Weighted Additive Spanners

Reyan Ahmed[1], Greg Bodwin[2], Faryad Darabi Sahneh[1], Stephen Kobourov[1], and Richard Spence[1]

[1] Department of Computer Science, University of Arizona
[2] Department of Computer Science, Georgia Institute of Technology

Abstract. A *spanner* of a graph G is a subgraph H that approximately preserves shortest path distances in G. Spanners are commonly applied to compress computation on metric spaces corresponding to weighted input graphs. Classic spanner constructions can seamlessly handle edge weights, so long as error is measured *multiplicatively.* In this work, we investigate whether one can similarly extend constructions of spanners with purely *additive* error to weighted graphs. These extensions are not immediate, due to a key lemma about the size of shortest path neighborhoods that fails for weighted graphs. Despite this, we recover a suitable amortized version, which lets us prove direct extensions of classic $+2$ and $+4$ unweighted spanners (both all-pairs and pairwise) to $+2W$ and $+4W$ weighted spanners, where W is the maximum edge weight. Specifically, we show that a weighted graph G contains all-pairs (pairwise) $+2W$ and $+4W$ weighted spanners of size $O(n^{3/2})$ and $O(n^{7/5})$ ($O(np^{1/3})$ and $O(np^{2/7})$) respectively. For a technical reason, the $+6$ unweighted spanner becomes a $+8W$ weighted spanner; closing this error gap is an interesting remaining open problem. That is, we show that G contains all-pairs (pairwise) $+8W$ weighted spanners of size $O(n^{4/3})$ ($O(np^{1/4})$).

Keywords: Additive spanner · Pairwise spanner · Shortest-path neighborhood

1 Introduction

An $f(\cdot)$-*spanner* of an undirected graph $G = (V, E)$ with $|V| = n$ nodes and $|E| = m$ edges is a subgraph H which preserves pairwise distances in G up to some error prescribed by f; that is,

$$\text{dist}_H(s,t) \leq f(\text{dist}_G(s,t)) \text{ for all nodes } s,t \in V.$$

Spanners were introduced by Peleg and Schäffer [25] in the setting with multiplicative error of type $f(d) = cd$ for some positive constant c. This setting was quickly resolved, with matching upper and lower bounds [4] on the sparsity of a spanner that can be achieved in general. At the other extreme are (purely) c-*additive spanners* (or $+c$ spanners), with error of type $f(d) = d + c$. More generally, if $f(d) = \alpha d + \beta$, we say that H is an (α, β)-*spanner*. Intuitively, additive error is much stronger than multiplicative error; most applications involve shrinking enormous input graphs that are too large to analyze directly, and so it is appealing to avoid error that scales with distance.

Additive spanners were thus initially considered perhaps too good to be true, and they were discovered only for particular classes of input graphs [22]. However, in a surprise to the area, a seminal paper of Aingworth, Chekuri, Indyk, and Motwani [3] proved that nontrivial additive spanners actually exist *in general*: every n-node undirected unweighted graph has a 2-additive spanner on $O(n^{3/2})$ edges. Subsequently, more interesting constructions of additive spanners were found: there are 4-additive spanners on $O(n^{7/5})$ edges [7, 11] and 6-additive spanners on $O(n^{4/3})$ edges [5, 21]. There are also natural generalizations of these results to the *pairwise* setting, where one is given $G = (V, E)$ and a set of demand pairs $P \subseteq V \times V$, where only distances between node pairs $(s, t) \in P$ need to be approximately preserved in the spanner [6, 8, 10, 12, 19, 20].

Despite the inherent advantages of additive error, multiplicative spanners have remained the more well-known and well-applied concept elsewhere in computer science. There seem to be two reasons for this:

1. Abboud and Bodwin [1] (see also [18]) give examples of graphs that have no c-additive spanner on $O(n^{4/3-\varepsilon})$ edges, for any constants $c, \varepsilon > 0$. Some applications call for a spanner on a near-linear number of edges, say $O(n^{1+\varepsilon})$, and hence these must abandon additive error if they need theoretical guarantees for every possible input graph. However, there is some evidence that many graphs of interest bypass this barrier; e.g. graphs with good expansion or girth properties [5].
2. Spanners are often used to compress metric spaces that correspond to *weighted* input graphs. This includes popular applications in robotics [9, 14, 23, 28], asynchronous protocol design [26], etc., and it incorporates the extremely well-studied case of Euclidean spaces which have their own suite of applications (see book [24]). Current constructions of multiplicative spanners can handle edge weights without issue, but purely additive spanners are known for unweighted input graphs only.

Addressing both of these points, Elkin et al. [16] (following [15]) recently provided constructions of *near-additive* spanners for weighted graphs. That is, for any fixed $\varepsilon, t > 0$, every n-node graph $G = (V, E, w)$ has a $(1 + \varepsilon, O(W))$-spanner on $O(n^{1+1/t})$ edges, where W is the maximum edge weight.[3] This extends a classic unweighted spanner construction of Elkin and Peleg [17] to the weighted setting. Additionally, while not explicitly stated in their paper, their method can be adapted to a $+2W$ purely additive spanner on $O(n^{3/2})$ edges (extending [3]).

