Towards Optimal Integrality Gaps for Hypergraph Vertex Cover in the Lovász-Schrijver Hierarchy

Iannis Tourlakis

Princeton University, Department of Computer Science, Princeton, NJ 08544
itourlak@cs.princeton.edu

Abstract. "Lift-and-project" procedures, which tighten linear relaxations over many rounds, yield many of the celebrated approximation algorithms of the past decade or so, even after only a constant number of rounds (e.g., for MAX-CUT, MAX-3SAT and SPARSEST-CUT). Thus proving super-constant round lowerbounds on such procedures may provide evidence about the inapproximability of a problem.
We prove an integrality gap of $k - \epsilon$ for linear relaxations obtained from the trivial linear relaxation for k-uniform hypergraph VERTEX COVER by applying even $\Omega(\log \log n)$ rounds of Lovász and Schrijver's LS lift-and-project procedure. In contrast, known PCP-based results only rule out $k - 1 - \epsilon$ approximations. Our gaps are tight since the trivial linear relaxation gives a k-approximation.

1 Introduction

Several papers [2, 9, 5, 1] have appeared over the past few years studying the quality of linear programming (LP) and semidefinite programming (SDP) relaxations derived by so-called "lift-and-project" procedures. These procedures (for examples, see Lovász and Schrijver [12] and Sherali and Adams [13]) start with an initial relaxation and add, in a controlled manner, more and more inequalities satisfied by all integral solutions in a bid to obtain tighter and hence, hopefully, better relaxations.

One motivation for studying lift-and-project procedures is the large gap that remains for some problems (such as VERTEX COVER) between approximation ratios achieved by known algorithms and ratios ruled out by PCPs. On the one hand, lift-and-project methods may lead to improved approximation algorithms for such problems. For example, the celebrated SDP-based algorithms of [11] for MAX-CUT and of [4] for SPARSEST CUT are efficiently derived by lift-and-project methods. On the other hand, ruling out good approximation algorithms based on lift-and-project methods might give evidence about the problem's true inapproximability.

In this paper we concentrate on proving integrality gaps for relaxations obtained by the LS lift-and-project technique defined by Lovász and Schrijver [12] (definitions in Section 2). The problem we study is VERTEX COVER on k-uniform hypergraphs. Progress on graph-related problems is sometimes made by considering hypergraph analogues, motivating our study of hypergraph VERTEX COVER. As with graph VERTEX COVER, there is a gap between the factors achieved by known approximation algorithms and those ruled out by PCPs: While PCP-based hardness results for k-uniform hypergraph VERTEX COVER rule out $k-1-\epsilon$ polynomial-time approximations [7], only $k-o(1)$ approximation algorithms are known. Our main result is that for all $\epsilon > 0$ and all sufficiently large n there exist k-uniform hypergraphs on n vertices for which $\Omega(\log\log n)$ rounds of LS are necessary to obtain a relaxation for VERTEX COVER with integrality gap less than $k-\epsilon$.

1.1 Related Work

Alekhnovich et al. [1] also considered integrality gaps of LS tightenings for k-uniform hypergraph VERTEX COVER proving that a gap of $k-1-\epsilon$ remains after a linear number of even LS_+ tightenings (LS_+, defined in Section 2, is a stronger version of LS where positive semidefinite constraints are also added). However, the work most closely related to ours is the paper of Arora et al. [2]. They show that an integrality gap of $2-\epsilon$ remains for graph VERTEX COVER even after $\Omega(\sqrt{\log n})$ rounds of LS. However, obtaining an analogous integrality gap of $k-\epsilon$ for k-uniform hypergraph VERTEX COVER for a non-constant number of rounds remained open since their analysis was especially tailored for graphs. The main technical contribution of this paper is adapting the proof framework of [2] to reason about hypergraphs. We discuss this further after some necessary technical definitions in Section 2 below.

Other recent papers studying Lovász-Schrijver tightenings are [14, 6, 10, 9, 5, 1]. Stephen and Tunçel [14] show that $\Omega(\sqrt{n})$ rounds of even the stronger LS_+ procedure are required to derive some simple inequalities for the MATCHING polytope. Cook and Dash [6] and Goemans and Tunçel [10] both show how for some relaxations n rounds of LS_+ are required to derive some simple inequalities. Feige and Krauthgamer [9] show that a large gap remains for INDEPENDENT SET after $\Omega(\log n)$ rounds of LS_+, while Buresh-Oppenheim et al. [5] show that the gap for MAX-kSAT, $k \geq 5$ remains $(2^k-1)/2^k-\epsilon$ even after $\Omega(n)$ rounds of LS_+. In addition to their results for hypergraph VERTEX COVER mentioned above, Alekhnovich et al. [1] extend the results of [5] to MAX-3SAT as well as proving an integrality gap of $(1-\epsilon)\ln n$ after $\Omega(n)$ rounds of LS_+ for SET COVER.

2 Lovász-Schrijver Liftings and Our Methodology

We will use the VERTEX COVER problem to explain relaxations and how to tighten them using LS liftings. The integer program (IP) characterization for a k-uniform hypergraph $G=(V,E)$ is: minimize $\sum_{i\in V} v_i$ such that $v_{j_1}+\ldots+v_{j_k} \geq 1$ for all hyperedges $\{j_1,\ldots,j_k\}$, where $v_i \in \{0,1\}$. The *integer hull* I is the convex hull of all solutions to this program. The standard LP relaxation is to allow $0 \leq v_i \leq 1$. Clearly the value of the LP is at most that of the IP. Let P be the convex hull of all solutions to the LP in $[0,1]^n$. A linear relaxation is *tightened* by adding constraints that also hold for the integral hull; in general this gives some polytope P' such that $I \subseteq P' \subseteq P$. The quality of a linear relaxation is measured by the ratio $\frac{\text{optimum value over } I}{\text{optimum value over } P}$, called its *integrality gap*.

