Lasserre Integrality Gaps for Graph Spanners and Related Problems

Michael Dinitz[1(✉)], Yasamin Nazari[1], and Zeyu Zhang[2]

[1] Johns Hopkins University, Baltimore, MD, USA
{mdinitz,ynazari}@jhu.edu
[2] Facebook, Inc., Menlo Park, CA, USA

Abstract. There has been significant recent progress on algorithms for approximating graph spanners, i.e., algorithms which approximate the best spanner for a given input graph. Essentially all of these algorithms use the same basic LP relaxation, so a variety of papers have studied the limitations of this approach and proved integrality gaps for this LP. We extend these results by showing that even the strongest lift-and-project methods cannot help significantly, by proving polynomial integrality gaps even for $n^{\Omega(\varepsilon)}$ levels of the Lasserre hierarchy, for both the directed and undirected spanner problems. We also extend these integrality gaps to related problems, notably Directed Steiner Network and Shallow-Light Steiner Network.

1 Introduction

A spanner is a subgraph which approximately preserves distances. Formally, a t-spanner of a graph G is a subgraph H such that $d_H(u, v) \leq t \cdot d_G(u, v)$ for all $u, v \in V$ (where d_H and d_G denote shortest-path distances in H and G respectively). Since H is a subgraph it is also the case that $d_G(u, v) \leq d_H(u, v)$, and thus a t-spanner preserves all distances up to a multiplicative factor of t, which is known as the *stretch*. Graph spanners originally appeared in the context of distributed computing [26,27], but have since been used as fundamental building blocks in applications ranging from routing in computer networks [30] to property testing of functions [5] to parallel algorithms [20].

Most work on graph spanners has focused on tradeoffs between various parameters, particularly the size (number of edges) and the stretch. Most notably, a seminal result of Althöfer et al. [1] is that every graph admits a $(2k-1)$-spanner with at most $n^{1+1/k}$ edges, for every integer $k \geq 1$. This tradeoff is also known to be tight, assuming the Erdős girth conjecture [19], but extensions to this fundamental result have resulted in an enormous literature on graph spanners.

Alongside this work on tradeoffs, there has been a line of work on *optimizing* spanners. In this line of work, we are usually given a graph G and a value t,

Supported in part by NSF award CCF-1909111.

C. Kaklamanis and A. Levin (Eds.): WAOA 2020, LNCS 12806, pp. 97–112, 2021.
https://doi.org/10.1007/978-3-030-80879-2_7

and are asked to find the t-spanner of G with the fewest number of edges. If G is undirected then this is known as Basic t-Spanner, while if G is directed then this is known as Directed t-Spanner. The best known approximation for Directed t-Spanner is an $O(n^{1/2})$-approximation [4], while for Basic t-Spanner the best known approximations are $O(n^{1/3})$ when $t = 3$ [4] and when $t = 4$ [17], and $O\left(n^{\frac{1}{\lfloor (t+1)/2 \rfloor}}\right)$ when $t > 4$. Note that this approximation for $t > 4$ is directly from the result of [1] by using the trivial fact that the optimal solution is always at least $n-1$ (in a connected graph). So this result is in a sense "generic", as both the upper bound and the lower bound are universal rather than applying to the particular input graph.

One feature of the algorithms of [4,17], as well as earlier work [14] and extensions to related settings (such as approximating fault-tolerant spanners [15,17] and minimizing the maximum or norm of the degrees [8,9]), is that they all use some variant of the same basic LP: a flow-based relaxation originally introduced for spanners by [14]. The result of [4] uses a slightly different LP (based on cuts rather than flows), but it is easy to show that the LP of [4] is no stronger than the LP of [14].

The fact that for Basic t-Spanner we cannot do better than the "generic" bound when $t > 4$, as well as the common use of a standard LP relaxation, naturally gives rise to a few questions. Is it possible to do better than the generic bound when $t > 4$? Can this be achieved with the basic LP? Can we improve our analysis of this LP to get improvements for Directed t-Spanner? In other words: what is the power of convex relaxations for spanner problems? It seems particularly promising to use lift-and-project methods to try for stronger LP relaxations, since one of the very few spanner approximations that uses a different LP relaxation was the use of the Sherali-Adams hierarchy to give an approximation algorithm for the Lowest Degree 2-Spanner problem [12].

It has been known since [18,21] that Directed t-Spanner does not admit an approximation better than $2^{\log^{1-\varepsilon} n}$ for any constant $\varepsilon > 0$ (under standard complexity assumptions), and it was more recently shown [13] that Basic t-Spanner cannot be approximated any better than $2^{(\log^{1-\varepsilon} n)/t}$ for any constant $\varepsilon > 0$. Thus no convex relaxation can do better than these bounds. But it is possible to prove stronger integrality gaps: it was shown in [14] that the integrality gap of the basic LP for Directed t-Spanner is at least $\tilde{\Omega}(n^{\frac{1}{3}-\varepsilon})$, while in [17] it was shown that the basic LP for Basic t-Spanner has an integrality gap of at least $\Omega(n^{\frac{2}{(1+\varepsilon)(t+1)+4}})$, nearly matching the generic upper bound (particularly for large t). But this left open a tantalizing prospect: perhaps there are stronger relaxations which could be used to get improved approximation bounds. Of course, the hardness of approximation results prove a limit to this. But even with the known hardness results and integrality gaps, it is possible that there is, say, an $O(n^{1/1000})$-approximation for Directed t-Spanner and an $O(n^{1/(1000t)})$-approximation for Basic t-Spanner that uses more advanced relaxations. It is also possible that a more complex relaxation, which perhaps cannot be solved in polynomial time, could lead to better approximations (albeit not in polynomial time).

1.1 Our Results and Techniques

This is the problem which we investigate: can we design stronger relaxations for spanners and related problems? While we cannot rule out all possible relaxations, we show that an extremely powerful lift-and-project technique, the Lasserre hierarchy [22] applied to the basic LP, does not give relaxations which are significantly better than the basic LP. This is true despite the fact that Lasserre is an SDP hierarchy rather than an LP hierarchy, and despite the fact that we allow a polynomial number of levels in the hierarchy (even though only a constant number of levels can be solved in polynomial time). And since the Lasserre hierarchy is at least as strong as other hierarchies such as the Sherali-Adams hierarchy [29] and the Lovasz-Schrijver hierarchy [24], our results also imply integrality gaps for these hierarchies.

In other words: *we show that even super-polynomial time algorithms which are based on any of the most well-known lift-and-project methods cannot approximate spanners much better than the current approximations based on the basic LP.*

Slightly more formally, we first rewrite the basic LP in a way that is similar to [4] but is equivalent to the stronger original formulation [14]. This makes the Lasserre lifts of the LP easier to reason about, thanks to the new structure of this formulation. We then consider the Lasserre hierarchy applied to this LP, and prove the following theorems.

