SIAM J. DISCRETE MATH.
Vol. 35, No. 3, pp. 2129–2144

LOWER BOUNDS ON SPARSE SPANNERS, EMULATORS, AND DIAMETER-REDUCING SHORTCUTS*

SHANG-EN HUANG† AND SETH PETTIE‡

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 prove 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 [*Proceedings of the* 14*th Annual ACM-SIAM Symposium on Discrete Algorithms* (SODA), Baltimore, MD, 2003] lower bound of $\Omega(n^{1/17})$. By combining these constructions with Abboud and Bodwin's [*J. ACM*, 64 (2017), 28] edge-splitting technique, we get additive stretch lower bounds of $+\Omega(n^{1/11})$ 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 for both spanners and emulators.

Key words. additive spanners, emulators, shortcutting directed edges

AMS subject classification. 05C12

DOI. 10.1137/19M1306154

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:

$$\mathrm{dist}_G(u, v) \le \mathrm{dist}_H(u, v) \le \mathrm{dist}_G(u, v) + \beta,$$

where β may depend on n, the number of nodes in G.

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 [19], and Coppersmith and Elkin [14]; see also [1, 2]. Given such a "base graph," derived graphs can be obtained through a variety of graph products such as the *alternation* product discovered independently by Hesse [19] 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, 19, 14, 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 $\mathrm{dist}_G(u, v)$ over all pairs $(u, v) \in E^*$. Thorup [24] conjectured that it is possible to reduce the diameter

*Received by the editors December 11, 2019; accepted for publication (in revised form) June 7, 2021; published electronically September 20, 2021.
https://doi.org/10.1137/19M1306154
Funding: This work was supported by NSF grants CCF-1514383, CCF-1637546, and CCF-1815316.
†Computer Science and Engineering, University of Michigan, Ann Arbor, MI 48109 USA (sehuang@umich.edu).
‡Electrical Engineering and Computer Sciences, University of Michigan, Ann Arbor, MI 48109 USA (pettie@umich.edu).

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 [18]	$\tilde{O}(n)$	$\tilde{O}(n^{2/3})$	$\tilde{O}(m)$
Liu, Jambulapati, and Sidford [22]	$\tilde{O}(nk)$	$\tilde{O}(n^{1/2+O(1/\log k)})$	$\tilde{O}(mk)$
Hesse [19]	$O(mn^{1/17})$	$\Omega(n^{1/17})$	—
New	$O(n)$	$\Omega(n^{1/6})$	—
	$O(m)$	$\Omega(n^{1/11})$	—

of any digraph to 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 poly$(\log n)$. This conjecture was confirmed for a couple special graph classes [24, 25] but refuted in general by Hesse [19], 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] Recently, Fineman [18] 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 $\tilde{O}(n^{2/3})$. The diameter is further reduced by Liu, Jambulapati, and Sidford [22] (and parallelized by Cao, Fineman, and Russell [12]) to $\tilde{O}(n^{1/2+o(1)})$ by slightly modifying Fineman's algorithm.

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, et al. [3] (see also [15, 17, 6, 21, 16]), Chechik [13] (see also [9]), and Baswana, et al. [6] (see also [27, 21]). The sparsest of these [6] has a size $O(n^{4/3})$ and stretch +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 a size $\Omega(n^{4/3-o(1)})$ and that any $\Omega(n^{4/3-\epsilon})$-size spanner has an additive stretch $+\Omega(n^\delta)$, $\delta = \delta(\epsilon)$. On the upper bound side, Pettie [23] showed that $O(n)$-size spanners could have additive stretch $+\tilde{O}(n^{9/16})$, and Bodwin and Williams [11] improved this to $O(\sqrt{n})$ for $O(n)$-size

[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. $\delta > 0$ is a constant.

Citation	**Spanner size**	**Additive stretch**	**Remarks**
Aingworth, et al. [3]	$O(n^{3/2})$	2	See also [15, 17, 6, 21]
Bodwin [9]	$O(n^{7/5})$	4	See also [13]
Baswana, et al. [6]	$O(n^{4/3})$	6	See also [27, 21]
Pettie [23]	$O(n^{1+\epsilon})$	$O(n^{9/16-7\epsilon/8})$	$0 \leq \epsilon$
Chechik [13]	$O(n^{20/17+\epsilon})$	$O(n^{4/17-3\epsilon/2})$	$0 \leq \epsilon$
Bodwin and Williams [10]	$O(n^{1+\epsilon})$	$O(n^{1/2-\epsilon/2})$ $O(n^{2/3-5\epsilon/3})$	$0 \leq \epsilon$
Bodwin and Williams [11]	$O_\delta(n^{1+\epsilon})$	$O(n^{3/7+\delta-9\epsilon/7})$	$\epsilon \in \{0\} \cup [2/13, 1/3)$
		$O(n^{3/7+\delta-5\epsilon/7})$	$\epsilon \in (0, 2/15]$
		$O(n^{1+\delta-5\epsilon})$	$\epsilon \in (2/15, 2/13)$
Abboud and Bodwin [1]	$O(n^{4/3-\epsilon})$	$\Omega(n^\delta)$	$\delta = \delta(\epsilon)$
	$O(n)$	$\Omega(n^{1/22})$	
New	$O(n)$	$\Omega(n^{1/11})$	

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 a stretch $+\Omega(n^{1/22})$. Using our lower bound for shortcuttings as a starting place, we improve [1] by giving an $+\Omega(n^{1/11})$ stretch lower bound for $O(n)$-size spanners. See Table 2.

