Rank Bounds and Integrality Gaps for Cutting Planes Procedures

Josh Buresh-Oppenheim*
Computer Science Department
University of Toronto
Toronto, ON M5S 3G4
bureshop@cs.toronto.edu

Nicola Galesi*
Universitat Politecnica de Catalunya
galesi@lsi.upc.es

Shlomo Hoory*
Computer Science Department
University of Toronto
Toronto, ON M5S 3G4
shlomoh@cs.toronto.edu

Avner Magen*
Computer Science Department
University of Toronto
Toronto, ON M5S 3G4
avner@cs.toronto.edu

Toniann Pitassi*
Computer Science Department
University of Toronto
Toronto, ON M5S 3G4
toni@cs.toronto.edu

Abstract

We present a new method for proving rank lower bounds for Cutting Planes and several procedures based on lifting due to Lovász and Schrijver when viewed as proof systems for unsatisfiability. We apply this method to obtain the following new results: First, we prove near-optimal rank bounds for Cutting Planes and Lovász-Schrijver proofs for several prominent unsatisfiable CNF examples, including random kCNF formulas and the Tseitin graph formulas. It follows from these lower bounds that a linear number of rounds of CP or LS procedures when applied to relaxations of integer linear programs is not sufficient for reducing the integrality gap. Secondly, we give unsatisfiable examples that have constant rank CP and LS proofs but that require linear rank Resolution proofs. Thirdly, we give examples where the CP rank is $O(\log n)$ but the LS rank is linear. Finally, we address the question of size versus rank. We show that, for both proof systems, rank does not accurately reflect proof size. That is, there are examples with polynomial-size CP/LS proofs, but requiring linear rank.

1. Introduction

Integer linear programming is the problem of optimizing a linear objective function over the integral points of a given (bounded or unbounded) polyhedron. In his seminal paper, Khachian [22] proposed the ellipsoid method for (nonintegral) linear programming, showing that the optimization problem over a polytope is polytime. The additional integrality constraints change the complexity of the problem dramatically: it is well-known that general integer LP is NP-hard. In both the unrestricted and the integral cases, one can also look at feasibility problems instead of at optimization problems. Here, the question is whether a given polytope is empty. The feasibility problem is closely related to the linear optimization problem, and here too the nonintegral version is easy while the integral one is NP-complete.

Cutting planes methods for integer linear programming are instrumental in bridging the gap between the true, computationally complex structures (the integral solutions to the problem, or, rather, their convex hull) and their relaxed counterpart, which are generally simple polytopes that contain the convex hull of the integral solutions but also contain other, extraneous nonintegral points. These are methods in which the initial, relaxed polytope P is transformed through a sequence of ever-decreasing (contained) polytopes to the integral hull of P, ie the smallest polytope containing the integral points of P. In this sequence, a polytope is produced from its predecessor by using the integrality constraint locally. A simple example of this kind of reasoning is that if one knows that a certain coordinate is at least β, then a stronger conclusion, that this coordinate is at least $\lceil\beta\rceil$, is valid for the integral hull of P. For optimization problems this sequence of polytopes produces a set of optimal values

*Funded in part by a grant from NSERC

that get closer and closer to the desired optimal integral solution, and for feasibility problems, the sequence terminates with the empty polytope if and only if the initial polytope contained no integral points.

From the complexity standpoint, there are three important desirable properties of the above sequence: (i) the local operations transforming a polytope to its successor are efficient (ii) the length of the sequence is small, and (iii) there is an efficient algorithm producing the sequence. In the feasibility integer programming problem, a proof can be associated with the above sequence of polytopes. The local steps are applied by the underlying rules of the system. Properties (i) and (ii) guarantee a small size proof, while (iii) guarantees automatizability of the proof system.

In this paper, we study several prominent cutting planes methods: Chvátal-Gomory cuts [6, 15], and a collection of matrix cut operations defined by Lovász and Schrijver [24]. These branch-and-cut methods for integer programs are currently among the most important techniques for solving a range of NP-hard 0/1 optimization problems. There are two standard complexity measures of interest for these procedures: rank and size. The size is the total number of cut operations that must be applied and the rank is the total number of rounds of cut operations that must be applied. In the terminology of proofs, the rank is therefore the minimum possible depth of a proof of unsatisfiability in the corresponding proof system.

Superpolynomial lower bounds on size for a cutting planes method are important since they show that *any* algorithm for satisfiability that produces a cutting planes proof will not be polynomial-time. Superpolynomial size lower bounds are known for the Chvátal-Gomory cutting planes method [25]. There are three distinct types of matrix cuts defined by Lovász and Schrijver, LS_0, LS and LS_+. Exponential lower bounds have been proven for LS_0 [9, 10]. For LS and LS_+, no nontrivial size bounds are known.

Rank is another natural measure that has been studied and captures the amount of inherent sequentialism in a proof. In some proof systems, there is a natural rank-based procedure for generating a proof which is practical in certain cases. For example, a rank-based method for Resolution is the familiar Davis-Putnam procedure, and a rank-based method for the Polynomial Calculus is a variation on the Gröbner basis algorithm. In both of these cases, it is important that it can be determined if there is a d round/rank derivation in time at most $n^{O(d)}$. It turns out that matrix cut systems have a somewhat similar property and therefore rank is a particularly interesting measure in this case. In [24] it was shown that for any polytope P, if one could optimize over P efficiently, then there is an algorithm for optimizing over P' efficiently, where P' is the polytope obtained by applying one "round" of commutative (LS) or noncommutative (LS_0) matrix cuts. Using similar arguments it can be shown that the same is true when considering the feasibility question rather than optimization. It follows that there is a deterministic algorithm that can "search through" all LS proofs of rank d in time $n^{O(d)}$. While this holds for other proof systems such as Resolution, it is less obvious here because the number of faces in the rank-r polytope is not easily bounded, even for small r.

Recently some limitations on the rank-based application of the LS procedure to the problem of approximating vertex cover [1] were shown. In this paper, we study limitations of all the above-mentioned cutting planes methods both in the case of unsatisfiable CNF formulas and optimization problems. We present a new method for proving rank lower bounds that applies to both Chvátal-Gomory cutting planes and matrix-cut proof systems. This method can be viewed as a game which produces a tree of (nonintegral) points in the polytope, whose depth is a lower bound on the rank of the polytope in all of the above proof systems. This game allows us to prove asymptotically tight rank bounds for many classes of unsatisfiable boolean formulas, especially those which contain some measure of expansion, like random kCNFs and the Tseitin principle on expander graphs. Prior to our result, the only rank bounds for unsatisfiable boolean examples were for the clique-coclique [25] formulas in Chvátal-Gomory cutting planes, and for the PHP in LS [16]. We then supply a particular optimization problem where cutting planes procedures are not helpful in the sense that after linearly-many rounds, the large integrality gap of the relaxation of the problem does not change at all. To the best of our knowledge there are no results of this form (see also [1, 12]) that give hardness for more than a logarithmic-number of rounds. Next, we give examples separating the LS-, CP-, and Resolution-rank, and examples with polynomial-size Resolution/CP/LS proofs, that require large rank.