Theorem 1. *For every constant $0 < \varepsilon < 1$ and sufficiently large n, the integrality gap of the $n^{\Omega(\varepsilon)}$-th level Lasserre SDP for* Directed $(2k-1)$-Spanner *is at least* $\left(\frac{n}{k}\right)^{\frac{1}{18}-\Theta(\varepsilon)}$.

Theorem 2. *For every constant $0 < \varepsilon < 1$ and sufficiently large n, the integrality gap of the $n^{\Omega(\varepsilon)}$-th level Lasserre SDP for* Basic $(2k-1)$-Spanner *is at least* $\frac{1}{k} \cdot \left(\frac{n}{k}\right)^{\min\left\{\frac{1}{18}, \frac{5}{32k-6}\right\}-\Theta(\varepsilon)} = n^{\Theta\left(\frac{1}{k}-\varepsilon\right)}$.

While the constant in the exponent is different, Theorem 2 is similar to [17] in that it shows that the integrality gap "tracks" the trivial approximation from [1] as a function of k. Thus for undirected spanners, even using the Lasserre hierarchy cannot give too substantial an improvement over the trivial greedy algorithm.

At a very high level, we follow the approach to building spanner integrality gaps of [14,17]. They started with random instances of the Unique Games problem, which are known to not admit any good solutions [6]. They then used these Unique Games instances to build spanner instances with the property that every spanner had to be large (or else the Unique Games instance would have had a good solution), but by "splitting flow" the LP could be very small.

In order to apply this framework to the Lasserre hierarchy, we need to make a number of changes. First, since Unique Games can be solved reasonably well by Lasserre [3,7], starting with a random instance of Unique Games will not work. Instead, we start with a more complicated problem known as Projection Games (the special case of Label Cover in which all the edge relations are functions). Hardness reductions for spanners have long used Projection Games [13,18,21], but previous integrality gaps [14,17] have used Unique Games since it was sufficient and was easier to work with. But we will need to overcome the complications that come from Projection Games.

Fortunately, an integrality gap for the Lasserre hierarchy for Projection Games was recently given by [10,25] (based on an integrality gap for CSPs from [31]), so we can use this as our starting point and try to plug it into the integrality gap framework of [14,17] to get an instance of either directed or undirected spanners. Unfortunately, the parameters and structure that we get from this are different enough from the parameters used in the integrality gap of the basic LP that we cannot use [14,17] as a black box. We need to reanalyze the instance using different techniques, even for the "easy" direction of showing that there are no good integral solutions. In order to do this, we also need some additional properties of the gap instance for Projection Games from [10] which were not stated in their original analysis. So we cannot even use [10] as a black box.

The main technical difficulty, though, is verifying that there is a "low-cost" fractional solution to the resulting SDP. For the basic LP this is straightforward, but for Lasserre we need to show that the associated slack moment matrices are all PSD. This turns out to be surprisingly tricky, because the slack moment matrices for the spanner SDP are much more complicated than the constraints of the Projection Game SDP. But by decomposing these matrices carefully we can show that each matrix in the decomposition is PSD. At a high level, we decompose the slack moment matrices as a summation of several matrices in a way that allows us to use the consistency properties of the feasible solution to the Projection Games instance in [10] to show that the overall sum is PSD.

Doing this requires us to use some properties of the feasible fractional solution provided by [10], some of which we need to prove as they were not relevant in the original setting. In particular, one important property which makes our task much easier is that their fractional solution actually satisfies *all* of the edges in the Projection Games instance. That is, their integrality gap is in a particular "place": the fractional solution has value 1 while every integral solution has much smaller value. This is enormously useful to us since spanners and the other network design problems we consider are minimization problems rather than maximization problems like Projection Games, as it essentially allows us to use essentially "the same" fractional solution (as it will also be feasible for the minimization version since it satisfies all edges). Technically, we end up combining this fact about the fractional solution of [10] with several properties of the Lasserre hierarchy to infer some more refined structural properties of the derived fractional solution for spanners, allowing us to argue that they are feasible for the Lasserre lifts.

Extensions. A number of other network design problems exhibit behavior that is similar to spanners, and we can extend our integrality gaps to these problems. In particular, we give a new integrality gap for Lasserre for Directed Steiner Network (DSN) (also called Directed Steiner Forest) and Shallow-Light Steiner Network (SLSN) [2]. In DSN we are given a directed graph $G = (V, E)$ (possibly with weights) and a collection of pairs $\{(s_i, t_i)\}_{i \in [p]}$, and are asked to find the cheapest subgraph such that there is an s_i to t_i path for all $i \in [p]$. In SLSN the graph is undirected, but each s_i and t_i is required to be connected within a global distance bound L. The best known approximation for DSN is an $O(n^{3/5+\varepsilon})$-approximation for arbitrarily small constant $\varepsilon > 0$ [11], which uses a standard flow-based LP relaxation. We can use the ideas we developed for spanners to also give integrality gaps for the Lasserre lifts of these problems. We provide the theorems here; details and proofs can be found in the full version [16].

Theorem 3. *For every constant $0 < \varepsilon < 1$ and sufficiently large n, the integrality gap of the $n^{\Omega(\varepsilon)}$-th level Lasserre SDP for* Directed Steiner Network *is at least $n^{\frac{1}{16}-\Theta(\varepsilon)}$.*

Theorem 4. *For every constant $0 < \varepsilon < 1$ and sufficiently large n, the integrality gap of the $n^{\Omega(\varepsilon)}$-th level Lasserre SDP for* Shallow-Light Steiner Network *is at least $n^{\frac{1}{16}-\Theta(\varepsilon)}$.*

2 Preliminaries: Lasserre Hierarchy

The Lasserre hierarchy lifts a polytope to a higher dimensional space, and then optionally projects this lift back to the original space in order to get tighter relaxations. The standard characterization for Lasserre (when the base polytope includes the hypercube constraints that every variable is in $[0, 1]$) is as follows [22,23,28]:

Definition 1 (Lasserre Hierarchy). *Let $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$, and define the polytope $K = \{x \in \mathbb{R}^n : Ax \geq b\}$. The r-th level of the Lasserre hierarchy $L_r(K)$ consists of the set of vectors $y \in [0, 1]^{\mathscr{P}([n])}$ (where $\mathscr{P}$ denotes the power set) that satisfy the following constraints:*

$$y_\varnothing = 1, \qquad M_{r+1}(y) := (y_{I \cup J})_{|I|,|J| \leq r+1} \succeq 0,$$

$$\forall \ell \in [m]: \quad M_r^\ell(y) := \left(\sum_{i=1}^{n} A_{\ell i} y_{I \cup J \cup \{i\}} - b_\ell y_{I \cup J} \right)_{|I|,|J| \leq r} \succeq 0$$

The matrix M_{r+1} is called the base moment matrix, and the matrices M_r^ℓ are called slack moment matrices.

Let us review (see, e.g., [28]) multiple helpful properties that we will use later. We include proofs in the full version [16] for completeness.

Claim. If $M_r(y) \succcurlyeq 0$, $|I| \leq r$, and $y_I = 1$, then $y_{I \cup J} = y_J$ for all $|J| \leq r$.