Additive emulators. Dor, Halperin, and Zwick [15] 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 [10, 11] pointed out that some spanner construtions [6] imply emulator bounds and gave new constructions of emulators with a size $O(n)$ and stretch $+O(n^{1/3})$ and with a size $O(n^{1+o(1)})$ and stretch $+O(n^{3/11})$.[3] Here we observe that Pettie's [23] $+\tilde{O}(n^{9/16})$ spanner, when turned into an $O(n)$-size emulator, has a stretch $+\tilde{O}(n^{1/4})$, which is slightly better than the linear-size emulators found in [6, 10, 11]. We improve Abboud and Bodwin's [1] lower bound and show that any $O(n)$-size emulator has an 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 possibility that emulators may be strictly more powerful than spanners. For example, at sparsity $O(n^{4/3})$, the best known emulators [15] are slightly better than spanners [6]. Below the 4/3 threshold the best *sublinear additive* emulators [26, 20] have a size $O(n^{1+\frac{1}{2^{k+1}-1}})$ and stretch function $d + O(d^{1-1/k})$.[4] Abboud, Bodwin, and Pettie [2]

[3]This last result is a consequence of [11, Thm. 5] and the fact that any pair set $P \subset V^2$ has a pair-wise emulator with size $|P|$.

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

TABLE 3
Upper and lower bounds on additive emulators. Emulators with sublinear additive stretch [26, 20, 2] *are not shown.*

Citation	**Emulator size**	**Additive stretch**	**Remarks**
Aingworth, et al. [3]	$O(n^{3/2})$	2	See also [15, 17, 6, 21]
Dor, Halperin, and Zwick [15]	$O(n^{4/3})$	4	
Baswana, et al. [6]	$O(n^{1+\epsilon})$	$O(n^{1/2-3\epsilon/2})$	(not claimed in [6])
Bodwin and Williams [10]	$O(n^{1+\epsilon})$	$O(n^{1/3-2\epsilon/3})$	
Pettie [23]	$O(n^{1+\epsilon})$	$\tilde{O}(n^{1/4-3\epsilon/4})$	(not claimed in [23])
Abboud and Bodwin [1]	$O(n)$	$\Omega(n^{1/22})$	
New	$O(n)$	$\Omega(n^{1/18})$	

showed that this tradeoff is optimal for emulators, but the best known sublinear additive spanners [23, 13] are polynomially worse.

There are a certain range of parameters where emulators are known to be polynomially sparser than spanners. For pairwise distance preservers with the specified set P of vertex pairs, Coppersmith and Elkin [14] showed that whenever $\omega(n^{1/2}) = |P| = o(n^{2-o(1)})$, any pairwise distance preserver has an $\omega(n + |P|)$ lower bound, which is worse than the trivial distance preserving emulator with size $|P|$. A similar separation holds for sourcewise distance preservers, where the goal is to exactly preserve distances between all vertex pairs in $S \subset V$. A trivial sourcewise *emulator* has size $|S|^2$, e.g., $O(n)$ for $|S| = \sqrt{n}$, but sourcewise *spanners* with a size $O(n)$ only exist for $|S| = O(n^{1/4})$ [8].

Organization. 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.

Existentially, the best-known upper bound on $O(n)$-size shortcut sets is the trivial $\tilde{O}(\sqrt{n})$ bound. Theorem 2.1 shows that we cannot go below $\Omega(n^{1/6})$.

THEOREM 2.1. *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 2.1. 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. In the calculations below we treat 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 $a \in B_d(R + kr)$. 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

$$n \approx \eta_d R^d \left(1^d + \left(1 + \frac{r}{R}\right)^d + \cdots + \left(1 + \frac{rD}{R}\right)^d\right)$$
$$= \eta_d R^d \left(1^d + \left(1 + \frac{1}{dD}\right)^d + \cdots + \left(1 + \frac{1}{d}\right)^d\right) \;=\; \Theta\left(R^d D\right) \quad \text{(by definition of } R\text{)},$$

where $\eta_d = \frac{1}{\sqrt{\pi d}}(\frac{2\pi e}{d})^{d/2}$ is the ratio of volume between a d-dimensional ball of radius R to a d-dimensional cube of side length R.

Edges. Define $\mathcal{V}_d(r)$ to be the set of all lattice points at the corners of the convex hull of $B_d(r)$. (This excludes points that happen to lie on the boundary but in the interior of one of its faces.) We treat elements of $\mathcal{V}_d(r)$ as vectors. For each layer-k vertex (a, k), $k \in \{0, \dots, D-1\}$, and for 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,

$$(a, 0) \to (a + v, 1) \to \cdots \to (a + Dv, D).$$

Lemma 2.2 shows that this path is unique. It was first proved by Hesse [19] and independently by Coppersmith and Elkin [14]. (Both proofs are inspired by Behrend's [7] construction of arithmetic progression-free sets, which uses ℓ_2 balls rather than convex hulls.)

Lemma 2.2 (cf. [19, 14]). *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. For the first claim, 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 $k_1 v_1 + k_2 v_2 + \cdots + k_s v_s$, where $k_1 + \cdots + k_s = D$ and $k_i \geq 0$. This implies that v is a nontrivial convex combination of the vectors in $\mathcal{V}_d(r)$, which contradicts the fact that $\mathcal{V}_d(r)$ is a strictly convex set.

The second claim follows from the fact that any *edge* in the unique x_1-to-y_1 path uniquely identifies both x_1 and y_1.

For the last claim, we can express the number of critical pairs as $|P| = |B_d(R)| \cdot |\mathcal{V}_d(r)|$. From Bárány and Larman [5], for any constant dimension d, we have $|\mathcal{V}_d(r)| = \Theta(r^{d\frac{d-1}{d+1}})$. ◻

Lemma 2.3. *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 2.2, 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 2.1. By Lemma 2.3, if $|P| = \Omega(n)$, then any shortcut set that makes the diameter $< D$ has a 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 and using the fact that $n = \Theta(R^d D)$, we have

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

Therefore, the diameter is $D = O(n^{1/(1+d+\frac{d+1}{d-1})})$. 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 $\Theta(n^{1/6})$. ∎

Theorem 2.1 is indifferent between $d = 2$ and $d = 3$, but that is only because the size of the shortcut set is precisely $O(n)$. When we allow it to be $O(n^{1+\epsilon})$, for $\epsilon > 0$, there is generally one optimum dimension.

