

A HIERARCHY OF LOWER BOUNDS FOR SUBLINEAR ADDITIVE SPANNERS*

AMIR ABOUD[†], GREG BODWIN[†], AND SETH PETTIE[‡]

Abstract. Spanners, emulators, and approximate distance oracles can be viewed as *lossy* compression schemes that represent an unweighted graph metric in small space, say $\tilde{O}(n^{1+\delta})$ bits. There is an inherent tradeoff between the sparsity parameter δ and the *stretch function* f of the compression scheme, but the qualitative nature of this tradeoff has remained a persistent open problem. It has been known for some time that when $\delta \geq 1/3$ there are schemes with constant *additive stretch* (distance d is stretched to at most $f(d) = d + O(1)$), and recent results of Abboud and Bodwin show that when $\delta < 1/3$ there are no such schemes. Thus, to get practically efficient graph compression with $\delta \rightarrow 0$ we must pay superconstant additive stretch, but exactly how much do we have to pay? In this paper we show that the lower bound of Abboud and Bodwin is just the first step in a *hierarchy* of lower bounds that characterize the asymptotic behavior of the optimal stretch function f for sparsity parameter $\delta \in (0, 1/3)$. Specifically, for any integer $k \geq 2$, any compression scheme using $O(n^{1+\frac{1}{2^k-1}-\epsilon})$ bits has a *sublinear additive stretch function* $f: f(d) = d + \Omega(d^{1-\frac{1}{k}})$. This lower bound matches Thorup and Zwick’s (2006) construction of sublinear additive *emulators*. It also shows that Elkin and Peleg’s $(1+\epsilon, \beta)$ -spanners have an essentially optimal tradeoff between δ , ϵ , and β , and that the sublinear additive spanners of Pettie (2009) and Chechik (2013) are not too far from optimal. To complement these lower bounds we present a new construction of $(1+\epsilon, O(k/\epsilon)^{k-1})$ -spanners with size $O((k/\epsilon)^{h_k} kn^{1+\frac{1}{2^{k+1}-1}})$, where $h_k < 3/4$. This size bound improves on the spanners of Elkin and Peleg (2004), Thorup and Zwick (2006), and Pettie (2009). According to our lower bounds neither the size nor stretch function can be substantially improved. Our lower bound technique exhibits several interesting degrees of freedom in the framework of Abboud and Bodwin. By carefully exploiting these freedoms, we are able to obtain lower bounds for several related combinatorial objects. We get lower bounds on the size of (β, ϵ) -hopsets, matching Elkin and Neiman’s construction (2016), and lower bounds on *shortcutting sets* for digraphs that preserve the transitive closure. Our lower bound simplifies Hesse’s (2003) refutation of Thorup’s conjecture (1992), which stated that adding a linear number of shortcuts suffices to reduce the diameter to polylogarithmic. Finally, we show matching upper and lower bounds for graph compression schemes that work for graph metrics with girth at least $2\gamma + 1$. One consequence is that Baswana et al.’s (2010) additive $O(\gamma)$ -spanners with size $O(n^{1+\frac{1}{2\gamma+1}})$ cannot be improved in the exponent.

Key words. graph theory, spanners, lower bounds, hopsets, graph sketching

AMS subject classifications. 68R10, 68R05

DOI. 10.1137/16M1105815

1. Introduction. *Spanners* [48], *emulators* [27, 59], and *approximate distance oracles* [58] can be viewed as kinds of compression schemes that approximately encode the distance metric of a (dense) undirected input graph $G = (V, E)$ in small space, where the notion of *approximation* is captured by a nondecreasing *stretch function* $f: \mathbb{N} \rightarrow \mathbb{N}$.

Spanners. An $f(d)$ -*spanner* $G' = (V, E')$ is a subgraph of G for which $\text{dist}_{G'}(u, v)$

*Received by the editors November 30, 2016; accepted for publication (in revised form) September 14, 2018; published electronically December 4, 2018. An extended abstract of this paper was presented at the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2017), in Barcelona, Spain.

<http://www.siam.org/journals/sicomp/47-6/M110581.html>

Funding: This research was supported by NSF grants CCF-1217338, CNS-1318294, CCF-1417238, CCF-1514339, CCF-1514383, CCF-1637546, and BSF grant 2012338.

[†]Department of Computer Science, Stanford University, Stanford, CA 94305 (abboud@cs.stanford.edu, gbodwin@cs.stanford.edu).

[‡]EECS, University of Michigan, Ann Arbor, MI 48109 (pettie@umich.edu).

is at most $f(\text{dist}_G(u, v))$. An (α, β) -spanner is one with stretch function $f(d) = \alpha d + \beta$. Notable special cases include *multiplicative α -spanners* [48, 8, 31, 58, 10, 9], when $\beta = 0$, and *additive β -spanners* [6, 27, 31, 59, 9, 64, 20, 40], when $\alpha = 1$. See [31, 9, 59, 49, 20, 45] for “mixed” spanners with $\alpha > 1, \beta > 0$.

Emulators. An $f(d)$ -emulator (also called a *Steiner spanner* [8]) is a *weighted* graph $G' = (V' \supseteq V, E', w')$ such that for each $u, v \in V$, $\text{dist}_{G'}(u, v) \in [\text{dist}_G(u, v), f(\text{dist}_G(u, v))]$. In other words, one is allowed to add Steiner points ($V' \setminus V$) and long-range (weighted) edges $(u, v) \in E' \setminus E$ such that distances are non-contracting.

(Unconstrained) distance oracles. For our purposes, an $f(d)$ -approximate distance oracle using space s is a bit string in $\{0, 1\}^s$ such that given $u, v \in V$, an estimate $\widetilde{\text{dist}}(u, v) \in [\text{dist}_G(u, v), f(\text{dist}_G(u, v))]$ can be computed by examining only the bit string. Note: the term “oracle” was used in [58] to indicate that $\widetilde{\text{dist}}(u, v)$ is computed in constant time [46, 3, 21]. Later work considered distance oracles with nonconstant query time [50, 5, 4, 32]. In this paper we make no restrictions on the query time at all. Thus, for our purposes distance oracles generalize spanners, emulators, and related objects.

In this paper we establish essentially optimal tradeoffs between the size of the compressed graph representation and the asymptotic behavior of its stretch function f . In order to put our results in context we must recount the developments of the last 30 years that investigated multiplicative, additive, (α, β) , and sublinear additive stretch functions.

1.1. Multiplicative stretch. Historically, the first notion of stretch studied in the literature was purely multiplicative stretch. Althöfer et al. [8] quickly settled the problem by showing that any graph contains an α -spanner with at most $m_{\alpha+2}(n)$ edges, and that the claim is false for $m_{\alpha+2}(n) - 1$. Here $m_g(n)$ is the maximum number of edges in a graph with n vertices and girth g . The upper bound of [8] follows directly from the observation that a natural greedy construction never closes a cycle with length at most $\alpha + 1$; the lower bound follows from the fact that no strict subgraph of a graph with girth $\alpha + 2$ is an α -spanner.¹ It has been conjectured [33, 17, 15] that the trivial upper bound $m_{2k+1}(n), m_{2k+2}(n) = O(n^{1+1/k})$ is sharp up to the leading constant, but this *Girth Conjecture* has only been proved for $k = 1$ (trivial) and $k \in \{2, 3, 5\}$ [18, 34, 51, 62, 60, 12, 41]. See [42, 43, 63] for lower bounds on $m_g(n)$.

1.2. Additive stretch. The Girth Conjecture implies that a spanner with size $O(n^{1+1/k})$ must stretch some pair of adjacent vertices at original distance $d = 1$ to distance $2k - 1$. If “stretch” is defined a priori to be multiplicative, then such $(2k - 1)$ -spanners are optimal. However, there is no reason to believe that $f(d) = (2k - 1)d$ is an optimal stretch *function* for size $O(n^{1+1/k})$. The girth argument could also be interpreted as lower bounding additive stretch or (α, β) -stretch. In general, the Girth Conjecture only implies that (α, β) -spanners with size $O(n^{1+1/k})$ have $\alpha + \beta \geq 2k - 1$.

Aingworth et al. [6] gave a construction of an additive 2-spanner with size $\tilde{O}(n^{3/2})$, which is optimal in the sense that neither the additive stretch 2 nor exponent $3/2$ can be unilaterally improved.² This result raised the tantalizing possibility that there

¹Removing any edge stretches the distance between its endpoints from 1 to at least $\alpha + 1$. Moreover, since every graph contains a bipartite subgraph with at least half the edges, $m_{2k+1} \leq 2m_{2k+2}(n)$ for every k . Thus, there are $(2k - 1)$ -spanners with size $O(m_{2k+2}(n))$.

²Moreover, later results of Bollobás, Coppersmith, and Elkin [16] show that for spanner size $O(n^{3/2})$, the stretch *function* $f(d) = d + 2$ is optimal for $1 \leq d \leq \Theta(\sqrt{n})$. See [31, 59, 9, 40] for constructions of additive 2-spanners with size $O(n^{3/2})$.

exist arbitrarily sparse additive spanners. Dor, Halperin, and Zwick [27] observed that additive 4-emulators exist with size $\tilde{O}(n^{4/3})$; i.e., the emulator introduces weighted edges connecting distant vertex pairs. Baswana et al. [9] constructed additive 6-spanners with size $O(n^{4/3})$, and Chechik [20] constructed additive-4 spanners with size $\tilde{O}(n^{7/5})$. See [64, 40, 31, 59, 27, 9] for other constructions of additive 2- and 6-spanners.

The “4/3” exponent proved to be very resilient for both emulators and spanners with additive stretch. This led to a line of work establishing additive spanners below the $n^{4/3}$ threshold with stretch polynomial in n [16, 9, 49, 20, 13]. The additive spanners of Bodwin and Vassilevska Williams [14] with stretch function $f(d) = d + n^\epsilon$ have size that is the minimum of $O(n^{\frac{4}{3} - \frac{7\epsilon}{9} + o(1)})$ and $O(n^{\frac{5}{4} - \frac{5\epsilon}{12} + o(1)})$.

1.3. Sublinear additive stretch. Elkin and Peleg [31] showed that the “4/3 barrier” could also be broken by tolerating $1 + \epsilon$ multiplicative stretch. In particular, for any integer κ and real $\epsilon > 0$, there are $(1 + \epsilon, \beta)$ -spanners with size $O(\beta n^{1+1/\kappa})$, where $\beta = O(\epsilon^{-1} \log \kappa)^{\log \kappa}$. The construction algorithm and size-bound both depend on ϵ . Thorup and Zwick [59] gave a surprisingly simple construction of an $O(kn^{1+\frac{1}{2k+1-1}})$ -size emulator with $(1 + \epsilon, O(k/\epsilon)^{k-1})$ -type stretch.

Thorup and Zwick’s emulator has the special property that its stretch holds for every $\epsilon > 0$ *simultaneously*; i.e., it can be selected as a function of d . Judiciously choosing $\epsilon = k/d^{\frac{1}{k}}$ leads to an emulator with a *sublinear additive stretch* function $f(d) = d + O(kd^{1-\frac{1}{k}} + 3^k)$.³ Thorup and Zwick also showed that this same stretch function also applies to their earlier [58] construction of multiplicative $(2k + 1)$ -spanners with size $O(kn^{1+\frac{1}{k+1}})$. Pettie [49] gave a construction of sublinear additive spanners whose size-stretch tradeoff is closer to the Thorup–Zwick emulators. For stretch function $d + O(kd^{1-\frac{1}{k}} + 3^k)$ the size is $O(kn^{1+\frac{(3/4)^{k-2}}{7-2(3/4)^{k-2}}})$, which is always $o(n^{1+(3/4)^{k+3}})$ for any fixed k . At their sparsest, Thorup and Zwick’s emulators [59] and Pettie’s spanners [49] have size $O(n \log \log n)$ and stretch $f(d) = d + O(\log \log n) \cdot d^{1-\Theta(1/\log \log n)} + (\log n)^{\log_2 3}$. Pettie [49] gave an even sparser $(1 + \epsilon, O(\epsilon^{-1} \log \log n)^{\log \log n})$ -spanner with size $O(n \log \log(\epsilon^{-1} \log \log n))$.

1.4. Lower bounds. Woodruff proved that any $k^{-1}n^{1+1/k}$ -size spanner with stretch function f must have $f(k) \geq 3k$. As a corollary, additive $(2k - 2)$ -spanners must have size $\Omega(k^{-1}n^{1+1/k})$, independent of the status of the Girth Conjecture. Bollobás, Coppersmith, and Elkin [16] showed that if the stretch f is such that $f(d) = d$ for $d \geq D$, then $\Omega(n^2/D)$ -size is necessary and sufficient for spanners and emulators.

In a recent surprise, Abboud and Bodwin [2] proved that no additive β -spanners, emulators, or distance oracles exist with $\beta = O(1)$ and exponent less than 4/3. More precisely, any construction of these three objects with additive $\beta = O(1)$ stretch has size $\Omega(n^{4/3}/2^{O(\sqrt{\log n})})$ and any construction with size $O(n^{4/3-\epsilon})$ has additive stretch $\beta = n^\delta$ for some $\delta = \delta(\epsilon)$. This result explained why all prior additive spanner constructions had a strange transition at 4/3 [27, 59, 9, 20, 14, 40, 64], but it

³The Thorup–Zwick emulator can easily be converted to a $(1 + \epsilon, \beta)$ -spanner by replacing weighted edges with paths up to length β . A careful analysis shows the size of the resulting spanner can be made $O((k/\epsilon)^{O(1)} n^{1+\frac{1}{2k+1-1}})$ (see section 3) which would slightly improve on [31]. Elkin [28] has stated that with minor changes, the Elkin–Peleg [31] spanners can also be expressed as $(1 + \epsilon, O(k/\epsilon)^{k-1})$ -spanners with size $O((k/\epsilon)^{O(1)} n^{1+\frac{1}{2k+1-1}})$. We state these bounds in Figure 1 rather than those of [31] in order to facilitate easier comparisons with subsequent constructions [59, 20, 49], and the new constructions of section 3.

Stretch Function

Citation		$d + O(\sqrt{d})$	$d + O(d^{\frac{2}{3}})$	$d + O(d^{\frac{3}{4}})$	$d + O(kd^{1-\frac{1}{k}})$
		or $(1 + \epsilon, O(\frac{1}{\epsilon}))$	or $(1 + \epsilon, O(\frac{1}{\epsilon})^2)$	or $(1 + \epsilon, O(\frac{1}{\epsilon})^3)$	or $(1 + \epsilon, O(\frac{k}{\epsilon})^{k-1})$
Elkin & Peleg	SPAN.	$O(\epsilon^{-O(1)} n^{\frac{8}{7}})$	$O(\epsilon^{-O(1)} n^{\frac{16}{15}})$	$O(\epsilon^{-O(1)} n^{\frac{32}{31}})$	$O\left(\left(\frac{k}{\epsilon}\right)^{O(1)} n^{1+\frac{1}{2^{k+1}-1}}\right)$
Thorup & Zwick	EMUL.	$O(n^{\frac{8}{7}})$	$O(n^{\frac{16}{15}})$	$O(n^{\frac{32}{31}})$	$O(kn^{1+\frac{1}{2^{k+1}-1}})$
	SPAN.	$O(n^{\frac{4}{3}})$	$O(n^{\frac{5}{4}})$	$O(n^{\frac{6}{5}})$	$O(kn^{1+\frac{1}{k+1}})$
Pettie	SPAN.	$O(n^{\frac{6}{5}})$	$O(n^{\frac{25}{22}})$	$O(n^{\frac{103}{94}})$	$O\left(kn^{1+\frac{(3/4)^{k-2}}{7-2(3/4)^{k-2}}}\right)$
Chechik	SPAN.	$\tilde{O}(n^{\frac{20}{17}})$			
New	SPAN.	$O(\epsilon^{-\frac{2}{7}} n^{\frac{8}{7}})$	$O(\epsilon^{-\frac{7}{15}} n^{\frac{16}{15}})$	$O(\epsilon^{-\frac{18}{31}} n^{\frac{32}{31}})$	$O\left(\left(\frac{k}{\epsilon}\right)^h kn^{1+\frac{1}{2^{k+1}-1}}\right)$
New Lower Bounds	ALL	$\Omega(n^{\frac{4}{3}-o(1)})$	$\Omega(n^{\frac{5}{4}-o(1)})$	$\Omega(n^{\frac{16}{15}-o(1)})$	$\Omega(n^{1+\frac{1}{2^{k+1}-1}-o(1)})$

FIG. 1. A summary of spanners and emulators with $(1+\epsilon, O(k/\epsilon)^{k-1})$ -type stretch and sublinear additive stretch $d + O(kd^{1-\frac{1}{k}})$. Constructions with the first type of stretch also have sublinear additive stretch $d + O(kd^{1-\frac{1}{k}})$, but only for d up to about $(1/\epsilon)^k$. Note: The new lower bounds do not contradict the upper bounds; the lower bounds are for stretch functions with smaller leading constants in the $O(k/\epsilon)^{k-1}$ and $O(kd^{1-\frac{1}{k}})$ terms. In the last cell of the table, $h = \frac{3 \cdot 2^{k-1} - (k+2)}{2^{k+1}-1} < 3/4$, which improves the dependence on ϵ that can be obtained from modified versions of existing constructions [31, 59].

did not suggest what the *optimal* stretch function should be for sparsity $n^{1+\delta}$ when $\delta \in [0, 1/3)$.

1.5. New results. See Figure 1 for a comparison of these results to prior work.

Spanner/emulator/distance oracle lower bounds. Our main result is a hierarchy of lower bounds for spanners, emulators, and distance oracles, which shows that tradeoffs offered by Thorup and Zwick's sublinear additive emulators [59] and Elkin and Peleg's $(1 + \epsilon, \beta)$ -spanners cannot be substantially improved. Building on Abboud and Bodwin's [2] $\Omega(n^{4/3-o(1)})$ lower bounds for additive spanners, we prove that for every integer $k \geq 2$ and $d < n^{o(1)}$, there is a graph \mathcal{H}_k on n vertices and $n^{1+\frac{1}{2^k-1}-o(1)}$ edges such that any spanner with size $n^{1+\frac{1}{2^k-1}-\epsilon}$, $\epsilon > 0$, stretches vertices at distance d to at least $d + c_k d^{1-\frac{1}{k}}$ for a constant $c_k = \Theta(1/k)$. More generally, we exhibit graph families that cannot be *compressed* into distance oracles on $n^{1+\frac{1}{2^k-1}-\epsilon}$ bits such that distances can be recovered below this error threshold. The consequences of this construction are that the existing sublinear additive emulators [59], sublinear additive spanners [49, 20], and $(1 + \epsilon, \beta)$ -spanners [31, 59, 49] are, to varying degrees, close to optimal. Specifically:

- The $(d + O(kd^{1-\frac{1}{k}} + 3^k))$ -emulator [59] with size $O(n^{1+\frac{1}{2^{k+1}-1}})$ cannot be improved by more than a constant factor (roughly k^2) in the stretch $O(kd^{1-\frac{1}{k}})$, or by a $o(1)$ in the exponent $1 + \frac{1}{2^{k+1}-1}$. (Our results say nothing about the optimality of the 3^k term when d is small.)
- The sublinear additive *spanners* of Pettie [49] and Chechik [20] probably

have suboptimal exponents, but not by much. For example, the exponent of Chechik’s [20] $\tilde{O}(n^{20/17})$ -size $(d + O(\sqrt{d}))$ -spanner is within 0.034 of optimal and the exponent of Pettie’s [49] $O(n^{25/22})$ -size $(d + O(d^{2/3}))$ -spanner is within 0.07 of optimal.

