

# The Norms of Graph Spanners

**Eden Chlamtáč**

Ben Gurion University of the Negev, Beersheva, Israel

**Michael Dinitz**

Johns Hopkins University, Baltimore, MD, USA

**Thomas Robinson**

Ben Gurion University of the Negev, Beersheva, Israel

---

## Abstract

A  $t$ -spanner of a graph  $G$  is a subgraph  $H$  in which all distances are preserved up to a multiplicative  $t$  factor. A classical result of Althöfer et al. is that for every integer  $k$  and every graph  $G$ , there is a  $(2k - 1)$ -spanner of  $G$  with at most  $O(n^{1+1/k})$  edges. But for some settings the more interesting notion is not the number of edges, but the degrees of the nodes. This spurred interest in and study of spanners with small maximum degree. However, this is not necessarily a robust enough objective: we would like spanners that not only have small maximum degree, but also have “few” nodes of “large” degree. To interpolate between these two extremes, in this paper we initiate the study of graph spanners with respect to the  $\ell_p$ -norm of their degree vector, thus simultaneously modeling the number of edges (the  $\ell_1$ -norm) and the maximum degree (the  $\ell_\infty$ -norm). We give precise upper bounds for all ranges of  $p$  and stretch  $t$ : we prove that the greedy  $(2k - 1)$ -spanner has  $\ell_p$  norm of at most  $\max(O(n), O(n^{\frac{k+p}{kp}}))$ , and that this bound is tight (assuming the Erdős girth conjecture). We also study universal lower bounds, allowing us to give “generic” guarantees on the approximation ratio of the greedy algorithm which generalize and interpolate between the known approximations for the  $\ell_1$  and  $\ell_\infty$  norm. Finally, we show that at least in some situations, the  $\ell_p$  norm behaves fundamentally differently from  $\ell_1$  or  $\ell_\infty$ : there are regimes ( $p = 2$  and stretch 3 in particular) where the greedy spanner has a provably superior approximation to the generic guarantee.

**2012 ACM Subject Classification** Theory of computation → Sparsification and spanners

**Keywords and phrases** spanners, approximations

**Digital Object Identifier** 10.4230/LIPIcs.ICALP.2019.40

**Category** Track A: Algorithms, Complexity and Games

**Related Version** A full version of the paper is available at <https://arxiv.org/abs/1903.07418>.

**Funding** *Eden Chlamtáč*: Supported in part by ISF grant 1002/14.

*Michael Dinitz*: Supported in part by NSF awards CCF-1464239 and CCF-1535887.

*Thomas Robinson*: Supported in part by ISF grant 1002/14.

## 1 Introduction

Graph spanners are subgraphs which approximately preserve distances. Slightly more formally, given a graph  $G = (V, E)$  (possibly with lengths on the edges), a subgraph  $H$  of  $G$  is a  $t$ -spanner of  $G$  if  $d_G(u, v) \leq d_H(u, v) \leq t \cdot d_G(u, v)$  for all  $u, v \in V$ , where  $d_G$  denotes shortest-path distances in  $G$  (and  $d_H$  in  $H$ ). The value  $t$  is called the *stretch* of the spanner.

Graph spanners were originally introduced in the context of distributed computing [27, 26], but have since proved to be a fundamental building block that is useful in a variety of applications, from property testing [7] to network routing [28]. When building spanners there are many objectives which we could try to optimize, but probably the most popular is the number of edges (the *size* or the *sparsity*). Not only is sparsity important in many applications, it also admits a beautiful tradeoff with the stretch, proved by Althöfer et al. [2]:



© Eden Chlamtáč, Michael Dinitz, and Thomas Robinson;

licensed under Creative Commons License CC-BY

46th International Colloquium on Automata, Languages, and Programming (ICALP 2019).

Editors: Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi;

Article No. 40; pp. 40:1–40:15



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



► **Theorem 1** ([2]). *For every integer  $k \geq 1$  and every weighted graph  $G = (V, E)$  with  $|V| = n$ , there is a  $(2k - 1)$ -spanner  $H$  of  $G$  with at most  $O(n^{1+1/k})$  edges.*

While understanding the tradeoff between the size and the stretch was a seminal achievement, for many applications (particularly in distributed computing) we care not just about the size, but also about the *maximum degree*. Unfortunately, unlike the size, there is no possible tradeoff between the stretch and the maximum degree. This is trivial to see: if  $G$  is a star, then the only spanner of  $G$  with non-infinite stretch has maximum degree of  $n - 1$ . In general, if  $G$  has maximum degree  $\Delta$ , then all we can say is the trivial fact that  $G$  has a spanner with maximum degree at most  $\Delta$ . Nevertheless, given the importance of the maximum degree objective, there has been significant work on building spanners that minimize the maximum degree from the perspective of *approximation algorithms* [23, 10, 9]. From this perspective, we are given a graph  $G$  and stretch value  $t$  and are asked to find the “best”  $t$ -spanner of  $G$  (where “best” means minimizing the maximum degree).

While this has been an interesting and productive line of research, clearly there are problems with the maximum degree objective as well. For example, if it is unavoidable for there to be some node of large degree  $d$ , the maximum degree objective allows us to make every other vertex also of degree  $d$ , with no change in the objective function. But clearly we would prefer to have fewer high-degree nodes if possible!

So we are left with a natural question: can we define a notion of “cost” of a spanner which discourages very high degree nodes, but if there are high degree nodes, still encourages the rest of the nodes to have small degree? There is of course an obvious candidate for such a cost function: the  $\ell_p$  norm of the degree vector. That is, given a spanner  $H$ , we can define  $\|H\|_p$  to be the  $\ell_p$ -norm of the  $n$ -dimensional vector in which the coordinate corresponding to a node  $v$  contains the degree of  $v$  in  $H$ . Then  $\|H\|_1$  is just (twice) the total number of edges, and  $\|H\|_\infty$  is precisely the maximum degree. Thus the  $\ell_p$ -norm is an interpolation between these two classical objectives. Moreover, for  $1 < p < \infty$ , this notion of cost has precisely the properties that we want: it encourages low-degree nodes rather than high-degree nodes, but if high-degree nodes are unavoidable it still encourages the rest of the nodes to be as low-degree as possible. These properties, of interpolating between the average and the maximum, are why the  $\ell_p$ -norm has appeared as a popular objective for a variety of problems, ranging from clustering (the famous  $k$ -means problem [22, 24]), to scheduling [4, 3, 1], to covering [21].

## 1.1 Our Results and Techniques

In this paper we initiate the study of graph spanners under the  $\ell_p$ -norm objective. We prove a variety of results, giving upper bounds, lower bounds, and approximation guarantees. Our main result is the analog of Theorem 1 for the  $\ell_p$ -norm objective, but we also characterize universal lower bounds as part of an effort to understand the generic approximation ratio for the related optimization problem. We also show that in some ways the  $\ell_p$ -norm can behave fundamentally differently than the traditional  $\ell_1$  or  $\ell_\infty$  norms, by proving that the greedy algorithm can have an approximation ratio that is *strictly better* than the generic guarantee, unlike the  $\ell_1$  or  $\ell_\infty$  settings.