Corollary 2.4. *Fix an $\epsilon \in [0, 1)$, and let d be an integer 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 $\Omega(n^{\left(1-\frac{d+1}{d-1}\epsilon\right)/\left(1+d+\frac{d+1}{d-1}\right)})$. 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 [19] and rediscovered by Abboud and Bodwin [1].

Let G_1 and G_2 be two copies of $G_{(d,r,D)}$ where each of them has $D+1$ layers. The product graph $G_1 \otimes G_2$ is defined below.

Vertices. The product graph has $2D+1$ vertex layers numbered $0, \ldots, 2D$. The vertex set of layer i is $\{(x, y, i) \mid x \in B_d(R + \lceil \frac{i}{2} \rceil r), y \in B_d(R + \lfloor \frac{i}{2} \rfloor r)\}$. Since we set $R = drD$, the total number of vertices is $\Theta\left(R^{2d} D\right)$.

Edges. Let (x, y, i) be a vertex in layer i. If i is even, then for every vector $v \in \mathcal{V}_d(r)$ we include an edge $((x, y, i), (x + v, y, i + 1))$. If i is odd, then for every vector $w \in \mathcal{V}_d(r)$ we include an edge $((x, y, i), (x, y + w, i + 1))$. The total number of edges in the product graph is then $\Theta(R^{2d} D r^{d\frac{d-1}{d+1}})$.

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

$$P = \{((a, b, 0), (a + Dv, b + Dw, 2D)) \mid a, b \in B_d(R); v, w \in \mathcal{V}_d(r)\}.$$

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 2.5. *For any $a,b \in B_d(R)$ and $v,w \in \mathcal{V}_d(r)$, 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 2.2 that

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

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 2.6. *Consider two pairs (x_1,y_1) and $(x_2,y_2) \in P$. Let P_1 and P_2 be the unique 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 *nonadjacent* 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 2.7. *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 a distance of at least 2. By Lemma 2.6, 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

$$|P| = \Theta\left(R^{2d}|\mathcal{V}_d(r)|^2\right) = \Theta\left(R^{2d} r^{2d\frac{d-1}{d+1}}\right).$$

THEOREM 2.8. *There exists a directed graph G with n vertices and m edges such that for any shortcut set E' with a size $O(m)$, the graph $(V,E\cup E')$ has a diameter $\Omega(n^{1/11})$.*

Proof. If we set $|P| = \Omega(m)$, by Lemma 2.7, any shortcut set E' with $O(m)$ shortcuts has a diameter $\Omega(D)$. In order to ensure $|P| = \Omega(m)$, it suffices to set $r^{d\frac{d-1}{d+1}} \geq D$. Hence,

$$\begin{aligned} n &= \Theta(R^{2d}D) \\ &= \Theta\left(r^{2d}D^{2d+1}\right) && (R = drD) \\ &= \Omega\left(D^{2\frac{d+1}{d-1}}D^{2d+1}\right) && \text{(plugging in relation between } r \text{ and } d, D\text{)} \\ &= \Omega\left(D^{2\frac{d+1}{d-1}+2d+1}\right). \end{aligned}$$

The exponent is minimized when d is either 2 or 3, so we get $n = \Omega(D^{11})$ and hence $D = O(n^{1/11})$. In particular, by setting $d=2$ we have $D = \Theta(n^{1/11})$, $r = \Theta(n^{3/22})$, and $R = \Theta(n^{5/22})$, and by setting $d=3$ we have $D = \Theta(n^{1/11})$, $r = \Theta(n^{2/33})$, and $R = \Theta(n^{5/33})$. ◻

3. Lower bounds on additive spanners and emulators. We now establish better bounds on $O(n)$-size additive spanners and emulators. In section 3.1 we give an $+\Omega(n^{1/13})$ stretch lower bound on spanners. Using a different construction, we improve this in section 3.2 to $+\Omega(n^{1/11})$. In section 3.3 we show how to adapt the $+\Omega(n^{1/13})$-spanner lower bound from section 3.1 to prove that $O(n)$-size emulators have stretch $+\Omega(n^{1/18})$.

Recall the definition of additive spanners and emulators.

DEFINITION 3.1. *Let $G = (V, E)$ be an unweighted undirected graph. A* subgraph *$H = (V, E')$, $E' \subseteq E$, is said to be a* spanner for G with additive stretch β *if for any two vertices $u, v \in V$,*

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

A weighted graph $H = (V, E', w)$ is an emulator for G with additive stretch β *if*

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

Observe that we can assume without loss of generality that if $(u, v) \in E'$, then $w(u, v) = \mathrm{dist}_G(u, v)$.

3.1. $O(n)$-size spanners. By combining the technique of Abboud and Bodwin [1] with the graphs constructed in section 2.2, we improve the $+\Omega(n^{1/22})$ lower bound of [1] to $+\Omega(n^{1/13})$ for $O(n)$-size spanners.

THEOREM 3.2. *There exists an undirected graph G with n vertices such that any spanner for G 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 $G_0 = G_{(d,r,D)} \otimes G_{(d,r,D)}$, then modify it in the *edge subdivision step* and the *biclique replacement step* to obtain G.

The edge subdivision 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 now have degree 2.

