Stabbing Planes

Paul Beame
University of Washington
beame@cs.washington.edu

Noah Fleming*
University of Toronto
noahfleming@cs.toronto.edu

Russell Impagliazzo
University of California, San Diego
russell@cs.ucsd.edu

Antonina Kolokolova[†]
Memorial University of Newfoundland
kol@mun.ca

Denis Pankratov[†]
University of Toronto
denisp@cs.toronto.edu

Toniann Pitassi[†]
University of Toronto
toni@cs.toronto.edu

Robert Robere[†]
University of Toronto
robere@cs.toronto.edu

October 6, 2017

Abstract

We introduce and develop a new semi-algebraic proof system, called Stabbing Planes that is in the style of DPLL-based modern SAT solvers. As with DPLL, there is only one rule: the current polytope can be subdivided by branching on an inequality and its "integer negation." That is, we can (nondeterministically choose) a hyperplane $ax \geq b$ with integer coefficients, which partitions the polytope into three pieces: the points in the polytope satisfying $ax \geq b$, the points satisfying $ax \leq b - 1$, and the middle slab $b - 1 < ax < b$. Since the middle slab contains no integer points it can be safely discarded, and the algorithm proceeds recursively on the other two branches. Each path terminates when the current polytope is empty, which is polynomial-time checkable. Among our results, we show somewhat surprisingly that Stabbing Planes can efficiently simulate Cutting Planes, and moreover, is strictly stronger than Cutting Planes under a reasonable conjecture. We prove linear lower bounds on the *rank* of Stabbing Planes refutations, by adapting a lifting argument in communication complexity.

1 Introduction

While defined in terms of non-deterministic algorithms for the Tautology problem, proof complexity has also provided indispensible tools for understanding deterministic algorithms for search problems, and in particular, for Satisfiability algorithms. Many algorithms for search can be classified according to the types of reasoning they implicitly use for case-analysis and pruning unpromising branches. Particular families of search algorithms can be characterized by *formal proof systems*; the size of proofs in these formal proof system, the time of the non-deterministic algorithm, captures the time taken on the instance by an *ideal implementation* of the search algorithm. This allows us to factor understanding the power of search algorithms of a given type into two questions:

1. How powerful is the proof system? For which kinds of input are there small proofs?

2. How close can actual implementations of the search method come to the ideal non-deterministic algorithm?

As an illustrative example, let us recall the *DPLL* algorithm [9, 10], which is one of the simplest algorithms for SAT and forms the basis of modern conflict-driven clause learning SAT solvers. Let $\mathcal{F} = C_1 \wedge C_2 \wedge \cdots \wedge C_m$ be a CNF formula over variables $x_1, x_2, \ldots, x_n$. A DPLL search tree for solving the SAT problem for $\mathcal{F}$ is constructed as follows. Begin by choosing a variable x_i (non-deterministically, or via some heuristic), and then recurse on the formulas $\mathcal{F} \restriction x_i = 0$, $\mathcal{F} \restriction x_i = 1$. If at any point we have found a satisfying assignment, the algorithm outputs *SAT*. Otherwise, if we have falsified every literal in some clause C, then we record the clause and halt the recursion. If every recursive branch ends up being labelled with a clause and a falsifying assignment, then the original formula $\mathcal{F}$ is unsatisfiable and one can take the tree as a proof of this fact; in fact, such a DPLL tree is equivalent to a *tree-like* Resolution refutation of the formula $\mathcal{F}$.

Modern SAT solvers still have a DPLL algorithm at the core (now with a highly tuned *branching heuristic* that chooses the "right" order for variables and assignments to recurse on in the search tree), but extends the basic recipe in two ways:

1. Smart handling of *unit clauses*: if $\mathcal{F}$ contains a clause C with a single variable under the current partial assignment, then we can immediately assign the variable so that the clause is satisfied.

2. *Clause learning* to speed up search: if a partial assignment ρ falsifies a clause, then the algorithm derives a new clause C_ρ by a resolution proof that "explains" this conflict, and adds the new clause to the formula $\mathcal{F}$.

It is the synergy between these three mechanisms — branching heuristics, unit propagation, and clause learning — that results in the outstanding performance of modern SAT solvers. In other words, while these algorithms are all formalizable in the same simple proof system, the sophistication of modern SAT-solvers comes from the attempt to algorithmically find small proofs when they exist. In many ways, the simplicity of the proof system enables this sophistication in proof-search methods.

In this work, we introduce a natural generalization of the DPLL-style branching algorithm to reasoning over integer-linear inequalities, formalized as a new semi-algebraic proof system that we call the *Stabbing Planes* (SP) system. We will give a more detailed description later, but intuitively, Stabbing Planes has the same branching structure as DPLL, but generalizes branching on single variables to branching on linear inequalities over the variables. We feel the closeness to DPLL makes Stabbing Planes a better starting point for understanding search algorithms based on linear inequalities, such as most integer linear programming(ILP) based solvers, than established proof systems using such inequalities such as Cutting Planes proofs.

We compare the power of Stabbing Planes proofs to these other proof systems. We show that one of these systems, Kraijicek's system tree-like $R(CP)$, is polynomially equivalent to Stabbing Planes (Theorem 4.2). However, the new formulation as Stabbing Planes proofs both gives greater motivation to studying $R(CP)$ and greatly clarifies the power of the proof system. Our main results about this system are:

1. Stabbing planes has quasi-polynomial size and poly-log rank proofs of any tautology provable using linear algebra over a constant modulus. In particular, this is true for the Tseitin graph tautologies, that are very frequently used to prove lower bounds for other proof systems. (Theorem 3.1)

2. Stabbing planes can simulate tree-like Cutting Planes proofs with only constant factor increases in size and rank (Theorem 4.4)

3. Stabbing planes can simulate general Cutting Planes proofs with a polynomial increase in size (Theorem 4.5)

4. Lower bounds on real communication protocols imply rank lower bounds for Stabbing Planes proofs (Lemma 5.3)

5. Stabbing planes proofs cannot be balanced (Theorem 5.7)

Together, these show that Stabbing Planes is at least as strong as established proof systems using inequalities, and possibly much stronger. So the proof system combines strength as a proof system with a simple branching structure that raises the possibility of elegant algorithms based on this proof system.

We now give a more precise description of the proof system. Let us formalize the system in stages. First, observe that the setting is quite different: we are given a system $A_1 x \geq b_1, A_2 \cdot x \geq b_2, \ldots, A_m \cdot x \geq b_m$ of integer-linear inequalities over real-valued variables $x_1, x_2, \ldots, x_n$ (for simplicity we will always assume that the inequalities $0 \leq x_i \leq 1$ are present for each variable x_i), and we seek to prove that no $\{0, 1\}$-solution exists. One can immediately see that the basic DPLL algorithm immediately works in this setting with little modification: one can still query variables and assign them to $\{0, 1\}$ values; now we label leaves of the search tree with any inequality $a_i \cdot x \geq b_i$ in the system that is falsified by the sequence of assignments made on the path from the root to the leaf. If every leaf ends up being labelled with a falsified inequality, then the tree certifies that the system of inequalities has no $\{0, 1\}$-solutions.

3

However, with the expanded domain we can consider the DPLL tree *geometrically*. To be more specific, imagine replacing each $\{0, 1\}$ query to a variable x_i in the decision tree with two "inequality queries" $x_i \leq 0$ and $x_i \geq 1$. Each node u in the tree after this replacement is now naturally associated with a polyhedral set $\mathcal{P}_u$ of points satisfying each of the the input inequalities *and* each of the inequalities on the path to this node. Since we began with a DPLL refutation, it is clear that for any leaf ℓ the polyhedral set $\mathcal{P}_\ell$ associated with the leaf is empty, as any $\{0, 1\}$ solution would have survived each of the inequalities queried on some path in the tree and thus would exist in one of the polyhedral sets at the leaves.

The stabbing planes system is then the natural generalization of the previous object: an SP refutation consists of a generalized DPLL refutation where each node is labelled with an *arbitrary* integral linear inequality $Ax \geq b$ (that is, the vector A and the parameter b are both integral), and the outgoing edges are labelled with the inequalities $Ax \geq b$ and $Ax \leq b - 1$. Clearly, any integral vector $x \in \mathbb{Z}^n$ will satisfy at least one of the inequalities labelling the outgoing edges, and so if the polyhedral set at each leaf (again, obtained by intersecting the original system with the inequalities on the path to the leaf) is empty then we have certified that the original system of inequalities has no integral solutions. (See Figure 1 for a simple example.) The main innovation of Stabbing Planes is its simplicity: refutations are simply


Figure 1: A partial SP refutation and the result on the unit square. The shaded areas are "removed" from the polytope, and we recurse on each side.

decision trees that query linear inequalities. Note that the more obvious extension of DPLL to linear inequalities would branch on $Ax \geq b$ and its *actual* negation, $Ax < b$. However with this branching rule, we would have to add additional rules in order to have completeness. By branching on an inequality and its "integer negation", we are able to get by with just one rule analogous to the resolution rule in DPLL.

From the perspective of SAT solving, even though tree-like Resolution and the search for satisfying assignments encapsulated by DPLL are equivalent, it is the search point of view of DPLL that has led to the major advances in SAT algorithms now found in modern conflict-directed clause learning (CDCL) SAT solvers. A natural hypothesis is that it is much easier to invent useful heuristics in the language of query-based algorithms, as opposed to algorithms based on the resolution rule. Stabbing Planes offers a similar benefit with respect to reasoning about inequalities.

With the exception of mixed integer programming (MIP) solvers (such as CPLEX [19]), current solvers that, like Stabbing Planes, manipulate integer linear inequalities over Boolean variables are generally built on the same backtracking-style infrastructure as DPLL and CDCL SAT solvers but maintaining information as integer linear inequalities as opposed to clausal forms. The solvers are known as *pseudoBoolean* solvers and have been the subject of considerable effort and development.

PseudoBoolean solvers work very well at handling the kinds of symmetric counting problems associated with, for example, the pigeonhole principle (PHP), which is known to be hard for CDCL SAT solvers, as well as other problems where the input constraints are much more succinctly and naturally expressed in inequality rather than clausal form. Innovations in pseudoBoolean solvers include use of normal forms for expressing constraints, techniques to generalize fast unit propogation and watched literals from DPLL to the analogue for integer linear inequalities, as well as methods to learn from conflicts and simplify learned constraints when integer coefficients from derived inequalities get too large [4, 30]. Despite all of this, even for the best pseudoBoolean solvers, the benefits of expressibility are usually not enough compensation for the added costs of manipulating and deriving new inequalities and they outperform CDCL solvers only in very limited cases in practice [4].

A key limitation of these pseudoBoolean solvers, which likely constrains their effectiveness, is the fact that all branching is based on assigning values to individual variables; i.e., dividing the problem by slabs parallel to one of the coordinate axes. Stabbing Planes eliminates this constraint on the search and allows one to apply a divide and conquer search based on arbitrary integer linear constraints that are not necessarily aligned with one of these coordinate axes. This opens up the space of algorithmic ideas considerably and should allow future pseudoBoolean solvers to take fuller advantage of the expressibility of integer linear constraints. For example, a Stabbing Planes search could choose to branch on a linear inequality that is derived from the geometric properties of the rational hull of the current constraints by, say, splitting the volume, or by doing a balanced split at a polytope vertex, since properties of the rational hull can be determined efficiently. Such operations could be potentially be done in conjunction with solvers such as CPLEX to obtain the best of both kinds of approaches.

Beyond the prospect of Stabbing Planes yielding improved backtracking search for pseudoBoolean solvers, Stabbing Planes should allow the same kind of learning of inequalities from conflicts that is being done in existing pseudoBoolean solvers, and hence get the benefits of both. In this work we do not focus on the theoretical benefits of learning from conflicts because we already can show considerable theoretical benefit from the more general branching alone and because the theoretical benefits of the restricted kinds of learned linear inequalities from conflicts available even in existing solvers are not at all clear.

From a proof complexity perspective, the SP system turns out to be polynomially equivalent to the semi-algebraic proof system *tree-like* R(CP), introduced by Krajíček [24]. Roughly speaking one can think of R(CP) as a mutual generalization of both Cutting Planes and Resolution — the lines of an R(CP) proof are clauses of integer linear inequalities, and in a single step one can take two clauses and either apply a cutting-planes rule to a single inequality in each clause or apply a resolution-style "cut". However, even though SP turns out to be equivalent to a system already in the literature, this new perspective turns out to be quite useful: we show that SP has quasi-polynomial size refutations of the Tseitin principle, and also that SP can polynomially-simulate Cutting Planes (neither result was previously known to hold for tree-like R(CP)).