The goal of this paper is to investigate whether or not all the other constructions of spanners with purely additive error extend similarly to weighted input graphs. As we will discuss shortly, there is a significant barrier to a direct extension of the method from [16]. However, we prove that this barrier can be

[3] Their result is actually a little stronger: W can be the maximum edge weight on the shortest path between the nodes being considered.

Unweighted		Weighted	
Stretch	Size	Stretch	Size
+2	$O(n^{3/2})$ [3]	$+2W$	$O(n^{3/2})$ [**this paper**], [16]
+4	$O(n^{7/5})$ [7,11]	$+4W$	$O(n^{7/5})$ [**this paper**]
+6	$O(n^{4/3})$ [5,21,30]	$+6W$	?
$+c$	$\Omega(n^{4/3-\varepsilon})$ [1,18]	$+8W$	$O(n^{4/3})$ [**this paper**]

Table 1: Table of additive spanner constructions for unweighted and weighted graphs, where W denotes the maximum edge weight.

overcome with some additional technical effort, thus leading to the following constructions. In these theorem statements, all edges have (not necessarily integer) edge weights in $(0, W]$. Let $p = |P|$ denote the number of demand pairs and $n = |V|$ the number of nodes in G.

Theorem 1. *For any $G = (V, E, w)$ and demand pairs P, there is a $+2W$ pairwise spanner with $O(np^{1/3})$ edges. In the all-pairs setting $P = V \times V$, the bound improves to $O(n^{3/2})$.*

Theorem 2. *For any $G = (V, E, w)$ and demand pairs P, there is a $+4W$ pairwise spanner with $O(np^{2/7})$ edges. In the all-pairs setting $P = V \times V$, the bound improves to $O(n^{7/5})$.*

These two results exactly match previous ones for unweighted graphs [3,11, 19,20], with $+2W$ ($+4W$) in place of $+2$ ($+4$). Theorem 1 is partially tight in the following sense: it implies that $O(n^{3/2})$ edges are needed for a $+2W$ spanner when $p = O(n^{3/2})$, and neither of these values can be unilaterally improved. Relatedly, Theorem 2 implies that $O(n^{7/5})$ edges are needed for a $+4W$ spanner when $p = O(n^{7/5})$; it may be possible to improve this, but it would likely imply an improved $+4W$ all-pairs spanner over [11] which will likely be hard to achieve (see discussion in [7]).

Our next two results are actually a bit weaker than the corresponding unweighted ones [13,19,27]: for a technical reason, we take on slightly more error in the weighted setting (the corresponding unweighted results have +6 and +2 error respectively).

Theorem 3. *For any $G = (V, E, w)$ and demand pairs P, there is a $+8W$ pairwise spanner with $O(np^{1/4})$ edges. In the all-pairs setting $P = V \times V$, the bound improves to $O(n^{4/3})$.*

Theorem 4. *For any $G = (V, E, w)$ and demand pairs $P = S \times S$, there is a $+4W$ pairwise spanner with $O(n|S|^{1/2})$ edges.*

We summarize our main results in Table 1, contrasted with known results for unweighted graphs.

1.1 Technical Overview: What's Harder With Weights?

There is a key point of failure in the known constructions of unweighted additive spanners when one attempts the natural extension to weighted graphs. To explain, let us give some technical background. Nearly all spanner constructions start with a *clustering* or *initialization* step: taking the latter exposition [21], a d-*initialization* of a graph G is a subgraph H obtained by choosing d arbitrary edges incident to each node, or all incident edges to a node of degree less than d. After this, many additive spanner constructions leverage the following key fact (the one notable exception is the +2 all-pairs spanner, which is why one can recover the corresponding weighted version from prior work):

Lemma 1 ([11, 13, 19, 20], etc.). *Let G be an undirected unweighted graph, let π be a shortest path, and let H be a d-initialization of G. If π is missing ℓ edges in H, then there are $\Omega(d\ell)$ different nodes adjacent to π in H.*

Proof. For each missing edge $(u, v) \in \pi$, by construction both u and v have degree at least d in H (otherwise, $\deg_H(u) < d$, in which edge (u, v) is added in the d-initialization H). By the triangle inequality, any given node is adjacent to at most three nodes in π. Hence, adding together the $\geq d$ neighbors of each of the ℓ missing edges, we count each node at most three times so the number of nodes adjacent to π is still $\Omega(d\ell)$. □

The difficulty of the weighted setting is largely captured by the fact that Lemma 1 fails when G is edge-weighted. As a counterexample, let π be a shortest path consisting of $\ell + 1$ nodes and ℓ edges of weight ε. Additionally, consider d nodes, each connected to every node along π with an edge of weight $W > \varepsilon\ell$. A candidate d-initialization H consists of selecting every edge of weight W. In this case, all ℓ edges in π are missing in H, but there are still only $d \neq \Omega(d\ell)$ nodes adjacent to π in H.