The biclique replacement step. Consider a vertex u in G_E that comes from one of the interior layers of G_0, i.e., layers $1, \ldots, 2D-1$, not 0 or $2D$. Note that u has degree at most 2δ with at most $\delta = \Theta(r^{d\frac{d-1}{d+1}})$ edges leading to the preceding layer and exactly δ edges leading to the following layer. We replace each such u with a complete bipartite clique $K_{\delta,\delta}$, where each biclique vertex becomes attached to one nonbiclique 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, edges, and critical pairs in G is

$$n = \Theta\left(R^{2d} D^2 \delta\right), \tag{3.1}$$

$$m = \Theta\left(R^{2d} D(D\delta + \delta^2)\right), \tag{3.2}$$

$$|P| = \Theta\left(R^{2d} \delta^2\right). \tag{3.3}$$

Lemma 3.3 is key to relating the size of the spanner with the pair set P.

LEMMA 3.3. *Every biclique edge belongs to at most one critical path.*

Proof. Every biclique has δ vertices on one side and δ vertices on the other side. Each vertex on the left side corresponds to a vector $v \in \mathcal{V}_d(r)$, and each vertex on the right side also corresponds to a vector $w \in \mathcal{V}_d(r)$. Each biclique edge uniquely determines a pair of vectors (v, w) and hence exactly one critical pair in P. ∎

LEMMA 3.4. *Every spanner of G with additive stretch $+(2D-1)$ must contain at least $D|P|$ biclique edges.*

Proof. For the sake of contradiction suppose there exists a spanner H containing at most $D|P|-1$ biclique edges. For any critical pair $(u,v) \in P$, let $P_{(u,v)}$ be the unique shortest path from u to v in G, and let $P'_{(u,v)}$ be the shortest path from u to v in H. By the pigeonhole principle there exists a pair $(x,y) \in P$ such that at least D biclique edges along $P_{(x,y)}$ are *missing* 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 biclique 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 3.2. The goal is to have parameters set up so that $D|P| = \Omega(n)$, and then apply Lemma 3.4. By comparing (3.1) with (3.3), it suffices to set $\delta \geq D$. We can express the number of vertices in terms of D as follows:

$$\begin{aligned} n &= \Theta\left(R^{2d}D^2\delta\right) \\ &= \Omega\left((rD)^{2d}D^3\right) && (\delta \geq D) \\ &= \Omega\left(\left(\delta^{\frac{d+1}{d(d-1)}}D\right)^{2d}D^3\right) && (\text{by definition of } \delta) \\ &= \Omega\left(D^{2\frac{d+1}{d-1}+2d+3}\right) && (\delta \geq D). \end{aligned}$$

The exponent is minimized when d is either 2 or 3, so $n = \Omega(D^{13})$, and hence the additive stretch is $D = O(n^{1/13})$. When $d=2$ we have $D = \Theta(n^{1/13})$, $r = \Theta(n^{3/26})$, and $R = \Theta(n^{5/26})$, and when $d=3$ we have $D = \Theta(n^{1/13})$, $r = \Theta(n^{2/39})$, and $R = \Theta(n^{5/39})$. ∎

COROLLARY 3.5. *Fix an $\epsilon \in [0, 1/3)$, and let d be an integer 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 a additive stretch $+\Omega(n^{(1-\frac{3d+1}{d-1}\epsilon)/(3+2d+2\frac{d+1}{d-1})})$. In particular, by setting $d=3$ the additive stretch becomes $\Omega(n^{\frac{1}{13}-\frac{5}{13}\epsilon})$.*

Proof. By the above construction, we may express n and $|P|$ as functions to D, δ, and ϵ:

$$\begin{aligned} n &= \Theta\left(R^{2d}D^2\delta\right) = \Theta\left(D^{2+2d}\delta^{2\frac{d+1}{d-1}+1}\right), \\ |P| &= \Theta\left(R^{2d}\delta^2\right) = \Theta\left(D^{2d}\delta^{2\frac{d+1}{d-1}+2}\right). \end{aligned}$$

To show that $D|P| \geq \Omega(n^{1+\epsilon})$ it suffices to set

$$\begin{aligned}
&\delta^{\left(2\frac{d+1}{d-1}+2\right)-(1+\epsilon)\left(2\frac{d+1}{d-1}+1\right)} = D^{(2d+2)(1+\epsilon)-(2d+1)} \\
\iff &\delta^{\left(2\frac{d+1}{d-1}+1\right)(-\epsilon+1/(2\frac{d+1}{d-1}+1))} = D^{(2d+2)\epsilon+1} \\
\iff &\delta^{\left(2\frac{d+1}{d-1}+1\right)(-\epsilon+1/(\frac{3d+1}{d-1}))} = D^{(2d+2)\epsilon+1} \\
\iff &\delta^{2\frac{d+1}{d-1}+1} = D^{((2d+2)\epsilon+1)\frac{(3d+1)/(d-1)}{1-(3d+1)/(d-1)\epsilon}}, \qquad \left(\text{meaningful if} \epsilon < \frac{d-1}{3d+1}\right).
\end{aligned}$$

Now, we are able to establish that there exists a graph with the following relation between n and D,

$$\begin{aligned}
n &= \Theta\left(D^{2+2d}\delta^{2\frac{d+1}{d-1}+1}\right) \\
&= \Theta\left(D^{2+2d+((2d+2)\epsilon+1)\frac{(3d+1)/(d-1)}{1-(3d+1)/(d-1)\epsilon}}\right) \\
&= \Theta\left(D^{\left(3+2d+2\frac{d+1}{d-1}\right)/\left(1-\frac{3d+1}{d-1}\epsilon\right)}\right),
\end{aligned}$$

which concludes the proof. ∎

3.2. An improved $O(n)$-size spanner lower bound. The construction from section 3.1 is versatile, inasmuch as it extends to polynomial densities (Corollary 3.5) and emulator lower bounds (section 3.3). However, it can be improved slightly for the specific case of $O(n)$-size additive spanners. It turns out that the the Cartesian product step (generating G_0 from $G_{(d,r,D)} \otimes G_{(d,r,D)}$) is inefficient and that we can do better with a simple replacement step.

