

Lower Bounds on Sparse Spanners, Emulators, and Diameter-reducing shortcuts*

Shang-En Huang

University of Michigan, USA
sehuang@umich.edu

Seth Pettie

University of Michigan, USA
pettie@umich.edu

Abstract

We prove better lower bounds on additive spanners and emulators, which are lossy compression schemes for *undirected* graphs, as well as lower bounds on *shortcut sets*, which reduce the diameter of *directed* graphs. We show that any $O(n)$ -size shortcut set cannot bring the diameter below $\Omega(n^{1/6})$, and that any $O(m)$ -size shortcut set cannot bring it below $\Omega(n^{1/11})$. These improve Hesse’s [16] lower bound of $\Omega(n^{1/17})$. By combining these constructions with Abboud and Bodwin’s [1] edge-splitting technique, we get additive stretch lower bounds of $+\Omega(n^{1/13})$ for $O(n)$ -size spanners and $+\Omega(n^{1/18})$ for $O(n)$ -size emulators. These improve Abboud and Bodwin’s $+\Omega(n^{1/22})$ lower bounds.

2012 ACM Subject Classification Theory of computation → Sparsification and spanners

Keywords and phrases additive spanners, emulators, shortcutting directed graphs

Digital Object Identifier 10.4230/LIPIcs.SWAT.2018.26

Acknowledgements Thanks to the reviewers who point out the detailed comparison between spanners and emulators. We also thank the reviewers who carefully and explicitly verified our constructions, and give us some insightful comments.

1 Introduction

A *spanner* of an undirected unweighted graph $G = (V, E)$ is a subgraph H that approximates the distance function of G up to some *stretch*. An *emulator* for G is defined similarly, except that H need not be a subgraph, and may contain *weighted* edges. In this paper we consider only *additive* stretch functions:

$$\text{dist}_G(u, v) \leq \text{dist}_H(u, v) \leq \text{dist}_G(u, v) + \beta,$$

where β may depend on n .

Graph compression schemes (like spanners and emulators) are related to the problem of *shortcutting* digraphs to reduce diameter, inasmuch as lower bounds for both objects are constructed using the same suite of techniques. These lower bounds begin from the construction of graphs in which numerous pairs of vertices have shortest paths that are *unique*, *edge-disjoint*, and relatively *long*. Such graphs were independently discovered by Alon [4], Hesse [16], and Coppersmith and Elkin [12]; see also [1, 2]. Given such a “base graph,” derived graphs can be obtained through a variety of graph products such as the

* This work was supported by NSF grants CCF-1514383 and CCF-1637546.



■ **Table 1** Upper and Lower bounds on shortcutting sets. The lower bounds are existential, and independent of computation time.

Citation	Shortcut Set Size	Diameter	Computation Time
Folklore/trivial	$O(n)$	$\tilde{O}(\sqrt{n})$	$O(m\sqrt{n})$
	$O(m)$	$\tilde{O}(n/\sqrt{m})$	$O(m^{3/2})$
Fineman [15]	$\tilde{O}(n)$	$\tilde{O}(n^{2/3})$	$\tilde{O}(m)$
Hesse [16]	$O(mn^{1/17})$	$\Omega(n^{1/17})$	—
new	$O(n)$	$\Omega(n^{1/6})$	—
	$O(m)$	$\Omega(n^{1/11})$	—

alternation product discovered independently by Hesse [16] and Abboud and Bodwin [1] and the *substitution* product used by Abboud and Bodwin [1] and developed further by Abboud, Bodwin, and Pettie [2].

In this paper we apply the techniques developed in [4, 16, 12, 1, 2] to obtain better lower bounds on shortcutting sets, additive spanners, and additive emulators.

Shortcutting Sets

Let $G = (V, E)$ be a directed graph and $G^* = (V, E^*)$ its transitive closure. The *diameter* of a digraph G is the maximum of $\text{dist}_G(u, v)$ over all pairs $(u, v) \in E^*$. Thorup [20] conjectured that it is possible to reduce the diameter of any digraph to $\text{poly}(\log n)$ by adding a set $E' \subseteq E^*$ of at most $m = |E|$ *shortcuts*, i.e., $G' = (V, E \cup E')$ would have diameter $\text{poly}(\log n)$. This conjecture was confirmed for a couple special graph classes [20, 21], but refuted in general by Hesse [16], who exhibited a graph with $m = \Theta(n^{19/17})$ edges and diameter $\Theta(n^{1/17})$ such that any diameter-reducing shortcutting requires $\Omega(mn^{1/17})$ shortcuts. More generally, there exist graphs with $m = n^{1+\epsilon}$ edges and diameter n^δ , $\delta = \delta(\epsilon)$, that require $\Omega(n^{2-\epsilon})$ shortcuts to make the diameter $o(n^\delta)$; see Abboud, Bodwin, and Pettie [2, §6] for an alternative proof of this result.

On the upper bound side, it is trivial to reduce the diameter to $\tilde{O}(\sqrt{n})$ with $O(n)$ shortcuts or diameter $\tilde{O}(n/\sqrt{m})$ with $O(m)$ shortcuts.¹ Unfortunately, the trivial shortcutting schemes are not efficiently constructible in near-linear time. In some applications of shortcuttings, efficiency of the construction is just as important as reducing the diameter. For example, a longstanding problem in parallel computing is to *simultaneously* achieve time and work efficiency in computing reachability.² Very recently, Fineman [15] proved that an $\tilde{O}(n)$ -size shortcut set can be computed in near-optimal work $\tilde{O}(m)$ (and $\tilde{O}(n^{2/3})$ parallel time) that reduces the diameter to $\tilde{O}(n^{2/3})$.