- When $\epsilon \geq 1/n^{o(1)}$, the existing constructions of $(1 + \epsilon, O(k/\epsilon)^{k-1})$ -spanners [31, 59, 49] with size $O((k/\epsilon)^{O(1)} n^{1 + \frac{1}{2k+1-1}})$ cannot be substantially improved in either the additive $O(k/\epsilon)^{k-1}$ term or the exponent $1 + \frac{1}{2k+1-1}$. This follows from the fact that any spanner with stretch of type $(1 + \hat{\epsilon}, O(k/\hat{\epsilon})^{k-1})$ for every $\hat{\epsilon} \geq \epsilon$ functions as a $(d + O(kd^{1-\frac{1}{k}}))$ -spanner for distances $d \leq O(k/\epsilon)^k$. However, there is no reason to believe that the *size* of such $(1 + \epsilon, \beta)$ -spanners must depend on ϵ , as it does in the current constructions.

There is an interesting new hierarchy of *phase transitions* in the interplay between our lower bounds and previous upper bounds [59]. Let C be a sufficiently large constant and c be a sufficiently small constant. If one wants a graph compression scheme with stretch $f(d) = d + C\sqrt{d}$, then one needs only $\tilde{O}(n^{8/7})$ bits of space to store an emulator [59]. However, if we want a slightly improved stretch $f(d) = d + c\sqrt{d}$, then, by our lower bound, the space requirement leaps to $\Omega(n^{4/3-o(1)})$. In general, the optimal space for stretch function $f(d) = d + c'd^{1-1/k}$ takes a polynomial jump as we shift c' from some sufficiently large constant $O(k)$ to a sufficiently small constant $\Omega(1/k)$.

An important take-away message from our work is that the sublinear additive stretch functions of type $f(d) = d + O(d^{1-1/k})$ used by Thorup and Zwick [59] are exactly of the “right” form. For example, such plausible-looking stretch functions as $f(d) = d + O(d^{1/3})$ and $f(d) = d + O(d^{2/3}/\log d)$ could only exist in the narrow bands not covered by our lower bounds: between space $n^{4/3-o(1)}$ and $n^{4/3}$ and between space $n^{8/7-o(1)}$ and $n^{8/7}$.

Spanner upper bounds. To complement our lower bounds we provide new upper bounds on the sparsity of spanners with stretch of type $(1 + \hat{\epsilon}, O(k/\hat{\epsilon})^{k-1})$, which holds for every $\hat{\epsilon} \geq \epsilon$. Our new spanners have size $O((k/\epsilon)^h kn^{1 + \frac{1}{2k+1-1}})$, where $h = \frac{3 \cdot 2^{k-1} - (k+2)}{2k+1-1} < 3/4$. This construction improves on the bounds that can be derived from [59, 31, 49] in the dependence on ϵ .⁴ For example, suppose we want something that behaves like a $(d + O(\sqrt{d}))$ -spanner for all d up to some threshold D . Our construction shows that size $O(D^{1/7} n^{8/7})$ suffices, which improves on Chechik’s $(d + O(\sqrt{d}))$ -spanner so long as $D < n^{4/17}$.

Hopset lower bounds. Hopsets are fundamental objects that are morally similar to emulators. They were explicitly defined by Cohen [23] but used implicitly in many earlier works [61, 39, 22, 53]. Let $G = (V, E, w)$ be an arbitrary undirected *weighted* graph and $H \subset \binom{V}{2}$ be a set of edges called the *hopset*. In the united graph $G' = (V, E \cup H, w)$, the weight of an edge $(u, v) \in H$ is the length of the shortest path in G between u and v . Define the β -limited distance in G' , denoted $\text{dist}_{G'}^{(\beta)}(u, v)$, to be the length of the shortest path from u to v that uses at most β edges in G' .⁵ We

⁴No bounds of this type are stated explicitly in [59] or [31]. In order to get a bound of this type—with the $1 + \frac{1}{2k+1-1}$ exponent and some $\text{poly}(1/\epsilon)$ dependence on ϵ —one must only adjust the sampling probabilities of [59]; however, adapting [31] requires slightly more significant changes [28].

⁵Note that whereas $\text{dist} = \text{dist}^{(\infty)}$ is metric, $\text{dist}^{(\beta)}$ does not necessarily satisfy the triangle inequality for finite β .

call H a (β, ϵ) -hopset, where $\beta \geq 1, \epsilon > 0$, if, for any $u, v \in V$, we have

$$\text{dist}_{G'}^{(\beta)}(u, v) \leq (1 + \epsilon) \text{dist}_G(u, v).$$

There is clearly some three-way tradeoff between β , ϵ , and $|H|$. Elkin and Neiman [29] recently showed that any graph has a (β, ϵ) -hopset with size $\tilde{O}(n^{1+1/\kappa})$, where $\beta = O(\frac{\log \kappa}{\epsilon})^{\log \kappa}$. (In this problem “ κ ” is not a natural parameter. It is likely that Elkin and Neiman’s tradeoff could be more precisely stated as follows: for any positive integer k and real $\epsilon > 0$, there is an $\tilde{O}(n^{1+\frac{1}{2k+1-1}})$ size (β, ϵ) -hopset with $\beta = O(k/\epsilon)^k$. Huang and Pettie [36] showed that Thorup and Zwick’s emulator [59] is actually a (β, ϵ) -hopset with size $O(n^{1+\frac{1}{2k+1-1}})$ and $\beta = O(k/\epsilon)^k$.)

In this work we show that any construction of (β, ϵ) -hopsets with worst-case size $n^{1+\frac{1}{2k-1}-\delta}$, where $k \geq 1$ is an integer and $\delta > 0$, must have $\beta = \Omega(\frac{1}{\epsilon})^k$.⁶ For example, hopsets with $\beta = o(1/\epsilon)$ must have size $\Omega(n^{2-o(1)})$ and those with $\beta = o(1/\epsilon^2)$ must have size $\Omega(n^{4/3-o(1)})$. This essentially matches the tradeoff of [29, 36] up to a constant in β that depends on k .

Lower bounds on shortcutting digraphs. In 1992, Thorup [55] conjectured that the diameter of any directed graph $G = (V, E)$ could be drastically reduced with a small number of *shortcuts*. In particular, there exists another directed graph $G' = (V, E')$ with $|E'| = O(|E|)$ and the same transitive closure relation as G (\rightsquigarrow), such that if $u \rightsquigarrow v$, then there is a poly($\log n$)-length path from u to v in G' . Thorup’s conjecture was confirmed for trees [55, 57, 19] and planar graphs [56] but finally refuted by Hesse [35] for general graphs. In this paper we give a simpler 1-page proof of Hesse’s refutation by modifying our spanner lower bound construction. In particular, we prove that for any $\epsilon > 0$, there is a digraph with size $n^{1+\epsilon}$ and diameter n^δ , $\delta = \delta(\epsilon)$, such that reducing the diameter to $o(n^\delta)$ requires at least $n^{2-\epsilon}$ shortcuts.

Spanners for high-girth graphs. Our lower bounds apply to the class of *all* undirected graph metrics. Baswana et al. [9] gave sparser spanners for a *restricted* class of graph metrics. Specifically, graphs with girth at least $2\gamma + 1$ contain additive 6γ -spanners with size $O(n^{1+\frac{1}{2\gamma+1}})$. We adapt our lower bound construction to prove that the exponent $1 + \frac{1}{2\gamma+1}$ is optimal, assuming the Girth Conjecture, and more generally we give lower bounds on compression schemes for the class of graphs with girth at least $2\gamma + 1$. Any scheme that uses $n^{1+\frac{1}{(\gamma+1)2^{k-1}-1}-\epsilon}$ bits must have stretch $f(d) \geq d + \Omega(d^{1-1/k})$ for any $d < n^{o(1)}$. We also give new constructions of emulators and spanners for girth- $(2\gamma + 1)$ graphs that shows that the exponent $1 + \frac{1}{(\gamma+1)2^{k-1}-1}$ is the best possible.

1.6. Related work. Much of the recent work on spanners has focused on preserving or approximating distances between specified *pairs* of vertices. See [25, 1, 2] for lower bounds on pairwise spanners and [25, 49, 26, 38, 37, 1, 45, 52] for upper bounds. Pairwise spanners have proven to be useful tools for constructing (sublinear) additive spanners; see [49, 20, 13].

The space/stretch tradeoffs offered by the best distance oracles [21, 46, 47, 3, 5, 4, 32] are strictly worse than those of the best spanners and emulators, even though distance oracles are entirely *unconstrained* in how they encode the graph metric. This

⁶The leading constant hidden by the Ω is roughly $1/(k2^k)$ but can probably be made $1/\text{poly}(k)$ with a more careful analysis. Note, however, that because of the “ $-o(1)$ ” in the exponent, the lower bound is vacuous when k is not constant.

is primarily due to the requirement that distance oracles respond to queries quickly. There are both unconditional [54] and conditional [24, 46, 47] lower bounds, suggesting that distance oracles with reasonable query time cannot match the best spanners or emulators.

1.7. Organization. In section 2 we generalize Abboud and Bodwin’s construction [2] to give a spectrum of lower bounds against graph compression schemes with sublinear additive stretch and $(1 + \epsilon, \beta)$ -stretch. In section 3 we combine ideas from Thorup and Zwick’s emulators [59] and Pettie’s spanners [49] to attain a new bound on sparse $(1 + \epsilon, \beta)$ -spanners. In section 4 we prove tight bounds on (β, ϵ) -hopsets. In section 5 we generalize the construction of section 2 to give stretch-sparseness lower bounds on the class of graphs with girth at least $2\gamma + 1$. Matching upper bounds for graphs of girth $2\gamma + 1$ are given in section 5.1. In section 6 we give a simpler refutation of Thorup’s shortcutting conjecture. In section 7 we highlight some remaining open problems.

2. The lower bound construction. The graphs in this section are parameterized by an integer $\ell \geq 2$, which determines the length of the shortest paths that will be hardest to approximate. Each graph has a layered structure, consisting of a layer of *input ports*, some number of interior layers, and a layer of *output ports*. In any given graph construction, p is the number of input/output ports.

Sections 2.1 and 2.2 present the *first base graph* $\dot{B}[p]$ and *second base graph* $\ddot{B}[p]$, and section 2.3 proceeds to combine them recursively to form a hierarchy of graphs $(\mathcal{H}_k[p])$ parameterized by k . The graphs $\dot{B}[p]$, $\ddot{B}[p]$, and $\mathcal{H}_2[p]$ follow the construction of Abboud and Bodwin [2].

2.1. The first base graph. Let $\dot{B}[p] = (L_0 \cup \dots \cup L_\ell, E)$ be an $(\ell + 1)$ -layer graph with the following properties:

- $\dot{B}[p]$ has p vertices per layer, and all edges connect vertices in adjacent layers.
- Each edge e is assigned a label $\text{label}(e) \in \mathcal{L}[p]$. (In the construction we have in mind, the label-set $\mathcal{L}[p]$ is a relatively dense “average-free” subset of $\{1, \dots, p/\ell\}$; see below for details.) For any vertex u , the edges connecting u to the previous layer have distinct labels, and the edges connecting u to the subsequent layer have distinct labels.
- Let $\mathcal{P}(\dot{B}[p]) \subset L_0 \times L_\ell$ be a set of pairs of input/output ports. Each $(u_0, u_\ell) \in \mathcal{P}(\dot{B}[p])$ has the property that there exists a *unique* shortest path $(u_0, u_1, \dots, u_\ell)$. Moreover, $\text{label}(u_0, u_1) = \dots = \text{label}(u_{\ell-1}, u_\ell)$, which implies that any two of these paths are edge disjoint. The edge set E is precisely the union of these paths over all pairs in $\mathcal{P}(\dot{B}[p])$.

These properties imply that the number of vertices and edges in $\dot{B}[p]$ is $n[p] \stackrel{\text{def}}{=} (\ell + 1)p$ and $m[p] \stackrel{\text{def}}{=} |E(\dot{B}[p])| = \ell \cdot |\mathcal{P}(\dot{B}[p])|$.

Refer to [7, 2] for constructions of $\dot{B}[p]$ satisfying these requirements, or to [25] for a construction without the layered structure. For the sake of completeness we give a short sketch of how $\dot{B}[p]$ is constructed using average-free sets [7, 2]. Let $\mathcal{L}[p] \subset \{1, \dots, \lfloor p/\ell \rfloor\}$ be an ℓ -average-free set, i.e., one for which the equation

$$\ell \cdot x_0 = x_1 + x_2 + \dots + x_\ell, \quad \text{where } x_0, x_1, \dots, x_\ell \in \mathcal{L}[p],$$

has no solutions, except the trivial $x_0 = x_1 = \dots = x_\ell$. Let $u_{i,j}$ denote the j th vertex in L_i . The edge set consists of

$$E = \{(u_{i,j}, u_{i+1,j'}) \mid i \in [0, \ell) \text{ and } (j' - j) \bmod p \in \mathcal{L}[p]\},$$

with label $(u_{i,j}, u_{i+1,j'}) = (j' - j) \bmod p$. The pair set consists of

$$\mathcal{P}(\dot{B}[p]) = \{(u_{0,j}, u_{\ell,(j+\ell x) \bmod p}) \mid \text{for all } j \in \{0, \dots, p-1\} \text{ and } x \in \mathcal{L}[p]\}.$$

The ℓ -average-free property of $\mathcal{L}[p]$ ensures that $(u_{0,j}, u_{1,j+x}, u_{2,j+2x}, \dots, u_{\ell,j+\ell x})$ is the unique shortest path between its endpoints.

2.2. The second base graph. Roughly speaking, $\ddot{B}[p]$ is obtained by taking a certain product of two copies of $\dot{B}[\sqrt{p}]$.⁷ Let $L_0^0 \cup \dots \cup L_{\ell}^0$ and $L_0^1 \cup \dots \cup L_{\ell}^1$ be the vertex sets of copies $\dot{B}^0[\sqrt{p}]$ and $\dot{B}^1[\sqrt{p}]$, each with respective pair-sets \mathcal{P}^0 and \mathcal{P}^1 . $\ddot{B}[p]$ is a layered graph with vertex set $\ddot{L}_0 \cup \dots \cup \ddot{L}_{2\ell}$ where $\ddot{L}_i = L_{i/2}^0 \times L_{i/2}^1$ when i is even and $\ddot{L}_i = L_{\lceil i/2 \rceil}^0 \times L_{\lfloor i/2 \rfloor}^1$ when i is odd. Vertices in \ddot{B} are identified with vertex pairs from $V(\dot{B}^0) \times V(\dot{B}^1)$. When i is even, an edge $((u, v), (u', v'))$ exists between layers \ddot{L}_i and \ddot{L}_{i+1} iff $(u, u') \in E(\dot{B}^0)$. Similarly, when i is odd, an edge $((u, v), (u', v'))$ exists between layers \ddot{L}_i and \ddot{L}_{i+1} iff $(v, v') \in E(\dot{B}^1)$. An edge in \ddot{B} inherits the label of the corresponding edge in \dot{B} , so the label set for $\ddot{B}[p]$ is $\mathcal{L}[\sqrt{p}]$. The pair-set for \ddot{B} is defined to be

$$\mathcal{P}(\ddot{B}[p]) = \{((u_0, v_0), (u_{\ell}, v_{\ell})) \mid (u_0, u_{\ell}) \in \mathcal{P}^0 \text{ and } (v_0, v_{\ell}) \in \mathcal{P}^1\}.$$

Observe that any length- 2ℓ path from layer \ddot{L}_0 to $\ddot{L}_{2\ell}$ corresponds to picking edges alternately from two paths, one from L_0^0 to L_{ℓ}^0 in \dot{B}^0 and one from L_0^1 to L_{ℓ}^1 in \dot{B}^1 . Lemma 2.1 summarizes the relevant properties of \ddot{B} and $\mathcal{P}(\ddot{B})$.

LEMMA 2.1. *Let $\xi_{\ell}(p)$ be a nondecreasing function of p such that $|\mathcal{L}[p]| \geq p/\xi_{\ell}(p)$, $|\mathcal{L}[p]| \leq p/2$, and $|\mathcal{P}(\dot{B}[p])| \geq p^2/\xi_{\ell}(p)$. The graph $\ddot{B} = \ddot{B}[p]$ has the following properties.*

1. *It has $\ddot{n}[p] \leq (2\ell + 1)p$ vertices and $\ddot{m}[p] \geq (1 - o(1))2\ell p^{3/2}/\xi_{\ell}(\sqrt{p})$ edges.*
2. *The vertices of each pair in $\mathcal{P}(\ddot{B}[p])$ are connected by a unique shortest path in $\ddot{B}[p]$, whose edge labels alternate between two labels in $\mathcal{L}[\sqrt{p}]$.*
3. *By definition, $|\mathcal{P}(\ddot{B}[p])| = (|\mathcal{P}(\dot{B}[\sqrt{p}])|)^2 \geq p^2/(\xi_{\ell}(\sqrt{p}))^2$.*

Proof. Part 1. Each layer of \ddot{B} contains $(\sqrt{p})^2$ vertices; there is no harm in adding dummy vertices to round it up to p . There are at least $\ell\sqrt{p}^2/\xi_{\ell}(\sqrt{p})$ edges in each of \dot{B}^0 and \dot{B}^1 , and each edge of \dot{B}^0, \dot{B}^1 is duplicated \sqrt{p} times in the construction of \ddot{B} . Parts 2 and 3. This part of the proof follows directly from the construction of \ddot{B} , and that $\mathcal{P}(\dot{B}[\sqrt{p}])$ has unique shortest paths in \dot{B} . \square

A standard extension of Behrend's construction [11] of progression-free sets (see [2, Appendix]) shows that $\xi_{\ell}(p) = 2^{O(\sqrt{\log p \log \ell})}$, so if $\ell = p^{o(1)}$ then $\xi_{\ell}(p) = p^{o(1)}$ as well, and if $\ell = p^{\epsilon}$ for an $\epsilon > 0$ then $\xi_{\ell}(p) = p^{\delta}$ for some $\delta = \delta(\epsilon) > \epsilon$. We are most interested in the cases when $\ell, \xi_{\ell}(p) = p^{o(1)}$.

2.3. A recursive construction. In this section we construct a hierarchy $\{(\mathcal{H}_k, \mathcal{P}_k)\}_{k \geq 1}$ of hard graphs (\mathcal{H}_k) and corresponding pair-sets (\mathcal{P}_k) such that each pair in \mathcal{P}_k has a unique shortest path in \mathcal{H}_k . We will show that, for any $k \geq 2$ and sufficiently small constant c_k , any spanner of \mathcal{H}_k with stretch function $f(d) = d + c_k d^{1-\frac{1}{k}} + \tilde{O}(1)$ must include at least $|\mathcal{P}_k|$ edges. Each $\mathcal{H}_k[p]$ is a layered graph with p input ports, p output ports, and some number of interior layers. In other words, the

⁷Here we let \sqrt{p} be short for $\lfloor \sqrt{p} \rfloor$. Ignoring issues of integrality only introduces $1 + o(1)$ factors in all the bounds.

first layer (“input ports”) and last layer (“output ports”) have size p each while the interior layers may have different sizes, and each node pair in \mathcal{P}_k is composed of one input port and one output port. Let $\overline{\mathcal{H}}_k[p]$ denote the graph with the same topology as $\mathcal{H}_k[p]$ but with layers reversed; that is, the roles of input and output ports are swapped.

The base case. The base case graph $\mathcal{H}_1[p] = (\{1, \dots, 2p\}, \{1, \dots, p\} \times \{p + 1, \dots, 2p\})$ is a complete bipartite graph on $2p$ vertices and its corresponding pair-set $\mathcal{P}_1[p] = \{1, \dots, p\} \times \{p + 1, \dots, 2p\}$ has size p^2 .

The inductive case. Let us first give a very informal overview of the construction and then discuss how we plan to prove its correctness. The goal is to produce a new graph \mathcal{H}_k that contains within it many copies of \mathcal{H}_{k-1} . The shortest path $P_{s,t}$ for each $(s, t) \in \mathcal{P}_k$ joins an input port s to an output port t , in \mathcal{H}_k , and meanders through many copies of \mathcal{H}_{k-1} . When $P_{s,t}$ goes through a copy of \mathcal{H}_{k-1} , it enters and exits it at a particular input/output port pair, say (x, y) . We hope that $(x, y) \in \mathcal{P}_{k-1}$ (a success); if this holds for *all* the copies of \mathcal{H}_{k-1} intersected by $P_{s,t}$, then any aggressive sparsification of these copies will introduce a significant additive error in each copy. Unfortunately, while $|\mathcal{P}_{k-1}|$ is large, it is not *that* large. Only a tiny $o(1)$ -fraction of the set of input/output port pairs of \mathcal{H}_{k-1} appears in \mathcal{P}_{k-1} . Thus, if $P_{s,t}$ walks into and out of each \mathcal{H}_{k-1} through random ports, it is likely to miss the pairs in \mathcal{P}_{k-1} (a failure).