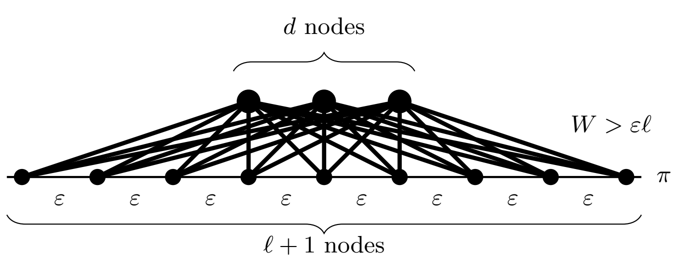

Fig. 1: A counterexample to Lemma 1 for weighted graphs.

The fix, as it turns out, is simple in construction but involved in proof. We simply replace initialization with *light initialization*, where one must specifically add the lightest d edges incident to each node. With this, the proof of Lemma 1 is still not trivial: it remains possible that an external node can be adjacent to arbitrarily many nodes along π, so a direct counting argument fails. However, we show that such occurrences can essentially be amortized against the rising and falling pattern of missing edge weights along π. This leads to a proof that *on average* an external node is adjacent to $O(1)$ nodes in π, which is good enough to push the proof through. We consider this weighted extension of Lemma 1 to be the main technical contribution of this work, and we are hopeful that it may be of independent interest as a structural fact about shortest paths in weighted graphs.

2 Neighborhoods of Weighted Shortest Paths

Here we introduce the extension of Lemma 1. Following the technique in [21], define a *d-light initialization* of a weighted graph $G = (V, E, w)$ to be a subgraph H obtained by including the d lightest edges incident to each node (or all edges incident to a node of degree less than d). Ties between edges of equal weight are broken arbitrarily; for clarity we assume this occurs in the background so that we can unambiguously refer to "the lightest d edges" incident to a node. We prove the weighted analogue of Lemma 1.

Theorem 5. *If H is a d-light initialization of an undirected weighted graph G, and there is a shortest path π in G that is missing ℓ edges in H, then there are $\Omega(d\ell)$ nodes adjacent to π in H.*

We give some definitions and notation which will be useful in the proof of Theorem 5. Let s and t be the endpoints of a shortest path π, and let $M := \pi \setminus E(H)$ be the set of edges in π currently missing in H so that $|M| = \ell$. For convenience we consider these edges to be *oriented* from s to t, so we write $(u, v) \in M$ to mean that $\mathrm{dist}_G(s, u) < \mathrm{dist}_G(s, v)$ and $\mathrm{dist}_G(u, t) > \mathrm{dist}_G(v, t)$. Suppose the edges in M are labeled in order $e_1, e_2, \ldots, e_\ell$ where $e_i = (u_i, v_i)$, and let w_i denote the weight of edge e_i. Given $u \in V$, let $N^*(u)$ denote the *d-neighborhood* of u as follows:

$$N^*(u) := \{v \in V \ \mid \ (u, v) \text{ is one of the lightest } d \text{ edges incident to } u\}.$$

We will show that the size of the union of the d-neighborhoods of the nodes u_1, $\ldots$, u_ℓ is $\Omega(d\ell)$, that is

$$\left| \bigcup_{(u,v)\in M} N^*(u) \right| = \Omega(d\ell)$$

noting that the above set is a subset of all nodes adjacent to π. In particular, the above set may not contain nodes v' connected to $u \in \pi$ by an edge that is 1) among the d lightest incident to v', 2) *not* among the d lightest incident to u. However, the above set necessarily contains all nodes v' which are connected to some u_i or v_i by an edge among the d lightest incident to u_i or v_i. We remark that if the d-neighborhoods $N^*(u_1)$, $N^*(u_2)$, $\ldots$, $N^*(u_\ell)$ are pairwise disjoint, then $|\bigcup_{(u,v)\in M} N^*(u)| = d\ell$, which immediately implies there are at least $d\ell$ nodes adjacent to π in H. Hence for the remainder of the proof, we assume there exist i and k with $1 \le i < k \le \ell$ such that $N^*(u_i) \cap N^*(u_k)$ is nonempty. We use the convention that if a and b are integers with $b < a$, then $\sum_{i=a}^{b} f(i) = 0$. The following lemma holds (see Figure 2):

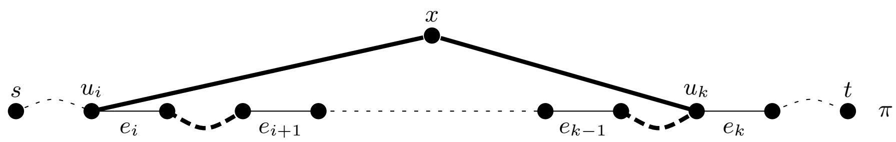

Fig. 2: Illustration of Lemma 2. The bold dashed curves represent subpaths in H.

