

Low Distortion Spanners

Seth Pettie

The University of Michigan

Abstract. A *spanner* of an undirected unweighted graph is a subgraph that approximates the distance metric of the original graph with some specified accuracy. Specifically, we say $H \subseteq G$ is an f -spanner of G if any two vertices u, v at distance d in G are at distance at most $f(d)$ in H . There is clearly some tradeoff between the sparsity of H and the distortion function f , though the nature of this tradeoff is still poorly understood.

In this paper we present a simple, modular framework for constructing sparse spanners that is based on interchangeable components called *connection schemes*. By assembling connection schemes in different ways we can recreate the additive 2- and 6-spanners of Aingworth et al. and Baswana et al. and improve on the $(1 + \epsilon, \beta)$ -spanners of Elkin and Peleg, the sublinear additive spanners of Thorup and Zwick, and the (non constant) additive spanners of Baswana et al. Our constructions rival the simplicity of all comparable algorithms and provide substantially better spanners, in some cases reducing the density *doubly* exponentially.

1 Introduction

An f -spanner of an undirected, unweighted graph G is a subgraph H such that

$$\delta_H(u, v) \leq f(\delta_G(u, v))$$

holds for every pair of vertices u, v , where δ_H is the distance metric w.r.t. H . The premier open problem in this area is to understand the necessary tradeoffs between the sparsity of H and the distortion function f . The problem of finding a sparse spanner is one in the wider area of *metric embeddings*, where *distortion* is almost universally defined to be multiplicative, of the form $f(d) = t \cdot d$ for some $t \geq 1$. Spanners, however, can possess substantially stronger properties. The recent work of Elkin and Peleg [11] and Thorup and Zwick [22] shows that the multiplicative distortion $f(d)/d$ can tend toward 1 as d increases; in this situation the nature of the tradeoff is between the sparsity of the spanner and the *rate* of convergence. It is unknown whether this type of tradeoff is the best possible or whether there exist arbitrarily sparse *additive* spanners, where $f(d) = d + O(1)$ and the tradeoff is between sparsity and the constant hidden in the $O(1)$ term.

Applications. The original application of spanners was in the efficient simulation of synchronized protocols in unsynchronized networks [3, 17]. Thereafter spanners were used in the design of low-stretch routing schemes using small routing tables (see [9, 19, 21] and the references therein), computing almost shortest paths in distributed networks [12], and in approximation algorithms for geometric spaces [15]. A recent application of spanners is in the design of approximate distance oracles and labeling schemes for arbitrary metrics; see [23, 4] for further references. Tree spanners have found a number of uses in recent years, such as solving diagonally dominant linear systems [20] and various approximation algorithms [13]. (Tree spanners cannot have any non-trivial distortion in the worst case so weaker notions are used, such as average distortion and expected distortion over a distribution of spanning trees.) In all the applications cited above the quality of the solution is directly related to the quality of the underlying spanners.

Sparseness-Distortion Tradeoffs. It was observed early on [16, 2] that a spanner has multiplicative distortion t if and only if $f(1) = t$, that is, if the distance between adjacent vertices in G is at most t in the spanner H . Althöfer et al. [2] proved that the sparsest multiplicative t -spanner has precisely $m_{t+2}(n)$ edges, where $m_g(n)$ is the maximum number of edges in a graph with n vertices and girth at least g .¹ The upper bound follows from a trivial greedy algorithm (similar to Kruskal’s minimum spanning tree algorithm) and the lower bound is also simple. In any graph with girth $t + 2$, removing any edge shifts the distance of its endpoints from 1 to at least $t+1$. Thus, the only multiplicative t -spanner is the graph itself. It is easy to show that $m_{2k+1}(n)$ and $m_{2k+2}(n)$ are $O(n^{1+1/k})$ and it has been conjectured by Erdős and others (see [24, 23]) that this bound is asymptotically tight. However, it has only been proved for $k = 1, 2, 3$, and 5; see [24, 23] for a longer discussion on the girth conjecture. The tradeoff between sparseness and $f(1)$ is fully understood inasmuch as it amounts to proving the girth conjecture. The only other situation that is understood to a similar degree is the threshold D beyond which f is isometric, i.e., where $f(d) = d$, for all $d \geq D$. Bollobas et al. [6] showed that these so called distance preservers have $\Theta(n^2/D)$ edges. The only known lower bound for an intermediate distance was given recently by Woodruff [25], who showed that $f(k) < 3k$ holds only if the spanner has $\Omega(k^{-1}n^{1+1/k})$ edges.² For this size spanner the best upper bound on $f(k)$ is $k^{\log 3}$, which we show in Section 3.