We also investigate the relationship between SP refutations and communication complexity. Given an unsatisfiable CNF $\mathcal{F}$ and any partition of the variables (X, Y) of $\mathcal{F}$ into two sets, one can consider the following two-party search problem $\mathsf{Search}_{X,Y}(\mathcal{F})$: Alice receives an assignment to the X-variables, Bob receives an assignment to the Y-variables, and they must communicate and output a falsified clause of $\mathcal{F}$ under their joint assignment. At this time *all* strong lower bound results for Cutting Planes refutations essentially follow from studying the communication complexity of $\mathsf{Search}_{X,Y}(\mathcal{F})$. For instance, it is known that:

- A depth-d Cutting Planes refutation of $\mathcal{F}$ yields a d-round real communication protocol for $\mathsf{Search}_{X,Y}(\mathcal{F})$ [25].

- A length-L tree-like Cutting Planes refutation of $\mathcal{F}$ yields an $O(\log L)$-round real communication protocol for $\mathsf{Search}_{X,Y}(\mathcal{F})$ [5, 25, 28].

- A length-L, space s Cutting Planes refutation of $\mathcal{F}$ yields an $O(s \log L)$-round real communication protocol for $\mathsf{Search}_{X,Y}(\mathcal{F})$ [11].

- A length-L Cutting Planes refutation of $\mathcal{F}$ yields a size L real communication game for $\mathsf{Search}_{X,Y}(\mathcal{F})$ [16].

Each of these results has been used to derive strong lower bounds on Cutting Planes refutations by proving the corresponding lower bound against the search problem [5, 11, 13, 17, 25, 28]. Furthermore, the above lower bound techniques hold even for the stronger *semantic* Cutting Planes system (the lines of which are integer linear inequalities, and from two integer linear inequalities we are allowed to make *any* sound deduction over integer points) [12]. This makes the known lower bounds much stronger, and it is quite surprising that all one needs to exploit for strong lower bounds is that the lines are linear inequalities (rather than exploiting some weakness of the deduction rules). However, this strength also illustrates a weakness of current techniques, as once the lines of a proof system $\mathcal{P}$ become expressive enough, semantic proof techniques (i.e. ones that work for the semantic version of the proof systems) completely break down since every tautology has a short semantic proof. Therefore, it is of key importance to develop techniques which truly exploit the "syntax" of proof systems, and not just the expressive power of the lines.

Hence, it is somewhat remarkable that we are able to show that each of the simulation results above still hold if we replace real communication protocols with SP refutations. That is, we show

- A depth-d Cutting Planes refutation of $\mathcal{F}$ yields a depth-d SP refutation of $\mathcal{F}$.

- A length-L tree-like Cutting Planes refutation of $\mathcal{F}$ yields a depth $O(\log L)$ SP refutation of $\mathcal{F}$.

- A length-L, space s Cutting Planes refutation of $\mathcal{F}$ yields a depth $O(s \log L)$ SP-refutation of $\mathcal{F}$.

- A length-L Cutting Planes refutation of $\mathcal{F}$ yields a size $O(L)$ SP-refutation of $\mathcal{F}$.

Since SP is a syntactic system this further motivates studying its depth- and size-complexity. We can use semantic techniques to get some lower bounds for SP: we show that a size-S and depth-d SP refutation yields a real communication protocol with cost $O(d)$ and for which the protocol tree has size $O(S \cdot n)$. This simulation yields new proofs of some depth lower bounds already known in the literature; however, these lower bounds are complemented by showing that neither SP refutations nor real communication protocols can be balanced. This should be viewed in a positive light: the depth- and size-complexity problems are truly different for SP, and furthermore, one seemingly cannot obtain size lower bounds for SP by proving depth lower bounds for real communication protocols (in contrast to, say, tree-like Cutting Planes). In sum, SP appears to be a very good candidate for a proof system on the "boundary" where current techniques fail to prove strong size lower bounds.

The rest of the paper is outlined as follows. After some preliminaries in Section 2, we give a simple refutation of the Tseitin problem in SP in Section 3. In Section 4, we prove a raft of simulation and equivalence results for SP — showing it is equivalent to R(CP), relating it to Cutting Planes in various measures such as depth, length, and space, and showing how an SP proof yields a real communication protocol for the canonical search problem. Finally, in Section 5, we prove depth lower bounds for SP and some impossibility results for balancing.

2 Preliminaries

Before we define the new proof system formally, we need to make a few general definitions that are relevant to semi-algebraic proof systems.

An *integer linear inequality* (or simply a *linear inequality*) in the variables $x = x_1, \ldots, x_n$ is $Ax \geq b$, where $A \in \mathbb{Z}^n$ and $b \in \mathbb{Z}$. A system of linear inequalities $\mathcal{F}$ is *unsatisfiable* if there is no Boolean assignment $\alpha \in \{0,1\}^n$ which simultaneously satisfies every inequality in $\mathcal{F}$. We sometimes refer to inequalities as lines and write $L \equiv Ax \geq b$. The integer negation of a line L is the inequality $\neg L \equiv Ax \leq b - 1$.

An unsatisfiable formula in a conjuctive normal form (CNF) defines an unsatisfiable system of linear inequalities $\mathcal{F}$ in a natural way. A clause $\bigvee_{i=1}^{k} x_i \vee \bigvee_{i=1}^{l} \neg x_i$, is translated into the inequality $\sum_{i=1}^{k} x_i + \sum_{i=1}^{l}(1 - x_i) \geq 1$, and $\mathcal{F}$ is the set of translations of all clauses. We assume that $\mathcal{F}$ always contains the axioms $x_i \geq 0$ and $-x_i \geq -1$ for all variables x_i, as we are interested in propositional proof systems for refuting unsatisfiable Boolean formulas.

Definition 2.1. A *propositional proof system* $\mathcal{P}$ is a non-deterministic polynomial time Turing machine (TM) deciding the language of unsatisfiable CNF formulas. Given an unsatisfiable CNF, the NP-certificate is called *the proof* or *the refutation*.

To compare the strength of proof systems, one typically uses the notion of polynomial simulation.

Definition 2.2. Let $\mathcal{P}_1$ and $\mathcal{P}_2$ be two proof systems. We say that $\mathcal{P}_1$ *polynomially simulates* $\mathcal{P}_2$ if for every unsatisfiable formula $\mathcal{F}$, the shortest refutation of $\mathcal{F}$ in $\mathcal{P}_1$ is at most polynomially longer than the shortest refutation in $\mathcal{P}_2$. $\mathcal{P}_1$ is *strictly stronger* than $\mathcal{P}_2$ if $\mathcal{P}_1$ polynomially simulates $\mathcal{P}_2$, but the converse does not hold. Finally, we say that $\mathcal{P}_1$ and $\mathcal{P}_2$ are *incomparable* if neither can polynomially simulate the other.

We now describe the proof system Stabbing Planes, which will be the central object of study of this paper.

Definition 2.3. Let $\mathcal{F}$ be an unsatisfiable system of linear integral inequalities. A *Stabbing Planes (*SP*) refutation* of $\mathcal{F}$ is a threshold decision tree — a directed binary tree in which each edge is labelled with a linear integral inequality. If the right outgoing edge of a certain node is labelled with $Ax \geq b$, then the left outgoing edge has to be labelled with its integer negation, $Ax \leq b - 1$. We refer to Ax (or the pair of inequalities $Ax \leq b - 1, Ax \geq b$) as *the query* corresponding to the node. The *slab* corresponding to the query is $\{x^* \in \mathbb{R}^n \mid b - 1 < Ax^* < b\}$.

Let the set of all paths from the root to a leaf in the tree be denoted by $\{p_1, \ldots, p_\ell\}$. Each leaf i is labelled with a non-negative linear combination of inequalities in $\mathcal{F}$ with the inequalities along the path p_i that yields $0 \geq 1$.

The *size* of a SP refutation is the number of bits needed to represent every inequality in the refutation. The *length* of a SP refutation is the number of nodes in the threshold tree. The size (length) of refuting a system of linear inequalities $\mathcal{F}$ in SP is the minimum size (length) of any SP refutation of $\mathcal{F}$. The *rank* or *depth* of a SP refutation $\mathcal{P}$ is the longest root-to-leaf path in the threshold tree of $\mathcal{P}$. The rank (depth) of refuting $\mathcal{F}$ in SP is the minimum rank (depth) over all SP refutations of $\mathcal{F}$.

Refutations in SP have an intuitive geometric interpretation: each step of a refutation can be viewed as nondeterministically removing a slab from the solution space and recursing on the resulting polytopes on both sides of the slab. The aim is to recursively cover the solution space with slabs until every feasible point within this polytope is removed. An example of this can be seen in Figure 1 in the introduction. In particular, the polytope at any step of the recursion is empty if and only if there exists a convex combination of the axioms and inequalities labelling the corresponding root-to-leaf path in the refutation equivalent to $0 \geq 1$. This is summarized in the following fact which follows directly from the Farkas' lemma. The "moreover" part of the following fact is an application of Carathéodory's theorem, and will be useful for technical reasons later in the paper. We refer the interested reader to [31] for some background on polytope theory.

Fact 2.4. *Let $\mathcal{F} = \{A_1 x \geq b_1, \ldots, A_m x \geq b_m\}$ be a system of integer linear inequalities. The polytope defined by $\mathcal{F}$ is empty if and only if there is a non-negative (rational) linear combination of the inequalities of $\mathcal{F}$ which evaluates to $0 \geq 1$. Moreover, the non-negative linear combination can be taken to be supported on $\leq n$ of the inequalities from $\mathcal{F}$, where n is the dimension of the space to which x belongs.*

It is straightforward to see that SP is a sound and complete proof system. Completeness follows from a simple observation that SP polynomially simulates DPLL. To see that SP is sound, let $\mathcal{R}$ be a SP refutation of some formula $\mathcal{F}$. Observe that for any node in $\mathcal{R}$ with outgoing edges labelled $Ax \geq b$ and $Ax \leq b - 1$, any $0 - 1$ assignment to the variables $\alpha \in \{0, 1\}^n$ must satisfy exactly one of the two inequalities. Therefore, if a Boolean solution $\alpha \in \{0, 1\}^n$ satisfies $\mathcal{F}$, then for at least one of the leaves of $\mathcal{R}$, one cannot derive $0 \geq 1$. This follows by Fact 2.4 because the polytope formed by the inequalities labelling root-to-leaf path is non-empty (α lies in this polytope) .

Next we recall a well-known and extensively-studied proof system: Cutting Planes (CP). For an introduction to Cutting Planes, we refer an interested reader to Chapter 19 in [21].

Definition 2.5. Let $\mathcal{F}$ be an unsatisfiable system of linear inequalities. A *Cutting Planes* (CP) *refutation* of $\mathcal{F}$ is a sequence of linear inequalities $\{L_1, \ldots, L_\ell\}$ such that $L_\ell = 0 \geq 1$ and each L_i is either an axiom $\in \mathcal{F}$ or is obtained from previous lines via one of the following inference rules. Let α, β be positive integers.

$$\text{Linear Combination: } \frac{Ax \geq a \quad Bx \geq b}{(\alpha A + \beta B)x \geq \alpha a + \beta b} \quad \text{Division: } \frac{\alpha Ax \geq b}{Ax \geq \left\lceil \frac{b}{\alpha} \right\rceil}$$

We refer to ℓ as the *length* of the refutation. The *length* of refuting $\mathcal{F}$ in CP is the minimum length of a CP refutation of $\mathcal{F}$.

The directed acyclic graph (DAG) $G = (V, E)$ associated with a CP refutation $\{L_1, \ldots, L_\ell\}$ is defined as follows. We have $V = \{L_1, \ldots, L_\ell\}$ and $(u, v) \in E$ if and only if the line labelling v was derived by an application of an inference rule involving the line labelling u. Without loss of generality, we may assume that there is only one vertex with out degree 0, which we call the root. The root of G is labelled with L_ℓ and the leaves are labelled with the axioms.

The *rank* or *depth* of the refutation is the length of the longest root-to-leaf path in G. The rank of refuting an unsatisfiable system of linear inequalities $\mathcal{F}$ is the minimum rank of any refutation of $\mathcal{F}$ in the given proof system. Finally, *tree-like* CP is defined by restricting proofs to be such that the underlying graph G is a tree.

On the Issue of Weights. A well-known theorem of Muroga shows that for any integer linear inequality $Ax \geq b$ separating two subsets $\mathcal{U}, \mathcal{V} \subseteq \{0, 1\}^n$ of 0/1 vectors, there is another linear integer inequality $A'x \geq b'$ separating the same set of vectors and also satisfying $||A||_\infty \leq 2^{\mathsf{poly}(n)}$ [26]. In the same vein, Cook, Coullard, and Turán showed that any Cutting Planes refutation of length ℓ can be transformed into a refutation of length at most polynomially larger and in which each coefficient has magnitude $2^{\mathsf{poly}(\ell,n)}$ [8]. This implies that the size of any CP proof (measured by the length of its encoding in bits) is polynomially related to its

length, and thus we may study the length of cutting planes proofs without loss of generality for the purpose of upper and lower bounds.