By its nature, the proof of Theorem 3.6 needs to explictly keep track of the leading absolute constant in the size of the spanner; i.e., it has at most $c_0 n = O(n)$ edges. (In contrast, the proof of Theorem 3.2 can easily accommodate any $O(n)$-size bound by tweaking r, R, D by constant factors.) Although c_0 and c (see Lemma 3.7) will eventually be set to constants, we treat them as parameters; all asymptotic notation hides constants that are *independent* of c_0, c.

THEOREM 3.6. *For any parameter $c_0 > 1$ and sufficiently large n there exists an undirected n-vertex graph G such that any spanner for G with at most $c_0 n$ edges has $+\Omega(n^{1/11}c_0^{-18/11})$ additive stretch.*

In Lemmas 3.7 and 3.8 we construct the *inner* and *outer* graphs, then discuss how to combine them using a substitution product.

LEMMA 3.7 (Inner graph construction). *Fix a parameter $c > 1$. There exists sufficiently large q, L such that $q = \Theta(L^2 c^6)$ and a graph $G_I = (V_I, E_I)$ with a set of critical pairs $P_I \subseteq V_I \times V_I$ satisfying the following conditions.*

1. *$|V_I| \leq qL$.*
2. *$|P_I| \geq q$.*
3. *For all $(u, v) \in P_I$, the shortest path between u and v is unique and has a length Lc.*
4. *For all $(u_1, v_1), (u_2, v_2) \in P$, the unique shortest paths between u_1 and v_1 and between u_2 and v_2 are edge-disjoint. Moreover, $\text{dist}_{G_I}(u_1, v_2) \geq Lc$, and $\text{dist}_{G_I}(u_2, v_1) \geq Lc$. (As a consequence, $|E_I| \geq cqL$.)*

Proof. We use almost the same construction as in Theorem 2.1 except that the graph will be undirected and we will pay closer attention to the density. In this proof

$d = \Theta(c)$ represents the density of the graph, not the geometric dimension. We will ultimately choose $d = 12.97c$. The graph G_I we construct consists of $Lc + 1$ layers, numbered by $0, 1, \ldots, Lc$.

Recall that $\mathcal{V}_2(r)$ is the set of all lattice points at the corners of the convex hull of $B_2(r)$. Let $\eta_2 = \pi = \Theta(1)$ be the ratio between the area of a circle of a radius 1 and the area of the unit square. In this case we know that for large r, $(\eta_2 - 0.1)r^2 \leq |B_2(r)| \leq (\eta_2 + 0.1)r^2$. We do not always have $|B_2(r)| = \eta_2 r^2$ because we are working on a lattice. In this sense η_2 is the limiting value; i.e., $\eta_2 = \lim_{r\to\infty} |B_2(r)|/r^2$. In addition, let $\xi_d = \Theta(1)$ be such that $|\mathcal{V}_2(\xi_d d^{3/2})| \geq d/(\eta_2 - 0.1)$. (It follows from [5] that we can choose $\xi_d = (3/(\pi - 0.1))^{3/2} \approx 0.979 = \Theta(1)$.) On the kth layer, the vertices are labelled by (a, k) where $a \in B_2(\sqrt{q/d} + k\xi_d d^{3/2})$. For each layer-$k$ vertex (a, k), $k \in \{0, \ldots, Lc-1\}$ and each vector $v \in \mathcal{V}_2(\xi_d d^{3/2})$, we connect an (undirected) edge between (a, k) and $(a + v, k + 1)$.

By choosing $\sqrt{q/d} = (Lc)\xi_d d^{3/2}$, we have $q = L^2c^2\xi_d^2 d^4 = \Theta(L^2c^6)$, and the total number of vertices in G can be upper bounded by the number of layers times the size of the last layer:

$$(Lc)|B_2(2\sqrt{q/d})| \leq (Lc)(\eta_2 + 0.1)(2\sqrt{q/d})^2 \leq 12.97qLc/d.$$

Thus, condition 1 is satisfied whenever $d \geq 12.97c$.

Define

$$P_I = \left\{ ((a, 0), (a + (Lc)v, (Lc))) \;\middle|\; a \in B_2(\sqrt{q/d}) \text{ and } v \in \mathcal{V}_2(\xi_d d^{3/2}) \right\}.$$

We have that $|P_I| = |B_2(\sqrt{q/d})| \cdot |\mathcal{V}_2(\xi_d d^{3/2})| \geq (\eta_2 - 0.1)(q/d) \cdot d/(\eta_2 - 0.1) = q$, so condition 2 is satisfied.

Now, for each pair of vertices $((a, 0), (a + (Lc)v, (Lc)))$ in P, there is an unique shortest path from $(a, 0)$ to $(a + (Lc)v, (Lc))$ by Lemma 2.2. Moreover, since the graph is a layered graph, any path from a vertex in the 0th layer to *any* vertex in the (Lc)th layer has length at least Lc, satisfying conditions 3 and 4. $\square$

Again, we use a similar construction to Theorem 2.1 to obtain our outer graph.