In this paper we prove that $O(n)$ -size shortcut sets cannot reduce the diameter below $\Omega(n^{1/6})$, and that $O(m)$ -size shortcut sets cannot reduce it below $\Omega(n^{1/11})$. See Table 1.

Additive Spanners

Additive spanners with constant stretches were discovered by Aingworth, Checkuri, Indyk, and Motwani [3] (see also [13, 14, 5, 18]), Chechik [11], and Baswana, Kavitha, Mehlhorn, and Pettie [5] (see also [23, 18]). The sparsest of these [5] has size $O(n^{4/3})$ and stretch

¹ Pick a set S of \sqrt{n} or \sqrt{m} vertices uniformly at random, and include $S^2 \cap E^*$ as shortcuts.

² This is the notorious *transitive closure bottleneck*.

■ **Table 2** Upper and lower bounds on additive spanners.

Citation	Spanner Size	Additive Stretch	Remarks
Aingworth, Chekuri, Indyk, and Mowani [3]	$O(n^{3/2})$	2	See also [13, 14, 5, 18]
Chechik [11]	$\tilde{O}(n^{7/5})$	4	
Baswana, Kavitha, Mehlhorn, and Pettie [5]	$O(n^{4/3})$	6	See also [23, 18]
Pettie [19]	$O(n^{1+\epsilon})$	$O(n^{9/16-7\epsilon/8})$	$0 \leq \epsilon$
Chechik [11]	$O(n^{20/17+\epsilon})$	$O(n^{4/17-3\epsilon/2})$	$0 \leq \epsilon$
Bodwin and Williams [9]	$O(n^{1+\epsilon})$	$O(n^{1/2-\epsilon/2})$ $O(n^{2/3-5\epsilon/3})$	$0 \leq \epsilon$
Bodwin and Williams [8]	$O(n^{1+o(1)+\epsilon})$	$O(n^{3/7-\epsilon})$	$0 \leq \epsilon \leq 6/49$
		$O(n^{3/5-12\epsilon/5})$	$6/49 \leq \epsilon \leq 2/13$
		$O(n^{3/7-9\epsilon/7})$	$2/13 \leq \epsilon < 1/3$
Abboud and Bodwin [1]	$O(n^{4/3-\epsilon})$	$\Omega(n^\delta)$	$\delta = \delta(\epsilon)$
	$O(n)$	$\Omega(n^{1/22})$	
new	$O(n)$	$\Omega(n^{1/13})$	

+6. Abboud and Bodwin [1] showed that the $4/3$ exponent could not be improved, in the sense that any $+n^{o(1)}$ spanner has size $\Omega(n^{4/3-o(1)})$, and that any $\Omega(n^{4/3-\epsilon})$ -size spanner has additive stretch $+ \Omega(n^\delta)$, $\delta = \delta(\epsilon)$. On the upper bound side, Pettie [19] showed that $O(n)$ -size spanners could have additive stretch $+ \tilde{O}(n^{9/16})$, and Bodwin and Williams [8] improved this to $O(\sqrt{n})$ for $O(n)$ -size spanners and $O(n^{3/7})$ for $O(n^{1+o(1)})$ -size spanners. Abboud and Bodwin [1] extended their lower bound to $O(n)$ -size spanners, showing that they require stretch $+ \Omega(n^{1/22})$. Using our lower bound for shortcuttings as a starting place, we improve [1] by giving an $+ \Omega(n^{1/13})$ stretch lower bound for $O(n)$ -size spanners. See Table 2.

Additive Emulators

Dor, Halperin, and Zwick [13] were the first to explicitly define the notion of an *emulator*, and gave a $+4$ emulator with size $O(n^{4/3})$. Abboud and Bodwin's [1] lower bound applies to emulators, i.e., we cannot go below the $4/3$ threshold without incurring polynomial additive stretch. Bodwin and Williams [9, 8] pointed out that some spanner constructions [5] imply emulator bounds, and gave new constructions of emulators with size $O(n)$ and stretch $+O(n^{1/3})$, and with size $O(n^{1+o(1)})$ and stretch $+O(n^{3/11})$.³ Here we observe that Pettie's [19] $+ \tilde{O}(n^{9/16})$ spanner, when turned into an $O(n)$ -size emulator, has stretch $+ \tilde{O}(n^{1/4})$, which is slightly better than the linear size emulators found in [5, 9, 8]. We improve Abboud and Bodwin's [1] lower bound and show that any $O(n)$ -size emulator has additive stretch $+ \Omega(n^{1/18})$. See Table 3.

Our emulator lower bounds are polynomially weaker than the spanner lower bounds. Although neither bound is likely sharp, this difference reflects the rule that emulators are probably more powerful than spanners. For example, at sparsity $O(n^{4/3})$, the best known emulators [13] are slightly better than spanners [5]. Below the $4/3$ threshold the best *sublinear*

³ This last result is a consequence of [8, Thm. 5] and the fact that any pair set $P \subset V^2$ has a pair-wise emulator with size $|P|$.

■ **Table 3** Upper and lower bounds on additive emulators. Emulators with sublinear additive stretch [22, 17, 2] are not shown.