An analogous result is not known for SP and therefore we currently must make the distinction between its size and length. Fortunately, all of our results hold in the best possible scenario; our upper bounds are low weight (polynomial-length); the simulations are length preserving, and our lower bounds hold for any weight.

3   Motivating Example: SP **Refutations of Tseitin Formulas**

Tseitin contradictions are among the most well-studied unsatisfiable formulas in proof complexity, and are the quintessential formulas that are believed to be hard for CP [21]. Despite the fact that exponential lower bounds for CP are known for many natural families of formulas (including recent lower bounds for random $O(\log n)$-CNF formulas), there are no nontrivial lower bounds known for the Tseitin contradictions, and for good reason: the only known lower bound method for CP reduces the problem of refuting a formula in CP to a monotone circuit problem, for which the corresponding monotone circuit problem for Tseitin is easy.

In this section, we demonstrate the power of Stabbing Planes by showing that there exists a shallow quasi-polynomial size SP refutation of the Tseitin formulas. This, together with our simulation results from Section 4, strongly suggests that SP is strictly more powerful than CP. In addition, we immediately conclude that SP is provably more powerful than CP in terms of depth.

Tseitin contradictions are any unsatisfiable family of mod-2 equations subject to the constraint that every variable occurs in exactly two equations. An instance of Tseitin, denoted $\text{Tseitin}(G, \ell)$ is defined by a connected undirected graph $G = (V, E)$ and a node labelling $\ell \in \{0, 1\}^V$ of odd total weight: $\sum_{v \in V} \ell_v = 1 \mod 2$. For each edge $e \in E$ there is a variable x_e in $\text{Tseitin}(G, \ell)$. Corresponding to each vertex $v \in V$ is the equation

$$\sum_{e \ni v} x_e \equiv \ell_v \mod 2$$

stating that the sum of the variables x_e incident with v is ℓ_v modulo 2. Summing up all of these equations modulo 2, the left-hand-side sums to zero since every variable occurs exactly twice, but the right-hand-side is one, since the node labelling is odd, and therefore the equations are unsatisfiable. When G has degree D, we can express $\text{Tseitin}(G, \ell)$ as a D-CNF formula containing $|V| \cdot 2^{D-1}$ clauses.

The obvious way to refute $\text{Tseitin}(G, \ell)$ under an assignment x is to find a vertex w for which the corresponding vertex equation is falsified. This can be achieved by the following divide-and-conquer procedure, which maintains a set $U \subseteq V$ such that $w \in U$. The process begins by setting $U = V$. Then, V is partitioned arbitrarily into two sets V_1, V_2 of roughly the same size. Query x_e for all edges e crossing the cut (V_1, V_2), and suppose that the sum of all such x_e is odd (the case when it is even is similar). We know that either $\sum_{v \in V_1} \ell_v$ or $\sum_{v \in V_2} \ell_v$ is even: if the first sum is even then the Tseitin formula restricted to V_1 contains a contradiction, and otherwise the formula restricted to V_2 contains a contradiction. In either case, we can remove roughly half of the graph and recurse.

By keeping track of a few more variables, we can repeat this procedure recursively until $|U| = 1$. Since we reduce the size of U by half each time, this procedure results in the recursion depth logarithmic in $|V|$. It turns out that this procedure can be realized in Stabbing Planes, where recursion depth roughly corresponds to the depth of the refutation. This results in a quasi-polynomial size refutation.

Theorem 3.1. *Let $G = (V, E)$ be an undirected graph, and let ℓ be a $\{0, 1\}$ vertex labelling with odd total weight. Then $\mathrm{Tseitin}(G, \ell)$ has an* SP *refutation of size $n^{O(\log n + D/\log n)}$ and rank $O(D + \log^2 n)$, where $n = |V|$ and D is the maximum degree in G.*

Proof. If $U \subseteq V$ is a set of vertices, then let $E(U) = \{uv \in E \mid u, v \in U\}$, and $\mathrm{Cut}(U) = \{uv \in E \mid u \in U, v \in \overline{U}\}$. Similarly, if $U_1, U_2 \subseteq V$ are disjoint then we let $\mathrm{Cut}(U_1, U_2) = \{uv \in E \mid u \in U_1, v \in U_2\}$. We construct the SP refutation recursively. During the recursion we maintain a set U of current vertices (initially $U = V$). At each recursive step, we split U into two halves U_1 and U_2 and query the total weight k of the edges crossing (U_1, U_2) via SP inequalities. Knowing k, a few additional queries allows us to determine which of U_1 or U_2 contains a contradiction, and we then recurse on the corresponding set of vertices.

More formally, we construct a proof while maintaining the following invariant: for the current subset of vertices $U \subseteq V$, we have queried linear inequalities implying that

$$\sum_{e \in \mathrm{Cut}(U)} x_e = k$$

for some $0 \le k \le |\mathrm{Cut}(U)|$ such that $k \not\equiv \sum_{v \in U} \ell_v \bmod 2$. Note that this invariant ensures that our Tseitin instance restricted to the edges incident on U is unsatisfiable, since summing up all vertex constraints within U yields

$$\sum_{e \in \mathrm{Cut}(U)} x_e + \sum_{e \in E(U)} 2x_e \equiv k \not\equiv \sum_{v \in U} \ell_v \pmod 2.$$

Initialization. Initially we have $U = V$ and the invariant clearly holds.

Recursive Step. Let U be the current set of vertices. By the invariant we know that $\sum_{v \in \mathrm{Cut}(U)} x_e = k$ for some $k \not\equiv \sum_{v \in U} \ell_v \bmod 2$. Partition U into two halves U_1 and U_2 arbitrarily, subject to $|U_1| = \lfloor |U|/2 \rfloor$. We first determine the value of the edges going between U_1 and U_2 by querying

$$\sum_{e \in \mathrm{Cut}(U_1, U_2)} x_e \ge \beta$$

for $\beta = 1, \dots, |\mathrm{Cut}(U_1, U_2)|$. To each leaf of this tree we attach a second binary search tree for determining the value $|\mathrm{Cut}(U_1)|$ by querying $\sum_{e \in \mathrm{Cut}(U_1)} x_e \ge \gamma$ for $\gamma = 1, \dots, |\mathrm{Cut}(U_1)|$. After these queries, at each leaf of the "combined" tree we will have

$$\sum_{e \in \mathrm{Cut}(U_1)} x_e = \gamma \qquad\qquad \text{for some } 0 \le \gamma \le |\mathrm{Cut}(U_1)|,$$

$$\sum_{e \in \mathrm{Cut}(U_1, U_2)} x_e = \beta \qquad\qquad \text{for some } 0 \le \beta \le |\mathrm{Cut}(U_1, U_2)|.$$

11

Furthermore, since $|\text{Cut}(U_1)| + |\text{Cut}(U_2)| = |\text{Cut}(U)| + 2|\text{Cut}(U_1, U_2)|$, we will have

$$\sum_{e \in \text{Cut}(U_2)} x_e = \delta \qquad\qquad \text{for some } 0 \le \delta \le |\text{Cut}(U_2)|,$$

where $\delta + \gamma = k + 2\beta$.

For any leaf of this tree where $\delta > |\text{Cut}(U_2)|$, we can derive a contradiction by summing the axioms $-x_e \ge -1$ for all $e \in \text{Cut}(U_2)$ with $\sum_{e \in \text{Cut}(U_2)} x_e \ge \delta$. Otherwise, for the remaining leaves observe that

$$\delta + \gamma \equiv k \not\equiv \sum_{v \in U_1} \ell_v + \sum_{v \in U_2} \ell_v \pmod 2.$$

Thus, exactly one of the following cases holds:

1. If $\gamma \not\equiv \sum_{v \in U_1} \ell(v) \bmod 2$, then recurse on U_1.

2. Otherwise, $\delta \not\equiv \sum_{v \in U_2} \ell(v) \bmod 2$. Recurse on U_2.

Termination. Our recursion terminates when U contains a single vertex v. By the invariant, we have derived $\sum_{e \in \text{Cut}(v)} x_e \equiv k \not\equiv \ell(v) \bmod 2$ for some $0 \le k \le |\text{Cut}(\{v\})|$. The axioms of $\text{Tseitin}(G, \ell)$ rule out Boolean assignments to the variables x_e for $e \in \text{Cut}(\{v\})$, which contradict $\ell(v)$; these axioms do not prohibit incorrect fractional assignments. Therefore, to derive a contradiction, we still need to enforce that the variables x_e for $e \in \text{Cut}(\{v\})$ take $\{0, 1\}$ values. We achieve this by querying all variables x_e for $e \in \text{Cut}(\{v\})$ via SP inequalities $x_e \ge 1, x_e \le 0$. This results in a complete binary tree of depth $\le D$. Clearly, $0 \ge 1$ is immediately obtained at the leaves that disagree with $\sum_{e \in \text{Cut}(\{v\})} x_e = k$. At the leaves that agree with $\sum_{e \in \text{Cut}(\{v\})} x_e = k$, the inequality $0 \ge 1$ immediately follows from the assignment to the edges incident to v and one of the axioms of Tseitin.

Finally, we analyze the size and rank of the constructed SP proof. In each recursive step, we make $O((nD)^2)$ queries to determine weights of edges crossing the two cuts. Each recursive step is computed by a pair of binary trees, each of depth at most $\log(nD)$. Our recursion terminates in $\log n$ rounds because we halve the number of current vertices in each step. Once the recursion terminates, we query the variables corresponding to edges incident to the single remaining vertex — this contributes 2^D factor to size and increases depth by at most D. Overall, the SP proof has size $n^{O(\log n + D/\log n)}$ and rank $O(D + \log^2 n)$. $\qquad\square$

Combining this upper bound with a know lower bound on the rank of refutations in CP for the Tseitin formulas allows us to separate SP and CP in terms of rank.

Corollary 3.2. SP *is strictly stronger than* CP *with respect to proof rank.*

Proof. By Theorem 4.4, any Cutting Planes refutation of rank r can be converted into a SP refutation of rank $O(r)$. Buresh-Oppenheim et al. proved $\Omega(n)$ lower bound on the rank of Cutting Planes of the Tseitin formulas on constant-degree expander graphs [7], while Theorem 3.1 shows that SP can refute such Tseitin formulas in rank $O(\log^2 n)$. $\qquad\square$

4 Simulation Theorems

In this section, we prove simulation theorems relating the SP proof system to other similar proof systems in the literature. We begin by showing that SP is polynomially equivalent to the tree-like R(CP) system (introduced by Krajicek in [24]), which can be thought of as tree-like Resolution with clauses of inequalities and allowing CP rules. Since tree-like R(CP) simulates tree-like CP, the natural question is whether SP (and consequently R(CP)) can simulate general CP. We answer this question positively by providing two simulations. First of all, we observe that SP can depth-simulate CP. This simulation, while preserving depth of the proof, can lead to an exponential increase in the size. Thus, by a different simulation we show that SP can size-simulate CP. This time around, while the simulation preserves the size of a CP refutation, it can significantly increase the depth. It is an interesting open question whether there is a simulation of CP by SP that can simultaneously preserve depth and size of CP refutations. To complete the picture, we note that general R(CP) can trivially simulate tree-like R(CP) (and consequently SP) and CP. We also show that tree-like CP refutations can be efficiently converted into balanced (logarithmic-depth) SP refutations — this shows that tree-like CP refutations, which cannot in general be balanced, *can* be balanced in SP. Figure 2 is a graphical depiction of the relative strengths of various proof systems related to SP.
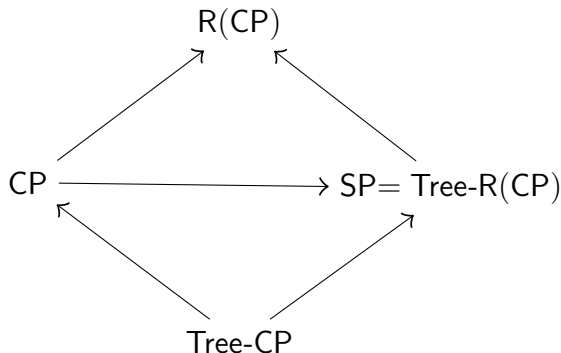
R(CP)

CP $\longrightarrow$ SP$=$ Tree-R(CP)

Tree-CP

Figure 2: Relationships between proof systems considered here. SP's equivalence to tree-like R(CP), as well as SP (and Tree-R(CP)) simulating CP are the new relationships proved in this paper.