Lovász and Schrijver [12] present several "lift-and-project" techniques for deriving tighter and tighter relaxations of a 0-1 integer program. These procedures take an n dimensional relaxed polytope, "lift" it to n^2 dimensions, add new constraints to the high-dimensional polytope, and then project back to the original space. Three progressively stronger such techniques, LS_0, LS and LS_+ are given in [12]; we will focus on LS, but also define LS_+ below for completeness.

The notation uses homogenized inequalities. Let G be a k-uniform hypergraph and let $\mathrm{VC}(G)$ be the cone in $\Re^{n+1}$ that contains $(x_0, x_1, \ldots, x_n)$ iff it satisfies the edge constraints $x_{i_1} + \ldots + x_{i_k} \geq x_0$ for each edge $\{i_1, \ldots, i_k\} \in G$. All cones in what follows will be in $\Re^{n+1}$ and we will be interested in the slice cut by the hyperplane $x_0 = 1$. Denote by $N^r(\mathrm{VC}(G))$ and $N_+^r(\mathrm{VC}(G))$ the feasible cones of all inequalities obtained from r rounds of the LS and LS_+ lifting procedures, respectively. The following lemma characterizes the effect of one round of the N and N_+ operators. (In the remainder of the paper we index columns and rows starting from 0 rather than 1, and e_i denotes the ith unit vector.)

Lemma 1 ([12]). *Let Q be a cone in $\Re^{n+1}$. Then $y \in \Re^{n+1}$ is in $N(Q)$ iff there exists a symmetric matrix $Y \in \Re^{(n+1)\times(n+1)}$ such that:*

1. *$Ye_0 = diag(Y) = y$,*
2. *$Ye_i, Y(e_0 - e_i) \in Q$, for all $i = 1, \ldots, n$.*

Moreover, $y \in N_+(Q)$ iff Y is in addition positive semidefinite.

To prove that $y \in N^{r+1}(Q)$, we have to construct a specific matrix Y and prove that the $2n$ vectors defined in Lemma 1 are in $N^r(Q)$. Such a matrix Y is called a *protection matrix* for y since it "protects" y for one round. The points y we protect will always have $y_0 = 1$ in which case condition 2 above is equivalent to the following condition (below P is the projection of Q along $x_0 = 1$):

2′. *For each i such that $x_i = 0$, $Ye_i = 0$; for each i such that $x_i = 1$, $Ye_i = y$; Otherwise Ye_i/x_i and $Y(e_0 - e_i)/(1 - x_i)$ are both in P.*

The simplest possible Y has $Y_{ij} = y_i y_j$ except along the diagonal where $Y_{ii} = y_i$. Then Ye_i/x_i and $Y(e_0 - e_i)/(1 - x_i)$ are identical to y except in the ith coordinate which are now 1 and 0, respectively. This matrix was used in [10, 5].

However, such a protection matrix does not work for us: To prove gaps of $k - \epsilon$ on k-uniform hypergraphs we will protect the point $y_\gamma = (1, \frac{1}{k} + \gamma, \ldots, \frac{1}{k} + \gamma)$ where $\gamma > 0$ is an arbitrarily small constant. For y_γ, the vector $Y(e_0 - e_i)/(1 - x_i)$ given by the simple protection matrix is not guaranteed to be in $\mathrm{VC}(G)$: if coordinate i is changed to 0 in y_γ, we violate all edge constraints involving the ith vertex. So more sophisticated protection matrices are needed.

For the case when $k = 2$ (i.e., for graphs), Arora et al. [2] used LP duality to show that appropriate protection matrices exist for at least $\Omega(\sqrt{\log n})$ rounds. In particular, they used a special combinatorial form of Farkas's lemma specifically applicable to the constraints involved in defining protection matrices for graphs.

We will also reduce the problem of showing our protection matrices exists to the feasibility of linear programs; however, we will need a much more complicated form of Farkas's lemma than in the graph case. As such, we will only be able to carry out our arguments for $O(\log\log n)$ rounds.

3 Lowerbounds for Hypergraph Vertex Cover

Theorem 1. *For all $k \geq 2$, $\epsilon > 0$ there exist constants $n_0(k,\epsilon), \delta(k,\epsilon) > 0$ s.t. for every $n \geq n_0(k,\epsilon)$ there exists a k-uniform hypergraph G on n vertices for which the integrality gap of $N^r(\mathrm{VC}(G))$ is at least $k-\epsilon$ for all $r \leq \delta(k,\epsilon)\log\log n$.*

Theorem 1 will follow from the following two theorems.

Theorem 2. *For all $k \geq 2$ and any $\epsilon > 0$ there exists an $n_0(k,\epsilon)$ such that for every $n \geq n_0(k,\epsilon)$ there exist k-uniform hypergraphs with n vertices and $O(n)$ hyperedges having $\Omega(\log n)$ girth but no independent set of size greater than ϵn.*

Theorem 3. *Fix $\gamma > 0$ and let $y_\gamma = (1, \frac{1}{k}+\gamma, \frac{1}{k}+\gamma, \ldots, \frac{1}{k}+\gamma)$. Let G be a k-uniform hypergraph such that $\mathrm{girth}(G) \geq \frac{20}{\gamma} r5^r$. Then $y_\gamma \in N^r(\mathrm{VC}(G))$.*

Theorem 2 is an easy extension to hypergraphs of a result proved by Erdős [8]. The remainder of this section will be devoted to proving Theorem 3.