Citation	Emulator Size	Additive Stretch	Remarks
Aingworth, Chekuri, Indyk, and Mowani [3]	$O(n^{3/2})$	2	See also [13, 14, 5, 18]
Dor, Halperin, and Zwick [13]	$O(n^{4/3})$	4	
Baswana, Kavitha, Mehlhorn and Pettie [5]	$O(n^{1+\epsilon})$	$O(n^{1/2-3\epsilon/2})$	(not claimed in [5])
Bodwin and Williams [9]	$O(n^{1+\epsilon})$	$O(n^{1/3-2\epsilon/3})$	
Bodwin and Williams [8]	$O(n^{1+o(1)+\epsilon})$	$O(n^{3/11-9\epsilon/11})$	(conseq. of [8, Thm. 5])
Pettie [19]	$O(n^{1+\epsilon})$	$\tilde{O}(n^{1/4-3\epsilon/4})$	(not claimed in [19])
Aboud and Bodwin [1]	$O(n)$	$\Omega(n^{1/22})$	
new	$O(n)$	$\Omega(n^{1/18})$	

additive emulators [22, 17] have size $O(n^{1+\frac{1}{2k+1-i}})$ and stretch function $d + O(d^{1-1/k})$.⁴ Aboud, Bodwin, and Pettie [2] showed that this tradeoff is optimal for emulators, but the best sublinear additive spanners [19, 11] are polynomially worse.

There are a certain range of parameters where the emulators are polynomially better than the spanners. On pairwise distance preservers, Bodwin [7] showed that whenever $\omega(n^{1/2}) = |P| = o(n^{2-o(1)})$, any pairwise distance preserver has an $\omega(n + |P|)$ lower bound, creating a gap comparing to an $\Theta(|P|)$ emulator.

There is also another situation where emulators are provable superior: a source-wise distance preserver for $S \subset V$ maintains distances between S -vertices without stretch. A trivial source-wise emulator has size $|S|^2$, e.g., $O(n)$ for $|S| = \sqrt{n}$, but in [12, 7] source-wise spanners with size $O(n)$ only exist for $|S| = O(n^{1/4})$.

Outline

In Section 2 we present diameter lower bounds for shortcut sets of size $O(n)$ and $O(m)$. Section 3 modifies the construction to give lower bounds on additive spanners and additive emulators. We conclude with some remarks in Section 4.

2 Lower Bounds on Shortcutting Digraphs

2.1 Using $O(n)$ Shortcuts

► **Theorem 1.** *There exists a directed graph G with n vertices, such that for any shortcut set E' with size $O(n)$, the graph $(V, E \cup E')$ has diameter $\Omega(n^{1/6})$.*

The remainder of Section 2.1 constitutes a proof of Theorem 1. We begin by defining the vertex set and edge set of G , and its *critical pairs*.

Vertices

The vertex set of G is partitioned into $D + 1$ layers numbered 0 through D . Define $B_d(\rho)$ to be the set of all lattice points in \mathbb{Z}^d within Euclidean distance ρ of the origin. Here we treat

⁴ I.e., vertices initially at distance d are stretched to $d + O(d^{1-1/k})$.

d as a constant. For each $k \in \{0, \dots, D\}$, layer- k vertices are identified with lattice points in $B_d(R + kr)$, where r, R are parameters of the construction. A vertex can be represented by a pair (a, k) , where $a \in B_d(R + rk)$. We want the size of all layers to be the same, up to a constant factor. To that end we fix $R = drD$, so the total number of vertices is

$$\begin{aligned} n &\approx \eta_d R^d \left(1^d + \left(1 + \frac{r}{R}\right)^d + \dots + \left(1 + \frac{rD}{R}\right)^d \right) \\ &= \eta_d R^d \left(1^d + \left(1 + \frac{1}{dD}\right)^d + \dots + \left(1 + \frac{1}{d}\right)^d \right) = \Theta(R^d D) \quad (\text{By definition of } R) \end{aligned}$$

where $\eta_d = \frac{1}{\sqrt{2\pi d}} \left(\frac{2\pi e}{d}\right)^{d/2}$ is the ratio of volume between a d -dimensional ball and a d -dimensional cube.

Edges

Define $\mathcal{V}_d(r)$ to be the set of all lattice points at the corners of the convex hull of $B_d(r)$. We treat elements of $\mathcal{V}_d(r)$ as vectors. For each layer- k vertex (a, k) , $k \in \{0, \dots, D-1\}$, and each vector $v \in \mathcal{V}_d(r)$, we include a directed edge $((a, k), (a + v, k + 1))$. All edges in G are of this form.

Critical Pairs

The critical pair set is defined to be

$$P = \{((a, 0), (a + Dv, D)) \mid a \in B_d(R) \text{ and } v \in \mathcal{V}_d(r)\}$$

Each such pair has a corresponding path of length D , namely $(a, 0) \rightarrow (a + v, 1) \rightarrow \dots \rightarrow (a + Dv, D)$. Lemma 2 shows that this path is unique. It was first proved by Hesse [16] and independently by Coppersmith and Elkin [12]. (Both proofs are inspired by Behrend's [6] construction of arithmetic progression-free sets, which uses ℓ_2 balls rather than convex hulls.)

► **Lemma 2.** (cf. [16, 12]) *The set of critical pairs P have the following properties:*

- For all $(x, y) \in P$, there is a unique path from x to y in G .
- For any two distinct pairs (x_1, y_1) and $(x_2, y_2) \in P$, their unique paths share no edge and at most one vertex.
- $|P| = \Theta(R^d r^{d \frac{d-1}{d+1}})$.