The rest of the paper is organized as follows. In Section 2 we define the Resolution/CP/LS proof systems, and give some background. In Section 3 we provide a general scheme for proving rank lower bounds. In Section 4 we prove rank lower bounds when the constraints are expanding. Section 5 deals with integrality gaps that are based on our rank lower bounds. Section 6 gives various separation examples for LS-, CP-, and Resolution-rank. Section 7 gives an example where both the Resolution/CP/LS proof size and rank are polynomial. In Section 8 we describe an algorithm for showing the unsatisfiability of formulas of LS-rank d in time $n^{O(d)}$ based on the results of [24].

2. Definitions and Background

Resolution: Resolution proofs work with clauses, viewed as sets of literals. If C and D are sets of literals, then the clause $(C \vee D)$ is derivable from the clauses $(x \vee C)$ and

$(\neg x \vee D)$ by the Resolution rule. A resolution refutation of a CNF formula f is a sequence of clauses $C_1, \ldots, C_q$ such that each clause is either a clause of f, or follows from two previous clauses by the resolution rule, and the final clause, C_q, is the empty clause. Let S be a resolution refutation of a CNF formula f, represented as a directed acyclic graph (with nodes corresponding to clauses). The *size* of S is the number of clauses in S; the *depth* or *rank* of S is the depth of the directed acyclic graph. The resolution *size* (or depth) of f is the minimal size (depth) over all resolution refutations of f. S is *tree-like* if the directed acyclic graph is a tree.

Proof systems based on linear programming: We describe several proof systems for systems of linear inequalities where the values of the variables are restricted to be boolean. In these proof systems, we begin with a polytope P defined by linear inequalities associated with the logical formulation of the problem. In the more common case of CNF-formulas we convert clauses to inequalities in the obvious way, eg $x_1 \vee \neg x_2 \vee x_3$ is converted to $x_1 + (1-x_2) + x_3 \geq 1$. Notice that the $0/1$ solutions to these inequalities are exactly the satisfying boolean assignments to the formula. Relaxing to $0 \leq x_i \leq 1$ makes the set of solutions a polytope whose integral points are the solutions to the original problem.

Two classes of cutting planes methods useful for either general integer programs or specifically for $0-1$ problems are Gomory-Chvátal cutting planes [6, 14], and matrix cuts, defined by Lovász and Schrijver [24]. We begin by describing Gomory-Chvátal cutting planes. This proof system is referred to in the literature as simply Cutting Planes (CP). Consider the following two rules: (1) (Linear combinations) From $f_1, \ldots, f_k$ derive $\sum_{i=1}^{k} \lambda_i f_i$, where λ_i are positive rational constants; (2) (Rounding) From $f - \lambda$ derive $f - \lceil \lambda \rceil$, provided that the coefficients of f are integers. Without loss of generality, we can assume that a rounding operation is always applied after every application of rule (1), and thus we can merge (1) and (2) into a single rule, called a *Chvátal-Gomory cut*.

Definition 1. *A Cutting Planes (CP) refutation for $f = f_1, \ldots, f_m$ is a sequence of linear inequalities, $g_1, \ldots, g_q$ such that each g_i is either an inequality from f, or an axiom ($x \geq 0$ or $1 - x \geq 0$), or follows from previous inequalities by a Chvátal-Gomory cut, and the final inequality g_q is $0 \geq 1$.*

There are several cutting planes proof systems defined by Lovász and Schrijver, collectively referred to as matrix cuts. These system allows one to "lift" the linear inequalities to degree-two polynomials, and then project back using the fact that $x^2 = x$ for $x \in \{0, 1\}$ to linear inequalities. To see that the definitions below are equivalent to the original definitions of Lovász and Schrijver ([24]), see [9].

Definition 2. *Given a polytope $P \subseteq [0,1]^n$ defined by $a_i x \geq b_i$ for $i = 1, 2, \ldots, m$:*

(1) An inequality $d - \langle c, x \rangle \geq 0$ is called an N-cut for P if

$$\begin{aligned} d - \langle c, x \rangle &= \sum_{i,j} \alpha_{ij} (b_i - \langle a_i, x \rangle) x_j \\ &+ \sum_{ij} \beta_{ij} (b_i - \langle a_i, x \rangle)(1 - x_j) \\ &+ \sum_{j} \lambda_j (x_j^2 - x_j), \end{aligned}$$

where $\alpha_{ij}, \beta_{ij} \geq 0$ and $\lambda_j \in R$ for $i = 1, \ldots, m$, $j = 1, \ldots, n$.

(2) A weakening of N-cuts, called N_0-cuts can be obtained if when simplifying to the linear term $d - \langle c, x \rangle$, we view $x_i x_j$ as distinct from $x_j x_i$.

(3) An inequality $d - \langle c, x \rangle$ is called an N_+-cut if

$$\begin{aligned} d - \langle c, x \rangle &= \sum_{i,j} \alpha_{ij} (b_i - \langle a_i, x \rangle) x_j \\ &+ \sum_{ij} \beta_{ij} (b_i - \langle a_i, x \rangle)(1 - x_j) \\ &+ \sum_{j} \lambda_j (x_j^2 - x_j) + \sum_{k} (g_k + \langle h_k, x \rangle)^2, \end{aligned}$$

where again $\alpha_{ij}, \beta_{ij} \geq 0$, $\lambda_j \in R$ for $i = 1, \ldots, m$, $j = 1, \ldots, n$ and $g_k + \langle h_k, x \rangle$ is a linear function for $k = 1, \ldots, n+1$.

The operators N, N_0 and N_+ are called the *commutative*, *non-commutative* and *semidefinite* operators, respectively. All three are collectively called *matrix-cut* operators.

Definition 3. *A Lovász-Schrijver (LS) refutation for f is a sequence of inequalities $g_1, \ldots, g_q$ such that each g_i is either an inequality from f or follows from previous inequalities by an N-cut as defined above, and such that the final inequality is $0 \geq 1$. Similarly, a LS_0 refutation uses N_0-cuts and LS_+ uses N_+-cuts.*

Definition 4. *Let $\mathcal{P}$ be one of the proof systems CP, LS, LS_0 or LS_+. Let S be a $\mathcal{P}$-refutation of f, viewed as a directed acyclic graph. The inequalities in S are represented with all coefficients in binary notation. The size of S is the sum of the sizes of all inequalities in S; the rank of S is the depth of the underlying directed acyclic graph. For a set of boolean inequalities f, the $\mathcal{P}$-size of f is the minimal size over all $\mathcal{P}$ refutations of f. Similarly the $\mathcal{P}$-rank of f is the minimal rank over all $\mathcal{P}$-refutations of f.*

Note that in our definition of these cutting planes systems, we can derive a new inequality from any number of previous inequalities in one step, whereas for Resolution, we are restricted to fanin-two However, in light of Caratheodory's theorem, we can assume wlog that the fanin in is at most $n+1$ in CP and n^2+n+1 in LS, and so the rank and size would not increase significantly if instead our proof systems were defined to have fanin 2.

