

Better Distance Preservers and Additive Spanners

GREG BODWIN, University of Michigan EECS, USA

VIRGINIA VASSILEVSKA WILLIAMS, MIT EECS, USA

We study two popular ways to sketch the shortest path distances of an input graph. The first is *distance preservers*, which are sparse subgraphs that agree with the distances of the original graph on a given set of demand pairs. Prior work on distance preservers has exploited only a simple structural property of shortest paths, called *consistency*, stating that one can break shortest path ties such that no two paths intersect, split apart, and then intersect again later. We prove that consistency alone is not enough to understand distance preservers, by showing both a lower bound on the power of consistency and a new general upper bound that polynomially surpasses it. Specifically, our new upper bound is that any p demand pairs in an n -node undirected unweighted graph have a distance preserver on $O(n^{2/3}p^{2/3} + np^{1/3})$ edges. We leave a conjecture that the right bound is $O(n^{2/3}p^{2/3} + n)$ or better.

The second part of this paper leverages these distance preservers in a new construction of *additive spanners*, which are subgraphs that preserve all pairwise distances up to an additive error function. We give improved error bounds for spanners with relatively few edges; for example, we prove that all graphs have spanners on $O(n)$ edges with $+O(n^{3/7+\epsilon})$ error. Our construction can be viewed as an extension of the popular path-buying framework to clusters of larger radii.

CCS Concepts: • **Mathematics of computing** → **Paths and connectivity problems**; **Extremal graph theory**; **Graph algorithms**;

Additional Key Words and Phrases: Spanners, distance preservers, sparsification

ACM Reference format:

Greg Bodwin and Virginia Vassilevska Williams. 2021. Better Distance Preservers and Additive Spanners. *ACM Trans. Algorithms* 17, 4, Article 36 (October 2021), 24 pages.
<https://doi.org/10.1145/3490147>

1 INTRODUCTION

How much can graphs be compressed while keeping their distance information roughly intact? This question falls within the scope of both metric embeddings and graph theory and is fundamental to our understanding of the metric properties of graphs. When the compressed version of the graph must be a subgraph, it is called a *spanner*. Spanners have a multitude of applications, essentially everywhere shortest path information needs to be compressed while still allowing for graph algorithms to be run, and they have been intensively studied theoretically [2–5, 7, 15–17, 19, 28–30]. The quality of a spanner is measured by the tradeoff between its sparsity

Preliminary version appeared in SODA 2016. Research performed while authors were employed by Stanford University. Authors' addresses: G. Bodwin, University of Michigan EECS, 2260 Hayward St, Ann Arbor, MI 48109, USA; email: bodwin@umich.edu; V. V. Williams, MIT EECS, 32 Vassar St, Cambridge, MA 02139, USA; email: virgi@mit.edu. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

1549-6325/2021/10-ART36 \$15.00

<https://doi.org/10.1145/3490147>

and its accuracy in preserving distances. There are several different versions of spanners; we will discuss two of them below.

1.1 Distance Preservers

One possible version of the spanner problem is *distance preservers*, in which we must preserve *some* of the pairwise distances *exactly*. All graphs in this paper are undirected and unweighted.

Definition 1 (Distance Preservers [12]). Given a graph $G = (V, E)$ and a set $P \subseteq V \times V$ of demand pairs, we say that a subgraph $H \subseteq G$ is a *distance preserver* of G, P if $\text{dist}_H(u, v) = \text{dist}_G(u, v)$ for all $(u, v) \in P$.

Distance preservers are fundamental combinatorial objects: they generalize BFS trees (which are distance preservers for the special case $P = \{s\} \times V$), they have applications to many other types of graph sketches [1–3, 6, 10, 14, 18, 21, 27], and they have been recently popular objects of study in their own right [6, 8, 9, 12–14, 20, 22, 25]. Since a distance preserver must contain the edges of a shortest path for each demand pair, without loss of generality constructing a distance preserver is equivalent to choosing a *tiebreaking scheme*, i.e., a way to select one of the possibly many shortest paths for each demand pair, and then overlaying the chosen shortest paths to obtain the preserver. A basic property a tiebreaking scheme can have is *consistency*, meaning that the intersection of any two chosen shortest paths π_1, π_2 is a single (possibly empty) continuous subpath of each. One major reason consistency is fundamental is because it is the property that gives rise to BFS trees: that is, for demand pairs $\{s\} \times V$, one gets a distance-preserving tree if and only if one overlays shortest paths according to a consistent tiebreaking scheme. Another reason consistency is fundamental is that it gives the current state-of-the-art upper bounds for distance preservers in general: given p demand pairs in an n -node graph, a consistent tiebreaking scheme will yield a distance preserver on $O(\min\{n + n^{1/2}p, np^{1/2}\})$ edges [12]. This result is even partially tight in the following sense: it implies that $p = O(\sqrt{n})$ demand pairs have a distance preserver on $O(n)$ edges, and this corollary cannot possibly be improved by increasing the range to any $p = \omega(n^{1/2})$ while still paying only $O(n)$ edges [12].

Given all this, it is natural to ask if we actually need to investigate any deeper properties of shortest paths, or if consistency alone will suffice to understand distance preservers. In this work, we answer this question with the following two new results:

THEOREM 1. *For any integer $p = O(n^2)$, there is an infinite family of n -node graphs, sets of p demand pairs, and consistent tiebreaking schemes for which the resulting distance preserver has $\Omega(\min\{n + n^{1/2}p, np^{1/2}\})$ edges.*

THEOREM 2. *Every n -node graph and set of p demand pairs has a distance preserver on $O(n^{2/3}p^{2/3} + np^{1/3})$ edges.*

See Figure 1 for a picture of these bounds. In particular, Theorem 1 implies that the analysis in [12] of *consistent tiebreaking* was exactly tight, but Theorem 2 shows that *different tiebreaking* improves on it polynomially – so consistency alone is not enough. The $+np^{1/3}$ term in Theorem 2 arises for purely technical reasons; it is unlikely to be really needed, but we have been unable to remove it.

CONJECTURE 3. *Every n -node graph and set of p demand pairs has a distance preserver on $O(n^{2/3}p^{2/3} + n)$ edges.*

1.2 Additive Spanners

Another type of spanner is when *all* pairwise distances must be preserved *up to an error term*:




Fig. 1. The state-of-the-art asymptotic tradeoffs for distance preservers. The upper bounds of [12] are in solid blue, of Theorem 2 in solid green, and of Conjecture 3 in dotted green. The dashed red curve is, roughly, a lower bound from [12]. Some subpolynomial improvements in [9] are not pictured here.

Definition 2 (Additive Spanners). Given a graph $G = (V, E)$, a subgraph $H \subseteq G$ is a $+\beta$ (additive) spanner of a graph G if $\text{dist}_H(u, v) \leq \text{dist}_G(u, v) + \beta$ for all $u, v \in V$.

There are many interesting constructions with β as an absolute constant: all n -node graphs have $+2$ spanners on $O(n^{3/2})$ edges [4], $+4$ spanners on $\tilde{O}(n^{7/5})$ edges [15], and $+6$ spanners on $O(n^{4/3})$ edges [7] (see also [23, 24, 32]). But unfortunately, this setting then hits a dead-end: there is *not* a general construction of $+\beta$ spanners on $O(n^{4/3-\varepsilon})$ edges, for any choice of absolute constants $\beta, \varepsilon > 0$ [2] (see also [1, 21, 31]). In particular, to obtain additive spanners in the regime $O(n^{4/3-\varepsilon})$, one essentially has to settle for $\beta = \text{poly}(n)$. Many applications call for spanners of linear size ($O(n)$ edges, or maybe a little more), so the additive error β_0 attainable by a spanner on $O(n)$ edges has attracted particular research interest. Currently, it is known that

$$\Omega(n^{1/11}) \leq \beta_0 \leq \tilde{O}(n^{1/2})$$

where the upper bound is by prior work of the authors [10] and the lower bound is by Lu [26] (see also [6, 15, 27]). We improve polynomially on the upper bound:

THEOREM 4. Assuming Conjecture 3, for any $\varepsilon > 0$ and $1 \leq \mathcal{E} \leq n^{1/3}$, every n -node graph has a $+O(n^{3/7+\varepsilon} \mathcal{E}^{-9/7})$ additive spanner on $O_\varepsilon(n\mathcal{E})$ edges. Moreover, this result holds unconditionally when $\mathcal{E} = O(1)$ or $\mathcal{E} = \Omega(n^{2/13})$.

See Figure 2 for a picture of these bounds. In particular, the new (unconditional) upper bound is $\beta_0 \leq n^{3/7+\varepsilon}$.

Technical Overview of Spanner Construction. Here we give a brief technical overview of our spanner construction, aimed at a reader with some familiarity with the area. Theorem 4 is based in part on the popular path-buying framework used in many previous spanner constructions, introduced in [7]. For background, these constructions begin with a *clustering phase*, in which (1) we add all edges incident on “low-degree” nodes to the spanner, and (2) the “high-degree” nodes are partitioned into clusters of radius 1 each. Then they have a *path-buying phase*, where some shortest paths are added to the graph with the goal of connecting the existing clusters.




Fig. 2. The state-of-the-art asymptotic tradeoffs for additive spanners. The upper part of the solid blue curve is an upper bound of [10], the lower solid blue line is the upper bound of [15], and the upper bound of Theorem 4 is in solid green, dotted in the regime where it is conditional on Conjecture 3. The dashed red line is, roughly, a lower bound from [26].

To obtain even sparser spanners at the cost of more error, it is a natural idea to try to apply the same framework, but using clusters of radius $r \gg 1$. Our key idea is that distance preservers play a key role in the proper extension of this method. Instead of differentiating between *high-degree* and *low-degree* nodes, we propose an analogous dichotomy between high-radius clusters that are *large* – meaning that they contain many nodes – and clusters that are *bottlenecked*, meaning they can be separated from the rest of the graph by removing a relatively small set of nodes. The idea nearly reduces to the high-degree/low-degree dichotomy in the extreme case of our construction where $r = 1$, but it gains efficiency for larger r . It turns out that bottlenecked clusters, analogous to low-degree nodes, can be handled cleanly with an application of distance preservers: one only has to add a distance preserver between pairs of nodes in the separating node set, and then for any shortest path, a subpath passing in and out of a bottlenecked cluster will already have all of its edges contained in a spanner.