Proof. ■ Let $x = (a, 0)$ and $v \in \mathcal{V}_d(r)$ be the vector for which $y = (a + Dv, D)$. One path from x to y exists by construction. Let $\mathcal{V}_d(r) = \{v_1, v_2, \dots, v_s\}$. Suppose there exists another path from x to y . It must have length D because all edges join consecutive layers. Every edge on this path corresponds to a vector v_i , which implies that Dv can be represented as a linear combination $k_1 v_1 + k_2 v_2 + \dots + k_s v_s$, where $k_1 + \dots + k_s = D$ and $k_i \geq 0$. This implies that v is a non-trivial convex combination of the vectors in $\mathcal{V}_d(r)$, which contradicts the fact that $\mathcal{V}_d(r)$ is a strictly convex set.

- Observe that any edge in the unique (x_1, y_1) path uniquely identifies both x_1 and y_1 .
- $|P| = |B_d(R)| \cdot |\mathcal{V}_d(r)|$. From Bárány and Larman [10], for any constant dimension d , we have $|\mathcal{V}_d(r)| = \Theta(r^{d \frac{d-1}{d+1}})$. ◀

► **Lemma 3.** *Let E' be a shortcut set for $G = (V, E)$. If the diameter of $G' = (V, E \cup E')$ is strictly less than D , then $|E'| \geq |P|$.*

Proof. Every path in G' corresponds to some path in G . However, for pairs in P , there is only one path in G , hence, any shortcut in E' useful for a pair $(x, y) \in P$ must have both endpoints on the unique x - y path in G . By Lemma 2, two such paths for pairs in P share no common edges, hence each shortcut can only be useful for at most one pair in P . If $|E'| < |P|$ then some pair $(x, y) \in P$ must still be at distance D in G' . ◀

Proof of Theorem 1. By Lemma 3, if $|P| = \Omega(n)$, then any shortcut set that makes the diameter $< D$ has size $\Omega(n)$. In order to have $|P| = \Omega(n)$, it suffices to let $r^{d \frac{d-1}{d+1}} \geq D$. This implies $r \geq D^{\frac{d+1}{d(d-1)}}$. From the construction, by fixing d as a constant, we have

$$n = \Theta(R^d D) = \Theta((rD)^d D) = \Omega(D^{1+d+\frac{d+1}{d-1}}).$$

Therefore, the diameter is $D = O\left(n^{1/(1+d+\frac{d+1}{d-1})}\right)$. We can maximize $D = \Theta(n^{1/6})$ in one of two ways, by setting $d = 2$, $r = \Theta(n^{1/4})$, and $R = \Theta(n^{5/12})$, or $d = 3$, $r = \Theta(n^{1/9})$, and $R = \Theta(n^{5/18})$. In either case, the construction leads to a graph with very similar structure: the number of vertices in each layer is $\Theta(n^{5/6})$, and the out degrees of each vertex are $\Theta(n^{1/6})$. ◀

► **Corollary 4.** Fix an $\epsilon \in [0, 1)$ and let d be such that $\epsilon \in [0, \frac{d-1}{d+1}]$. There exists a directed graph G with n vertices, such that for any shortcut set E' with $O(n^{1+\epsilon})$ shortcuts, the graph $(V, E \cup E')$ has diameter $\Omega(n^{(1-\frac{d+1}{d-1}\epsilon)/(1+d+\frac{d+1}{d-1})})$. In particular, by setting $d = 3$ the diameter lower bound becomes $\Omega(n^{\frac{1}{6}-\frac{1}{3}\epsilon})$.

Proof. In order to have $|P| > n^{1+\epsilon}$, it suffices to let $r^{d \frac{d-1}{d+1}} \geq Dn^\epsilon$. Hence, we have

$$\begin{aligned} n^{1-\frac{d+1}{d-1}\epsilon} &= \Theta(R^d D n^{-\frac{d+1}{d-1}\epsilon}) \\ &= \Omega(r^d D^{1+d} n^{-\frac{d+1}{d-1}\epsilon}) && (R = \Theta(rD)) \\ &= \Omega(D^{1+d+\frac{d+1}{d-1}}) && (r^d \geq (Dn^\epsilon)^{\frac{d+1}{d-1}}) \end{aligned}$$

2.2 Using $O(m)$ Shortcuts

Let $G_{(d,r,D)}$ denote the layered graph constructed in Section 2.1 with parameters d, D, r , and $R = drD$, and let P_G be its critical pair set. The total number of edges $m = \Theta(n|\mathcal{V}_d(r)|)$ is always larger than $|P_G| = \Theta(\frac{n}{D}|\mathcal{V}_d(r)|)$ by a factor of D . In order to get a lower bound for $O(m)$ shortcuts, we use a Cartesian product combining two such graphs layer by layer, forming a sparser graph. This transformation was discovered by Hesse [16] and rediscovered by Abboud and Bodwin [1].

Let $G_1 = G_{(d_1,r_1,D)}$ and $G_2 = G_{(d_2,r_2,D)}$ be two graphs with the same number of vertex layers $(D + 1)$. The product graph $G_1 \otimes G_2$ is defined below.

Vertices

The product graph has $2D + 1$ vertex layers numbered $0, \dots, 2D$. The vertex set of layer i is $\{(x, y, i) \mid x \in B_{d_1}(R_1 + \lceil \frac{i}{2} \rceil r_1), y \in B_{d_2}(R_2 + \lfloor \frac{i}{2} \rfloor r_2)\}$. Since we set $R_j = d_j r_j D$, the total number of vertices is $\Theta\left(R_1^{d_1} R_2^{d_2} D\right)$.

Edges

Let (x, y, i) be a vertex in layer i . If i is even, then for every vector $v \in \mathcal{V}_{d_1}(r_1)$ we include an edge $((x, y, i), (x + v, y, i + 1))$. If i is odd, then for every vector $w \in \mathcal{V}_{d_2}(r_2)$, we include an edge $((x, y, i), (x, y + w, i + 1))$. The total number of edges in the product graph is then $\Theta\left(R_1^{d_1} R_2^{d_2} D \left(r_1^{\frac{d_1-1}{d_1+1}} + r_2^{\frac{d_2-1}{d_2+1}}\right)\right)$.