Definition 5. *A refutation system A p-simulates a refutation system B (over the same language) if for every $x \in L$, the length of the shortest refutation of x in A is bounded by a polynomial in the length of the shortest proof of x in B.*

By definition LS p-simulates LS_0 and LS_+ p-simulates LS, and these simulations are rank preserving. Moreover for unsatisfiable CNF formulas, CP, LS_0, LS and LS_+ can all p-simulate Resolution and this simulation is rank-preserving [8]. It has also been shown that CP can p-simulate small-weight LS_0 [20]. In terms of negative results for simulations, the propositional pigeonhole principle (PHP) provides a family of unsatisfiable CNF examples requiring exponential-size Resolution proofs [19] but with polynomial-size CP, LS_0, LS and LS_+ proofs [8]. For CP and LS_0, exponential size lower bounds for one specific family of boolean examples are known [25, 10]. For LS and LS_+, no superpolynomial lower bounds are known.

Now let us review what is known with respect to rank. Any system of linear inequalities has a rank n LS proof. For CP, the rank of any polytope in the unit cube is at most $O(n^2 \log n)$, and moreover there are examples requiring CP-rank more than n [11]. However for unsatisfiable examples, the CP-rank is at most n [4]. For CP, linear rank bounds for unsatisfiable CNF examples were first obtained in [7]; however, these examples have exponentially-many faces (inequalities) and thus the rank is still small in the input size. Linear rank bounds for CP (as a function of the input size) for unsatisfiable CNF examples were first proven in [21], and also follow from the size bounds [25]. For LS, linear rank lower bounds for PHP were proven in [16]. In summary, the only known high-rank, unsatisfiable CNF examples were the clique-coclique formulas for CP and the PHP for LS. In this paper, we prove rank bounds for all of these proof systems for several sets of boolean inequalities satisfying certain combinatorial conditions.

3. Proving Rank Lower Bounds

Instead of looking at the halfspaces that can be generated syntactically by any of the these proof systems from the axioms, we look at their dual: the set of points that are still valid for all these halfspaces at a given stage of the proof.

For a given set of linear inequalities L that define a polytope in $\mathbb{R}^n$, let $P_L = P_L^{(0)}$ be that polytope. Given a proof system (CP, LS_0, LS, LS_+), let $P_L^{(i)}$ denote the polytope defined by all inequalities that can be derived in depth i from the initial inequalities in that proof system. Clearly $P_L^{(i+1)} \subset P_L^{(i)}$. The *rank* of L (or P_L) is the minimal i such that $P_L^{(i)}$ is the convex hull of the integral points in P_L. The *rank* of a point $x \in \mathbb{R}^n$ with respect to P_L is the minimal i such that $x \notin P_L^{(i)}$. That the rank of any bounded polytope in any of these proof systems is finite is a well-known fact ([15, 6, 24]). Note that, if P contains no integral points, then the rank of the polytope is the maximal rank of its points.

In what follows, we give methods for proving rank lower bounds for many natural, polysize sets of contradictory linear inequalities. These lower bounds follow by characterizing some of the points in $P_L^{(i)}$ that survive in $P_L^{(i+1)}$. We call these characterizations "protection lemmas," because they argue that certain points are protected from removal in the next round provided certain points survived the current round. These sorts of lemmas have been used in the past to prove rank lower bounds for specific polytopes in specific proof systems ([7, 13]). We develop a common protection lemma that works for many examples in any of the proof systems. Moreover, we define a simple, two-player game that uses this common protection lemma to establish lower bounds.

Protection Lemmas: For $x \in \mathbb{R}^n$, $e \in \{1,\ldots,n\}$, and $a \in \mathbb{R}$, we denote by $x^{(e,a)}$ the point that is the same as x except that the e-th coordinate has value a. For $x \in \mathbb{R}^n$, we denote by $E(x)$ the set of coordinates on which x is non-integral.

CP: Let

$$P' \quad = \quad \{x \in P : \langle a,x\rangle \geq \lceil b \rceil \text{ whenever } a \in \mathbb{Z}^n, b \in \mathbb{R}, \text{ and } \langle a,y\rangle \geq b \text{ for all } y \in P\}.$$

Then, for $i \geq 0$, we have $P^{(i+1)} = (P^{(i)})'$.

Lemma 6 (CP Lemma). *The following holds for CP: Let P be a bounded polytope in $\mathbb{R}^n$. Let $x \in \frac{1}{2}\mathbb{Z}^n$, and let $E = E(x)$ be partitioned into sets $E_1, E_2 \ldots, E_t$. Suppose that for every $j \in \{1,2,\ldots,t\}$ we can represent x as an average of vectors in $P^{(k)}$ that are 0-1 on E_j and agree with x elsewhere. Then $x \in P^{(k+1)}$.*

Proof. Assume for contradiction that $x \notin P^{(k+1)}$. Then there is a vector $a \in \mathbb{Z}^n$ and a non integral scalar b, such that $\langle a,y\rangle \geq b$ for all $y \in P^{(k)}$ and $\langle a,x\rangle < \lceil b \rceil$. Clearly $x \in P^{(k)}$, being an average of points in that polytope. So $\langle a,x\rangle \geq b$ and it follows that $\langle a,x\rangle$ must be in $\frac{1}{2}+\mathbb{Z}$. Thus $\sum_{e\in E(x)} a_e$ must be odd, and since $\sum_{e\in E(x)} a_e = \sum_i \sum_{e\in E_i} a_e$, there is a j such that $\sum_{e\in E_j} a_e$ is odd. Consider the set of vectors $V \subset P^{(k)}$ that average to x and that differ from x exactly on E_j where they take 0/1 values. Since $\sum_{e\in E_j} a_e$ is odd we can see that $\langle a,v\rangle$ are integrals for all $v \in V$. But then $\langle a,v\rangle \geq \lceil b \rceil$. Since x is an average of the $v \in V$, we also get $\langle a,x\rangle \geq \lceil b \rceil$. Contradiction. □

LS_0: Let

$$P' = \cap_{i=1}^n \text{conv}(P \cap \{x_i = 0\}, P \cap \{x_i = 1\}), \qquad (1)$$

where $\{x_i = a\}$ is an abbreviation for $\{x \in \mathbb{R}^n : x_i = a\}$. [24] proves that for $i \geq 0$, $P^{(i+1)} = (P^{(i)})'$.