Theorem 4 is obtained by plugging the distance preservers of Conjecture 3 into this large/bottlenecked framework, and then proceeding through a reasonable generalization of the path-buying method of [7] with a few extra optimizations. Conjecture 3 is confirmed in the settings $p = O(n^{1/2})$ and $p = \Omega(n)$, and this corresponds to the ranges $\mathcal{E} = O(1)$ and $\mathcal{E} = \Omega(n^{2/13})$ in which Theorem 4 holds unconditionally. One can also obtain a weaker yet unconditional version of Theorem 4 in the regime $\omega(1) \leq \mathcal{E} \leq o(n^{2/13})$ by plugging the unconditional distance preserver bounds known in the intermediate regime. These bounds are given in Section 3.4.

2 DISTANCE PRESERVERS

We will be concerned only with existential distance preserver bounds in this paper and not with construction runtime. We note that any existential bound is automatically algorithmic: that is, if $e^*(n, p)$ is the maximum number of edges needed for a distance preserver of p demand pairs in an n -node graph, then there is a folklore algorithm that runs in polynomial time and always outputs

a distance preserver on $\leq e^*(n, p)$ edges. This algorithm is simply to consider the edges of G in any order and delete the current edge if doing so does not change the distance between any given demand pair. By construction, the final graph H will be the unique distance preserver of itself, so it cannot have more than $e^*(n, p)$ edges.

As discussed in the introduction, it is nice theoretically to reduce the construction of distance preservers to *shortest path tiebreaking*:

Definition 3 (Tiebreaking Schemes). A tiebreaking scheme on a graph G is a map π from node pairs (u, v) to a shortest $u \rightsquigarrow v$ path $\pi(u, v)$ in G .

A subgraph is a distance preserver if it contains the edges of a shortest path for each demand pair, so w.l.o.g. we may assume that any distance preserver is constructed by selecting a tiebreaking scheme π and then overlaying the chosen shortest paths for each demand pair. In other words, given a graph G and demand pairs P , the distance preserver H associated to a given tiebreaking scheme π is the one where an edge e is included in H if there is $(u, v) \in P$ with $e \in \pi(u, v)$.

2.1 Limitations of Consistent Tiebreaking

Definition 4 (Consistency). A tiebreaking scheme π is *consistent* if, for any nodes w, x, y, z , we have that $\pi(w, x) \cap \pi(y, z)$ is a (possibly empty) continuous subpath of both $\pi(w, x)$ and $\pi(y, z)$.

It is folklore that every graph has a consistent tiebreaking scheme. For example, one is obtained by *random reweighting*, in which we add a tiny random variable to each edge weight to break ties with probability 1. In the reweighted graph G' all shortest paths are unique, and thus there is only one tiebreaking scheme π over G' . For any two paths $\pi(w, x), \pi(y, z)$ their intersection (if nonempty) is the unique shortest path between its endpoints, which is thus a continuous subpath of both $\pi(w, x)$ and $\pi(y, z)$. Hence π is consistent in G' , and so it is also consistent in G when the original unit edge weights are restored. Coppersmith and Elkin proved:

THEOREM 5 ([12]). For any p demand pairs and consistent tiebreaking scheme π in an n -node graph, the associated distance preserver has $O(\min\{n^{1/2}p + n, np^{1/2}\})$ edges.

Our first new result is that this is tight for consistency. In particular:

THEOREM 6 (SEE FIGURE 3). For any $p = O(n^2)$, there is an infinite family of n -node graphs G , consistent tiebreaking schemes π , and p demand pairs so that the associated distance preserver has $\Omega(\min\{n^{1/2}p + n, np^{1/2}\})$ edges.

PROOF. Let q be a prime and let G be the graph whose nodes are the elements of the field F_q^2 and whose edge set contains exactly the node pairs whose first coordinate differs by 1, not including “wraparound” edges of the form $((0, x), (q - 1, y))$. The demand pairs P are exactly the node pairs $((0, x), (q - 1, y))$ for any $x, y \in F_q$. By a *line* in F_q^2 we will mean a set of points that can be described by

$$\{(x, mx + b) \mid x \in F_q\}$$

for some fixed $m, b \in F_q$ (note this does not include “vertical” lines with fixed first coordinate). Each demand pair (u, v) uniquely determines a line in F_q^2 passing through both points. Define $\pi(u, v)$ as the path containing exactly the nodes on this line. Note that $\text{dist}(u, v) = q - 1$ since the first coordinate of two nodes only differs by 1 over an edge; thus, each $\pi(u, v)$ is indeed a shortest path since it has exactly $q - 1$ edges, so π is a tiebreaking scheme. Additionally, π is consistent, since the paths $\pi(u, v)$ correspond to lines in F_q^2 and any two lines intersect on only 0 or 1 nodes.

The distance preserver associated to π is G itself, since every edge of G similarly determines a line in F_q^2 which is thus equal to $\pi(u, v)$ for exactly one demand pair (u, v) . The graph G has $q^2 =: n$




Fig. 3. The construction described in Theorem 6, with $q = 7$. For clarity, only the first 4 out of the 49 shortest paths $\pi(u, v)$ are pictured here.

nodes and $(q - 1)q^2 = \Theta(n^{3/2})$ edges, and we have $|P| = q^2 = n$ demand pairs, which gives our claimed lower bound in the special case when $p = n$. To obtain our lower bound for other values of p :

- In the regime $n \leq p \leq n^2$, we modify our construction by taking only its first few layers. That is: letting $k \leq q - 1$ be a parameter of the construction, we only let the first coordinate range from 0 to k but everything else remains the same (essentially, the nodes of G are now $[k] \times F_q$). This gives a lower bound on $n := kq$ nodes, $(k - 1)q^2 = \Theta(n^2/k)$ edges, and $q^2 = n^2/k^2$ demand pairs. The lower bound thus follows by choosing $k = \Theta(n/\sqrt{p})$.
- In the regime $\sqrt{n} \leq p \leq n$, we modify our construction by arbitrarily choosing p demand pairs to keep. We then discard the rest, as well as all the edges in G that appear in $\pi(s, t)$ for any discarded demand pair (s, t) . The resulting construction has n nodes, p pairs, and since each surviving demand pair has $q - 1 = \Theta(\sqrt{n})$ distinct edges in its path, the lower bound has $\Theta(p\sqrt{n})$ edges.


Finally, the lower bound of $\Omega(n)$ in the regime $p \leq \sqrt{n}$ is trivial (let G be a big long path, and take the endpoints of this path as one of the demand pairs). $\square$

2.2 Beyond Consistent Tiebreaking


We will next prove our new upper bound on distance preservers. We first develop two technical ingredients. In the sequel it will be convenient to interpret demand pairs $(s, t) \in P$ to be ordered (e.g. (s, t) rather than (t, s)), and we will write $(u, v) \in \pi(s, t)$ to mean that the nodes u, v appear adjacently in that order in $\pi(s, t)$; that is, this notation implies that u is closer to s and v is closer to t .

Definition 5 (Branching Events [12]). Given a graph G , a set of demand pairs P , and a tiebreaking scheme π , a *branching event* is a pair of distinct edges $(u, v), (w, v)$ that share an endpoint node v , such that there are two distinct demand pairs $p, p' \in P$ with $(u, v) \in \pi(p) \setminus \pi(p')$ and $(w, v) \in \pi(p') \setminus \pi(p)$.

Due to our orientation of edges, this technically means (for example) that a directed tree rooted at a node s could have 0 branching events if the tree paths are all directed away from s , or $\Omega(n^2)$ branching events if the paths are directed towards s . So branching events depend on the ordering of demand pairs $((s, t)$ vs. $(t, s))$. One could just as easily define branching events to be two oriented



(a) If π routes $2d + 2$ paths for demand pairs starting with w to branch with another path $\pi(x, y)$...



(b) ... then there is an alternate shortest $w \rightsquigarrow v_j$ path for one of these nodes, which uses a different edge entering v_j .

Fig. 4. The pigeonhole technique used in Lemma 8.

edges leaving the same node $((v, u), (v, w))$, as opposed to two oriented edges entering the same node as we've done here; this choice is somewhat arbitrary.

LEMMA 7 ([12]). *If a tiebreaking scheme π for G, P has b branching events, then the associated distance preserver has $O((nb)^{1/2} + n)$ edges.*

PROOF. Let $H = (V, E)$ be the associated distance preserver. Pick an arbitrary total ordering of the set P and label each edge $e \in E$ with the first demand pair $(s, t) \in P$ in the ordering so that $e \in \pi(s, t)$. Since we assign e the orientation from s towards t , we can say that each node $v \in V$ is the endpoint for $\geq \binom{\deg_{in}(v)}{2}$ different branching events. We then compute:

$$\begin{aligned}
 |E| &= \sum_{v \in V} \deg_{in}(v) \\
 &= O(n) + \sum_{\{v \in V \mid \deg_{in}(v) \geq 2\}} \deg_{in}(v) \\
 &= O(n) + O\left(\left(n \cdot \sum_{\{v \in V \mid \deg_{in}(v) \geq 2\}} \deg_{in}(v)^2\right)^{1/2}\right) && \text{Cauchy-Schwarz} \\
 &= O(n) + O\left(\left(n \cdot \sum_{v \in V} \binom{\deg_{in}(v)}{2}\right)^{1/2}\right) \\
 &= O(n) + O\left((nb)^{1/2}\right). && \square
 \end{aligned}$$

LEMMA 8 (SEE FIGURE 4). *Let G be an n -node graph and let S be a subset of $|S| = \sigma$ nodes of diameter d (meaning for all $s_1, s_2 \in S$ we have $\text{dist}(s_1, s_2) \leq d$). Let $P \subseteq S \times V$ be a set of $|P| = p$ demand pairs. Then there is a distance preserver of G, P on*

$$O\left((np\sigma d)^{1/2} + n\right)$$

edges.

PROOF. Let π be a tiebreaking scheme whose associated distance preserver H has as few edges as possible. Every edge e can then be labeled with a demand pair $(s, t) \in P$ so that every $s \rightsquigarrow t$ shortest path in H includes e (if no such demand pair (s, t) exists, then we can modify π to make all paths avoid e , thus removing e from H). Each branching event (e, e') is then labeled with a pair of demand pairs $p, p' \in P$, inheriting the labels of e, e' .