Lemma 2. *Let π be a shortest path, let $x \in V$ be a node such that $x \in N^*(u_i) \cap N^*(u_k)$ for some $1 \le i < k \le \ell$, and consider the edges $e_i, \ldots, e_k \in M$ with weights $w_i, \ldots, w_k$. Then*

$$w_k \ge \sum_{i'=i+1}^{k-1} w_{i'}.$$

Proof. Consider the subpath of π from u_i to u_k, denoted $\pi[u_i \rightsquigarrow u_k]$. We have

$$\begin{aligned} \sum_{i'=i}^{k-1} w_{i'} &\le \text{length}\,(\pi[u_i \rightsquigarrow u_k]) \\ &\le w(u_i, x) + w(x, u_k) && (\pi[u_i \rightsquigarrow u_k] \text{ is a shortest path}) \\ &\le w_i + w_k \end{aligned}$$

where the last inequality follows from the fact that edges $(u_i, x), (x, u_k)$ are among the d lightest edges incident to u_i and u_k respectively (since $x \in N^*(u_i) \cap N^*(u_k)$), but e_i and e_k are not, since they are omitted from H. Lemma 2 follows by subtracting w_i from both sides of the above inequality. □

For the next part, for edge $e \in M$, say that e is *pre-heavy* if its weight is strictly greater than the preceding edge in M, and/or *post-heavy* if its weight is strictly greater than the following edge in M. For notational convenience, if an edge is not pre-heavy, we say the edge is pre-light. Similarly, if an edge is not post-heavy, we say the edge is post-light. By convention, the first edge $e_1 \in M$ is pre-light and the last edge $e_\ell \in M$ is post-light. We state the following simple lemma; recall that $|M| = \ell$.

Lemma 3. *Either more than $\frac{\ell}{2}$ edges in M are pre-light, or more than $\frac{\ell}{2}$ edges in M are post-light.*

Proof. Let S_1 be the set of edges in M which are pre-light, and let S_2 be the set of edges in M which are post-light. Note that $e_1 \in S_1$ and $e_\ell \in S_2$. For each of the $\ell - 1$ pairs of consecutive edges (e_i, e_{i+1}) in M where $i = 1, \dots, \ell - 1$, it is immediate by definition that either $e_i \in S_2$ or $e_{i+1} \in S_1$ (or both if $w_i = w_{i+1}$). These statements imply $|S_1| + |S_2| \geq \ell + 1$, so at least one of S_1 or S_2 has cardinality at least $\frac{\ell+1}{2} > \frac{\ell}{2}$. □

In the sequel, we assume without loss of generality that more than $\frac{\ell}{2}$ edges in M are pre-light; the other case is symmetric by exchanging the endpoints s and t of π. We can now say the point of the previous two lemmas: together, they imply that *most* edges $(u, v) \in M$ have mostly non-overlapping d-neighborhoods $N^*(u)$. That is:

Lemma 4. *Let π be a shortest path. For any node $x \in V$, there exist at most three nodes u along π such that $x \in N^*(u)$ and edge $(u, v) \in M$ is pre-light.*

Proof. Suppose for sake of contradiction there exist four nodes u_i, u_a, u_b, u_k with $1 \leq i < a < b < k \leq \ell$ such that x belongs to the d-neighborhoods of u_i, u_a, u_b, and u_k, and the edges (u_i, v_i), (u_a, v_a), (u_b, v_b), and (u_k, v_k) are pre-light. In particular, we have $k \geq i + 3$ and $x \in N^*(u_i) \cap N^*(u_k)$. By Lemma 2 and the above observation, we have

$$w_k \geq \sum_{i'=i+1}^{k-1} w'_i = w_{i+1} + \ldots + w_{k-1} \geq w_{i+1} + w_{k-1}$$

By assumption, $e_k = (u_k, v_k)$ is pre-light, so $w_{k-1} \geq w_k$, and the above inequality implies $w_k \geq w_{i+1} + w_{k-1} \geq w_{i+1} + w_k$, or $w_{i+1} = 0$. Since edge weights are strictly positive, we have contradiction, proving Lemma 4. □

Finally, define set X^* as follows:

$$X^* := \bigcup_{\substack{(u,v) \in M \\ \text{is pre-light}}} N^*(u).$$

By Lemma 3 and the above pre-heavy assumption, there are more than $\frac{\ell}{2}$ pre-light edges (u, v), so the multiset containing all d-neighborhoods $N^*(u)$ contains more than $\frac{d\ell}{2}$ nodes. By Lemma 4, any given node is contained in at most three of these d-neighborhoods, implying $|X^*| > \dfrac{d\ell}{6}$. Since X^* is a subset of $|\bigcup_{(u,v) \in M} N^*(u)|$, we conclude that there are $\Omega(d\ell)$ nodes adjacent to π in H. proving Theorem 5.

3 Spanner Constructions

We show how Theorem 5 can be used to construct additive spanners on edge-weighted graphs. These constructions are not significant departures from prior work; the main difference is applying Theorem 5 in the right place.