The following lemma is immediate from equation (1):

Lemma 7 (LS$_0$ Lemma). *The following holds for LS$_0$: Let $P \subset [0,1]^n$ be a polytope, and x be a point in P. Then, if for any $i \in E(x)$ there is a set of points $S_i \subset P^{(k)}$ with i-th coordinate in $\{0,1\}$ such that $x \in conv(S_i)$, then $x \in P^{(k+1)}$.*

LS and LS$_+$: The following lemma is from [13]:

Lemma 8 (LS/LS$_+$ Lemma). *The following holds for LS and LS$_+$: Let $P \subset [0,1]^n$ be a polytope, and x be a point in P. If, for any $i \in E(x)$, $x^{(i,0)}, x^{(i,1)} \in P^{(k)}$, then $x \in P^{(k+1)}$.*

A Game: Lemmas 6, 7 and 8 all conclude the same thing from different hypotheses. We now state a protection lemma that holds for all of the proof systems because it uses a hypothesis that is weaker than any of those in the previous protection lemmas:

Lemma 9 (Game Lemma). *The following holds for CP, LS$_0$, LS and LS$_+$: Let $P \subset [0,1]^n$ be a polytope, and $x \in \frac{1}{2}\mathbb{Z}^n \cap P$. If, for any $i \in E(x)$, $x^{(i,0)}, x^{(i,1)} \in P^{(k)}$, then $x \in P^{(k+1)}$.*

This lemma gives us the following Prover-Adversary game for showing a lower-bound on the rank of a point $w \in \frac{1}{2}\mathbb{Z}^n$ with respect to P. We think of the Prover as trying to show that w has high rank, while the Adversary is trying to foil that proof. The game proceeds in rounds. During each round, there is a current point $x \in \frac{1}{2}\mathbb{Z}^n$, whose initial value is w. At each round, the Prover chooses between the two following types of moves:

1. *Prover-move*: The Prover generates a set of points Y such that x is a convex combination of those points. The Adversary selects one point $y \in Y$ to be the new x.

2. *Adversary-move*: The Adversary selects a coordinate e such that x_e is $\frac{1}{2}$ and a value $a \in \{0,1\}$. The new x is $x^{(e,a)}$.

The game ends when x is no longer in P. The Prover gets one point for each Adversary-move.

Lemma 10. *If the Prover has a strategy to earn m points against any adversary, then the (CP, LS$_0$, LS, or LS$_+$)-rank of w with respect to P is at least m.*

Proof. By induction on r, the maximum number of rounds in the strategy. If $r = 0$, then $m = 0$, but the rank of w can never be less than 0. For arbitrary $r > 0$, the Prover can start by making a Prover-move or an Adversary-move. If it is a Prover-move, then the Prover presents Y and, no matter which $y \in Y$ the Adversary chooses, the Prover has a strategy to earn m points. By induction, each $y \in Y$ has rank at least m. By convexity, the rank of w, which is a convex combination of points in Y, is at least m. If it is an Adversary-move, then, no matter which e and a the Adversary chooses, the Prover has a strategy to earn $m-1$ points. By induction, $w^{(e,a)}$ has rank at least $m-1$ for all possible (e,a), so by Lemma 9, w has rank at least m. $\square$

4. Expanding Constraints

In what follows, we deal with F, a set of mod-2 equations over n variables. That is, each equation in F is of the form $\sum_{i \in S} x_i \equiv a \pmod 2$, where $S \subset [n]$ and $a \in \{0,1\}$. Notice that each such equation can be represented by the conjunction of $2^{|S|-1}$ clauses, each of which can be represented as a linear inequality. We denote by P_F the polytope bounded by these inequalities and by the inequalities $0 \le x_i \le 1$.

Let G_F the bipartite graph from the set F to the set of variables where each equation is connected to the variables it contains. We prove a rank lower bound for P_F as a function of the expansion of G_F.

Definition 11. *Let G be a bipartite graph from V to U. The boundary of a set $X \subset V$ is $\partial X \stackrel{d}{=} \{u \in U : |\Gamma(u) \cap X| = 1\}$. G is an (r,ε)-boundary expander if for all subsets $X \subset V$ where $|X| \le r$, we have $|\partial X| \ge \varepsilon |X|$. The boundary expansion of a set $X \subset V$ is the value $|\partial X|/|X|$.*

The reason that we require G_F to be a good expander is that it allows us to satisfy subsets of F:

Lemma 12. *Consider a set F of m mod-2 equations over n variables. Assume that for any variable v and any value $a \in \{0,1\}$, there is a solution to F where v assumes the value a. Then all the 0-1 solutions to F average to the all-$\frac{1}{2}$ assignment.*

Proof. Let $S_{v,a}$ be a solution to F in which variable v is set to a. It is easy to see that the mapping $S \mapsto S + S_{v,1} - S_{v,0}$ is a one to one mapping from solutions with $v = 0$ onto solutions with $v = 1$, Therefore the average over all solutions to F is $\frac{1}{2}$ on v. $\square$

Lemma 13. *Let F be a set of m mod-2 equations over n variables. Assume G_F is an (m,δ)-boundary expander for any $\delta > 0$. Then F has a 0-1 solution.*

Proof. There exists some variable v_1 in $\partial(F)$. Assume v_1 is connected to equation f_1. For some $1 \le i \le m$, assume we have pairings $(f_1, v_1), \ldots, (f_i, v_i)$. Now, there must be some $v_{i+1} \notin \{v_1, \ldots, v_i\}$ in $\partial(F \setminus \{f_1, \ldots, f_i\})$. It is connected to some equation f_{i+1} in $F \setminus \{f_1, \ldots, f_i\}$. Add (f_{i+1}, v_{i+1}) to the set of pairs. Eventually we have the set of pairs $(f_1, v_1), \ldots, (f_m, v_m)$. To satisfy F, set all variables not in $\{v_1, \ldots, v_m\}$ arbitrarily. Now, for $i = m$ to 1, set v_i so that it satisfies equation f_i (notice that in this order, v_i is the last unassigned variable of f_i). $\square$

We now use the game to show a rank lower bound for expanding sets of equations. For $x \in \frac{1}{2}\mathbb{Z}^n$, let $G_F(x)$ be the subgraph of G_F induced by the set of variables $E(x)$ and the set of equations connected to those variables.

Theorem 14. *Let $\varepsilon > 0$ and let $w \in \frac{1}{2}\mathbb{Z}^n$. If $G_F(w)$ is an $(r, 2+\varepsilon)$-boundary expander, then w has (CP, LS_0, LS, LS_+)-rank at least $r\varepsilon$ with respect to P_F.*

Proof. We start the game with $x = w$. Clearly $x \in P_F$ since each clause expressing F must contain at least two literals set to $\frac{1}{2}$ by the expansion requirement. Let $\Gamma_x(X)$ be the neighbor set of $X \subset F$ in $G_F(x)$. Let ℓ initially be set to r. The Prover's strategy is as follows:

1. Let the Adversary move as long as all subsets $X \subset F$ in $G_F(x)$ of size at most ℓ have boundary expansion > 2 in $G_F(x)$. Note that after such a move we have $x \in P_F$ since all equations in $G_F(x)$ have degree at least 2.

2. Let B be a maximal subset of equations in $G_F(x)$ with boundary expansion ≤ 2 such that $|B| \leq \ell$. Now the Prover moves. Let Y be the sets of all assignments satisfying B that are $0-1$ on $\Gamma_x(B)$ and that agree with x elsewhere. To see that Y is nonempty and that it does indeed average to x, consider an arbitrary variable v in $\Gamma_x(B)$ and an arbitrary value $a \in \{0,1\}$. By Lemma 12 it is enough to show that there is a point in Y in which v is set to a. Notice that B still has boundary-expansion greater than 0 on the graph $G_F(x)$ minus v, and so Lemma 13 implies that, regardless of the setting of v, there exists a $0-1$ assignment on $\Gamma_x(B) \setminus \{v\}$ satisfying B. The Adversary selects one $y \in Y$ to be the new x.