We then turn to the question of space-time simulations. Recall, that a proof system can be thought of as a non-deterministic Turing machine. The notion of space of CP refutations intuitively corresponds to the minimum size of the work tape of such a non-deterministic Turing machine that is required to carry out the computation. In this analogy, the notion of length of CP refutations corresponds to the running time of the given TM. We show that general CP refutations that use length ℓ and space s can be turned into depth $O(s \log \ell)$ SP refutations. Thus, sufficiently strong lower bounds on the depth of SP refutations lead to time-space tradeoffs for CP.

Equivalence of SP **with tree-like** R(CP). Here we show the SP system is polynomially equivalent to the R(CP) proof system. We begin by formally defining the R(CP) proof system.

Definition 4.1. The R(CP) proof system is a syntactic proof system defined as follows. The lines of the R(CP) system are disjunctions of integer linear inequalities $\Gamma = L_1 \vee L_2 \vee \cdots \vee L_\ell$, and the lines are equipped with the following deductive rules. Let Γ be an arbitrary disjunction of integer linear inequalities, let $Ax \geq b, Cx \geq d$ be arbitrary integer linear inequalities, and let α, β be any positive integers.

$$\text{Linear Combination: } \frac{(Ax \geq b) \vee \Gamma \qquad (Cx \geq d) \vee \Gamma}{(\alpha A + \beta C)x \geq (\alpha b + \beta d) \vee \Gamma} \qquad \text{Division: } \frac{(\alpha Ax \geq b) \vee \Gamma}{(Ax \geq \lceil b/\alpha \rceil) \vee \Gamma}$$

$$\text{Axiom Introduction: } \frac{}{(Ax \geq b) \vee (Ax \leq b - 1)} \qquad \text{Weakening: } \frac{\Gamma}{(Ax \geq b) \vee \Gamma}$$

$$\text{Cut: } \frac{(Ax \geq b) \vee \Gamma \qquad (Ax \leq b - 1) \vee \Gamma}{\Gamma} \qquad \text{Elimination: } \frac{(0 \geq 1) \vee \Gamma}{\Gamma}$$

An R(CP) proof is *tree-like* if the proof DAG is a tree.

Theorem 4.2. *Let $\mathcal{C}$ be an unsatisfiable CNF, and let $C_1, C_2, \ldots, C_m$ be the representation of $\mathcal{C}$ as an integer-linear system of inequalities. For any* SP *refutation of $\mathcal{C}$ with size s and depth d there is a tree-like* R(CP) *refutation of $\mathcal{C}$ of size $O(s(d^2 + dm))$ and width $d + 1$.*

Proof. Let $\mathcal{T}$ be the SP refutation of $\mathcal{C}$, and consider any path p in the tree $\mathcal{T}$. Let $L_1, L_2, \ldots, L_t$ be the sequence of inequalities on p. We first show how to derive the clause

$$\neg L_1 \vee \neg L_2 \vee \cdots \vee \neg L_t$$

efficiently in R(CP) from $\mathcal{C}$. Begin by using the Axiom Introduction rule to introduce the lines $L_i \vee \neg L_i$ for each $i = 1, 2, \ldots, t$. Then, for each i, repeatedly apply the Weakening rule to the line $L_i \vee \neg L_i$ to deduce

$$L_i \vee \neg L_1 \vee \neg L_2 \cdots \vee \neg L_t,$$

and then for each input line C_i similarly apply the Weakening rule to deduce

$$C_i \vee \neg L_1 \vee \neg L_2 \vee \cdots \vee \neg L_t.$$

Since $\mathcal{T}$ is a Stabbing Planes refutation, there is a convex combination of $L_1, L_2, \ldots, L_t$, $C_1, \ldots, C_m$ equalling $0 \geq 1$. Furthermore, by Fact 2.4, we can assume without loss of generality that this convex combination contains at most a linear number of these inequalities. Finally, because any convex combination can be performed in tree-like Cutting Planes by repeatedly applying the "Linear Combination" rule, there is a linear size tree-like cutting planes

refutation of the system $L_1, L_2, \ldots, L_t, C_1, \ldots, C_m$. By simulating this proof in R(CP) on the lines

$$L_1 \vee \neg L_1 \vee \cdots \vee \neg L_t$$
$$L_2 \vee \neg L_1 \vee \cdots \vee \neg L_t$$
$$\vdots$$
$$L_t \vee \neg L_1 \vee \cdots \vee \neg L_t$$
$$C_1 \vee \neg L_1 \vee \cdots \vee \neg L_t$$
$$\vdots$$
$$C_m \vee \neg L_1 \vee \cdots \vee \neg L_t,$$

(that is, by applying the appropriate cutting planes rules to the first inequality in each line), we can deduce the line

$$(0 \geq 1) \vee \neg L_1 \vee \cdots \vee \neg L_t,$$

which can be simplified to

$$\neg L_1 \vee \neg L_2 \vee \cdots \vee \neg L_t$$

by the elimination rule. The proof of this path requires size $O(t^2 + tm + t + m) = O(d^2 + dm)$, where d is the depth of the tree $\mathcal{T}$, and clearly can be implemented as a tree-like proof.

Now we are nearly finished. In parallel, for each path p of the tree $\mathcal{T}$ construct the corresponding clause in short tree-like R(CP) as above. Applying the cut rule to the paths appropriately yields the empty clause in size $O(s(d^2 + dm))$. $\qquad\square$

Next, we prove the converse.

Theorem 4.3. *Let $\mathcal{C}$ be an unsatisfiable CNF, and let $C_1, C_2, \ldots, C_m$ be the representation of $\mathcal{C}$ as an integer linear system of equations. For any tree-like R(CP) proof of the disjunction $\neg C_1 \vee \neg C_2 \vee \cdots \vee \neg C_m$ with size s and depth d there is an SP refutation of $\mathcal{C}$ of size at most $2s$ and rank at most $2d$.*

Proof. Let $\mathcal{R}$ be the R(CP) proof of the disjunction, and we construct the SP refutation by structural induction. Consider any leaf of the proof $\mathcal{R}$ which, by assumption, is labelled with an input axiom $L \vee \neg L$ for some integer linear inequality L. It is easy to give a short SP refutation of $L, \neg L$: query the inequality $L, \neg L$ and refute each side of the tree appropriately.

By induction suppose that we have a tree-like R(CP) proof of a clause Γ. We break into cases depending on the last inference rule used to derive Γ.

Case 1. Linear Combination. Write Γ as the line

$$(\alpha A + \beta B)x \geq (\alpha a + \beta b) \vee \Delta,$$

which was deduced from lines

$$(Ax \geq a) \vee \Delta, \quad (Bx \geq b) \vee \Delta$$

15

by the Linear Combination rule. Let L_A be the inequality $Ax \geq a$, and L_B the inequality $Bx \geq b$. Begin the SP proof by making the query to the line L_A and $\neg L_A$ — on the branch labelled with $\neg L_A$, by induction we can construct a refutation of the clause $L_A \vee \Delta$. Then, on the branch labelled with L_A, branch on the inequalities L_B and $\neg L_B$. Again, on the $\neg L_B$ branch, we can apply induction to get an SP refutation of the clause $L_B \vee \Delta$. On the path labelled with L_A and L_B, we can immediately deduce a contradiction in stabbing planes from L_A, L_B, and $\neg(\alpha L_A + \beta L_B)$.

Case 2. Division. Write Γ as $(Ax \geq \lceil a/\alpha \rceil) \vee \Delta$, deduced from the line $(\alpha Ax \geq a) \vee \Delta$. Let $L_A \equiv \alpha Ax \geq a$. Begin the SP refutation by querying the inequalities L_A and $\neg L_A$. On the branch labelled $\neg L_A$, we can inductively construct a refutation of the clause $L_A \vee \Delta$. On the other branch, labelled with L_A, it is enough to observe that the intersection of L_A, which is $\alpha Ax \geq a$, and the inequality $Ax \leq \lceil a/\alpha \rceil - 1$, provided as an axiom to SP, is empty.

Case 3. Weakening. Write Γ as $L_A \vee \Delta$, deduced from Δ. This case is easy — by induction $\neg \Delta$ has a short SP refutation, which implies that $\neg L_A$ and $\neg \Delta$ has a short refutation.

Case 4. Cut. Suppose the line Γ was deduced from lines $L_A \vee \Gamma$ and $\neg L_A \vee \Gamma$ by the cut rule. This case is also straightforward — begin by querying the inequalities L_A and $\neg L_A$. On the branch labelled with L_A, we apply induction to construct a refutation of $\neg L_A \vee \Gamma$; symmetrically, on the branch labelled $\neg L_A$, apply induction to construct a refutation of $L_A \vee \Gamma$.

At worst, each line of the R(CP) proof is replaced with two inequality queries, of which at most two of the children are not immediately labelled with a convex combination equalling $0 \geq 1$, and so the size of the resulting tree is at most $2s$ and the depth is at most $2d$. □

Depth Simulation of Cutting Planes. Here we prove that there is a depth-preserving simulation of CP by SP.

Theorem 4.4. *For every Cutting Planes refutation of rank d, there is a* SP *refutation of the same tautology with rank at most $2d$. Moreover, if the* CP *refutation is tree-like of size s then the resulting* SP *refutation is of size $O(s)$ and rank $2d$.*

Proof. It is sufficient to prove the "moreover" part of the statement, since by recursive doubling, any CP refutation can be converted into a tree-like CP refutation where the rank remains the same, but the size may increase exponentially. Thus, from now on we assume that $\mathcal{R}_{\mathsf{CP}}$ is a size-s rank-d tree-like Cutting Planes refutation. We show that there is a size $O(s)$, rank $r \leq 2d$ SP refutation $\mathcal{R}_{\mathsf{SP}}$ of the same contradiction.

Let G be the graph (tree) associated with $\mathcal{R}_{\mathsf{CP}}$. The refutation $\mathcal{R}_{\mathsf{SP}}$ will be constructed from $\mathcal{R}_{\mathsf{CP}}$ by proceeding from the root of G to the leaves. In the process, we keep track of a subtree T in G, which we are left to simulate, and an associated current node N in $\mathcal{R}_{\mathsf{SP}}$, which we are constructing. Along the way the following invariant will be maintained: at every recursive step (N, T) such that $T \neq G$, if the root of the subtree T is labelled with the inequality $Cx \geq c$, then the edge leading to N in $\mathcal{R}_{\mathsf{SP}}$ is labelled with $Cx \leq c - 1$. Originally $T = G$ and $\mathcal{R}_{\mathsf{SP}}$ contains only a single node N. Consider the subtree T at the current recursive step in the proof, and we break into three cases.

Figure 3: A tree-like CP refutation invoking the non-negative linear combination inference rule and the the corresponding SP Refutation.

Case 1. The root of T labelled with $Cx \geq c$ has two children labelled with $Ax \geq a$ and $Bx \geq b$. Non-negative linear combinations are the only inferences in CP which take two premises, therefore, $Cx \geq c$ is a non-negative linear combination of $Ax \geq a$ and $Bx \geq b$. In the SP refutation $\mathcal{R}_{\mathsf{SP}}$ at the current node N, query $Ax \geq a$. On the branch labelled with $Ax \geq a$, query $Bx \geq b$ (see Figure 3). This sequence of queries results in three leaf nodes, labelled $N^{(1)}, N^{(2)}, N^{(3)}$ as in Figure 3. For the leaf node $N^{(1)}$ of the edge labelled with $Ax \leq a - 1$ in $\mathcal{R}_{\mathsf{SP}}$ let $T^{(1)}$ be the subtree rooted at the child of the root of T labelled with $Ax \geq a$. Recurse on $(N^{(1)}, T^{(1)})$. Similarly, for the leaf node $N^{(2)}$ of the pair of introduced edges labelled with $Ax \geq a$ and $Bx \leq b - 1$, let $T^{(2)}$ be the subtree rooted at the child of the root of T labelled with $Bx \geq b$. Recurse on $(N^{(2)}, T^{(2)})$.

For the final leaf node $N^{(3)}$ we can derive $0 \geq 1$. To see this, observe that if the current subtree T is G (i.e. the base case) then the root node of T is labelled with $0 \geq 1$. In this case, $Ax \geq a$ and $Bx \geq b$ are the premises used to derive $0 \geq 1$ by a non-negative linear combination in $\mathcal{R}_{\mathsf{CP}}$, and we can use this very non-negative combination at the leaf $N^{(3)}$. Otherwise, the root of the current subtree T is labelled with some inequality $Cx \geq c$. By the invariant, the edge leading to N in the refutation $\mathcal{R}_{\mathsf{SP}}$ we are constructing was labelled with $Cx \leq c - 1$. In the CP refutation $\mathcal{R}_{\mathsf{CP}}$, $Cx \geq c$ was derived by a non-negative linear combination of $Ax \geq a$ and $Bx \geq b$. Therefore, a non-negative combination of $Cx \leq c - 1$, $Ax \geq a$, and $Bx \geq b$ derives $0 \geq 1$, so, label $N^{(3)}$ with this combination.