### 1.1.1 Upper Bound

We begin by proving our main result: a universal upper bound (the analog of Theorem 1) for  $\ell_p$ -norm spanners. Recall the classical greedy algorithm for constructing a  $t$ -spanner  $H$  of a graph  $G = (V, E)$ . Consider the edges in nondecreasing order of edge length, and when considering edge  $\{u, v\}$ , add it to  $H$  if currently  $d_H(u, v) > t \cdot d_G(u, v)$ . We call  $H$  the *greedy*

$t$ -spanner of  $G$ . It is trivial to show that the greedy  $t$ -spanner has girth at least  $t + 2$ . This is the algorithm that was used to prove Theorem 1, and it has since received extensive study (see, e.g., [20, 8]) and will form the basis of our upper bound:

► **Theorem 2.** *Let  $k \geq 1$  be an integer, let  $G = (V, E)$  be a graph (possibly with lengths on the edges), and let  $H = (V, E_H)$  be the greedy  $(2k - 1)$ -spanner of  $G$ . Then  $\|H\|_p \leq \max\left(O(n), O\left(n^{\frac{k+p}{kp}}\right)\right)$  for all  $p \geq 1$ .*

In other words, if  $p \geq \frac{k}{k-1}$  then our upper bound is  $O(n)$ , and otherwise it is  $O\left(n^{\frac{k+p}{kp}}\right)$ . Clearly this interpolates between  $p = 1$  and  $p = \infty$ : when  $p = 1$  this is the same bound as Theorem 1, while if  $p = \infty$  this gives  $O(n)$  which is the only possible bound in terms of  $n$ . To get some more intuition for this bound, note that  $n^{\frac{k+p}{kp}}$  would be the  $\ell_p$ -norm of  $H$  if  $H$  had the size guaranteed by Theorem 1 and was also regular. So one way to think of this bound is that while the greedy spanner can be non-regular, its  $\ell_p$ -norm still acts as if it were regular.

It is also straightforward to prove that this bound is tight if we again assume the Erdős girth conjecture [19]; for completeness, we do this in the full version [12].

The proof of Theorem 1 from [2] is relatively simple: the greedy  $(2k - 1)$ -spanner has girth at least  $2k + 1$ , and any graph with more than  $n^{1+1/k}$  edges must have a cycle of length at most  $2k$ . Generalizing this to the  $\ell_p$ -norm is significantly more complicated, since it is not nearly as easy to show a relationship between the girth and the  $\ell_p$ -norm. But this is precisely what we do.

It turns out to be easiest to prove Theorem 2 for stretch 3: it just takes one more step beyond [2] to split the vertices of the high-girth graph (the spanner) into “low” and “high” degrees, and show that each vertex set does not contribute too much to the  $\ell_p$  norm. However for larger stretch values this approach does not work: the main lemma used for stretch 3 (Lemma 5) is simply false when generalized to larger stretch bounds. Instead, we need a much more involved decomposition into “low”, “medium”, and “high”-degree nodes. This decomposition is very subtle, since the categories are not purely about the degree, but rather about how the degree relates to expansion at some particular distances from the node. We also need to further decompose the “high”-degree nodes into sets determined by which distance level we consider the expansion of. We then separately bound the contribution to the  $p$ -norm of each class in the decomposition; for “low”-degree nodes this is quite straightforward, but for medium and high-degree nodes this requires some subtle arguments which strongly use the structure of large-girth graphs.

### 1.1.2 Universal Lower Bounds

To motivate our next set of results, consider the optimization problem of finding the “best”  $t$ -spanner of a given input graph. When “best” is the smallest  $\ell_1$ -norm this is known as the BASIC  $t$ -SPANNER problem [16, 5, 15, 18], and when “best” is the smallest  $\ell_\infty$ -norm this is the LOWEST-DEGREE  $t$ -SPANNER problem [23, 10, 9]. It is natural to consider this problem for the  $\ell_p$ -norm as well. It is also natural to consider how well the greedy algorithm (used to prove the upper bound of Theorem 2) performs as an approximation algorithm.

To see an obvious way of analyzing the greedy algorithm as an approximation algorithm, consider the  $\ell_1$ -norm. Theorem 1 implies that the greedy algorithm always returns a spanner of size at most  $O(n^{1+1/k})$ , while clearly *every* spanner must have size at least  $\Omega(n)$  (assuming that the input graph is connected). Thus we immediately get that the greedy algorithm is an  $O(n^{1/k})$ -approximation. By dividing a universal upper bound (an upper bound on the size of the greedy spanner that holds for every graph) by a universal lower bound (a lower bound on the size of *every* spanner in every graph), we can bound the approximation ratio in a way that is generic, i.e., that is essentially independent of the actual graph.

Now consider the  $\ell_\infty$ -norm. The generic approach seems to break down here: the universal upper bound is only  $\Theta(n)$  (as shown by the star graph), while the universal lower bound is only  $\Theta(1)$  (as shown by the path). So it seems like the generic guarantee is just the trivial  $\Theta(n)$ . But this is just because  $n$  is the wrong parameter in this setting: the correct parameterization is based on  $\Delta$ , the maximum degree of  $G$  (i.e.,  $\Delta = \|G\|_\infty$ ). With respect to  $\Delta$ , the greedy algorithm (or any algorithm) returns a spanner with maximum degree at most  $\Delta$ , while *any*  $t$ -spanner of a graph with maximum degree  $\Delta$  must have maximum degree at least  $\Omega(\Delta^{1/t})$  (assuming the graph is unweighted). So there is still a “generic” guarantee which implies that the greedy algorithm is an  $O(\Delta^{1-1/t}) \leq O(n^{1-1/t})$ -approximation.

This suggests that for  $1 < p < \infty$ , we will need to parameterize by *both* the number of nodes  $n$  and the  $\ell_p$ -norm  $\Lambda$  of  $G$ . We can define both universal upper bounds and universal lower bounds with respect to this dual parameterization:

$$\begin{aligned} \text{UB}_t^p(n, \Lambda) &= \max_{\substack{G=(V,E):|V|=n,\|G\|_p=\Lambda, \\ \text{and } G \text{ is connected}}} \min_{H: H \text{ is a } t\text{-spanner of } G} \|H\|_p \\ \text{LB}_t^p(n, \Lambda) &= \min_{\substack{G=(V,E):|V|=n,\|G\|_p=\Lambda, \\ \text{and } G \text{ is connected}}} \min_{H: H \text{ is a } t\text{-spanner of } G} \|H\|_p \end{aligned}$$

With this notation, we can define the generic guarantee  $g_t^p(n, \Lambda) = \text{UB}_t^p(n, \Lambda)/\text{LB}_t^p(n, \Lambda)$ , and if we want a guarantee purely in terms of  $n$  we can define the generic guarantee  $g_t^p(n) = \max_\Lambda g_t^p(n, \Lambda)$ . Our upper bound of Theorem 2 can then be restated as the claim that  $\text{UB}_{2k-1}^p(n, \Lambda) \leq \min \left\{ \Lambda, \max \left\{ O(n), O\left(n^{\frac{k+p}{kp}}\right) \right\} \right\}$  for all  $n, k, p, \Lambda$ . So in order to understand the generic guarantees  $g_{2k-1}^p(n, \Lambda)$  or  $g_{2k-1}^p(n)$ , we need to understand the universal lower bound quantity  $\text{LB}_{2k-1}^p(n, \Lambda)$ .

Surprisingly, unlike the  $\ell_1$  and  $\ell_\infty$  cases, the universal lower bound for other values of  $p$  is extremely complex. Understanding its value, and understanding the structure of the extremal graphs which match the bound given by  $\text{LB}_t^p(n, \Lambda)$ , are the most technically involved results in this paper. However, while the analysis and even the exact formulation of the lower bound is quite complex, it turns out to be easily computable from a simple linear program:

► **Theorem 3.** *There is an explicit linear program of size  $O(t)$  which calculates  $\text{LB}_t^p(n, \Lambda)$  for any  $t \in \mathbb{N}, p \geq 1$ . The bound given by the program is tight up to a factor of  $\log(n)^{O(t)}$ .*

Our linear program and the proof of Theorem 3 appear in the full version [12]. In fact, our linear program not only calculates a lower bound on the  $\ell_p$ -norm of any  $t$ -spanner, it also gives the parameters which define an extremal graph of  $\ell_p$ -norm  $\Lambda$  with a  $t$ -spanner whose  $\ell_p$ -norm matches this lower bound. While the structure of these extremal graphs is simple, the dependence of the parameters of these graphs on  $t$  and  $p$  is quite complex. Nevertheless, we give a complete explicit description of these graphs for every possible value of  $p, t$ .