   Set ℓ to $\ell - |B|$. If $\ell = 0$, stop the game. Otherwise, we argue that $x \in P_F$. Indeed, in that case $|B|$ is strictly smaller than ℓ, and it is always the case that any equation f not in B has at least two neighbors in $G_F(x)$ since otherwise $b \bigcup \{f\}$ would also have boundary-expansion at most 2 contradicting the maximality of B.

3. Repeat until the game is over. See figure 1 in the appendix for a snapshot.

Now we will show that the Prover always earns at least $r\varepsilon$ points. Assume the game ends after k rounds of the strategy. For any round $i \leq k$, let B_i be the set of vertices designated in step 2 and let $S = \bigcup_{j=1}^{k} B_j$. The size of S is r, so S had a boundary of size at least $(2+\varepsilon)r$ in G_F. At the end of the game, S has no boundary (in fact it has no neighbors) in $G_F(x)$. At most $2r$ of these boundary nodes were removed by the Prover: at the beginning of step 2 of round i, B_i has at most $2|B_i|$ boundary nodes and every boundary node of S is a boundary node for exactly one B_i. Hence at least εr of S's original boundary nodes were removed by the Adversary. By Lemma 10, w has the required rank. $\square$

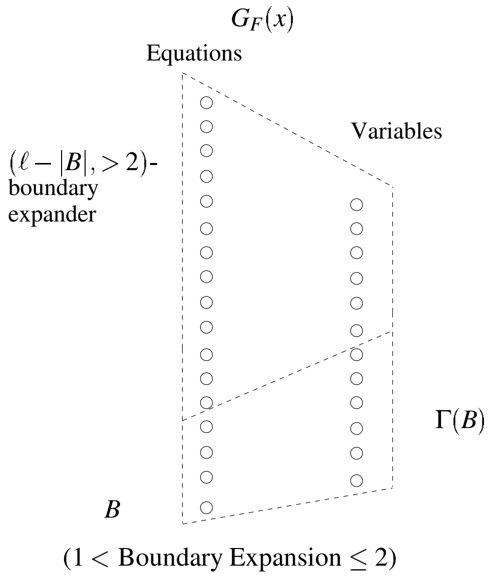

Figure 1. A snapshot of the strategy for Theorem 14.

It turns out that many common formulas are examples of boundary-expanding mod-2 equations.

Definition 15. *The Tseitin tautology for an odd-size graph $G = (V, E)$, denoted $TS(G)$, is the following: there exists no 0-1 edge assignment $\phi : E \to \{0,1\}$, such that for every vertex $v \in V$*

$$\sum_{u \in \Gamma(v)} \phi((v,u)) \equiv 1 \pmod 2.$$

Definition 16. *There are $2\binom{n}{k}$ linear, mod-2 equations over n variables that contain exactly k different variables. Let $\mathcal{M}_m^{k,n}$ be the probability distribution induced by choosing m of these equations uniformly and independently. There are $2^k\binom{n}{k}$ clauses over n variables that contain exactly k different variables. Let $\mathcal{N}_m^{k,n}$ be the probability distribution induced by choosing m of these clauses uniformly and independently.*

Theorem 14 enables us to prove the main result of this paper:

Corollary 17. *The following holds for CP, LS_0, LS and LSP:*
(1) The Tseitin tautology on a graph H has rank at least $(c-2)n/2$, where c is the edge-expansion of H;
(2) $F \sim \mathcal{M}_m^{k,n}$ for $k \geq 5$ and m linear in n requires rank $\Omega(n)$ with high probability;
(3) $C \sim \mathcal{N}_m^{k,n}$ for $k \geq 5$ and m linear in n requires rank $\Omega(n)$ with high probability.

Proof. Throughout, let w be the all $\frac{1}{2}$ point. (1) The edge-expansion of a graph $H=(V,E)$ is the density of the sparsest cut:

$$\min_{S\subset V, |S|\leq |V|/2} \frac{e(S,V\setminus S)}{|S|}.$$

Let w be the all $\frac{1}{2}$ point. It is easy to see that $G_{TS(H}(w)$ is an $(n/2,c)$-boundary-expander.

(2) It is well-known that $G_F(w)$ is an excellent expander: for any constant Δ, ε,k, where $m=\Delta n$, there exists a constant $\alpha>0$ such that $G_F(w)$ is almost always an $(\alpha n, k-1-\varepsilon)$-expander. Every (r,δ) bipartite expander graph on (V,U) where V has maximal degree d is an $(r, 2\delta-d)$-boundary-expander. Hence $G_F(w)$ is an $(\alpha n, k-2-2\varepsilon)$-boundary-expander. For $k\geq 5$ and small ε, the boundary-expansion is more than 2.

(3) $G_C(w)$, the bipartite graph associated with the clauses of C, is the same as $G_F(w)$ for random F. Generate C' by adding, for each $c\in C$, the following clauses: if c has an even (odd) number of positive literals, all clauses on the same variables as c that have an even (odd) number of positive literals. Clearly w's rank with respect to P_C is at least its rank with respect to $P_{C'}$, but C' is a set of $|C|$ mod-2 equations such that $G_{C'}(w)$ is an $(\alpha n, k-2-2\varepsilon)$-boundary expander (with high probability, given $\Delta, \varepsilon, k, \alpha$ as in (2). □

5. Integrality gaps from rank lower bounds

The problem MAX-k-SAT (MAX-k-XOR-SAT) is the following: given a set of k-clauses (mod-2 equations), determine the maximum number of clauses (equations) that can be satisfied simultaneously. This problem is well-studied in the theory of approximation algorithms and it is known that it cannot be well-approximated in polynomial time if $P\neq NP$. Here we show inapproximation results (that are unconditional) for a restricted class of approximation algorithms that involve CP- or LS-relaxations of a linear program. These algorithms are not necessarily polytime. Similar results have been shown for LS-relaxations of vertex cover ([1]) and maximum independent set ([12]). The former shows that a large integrality gap remains after $\Omega(\sqrt{\log n})$ rounds of LS and the latter, $\Omega(\log n)$ rounds.

Given a set of k-mod-2 equations $F=\{f_1,\ldots,f_m\}$ over variables $x_1,\ldots,x_n$, add a new set of variables $y_1,\ldots,y_m$. For each f_i: $\sum_{j\in I_i} x_j \equiv a_i \pmod 2$, let f_i' be the equation $y_i+\sum_{j\in I_i} x_j \equiv a_i+1 \pmod 2$. Let F' be the set of f_i''s. If y_i is 1, then f_i' is satisfied if and only if f_i is satisfied. Hence we want to maximize the linear function $\sum_{i=1}^m y_i$ over the constraints F' within the boolean cube. Call this linear program L_F. An r-round CP- (respectively, LS_0-, LS-, LS_+-) relaxation of L_F (or any linear program) is a linear program with the same optimization function but with any additional constraints that can be generated in depth r from the original constraints using CP (respectively, LS_0, LS, LS_+).

