U.K.; New York, 2001.

30 Gal Yehuda and Amir Yehudayoff. A lower bound for essential covers of the cube. *CoRR*, abs/2105.13615, 2021. `arXiv:2105.13615`.

31 Gal Yehuda and Amir Yehudayoff. A lower bound for essential covers of the cube. *arXiv preprint*, 2021. `arXiv:2105.13615`.

A Lower bound for the Least Ordering Principle

▶ **Definition 42.** *Let $n \in \mathbb{N}$. The* Least Ordering Principle, LOP_n*, is the following set of unsatisfiable linear inequalities over the variables $P_{i,j}$ ($i \neq j \in [n]$):*

$$P_{i,j} + P_{j,i} = 1 \quad \text{for all } i \neq j \in [n]$$
$$P_{i,k} - P_{i,j} - P_{j,k} \geq 1 \text{ for all } i \neq j \neq k \in [n]$$
$$\sum_{i=1, i\neq j}^{n} P_{i,j} \geq 1 \text{ for all } j \in [n]$$

▶ **Lemma 43.** *For any $X \subseteq [n]$ of size at most $n-3$, there is an admissible point for LOP_n integer on any edge mentioning an element in X.*

Proof. Let $\preceq$ be any total order on the elements in X. Our admissible point x will be

$$x(P_{i,j}) = \begin{cases} 1 & \text{if } i,j \in X \text{ and } i \preceq j \text{, or if } i \notin X, j \in X \\ 0 & \text{if } i,j \in X \text{ and } j \preceq i \text{, or if } i \in X, j \notin X \\ 1/2 & \text{otherwise (if } i,j \notin X). \end{cases}$$

The existential axioms $\sum_{i=1,i\neq j}^{n} P_{i,j}$ are always satisfied - if $j \in X$ then there is some $i \notin X$ with $P_{i,j} = 1$, and otherwise there are at least two distinct $i,k \neq j \in X$ with $P_{i,j}, P_{k,j} = 1/2$. For the transitivity axioms $P_{i,k} - P_{i,j} - P_{j,k} \geq 1$, note that if 2 or more of i,j,k are not in X there are at least 2 variables set to $1/2$, and otherwise it is set in a binary fashion to something consistent with a total order. ◀

We will assume that a SP refutation $\mathcal{T}$ of LOP_n only involves variables $P_{i,j}$ where $i < j$ - this is without loss of generality as we can safely set $P_{j,i}$ to $1 - P_{i,j}$ whenever $i > j$, and will often write $P_{\{i,j\}}$ for such a variable. We consider the underlying graph of the support of a query, i.e. an undirected graph with edges $\{i,j\}$ for every variable $P_{\{i,j\}}$ that appears with non-zero coefficient in the query.

For some function $f(n)$, we say the query is $f(n)$-*wide* if the smallest edge cover of its graph has at least $f(n)$ nodes . A query that is not $f(n)$-wide is $f(n)$-*narrow*. The next lemma works much the same as Theorem 18.

▶ **Lemma 44.** *Fix $\epsilon > 0$ and suppose we have some* SP *refutation $\mathcal{T}$ of* LOP_n*, where $|\mathcal{T}| \leq n^{\frac{1-\epsilon}{4}}$. Then, if n is large enough, we can find some* SP *refutation $\mathcal{T}'$ of* LOP_{cn}*, where c is a positive universal constant that may be taken arbitrarily close to 1, $\mathcal{T}'$ contains only $n^{3/4}$-wide queries, and $|T'| \leq |\mathcal{T}|$.*

Proof. We iteratively build up an initially empty restriction ρ. At every stage ρ imposes a total order on some subset $X \subseteq [n]$ and places the elements in X above the elements not in X. So ρ sets every edge not contained entirely in $[n] \setminus X$ to something binary, and $\mathsf{LOP}_n\!\restriction_\rho = \mathsf{LOP}_{n-|X|}$ (up to a renaming of variables).