3.1 Proof of Theorem 3

As in previous LS lowerbounds, the theorem will be proved by induction where the inductive hypothesis will require some set of vectors to be in $N^m(\mathrm{VC}(G))$ for $m \leq r$. These vectors will be essentially all-$(\frac{1}{k}+\gamma)$ except possibly for a few small neighbourhoods where they can take arbitrary nonnegative values so long as the edge constraints for G are satisfied. The exact characterization is given by the following definition similar to an analogous definition in [2]:

Definition 1. *Let $S \subseteq [n]$, R be a positive integer, and $\gamma > 0$. A nonnegative vector $(\alpha_0, \alpha_1, \ldots, \alpha_n)$ with $\alpha_0 = 1$ is an (S,R,γ)-vector if the entries satisfy the edge constraints and if $\alpha_j = \frac{1}{k}+\gamma$ for each $j \notin \cup_{w\in S} Ball(w,R)$. Here $Ball(w,R)$ denotes the set of vertices within distance R of w in the graph.*

Let $R_r = 0$ and let $R_m = 5R_{m+1} + 1/\gamma$ for $0 \leq m < r$. Note $R_m \leq \mathrm{girth}(G)/20$.

To prove Theorem 3 we prove the following inductive claim. The theorem is a subcase of $m = r$:

Inductive Claim for $N^m(\mathrm{VC}(G))$: *For every set S of $r-m$ vertices, every (S, R_m, γ)-vector is in $N^m(\mathrm{VC}(G))$.*

The base case $m = 0$ is trivial since (S, R_m, γ)-vectors satisfy the edge constraints for G. In the remainder of this section we prove the Inductive Claim for $m+1$ assuming truth for m.

To that end, let α be an (S, R_{m+1}, γ)-vector where $|S| = r - m - 1$. By induction, every $(S \cup \{i\}, R_m, \gamma)$-vector is in $N^m(\mathrm{VC}(G))$. So to prove that $\alpha \in N^{m+1}(\mathrm{VC}(G))$ it suffices by Lemma 1 to exhibit an $(n+1) \times (n+1)$ symmetric matrix Y such that:

A. $Ye_0 = diag(Y) = \alpha$,
B. For each i such that $\alpha_i = 0$, $Ye_i = 0$; for each i such that $\alpha_1 = 1$, $Ye_0 = Ye_i$; otherwise, Ye_i/α_i and $Y(e_0 - e_i)/(1 - \alpha_i)$ are $(S \cup \{i\}, R_m, \gamma)$-vectors.

As was done in [2] and [12], we will write these conditions as a linear program and show that the program is feasible, proving the existence of Y. Our notation will assume symmetry, namely Y_{ij} represents $Y_{\{i,j\}}$.

Condition **A** requires:

$$Y_{kk} = \alpha_k \qquad \forall k \in \{1, \ldots, n\}\,. \tag{1}$$

Condition **B** requires first of all that Ye_j/α_j and $Y(e_0 - e_j)/(1 - \alpha_j)$ satisfy the edge constaints: For all $i \in \{1, \ldots, n\}$ and all $\{j_1, \ldots, j_k\} \in E$:

$$\alpha_i \le Y_{ij_1} + \ldots + Y_{ij_\ell} \le \alpha_i + (\alpha_{j_1} + \ldots + \alpha_{j_k} - 1). \tag{2}$$

Vertices i, t are called a *distant pair* if $t \notin \cup_{w \in S \cup \{i\}} Ball(w, 5R_{m+1} + 1/\gamma)$. (Note then that $\alpha_t = \frac{1}{k} + \gamma$.) Condition **B** requires that for such a pair, the tth coordinates of Ye_i/α_i and $Y(e_0 - e_i)/(1 - \alpha_i)$ are $\frac{1}{k} + \gamma$. In particular,

$$Y_{it} = \alpha_i \alpha_t = \alpha_i \left(\frac{1}{k} + \gamma \right). \tag{3}$$

Note that distant pairs have the property that every path in G that connects them contains at least $4R_m + 1/\gamma - 1$ hyperedges each of which α oversatisfies by $k\gamma$, and at most R_m hyperedges which α does *not* oversatisfy by $k\gamma$.

Finally, condition **B** requires that Ye_i/α_i, $Y(e_0 - e_i)/(1 - \alpha_i)$ are in $[0,1]^{n+1}$:

$$0 \le Y_{ij} \le \alpha_i, \qquad \forall i, j \in \{1, \ldots, n\}, i \ne j \tag{4}$$

$$-Y_{ij} \le 1 - \alpha_i - \alpha_j, \qquad \forall i, j \in \{1, \ldots, n\}, i \ne j \tag{5}$$

The above constraints are equivalent to the following four constraint families:

$$Y_{ij} \le \beta(i,j), \qquad \forall i, j \in \{1, \ldots, n\} \tag{6}$$

$$-Y_{ij} \le \delta(i,j), \qquad \forall i, j \in \{1, \ldots, n\} \tag{7}$$

$$Y_{ij_1} + \ldots + Y_{ij_k} \le a(i, j_1, \ldots, j_k), \qquad \forall \{j_1, \ldots, j_k\} \in E \tag{8}$$

$$-Y_{ij_1} - \ldots - Y_{ij_k} \le b(i, j_1, \ldots, j_k), \qquad \forall \{j_1, \ldots, j_k\} \in E \tag{9}$$

