

On additive spanners in weighted graphs with local error

Reyan Ahmed¹, Greg Bodwin², Keaton Hamm³, Stephen Kobourov¹, and Richard Spence¹

¹ Department of Computer Science, University of Arizona

² Department of Computer Science, University of Michigan

³ Department of Mathematics, University of Texas at Arlington

Abstract. An *additive $+\beta$ spanner* of a graph G is a subgraph which preserves distances up to an additive $+\beta$ error. Additive spanners are well-studied in unweighted graphs but have only recently received attention in weighted graphs [Elkin et al. 2019 and 2020, Ahmed et al. 2020]. This paper makes two new contributions to the theory of weighted additive spanners.

For weighted graphs, [Ahmed et al. 2020] provided constructions of sparse spanners with *global* error $\beta = cW$, where W is the maximum edge weight in G and c is constant. We improve these to *local* error by giving spanners with additive error $+cW(s, t)$ for each vertex pair (s, t) , where $W(s, t)$ is the maximum edge weight along the shortest s - t path in G . These include pairwise $+(2+\varepsilon)W(\cdot, \cdot)$ and $+(6+\varepsilon)W(\cdot, \cdot)$ spanners over vertex pairs $\mathcal{P} \subseteq V \times V$ on $O_\varepsilon(n|\mathcal{P}|^{1/3})$ and $O_\varepsilon(n|\mathcal{P}|^{1/4})$ edges for all $\varepsilon > 0$, which extend previously known unweighted results up to ε dependence, as well as an all-pairs $+4W(\cdot, \cdot)$ spanner on $\tilde{O}(n^{7/5})$ edges.

Besides sparsity, another natural way to measure the quality of a spanner in weighted graphs is by its *lightness*, defined as the total edge weight of the spanner divided by the weight of an MST of G . We provide a $+\varepsilon W(\cdot, \cdot)$ spanner with $O_\varepsilon(n)$ lightness, and a $+(4+\varepsilon)W(\cdot, \cdot)$ spanner with $O_\varepsilon(n^{2/3})$ lightness. These are the first known additive spanners with nontrivial lightness guarantees. All of the above spanners can be constructed in polynomial time.

1 Introduction

Given an undirected graph $G(V, E)$, a *spanner* is a subgraph H which approximately preserves distances in G up to some error. Spanners are an important primitive in the literature on network design and shortest path algorithms, with applications in motion planning in robotics [9, 15, 28, 32], asynchronous protocol design [30], approximate shortest path algorithms [16], and much more; see survey [3]. One general goal in research on spanners is to minimize the size or sparsity of the spanner (as measured by the number of edges $|E(H)|$), given some error by which distances can be distorted. For weighted graphs, another desirable goal is to minimize the *lightness*, defined as the total weight of the spanner divided by the weight of a minimum spanning tree (MST) of G .

Spanners were introduced in the 1980s by Peleg and Schäffer [29], who first considered *multiplicative* error. A subgraph H is a (multiplicative) k -spanner of $G(V, E)$ if $d_H(s, t) \leq k \cdot d_G(s, t)$ for all vertices $s, t \in V$, where $d_G(s, t)$ is the distance between s and t in G . Since H is a subgraph, we also have $d_G(s, t) \leq d_H(s, t)$ by definition, so the distances in H approximate those of G within a multiplicative factor of k , sometimes called the stretch factor. Althöfer et al. [6] showed that all n -vertex graphs have multiplicative $(2k - 1)$ -spanners on $O(n^{1+1/k})$ edges with $O(n/k)$ lightness, and this edge bound is the best possible assuming the girth conjecture by Erdős [21] from extremal combinatorics. Meanwhile, this initial lightness bound has been repeatedly improved in follow-up work, and the optimal bound still remains open [11, 13, 19, 22, 27].

Multiplicative spanners are extremely well applied in computer science. However, they are typically applied to very large graphs where it may be undesirable to take on errors that scale with the (possibly very large) distances in the input graph. A more desirable error is *additive* error which does not depend on the original graph distances at all:

Definition 1 (Additive $+\beta$ spanner). *Given a graph $G(V, E)$ and $\beta \geq 0$, a subgraph H is a $+\beta$ spanner of G if*

$$d_G(s, t) \leq d_H(s, t) \leq d_G(s, t) + \beta \tag{1}$$

for all vertices $s, t \in V$.

A pairwise $+\beta$ spanner is a subgraph H for which (1) only needs to hold for specific vertex pairs $\mathcal{P} \subseteq V \times V$ given on input, and a subsetwise spanner is a pairwise spanner with $\mathcal{P} = \mathcal{S} \times \mathcal{S}$ for some $\mathcal{S} \subseteq V$ (if $\mathcal{P} = V \times V$, these are sometimes called all-pairs spanners, for clarity). It is known that all unweighted graphs $G(V, E)$ with $|V| = n$ have (all-pairs) $+2$ spanners on $O(n^{3/2})$ edges [5, 26], $+4$ spanners on $\tilde{O}(n^{7/5})$ edges [8, 12], and $+6$ spanners on $O(n^{4/3})$ edges [7, 26, 33]. On the negative side, there exist graphs which have no $+\beta$ spanner on $O(n^{4/3-\epsilon})$ edges even for arbitrarily large constant β [1]. This presents a barrier to using additive spanners in applications where a very sparse subgraph, say on $O(n^{1.001})$ edges, is needed. However, the lower bound construction in [1] is rather pathological; tradeoffs do continue for certain natural classes of graphs with good girth or expansion properties [7], and recent experimental work [4] showed that tradeoffs seem to continue for graphs constructed from common random graph models.

A more serious barrier preventing the applicability of additive spanners is that classic constructions only apply to *unweighted* graphs, while many naturally-occurring metrics are not expressible by a unit-weight graph. To obtain additive spanners of weighted graphs $G = (V, E, w)$ where $w : E \rightarrow \mathbb{R}^+$, the error term $+\beta$ needs to scale somehow with the edge weights of the input graph. Prior work [2] has considered *global* error of type $\beta = cW$, where $W = \max_{e \in E} w(e)$ is the maximum edge weight in the input graph and c is a constant. However, a more desirable paradigm studied by Elkin, Gitlitz, and Neiman [17, 18] is to consider *local* error in terms of the maximum edge weight along a shortest s - t path:

Definition 2 (Local $+cW(\cdot, \cdot)$ spanner). *Given a graph $G(V, E)$, subgraph H is a (local) $+cW(\cdot, \cdot)$ spanner if $d_H(s, t) \leq d_G(s, t) + cW(s, t)$ for all $s, t \in V$, where $W(s, t)$ is the maximum edge weight along a shortest path $\pi(s, t)$ in G .⁴*

It is often the case that $W(s, t) \ll W$ for many vertex pairs (s, t) in which a $+cW(\cdot, \cdot)$ spanner has much less additive error for such vertex pairs. Additionally, a $+cW(\cdot, \cdot)$ spanner is also a multiplicative $(c + 1)$ -spanner, whereas a $+cW$ spanner can have unbounded multiplicative stretch. This relationship between additive and multiplicative stretch is thematic in the area [7, 20].

Sparse Local Additive Spanners. Weighted additive spanners were first studied by Elkin et al. [17], who gave a local $+2W(\cdot, \cdot)$ spanner with $O(n^{3/2})$ edges, as well as a “mixed” spanner with a similar error type. Ahmed, Bodwin, Sahneh, Kobourov, and Spence [2] gave a comprehensive study of weighted additive spanners, including a *global* $+4W$ spanner with $\tilde{O}(n^{7/5})$ edges and $+8W$ spanner with $O(n^{4/3})$ edges, analogous to the previously-known unweighted constructions. The $+6$ unweighted error vs. $+8W$ global weighted error left a gap to be closed; this was mostly closed in a recent follow-up work of Elkin et al. [18], who gave a *local* all-pairs $+(6 + \varepsilon)W(\cdot, \cdot)$ spanner on $O_\varepsilon(n^{4/3})$ edges⁵ by generalizing the $+6$ spanner by Knudsen [26], and a subsetwise $+(2 + \varepsilon)W(\cdot, \cdot)$ spanner on $O_\varepsilon(n\sqrt{|\mathcal{S}|})$ edges [18] (following a similar $+2$ subsetwise spanner in unweighted graphs by Elkin (unpublished), later published in [14, 31]). Our first contribution is the improvement of several remaining known constructions of weighted additive spanners from global to local error.