It is perfectly consistent with the girth conjecture and Woodruff’s lower bound [25] that there are spanners with size $O(n^{1+1/k})$ and constant additive distortion $f(d) = d + 2k - 2$, though little progress has been made in proving or disproving their existence. Aingworth et al. [1] (see also [10, 11, 22]) showed that there are additive 2-spanners with size $O(n^{3/2})$, which is optimal, and Baswana et al. [5] gave an additive 6-spanner with size $O(n^{4/3})$. Below the $O(n^{4/3})$ thresh-

¹ Girth is the length of the shortest cycle.

² Woodruff expressed this result as a lower bound on the *additive* distortion.

old the best known tradeoff is quite weak; it is shown in [5] that there is an $O(n^{1+\epsilon})$ -sized spanner with $f(d) = d + O(n^{1-3\epsilon})$, for any $\epsilon \in (0, 1/3)$.

One nice property of additive spanners is that $f(d)/d$ quickly tends toward 1 as d increases. Elkin and Peleg [11] and Thorup and Zwick [22] have shown that this property can be achieved without directly addressing the problem of guaranteeing a constant additive distortion. Elkin and Peleg [11] define an (α, β) -spanner to be one with distortion $f(d) = \alpha d + \beta$. They show the existence of $(1 + \epsilon, \beta)$ -spanners with size $O(\beta n^{1+1/k})$, where β is roughly $(\epsilon^{-1} \log k)^{\log k}$. Thorup and Zwick [22] gave a remarkably simple spanner construction with similar but incomparable properties. They showed that there is a $O(kn^{1+1/k})$ -size $(1 + \epsilon, O(\lceil 1 + \frac{2}{\epsilon} \rceil^{k-2}))$ -spanner, which holds for *all* ϵ simultaneously. When ϵ^{-1} is chosen to be $\Theta(d^{1/(k-1)})$ the distortion function is $f(d) = d + O(d^{1-1/(k-1)} + 2^k)$. Notice that the β of the Thorup-Zwick spanner is exponentially larger than that of Elkin and Peleg.

Our Results. In this paper we present a simple, modular framework for constructing low distortion spanners that generalizes much of the recent work on additive and (α, β) -spanners. In our framework a spanner is expressed as a list of *connection schemes*, which are essentially interchangeable components that can be combined in various ways. This framework simplifies the construction of spanners and greatly simplifies their analysis. Once the list of connection schemes is fixed the size and distortion of the spanner follow from some straightforward linear recurrences. In our framework it is possible to succinctly express the additive 2-spanners of [1, 11, 22] and the additive 2- and 6-spanners of [5], as well as the additive 4-spanner suggested in [7]. By properly combining connection schemes we can simultaneously improve the sparseness and distortion of both the Elkin-Peleg and Thorup-Zwick spanners.

One nice feature of our framework is that it is possible to obtain *linear* size spanners with relatively good distortion. Previous to this work the only linear size spanners [2, 14] had $O(\log n)$ multiplicative distortion. (The Elkin-Peleg spanners always have $\Omega(n(\epsilon^{-1} \log \log n)^{\log \log n})$ edges. The size of the Thorup-Zwick spanners is $\Omega(n \log n)$, though at this sparsity the guaranteed distortion is quite weak.) We can construct an $O(n)$ -size $(5 + \epsilon, \beta)$ -spanner, where $\epsilon > 0$ is constant and $\beta = \text{polylog}(n)$, as well as an additive $\tilde{O}(n^{9/16})$ -spanner. Under relatively mild assumptions we can actually push the density *and* multiplicative distortion arbitrarily close to 1. For graphs with *quadratic expansion* there are $(1 + \epsilon, \beta)$ -spanners with $(1 + \epsilon)n$ edges, for any $\epsilon > 0$. By quadratic expansion we mean that the number of vertices within distance d of any vertex is $\Omega(d^2)$.