Interestingly, despite the fact that  $\text{LB}_t^p(n, \Lambda)$  is fundamentally a question of extremal graph theory (although as discussed our motivation is the generic guarantee on approximation algorithms), our techniques are in some ways more related to approximation algorithms. We give a linear program which computes the LB function, and we reason about it by explicitly constructing dual solutions. This is, to the best of our knowledge, the first time that structural bounds on spanners (as opposed to approximation bounds) have been derived using linear programs. Moreover, the structure of the extremal graphs is fundamentally related to a quantity which we call the  *$p$ -log density* of the input graph. This is a generalization of the notion of “log-density”, which was introduced as the fundamental parameter when designing approximation algorithms for the DENSEST  $k$ -SUBGRAPH (DkS) problem [6], and has since proved useful in many approximation settings (see, e.g., [10, 11, 14, 13]).

### 1.1.3 Greedy Can Do Better Than The Generic Bound

As discussed, when  $p = 1$  or  $p = \infty$ , the approximation ratio of the greedy algorithm can be bounded by the generic guarantee. But it turns out that the connection is actually even closer: when  $p = 1$  and  $p = \infty$ , for every  $n$  and  $\Lambda$  the approximation ratio of the greedy algorithm is *equal* to the generic guarantee  $g_{2k-1}^p(n, \Lambda)$ . In other words, greedy is no better than generic in the traditional settings (we prove this for completeness, but it is essentially folklore). In fact, for the  $\ell_1$  objective, giving *any* approximation algorithm which is better than the generic guarantee  $g_{2k-1}^1(n)$  is a long-standing open problem [18] which has only been accomplished for stretch 3 [5] and stretch 4 [18], while for the  $\ell_\infty$  objective such an improvement was only shown recently [9] (and not with the greedy algorithm).

We show that, at least in some regimes of interest,  $\ell_p$ -norm graph spanners exhibit fundamentally different behavior from  $\ell_1$  and  $\ell_\infty$ : the greedy algorithm has approximation ratio which is *better* than the generic guarantee, even though the universal upper bound is proved via the greedy algorithm! In particular, we consider the regime of stretch 3,  $p = 2$ , and  $\Lambda = \Theta(n)$ . This is a very natural regime, since  $p = 2$  is the most obvious and widely-studied norm other than  $\ell_1$  and  $\ell_\infty$ , and stretch 3 is the smallest value for which nontrivial sparsification can occur.

Our theorems about UB and LB imply that  $g_3^2(n) = g_3^2(n, n) = \Theta(\sqrt{n})$ . But we show that in this setting (and in fact for any  $\Lambda$  as long as  $p = 2$  and the stretch is 3) the greedy algorithm is an  $O(n^{63/128})$ -approximation. Thus we show that, unlike  $\ell_1$  and  $\ell_\infty$ , for  $p = 2$  the greedy algorithm provides an approximation guarantee that is strictly better than the generic bound, both for specific values of  $\Lambda$  and when considering the worst case  $\Lambda$ .

## 1.2 Outline

We begin in Section 2 with some basic definitions and preliminaries. In order to illustrate the basic concepts in a simpler and more understandable setting, we then focus in Section 3 on the special case of stretch 3: we prove the stretch-3 version of Theorem 1 in Section 3.1, and then show that the greedy algorithm has approximation ratio better than the generic guarantee in Section 3.2. We then prove our upper and lower bounds in full generality: the upper bound (i.e., the proof of Theorem 2) in Section 4, and then our universal lower bound in Section 5. Due to space constraints, all missing proofs can be found in the appendices.

## 2 Definitions and Preliminaries

Let  $G = (V, E)$  be a graph, possibly with lengths on the edges. For any vertex  $u \in V$ , we let  $d(u)$  denote the degree of  $u$  and let  $N(u)$  denote the neighbors of  $u$ . We will also generalize this notation slightly by letting  $N_i(u)$  denote the set of vertices that are exactly  $i$  hops away from  $u$  (i.e., their distance from  $u$  if we ignore lengths is exactly  $i$ ), and we let  $d_i(u) = |N_i(u)|$ . Note that by definition,  $N_0(u) = \{u\}$  and  $d_0(u) = 1$  for all  $u \in V$ . We will sometimes use  $B(v, r) = \cup_{i=0}^r N_i(v)$  to denote the ball around  $v$  of radius  $r$ .

We let  $d_G : V \times V \rightarrow \mathbb{R}_{\geq 0}$  denote the shortest-path distances in  $G$ . A subgraph  $H = (V, E_H)$  of a graph  $G = (V, E)$  is a  $t$ -spanner of  $G$  if  $d_H(u, v) \leq t \cdot d_G(u, v)$  for all  $u, v \in E$ . Recall that  $\|\vec{x}\|_p = (\sum_{i=1}^n x_i^p)^{1/p}$  for any  $p \geq 1$  and  $\vec{x} \in \mathbb{R}^n$ . To measure the “cost” of a spanner, for any graph  $G = (V, E)$ , let  $\vec{d}_G$  denote the vector of degrees in  $G$  and for any  $p \geq 1$ , let  $\|G\|_p = \|\vec{d}_G\|_p$ . For any subset  $S \subseteq V$ , we let  $\|S\|_p$  denote the  $\ell_p$  norm of the vector obtained from  $\vec{d}_G$  by removing the coordinate of every node not in  $S$  (note that we do not remove the nodes from the graph, i.e.,  $\|S\|_p$  is the norm of the degrees in  $G$  of the nodes in  $S$ , not in the subgraph induced by  $S$ ).

### 3 Warmup: Stretch 3

We begin by analyzing the special case of stretch 3, particularly for the  $\ell_2$ -norm. More specifically, we will focus on bounding  $\text{UB}_3^p(n, \Lambda)$ . This is one of the simplest cases, but demonstrates (at a very high level) the outlines of our upper bound. Moreover, in this particular case we can prove that the greedy algorithm performs better than the generic guarantee, showing a fundamental difference between the  $\ell_2$  norm and the more traditional  $\ell_1$  and  $\ell_\infty$  norms.

#### 3.1 Upper Bound

Recall that *greedy spanner* is the spanner obtained from the obvious greedy algorithm: starting with an empty graph as the spanner, consider the edges one at a time in nondecreasing length order, and add an edge if the current spanner does not span it (within the given stretch requirement). It is obvious that when run with stretch parameter  $t$  this algorithm does indeed return a  $t$ -spanner, and moreover it will return a  $t$ -spanner that has girth at least  $t + 2$  (if there is a  $(t + 1)$ -cycle then the algorithm would not have added the final edge).

Our main goal in this section will be to prove the following theorem.

► **Theorem 4.** *Let  $G = (V, E)$  be a graph and let  $H = (V, E_H)$  be the greedy 3-spanner of  $G$ . Then  $\|H\|_p \leq \max(O(n), O(n^{(2+p)/(2p)}))$  for all  $p \geq 1$ .*

In other words, when  $1 \leq p \leq 2$  the greedy 3-spanner  $H$  has  $\|H\|_p \leq O(n^{(2+p)/(2p)})$ , and when  $p \geq 2$  we get that that  $\|H\|_p \leq O(n)$ .

To prove this theorem, we will first show that nodes with “large” degree cannot be incident with too many edges in any graph of girth at least 5 (like the greedy 3-spanner). This is the most important step, since for  $p > 1$  the  $p$ -norm of a graph gives greater “weight” to nodes with larger degree.

► **Lemma 5.** *Let  $G = (V, E)$  be a graph with girth at least 5. Then  $\sum_{v \in V: d(v) \geq 2\sqrt{n}} d(v) \leq 2n$ .*

**Proof.** Suppose for the sake of contradiction that these vertices have total degree greater than  $2n$ , and let  $\{v_1, \dots, v_{\ell+1}\}$  be a minimal set with this property. That is, all these vertices have degree at least  $2\sqrt{n}$ , and furthermore  $\sum_{i=1}^{\ell} d(v_i) \leq 2n < \sum_{i=1}^{\ell+1} d(v_i)$ .