Theorem 1. *Let $\varepsilon > 0$. Then every weighted graph G and set $\mathcal{P} \subseteq V \times V$ of vertex pairs has:*

1. *a deterministic pairwise $+(2 + \varepsilon)W(\cdot, \cdot)$ spanner on $O_\varepsilon(n|\mathcal{P}|^{1/3})$ edges,*
2. *a deterministic pairwise $+(6 + \varepsilon)W(\cdot, \cdot)$ spanner on $O_\varepsilon(n|\mathcal{P}|^{1/4})$ edges,*
3. *a pairwise $+2W(\cdot, \cdot)$ spanner on $O(n|\mathcal{P}|^{1/3})$ edges,*
4. *a pairwise $+4W(\cdot, \cdot)$ spanner on $O(n|\mathcal{P}|^{2/7})$ edges, and*
5. *an all-pairs $+4W(\cdot, \cdot)$ spanner on $\tilde{O}(n^{7/5})$ edges.*

Unweighted versions of these results were proved in [8, 23, 24], and weighted versions with global error but without ε dependence were proved in [2]. This theorem is the first to provide versions with local error. Moreover, the first two pairwise constructions are deterministic, unlike the randomized constructions from [3]. Together with [17, 18], the above results complete the task of converting unweighted additive spanners to weighted additive spanners with local error; see Tables 1 and 2.

⁴ If there are multiple shortest s - t paths, then we break ties consistently so that subpaths of shortest paths are also shortest paths.

⁵ We use $O_\varepsilon(f(n))$ as shorthand for $O(\text{poly}(\frac{1}{\varepsilon})f(n))$.

Unweighted			Weighted		
$+\beta$	Size	Ref.	$+\beta$	Size	Ref.
+2	$O(n^{3/2})$	[5, 7, 14, 26]	$+2W(\cdot, \cdot)$	$O(n^{3/2})$	[17]
+4	$\tilde{O}(n^{7/5})$	[12]	$+4W(\cdot, \cdot)$	$\tilde{O}(n^{7/5})$	[this paper]
+6	$O(n^{4/3})$	[7, 26]	$+(6 + \varepsilon)W(\cdot, \cdot)$	$O_\varepsilon(n^{4/3})$	[18]
$+n^{o(1)}$	$\Omega(n^{4/3-\varepsilon})$	[1]			

Table 1: All-pairs additive spanner constructions for unweighted and weighted graphs.

Unweighted				Weighted		
Type	$+\beta$	Size	Ref.	$+\beta$	Size	Ref.
Subset	+2	$O(n\sqrt{ \mathcal{S} })$	[14]	$+(2 + \varepsilon)W(\cdot, \cdot)$	$O_\varepsilon(n\sqrt{ \mathcal{S} })$	[18]
Pairwise	+2	$O(n \mathcal{P} ^{1/3})$	[10, 24]	$+2W(\cdot, \cdot)$	$O(n \mathcal{P} ^{1/3})$	[this paper]
Pairwise	+4	$\tilde{O}(n \mathcal{P} ^{2/7})$	[23]	$+4W(\cdot, \cdot)$	$O(n \mathcal{P} ^{2/7})$	[?]
Pairwise	+6	$O(n \mathcal{P} ^{1/4})$	[23]	$+(6 + \varepsilon)W(\cdot, \cdot)$	$O_\varepsilon(n \mathcal{P} ^{1/4})$	[this paper]

Table 2: Pairwise and subsetwise (purely) additive spanner constructions for unweighted and weighted graphs.

Lightweight Local Additive Spanners. All of the aforementioned results are in terms of the number of edges $|E(H)|$ of the spanner. If minimizing the total edge weight is more desirable than constructing a sparse spanner, a natural problem is to construct *lightweight* spanners. Given a connected graph $G = (V, E)$ with positive edge weights $w : E \rightarrow \mathbb{R}^+$, the *lightness* of a subgraph H is defined by

$$\text{lightness}(H) := \frac{w(H)}{w(\text{MST}(G))} \quad (2)$$

where $w(H)$ and $w(\text{MST}(G))$ are the sum of edge weights in H and an MST of G , respectively. Section 4 highlights why none of the aforementioned sparse spanners have good lightness guarantees. Our second contribution is the following:

Theorem 2. *Let $\varepsilon > 0$. Then every weighted graph G has:*

1. *a deterministic all-pairs $+\varepsilon W(\cdot, \cdot)$ spanner with $O_\varepsilon(n)$ lightness, and*
2. *a deterministic all-pairs $+(4 + \varepsilon)W(\cdot, \cdot)$ spanner with $O_\varepsilon(n^{2/3})$ lightness.*

To the best of our knowledge, these are the first nontrivial lightness results known for additive spanners. For comparison on the first result, it is easy to show that every graph G has an all-pairs distance preserver H with $\text{lightness}(H) = O(n^2)$. It follows from the seminal work of Khuller, Raghavachari, and Young [25] on shallow-light trees that every graph G has a subgraph H that preserves distances up to a $(1 + \varepsilon)$ multiplicative factor, with lightness $O_\varepsilon(n)$. Our first result implies that the same lightness bound (up to the specifics of the ε dependence) can be achieved with *additive* error. Theorem 2.1 strictly

strengthens this consequence of [25], since local $+\varepsilon W(\cdot, \cdot)$ error implies multiplicative $(1 + \varepsilon)$ stretch as previously mentioned. Additionally, Theorem 2.1 is tight, in the sense that an unweighted complete graph K_n has lightness $\Theta(n)$ and no nontrivial $+\varepsilon W(\cdot, \cdot)$ spanner. Table 3 summarizes these lightness results compared with those in unweighted graphs.

Unweighted		Weighted		
$+\beta$	Lightness	$+\beta$	Lightness	Ref.
0	$O(n)$	$+\varepsilon W(\cdot, \cdot)$	$O_\varepsilon(n)$	[this paper]
+2	$O(n^{1/2})$?	
+4	$\tilde{O}(n^{2/5})$	$+(4 + \varepsilon)W(\cdot, \cdot)$	$O_\varepsilon(n^{2/3})$	[this paper]
+6	$O(n^{1/3})$?	

Table 3: Lightweight all-pairs additive spanners in unweighted and weighted graphs.

The lightweight spanners are based on a new initialization technique which we call *d-lightweight initialization*, in which an initial set of lightweight edges is added to the spanner starting from the MST of the input graph. Nearly all of the above spanners have a common theme in the construction and analysis: add an initial set of edges oblivious to the distances in the graph, then add shortest paths for any vertex pairs which do not satisfy inequality (1). The size or lightness bounds are then analyzed by determining how many pairs of nearby vertices there are whose distances sufficiently improve upon adding a shortest path; this method was also used by Elkin et al. [18]. Note that such improvements in unweighted graphs must be by at least 1; in weighted graphs, distances may improve by arbitrarily small amount which leads to the ε dependence. We leave as open questions the lightness bounds for $+2W$ and $+6W$ spanners (with or without ε dependence), whether the lightness for the $+(4 + \varepsilon)W(\cdot, \cdot)$ spanner can be improved to $\tilde{O}(n^{2/5})$, and whether we can construct additive spanners in weighted graphs which are simultaneously sparse *and* lightweight.

2 Preliminaries

Many additive spanner constructions begin with either a clustering [7, 14, 23] or initialization phase [2, 18, 26], where an initial set of edges is added to the spanner oblivious to distances or vertex pairs \mathcal{P} in the graph; after this phase, additional edges or paths are added so that the resulting subgraph is a valid spanner. Experimental results [4] suggest that initialization is preferred over clustering in terms of runtime and spanner size, and all constructions in this paper are initialization-based. For weighted graphs, a *d-light initialization* of G is a subgraph obtained by selecting the d lightest edges incident to every vertex, or all edges if the degree is less than d . We exploit the following lemma from [2]:

Lemma 1 ([2]). *Let H be a d -light initialization of a weighted graph G , and let $\pi(s, t)$ be a shortest path in G . If there are ℓ edges of $\pi(s, t)$ absent from H , then there is a set N of $\frac{d\ell}{6} = \Omega(d\ell)$ vertices, where each vertex in N is adjacent to a vertex on $\pi(s, t)$, connected via an edge of weight at most $W(s, t)$.*

We refer to vertices in N as the d -light neighbors of $\pi(s, t)$. The fact that d -light neighbors are connected to $\pi(s, t)$ via light edges of weight $\leq W(s, t)$ was not explicitly stated in [2] but follows directly from the proof, as N is constructed by taking the d lightest edges incident to vertices on $\pi(s, t)$ which are incident to a missing edge of weight at most $W(s, t)$. After d -light initialization, additional edges or paths are added to H in order to “satisfy” the remaining unsatisfied vertex pairs⁶. We will use the standard method of $+\beta$ spanner completion, where we iterate over each vertex pair (s, t) in nondecreasing order of maximum weight $W(s, t)$ (then by nondecreasing distance $d_G(s, t)$ in case of a tie) and add $\pi(s, t)$ to the spanner if (s, t) is unsatisfied.

3 Sparse local additive spanners

Instead of $+cW$ additive error considered in [2], we consider $+cW(\cdot, \cdot)$ local error.