2 Notation and Overview

Throughout the paper $G = (V, E)$ denotes the input graph. We denote by $\delta_H(u, v)$ and $P_H(u, v)$ the distance from u to v in H and the associated shortest path, respectively. In general there are many shortest paths between two

vertices. We insist that if $x, y \in P_H(u, v)$ then $P_H(x, y) \subseteq P_H(u, v)$. Whenever H is omitted it is assumed to be G . Our spanner constructions all refer to vertex sets V_0, V_1, \dots, V_o , where $V_0 = V$ and V_i is derived by sampling V_{i-1} with probability q_i/q_{i-1} , where $1 = q_0 > q_1 > \dots > q_o$. Thus, the expected size of V_i is nq_i . Let $p_i(v)$ be the closest vertex in V_i to v , breaking ties arbitrarily, and let $\text{rad}_i(v) = \delta(v, p_i(v))$. If $i = o + 1$ then $p_{o+1}(v)$ is non-existent and $\text{rad}_{o+1}(v) = \infty$ by definition. Let $\text{Ball}(v, r) = \{u : \delta(v, u) < r\}$. We define $\mathcal{B}_i^\epsilon(v) = \text{Ball}(v, \epsilon \cdot \text{rad}_{i+1}(v))$ and $\mathcal{B}_i^-(v) = \text{Ball}(v, \text{rad}_{i+1}(v) - 1)$, where ϵ is taken to be 1 if omitted. Let $\bar{\mathcal{B}}_i^x(v) = \mathcal{B}_i^x(v) \cup \{p_{i+1}(v)\}$, where x is ‘-’ or some ϵ .

Connection			
Scheme	Connected Pairs	Distortion	Expected Size
A	$V_i \times \bar{\mathcal{B}}_i(\cdot)$	exact	nq_i/q_{i+1}
B	$V_i \times \bar{\mathcal{B}}_i^-(\cdot)$	$d + 2(\log d + 1)$	$n\sqrt{q_i/q_{i+1}}$
	$V_i \times \bar{\mathcal{B}}_i^{1/2}(\cdot) \cap V_i$	$d + 2$	$n\sqrt{q_i/q_{i+1}}$
	$V_o \times V_o$	$d + 2$	$n\sqrt{nq_o}$
C	$V_i \times \bar{\mathcal{B}}_i^{1/3}(\cdot) \cap V_i$	exact	$n + nq_i^2/q_{i+1}^{3/2}$
	$V_o \times V_o$	exact	$n + n^{5/2}q_o^2$
D (r)	$V_i \times \bar{\mathcal{B}}_i(\cdot) \cap \text{Ball}(\cdot, r) \cap V_i$	exact	nrq_i^2/q_{i+1}
x	$V_i \times p_{i+1}(\cdot)$	exact	n

Fig. 1. The connection schemes. Here $0 \leq i \leq o$. Schemes **B** and **C** have slightly stronger guarantees at $i = o$. The notation $V_i \times \bar{\mathcal{B}}_i(\cdot)$ is short for $\{(u, v) : u \in V_i, v \in \bar{\mathcal{B}}_i(u)\}$.

In Section 4 we describe five connection schemes called **A**, **B**, **C**, **D**, and **x**. In our framework a spanner can be expressed by choosing an *order* o and a list of the connection schemes employed at each level. For instance, in our compact notation the spanner **ABB** employs scheme **A** at level zero and scheme **B** at levels 1 and 2, where in this case $o = 2$. When a connection scheme is employed at level i it returns a subgraph that connects each $v \in V_i$ to some subset of the vertices in $\bar{\mathcal{B}}_i(v)$; the particulars depend on the scheme used. The overall properties of the spanner are determined by the sequence of connection schemes and, in general, a larger order o leads to a sparser spanner with higher distortion. Figure 1 lists the specifications for the different schemes and Figure 2 lists some of the interesting spanners that can be generated from $\{\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}, \mathbf{x}\}^*$.