Case 2: The root of T, labelled with $Ax \geq a$, has a single child derived by an application of the division rule from $Bx \geq b$. Note that by the invariant, the edge leading to the current node is labelled with $Ax \leq a - 1$. Thus, at this current node we query $Bx \geq b$. Let $N^{(1)}$ be the leaf node labelled with $Bx \leq b - 1$ and $N^{(2)}$ the leaf labelled with $Bx \geq b$. At $N^{(1)}$, we let $T^{(1)}$ be the subtree of T rooted at the child of the root of T and recurse on $(N^{(1)}, T^{(1)})$. On the other hand, at $N^{(2)}$ we can derive $0 \geq 1$ by a non-negative linear combination. This follows like so: by the division rule of CP the inequality $Ax \geq a$ is exactly the inequality $\frac{B}{d}x \geq \lceil \frac{b}{d} \rceil$ for some $d \in \mathbb{Z}_{\geq 0}$ dividing all entries in B. Therefore, subtracting $dAx \leq d(a - 1)$ from $Bx \geq b$, we derive $0 \geq b - \lceil \frac{b}{d} \rceil + d \geq 1$, and so we label $N^{(2)}$ with this linear combination.

17

Case 3: T is a single node — a leaf of the CP refutation labelled with some axiom $Ax \geq a$. By the invariant, the edge leading to N in the SP refutation $\mathcal{R}_{\mathsf{SP}}$ is labelled $Ax \leq a - 1$. We derive $0 \geq 1$ at N by subtracting $Ax \leq a - 1$ from the axiom $Ax \geq a$.

To see that the SP refutation we have constructed has rank at most twice that of the tree-like Cutting Planes refutation, observe that non-negative linear combinations are the only inference rule of Cutting Planes which this construction requires depth 2 to simulate, while all other rules require depth 1.

To measure the size, note that every CP rule with a single premise is simulated in SP by a single query, where one of the outgoing edges of that query is immediately labelled with $0 \geq 1$. Each rule with two premises (case 1) is simulated by two queries in the SP refutation, where one of the three outgoing edges is immediately labelled with $0 \geq 1$. Each of these queries branch only on the inequalities belonging to $\mathcal{R}_{\mathsf{CP}}$. Therefore, the size of the SP refutation $\mathcal{R}_{\mathsf{SP}}$ is $O(s)$. □

Size Simulation of Cutting Planes. Next, we show that SP size-simulates CP.

Theorem 4.5. SP *polynomially simulates* CP.

Proof. Let $\mathcal{R} = \{A_1 x \geq a_1, A_2 x \geq a_2, \ldots, A_m x \geq a_m\}$ be a CP refutation of an unsatisfiable set of integer linear inequalities $\mathcal{F}$. We construct a SP refutation of $\mathcal{F}$ line by line, following $\mathcal{R}$. Our SP refutation is a tree where the right-most path is of length $m + 1$ with edges labelled $A_1 \geq b_1, \ldots, A_m \geq b_m$. The left child of node $i \leq m$ along this path is labelled with $0 \geq 1$, which is derived as a non-negative linear combination of $A_j x \geq a_j$ for $j < i$, $A_i x \leq a_i - 1$, and $\mathcal{F}$. The last node in the path is also labelled with $0 \geq 1$. See Figure 4.


Figure 4: The SP simulation of a CP refutation.

Since $A_m x \geq a_m \equiv 0 \geq 1$, the last node is trivially labelled with $0 \geq 1$. Thus, we only need to show that the left child of every node can be legally labelled with $0 \geq 1$.

18

1. $A_i x \geq a_i$ is an axiom. We can derive $0 \geq 1$ by subtracting $A_i x \leq a_i - 1$ from $A_i x \geq a_i \in \mathcal{F}$.

2. $A_i x \geq a_i$ is a non-negative combination of two previous inequalities, i.e., $A_i x \geq a_i$ is

$$\alpha A_{j_1} x + \beta A_{j_2} x \geq \alpha a_{j_1} + \beta a_{j_2}$$

   for some $j_1, j_2 < i$ and $\alpha, \beta \in \mathbb{Z}_{\geq 0}$. We can derive $0 \geq 1$ by subtracting $A_i x \leq a_i - 1$ from the non-negative linear combination of $A_{j_1} x \geq a_{j_1}$ and $A_{j_2} x \geq a_{j_2}$ used to derive $A_i x \geq a_i$.

3. $A_i x \geq a_i$ is obtained by an application of the division rule to $A_j x \geq a_j$ for some $j < i$, i.e., $A_i x \geq a_i$ is

$$\frac{A_j}{c} x \geq \left\lceil \frac{a_j}{c} \right\rceil$$

   where $c \in \mathbb{N}$ divides every entry in A_j. On the path to this node in our SP refutation we queried $A_j x \geq a_j$. Dividing this inequality by c and subtracting $A_i x \leq a_i - 1$, we obtain $0 \geq a_j / c - (\lceil a_j / c \rceil - 1)$. This gives us $0 \geq \frac{a_j}{c} + 1 - \left\lceil \frac{a_j}{c} \right\rceil$. Since the right-hand side is strictly positive this can be normalized to give $0 \geq 1$. $\square$

Tree-Like CP **and Balanced** SP**.** It is known that CP refutations cannot be balanced (i.e. size-s refutations being turned into size $O(s)$ depth $O(\log s)$ refutation) in CP. Here, we show that CP proofs can be turned into balanced SP proofs. More specifically, we prove the following.

Theorem 4.6. *Suppose there is a size s tree-like* CP *refutation of a set of linear integer inequalities $\mathcal{F}$. Then there is a size s depth $O(\log s)$* SP *refutation of $\mathcal{F}$.*

Proof. The construction is recursive. Let T be the tree corresponding to $\mathcal{P}$.

   Base case. $|T| = O(1)$, then we can use one of the previous simulation theorems to create an SP refutation of $\mathcal{P}$.

   Recursive step. Let v be a node in T such that the subtree T_v rooted at v satisfies $|T|/3 \leq |T_v| \leq 2|T|/3$, which must exist since the size measure is additive. Let $Bx \geq b$ be the line in $\mathcal{P}$ corresponding to the node v. Our SP simulation starts by querying Bx. We continue recursively depending on the outcome of the query. If $Bx \geq b$ then we apply the recursive construction to $T \setminus T_v$, where we can think of $Bx \geq b$ as a new axiom. Otherwise, if $Bx \leq b-1$ (which contradicts the the input set of inequalities $\mathcal{F}$) we apply the recursive construction to T_v. The size is clearly preserved, and the depth of the proof becomes logarithmic, since we are reducing the size of the proof to be simulated by a constant factor on each branch of a query. $\square$

Bounded Depth and Space CP **yields Balanced** SP **Proofs** In order to talk about the *space of refuting* a given CNF, one needs to generalize the notion of lines to the notion of *configurations*. Configurations model the state of the working tape of the non-deterministic TM in the definition of a proof system. More formally, we have the following.

Definition 4.7. Let $\mathcal{P}$ be a proof system, we consider a refutation of an unsatisfiable formula $\mathcal{F}$ as a sequence of *configurations* $\mathbb{D}_1, \ldots, \mathbb{D}_k$, where each configuration $\mathbb{D}_i$ is a set of lines of $\mathcal{P}$ such that $\mathbb{D}_k = \emptyset$ and each $\mathbb{D}_i$ follows $\mathbb{D}_{i-1}$ by one of the derivation steps. Let $\mathbb{D}'_{i-1}$ is $\mathbb{D}_{i-1} \setminus \{L_{i-1,1}, \ldots L_{i-1,p}\}$ where $\{L_{i-1,1}, \ldots L_{i-1,p}\}$ is an arbitrary (possibly empty) subset of inequalities contained in $\mathbb{D}_{i-1}$.

- *Axiom Download*: $\mathbb{D}_i = \mathbb{D}'_{i-1} \cup \{L_j\}$ where L_j is one of the lines in $\mathcal{F}$ or one of the axioms of the proof system Π.

- *Inference*: $\mathbb{D}_i = \mathbb{D}'_{i-1} \cup \{L\}$ for some line L inferred by one of the inference rules of Π applied to the set of inequalities in $\mathbb{D}_{i-1}$.

The *space* of a refutation $R = \mathbb{D}_1, \ldots, \mathbb{D}_k$ is $Space(R) = \max_{i \in [k]} |\mathbb{D}_i|$ and the space of refuting $\mathcal{F}$ in a proof system $\mathcal{P}$ is $Space(\mathcal{F}) = \min_{R \text{ - refutation of } \mathcal{F}} Space(R)$.

Next, we show that strong depth lower bounds on SP refutations can lead to time-space tradeoffs for CP.

Theorem 4.8. *Suppose that we have a length ℓ, space s CP refutation of an unsatisfiable set of linear integral inequalities. Then there is a depth $O(s \log \ell)$ SP refutation of the same set of linear integral inequalities.*

Proof. Let $\mathbb{D}_1, \ldots, \mathbb{D}_\ell$ be configurations of the given CP refutation. Then $\mathbb{D}_\ell = \{0 \geq 1\}$. The proof follows by applying the following main observation with $i = \ell$ to obtain an SP tree where each leaf is labelled with $0 \geq 1$.

Main observation. Let $\mathbb{D}_1, \ldots, \mathbb{D}_\ell$ be configurations in a CP refutation such that $|\mathbb{D}_i| \leq s$ for $i \in [\ell]$. For every i there is a SP tree of queries of depth $O(s \log i)$ such that there is exactly one path along which we know all $L \in \mathbb{D}_i$ and other paths end in leaves labelled $0 \geq 1$.

In the rest, we prove the main observation. The construction is recursive. Let $c \geq 1$ be the hidden constant from the big-Oh notation (it will be clear from the proof that it exists).

Base case. $\mathbb{D}$ is a singleton set containing an axiom. Take the tree to be a single node. It clearly satisfies the statement of the lemma.

Recursive step. Assume for simplicity that i is divisible by 2. The SP tree starts with a complete binary tree in which every line from $\mathbb{D}_{i/2}$ is queried. The depth of this tree is $\leq s$. Exactly one path P is labelled with lines from $\mathbb{D}_{i/2}$. All other paths contain at least one label that is the *negation* of a line from $\mathbb{D}_{i/2}$. To finish the construction, we treat these two cases separately.

In the case that a path contains a negation of a line from $\mathbb{D}_{i/2}$, we attach to its leaf the SP tree we obtain recursively by running our construction on $\mathbb{D}_1, \ldots, \mathbb{D}_{i/2}$. We know that the attached tree has all, but one, leaves labelled $0 \geq 1$. For the one leaf that is not labelled with $0 \geq 1$, we know all lines from $\mathbb{D}_{i/2}$ on the path to that leaf. Since we attached this tree to a path along which we know the negation of a line from $\mathbb{D}_{i/2}$, we can immediately label this leaf by $0 \geq 1$. The overall depth that we get in this case is s for the initial tree and $cs \log(i/2)$ for the tree obtained recursively, i.e., $s + cs(\log i - 1) = s + cs \log i - cs \leq cs \log i$, since $c \geq 1$.

20

In case of the path P note that $\mathbb{D}_{i/2+1}, \ldots, \mathbb{D}_i$ can be viewed as configurations of a refutation of the original set of unsatisfiable linear integral inequalities together with $\mathbb{D}_{i/2}$ treated as additional axioms. In our SP construction, at the leaf of P we also know all inequalities $\mathbb{D}_{i/2}$, and so can treat them as axioms. Thus, we can apply the recursive construction to $\mathbb{D}_{i/2+1}, \ldots, \mathbb{D}_i$. The calculation of the overall depth is exactly the same as in the previous case. The big-Oh is needed to take care of rounding issues when i is not divisible by 2. □

5 Impossibility Results

In this section, we prove near-optimal lower bounds on SP rank via reductions to randomized and real communication complexity. We then tackle the harder problem of proving unrestricted superpolynomial size lower bounds for SP. Although we are unable to prove such lower bounds we explain why current approaches fail. Essentially all lower bounds for CP have been obtained by reducing to a communication complexity problem; in the case of tree-like CP, the reduction is to the communication complexity of a corresponding search problem. For more general dag-like CP, the reduction is to the size of "communication games" [13, 17] (communication games are a dag-like model of communication that gives an equivalence between communication *size* and monotone circuit size, analogous to the famous equivalence between communication *depth* and monotone formula size). Although tree-like CP proofs cannot be balanced in general, communication protocols (both deterministic and randomized) can be balanced, and thus tree-like CP lower bounds follow from communication complexity lower bounds. Similarly, we show that it is not possible to balance SP refutations, and thus we cannot in general obtain size lower bounds directly from rank lower bounds. Moreover, we show that SP refutations imply real communication protocols , and unlike ordinary communication protocols, we show that real protocols *cannot* in general be balanced. This rules out proving length lower bounds on SP refutations from (real) communication complexity lower bounds.