3.1 Pairwise $+(2 + \varepsilon)W(\cdot, \cdot)$ and $+(6 + \varepsilon)W(\cdot, \cdot)$ spanners

For the pairwise $+(2 + \varepsilon)W(\cdot, \cdot)$ and $+(6 + \varepsilon)W(\cdot, \cdot)$ spanners (Theorem 1.1-2), we describe a deterministic construction. The analysis behind the edge bounds uses a set-off and improving strategy also used in [18]. The constructions involve one additional step of adding a certain number of edges along every vertex pair’s shortest path before spanner completion.

Let $\ell, d \geq 1$ be parameters which are defined later. Let H be a d -light initialization. Then for each vertex pair $(s, t) \in \mathcal{P}$, consider the shortest path $\pi(s, t)$ in G and add the first ℓ missing edges and the last ℓ missing edges to H (if $\pi(s, t)$ is missing at most 2ℓ edges, all missing edges from $\pi(s, t)$ are added to H). We remark that if (s, t) is already satisfied, we can skip this step for the pair (s, t) .

After this phase, we perform $+(2 + \varepsilon)W(\cdot, \cdot)$ or $+(6 + \varepsilon)W(\cdot, \cdot)$ spanner completion as follows: for each pair $(s, t) \in \mathcal{P}$ sorted in nondecreasing order of maximum weight $W(s, t)$, if (s, t) is still unsatisfied, add the remaining edges from the path $\pi(s, t)$ to H . This construction clearly outputs a valid pairwise spanner, and $O(nd + \ell|\mathcal{P}|)$ edges are added in the “distance-oblivious” phase before spanner completion. It remains to determine the number of edges added in spanner completion.

Let $(s, t) \in \mathcal{P}$ be a vertex pair for which $\pi(s, t)$ is added to H during spanner completion. Observe that after d -light initialization, the first ℓ missing edges and the last ℓ missing edges from $\pi(s, t)$ do not overlap; otherwise no remaining edges from $\pi(s, t)$ would have been added to H . Let $u_1v_1, \dots, u_\ell v_\ell$ denote the first ℓ missing edges on $\pi(s, t)$ which are added after d -light initialization, and

⁶ A vertex pair (s, t) is *satisfied* if the spanner inequality (1) holds for that pair.

let $u'_1 v'_1, \dots, u'_\ell v'_\ell$ denote the last ℓ missing edges, where u_i (or u'_i) is closer to s than v_i (or v'_i); see Fig. 1 for illustration.




Fig. 1: Illustration of Lemma 2 (left) and Lemma 3 (right) with $\ell = 2$. Note that $s = u_1$ and $t = v'_2$ in this example. By adding $\pi(s, t)$, at least one of the pairs' (a, x) or (x, b) distance improves by at least $\frac{\varepsilon W(s, t)}{2}$. Note that a, x, b are not necessarily distinct.

We will refer to the set $\{u_1, \dots, u_\ell\}$ as the *prefix* and the set $\{v'_1, \dots, v'_\ell\}$ as the *suffix*. Consider the shortest paths $\pi(s, v_\ell)$ and $\pi(u'_1, t)$. By Lemma 1, there are $\Omega(d\ell)$ d -light neighbors which are adjacent to a vertex in the prefix and suffix, respectively, connected by an edge of weight at most $W(s, t)$.

Consider the subpath $\pi(v_\ell, u'_1)$; suppose $z \geq 1$ edges of $\pi(s, t)$ are added during spanner completion. By Lemma 1 again, there are $\Omega(dz)$ d -light neighbors which are adjacent to a vertex on $\pi(v_\ell, u'_1)$. In the following lemmas, denote by H_0 and H_1 the spanner immediately before and after $\pi(s, t)$ is added, respectively.

Lemma 2. *Let $(s, t) \in \mathcal{P}$ be such that $\pi(s, t)$ is added to H during $+(2 + \varepsilon)W(\cdot, \cdot)$ spanner completion. Let a and b be vertices in the prefix and suffix respectively. Let x be a d -light neighbor of the path $\pi(v_\ell, u'_1)$. Then both of the following hold:*

1. $d_{H_1}(a, x) \leq d_G(a, x) + 2W(s, t)$ and $d_{H_1}(b, x) \leq d_G(b, x) + 2W(s, t)$
2. $d_{H_0}(a, x) - d_{H_1}(a, x) > \frac{\varepsilon W(s, t)}{2}$ or $d_{H_0}(b, x) - d_{H_1}(b, x) > \frac{\varepsilon W(s, t)}{2}$.

Proof. 1. We have by the triangle inequality

$$\begin{aligned} d_{H_1}(a, x) &\leq d_G(a, y) + w(xy) \leq d_G(a, y) + W(s, t) \\ &\leq [d_G(a, x) + w(xy)] + W(s, t) \\ &\leq d_G(a, x) + 2W(s, t) \end{aligned} \quad (3)$$

where we have used the fact that x is a d -light neighbor, so $w(xy) \leq W(s, t)$ by Lemma 1. Similarly, $d_{H_1}(b, x) \leq d_G(b, x) + 2W(s, t)$.

2. The vertices a and b are on the prefix and suffix respectively. Suppose otherwise $d_{H_0}(a, x) - d_{H_1}(a, x) \leq \frac{\varepsilon W(s, t)}{2}$ and $d_{H_0}(b, x) - d_{H_1}(b, x) \leq \frac{\varepsilon W(s, t)}{2}$. Consider a shortest s - t path in H_0 if it exists (if no such path exists, then $d_{H_0}(s, x) = \infty$ or $d_{H_0}(t, x) = \infty$, and 2. follows). Since a is on the prefix and

the first ℓ edges of $\pi(s, t)$ are added before completion, we necessarily have $d_{H_0}(s, a) = d_G(s, a)$ and symmetrically $d_{H_0}(b, t) = d_G(b, t)$. Then

$$\begin{aligned}
d_{H_0}(s, t) &\leq d_G(s, a) + d_{H_0}(a, x) + d_{H_0}(x, b) + d_G(b, t) \\
&\leq d_G(s, a) + \left[d_{H_1}(a, x) + \frac{\varepsilon W(s, t)}{2} \right] \\
&\quad + \left[d_{H_1}(x, b) + \frac{\varepsilon W(s, t)}{2} \right] + d_G(b, t) \\
&\stackrel{(3)}{\leq} d_G(s, a) + d_G(a, y) + d_G(y, b) + d_G(b, t) + (2 + \varepsilon)W(s, t) \\
&= d_G(s, t) + (2 + \varepsilon)W(s, t)
\end{aligned}$$

contradicting that $\pi(s, t)$ was added during $+(2 + \varepsilon)W(\cdot, \cdot)$ spanner completion. \square

For the pairwise $+(6 + \varepsilon)W(\cdot, \cdot)$ spanner, we consider arbitrary d -light neighbors a and b of the prefix and suffix, and similarly consider vertex pairs (a, x) , (b, x) whose distances sufficiently improve:

Lemma 3. *Let $(s, t) \in \mathcal{P}$ be such that $\pi(s, t)$ is added to H during $+(6 + \varepsilon)W(\cdot, \cdot)$ spanner completion. Let a and b be d -light neighbors adjacent to vertices u_i and v'_j in the prefix and suffix, respectively. Let x be a d -light neighbor of the path $\pi(v_\ell, u'_1)$. Then both of the following hold:*

1. $d_{H_1}(a, x) \leq d_G(a, x) + 4W(s, t)$ and $d_{H_1}(x, b) \leq d_G(x, b) + 4W(s, t)$
2. $d_{H_0}(a, x) - d_{H_1}(a, x) > \frac{\varepsilon W(s, t)}{2}$ or $d_{H_0}(x, b) - d_{H_1}(x, b) > \frac{\varepsilon W(s, t)}{2}$.

Proof. 1. Similar to Lemma 2, we have by the triangle inequality

$$\begin{aligned}
d_{H_1}(a, x) &\leq w(au_i) + d_G(u_i, y) + w(xy) \\
&\leq d_G(u_i, y) + 2W(s, t) \\
&\leq [w(u_i a) + d_G(a, x) + w(xy)] + 2W(s, t) \\
&\leq d_G(a, x) + 4W(s, t)
\end{aligned} \tag{4}$$

where we have used the fact that a and x are d -light neighbors, so $w(u_i a) \leq W(s, t)$ and $w(xy) \leq W(s, t)$. Similarly, $d_{H_1}(x, b) \leq d_G(x, b) + 4W(s, t)$.