LEMMA 3.8 (Outer graph Construction, three-dimensional version). *For any given $q, L \in \mathbb{N}$, there exists an undirected graph $G_0 = (V_0, E_0)$ with a set of critical pairs $P \subseteq V_0 \times V_0$ satisfying*

1. *G_0 is a layered graph with $L + 1$ layers. Each vertex in G_0 connects to at most q vertices in the next layer and at most q vertices in the previous layer.*
2. $|V_0| = \Theta(L^4q^2)$.
3. $|P| = \Theta(L^3q^3)$.
4. *For all $(u, v) \in P$, the shortest path between u and v (denoted by P_{uv}) is unique. Moreover, P_{uv} has a length exactly L.*
5. *For all $(u_1, v_1), (u_2, v_2) \in P$, the unique shortest paths between u_1 and v_1 and between u_2 and v_2 are edge-disjoint.*

Proof. Consider the following $(L + 1)$-layer graph. Vertices in the kth layer are identified with points in the three-dimensional integer lattice inside the ball of radius $Lr + kr$ around the origin. Here r is the minimum value such that $|\mathcal{V}_3(r)| \geq q$. From Bárány and Larman [5] we have $r = \Theta(q^{2/3})$.

We label each vertex with its coordinate and its layer number: $(a, k) \in B_3(Lr + kr) \times [L + 1]$. Fix an arbitrary subset $\mathcal{V}_3'(r) \subseteq \mathcal{V}_3(r)$ of any q vectors. For each vertex (a,k) in the kth layer $(0 \le k < L)$, and for every vector $x \in \mathcal{V}'_3(r)$, the edge $((a,k),(a+x,k+1))$ is added to the graph.

It is straightforward to check that $|V_0| \approx \sum_{k=0}^{L} \eta_3(2(Lr+kr))^3 = \Theta(L^4q^2)$. For each vector $v \in \mathcal{V}'_3(r)$ and each layer-0 vertex $(a,0)$, the vertex pair $((a,0),(a+Lv,L))$ is added to the critical pair set P; hence $|P| = \Theta((Lr)^3q) = \Theta(L^3q^3)$. By the same argument as in the proof of Lemma 2.2, there is exactly one shortest path of length L connecting $(a,0)$ and $(a+Lv,L)$. Moreover, no edge belongs to more than one critical path. ∎

Recall that we are aiming for lower bounds against spanners with a size c_0n. Once c_0 is fixed, we choose a $c = \Theta(c_0)$ and invoke Lemma 3.7 to construct an inner graph G_I with parameters q, L. Once q, L are fixed we invoke Lemma 3.8 to build the outer graph G_0. Our final graph G is formed from G_0, G_I through the *inner graph replacement step* and the *edge subdivision step*, as follows.

Inner graph replacement step. For every vertex $(a,k) \in G_0$ $(0 < k < L)$, we replace (a,k) with a copy of inner graph G_I as follows.

Recall that the critical pair set for G_I has a size q. We regard the sources of these q pairs to be *input ports* and the sinks to be *output ports.* Let $G_{I,(a,k)}$ be the copy of G_I substituted for (a,k) in the outer graph. For each $v_i \in \mathcal{V}'_3(r) = \{v_1, v_2, \dots, v_q\}$ and each critical path of G_0 passing through $(a-v_i,k-1),(a,k),(a+v_i,k+1)$, we reattach the vertex $(a-v_i,k-1)$ to the ith input port of $G_{I,(a,k)}$ and reattach the vertex $(a+v_i,k+1)$ to the ith output port of $G_{I,(a,k)}$. Let G^* be the result of this process.

The edge subdivision step. Every edge in G^* that was inherited from G_0 (i.e., not inside any copy of G_I) is subdivided into a path of $L/2$ edges. The outcome of this process is G.

Observe that for every critical pair (x,y) from G_0, there is a unique shortest path between x and y in G of length $\frac{1}{2}L^2 + (L-1)Lc$, where $\frac{1}{2}L^2$ edges come from the subdivision step and the remaining ones come from copies of G_I. Moreover, any two unique shortest paths are edge-disjoint.

LEMMA 3.9. *Every spanner of G with an additive stretch $+(L-2)$ contains at least $(\frac{1}{2}L^2 + cL(\frac{L-1}{2}))|P|$ edges.*

Proof. Suppose there exists a spanner H of G with an additive stretch $+(L-2)$ but has strictly less than $(\frac{1}{2}L^2+cL(\frac{L-1}{2}))|P|$ edges. By the pigeonhole principle there must exist a critical pair $(x,y) \in P$ with the unique shortest path $P_{(x,y)}$ that is in one of the following two cases: (1) H is missing an edge in $P_{(x,y)}$ introduced in the edge subdivision step, or (2) H is missing at least one critical edge from $P_{(x,y)}$ in at least half $((L-1)/2)$ of the copies of G_I along $P_{(x,y)}$.

Let $P'_{(x,y)}$ be a shortest path connecting x and y in H. If (1) holds, then $P'_{(x,y)}$ traverses at least two more subdivided edges than $P_{(x,y)}$ and at least the same number of copies of G_I; hence $|P'_{(x,y)}| \ge |P_{(x,y)}| + L$, which is a contradiction.

Suppose now (2) holds; let $Q'_{(x,y)}$ on G_0 be the result of $P'_{(x,y)}$ by contracting inner graphs on H back into vertices and by replacing all subdivided edges back into single edges. Let $Q_{(x,y)}$ be the result of $P_{(x,y)}$ through the same process. Depending on whether $Q_{(x,y)} \neq Q'_{(x,y)}$, we have two cases:

- If $Q_{(x,y)} \neq Q'_{(x,y)}$, then since $Q_{(x,y)}$ is the unique shortest path on G_0 by Lemma 3.8, $|Q'_{(x,y)}| \ge |Q_{(x,y)}| + 2$. This implies that $|P'_{(x,y)}| \ge |P_{(x,y)}| + L$, which is a contradiction.