3.1 Subset and Pairwise Spanners

Definition 1 (Pairwise/Subset Additive Spanners). *Given a graph $G = (V, E, w)$ and a set of demand pairs $P \subseteq V \times V$, a subgraph $H = (V, E_H \subseteq E, w)$ is a $+c$* pairwise spanner *of G, P if*

$$\mathrm{dist}_H(s,t) \leq \mathrm{dist}_G(s,t) + c \text{ for all } (s,t) \in P.$$

When $P = S \times S$ for some $S \subseteq V$, we say that H is a $+c$ subset spanner *of G, S.*

In the following results, all graphs G are undirected and connected with (not necessarily integer) edge weights in the interval $(0, W]$, where W is the maximum edge weight. Let $|V| = n$, let $p = |P|$ denote the number of demand pairs (for pairwise spanners), and let $\sigma = |S|$ denote the number of sources (for subset spanners).

Theorem 6. *Any n-node graph $G = (V, E, w)$ with source nodes $S \subseteq V$ has a $+4W$ subset spanner with $O(n\sigma^{1/2})$ edges.*

Proof. The construction of the $+4W$ subset spanner H is as follows, essentially following [21]. Let d be a parameter of the construction, and let H be a d-light initialization of G. Then, while there are nodes $s, t \in S$ such that $\mathrm{dist}_H(s,t) > \mathrm{dist}_G(s,t) + 4W$, choose any $s \rightsquigarrow t$ shortest path $\pi(s,t)$ in G and add all its edges to H. It is immediate that this algorithm terminates with H a $+4W$ subset spanner of G, so we now analyze the number of edges $|E_H|$ in the final subgraph H.

At any point in the algorithm, say that an ordered pair of nodes $(s, v) \in S \times V$ is *near-connected* if there exists v' adjacent to v in H such that $\mathrm{dist}_H(s, v') = \mathrm{dist}_G(s, v')$. We then have the following observation

$$\mathrm{dist}_H(s,v) \leq \mathrm{dist}_H(s,v') + W = \mathrm{dist}_G(s,v') + W. \tag{1}$$

When nodes $s, t \in S$ with shortest path $\pi(s,t)$ are considered in the construction, there are two cases:

1. If there are two nodes v', v'' adjacent in H to a node $v \in \pi(s,t)$, and the pairs (s, v) and (t, v) are near-connected, then we have by triangle inequality and (1):

$$\begin{aligned}\mathrm{dist}_H(s,t) &\leq \mathrm{dist}_H(s,v) + \mathrm{dist}_H(t,v)\\ &\leq (\mathrm{dist}_G(s,v') + W) + (\mathrm{dist}_G(t,v'') + W)\\ &= \mathrm{dist}_G(s,v') + \mathrm{dist}_G(t,v'') + 2W\\ &\leq \mathrm{dist}_G(s,v) + \mathrm{dist}_G(t,v) + 4W\\ &= \mathrm{dist}_G(s,t) + 4W.\end{aligned}$$

where the last equality follows from the optimal substructure property of shortest paths. In this case, the path $\pi(s,t)$ is not added to H.

2. Otherwise, suppose there is no node v' adjacent in H to a node $v \in \pi(s,t)$ where (s,v) and (t,v) are near-connected. After adding the path $\pi(s,t)$ to H, every such node v' becomes near-connected to both s and t. If there are ℓ edges in $\pi(s,t)$ currently missing in H, then by Theorem 5 we have $\Omega(\ell d)$ nodes adjacent to $\pi(s,t)$, so $\Omega(\ell d)$ node pairs in $S \times V$ go from not near-connected to near-connected. Since there are σn node pairs in $S \times V$, we add a total of $O(\sigma n/d)$ edges to H in this case.

Putting these together, the final size of H is $|E_H| = O\left(nd + \frac{\sigma n}{d}\right)$. Setting $d := \sqrt{\sigma}$ proves Theorem 6. □

We now give our constructions for pairwise spanners. The following lemma will be useful:

Lemma 5 ([7]). *Let $a, b > 0$ be absolute constants, and suppose there is an algorithm that, on input G, P, produces a subgraph H on $O(n^a|P|^b)$ edges satisfying*

$$\mathrm{dist}_H(s,t) \leq \mathrm{dist}_G(s,t) + c$$

for at least a constant fraction of the demand pairs $(s,t) \in P$. Then there is a $+c$ pairwise spanner H' of G, P on $O(n^a|P|^b)$ edges.

Using the slack to satisfy only a constant fraction of the demand pairs, we have the following proofs.