Claim. If $M_r(y) \succcurlyeq 0$, $|I| \leq r$, and $y_I = 0$, then $y_{I \cup J} = 0$ for all $|J| \leq r$.

Lemma 1. *If $M_{r+1}(y) \succeq 0$ then $M^{i,1}(y) = (y_{I \cup J \cup \{i\}})_{|I|,|J| \leq r} \succeq 0$ and $M^{i,0}(y) = (y_{I \cup J} - y_{I \cup J \cup \{i\}})_{|I|,|J| \leq r} \succeq 0$ for all $i \in [n]$.*

3 Projection Games: Background and Previous Work

In this section we discuss the Projection Games problem, its Lasserre relaxation, and the integrality gap that was recently developed for it [10] which form the basis of our integrality gaps for spanners and related problems. We begin with the problem definition.

Definition 2 (Projection Games). *Given a bipartite graph (L, R, E), a (label) alphabet set Σ, and projections (functions) $\pi_e : \Sigma \to \Sigma$ for each $e \in E$, the objective is to find a label assignment $\alpha : L \cup R \to \Sigma$ that maximizes the number of edges $e = (v_L, v_R) \in E$ where $\pi_e(\alpha(v_L)) = \alpha(v_R)$. We refer to these as satisfied edges.*

We sometimes use relation notation for the projection π_e, e.g., we use $(\sigma_1, \sigma_2) \in \pi_e$. Projection Games is the standard Label Cover problem, but where every relation is required to be a projection (and hence we inherit the relation notation when useful). Similarly, if we further restrict every function π_e to be a bijection, then we have the Unique Games problem. So Projection Games lies between Unique Games and Label Cover.

The basis of our integrality gaps is the integrality gap instance recently shown by [10] for Lasserre relaxations of Projection Games. We first formally define this SDP. For every $\Psi \subseteq (L \cup R) \times \Sigma$ we will have a variable y_Ψ. Then the r-th level Lasserre SDP for Projection Games is the following:

$$\mathrm{SDP}^r_{Proj}: \qquad \begin{aligned} \max & \sum_{(v_L,v_R) \in E, (\sigma_L,\sigma_r) \in \pi_{(u,v)}} y_{(v_L,\sigma_L),(v_R,\sigma_R)} \\ s.t.\ & y_\varnothing = 1 \\ & M_r(y) := \left(y_{\Psi_1 \cup \Psi_2}\right)_{|\Psi_1|,|\Psi_2| \leq r} \succcurlyeq 0 \\ & M_r^v(y) := \left(\sum_{\sigma \in \Sigma} y_{\Psi_1 \cup \Psi_2 \cup \{(v,\sigma)\}} - y_{\Psi_1 \cup \Psi_2}\right)_{|\Psi_1|,|\Psi_2| \leq r} = \mathbf{0} \ \forall v \in V \end{aligned}$$

It is worth noting that for simplicity we are not using the original presentation of this SDP given by [10]: they used a vector inner product representation. But it can be shown (and we do so in the full version [16]) that these representations are equivalent. To prove their integrality gap, [10] gave a Projection Games instance with the following properties. One of the properties is not proven in their paper, but is essentially trivial. We give a proof of this property, as well as a discussion of how the other properties follow from their construction, in the full version [16].

Lemma 2. *For any constant $0 < \varepsilon < 1$, there exists an instance of* Projection Games $(L, R, E_{Proj}, \Sigma, (\pi_e)_{e \in E_{Proj}})$ *with the following properties:*

1. $\Sigma = \left[n^{\frac{3-3\varepsilon}{5}}\right]$, $R = \{x_1, \ldots, x_n\}$, $L = \{c_1, \ldots, c_m\}$, *where* $m = n^{1+\varepsilon}$.
2. *There exists a feasible solution* $\mathbf{y}^*$ *for the* $r = n^{\Omega(\varepsilon)}$*-th level* SDP^r_{Proj}*, such that for all* $\{c_i, x_j\} \in E_{Proj}$, $\sum_{(\sigma_L, \sigma_R) \in \pi_{(c_i, x_j)}} y_{(c_i, \sigma_L), (x_j, \sigma_R)} = 1$.
3. *At most* $O\left(\frac{n^{1+\varepsilon} \ln n}{\varepsilon}\right)$ *edges can be satisfied.*
4. *The degree of vertices in* L *is* $K = n^{\frac{1-\varepsilon}{5}} - 1$, *and the degree of vertices in* R *is at most* $2Kn^{\varepsilon}$.

4 Lasserre Integrality Gap for Directed $(2k-1)$-Spanner

In this section we prove our main result for the Directed $(2k-1)$-Spanner problem: a polynomial integrality gap for polynomial levels of the Lasserre hierarchy.

4.1 Spanner LPs and Their Lasserre Lifts

The standard flow-based LP for spanners (including both the directed and basic $(2k-1)$-spanner problems) was introduced by [14], and has subsequently been used in many other spanner problems [4,8,15]. However, it is an awkward fit with Lasserre, most notably since it has an exponential number of variables (and hence so too do all the Lasserre lifts). We instead will work with an equivalent LP which appeared implicitly in [14] and is slightly stronger than the "antispanner" LP of [4].

First we review the standard LP, as stated in [14], and then describe the new LP. Let $\mathcal{P}_{u,v}$ denote the set of all stretch-$(2k-1)$ paths from u to v.

$$
\begin{aligned}
\text{LP}^{Flow}_{Spanner}: \quad \min & \sum_{e \in E} x_e \\
s.t. & \sum_{P \in \mathcal{P}_{u,v}: e \in P} f_P \le x_e && \forall (u,v) \in E, \forall e \in E \\
& \sum_{P \in \mathcal{P}_{u,v}} f_P \ge 1 && \forall (u,v) \in E \\
& x_e \ge 0 && \forall e \in E \\
& f_P \ge 0 && \forall (u,v) \in E, P \in \mathcal{P}_{u,v}
\end{aligned}
$$

Since the f_P variables do not appear in the objective function, we can project the polytope defined by $\text{LP}^{Flow}_{Spanner}$ onto the x_e variables and use the same objective function to get an equivalent LP but with only the x_e variables. To define this LP more formally, let $\mathcal{Z}^{u,v} = \{\mathbf{z} \in [0,1]^{|E|} : \sum_{e \in P} z_e \ge 1 \ \forall P \in \mathcal{P}_{u,v}\}$ be the polytope bounded by $0 \le z_e \le 1$ for all $e \in E$ and $\sum_{e \in P} z_e \ge 1$ for all $P \in \mathcal{P}_{u,v}$. Then we will use the following LP relaxation as our starting point, which informally says that every "stretch-$(2k-1)$ fractional cut" must be covered, and hence in an integral solution there is a stretch-$(2k-1)$ path for all edges (and so corresponds to a $(2k-1)$-spanner).

$$\begin{aligned}\mathrm{LP}_{Spanner}: \quad \min & \sum_{e \in E} x_e \\ s.t. & \sum_{e \in E} z_e x_e \geq 1 \ \forall (u,v) \in E, \text{ and } \mathbf{z} \in \mathcal{Z}^{u,v} \\ & x_e \geq 0 \qquad \forall e \in E\end{aligned}$$

While as written there are an infinite number of constraints, it is easy to see by convexity that we need to only include the (exponentially many) constraints corresponding to vectors $\mathbf{z}$ that are vertices in the polytope $\mathcal{Z}^{u,v}$, for each $(u,v) \in E$. Thus there are only an exponential number of constraints, but for simplicity we will analyze this LP as if there were constraints for all possible $\mathbf{z}$. We show in the full version [16] that $\mathrm{LP}_{Spanner}$ and $\mathrm{LP}_{Spanner}^{Flow}$ are equivalent.