Here (1) $\beta(i,j) = \alpha_i \alpha_j$ if i, j is a distant pair and $\beta(i,j) = \min(\alpha_i, \alpha_j)$ otherwise; (2) $\delta(i,j) = -\alpha_i$ if $i = j$; $\delta(i,j) = -\alpha_i \alpha_j$ if i, j is a distant pair; and $\delta(i,j) = 1 - \alpha_i - \alpha_j$ otherwise; (3) $a(i, j_1, \ldots, j_k) = \alpha_i + (\alpha_{j_1} + \ldots + \alpha_{j_k} - 1)$; and (4) $b(i, j_1, \ldots, j_k) = -\alpha_i$. Note that $\beta(i,j) + \delta(i,j) \ge 0$ always, since $\alpha \in [0,1]^{n+1}$,

To prove the consistency of constraints (6)–(9), we will use a special combinatorial version of Farkas's Lemma similar in spirit to that used in [12] and [2]. Before giving the exact form, we require some definitions.

Definition 2. *A* tiling (W, P, N) for G *is a connected k-uniform hypergraph $H = (W, P, N)$ on vertices W and two disjoint k-edge sets P and N such that:*

1. *Each vertex in W is labelled by Y_{ij}, $i, j \in \{1, \ldots, n\}$. Note that distinct vertices need not have different labels.*
2. *Each vertex belongs to at most one edge in P and at most one edge in N (in particular, all edges in P are mutually disjoint, as are all edges in N).*
3. *All edges in $P \cup N$ are of the form $\{Y_{ij_1}, \ldots, Y_{ij_k}\}$ where $\{j_1, \ldots, j_k\} \in E$.*

The edges of a tiling are called tiles. *Vertices in W not incident to any tile in P are called* unmatched negative *vertices; vertices in W not incident to any tile in N are called* unmatched positive *vertices. Let U_N and U_P denote the sets of unmatched negative and unmatched positive vertices, respectively. A vertex labelled Y_{ii} is called* diagonal. *Denote the set of unmatched positive diagonal vertices in W by U_{PD}. A vertex labelled Y_{ij} is called a* distant pair *if $\{i, j\}$ are a distant pair. Given a tile $\{Y_{ij_1}, \ldots, Y_{ij_k}\}$ in P or N, call $\{j_1, \ldots, j_k\} \in E$ its* bracing edge *and i its* bracing node. *An edge $\{j_1, \ldots, j_k\}$ in G is called* overloaded *if $\alpha_{j_1} = \ldots = \alpha_{j_k} = \frac{1}{k} + \gamma$. A tile $\{Y_{ij_1}, \ldots, Y_{ij_k}\}$ in H is* overloaded *if its bracing edge is overloaded.*

Given a tiling $H = (W, P, N)$ for G, define the following sums:

$$S_1^H = \sum_{\{Y_{ij_1}, \ldots, Y_{ij_k}\} \in P} a(i, j_1, \ldots, j_k) + \sum_{\{Y_{ij_1}, \ldots, Y_{ij_k}\} \in N} b(i, j_1, \ldots, j_k), \tag{10}$$

$$S_2^H = \sum_{Y_{ij} \in U_P} \delta(i, j) + \sum_{Y_{ij} \in U_N} \beta(i, j). \tag{11}$$

Finally, let $S^H = S_1^H + S_2^H$.

Lemma 2 (Special case of Farkas's Lemma). *Constraints (6)–(9) are unsatisfiable iff there exists a tiling $H = (W, P, N)$ for G such that $S^H < 0$.*

Proof. Note first that by the general form of Farkas's lemma constraints (6)–(9) are unsatisfiable iff there exists a positive rational linear combination of them where the LHS is 0 and the RHS is negative.

Now suppose that there exists a tiling $H = (W, P, N)$ such that $S^H < 0$. Consider the following linear integer combination of the constraints: (1) For each tile $e = \{Y_{ij_1}, \ldots, Y_{ij_k}\} \in H$, if $e \in P$, add the constraint $Y_{ij_1} + \ldots + Y_{ij_k} \leq a(i, j_1, \ldots, j_k)$; if $e \in N$, add the constraint $-Y_{ij_1} - \ldots - Y_{ij_k} \leq b(i, j_1, \ldots, j_k)$; (2) For each $v \in U_N$ labelled Y_{ij}, add the constraint $Y_{ij} \leq \beta(i, j)$; (3) For each $v \in U_P$ labelled Y_{ij}, add the constraint $-Y_{ij} \leq \delta(i, j)$. But then, for this combination of constraints the LHS equals 0 while the RHS equals $S^H < 0$. So by Farkas's lemma the constraints are unsatisfiable.

Now assume on the other hand that the constraints are unsatisfiable. So there exists a positive rational linear combination of the constraints such that the LHS is 0 and the RHS is negative. By clearing out denominators, we can assume that the linear combination has *integer* coefficients. Hence, as $\beta(i, j) + \delta(i, j) \geq 0$

always, our combination must contain constraints of type (8) and (9). Moreover, since the LHS is 0, for each Y_{ij} appearing in the integer combination there must be a corresponding occurrence of $-Y_{ij}$. But then, it is easy to see that the constraints in the integer linear combination can be grouped into a set of tilings $\{H_i = (W_i, P_i, N_i)\}$ such that the RHS of the linear combination equals $\sum_i S^{H_i}$. Since the RHS is negative, it must be that at least one of the tilings H in the set is such that $S^H < 0$. The lemma follows.