Because  $G$  has girth at least 5, any two vertices  $v_i, v_j$  in this set have at most one common neighbor. That is,  $|N(v_i) \cap N(v_j)| \leq 1$ . Thus, for every  $j \in [\ell + 1]$ , the number of “new” neighbors contributed by  $N(v_j)$  is  $|N(v_j) \setminus (\bigcup_{i=1}^{j-1} N(v_i))| \geq |N(v_j)| - \sum_{i=1}^{j-1} |N(v_i) \cap N(v_j)| \geq d(v_j) - (j - 1) \geq d(v_j) - \ell$ .

On the other hand, we have  $2n \geq \sum_{i=1}^{\ell} d(v_i) \geq \ell \cdot 2\sqrt{n}$ , and so we have  $\ell \leq \sqrt{n}$ . Thus, every  $v_j$  contributes at least  $d(v_j) - \ell \geq d(v_j) - \sqrt{n} \geq d(v_j)/2$  new neighbors, and so we get  $n \geq \left| \bigcup_{j=1}^{\ell+1} N(v_j) \right| = \sum_{j=1}^{\ell+1} \left| N(v_j) \setminus \left( \bigcup_{i=1}^{j-1} N(v_i) \right) \right| \geq \sum_{j=1}^{\ell+1} d(v_j)/2$ , which contradicts our assumption that  $\sum_{j=1}^{\ell+1} d(v_j) > 2n$ . ◀

We can now prove Theorem 4.

**Proof of Theorem 4.** Let  $V_{low} = \{v \in V : d(v) \leq 2\sqrt{n}\}$ , and let  $V_{high} = \{v \in V : d(v) > 2\sqrt{n}\}$ . Since  $H$  has girth at least 5, we can apply Lemma 5. So using this lemma and



standard algebraic inequalities, we get that

$$\begin{aligned} \|H\|_p &= \left( \sum_{v \in V_{low}} d(v)^p + \sum_{v \in V_{high}} d(v)^p \right)^{\frac{1}{p}} \leq \left( \sum_{v \in V_{low}} d(v)^p \right)^{\frac{1}{p}} + \left( \sum_{v \in V_{high}} d(v)^p \right)^{\frac{1}{p}} \\ &\leq \left( \sum_{v \in V_{low}} (2\sqrt{n})^p \right)^{\frac{1}{p}} + \sum_{v \in V_{high}} d(v) \leq (n \cdot 2^p n^{p/2})^{\frac{1}{p}} + \sum_{v \in V_{high}} d(v) \leq 2n^{\frac{2+p}{2p}} + 2n, \end{aligned}$$

which implies the theorem.  $\blacktriangleleft$

It is easy to show that the above bound is tight: for every  $p \geq 1$  there are graphs in which every 3-spanner has size at least  $\max(\Omega(n), \Omega(n^{\frac{2+p}{2p}}))$ . In fact, we can generalize slightly to also account for different values of  $\Lambda$ . Theorem 2 can be interpreted as claiming that  $UB_3^p(n, \Lambda) \leq O(\min(\max(n, n^{\frac{2+p}{2p}}), \Lambda))$ . In the full version [12] we show that this is tight:  $UB_3^p(n, \Lambda) \geq \Omega(\min(\max(n, n^{\frac{2+p}{2p}}), \Lambda))$  for all  $p \geq 1$  and  $\Omega(n^{1/p}) \leq \Lambda \leq O(n^{\frac{1+p}{p}})$ .

### 3.2 Greedy vs Generic

It is not hard to show that in the traditional settings in which spanners have been studied, the  $\ell_1$  and  $\ell_\infty$  norms, the greedy algorithm does no better than the generic guarantee, for all relevant parameter regimes. In slightly more detail, for  $\ell_\infty$  it is relatively easy to show that  $UB_t^\infty(n, \Lambda) = \Theta(\Lambda)$ , while  $LB_t^\infty(n, \Lambda) = \Theta(\Lambda^{1/t})$ . Thus the generic guarantee  $g_t^\infty(n, \Lambda) = \Theta(\Lambda^{1-\frac{1}{t}})$ , and moreover we can build graphs in which the approximation ratio of the greedy algorithm is also  $\Theta(\Lambda^{1-\frac{1}{t}})$ . Similarly, for the  $\ell_1$ -norm, classical results on spanners imply that  $UB_{2k-1}^1(n, \Lambda) = \Theta(\min(n^{1+\frac{1}{k}}, \Lambda))$  and  $LB_{2k-1}^1(n, \Lambda) = \Theta(n)$ , so the generic guarantee is  $g_{2k-1}^1(n, \Lambda) = \Theta(\min(n^{1+\frac{1}{k}}, \Lambda)/n)$  and there are graphs for all parameter regimes where this is the approximation ratio achieved by greedy.

We show that the behavior of the greedy spanner in intermediate  $\ell_p$ -norms is fundamentally different: in some parameter regimes of interest, greedy *outperforms* the generic guarantee!

To demonstrate this, consider the regime of stretch 3 with the  $\ell_2$  norm and with  $\Lambda = n$  in unweighted graphs. In this regime, the results of Section 3.1 imply that  $UB_3^2(n, n) = \Theta(n)$ . On the other hand, our results on the universal lower bound from Section 5 (Corollary 19 in particular) directly imply that  $LB_3^2(n, n) = \tilde{\Theta}(\sqrt{n})$ . Thus the generic guarantee is  $g_3^2(n, n) = \tilde{\Theta}(\sqrt{n})$ , and this is the worst case over  $\Lambda$  and thus  $g_3^2(n) = \tilde{\Theta}(\sqrt{n})$ . However, we show that the greedy algorithm is a strictly better approximation, even without parameterizing by  $\Lambda$ .

► **Theorem 6.** *The greedy algorithm is an  $O(n^{63/128})$ -approximation for the problem of computing the 3-spanner with smallest  $\ell_2$ -norm (where the input is an unweighted graph).*

To prove this, let  $G = (V, E)$  be a graph with  $|V| = n$ , let  $H$  be the greedy 3-spanner of  $G$ , and let  $H^*$  be a 3-spanner of  $G$  with minimum  $\ell_2$ -norm. Let  $\alpha = \log_n \|H^*\|_2$ , so  $\|H^*\|_2 = n^\alpha$ ; note that  $\alpha \geq 1/2$ . We first prove a lemma which uses  $\|H^*\|_2$  to bound neighborhoods.

► **Lemma 7.**  $|B_{H^*}(v, r)| \leq n^{2\alpha(1-2^{-r})}$  for all  $v \in V$  and  $r \in \mathbb{N}$ .

**Proof.** We use induction on  $r$ . For the base case  $r = 1$ , since  $\|H^*\|_2 = n^\alpha$  we know that  $v$  has degree less than  $n^\alpha$ , and thus  $|B_{H^*}(v, 1)| \leq n^\alpha = n^{2\alpha(1-2^{-1})}$ .

Now suppose that the theorem is true for some integer  $r - 1$ . Let  $|B_{H^*}(v, r - 1)| = n^\gamma \leq n^{2\alpha(1-2^{-(r-1)})}$  (by induction). Since  $\|H^*\|_2 = n^\alpha$ , the average degree (in  $H^*$ ) of the nodes in  $B_{H^*}(v, r - 1)$  is at most  $n^{\alpha-(\gamma/2)}$ . Thus we get that  $|B_{H^*}(v, r)| \leq n^\gamma \cdot n^{\alpha-(\gamma/2)} = n^{\alpha+(\gamma/2)} \leq n^{\alpha+\alpha(1-2^{-(r-1)})} = n^{2\alpha(1-2^{-r})}$ , as claimed.  $\blacktriangleleft$

Using this lemma, we can now prove Theorem 6.