The problem with this approach is not the random assignment of input/output ports but the *independence* across copies of \mathcal{H}_{k-1} . We solve this problem by *correlating* the success or failure events associated with $P_{s,t}$. That is, we ensure that $P_{s,t}$ either enters/leaves *every* copy of \mathcal{H}_{k-1} along a pair in \mathcal{P}_{k-1} , or it enters/leaves *no* copy of \mathcal{H}_{k-1} using a pair in \mathcal{P}_{k-1} . Thus, many of the potential pairs are useless and may be discarded, but some of the pairs (s, t) must accumulate lots of error at each copy of \mathcal{H}_{k-1} that $P_{s,t}$ touches.

We now give this argument more formally. When $k \geq 2$ we construct $\mathcal{H}_k[p]$ from $\mathcal{H}_{k-1}[\cdot]$ and $\check{B} = \check{B}[p]$ as follows. Let the label-set of \check{B} be $\mathcal{L} = \mathcal{L}[\sqrt{p}]$ and $p' = |\mathcal{L}|$. Let $\mathcal{H}_{k-1}, \overline{\mathcal{H}}_{k-1}$ be the standard and reversed copies of $\mathcal{H}_{k-1}[p']$ and $\pi : \mathcal{L} \rightarrow \{1, \dots, p'\}$ be a *port assignment permutation* selected uniformly at random.

Recall that \check{B} consists of layers $\check{L}_0, \dots, \check{L}_{2\ell}$. Layers \check{L}_0 and $\check{L}_{2\ell}$ become the input and output ports of \mathcal{H}_k and are left as is. For each vertex u in an interior layer \check{L}_i , we replace u with a graph $\mathcal{H}(u)$, which is a copy of \mathcal{H}_{k-1} if i is odd and $\overline{\mathcal{H}}_{k-1}$ if i is even. For each former edge $(u, u') \in \check{L}_i \times \check{L}_{i+1}$ in \check{B} with $\text{label}(u, u') = a$, we replace it with a path of length $(2\ell - 1)^{k-1}$ connecting the $\pi(a)$ th output port of $\mathcal{H}(u)$ (or leave it at u if $i = 0$) and the $\pi(a)$ th input port of $\mathcal{H}(u')$ (or leave it at u' if $i + 1 = 2\ell$). The resulting graph is $\mathcal{H}_k[p]$; see Figure 2 for a diagram. It remains to define the new pair-set $\mathcal{P}_k[p]$.

Let $(u_0, u_{2\ell}) \in \mathcal{P}(\check{B})$ be one of the pairs in \check{B} , and suppose the edges on the unique shortest path from u_0 to $u_{2\ell}$ alternate between labels a and b . The corresponding path $Q_{(u_0, u_{2\ell})}$ in \mathcal{H}_k passes through some $\mathcal{H}(u_1), \mathcal{H}(u_2), \dots, \mathcal{H}(u_{2\ell-1})$, where $\mathcal{H}(u_1), \mathcal{H}(u_3), \dots$ are copies of $\mathcal{H}_{k-1}[p']$ and $\mathcal{H}(u_2), \mathcal{H}(u_4), \dots$ are copies of $\overline{\mathcal{H}}_{k-1}[p']$.

By construction, $Q_{(u_0, u_{2\ell})}$ enters $\mathcal{H}(u_i)$ at the $\pi(a)$ th input port and leaves at the $\pi(b)$ th output port, if i is odd, or the reverse if i is even. Up to reversal, the input/output terminals through each $\mathcal{H}(u_i)$ are identical for all $i \in [1, 2\ell - 1]$. The pair-set $\mathcal{P}_{k-1}[p]$ consists of all $(u_0, u_{2\ell}) \in \mathcal{P}(\check{B})$ (whose unique shortest path in \check{B} is labeled with, say, a, b) for which $(\pi(a), \pi(b)) \in \mathcal{P}_{k-1}[p']$.




FIG. 2. The edge-labels of a shortest path $(u_0, u_1, \dots, u_{2\ell})$ for a pair $(u_0, u_{2\ell}) \in \mathcal{P}(\ddot{B})$ always alternate between some $a, b \in \mathcal{L}$. To form \mathcal{H}_k we substitute for each vertex $u \in \ddot{B}$ a graph $\mathcal{H}(u)$, which is a copy of either \mathcal{H}_{k-1} or $\overline{\mathcal{H}}_{k-1}$, depending on whether u appears in an odd or even numbered layer, respectively; then we replace each edge (u, u') with a path of length $(2\ell - 1)^{k-1}$. The endpoints of this path are the $(\pi(\text{label}(u, u')))$ th output/input ports of $\mathcal{H}(u)$ and $\mathcal{H}(u')$. If $(\pi(a), \pi(b))$ is in \mathcal{P}_{k-1} , the pair set of \mathcal{H}_{k-1} , then there is a unique shortest path in \mathcal{H}_{k-1} from its $\pi(a)$ th input port to its $\pi(b)$ th output port.

LEMMA 2.2. The expected size of \mathcal{P}_k is $\frac{p^2}{(\xi_\ell(\sqrt{p}))^2} \cdot \frac{|\mathcal{P}_{k-1}|}{(p')^2}$. Assuming $\xi_\ell(\cdot)$ is a non-decreasing function for all ℓ , the expected size of \mathcal{P}_k is on the order of $p^2 / (\xi_\ell(\sqrt{p}))^{2(k-1)}$.

Proof. By definition of $\xi_\ell(\cdot)$ and the construction of $\ddot{B}[p]$, there are $p^2 / (\xi_\ell(\sqrt{p}))^2$ candidate pairs in $\mathcal{P}(\ddot{B}[p])$, each of which, say $(u_0, u_{2\ell})$, is associated with two alternating labels, say a and b . Since a uniformly random input/output pair $(\pi(a), \pi(b))$ is in \mathcal{P}_{k-1} with probability $|\mathcal{P}_{k-1}| / (p')^2$, $(u_0, u_{2\ell})$ is retained in \mathcal{P}_k with exactly this probability.

The second part of the proof is by induction. Let $\xi = \xi_\ell(\sqrt{p})$. When $k = 1$ we have $|\mathcal{P}_1| = p^2 = p^2 / \xi^0$. By the inductive hypothesis and the fact that ξ_ℓ is nondecreasing, $|\mathcal{P}_{k-1}| \geq (p')^2 / \xi^{2(k-2)}$ (in expectation) and so $|\mathcal{P}_k| \geq (p^2 / \xi^2) \cdot (1 / \xi^{2(k-2)}) = p^2 / \xi^{2(k-1)}$ (in expectation). \square

LEMMA 2.3. If $(u_0, u_{2\ell}) \in \mathcal{P}_k$, then there is a unique shortest path from u_0 to $u_{2\ell}$ in \mathcal{H}_k which passes through exactly $(2\ell - 1)^{k-1}$ complete bipartite graphs (copies of $\mathcal{H}_1[\cdot]$) and has $\text{dist}_{\mathcal{H}_k}(u_0, u_{2\ell}) = (2(k - 1)\ell + 1)(2\ell - 1)^{k-1}$, which is exactly the distance from the first to last layer of \mathcal{H}_k .

Proof. We prove the claim by induction. Let d_k be the length of shortest paths for \mathcal{P}_k in \mathcal{H}_k . Uniqueness of shortest paths in \mathcal{H}_1 is obvious, and d_1 is clearly 1. In the construction of \mathcal{H}_k , shortest paths for pairs in $\mathcal{P}(\ddot{B}[p])$ traverse 2ℓ edges via $2\ell - 1$ interior layers. Each edge becomes a path of length $(2\ell - 1)^{k-1}$ and each interior vertex becomes a copy of \mathcal{H}_{k-1} . Suppose a pair $(u_0, u_{2\ell}) \in \mathcal{P}(\ddot{B}[p])$ is preserved in \mathcal{P}_k and let $(u_0, u_1, \dots, u_{2\ell})$ be the shortest path in \ddot{B} . By the definition of \mathcal{P}_k and the inductive hypothesis, there are unique shortest paths from the given input port to the given output port in $\mathcal{H}(u_1), \mathcal{H}(u_2), \dots, \mathcal{H}(u_{2\ell-1})$. Each passes through $(2\ell - 1)^{k-2}$ complete bipartite graphs and each has length exactly d_{k-1} . Any alternative shortest path would have to visit consecutive layers without backing up to earlier layers and would therefore visit $u_0, \mathcal{H}(u'_1), \dots, \mathcal{H}(u'_{2\ell-1}), u_{2\ell}$ for some $(u'_1, \dots, u'_{2\ell-1}) \neq (u_1, \dots, u_{2\ell-1})$. This, however, violates Lemma 2.1(2) on the uniqueness of shortest paths between

pairs in $\mathcal{P}(\ddot{B})$. We have the recurrence

$$\begin{aligned} d_1 &= 1, \\ d_k &= (2\ell - 1)d_{k-1} + 2\ell(2\ell - 1)^{k-1}, \end{aligned}$$

which has the closed form $d_k = (2(k - 1)\ell + 1)(2\ell - 1)^{k-1}$. □

DEFINITION 2.4. *An edge e is critical for a pair $(u_0, u_{2\ell}) \in \mathcal{P}_k$ if it lies in a copy of \mathcal{H}_1 (a complete bipartite graph), and it is on the unique shortest u_0 - $u_{2\ell}$ path.*

LEMMA 2.5. *Let \mathcal{H}'_k be \mathcal{H}_k , with all critical edges for $(u_0, u_{2\ell})$ removed. Then $\text{dist}_{\mathcal{H}'_k}(u_0, u_{2\ell}) \geq \text{dist}_{\mathcal{H}_k}(u_0, u_{2\ell}) + 2(2\ell - 1)^{k-1}$.*

Proof. Let $(u_0, u_1, \dots, u_{2\ell})$ be the unique shortest path from u_0 to $u_{2\ell}$ in \ddot{B} . In \mathcal{H}'_k , if we take a path that does *not* pass through $\mathcal{H}(u_1), \dots, \mathcal{H}(u_{2\ell-1})$, then at some point it must move back to an earlier layer (from, say, layer i to layer $i - 1$) before advancing forward again (from layer $i - 1$ to i , and onward to 2ℓ). Since each edge in \ddot{B} corresponds to a path of length $(2\ell - 1)^{k-1}$ in \mathcal{H}_k , such a detour increases the path length by at least $2(2\ell - 1)^{k-1}$. On the other hand, if we do take a path passing through $\mathcal{H}(u_1), \dots, \mathcal{H}(u_{2\ell-1})$, then it must use the same input/output ports as the unique shortest u_0 - $u_{2\ell}$ path. By the inductive hypothesis, the additive stretch inside each of these subgraphs must be $2(2\ell - 1)^{k-2}$. (This is true when $k = 2$ as well, since in this case $\mathcal{H}(u_1), \dots, \mathcal{H}(u_{2\ell-1})$ are complete bipartite graphs, and removing the critical edge increases the distance from 1 to 3 in each one.) Thus, the total additive stretch is at least $(2\ell - 1) \cdot 2(2\ell - 1)^{k-2} = 2(2\ell - 1)^{k-1}$. □

LEMMA 2.6. *The shortest paths in \mathcal{H}_k for two pairs in \mathcal{P}_k share no critical edges. As a consequence, any spanner of \mathcal{H}_k with fewer than $|\mathcal{P}_k|$ edges must stretch some pair of vertices at distance $d = (2(k - 1)\ell + 1)(2\ell - 1)^{k-1}$ by an additive $2(2\ell - 1)^{k-1} \approx \frac{2}{(k-1)^{1-1/k}} \cdot d^{1-\frac{1}{k}}$.*

Proof. The proof is by induction on k ; it clearly holds when $k = 1$. For $k \geq 2$, each pair $(u_0, u_{2\ell}) \in \mathcal{P}_k$ is identified with a pair of labels $a, b \in \mathcal{L}$, which determines the input/output ports of $\mathcal{H}(u_1), \dots, \mathcal{H}(u_{2\ell-1})$ used in the shortest u_0 - $u_{2\ell}$ path. No other pair $(u'_0, u'_{2\ell})$ whose shortest path intersects some $\mathcal{H}(u_i)$ can be associated with the same two labels; hence it must enter and exit $\mathcal{H}(u_i)$ at different input/output ports than $(u_0, u_{2\ell})$. By the inductive hypothesis, $(u_0, u_{2\ell})$ and $(u'_0, u'_{2\ell})$ share no critical edges in any $\mathcal{H}(u_i)$ and therefore no critical edges in \mathcal{H}_k . □

We now analyze the size and density of $\mathcal{H}_k[p]$. Let $n_k[p]$ and $m_k[p]$ be the number of vertices and edges in $\mathcal{H}_k[p]$. The construction of \mathcal{H}_k gives the following recursive definition:

$$\begin{aligned} \dot{n}[p] &= (\ell + 1)p, & \dot{m}[p] &\geq \frac{\ell p^2}{\xi_\ell(p)}, \\ \ddot{n}[p] &= (2\ell + 1)p, & \ddot{m}[p] &\geq \frac{2\ell p^{3/2}}{\xi_\ell(\sqrt{p})}, \quad \ddot{m}[p] \leq 2\ell p^{3/2}, \\ n_1[p] &= 2p, & m_1[p] &= p^2, \\ n_k[p] &= 2p + (2\ell - 1)p \cdot n_{k-1} [|\mathcal{L}[\sqrt{p}]|] + (2\ell - 1)^{k-1} \cdot \dot{m}[p], & m_k[p] &\geq (2\ell - 1)p \cdot m_{k-1} [|\mathcal{L}[\sqrt{p}]|]. \end{aligned}$$

LEMMA 2.7. *For all $k \geq 1$ and sufficiently large p , $n_k[p] \leq c(2\ell)^k p^{2-\frac{1}{2^{k-1}}}$ and $m_k[p] \geq (2\ell - 1)^{k-1} p^2 / \xi^{2(k-1)}$, where $\xi = \xi_\ell(\sqrt{p})$.*

Proof. The bounds clearly hold when $k = 1$. Assuming the claim holds inductively for n_{k-1} , we have

$$\begin{aligned} n_k[p] &\leq (2\ell)p \cdot n_{k-1} \lceil \sqrt{p}/2 \rceil + (2\ell - 1)^{k-1} \ddot{m}[p] && (|\mathcal{L}[\sqrt{p}]| \leq \sqrt{p}/2) \\ &\leq (2\ell)p \cdot \left[c(2\ell)^{k-1} (\sqrt{p}/2)^{2 - \frac{1}{2^{k-2}}} \right] + (2\ell - 1)^{k-1} 2\ell p^{3/2} \\ &< (c/2^{2 - \frac{1}{2^{k-2}}}) (2\ell)^k p^{2 - \frac{1}{2^{k-1}}} + (2\ell)^k p^{3/2} \\ &< c(2\ell)^k p^{2 - \frac{1}{2^{k-1}}} && (\text{for, say, } c = 2) \end{aligned}$$

and assuming the claim holds inductively for m_{k-1} , we have

$$\begin{aligned} m_k[p] &\geq (2\ell - 1)p \cdot m_{k-1} \lceil \sqrt{p}/\xi \rceil && (|\mathcal{L}[\sqrt{p}]| \geq \sqrt{p}/\xi_\ell(\sqrt{p})) \\ &\geq (2\ell - 1)p \cdot (2\ell - 1)^{k-2} (\sqrt{p}/\xi)^2 / \xi^{2(k-2)} && (\xi_\ell(\cdot) \text{ nondecreasing}) \\ &\geq (2\ell - 1)^{k-1} p^2 / \xi^{2(k-1)}. \end{aligned}$$

We are mainly interested in cases in which ℓ is not too large, i.e., $\ell, \xi = (n_k[p])^{o(1)}$. In this case the density of $\mathcal{H}_k[p]$ is $m_k[p]/n_k[p] \geq (n_k[p])^{\frac{1}{2^{k-1}} - o(1)}$. \square

THEOREM 2.8 (sublinear additive spanner lower bounds). *For any integer $k \geq 2$ and a sufficiently small constant $c_k = O(1/k)$, any spanner construction with stretch function bounded by $f(d) \leq d + c_k d^{1 - \frac{1}{k}} + \tilde{O}(1)$ has size $\Omega(n^{1 + \frac{1}{2^{k+1}-1} - o(1)})$ in the worst case.*

Proof. Let $\mathcal{H}_k[p]$ be the input graph with respect to some sufficiently large $\ell = (\log p)^{O(1)}$. For this parameterization $\xi_\ell(p) = 2^{O(\sqrt{\log p \log \log p})} = p^{o(1)}$, we have that $m_k[p]$ and the size of the pair-set \mathcal{P}_k are both $n^{1 + \frac{1}{2^{k-1}} - o(1)}$, where $n = n_k[p]$. Any spanner with size less than $|\mathcal{P}_k|$ must stretch some pair at distance $d_k = (2(k-1)\ell + 1)(2\ell - 1)^{k-1}$ to $d_k + 2(2\ell - 1)^{k-1}$, which is strictly greater than $d_k + c_k d^{1 - \frac{1}{k}} + \tilde{O}(1)$ when $c_k < 2/(k-1)^{1-1/k}$ is sufficiently small and ℓ sufficiently large to make the $\tilde{O}(1)$ error comparatively negligible. \square

Remark 2.9. Since the diameter of $\mathcal{H}_k[p]$ is $O(d_k)$, any emulator for $\mathcal{H}_k[p]$ on the same vertex set (i.e., without Steiner points) can be converted to a spanner with at most an $O(d_k) = n^{o(1)}$ blowup in the number of edges. Thus, Theorem 2.8 applies to this class of emulators. The argument breaks down for (Steiner) emulators since we can preserve all distances for pairs in \mathcal{P}_k with just $O(n)$ edges, simply by replacing all bipartite cliques (copies of \mathcal{H}_1) with stars. See Theorem 2.11 for a lower bound that applies to emulators with Steiner points.

THEOREM 2.10 $((1 + \epsilon, \beta)$ -spanner lower bounds). *Any $(1 + \epsilon, \beta)$ -spanner construction with worst-case size at most $n^{1 + \frac{1}{2^{k+1}-1} - \delta}$, $\delta > 0$, has $\beta = \Omega(\frac{1}{\epsilon^{(k-1)}})^{k-1}$.*

Proof. Let $\mathcal{H}_k[p]$ be the input graph with respect to an $\ell \approx \frac{1}{\epsilon}$ to be chosen shortly. Any spanner with size $n^{1 + \frac{1}{2^{k+1}-1} - \delta} < |\mathcal{P}_k|$ stretches a pair at distance $d_k = (2(k-1)\ell + 1)(2\ell - 1)^{k-1}$ to $d_k + 2(2\ell - 1)^{k-1} = d_k(1 + \frac{2}{2(k-1)\ell + 1})$. We choose $\ell \geq 2$ to be minimal such that $\frac{1}{2(k-1)\ell + 1} \leq \epsilon$, that is, $\ell = \lceil \frac{1-\epsilon}{2\epsilon(k-1)} \rceil$ and the additive stretch is roughly $2\epsilon d_k$. In order for this to be a $(1 + \epsilon, \beta)$ -spanner we would need

$$\beta = \Omega(\epsilon d_k) = \Omega((2\ell - 1)^{k-1}) = \Omega\left(\left(2\lceil \frac{1-\epsilon}{2\epsilon(k-1)} \rceil - 1\right)^{k-1}\right). \quad \square$$

Theorem 2.10 shows that the existing $(1 + \epsilon, O(k/\epsilon)^{k-1})$ -spanners with size $O((k/\epsilon)^{O(1)}n^{1+\frac{1}{2k+1-1}})$ are optimal in the following sense. If k is constant, then we cannot improve β by more than a constant factor $\approx (k^2)^{k-1}$ without increasing the exponent to $1 + \frac{1}{2k-1} - o(1)$. Moreover, any constant reduction in the exponent increases β to $\Theta(1/(k\epsilon))^k$. Once again, the argument of Theorem 2.10 applies to $(1 + \epsilon, \beta)$ -emulators that do not use Steiner points.