Critical Pairs

By combining two graphs, we are able to construct a larger set of critical pairs, as follows.

$$P = \{((a, b, 0), (a + Dv, b + Dw, 2D)) \mid a \in B_{d_1}(R_1), b \in B_{d_2}(R_2), v \in \mathcal{V}_{d_1}(r_1), w \in \mathcal{V}_{d_2}(r_2)\}$$

In other words, a pair in P can be viewed as the product of two pairs $((a, 0), (a + Dv, D)) \in P_{G_1}$ and $((b, 0), (b + Dw, D)) \in P_{G_2}$.

► **Lemma 5.** *For any $a \in B_{d_1}(R_1)$, $b \in B_{d_2}(R_2)$, $v \in \mathcal{V}_{d_1}(r_1)$ and $w \in \mathcal{V}_{d_2}(r_2)$, there is a unique path from $(a, b, 0)$ to $(a + Dv, b + Dw, 2D)$.*

Proof. Every path in $G_1 \otimes G_2$ from layer 0 to layer $2D$ corresponds to two paths from layers 0 to D in G_1 and G_2 , respectively. It follows from Lemma 2 that

$$(a, b, 0) \rightarrow (a + v, b, 1) \rightarrow (a + v, b + w, 2) \rightarrow \cdots \rightarrow (a + Dv, b + Dw, 2D)$$

is a unique path in $G_1 \otimes G_2$. ◀

In $G_1 \otimes G_2$ it is no longer true that pairs in P have edge-disjoint paths. They may intersect at just one edge.

► **Lemma 6.** *Consider two pairs (x_1, y_1) and $(x_2, y_2) \in P$. Let P_1 and P_2 be the unique shortest paths in the combined graph from x_1 to y_1 and from x_2 to y_2 . Then, $P_1 \cap P_2$ contains at most one edge.*

Proof. Any two *non-adjacent* vertices on the unique x_1 - y_1 path uniquely identify x_1 and y_1 . Thus, two such paths can intersect in at most 2 (consecutive) vertices, and hence one edge. ◀

► **Lemma 7.** *Let E' be a shortcut set on $G = (V, E)$. If the diameter of $(V, E \cup E')$ is strictly less than $2D$, then $|E'| \geq |P|$.*

Proof. Assume the diameter of $(V, E \cup E')$ is strictly less than $2D$. Every useful shortcut connects vertices that are at distance at least 2. By Lemma 6, such a shortcut can only be useful for one pair in P . Thus, if the diameter of $(V, E \cup E')$ is less than $2D$, $|E'| \geq |P|$. ◀

By construction, the size of $|P|$ is

$$|P| = \Theta\left(R_1^{d_1} R_2^{d_2} |\mathcal{V}_{d_1}(r_1)| |\mathcal{V}_{d_2}(r_2)|\right) = \Theta\left(R_1^{d_1} R_2^{d_2} r_1^{\frac{d_1-1}{d_1+1}} r_2^{\frac{d_2-1}{d_2+1}}\right).$$

► **Theorem 8.** *There exists a directed graph G with n vertices and m edges such that for any shortcut set E' with size $O(m)$, the graph $(V, E \cup E')$ has diameter $\Omega(n^{1/11})$.*

26:8 Lower Bounds on Sparse Spanners

Proof. If we set $|P| = \Omega(m)$, by Lemma 7, any shortcut set E' with $O(m)$ shortcuts has diameter $\Omega(D)$. In order to ensure $|P| = \Omega(m)$, it suffices to set $r_1^{\frac{d_1-1}{d_1+1}} \geq r_2^{\frac{d_2-1}{d_2+1}} \geq D$. Hence,

$$\begin{aligned}
 n &= \Theta(R_1^{d_1} R_2^{d_2} D) \\
 &= \Theta\left(r_1^{d_1} r_2^{d_2} D^{d_1+d_2+1}\right) && (R_j = d_j r_j D) \\
 &= \Omega\left(D^{\frac{d_1+1}{d_1-1} D^{\frac{d_2+1}{d_2-1}} D^{d_1+d_2+1}}\right) && (\text{plugging in relation between } r_j \text{ and } d_j, D) \\
 &= \Omega\left(D^{\frac{d_1+1}{d_1-1} + \frac{d_2+1}{d_2-1} + d_1 + d_2 + 1}\right)
 \end{aligned}$$

The exponent is minimized when d_1 and d_2 are either 2 or 3. By setting $d_1 = d_2 = 2$, we get $n = \Omega(D^{11})$ and hence $D = O(n^{1/11})$. In this construction we have $d_1 = d_2 = 2$, $D = \Theta(n^{1/11})$, $r_1 = r_2 = \Theta(n^{3/22})$ and $R_1 = R_2 = \Theta(n^{5/22})$. ◀

3 Lower Bounds on Additive Spanners and Emulators

3.1 $O(n)$ -sized Spanners

► **Definition 9.** Let $G = (V, E)$ be an (unweighted) undirected graph. A subgraph $H = (V, E' \subseteq E)$ is said to be an *spanner with additive stretch* β if for any two vertices $x, y \in V$, $\text{dist}_G(x, y) \leq \text{dist}_H(x, y) \leq \text{dist}_G(x, y) + \beta$.