Using Definition 1, we can write the r-th level Lasserre SDP of $\mathrm{LP}_{Spanner}$ (which is the SDP we consider in our integrality gap results) as follows. We let $\mathrm{SDP}_{Spanner}^{r}$ denote this SDP.

$$\begin{aligned}&\min \textstyle\sum_{e \in E} y_e \\ &s.t.\ y_{\varnothing} = 1 \\ &\quad M_{r+1}(y) := (y_{I \cup J})_{|I|,|J| \leq r+1} \succcurlyeq 0 \\ &\quad M_r^{\mathbf{z}}(y) := \left(\sum_{e \in E} z_e y_{I \cup J \cup \{e\}} - y_{I \cup J} \right)_{|I|,|J| \leq r} \succcurlyeq 0 \ \forall (u,v) \in E, \mathbf{z} \in \mathcal{Z}^{u,v}\end{aligned}$$

4.2 Spanner Instance

We now formally define the instance of DIRECTED $(2k-1)$-SPANNER that we will analyze to prove the integrality gap. We follow the framework of [14], who showed how to use the hardness framework of [18,21] to prove integrality gaps for the basic flow LP. We start with a different instance (integrality gaps instances for PROJECTION GAMES rather than random instances of UNIQUE GAMES), and also change the reduction in order to obtain a better dependency on k.

Roughly speaking, given a PROJECTION GAMES instance, we start with the "label-extended" graph. For each original vertex in the projection game, we create a group of vertices in the spanner instance of size $|\Sigma|$, so each vertex in the group can be thought as a label assignment for the PROJECTION GAMES vertex. We add paths between these groups corresponding to each function π_e (we add a path if the associated assignment satisfies the PROJECTION GAMES edges). We add many copies of the PROJECTION GAMES graph itself as the "outer graph", and then connect each PROJECTION GAMES vertex to the group associated with it. The key point is to prove that any integral solution must contain either many outer edges or many "connection edges" (in order to span the outer edges), while the fractional solution can buy connection and inner edges fractionally and simultaneously span all of the outer edges. A schematic version of this graph is presented in Fig. 1.

More formally, given the PROJECTION GAMES instance $(L = \cup_{i \in [m]} \{c_i\}, R = \cup_{i \in [n]} \{x_i\}, E_{Proj}, \Sigma = [n^{\frac{3-3\varepsilon}{5}}], (\pi_e)_{e \in E_{Proj}})$ from Lemma 2, we create a directed graph $G = (V, E)$ as follows (note that K is the degree of the vertices in L).

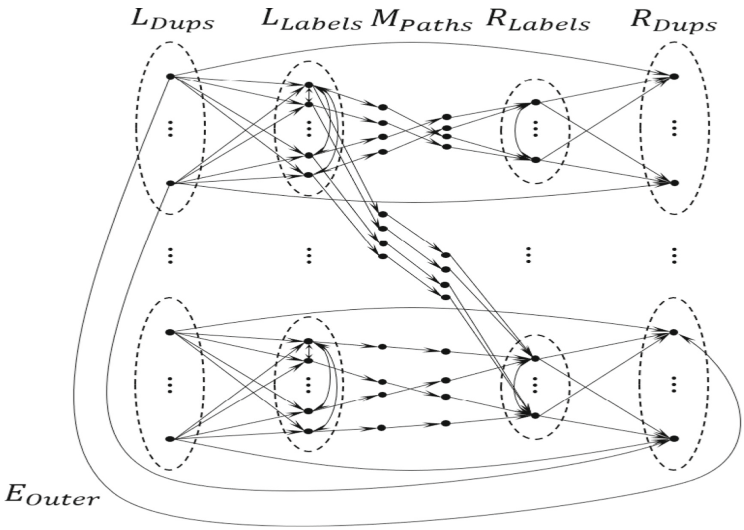

Fig. 1. Directed $(2k-1)$-Spanner instance

For every $c_i \in L$, we create $|\Sigma| + kK|\Sigma|$ vertices: $c_{i,\sigma}$ for all $\sigma \in \Sigma$ and c_i^l for all $l \in [kK|\Sigma|]$. We also create edges $(c_i^l, c_{i,\sigma})$ for each $\sigma \in \Sigma$ and $l \in [kK|\Sigma|]$. We call this edge set E_L. For every $x_i \in R$, we create $|\Sigma|+kK|\Sigma|$ vertices: $x_{i,\sigma}$ for $\sigma \in \Sigma$ and x_i^l for $l \in [kK|\Sigma|]$. We also create an edge $(x_i^l, x_{i,\sigma})$ for each $\sigma \in \Sigma$ and $l \in [kK|\Sigma|]$. We call this edge set E_R. For every $e = \{c_i, x_j\} \in E_{Proj}$, we create edges (c_i^l, x_j^l) for each $l \in [kK|\Sigma|]$. We call this edge set E_{Outer}. For each $e = \{c_i, x_j\} \in E_{Proj}$ and $(\sigma_L, \sigma_R) \in \pi_{i,j}$, we create vertices $w_{i,j,\sigma_L,\sigma_R,t}$ for $t \in [2k-4]$ and edges $(c_{i,\sigma_L}, w_{i,j,\sigma_L,\sigma_R,1}), (w_{i,j,\sigma_L,\sigma_R,1}, w_{i,j,\sigma_L,\sigma_R,2}), \ldots, (w_{i,j,\sigma_L,\sigma_R,2k-4}, x_{j,\sigma_R})$. We call this edge set E_M. Finally, for technical reasons we need some other edges E_{LStars} and E_{RStars} inside each of the groups groups.