THEOREM 2.11 (distance oracle/emulator lower bounds). *Consider any data structure for the class of n -vertex undirected graphs that answers approximate distance queries. If its stretch function is*

- $f(d) \leq d + c_k d^{1-\frac{1}{k}} + \tilde{O}(1)$ for an integer k and a sufficiently small constant $c_k < 2/(k-1)^{1-1/k}$, or
- $f(d) \leq (1 + \epsilon)d + \beta$ where $\beta = o((\frac{1}{\epsilon(k-1)})^{k-1})$,

then on some graph, the data structure occupies at least $n^{1+\frac{1}{2k-1}-o(1)}$ bits of space.

Proof. The following proof strategy was employed by Althöfer et al. [8] to bound the size of emulators. It was also used by Matoušek [44] for bounding low-distortion embeddings into l_∞^d , and by Thorup and Zwick [58] and Abboud and Bodwin [2] to bound the size of approximate distance data structures.

Fix a graph $\mathcal{H}_k = \mathcal{H}_k[p]$ with pair-set \mathcal{P}_k . For any subset $\mathcal{P}' \subseteq \mathcal{P}_k$ let $G(\mathcal{P}')$ be obtained from \mathcal{H}_k by removing all the critical edges for each pair in \mathcal{P}' , and let $\mathcal{G} = \{G(\mathcal{P}') \mid \mathcal{P}' \subseteq \mathcal{P}_k\}$ be the class of $2^{|\mathcal{P}_k|}$ such graphs. Fix any two graphs $G_A, G_B \in \mathcal{G}$. There must exist some pair $(u, v) \in \mathcal{P}_k$ such that G_A contains all of the critical edges for (u, v) whereas G_B contains none of them. By Lemma 2.5 we have

$$\begin{aligned} \text{dist}_{G_A}(u, v) &= d_A \stackrel{\text{def}}{=} (2(k-1)\ell + 1)(2\ell - 1)^{k-1}, \\ \text{dist}_{G_B}(u, v) &\geq d_B \stackrel{\text{def}}{=} (2(k-1)\ell + 1)(2\ell - 1)^{k-1} + 2(2\ell - 1)^{k-1}. \end{aligned}$$

If $f(d_A) < d_B$, then no single data structure (bit string) can be used to encode two distinct graphs $G_A, G_B \in \mathcal{G}$, implying that the data structures for \mathcal{G} must use at least $\log_2(2^{|\mathcal{P}_k|})$ bits on average. If $\ell = (\log p)^{O(1)} = (\log n)^{O(1)}$ is sufficiently large and c_k sufficiently small, then

$$f(d_A) = d_A + c_k d_A^{1-\frac{1}{k}} + \tilde{O}(1) < d_B \approx d_A + \frac{2}{(k-1)^{1-1/k}} \cdot d_A^{1-\frac{1}{k}}.$$

For these parameters $|\mathcal{P}_k| = n^{1+\frac{1}{2k-1}-o(1)}$, where the $n^{-o(1)}$ factor is $2^{-O(\sqrt{\log n \log \log n})}$. We extend this argument to the case of $(1 + \epsilon, \beta)$ -type stretch using an identical argument to the one given in Theorem 2.10. \square

3. New upper bounds on $(1 + \epsilon, \beta)$ -spanners. Thorup and Zwick [59] gave a very simple randomized construction of an emulator with size $O(kn^{1-\frac{1}{2k+1-1}})$ and stretch function $f(d) = d + O(kd^{1-1/k} + 3^k)$. Alternatively, one can view this as a $(1 + \epsilon, O(k/\epsilon)^{k-1})$ -emulator for every $\epsilon > 0$, where the optimal choice of ϵ , as a function of d , is $\epsilon = \Theta(k/d^{1/k})$.

3.1. The Thorup–Zwick emulator. The Thorup–Zwick emulator is parameterized by an integer $k \geq 2$. Let $G = (V, E)$ be the input graph. One samples vertex sets $V = V_0 \supset V_1 \supset V_2 \cdots \supset V_k$ where vertices in V_i are promoted to V_{i+1} with probability q_{i+1}/q_i , so $\mathbb{E}[|V_i|] = q_i n$. Define $p_i(u)$ to be the closest V_i -vertex to u , breaking ties in a consistent manner. Define $\mathcal{B}(u, r) = \{v \mid \text{dist}(u, v) \leq r\}$ to be

the set of vertices inside the radius- r ball centered at u and let $\mathcal{B}_i(u)$ be short for $\mathcal{B}(u, \text{dist}(u, p_i(u)) - 1)$. For $i \geq k + 1$, $p_i(u)$ does not exist and $\mathcal{B}_i(u)$ is the entire graph, by definition. The emulator edge set consists of $E_0 \cup E_1 \cup \dots \cup E_k$, where E_i is defined as follows:

$$E_i = \{(u, v) \mid u, v \in V_i \text{ and } v \in \mathcal{B}_{i+1}(u)\} \cup \{(u, p_{i+1}(u)) \mid u \in V\}.$$

The length of all emulator edges is precisely the distance between their endpoints in G . Since $|\mathcal{B}_{i+1}(u)|$ is q_{i+1}^{-1} in expectation, the expected number of edges contributed by E_i is $n + nq_i^2/q_{i+1}$ for $i < k$ and is $(nq_k)^2$ when $i = k$. Setting $q_i = n^{-\frac{2^i-1}{2^{k+1}-1}}$ makes the size of the emulator $O(kn^{1+\frac{1}{2^{k+1}-1}})$ in expectation. In order to obtain a $d + O(kd^{1-\frac{1}{k}})$ -type stretch bound for all distances $d \leq D$, it actually suffices to restrict E_i to pairs at distance at most $(r+2)^i$, where $r = D^{1/k}$. Letting $P(u, v)$ be any shortest path from u to v , the subgraph $S_{TZ}(k, r) = (V, E'_0 \cup E'_1 \cup \dots \cup E'_k)$ is a spanner, where

$$E'_i = \bigcup_{\substack{(u,v) \in E_i : \\ v \in \mathcal{B}(u, (r+2)^i)}} P(u, v).$$

As we show in sections 3.2 and 3.3, the spanner $S_{TZ}(k, r)$ behaves exactly like the emulator for all distances up to D ; i.e., it has stretch function $d + O(kd^{1-\frac{1}{k}})$ for all sufficiently large $d \leq D$.⁸ However, choosing the optimum sampling probabilities as a function of r, k, n is no longer trivial. Since each path in E_i contributes $(r+2)^i$ edges, the spanner size is on the order of kn (for paths of the form $P(u, p_i(u))$) plus

$$\frac{n}{q_1} + \frac{nq_1^2 r}{q_2} + \frac{nq_2^2 r^2}{q_3} + \dots + \frac{nq_{k-1}^2 r^{k-1}}{q_k} + (nq_k)^2 r^k.$$

Assuming this sum is minimized when E'_0, E'_1, \dots, E'_k contribute equally, we have the following equalities:

$$\begin{aligned} q_2 &= rq_1^3 && \text{(balancing } E'_0 \text{ and } E'_1) \\ q_3 &= r^2 q_2^2 q_1 && \text{(balancing } E'_0 \text{ and } E'_2) \\ &\vdots && \\ q_k &= r^{k-1} q_{k-1}^2 q_1. && \text{(balancing } E'_0 \text{ and } E'_{k-1}) \end{aligned}$$

If q_i is constrained to be of the form $n^{-g(i)} r^{-h(i)}$, these equalities are satisfied when

$$\begin{aligned} g(i) &= 2g(i-1) + g(1) && \text{(for } i \geq 2) \\ &= (2^i - 1)g(1) && \text{(by induction)} \\ \text{and } h(i) &= 2h(i-1) + h(1) - (i-1) && \text{(for } i \geq 2) \\ &= (2^i - 1)h(1) - 2^i + (i+1). && \text{(by induction)} \end{aligned}$$

⁸Thorup and Zwick [59, p. 809] also noted that their emulator can be converted to a spanner, but their sketch of how to do this was incorrect.

So $q_k = n^{-(2^k-1)g(1)}r^{-(2^k-1)h(1)+2^k-(k+1)}$. Plugging this equality into the expression for $|E'_k|$ and balancing with $|E'_0|$, we have

$$|E'_k| = (nq_k)^2r^k = n^{2-2(2^k-1)g(1)}r^{-2[(2^k-1)h(1)-2^k+(k+1)]+k} = n^{1+g(1)}r^{h(1)} = |E'_0|,$$

which is minimized when

$$g(1) = \frac{1}{2^{k+1} - 1},$$

$$h(1) = \frac{2^{k+1} - (k + 2)}{2^{k+1} - 1}.$$

For example, when $k = 2$ and $h(1) = 4/7$ this leads to a $d + O(\sqrt{d})$ -spanner for distances $d \leq D = r^2$ having size $O(r^{4/7}n^{8/7}) = O(D^{2/7}n^{8/7})$. Since $h(1)$ is strictly less than 1 for any fixed k , the spanner size is always $o(rkn^{1+\frac{1}{2^{k+1}-1}}) = o(D^{1/k}kn^{1+\frac{1}{2^{k+1}-1}})$.

3.2. Even sparser $(1 + \epsilon, \beta)$ -spanners. In order to form an even sparser spanner we substitute for E'_1 a subgraph whose size has no dependence on r but preserves the relevant distances well enough, up to an additive +2 error. The following theorem is proved using the same path-buying algorithm for constructing additive 6-spanners [9, 49]. The algorithm begins with the subgraph E'_0 and supplements it with an \tilde{E}_1 to guarantee +2 stretch for each $u, v \in V_1$ that were connected by an edge in E_1 .

THEOREM 3.1 (see [9] and [49]). *Suppose V_1, V_2 are sampled with probability q_1 and q_2 , with $q_2 < q_1$. Then there is an edge-set \tilde{E}_1 with expected size $O(nq_1^2/q_2)$ such that if $u, v \in V_1$ and $v \in \mathcal{B}_2(u)$, then*

$$\text{dist}_{E'_0 \cup \tilde{E}_1}(u, v) \leq \text{dist}(u, v) + 2.$$

Proof (sketch). We assume the reader is familiar with the path-buying algorithm and its analysis [9]. Let $\mathcal{P} \subset \binom{V_1}{2}$ be the pairs for which we are guaranteeing good stretch, i.e., $\{u, v\} \in \mathcal{P}$ if $u \in \mathcal{B}_2(v)$ or $v \in \mathcal{B}_2(u)$. Since $|\mathcal{B}_2(u)|$ is $1/q_2$ in expectation, $|\mathcal{P}|$ is $O(nq_1^2/q_2)$ in expectation. For each $\{u, v\} \in \mathcal{P}$ we evaluate $P(u, v)$ and buy it (set $\tilde{E}_1 \leftarrow \tilde{E}_1 \cup P(u, v)$) if its current value exceeds its cost. The value is the number of pairs $\{x, y\} \in \mathcal{P}$ with x, y adjacent to $P(u, v)$ for which $\text{dist}_{E'_0 \cup \tilde{E}_1 \cup P(u, v)}(x, y) < \text{dist}_{E'_0 \cup \tilde{E}_1}(u, v)$. It is argued by the pigeonhole principle that any path $P(u, v)$ not bought has $\text{dist}_{E'_0 \cup \tilde{E}_1}(u, v) \leq \text{dist}(u, v) + 2$, and that each pair in \mathcal{P} is charged for $O(1)$ edges in \tilde{E}_1 . \square

We sample vertex sets $V = V_0 \supset V_1 \supset \dots \supset V_k$ as before and construct the spanner $S(k, r)$ with edge set $E'_0 \cup \tilde{E}_1 \cup E'_2 \cup \dots \cup E'_k$, where \tilde{E}_1 is the edge set from Theorem 3.1. The expected size of the entire spanner is therefore

$$\frac{n}{q_1} + \frac{nq_1^2}{q_2} + \frac{nq_2^2r^2}{q_3} + \dots + \frac{nq_{k-1}^2r^{k-1}}{q_k} + (nq_k)^2r^k.$$

Letting $q_i = n^{-\frac{2^i-1}{2^{k+1}-1}}r^{-h(i)}$, we balance the contribution of $E'_0, \tilde{E}_1, E'_2, \dots, E'_k$ by having h satisfy the following:

$$h(2) = 3h(1) \quad (\text{balancing } \tilde{E}_1 \text{ and } E'_0)$$

$$\text{and for } i \geq 3, h(i) = 2h(i-1) + h(1) - (i-1) \quad (\text{balancing } E'_{i-1} \text{ and } E'_0)$$

$$= (2^i - 1)h(1) - 3 \cdot 2^{i-2} + (i+1). \quad (\text{by induction})$$

Following similar calculations, it follows that the spanner size is minimized when

$$h(1) = \frac{3 \cdot 2^{k-1} - (k+2)}{2^{k+1} - 1}.$$

We shall prove shortly that this spanner is, indeed, a $d + O(kd^{1-\frac{1}{k}} + 3^k)$ -spanner. For example, when $k = 1$ we have $h(1) = 2/7$, so it is a $d + O(\sqrt{d})$ -spanner for all $d \leq D \leq r^2$ with size $O(r^{2/7}n^{8/7}) = O(D^{1/7}n^{8/7})$. For any fixed k , $h(1) < 3/4$, so the spanner has size $o(D^{\frac{3}{4k}}kn^{1+\frac{1}{2^{k+1}-1}})$.

Remark 3.2. We were able to substitute \tilde{E}_1 for E'_1 without disturbing the exponent $1 + \frac{1}{2^{k+1}-1}$ of the spanner, but only because the path-buying algorithm buys $O(nq_1^2/q_2)$ additional edges when initialized with the edge set E'_0 . In general we can use [49, Thm. 4.2] to substitute an \tilde{E}_i for E'_i , but its size is $O(n\sqrt{q_i/q_{i+1}})$. This improves the exponent attached to r but worsens the exponent attached to n . For example, balancing E'_2 and E'_0 lets us put $q_3 = (q_1)^7 r^{O(1)}$, whereas balancing \tilde{E}_2 and E'_0 forces $q_3 = (q_1)^5$.

3.3. Stretch analysis. See Figure 3 for a visualization of the argument that follows. We analyze the stretch of the spanner $S = S(k, r)$ with edge set $E'_0 \cup \tilde{E}_1 \cup E'_2 \cup \dots \cup E'_k$. We will first consider two vertices u, v at distance at most ℓ^i for some integers $\ell \geq 2, i \geq 0$. We will assume for the time being that $r = \infty$ and calculate specific quantities related to the spanner distance $\text{dist}_S(u, v)$ without considering the constraints imposed by a finite r . Once these quantities are calculated, it will be clear that the analysis goes through so long as $\ell \leq r$. The pair u, v can be either *complete* or *incomplete* (or both), as explained in the following definition.

DEFINITION 3.3. Define $\{C_\ell(i), I_\ell(i)\}_{\ell \in [2, r], i \geq 0}$ to be integers such that for all u, v with $\text{dist}(u, v) \leq \ell^i$, at least one of the following inequalities holds. Here $S = S(k, r)$ is the spanner.

$$\begin{aligned} \text{dist}_S(u, v) &\leq \text{dist}(u, v) + C_\ell(i), && \text{("}u \cdots v \text{ is complete")} \\ \text{dist}_S(u, p_{i+1}(u)) &\leq I_\ell(i). && \text{("}u \cdots v \text{ is incomplete")} \end{aligned}$$

LEMMA 3.4. The following values for $\{C_\ell(i), I_\ell(i)\}_{\ell \in [2, r], i \geq 0}$ satisfy Definition 3.3.

$$\begin{aligned} C_\ell(0) &= 0 && \text{for all } \ell, \\ I_\ell(0) &= 1 && \text{for all } \ell, \\ C_\ell(1) &= 6 && \text{for all } \ell, \\ I_\ell(1) &= \ell + 3 && \text{for all } \ell, \\ C_\ell(i) &= \min \begin{cases} \ell \cdot C_\ell(i-1) \\ (\ell-1) \cdot C_\ell(i-1) + 4 \cdot I_\ell(i-1) \end{cases} && \text{for all } \ell \text{ and } i \geq 2, \\ I_\ell(i) &= \ell^i + 3 \cdot I_\ell(i-1) && \text{for all } \ell \text{ and } i \geq 2. \end{aligned}$$

Proof. In the base case ($i = 0$), we have $\ell^0 = 1$, so u and v are adjacent in the input graph. If $(u, v) \in E'_0$ then $\text{dist}_S(u, v) = 1$, and if $(u, v) \notin E'_0$ then it must be that $\text{dist}(u, p_1(u)) = 1$, so $C_\ell(0) = 0, I_\ell(0) = 1$ satisfy Definition 3.3 for all ℓ .

When $i > 0$, partition the shortest path from u to v into at most ℓ segments with length ℓ^{i-1} , and let u_j be the vertex on the path at distance $j\ell^{i-1}$ from u . For the sake of simplicity, assume $\text{dist}(u, v) = \ell^i$, so $v = u_\ell$. Each segment from u_j to u_{j+1} is classified as either *complete* or *incomplete*. If all segments are complete, then




FIG. 3. The shortest path from $u = u_0$ to $v = u_\ell$ has length ℓ^i ; it is partitioned into segments of length ℓ^{i-1} . A segment (u_j, u_{j+1}) is complete if S contains a path of length $\ell^{i-1} + C_\ell(i-1)$ from u_j to u_{j+1} and incomplete if the distance from u_j to $p_i(u_j)$ is at most $I_\ell(i-1)$. If all segments are complete (not depicted), then $\text{dist}_S(u_0, u_\ell) \leq \text{dist}(u_0, u_\ell) + \ell \cdot C_\ell(i-1)$. If only the first s segments and last s' segments are complete and $p_i(u_{\ell-s'})$ lies in the ball $\mathcal{B}_{i+1}(p_i(u_s))$, then S contains a path from u_0 to u_ℓ with length $\text{dist}(u_0, u_\ell) + (s + s')C_\ell(i-1) + 4I_\ell(i-1)$. On the other hand, if $p_i(u_{\ell-s'})$ lies outside $\mathcal{B}_{i+1}(p_i(u_s))$, then this gives a bound on the distance from $p_i(u_s)$ to $p_{i+1}(p_i(u_s))$ and therefore a bound on $\text{dist}(u_0, p_{i+1}(u_0))$. From these cases we derive recursive expressions for $C_\ell(i)$ and $I_\ell(i)$.