So to show that the constraints for the matrix Y are consistent and thus complete the proof of the Inductive Claim for $m+1$ (and complete the proof of Theorem 3), we will show that $S^H \geq 0$ for any tiling $H = (W, P, N)$ for G. To that end, fix a tiling $H = (W, P, N)$ for G. Our analysis divides into three cases depending on the size of U_{PD}, the set of unmatched positive diagonal vertices in H. In the first (and easiest) case, $|U_{PD}| = 0$; in the second, $|U_{PD}| \geq 2$; and in the final, $|U_{PD}| = 1$. We will show that $S^H \geq 0$ in all these cases . To reduce clutter, we drop the superscript H from S_1^H, S_2^H and S^H. In what follows, let C be the subgraph of G induced by the bracing edges of all tiles in H.

We first note two easy facts about H used below:

Proposition 1. *1. Suppose H contains a diagonal vertex. Then C is connected. Moreover, for any vertex labelled Y_{ij} in H, there exists a path p in H such that the bracing edges corresponding to the tiles in p form a path p' from i to j in C.*
2. The distance between any two diagonal vertices in H is at least $\mathrm{girth}(G)/2$.

Proof. Part (2) is sketched in the Appendix. We leave part (1) as an easy exercise.

Case 1: No unmatched positive diagonal vertices

Consider the following sum:

$$S_2' = \sum_{Y_{ij} \in U_N} \alpha_i \alpha_j - \sum_{Y_{ij} \in U_P} \alpha_i \alpha_j. \tag{12}$$

Note that since $\alpha \in [0,1]^{n+1}$, it follows that $-\alpha_i\alpha_j \leq 1 - \alpha_i - \alpha_j$ and $\alpha_i \alpha_j \leq \alpha_i$. So since there are no unmatched positive diagonal vertices in the tiling, $\alpha_i\alpha_j \leq \beta(i,j)$ and $-\alpha_i\alpha_j \leq \delta(i,j)$ for all unmatched vertices labelled Y_{ij} in the tiling, and hence, $S_2 \geq S_2'$. So to show $S \geq 0$ in this case, it suffices to show $S_1 + S_2' \geq 0$.

To that end, consider the following sum:

$$\sum_{\{Y_{ij_1},\ldots,Y_{ij_k}\} \in P} (-\alpha_i\alpha_{j_1} - \ldots - \alpha_i\alpha_{j_k}) + \sum_{\{Y_{ij_1},\ldots,Y_{ij_k}\} \in N} (\alpha_i\alpha_{j_1} + \ldots + \alpha_i\alpha_{j_k}). \tag{13}$$

By properties (2) and (3) of a tiling, it follows that (13) telescopes and equals S_2'. But then, to show that $S \geq 0$ in this case it suffices to show that for each $\{Y_{ij_1}, \ldots, Y_{ij_k}\} \in P$,

$$a(i, j_1, \ldots, j_k) - [\alpha_i(\alpha_{j_1} + \ldots + \alpha_{j_k})] = (1 - \alpha_i)(\alpha_{j_1} + \ldots + \alpha_{j_k} - 1), \tag{14}$$

is nonnegative (the first term comes from the term in S_1 for the tile, and the second from the term for the tile in (13)), and for each $\{Y_{ij_1}, \ldots, Y_{ij_k}\} \in N$,

$$b(i, j_1, \ldots, j_k) + [\alpha_i(\alpha_{j_1} + \ldots + \alpha_{j_k})] = \alpha_i(\alpha_{j_1} + \ldots + \alpha_{j_k} - 1), \tag{15}$$

is nonnegative (again, the first term comes from S_1 and the second from (13)). But (14) and (15) are both nonnegative since the bracing edges for all tiles are in G and α satisfies the edge constraints for G. Hence, $S \geq 0$ in this case.

Case 2: At least 2 unmatched positive diagonal vertices

We first define some notation to refer to quantities (14) and (15) which will be used again in this case: For a tile $e = \{Y_{ij_1}, \ldots, Y_{ij_k}\}$, define $\zeta(e)$ to be $(1 - \alpha_i)(\alpha_{j_1} + \ldots + \alpha_{j_k} - 1)$ if $e \in P$ and be $\alpha_i(\alpha_{j_1} + \ldots + \alpha_{j_k} - 1)$ if $e \in N$.

In Case 1 we showed that $S \geq 0$ by (1) defining a sum S_2' such that $S \geq S_1 + S_2'$, then (2) noting that $S_1 + S_2' = \sum_{e \in H} \zeta(e)$, and finally (3) showing that $\zeta(e) \geq 0$ for all tiles e in H. Unfortunately, in the current case, since H contains unmatched positive diagonal vertices, it is no longer true that $S \geq S_1 + S_2'$. However, it is easy to see that $S \geq (S_1 + S_2') - \sum_{Y_{ii} \in U_{PD}} (\alpha_i - \alpha_i^2)$. In particular, $S \geq \sum_{e \in H} \zeta(e) - \sum_{Y_{ii} \in U_{PD}} (\alpha_i - \alpha_i^2)$. So since $\zeta(e) \geq 0$ for all tiles, to show that $S \geq 0$ for the current case, it suffices to show that for "many" tiles e in H, $\zeta(e)$ is sufficiently large so that $\sum_{e \in H} \zeta(e) \geq \sum_{Y_{ii} \in U_{PD}} (\alpha_i - \alpha_i^2)$.