Theorem 18. *Let $k\geq 5$. For any constant $\varepsilon>0$, there are constants $\Delta,\beta>0$ such that if $F\sim \mathcal{M}^{k,n}_{\Delta n}$ then the integrality gap of any βn-round CP- (resp., LS_0-, LS-, LS_+-) relaxation of L_F is at least $2-\varepsilon$ with high probability.*

Proof. Given ε, fix $\Delta\geq (8-4\varepsilon+\varepsilon^2)/\varepsilon^2$. It is not hard to see, using a Chernoff bound and a union bound, that, with high probability, no boolean assignment satisfies more than a $1/(2-\varepsilon)$ fraction of F's equations. On the other hand, consider an assignment w that sets the variables $y_1,\ldots,y_{\Delta n}$ to 1 and sets $x_1,\ldots,x_n$ to $\frac{1}{2}$. Clearly, w satisfies all of F's equations. Furthermore, it is well-known that $G_{F'}(w)$ is almost surely an $(\alpha n, 2+\delta)$-boundary expander for some $\alpha,\delta>0$ that depend on Δ. Let $\beta=\alpha\delta$. Hence, by Theorem 14, w remains a feasible solution for any βn-round CP- (resp., LS_0-, LS-, LS_+-) relaxation of L_F. □

We can form a linear program L_C for a set of k-clauses C in an analogous manner. Similarly, for any $k\geq 5$ and any $\varepsilon>0$, there exists $\Delta,\beta>0$ such that if $C\sim \mathcal{N}^{k,n}_{\Delta n}$, then the integrality gap of any βn-round relaxation of L_C is at least $\frac{2^k}{2^k-1}-\varepsilon$ with high probability.

6. Separating CP, LS and Resolution ranks

We consider the following generalization of PHP_n, the Pigeonhole Principle on $n+1$ pigeons and n holes, first suggested in [3]. Let $G=(U,V,E)$ be a bipartite graph, where $|U|=n+1$ and $|V|=n$. The tautology $\mathrm{PHP}(G)$ is the statement that G doesn't have a perfect matching. The formal statement of this is (1) For each $i\in U$, $\sum_{j\in\Gamma(i)} x_{j,i}\geq 1$; (2) For all $j\in V$, $i,i'\in\Gamma(j)$, such that $i\neq i'$, $x_{j,i}+x_{j,i'}\leq 1$. The standard PHP_n is just $\mathrm{PHP}(K_{n+1,n})$, where $K_{n+1,n}$ is the complete $n+1,n$ bipartite graph.

In this section we show the following separations: (1) PHP_n has a LS-rank of n but a CP-rank of $O(\log n)$; (2) For an expander graph G with degrees at most d. the Resolution refutation of $\mathrm{PHP}(G)$ has rank of $\Omega(n)$, but both its LS-rank and CP-rank are $O(d)$.

The Resolution rank lower bound is proven in [3]. The LS, CP rank upper bound follows from the following reasoning. As observed in [16], it is possible to derive in both systems $\sum_{i\in\Gamma(j)} X_{j,i}\leq 1$ for all $j\in V$, in rank $O(d)$. The point is that the polytope, defined by adding these new inequalities is the empty polytope, and therefore we can get the desired contradiction in one LS or CP step.

These separation results are particularly interesting in light of the fact that in all three systems (CP, LS and Resolution) one can easily prove this tautology with polynomial size tree-like proofs. An interesting open question is whether rank is always at most logarithmic in the size of the smallest tree-like CP proof. By the above, the answer to the analogous question for Resolution and LS is negative.

For the separation result of the CP and LS ranks, we start with the upper bound on the CP-rank of PHP_n. This result was proved independently by [2].

Theorem 19. *The CP-rank of PHP_n is $O(\log n)$.*

Proof. For a subset $S \in \{1,2,\ldots,n\}$ and $1 \leq j \leq n-1$ let $f_{S,j}$ be the inequality $\sum_{i\in S} x_{ij} \leq 1$. We claim that it is possible to deduce from $f_{S,j}$ for every S of size k any $f_{T,j}$ with T of size $< 2k$ in one Chvátal cut. In other words, if $f_{S,j}$ are valid for $\mathrm{PHP}^{(r)}$ for every S of size k and every j, then $f_{T,j}$ is valid for $\mathrm{PHP}^{(r+1)}$ for every T of size $< 2k$. This means that for all j, $\sum_{i=1}^{n} x_{ij} \leq 1$ is valid for $\mathrm{PHP}^{(O(\log n))}$. On the other hand, no solution that satisfies these inequalities can satisfy all the axioms $\sum_{j=1}^{n-1} x_{ij} \geq 1$ for every i. Therefore $\mathrm{PHP}^{(O(\log n))} = \emptyset$, and the Chvátal-rank of PHP_n is $O(\log n)$. To see the claim, take any j and T of size $l < 2k$, and sum up with coefficients $1/\binom{l-1}{k-1}$ the inequalities $f_{S,j}$ over all subsets $S \subset T$ of size k. After rounding the deduced inequality is

$$\sum_{i\in T} x_{i,j} \leq \left\lfloor \frac{\binom{l}{k}}{\binom{l-1}{k-1}} \right\rfloor = \lfloor l/k \rfloor \leq 1, \qquad (2)$$

namely, $f_{T,j}$. A good way to think of (2) is that when using the symmetric sum, we only care about the average threshold for a single variable. In f_S, j it is $1/|S|$, and so basically all we do is take the threshold $x_i \leq 1/|S|$ and turn it into $\sum_{i\in T} x_i \leq |T|/|S|$, and if $|T < 2|S|$ we get $\sum_{i\in T} \leq \lceil |T|/|S| \rceil \leq 1$. □

In fact, this bound is tight by [2]. In light of the fact that LS_+ has constant-rank proofs of the PHP [17], LS_+ is separated from CP with respect to rank.

A linear lower bound for the LS-rank of PHP_n was given by [16]. We will give a proof for the LS_0-rank, which we think is simpler and more illuminating.

Theorem 20. *The LS_0-rank of PHP_n is $n-1$.*

Proof. PHP_2 consists of a single point, and its LS_0-rank is therefore 1. The proof proceeds by induction on n. By Lemma 7, for PHP_n, the LS_0-rank of the all $1/n$ matrix is 1 more than the minimal rank of the following matrices $l = 1,\ldots,n$. (Note that for all of them $x_{1,1}$ is in $\{0,1\}$). Let $x_{i,j}$ equal 1 if $i=1$ and $j=1$; equal 0 if $i=1$ or $j=1$ but not both; and equal $\frac{1}{n-1}$ otherwise. All those matrices are just the all $\frac{1}{n-1}$ matrices for PHP_{n-1}. By induction their LS_0-rank is at least $n-2$. □

The PHP has polynomial-size (tree-like) LS_0 proofs. The fact that LS requires rank $\Omega(n)$ for the PHP shows that for both LS and LS_0 proofs, large rank is not a good indicator of large size. Our rank upper bound for PHP in CP and the fact that LS_+ has constant-rank proofs of PHP leaves the question open for these two proof systems. In the next section, we give a different formula which shows that CP can have large rank and small size.