$\text{dist}_S(u, v) \leq \text{dist}_S(u, v) \leq \text{dist}(u, v) + \ell \cdot C_\ell(i-1)$. If there is at least one incomplete segment, let there be s complete segments on a prefix of the path and s' complete segments on a suffix of the path, where $s + s' \leq \ell - 1$. It follows that

$$\begin{aligned} \text{dist}_S(u, p_i(u_s)) &\leq \text{dist}_S(u, u_s) + \text{dist}_S(u_s, p_i(u_s)) \\ &\leq \text{dist}(u, u_s) + s \cdot C_\ell(i-1) + I_\ell(i-1), \\ \text{dist}_S(v, p_i(u_{\ell-s'})) &\leq \text{dist}_S(v, u_{\ell-s'}) + \text{dist}_S(u_{\ell-s'}, p_i(u_{\ell-s'})) \\ &\leq \text{dist}(v, u_{\ell-s'}) + s' \cdot C_\ell(i-1) + I_\ell(i-1). \end{aligned}$$

If $p_i(u_{\ell-s'}) \notin \mathcal{B}_{i+1}(p_i(u_s))$, then

$$\begin{aligned} \text{dist}_S(u, p_{i+1}(u)) &\leq \text{dist}(u, p_i(u_s)) + \text{dist}(p_i(u_s), p_i(u_{\ell-s'})) \\ &\leq (\ell - s')\ell^{i-1} + 3I_\ell(i) \\ &\leq \ell^i + 3I_\ell(i) \end{aligned} \quad \text{worst case when } s' = 0$$

and the path from u to v is incomplete. On the other hand, if $p_i(u_{\ell-s'}) \in \mathcal{B}_{i+1}(p_i(u_s))$, then S contains a shortest (or nearly shortest, if $i = 1$) path from $p_i(u_s)$ to $p_i(u_{\ell-s'})$, so

$$\begin{aligned} \text{dist}_S(u, v) &\leq \text{dist}_S(u, p_i(u_s)) + \text{dist}_S(p_i(u_s), p_i(u_{\ell-s'})) + \text{dist}_S(p_i(u_{\ell-s'}), v) \\ &\leq [s(\ell^{i-1} + C_\ell(i-1)) + I_\ell(i-1)] \quad \text{from } u \text{ to } p_i(u_s) \\ &\quad + [(\ell - s - s')\ell^{i-1} + 2I_\ell(i-1) \{+2\}] \quad \text{from } p_i(u_s) \text{ to } p_i(u_{\ell-s'}) \\ &\quad + [s'(\ell^{i-1} + C_\ell(i-1)) + I_\ell(i-1)] \quad \text{from } p_i(u_{\ell-s'}) \text{ to } v \\ &\leq \text{dist}(u, v) + (\ell - 1)C_\ell(i-1) + 4I_\ell(i-1) \{+2\} \quad \text{worst case when } s + s' = \ell - 1 \end{aligned}$$

where the $\{+2\}$ is only present if $i = 1$. We satisfy Definition 3.3 by setting $C_\ell(1) = 6$, $I_\ell(1) = \ell + 3$, and, for $i \geq 2$, $I_\ell(i) = \ell^i + 3I_\ell(i-1)$ and $C_\ell(i)$ is the maximum of $\ell \cdot C_\ell(i-1)$ and $(\ell - 1)C_\ell(i-1) + 4I_\ell(i-1)$. \square

We now find closed form bounds for $C_\ell(i)$ and $I_\ell(i)$.

LEMMA 3.5. *The values defined inductively in Lemma 3.4 satisfy the following (in)equalities:*

$$\begin{aligned} I_2(i) &= 3^{i+1} - 2^{i+1}, \\ C_2(i) &\leq 3^{i+1}, \\ I_3(i) &= (i+1)3^i, \\ C_3(i) &\leq 4i3^i. \end{aligned}$$

Define $c_\ell = \ell/(\ell-3)$. For all $\ell \geq 4$ and $i \geq 1$,

$$\begin{aligned} I_\ell(i) &\leq c_\ell \ell^i, \\ C_\ell(i) &\leq \min \left\{ \begin{array}{l} 4c_\ell \ell^i, \\ (4c_\ell i + 2)\ell^{i-1}. \end{array} \right. \end{aligned}$$

Proof. All bounds are established by induction on i . The cases when $\ell \in \{2, 3\}$ are left as an exercise. When $\ell \geq 4$ the base cases $i \in \{0, 1\}$ clearly hold. For incomplete paths and $i \geq 2$ we have

$$\begin{aligned} I_\ell(i) &= \ell^i + 3 \cdot I_\ell(i-1) && \text{(by definition)} \\ &\leq \ell^i(1 + 3c_\ell/\ell) \leq c_\ell \ell^i, && \text{(induction hypothesis, } c_\ell = \frac{\ell}{\ell-3}\text{)} \end{aligned}$$

and for complete paths we have two cases:

$$\begin{aligned} C_\ell(i) &= (\ell-1)C_\ell(i-1) + 4I_\ell(i-1) && \text{(by definition)} \\ &\leq (\ell-1)4c_\ell \ell^{i-1} + 4c_\ell \ell^{i-1} && \text{(1st induction hypothesis)} \\ &= 4c_\ell \ell^i \\ \text{and } &\leq (\ell-1)(4c_\ell(i-1) + 2)\ell^{i-2} + 4c_\ell \ell^{i-1} && \text{(2nd induction hypothesis)} \\ &\leq (4c_\ell i + 2)\ell^{i-1}. \quad \square \end{aligned}$$

Observe that when we check whether $p_i(u_{\ell-s'}) \in \mathcal{B}_{i+1}(p_i(u_s))$, $i \geq 2$, the distance between $p_i(u_{\ell-s'})$ and $p_i(u_s)$ is maximized when $s = s' = 0$; it is at most

$$\ell^i + 2I_\ell(i-1) = \ell^i + 2c_\ell \ell^{i-1} < (\ell+2)^i.$$

Thus, as long as $\ell \leq r$, the criterion $p_i(u_{\ell-s'}) \in \mathcal{B}(p_i(u_s), (r+2)^i)$ will also hold. This retroactively justifies the constraint $\ell \leq r$ in Lemma 3.4. Theorem 3.6 summarizes the size-stretch tradeoffs for various ranges of d . Using more careful random sampling, Huang and Pettie [36] removed a factor k from the size bound of Theorem 3.6 and other related constructions.

THEOREM 3.6. *The spanner $S(k, r)$ has size $O(r^h k n^{1+\frac{1}{2^{k+1}-1}})$, where $h = \frac{3 \cdot 2^{k-1} - (k+2)}{2^{k+1}-1} < 3/4$. Its stretch changes as a function of the distance d being approximated.*

- For $d \geq 2^k$ it is a multiplicative $O((3/2)^k)$ -spanner.
- For $d \geq 3^k$ it is a multiplicative $O(k)$ -spanner.
- For $d \geq \ell^k$, $\ell \in [4, k]$, it is a multiplicative $(5 + O(1/\ell))$ -spanner, and when $\ell \in [k, r]$, it is a multiplicative $(1 + (4k + O(1))/\ell)$ -spanner.

$S(k, r)$ is a $(1 + \epsilon, ((4k + O(1))/\epsilon)^{k-1})$ -spanner for every ϵ such that $(4k + O(1))/\epsilon < r$. Its stretch function can also be expressed as $f(d) = d + (4 + o(1))kd^{1-\frac{1}{k}} + 3^k$ for all $d \leq r^k$, and $f(d) = d + (4 + o(1))kd/r$ for larger d , where the $o(1)$ s go to zero as d increases.

Proof. Let $\text{dist}(u, v) \geq 2^k$ be the distance to be approximated and $\ell = \lfloor d^{\frac{1}{k}} \rfloor$, where $\ell \leq r$. Partition the shortest path $P(u, v)$ into intervals of length precisely ℓ^{k-1} , with at most one shorter interval. Since $d < (\ell + 1)^k$, there are between ℓ and $\lfloor (\ell + 1)(1 + 1/\ell)^{k-1} \rfloor$ intervals. If all intervals are complete, then $\text{dist}_S(u, v) \leq d + \lceil \frac{d}{\ell^{k-1}} \rceil C_\ell(k-1)$. If at least one is incomplete, then $\text{dist}_S(u, v) \leq d + (\lceil \frac{d}{\ell^{k-1}} \rceil - 1)C_\ell(k-1) + 4I_\ell(k-1)$. If $\ell \in [3, k-1]$, then according to Lemma 3.5, $C_\ell(k-1) = 4I_\ell(k-1)$, and so our claimed upper bound holds in either case. If $\ell \geq k$ or $\ell = 2$, then $C_\ell(k-1) < 4I_\ell(k-1)$, so the second case is worse. When $\ell = 2$ we have

$$\begin{aligned} & d + \left(\lceil \frac{d}{2^{k-1}} \rceil - 1 \right) C_2(k-1) + 4I_2(k-1) \\ & < d + \left(\lceil \frac{d}{2^{k-1}} \rceil - 1 \right) 3^k + 4 \cdot 3^k \\ & \leq d(1 + 3(3/2)^{k-1}) + 4 \cdot 3^k. \end{aligned}$$

So $S(k, r)$ is a multiplicative $O((3/2)^{k-1})$ -spanner for $d \geq 2^k$. This is a nontrivial multiplicative stretch. Traditional multiplicative stretch spanners with size $n^{1+\frac{1}{2k+1-1}}$ stretch some pairs by a factor of $2^{k+2} - 3$. When $\ell = 3$ we have

$$\begin{aligned} & d + \left(\lceil \frac{d}{3^{k-1}} \rceil - 1 \right) C_3(k-1) + 4I_3(k-1) \\ & < d + \left(\lceil \frac{d}{3^{k-1}} \rceil - 1 \right) 4(k-1)3^{k-1} + 4k3^{k-1} \\ & \leq d(1 + 4(k-1)/3 + 4k/3). \end{aligned} \tag{since $3^{k-1} \leq d/3$ }$$

Thus $S(k, r)$ functions as a multiplicative $O(k)$ -spanner when $d \geq 3^k$. When $\ell \in [4, k]$,

$$\begin{aligned} & d + \left(\lceil \frac{d}{\ell^{k-1}} \rceil - 1 \right) C_\ell(k-1) + 4I_\ell(k-1) \\ & < d + \left(\lceil \frac{d}{\ell^{k-1}} \rceil - 1 \right) 4c_\ell \ell^{k-1} + 4c_\ell \ell^{k-1} \\ & < d(1 + 4c_\ell + 4c_\ell/\ell) \\ & = d \left(1 + \frac{4(\ell+1)}{\ell-3} \right) = (5 + O(\frac{1}{\ell}))d. \end{aligned} \tag{since $\ell^{k-1} \leq d/\ell, c_\ell = \ell/(\ell-3)$ }$$

The multiplicative stretch of $S(k, r)$ tends to 5 as d increases from 3^k to $(k-1)^k$. When $\ell \geq k$ we have

$$\begin{aligned} & d + \left(\lceil \frac{d}{\ell^{k-1}} \rceil - 1 \right) C_\ell(k-1) + 4I_\ell(k-1) \\ & \leq d + \left(\lceil \frac{d}{\ell^{k-1}} \rceil - 1 \right) (4c_\ell(k-1) + 2)\ell^{k-2} + 4c_\ell \ell^{k-1} \\ & \leq d \left(1 + \frac{4c_\ell k + 2}{\ell} \right), \end{aligned} \tag{since $\ell^{k-1} \leq d/\ell$ }$$

When $d \geq \ell^k \geq k^k$ the multiplicative stretch is $1 + (1 + o(1))4k/\ell$, where the $o(1) = O(1/\ell)$ tends to zero as ℓ increases. When $\ell \geq (4c_\ell k + 2)/\epsilon$ the multiplicative stretch becomes $1 + \epsilon$.

One may confirm that by setting $\ell = \lfloor d^{\frac{1}{k}} \rfloor$, in all the cases above the stretch function of $S(k, r)$ can be expressed as $f(d) = d + O(kd^{1-\frac{1}{k}} + 3^k)$ for $\ell^k \leq d \leq r^k$ and $f(d) = d + O(kd/r)$ for $d \geq r^k$. The leading constants in the terms $O(kd^{1-\frac{1}{k}})$ and $O(kd/r)$ tend to 4 as d increases. \square

Setting $r = (4k + O(1))/\epsilon$, we obtain a $(1 + \epsilon, O(k/\epsilon)^{k-1})$ -spanner with size $O((k/\epsilon)^k n^{1+\frac{1}{2k+1-1}})$. This spanner is sparsest when $\epsilon > 0$ is a fixed constant and $k = \log_2 \log_2 n - O(1)$: it is then a $(1 + \epsilon, ((4 + o(1)) \log \log n)^{\log \log n - O(1)})$ -spanner with size $O(n(\log \log n)^{7/4})$. When $k = \log \log n$ it is possible to reduce the size of this spanner to $O(kn + nr^{3/4}) = O(n(\log \log n + (\epsilon^{-1} \log \log n)^{3/4}))$. The kn term reflects the cost of the paths $\{P(u, p_i(u))\}_{u \in V, i \in [1, k]}$. Rather than equalize the remaining contribution of E'_0, \dots, E'_k , one chooses the sampling probabilities such that $|\tilde{E}_1|$ and $|E'_0|$ are balanced and $|E'_2|, |E'_3|, \dots, |E'_k|$ decay geometrically.

Even sparser $(1 + \epsilon, \beta)$ -spanners are known, but they have slightly worse tradeoffs. Pettie [49] constructed a $(1 + \epsilon, O(\epsilon^{-1} \log \log n)^{\log \log n})$ -spanner with size $O(n \log \log(\epsilon^{-1} \log \log n))$.

4. Lower bounds for hopsets. In this section, we show lower bounds on the tradeoffs between β and ϵ in (β, ϵ) -hopsets, subject to an upper bound on the number of edges in the hopset. We begin by making some minor modifications to the construction of the lower bound graphs $\{\mathcal{H}_k\}_k$ from section 2 and then prove lower bounds on hopsets for \mathcal{H}_k .

4.1. A new construction of \mathcal{H}_k . In the base case $k = 1$, redefine $\mathcal{H}_1[p]$ to be a copy of $\tilde{B}[p]$ on $\ell + 1$ layers (rather than a biclique), each edge of which has unit length. Naturally \mathcal{P}_1 is $\mathcal{P}(\tilde{B}[p])$. Rather than have p^2 edges and p^2 pairs in its pair-set, the new $\mathcal{H}_1[p]$ has $p^{2-o(1)}$ edges and $p^{2-o(1)}$ pairs in \mathcal{P}_1 when $\ell = p^{o(1)}$. The graph $\mathcal{H}_k[p]$ is formed as before, by taking a copy of $\tilde{B}[p]$ and replacing each vertex in an interior layer with a standard or reversed copy of $\mathcal{H}_{k-1}[p']$, where $p' = p/\xi_\ell(\sqrt{p})$. Rather than subdivide edges of $\tilde{B}[p]$ into paths of length $(2\ell - 1)^{k-1}$, we leave them as is, but give them weight $(2\ell - 1)^{k-1}$. The construction of \mathcal{P}_k from \mathcal{P}_{k-1} is exactly as in section 2. When $\ell = p^{o(1)}$, the size of $\mathcal{H}_k[p]$ and $\mathcal{P}_k[p]$ only differ from the old $\mathcal{H}_k[p]$ and $\mathcal{P}_k[p]$ (from section 2) by $p^{o(1)}$ factors.

LEMMA 4.1. *If $(u, v) \in \mathcal{P}_k$, then there is a unique shortest path from u to v in \mathcal{H}_k . If $k = 1$ the path has length exactly ℓ , and if $k \geq 2$ the path has length exactly $(2k - 1)\ell(2\ell - 1)^{k-1}$ and passes through $2\ell - 1$ copies of \mathcal{H}_{k-1} .*

Proof. The proof follows the same lines as Lemma 2.3. Let d_k be the distance between the input ports and output ports in \mathcal{H}_k . Then

$$\begin{aligned} d_1 &= \ell, \\ d_k &= (2\ell - 1)d_{k-1} + 2\ell(2\ell - 1)^{k-1}, \end{aligned}$$

and the claim follows by induction on k . \square

4.2. Simplifying the hopset H . Consider a hopset H for \mathcal{H}_k . In order to simplify the arguments to come we will manipulate H so that it satisfies certain structural properties.

DEFINITION 4.2. *Let H be a hopset of \mathcal{H}_k .*




FIG. 4. Top: An order i short edge (u, v) joining two vertices in adjacent copies of \mathcal{H}_{i-1} , the edge joining these copies being (\hat{u}, \hat{v}) . Bottom: Replacing (u, v) with three hops $(u, \hat{u}), (\hat{u}, \hat{v}), (\hat{v}, v)$. If (order $i - 1$) edges $(u, \hat{u}), (\hat{v}, v)$ are still short, then they are processed recursively. Note: \hat{u} and \hat{v} are input/output ports in copies of \mathcal{H}_{i-1} ; only u and v may be contained in copies of \mathcal{H}_{i-2} .

1. An edge $(u, v) \in H$ has order i , $1 \leq i \leq k$, if u and v are contained in a single copy of \mathcal{H}_i within \mathcal{H}_k .
2. Suppose $(u, v) \in H$ has order i . If u and v are in adjacent copies of \mathcal{H}_{i-1} (or u is in a copy of \mathcal{H}_{i-1} and v is an adjacent input/output port of the copy of \mathcal{H}_i containing it), then (u, v) is short. Otherwise (u, v) is long.

Later it will be convenient to assume that H contains only long edges. Lemma 4.3 shows that short edges can be expunged from H without affecting β and $|H|$ by more than a constant factor.

LEMMA 4.3. Let H be a (β, ϵ) hopset for \mathcal{H}_k . Then there is an $(O(k\beta), \epsilon)$ hopset H' for \mathcal{H}_k containing only long edges, with $|H'| \leq 2|H|$.

Proof. Let $(u, v) \in H$ be an order i short edge connecting adjacent copies of \mathcal{H}_{i-1} , and let (\hat{u}, \hat{v}) be the edge joining these copies. See Figure 4. Replace (u, v) in H with edges $(u, \hat{u}), (\hat{v}, v)$. Any path formerly using (u, v) can now use three edges in its place: $(u, \hat{u}), (\hat{u}, \hat{v}), (\hat{v}, v)$. Observe that $(u, \hat{u}), (\hat{v}, v)$ are order $i - 1$ edges, which may be short order $i - 1$ edges. If (u, \hat{u}) and/or (\hat{v}, v) are short, recursively process them in the same way. Whereas processing (u, v) spawned two edges, processing (u, \hat{u}) or (\hat{v}, v) spawns a single edge since \hat{u} and \hat{v} are input/output ports of copies of \mathcal{H}_{i-1} , and not contained in any copy of \mathcal{H}_{i-2} . Thus, after this recursive process completes, each original short edge (u, v) is simulated by a path with at most $O(k)$ hops. \square

Henceforth we only consider hopsets for \mathcal{H}_k that contain only long edges.

4.3. Tradeoffs between β and ϵ . We next assign ownership of each long order i edge $(x, y) \in H$ to a pair in \mathcal{P}_k . Suppose that G_{i-1}^x and G_{i-1}^y are the copies of \mathcal{H}_{i-1} containing x and y , respectively. Let $P \in \mathcal{P}_k$ own (x, y) if the unique shortest path for P intersects both G_{i-1}^x and G_{i-1}^y . It is not obvious how to assign ownership over short edges. Lemma 4.4 motivates our procedure for expunging short edges by showing that each remaining long edge is owned by at most one pair in \mathcal{P}_k .

LEMMA 4.4. Each long edge $(x, y) \in H$ is owned by at most one pair in \mathcal{P}_k .

Proof. Suppose (x, y) has order i . Let G_i be the copy of \mathcal{H}_i containing x, y and G_{i-1}^x, G_{i-1}^y be the copies of \mathcal{H}_{i-1} within G_i containing x and y . Each pair in \mathcal{P}_k has a unique shortest path in \mathcal{H}_k ; if it intersects G_i , then it enters and exits G_i by a unique

(input port, output port) pair, which is included in \mathcal{P}_i . Thus, it suffices to prove that at most one pair in \mathcal{P}_i has a shortest path intersecting both G_{i-1}^x, G_{i-1}^y . Since (x, y) is long, G_{i-1}^x, G_{i-1}^y are not adjacent; i.e., the corresponding nodes \bar{x}, \bar{y} in \tilde{B} are at distance at least 2. In order for a path in \mathcal{P}_i to intersect G_{i-1}^x, G_{i-1}^y the edges on the path between \bar{x} and \bar{y} must be labeled alternately with two labels a, b . The triple \bar{x}, a, b uniquely determines the input port and output port in G_i and therefore uniquely determines a member of \mathcal{P}_k that owns (x, y) . If the shortest path between \bar{x} and \bar{y} is not labeled alternately with two labels a, b , then no pair in \mathcal{P}_k owns (x, y) . \square

If the size of the hopset H is strictly less than $|\mathcal{P}_k|$, then some pair in \mathcal{P}_k must not own any edges. Lemma 4.5 shows that for any pair with this property, it is impossible to get below additive error $2(\ell + 1)^{k-1}$ via a path having at most $(\ell - 1)^k$ hops.