The connection schemes **A**, **B**, **C**, **D** and **x** all produce subgraphs that connect certain pairs of vertices by shortest or almost shortest paths. The three features of a connection scheme we care about are the pairs of vertices to be connected, the guaranteed distortion, and the expected size of the subgraph as a function of the sampling probabilities. The properties of each of the connection schemes

Encoding	Distortion $f(d)$	Size	Notes
\mathbf{A}^2 or \mathbf{B}	$d + 2$	$O(n^{3/2})$	[1, 10, 11, 5, 22]
\mathbf{AC}	$d + 4$	$O(n^{3/2}), \Omega(n^{4/3})$	[7]
\mathbf{AB}	$d + 6$	$O(n^{4/3})$	[5]
\mathbf{A}^{o+1}	$d + O(d^{1-1/o} + 3^o)$	$O(on^{1+1/o})$	[22]
not appl.	$d + O(d^{1-1/o} + 2^o)$	$O(on^{1+1/o})$	[22]
\mathbf{AB}^2	$d + O(\sqrt{d})$	$O(n^{6/5})$	new
$\mathbf{AB}^2\mathbf{C}$	$d + O(d^{2/3})$	$O(n^{25/22})$	new
$\mathbf{AB}^2\mathbf{C}^{o-2}$	$d + O(od^{1-1/o})$	$O(on^{1+\frac{(3/4)^{o-2}}{7-2(3/4)^{o-2}}})$	new , $d > o^o$

Very Sparse Spanners			
not appl.	$O(d \log n)$	$O(n)$	[2, 14]
$\mathbf{AD}^{\log \log n}$	$(5 + \epsilon)d$	$O(n)$	new , $d > \epsilon^{-\log \log n}$
\mathbf{xCC}	$d + \tilde{O}(n^{9/16})$	$O(n)$	new
$\mathbf{xD}^{\log \log n}$	$(1 + \epsilon)d + \beta''$	$(1 + \epsilon)n$	new , see below
not appl.	$(1 + \epsilon)d + \beta$	$O(n\beta)$	[11]
$\mathbf{AC}^{O(\log \log \epsilon^{-1})}$	$(1 + \epsilon)d + \beta'$	$O(n \log \log \epsilon^{-1})$	new , $\epsilon < 1/\log \log n$
$\mathbf{D}^{\log \log n}$	$(1 + \epsilon)d + \beta'$	$O(n \log \log \epsilon^{-1})$	new , $\epsilon < 1/\log \log n$
$\mathbf{AC}^{O(\log \log n)}$	$d + \tilde{O}(d^{1-\frac{1}{O(\log \log n)}})$	$O(n \log \log n)$	new , $d > \text{polylog}(n)$

Fig. 2. Some of the spanners generated by $\{\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}(\cdot), \mathbf{x}\}^*$. Here β, β' , and β'' are all $O(\epsilon^{-1} \log \log n)^{\log \log n}$.

are given in Figure 1. (Notice that some of the connection schemes have slightly stronger properties when used at the highest level o .) Let us decipher a few of the lines in Figure 1. When \mathbf{A} is used at level i it returns a subgraph H_i such that for $v \in V_i$ and $u \in \tilde{\mathcal{B}}_i(v)$, $\delta_{H_i}(v, u) = \delta(v, u)$, and furthermore, the expected size of H_i is on the order of nq_i/q_{i+1} . (The notation $V_i \times \tilde{\mathcal{B}}_i(\cdot)$ is short for the set of pairs $\{(v, u) : v \in V_i, u \in \tilde{\mathcal{B}}_i(v)\}$.) Like \mathbf{A} , schemes \mathbf{C} and \mathbf{D} have no distortion but connect fewer pairs of vertices. For $v \in V_i$, scheme \mathbf{C} only connects the pair (v, u) if u is in *both* V_i and $\tilde{\mathcal{B}}_i^{1/3}(v)$. Scheme $\mathbf{D}(r)$ requires u to be in V_i , $\tilde{\mathcal{B}}_i(v)$, and $\text{Ball}(v, r)$, where r is a given parameter that influences the size of the subgraph. Scheme \mathbf{B} guarantees two grades of distortion. If u is in both $\tilde{\mathcal{B}}_i^{1/2}(v)$ and V_i the additive distortion is 2 and if u is in $\tilde{\mathcal{B}}_i^-(v)$ the additive distortion is $2(\log d + 1)$, where $d = \delta(v, u)$. Scheme \mathbf{x} simply connects every $v \in V_i$ to the nearest vertex $p_{i+1}(v) \in V_{i+1}$.