7. CP proofs with large rank and small size

In [2] it is shown that the size s of a CP proof of a tautology is $O(n^r)$ where n is the number of variables, and r is the CP-rank of the polytope associated with the tautology. Here we show an example where this bound is very far from being tight. Specifically, we show an example of a tautology which has a quadratic-size CP proof (in fact even a resolution proof with that size) and a linear CP-rank. It turns out that such a separation between size and rank can we witnessed by any formula that has polysize CP refutations, but requires exponential tree-like CP refutations ([2]).

The unsatisfiable formula we take is GT_n which is the negation of the property that every total ordering on n elements has a maximal element (alternatively, that a directed graph closed under transitivity and with no cycles of size two has a source node). The formula was introduced by [23] and is formulated using $n(n-1)$ variables. Stalmark and Bonet and Galesi ([26, 5]) show that GT_n has a polynomial refutation in the Resolution proof system, but requires width $\Omega(n)$. Since CP polynomially simulates resolution, there is a also a polynomial CP proof to the formula. In fact, a little tweaking of this refutation gives a rank $O(n)$ CP refutation, whereas the Resolution proof itself has rank $\Omega(n^2)$. It remains to show

Theorem 21. *The CP-rank and the LS_0-rank of the polyope associated with the GT_n is $\Theta(n)$.*

We associate a partial ordering $\prec$ on $[n]$ with a vector $x_\prec \in \{0,\frac{1}{2},1\}^{n(n-1)}$ by the assignment $x_{ij} = 0,1,\frac{1}{2}$ when i is smaller than, bigger than or incomparable to j respectively.

Definition 22. *A (partial) order $\prec$ is called s-scaled* if there is a partition of $[n]$ into sets $A_1, A_2, \ldots, A_s$, such that $\prec$ is a total ordering on any of the A_i's and is not defined between elements in different A_i's. Notice that we may look at a s-scaled order as a transitive graph that is the union of s complete (directional) graphs.

Claim 23. *If $\prec$ is s-scaled, $s > 2$ then $x_\prec$ remains after $s-3$ rounds of CP or LS_0 cuts.*

The claim immediately provides a lower bound of $n-2$ for the rank of P since the vector associated with the empty order (which is n-scaled) has that rank.

Proof. (of Claim 23) By induction on s. Suppose $\prec$ is 3-scaled. We need to show that $x_\prec \in P = P^{(0)}$. Transitivity inequalities clearly hold for three elements in the same A_i. A transitivity inequality that involves more than one A_i must contain at least two variables with value $\frac{1}{2}$ and therefore must be satisfied. The "no maximal element" inequalities also hold, because for every element there are at least two others to which it is not comparable, and the associated two halves alone satisfy the inequality. For a general s we let $x = x_{\prec}$. Notice that $E(x)$ is a set of all edges connecting different components of the graph when we associate $\prec$ with a graph which is a union of s complete graphs. We partition the edges in $E(x)$ to $\binom{s}{2}$ sets by the componenets they connect and argue that x and this partition satisfy the conditions of Lemma 6 with $k = s-4$. Indeed, for a choice of components A and B we denote by $\prec_A$ by the order which is the same as $\prec$ except all the elemens of A are bigger than those of B. Similarly we define $\prec_B$. It is easy to see that $x = (x_{\prec_A} + x_{\prec_B})/2$. Since $\prec_A, \prec_B$ are $(s-1)$-scaled we inductively have that $rank(x_{\prec_A}), rank(x_{\prec_B}) \geq s-3$, and by Lemma 6 $rank(x) \geq s-2$. Notice that since Lemma 6 is valid both to CP and to LS_0 'rank' here can be taken as CP-rank as well as LS_0-rank. $\square$

8. Automatizability of LS for small-rank CNF formulas

We use here the terms *strong separation oracle* and *face-complexity* see [18] for definitions.

Assume we are given a strong separation oracle for a polytope $P \subseteq [0,1]^n$ of face-complexity φ. Then, we show an algorithm for either LS or LS_0 proof systems, that checks if $P^{(r)}$ is empty with running time $\mathrm{poly}(n,\varphi)^r$. Note that for a polytope arising from CNF formulas, $\varphi = O(n)$, and consequently the running time is $n^{O(r)}$. The claim follows from the following lemmas:

Lemma 24. *A strong separation oracle for $P \subseteq [0,1]^n$ with face-complexity φ implies a strong separation oracle to $P^{(1)}$ with a polynomial running time.*

Lemma 25. *If a polytope $P \subseteq [0,1]^n$ has face-complexity φ, then $P^{(1)}$ has face-complexity bounded by $O(n^6 \cdot \varphi)$.*

Lemma 24 implies a strong separation oracle for $P^{(r)}$ with running time $\mathrm{poly}(n,\varphi)^r$. By Lemma 25 the face-complexity of $P^{(r)}$ is bounded by $\varphi \cdot n^{O(r)}$. Therefore, the claim follows using Theorem 6.4.9 from [18]. The result for LS_0 is obtained using very similar arguments.

Proof. (of lemma 24)

We are given a strong separation oracle for P, that given $x \in \mathbb{R}^n$ states either that $x \in P$ or otherwise supplies a hyperplane separating x from P.

Following the definition of [24], we move to the cone $\overline{P}$ in $\mathbb{R}^{n+1}$

$$\overline{P} = \{(a, a \cdot x_1, \ldots, a \cdot x_n) : a \geq 0 \text{ and } (x_1, \ldots, x_n) \in P\}.$$

It is easy to see that a strong separation oracle for P implies one for $\overline{P}$, and that the face-complexity of P and $\overline{P}$ are the same. We define a polytope $M(\overline{P})$ in $\mathbb{R}^{(n+1)^2}$ as the collection of $(n+1) \times (n+1)$ matrices Y satisfying (i) Y is symmetric, (ii) $Y_0 = \mathrm{diag}(Y)$, (iii) $Y_i \in \overline{P}$, (iv) $Y_0 - Y_i \in \overline{P}$, where we denote by $Y_0, \ldots Y_n$ the columns of Y, and by $\mathrm{diag}(Y)$ its diagonal.

$$P^{(1)} = \left\{ x \in \mathbb{R}^n : Y \in M(\overline{P}) \text{ and } Y_0 = \begin{pmatrix} 1 \\ x \end{pmatrix} \right\}.$$

Let $x \in \mathbb{R}^n$. Consider the following polytope $Q_{x,\overline{P}}$ in $\mathbb{R}^{(n+1)^2}$.

$$Q_{x,\overline{P}} = \left\{ Y \in M(\overline{P}) | Y_0 = \begin{pmatrix} 1 \\ x \end{pmatrix} \right\}.$$