LEMMA 4.5 (compare to Lemma 2.5). *Let H be a hopset for \mathcal{H}_k containing only long edges, and let $(u, v) \in \mathcal{P}_k$ be a pair that owns no edges in H . Then we have*

$$\text{dist}_{\mathcal{H}_k \cup H}^{((\ell-1)^k)}(u, v) \geq \text{dist}_{\mathcal{H}_k}(u, v) + 2(\ell + 1)^{k-1}.$$

Proof. The proof is by induction over k .

Base case. When $k = 1$, \mathcal{H}_1 is a layered bipartite graph, so we have either

$$\text{dist}_{\mathcal{H}_k \cup H}^{(\ell-1)}(u, v) = \text{dist}_{\mathcal{H}_k}(u, v)$$

or

$$\text{dist}_{\mathcal{H}_k \cup H}^{(\ell-1)}(u, v) \geq \text{dist}_{\mathcal{H}_k}(u, v) + 2,$$

so it suffices to rule out the former possibility. We have $\text{dist}_{\mathcal{H}_1}(u, v) = \ell$; thus, the shortest u - v path in $\mathcal{H}_k \cup H$ using at most $\ell - 1$ hops must include at least one edge in H . All edges in H have order 1. Since (u, v) owns no edge in H , there is no edge $(x, y) \in H$ with x, y on the unique shortest u - v path. It follows that $\text{dist}_{\mathcal{H}_1 \cup H}^{(\ell-1)}(u, v) \neq \text{dist}_{\mathcal{H}_1}(u, v)$ and the base case is complete.

Inductive step. We now argue the inductive step. Let U be a u - v path in $\mathcal{H}_k \cup H$ that uses at most $(\ell - 1)^k$ hops and P be the corresponding path in \mathcal{H}_k , i.e., the one obtained by replacing each H -edge in U with a shortest path between its endpoints. Finally, define \hat{P} to be the projection of P onto $\tilde{B}[p]$. We consider two cases depending on whether U uses at least one order k edge from H or not.

Inductive step: Case 1. Suppose that U includes an edge $(x, y) \in H$ of order k . In this case we do not need the inductive hypothesis. Since the pair (u, v) does not own (x, y) , this means that x (or y) is in a copy of \mathcal{H}_{k-1} that is disjoint from the unique shortest u - v path in \mathcal{H}_k . Thus, \hat{P} is not equal to the unique shortest u - v path in \tilde{B} . Since \tilde{B} is bipartite, the length of \hat{P} is at least $2 + \text{dist}_{\tilde{B}}(u, v)$. Each of these two edges has weight $(2\ell - 1)^{k-1}$ in \mathcal{H}_k , so P (and U) have length at least $\text{dist}_{\mathcal{H}_k}(u, v) + 2(2\ell - 1)^{k-1}$. This same analysis applies whenever \hat{P} is not identical to the shortest u - v path in \tilde{B} .

Inductive step: Case 2. Suppose that U contains no edges of H with order k . By the above analysis, we can restrict our attention to the case when $\hat{P} = (u = u_0, u_1, \dots, u_{2\ell} = v)$ is the shortest u - v path in \tilde{B} . Let $G(u_j)$ be the copy of \mathcal{H}_{k-1} substituted for u_j and U_j . It follows that U contains all the weighted edges joining consecutive $G(u_{j-1}), G(u_j)$, and some paths $\{U_j\}$ joining an input port and output port of $G(u_j)$. Moreover, these input/output port pairs must be in \mathcal{P}_{k-1} .

Partition the $\{U_j\}_{1 \leq j \leq 2\ell-1}$ based on whether their hop count is at most $(\ell-1)^{k-1}$ or at least $(\ell-1)^{k-1} + 1$. There can be at most $\lfloor \frac{(\ell-1)^k - 2\ell}{(\ell-1)^{k-1} + 1} \rfloor = \ell - 2$ subpaths in the second category, meaning at least $(2\ell - 1) - (\ell - 2) = \ell + 1$ of the subpaths in $\{U_j\}$ use at most $(\ell-1)^{k-1}$ hops. Applying the inductive hypothesis to these subpaths, we have

$$\text{dist}_{\mathcal{H}_k \cup H}^{((\ell-1)^k)}(u, v) \geq \text{dist}_{\mathcal{H}_k}(u, v) + (\ell + 1) \cdot 2(\ell + 1)^{k-2} = \text{dist}_{\mathcal{H}_k}(u, v) + 2(\ell + 1)^{k-1}. \square$$

We are finally ready to show the following.

THEOREM 4.6. *Fix a positive integer k and parameter $\epsilon > 1/n^{o(1)}$. Any construction of (β, ϵ) -hopsets with size $n^{1 + \frac{1}{2^{k-1}} - \delta}$, $\delta > 0$, has $\beta = \Omega_k(\frac{1}{\epsilon})^k$.*

Proof. Construct $\mathcal{H}_k[p]$ with respect to an $\ell = p^{o(1)}$ to be determined, so $|\mathcal{P}_k| = n^{1 + \frac{1}{2^{k-1}} - o(1)}$. Let H be a (β, ϵ) hopset for \mathcal{H}_k containing only long edges. If $|H| \leq n^{1 + \frac{1}{2^{k-1}} - \delta}$ for some $\delta > 0$, then $|H| < |\mathcal{P}_k|$, meaning some pair $(u, v) \in \mathcal{P}_k$ owns no H edges. By Lemma 4.5, we then have

$$\text{dist}_{\mathcal{H}_k \cup H}^{(\ell-1)^k}(u, v) = \text{dist}_{\mathcal{H}_k}(u, v) + 2(\ell + 1)^{k-1}.$$

By Lemma 4.1, $\text{dist}_{\mathcal{H}_k}(u, v) = (2k - 1)\ell(2\ell - 1)^{k-1}$. Thus the relative error $\hat{\epsilon}$ of any $(\ell - 1)^k$ -hop path is

$$\hat{\epsilon} = \frac{2(\ell + 1)^{k-1}}{(2k - 1)\ell(2\ell - 1)^{k-1}} > \frac{2}{(2k - 1)\ell 2^{k-1}}.$$

We choose ℓ as a function of k and ϵ so that $\ell < \frac{1}{2^{k-2}(2k-1)\epsilon}$, which implies $\hat{\epsilon} > \epsilon$. In order for H to be a (β, ϵ) -hopset for \mathcal{H}_k , it must be that $\beta > (\ell - 1)^k = \Omega_k(1/\epsilon)^k$. \square

Observe that Theorem 4.6 implies several interesting corollaries: any (β, ϵ) -hopset with $\beta = o(1/\epsilon)$ must have size $\Omega(n^{2-o(1)})$ and any such hopset with $\beta = o(1/\epsilon^2)$ must have size $\Omega(n^{4/3-o(1)})$.

Remark 4.7. The analysis of Lemma 4.5 still has some slack in it, which introduces the 2^{k-1} factor error in Theorem 4.6. A more careful analysis will most likely reduce the hop lower bound to $\beta = \Omega(1/(k\epsilon))^k$, mimicking the dependency on k from Theorems 2.10 and 2.11.

Remark 4.8. The construction of \mathcal{H}_k from this section was essentially the same as that of section 2, except in the base case. Had we used section 2's definition of $\mathcal{H}_1[p]$ (a $K_{p,p}$ biclique, rather than a copy of $\hat{B}[p]$), Theorem 4.6 would have arrived at a weaker lower bound on $\beta = \Omega_k(1/\epsilon)^{k-1}$.

5. Lower bounds on compressing high-girth graphs. The density of the graph \mathcal{H}_k constructed in section 2 comes exclusively from complete bipartite graphs (copies of \mathcal{H}_1); that is, \mathcal{H}_k has girth 4. This feature of the construction turns out to be absolutely essential. Baswana et al. [9] showed that the class of graphs with girth (length of the shortest cycle) larger than 4 contains additive spanners below the 4/3 threshold. For example, graphs with girth 5 contain additive 12-spanners with size $O(n^{6/5})$.

THEOREM 5.1 (see [9]). *For any integer $\gamma \geq 1$, any graph with girth at least $2\gamma + 1$ contains an additive 6γ -spanner on $O(n^{1 + \frac{1}{2\gamma+1}})$ edges.*

In this section, we extend our lower bound technique to show that the exponent of Theorem 5.1 is optimal. More generally, we establish a hierarchy of tradeoffs for sublinear additive graph compression schemes that depend on k and γ . From a technical point of view, this section highlights two degrees of freedom that were not used in section 2 or [2]. We use alternative base-case graphs (rather than bicliques \mathcal{H}_1) and form $\tilde{B}[p]$ from an *imbalanced* product of $\tilde{B}[p_1]$ and $\tilde{B}[p_2]$, where $p = p_1 p_2$. Our construction uses a slightly stronger, but equivalent, statement of the Girth Conjecture that asserts a lower bound on the degree rather than the total size.

CONJECTURE 5.2 (Girth Conjecture [33, 17, 15]). *For any integer $\gamma \geq 1$, there exists a graph with n vertices, girth $2\gamma + 2$, and minimum degree $\Omega(n^{1/\gamma})$.*

Our goal is to prove the following theorem. Observe that by setting $k = 2$, Theorem 5.3 implies that the exponent of Theorem 5.1 cannot be improved.

THEOREM 5.3. *Fix integers $\gamma \geq 1, k \geq 2$. Consider any data structure that answers approximate distance queries for the class of n -vertex undirected graphs with girth at least $2\gamma + 1$. Assuming the Girth Conjecture, if the stretch function of the data structure is*

$$f(d) < d + c_{k,\gamma} d^{1-1/k} \text{ for } c_{k,\gamma} \approx \frac{2}{(\gamma(k-1))^{1-1/k}} \text{ and } d \text{ sufficiently large,}$$

then on some graph the data structure occupies at least $\Omega(n^{1+\frac{1}{(\gamma+1)2^{k-1}-1}-o(1)})$ bits.

We remark that this theorem holds trivially for any superconstant γ (with a sufficiently large $n^{-o(1)}$ factor), since the lower bound becomes $\Omega(n)$. Thus, we treat γ as a fixed constant throughout this section. The remainder of this section constitutes a proof of Theorem 5.3. We first make a simple observation about the hypothesized graphs from the Girth Conjecture.

Observation 5.4. Fix a $\gamma \geq 1$. If G has girth $2\gamma + 2$ and minimum degree $\Omega(n^{1/\gamma})$, there are $\Omega(n^{d/\gamma})$ nodes at distance exactly d from any node u for any $0 \leq d \leq \gamma$.

Recall that the reason $\mathcal{H}_1[p]$ from section 2 was useful was because its edge-set was covered by unique, disjoint shortest paths between p input ports and p output ports. We will prove something analogous for high-girth graphs, but with these properties weakened slightly, in ways that have no adverse effect on the overall construction.

DEFINITION 5.5. *In a graph G , we say that a pair of nodes s, t requires an edge e if every shortest path from s to t includes e .*

LEMMA 5.6. *Fix integers $\gamma \geq 1$ and $0 \leq i < \gamma$. Assuming the Girth Conjecture, there is a graph $H = (V, E)$ on n vertices and $\Omega(n^{1+1/\gamma})$ edges with girth $2\gamma + 2$, and disjoint node subsets $S, T \subseteq V$ of sizes $|S| = \Omega(n^{i/\gamma}), |T| = \Omega(n^{(\gamma+1-i)/\gamma})$ such that each edge $e \in E$ is required by some $(s, t) \in S \times T$ with $\text{dist}(s, t) = \gamma$.*

Proof. Let H initially be any graph with minimum degree $\Omega(n^{1/\gamma})$ and girth $2\gamma + 2$. Sample node subsets S, T independently and uniformly at random of the appropriate size and let $P \subset S \times T$ be such that $(s, t) \in P$ if and only if $\text{dist}(s, t) = \gamma$. The shortest s - t path is unique, due to H 's girth, so all its edges are required by (s, t) . Discard from H all edges not required by any pair in P .

We will now prove that any particular edge remains in H with constant probability, so there exists some choice of S, T for which at least a constant fraction of the edges are retained. Let us consider an arbitrary edge (u, v) in H . If there exist nodes $s \in S, t \in T$ such that $\text{dist}_{H-\{(u,v)\}}(s, u) + \text{dist}_{H-\{(u,v)\}}(t, v) = \gamma - 1$, then $\text{dist}(s, t) = \gamma$ and the unique s - t shortest path requires (u, v) .

Let A be all vertices at distance exactly $\gamma - i$ from u in $H - \{(u, v)\}$ and B be all vertices at distance exactly $i - 1$ from v in $H - \{(u, v)\}$. By Observation 5.4, $|A| = \Omega(n^{\frac{\gamma-i}{\gamma}})$ and $|B| = \Omega(n^{\frac{i-1}{\gamma}})$. Since $|A||S|$ and $|B||T|$ are both $\Omega(n)$, with constant probability $A \cap S \neq \emptyset$ and $B \cap T \neq \emptyset$. \square

LEMMA 5.7. *Let $H = (V, E)$ be a graph and $S, T \subseteq V$ be node subsets as described in Lemma 5.6. Then there exists a set $\mathcal{P}_1^\gamma \subseteq S \times T$ of size $|\mathcal{P}_1^\gamma| = \Omega(n^{1+1/\gamma}/\gamma^3)$, as well as a mapping $\phi : \mathcal{P}_1^\gamma \rightarrow E$, with the following two properties:*

- For each $(s, t) \in \mathcal{P}_1^\gamma$, the pair (s, t) requires the edge $\phi(s, t)$.
- For each $(s, t) \in \mathcal{P}_1^\gamma$, the unique shortest path from s to t in H contains no edge e such that $e = \phi(s', t')$ for some $(s', t') \neq (s, t)$.

Proof. Let $P \subseteq S \times T$ be the set of s, t pairs for which $\text{dist}(s, t) = \gamma$. Since $P = \Theta(n^{1+1/\gamma})$, each node pair in P has distance γ , and H has $\Omega(n^{1+1/\gamma})$ edges, it follows that the average edge in H is required by $c = O(\gamma)$ different pairs in P . By Markov’s inequality, at most half the edges in H are required by more than $2c$ pairs; let F be the set of edges required by $2c$ or fewer node pairs in $S \times T$. We then have $|F| = \Omega(n^{1+1/\gamma})$.

We now build \mathcal{P}_1^γ and ϕ using the following process. Iterate through the pairs in P in any order. For each $(s, t) \in P$, let Q be the unique shortest path from s to t . If $F \cap Q = \emptyset$, i.e., if (s, t) requires no F -edges, then discard (s, t) from P . Otherwise, include (s, t) in \mathcal{P}_1^γ , set $\phi(s, t)$ to be any edge in $F \cap Q$, and discard from P any other pair that requires any edge in $F \cap Q$. Since $|Q| \leq \gamma$ and edges in F are required by at most $2c$ pairs, we discard $O(c\gamma)$ pairs for each (s, t) not discarded.

The necessary properties of \mathcal{P}_1^γ are immediate from the construction. To bound the size of $|\mathcal{P}_1^\gamma|$, first note that $\Omega(n^{1+1/\gamma}/\gamma)$ pairs in P require at least one edge in F , since $|F| = \Omega(n^{1+1/\gamma})$. Of these $\Omega(n^{1+1/\gamma}/\gamma)$ node pairs, each one added to \mathcal{P}_1^γ causes at most $O(c\gamma) = O(\gamma^2)$ to be discarded, so $|\mathcal{P}_1^\gamma| = \Omega(n^{1+1/\gamma}/\gamma^3)$. \square

The lower bound construction. The graph $\hat{B}[p]$ is defined exactly as before, and the parameter $\ell \geq 2$ is fixed throughout. The graph \hat{B} can now be formed from an imbalanced product. Construct $\hat{B}[p_1, p_2]$ from copies of $\hat{B}[p_1]$ and $\hat{B}[p_2]$ in exactly the same way that $\hat{B}[p]$ is constructed from two copies of $\hat{B}[\sqrt{p}]$. The number of vertex layers in $\hat{B}[p_1, p_2]$ is still $2\ell + 1$, and each layer contains $p_1 p_2$ vertices. However, a node in an internal layer has $|\mathcal{L}[p_1]|$ neighbors in the previous layer and $|\mathcal{L}[p_2]|$ neighbors in the next layer (or vice versa), so the density of $\hat{B}[p_1, p_2]$ is determined by $\max\{p_1, p_2\}$.

We define $\mathcal{H}_1^\gamma[p_1, p_2]$ to be a graph drawn from Lemma 5.6, with $n = n(p_1, p_2)$ vertices and input/output ports S, T selected with the following cardinality. When $\gamma \geq 3$ is odd,

$$|S| = |T| = p_1 = p_2 = n^{(\gamma+1)/(2\gamma)},$$

which, in Lemma 5.6, corresponds to choosing $i = \frac{\gamma+1}{2}$. When $\gamma \geq 2$ is even,

$$|S| = p_1 = n^{(\gamma+2)/(2\gamma)} \quad \text{and} \quad |T| = p_2 = n^{1/2},$$

which corresponds to picking $i = \frac{\gamma+2}{2}$ in Lemma 5.6. We define \mathcal{P}_1^γ to be the set of $\Omega(p_1 p_2)$ node pairs in $\mathcal{H}_1^\gamma[p_1, p_2]$ from Lemma 5.7.

We proceed as in Theorem 5.3, but with a few critical differences. Although \mathcal{P}_1^γ -paths through \mathcal{H}_1^γ have length γ , there could be “shortcuts” between input ports and output ports not covered by \mathcal{P}_1^γ ; the length of a shortcut might be as low as 1. When forming \mathcal{H}_k^γ we subdivide edges in \hat{B} as before, but we have to make these paths a factor γ longer to sufficiently penalize paths that attempt to deviate far from the

unique shortest path and thereby take advantage of many shortcuts elsewhere in \mathcal{H}_k^γ . The imbalanced product $\ddot{B}[p_1, p_2]$ is only used in the formation of $\mathcal{H}_2[p]$, and only when γ is even.

When γ is odd, $\mathcal{H}_2^\gamma[p]$ is constructed from $\ddot{B}[p] = \ddot{B}[\sqrt{p}, \sqrt{p}]$ and $\mathcal{H}_1^\gamma[\frac{\sqrt{p}}{\xi_\ell(\sqrt{p})}, \frac{\sqrt{p}}{\xi_\ell(\sqrt{p})}]$ exactly as in section 2, but replacing edges in \ddot{B} by paths of length $\gamma(2\ell - 1)$. When γ is even, to construct $\mathcal{H}_2^\gamma[p]$ we pick p_1, p_2 to have the “right” proportions such that $p_1 p_2 = p$. The right proportions are dictated by the function $\xi_\ell(\cdot)$ from the construction of \ddot{B} and Lemma 5.6. Let $p'_1 = |\mathcal{L}[p_1]| \geq p_1/\xi_\ell(p_1)$ and $p'_2 = |\mathcal{L}[p_2]| \geq p_2/\xi_\ell(p_2)$ be the number of edges connecting an internal node u in $\ddot{B}[p_1, p_2]$ to previous/subsequent layers. When forming \mathcal{H}_2^γ , each of these edges gets attached to a different input/output port of $\mathcal{H}_1^\gamma(u)$, so we need $p'_1 = (p'_2)^{(\gamma+2)/\gamma}$. In cases where $\xi_\ell(p) = p^{o(1)}$, we can ignore the distinction between p_1 and p'_1 and just set $p_1 = p^{\frac{\gamma+2}{2\gamma+2}}$ and $p_2 = p^{\frac{\gamma}{2\gamma+2}}$. When $k \geq 3$, $\mathcal{H}_k^\gamma[p]$ is constructed from $\ddot{B}[p]$ and $\mathcal{H}_{k-1}^\gamma[\cdot]$ as before but subdivides edges into paths of length $\gamma(2\ell - 1)^{k-1}$.

We now analyze the distances in \mathcal{H}_k^γ of pairs in \mathcal{P}_k^γ .