While there exists a $n^{3/4}$-narrow query $q \in \mathcal{T}\!\restriction_\rho$ we simply take its smallest edge cover, which has size at most $n^{3/4}$ by definition, and add its nodes in any fashion to the total order in ρ. Now all of the variables mentioned by $q \in \mathcal{T}\!\restriction_\rho$ are fully evaluated and q is redundant. We repeat this at most $n^{\frac{1-\epsilon}{4}}$ times (as $|\mathcal{T}| \leq n^{\frac{1-\epsilon}{4}}$ and each iteration renders at least one query in $\mathcal{T}$ redundant). At each stage we grow the domain of the restriction by at most $n^{3/4}$, so the domain of ρ is always bounded by $n^{1-\epsilon/4}$. We also cannot exhaust the tree $\mathcal{T}$ in this way, as otherwise $\mathcal{T}$ mentioned at most $n^{1-\epsilon/4} < n-3$ elements and by Lemma 43 there is an admissible point not falling in any slab of $\mathcal{T}$, violating Fact 4.

When this process finishes we are left with a $n^{3/4}$-wide refutation $\mathcal{T}'$ of $\mathsf{LOP}_{n-n^{1-\epsilon/4}}$. As ϵ was fixed we find that as n goes to infinity $n - n^{1-\epsilon/4}$ tends to n. ◀

▶ **Lemma 45.** *Let $d \leq (n-3)/2$. Given any disjoint set of pairs $D = \{\{l_1, r_1\}, \ldots, \{l_d, r_d\}\}$ (where WLOG $l_i < r_i$ in $[n]$ as natural numbers) and any binary assignment $b \in \{0,1\}^D$, the assignment x_b with*

$$x_b(P_{\{i,j\}}) = \begin{cases} b(\{l_k, r_k\}) & \text{if } \{i,j\} = \{l_k, r_k\} \in X \text{ for some } k \\ 1/2 & \text{otherwise} \end{cases}$$

is admissible.

Proof. The existential axioms $\sum_{i=1, i\neq j}^{n} P_{i,j}$ are always satisfied, as for any j there are at least $n-2$ $i \in [n]$ different from j with $P_{i,j} = 1/2$. For the transitivity axioms $P_{i,k} - P_{i,j} - P_{j,k} \geq 1$, note that due to the disjointness of D at least two variables on the left hand side are set to 1/2. ◀

▶ **Theorem 46.** *Fix some $\epsilon > 0$ and let $\mathcal{T}$ any* SP *refutation of* LOP_n*. Then, for n large enough, $|\mathcal{T}| \in \Omega(n^{\frac{1-\epsilon}{4}})$.*

Proof. Suppose otherwise - then, by Lemma 44, we can find some $\mathcal{T}'$ refuting LOP_{cn}, with $|\mathcal{T}'| \leq |\mathcal{T}|$, every query $n^{3/4}$-wide, and c independent of n. We greedily create a set of pairs D by processing the queries in $\mathcal{T}'$ one by one and choosing in each a matching of size $n^{1/2}$ disjoint from the elements appearing in D - this always succeeds, as at every stage $|D| \in O(n^{\frac{1-\epsilon}{4}} \cdot n^{1/2})$ and involves at most $O(2n^{\frac{3-\epsilon}{4}}) < n^{3/4} - n^{1/2}$ elements.

So by Lemma 45, after setting every edge not in D to 1/2, we have some set of linear polynomials $\mathcal{R} = \{a(x) = \mathbf{a}x - b - 1/2 : (\overline{a}, b) \in \mathcal{T}'\}$ covering the hypercube $\{0,1\}^D$, where every polynomial $p \in \mathcal{R}$ mentions at least $n^{1/2}$ edges. By Lemma 11 each such polynomial in $\mathcal{R}$ rules out at most $2^{|D|}/n^{1/4}$ points, and so we must have $|\mathcal{T}| \geq |T'| \geq |\mathcal{R}| \geq n^{1/4}$. ◀