More formally, $V = L_{Dups} \cup L_{Labels} \cup M_{Paths} \cup R_{Labels} \cup R_{Dups}$ and $E = E_L \cup E_{LStars} \cup E_M \cup E_{RStars} \cup E_R \cup E_{Outer}$, such that:

$$
\begin{aligned}
&L_{Labels} = \{c_{i,\sigma} \mid i \in |L|, \sigma \in \Sigma\}, && L_{Dups} = \{c_i^l \mid i \in |L|, l \in [kK|\Sigma|]\},\\
&R_{Labels} = \{x_{i,\sigma} \mid i \in |R|, \sigma \in \Sigma\}, && R_{Dups} = \{x_i^l \mid i \in |R|, l \in [kK|\Sigma|]\},\\
&M_{Paths} = \{w_{i,j,\sigma_L,\sigma_R,t} \mid \{c_i, x_j\} \in E_{Proj}, (\sigma_L, \sigma_R) \in \pi_{i,j}, t \in [2k-4]\} &&\\
&E_L^{i,l} = \{(c_i^l, c_{i,\sigma}) \mid \sigma \in \Sigma\}, && E_L^l = \cup_{i \in |L|} E_L^{i,l},\ E_L = \cup_{l \in [kK|\Sigma|]} E_L^l,\\
&E_R^{i,l} = \{(x_i^l, x_{i,\sigma}) \mid \sigma \in \Sigma\}, && E_R^l = \cup_{i \in |R|} E_R^{i,l},\ E_R = \cup_{l \in [kK|\Sigma|]} E_R^l,\\
&E_M^{i,j,\sigma_L,\sigma_R} = \{(c_{i,\sigma_L}, w_{i,j,\sigma_L,\sigma_R,1}), (w_{i,j,\sigma_L,\sigma_R,1}, w_{i,j,\sigma_L,\sigma_R,2}), &&\\
&\qquad (w_{i,j,\sigma_L,\sigma_R,2}, w_{i,j,\sigma_L,\sigma_R,3}), \ldots, (w_{i,j,\sigma_L,\sigma_R,2k-4}, x_{j,\sigma_R})\} &&\\
&E_M^{i,j} = \cup_{(\sigma_L,\sigma_R) \in \pi_{i,j}} E_M^{i,j,\sigma_L,\sigma_R} && E_M = \cup_{i,j:\{c_i,x_j\} \in E_{Proj}} E_M^{i,j}\\
&E_{Outer} = \{(c_i^l, x_j^l) \mid \{c_i, x_j\} \in E_{Proj}, l \in [kK|\Sigma|]\} &&\\
&E_{LStars} = \{(c_{i,1}, c_{i,\sigma}), (c_{i,\sigma}, c_{i,1}) \mid i \in |L|, \sigma \in \Sigma \setminus \{1\}\}, &&\\
&E_{RStars} = \{(x_{i,1}, x_{i,\sigma}), (x_{i,\sigma}, x_{i,1}) \mid i \in |R|, \sigma \in \Sigma \setminus \{1\}\}, &&
\end{aligned}
$$

Note that if $k < 3$ then then there is no vertex set M_{Paths}, but only E_M which directly connects c_{i,σ_L} and x_{j,σ_R} for each $\{c_i, x_j\} \in E_{Proj}$ and $(\sigma_L, \sigma_R) \in \pi_{i,j}$.

4.3 Fractional Solution

We now provide a low-cost feasible vector solution for the r-th level Lasserre lift of the spanner instance described above. Slightly more formally, we define values $\{y'_S : S \subseteq E\}$ and show that they form a feasible solution for the r-th level Lasserre lift $\mathrm{SDP}^r_{Spanner}$, and show that the objective value is $O(|V|)$. We start with a feasible solution $\{y^*_\Psi : \Psi \subseteq (L \cup R) \times \Sigma\}$ to the $(r+2)$-th level Lasserre lift $\mathrm{SDP}^{r+2}_{Proj}$ for the Projection Games instance (based on Lemma 2) that was used to construct our directed spanner instance, and adapt it to a solution for the spanner problem. We first define a function $\Phi : E \setminus E_{Outer} \to \mathscr{P}((L \cup R) \times \Sigma)$ (where $\mathscr{P}$ indicates the power set) as follows:

$$\Phi(e) = \begin{cases} \varnothing, & \text{if } e \in E_{LStars} \cup E_M \cup E_{RStars} \\ \{(c_i, \sigma)\}, & \text{if } e \in E_L \text{ and } e \text{ has an endpoint } c_{i,\sigma} \in L_{labels} \\ \{(x_i, \sigma)\}, & \text{if } e \in E_R \text{ and } e \text{ has an endpoint } x_{i,\sigma} \in R_{labels} \end{cases}$$

We then extend the definition of Φ to $\mathscr{P}(E \setminus E_{Outer}) \to \mathscr{P}((L \cup R) \times \Sigma)$ by setting $\Phi(S) = \cup_{e \in S} \Phi(e)$. Next, we define the solution $\{y'_S \mid S \subseteq E\}$: For any set S containing any edge in E_{Outer}, we define $y'_S = 0$, otherwise, let $y'_S = y^*_{\Phi(S)}$. Note that based on how we defined the function Φ, for all edges in $E_{LStars} \cup E_M \cup E_{RStars}$ we have $y'_S = y^*_\varnothing = 1$. In other words, these edges will be picked integrally in our feasible solution. At a very high level, what we are doing is fractionally buying edges in E_L, E_R and integrally buying edges in E_M in order to span edges in E_{Outer}. The edges in E_{LStars} and E_{RStars} are used to span edges in E_L and E_R. For the cost, a straightforward calculation (which appears in the full version [16]) implies the following.

Lemma 3. *The objective value of* $\mathbf{y}'$ *in* $SDP^r_{Spanner}$ *is* $O(|V|)$.

Feasibility. We now show that the described vector solution $\mathbf{y}'$ is feasible for the r-th level of Lasserre, i.e., that all the moment matrices defined in $SDP^r_{Spanner}$ are PSD. This is the most technically complex part of the analysis, particularly for the slack moment matrices for edges in E_{outer}.

We first prove that the base moment matrix in $SDP^r_{Spanner}$ is PSD for our solution $\mathbf{y}'$ by using the fact that the base moment matrix in SDP^{r+2}_{Proj} is PSD for the Projection Games solution $\mathbf{y}^*$ (the proof be can found in the full version [16]).

Theorem 5. *The base moment matrix* $M_{r+1}(\mathbf{y}') = (y'_{I \cup J})_{|I|,|J| \le r+1}$ *is PSD.*

Showing that the slack moment matrices of our spanner solution are all PSD is more subtle and requires a case by case analysis. We divide this argument into three cases (given by the next three theorems), from simplest to most complex, depending on the edge in the underlying linear constraint.

Theorem 6. *Slack moment matrix* $M_r^{\mathbf{z}}(\mathbf{y}') = \left(\sum_{e \in E} z_e y'_{I\cup J\cup\{e\}} - y'_{I\cup J}\right)_{|I|,|J|\le r}$ *is PSD for all* $(u,v) \in E_{LStars} \cup E_M \cup E_{RStars}$ *and* $\mathbf{z} \in \mathcal{Z}^{u,v}$.

This first case corresponds to pairs (u,v) for which we assigned $y'_{(u,v)} = 1$, so is comparatively simple to analyze via basic properties of the Lasserre hierarchy. The proof is included in the full version [16].