SP **Refutations Imply Communication Protocols.** Real communication protocols were introduced by Krajíček [25]. In this model, the players are allowed to communicate by sending real-valued functions of their inputs to a referee who announces their comparison.

Definition 5.1. A *real communication protocol* is a full binary tree in which every non-leaf node v is labeled with a pair of functions $a^v : \mathcal{X} \to \mathbb{R}, b^v : \mathcal{Y} \to \mathbb{R}$, and the leaves are labelled with elements in $\mathcal{Z}$. Two players, Alice and Bob, receive inputs from $\mathcal{X} \times \mathcal{Y}$, with Alice receiving $x \in \mathcal{X}$ and Bob receiving $y \in \mathcal{Y}$. Beginning at the root, the players traverse the tree as follows: at each node, they send real values $a^v(x)$ and $b^v(y)$ to a "referee" who returns (to both of them) a bit indicating the result of the comparison $a^v(x) \geq b^v(y)$; the players proceed to the left child if $a^v \geq b^v$, and to the right child otherwise. Once they reach a leaf, the protocol halts, and the players output the value in $\mathcal{Z}$ labelling the leaf; it computes a function $f : \mathcal{X} \times \mathcal{Y} \to \mathcal{Z}$ in the natural way.

The *cost* of a real protocol is the depth of the tree, or equivalently the maximum number of rounds of communications with the referee over any input (x, y), and the *size* is the number

of nodes in the protocol. Similarly, the cost (size) of computing a function f is the smallest cost (size) real protocol computing f.

Krajíček showed that from a low-rank CP refutation of an unsatisfiable CNF, one can obtain a real communication protocol for solving a related search problem [25]. We describe this search problem next.

Definition 5.2. Let $\mathcal{F} = C_1 \wedge C_2 \wedge \cdots \wedge C_m$ be an unsatisfiable CNF and (X, Y) be a partition of the variables. The relation $\mathsf{Search}_{X,Y}(\mathcal{F}) \subseteq \{0,1\}^X \times \{0,1\}^Y \times [m]$ consists of all triples (x, y, i) such that the total assignment $z = (x, y)$ to all of the variables of C_i falsifies the clause i.

The search problem is the natural interpretation of a refutation in the setting of communication. Indeed, essentially every lower bound for CP has been proved by reducing to the communication complexity of the search problem. In a similar manner, we show that SP refutations can be turned into both randomized and real protocols for the search problem which preserve the rank of the refutation.

Lemma 5.3. *Let $\mathcal{F}$ be an unsatisfiable CNF formula and (X, Y) be any partition of the variables. Any SP refutation of $\mathcal{F}$ of rank r implies a real communication protocol of cost $O(r + \log n)$ and an $O(r \log n + \log^2 n)$ randomized bounded-error protocol for solving $\mathsf{Search}_{X,Y}(\mathcal{F})$.*

Proof. Let $\mathcal{P}$ be a rank r SP refutation of $\mathcal{F}$ and let $(x, y) \in \{0,1\}^{|X|} \times \{0,1\}^{|Y|}$ be any assignment to the variables of $\mathcal{F}$; Alice is given x and Bob is given y. Our protocol will traverse from the root of $\mathcal{P}$ to the leaves in search of a clause that is falsified by (x, y). At each node in $\mathcal{P}$ that the players visit, labeled with a query $(Ax \leq b - 1, Ax \geq b)$, they will evaluate their input on $Ax \geq b$. Because the assignment (x, y) is integral, it will satisfy exactly one of the inequalities

$$Ax \leq b - 1, \qquad Ax \geq b, \tag{1}$$

and will falsify the other. Alice and Bob will then continue to the child of the current node which is reached by traversing the edge labeled with the inequality in (1) that is satisfied by (x, y). This protocol will continue for at most r iterations until the players reach a leaf of the refutation $\mathcal{P}$.

Every leaf of the protocol is labelled with a convex combination of the axioms of $\mathcal{F}$ along with the inequalities labelling the path leading to this leaf, which evaluates to $0 \geq 1$. Because we have maintained that (x, y) satisfies each of the inequalities along the path leading to this leaf, it must falsify one of the axioms of $\mathcal{F}$ used in the convex combination. If this was not the case, the polytope formed by the inequalities in the convex combination would contain a feasible point, and by Fact 2.4, a convex combination equalling $0 \geq 1$ would not exist.

Once at a leaf, Alice and Bob can communicate in at most $O(\log n)$ rounds to find the clause of $\mathcal{F}$ that is falsified under (x, y). Using Fact 2.4, we may assume that this is a convex combination of $\ell \leq n + 2$ linear inequalities. Denote these inequalities by $A_1 x \geq b_1, \ldots, A_\ell x \geq b_\ell$, and their coefficients in the convex combination be $c_1, \ldots, c_\ell$. Alice and

Bob will repeat the following procedure, maintaining that the current inequality is always less than or equal to it's constant term. Divide the set of inequalities into two halves, and define the threshold function

$$L := \left(\sum_{i=1}^{\ell/2} c_i (A_i x - b_i) \right) - \left(\sum_{j=\ell/2+1}^{\ell} c_j (A_j x - b_j) \right) \geq 0.$$

The players will run the protocol for deciding L on assignment (x, y). If (x, y) falsifies L, they recurse on the subset of inequalities $A_1 x \geq b_1, \ldots, A_{\ell/2} x \geq b_{\ell/2}$, otherwise, they recurse on the other half.

Since the original convex combination evaluated to $0 \geq 1$, at least one of $\sum_{i=1}^{\ell/2} c_i (A_i x - b_i)$ and $\sum_{j=\ell/2+1}^{\ell} c_j (A_j x - b_j)$ must be strictly less than 0 on the assignment (x, y). Repeating this process will converge to an inequality which is violated on (x, y). Because we have ensured that the only inequalities that (x, y) falsifies in this convex combination are axioms of $\mathcal{F}$, this process will solve the search problem. Moreover, because this convex combination contains at most $n + 2$ linear inequalities, this process will terminate in $O(\log n)$ rounds.

In the model of real communication, any linear inequality $Ax \geq b$ can be evaluated in a single bit of real communication; Alice sends Ax to the referee, while Bob sends $b - Ay$ (Here, y is treated as a vector of length n having 0s in coordinates corresponding to x, and similarly for x). Therefore, this leads to a $O(r + \log n)$ real communication protocol for $\mathsf{Search}_{X,Y}(\mathcal{F})$.

Next, we adapt the above procedure to produce a randomized communication protocol by showing that any inequality in the SP refutation can be computed in low communication. To do this, Alice and Bob run the ε-error $O(\log m + \log \varepsilon^{-1})$-bit protocol of Nisan [27] for deciding an m-bit linear inequality. By the well-known result of Muroga [26], any inequality on n Boolean variables only requires coefficients representable by $O(n \log n)$ bits (recall that Alice and Bob's input will always be a Boolean assignment and so this suffices). Because there are at most $r + \log n$ inequalities evaluated along any root-to-leaf path in the refutation, the protocol is repeated at most $r + \log n$ many times. By a union bound, we require $\varepsilon < c/(\log n + r)$, where c is some constant bound on the error that we allow. Therefore, every inequality can be computed in $O(\log n + \log r)$ many bits to compute, giving a $O(r \log n + \log^2 n)$ bounded-error randomized protocol for $\mathsf{Search}_{X,Y}(\mathcal{F})$.

$\square$

Rank Lower Bounds For SP. To take advantage of Lemma 5.3, we need to find some candidate formulas on which to prove rank lower bounds and then study the search problem obtained from applying this transformation. We do so for both the Tseitin formulas and a variant of the pebbling contradictions, a reformulation of the classical black pebbling games as an unsatisfiable 3-CNF formula, originally introduced by Ben-Sasson et al. [2, 3].

The black pebbling game can be phrased as a contradictory 3-CNF as follows: Let G be a DAG with a set of source nodes $S \subseteq V(G)$ (having fanin 0), a unique sink node t (with fanin 2 and fanout 0), and the remaining nodes each having fanin exactly 2. The pebbling contradictions Peb_G consists of the following $n + 1$ clauses over variables $v \in V(G)$:

- sink axiom: a single clause $\neg t$,

- source axioms: a clause s for every source $s \in S$.

- pebbling axioms: a clause $\neg u \vee \neg v \vee w$ for every $w \in V \setminus S$ with immediate children u, v.

Unfortunately, both the pebbling contradictions and the Tseitin formulas have short SP refutations. In particular, for any graph G, the polytope formed by the constraints of Peb_G is empty and therefore a nonnegative combination of the constraints yielding $0 \geq 1$ exists, this is a valid rank-1 SP refutation. For Tseitin, this follows from the poly-logarithmic rank upper bound in Theorem 3.1. Therefore, we modify these formulas to make them harder to solve.

A standard technique for amplifying the hardness of computing some function $f : \mathcal{X}^n \to \mathcal{Z}$ is by *lifting* that function. This is done by obscuring the input variables by replacing them each by a small function $g : \mathcal{A} \to \mathcal{X}$ known as a *gadget*, which must be evaluated before learning the input to the original function. For an input $\alpha \in \mathcal{A}^n$, the function f *lifted* by gadget g is then
$$(f \circ g^n)(\alpha) = (g(\alpha_1), \ldots, g(\alpha_n)).$$
The intuition is that this lifted function $f \circ g^n$ should be much harder than the original because the players must first evaluate the gadget $g(\alpha_i)$ to learn each bit of the actual input to the function f. Furthermore, intuition says that if the gadget is sufficiently difficult to compute, then the model will be reduced to using much more rudimentary methods to evaluate the lifted function.

The standard hard-to-compute gadget is the *pointer* or *index* gadget, $\mathsf{IND}_\ell : [\ell] \times \{0, 1\}^\ell \to \{0, 1\}$. For an input $(x, y) \in [\ell] \times \{0, 1\}^\ell$, x is a $\log \ell$-bit string encoding a pointer into the ℓ-bit string $y \in \{0, 1\}^\ell$. The output of $\mathsf{IND}_\ell(x, y)$ is $y[x]$, the x-th bit in the string y. This is most often applied in communication complexity, where typically the variable partition between the players is that Alice is given $x \in [\ell]$ and Bob is given $y \in \{0, 1\}^\ell$. In any standard model of communication, for this partition of the variables, it is difficult to imagine any communication protocol which could compute the index gadget with significant advantage over the trivial protocol; sending every bit of Alice's pointer x to Bob.

Raz and McKenzie formalized this intuition, in what has become known as a lifting theorem [15, 29]. They show that deterministic communication protocols cannot compute any function f lifted by the index gadget significantly better than simply mimicking a decision tree computing f, and performing the trivial protocol for evaluating the index gadget every time a bit of the input to f is needed.

Lifting theorems for real communication were originally proved by Bonet et al. [6] based on the techniques of Raz-McKenzie. Their theorem lifts lower bounds on the decision tree complexity of a function f to lower bounds on the cost of real communication protocols computing $f \circ \mathsf{IND}_\ell^n$. The decision tree complexity $\mathsf{DT}(f)$ of a function f is simply the minimum depth need by any decision tree to compute f. We use a simplified lifting theorem for real communication by de Rezende et al. [11], which we state next.

Theorem 5.4. *(de Rezende et al. [11]) Let f be a function with domain $\{0, 1\}^n$ and let $\ell = n^4$. If there is a real communication protocol of cost c that solves $f \circ \mathsf{IND}_\ell^n$ where Alice is given*

24

$x \in [\ell]^n$ and Bob is given $y \in \{0,1\}^{n\ell}$, then there is a decision tree solving f using $O(c/\log \ell)$ queries.