We require the following lemma:

Lemma 3. *If there are $\ell \geq 2$ diagonal vertices in H, then there exist ℓ disjoint paths of length* $\mathrm{girth}(G)/4 - 2$ *in H which do not involve any diagonal vertices.*

Proof. Let Q be a tree in H that spans all the diagonal vertices in H (such a tree exists since H is connected). For each diagonal vertex $u \in H$ let $B(u)$ be the tiles in Q with distance at most $\mathrm{girth}(G)/4 - 1$ from u. Then for all diagonal vertices $u, v \in H$, the balls $B(u)$ and $B(v)$ must be disjoint: otherwise, the distance between u and v would be less than $\mathrm{girth}(G)/2$, contradicting part (2) of Proposition 1. Since Q is a spanning tree, for each diagonal vertex $u \in H$, there exists a path of length at least $\mathrm{girth}(G)/4 - 2$ in $B(u)$.

By the lemma, for each vertex $v \in U_{PD}$ there exists a path p_v in H of length $\mathrm{girth}(G)/4 - 2$ such that for distinct $u, v \in U_{PD}$ the paths p_u and p_v share no tiles. We will show (provided that n is sufficiently large) that for each $v \in U_{PD}$, $\sum_{e \in p_v} \zeta(e) \geq 1$. Since $\alpha_i - \alpha_i^2 \leq \frac{1}{4}$, it will follow that $S \geq 0$, completing the proof for this case.

So fix $v \in U_{PD}$. Since $R^m \leq \mathrm{girth}(G)/20$, p_v contains at least $\mathrm{girth}(G)/20$ disjoint pairs of adjacent overloaded tiles (i.e., $\alpha_j = \frac{1}{k} + \gamma$ for all vertices j in the bracing edges of the two tiles in the pair). Let $e_1 = \{Y_{qr_1}, \ldots, Y_{qr_k}\} \in P$ and $e_2 = \{Y_{st_1}, \ldots, Y_{st_k}\} \in N$ be such a pair. Note that $\zeta(e_1) = (1 - \alpha_i)k\gamma$ and $\zeta(e_2) = \alpha_s k\gamma$. Now, either $s = r_\ell$ for some $\ell \in [k]$, or $s = q$. If $s = r_\ell$, then $\alpha_s = \frac{1}{k} + \gamma$ and $\zeta(e_1) + \zeta(e_2) \geq k\gamma(\frac{1}{k} + \gamma)$. If instead $s = q$, then $\zeta(e_1) + \zeta(e_2) \geq k\gamma$. In either case, $\zeta(e_1) + \zeta(e_2) \geq k\gamma(\frac{1}{k} + \gamma)$. Hence, summing over all girth$(G)/20$ disjoint pairs of adjacent overloaded tiles, it follows that $\sum_{e\in p_v}\zeta(e) \geq k\gamma(\frac{1}{k}+\gamma)\,\mathrm{girth}(G)/20$. As desired, the latter is indeed greater than 1 for large n since girth$(G) = \Omega(\log n)$.

Case 3: Exactly 1 unmatched positive diagonal vertex

The argument in Case 2 actually rules out the possibility that the tiling contains two or more diagonal vertices, unmatched or not. Hence, we will assume that our tiling contains exactly one diagonal vertex labelled without loss of generality by Y_{11}. We consider two subcases.

Subcase 1: At least one of the unmatched vertices is a distant pair:
Arguing as in Case 2, we have in this subcase that $S = [\sum_{e\in H}\zeta(e)] - (\alpha_1 - \alpha_1^2)$ where, moreover, $\zeta(e) \geq 0$ for all tiles $e \in H$. Hence, to show that $S \geq 0$ in this subcase, it suffices to show that there exists a subset $H' \subseteq H$ such that $\sum_{e\in H'}\zeta(e) \geq \frac{1}{4} \geq \alpha_1 - \alpha_1^2$.

To that end, let v be an unmatched vertex in H labelled Y_{ij} where $\{i,j\}$ is a distant pair. By part (1) of Proposition 1 there is a path p in H such that the bracing edges corresponding to the tiles in p form a path q in C connecting vertices i and j. Since i, j are distant, q must contain, by definition, at least $4R_{m+1} + 1/\gamma$ overloaded edges. Moreover, it is not hard to see that there exists a sub-path p' of p of length at most $5R_{m+1} + 1/\gamma$ such that the bracing edges of the tiles in p' include all these overloaded edges. Hence p' must contain at least $1/\gamma$ disjoint pairs of overloaded tiles. But then, arguing as in Case 2, it follows that $\sum_{e\in p}\zeta(e) \geq (1/\gamma)(k\gamma(\frac{1}{k}+\gamma)) \geq \frac{1}{4}$, and hence that $S \geq 0$ in this subcase.

Subcase 2: None of the unmatched vertices is a distant pair:
Note first that we can assume that H contains no cycles: If it does, then it is easy to see that C must also contain a cycle, and hence, we can use the ideas from Case 2 to show that $S \geq 0$.

So assume H has no cycles and define a tree T as follows: There is a node in T for each tile in H. The root $root(T)$ corresponds to the tile containing Y_{11}. There is an edge between two nodes in T iff the tiles corresponding to the nodes share a vertex. Note that T is a tree since H is acyclic. For a node $v \in T$, let $Tile(v)$ denote its corresponding tile in H. Finally, for a node $v \in T$ we abuse notation and say $v \in P$ (resp., $v \in N$) if $Tile(v) \in P$ (resp., $Tile(v) \in N$).