By definition $x \in P^{(1)}$ if and only if $Q_{x,\overline{P}}$ is not empty. We first argue that $Q_{x,\overline{P}}$ has a separation oracle. To see that, observe that $Q_{x,\overline{P}}$ is an intersection of $O(n^2)$ halfspaces and hyperplanes, and $O(n)$ projection-preimages of P. Since the face complexity of $Q_{x,\overline{P}}$ is bounded by φ, we can apply [18] Theorem 6.4.9 to obtain an algorithm that checks whether $Q_{x,\overline{P}}$ is empty, and consequently whether $x \in P^{(1)}$. Assume now that $x \notin P^{(1)}$. Along the above run of the algorithm (ending with the conclusion $Q_{x,\overline{P}} = \emptyset$), the separation oracle for P has been invoked a polynomial number of times, resulting in a polynomial number of halfspaces containing P. Let R be the intersection of those halfspaces. The crucial point to note here is that $Q_{x,R} = \emptyset$. This is since $Q_{x,R}$ and $Q_{x,\overline{P}}$ are indistinguishable to this run of the algorithm.

Let $(a_j, \cdot) \geq b_j$ be the halfspaces defining R. By the duality theorem, there is a positive combination $\vec{\alpha}$ of the inequalities $(a_j, Y_i) \geq b_j$ and $(a_j, Y_0 - Y_i) \geq b_j$ plus a combination of the inequalities of $M(\overline{P})$, such that (i) the coefficient vector of the Y variables is 0 and (ii) the constant term is of the form $\sum \alpha_i x_i \geq b > 0$. On the other hand, if $x \in P^{(1)}$ then $Q_{x,R}$ is not empty and so the same combination cannot lead to a contradiction and so $\sum \alpha_i x_i \leq 0$. This provides the desired separation. The only thing left to is to find the combination (the vector of coefficients α) that leads to the above contradiction. Here we use the fact that R has a polynomial number of faces, and so to find the combination satisfying both (i) and (ii) above is nothing but solving a polynomial linear program. $\square$

Proof. (of lemma 25) The face-complexity of $M(\overline{P})$ is at most φ. By [18] Theorem 6.2.4 a, the vertex-complexity of $M(\overline{P})$ is at most $O(n^4\varphi)$. This bound also applies to the vertex-complexity of $\overline{P}^{(1)}$ since it is just a projection of $M(\overline{P})$. By Theorem 6.2.4 b, the face-complexity of $\overline{P}^{(1)}$ is $O(n^2 \cdot n^4\varphi)$, and the lemma follows. $\square$

References

[1] Arora, Bollobas, and Lovasz. Proving integrality gaps without knowing the linear program. In *FOCS: IEEE Symposium on Foundations of Computer Science (FOCS)*, 2002.

[2] A. Atserias, M. L. Bonet, and J. Levy. On chvátal rank and cutting planes proofs. Manuscript, 2003.

[3] E. Ben-Sasson and A. Wigderson. Short proofs are narrow – resolution made simple. In *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing*, pages 517–526, Atlanta, GA, May 1999.

[4] A. Bockmayr, F. Eisenbrand, M. Hartmann, and A. Schulz. On the chvatal rank of polytopes in the 0/1 cube. Technical Report 616, Technical University of Berlin, Department of Mathematics, Saarbruecken, Dec. 1998.

[5] M. Bonet and N. Galesi. A study of proof search algorithms for resolution and polynomial calculus. 1999.

[6] V. Chvátal. Edmonds polytopes and a hierarchy of combinatorial problems. *Discrete Mathematics*, 4, 1973.

[7] V. Chvátal, W. Cook, and M. Hartmann. On cutting-plane proofs in combinatorial optimization. *Linear Algebra and its Applications*, 114/115:455–499, 1989.

[8] W. Cook, C. R. Coullard, and G. Turan. On the complexity of cutting plane proofs. *Discrete Applied Mathematics*, 18:25–38, 1987.

[9] S. Dash. *On the matrix cuts of Lovász and Schrijver and their use in Integer Programming*. PhD thesis, Department of Computer Science, Rice University, Mar. 2001.

[10] S. Dash. An exponential lower bound on the length of some classes of branch-and-cut proofs. In *IPCO*, 2002.

[11] F. Eisenbrand and A. S. Schulz. Bounds on the Chvatal rank of polytopes in the $0/1$-cube. *Lecture Notes in Computer Science*, 1610:137–??, 1999.

[12] U. Feige and R. Krauthgamer. The probable value of lovász-schrijver relaxations for maximum independent set. *SIAM Journal on Computing*, 32(2):345–370, 2003.

[13] M. Goemans and L. Tunçel. When does the postive semidefiniteness constraint help in lifting procedures. *Mathematics of Operations Research*, 26:796–815, 2001.

[14] R. Gomory. Outline of an algorithm for integer solutions to linear programs. *Bulletin of the American Mathematical Society*, 64:275–278, 1958.

[15] R. E. Gomory. Solving linear programming problems in integers. In R. Bellman and M. Hall, Jr., editors, *Combinatorial Analysis*, pages 211–215, Providence, RI, 1960. Symposia in Applied Mathematics X, American Mathematical Society.

[16] D. Grigoriev, E. A. Hirsch, and D. V. Pasechnik. Complexity of semi-algebraic proofs. In *Symposium on Theoretical Aspects of Computer Science*, pages 419–430, 2002.

[17] D. Grigoriev, E. A. Hirsch, and D. V. Pasechnik. Exponential lower bound for static semi-algebraic proofs. *Lecture notes in computer science*, 2380:257–268, 2002.

[18] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric algorithms and combinatorial optimization*, volume 2 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, second edition, 1993.

[19] A. Haken. The intractability of resolution. *Theoretical Computer Science*, 39:297–305, 1985.

[20] E. Hirsch and A. Kojevnikov. Several notes on the power of gomory-chvatal cuts. Technical Report TR03-012, ECCC, 2003.

[21] R. Impagliazzo, T. Pitassi, and A. Urquhart. Upper and lower bounds on tree-like cutting planes proofs. In *Proceedings from Logic in Computer Science*, 1994.

[22] L. G. Khachian. A polynomial time algorithm for linear programming. *Doklady Akademii Nauk SSSR, n.s.*, 244(5):1093–1096, 1979. English translation in *Soviet Math. Dokl. 20*, 191–194.

[23] B. Krishnamurthy. Short proofs for tricky formulas. *Acta Informatica*, 22:253–275, 1985.

[24] L. Lovász and A. Schrijver. Cones of matrices and set-functions and 0-1 optimization. *SIAM J. Optimization*, 1(2):166–190, 1991.

[25] P. Pudlák. Lower bounds for resolution and cutting plane proofs and monotone computations. *Journal of Symbolic Logic*, 62(3):981–998, Sept. 1997.

[26] G. Stalmark. Short resolution proofs for a sequence of tricky formulas. *Acta Informatica*, 33:277–280, 1996.