2. Suppose otherwise $d_{H_0}(a, x) - d_{H_1}(a, x) \leq \frac{\varepsilon W(s, t)}{2}$ and $d_{H_0}(x, b) - d_{H_1}(x, b) \leq \frac{\varepsilon W(s, t)}{2}$. Consider a shortest s - t path in H_0 if it exists (if no such path exists,

then $d_{H_0}(s, x) = \infty$ or $d_{H_0}(t, x) = \infty$ and 2. follows). Then

$$\begin{aligned}
 d_{H_0}(s, t) &\leq d_G(s, u_i) + w(u_i a) + d_{H_0}(a, x) + d_{H_0}(x, b) + w(bv'_i) + d_G(v'_i t) \\
 &\leq d_G(s, u_i) + W(s, t) + \left[d_{H_1}(a, x) + \frac{\varepsilon W(s, t)}{2} \right] \\
 &\quad + \left[d_{H_1}(x, b) + \frac{\varepsilon W(s, t)}{2} \right] + W(s, t) + d_G(v'_i t) \\
 &\leq d_G(s, u_i) + d_{H_1}(a, x) + d_{H_1}(x, b) + d_G(v'_i, t) + (2 + \varepsilon)W(s, t) \\
 &\stackrel{(4)}{\leq} d_G(s, u_i) + [d_G(u_i, y) + 2W(s, t)] \\
 &\quad + [d_G(y, v'_i) + 2W(s, t)] + d_G(v'_i, t) + (2 + \varepsilon)W(s, t) \\
 &= d_G(s, t) + (6 + \varepsilon)W(s, t)
 \end{aligned}$$

contradicting that $\pi(s, t)$ was added during $+(6 + \varepsilon)W(\cdot, \cdot)$ spanner completion. \square

Lemma 4. *By setting $d = |\mathcal{P}|^{1/3}$ and $\ell = n/|\mathcal{P}|^{2/3}$, the pairwise $+(2 + \varepsilon)W(\cdot, \cdot)$ construction outputs a subgraph H with $|E(H)| = O_\varepsilon(n|\mathcal{P}|^{1/3})$.*

Proof. In the distance-oblivious phase, we add $O(nd + \ell|\mathcal{P}|) = O(n|\mathcal{P}|^{1/3})$ edges to H . A vertex pair (v, x) is *set-off* if it is the first time that $d_H(v, x) \leq d_G(v, x) + (2 + \varepsilon)W(s, t)$ and is *improved* if its distance in H decreases by at least $\frac{\varepsilon W(s, t)}{2}$. Suppose adding $\pi(s, t)$ during $+(2 + \varepsilon)W(\cdot, \cdot)$ spanner completion adds $z \geq 1$ additional edges. Let x be a d -light neighbor of $\pi(v_\ell, u'_1)$ and let a, b be vertices in the prefix and suffix. By Lemma 1, there are $\Omega(dz)$ vertices x adjacent to $\pi(v_\ell u'_1)$. By Lemma 2, both of the pairs (a, x) , (b, x) are set-off if not already, and at least one of the pairs is improved upon adding $\pi(s, t)$. Since there are $\Omega(\ell)$ choices for a or b , this gives $\Omega(dz \times \ell) = \Omega(nz/|\mathcal{P}|^{1/3})$ improvements upon adding z edges to H .

Once a pair (v, x) is set-off, it can only be improved $O(\frac{1}{\varepsilon})$ times; this follows since pairs are ordered by their maximum weight, so any improvement is by at least $\frac{\varepsilon W(s, t)}{2}$. If Z total edges are added during spanner completion, then the number of improvements is $\Omega(d\ell Z)$. There are $O(n^2)$ vertex pairs and once set-off, each vertex pair is improved $O(\frac{1}{\varepsilon})$ times, in which we have $\Omega(d\ell Z) = O(\frac{n^2}{\varepsilon})$. Since $d = |\mathcal{P}|^{1/3}$ and $\ell = n/|\mathcal{P}|^{2/3}$, we obtain $Z = O(\frac{1}{\varepsilon}n|\mathcal{P}|^{1/3})$. Altogether we obtain $|E(H)| = O(\frac{1}{\varepsilon}n|\mathcal{P}|^{1/3})$. \square

Lemma 5. *By setting $d = |\mathcal{P}|^{1/4}$ and $\ell = n/|\mathcal{P}|^{3/4}$, the pairwise $+(6 + \varepsilon)W(\cdot, \cdot)$ construction outputs a subgraph H with $|E(H)| = O_\varepsilon(n|\mathcal{P}|^{1/4})$.*

Proof. This is the same as Lemma 4 except that a and b are d -light neighbors of the subpaths containing the first and last ℓ missing edges. Then if $z \geq 1$ edges in $\pi(s, t)$ are added during $+(6 + \varepsilon)W(\cdot, \cdot)$ spanner completion, we obtain $\Omega(dz \times d\ell) = \Omega(d^2\ell z) = \Omega(nz/|\mathcal{P}|^{1/4})$ improvements upon adding z edges to H .

If Z total edges are added during $+(6 + \varepsilon)W(\cdot, \cdot)$ spanner completion, then similar reasoning gives

$$\Omega(d^2 \ell Z) = O\left(\frac{n^2}{\varepsilon}\right)$$

in which $Z = O(\frac{1}{\varepsilon}n|\mathcal{P}|^{1/4})$ and $|E(H)| = O(\frac{1}{\varepsilon}n|\mathcal{P}|^{1/4})$. \square

Proof (Lemma 6). Since x is a d -light neighbor of $\pi(s, t)$, we have $w(xy) \leq W(s, t)$. Using the triangle inequality:

$$\begin{aligned} d_G(s, x) + d_G(x, t) &\leq [d_G(s, y) + w(xy)] + [w(xy) + d_G(y, t)] \\ &\leq d_G(s, t) + 2w(xy) \\ &\leq d_G(s, t) + 2W(s, t) \end{aligned}$$

completing the proof. \square

Lemmas 4 and 5 imply Theorem 1.1 and 1.2 respectively. Further, by setting $W = 1$ and $\varepsilon = 0.5$, these results imply pairwise $+2$ and $+6$ spanners of size $O(n|\mathcal{P}|^{1/3})$ and $O(n|\mathcal{P}|^{1/4})$ in *unweighted* graphs (as a $+2.5$ spanner of an unweighted graph is also a $+2$ spanner). These edge bounds match those of existing pairwise $+2$ spanners [10, 24] and pairwise $+6$ spanners [23].

3.2 Pairwise $+2W(\cdot, \cdot)$ and $+4W(\cdot, \cdot)$ spanners

We prove Theorem 1.3 and 1.4: every weighted graph has a pairwise $+2W(\cdot, \cdot)$ spanner and $+4W(\cdot, \cdot)$ spanner on $O(n|\mathcal{P}|^{1/3})$ edges and $O(n|\mathcal{P}|^{2/7})$ edges respectively. This removes the ε dependence from Theorem 1.1 and uses local error instead of global W error as in [2]. First, we need the following simple lemma:

Lemma 6. *Let H be a d -light initialization, and let $s, t \in V$. Let x be a d -light neighbor of $\pi(s, t)$ connected to a vertex $y \in \pi(s, t)$ in H . Consider a shortest path tree in G rooted at x . Then the distance from s to t in this tree is at most $d_G(s, t) + 2W(s, t)$.*

Proof. Since x is a d -light neighbor of $\pi(s, t)$, we have $w(xy) \leq W(s, t)$. Using the triangle inequality:

$$\begin{aligned} d_G(s, x) + d_G(x, t) &\leq [d_G(s, y) + w(xy)] + [w(xy) + d_G(y, t)] \\ &\leq d_G(s, t) + 2w(xy) \\ &\leq d_G(s, t) + 2W(s, t) \end{aligned}$$

completing the proof. \square

Proof (Theorem 1.3). The main idea is to first show the existence of a pairwise $+2W(\cdot, \cdot)$ spanner with slack on $O(n|\mathcal{P}|^{1/3})$ which satisfies at least a constant fraction of the vertex pairs. Then, we apply a lemma from [8] (Lemma 7) to show the existence of a pairwise spanner. Let $\ell, d \geq 1$ be parameters, and let H be a d -light initialization ($O(nd)$ edges). We iterate over each pair $(s, t) \in \mathcal{P}$.

If $d_H(s, t) \leq d_G(s, t) + 2W(s, t)$, then (s, t) is already satisfied and we can do nothing. If $d_H(s, t) > d_G(s, t) + 2W(s, t)$ and $\pi(s, t)$ has fewer than ℓ missing edges in H , add all missing edges in $\pi(s, t)$ to H ($O(\ell|\mathcal{P}|)$ edges). Otherwise (s, t) is unsatisfied and $\pi(s, t)$ is missing more than ℓ edges. In this case, we do not add any edges yet.