Recursively define a function t on the nodes of T as follows: Let v be a node in T and suppose $Tile(v) = \{Y_{ij_1}, \ldots, Y_{ij_k}\}$. If $v \in P$, then

$$t(v) = a(i, j_1, \ldots, j_k) + \sum_{v'\in Child(v)} t(v') + \sum_{\text{unmatched } Y_{ij_\ell} \text{ in } Tile(v)} \delta(i, j_\ell). \quad (16)$$

If instead $v \in N$, then

$$t(v) = b(i, j_1, \ldots, j_k) + \sum_{v'\in Child(v)} t(v') + \sum_{\text{unmatched } Y_{ij_\ell} \text{ in } Tile(v)} \beta(i, j_\ell). \quad (17)$$

A simple induction now shows that $t(root(T)) = S$ (recall that $root(T)$ corresponds to the tile containing Y_{11}). Hence, to show that $S \geq 0$ in this subcase it suffices to show that $t(root(T)) \geq 0$. This will follow from the following lemma:

Lemma 4. *Given $v \in T$, let $e = Tile(v) = \{Y_{ij_1}, \ldots, Y_{ij_k}\}$. Assume, without loss of generality, that $Y_{ij_1}, \ldots, Y_{ij_\ell}$ (where $0 \leq \ell \leq k$) are the labels of the vertices in e shared with the tile corresponding to v's parent (where $\ell = 0$ iff $v = root(T)$). Then,*

1. *If $v = root(T)$, then $t(v) = t(root(T)) \geq 0$; otherwise,*
2. *If $v \in P$, then $t(v) \geq \min(\alpha_i, \sum_{r=1}^{\ell} \alpha_{j_r})$;*
3. *If instead $v \in N$, then $t(v) \geq \min(0, 1 - \alpha_i - \sum_{r=1}^{\ell} \alpha_{j_r})$.*

Proof. The proof is by induction on the size of the subtree at v. For the base case, let v be a leaf. Hence, the vertices labelled by $Y_{ij_{\ell+1}}, \ldots, Y_{ij_k}$ in e are all unmatched non-distant pair vertices (since H contains no unmatched distant pair vertices). Suppose $v \in P$. By equation (16),

$$t(v) = \alpha_i + \sum_{r=1}^{k} \alpha_r - 1 + \sum_{r=\ell+1}^{k} \delta(i, j_r). \tag{18}$$

There are two subcases to consider. In the first subcase, $j_r \neq i$ for all $r = \ell + 1, \ldots, k$. Then $\delta(i, j_r) = 1 - \alpha_i - \alpha_r$ for all $r = \ell + 1, \ldots, k$ and it follows that $t(v) = \sum_{r=1}^{\ell} \alpha_r + (k - \ell - 1)(1 - \alpha_i) \geq \sum_{r=1}^{\ell} \alpha_r$. In the second subcase, $j_r = i$ for some r, say $r = k$. This can only happen for one r since H only contains one diagonal vertex, namely Y_{11}. In particular, v must be $root(T)$. Then $\delta(i, j_k) = -\alpha_i$ and $\delta(i, j_r) = 1 - \alpha_i - \alpha_r$ for all $r = \ell + 1, \ldots, k - 1$. Since α satisfies the edge constraints for G, it follows that $t(v) = t(root(T)) \geq 0$.

Suppose instead that $v \in N$. By equation (17), $t(v) = -\alpha_i + \sum_{r=\ell+1}^{k} \beta(i, j_r)$. Now, $\beta(i, j_r) \geq \min(\alpha_i, \alpha_{j_r})$ for all $r = \ell + 1, \ldots, k$. If $\beta(i, j_r) = \alpha_i$ for some j_r, then $t(v) \geq 0$ as desired. So assume $\beta(i, j_r) = \alpha_{j_r}$ for all $r = \ell + 1, \ldots, k$. Since α satisfies the edge constraints for G, $\sum_{r=1}^{k} \alpha_{j_r} \geq 1$. But then, $t(v) \geq 1 - \alpha_i - \sum_{r=1}^{\ell} \alpha_{j_r}$, completing the proof for the base case.

For the inductive step, let v be any node in T and assume the lemma holds for all children of v. Assume first that $v \in P$ and hence, all children of v are in N (by definition of a tiling). By the induction hypothesis, there are two possibilities: either (a) $t(v') \geq 0$ for all children v' of v, and there are no unmatched vertices in $Tile(v)$, or (b) there exists some unmatched vertex in $Tile(v)$ or $t(v') \geq 1 - \alpha_i - \sum_{s=1}^{t} \alpha_{j_{r_s}}$ for some child v' of v. In case (a), it follows from (16) and from from the fact that α satisfies the edge constraints for G that $t(v) \geq \alpha_i$. In case (b), using the arguments from the base case, it follows that $t(v) \geq 0$ when $v = root(T)$, and $t(v) \geq \sum_{r=1}^{\ell} \alpha_{j_r}$ when $v \neq root(T)$.

Now assume that $v \in N$. The arguments from the above case when $v \in P$, as well as the arguments from the base case, can now be adapted to show that $t(v) \geq \min(0, 1 - \alpha_i - \sum_{r=1}^{\ell} \alpha_{j_r})$. Note that in in this case $v \neq root(T)$ since $root(T) \in P$. The inductive step, and hence the lemma, now follow.

So $S \geq 0$ in Case 3 also, and the Inductive Claim now follows for $m + 1$.

4 Discussion

The integrality gap of $2-\epsilon$ obtained in [2] for graph VERTEX COVER holds for $\Omega(\sqrt{\log n})$ rounds of LS (improved to $\Omega(\log n)$ rounds in [3]). We conjecture that our integrality gaps in the hypergraph case should also hold for at least $\Omega(\log n)$ rounds. Indeed, examining the proof of Theorem 3, it can be shown that by redefining the recursive definition of R_m to be $R_m = R_{m+1} + 1/\gamma$, then all cases considered in the proof except Subcase 1 of Case 3 can be argued for $\Omega(\log n)$ rounds. While it can be argued that $S \geq 0$ for $\Omega(\sqrt{\log n})$ rounds for graphs in this subcase (and in fact for $\Omega(\log n)$ rounds if the definition of (S, R, γ)-vectors is further refined as in [3]), a proof for hypergraphs eludes us.