We will prove this lemma by showing that there are only $O(p\sigma d)$ total branching events in H ; the lemma then follows by plugging $b = O(p\sigma d)$ into Lemma 7. Suppose towards contradiction

that we have $> p\sigma(2d+1)$ branching events in H ; then by the pigeonhole principle there is a node $w \in S$, a demand pair $(x, y) \in P$, and (at least) $2d+2$ branching events $\{b_1, \dots, b_{2d+2}\}$ so that each branching event b_i has a label of the form $(x, y), (w, z_i)$ (i.e. x, y, w are the same in the labels of all these branching events, and only z_i varies). Let v_i be the node shared by the two edges in the branching event b_i , and we will assume for convenience that the nodes $\{v_i\}$ are ordered by increasing $\text{dist}(x, v_i)$. By the triangle inequality and the diameter of S , for any $1 \leq i \leq 2d+2$ we have

$$d \geq \text{dist}(w, x) \geq \text{dist}(w, v_i) - \text{dist}(x, v_i) \geq -\text{dist}(w, x) \geq -d.$$

Hence, by the pigeonhole principle, there are $1 \leq i < j \leq 2d+2$ with

$$\text{dist}(w, v_i) - \text{dist}(x, v_i) = \text{dist}(w, v_j) - \text{dist}(x, v_j)$$

and so

$$\text{dist}(w, v_i) + \text{dist}(x, v_j) - \text{dist}(x, v_i) = \text{dist}(w, v_j).$$

Note that x, v_i, v_j all lie on $\pi(x, y)$ in that order, and so $\text{dist}(x, v_j) - \text{dist}(x, v_i) = \text{dist}(v_i, v_j)$. Therefore

$$\text{dist}(w, v_i) + \text{dist}(v_i, v_j) = \text{dist}(w, v_j).$$

This equation implies that v_i lies on a shortest $w \rightsquigarrow v_j$ path, and so we can form a shortest $w \rightsquigarrow v_j$ path by concatenating $\pi(w, v_i)$ with $\pi(x, y)[v_i \rightsquigarrow v_j]$. This concatenated path enters v_j along an edge in $\pi(x, y)$. Thus, there exists a shortest $w \rightsquigarrow v_i$ path that uses an edge in $\pi(x, y)$ to enter v_i . Hence it is *not* the case that every shortest $w \rightsquigarrow v_i$ path includes an edge e used in the branching event b_i . This contradicts the labeling of our branching events. $\square$

In order to convert this lemma to an upper bound on distance preservers in general, we will use the following standard lemma in graph theory:

LEMMA 9. *Let $G = (V, E)$ be a nonempty graph with average degree $\mathcal{E}$. Then there exists a node v that has $\Omega(\mathcal{E})$ neighbors of degree $\Omega(\mathcal{E})$ each.*

PROOF. While there exists a node of degree $\leq \mathcal{E}/4$, remove that node and all of its incident edges from G . Letting n be the initial number of nodes in G , the total number of edges removed in this way is $\leq n\mathcal{E}/4 \leq |E|/2$, so when this process halts G still has at least $|E|/2$ edges, so it is nonempty and it has minimum degree $\Omega(\mathcal{E})$. So we can pick any node in the remaining graph. $\square$

THEOREM 10. *Every n -node graph and set of $|P| = p$ demand pairs has a distance preserver on $O(n^{2/3}p^{2/3} + np^{1/3})$ edges.*

PROOF. Let $\mathcal{E}$ be an integer parameter that we will choose later. Let π' be any tiebreaking scheme and let $H' = \pi'(P)$ be the associated distance preserver. Repeat the following process until the average degree in H' is $\leq \mathcal{E}$.

Using Lemma 9, we can find a node v in H' and a set S containing $|S| = \Omega(\mathcal{E})$ neighbors of v which all have degree $\Omega(\mathcal{E})$. Hence there is a set $Q \subseteq P$ of $|Q| = \mathcal{E}^2$ demand pairs so that, for every $(x, y) \in Q$, the path $\pi'(x, y)$ contains a node in S . Letting $s_{x,y} \in S$ be a node so that $s_{x,y} \in \pi'(x, y)$, we may split each demand pair (x, y) into two demand pairs $(s_{x,y}, x), (s_{x,y}, y)$. Now, using Lemma 8 (with parameters $d = 2, \sigma \leq \mathcal{E}, p = 2\mathcal{E}^2$) we can build a distance preserver of the demand pairs in Q on

$$O\left(\sqrt{n|Q||S|} + n\right) = O(n^{1/2}\mathcal{E}^{3/2} + n)$$

edges.

We can now remove the (original) demand pairs in Q from P , and repeat. Since we remove $\Theta(\mathcal{E}^2)$ demand pairs from P in each round, we repeat only $O(p/\mathcal{E}^2)$ times in total. Thus, if we union

together the distance preservers computed in each round and also the final leftover preserver H' , the total size is

$$\underbrace{O\left(\frac{p}{\mathcal{E}^2}\right)}_{\text{number of rounds}} \cdot \underbrace{O(n^{1/2}\mathcal{E}^{3/2} + n)}_{\text{size in each round}} + \underbrace{O(n\mathcal{E})}_{\text{size of } H'} = O(pn^{1/2}\mathcal{E}^{-1/2} + np\mathcal{E}^{-2} + n\mathcal{E})$$

edges. Choosing $\mathcal{E} = p^{2/3}n^{-1/3}$ gives

$$|E| = O(n^{2/3}p^{2/3} + n^{5/3}p^{-1/3}),$$

which is $O(n^{2/3}p^{2/3})$ when $p = \Omega(n)$. Choosing $\mathcal{E} = p^{1/3}$ gives

$$|E| = O(np^{1/3} + n^{1/2}p^{5/6}),$$

which is $O(np^{1/3})$ when $p = O(n)$. □

It seems that the $+np^{1/3}$ term arises in this argument for a purely technical reason. The culprit is essentially the $+n$ in Lemma 7, which is amplified over the union bound of the separate preservers built for each set Q in the above proof. Of course the $+n$ in Lemma 7 can't be removed in general, as shown by the counterexample where G is a big long path, which needs $\Omega(n)$ edges for a distance preserver despite having $b = 0$ branching events. However, such graphs generally admit very good distance preservers when many demand pairs are considered. We thus conjecture that with more care this problem can be avoided.

CONJECTURE 11. *Every n -node graph and p demand pairs have a distance preserver on $O(n^{2/3}p^{2/3} + n)$ edges.*

3 ADDITIVE SPANNERS

We now give our additive spanner construction. As discussed previously, our bounds are parameterized on an arbitrary $\varepsilon > 0$. We will use standard $O_\varepsilon(\cdot)$ notation to hide multiplicative factors that depend only on the choice of ε (but not on n), and we will use notation $\widehat{O}(\cdot)$ to hide factors of the form $n^{O(\varepsilon)}$. (We also use $\widehat{\Theta}, \widehat{\Omega}$ notation defined analogously.) We will prove:

THEOREM 12. *Assuming Conjecture 11, for any $\varepsilon > 0$ and $1 \leq \mathcal{E} \leq n^{1/3}$, every n -node graph $G = (V, E)$ has an additive spanner with*

$$+\widehat{O}(n^{3/7}\mathcal{E}^{-9/7})$$

error and $O_\varepsilon(n\mathcal{E})$ edges.

At the end, we will also give versions of this theorem with slightly worse error bounds that hold unconditionally. We let β be a parameter of the construction, which roughly controls the spanner error (in the end we will set $\beta \approx n^{3/7}\mathcal{E}^{-9/7}$). We will use the following notations:

$$\begin{aligned} B(v, r) &:= \{u \in V \mid \text{dist}(u, v) \leq r\}, \text{ and} \\ B_=(v, r) &:= \{u \in V \mid \text{dist}(u, v) = r\}. \end{aligned}$$

We view our approach as an extension of the classic path-buying method from [7] to clusters of larger radius; our construction reduces to something very similar to the one in [7] when the parameters are set such that $\beta = O(1)$. Although familiarity with [7] is not required to read the following proofs, it may be helpful, and we will occasionally pause to explain analogs to this argument in [7] for the reader who is familiar with this prior work.

3.1 Clustering

The construction in [7] begins with a *clustering*, in which the nodes of the graph are either “low-degree” or “high-degree;” they accept all edges incident to low-degree nodes into the spanner, and the high-degree nodes are covered by a small number of clusters of diameter 2 each. In our setting where $\text{poly}(n)$ spanner error is allowed ([7] only allows +6 error), it is natural to attempt an analogous method, using clusters of radius $\text{poly}(n)$ instead of 2. We set up an analogous clustering in the following two lemmas. Our “bottlenecked clusters” are analogous to the low-degree nodes from [7], while the “large clusters” are analogous to the clusters that cover high-degree nodes.

LEMMA 13. *There exists a set of k nodes $\{v_1, \dots, v_k\}$ and k associated integers $\{r_1, \dots, r_k\}$, where each $r_i = \widehat{\Theta}(\beta)$, satisfying the following two properties:*

- (Coverage) *For each $v \in V$, there is $1 \leq i \leq k$ such that $v \in B(v_i, r_i)$.*
- (Non-Overlap) $\sum_{i=1}^k |B(v_i, 2r_i)| = O_\varepsilon(n)$.

(Note the different radii, r_i vs. $2r_i$, in the coverage/non-overlap properties.)

PROOF. First, for every node $v \in V$ we compute a *potential radius* r'_v as follows. This will ultimately be half the radius of the cluster centered at v , if v is chosen as a cluster center. Let c be a parameter and initialize $r'_v \leftarrow \beta$. If

$$c |B(v, r'_v)| \geq |B(v, 4r'_v)|,$$

then finalize r'_v at its current value. Otherwise, set $r'_v \leftarrow 4r'_v$ and repeat. Since $|B(v, r'_v)| \leq n$ always, and this quantity grows by at least $\cdot c$ in each iteration, we repeat $\leq \log_c n$ times. Thus we have

$$r'_v \leq \beta 4^{\log_c n} = \beta n^{\frac{1}{\log_4 c}}.$$

We choose $c = \Theta(4^{1/\varepsilon})$, and so $r'_v \leq \beta n^\varepsilon$. Our next step is to select the nodes that will be cluster centers. Sort the nodes $v \in V$ descendingly by value of r'_v . We repeat the following process until we are out of nodes. Remove the first remaining node v from the list and add it to the list of selected cluster centers. Then, for each remaining node u in the list with $B(u, r'_u) \cap B(v, r'_v) \neq \emptyset$, delete u from the list.