In Section 3 we show how connection schemes can be composed in various ways to yield spanners with different sparseness-distortion tradeoffs. The construction and analysis of these spanners is inspired by the *distance emulators* of Thorup and Zwick [22]. In Section 4 we present the algorithms behind schemes \mathbf{A} and \mathbf{C} . See [18] for a description of the other schemes.

3 Modular Spanner Construction

In our framework a spanner is expressed as a finite sequence of connection schemes. For instance, the spanner **ABB** consists of the edge sets H_0, H_1 , and H_2 , where H_0 is the subgraph returned by the connection scheme **A** applied to the zeroth level, and H_1 and H_2 are the subgraphs returned by applying **B** to levels 1 and 2. The size of the spanner depends solely on the sampling probabilities and there is typically one optimal choice of probabilities. For instance, for the spanner **ABB** the expected size is asymptotically $n/q_1 + n\sqrt{q_1/q_2} + n\sqrt{nq_2}$, which is optimized at $q_1 = n^{-1/5}, q_2 = n^{-3/5}$. The distortion of the spanner is analyzed by solving some linear recurrences. The derivation of these recurrences is sketched below and formally proved in Lemma 1. Lemma 2 solves these recurrences for the class of spanners that use schemes **A, B**, and **C**. These three schemes are sufficient to prove our strongest sparseness-distortion trade-offs (Theorem 1) but they are ultimately incapable of generating spanners with $o(n \log \log n)$ edges. See the full version [18] for constructions of sparser spanners.

Suppose we have a spanner H defined by a finite string $\tau \in \{\mathbf{A}, \mathbf{B}, \mathbf{C}\}^{o+1}$. Let $\tau(j)$ be the j th character of τ , for some $0 \leq j \leq o$, and let $H_j \subseteq H$ be the subgraph returned by the connection scheme $\tau(j)$ at level j . Let v and v' be two vertices at distance at most Δ^j in the original graph, where $\Delta \geq 2$ is an integer. To get from v to v' in H we divide up $P(v, v')$ into segments of length Δ^{j-1} ; see Figure 3. Let v_ℓ be the first vertex in the ℓ th segment. Using only the subgraph $H_0 \cup \dots \cup H_{j-1}$ we try to take short hops from $v = v_0$ to v_1 , from v_1 to v_2 , and so on. A *short* hop is one whose length is at most S_Δ^{j-1} . If a short hop exists we call the associated segment *successful*. We will see shortly that S_Δ^{j-1} is usually $\Delta^{j-1} + O(j\Delta^{j-2})$, which means that $S_\Delta^{j-1}/\Delta^{j-1}$ tends to 1 as Δ (and $\delta(v, v')$) increases. Not all segments will be successful. If we encounter a *failed* segment $v_\ell \dots v_{\ell+1}$ we require that $\delta_H(v_\ell, p_j(v_\ell)) = \delta(v_\ell, p_j(v_\ell)) \leq F_\Delta^{j-1}$, where F_Δ^{j-1} is usually $O(\Delta^{j-1})$. Suppose at least one segment fails and let z and z' be, respectively, the first vertex of the first failed segment and the last vertex of the last failed segment. Let s and s' be the number of segments between v and z , and between v' and z' , respectively. By the definition of failure $\delta(z, p_j(z))$ and $\delta(z', p_j(z'))$ are at most F_Δ^{j-1} and by the triangle inequality $\delta(p_j(z), p_j(z')) \leq 2F_\Delta^{j-1} + (\Delta - s - s')\Delta^{j-1}$. If we could get from $p_j(z)$ to $p_j(z')$ by a shortest (or almost shortest) path in H_j then we would declare the whole path $v \dots v'$ a success. Whether this is possible depends on whether $p_j(z')$ lies within $\mathcal{B}_j(p_j(z))$. If $\tau(j) = \mathbf{B}$ we actually require that $p_j(z')$ be within $\mathcal{B}_j^{1/2}(p_j(z))$ and if $\tau(j) = \mathbf{C}$ we would require that $p_j(z') \in \mathcal{B}_j^{1/3}(p_j(z))$. In any case, if there is *not* a short path from $p_j(z)$ to $p_j(z')$ we have an upper bound on $\text{rad}_{i+1}(p_j(z)) = O(\delta(p_j(z), p_j(z')))$, and, therefore, an upper bound on $\text{rad}_{i+1}(v)$ as well. In this way we can bound S_Δ^j and F_Δ^j in terms of S_Δ^{j-1} , F_Δ^{j-1} , and $\tau(j)$. Formally, we define S and F with respect to some spanner as follows.