After iterating over each vertex pair, consider the following randomized process: for each vertex in V , select it with probability $\frac{6}{d\ell}$. If more than $\frac{6n}{d\ell}$ vertices are selected, repeat this process. For each selected vertex, add a shortest path tree rooted at that vertex to H . Note that at most $\frac{6n}{d\ell}$ vertices are selected, so at most $\frac{6n(n-1)}{d\ell} = O\left(\frac{n^2}{d\ell}\right)$ edges are added in this process.

By Lemma 1, every unsatisfied pair (s, t) has at least $\frac{d\ell}{6}$ d -light neighbors. Then with probability at least $1 - \left(1 - \frac{6}{d\ell}\right)^{d\ell/6}$, we will have selected at least one d -light neighbor of $\pi(s, t)$. This probability is at least a constant, since $1 - \frac{1}{e} \leq 1 - \left(1 - \frac{1}{x}\right)^x \leq 1$ for all $x \geq 1$. By Lemma 6, adding a shortest path tree at any d -light neighbor of $\pi(s, t)$ satisfies the pair (s, t) .

Under this randomized process, every unsatisfied vertex pair is satisfied with at least constant probability α , so by linearity of expectation, the expected number of satisfied vertex pairs is at least $\alpha|\mathcal{P}|$. Then there necessarily exists a subset of $\frac{6n}{d\ell}$ vertices such that adding a shortest path tree rooted at each vertex satisfies $\alpha|\mathcal{P}|$ vertex pairs.

The number of edges in this spanner with slack is at most $k\left(nd + \ell|\mathcal{P}| + \frac{n^2}{\ell d}\right)$ for some constant k ; set $\ell := \frac{n}{|\mathcal{P}|^{2/3}}$ and $d := |\mathcal{P}|^{1/3}$ to obtain the desired edge bound of $kn|\mathcal{P}|^{1/3} = O(n|\mathcal{P}|^{1/3})$. This only implies the existence of a $+2W(\cdot, \cdot)$ spanner which satisfies a constant fraction α of vertex pairs; to finish, we use the following lemma:

Lemma 7 ([8]). *Let $a, b, c, k, \alpha > 0$ be constants (independent of n or $|\mathcal{P}|$) with $\alpha \leq 1$, let G be an n -vertex graph, and let p^* be a parameter. Suppose that:*

- *There is an algorithm which, on input G, \mathcal{P} with $|\mathcal{P}| \leq p^*$, returns a pairwise spanner on $O(n^a)$ edges*
- *There is an algorithm which, on input G, \mathcal{P} with $|\mathcal{P}| \geq p^*$, returns a spanner with slack on $\leq kn^b|\mathcal{P}|^c$ edges which satisfies at least $\alpha|\mathcal{P}|$ vertex pairs*

Then there is a pairwise spanner of G, \mathcal{P} on $O(n^a + n^b|\mathcal{P}|^c)$ edges.

Proof (Lemma 7). While the number of unsatisfied vertex pairs $|\mathcal{P}|$ is greater than p^* , compute a spanner with slack on $\leq kn^b|\mathcal{P}|^c$ edges, then remove the satisfied vertex pairs from \mathcal{P} . Repeat this process until $|\mathcal{P}| \leq p^*$, then compute a pairwise spanner on $O(n^a)$ edges. The union of these computed spanners is a pairwise spanner of G over the initial set of vertex pairs \mathcal{P} .

Let $p_1 = |\mathcal{P}|$ denote the number of vertex pairs initially, and for $i \geq 1$, let p_i denote the number of unsatisfied vertex pairs immediately before the i^{th} iteration. Since each spanner with slack satisfies at least a constant fraction α of the unsatisfied vertex pairs, we have $p_i \leq (1 - \alpha)p_{i-1}$, which implies $p_i \leq$

$p_1(1 - \alpha)^{i-1}$ for all $i \geq 1$. Let $H_i = (V_i, E_i)$ denote the spanner with slack computed on the i^{th} iteration, so that $|E_i| \leq kn^b p_i^c$. By a simple union bound, we have

$$\begin{aligned} \left| \bigcup_{i \geq 1} E_i \right| &\leq \sum_{i \geq 1} |E_i| \leq \sum_{i \geq 1} kn^b p_i^c \\ &\leq kn^b \sum_{i \geq 1} (p_1(1 - \alpha)^{i-1})^c \\ &= kn^b p_1^c \sum_{i \geq 1} ((1 - \alpha)^c)^{i-1} \\ &= O(n^b |\mathcal{P}|^c) \end{aligned}$$

since $\sum_{i \geq 1} ((1 - \alpha)^c)^{i-1}$ is a geometric series with common ratio $(1 - \alpha)^c < 1$. The final pairwise spanner on $\leq p^*$ pairs adds $O(n^a)$ edges, completing the proof. \square

Lemma 7 implies the existence of a pairwise $+2W(\cdot, \cdot)$ spanner on $O(n|\mathcal{P}|^{1/3})$ edges. \square

We now show that every weighted graph G has a pairwise $+4W(\cdot, \cdot)$ spanner on $O(n|\mathcal{P}|^{2/7})$ edges.

Lemma 8. *Let H be a d -light initialization and let $s, t \in V$. Let v_1 and v_2 be two different d -light neighbors of $\pi(s, t)$, connected to two (not necessarily distinct) vertices $u_1, u_2 \in \pi(s, t)$ respectively. Consider adding the missing edges along the shortest paths from s - u_1 , v_1 - v_2 , and u_2 - t to H . Then $d_H(s, t) \leq d_G(s, t) + 4W(s, t)$.*

Proof. See Fig. 2 for illustration. Since v_1 and v_2 are d -light neighbors of u_1, u_2 respectively, we have $w(u_1 v_1) \leq W(s, t)$ and $w(u_2 v_2) \leq W(s, t)$. By triangle inequality, we have

$$\begin{aligned} d_H(s, t) &\leq d_H(s, u_1) + w(u_1 v_1) + d_H(v_1, v_2) + w(u_2 v_2) + d_H(u_2, t) \\ &\leq d_G(s, u_1) + W(s, t) + d_G(v_1, v_2) + W(s, t) + d_G(u_2, t) \\ &= d_G(s, u_1) + d_G(v_1, v_2) + d_G(u_2, t) + 2W(s, t) \end{aligned}$$

Again by triangle inequality, $d_G(v_1, v_2) \leq w(u_1 v_1) + d_G(u_1, u_2) + w(u_2 v_2) \leq d_G(u_1, u_2) + 2W(s, t)$; hence we have

$$\begin{aligned} &\leq d_G(s, u_1) + d_G(u_1, u_2) + d_G(u_2, t) + 4W(s, t) \\ &= d_G(s, t) + 4W(s, t) \end{aligned}$$

as desired. \square

Proof (Theorem 1.4). As in Theorem 1.3, we use two parameters $d, \ell \geq 1$. Let H be a d -light initialization. We iterate over each pair $(s, t) \in \mathcal{P}$. If (s, t) is already satisfied, do nothing; otherwise, if $\pi(s, t)$ is missing at most ℓ edges in H , add the shortest path $\pi(s, t)$ to H . This adds at most $nd + \ell p$ edges to H .

Construct a random sample R_1 of vertices by selecting each vertex with probability $\frac{6d}{n}$; repeat if more than $6d$ vertices were selected. Add a shortest path tree rooted at each $v \in R_1$. We claim that if (s, t) is unsatisfied and there are at least $\frac{n}{d^2}$ missing edges on $\pi(s, t)$, then with at least a constant probability, we will have selected a d -light neighbor of $\pi(s, t)$. Since $\pi(s, t)$ has at least $\frac{n}{d^2}$ missing edges, Lemma 1 implies there are at least $\frac{n}{6d}$ d -light neighbors of $\pi(s, t)$. Using a similar reasoning to Theorem 1.3, with at least a constant probability, we will have selected at least one d -light neighbor for R_1 . By Lemma 6, adding the shortest path tree rooted at this neighbor satisfies the pair (s, t) . This sample R_1 adds $O(dn)$ edges to H .

Lastly, there may be vertex pairs (s, t) such that (s, t) is unsatisfied, and $\pi(s, t)$ is missing greater ℓ but fewer than $\frac{n}{d^2}$ edges. For each such pair, add the first ℓ missing edges and the last ℓ missing edges to H ($O(\ell p)$ edges). Let the *prefix* denote the set of vertices on $\pi(s, t)$ which are incident to one of the first ℓ missing edges which were added to H ; define the *suffix* similarly.




Fig. 2: Illustration of the $+4W(\cdot, \cdot)$ randomized construction, with edges incident to the prefix and suffix bolded ($\ell = 2$ in this example).