Theorem 7. *Any graph G with demand pairs P has a $+2W$ pairwise spanner with $O(np^{1/3})$ edges.*

Proof. Let d and ℓ be parameters of the construction, and let H be a d-light initialization of G. For each demand pair $(s,t) \in P$ whose shortest path $\pi(s,t)$ is missing at most ℓ edges in H, add all edges in $\pi(s,t)$ to H. By Theorem 5, any remaining demand pair $(s,t) \in P$ has $\Omega(d\ell)$ nodes adjacent to $\pi(s,t)$. Let R be a random sample of nodes obtained by including each one independently with probability $1/(\ell d)$; thus, with constant probability or higher, there exists $r \in R$ and $v \in \pi(s,t)$ such that nodes r and v are adjacent in H. Add to H a shortest path tree rooted at each $r \in R$. We then compute:

$$\begin{aligned}
\mathrm{dist}_H(s,t) &\leq \mathrm{dist}_H(s,r) + \mathrm{dist}_H(r,t) \\
&= \mathrm{dist}_G(s,r) + \mathrm{dist}_G(r,t) \\
&\leq \mathrm{dist}_G(s,v) + \mathrm{dist}_G(v,t) + 2W \\
&= \mathrm{dist}_G(s,t) + 2W.
\end{aligned}$$

The distance for each pair $(s,t) \in P$ is approximately preserved in H with at least a constant probability, which is sufficient for Lemma 5. The number of edges in the final subgraph H is

$$|E(H)| = O(nd + \ell p + n^2/(\ell d));$$

setting $\ell = n/p^{2/3}$ and $d = p^{1/3}$ proves Theorem 7. □

Theorem 8. *Any graph G with demand pairs P has a $+4W$ pairwise spanner with $O(np^{2/7})$ edges.*

Proof. Let d and ℓ be parameters of the construction, and let H be a d-light initialization of G. For each demand pair $(s,t) \in P$ whose shortest path $\pi(s,t)$ is missing at most ℓ edges in H, add all edges in $\pi(s,t)$ to H. To handle each $(s,t) \in P$ whose shortest path $\pi(s,t)$ is missing at least n/d^2 edges in H, we let R_1 be a random sample of nodes obtained by including each node independently with probability d^2/n, then add a shortest path tree rooted at each $r \in R_1$ to H. By an identical analysis to Theorem 7, for each such pair, with constant probability or higher we have

$$\mathrm{dist}_H(s,t) \leq \mathrm{dist}_G(s,t) + 2W.$$

Finally, we consider the "intermediate" pairs $(s,t) \in P$ whose shortest path $\pi(s,t)$ is missing more than ℓ but fewer than n/d^2 edges in H. We add the first and last ℓ missing edges in $\pi(s,t)$ to the spanner; we will refer to the *prefix* (resp. *suffix*) of $\pi(s,t)$ to mean the shortest prefix (suffix) containing these ℓ missing edges. By Theorem 5, there are $\Omega(\ell d)$ nodes adjacent to the prefix and $\Omega(\ell d)$ nodes adjacent to the suffix. Let R_2 be a random sample of nodes obtained by including each node with probability $1/(\ell d)$, and for each pair $r, r' \in R_2$, add to H all edges in the shortest $r \rightsquigarrow r'$ path in G among the paths that are missing at most n/d^2 edges (ignore any pair r, r' if no such path exists). With constant probability or higher, we sample r, r' adjacent to nodes v, v' in the prefix, suffix respectively, in which case we have:

$$\begin{aligned}
\mathrm{dist}_H(s,t) &\leq \mathrm{dist}_H(s,v) + \mathrm{dist}_H(v,v') + \mathrm{dist}_H(v',t) \\
&= \mathrm{dist}_G(s,v) + \mathrm{dist}_H(v,v') + \mathrm{dist}_G(v',t) \\
&\leq \mathrm{dist}_G(s,v) + \mathrm{dist}_H(r,r') + 2W + \mathrm{dist}_G(v',t).
\end{aligned}$$

Notice that $\mathrm{dist}_H(r,r') \leq 2W + \mathrm{dist}_G(v,v')$, due to the existence of the path $r \circ \pi(s,t)[v,v'] \circ r'$ which is indeed missing $\leq n/d^2$ edges. Thus we may continue:

$$\begin{aligned}
&\leq \mathrm{dist}_G(s,v) + \mathrm{dist}_G(v,v') + 4W + \mathrm{dist}_G(v',t) \\
&= \mathrm{dist}_G(s,t) + 4W.
\end{aligned}$$

The distance for each pair $(s,t) \in P$ is approximately preserved in H with at least constant probability, which again suffices by Lemma 5, and the number of edges in H is

$$|E(H)| = O\left(nd + p\ell + n^3/(\ell^2 d^4)\right).$$

Setting $\ell = n/p^{5/7}$ and $d = p^{2/7}$ completes the proof of Theorem 8. ◻

Theorem 9. *Any graph G with demand pairs P has a $+8W$ pairwise spanner containing $O(np^{1/4})$ edges.*