**Proof of Theorem 6.** Lemma 7 implies that  $|B_{H^*}(v, 6)| \leq n^{(63/32)\alpha}$  for all  $v \in V$ . Since  $H^*$  is a 3-spanner of  $G$ , every vertex in  $B_G(v, 2)$  must be in  $B_{H^*}(v, 6)$ , and thus  $|B_G(v, 2)| \leq n^{(63/32)\alpha}$ . Now we can use this to bound the number of 2-paths in  $H$ . Let  $P_2(H)$  denote the number of paths of length 2 in  $H$ . Since  $H$  is the greedy 3-spanner of  $G$  it must have girth at least 5. This means that every path of length 2 in  $H$  which starts from  $v$  must have a different other endpoint: there cannot be two different paths of the form  $v - w - u$  and  $v - x - u$  in  $H$ , or else  $H$  would have girth at most 4. Thus the number of 2-paths in  $H$  which start from  $v$  is bounded by  $|B_H(v, 2)| \leq |B_G(v, 2)| \leq n^{(63/32)\alpha}$ , and thus  $P_2(H) \leq n^{1+(63/32)\alpha}$ .

On the other hand, note that instead of counting 2-paths in  $H$  by their starting vertex, we could instead count them by their middle vertex. The number of 2-paths where  $v$  is the *middle* node is  $\binom{d_H(v)}{2} \geq d_H(v)^2/4$ , and thus  $P_2(H) \geq \sum_{v \in V} (d_H(v)^2/4) = \|H\|_2^2/4$ . Combining these two inequalities implies that  $\|H\|_2 \leq 4n^{\frac{1}{2} + \frac{63}{64}\alpha}$ , and hence the greedy spanner has approximation ratio of at most  $\frac{\|H\|_2}{\|H^*\|_2} \leq \frac{4n^{\frac{1}{2} + \frac{63}{64}\alpha}}{n^\alpha} = 4n^{\frac{1}{2} - \frac{1}{64}\alpha} \leq 4n^{\frac{1}{2} - \frac{1}{128}} = 4n^{63/128}$ . ◀

#### 4 Upper Bound: General Stretch

We now want to generalize the bounds from Section 3 to hold for larger stretch ( $2k - 1$  in particular) in order to prove Theorem 2. A natural approach would be an extension of the stretch 3 analysis: if in Lemma 5 we replaced the bound of  $2\sqrt{n}$  with  $2n^{1/k}$ , then the proof of Theorem 4 could easily be extended to prove Theorem 2. Unfortunately this is impossible: there are graphs of girth at least  $2k + 1$  where it is not true that the number of edges incident with nodes of degree at least  $2n^{1/k}$  is at most  $O(n)$ . This can be seen from, e.g., [25] for  $k = 3$ .

So we cannot just break the vertices into “high-degree” and “low-degree” as we did for stretch 3. Instead, our decomposition is more complicated. We will still have low-degree nodes, which can be analyzed trivially. But our definition of “high” will actually be parameterized by a distance  $j$ , and we will define a node to be “high-degree” at distance  $j$  if its degree is large relative to the expansion of its neighborhood at approximately distance  $j$ . We will also introduce a new type of “medium-degree” node. In Section 4.1 we define this decomposition and prove that it is a full decomposition of  $V$ , and then in Sections 4.2 and 4.3 we show that no part in this decomposition can contribute too much to the overall cost.

First, though, we make one simple observation that will allow us to simplify notation by only considering one particular value of  $p$ . While we could analyze general values of  $p$  as we did for stretch 3 in Section 3.1, it is actually sufficient to prove the bound for the special case of  $k$  and  $p$  where the two terms in the maximum are equal, i.e., when  $\frac{k+p}{kp} = 1$ .

► **Lemma 8.** *Let  $k \geq 1$  be an integer, let  $G = (V, E)$  be a graph, and let  $H = (V, E_H)$  be the greedy  $(2k - 1)$ -spanner of  $G$ . If  $\|H\|_{p'} = O(n)$  for  $p' = k/(k - 1)$  then  $\|H\|_p \leq \max\left(O(n), O\left(n^{\frac{k+p}{kp}}\right)\right)$  for all  $p \geq 1$ .*

**Proof.** First note that  $p' = k/(k - 1)$  if and only if  $\frac{k+p'}{kp'} = 1$ . So we break into two cases, one for  $p > p'$  and one for  $1 \leq p < p'$ . For the first case, where  $p > p'$ , the result follows simply because of the monotonicity of  $p$ -norms:  $\|H\|_p \leq \|H\|_{p'} = O(n) = \max(O(n), O(n^{\frac{k+p}{kp}}))$ .

For the second case, where  $1 \leq p < p'$ , let  $q$  be the value such that  $1 \leq p \leq p'$  and  $\frac{1}{p'} + \frac{1}{q} = \frac{1}{p}$ . Recall that  $\vec{d}_H$  is the degree vector of  $H$ . Then Hölder’s inequality implies that  $\|\vec{d}_H\|_p \leq \|\vec{d}_H\|_{p'} \|1\|_q = n^{\frac{1}{p} - \frac{1}{p'}} \|\vec{d}_H\|_{p'}$ . Since by assumption we have  $\|\vec{d}_H\|_{p'} \leq O(n)$ , this implies that  $\|H\|_p \leq O(n^{1 + \frac{1}{p} - \frac{1}{p'}}) = O(n^{\frac{1}{p} - \frac{k-1}{k} + 1}) = O(n^{\frac{k+p}{kp}})$ , as claimed. ◀



## 4.1 Graph Decomposition

Recall that  $d_i(v)$  denotes the number of vertices at distance exactly  $i$  from  $v$ . This will let us define the following vertex sets.

► **Definition 9.** *Let  $G = (V, E)$  be a graph of girth at least  $2k + 1$ , with  $k \geq 3$ . Then define*

$$\begin{aligned} V_{low} &= \{v \in V : d_1(v) \leq n^{1/k}\} \\ V_{med} &= \{v \in V : n^{(k-2)/(k-1)} d_1(v)^{1/(k-1)} \leq d_{k-1}(v)\} \\ V_{high,j} &= \{v \in V : d_{k-2j-1}(v) \leq n^{1/(k-1)} d_{k-2j-3}(v) d_1(v)^{(k-2)/(k-1)}\}, \end{aligned}$$

where  $0 \leq j \leq \lfloor (k-3)/2 \rfloor$ .

It is not hard to see that this notion of high still corresponds to a deviation from regularity, as in the stretch 3 setting; the difference is that this deviation is relative to the size of the neighborhood at distance  $k-2j-1$  vs the neighborhood at distance  $k-2j-3$ .

As we will see in Sections 4.2 and 4.3, analyzing the contribution of  $V_{high,j}$  to the  $p$ -norm of the greedy spanner is in some sense the “main” technical step: analyzing  $V_{low}$  is straightforward, and analyzing  $V_{med}$ , while nontrivial, turns out to be easier than the case for  $V_{high,j}$ . Before we do this, though, we will show that we have a full decomposition of  $V$ :

► **Theorem 10.** *Let  $G = (V, E)$  be a graph of girth at least  $2k + 1$ , with  $k \geq 3$ . Then  $V = V_{low} \cup V_{med} \cup (\cup_{0 \leq j \leq \lfloor (k-3)/2 \rfloor} V_{high,j})$ .*

**Proof.** We prove the case when  $k$  is odd. The even case is similar. Assume that  $v \notin \cup_{0 \leq j \leq \lfloor (k-3)/2 \rfloor} V_{high,j}$ . Then by the definition of  $V_{high,j}$ , we know that  $d_{k-2j-1}(v) > n^{1/(k-1)} d_{k-2j-3}(v) d_1(v)^{(k-2)/(k-1)}$  for all  $j$ . Then a straightforward induction on  $j$  implies that

$$d_{k-1}(v) > n^{1/2} d_1(v)^{(k-2)/2}. \quad (1)$$

If further we assume that  $v \notin V_{low}$ , then  $d_1(v) > n^{1/k}$ , and thus

$$(d_1(v))^{k(k-3)/(2(k-1))} > n^{(k-3)/(2(k-1))}. \quad (2)$$

Finally, assuming that  $v \notin V_{med}$  implies that

$$n^{(k-2)/(k-1)} (d_1(v))^{1/(k-1)} > d_{k-1}(v). \quad (3)$$

If we then multiply inequalities (1), (2) and (3), after some elementary algebra we find that  $1 > 1$ , a contradiction. Thus  $v \in V_{low} \cup V_{med} \cup (\cup_{0 \leq j \leq \lfloor (k-3)/2 \rfloor} V_{high,j})$ . ◀

## 4.2 Structural Lemmas for High-Girth Graphs

With Theorem 10 in hand, it remains to bound the contribution to the  $p$ -norm of the spanner of these different vertex sets. In order to do this, we start with a few useful lemmas (proofs can be found in the full version [12]). We first give a simple lemma: if the girth is large enough, then the neighborhoods around a node can be bounded by the neighborhoods around its neighbors.