- Otherwise, $Q_{(x,y)} = Q'_{(x,y)}$. In this case for every inner graph that has a missing edge on $P_{(x,y)}$, $P'_{(x,y)}$ traverses at least two more edges. Since there are at least $(L-1)/2$ such inner graphs, $|P'_{(x,y)}| \geq |P_{(x,y)}| + (L-1)$. In either case, $P'_{(x,y)}$ has at least $+(L-1)$ additive stretch so H cannot be a $+(L-2)$ spanner. ∎

Proof of Theorem 3.6. Given the density parameter c_0, we will choose a larger parameter $c = \Theta(c_0)$ (defined precisely below) and construct the inner graph G_I (Lemma 3.7) with at most qL vertices, at least cqL edges, and q critical pairs, with $q = \Theta(L^2c^6)$. Once q, L are fixed, we construct the outer graph G_0 using Lemma 3.8. After the replacement and subdivision steps, G has $|V| = \Theta(L^5q^3)$ vertices and $|P| = \Theta(L^3q^3)$ critical pairs. This implies that there is some absolute constant $\lambda > 1$ such that $\lambda L(\frac{L-1}{2})|P| \geq |V|$.

Now, by Lemma 3.9, any spanner of G with at most $cL(\frac{L-1}{2})|P|$ edges has additive stretch at least $+(L-1)$. We choose $c = \lambda c_0$, so $cL(\frac{L-1}{2})|P| \geq c_0|V|$. Therefore, any spanner of G with at most $c_0|V|$ edges has an additive stretch of at least $+(L-1)$. Since $q = \Theta(L^2c^6) = \Theta(L^2c_0^6)$, it follows that $|V| = \Theta(L^{11}c_0^{18})$. Thus, we conclude that any spanner of G with c_0n edges has an additive stretch $+\Omega(|V|^{1/11}c_0^{-18/11})$. ∎

Remark 3.10. It follows from Theorem 3.6 that any $\Theta(n^{1+\epsilon})$-size spanner has $+\Omega(n^{\frac{1}{11}-\frac{18}{11}\epsilon})$ stretch, which is only better than the $+\Omega(n^{\frac{1}{13}-\frac{5}{13}\epsilon})$ bound of Corollary 3.5 when $\epsilon < 2/181$ is quite small.

Remark 3.11. The construction from Theorem 3.6 cannot be easily translated into an emulator lower bound. The reason is that the number of critical pairs is always sublinear in the number of vertices. A distance-preserving emulator of linear size always exists in this type of construction.

3.3. $O(n)$-size emulators. The difference between emulators and spanners is that emulators can use weighted edges not present in the original graph. In this section, our lower bound graph G is constructed exactly as in section 3.1 but with different numerical parameters.

THEOREM 3.12. *There exists an undirected graph G with n vertices such that any emulator with $O(n)$ edges has $+\Omega(n^{1/18})$ additive stretch.*

Before proving Theorem 3.12 we first argue that the size of low-stretch emulators is tied to the number of critical pairs $|P|$ for G.

LEMMA 3.13. *Every emulator for G with additive stretch $+(2D-1)$ 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 $\mathrm{dist}_G(u,v)$. (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 3.4.

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 $(x,y) \in P$, $\mathrm{dist}_{H'}(x,y) = \mathrm{dist}_H(x,y)$, and H' is a spanner with at most $n + 2D|H|$ edges. In particular, it has at most $2D|H|$ biclique edges since each weighted edge in some $P_{(x,y)}$ contributes at most $2D$ biclique edges to H'. By Lemma 3.4, the number of biclique edges in H' is at least $D|P|$, hence $|H| \geq |P|/2$. ∎

Proof of Theorem 3.12. In order to get $|P| = \Omega(n)$, it suffices to set $\delta \geq D^2$. Now, we have

$$\begin{aligned} n &= \Theta\left(R^{2d} D^2 \delta\right) \\ &= \Omega\left((rD)^{2d} D^4\right) && (\delta \geq D^2) \\ &= \Omega\left(\left(\delta^{\frac{d+1}{d(d-1)}} D\right)^{2d} D^4\right) && (\text{by definition of } \delta) \\ &= \Omega\left(D^{4\frac{d+1}{d-1}+2d+4}\right) && (\delta \geq D^2). \end{aligned}$$

The exponent is minimized when $d = 3$. This implies $n = \Omega(D^{18})$, and hence $D = O(n^{1/18})$. These parameters can be achieved asymptotically by setting $D = \Theta(n^{1/18})$, $\delta = D^2$, $r = \Theta(n^{2/27})$, and $R = \Theta(n^{7/54})$. ∎

COROLLARY 3.14. *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 an additive stretch $+\Omega(n^{(1-\frac{3d+1}{d-1}\epsilon)/(4+2d+4\frac{d+1}{d-1})})$. In particular, by setting $d = 3$ the additive stretch lowerbound becomes $\Omega(n^{\frac{1}{18}-\frac{5}{18}\epsilon})$.*

Proof. The proof is the similar to the proof of Corollary 3.5. By setting

$$\delta = D^{(2+(2d+2)\epsilon)/\left(1-\frac{3d+1}{d-1}\epsilon\right)},$$

we have that

$$\begin{aligned} n &= \Theta\left(R^{2d} D^2 \delta\right) = \Theta\left(D^{\left(4+2d+4\frac{d+1}{d-1}\right)/\left(1-\frac{3d+1}{d-1}\epsilon\right)}\right), \text{ and} \\ |P| &= \Theta\left(R^{2d} \delta^2\right) = \Theta\left(D^{\left(4+2d+4\frac{d+1}{d-1}+\left(4+2d+4\frac{d+1}{d-1}\right)\epsilon\right)/\left(1-\frac{3d+1}{d-1}\epsilon\right)}\right). \end{aligned}$$

Now we do have $|P| = \Theta(n^{1+\epsilon})$ which completes the proof. ∎

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 3.15. *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 an additive error $+\tilde{\Omega}(n^{1/18})$.*

Proof. For each subset $T \subseteq P$ construct the graph G_T by removing all biclique edges from G that are on the critical paths of pairs in T. Because all biclique 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 $T \backslash T'$. Since $\mathrm{dist}_{G_T}(x, y) \geq \mathrm{dist}_{G_{T'}}(x, y) + 2D$, the additive stretch of any such scheme must be at least $2D$. Alternatively, any scheme with stretch $2D - 1$ must use bitstrings of a length of at least length $|P|$.