Proof. Let ℓ, d be parameters of the construction and let H be a d-light initialization of G. For each $(s,t) \in P$ whose shortest path $\pi(s,t)$ is missing $\leq \ell$ edges in H, add all edges in $\pi(s,t)$ to H. Otherwise, like before, we add the first and last ℓ missing edges of $\pi(s,t)$ to H (prefix and suffix). Then, randomly sample a set R by including each node with probability $1/(\ell d)$, and use Theorem 6 to add a $+4W$ subset spanner on the nodes in R. By Theorem 5, the prefix and suffix each have $\Omega(\ell d)$ adjacent nodes. Thus, with constant probability or higher, we sample $r, r' \in R$ adjacent to v, v' in the added prefix and suffix respectively. We then compute:

$$\begin{aligned}
\mathrm{dist}_H(s,t) &\leq \mathrm{dist}_H(s,v) + \mathrm{dist}_H(v,v') + \mathrm{dist}_H(v',t) \\
&\leq \mathrm{dist}_G(s,v) + \mathrm{dist}_H(v,v') + \mathrm{dist}_G(v',t) \\
&\leq \mathrm{dist}_G(s,v) + \mathrm{dist}_H(r,r') + 2W + \mathrm{dist}_G(v',t) \\
&\leq \mathrm{dist}_G(s,v) + \mathrm{dist}_G(r,r') + 6W + \mathrm{dist}_G(v',t) \\
&\leq \mathrm{dist}_G(s,v) + \mathrm{dist}_G(v,v') + 8W + \mathrm{dist}_G(v',t) \\
&= \mathrm{dist}_G(s,t) + 8W.
\end{aligned}$$

Again, the distance for each pair $(s,t) \in P$ is approximately preserved in H with at least constant probability, which suffices by Lemma 5. The number of edges in H is

$$|E(H)| = O\left(nd + p\ell + n^{3/2}/\sqrt{\ell d}\right).$$

Setting $\ell = n/p^{3/4}$ and $d = p^{1/4}$ completes the proof of Theorem 9. □

4 All-pairs Additive Spanners

We now turn to the *all-pairs* setting, i.e., demand pairs $P = V \times V$. We use the following lemma from [7]:

Lemma 6 ([7]). *Let G be a graph, and suppose one can choose a function π that associates each node pair to a path between them with the following properties:*

- *for all (s,t) the length of the path $\pi(s,t)$ (i.e., the sum of its edge weights) is $\leq \mathrm{dist}_G(s,t) + k$,*
- *π depends only on the input graph G and the number of demand pairs $|P|$ (but not otherwise on the contents of P), and*
- *for some parameter p^* and any $|P| \geq p^*$ demand pairs, we have*

$$\left| \bigcup_{(s,t)\in P} \pi(s,t) \right| < |P|.$$

Then there is an all-pairs k-additive spanner of G containing $\leq p^*$ edges.

Notice that all the above pairwise spanner constructions are *demand-oblivious* – that is, the approximate shortest paths analyzed in order to preserve each demand pair in the spanner depend on the random bits of the construction, which in turn depend on the number of demand pairs $|P|$, but they do not otherwise depend on the contents of P. See [7] for more discussion of this property. Thus we may apply Lemma 6 as follows. For the $+2W$ pairwise bound of $O(np^{1/3})$ provided in Theorem 7, we note that the bound is $< p$ for $p = \Omega(n^{3/2})$ demand pairs (and a sufficiently large constant in the Ω). Hence, taking $p^* = \Theta(n^{3/2})$, Lemma 6 says:

Theorem 10. *Every graph has a $+2W$ spanner on $O(n^{3/2})$ edges.*

Identical logic applied to Theorems 8 and 9 gives:

Theorem 11. *Every n-node graph has a $+4W$ spanner on $O(n^{7/5})$ edges.*

Theorem 12. *Every n-node graph has a $+8W$ additive spanner on $O(n^{4/3})$ edges.*

5 Conclusions and Open Problems

We have shown that most important unweighted additive spanner constructions have natural weighted analogues. At present, the exceptions are the $+4W$ subset spanner on $O(n|S|^{1/2})$ edges (which should probably have only $+2W$ error) and the $+8W$ all-pairs/pairwise spanners (which should probably have only $+6W$ error). Closing these error gaps is an interesting open problem. It would also be interesting to obtain weighted analogues of related concepts, most notably, the Thorup-Zwick emulators [29], which are optimal [2] in essentially the same way that the 6-additive spanner on $O(n^{4/3})$ edges is optimal.

Finally, as mentioned earlier, it would be interesting to find constructions of *purely* additive spanners parametrized by some other statistic besides the maximum edge weight W; a natural parameter is $W(u, v)$, the maximum edge weight along a shortest u-v path.

References

1. Amir Abboud and Greg Bodwin. The 4/3 additive spanner exponent is tight. *Journal of the ACM (JACM)*, 64(4):1–20, 2017.
2. Amir Abboud, Greg Bodwin, and Seth Pettie. A hierarchy of lower bounds for sublinear additive spanners. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 568–576. Society for Industrial and Applied Mathematics, 2017.
3. Donald Aingworth, Chandra Chekuri, Piotr Indyk, and Rajeev Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM Journal on Computing*, 28:1167–1181, 04 1999.
4. Ingo Althöfer, Gautam Das, David Dobkin, and Deborah Joseph. Generating sparse spanners for weighted graphs. In *Scandinavian Workshop on Algorithm Theory (SWAT)*, pages 26–37, Berlin, Heidelberg, 1990. Springer Berlin Heidelberg.
5. Surender Baswana, Telikepalli Kavitha, Kurt Mehlhorn, and Seth Pettie. Additive spanners and (α, β)-spanners. *ACM Transactions on Algorithms (TALG)*, 7(1):5, 2010.