The final radius associated with each selected cluster center u will be $r_u := 2r'_u$. To verify the coverage property: let $v \in V$. If v is chosen as a cluster center we immediately have $\text{dist}(v, v) = 0 \leq 2r'_v$. Otherwise, we discarded v during the construction, so there is a cluster center u considered before v with $B(u, r'_u) \cap B(v, r'_v) \neq \emptyset$. Since $r'_u \geq r'_v$, this implies $\text{dist}(v, u) \leq 2r'_u = r_u$. For the non-overlap property, we compute:

$$\begin{aligned} \sum_{i=1}^k |B(v_i, 2r_i)| &= \sum_{i=1}^k |B(v_i, 4r'_i)| \\ &\leq \sum_{i=1}^k c |B(v_i, r'_i)| \\ &\leq cn && \text{since } \{B(v_i, r'_i)\} \text{ are disjoint} \\ &= O_\varepsilon(n). \end{aligned} \quad \square$$


We can classify balls as in Lemma 13 into two types as follows:

LEMMA 14 (SEE FIGURE 5). *For all nodes v and positive integers r , we have either:*

- (Large) $|B(v, 2r)| = \Omega(r^{4/3} \mathcal{E})$, or
- (Bottlenecked) *there exists $r < r^* \leq 2r$ so that $|B(v, r^*)| \leq |B(v, r^*)|^{1/4} \mathcal{E}^{3/4}$.*



(a) In a “large” ball, every ring around the core contains many nodes, so the cluster contains many nodes overall.



(b) In a “bottlenecked” ball, at least one of the rings around the core (its “bottleneck layer”) contains few nodes. This small layer will be exploited in our construction.

Fig. 5. The two types of clusters used in our construction of additive spanners.

PROOF. Suppose that $B(v, 2r)$ is not bottlenecked. We have $|B(v, r)| \geq 1$ (conservatively), and we also have the relationship

$$\begin{aligned} |B(v, r^*)| &= |B(v, r^* - 1)| + |B_=(v, r^*)| \\ &\geq |B(v, r^* - 1)| + |B(v, r^* - 1)|^{1/4} \mathcal{E}^{3/4} \end{aligned}$$

for all $r < r^* \leq 2r$ (since $|B_=(v, r^*)| \geq |B(v, r^* - 1)|^{1/4} \mathcal{E}^{3/4}$). Phrased another way, if we let $x_j := |B(v, r + j)|$, then we have a recurrence with initial condition $x_0 \geq 1$, and recurrence relation

$$x_j \geq x_{j-1} + x_{j-1}^{1/4} \mathcal{E}^{3/4}.$$

The general solution to a recurrence of this form is $x_j = \Omega(j^{4/3} \mathcal{E})$, and plugging in $j = r$ thus gives $|B(v, 2r)| = \Omega(r^{4/3} \mathcal{E})$, implying that the cluster is large. $\square$

In the following, we let $\{v_1, \dots, v_k\}, \{r_1, \dots, r_k\}$ be selected as in Lemma 13. We classify each pair v_i, r_i into *large* or *bottlenecked* as in Lemma 14 (some balls can satisfy both criteria, in which case we can classify as either large or bottlenecked, it doesn't matter). If it is large, then we refer to the ball $B(v_i, 2r_i)$ as a *large cluster*, and $B(v_i, r_i)$ is its *core*. If it is bottlenecked, then we refer to the ball $B(v_i, r_i^*)$ as a *bottlenecked cluster*, and $B_=(v_i, r_i^*)$ is its *bottleneck layer*. A node $x \in B(v_i, r_i^* - 1)$ is *bottlenecked* by the cluster $B(v_i, r_i^*)$. Note that a node on the bottleneck layer itself is not generally bottlenecked (although it can be bottlenecked if it is also contained in a different cluster). In either case, v_i is the *center node* of the cluster.

3.2 Initialization Phase

Like many spanner constructions, ours can be divided into an *initialization* and *completion* phase. The completion phase is the part that enforces the final spanner inequality. The initialization phase adds some edges to the graph beforehand, whose purpose is purely to assist in the argument that the completion phase doesn't add too many edges to the spanner.

In this part we will state the initialization phase (only) and provide some relevant analysis before moving on to the completion phase. First off, our initialization phase actually includes an additive spanner on the inside. We will use the following spanners from prior work internally:

THEOREM 15 ([11]). *For any $\mathcal{E} \geq 1$, every n -node graph has a spanner on $O(n\mathcal{E})$ edges with additive error $+O(n^{1/2}\mathcal{E}^{-1/2})$.*

The spanners in Theorem 15 are state-of-the-art only in some range of parameters, when $\mathcal{E}$ is not too large. If one uses other spanner constructions instead and rebalances parameters, from [11] or [15], one can optimize the quality of our spanner for certain larger values of $\mathcal{E}$. We will not do so, since using piecewise internal spanners introduces considerable additional complexity in the analysis, and the spanners from Theorem 15 give the best possible results when $\mathcal{E} = O(1)$ which we consider to be the most interesting setting. We now state our initialization phase:

- (1) **(Initialization Phase)** The spanner H is initially empty. Then, we compute a clustering as described in the previous section, and:
- (a) For any cluster C of either type, add a shortest path tree rooted at its center node v to all nodes in C .
 - (b) For each bottlenecked cluster $B(v, r^*)$, construct a distance preserver with demand pairs $B_=(v, r^*)^2$ in the subgraph induced on $B(v, r^*)$, and add its edges to H .
 - (c) For each large cluster $B(v, 2r)$, construct a spanner using Theorem 15 in the subgraph induced on $B(v, 2r)$, and add its edges to H .
- (2) *(Then perform the Completion Phase, described later)*

The distance preservers on bottlenecked clusters are analogous to the step in [7] in which we add all edges to low-degree nodes. This roughly lets us “ignore” bottlenecked nodes, as we ignore low-degree nodes in [7], in a sense we will make precise shortly. The spanners on large clusters do not have an analogy in [7]; this is an optimization that is only useful in the setting of large-radius clusters. The shortest path trees enforce that, if two nodes x, y are in the same cluster C with center node v , then we have

$$\text{dist}_H(x, y) \leq \text{dist}_H(x, v) + \text{dist}_H(v, y) = \widehat{O}(\beta);$$

this inequality will occasionally be useful. Let us next count the number of edges added to the spanner in the initialization phase.

LEMMA 16 (INITIALIZATION EDGE BOUND). *Assuming Conjecture 11, only $O_\epsilon(n\mathcal{E})$ edges are added to H during the initialization phase.*

PROOF. For each bottlenecked cluster $B(v, r^*)$, the associated distance preserver has

$$|B_=(v, r^*)^2| = |B_=(v, r^*)|^2 \leq |B(v, r^*)|^{1/2} \mathcal{E}^{3/2}$$

demand pairs. Using Conjecture 11, the number of edges in the distance preserver is

$$O\left(|B(v, r^*)|^{2/3} \left(|B(v, r^*)|^{1/2} \mathcal{E}^{3/2}\right)^{2/3} + |B(v, r^*)|\right) = O(|B(v, r^*)|\mathcal{E}).$$

For each large cluster $|B(v, 2r)|$, since we apply Theorem 15 with parameter $\mathcal{E}$, we add $O(\mathcal{E}|B(v, 2r)|)$ edges to the cluster. We also add a shortest path tree to both types of clusters, but this does not change either estimate. By a union bound, the total number of edges added over




Fig. 6. If a shortest path $\pi(s, t)$ contains a missing edge inside a bottlenecked cluster, then due to the distance preservers added in the initialization phase, we can reroute a subpath of $\pi(s, t)$ to remove this missing edge. This observation is applied repeatedly to prove Lemma 17.

all clusters is thus

$$\begin{aligned} \sum_{C \text{ is a cluster}} O(\varepsilon|C|) &= O\left(\varepsilon \cdot \sum_{C \text{ is a cluster}} |C|\right) \\ &= O_\varepsilon(n\varepsilon), \end{aligned}$$

where this last equality follows from the non-overlap property of Lemma 13. $\square$

Before we state our completion phase, we need a technical lemma describing the structure of shortest paths in the post-initialization spanner. We think of Lemma 18 as a more technically-laden version of the trivial fact used in [7] that, after we add all edges to low-degree nodes, no missing edges in a shortest path are incident on low-degree nodes; its role in our overall proof is essentially as a tool that lets us focus on large clusters and ignore bottlenecked clusters. In the following, a *missing edge* is an edge in the original graph G that is not contained in the current spanner (here, meaning after the initialization phase is complete).

LEMMA 17 (SEE FIGURE 6). *For any two non-bottlenecked nodes s, t , there is a shortest path $\pi(s, t)$ in G such that no missing edge is contained in a bottlenecked cluster.*

PROOF. Let $\pi(s, t)$ be an arbitrary shortest path, let $B(v, r^*)$ be a bottlenecked cluster, and suppose $\pi(s, t)$ has a missing edge $e \in B(v, r^*)$. Let $\pi(s, t)[x \rightsquigarrow y]$ be the longest contiguous subpath contained in $B(v, r^*)$ that includes the missing edge e . Since s, t are not bottlenecked, this subpath must terminate in the bottleneck layer; that is, $x, y \in B_=(v, r^*)$. Therefore, we added a shortest $x \rightsquigarrow y$ path contained in $B(v, r^*)$ as part of a distance preserver in the initialization phase. We can thus reroute the subpath $\pi(s, t)[x \rightsquigarrow y]$ to coincide with this previously-added shortest $x \rightsquigarrow y$ path. This changes at least one missing edge e from $\pi(s, t)$, and it does not introduce any new missing edges to $\pi(s, t)$. Thus we can repeat this process until the lemma holds. $\square$

LEMMA 18. *For any shortest path $\pi(s, t)$ as in Lemma 17, there is a collection of edge-disjoint subpaths Q of $\pi(s, t)$ such that:*

- Every missing edge in $\pi(s, t)$ is contained in a subpath $q \in Q$,
- Every subpath in Q except possibly for the last one has length $\widehat{\Omega}(\beta)$, and
- Each subpath $q \in Q$ may be “charged” to a large cluster C such that $q \subseteq C$, and each large cluster is charged for at most one subpath in Q .