By combining the technique of Abboud and Bodwin [1] with the graphs constructed in Section 2.2, we obtain a substantially better lower bound on $O(n)$ -size additive spanners.

► **Theorem 10.** *There exists an undirected graph G with n vertices, such that any spanner with $O(n)$ edges has $+\Omega(n^{1/13})$ additive stretch.*

In this section we regard $G_{(d,r,D)}$ to be an *undirected* graph. We begin with the undirected graph $G_0 = G_{(d_1, r_1, D)} \otimes G_{(d_2, r_2, D)}$, then modify it in the *edge expansion step* and the *clique replacement step* to obtain G .

The Edge Expansion Step

Every edge in G_0 is subdivided into D edges, yielding G_E . This step makes the graph very sparse since most of the vertices in G_E have degree 2.

The Clique Replacement Step

Consider a vertex u in G_E that comes from one of the interior layers of G_0 , i.e., layers $1, \dots, 2D - 1$, not 0 or $2D$. Note that u has degree $\delta_1 + \delta_2$, with $\delta_1 = \Theta\left(r_1^{\frac{d_1-1}{d_1+1}}\right)$ edges leading to the preceding layer and $\delta_2 = \Theta\left(r_2^{\frac{d_2-1}{d_2+1}}\right)$ edges leading to the following layer (or vice versa). We replace each such u with a complete bipartite clique K_{δ_1, δ_2} , where each clique vertex becomes attached to one non-clique edge formerly attached to u . The final graph is denoted by G .

Critical Pairs

The set P of *critical pairs* for G is identical to the set of critical pairs for G_0 . For each $(x, y) \in P$, the unique x - y path in G is called a *critical path*.

From the construction, the number of vertices in G is then

$$n = \Theta \left(R_1^{d_1} R_2^{d_2} D^2 (\delta_1 + \delta_2) \right). \quad (1)$$

The number of edges in G is now

$$m = \Theta \left(R_1^{d_1} R_2^{d_2} D (D\delta_1 + D\delta_2 + \delta_1\delta_2) \right). \quad (2)$$

The size of P is

$$|P| = \Theta \left(R_1^{d_1} R_2^{d_2} \delta_1 \delta_2 \right). \quad (3)$$

Lemma 11 is key to relating the size of the spanner with the pair set P .

► **Lemma 11.** *Every clique edge belongs to at most one critical path.*

Proof. Every clique has δ_1 vertices on one side and δ_2 vertices on the other side. Each vertex on the δ_1 side corresponds to a vector $v \in \mathcal{V}_{d_1}(r_1)$ and each vertex on the δ_2 side corresponds to a vector $w \in \mathcal{V}_{d_2}(r_2)$. Each clique edge uniquely determines a pair of vectors (v, w) , and hence exactly one critical pair in P . ◀

► **Lemma 12.** *Every spanner of G with additive stretch $+(2D - 1)$ must contain at least $D|P|$ clique edges.*

Proof. For the sake of contradiction suppose there exists a spanner H containing at most $D|P| - 1$ clique edges. By the pigeonhole principle there exists a pair $(x, y) \in P$ such that at least D clique edges are *missing* in H .

Let $P_{(x,y)}$ be the unique shortest path from x to y in G , and let $P'_{(x,y)}$ be a shortest path from x to y in H . Since G_0 is formed from G by contracting all bipartite cliques and replacing subdivided edges with single edges, we can apply the same operations on $P'_{(x,y)}$ to get a path $P''_{(x,y)}$ in G_0 . We now consider two cases:

- If $P''_{(x,y)}$ is the unique shortest path from x to y in G_0 , then $P'_{(x,y)}$ suffers at least a $+2$ stretch on each of the D missing clique edges, so $|P'_{(x,y)}| \geq |P_{(x,y)}| + 2D$.
- If $P''_{(x,y)}$ is not the unique shortest path from x to y in G_0 , then it must traverse at least two more edges than the shortest x - y path in G_0 (because G_0 is bipartite), each of which is subdivided D times in the formation of G . Thus $|P'_{(x,y)}| \geq |P_{(x,y)}| + 2D$.

In either case, $P'_{(x,y)}$ has at least $+2D$ additive stretch and H cannot be a $+(2D - 1)$ spanner. ◀

Proof of Theorem 10. The goal is to have parameters set up so that $D|P| = \Omega(n)$, so that we can apply Lemma 12. Without loss of generality $\delta_1 \geq \delta_2$. By comparing (1) with (3), it suffices to set $\delta_1 \geq \delta_2 \geq D$. We can express the number of vertices in terms of D as follows:

$$\begin{aligned} n &= \Theta \left(R_1^{d_1} R_2^{d_2} D^2 \delta_1 \right) \\ &= \Omega \left((r_1 D)^{d_1} (r_2 D)^{d_2} D^3 \right) && (\delta_1 \geq \delta_2 \geq D) \\ &= \Omega \left(\left(\delta_1^{\frac{d_1+1}{d_1(d_1-1)}} D \right)^{d_1} \left(\delta_2^{\frac{d_2+1}{d_2(d_2-1)}} D \right)^{d_2} D^3 \right) && (\text{by definition of } \delta_1 \text{ and } \delta_2) \\ &= \Omega \left(D^{\frac{d_1+1}{d_1-1} + d_1 + \frac{d_2+1}{d_2-1} + d_2 + 3} \right) && (\delta_1 \geq \delta_2 \geq D) \end{aligned}$$