6. Greg Bodwin. Linear size distance preservers. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 600–615. Society for Industrial and Applied Mathematics, 2017.
7. Greg Bodwin. A note on distance-preserving graph sparsification. *arXiv preprint arXiv:2001.07741*, 2020.
8. Greg Bodwin and Virginia Vassilevska Williams. Better distance preservers and additive spanners. In *Proceedings of the Twenty-seventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 855–872. Society for Industrial and Applied Mathematics, 2016.
9. Leizhen Cai and J. Mark Keil. Computing visibility information in an inaccurate simple polygon. *International Journal of Computational Geometry & Applications*, 7(6):515–538, 1997.
10. Hsien-Chih Chang, Pawel Gawrychowski, Shay Mozes, and Oren Weimann. Near-Optimal Distance Emulator for Planar Graphs. In *Proceedings of 26th Annual European Symposium on Algorithms (ESA 2018)*, volume 112, pages 16:1–16:17, 2018.
11. Shiri Chechik. New additive spanners. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms (SODA)*, pages 498–512. Society for Industrial and Applied Mathematics, 2013.
12. Don Coppersmith and Michael Elkin. Sparse sourcewise and pairwise distance preservers. *SIAM Journal on Discrete Mathematics*, 20(2):463–501, 2006.
13. Marek Cygan, Fabrizio Grandoni, and Telikepalli Kavitha. On pairwise spanners. In *Proceedings of 30th International Symposium on Theoretical Aspects of Computer Science (STACS 2013)*, volume 20, pages 209–220, 2013.
14. Andrew Dobson and Kostas E Bekris. Sparse roadmap spanners for asymptotically near-optimal motion planning. *The International Journal of Robotics Research*, 33(1):18–47, 2014.
15. Michael Elkin. Computing almost shortest paths. *ACM Transactions on Algorithms (TALG)*, 1(2):283–323, 2005.
16. Michael Elkin, Yuval Gitlitz, and Ofer Neiman. Almost shortest paths and PRAM distance oracles in weighted graphs. *arXiv preprint arXiv:1907.11422*, 2019.
17. Michael Elkin and David Peleg. $(1+\epsilon, \beta)$-spanner constructions for general graphs. *SIAM Journal on Computing*, 33(3):608–631, March 2004.
18. Shang-En Huang and Seth Pettie. Lower bounds on sparse spanners, emulators, and diameter-reducing shortcuts. In *Proceedings of 16th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT)*, pages 26:1–26:12, 2018.
19. Telikepalli Kavitha. New pairwise spanners. *Theory of Computing Systems*, 61(4):1011–1036, Nov 2017.
20. Telikepalli Kavitha and Nithin M. Varma. Small stretch pairwise spanners and approximate d-preservers. *SIAM Journal on Discrete Mathematics*, 29(4):2239–2254, 2015.
21. Mathias Bæk Tejs Knudsen. Additive spanners: A simple construction. In *Scandinavian Workshop on Algorithm Theory (SWAT)*, pages 277–281. Springer, 2014.
22. Arthur Liestman and Thomas Shermer. Additive graph spanners. *Networks*, 23:343 – 363, 07 1993.
23. James D Marble and Kostas E Bekris. Asymptotically near-optimal planning with probabilistic roadmap spanners. *IEEE Transactions on Robotics*, 29(2):432–444, 2013.
24. Giri Narasimhan and Michiel Smid. *Geometric Spanner Networks*. Cambridge University Press, New York, NY, USA, 2007.

25. David Peleg and Alejandro A. Schäffer. Graph spanners. *Journal of Graph Theory*, 13(1):99–116, 1989.
26. David Peleg and Eli Upfal. A trade-off between space and efficiency for routing tables. *Journal of the ACM (JACM)*, 36(3):510–530, 1989.
27. Seth Pettie. Low distortion spanners. *ACM Transactions on Algorithms (TALG)*, 6(1):7, 2009.
28. Oren Salzman, Doron Shaharabani, Pankaj K Agarwal, and Dan Halperin. Sparsification of motion-planning roadmaps by edge contraction. *The International Journal of Robotics Research*, 33(14):1711–1725, 2014.
29. Mikkel Thorup and Uri Zwick. Spanners and emulators with sublinear distance errors. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 802–809. Society for Industrial and Applied Mathematics, 2006.
30. David P Woodruff. Additive spanners in nearly quadratic time. In *International Colloquium on Automata, Languages, and Programming*, pages 463–474. Springer, 2010.