Construct a second random sample R_2 by sampling each vertex with probability $\frac{6}{d\ell}$; repeat if more than $\frac{6n}{d\ell}$ vertices are sampled. Then, for every two vertices $v_1, v_2 \in R_2$, add all edges among the shortest v_1 - v_2 path among all paths which are missing at most $\frac{n}{d^2}$ edges; this adds $O\left(\left(\frac{n}{d\ell}\right)^2 \cdot \frac{n}{d^2}\right) = O\left(\frac{n^3}{\ell^2 d^4}\right)$ edges. By Lemma 1, there are $\frac{d\ell}{6} = \Omega(d\ell)$ d -light neighbors of the prefix, and $\Omega(d\ell)$ d -light neighbors of the suffix. Then with at least a constant probability, we will have sampled in R_2 at least one d -light neighbor v_1 of the prefix and one d -light neighbor v_2 of the suffix of $\pi(s, t)$. By adding a shortest path from v_1 to v_2 containing at most $\frac{n}{d^2}$ edges along with the first and last ℓ edges of $\pi(s, t)$, this satisfies the pair (s, t) by Lemma 8.

The number of edges in H is $|E(H)| = O(nd + |\mathcal{P}|\ell + n^3/(\ell^2 d^4))$. Set $\ell := n/|\mathcal{P}|^{5/7}$ and $d := |\mathcal{P}|^{2/7}$ to obtain the desired edge bound of $O(n|\mathcal{P}|^{2/7})$. As in Theorem 1.3, this only ensures the existence of a $+4W(\cdot, \cdot)$ spanner with slack; apply Lemma 7 to prove Theorem 1.4. \square

3.3 All-pairs $+cW(\cdot, \cdot)$ spanners

We now discuss the all-pairs setting, to prove Theorem 1.5. Our construction is essentially the same as the pairwise construction given above, but the analysis is slightly different. The parameter ℓ no longer plays a role, and we skip the step where we consider pairs (s, t) such that $\pi(s, t)$ is missing at most ℓ edges. But the rest of the construction is the same:

- We start the spanner H as a d -light initialization,
- We take a random sample R_1 by selecting each vertex with probability $\frac{6d}{n}$ and add a shortest path tree rooted at each sampled vertex, and then
- We take a random sample R_2 by including each vertex with probability $\frac{6}{d}$. For any two nodes $s, t \in R_2$, let $\pi(s, t)$ be the shortest $s \rightsquigarrow t$ path among all paths that are missing $\leq \frac{n}{d^2}$ edges in the initialization (if no such path exists, then ignore the pair s, t). Then we add all edges in $\pi(s, t)$ to the spanner.

The total number of edges in this construction is $O(nd + \frac{n^3}{d^4})$. Setting $d = n^{2/5}$ gives an edge bound of $O(n^{7/5})$. Finally, by the same analysis as above, each demand pair is satisfied with constant probability. So if we repeat the construction $C \log n$ times for a large enough constant C , and we union together the $O(\log n)$ spanners computed in each round, then with high probability all demand pairs are satisfied. The final spanner thus has $O(n^{7/5} \cdot \log n) = \tilde{O}(n^{7/5})$ edges in total.

4 Lightweight local additive spanners

In this section, we prove Theorem 2 by constructing *lightweight* additive spanners. The previous spanner algorithms based on d -light initialization can produce spanners with poor lightness: Let G be the graph obtained by taking $K_{\frac{n}{2}, \frac{n}{2}}$ with all edges of weight W , then adding two paths of weight 0 connecting the vertices within each bipartition. Then $w(\text{MST}(G)) = W$, while d -light initialization already adds $\Omega(Wnd)$ weight to the spanner H , or $\Omega(nd)$ lightness.

In order to construct lightweight additive spanners, we introduce a new initialization technique called *d -lightweight initialization*, which adds edges of total weight at most d for each vertex, starting from the MST of G . We first perform the following simple modifications to the input graph G in order:

1. Scale the edge weights of G linearly so that the weight of the MST is $\frac{n}{2}$.
2. Remove all edges of weight $\geq n$ from G .

Note that step 1 also scales $W(s, t)$ for vertex pairs (s, t) , but the validity of an additive spanner or spanner path is invariant to scaling. Step 2 does not change the shortest path metric or the maximum edge weights $W(s, t)$; since the MST has weight $\frac{n}{2}$, this implies $W(s, t) \leq d_G(s, t) \leq \frac{n}{2}$.

Given G and $d \geq 0$, the *d -lightweight initialization* is a subgraph H defined as follows: let H be an MST of G . If $d > 0$, then for each vertex $v \in V$,

consider all incident edges to v which have not already been added to H , in nondecreasing order of weight. Add these incident edges one by one to H until the next edge causes the total edge weight added corresponding to v to be greater than d . This subgraph H has $O(n + nd)$ total edge weight, and $O(1 + d)$ lightness. Note that 0-lightweight initialization is simply the MST of G . After d -lightweight initialization, perform $+cW(\cdot, \cdot)$ spanner completion by iterating over each vertex pair (s, t) in nondecreasing order of the maximum weight $W(s, t)$ (then by nondecreasing distance $d_G(s, t)$ if two pairs have the same maximum weight), and adding $\pi(s, t)$ to H if (s, t) is unsatisfied.

4.1 Lightweight initialization and neighborhoods of shortest paths

In order to prove the desired lightness bounds, we consider a subdivided graph obtained as follows: for each MST edge e of $\text{MST}(G)$, subdivide e into $\lceil w(e) \rceil$ edges of weight $\frac{w(e)}{\lceil w(e) \rceil}$. Recall that $w(\text{MST}(G)) = \frac{n}{2}$, so this adds $O(n)$ vertices, all of which are on MST edges. Let G' denote this subdivided graph. We do not modify the maximum edge weights $W(s, t)$ for each pair $(s, t) \in V(G)$, even if such edges are subdivided in G' . We remark that the MST of G' also has weight $\frac{n}{2}$, and all MST edges have weight at most 1.

Lemma 9. *Let $e = uv$ be an edge in $E(G')$ which is not contained in the d -lightweight initialization H . Then there are $\Omega(d^{1/2})$ vertices x in G' such that $d_H(u, x) \leq w(e)$. Moreover, if $\varepsilon' \in (0, 1]$, there are $\Omega(\varepsilon'w(e))$ vertices x such that $d_H(u, x) \leq \varepsilon'w(e)$.*

Proof. If $w(e) \geq d^{1/2}$, then consider the set N of vertices (including u) such that each $x \in N$ is connected to u by a path of $\leq \lfloor w(e) \rfloor$ MST edges. Since all MST edges have weight ≤ 1 , all vertices $x \in N$ satisfy $d_H(u, x) \leq w(e)$. Then $|N| = \max\{1, \lfloor w(e) \rfloor\} \geq \frac{w(e)}{2} = \Omega(d^{1/2})$, proving this case. Note that $w(e) \leq n$ by step 2 above. Moreover given $\varepsilon' \in (0, 1]$, consider vertices x which are at most $\lfloor \varepsilon'w(e) \rfloor$ MST edges away from u . There are $\max\{1, \lfloor \varepsilon'w(e) \rfloor\} \geq \frac{\varepsilon'w(e)}{2} = \Omega(\varepsilon'w(e))$ such vertices, all of distance $\leq \varepsilon'w(e)$ from u .

If $w(e) < d^{1/2}$, then every edge added when u is considered in the d -lightweight initialization has weight less than $d^{1/2}$, since edges are added in nondecreasing weight. We thus add at least $d^{1/2} - 1$ edges corresponding to u . By construction, each such edge is lighter than e , so for any vertex x that is the endpoint of one of these edges ux , we have $d_H(u, x) \leq w(ux) \leq w(e)$. Lastly by considering $x = u$, we obtain $d^{1/2}$ vertices x in G' . \square

We now prove a lightweight analogue of Lemma 1:

Lemma 10. *Let H be a d -lightweight initialization, let $\pi(s, t)$ be a shortest s - t path in G' , let z be the total weight of the edges in $\pi(s, t)$ which are absent in H , and let $\varepsilon' \in (0, 1]$. Then there is a set N of at least $\frac{\varepsilon'z}{10} = \Omega(\varepsilon'z)$ vertices in G' which are of distance at most $\varepsilon'W(s, t)$ in H from some vertex in $\pi(s, t)$.*

Proof. Let $M_0 = \{e_1, \dots, e_\ell\}$ be the set of ℓ edges in $\pi(s, t)$ which are absent in H , where z is the total weight of these edges. We denote the endpoints of e_i by u_i and v_i where u_i is closer to s . We will first construct a subset $M_1 \subseteq M_0$ of edges which satisfies two properties: i) the neighborhoods of nearby MST vertices from each edge in M_1 are pairwise disjoint, and ii) the total edge weight in M_1 is $\Omega(z)$.