26:10 Lower Bounds on Sparse Spanners

The exponent is minimized when d_1 and d_2 are either 2 or 3. By plugging in $d_1 = d_2 = 2$, we get $n = \Omega(D^{13})$ and hence the additive stretch $D = O(n^{1/13})$. This admits a construction with parameters $d_1 = d_2 = 2$, $D = \Theta(n^{1/13})$, $r = \Theta(n^{3/26})$ and $R = \Theta(n^{5/26})$. ◀

► **Corollary 13.** Fix an $\epsilon \in [0, 1/3)$ and let d be such that $\epsilon \in [0, \frac{d-1}{3d+1}]$. There exists a graph G with n vertices such that any spanner $H \subseteq G$ with $O(n^{1+\epsilon})$ edges has additive stretch $+\Omega\left(n^{(1-\frac{3d+1}{d-1}\epsilon)/(3+2d+2\frac{d+1}{d-1})}\right)$. In particular, by setting $d = 3$ the additive stretch becomes $\Omega(n^{\frac{1}{13}-\frac{5}{13}\epsilon})$.

3.2 $O(n)$ -sized Emulators

► **Definition 14.** Let $G = (V, E)$ be an (unweighted) undirected graph. A weighted graph $H = (V, E', w)$ is said to be an *emulator with additive stretch* β if for any two vertices $x, y \in V$, $\text{dist}_G(x, y) \leq \text{dist}_H(x, y) \leq \text{dist}_G(x, y) + \beta$.

The difference between emulators and spanners is that emulators can use weighted edges not present in G . The lower bound graph we use is constructed exactly as in Section 3.1, but with different numerical parameters.

► **Theorem 15.** There exists an undirected graph G with n vertices such that any emulator with $O(n)$ edges has $+\Omega(n^{1/18})$ additive stretch.

► **Lemma 16.** Every emulator with additive stretch $+(2D - 1)$ on G , requires at least $|P|/2$ edges.

Proof. Let H be an emulator with additive stretch $+(2D - 1)$. Without loss of generality, we may assume that any $(u, v) \in E(H)$ has weight precisely $\text{dist}_G(u, v)$. (It is not allowed to be smaller, and it is unwise to make it larger.) We proceed to convert H into a spanner H' that has the same stretch $+(2D - 1)$ on all pairs in P , then apply Lemma 12.

Initially H' is empty. Consider each $(x, y) \in P$ one at a time. Let $P_{(x,y)}$ be the shortest path in H and $P'_{(x,y)}$ be the corresponding path in G . Include the entire path $P'_{(x,y)}$ in H' . After this process is complete, for any $(x, y) \in P$, $\text{dist}_{H'}(x, y) = \text{dist}_H(x, y)$, and H' is a spanner with at most $n + 2D|H|$ edges. In particular, it has at most $2D|H|$ clique edges since each weighted edge in some $P_{(x,y)}$ contributes at most $2D$ clique edges to H' . By Lemma 12, the number of clique edges in H' is at least $D|P|$, hence $|H| \geq |P|/2$. ◀

Proof of Theorem 15. In order to get $|P| = \Omega(n)$, it suffices to set $\delta_1 \geq \delta_2 \geq D^2$.

Now, we have

$$\begin{aligned}
 n &= \Theta\left(R_1^{d_1} R_2^{d_2} D^2 \delta_1\right) \\
 &= \Omega\left((r_1 D)^{d_1} (r_2 D)^{d_2} D^4\right) && (\delta_1 \geq \delta_2 \geq D^2) \\
 &= \Omega\left(\left(\delta_1^{\frac{d_1+1}{d_1(d_1-1)}} D\right)^{d_1} \left(\delta_2^{\frac{d_2+1}{d_2(d_2-1)}} D\right)^{d_2} D^4\right) && (\text{by definition of } \delta_1 \text{ and } \delta_2) \\
 &= \Omega\left(D^{2\frac{d_1+1}{d_1-1} + d_1 + 2\frac{d_2+1}{d_2-1} + d_2 + 4}\right) && (\delta_1 \geq \delta_2 \geq D^2)
 \end{aligned}$$

The exponent is minimized when $d_1 = d_2 = 3$. This implies $n = \Omega(D^{18})$. Thus, we have the additive stretch $D = O(n^{1/18})$. ◀

► **Corollary 17.** Fix an $\epsilon \in [0, 1/3)$ and let d be such that $\epsilon \in [0, \frac{d-1}{3d+1}]$. There exists a graph G with n vertices such that any emulator H with $O(n^{1+\epsilon})$ edges has additive stretch $+\Omega\left(n^{(1-\frac{3d+1}{d-1}\epsilon)/(4+2d+2\frac{d+1}{d-1})}\right)$. In particular, by setting $d = 3$ the additive stretch lowerbound becomes $\Omega(n^{\frac{1}{18}-\frac{5}{18}\epsilon})$.

Using the same proof technique as in [1, 2], it is possible to extend our emulator lower bound to *any* compressed representation of graphs using $\tilde{O}(n)$ bits.

► **Theorem 18.** Consider any mapping from n -vertex graphs to $\tilde{O}(n)$ -length bitstrings. Any algorithm for reconstructing an approximation of dist_G , given the bitstring encoding of G , must have additive error $+\tilde{\Omega}(n^{1/18})$.

Proof. For each subset $T \subseteq P$ construct the graph G_T by removing all clique edges from G that are on the critical paths of pairs in T . Because all clique edges are missing, for all $(x, y) \in T$ we have $d_{G_T}(x, y) \geq d_G(x, y) + 2D$. On the other hand, for all $(x, y) \notin T$, $d_{G_T}(x, y) = d_G(x, y)$.