Theorem 7. *Slack moment matrix* $M_r^{\mathbf{z}}(\mathbf{y}') = \left(\sum_{e \in E} z_e y'_{I\cup J\cup\{e\}} - y'_{I\cup J}\right)_{|I|,|J|\le r}$ *is PSD for every* $(u,v) \in E_L \cup E_R$ *and* $\mathbf{z} \in \mathcal{Z}^{u,v}$.

This second case corresponds to the edges in E_L and E_R, and is a bit more complex since these edges are only bought fractionally in our solution. At a high-level, we first partition the terms of the summation $\sum_{e\in E} z_e (y'_{I\cup J\cup\{e\}})_{|I|,|J|\le r}$ depending on which middle edge e is included in the path that spans (u,v), and then show that each part in the partition is PSD even after subtracting an auxiliary matrix. Finally, we show that the sum of these auxiliary matrices is PSD after subtracting $(y'_{I\cup J})_{|I|,|J|\le r}$. Each step can be derived from the fact that $\mathbf{y}^*$ is a feasible solution of $\mathrm{SDP}^{r+2}_{Proj}$. The complete proof is included in the full version [16].

Now we move to the main technical component of our integrality gap analysis: proving that the slack moment matrices corresponding to outer edges are PSD.

Theorem 8. *Slack moment matrix* $M_r^{\mathbf{z}}(\mathbf{y}') = \left(\sum_{e \in E} z_e y'_{I\cup J\cup\{e\}} - y'_{I\cup J}\right)_{|I|,|J|\le r}$ *is PSD for every* $(u,v) \in E_{outer}$ *and* $\mathbf{z} \in \mathcal{Z}^{u,v}$.

Proof. In this case the edge (u,v) is not included in the solution, and is thus only spanned using other (inner) edges. The main difficulty here is that the spanning path is longer, and each hop of the path contributes to the solution differently. To prove this theorem, we have to again partition the summation $\sum_{e\in E} z_e (y'_{I\cup J\cup\{e\}})_{|I|,|J|\le r}$, while subtracting and adding a more carefully designed set of auxiliary matrices. Note that in this argument we crucially use the fact that the SDP solution of [10] satisfies all of the demands (property 2. in Lemma 2).

We first show how to decompose $M_r^{\mathbf{z}}(\mathbf{y}')$ as the sum of several simpler matrices. This will let us reason about each matrix differently based on the assigned values and their connection to the Projection Games constraints. We will then explain why each of these matrices is PSD. Observe that for each $(u,v) = (c_i^l, x_j^l) \in E_{outer}$, the set of stretch-$(2k-1)$ paths consist of the outer edge, or one of the paths that go through some labels (σ_L, σ_R). It is not hard to see that any other path connecting such pairs has length larger than $(2k-1)$. More formally:

Claim. For every pair $(c_i^l, x_j^l) \in E_{Outer}$, the length $(2k-1)$ paths from c_i^l to x_j^l are: the path consisting of only the edge (c_i^l, x_j^l), or the paths consisting of edges $\{(c_i^l, c_{i,\sigma_L})\} \cup E_M^{i,j,\sigma_L,\sigma_R} \cup \{(x_{j,\sigma_R}, x_j^l)\}$ for some $(\sigma_L, \sigma_R) \in \pi_{i,j}$.

We use this observation, and the fact that $y_e = 0$ for all $e \in E_{Outer}$ to break $M_r^{\mathbf{z}}(\mathbf{y}')$ into several pieces, and argue that each piece is PSD.

$$\left(\sum_{e\in E} z_e y'_{I\cup J\cup\{e\}} - y'_{I\cup J}\right)_{|I|,|J|\le r} \tag{1}$$

$$= \sum_{e\in E} z_e \left(y^*_{\Phi(I\cup J)\cup\Phi(e)}\right)_{|I|,|J|\le r} - \left(y^*_{\Phi(I\cup J)}\right)_{|I|,|J|\le r}$$

$$= \sum_{\sigma_L\in\Sigma} z_{(c_i^l,c_{i,\sigma_L})} \left(y^*_{\Phi(I\cup J)\cup\{(c_i,\sigma_L)\}}\right)_{|I|,|J|\le r}$$

$$+ \sum_{(\sigma_L,\sigma_R)\in\pi_{i,j}} \sum_{e\in E_M^{i,j,\sigma_L,\sigma_R}} z_e \left(y^*_{\Phi(I\cup J)\cup\varnothing}\right)_{|I|,|J|\le r}$$

$$+ \sum_{\sigma_R\in\Sigma} z_{(x_{j,\sigma_R},x_j^l)} \left(y^*_{\Phi(I\cup J)\cup\{(x_j,\sigma_R)\}}\right)_{|I|,|J|\le r}$$

$$+ \sum_{e\in E\setminus(E_L^{i,l}\cup E_M^{i,j}\cup E_R^{j,l})} z_e \left(y'_{I\cup J\cup\{e\}}\right)_{|I|,|J|\le r} - \left(y^*_{\Phi(I\cup J)}\right)_{|I|,|J|\le r}$$

$$= \sum_{\sigma_L\in\Sigma} z_{(c_i^l,c_{i,\sigma_L})}\left(y^*_{\Phi(I\cup J)\cup\{(c_i,\sigma_L)\}} - \sum_{\sigma_R:(\sigma_L,\sigma_R)\in\pi_{i,j}} y^*_{\Phi(I\cup J)\cup\{(c_i,\sigma_L),(x_j,\sigma_R)\}}\right)_{|I|,|J|\le r} \tag{2}$$

$$+ \sum_{(\sigma_L,\sigma_R)\in\pi_{i,j}} \sum_{e\in E_M^{i,j,\sigma_L,\sigma_R}} z_e\left(y^*_{\Phi(I\cup J)} - y^*_{\Phi(I\cup J)\cup\{(c_i,\sigma_L),(x_j,\sigma_R)\}}\right)_{|I|,|J|\le r} \tag{3}$$

$$+ \sum_{\sigma_R\in\Sigma} z_{(x_{j,\sigma_R},x_j^l)}\left(y^*_{\Phi(I\cup J)\cup\{(x_j,\sigma_R)\}} - \sum_{\sigma_L:(\sigma_L,\sigma_R)\in\pi_{i,j}} y^*_{\Phi(I\cup J)\cup\{(c_i,\sigma_L),(x_j,\sigma_R)\}}\right)_{|I|,|J|\le r} \tag{4}$$

$$+ \sum_{e\in E\setminus(E_L^{i,l}\cup E_M^{i,j}\cup E_R^{j,l})} z_e \left(y'_{I\cup J\cup\{e\}}\right)_{|I|,|J|\le r} \tag{5}$$

$$- \left(y^*_{\Phi(I\cup J)} - \sum_{(\sigma_L,\sigma_R)\in\pi_{i,j}} y^*_{\Phi(I\cup J)\cup\{(c_i,\sigma_L),(x_j,\sigma_R)\}}\right)_{|I|,|J|\le r} \tag{6}$$

$$+ \sum_{(\sigma_L,\sigma_R)\in\pi_{i,j}} \left(z_{(c_i^l,c_{i,\sigma_L})} + \sum_{e\in E_M^{i,j,\sigma_L,\sigma_R}} z_e + z_{(x_{j,\sigma_R},x_j^l)} - 1\right) \tag{7}$$

$$\times \left(y^*_{\Phi(I\cup J)\cup\{(c_i,\sigma_L),(x_j,\sigma_R)\}}\right)_{|I|,|J|\le r} \tag{8}$$

To prove the third equality, we have subtracted and added the sum over $(\sigma_L, \sigma_R) \in \pi_{i,j}$, and then partitioned it over the other sums.