Definition 1. Let H be a spanner defined by some finite string $\tau \in \{\mathbf{A}, \mathbf{B}, \mathbf{C}\}^*$. We define S_Δ^j and F_Δ^j to be minimal such that for any two vertices v, v' at distance




Fig. 3. The vertices v and v' are at distance at most Δ^j . The vertices z and z' are, respectively, the first on the first failed segment and last on the last failed segment.

at most Δ^j , where $0 \leq j \leq o$, at least one of the following inequalities holds:

$$\delta_H(v, v') \leq S_{\Delta}^j \quad \text{or} \quad \delta_H(v, V_{j+1}) \leq F_{\Delta}^j$$

Notice that, despite the terminology, a path may both succeed and fail. Any two vertices at distance at most Δ^o must be connected in H by a path of length at most S_{Δ}^o , that is, every such path must be a success. Such a path cannot fail because V_{o+1} does not exist, and, therefore $\delta_H(v, V_{o+1})$ is undefined. This simply reflects the fact that in our connection schemes, all vertices in V_o are connected by (nearly) shortest paths H .

Lemma 1 shows that S and F are bounded by some straightforward recurrences. It only considers spanners that employ scheme \mathbf{A} at the zeroth level, which is generally the wisest choice.

Lemma 1. Consider a spanner defined by $\tau = \mathbf{A} \cdot \{\mathbf{A}, \mathbf{B}, \mathbf{C}\}^o$. Then $S_{\Delta}^o = F_{\Delta}^o = 1$ holds for all Δ and:

$$F_{\Delta}^j \leq \begin{cases} 3F_{\Delta}^{j-1} + \Delta^j & \text{for } \tau(j) = \mathbf{A} \\ 5F_{\Delta}^{j-1} + 2\Delta^j & \text{for } \tau(j) = \mathbf{B} \\ 7F_{\Delta}^{j-1} + 3\Delta^j & \text{for } \tau(j) = \mathbf{C} \end{cases}$$

$$S_{\Delta}^j \leq \max \text{ of } \Delta S_{\Delta}^{j-1} \text{ and}$$

$$\begin{cases} (\Delta - 1)S_{\Delta}^{j-1} + 4F_{\Delta}^{j-1} + \Delta^{j-1} & \text{for } \tau(j) \in \{\mathbf{A}, \mathbf{C}\} \\ (\Delta - 1)S_{\Delta}^{j-1} + 4F_{\Delta}^{j-1} + \Delta^{j-1} + 2 & \text{for } \tau(j) = \mathbf{B} \end{cases}$$

Proof. For the base case, consider any adjacent v, v' in G . If the edge (v, v') is in H_0 (returned by \mathbf{A} at level 0) then $\delta(v, v') = 1 = S_{\Delta}^0$. If not then, by the definition of \mathbf{A} , $v' \notin \mathcal{B}_0(v)$ and $\delta_{H_0}(v, p_1(v)) = 1 = F_{\Delta}^0$. Let v, v', z, z', s , and s' be as in the above discussion; see Figure 3. If the spanner does not contain a short