There are $2^{|P|}$ such graphs. If we represent all such graphs with bitstrings of length $|P| - 1$ then by the pigeonhole principle two such graphs G_T and $G_{T'}$ are mapped to the same bitstring. Let (x, y) be any pair in $T \setminus T'$. Since $\text{dist}_{G_T}(x, y) \geq \text{dist}_{G_{T'}}(x, y) + 2D$, the additive stretch of any such scheme must be at least $2D$. Alternatively, any scheme with stretch $2D - 1$ must use bitstrings of length at least $|P|$.

Now, by setting $d = 3$ with $D = \tilde{\Theta}(n^{1/18})$, $r_1 = r_2 = \tilde{\Theta}(n^{2/27})$ and $R_1 = R_2 = \tilde{\Theta}(n^{7/54})$, we have $|P| = \tilde{\Theta}(n)$. Thus any $\tilde{O}(n)$ -length encoding must recover approximate distances with stretch $+\tilde{\Omega}(n^{1/18})$. ◀

4 Conclusion

Our constructions, like [1, 12, 2, 16], are based on looking at the convex hulls of integer lattice points in \mathbb{Z}^d lying in a ball of some radius. Whereas Theorems 15 and 18 hold for $d = 3$, Theorems 1, 8, and 10 are indifferent between dimensions $d = 2$ and $d = 3$, but that is only because d must be an integer.

Suppose we engage in a little magical thinking, and imagine that there are integer lattices in any *fractional* dimension, and moreover, that some analogue of Bárány and Larman's [10] bound holds in these lattices. If such objects existed then we could obtain slightly better lower bounds. For example, setting $d = 1 + \sqrt{2}$ in the proof of Theorem 1, we would conclude that any $O(n)$ -size shortcut set cannot reduce the diameter below $\Omega(n^{1/(3+2\sqrt{2})})$, which is an improvement over $\Omega(n^{1/6})$ as $3 + 2\sqrt{2} < 5.83$.

For near-linear size spanners and emulators there are still large gaps between the best lower and upper bounds on additive stretch: $[n^{1/13}, n^{3/7}]$ in the case of spanners and $[n^{1/18}, n^{1/4}]$ in the case of emulators. None of the existing lower or upper bound techniques seem up to the task of closing these gaps entirely.

References

- 1 Amir Abboud and Greg Bodwin. The 4/3 additive spanner exponent is tight. *J. ACM*, pages 28:1–28:20, 2017.
- 2 Amir Abboud, Greg Bodwin, and Seth Pettie. A hierarchy of lower bounds for sublinear additive spanners. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2017.

- 3 Donald Aingworth, Chandra Chekuri, Piotr Indyk, and Rajeev Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM Journal on Computing*, 28(4):1167–1181, 1999.
- 4 Noga Alon. Testing subgraphs in large graphs. *Random Structures & Algorithms*, 21(3-4):359–370, 2002.
- 5 Surender Baswana, Telikepalli Kavitha, Kurt Mehlhorn, and Seth Pettie. Additive spanners and (α, β) -spanners. *ACM Transactions on Algorithms*, 7(1):5, 2010.
- 6 Felix Behrend. On sets of integers which contain no three terms in arithmetic progression. *Proc. Nat. Acad. Sci.*, 32:331–332, 1946.
- 7 Greg Bodwin. Linear size distance preservers. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 600–615, 2017.
- 8 Greg Bodwin and Virginia Vassilevska Williams. Better distance preservers and additive spanners. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2016.
- 9 Gregory Bodwin and Virginia Vassilevska Williams. Very sparse additive spanners and emulators. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science (ITCS)*, pages 377–382, 2015.
- 10 Imre Bárány and David G. Larman. The convex hull of the integer points in a large ball. *Mathematische Annalen*, 312(1):167–181, 1998.
- 11 Shiri Chechik. New additive spanners. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 498–512, 2013.
- 12 Don Coppersmith and Michael Elkin. Sparse sourcewise and pairwise distance preservers. *SIAM Journal on Discrete Mathematics*, pages 463–501, 2006.
- 13 Dorit Dor, Shay Halperin, and Uri Zwick. All-pairs almost shortest paths. *SIAM Journal on Computing*, 29(5):1740–1759, 2000.
- 14 Michael Elkin and David Peleg. $(1 + \epsilon, \beta)$ -spanner constructions for general graphs. *SIAM Journal on Computing*, 33(3):608–631, 2004.
- 15 Jeremy T. Fineman. Nearly work-efficient parallel algorithm for digraph reachability. In *Proceedings of the 50th Annual ACM Symposium on Theory of Computing (STOC)*, 2018.
- 16 William Hesse. Directed graphs requiring large numbers of shortcuts. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2003.
- 17 Shang-En Huang and Seth Pettie. Thorup-Zwick emulators are universally optimal hopsets. *CoRR*, abs/1705.00327, 2017. [arXiv:1705.00327](https://arxiv.org/abs/1705.00327).
- 18 Mathias Bæk Tejs Knudsen. Additive spanners: A simple construction. In *Scandinavian Workshop on Algorithm Theory (SWAT)*, pages 277–281, 2014.
- 19 Seth Pettie. Low distortion spanners. *ACM Transactions on Algorithms*, 6(1):7, 2009.
- 20 Mikkel Thorup. On shortcutting digraphs. In *International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 205–211. Springer, 1992.
- 21 Mikkel Thorup. Shortcutting planar digraphs. *Combinatorics, Probability and Computing*, 4(3):287–315, 1995.
- 22 Mikkel Thorup and Uri Zwick. Spanners and emulators with sublinear distance errors. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2006.
- 23 David P. Woodruff. Lower bounds for additive spanners, emulators, and more. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2006.