PROOF. Initially $Q \leftarrow \emptyset$. Repeat the following process until no longer possible: let (x, y) be the first missing edge along $\pi(s, t)$ that is not currently contained in any subpath in Q . By the coverage property of Lemma 13, we have $x \in B(v, r)$ for some center node v with associated radius r . Let q be the subpath of $\pi(s, t)$ that starts with the edge (x, y) and extends as far as possible under the constraint $q \subseteq B(v, 2r)$. Add q to Q , then repeat until all missing edges along $\pi(s, t)$ are covered by subpaths in Q .

It is immediate from the construction that our subpaths are edge-disjoint and that they cover the missing edges of $\pi(s, t)$. We next verify the subpath length property: a selected subpath q starts at a node $x \in B(v, r)$, and it ends when it cannot be extended further, either because it reaches the end of $\pi(s, t)$ (in which case it is the last subpath) or because it reaches $B_=(v, 2r)$ (in which case it has length at least $r = \widehat{\Omega}(\beta)$).

Next, we verify the subpath charging property. For each missing edge (x, y) , we have $x \in B(v, r)$. Notice that v cannot be the center of a *bottlenecked* cluster, since otherwise the missing edge (x, y) would be contained in this bottlenecked cluster, but by Lemma 17 no missing edge in $\pi(s, t)$ is in a bottlenecked cluster. Thus v is the center node of a *large* cluster, and we can charge the associated subpath q to this large cluster, since we have $q \subseteq B(v, 2r)$.

Finally, we verify that each large cluster is only charged for one subpath in Q . Suppose for contradiction that a large cluster $B(v, 2r)$ is charged for two different subpaths q_1, q_2 , due to two different missing edges $(x_1, y_1), (x_2, y_2)$ with $x_1, x_2 \in B(v, r)$. That means there is a node $z \in \pi(s, t)[x_1 \rightsquigarrow y_2]$ with $z \notin B(v, 2r)$, since otherwise the subpath q_1 would extend until it covers (x_2, y_2) as well. So we have

$$\text{dist}_G(x_1, x_2) = \text{dist}_G(x_1, z) + \text{dist}_G(z, x_2) > 2r,$$

since $x_1, x_2 \in B(v, r)$ but $z \notin B(v, 2r)$. However, by the triangle inequality, we also have

$$\text{dist}_G(x_1, x_2) \leq \text{dist}_G(x_1, v) + \text{dist}_G(v, x_2) \leq 2r,$$

completing the contradiction. □

3.3 Completion Phase

We now state the completion phase of the algorithm:

(1) (Perform the Initialization Phase, described previously)

(Completion Phase:)

(2) Let c be a sufficiently large positive constant. For each pair of non-bottlenecked nodes s, t , if currently

$$\text{dist}_H(s, t) > \text{dist}_G(s, t) + \beta n^{c\epsilon} :$$

(a) Let $\pi(s, t), Q$ be a shortest path and collection of subpaths as in Lemma 18.

(b) For each subpath $q \in Q$, letting x, y be its endpoints, add (x, y) as a **weighted** edge to the spanner H , with weight $\text{dist}_G(x, y)$. Charge (x, y) to the large cluster to which we charged q .

(3) For each large cluster C , let P_C be the set of weighted edges charged to C . Replace the edges P_C with a distance preserver on the subgraph induced on C , with demand pairs P_C .

(4) Return the spanner H .

One can imagine our completion phase as following the greedy completion strategy from [7] or its update in [23]: while there is a pair of non-bottlenecked nodes s, t that violate the spanner inequality, add a shortest path $\pi(s, t)$ to the spanner to correct this. But there are two key differences. First, while [7] uses a “path-buying” argument to directly bound the number of edges added to each cluster, we use path-buying to bound the *number of subpaths* that get “charged” to each large cluster. We will then apply distance preserver bounds at the end to convert a bound on number of subpaths per cluster to a bound on number of edges per cluster.

Second, in order to apply these distance preserver bounds, we need to be very careful about our choice of shortest subpaths through each large cluster. That motivates the use of our intermediate weighted edges: these act as “placeholders” for the specific shortest path through a large cluster that will ultimately be selected; they record that we have committed to enforcing $\text{dist}_H(x, y) = \text{dist}_G(x, y)$ in the final spanner, even though we haven’t yet picked the shortest path that will achieve this. We resolve our weighted edges into shortest paths at the end: this is a way to delay our choice of shortest $x \rightsquigarrow y$ path until we know the *entire* set of shortest paths in a large cluster that will need to be selected, which lets us use a distance preserver construction. If we tried to simplify our algorithm by avoiding weighted edges, instead having it commit to an entire shortest path $\pi(s, t)$ right when the node pair s, t is considered, then we would only be able to use “online” distance preserver bounds, where the shortest path for each demand pair is selected before the next demand pair arrives.

We will now begin to formally prove properties of our construction. We begin with correctness of the output spanner:

LEMMA 19 (SPANNER CORRECTNESS - SEE FIGURE 7). *For all nodes s, t , we have $\text{dist}_H(s, t) \leq \text{dist}_G(s, t) + \widehat{O}(\beta)$.*

PROOF. Let $\pi(s, t)$ be a shortest $s \rightsquigarrow t$ path as in Lemma 17. Let s' be the last node on $\pi(s, t)$ such that there exists a bottlenecked cluster C with $s, s' \in C$, or if s is not contained in any bottlenecked cluster then let $s' := s$. In either case, since we added a shortest path tree rooted at the center node of C , the distance between s, s' is at most twice the radius of C ; that is,

$$\text{dist}_H(s, s') = \widehat{O}(\beta).$$

Similarly, we let t' be the first node on $\pi(s, t)$ such that there exists a bottlenecked cluster C with $t, t' \in C$, or if t is not contained in any bottlenecked cluster then let $t' := t$. For the same reason we have

$$\text{dist}_H(t, t') = \widehat{O}(\beta).$$

Next, let s'', t'' be the first, last non-bottlenecked nodes along the subpath $\pi(s, t)[s' \rightsquigarrow t']$. Notice that all edges on the $s' \rightsquigarrow s''$ and $t' \rightsquigarrow t''$ subpaths are contained in a bottlenecked cluster, but by choice of s', t' these bottlenecked clusters do not also contain s, t . Thus, by Lemma 17, every edge along the $s' \rightsquigarrow s''$ and $t' \rightsquigarrow t''$ subpaths is already contained in the spanner, and so we have

$$\text{dist}_H(s', s'') = \text{dist}_G(s', s'') \text{ and } \text{dist}_H(t', t'') = \text{dist}_G(t', t'').$$

Finally, in the completion phase, since s'', t'' are non-bottlenecked, we ensure that

$$\text{dist}_H(s'', t'') \leq \text{dist}_G(s'', t'') + \widehat{O}(\beta).$$




Fig. 7. The organization of the subpaths analyzed in our proof of spanner correctness (Lemma 19).

Putting it all together, we have:

$$\begin{aligned}
 \text{dist}_H(s, t) &\leq \text{dist}_H(s, s') + \text{dist}_H(s', s'') + \text{dist}_H(s'', t'') + \text{dist}_H(t'', t') + \text{dist}_H(t', t) \\
 &= \text{dist}_H(s', s'') + \text{dist}_H(s'', t'') + \text{dist}_H(t'', t') + \widehat{O}(\beta) \\
 &= \text{dist}_G(s', s'') + \text{dist}_H(s'', t'') + \text{dist}_G(t'', t') + \widehat{O}(\beta) \\
 &= \text{dist}_G(s', s'') + \text{dist}_G(s'', t'') + \text{dist}_G(t'', t') + \widehat{O}(\beta) \\
 &= \text{dist}_G(s', t') + \widehat{O}(\beta) \\
 &\leq \text{dist}_G(s, t) + \widehat{O}(\beta).
 \end{aligned}$$

□

Our goal is now to control the number of edges added in the completion phase. As discussed previously, we generally follow the path-buying method from [7], with one new optimization for our setting. In [7], the strategy is roughly to argue that each time a shortest path $\pi(s, t)$ is added, for each cluster C that is hit by $\pi(s, t)$, we improve the spanner distance from C to either the first or the last cluster hit by $\pi(s, t)$. This is used to control the number of times any given cluster is hit by an added shortest path. Our strategy is similar, but we can gain a bit by considering not just the first and last large clusters hit by $\pi(s, t)$, but rather the *first few* and *last few* large clusters hit by $\pi(s, t)$. This finally reveals the purpose of the spanners we added to large clusters in the initialization phase: essentially, they allow us to traverse a few extra clusters at the start and end and stay within our error budget.

Let us next formalize what we mean by “first few” and “last few” large clusters. For a shortest path $\pi(s, t)$ with associated subpath list Q selected in the completion phase, we define the *prefix* $Q_{\text{pre}} \subseteq Q$ to be the minimal prefix of subpaths in Q that satisfies

$$\sum_{\substack{C \text{ large cluster} \\ \text{some } q \in Q_{\text{pre}} \text{ charged to } C}} |C| \geq \beta^{5/3} \mathcal{E}.$$

To avoid confusion around the word “prefix:” note that Q_{pre} is **not** obtained by selecting some prefix of each subpath in Q ; rather, each subpath in Q is entirely in or entirely out of Q_{pre} . We are taking the first few subpaths in Q used by the path being analyzed, until they have sufficiently large total size. Similarly, we define the *suffix* Q_{suf} as the minimal suffix of subpaths in Q that satisfies

$$\sum_{\substack{C \text{ large cluster} \\ \text{some } q \in Q_{\text{suf}} \text{ charged to } C}} |C| \geq \beta^{5/3} \mathcal{E},$$

and we again caution that this is **not** obtained by taking a suffix of each path in Q , but rather, taking all of the last few subpaths in the ordering of Q .

If no prefix/suffix of Q achieves this inequality, then we can define $Q_{\text{pre}} = Q_{\text{suf}} = Q$, but it turns out we will not encounter this situation in our algorithm so it won't be too important (Lemma 21 to follow implies that, if the prefix and suffix overlap, then the node pair s, t is already well-spanned and so it won't be selected in the completion phase). It is technically possible for a node to be contained in several large clusters charged by prefix/suffix subpaths. However, we can bound the frequency with which this happens, leading to a bound on the number of *distinct* nodes in these prefix/suffixes. Let A_{pre} be the set of nodes contained in any large cluster C for which some $q \in Q_{\text{pre}}$ is charged to C , and similarly let A_{suf} be the set of nodes contained in any large cluster C for which some $q \in Q_{\text{suf}}$ is charged to C .