LEMMA 5.8 (compare to Lemma 2.3). *Fix a $(u_0, u_{2\ell}) \in \mathcal{P}_k^\gamma$ whose unique shortest path in \ddot{B} is $(u_0, u_1, \dots, u_{2\ell})$. The following hold.*

- *There is a unique u_0 - $u_{2\ell}$ shortest path in \mathcal{H}_k^γ . It has length $\gamma(2(k - 1)\ell + 1)(2\ell - 1)^{k-1}$.*
- *Any path from u_0 to $u_{2\ell}$ in \mathcal{H}_k^γ that intersects some $\mathcal{H}_{k-1}^\gamma(u')$, $u' \notin \{u_1, \dots, u_{2\ell-1}\}$, is at least $2(2\ell - 1)^{k-1}$ longer than the shortest path.*

Proof. The proof is by induction. When $k = 1$, $(u_0, u_{2\ell}) \in \mathcal{P}_1^\gamma$ implies that $\text{dist}_{\mathcal{H}_1^\gamma}(u_0, u_{2\ell}) = \gamma$ (by Lemma 5.6), and by the girth of \mathcal{H}_1^γ the path is unique.

We now turn to the inductive step. By the inductive hypothesis, there is only one shortest path that passes through $\mathcal{H}_{k-1}^\gamma(u_1), \dots, \mathcal{H}_{k-1}^\gamma(u_{2\ell-1})$, and it has length

$$d_k^\gamma = (2\ell - 1)d_{k-1}^\gamma + (2\ell) \cdot \gamma(2\ell - 1)^{k-1},$$

which has a closed form solution $d_k^\gamma = \gamma(2(k - 1)\ell + 1)(2\ell - 1)^{k-1}$.

By Lemma 2.3, any path from an input port of \mathcal{H}_k^γ to an output port of \mathcal{H}_k^γ passes through $(2\ell - 1)^{k-1}$ copies of \mathcal{H}_1^γ . Thus, the *minimum* length of such a path is exactly $d_k^\gamma - (\gamma - 1)(2\ell - 1)^{k-1}$. Consider a path that passes through some $\mathcal{H}_{k-1}^\gamma(u')$, where $u' \notin \{u_1, \dots, u_{2\ell-1}\}$. Since the shortest u_0 - $u_{2\ell}$ path in \ddot{B} is unique and \ddot{B} is bipartite, this path traverses at least two additional subdivided edges, each of length $\gamma(2\ell - 1)^{k-1}$. The length of such a path is therefore at least

$$d_k^\gamma - (\gamma - 1)(2\ell - 1)^{k-1} + 2 \cdot \gamma(2\ell - 1)^{k-1},$$

which is at least $d_k^\gamma + 2(2\ell - 1)^{k-1}$. Thus the shortest u_0 - $u_{2\ell}$ path in \mathcal{H}_k^γ is unique. \square

Since paths through \mathcal{H}_1^γ overlap, we need to update the definition of a “critical” edge.

DEFINITION 5.9 (compare to Definition 2.4). *An edge e is critical for a pair $(u_0, u_{2\ell}) \in \mathcal{P}_k^\gamma$ if it lies in a copy of \mathcal{H}_1^γ , the unique shortest u_0 - $u_{2\ell}$ path in \mathcal{H}_k^γ enters and leaves that copy of \mathcal{H}_1^γ by some pair $(s, t) \in \mathcal{P}_1^\gamma$, and we have $\phi(s, t) = e$. (Thus (s, t) requires e in \mathcal{H}_1^γ , and so $(u_0, u_{2\ell})$ requires e in \mathcal{H}_k^γ .)*

LEMMA 5.10. *Let $\tilde{\mathcal{H}}_k^\gamma$ be \mathcal{H}_k^γ with all critical edges for $(u_0, u_{2\ell})$ removed. Then $\text{dist}_{\tilde{\mathcal{H}}_k^\gamma}(u_0, u_{2\ell}) \geq \text{dist}_{\mathcal{H}_k^\gamma}(u_0, u_{2\ell}) + 2(2\ell - 1)^{k-1}$.*

Proof. We proceed by induction. In the base case of $k = 1$, suppose a critical edge is removed for $(u_0, u_{2\ell})$. Since $\text{dist}_{\mathcal{H}_1^\gamma}(u_0, u_{2\ell}) = \gamma$ and \mathcal{H}_1^γ has girth $2\gamma + 2$, $\text{dist}_{\tilde{\mathcal{H}}_1^\gamma}(u_0, u_{2\ell}) \geq \gamma + 2$. For the inductive step, let Q be the shortest path from u_0 to $u_{2\ell}$ in \mathcal{H}_k^γ , and let \tilde{Q} be the shortest path in the graph $\tilde{\mathcal{H}}_k^\gamma$. Suppose first that \tilde{Q} traverses the exact same copies of \mathcal{H}_{k-1}^γ that Q traverses. In this case the claim follows from the inductive hypothesis: we accumulate $2(2\ell - 1)^{k-2}$ additive stretch in each of the $2\ell - 1$ copies of \mathcal{H}_{k-1}^γ traversed. If \tilde{Q} deviates and intersects some other copy of \mathcal{H}_{k-1}^γ , then by Lemma 5.8, the additive stretch is at least $2(2\ell - 1)^{k-1}$. \square

LEMMA 5.11. *Let $\tilde{\mathcal{H}}_k^\gamma$ be \mathcal{H}_k^γ with the critical edges for all pairs in \mathcal{P}_k^γ except $(u_0, u_{2\ell})$ removed. Then $\text{dist}_{\tilde{\mathcal{H}}_k^\gamma}(u_0, u_{2\ell}) = \text{dist}_{\mathcal{H}_k^\gamma}(u_0, u_{2\ell})$.*

Proof. Let s, t be the input/output ports used by shortest paths from u_0 to $u_{2\ell}$ in any internal copy of \mathcal{H}_1^γ . By Lemma 5.7, the unique shortest path from s to t in \mathcal{H}_1^γ does not include any edge e for which $\phi(s', t') = e, (s', t') \neq (s, t)$. It follows that the distance from s to t is the same in \mathcal{H}_k^γ and $\tilde{\mathcal{H}}_k^\gamma$. \square

The final piece of the proof is exactly identical to the lower bound argument in section 2. In particular, we define a family of $2^{|\mathcal{P}_k^\gamma|}$ graphs by keeping/removing the critical edges for each pair in \mathcal{P}_k^γ in all possible combinations. By Lemmas 5.10 and 5.11, any two of these graphs will disagree on a pairwise distance $(u_0, u_{2\ell})$ by an additive $2(2\ell - 1)^{k-1}$. By Lemma 5.8, we have $\text{dist}_{\mathcal{H}_k^\gamma}(u_0, u_{2\ell}) = d = \gamma(2(k - 1)\ell + 1)(2\ell - 1)^{k-1}$. Thus, the additive stretch $2(2\ell - 1)^{k-1}$ is roughly $\frac{2}{(\gamma(k-1))^{1-1/k}} \cdot d^{1-1/k}$. If the stretch function of the distance oracle is $f(d) \leq d + cd^{1-1/k}$ for sufficiently large d and a sufficiently small constant $c < \frac{2}{(\gamma(k-1))^{1-1/k}}$, then it cannot map any of these graphs to the same bit-string. This gives the stretch part of the lower bound claimed in Theorem 5.3. It remains only to compute the size of this graph family. We have $2^{|\mathcal{P}_k^\gamma|}$ distinct graphs, so we need to obtain a lower bound on $|\mathcal{P}_k^\gamma|$. In order to avoid tedious calculations let us assume that $\ell = p^{o(1)}$, so $\xi_\ell(p) = p^{o(1)}$ as well. In particular, $\dot{n}[p], \ddot{n}[p] = p^{1+o(1)}, \ddot{m}[p_1, p_2] = p_1 p_2 (p_1 + p_2)^{1-o(1)}$, and $|\mathcal{P}(\tilde{B}[p_1, p_2])| = (p_1 p_2)^{2-o(1)}$. Letting $n_1^\gamma[p_1, p_2]$ be the number of vertices in $\mathcal{H}_1^\gamma[p_1, p_2]$ and $n_k^\gamma[p]$ be the number of vertices in $\mathcal{H}_k^\gamma[p]$, we have

$$\begin{aligned} n_1^\gamma[p_1, p_2] &= (p_1 p_2)^{\gamma/(\gamma+1)}, \\ n_2^\gamma[p] &= p^{1+o(1)} \cdot n_1^\gamma[p_1^{1-o(1)}, p_2^{1-o(1)}] + (\ddot{m}[p_1, p_2])^{1+o(1)}, \end{aligned}$$

where $p_1, p_2 = \sqrt{p}$ if γ is odd and $p_1 = p^{\frac{\gamma+2}{2\gamma+2}}, p_2 = p^{\frac{\gamma}{2\gamma+2}}$ if γ is even. When $k \geq 3$,

$$n_k^\gamma[p] = p^{1+o(1)} \cdot n_{k-1}^\gamma[p^{1/2-o(1)}] + (\ddot{m}[p])^{1+o(1)}.$$

Whether γ is even or odd, $n_1^\gamma[p_1^{1-o(1)}, p_2^{1-o(1)}] = (p_1 p_2)^{\gamma/(\gamma+1)-o(1)}$. The density of $\tilde{B}[p_1, p_2]$ is maximized when p_1 and p_2 are most imbalanced. This occurs when $\gamma = 2, p_1 = p^{2/3}$, and $p_2 = p^{1/3}$, making $\ddot{m}[p_1, p_2] = p^{5/3-o(1)}$. Thus, for any $\gamma \geq 2$, $n_2^\gamma[p] = p^{2-\frac{1}{\gamma+1}+o(1)}$. By induction on k , $n_k^\gamma[p] = p^{2-\frac{1}{(\gamma+1)2^{k-2}}+o(1)}$.

By Lemma 5.7, $|\mathcal{P}_1^\gamma[p_1, p_2]| = \Omega(p_1 p_2)$. The same inductive proof from section 2 shows that for any $k \geq 2$, $|\mathcal{P}_k^\gamma[p]| = p^{2-o(1)}$. Expressed in terms of $n = n_k^\gamma[p], p^{2-o(1)}$ is $n^{1+\frac{1}{(\gamma+1)2^{k-1}-1}-o(1)}$. Theorem 5.3 follows.

5.1. Matching upper bounds. The subgraph E'_0 from section 3 can be viewed as a *radius-1 clustering* of the graph, obtained from the following procedure. First,

cluster centers $V_1 \subset V$ are sampled with probability q_1 . Each vertex incident to V_1 is *clustered* and joins the cluster of one such adjacent V_1 vertex. E'_0 contains a star spanning each cluster and all edges incident to unclustered vertices, which number $O(n/q_1)$ in expectation. Baswana et al. [9] observed that in graphs with girth at least $2\gamma+1$, this procedure can be generalized to compute a radius- γ clustering with similar properties.

THEOREM 5.12 (see [9]). *Let $G = (V, E)$ be a graph with girth at least $2\gamma + 1$. Fix $q_1 < 1$ and let $V_1 \subset V$ be obtained by sampling each element of V with probability q_1 . Any $v \in V$ with $\text{dist}(v, V_1) \leq \gamma$ is clustered and joins the cluster of the closest V_1 vertex, breaking ties consistently. Let E_0^γ contain a radius- γ tree on each cluster and all edges incident to unclustered vertices. In expectation $|E_0^\gamma| = O(n/q_1^{1/\gamma})$.*

We can use the E_0^γ from Theorem 5.12 in lieu of E_0 in the construction of Thorup-Zwick emulators [59]. The total size of the emulator $E^\gamma = E_0^\gamma \cup E_1 \cup \dots \cup E_k$ is then on the order of

$$\frac{n}{q_1^{1/\gamma}} + \frac{nq_1^2}{q_2} + \dots + \frac{nq_{k-1}^2}{q_k} + (nq_k)^2.$$

If we write q_i as $n^{-g(i)}$, $g(i)$ must satisfy the following.

$$\begin{aligned} g(i) &= 2g(i-1) + \frac{g(1)}{\gamma} && \text{(balancing } E_0^\gamma \text{ and } E_{i-1}, \text{ for } i \in [2, k]) \\ &= (\gamma+1)(2^{i-1} - 1)\frac{g(1)}{\gamma}. && \text{(by induction)} \end{aligned}$$

Balancing the size of E_k and E_0^γ , we have

$$|E_k| = (nq_k)^2 = n^{2 - ((\gamma+1)2^k - 2)\frac{g(1)}{\gamma}} = n^{1 + \frac{g(1)}{\gamma}} = |E_0^\gamma|,$$

which is satisfied when $\frac{g(1)}{\gamma} = \frac{1}{(\gamma+1)2^k - 1}$, implying the size of the emulator is $O(kn^{1 + \frac{1}{(\gamma+1)2^k - 1}})$. The analysis of the emulator proceeds exactly as in section 3, by bounding the quantities $C_\ell(i)$ and $I_\ell(i)$ inductively. Substituting E_0^γ for E_0 only affects the following base cases.

$$\begin{aligned} C_\ell(0) &= 0 && \text{for all } \ell, \\ I_\ell(0) &= \gamma && \text{for all } \ell. \end{aligned}$$

This is justified by Theorem 5.12. Any path with length $\ell^0 = 1$ is a single edge, say (u, v) . If u is unclustered in E_0^γ , then $(u, v) \in E^\gamma$, $\text{dist}_{E^\gamma}(u, v) = 1$, and (u, v) is *complete*. On the other hand, if u is clustered, then $\text{dist}_{E^\gamma}(u, p_1(u)) \leq \gamma$ and (u, v) is *incomplete*. With these base cases it is straightforward to show that the stretch function for E^γ is $f(d) = d + O(\gamma kd^{1-1/k})$. For example, when $\gamma = k = 2$ we see that every girth-5 graph has a $(d + O(\sqrt{d}))$ -emulator with size $O(n^{12/11})$.

This emulator can be converted to a $(1 + \epsilon, O(\gamma k/\epsilon)^{k-1})$ -spanner by applying the same transformations from section 3, using Theorem 5.13 in lieu of Theorem 3.1.

THEOREM 5.13. *Let $G = (V, E)$, q_1, V_1 , and E_0^γ be as in Theorem 5.12. Suppose V_2 is obtained by sampling each element of V with probability q_2 , where $q_2 < q_1$. There is an edge-set E_1^γ with expected size $O(\gamma^2 n q_1^2 / q_2)$ such that for $u, v \in V_1$ and $v \in \mathcal{B}_2(u)$,*

$$\text{dist}_{E_0^\gamma \cup E_1^\gamma}(u, v) \leq 2\gamma.$$

Remark 5.14. The γ^2 factor arises from two parts of the path-buying algorithm's analysis that depend on the cluster radii. The *cost* of a path (number of missing edges) is at most the number of clusters touching the path divided by the cluster diameter, 2γ . Once a cluster-pair is charged we have their correct distance to within $+O(\gamma)$. The path-buying algorithm only charges this cluster pair again when the distance improves, so at most $O(\gamma)$ times.

The base-case values for $C_\ell(i)$ and $I_\ell(i)$ are updated as follows. For all ℓ ,

$$\begin{aligned} C_\ell(0) &= 0, \\ I_\ell(0) &= \gamma, \\ C_\ell(1) &= 6\gamma, \\ I_\ell(1) &= \ell + 3\gamma. \end{aligned}$$

It is easy to check that these base cases increase $I_\ell(i) - \ell^i$ by a factor of γ . It is for this reason that we use a slightly larger threshold $(r + 2\gamma)^i$ when forming E'_i in the following construction. The spanner $S(k, r, \gamma)$ has the edge-set $E_0^\gamma \cup \tilde{E}_1^\gamma \cup E'_2 \cup \dots \cup E'_k$, where E'_i is obtained by replacing each (weighted) pair $(u, v) \in E_i$ with a shortest path $P(u, v)$, assuming $\text{dist}(u, v)$ is sufficiently short.

$$E'_i = \bigcup_{\substack{(u,v) \in E_i : \\ \text{dist}(u,v) \leq (r+2\gamma)^i}} P(u, v).$$

It follows that the size of the spanner is on the order of⁹

$$\frac{n}{q_1^{1/\gamma}} + \frac{nq_1^2}{q_2} + \frac{nq_2^2 r^2}{q_3} \dots \frac{nq_{k-1}^2 r^{k-1}}{q_k} + (nq_k)^2 r^k.$$

Writing $q_i = n^{-g(i)} r^{-h(i)}$, $g(i)$ satisfies the same recurrence as before and $h(i)$ satisfies the following.

$$\begin{aligned} h(2) &= 2h(1) + \frac{h(1)}{\gamma}, && \text{(balancing } E_0^\gamma \text{ and } \tilde{E}_1^\gamma) \\ h(i) &= 2h(i-1) + \frac{h(1)}{\gamma} - (i-1) && \text{(balancing } E_0^\gamma \text{ and } E'_{i-1}, i \in [3, k]) \\ &= ((\gamma + 1)2^{i-1} - 1) \frac{h(1)}{\gamma} - 3 \cdot 2^{i-2} + (i+1). && \text{(by induction, for } i \in [3, k]) \end{aligned}$$

Finally, we balance E_0^γ and E'_k ,

$$\begin{aligned} |E'_k| &= (nq_k)^2 r^k = n^{2 - ((\gamma+1)2^k - 2) \frac{g(1)}{\gamma}} r^{k - [((\gamma+1)2^k - 2) \frac{h(1)}{\gamma} - 3 \cdot 2^{k-1} + 2(k+1)]} = n^{1 + \frac{g(1)}{\gamma}} r^{1 + \frac{h(1)}{\gamma}} \\ &= |E_0^\gamma|, \end{aligned}$$

by setting $g(1), h(1)$ as follows:

$$\begin{aligned} \frac{g(1)}{\gamma} &= \frac{1}{(\gamma + 1)2^k - 1}, \\ \frac{h(1)}{\gamma} &= \frac{3 \cdot 2^{k-1} - (k + 2)}{(\gamma + 1)2^k - 1}. \end{aligned}$$

⁹For simplicity we treat the γ^2 factor in $|\tilde{E}_1^\gamma|$ as a constant.

Thus, the size of the resulting spanner is $O(r^h kn^{1+\frac{1}{(\gamma+1)2^{k-1}}})$, where $h = \frac{3 \cdot 2^{k-1} - (k+2)}{(\gamma+1)2^{k-1}}$. For example, setting $k = \gamma = 2$ and $r = \sqrt{D}$, this shows that every graph with girth 5 contains a subgraph that functions like a $(d + O(\sqrt{d}))$ -spanner for all $d \leq D$, with size $O(r^{\frac{2}{11}} n^{\frac{12}{11}}) = O(D^{\frac{1}{11}} n^{\frac{12}{11}})$.

Theorem 5.15 summarizes our emulator and spanner constructions for high-girth graphs. The size can be further reduced by a factor of k using the sampling technique of Huang and Pettie [36].

THEOREM 5.15. *Let G be a graph with girth at least $2\gamma + 1$. There is an additive 4γ -emulator and additive 6γ -spanner for G with size $O(n^{1+\frac{1}{2\gamma+1}})$. For any integer $k \geq 2$, there is a $(d + O(\gamma kd^{1-1/k}))$ -emulator for G with size $O(kn^{1+\frac{1}{(\gamma+1)2^{k-1}}})$ and a $(1 + \epsilon, O(\gamma k/\epsilon)^{k-1})$ -spanner for G with size $O((\gamma k/\epsilon)^h kn^{1+\frac{1}{(\gamma+1)2^{k-1}}})$, where $h = \frac{3 \cdot 2^{k-1} - (k+2)}{(\gamma+1)2^{k-1}} < \frac{3}{2(\gamma+1)}$.*

6. Lower bounds on shortcutting digraphs. In this section we consider directed unweighted graphs $G = (V, E)$. Let $u \rightsquigarrow v$ be the reachability (transitive closure) relation for G , indicating a directed path from u to v . In 1992 Thorup [55] conjectured that for any directed graph $G = (V, E)$ there exists another $G' = (V, E')$ such that (i) G and G' have the same reachability relation (\rightsquigarrow), (ii) $|E'| \leq 2|E|$, and (iii) every $u \rightsquigarrow v$ is witnessed in G' by a directed path with length $\text{poly}(\log n)$; this is called the *diameter* of G' . Thorup's conjecture was confirmed for trees [55, 57, 19] and planar graphs [56], but was disproved in a strong form by Hesse [35], who showed that there exists a G with $n^{1+\epsilon}$ edges and diameter $n^{\delta(\epsilon)}$, such that any shortcutting G' with diameter $o(n^\delta)$ requires $\Omega(n^{2-\epsilon})$ edges. In this section we give a simpler proof of Hesse's result—a refutation of Thorup's conjecture—by generalizing the construction of \dot{B} from section 2.