► **Lemma 11.** *Let  $G = (V, E)$  have girth at least  $2k+1$  with  $k \geq 2$ . Then  $\sum_{w \in N_1(v)} d_{k-1}(w) \leq d_k(v) + d_1(v) d_{k-2}(v)$  for all  $v \in V$ .*

## 40:10 The Norms of Graph Spanners

With this lemma in hand, we will now prove a more complicated technical lemma which will likewise hold for all high-girth graphs. For a given  $v, w$  with  $v \in N(w)$ , we can consider the fraction of the  $k$ -neighborhood of  $w$  which is also contained in the  $(k-1)$ -neighborhood of  $v$ . Then if we sum this fraction over all neighbors  $v$  of  $w$ , we would of course get 1 since the girth constraint would imply that any two neighbors of  $v$  cannot both be first hops on paths to the same node in  $N_k(w)$ . But what if we consider the slightly different ratio of  $d_{k-1}(v)/d_k(w)$ ? This is notably different since it includes in the numerator not just  $N_{k-1}(v) \cap N_k(w)$ , but also  $N_{k-1}(v) \cap N_{k-2}(w)$ . It will prove useful for us to reason about these values, so we show that “on average” they behave approximately the same: if we sum over the neighbors of any given node then these fractions can add up to something quite large (not 1), but overall they only add up to  $O(n)$ .

► **Lemma 12.** *Let  $k \geq 1$  be an integer, and let  $G = (V, E)$  have girth at least  $2k + 1$  and minimum degree at least 4. Then  $\sum_{w \in V} \sum_{v \in N(w)} \frac{d_{k-1}(v)}{d_k(w)} \leq 2n$ .*

The proof of this lemma is quite technical, but is done with an induction on  $k$  and careful use of the arithmetic-harmonic mean inequality. While Lemma 12 is the main structural result that we will use to bound the “high” degree nodes, the following corollary makes it slightly simpler to use.

► **Corollary 13.** *Let  $k \geq 2$  be an integer, and let  $G = (V, E)$  have girth at least  $2k + 1$  and minimum degree at least 4. Then  $\sum_{v \in V} \frac{(d_1(v))^2 d_{k-2}(v)}{d_k(v) + d_1(v) d_{k-2}(v)} \leq 2n$ .*

### 4.3 Proving Theorem 2

We can now finally prove Theorem 2 by analyzing the contribution of the different sets in the decomposition to any graph of girth at least  $2k + 1$  (in particular, the greedy  $(2k-1)$ -spanner). All missing proofs can be found in the full version [12].

The analysis of the low nodes is straightforward, while the analysis of the medium nodes is slightly more complex. But the main difficulty is in the high nodes.

► **Lemma 14.** *Let  $k \geq 2$  be an integer, let  $p = \frac{k}{k-1}$ , and let  $G = (V, E)$  have girth at least  $2k + 1$ . Then  $\|V_{low}\|_p \leq n$ .*

We next bound the medium nodes.

► **Lemma 15.** *Let  $k \geq 2$  be an integer, let  $p = \frac{k}{k-1}$ , and let  $G = (V, E)$  have girth at least  $2k + 1$ . Then  $\|V_{med}\|_p \leq n$ .*

We now bound the high nodes, with one degree assumption which we will later remove.

► **Lemma 16.** *Let  $G = (V, E)$  be a graph of girth at least  $2k + 1$  with  $k \geq 3$ . Further assume that the graph has minimum degree at least 4. Then  $\|V_{high,j}\|_{k/(k-1)} = O(n)$  for all  $0 \leq j \leq \lfloor (k-3)/2 \rfloor$ .*

**Proof.** We will break the high nodes into the following two sets:

$$\begin{aligned} V'_{high,j} &= \{v \in V_{high,j} : d_{k-2j-1}(v) \geq d_{k-2j-3}(v)d_1(v)\} \\ V''_{high,j} &= \{v \in V_{high,j} : d_{k-2j-1}(v) < d_{k-2j-3}(v)d_1(v)\}. \end{aligned}$$

Obviously  $V_{high,j} = V'_{high,j} \cup V''_{high,j}$ , so we can bound each of the two sets separately. For the first set, we get that

$$\begin{aligned} \|V'_{high,j}\|_{\frac{k}{k-1}} &= \left( \sum_{v \in V'_{high,j}} (d_1(v))^{\frac{k}{k-1}} \right)^{\frac{k-1}{k}} \leq \left( \sum_{v \in V'_{high,j}} \frac{n^{\frac{1}{k-1}} d_1(v)^2 d_{k-2j-3}(v)}{d_{k-2j-1}(v)} \right)^{\frac{k-1}{k}} \\ &\leq \left( 2 \sum_{v \in V'_{high,j}} \frac{n^{\frac{1}{k-1}} d_1(v)^2 d_{k-2j-3}(v)}{d_{k-2j-1}(v) + d_1(v) d_{k-2j-3}(v)} \right)^{\frac{k-1}{k}} \leq 4n. \end{aligned}$$

The first inequality is from the definition of  $V_{high}$ , the second is from the definition of  $V'_{high}$ , and the final inequality is from Corollary 13.

To analyze  $V''_{high}$ , note that  $d_{k-2j-1}(v) + d_1(v) d_{k-2j-3}(v) < 2d_1(v) d_{k-2j-3}(v)$  by definition for all  $v \in V''_{high,j}$ . Combining this with Corollary 13 implies that

$$\begin{aligned} \|V''_{high,j}\|_{k/(k-1)} &\leq \|V''_{high,j}\|_1 = \sum_{v \in V''_{high,j}} d_1(v) = \sum_{v \in V''_{high,j}} \frac{d_1(v)^2 d_{k-2j-3}(v)}{d_1(v) d_{k-2j-3}(v)} \\ &\leq 2 \sum_{v \in V''_{high,j}} \frac{d_1(v)^2 d_{k-2j-3}(v)}{d_{k-2j-1}(v) + d_1(v) d_{k-2j-3}(v)} \leq 4n. \end{aligned}$$

Thus  $\|V_{high,j}\|_{k/(k-1)} \leq \|V'_{high,j}\|_{k/(k-1)} + \|V''_{high,j}\|_{k/(k-1)} \leq 8n$ . ◀

Putting this all together gives the following theorem.

► **Theorem 17.** *Let  $G = (V, E)$  have girth at least  $2k + 1$ ,  $k \geq 2$  and minimum degree at least 4. Then  $\|G\|_p \leq O(kn)$  for  $p = \frac{k}{k-1}$ .*

**Proof.** We know from Theorem 10 that  $V = V_{low} \cup V_{med} \cup (\cup_{0 \leq j \leq \lfloor (k-3)/2 \rfloor} V_{high,j})$  for  $k \geq 3$ . Thus  $\|G\|_p \leq \|V_{low}\|_p + \|V_{med}\|_p + \sum_{j=0}^{\lfloor (k-3)/2 \rfloor} \|V_{high,j}\|_p \leq O(kn)$ , where we used Lemmas 14, 15, 16, to bound the contribution of each set. If  $k = 2$  then  $V_{med} = V$  and the proof is similar (alternatively see Theorem 4). ◀

We can now remove the degree assumption and the restriction to  $p = \frac{k}{k-1}$ , to finally prove Theorem 2.