LEMMA 20 (PREFIX/SUFFIX SIZE). $|A_{\text{pre}}|, |A_{\text{suf}}| = \widehat{\Omega}(\beta^{5/3}\mathcal{E})$.

PROOF. Let v be a node in A_{pre} or A_{suf} . Let x, y be the first, last nodes on $\pi(s, t)$ that are respectively contained in large clusters C_x, C_y with $v \in C_x, v \in C_y$. Every large cluster has radius $\widehat{O}(\beta)$, and so we have

$$\text{dist}_G(x, v), \text{dist}_G(y, v) = \widehat{O}(\beta).$$

So by the triangle inequality, we have $\text{dist}_G(x, y) = \widehat{O}(\beta)$. From Lemma 18, every subpath in Q has length $\widehat{O}(\beta)$, except possibly for the last one. Hence there can only be $\widehat{O}(1)$ total subpaths in Q that contain nodes between x and y , and so we can charge only $\widehat{O}(1)$ different large clusters that contain v . Thus v is counted only $\widehat{O}(1)$ times in the sum used to define the prefix and suffix, leading to the claimed bound. $\square$

When we say the prefix/suffix can be traversed within our error budget, we mean the following lemma:

LEMMA 21 (PREFIX/SUFFIX DISTANCE ERROR). *For any $x \in A_{\text{pre}}, y \in A_{\text{suf}}$, we have*

$$\text{dist}_H(s, x) \leq \text{dist}_G(s, x) + \widehat{O}(\beta) \text{ and } \text{dist}_H(y, t) \leq \text{dist}_G(y, t) + \widehat{O}(\beta).$$

PROOF. We will prove the inequality for the prefix here ($x \in A_{\text{pre}}$); the inequality for the suffix is essentially identical. Let C^* be the large cluster charged by a subpath $q^* \in Q_{\text{pre}}$ with $x \in C^*$, and let $x' \in \pi(s, t)$ be the first node on the subpath q^* . Let $\{q_1, \dots, q_k, q^*\}$ be the leading subpaths of Q_{pre} up to q^* , and let $\{C_1, \dots, C_k, C^*\}$ be the respective large clusters charged for these subpaths. Let x_i, y_i be the endpoints of each subpath q_i . We then have

$$\begin{aligned} \sum_{i=1}^{k-1} \text{dist}_H(x_i, y_i) &\leq \sum_{i=1}^{k-1} \text{dist}_G(x_i, y_i) + \widetilde{O}\left(\frac{|C_i|^{1/2}}{\mathcal{E}^{1/2}}\right) && \text{Spanners from initialization} \\ &\leq \sum_{i=1}^{k-1} \text{dist}_G(x_i, y_i) + \widetilde{O}\left(\frac{|C_i|}{\beta^{2/3}\mathcal{E}}\right) && \text{since } |C_i| = \Omega(\beta^{4/3}\mathcal{E}) \text{ by Lemma 14} \\ &\leq \left(\sum_{i=1}^{k-1} \text{dist}_G(x_i, y_i)\right) + \widetilde{O}\left(\frac{\beta^{5/3}\mathcal{E}}{\beta^{2/3}\mathcal{E}}\right) && \text{Def of prefix} \\ &\leq \left(\sum_{i=1}^{k-1} \text{dist}_G(x_i, y_i)\right) + \widetilde{O}(\beta). \end{aligned}$$

From Lemma 18, every edge along $\pi(s, t)[s \rightsquigarrow x]$ except those contained in $x_i \rightsquigarrow y_i$ subpaths is already contained in the spanner. So we have

$$\begin{aligned} \text{dist}_H(s, x') &\leq \text{dist}_H(s, x_1) + \left(\sum_{i=1}^k \text{dist}_H(x_i, y_i) \right) + \left(\sum_{i=1}^{k-1} \text{dist}_H(y_i, x_{i+1}) \right) + \text{dist}_H(y_k, x') \\ &= \text{dist}_G(s, x_1) + \left(\sum_{i=1}^k \text{dist}_G(x_i, y_i) + \widetilde{O}(\beta) \right) + \left(\sum_{i=1}^{k-1} \text{dist}_G(y_i, x_{i+1}) \right) + \text{dist}_G(y_k, x') \\ &= \text{dist}_G(s, x') + \widetilde{O}(\beta). \end{aligned}$$

Finally, we can complete the proof using the triangle inequality:

$$\begin{aligned} \text{dist}_H(s, x) &\leq \text{dist}_H(s, x') + \text{dist}_H(x', x) \\ &\leq (\text{dist}_G(s, x') + \widetilde{O}(\beta)) + \text{dist}_H(x', x) \\ &\leq \text{dist}_G(s, x) + \text{dist}_G(x, x') + \text{dist}_H(x', x) + \widetilde{O}(\beta) \\ &\leq \text{dist}_G(s, x) + \widehat{O}(\beta) \quad x, x' \text{ in same large cluster. } \quad \square \end{aligned}$$

In our next two lemmas, we will focus on a particular large cluster L that is charged by at least one subpath $q_L \in Q$. We say that L is *cluster-preserved* with a node x if there exists a node $\ell \in L$ and a node x' such that x, x' are in the same large cluster and $\text{dist}_H(\ell, x') = \text{dist}_G(\ell, x')$. We use cluster-preservation as a sort of potential function: we will argue that, when we charge a new weighted edge to L , we also cause L to become cluster-preserved with many new nodes. This limits the total number of weighted edges that may be charged to L over the entire completion phase, since L can only become newly cluster-preserved with n nodes in total.

LEMMA 22. *Before the node pair s, t is considered, either L is not cluster-preserved with any node in A_{pre} , or L is not cluster-preserved with any node in A_{suf} .*

PROOF. Let us first handle the case where the subpath q_L charging L is in neither the prefix nor the suffix. Suppose for contradiction that there are nodes $x \in A_{\text{pre}}, y \in A_{\text{suf}}$ and $\ell_x, \ell_y \in L$ with

$$\text{dist}_H(x, \ell_x) = \text{dist}_G(x, \ell_x) \text{ and } \text{dist}_H(y, \ell_y) = \text{dist}_G(y, \ell_y).$$

Let x', y' be nodes on $\pi(s, t)$ that share a large cluster with x, y respectively, and let $\ell \in q_L$. We then have

$$\begin{aligned} \text{dist}_H(s, t) &\leq \text{dist}_H(s, x) + \text{dist}_H(x, \ell_x) + \text{dist}_H(\ell_x, \ell_y) + \text{dist}_H(\ell_y, y) + \text{dist}_H(y, t) \\ &\leq \text{dist}_H(s, x) + \text{dist}_G(x, \ell_x) + \text{dist}_G(\ell_y, y) + \text{dist}_H(y, t) + \widehat{O}(\beta) \quad \ell_x, \ell_y \text{ in same cluster} \\ &\leq \text{dist}_G(s, x) + \text{dist}_G(x, \ell_x) + \text{dist}_G(\ell_y, y) + \text{dist}_G(y, t) + \widehat{O}(\beta) \quad \text{Lemma 21} \\ &\leq \text{dist}_G(s, x') + \text{dist}_G(x', \ell) + \text{dist}_G(\ell, y') + \text{dist}_G(y', t) + \widehat{O}(\beta) \end{aligned}$$

where the last inequality uses the triangle inequality together with the fact that the nodes (x, x') , (ℓ_x, ℓ, ℓ_y) , (y, y') each lie in same large cluster, which has radius $\widehat{O}(\beta)$. Since x', ℓ, y' all lie along $\pi(s, t)$, this implies

$$\text{dist}_H(s, t) \leq \text{dist}_G(s, t) + \widehat{O}(\beta).$$

But, by choice of large enough constant c , this means we wouldn't select the node pair s, t in the completion phase and add any corresponding weighted edges, since the distance inequality for this node pair would be already satisfied. This contradicts the assumption that s, t were selected in the completion phase.

The case where q_L is in the prefix follows an essentially identical analysis, which we will not repeat: in this case we do not need nodes x, x' , and we specifically argue that L is not cluster-preserved with any node in A_{suf} (or else the above inequality holds, again contradicting that s, t were selected). In the case where q_L is in the suffix, we instead do not need nodes y, y' , and we specifically argue that L is not cluster-preserved with any node in A_{pre} . $\square$

LEMMA 23. *After the node pair s, t is considered, L is cluster-preserved with every node in $A_{\text{pre}} \cup A_{\text{suf}}$.*

PROOF. Let $\pi(s, t)$ be the shortest path in G , and let $\pi^*(s, t)$ be the shortest path in the spanner H obtained from $\pi(s, t)$ by replacing each $x \rightsquigarrow y$ subpath in Q with the weighted edge (x, y) added to the spanner. (Notice that, from Lemma 18, every edge in $\pi(s, t)$ not contained in a subpath in Q is already in H from the initialization phase, so this does indeed describe a path contained in H .)