The integrality gaps of $k-1-\epsilon$ given in [1] for k-uniform hypergraph VERTEX COVER held not only for LS but also for LS_+ liftings, the stronger semidefinite version of Lovász-Schrijver liftings. As mentioned in the introduction, obtaining optimal integrality gaps for both graph and hypergraph VERTEX COVER for LS_+ remains a difficult open question.

Finally, can our techniques be applied to other problems? For instance, could our techniques be used to show that, say, even after $\log \log n$ rounds of LS the integrality gap of the linear relaxation for MAX-CUT remains larger than the approximation factor attained by the celebrated SDP-based Goemans-Williamson algorithm [11]?

5 Acknowledgements

I would like to thank Sanjeev Arora and Elad Hazan for many helpful discussions. I would also like to thank the anonymous referees for helpful comments.

References

1. M. Alekhnovich, S. Arora, and I. Tourlakis. Towards strong nonapproximability results in the Lovász-Schrijver hierarchy. In *Proceedings of the 37th Annual Symposium on the Theory of Computing*, New York, May 2005. ACM Press.
2. S. Arora, B. Bollobás, and L. Lovász. Proving integrality gaps without knowing the linear program. In *Proceedings of the 43rd Symposium on Foundations of Computer Science (FOCS-02)*, pages 313–322, Los Alamitos, Nov. 16–19 2002.
3. S. Arora, B. Bollobás, L. Lovász, and I. Tourlakis. Proving integrality gaps without knowing the linear program. Manuscript, 2005.
4. S. Arora, S. Rao, and U. Vazirani. Expander flows, geometric embeddings and graph partitioning. In *Proceedings of the thirty-sixth annual ACM Symposium on Theory of Computing (STOC-04)*, pages 222–231, New York, June 13–15 2004. ACM Press.
5. J. Buresh-Oppenheim, N. Galesi, S. Hoory, A. Magen, and T. Pitassi. Rank bounds and integrality gaps for cutting planes procedures. In *FOCS: IEEE Symposium on Foundations of Computer Science (FOCS)*, 2003.
6. W. Cook and S. Dash. On the matrix-cut rank of polyhedra. *Mathematics of Operations Research*, 26(1):19–30, 2001.

7. I. Dinur, V. Guruswami, S. Khot, and O. Regev. A new multilayered PCP and the hardness of hypergraph vertex cover. In ACM, editor, *Proceedings of the Thirty-Fifth ACM Symposium on Theory of Computing, San Diego, CA, USA, June 9–11, 2003*, pages 595–601, New York, NY, USA, 2003. ACM Press.
8. P. Erdős. Graph theory and probability. *Canadian Journal of Mathematics*, 11:34–38, 1959.
9. U. Feige and R. Krauthgamer. The probable value of the Lovász–Schrijver relaxations for maximum independent set. *SIAM Journal on Computing*, 32(2):345–370, Apr. 2003.
10. M. X. Goemans and L. Tunçel. When does the positive semidefiniteness constraint help in lifting procedures. *Mathematics of Operations Research*, 26:796–815, 2001.
11. M. X. Goemans and D. P. Williamson. .878-approximation algorithms for MAX CUT and MAX 2SAT. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on the Theory of Computing*, pages 422–431, Montréal, Québec, Canada, 23–25 May 1994.
12. L. Lovász and A. Schrijver. Cones of matrices and set-functions and 0-1 optimization. *SIAM Journal on Optimization*, 1(2):166–190, May 1991.
13. H. D. Sherali and W. P. Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM Journal on Discrete Mathematics*, 3(3):411–430, Aug. 1990.
14. T. Stephen and L. Tunçel. On a representation of the matching polytope via semidefinite liftings. *Mathematics of Operations Research*, 24(1):1–7, 1999.

A Proof sketch of part (2) of Proposition 1

Since H is connected, there exists a path between any two diagonal vertices in H. We will show that the subgraph of G induced by the bracing edges of the tiles in such path must contain a cycle. Part (2) will then follow.

To that end, let q be an arbitrary path in H comprised, in order, of tiles $e_1, \ldots, e_r$. Consider these tiles in order beginning with e_1. As long as the bracing node in successive tiles does not change, then the bracing edges of these tiles form a path q' in G. If the bracing node changes at some tile, say e_i, then the bracing edge for e_i starts a new path q'' in G. Moreover, some vertex $w \in G$ from the last edge visited in q' becomes the (new) bracing node for e_i. The bracing edges of the tiles following e_i now extend q'' until an edge e_j is encountered with yet a different bracing node. But then, the bracing edge for e_j must contain w. Hence the bracing edge for e_j must extend q'. Continuing this argument, we see that each time the bracing node switches, the tiles switch back and forth from contributing to q' or q''.

Each time the current tile e_ℓ contains some diagonal vertex labelled Y_{ii}, the bracing node i for e_ℓ is also contained in the bracing edge for e_ℓ. But then, since the bracing node belongs to, say, q' while the bracing edge belongs to q'', it follows that q' and q'' must intersect at vertex i in G. Hence, if a tile path q contains *two* diagonal vertices, q' and q'' must intersect twice. Part (2) follows.