The matrix in (3) is PSD by basic properties of Lasserre. The matrix in (5) is PSD because it is a principal submatrix of $M_{r+1}(\mathbf{y}')$, which is PSD by Theorem 5. The matrix in (8) is equal to $\left(y'_{I\cup J\cup\{(c_i^l,c_{i,\sigma_L}),(x_{j,\sigma_R},x_j^l)\}}\right)_{|I|,|J|\le r}$, and so it is also a principal submatrix of $M_{r+1}(\mathbf{y}')$ and is hence PSD. The coefficient in (7) is non-negative because $\sum_{e\in P} z_e \ge 1$ for all path $P \in \mathcal{P}_{c_i^l,x_j^l}$ by definition of $\mathbf{z}$.

We now argue that the matrices in (2), (4), and (6) are all-zero matrices, which will complete the proof. In order to show this, we need the following claim, which uses the fact that the fractional solution $\mathbf{y}^*$ to the Projection Games instance satisfies all of the demands.

Lemma 4. *$y^*_{\Psi\cup\{(c_i,\sigma_L),(x_j,\sigma_R)\}} = 0$ for all $\{c_i, x_j\} \in E_{Proj}$, $(\sigma_L, \sigma_R) \notin \pi_{i,j}$, $|\Psi| \leq 2r$.*

Proof. By Property 2. of Lemma 2, we know that $\sum_{(\sigma_L,\sigma_R)\in\pi_{i,j}} y^*_{(c_i,\sigma_L),(x_j,\sigma_R)} = 1$. By the slack moment matrix constraints in SDP^r_{Proj}, we know that $M^{c_i}_r(\mathbf{y}^*)$ and $M^{x_j}_r(\mathbf{y}^*)$ are both all-zero matrices. Hence,

$$\sum_{\sigma_L\in\Sigma}\sum_{\sigma_R\in\Sigma} y^*_{(c_i,\sigma_L),(x_j,\sigma_R)} = \sum_{\sigma_R\in\Sigma} y^*_{(x_j,\sigma_R)} = 1 = \sum_{(\sigma_L,\sigma_R)\in\pi_{i,j}} y^*_{(c_i,\sigma_L),(x_j,\sigma_R)}$$

In other words, since $y^*_{(c_i,\sigma_L),(x_j,\sigma_R)} \geq 0$ for all $\sigma_L \in \Sigma$ and $\sigma_R \in \Sigma$, it follows for all $(\sigma_L, \sigma_R) \notin \pi_{i,j}$ that $y^*_{(c_i,\sigma_L),(x_j,\sigma_R)} = 0$. Then basic properties of Lasserre imply $y^*_{\Psi\cup\{(c_i,\sigma_L),(x_j,\sigma_R)\}} = 0$.

Next, we use this lemma to complete the proof. We argue that all entries of the matrix in (2) are zero. This is because for any entry with index I and J we have:

$$\begin{aligned} y^*_{\Phi(I\cup J)\cup\{(c_i\sigma_L)\}} &= \sum_{\sigma_R\in\Sigma} y^*_{\Phi(I\cup J)\cup\{(c_i,\sigma_L),(x_j,\sigma_R)\}} \\ &= \sum_{\sigma_R:(\sigma_L,\sigma_R)\in\pi_{i,j}} y^*_{\Phi(I\cup J)\cup\{(c_i,\sigma_L),(x_j,\sigma_R)\}}, \end{aligned}$$

where the last equality follows from above Lemma 4. A similar argument implies that the matrix in (4) is also all-zero.

To prove the matrix in (6) is also all-zero, consider an entry with index I and J. Then

$$\begin{aligned} y^*_{\Phi(I\cup J)} &= \sum_{\sigma_L\in\Sigma} y_{\Phi(I\cup J)\cup\{(c_i,\sigma_L)\}} = \sum_{\sigma_L\in\Sigma}\sum_{\sigma_R\in\Sigma} y^*_{\Phi(I\cup J)\cup\{(c_i,\sigma_L),(x_j,\sigma_R)\}} \\ &= \sum_{(\sigma_L,\sigma_R)\in\pi_{i,j}} y^*_{\Phi(I\cup J)\cup\{(c_i,\sigma_L),(x_j,\sigma_R)\}}. \end{aligned}$$

Again, we have used Lemma 4 in the last equality.

Therefore, (3), (5), (8) are PSD, (2), (4), and (6) are all-zero, and (7) is non-negative, which proves the theorem.

4.4 Proof of Theorem 1

In order to prove Theorem 1, we now just need to analyze the integral optimum solution. In the full version [16] we prove the following.

Lemma 5. *The optimal $(2k-1)$-spanner of G has at least $nkK|\Sigma|\sqrt{K}$ edges.*

The proof of Theorem 1 then follows (after plugging in all parameters) from Lemma 5, the feasibility of the fractional solutions (Theorems 5, 6, 7, and 8), and the objective value of the fractional solution (Lemma 3).

To see this, remember from Lemma 2 that $m = n^{1+\varepsilon}$, $K = n^{\frac{1-\varepsilon}{5}} - 1$, $|\Sigma| = n^{\frac{3-3\varepsilon}{5}}$. Also

$$\begin{aligned}|V| &= |L_{Dups}| + |L_{Labels}| + |M_{Paths}| + |R_{Labels}| + |R_{Dups}| \\ &= mkK|\Sigma| + m|\Sigma| + (2k-4)mK|\Sigma| + n|\Sigma| + nkK|\Sigma| \\ &= O(mkK|\Sigma|) = O\left(kn^{\frac{9+\varepsilon}{5}}\right)\end{aligned}$$

Therefore, $\frac{nkK|\Sigma|\sqrt{K}}{O(|V|)} = \Omega\left(\frac{\sqrt{K}}{n^{\varepsilon}}\right) = \Omega\left(n^{\frac{1-11\varepsilon}{10}}\right) = \Omega\left(\left(\frac{|V|}{k}\right)^{\frac{1}{18}-\Theta(\varepsilon)}\right)$, proving the theorem.