Let $M_1 := \emptyset$. While M_0 is not empty, select an edge $e_i \in M_0$ of maximum weight and add it to M_1 . Then remove e_i from M_0 , as well as all other edges e_j such that $d_{G'}(u_i, u_j) \leq 2\varepsilon'w(e_i)$ from M_0 . When M_0 is empty, we will be left with a set M_1 of edges which are relatively far apart. For each edge $e_i = u_i v_i \in M_1$, add to N all vertices connected to u_i by a path of $\leq \varepsilon'w(e_i)$ MST edges (including u_i itself). Let N_i denote the set of vertices added to N corresponding to edge e_i . Since MST edges in G' have weight at most 1, we have $|N_i| \geq \frac{\varepsilon'w(e_i)}{2}$ by Lemma 9, and $d_H(u_i, x) \leq \varepsilon'w(e_i) \leq \varepsilon'W(s, t)$ for each $x \in N_i$.




Fig. 3: Illustration of Lemma 10 where MST edges are bolded and $\varepsilon' = \frac{1}{2}$. By selecting $e_i = u_i v_i$ of weight 5.2 and adding e_i to M_1 , the five MST vertices which are connected by at most $\lfloor \frac{1}{2} \times 5.2 \rfloor = 2$ MST edges are added to N , and $|N_i| = 5$. Note that $u_{i+1} = v_i$ in this example since e_i and e_{i+1} are adjacent.

We claim that the sets $N_i : e_i \in M_1$ are pairwise disjoint. Suppose otherwise there exists some vertex $x \in N_i \cap N_j$, and w.l.o.g. edge e_i was selected before edge e_j so that $w(e_i) \geq w(e_j)$. Since $x \in N_i$, we have $d_H(u_i, x) \leq \varepsilon'w(e_i)$ and similarly $d_H(u_j, x) \leq \varepsilon'w(e_j)$. Then by the triangle inequality:

$$\begin{aligned} d_{G'}(u_i, u_j) &\leq d_H(u_i, u_j) \leq d_H(u_i, x) + d_H(x, u_j) \\ &\leq \varepsilon'w(e_i) + \varepsilon'w(e_j) \leq 2\varepsilon'w(e_i). \end{aligned}$$

However this is a contradiction: once edge e_i is selected for M_1 , edge e_j is also removed from M_0 since the distance from u_i to u_j is at most $2\varepsilon'w(e_i)$, so e_j could not have been selected for M_1 .

We also claim that the total edge weight in M_1 is $\Omega(z)$. Whenever an edge e_i is selected for M_1 , we remove e_i and any nearby edges e_j with $d_{G'}(u_i, u_j) \leq 2\varepsilon'w(e_i)$ from M_0 . Since all nearby edges by definition are on $\pi(s, t)$, the total weight of the edges removed from M_0 when e_i is selected is at most $(1+4\varepsilon')w(e_i)$. Then the total edge weight in M_1 is at least $\frac{z}{1+4\varepsilon'} \geq \frac{z}{5} = \Omega(z)$, since $0 < \varepsilon' \leq 1$.

Let $N := \bigcup_{e_i \in M_1} N_i$. As shown previously, each N_i contains at least $\frac{\varepsilon' w(e_i)}{2}$ vertices. By the above claims, the neighborhoods N_i are pairwise disjoint and the sum of edge weights in M_1 is at least $\frac{z}{5}$, so we obtain

$$|N| = \sum_{e_i \in M_1} |N_i| \geq \sum_{e_i \in M_1} \frac{\varepsilon' w(e_i)}{2} \geq \frac{\varepsilon'}{2} \cdot \frac{z}{5} = \Omega(\varepsilon' z). \quad \square$$

4.2 Proof of Theorem 2

Using Lemmas 9 and 10, we can analyze the lightness of the above spanner constructions by considering the number of vertex pairs in the subdivided graph G' whose distances sufficiently improve upon adding $\pi(s, t)$, where $s, t \in V(G)$.

Lemma 11. *Let $(s, t) \in V(G)$ be a vertex pair such that $\pi(s, t)$ is added during $+\varepsilon W(\cdot, \cdot)$ spanner completion. Let $x \in N$ be a vertex in G' which is of distance $\leq \frac{\varepsilon}{4} W(s, t)$ from some vertex u in $\pi(s, t)$ in H . Let H_0 and H_1 denote the spanner before and after $\pi(s, t)$ is added. Then both of the following hold:*

1. $d_{H_1}(s, x) \leq d_{G'}(s, x) + \frac{\varepsilon}{2} W(s, t)$ and $d_{H_1}(t, x) \leq d_{G'}(t, x) + \frac{\varepsilon}{2} W(s, t)$
2. $d_{H_0}(s, x) - d_{H_1}(s, x) \geq \frac{\varepsilon}{4} W(s, t)$ or $d_{H_0}(t, x) - d_{H_1}(t, x) \geq \frac{\varepsilon}{4} W(s, t)$.

Similar to Lemma 2, statement 1. holds by the triangle inequality, and 2. can be proven by contradiction: if neither inequality was true, then $d_{H_0}(s, t) \leq d_G(s, t) + \varepsilon W(s, t)$, contradicting that $\pi(s, t)$ was added during spanner completion.




Fig. 4: Illustration of Lemma 11. By adding $\pi(s, t)$ to H during $+\varepsilon W(\cdot, \cdot)$ spanner completion, both pairs (s, x) and (t, x) are satisfied, and at least one of the pairs' distances improves by at least $\frac{\varepsilon}{4} W(u, v)$.

Proof (Theorem 2.1). By Lemma 10 with $\varepsilon' = \frac{\varepsilon}{4}$, there are $\Omega(\varepsilon z)$ vertices $x \in V(G')$ which are of distance $\leq \frac{\varepsilon}{4} W(s, t)$ from some vertex in $\pi(s, t)$, so adding path $\pi(s, t)$ of weight z improves $\Omega(\varepsilon z)$ vertex pairs. If Z is the total weight added during $+\varepsilon W(\cdot, \cdot)$ spanner completion, then there are $\Omega(\varepsilon Z)$ improvements.

Once a vertex pair is set-off, it is only improved a *constant* number of times (since any such pair (s, x) or (t, x) , once set-off, has error $\frac{\varepsilon}{2} W(s, t)$, and any improvement is by $\Omega(\varepsilon W(s, t))$). Then by considering the number of improvements, we obtain $\Omega(\varepsilon Z) = O(n^2) \implies Z = O(\frac{1}{\varepsilon} n^2)$. This result does not depend on d -lightweight initialization, so set $d = 0$. The total weight of the spanner H is $O(\frac{1}{\varepsilon} n^2)$; since $w(\text{MST}(G)) = \frac{n}{2}$, we obtain $\text{lightness}(H) = O_\varepsilon(n)$ as desired. \square

For the $+(4 + \varepsilon)W(\cdot, \cdot)$ spanner, the corresponding lemma is as follows:

Lemma 12. *Let $(u, v) \in V(G)$ be a vertex pair such that $\pi(u, v)$ is added during $+(4 + \varepsilon)W(\cdot, \cdot)$ spanner completion. Let a and b be d -lightweight neighbors of s and t in G' , respectively, such that $d_{H_0}(x, a) \leq W(s, t)$ and $d_{H_0}(y, b) \leq W(s, t)$. Let $x \in N$ be a vertex in G' which is of distance $\leq \frac{\varepsilon}{4}W(s, t)$ from some vertex y in $\pi(s, t)$. Then both of the following hold:*

1. $d_{H_1}(a, x) \leq d_{G'}(a, x) + (2 + \varepsilon)W(s, t)$ and $d_{H_1}(b, x) \leq d_{G'}(b, x) + (2 + \varepsilon)W(s, t)$
2. $d_{H_0}(a, x) - d_{H_1}(a, x) \geq \frac{\varepsilon}{4}W(s, t)$ or $d_{H_0}(b, x) - d_{H_1}(b, x) \geq \frac{\varepsilon}{4}W(s, t)$.

Again, this is proved using the same methods as in Lemmas 2, 3, and 11; see Fig. 5.




Fig. 5: Illustration of Lemma 12. By adding $\pi(s, t)$ to H during $+(4 + \varepsilon)W(\cdot, \cdot)$ spanner completion, both pairs (a, x) and (b, x) are satisfied, and at least one of the pairs' distances improves by at least $\frac{\varepsilon}{4}W(s, t)$.

Proof (Theorem 2.2). Let (s, t) be a vertex pair for which $\pi(s, t)$ is added during $+(4 + \varepsilon)W(\cdot, \cdot)$ spanner completion. Again, by Lemma 10 with $\varepsilon' = \frac{\varepsilon}{4}$, there are $\Omega(\varepsilon z)$ choices for x which are of distance $\leq \frac{\varepsilon}{4}W(s, t)$ from some vertex in $\pi(u, v)$.