**Proof of Theorem 2.** The case of  $k = 1$  is trivial since every graph  $H$  has  $\|H\|_p \leq O(n^{\frac{p+1}{p}})$ . For  $k \geq 2$ , by Lemma 8, we may assume that  $p = \frac{k}{k-1}$ . We will use induction on the number of vertices of degree less than 4. If  $H$  has no vertices with degree less than 4, then Theorem 17 implies Theorem 2. Otherwise, let  $v \in V$  be a vertex of degree at most 3, and let  $G' = G - v$  be the graph obtained by removing  $v$ . Then it is easy to see that  $\|\vec{d}_G - \vec{d}_{G'}\|_1 \leq 6$ , since one entry in the degree vector of value at most 3 gets removed and at most three other entries get decreased by 1. Thus we can use triangle inequality and monotonicity of norms to get that  $\|G\|_p - \|G'\|_p \leq \|\vec{d}_G - \vec{d}_{G'}\|_p \leq \|\vec{d}_G - \vec{d}_{G'}\|_1 \leq 6$ . Hence by the induction hypothesis we get that  $\|G\|_p \leq O(kn)$  as required. ◀

## 5 Universal Lower Bound

As stated in Theorem 3, our lower bound can be calculated by a simple linear program of size  $O(t)$  (where  $t$  is the stretch). We give this linear program formally in the full version [12]. The linear program assumes that the graph has a fairly regular structure. In particular, it

## 40:12 The Norms of Graph Spanners

assumes that the extremal  $t$ -spanner  $H$  is a layered graph with  $t + 1$  layers  $V_0, \dots, V_t$ , such that the subgraph induced on every two subsequent layers  $V_i, V_{i+1}$  is bipartite and biregular (in each side, all vertices have the same degree), and that the original extremal graph  $G$  (the graph whose spanner  $H$  achieves the lower bound) in addition has a biregular graph between  $V_0$  and  $V_t$  which contributes most of the  $p$ -norm of  $G$ , and is spanned by the layered graph  $H$ . Such a spanner  $H$  can be briefly described by the cardinalities of the layers  $V_i$  and the degrees of the bipartite graphs connecting every two consecutive layers.

As we show, this assumption is without loss of generality, in the sense that pruning any graph to obtain this structure can change the  $p$ -norm of the graph or its spanner by at most a polylogarithmic factor. The linear program captures the constraints that the parameters of a spanner with such a regular structure must satisfy. These constraints are also sufficient in the sense that given any solution to the linear program, we can construct a graph  $G$  and spanner  $H$  of this form with the parameters given by this LP solution.

In fact, the extremal spanners which match our lower bound have a fairly specific structure with consistent properties:

- The layers in the extremal can be partitioned into three sections: an initial section in which we have layers of decreasing size  $|V_0| \geq |V_1| \dots \geq |V_L|$ , a middle section consisting of equal size layers  $|V_L| = \dots = |V_{L+C}|$ , and a final section with layers of increasing size  $|V_{L+C}| \leq \dots \leq |V_{L+C+R}|$ . In some cases one of the first two sections may be missing.
- The bipartite graphs between every two consecutive layers in the spanner have the same contribution to the  $p$ -norm of the spanner.
- In addition to the edges in the spanner, the original graph also contains a biclique between the outer layers  $V_0$  and  $V_t$ , so that  $\|G\|_p = \Theta(|V_0|^{1/p}|V_t|)$ .

The structure of these spanners has the property that given the lengths of the three sections, we can derive the exact structure of the spanner, and hence the exact value of the lower bound. In our analysis, we focus on this specific family of graphs, and show that it suffices to describe our lower bound.

While the lower bound for  $p = 1$  or  $p = \infty$  is simple, it turns out that the lower bound for intermediate values of  $p$  is quite complex, and depends on the stretch  $t$ , the norm parameter  $p$ , and the  $p$ -norm of the input graph  $\Lambda$  in a highly non-trivial way. To identify the extremal spanners and prove their optimality, we look at the dual of our linear program, and for every graph in our family of candidate extremal spanners, examine whether there exists a dual solution which satisfies complementary slackness w.r.t. the primal LP solution corresponding to our spanner. With this approach, for every  $p, t, \Lambda$ , we are able to identify the exact constraints that the parameters of an optimal spanner from our family must satisfy, and give an explicit solution, which gives our lower bound.

As an example, our analysis identifies the lower bound for relatively low values of  $p$ .

► **Theorem 18.** *If  $t$  is even, then for all  $p \in [1, \varphi]$  (where  $\varphi = \frac{1+\sqrt{5}}{2}$  is the golden ratio),*

$$\text{LB}_t^p(n, \Lambda) = \tilde{\Theta} \left( \max \left\{ n^{1/p}, \Lambda^\alpha \right\} \right) \text{ for } \alpha = 1 / \left( (p+1) \left( 1 - ((p-1)/p)^{t/2} \right) \right).$$

*If  $t$  is odd, then for all  $p \in [1, 2]$ ,*

$$\text{LB}_t^p(n, \Lambda) = \tilde{\Theta} \left( \max \left\{ n^{1/p}, \Lambda^\beta \right\} \right) \text{ for } \beta = 1 / \left( 1 + p \left( 1 - ((p-1)/p)^{(t-1)/2} \right) \right).$$

► **Corollary 19.** *For all  $p \in [1, \varphi]$ , we have  $\text{LB}_2^p(n, \Lambda) = \tilde{\Theta} \left( \max \left\{ n^{1/p}, \Lambda^{p/(p+1)} \right\} \right)$ . For all  $p \in [1, 2]$ , we have  $\text{LB}_3^p(n, \Lambda) = \tilde{\Theta} \left( \max \left\{ n^{1/p}, \sqrt{\Lambda} \right\} \right)$ .*

Note that the dependence on  $n$  for this range of parameters is minimal. In fact, the only dependence on  $n$  is due to the fact that any connected  $n$ -vertex graph (such as the spanner of a connected  $n$ -vertex graph) must have  $p$ -norm at least  $n^{1/p}$ . If we remove the condition that the graph must be connected, the lower bounds in Theorem 18 become  $\tilde{\Theta}(\Lambda^\alpha)$  and  $\tilde{\Theta}(\Lambda^\beta)$ .

For higher values of  $p$ , the lower bound becomes more complex. In particular, the parameters which determine the extremal spanner depend not only on  $p$  and  $t$ , but also on the  $p$ -log density of the graph, which we define to be  $\log_n(\Lambda) = \log_n(\|G\|_p)$ . This parameter generalizes the notion of log-density, which is at the heart of several recent breakthroughs in approximation algorithms [6, 10, 11, 14, 13], in which log-density was used to mean  $p$ -log density for  $p = 1$  or  $p = \infty$ . As in that line of work, the structure and parameters of the graphs of interest here (the extremal spanners) is a function of the  $p$ -log density of our graph which does not depend on  $n$ . The complete technical details of our lower bound appear in the full version of the paper [12].