Let $z \in A_{\text{pre}} \cup A_{\text{suf}}$, and let C be the large cluster containing z that is charged by a prefix or suffix subpath. By construction there is a weighted edge $(x, y) \in \pi^*(s, t)$ with $x, y \in C$. Similarly, there is a weighted edge $(x', y') \in \pi^*(s, t)$ with $x', y' \in L$. Thus we have $\text{dist}_H(x, x') = \text{dist}_G(x, x')$, due to the shortest path $\pi^*(s, t)$. So L is cluster-preserved with z after $\pi^*(s, t)$ is added. $\square$

We now put the pieces together:

LEMMA 24 (SEE FIGURE 8). *Assuming Conjecture 11, the number of edges added to H during the completion phase is*

$$O_\epsilon(n) + \widehat{O}\left(\frac{n^{5/3}}{\beta^{14/9}\mathcal{E}}\right).$$

PROOF. Each time we charge a weighted edge to a large cluster L , by Lemmas 22 and 23, we also newly cluster-preserve it with all nodes in A_{pre} or with all nodes in A_{suf} . Using Lemma 20, we have

$$|A_{\text{suf}}|, |A_{\text{pre}}| = \widehat{\Omega}(\beta^{5/3}\mathcal{E}).$$

Thus, the total number of weighted edges charged to L is only

$$\widehat{O}\left(\frac{n}{\beta^{5/3}\mathcal{E}}\right).$$

Applying Conjecture 11, the number of edges added to L as part of the final distance preserver is

$$\begin{aligned} & O\left(|L| + |L|^{2/3} \widehat{O}\left(\frac{n}{\beta^{5/3}\mathcal{E}}\right)^{2/3}\right) \\ &= O\left(|L| + |L| (\beta^{4/3}\mathcal{E})^{-1/3} \widehat{O}\left(\frac{n^{2/3}}{\beta^{10/9}\mathcal{E}^{2/3}}\right)\right) \quad \text{since } |L| = \Omega(\beta^{4/3}\mathcal{E}) \text{ by Lemma 14} \\ &= O(|L|) + \widehat{O}\left(|L| \cdot \frac{n^{2/3}}{\beta^{14/9}\mathcal{E}}\right). \end{aligned}$$

Hence the total number of edges added in the completion phase is

$$\begin{aligned} & \sum_{L \text{ large cluster}} O(|L|) + \widehat{O}\left(|L| \cdot \frac{n^{2/3}}{\beta^{14/9}\mathcal{E}}\right) \\ &= O_\epsilon(n) + \widehat{O}\left(\frac{n^{5/3}}{\beta^{14/9}\mathcal{E}}\right) \quad \text{Non-overlap property of Lemma 13.} \quad \square \end{aligned}$$




Fig. 8. The “path-buying” part of our argument argues that each time we charge a new weighted edge to a large cluster L , we also cluster-preserve L with many new nodes (either A_{pre} or A_{suf}) for the first time, thus limiting the total number of weighted edges charged to L .

We now finish the proof with a parameter balance. Setting

$$\beta = \widehat{\Theta}(n^{3/7} \mathcal{E}^{-9/7})$$

means the number of edges added in the completion phase is

$$\begin{aligned} & O_{\varepsilon}(n) + O\left(\frac{n^{5/3}}{(n^{3/7} \mathcal{E}^{-9/7})^{14/9} \mathcal{E}}\right) \\ &= O_{\varepsilon}\left(n + \frac{n^{5/3}}{n^{2/3} \mathcal{E}^{-2} \mathcal{E}}\right) \\ &= O_{\varepsilon}(n \mathcal{E}), \end{aligned}$$

which completes the proof of Theorem 12.

3.4 Unconditional Extensions

Theorem 12 is phrased conditionally on Conjecture 11. Here we discuss unconditional extensions of our result, that rely only on known distance preserver bounds. In order to work generally with all possible distance preserver bounds at once, in the following we let a, b be absolute constants such that for any n -node graph and set of p demand pairs, there is a distance preserver on $O(n + n^a p^b)$ edges. Thus, [12] shows that we may take $a = 1/2, b = 1$ or $a = 1, b = 1/2$, and under Conjecture 11 we may take $a = b = 2/3$. We then change the following parts of the preceding argument:

- (1) Distance preserver bounds influence Lemma 14, where we classify clusters into large or bottlenecked. The specific choice of parameters used for bottlenecked clusters comes from our need to argue in Lemma 16 that each bottlenecked cluster C costs only $O(|C|\mathcal{E})$ edges for its distance preserver in the initialization. So, we need to redefine bottlenecked clusters in such a way that this still holds. Specifically, we say that a cluster is bottlenecked if there exists $r < r^* \leq 2r$ so that

$$|B_{\leq}(v, r^*)| \leq |B(v, r^*)|^{\frac{1-a}{2b}} \mathcal{E}^{\frac{1}{2b}}.$$

That way, the cost of the distance preserver for a bottlenecked cluster is

$$\begin{aligned}
 & O \left(|B(v, r^*)| + |B(v, r^*)|^a |B_=(v, r^*)|^{2b} \right) \\
 &= O \left(|B(v, r^*)| + |B(v, r^*)|^a \left(|B(v, r^*)|^{\frac{1-a}{2b}} \mathcal{E}^{\frac{1}{2b}} \right)^{2b} \right) \\
 &= O \left(|B(v, r^*)| + |B(v, r^*)|^a |B(v, r^*)|^{1-a} \mathcal{E} \right) \\
 &= O \left(|B(v, r^*)| \mathcal{E} \right),
 \end{aligned}$$

as desired.

- (2) This redefinition of bottlenecked clusters influences the associated size bound for large clusters. Revisiting the proof of Lemma 14, if $B(v, 2r)$ is not bottlenecked, then we can again describe its size using a recurrence where $x_j := |B(v, r + j)|$. The initial condition is $x_0 \geq 1$, and the recurrence relation is

$$x_j \geq x_{j-1} + x_{j-1}^{\frac{1-a}{2b}} \mathcal{E}^{\frac{1}{2b}}.$$

This recurrence solves asymptotically to

$$x_r = \Omega \left(r^{\frac{2b}{2b+a-1}} \mathcal{E}^{\frac{1}{2b+a-1}} \right).$$

So this is our new size lower bound for large clusters.

- (3) The definition of the prefix and suffix $Q_{\text{pre}}, Q_{\text{suf}}$ use a particular bound on the sum of cluster sizes. This bound is selected to push through Lemma 21, which lets us reach prefix and suffix nodes within small error. This is a function of both the spanner quality in Theorem 15 and the sizes of large clusters. Our new definition of the prefix $Q_{\text{pre}} \subseteq Q$ is the minimal prefix of subpaths satisfying

$$\sum_{\substack{C \text{ large cluster} \\ \text{some } q \in Q_{\text{pre}} \text{ charged to } C}} |C| \geq \beta^{\frac{3b+a-1}{2b+a-1}} \mathcal{E}^{\frac{2b+a}{4b+2a-2}}$$

and the suffix $Q_{\text{suf}} \subseteq Q$ is defined similarly. To re-prove Lemma 21 under this new definition of the prefix and suffix, we proceed through the proof exactly as before, and recompute the following intermediate step:

$$\begin{aligned}
 \sum_{i=1}^{k-1} \text{dist}_H(x_i, y_i) &\leq \sum_{i=1}^{k-1} \text{dist}_G(x_i, y_i) + \tilde{O} \left(|C_i|^{1/2} \mathcal{E}^{-1/2} \right) && \text{Spanners from initialization} \\
 &\leq \sum_{i=1}^{k-1} \text{dist}_G(x_i, y_i) + \tilde{O} \left(|C_i| \left(\beta^{\frac{2b}{2b+a-1}} \mathcal{E}^{\frac{1}{2b+a-1}} \right)^{-1/2} \mathcal{E}^{-1/2} \right) && \text{Large cl. size} \\
 &\leq \sum_{i=1}^{k-1} \text{dist}_G(x_i, y_i) + \tilde{O} \left(|C_i| \cdot \beta^{\frac{-b}{2b+a-1}} \mathcal{E}^{\frac{-2b-a}{4b+2a-2}} \right) \\
 &\leq \left(\sum_{i=1}^{k-1} \text{dist}_G(x_i, y_i) \right) + \tilde{O} \left(\left(\beta^{\frac{3b+a-1}{2b+a-1}} \mathcal{E}^{\frac{2b+a}{4b+2a-2}} \right) \beta^{\frac{-b}{2b+a-1}} \mathcal{E}^{\frac{-2b-a}{4b+2a-2}} \right) && \text{Def of prfx} \\
 &\leq \left(\sum_{i=1}^{k-1} \text{dist}_G(x_i, y_i) \right) + \tilde{O}(\beta).
 \end{aligned}$$

From there, the rest of the proof is the same as before.

- (4) Lemma 24 uses distance preserver bounds to convert our limit on the number of weighted edges charged to each large cluster to a bound on the number of edges in its associated distance preserver. Specifically, we get

$$|A_{\text{suf}}|, |A_{\text{pre}}| = \widehat{\Omega} \left(\beta^{\frac{3b+a-1}{2b+a-1}} \mathcal{E}^{\frac{2b+a}{4b+2a-2}} \right),$$

and so the number of weighted edges charged to any given large cluster L is

$$\widehat{O} \left(\frac{n}{\beta^{\frac{3b+a-1}{2b+a-1}} \mathcal{E}^{\frac{2b+a}{4b+2a-2}}} \right).$$

Applying our distance preserver bound of $O(n + n^a p^b)$, the number of edges added to L as part of the final distance preserver is

$$\begin{aligned} & O \left(|L| + |L|^a \cdot \widehat{O} \left(\frac{n}{\beta^{\frac{3b+a-1}{2b+a-1}} \mathcal{E}^{\frac{2b+a}{4b+2a-2}}} \right)^b \right) \\ &= O \left(|L| + |L| \cdot \left(\beta^{\frac{2b}{2b+a-1}} \mathcal{E}^{\frac{1}{2b+a-1}} \right)^{a-1} \cdot \widehat{O} \left(\frac{n}{\beta^{\frac{3b+a-1}{2b+a-1}} \mathcal{E}^{\frac{2b+a}{4b+2a-2}}} \right)^b \right) \quad \text{Large cl. size} \\ &= O(|L|) + \widehat{O} \left(|L| \cdot \left(\beta^{\frac{2ab-2b}{2b+a-1}} \mathcal{E}^{\frac{a-1}{2b+a-1}} \right) \cdot \left(\frac{n^b}{\beta^{\frac{3b^2+ab-b}{2b+a-1}} \mathcal{E}^{\frac{2b^2+ab}{4b+2a-2}}} \right) \right) \\ &= O(|L|) + \widehat{O} \left(|L| \cdot n^b \cdot \beta^{\frac{-3b^2+ab-b}{2b+a-1}} \mathcal{E}^{\frac{-2b^2-ab+2a-2}{4b+2a-2}} \right). \end{aligned}$$

Summing over all large clusters and applying the non-overlap property as before, we get that the total number of edges added in the completion phase is

$$O_\varepsilon(n) + \widehat{O} \left(n^{1+b} \cdot \beta^{\frac{-3b^2+ab-b}{2b+a-1}} \mathcal{E}^{\frac{-2b^2-ab+2a-2}{4b+2a-2}} \right).$$