Similar to [18], we observe that the first edge (say st_1) in $\pi(s, t)$ (starting from s) is absent in H_0 immediately before $\pi(s, t)$ is added. Suppose otherwise $st_1 \in E(H_0)$, then consider the pair (t_1, t) . Since $+(4 + \varepsilon)W(\cdot, \cdot)$ spanner completion considers all vertex pairs in nondecreasing $W(s, t)$ and then by distance $d_G(s, t)$, the pair (t_1, t) is already satisfied before considering (s, t) . Then $d_{H_0}(s, t) \leq w(st_1) + d_{H_0}(t_1, t) \leq w(st_1) + [d_G(t_1, t) + (4 + \varepsilon)W(s, t)] \leq d_G(s, t) + (4 + \varepsilon)W(s, t)$, contradicting that $\pi(s, t)$ was added to H . Symmetrically, the last edge in $\pi(s, t)$ is absent in H_0 .

By Lemma 9 and the above observation, we can establish there are $\Omega(d^{1/2})$ choices for a and b . Then for every choice of x, a, b , adding $\pi(s, t)$ sets off the pairs $(a, x), (b, x)$ if not already, and improves at least one pair's distance by at least $\frac{\varepsilon}{4}W(s, t)$. By Lemma 10, this gives $\Omega(\varepsilon d^{1/2} z)$ improvements upon adding z total edge weight in $\pi(s, t)$.

If Z is the total weight added during $+(4 + \varepsilon)W(\cdot, \cdot)$ spanner completion, then

$$\Omega(\varepsilon d^{1/2} Z) = O\left(\frac{n^2}{\varepsilon}\right) \implies Z = O_\varepsilon\left(\frac{n^2}{d^{1/2}}\right).$$

By setting $d := n^{2/3}$, we obtain that the total weight of the spanner H is $O_\varepsilon(n^{5/3})$; since $w(\text{MST}(G)) = \frac{n}{2}$, this implies $\text{lightness}(H) = O_\varepsilon(n^{2/3})$ as desired. \square

5 Conclusion

We showed that many of the global $+cW$ spanners from [2] are actually local $+cW(\cdot, \cdot)$ spanners by strengthening the analysis and Lemma 1. We additionally provided several additional pairwise and all-pairs spanners based on d -light initialization; this suggests d -light initialization can be used to construct several other types of additive spanners.

We then provided the first known lightweight additive spanner constructions, whose lightness does not depend on W . These lightness results are likely not optimal; we leave the questions of finding more optimized lightness guarantees, as well as lightness guarantees for pairwise and subsetwise spanners as future open problems. Further, another natural question is whether we can construct additive spanners in weighted graphs which are simultaneously sparse *and* lightweight.

Acknowledgements The authors wish to thank Michael Elkin, Faryad Darabi Sahneh, and the anonymous reviewers for their discussion and comments.

References

1. Abboud, A., Bodwin, G.: The $4/3$ additive spanner exponent is tight. *Journal of the ACM (JACM)* **64**(4), 1–20 (2017)
2. Ahmed, R., Bodwin, G., Darabi Sahneh, F., Kobourov, S., Spence, R.: Weighted additive spanners. In: Adler, I., Müller, H. (eds.) *Graph-Theoretic Concepts in Computer Science*. pp. 401–413. Springer (2020)
3. Ahmed, R., Bodwin, G., Sahneh, F.D., Hamm, K., Jebelli, M.J.L., Kobourov, S., Spence, R.: Graph spanners: A tutorial review. *Computer Science Review* **37**, 100253 (2020)
4. Ahmed, R., Bodwin, G., Sahneh, F.D., Hamm, K., Kobourov, S., Spence, R.: Multi-level weighted additive spanners. *arXiv preprint arXiv:2102.05831* (2021)
5. Aingworth, D., Chekuri, C., Indyk, P., Motwani, R.: Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM Journal on Computing* **28**, 1167–1181 (1999)
6. Althöfer, I., Das, G., Dobkin, D., Joseph, D.: Generating sparse spanners for weighted graphs. In: *Scandinavian Workshop on Algorithm Theory (SWAT)*. pp. 26–37. Springer Berlin Heidelberg, Berlin, Heidelberg (1990)
7. Baswana, S., Kavitha, T., Mehlhorn, K., Pettie, S.: Additive spanners and (α, β) -spanners. *ACM Transactions on Algorithms (TALG)* **7**(1), 5 (2010)
8. Bodwin, G.: A note on distance-preserving graph sparsification. *arXiv preprint arXiv:2001.07741* (2020)
9. Cai, L., Keil, J.M.: Computing visibility information in an inaccurate simple polygon. *International Journal of Computational Geometry & Applications* **7**(6), 515–538 (1997)
10. Censor-Hillel, K., Kavitha, T., Paz, A., Yehudayoff, A.: Distributed construction of purely additive spanners. In: Gavoille, C., Ilcinkas, D. (eds.) *Distributed Computing*. pp. 129–142. Springer (2016)
11. Chandra, B., Das, G., Narasimhan, G., Soares, J.: New sparseness results on graph spanners. In: *Proceedings of the eighth annual symposium on Computational geometry*. pp. 192–201. ACM (1992)

12. Chechik, S.: New additive spanners. In: Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms (SODA). pp. 498–512. Society for Industrial and Applied Mathematics (2013)
13. Chechik, S., Wulff-Nilsen, C.: Near-optimal light spanners. In: Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms (SODA). pp. 883–892. Society for Industrial and Applied Mathematics (2016)
14. Cygan, M., Grandoni, F., Kavitha, T.: On pairwise spanners. In: Proceedings of 30th International Symposium on Theoretical Aspects of Computer Science (STACS 2013). vol. 20, pp. 209–220 (2013)
15. Dobson, A., Bekris, K.E.: Sparse roadmap spanners for asymptotically near-optimal motion planning. *The International Journal of Robotics Research* **33**(1), 18–47 (2014)
16. Dor, D., Halperin, S., Zwick, U.: All-pairs almost shortest paths. *SIAM Journal on Computing* **29**(5), 1740–1759 (2000)
17. Elkin, M., Gitlitz, Y., Neiman, O.: Almost shortest paths and PRAM distance oracles in weighted graphs. arXiv preprint arXiv:1907.11422 (2019)
18. Elkin, M., Gitlitz, Y., Neiman, O.: Improved weighted additive spanners. arXiv preprint arXiv:2008.09877 (2020)
19. Elkin, M., Neiman, O., Solomon, S.: Light spanners. In: International Colloquium on Automata, Languages, and Programming. pp. 442–452. Springer (2014)
20. Elkin, M., Peleg, D.: $(1 + \epsilon, \beta)$ -spanner constructions for general graphs. *SIAM Journal on Computing* **33**(3), 608–631 (Mar 2004)
21. Erdős, P.: Extremal problems in graph theory. In: Proceedings of the Symposium on Theory of Graphs and its Applications. p. 2936 (1963)
22. Filtser, A., Solomon, S.: The greedy spanner is existentially optimal. In: Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing. pp. 9–17. ACM (2016)
23. Kavitha, T.: New pairwise spanners. *Theory of Computing Systems* **61**(4), 1011–1036 (2017)
24. Kavitha, T., Varma, N.M.: Small stretch pairwise spanners and approximate d -preservers. *SIAM Journal on Discrete Mathematics* **29**(4), 2239–2254 (2015)
25. Khuller, S., Raghavachari, B., Young, N.: Balancing minimum spanning trees and shortest-path trees. *Algorithmica* **14**(4), 305–321 (1995)
26. Knudsen, M.B.T.: Additive spanners: A simple construction. In: Scandinavian Workshop on Algorithm Theory (SWAT). pp. 277–281. Springer (2014)
27. Le, H., Solomon, S.: A unified and fine-grained approach for light spanners. arXiv preprint arXiv:2008.10582 (2020)
28. Marble, J.D., Bekris, K.E.: Asymptotically near-optimal planning with probabilistic roadmap spanners. *IEEE Transactions on Robotics* **29**(2), 432–444 (2013)
29. Peleg, D., Schäffer, A.A.: Graph spanners. *Journal of Graph Theory* **13**(1), 99–116 (1989)
30. Peleg, D., Upfal, E.: A trade-off between space and efficiency for routing tables. *Journal of the ACM (JACM)* **36**(3), 510–530 (1989)
31. Pettie, S.: Low distortion spanners. *ACM Transactions on Algorithms (TALG)* **6**(1), 7 (2009)
32. Salzman, O., Shaharabani, D., Agarwal, P.K., Halperin, D.: Sparsification of motion-planning roadmaps by edge contraction. *The International Journal of Robotics Research* **33**(14), 1711–1725 (2014)
33. Woodruff, D.P.: Additive spanners in nearly quadratic time. In: International Colloquium on Automata, Languages, and Programming. pp. 463–474. Springer (2010)