6.1. The construction. Recall that $\dot{B}[p]$ is parameterized by an integer $\ell \geq 2$. Its vertex set is partitioned into $\ell + 1$ layers of p vertices; each vertex has $p/\xi_\ell(p)$ edges leading to the next layer, each of which is assigned a distinct label from the set $\mathcal{L}[p]$. Here $\xi_\ell(p) = 2^{\Theta(\sqrt{\log \ell \log p})}$. The set $\mathcal{P} = \mathcal{P}(\dot{B}[p])$ consists of $p^2/\xi_\ell(p)$ pairs, each having a unique length- ℓ shortest path. Each element $(u, v) \in \mathcal{P}$ is generated by picking a vertex u in the first layer and a label $a \in \mathcal{L}[p]$; v is the vertex in the last layer reached by starting at u and repeatedly following edges labeled a . In this section we regard $\dot{B}[p]$ as being a directed acyclic graph, with all edges oriented toward the higher numbered layer.

Rather than form $\dot{B}[p]$ by taking the product of two copies of $\dot{B}[\sqrt{p}]$, we take the product of k copies of $\dot{B}[p^{1/k}]$. Let the layers of $\dot{B}[p^{1/k}]$ be L_0, \dots, L_ℓ . The vertex set of $\dot{B}[p]$ is partitioned into layers $\ddot{L}_0, \dots, \ddot{L}_{k\ell}$. If q is written $ik + j$, where $i \leq \ell, j < k$, \ddot{L}_q is the set $L_{i+1}^j \times L_i^{k-j}$. A directed edge $(\nu, \nu') \in \ddot{L}_q \times \ddot{L}_{q+1}$ exists if ν and ν' only differ in their j th component and $(\nu[j], \nu'[j])$ is in the edge-set of $\dot{B}[p^{1/k}]$. In other words, a path from layer \ddot{L}_0 to layer \ddot{L}_q in \dot{B} simulates k independent paths, from layer L_0 to L_i in $k - j$ copies of \dot{B} , and from layer L_0 to L_{i+1} in the remaining j copies. The pair-set $\mathcal{P}(\dot{B}[p])$ is defined as one might expect:

$$\mathcal{P}(\dot{B}[p]) = \{(\nu, \nu') \mid (\nu[j], \nu'[j]) \in \mathcal{P}(\dot{B}) \text{ for each } 0 \leq j < k\}.$$

Thus, for any $(\nu, \nu') \in \mathcal{P}(\dot{B}[p])$, $\nu \rightsquigarrow \nu'$ is witnessed by a unique path having length $k\ell$, and the labels along this path form a periodic sequence $(a_0, a_1, \dots, a_{k-1}, a_0, a_1, \dots)$ for some $(a_0, \dots, a_{k-1}) \in (\mathcal{L}[p^{1/k}])^k$. The size of the pair-set $\mathcal{P} = \mathcal{P}(\dot{B}[p])$ is $|\mathcal{P}| \geq (p^{2/k}/\xi)^k = p^2/\xi^k$, where $\xi = \xi_\ell[p^{1/k}]$.

LEMMA 6.1. *The diameter of $G = \ddot{B}[p]$ is $k\ell$. Any graph G' with the same transitive closure as G and diameter $k\ell/(k-1) - 1$ must have at least $|\mathcal{P}|$ edges.*

Proof. Let $(\nu, \nu') \in \mathcal{P}$, P be the unique ν - ν' path in G , and P' be a ν - ν' path in G' having length strictly shorter than $k\ell/(k-1)$. It must be that P' contains an edge that shortcuts at least k consecutive edges in P . However, because the edge-labels along P are periodic with length k , any length- k subpath of P uniquely identifies P . Thus, no edge of G' that shortcuts k or more edges in G can be used by two distinct pairs in \mathcal{P} . It follows that G' contains at least $|\mathcal{P}|$ edges or the diameter of G' is at least $k\ell/(k-1)$. \square

The number of vertices and edges in $\ddot{B}[p]$ is $n = (k\ell+1)p$ and $m = k\ell p^{1+1/k} / \xi_\ell(p^{1/k})$, respectively. Setting $\ell = p^{\delta/k}$ for some small $\delta > 0$ we have $\xi_\ell(p^{1/k}) = 2^{\Theta(\sqrt{\log \ell \cdot \log p^{1/k}})} = p^{\Theta(\sqrt{\delta/k})}$, so the density of the graph is roughly $p^{(1-\Theta(\sqrt{\delta}))/k} \approx n^{\frac{1-\Theta(\sqrt{\delta})}{k(1+\delta/k)}}$. By Lemma 6.1, in order to reduce the diameter to $p^{\delta/k} = O(n^{\frac{\delta}{k(1+\delta/k)}})$ we need to add $|\mathcal{P}| = p^2 / \xi^k = p^{2-\Theta(\sqrt{\delta})}$ shortcuts. By setting $\delta = O(\epsilon^2)$ to be sufficiently small and $k = \Omega(1/\epsilon)$, we arrive at the same conclusion of Hesse [35].

THEOREM 6.2 (see [35]). *For any $\epsilon > 0$ there exist a $\delta = \delta(\epsilon)$ and a directed graph G with n vertices, at most $n^{1+\epsilon}$ edges, and diameter n^δ with the following property. Any graph G' with the same transitive closure as G and diameter $o(n^\delta)$ must contain at least $n^{2-\epsilon}$ edges.*

7. Conclusion. In this paper, we characterized the optimal asymptotic behavior of sublinear additive stretch functions f for spanners, emulators, or any graph compression scheme. Roughly speaking, any representation using $n^{1+\frac{1}{2k-1}-\delta}$ bits (for any $\delta > 0$) must have stretch function $f(d) = d + \Omega(d^{1-\frac{1}{k}})$. Previous constructions of sublinear additive emulators [59] and $(1+\epsilon, \beta)$ -spanners ([31, 59] and the construction of section 3) show that neither the exponent $1 + \frac{1}{2k-1}$ nor additive stretch $\Omega(d^{1-\frac{1}{k}})$ can be improved for any k .

The main distinction between $(1+\epsilon, \beta)$ -spanners [31, 59, 49] and sublinear additive emulators/spanners [59, 49, 20] is that constructions of the former take ϵ as a parameter (which affects the size of the spanner), whereas the latter have $(1+\epsilon, \beta)$ -stretch for all ϵ ; that is, ϵ can be chosen in the analysis. An interesting open question is whether one can match the size-stretch tradeoff of Thorup and Zwick’s optimal emulators [59] with a spanner. (Constructions in [49, 20] are off, from [59] (and our lower bounds), by a polynomial factor.) It would be possible to construct such spanners given a pairwise spanner with a sublinear additive stretch function. For example, when $S \subset V$ with $|S| = \Omega(n^{4/7})$ and $P = S \times S$, does there exist a pairwise spanner for P with stretch $d + O(\sqrt{d})$ and size $O(|P|)$? If such an object existed, we would immediately have an optimal $(d + O(\sqrt{d}))$ -spanner with size $O(n^{8/7})$; see [59, 49].

Our lower bounds match the existing upper bounds in the distance regime $2^{\Omega(k)} \ll d < n^{o(1)}$, while they say nothing when $d = 2^{O(k)}$, and they are weaker when $d = n^{\Omega(1)}$. An interesting open problem is to understand the sparseness-stretch tradeoffs available when $d = O(2^k)$ is tiny (see [30, 9, 45]) and when $d = n^{\Omega(1)}$ is very large [16, 14].

Acknowledgments. We are grateful to Virginia Vassilevska Williams for useful technical discussions and for advice about the directions taken by this paper. We thank Michael Elkin for proposing the question of finding a lower bound hierarchy for mixed spanners.

REFERENCES

- [1] A. ABOUD AND G. BODWIN, *Error amplification for pairwise spanner lower bounds*, in Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM, New York, SIAM, Philadelphia, 2016, pp. 841–854, <https://doi.org/10.1137/1.9781611974331.ch60>.
- [2] A. ABOUD AND G. BODWIN, *The 4/3 additive spanner exponent is tight*, J. ACM, 64 (2017), 28.
- [3] I. ABRAHAM AND C. GAVOILLE, *On approximate distance labels and routing schemes with affine stretch*, in Proceedings of the 25th International Symposium on Distributed Computing (DISC), Springer, Berlin, 2011, pp. 404–415.
- [4] R. AGARWAL, *The space-stretch-time tradeoff in distance oracles*, in Proceedings of the 22nd Annual European Symposium on Algorithms (ESA), Lecture Notes in Comput. Sci. 8737, Springer, Berlin, 2014, pp. 49–60.
- [5] R. AGARWAL AND P. B. GODFREY, *Distance oracles for stretch less than 2*, in Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM, New York, SIAM, Philadelphia, 2013, pp. 526–538.
- [6] 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, <https://doi.org/10.1137/S0097539796303421>.
- [7] N. ALON, *Testing subgraphs in large graphs*, in Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science (FOCS), IEEE, Washington, DC, 2001, pp. 434–441.
- [8] I. ALTHÖFER, G. DAS, D. DOBKIN, D. JOSEPH, AND J. SOARES, *On sparse spanners of weighted graphs*, Discrete Comput. Geom., 9 (1993), pp. 81–100.
- [9] S. BASWANA, T. KAVITHA, K. MEHLHORN, AND S. PETTIE, *Additive spanners and (α, β) -spanners*, ACM Trans. Algorithms, 7 (2010), 5.
- [10] S. BASWANA AND S. SEN, *A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs*, Random Structures Algorithms, 30 (2007), pp. 532–563.
- [11] F. BEHREND, *On sets of integers which contain no three terms in arithmetic progression*, Proc. Natl. Acad. Sci., 32 (1946), pp. 331–332.
- [12] C. BENSON, *Minimal regular graphs of girth eight and twelve*, Canad. J. Math., 18 (1966), pp. 1091–1094.
- [13] G. BODWIN AND V. VASSILEVSKA WILLIAMS, *Very sparse additive spanners and emulators*, in Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science (ITCS), ACM, New York, 2015, pp. 377–382.
- [14] G. BODWIN AND V. VASSILEVSKA WILLIAMS, *Better distance preservers and additive spanners*, in Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM, New York, SIAM, Philadelphia, 2016, pp. 855–872, <https://doi.org/10.1137/1.9781611974331.ch61>.
- [15] B. BOLLOBÁS, *Extremal Graph Theory*, London Math. Soc. Monogr. 11, Academic Press, London, 1978.
- [16] B. BOLLOBÁS, D. COPPERSMITH, AND M. ELKIN, *Sparse distance preservers and additive spanners*, SIAM J. Discrete Math., 19 (2005), pp. 1029–1055, <https://doi.org/10.1137/S0895480103431046>.
- [17] J. BONDY AND M. SIMONOVITS, *Cycles of even length in graphs*, J. Combin. Theory Ser. B, 16 (1974), pp. 97–105.
- [18] W. G. BROWN, *On graphs that do not contain a Thomsen graph*, Canad. Math. Bull., 9 (1966), pp. 281–285.
- [19] B. CHAZELLE, *Computing on a free tree via complexity-preserving mappings*, Algorithmica, 2 (1987), pp. 337–361.
- [20] S. CHECHIK, *New additive spanners*, in Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM, New York, SIAM, Philadelphia, 2013, pp. 498–512, <https://doi.org/10.1137/1.9781611973105.36>.
- [21] S. CHECHIK, *Approximate distance oracles with improved bounds*, in Proceedings of the 47th Annual ACM on Symposium on Theory of Computing (STOC), ACM, New York, 2015, pp. 1–10.
- [22] E. COHEN, *Using selective path-doubling for parallel shortest-path computations*, J. Algorithms, 22 (1997), pp. 30–56.
- [23] E. COHEN, *Polylog-time and near-linear work approximation scheme for undirected shortest-paths*, J. ACM, 47 (2000), pp. 132–166.
- [24] H. COHEN AND E. PORAT, *On the Hardness of Distance Oracle for Sparse Graph*, preprint, <https://arxiv.org/abs/1006.1117>, 2010.

- [25] D. COPPERSMITH AND M. ELKIN, *Sparse sourcewise and pairwise distance preservers*, SIAM J. Discrete Math., 20 (2006), pp. 463–501, <https://doi.org/10.1137/050630696>.
- [26] M. CYGAN, F. GRANDONI, AND T. KAVITHA, *On pairwise spanners*, in Proceedings of the 30th International Symposium on Theoretical Aspects of Computer Science (STACS), Kiel, Germany, 2013, pp. 209–220.
- [27] D. DOR, S. HALPERIN, AND U. ZWICK, *All-pairs almost shortest paths*, SIAM J. Comput., 29 (2000), pp. 1740–1759, <https://doi.org/10.1137/S0097539797327908>.
- [28] M. ELKIN, *Personal communication*, 2013.
- [29] M. ELKIN AND O. NEIMAN, *Hopsets with constant hopbound, and applications to approximate shortest paths*, in Proceedings of the 57th IEEE Symposium on Foundations of Computer Science (FOCS), IEEE, Washington, DC, 2016, pp. 128–137.
- [30] M. ELKIN AND D. PELEG, $(1 + \epsilon, \beta)$ -spanner constructions for general graphs, in Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC), ACM, New York, 2001, pp. 173–182.
- [31] M. ELKIN AND D. PELEG, $(1 + \epsilon, \beta)$ -spanner constructions for general graphs, SIAM J. Comput., 33 (2004), pp. 608–631, <https://doi.org/10.1137/S0097539701393384>.
- [32] M. ELKIN AND S. PETTIE, *A linear-size logarithmic stretch path-reporting distance oracle for general graphs*, ACM Trans. Algorithms, 12 (2016), 50.
- [33] P. ERDŐS, *Extremal problems in graph theory*, in Theory of Graphs and Its Applications (Proc. Sympos. Smolenice, 1963), Publ. House Czechoslovak Acad. Sci., Prague, 1963, pp. 29–36.
- [34] P. ERDŐS, A. RÉNYI, AND V. T. SÓS, *On a problem of graph theory*, Studia Sci. Math. Hungar., 1 (1966), pp. 215–235.
- [35] W. HESSE, *Directed graphs requiring large numbers of shortcuts*, in Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM, New York, SIAM, Philadelphia, 2003, pp. 665–669.
- [36] S.-E. HUANG AND S. PETTIE, *Thorup-Zwick Emulators Are Universally Optimal Hopsets*, preprint, <https://arxiv.org/abs/1705.00327>, 2017.
- [37] T. KAVITHA, *New pairwise spanners*, in Proceedings of the 32nd International Symposium on Theoretical Aspects of Computer Science (STACS), Munich, Germany, 2015, pp. 513–526.
- [38] T. KAVITHA AND N. M. VARMA, *Small stretch pairwise spanners and approximate D -preservers*, SIAM J. Discrete Math., 29 (2015), pp. 2239–2254, <https://doi.org/10.1137/140953927>.
- [39] P. N. KLEIN AND S. SUBRAMANIAN, *A randomized parallel algorithm for single-source shortest paths*, J. Algorithms, 25 (1997), pp. 205–220.
- [40] M. B. T. KNUDSEN, *Additive spanners: A simple construction*, in Proceedings of the 14th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT), Copenhagen, Denmark, 2014, pp. 277–281.
- [41] F. LAZEBNIK AND V. A. USTIMENKO, *New examples of graphs without small cycles and of large size*, European J. Combin., 14 (1993), pp. 445–460.
- [42] F. LAZEBNIK, V. A. USTIMENKO, AND A. J. WOLDAR, *A new series of dense graphs of high girth*, Bull. Amer. Math. Soc., 32 (1995), pp. 73–79.
- [43] F. LAZEBNIK, V. A. USTIMENKO, AND A. J. WOLDAR, *A characterization of the components of the graphs $d(k, q)$* , Discrete Math., 157 (1996), pp. 271–283.
- [44] J. MATOUŠEK, *On the distortion required for embedding finite metric spaces into normed spaces*, Israel J. Math., 93 (1996), pp. 333–344.
- [45] M. PARTER, *Bypassing Erdős’ girth conjecture: Hybrid stretch and sourcewise spanners*, in Proceedings of the 41st International Colloquium Automata, Languages, and Programming (ICALP), Copenhagen, Denmark, 2014, pp. 608–619.
- [46] M. PĂTRAȘCU AND L. RODITTY, *Distance oracles beyond the Thorup–Zwick bound*, SIAM J. Comput., 43 (2014), pp. 300–311, <https://doi.org/10.1137/11084128X>.
- [47] M. PATRAȘCU, L. RODITTY, AND M. THORUP, *A new infinity of distance oracles for sparse graphs*, in Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS), IEEE, Washington, DC, 2012, pp. 738–747.
- [48] D. PELEG AND A. A. SCHAFFER, *Graph spanners*, J. Graph Theory, 13 (1989), pp. 99–116.
- [49] S. PETTIE, *Low distortion spanners*, ACM Trans. Algorithms, 6 (2010), 7.
- [50] E. PORAT AND L. RODITTY, *Preprocess, set, query!*, Algorithmica, 67 (2013), pp. 516–528.
- [51] I. REIMAN, *Über ein Problem von K. Zarankiewicz*, Acta. Math. Acad. Sci. Hungary, 9 (1958), pp. 269–273.
- [52] L. RODITTY, M. THORUP, AND U. ZWICK, *Deterministic constructions of approximate distance oracles and spanners*, in Proceedings of the 32nd International Colloquium on Automata, Languages, and Programming (ICALP), Lisbon, Portugal, 2005, pp. 261–272.
- [53] H. SHI AND T. H. SPENCER, *Time-work tradeoffs of the single-source shortest paths problem*, J. Algorithms, 30 (1999), pp. 19–32.

- [54] C. SOMMER, E. VERBIN, AND W. YU, *Distance oracles for sparse graphs*, in Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS), IEEE, Washington, DC, 2009, pp. 703–712.
- [55] M. THORUP, *On shortcutting digraphs*, in Proceedings of the 18th International Workshop on Graph Theoretic Concepts in Computer Science (WG), Wiesbaden-Naurod, Germany, 1992, pp. 205–211.
- [56] M. THORUP, *Shortcutting planar digraphs*, *Combin. Probab. Comput.*, 4 (1995), pp. 287–315.
- [57] M. THORUP, *Parallel shortcutting of rooted trees*, *J. Algorithms*, 23 (1997), pp. 139–159.
- [58] M. THORUP AND U. ZWICK, *Approximate distance oracles*, *J. ACM*, 52 (2005), pp. 1–24.
- [59] 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), ACM, New York, SIAM, Philadelphia, 2006, pp. 802–809.
- [60] J. TITS, *Sur la trichotomie et certains groupes qui s'en déduisent*, *Inst. Hautes Études Sci. Publ. Math.*, 1959 (1959), pp. 13–60.
- [61] J. D. ULLMAN AND M. YANNAKAKIS, *High-probability parallel transitive-closure algorithms*, *SIAM J. Comput.*, 20 (1991), pp. 100–125, <https://doi.org/10.1137/0220006>.
- [62] R. WENGER, *Extremal graphs with no C^4 s, C^6 s, or C^{10} s*, *J. Combin. Theory Ser. B*, 52 (1991), pp. 113–116.
- [63] A. WOLDAR AND V. USTIMENKO, *An application of group theory to extremal graph theory*, in *Group Theory*, Proceedings of the Ohio State-Denison Conference, World Scientific, River Edge, NJ, 1993, pp. 293–298.
- [64] D. P. WOODRUFF, *Additive spanners in nearly quadratic time*, in Proceedings of the 37th International Colloquium on Automata, Languages and Programming (ICALP), Bordeaux, France, 2010, pp. 463–474.