Our goal is now to is to combine this theorem with Lemma 5.3 in order to prove lower bounds on the rank of SP refutations of $\mathrm{Peb}_G \circ \mathsf{IND}_\ell^n$. Syntactically speaking though, $\mathrm{Peb}_G \circ \mathsf{IND}_\ell^n$ is not a valid input to our proof system. Therefore, we must show that the lifted function can indeed be phrased as a small CNF formula. The following encoding is due to Beame et al. [1]:

Let $\mathcal{F} = C_1 \vee \ldots \vee C_m$ be a CNF formula over variables $x_1, \ldots, x_n$. The CNF representing $\mathcal{F} \circ \mathsf{IND}_\ell^n$ is defined on new sets of variables $y_{i,j}$ and $z_{i,j}$ for all $i \in [n]$ and $j \in [\ell]$. This CNF has the following set of clauses

- Pointer clauses: for each $i \in [n]$, a clause

$$y_{i,1} \vee \ldots \vee y_{i,\ell}.$$

- $\mathcal{F}$-clauses: for each clause $C_i \in \mathcal{F}$, where $C_i = y_{i_1} \vee \ldots \vee y_{i_k} \vee \neg x_{i_{k+1}} \vee \ldots \vee \neg x_{i_s}$ and for every $(j_1, \ldots, j_n) \in [\ell]^n$, a clause

$$(y_{i_1,j_1} \to z_{i_1,j_1}) \vee \ldots \vee (y_{i_k,j_k} \to z_{i_k,j_k}) \vee (y_{i_{k+1},j_{k+1}} \to \neg z_{i_{k+1},j_{k+1}}) \vee \ldots \vee (y_{i_s,j_s} \to \neg z_{i_s,j_s}).$$

We will abuse notation and use $\mathcal{F} \circ \mathsf{IND}_\ell^n$ to denote the function, as well as it's CNF formulation, and use context to differentiate between the two.

A final subtlety that should be mentioned is that applying Theorem 5.4 to an SP refutation of $\mathrm{Peb}_G \circ \mathsf{IND}_\ell^n$ yields a protocol for $\mathsf{Search}_{X,Y}(\mathrm{Peb}_G \circ \mathsf{IND}_\ell^n)$ which is not in the correct form to apply Theorem 5.4 ($\mathsf{Search}_{X,Y}(\mathrm{Peb}_G \circ \mathsf{IND}_\ell^n)$ is a function of a lifted function, whereas Theorem 5.4 can only be applied to lifted functions). Luckily, this is not a significant issue; Huynh et al. [18] show that, for any unsatisfiable CNF $\mathcal{F}$, any real communication protocol for $\mathsf{Search}_{X,Y}(\mathcal{F} \circ \mathsf{IND}_\ell^n)$, where $X = [\ell]^n$ and $Y = \{0,1\}^{n\ell}$, implies a real communication protocol for $\mathsf{Search}_{X,Y}(\mathcal{F}) \circ \mathsf{IND}_\ell^n$ with the same parameters.

It is now straightforward to obtain lower bounds on the rank of SP refutations. For the lifted pebbling formulas, SP rank lower bounds follow from combining Lemma 5.3 and Theorem 5.4 with a lower bound on the complexity of decision trees solving Peb_G proved by de Rezende et al. [11].

Theorem 5.5. *There exists a graph G of indegree 2 on n vertices such that the unsatisfiable CNF formula $\mathrm{Peb}_G \circ \mathsf{IND}_\ell^n$, for $\ell = n^4$, on $n(\ell + \log \ell)$ variables requires rank $\Omega(\sqrt{n \log n})$ to refute in SP.*

Proof. Consider the pebbling formulas. de Rezende et al. [11] showed the existence of a graph G on n vertices with indegree 2 such that the decision tree complexity of outputting a falsified clause of the Peb_G formulas is $\Omega(\sqrt{n/\log n})$. Applying the real communication lifting theorem (Theorem 5.4) and combining this with the fact that shallow SP refutations give efficient protocols for the associated search problem (Lemma 5.3), proves the desired $\Omega(\sqrt{n \log n})$ lower bound on the rank of SP refutations of $\mathrm{Peb}_G \circ \mathsf{IND}_\ell^n$. $\qquad \square$

Finally, a similar technique can be applied to obtain a lower bound on the rank of SP refutations for a lifted variant of the Tseitin formulas. This follows from the lower bound on the randomized communication complexity of the search problem for the Tseitin formulas lifted by a small constant-size gadget, which was obtained by Göös and Pitassi [14]. In particular, they use the *versatile gadget*, $\mathsf{VER} : \mathbb{Z}_4 \times \mathbb{Z}_4 \to \{0, 1\}$, which is defines as

$$\mathsf{VER}(x, y) = 1 \iff x + y \pmod 4 \in \{2, 3\}.$$

Theorem 5.6. *(Göös and Pitassi [14]) There exists a constant-degree graph G on n vertices such that, if ℓ is any $\{0, 1\}$ vertex labelling with odd total weight and (X, Y) is any partition of the variables, any bounded-error randomized communication protocol for* $\mathsf{Search}_{X,Y}(\mathrm{Tseitin}(G, \ell) \circ \mathsf{VER}^n)$ *on $O(n)$ variables, requires $\Omega(n/\log n)$ bits of communication.*

Furthermore, Göös and Pitassi showed how, for any CNF formula $\mathcal{F}$, the composed function $\mathcal{F} \circ \mathsf{VER}^n$ can be encoded as a CNF formula: For each variable z_i in the $\mathcal{F}$, define new variables $\mathrm{vars}(z_i) = \{x_{i,1}, x_{i,2}, x_{i,3}, x_{i,4}\}$, where the $(x_{i,1}, x_{i,2})$ and $(x_{i,3}, x_{i,4})$ should be thought of as binary encodings of the pair of inputs in $\mathbb{Z}_4 \times \mathbb{Z}_4$ to VER. In particular, let α be a truth assignment and $\alpha \restriction_{\mathrm{vars}(z_i)}$ be its restriction to the variable set $\mathrm{vars}(z_i)$. We will abuse notation and interpret $\mathsf{VER}(\alpha \restriction_{\mathrm{vars}(z_i)})$ as applying the gadget VER to the pair of elements in $\mathbb{Z}_4 \times \mathbb{Z}_4$ that $\alpha \restriction_{\mathrm{vars}(z_i)}$ is the binary encoding of.

Let C be a clause in $\mathcal{F}$ and assume for simplicity that $C = z_1 \vee \neg z_2$; it will be clear that this definition will generalize to any arbitrary clause. For every truth assignment $\alpha \in \{0, 1\}^{4n}$ such that $\mathsf{VER}(\alpha \restriction_{\mathrm{vars}(z_1)}) = 0$ and $\mathsf{VER}(\alpha \restriction_{\mathrm{vars}(z_2)}) = 1$, define a new clause

$$C_\alpha = \left(\bigvee_{i=1}^{4} \neg x_{1,i}^\alpha \right) \vee \left(\bigvee_{i=1}^{4} \neg x_{2,i}^\alpha \right),$$

where $x_{i,j}^\alpha = x_{i,j}$ if $\alpha \restriction_{x_{i,j}} = 1$ and $x_{i,j}^\alpha = \neg x_{i,j}$ otherwise. These clauses simply state that the output of the gadgets which correspond to the variables occurring in the clause C cannot be a falsifying assignment to C. The CNF representation of $\mathcal{F} \circ \mathsf{VER}^n$ is the conjunction of the clauses C_α for every $C \in \mathcal{F} \circ \mathsf{VER}^n$.

Observe that if $\mathcal{F}$ contains m clauses, each of width at most w, then the CNF representation of $\mathcal{F} \circ \mathsf{VER}^n$ contains at most $m \cdot 2^{4w}$ clauses. The width of every clause in the Tseitin formulas are bounded by d, where d is the maximal degree of any vertex in the underlying graph. Using this fact, we are able to obtain near-optimal lower bounds on the rank of SP refutations by combining Theorem 5.6 with Lemma 5.4 in the same manner as the proof of Theorem 5.5.

Theorem 5.7. *There a constant-degree graph G on n vertices such that if ℓ is any $\{0, 1\}$ vertex labelling with odd total weight, the CNF formula $\mathrm{Tseitin}(G, \ell) \circ \mathsf{VER}^n$, on $O(n)$ variables and clauses, requires SP refutations of rank $\Omega(n/\log^2 n)$.*

The lower bound from Theorem 5.7 should be contrasted with the logarithmic-rank SP upper bound on Tseitin from Theorem 3.1.

SP **Refutations Cannot Be Balanced.** Optimistically, one could hope that the length and rank of SP refutations may be closely related and therefore that we could leverage these rank bounds to obtain lower bounds on the length of SP refutations. We answer this question negatively, showing that there exists a contradictory CNF which admits short refutations, but for which these refutations must be almost maximally deep. That is, we show that SP refutations cannot be balanced; an SP refutation of length S does not imply one of rank $O(\log S)$. This shows that in SP the rank of refutations is a distinct complexity measure from the length.

In order to obtain time-space tradeoffs, de Rezende et al. [11] proved Resolution upper bounds on the lifted pebbling contradictions. Combining this upper bound (which can be simulated efficiently in SP) with the lower bound from Theorem 5.5 exhibits a formula that requires small size, but near-maximal rank to refute in SP.

Theorem 5.8. SP *refutations cannot be balanced.*

Proof. Suppose that a SP refutation of length S implied the existence of a SP refutation of the same formula of rank $O(\log S)$. Let G be the graph from Theorem 5.5 on n vertices, and let Peb_G be the pebbling contradictions defined on this graph. It follows immediately from Lemmas 7.2 and 7.3 from de Rezende et al. [11] that for any graph of indegree 2 on n vertices, that there is a Resolution refutation of $\mathrm{Peb}_G \circ \mathsf{IND}_\ell^n$ of length $O(n\ell^3)$. Since SP can p-simulate Cutting Planes (and therefore Resolution), this implies a $\mathrm{poly}(n)$ upper bound on the same formula in SP. Under the assumption that SP refutation can be balanced, this would imply a SP refutation of depth $O(\log n)$ of $\mathrm{Peb}_G \circ \mathsf{IND}_\ell^n$, contradicting the lower bound from Corollary 5.5. $\qquad\square$

Although it is a well-known fact that tree-like Cutting Planes refutations cannot be balanced, Impagliazzo et al. [20] show that the randomized communication protocols for the search problem obtained from CP refutations can be balanced. Using this fact, they show that a length S tree-like Cutting Planes proof implies a depth $O(\log S)$ protocol for the search problem. This implies that communication cost lower bounds for the search problem imply length lower bounds for tree-like Cutting Planes refutations.

Optimistically one could hope that a similar approach could be applied to SP refutations. This is reinforced by the fact that the real communication protocols for the search problem obtained from SP refutations (Lemma 5.3) maintains the same topology as the refutation. That is, the cost and size of the resulting protocol are approximately equal to the rank and length of the proof (unfortunately, this is not the case for the randomized protocols obtained from SP refutations). Therefore, one might hope that even though SP cannot be balanced, the corresponding real communication protocols for the search problem can be balanced. Thus, lower bounds on the rank of real-communication protocols for the search problem would imply lower bounds on the size of SP proofs.

Corollary 5.9. *Any* SP *refutation of length S and rank r of an unsatisfiable formula $\mathcal{F}$ implies a real communication protocol of size $O(S \cdot n)$ and cost $O(r + \log n)$ for solving* $\mathsf{Search}_{X,Y}(\mathcal{F})$.

Proof. This follows from observing that the protocol obtained in Lemma 5.3 also preserves the topology (and therefore both the rank and the length) of the refutation.

$\qquad\square$

Real Communication Protocols Cannot Be Balanced. Analogous to Theorem 5.8 (showing that SP proofs cannot be balanced) in this section we will show that real communication protocols cannot be balanced. This should be contrasted with other standard models of communication such as randomized and deterministic, which can be balanced. In particular, we exhibit a function which has a real communication protocol of small size, but for which every real protocol must have high cost. Towards this end, we first prove lower bounds on the real communication complexity of the famous set disjointness function.

The set disjointness function DISJ_n is the canonical NP^{cc}-complete problem. Each player is given an n-bit string, interpreted as indicator vectors for an underlying set of n elements, and they are asked to determine whether their sets are disjoint. That is, the players aim to solve the function

$$\mathrm{DISJ}_n(x, y) = \bigvee_{i \in [n]} (x_i \wedge y_i).$$

To our knowledge, the only known technique for obtaining lower bounds on the real communication of any problem are via lifting theorems, reducing the task of proving lower bounds on lifted functions to the decision tree complexity of the un-lifted function. Although DISJ_n can be seen as a lifted function (the OR_n function lifted by the two-bit AND gadget), these real communication lifting theorems work only for super-constant sized gadgets, and therefore cannot be applied directly to DISJ_n. We circumvent this difficulty by exploiting the fact that DISJ_n is NP^{cc}-complete. To do so, we find a lifted function in NP^{cc} to which our simulation theorems can be applied. Consider the n-bit OR_n function composed with the index gadget,

$$\mathrm{OR} \circ \mathsf{IND}_\ell^n,$$

for some ℓ defined later.