Now, by setting $d = 3$ with $D = \tilde{\Theta}(n^{1/18})$, $r = \tilde{\Theta}(n^{2/27})$, and $R = \tilde{\Theta}(n^{7/54})$, 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, 14, 2, 19], 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 3.12 and 3.15 hold for $d = 3$, Theorems 2.1, 2.8, and 3.2 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 [5] 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 2.1, 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: $[n^{1/11}, n^{3/7}]$ 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.

Acknowledgments. Thanks to Greg Bodwin for inspiring the authors which resulted in improving sparse additive spanner lower bounds in section 3.2. We thank anonymous reviewers for their helpful comments and suggestions improving the clarity of this manuscript.

REFERENCES

[1] A. Abboud and G. Bodwin, *The 4/3 additive spanner exponent is tight*, J. ACM, 64 (2017), 28.

[2] A. Abboud, G. Bodwin, and S. Pettie, *A hierarchy of lower bounds for sublinear additive spanners*, SIAM J. Comput., 47 (2018), pp. 2203–2236.

[3] D. Aingworth, C. Chekuri, P. Indyk, and R. Motwani, *Fast estimation of diameter and shortest paths (without matrix multiplication)*, SIAM J. Comput., 28 (1999), pp. 1167–1181.

[4] N. Alon, *Testing subgraphs in large graphs*, Random Struct. Algorithms, 21 (2002), pp. 359–370.

[5] I. Bárány and D. G. Larman, *The convex hull of the integer points in a large ball*, Math. Ann., 312 (1998), pp. 167–181.

[6] S. Baswana, T. Kavitha, K. Mehlhorn, and S. Pettie, *Additive spanners and (α, β)-spanners*, ACM Trans. Algorithms, 7 (2010), 5.

[7] F. Behrend, *On sets of integers which contain no three terms in arithmetic progression*, Proc. Natl. Acad. Sci. USA, 32 (1946), pp. 331–332.

[8] G. Bodwin, *Linear size distance preservers*, in Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), Barcelona, Spain, 2017, pp. 600–615.

[9] G. Bodwin, *A Note on Distance-Preserving Graph Sparsification*, preprint, arXiv:2001.07741, 2020, https://arxiv.org/abs/2001.07741.

[10] G. Bodwin and V. V. Williams, *Very sparse additive spanners and emulators*, in Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science (ITCS), Rehovot, Israel, 2015, pp. 377–382.

[11] G. Bodwin and V. V. Williams, *Better distance preservers and additive spanners*, in Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), Arlington, VA, 2016.

[12] N. Cao, J. T. Fineman, and K. Russell, *Efficient construction of directed hopsets and parallel approximate shortest paths*, in Proccedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, Chicago, IL, 2020.

[13] S. Chechik, *New additive spanners*, in Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2013, New Orleans, LA, pp. 498–512.

[14] D. Coppersmith and M. Elkin, *Sparse sourcewise and pairwise distance preservers*, SIAM J. Discrete Math., 20 (2006), pp. 463–501.

[15] D. Dor, S. Halperin, and U. Zwick, *All-pairs almost shortest paths*, SIAM J. Comput., 29 (2000), pp. 1740–1759.

[16] M. Elkin and O. Neiman, *Near-additive spanners and near-exact hopsets, a unified view*, Bull. Eur. Assoc. Theor. Comput. Sci. EATCS, 130 (2020) pp. 1–24.

[17] M. Elkin and D. Peleg, *$(1+\epsilon,\beta)$-spanner constructions for general graphs*, SIAM J. Comput., 33 (2004), pp. 608–631.

[18] J. T. Fineman, *Nearly work-efficient parallel algorithm for digraph reachability*, in STOC 2018: Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, Los Angeles, CA, June 25-29, 2018, pp. 457–470.

[19] W. Hesse, *Directed graphs requiring large numbers of shortcuts*, in Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), Baltimore, MD, 2003.

[20] S. Huang and S. Pettie, *Thorup-Zwick emulators are universally optimal hopsets*, Inform. Process. Lett., 142 (2019), pp. 9–13.

[21] M. B. T. Knudsen, *Additive spanners: A simple construction*, in Proceedings of the Scandinavian Workshop on Algorithm Theory (SWAT), Copenhagen, Denmark, 2014, pp. 277–281.

[22] Y. P. Liu, A. Jambulapati, and A. Sidford, *Parallel reachability in almost linear work and square root depth*, in Proceedings of the 60th IEEE Annual Symposium on Foundations of Computer Science (FOCS), Baltimore, MD, 2019.

[23] S. Pettie, *Low distortion spanners*, ACM Trans. Algorithms, 6 (2009), p. 7.

[24] M. Thorup, *On shortcutting digraphs*, in International Workshop on Graph-Theoretic Concepts in Computer Science, Springer, 1992, pp. 205–211.

[25] M. Thorup, *Shortcutting planar digraphs*, Combin. Probab. Comput., 4 (1995), pp. 287–315.

[26] M. Thorup and U. Zwick, *Spanners and emulators with sublinear distance errors*, in Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), Miami, FL, 2006, pp. 802–809.

[27] D. P. Woodruff, *Lower bounds for additive spanners, emulators, and more*, in Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS), Berkeley, CA, 2006, pp. 389–398.