5 Undirected $(2k-1)$-Spanner

In order to extend these techniques to the undirected case (Theorem 2), we need to make a number of changes. This is the most technically complex result in this paper, so due to space constraints we defer all details to the full version [16]. First, for technical reasons we need to replace the middle paths E_M with edges, and instead add "outside" paths (as was done in [17]). More importantly, though, we have the same fundamental problem that always arises when moving from directed to undirected spanners: without directions on edges, there can be many more short cycles (and thus ways of spanning edges) in the resulting graph. In particular, if we directly change every edge on the integrality gap instance of Directed $(2k-1)$-Spanner in the previous section to be undirected, then there are more paths available to span the outer edges (more formally, the analog of Claim 4.3 is false), so the integral optimum might no longer be large.

This difficulty is fundamentally caused by the fact that the graph of the Projection Games instance from [10] that we use as our starting point might have short cycles. These turn into short spanning paths of outer edges in our spanner instance. In order to get around this, we first carefully subsample and prune to remove a selected subset of edges in E_{Proj}, causing the remaining graph to have large girth (at least $2k+2$) but without losing too much of its density or any of the other properties that we need. This is similar to what was done by [17] to prove an integrality gap for the base LP, but here we are forced to start with the instance of [10], which has a far more complicated structure than the random Unique Games instances used by [17]. This is the main technical difficulty, but once we overcome it we can use the same ideas as in Sect. 4 to prove Theorem 2.

Acknowledgements. The authors would like to thank Amitabh Basu for many helpful discussions.

References

1. Althöfer, I., Das, G., Dobkin, D., Joseph, D., Soares, J.: On sparse spanners of weighted graphs. Discrete Comput. Geom. **9**(1), 81–100 (1993)

2. Babay, A., Dinitz, M., Zhang, Z.: Characterizing demand graphs for (fixed-parameter) shallow-light steiner network. arXiv preprint arXiv:1802.10566 (2018)
3. Barak, B., Raghavendra, P., Steurer, D.: Rounding semidefinite programming hierarchies via global correlation. In: Proceedings of the 2011 IEEE 52Nd Annual Symposium on Foundations of Computer Science, pp. 472–481. FOCS 2011 (2011)
4. Berman, P., Bhattacharyya, A., Makarychev, K., Raskhodnikova, S., Yaroslavtsev, G.: Approximation algorithms for spanner problems and directed Steiner forest. Inf. Comput. **222**, 93–107 (2013)
5. Bhattacharyya, A., Grigorescu, E., Jung, K., Raskhodnikova, S., Woodruff, D.P.: Transitive-closure spanners. SIAM J. Comput. **41**(6), 1380–1425 (2012)
6. Charikar, M., Hajiaghayi, M., Karloff, H.J.: Improved approximation algorithms for label cover problems. Algorithmica **61**(1), 190–206 (2011)
7. Charikar, M., Makarychev, K., Makarychev, Y.: Near-optimal algorithms for unique games. In: Proceedings of the Thirty-eighth Annual ACM Symposium on Theory of Computing, pp. 205–214. STOC 2006 (2006)
8. Chlamtáč, E., Dinitz, M.: Lowest-degree k-spanner: approximation and hardness. Theor. Comput. **12**(15), 1–29 (2016)
9. Chlamtác, E., Dinitz, M., Robinson, T.: Approximating the Norms of Graph Spanners. In: Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019). Leibniz International Proceedings in Informatics (LIPIcs), vol. 145, pp. 11:1–11:22 (2019)
10. Chlamtác, E., Manurangsi, P., Moshkovitz, D., Vijayaraghavan, A.: Approximation algorithms for label cover and the log-density threshold. In: Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms. SODA 2017 (2017)
11. Chlamtáč, E., Dinitz, M., Kortsarz, G., Laekhanukit, B.: Approximating spanners and directed Steiner forest: upper and lower bounds. In: Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms. SODA 2017 (2017)
12. Chlamtáč, E., Dinitz, M., Krauthgamer, R.: Everywhere-sparse spanners via dense subgraphs. In: Proceedings of the 53rd IEEE Annual Symposium on Foundations of Computer Science, pp. 758–767. FOCS 2012 (2012)
13. Dinitz, M., Kortsarz, G., Raz, R.: Label cover instances with large girth and the hardness of approximating basic k-spanner. ACM Trans. Algorithms **12**(2), 25:1–25:16 (2015)
14. Dinitz, M., Krauthgamer, R.: Directed spanners via flow-based linear programs. In: Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing, pp. 323–332. STOC 2011 (2011)
15. Dinitz, M., Krauthgamer, R.: Fault-tolerant spanners: better and simpler. In: Proceedings of the 30th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing, pp. 169–178. PODC 2011 (2011)
16. Dinitz, M., Nazari, Y., Zhang, Z.: Lasserre integrality gaps for graph spanners and related problems. CoRR abs/1905.07468 (2019). http://arxiv.org/abs/1905.07468
17. Dinitz, M., Zhang, Z.: Approximating low-stretch spanners. In: Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms. SODA (2016)
18. Elkin, M., Peleg, D.: The hardness of approximating spanner problems. Theor. Comput. Syst. **41**(4), 691–729 (2007)
19. Erdős, P.: Extremal problems in graph theory. In: Theory Graphs and its Applications: Proceedings of the Symposium held in Smolenice 1963, pp. 29–36 (1964)
20. Friedrichs, S., Lenzen, C.: Parallel metric tree embedding based on an algebraic view on Moore-Bellman-Ford. J. ACM (JACM) **65**(6), 43 (2018)

21. Kortsarz, G.: On the hardness of approximating spanners. Algorithmica **30**(3), 432–450 (2001)
22. Lasserre, J.B.: An explicit exact SDP relaxation for nonlinear 0-1 programs. In: Aardal, K., Gerards, B. (eds.) IPCO 2001. LNCS, vol. 2081, pp. 293–303. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-45535-3_23
23. Laurent, M.: A comparison of the sherali-adams, lovász-schrijver, and lasserre relaxations for 0–1 programming. Math. Oper. Res. **28**(3), 470–496 (2003)
24. Lovász, L., Schrijver, A.: Cones of matrices and set-functions and 0–1 optimization. SIAM J. Optim. **1**(2), 166–190 (1991)
25. Manurangsi, P.: On approximating projection games. Master's thesis, Massachusetts Institute of Technology (2015)
26. Peleg, D., Schäffer, A.A.: Graph spanners. J. Graph Theor. **13**(1), 99–116 (1989)
27. Peleg, D., Ullman, J.D.: An optimal synchronizer for the hypercube. SIAM J. Comput. **18**(4), 740–747 (1989)
28. Rothvoß, T.: The Lasserre hierarchy in approximation algorithms. Lecture Notes for the MAPSP, pp. 1–25 (2013)
29. Sherali, H.D., Adams, W.P.: A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. SIAM J. Discret. Math. **3**(3), 411–430 (1990)
30. Thorup, M., Zwick, U.: Compact routing schemes. In: Proceedings of the Thirteenth Annual ACM Symposium on Parallel Algorithms and Architectures, pp. 1–10. ACM (2001)
31. Tulsiani, M.: CSP gaps and reductions in the Lasserre hierarchy. In: Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing, pp. 303–312. ACM (2009)