## 6 Future Work

In this paper we have initiated the study of graph spanners with cost defined by the  $\ell_p$ -norm of the degree vector, since this provides an interesting interpolation between the  $\ell_1$ -norm (only caring about the number of edges) and the  $\ell_\infty$ -norm (only caring about the maximum degree). But we have only scratched the surface: many of the hundreds of results on graph spanners can be extended or reexamined with respect to the  $\ell_p$ -norm. There are also some very interesting direct extensions of this paper that would be interesting to study. In particular, we showed that the approximation ratio achieved by the greedy algorithm is strictly better than the generic guarantee for the  $\ell_2$ -norm with stretch 3, unlike the  $\ell_1$  and  $\ell_\infty$  norms. This suggests further study of the greedy algorithm in general, but also suggests extending the recent line of work on approximation algorithms for graph spanners (mostly using convex relaxations and rounding) to general  $\ell_p$ -norms. The approaches taken for the  $\ell_1$ -norm in the past [16, 17, 5, 18] have been quite different from the approaches used for the  $\ell_\infty$ -norm [23, 10, 9]; is there a way of interpolating between them to get even better approximations for intermediate  $\ell_p$ -norms?

---

## References

- 1 Noga Alon, Yossi Azar, Gerhard J. Woeginger, and Tal Yadid. Approximation Schemes for Scheduling. In *Proceedings of the Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '97, 1997.
- 2 Ingo Althöfer, Gautam Das, David Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discrete Comput. Geom.*, 9(1):81–100, 1993. doi:10.1007/BF02189308.
- 3 Yossi Azar, Leah Epstein, Yossi Richter, and Gerhard J. Woeginger. All-Norm Approximation Algorithms. In Martti Penttonen and Erik Meineche Schmidt, editors, *Algorithm Theory – SWAT 2002*, 2002.
- 4 Nikhil Bansal and Kirk Pruhs. Server Scheduling in the  $L_p$  Norm: A Rising Tide Lifts All Boat. In *Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing*, STOC '03, pages 242–250, 2003.
- 5 Piotr Berman, Arnab Bhattacharyya, Konstantin Makarychev, Sofya Raskhodnikova, and Grigory Yaroslavtsev. Approximation algorithms for spanner problems and Directed Steiner Forest. *Inf. Comput.*, 222:93–107, 2013. doi:10.1016/j.ic.2012.10.007.

- 6 Aditya Bhaskara, Moses Charikar, Eden Chlamtac, Uriel Feige, and Aravindan Vijayaraghavan. Detecting high log-densities: an  $O(n^{1/4})$ -approximation for densest  $k$ -subgraph. In Leonard J. Schulman, editor, *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 201–210. ACM, 2010. doi:10.1145/1806689.1806719.
- 7 Arnab Bhattacharyya, Elena Grigorescu, Kyomin Jung, Sofya Raskhodnikova, and David P. Woodruff. Transitive-closure Spanners. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '09*, pages 932–941, Philadelphia, PA, USA, 2009. Society for Industrial and Applied Mathematics. URL: <http://dl.acm.org/citation.cfm?id=1496770.1496871>.
- 8 Barun Chandra, Gautam Das, Giri Narasimhan, and José Soares. New Sparseness Results on Graph Spanners. In *Proceedings of the Eighth Annual Symposium on Computational Geometry, SCG '92*, pages 192–201, New York, NY, USA, 1992. ACM. doi:10.1145/142675.142717.
- 9 Eden Chlamtác and Michael Dinitz. Lowest Degree  $k$ -Spanner: Approximation and Hardness. In Klaus Jansen, José D. P. Rolim, Nikhil R. Devanur, and Cristopher Moore, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2014)*, volume 28 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 80–95, Dagstuhl, Germany, 2014. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.APPROX-RANDOM.2014.80.
- 10 Eden Chlamtac, Michael Dinitz, and Robert Krauthgamer. Everywhere-Sparse Spanners via Dense Subgraphs. In *Proceedings of the 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science, FOCS '12*, pages 758–767, Washington, DC, USA, 2012. IEEE Computer Society. doi:10.1109/FOCS.2012.61.
- 11 Eden Chlamtác, Michael Dinitz, and Yury Makarychev. Minimizing the Union: Tight Approximations for Small Set Bipartite Vertex Expansion. In Philip N. Klein, editor, *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 881–899. SIAM, 2017. doi:10.1137/1.9781611974782.56.
- 12 Eden Chlamtác, Michael Dinitz, and Thomas Robinson. The Norms of Graph Spanners. *CoRR*, abs/1903.07418, 2019. arXiv:1903.07418.
- 13 Eden Chlamtác and Pasin Manurangsi. Sherali-Adams Integrality Gaps Matching the Log-Density Threshold. In Eric Blais, Klaus Jansen, José D. P. Rolim, and David Steurer, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2018, August 20-22, 2018 - Princeton, NJ, USA*, volume 116 of *LIPIcs*, pages 10:1–10:19. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018. doi:10.4230/LIPIcs.APPROX-RANDOM.2018.10.
- 14 Eden Chlamtác, Pasin Manurangsi, Dana Moshkovitz, and Aravindan Vijayaraghavan. Approximation Algorithms for Label Cover and The Log-Density Threshold. In Philip N. Klein, editor, *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 900–919. SIAM, 2017. doi:10.1137/1.9781611974782.57.
- 15 Michael Dinitz, Guy Kortsarz, and Ran Raz. Label Cover Instances with Large Girth and the Hardness of Approximating Basic  $k$ -Spanner. *ACM Trans. Algorithms*, 12(2):25:1–25:16, December 2015. doi:10.1145/2818375.
- 16 Michael Dinitz and Robert Krauthgamer. Directed Spanners via Flow-based Linear Programs. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing, STOC '11*, pages 323–332, New York, NY, USA, 2011. ACM. doi:10.1145/1993636.1993680.
- 17 Michael Dinitz and Robert Krauthgamer. Fault-tolerant Spanners: Better and Simpler. In *Proceedings of the 30th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing, PODC '11*, pages 169–178, New York, NY, USA, 2011. ACM. doi:10.1145/1993806.1993830.



- 18 Michael Dinitz and Zeyu Zhang. Approximating Low-stretch Spanners. In *Proceedings of the Twenty-seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '16, pages 821–840, Philadelphia, PA, USA, 2016. Society for Industrial and Applied Mathematics. URL: <http://dl.acm.org/citation.cfm?id=2884435.2884494>.
- 19 Paul Erdős. *Extremal problems in graph theory*, pages 29–36. Academia Praha, Czechoslovakia, 1964.
- 20 Arnold Filtser and Shay Solomon. The Greedy Spanner is Existentially Optimal. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*, PODC '16, pages 9–17, 2016.
- 21 Daniel Golovin, Anupam Gupta, Amit Kumar, and Kanat Tangwongsan. All-Norms and All- $L_p$ -Norms Approximation Algorithms. In *FSTTCS*, volume 2 of *LIPICs*, pages 199–210. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2008.
- 22 Tapas Kanungo, David M. Mount, Nathan S. Netanyahu, Christine D. Piatko, Ruth Silverman, and Angela Y. Wu. A local search approximation algorithm for k-means clustering. *Computational Geometry*, 28(2):89–112, 2004. Special Issue on the 18th Annual Symposium on Computational Geometry - SoCG2002. doi:10.1016/j.comgeo.2004.03.003.
- 23 Guy Kortsarz and David Peleg. Generating Low-Degree 2-Spanners. *SIAM J. Comput.*, 27(5):1438–1456, 1998. doi:10.1137/S0097539794268753.
- 24 S. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, March 1982. doi:10.1109/TIT.1982.1056489.
- 25 Stefan Neuwirth. The size of bipartite graphs with girth eight. *arXiv Mathematics e-prints*, page math/0102210, February 2001. arXiv:math/0102210.
- 26 David Peleg and Alejandro A. Schäffer. Graph spanners. *Journal of Graph Theory*, 13(1):99–116, 1989. doi:10.1002/jgt.3190130114.
- 27 David Peleg and Jeffrey D. Ullman. An optimal synchronizer for the hypercube. *SIAM J. Comput.*, 18:740–747, August 1989. doi:10.1137/0218050.
- 28 Mikkel Thorup and Uri Zwick. Compact Routing Schemes. In *Proceedings of the Thirteenth Annual ACM Symposium on Parallel Algorithms and Architectures*, SPAA '01, pages 1–10, 2001.