Lemma 5.10. $\mathrm{OR}_n \circ \mathsf{IND}_\ell^n \in \mathsf{NP}^{cc}$, *for any $\ell \leq 2^{polylog(n)}$.*

Proof. First, observe that the index gadget $\mathsf{IND}_\ell(x_i, y_i)$, for a single bit i of the input to the OR_n function, can be computed by a brute force protocol in $\log \ell$ bits of communication. Alice simply sends to Bob the $\log \ell$ bits of her input $x_i = x_{i,1}, \ldots, x_{i,\log \ell}$. Bob is then able to evaluate $\mathsf{IND}_\ell(x_i, y_i)$.

Now, consider the following NP^{cc} protocol for $\mathrm{OR}_n \circ \mathsf{IND}_\ell^n$: Alice and Bob are given as a proof, a $\log n$-bit string indicating the index $i \in [n]$ of the OR_n function where $\mathsf{IND}_\ell(x_i, y_i) = 1$. They then perform the brute force protocol to evaluate $\mathsf{IND}_\ell(x_i, y_i)$ and verify that the outcome is indeed 1. In total, this requires $\log \ell + \log n + 1 = polylog(n)$ bits of NP^{cc}-communication. $\square$

To obtain lower bounds on $\mathrm{OR}_n \circ \mathsf{IND}_\ell^n$, we appeal to the real communication simulation theorem (Theorem 5.4), reducing communication lower bounds for $\mathrm{OR}_n \circ \mathsf{IND}_\ell^n$ on the lifted problem to the well-known linear decision tree lower bounds on the OR_n function.

Lemma 5.11. *Let $\ell = n^4$. The cost of any real communication protocol computing $\mathrm{OR}_n \circ \mathsf{IND}_\ell^n$ is $\Omega(n \log \ell)$.*

Proof. Combining the $\Omega(n)$ lower bound on the decision tree complexity of computing the OR_n function with the simulation theorem of de Rezende et al. [11] proves the result. $\square$

Theorem 5.12. *The cost of any real communication protocol computing* DISJ_n *is* $\Omega((n \log n)^{1/5})$.

Proof. Let $\ell = n^4$. Consider the following reduction from $\mathrm{OR} \circ \mathsf{IND}_\ell^n$ to an instance of set disjointness. By Lemma 5.10, the NP^{cc}-complexity of $\mathrm{OR}_n \circ \mathsf{IND}_\ell^n$ is $\log \ell + \log n + 1$. That is, there exists a cover of the 1s of the communication matrix of $\mathrm{OR}_n \circ \mathsf{IND}_\ell^n$ with at most $2n\ell$ rectangles. Enumerating the 1-rectangles gives us an instance of set disjointness: on input (x, y) to $\mathrm{OR}_n \circ \mathsf{IND}_\ell^n$, Alice constructs the $2n\ell$-bit string which is the indicator vector $I_x(x)$ of the set of 1 rectangles which x belongs to, similarly Bob constructs $I_y(y)$ the same for y. Thus $\mathrm{OR}_n \circ \mathsf{IND}_\ell^n(x, y) = 1$ iff $\mathrm{DISJ}_{2n\ell}(I_x(x), I_y(y)) = 1$.

Therefore, the lower bound of $\Omega(n \log \ell)$ on the cost of computing $\mathrm{OR}_n \circ \mathsf{IND}_\ell^n$, from Lemma 5.11, implies a lower bound of $\Omega((n \log \ell)^{1/5})$ on the cost of any real communication protocol computing DISJ_n. $\square$

Corollary 5.13. *Real communication protocols cannot be balanced.*

Proof. We begin by giving a cost n, size $2n + 1$ real communication protocol for $\mathrm{DISJ}_n = \vee_{i=1}^n (x_i \wedge y_i)$. Sequentially from $i = 1, \ldots, n$, Alice and Bob solve $x_i \wedge y_i$. To do this, Alice sends x_i to the referee and Bob sending $2 - y_i$. Observe that $x_i \wedge y_i = 1$ iff $x_i \geq 2 - y_i$. This protocol contains exactly $2n$ nodes, one for each query to $x_i \wedge y_i$ for $i \in [n]$, one for each node announcing that $x_i = y_i = 1$ and a single node announcing that $x \cap y = \emptyset$.

Suppose by contradiction that one could balance real communication protocols. The size $2n+1$ protocol would therefore imply a cost $\log(2n+1)$ real protocol for DISJ_n, contradicting the lower bound from Theorem 5.12. $\square$

6 Conclusions

This paper introduces and develops the Stabbing Planes proof system as a natural extension of DPLL and pseudoBoolean solvers to handle a more expressive set of queries. Although it is equivalent to a tree-like version of a system already in the literature, this new perspective turns out to be quite useful for proving upper bounds. This paper is only a preliminary exploration of the SP proof system and leaves open many interesting problems from both a theoretical as well as a practical perspective.

1. As mentioned in the preliminaries, we do not have an analog to Cook et al. [8] for Stabbing Planes and so it is unknown whether for SP refutations, the length and size (number of bits) can be treated as the same measure. That is, is it possible to prove that any SP refutation of *length l* can be simulated by an SP planes refutation of size $\mathsf{poly}(l, n)$?

2. We have shown that that Cutting Planes refutations of small rank can be simulated by Stabbing Planes refutation of small rank, and that Cutting Planes refutations of small

size can also be simulated by SP refutations of small size. Can we simulate both rank and size efficiently? That is, can any CP refutation of rank r and size s be simulated by a SP refutation of rank $\mathsf{poly}(r)$ and size $\mathsf{poly}(s)$?

3. Prove superpolynomial lower bounds for SP. Krajíček [23] gave exponential lower bounds on the length of R(CP) refutations when both the width of the clauses, and the size of the coeficients appearing in the inequalities are sufficiently bounded. This was later improved by Kojevnikov [22] to remove the restriction on the size of the coeffi-cients for tree-like R(CP). In particular, to obtain any lower bound at all, the width of the clauses appearing in the R(CP) refutations must be bounded by $o(n/\log n)$. From Theorem 4.2, a size S and rank D SP refutation implies an R(CP) proof of size $O(S)$ and width $O(D)$. Therefore, this result is also a size lower bound for bounded-depth SP. Unfortunately, it appears that these techniques are fundamentally limited to be ap-plicable only to SP refutations with low depth, and so new techniques seem needed to overcome this barrier.

4. As mentioned in the introduction, we feel that SP has potential, in combination with state-of-the-art algorithms for SAT, for improved performance on certain hard instances, or possibly to solve harder problems such as maxSAT or counting satisfying assign-ments, possibly in conjunction with solvers such as CPLEX. The upper bound on the Tseitin example illustrates the kind of reasoning that SP is capable of: arbitrarily split-ting the solution space into sub-problems based on some measure of progress. This opens up the space of algorithmic ideas for solvers and should allow one to take fuller advantage of the expressibility of integer linear inequalities. For example, since geo-metric properties of the rational hull formed by the set of constraints can be determined efficiently, an SP-based solver could branch on linear inequalities representing some geometric properties of the rational hull. Therefore, it is an open problem to realize a SP based solvers or to implement SP-like branching in conjunction with current solvers such as CPLEX.

5. Separate CP and SP. It has been a long-standing conjecture that CP does not have short refutations of the Tseitin formulas. The intuition for this is that CP is unable to count $\mod 2$. On the other hand, Theorem 5.3 gives a quasi-polynomial upper bound on the Tseitin formulas in SP. Therefore, a natural approach to separating SP and CP is through resolving this conjecture for CP.

References

[1] Paul Beame, Trinh Huynh, and Toniann Pitassi. Hardness amplification in proof com-plexity. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 87–96, 2010.

[2] Eli Ben-sasson, Russell Impagliazzo, and Avi Wigderson. Near-optimal separation of treelike and general resolution. Technical report, Electronic Colloquium in Computation Complexity, 2000.

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

[4] Daniel Le Berre. Handling Pseudo-Boolean constraints in a CDCL solver: a practical survey. In *Dagstuhl Seminar 15171: Theory and Practice of SAT Solving*, April 2015. http://materials.dagstuhl.de/files/15/15171/15171.DanielLe

[5] M. Bonet, J. Esteban, N. Galesi, and J. Johannsen. On the Relative Complexity of Resolution Refinements and Cutting Planes Proof Systems. *SIAM Journal on Computing*, 30(5):1462–1484, January 2000.

[6] Maria Luisa Bonet, Juan Luis Esteban, Nicola Galesi, and Jan Johannsen. On the relative complexity of resolution refinements and cutting planes proof systems. *SIAM J. Comput.*, 30(5):1462–1484, 2000.

[7] 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.

[8] William Cook, Collette Coullard, and György Turán. On the complexity of cutting-plane proofs. *Discrete Applied Mathematics*, 18(1):25 – 38, 1987.

[9] Martin Davis, George Logemann, and Donald Loveland. A machine program for theorem-proving. *Commun. ACM*, 5(7):394–397, July 1962.

[10] Martin Davis and Hilary Putnam. A computing procedure for quantification theory. *J. ACM*, 7(3):201–215, 1960.

[11] Susanna F. de Rezende, Jakob Nordstrom, and Marc Vinyals. How limited interaction hinders real communication (and what it means for proof and circuit complexity). In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 295–304, 2016.

[12] Yuval Filmus, Pavel Hrubes, and Massimo Lauria. Semantic versus syntactic cutting planes. In *33rd Symposium on Theoretical Aspects of Computer Science, STACS 2016, February 17-20, 2016, Orléans, France*, pages 35:1–35:13, 2016.

[13] Noah Fleming, Denis Pankratov, Toniann Pitassi, and Robert Robere. Random cnfs are hard for cutting planes. *Electronic Colloquium on Computational Complexity (ECCC)*, 24:45, 2017.

[14] Mika Göös and Toniann Pitassi. Communication lower bounds via critical block sensitivity. *CoRR*, abs/1311.2355, 2013.

[15] Mika Göös, Toniann Pitassi, and Thomas Watson. Deterministic communication vs. partition number. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 1077–1088, 2015.

[16] Pavel Hrubes and Pavel Pudlák. A note on monotone real circuits. *Electronic Colloquium on Computational Complexity (ECCC)*, 24:48, 2017.

[17] Pavel Hrubes and Pavel Pudlák. Random formulas, monotone circuits, and interpolation. *Electronic Colloquium on Computational Complexity (ECCC)*, 24:42, 2017.

[18] Trinh Huynh and Jakob Nordstrom. On the virtue of succinct proofs: Amplifying communication complexity hardness to time-space trade-offs in proof complexity. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing*, STOC '12, pages 233–248, New York, NY, USA, 2012. ACM.

[19] IBM ILOG. The CPLEX optimizer. https://www-01.ibm.com/software/commerce/optimization/cplex-optimizer/.

[20] Russell Impagliazzo, Toniann Pitassi, and Alasdair Urquhart. Upper and lower bounds for tree-like cutting planes proofs. In *Proceedings of the Ninth Annual Symposium on Logic in Computer Science (LICS '94), Paris, France, July 4-7, 1994*, pages 220–228, 1994.

[21] Stasys Jukna. *Boolean function complexity : advances and frontiers*. Algorithms and combinatorics. Springer, Berlin, Heidelberg, 2012.

[22] Arist Kojevnikov. Improved lower bounds for tree-like resolution over linear inequalities. In *Theory and Applications of Satisfiability Testing - SAT 2007, 10th International Conference, Lisbon, Portugal, May 28-31, 2007, Proceedings*, pages 70–79, 2007.

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

[24] Jan Krajek. Discretely Ordered Modules as a First-Order Extension of the Cutting Planes Proof System. *The Journal of Symbolic Logic*, 63(4):1582–1596, 1998.

[25] Jan Krajek. Interpolation by a Game. *Mathematical Logic Quarterly*, 44(4):450–458, January 1998.

[26] Saburo Muroga. *Threshold logic and its applications*. Wiley-Interscience, 1972.

[27] Noam Nisan. The communication complexity of threshold gates. In *In Proceedings of Combinatorics, Paul Erdos is Eighty*, pages 301–315, 1994.

[28] Pavel Pudlák. Lower bounds for resolution and cutting plane proofs and monotone computations. *J. Symb. Log.*, 62(3):981–998, 1997.

[29] Ran Raz and Pierre McKenzie. Separation of the monotone NC hierarchy. *Combinatorica*, 19(3):403–435, 1999.

[30] Olivier Roussel and Vasco M. Manquinho. Pseudo-Boolean and cardinality constraints. In Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors, *Handbook of Satisfiability*, chapter 22, pages 695–733. IOS Press, 2009.

[31] Günter M. Ziegler. *Lectures on Polytopes*. Springer-Verlag, New York, 1995.