- (5) Finally, we need to revisit our setting of β in the final parameter balance. We set

$$\beta = \widehat{O} \left(n^{\frac{2b+a-1}{3b-a+1}} \mathcal{E}^{\frac{-2b-a-4}{6b-2a+2}} \right),$$

and so the number of edges added in the completion phase is

$$\begin{aligned} & O_\varepsilon(n) + \widehat{O} \left(n^{1+b} \cdot \left(n^{\frac{2b+a-1}{3b-a+1}} \mathcal{E}^{\frac{-2b-a-4}{6b-2a+2}} \right)^{\frac{-3b^2+ab-b}{2b+a-1}} \mathcal{E}^{\frac{-2b^2-ab+2a-2}{4b+2a-2}} \right) \\ &= O_\varepsilon(n) + \widehat{O} \left(n^{1+b} \cdot \left(n^{\frac{-3b^2+ab-b}{3b-a+1}} \mathcal{E}^{\frac{(-3b^2+ab-b)(-2b-a-4)}{(6b-2a+2)(2b+a-1)}} \right) \mathcal{E}^{\frac{-2b^2-ab+2a-2}{4b+2a-2}} \right) \\ &= O_\varepsilon(n) + \widehat{O} \left(n^{1+b} \cdot \left(n^{-b} \mathcal{E}^{\frac{-b(-2b-a-4)}{2(2b+a-1)}} \right) \mathcal{E}^{\frac{-2b^2-ab+2a-2}{4b+2a-2}} \right) \\ &= O_\varepsilon(n) + \widehat{O} \left(n \cdot \mathcal{E}^{\frac{2b^2+ab+4b}{4b+2a-2}} \cdot \mathcal{E}^{\frac{-2b^2-ab+2a-2}{4b+2a-2}} \right) \\ &= O_\varepsilon(n) + \widehat{O}(n\mathcal{E}), \end{aligned}$$

completing the edge bound.

Plugging in $a = b = 2/3$ gives error of

$$\beta = \widehat{O}(n^{3/7} \mathcal{E}^{-9/7}),$$

exactly as in the previous section. Plugging in $a = 1/2, b = 1$ for the other distance preserver bound in [12] gives

$$\beta = \widehat{O}(n^{3/7} \mathcal{E}^{-13/14}),$$

thus establishing that Theorem 12 holds unconditionally when $\mathcal{E} = O(1)$. We can also apply our bounds in Theorem 10, although we need to be a little more careful here since the respective settings of $a = 1, b = 1/3$ and $a = b = 2/3$ only hold on some range of parameters. The easiest way forward is to proceed through the previous proof, setting all parameters using *both* possible settings of a, b . For example, a bottlenecked cluster would be one that satisfies

$$|B_=(v, r^*)| \leq |B(v, r^*)|^{\frac{1-a}{2b}} \mathcal{E}^{\frac{1}{2b}}$$

for *both* choices $a = 1, b = 1/3$ and $a = b = 2/3$. The associated large cluster size bound is then

$$x_r = \Omega \left(\min_{(a=1, b=1/3), (a=b=2/3)} \left\{ r^{\frac{2b}{2b+a-1}} \mathcal{E}^{\frac{1}{2b+a-1}} \right\} \right),$$

and so on. In the end, the spanner error is

$$\beta = \widehat{O}(n^{3/7} \mathcal{E}^{-9/7} + n^{2/3} \mathcal{E}^{-17/6}).$$

The former term dominates in the parameter regime

$$\mathcal{E} = \widehat{\Omega}(n^{2/13}),$$

while the latter term dominates when

$$\mathcal{E} = \widehat{O}(n^{2/13}).$$

Thus, in particular, Theorem 12 holds unconditionally in the parameter regime $\mathcal{E} = \widehat{\Omega}(n^{2/13})$.

ACKNOWLEDGMENTS

We are grateful to several anonymous reviewers for useful comments and corrections that have improved this paper.

REFERENCES

- [1] Amir Abboud and Greg Bodwin. 2016. Error amplification for pairwise spanner lower bounds. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Society for Industrial and Applied Mathematics, 841–854.
- [2] Amir Abboud and Greg Bodwin. 2017. The 4/3 additive spanner exponent is tight. *Journal of the ACM (JACM)* 64, 4 (2017), 28:1–28:14.
- [3] Amir Abboud, Greg Bodwin, and Seth Pettie. 2017. A hierarchy of lower bounds for sublinear additive spanners. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Society for Industrial and Applied Mathematics, 568–576.
- [4] Donald Aingworth, Chandra Chekuri, Piotr Indyk, and Rajeev Motwani. 1999. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM J. Comput.* 28, 4 (1999), 1167–1181.
- [5] Ingo Althöfer, Gautam Das, David Dobkin, Deborah Joseph, and José Soares. 1993. On sparse spanners of weighted graphs. *Discrete & Computational Geometry* 9, 1 (1993), 81–100.
- [6] Béla Bollobás, Don Coppersmith, and Michael Elkin. 2005. Sparse distance preservers and additive spanners. *SIAM Journal on Discrete Mathematics* 19, 4 (2005), 1029–1055.
- [7] Surender Baswana, Telikepalli Kavitha, Kurt Mehlhorn, and Seth Pettie. 2010. Additive spanners and (α, β) -spanners. *ACM Transactions on Algorithms (TALG)* 7, 1 (2010), 5.
- [8] Greg Bodwin. 2019. On the structure of unique shortest paths in graphs. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2071–2089.
- [9] Greg Bodwin. 2021. New results on linear size distance preservers. *SIAM J. Comput.* 50, 2 (2021), 662–673.
- [10] Greg Bodwin and Virginia Vassilevska Williams. 2015. Very sparse additive spanners and emulators. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science (ITCS)*. ACM, 377–382.

- [11] Greg Bodwin and Virginia Vassilevska Williams. 2016. Better distance preservers and additive spanners. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Society for Industrial and Applied Mathematics, 855–872.
- [12] Don Coppersmith and Michael Elkin. 2006. Sparse sourcewise and pairwise distance preservers. *SIAM Journal on Discrete Mathematics* 20, 2 (2006), 463–501.
- [13] Marek Cygan, Fabrizio Grandoni, and Telikepalli Kavitha. 2013. On pairwise spanners. In *30th International Symposium on Theoretical Aspects of Computer Science (STACS 2013) (Leibniz International Proceedings in Informatics (LIPIcs))*, Natacha Portier and Thomas Wilke (Eds.), Vol. 20. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 209–220. <http://drops.dagstuhl.de/opus/volltexte/2013/3935>.
- [14] Hsien-Chih Chang, Pawel Gawrychowski, Shay Mozes, and Oren Weimann. 2018. Near-optimal distance emulator for planar graphs. In *26th Annual European Symposium on Algorithms (ESA 2018) (Leibniz International Proceedings in Informatics (LIPIcs))*, Yossi Azar, Hannah Bast, and Grzegorz Herman (Eds.), Vol. 112. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 16:1–16:17. <https://doi.org/10.4230/LIPIcs.ESA.2018.16>
- [15] Shiri Chechik. 2013. New additive spanners. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 498–512.
- [16] Dorit Dor, Shay Halperin, and Uri Zwick. 2000. All-pairs almost shortest paths. *Siam Journal on Computing (SICOMP)* 29, 5 (2000), 1740–1759.
- [17] Michael Elkin and David Peleg. 2004. $(1+\epsilon, \beta)$ -spanner constructions for general graphs. *SIAM J. Comput.* 33, 3 (2004), 608–631.
- [18] Michael Elkin and Seth Pettie. 2016. A linear-size logarithmic stretch path-reporting distance oracle for general graphs. *ACM Transactions on Algorithms (TALG)* 12, 4 (2016), 50.
- [19] Arnold Filtser and Shay Solomon. 2016. The greedy spanner is existentially optimal. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*. 9–17.
- [20] Kshitij Gajjar and Jaikumar Radhakrishnan. 2017. Distance-preserving subgraphs of interval graphs. In *25th Annual European Symposium on Algorithms (ESA 2017) (Leibniz International Proceedings in Informatics (LIPIcs))*, Kirk Pruhs and Christian Sohler (Eds.), Vol. 87. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 39:1–39:13. <https://doi.org/10.4230/LIPIcs.ESA.2017.39>
- [21] Shang-En Huang and Seth Pettie. 2018. Lower bounds on sparse spanners, emulators, and diameter-reducing shortcuts. In *16th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2018) (Leibniz International Proceedings in Informatics (LIPIcs))*, David Eppstein (Ed.), Vol. 101. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 26:1–26:12. <https://doi.org/10.4230/LIPIcs.SWAT.2018.26>
- [22] Telikepalli Kavitha. 2017. New pairwise spanners. *Theory of Computing Systems* 61, 4 (2017), 1011–1036.
- [23] Mathias Bæk Tejs Knudsen. 2014. Additive spanners: A simple construction. In *Scandinavian Workshop on Algorithm Theory*. Springer, 277–281.
- [24] Mathias Bæk Tejs Knudsen. 2017. Additive spanners and distance oracles in quadratic time. In *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017) (Leibniz International Proceedings in Informatics (LIPIcs))*, Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl (Eds.), Vol. 80. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 64:1–64:12. <https://doi.org/10.4230/LIPIcs.ICALP.2017.64>
- [25] Telikepalli Kavitha and Nithin M. Varma. 2015. Small stretch pairwise spanners and approximate d-preservers. *SIAM Journal on Discrete Mathematics* 29, 4 (2015), 2239–2254.
- [26] Kevin Lu. 2019. *New methods for approximating shortest paths*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- [27] Seth Pettie. 2009. Low distortion spanners. *ACM Transactions on Algorithms (TALG)* 6, 1 (2009), 7.
- [28] David Peleg and Jeffrey Ullman. 1989. An optimal synchronizer for the hypercube. *SIAM Journal on Computing (SICOMP)* 18, 4 (1989), 740–747.
- [29] David Peleg and Eli Upfal. 1989. A trade-off between space and efficiency for routing tables. *Journal of the ACM (JACM)* 36, 3 (1989), 510–530.
- [30] Mikkel Thorup and Uri Zwick. 2006. Spanners and emulators with sublinear distance errors. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Society for Industrial and Applied Mathematics, 802–809.
- [31] David P. Woodruff. 2006. Lower bounds for additive spanners, emulators, and more. In *Foundations of Computer Science, 2006. FOCS'06. 47th Annual IEEE Symposium on*. IEEE, 389–398.
- [32] David P. Woodruff. 2010. Additive spanners in nearly quadratic time. In *International Colloquium on Automata, Languages, and Programming*. Springer, 463–474.

Received June 2020; revised June 2021; accepted August 2021