path from $p_j(z)$ to $p_j(z')$ (failure) then we can conclude that $p_j(z') \notin \mathcal{B}_j(p_j(z))$ if $\tau(j) = \mathbf{A}$, that $p_j(z') \notin \mathcal{B}_j^{1/2}(p_j(z))$ if $\tau(j) = \mathbf{B}$, and that $p_j(z') \notin \mathcal{B}_j^{1/3}(p_j(z))$ if $\tau(j) = \mathbf{C}$. It follows that:

$$\delta(p_j(z), p_{j+1}(p_j(z))) \leq \begin{cases} 2\mathbf{F}_\Delta^{j-1} + (\Delta - s - s')\Delta^{j-1} & \text{if } \tau(j) = \mathbf{A} \\ 2(2\mathbf{F}_\Delta^{j-1} + (\Delta - s - s')\Delta^{j-1}) & \text{if } \tau(j) = \mathbf{B} \\ 3(2\mathbf{F}_\Delta^{j-1} + (\Delta - s - s')\Delta^{j-1}) & \text{if } \tau(j) = \mathbf{C} \end{cases}$$

The distance from v to V_{j+1} is at most $\delta(v, p_{j+1}(p_j(z)))$, which we bound as:

$$\begin{aligned} \delta(v, p_{j+1}(v)) &\leq \delta(v, z) + \delta(z, p_j(z)) + \delta(p_j(z), p_{j+1}(p_j(z))) \\ &\leq s\Delta^{j-1} + \mathbf{F}_\Delta^{j-1} + t(2\mathbf{F}_\Delta^{j-1} + (\Delta - s - s')\Delta^{j-1}) \\ &\quad \{t = 1, 2, 3 \text{ depending on } \tau(j)\} \\ &\leq (s + t(\Delta - s - s'))\Delta^{j-1} + (2t + 1)\mathbf{F}_\Delta^{j-1} \\ &\leq (2t + 1)\mathbf{F}_\Delta^{j-1} + t\Delta^j \quad \{\text{worst case is } s = s' = 0\} \end{aligned}$$

We obtain the claimed bounds on \mathbf{F}_Δ^j by setting $t = 1, 2$, and 3 when $\tau(j)$ is, respectively, \mathbf{A} , \mathbf{B} , and \mathbf{C} . This covers the case when the path $v \dots v'$ is a failure. One way for it to be a success is if each of the Δ segments is a success, that is, if z and z' do not exist. In general there will be some failed segments and we can only declare the path successful if there is a *short* path from $p_j(z)$ to $p_j(z')$. We demand a shortest path if $\tau(j) \in \{\mathbf{A}, \mathbf{C}\}$ and tolerate an additive error of 2 if $\tau(j) = \mathbf{B}$. We can now bound \mathbf{S}_Δ^j as follows:

$$\begin{aligned} \delta_H(v, v') &\leq \max\{\Delta\mathbf{S}_\Delta^{j-1}, \delta_H(v, z) + \delta_H(z, p_j(z)) + \delta_H(p_j(z), p_j(z')) \\ &\quad + \delta_H(p_j(z'), z') + \delta_H(z', v')\} \\ &\leq \max\{\Delta\mathbf{S}_\Delta^{j-1}, (s + s')\mathbf{S}_\Delta^{j-1} + 4\mathbf{F}_\Delta^{j-1} + (\Delta - s - s')\Delta^{j-1} [+2]\} \\ &\leq \max\{\Delta\mathbf{S}_\Delta^{j-1}, (\Delta - 1)\mathbf{S}_\Delta^{j-1} + 4\mathbf{F}_\Delta^{j-1} + \Delta^{j-1} [+2]\} \end{aligned}$$

where the “[+2]” is only present if $\tau(j) = \mathbf{B}$.

Lemma 2 solves these recurrences for spanners that use schemes \mathbf{A} , \mathbf{B} , & \mathbf{C} .

Lemma 2. (ABC Spanners) *Consider any spanner H defined by $\tau \in \mathbf{A} \cdot \{\mathbf{A}, \mathbf{B}, \mathbf{C}\}^o$. If $\Delta \geq 8$ and $c = 3\Delta/(\Delta - 7)$ then:*

$$\mathbf{F}_\Delta^j \leq c\Delta^j \quad \mathbf{S}_\Delta^j \leq \begin{cases} \Delta^j + 4cj\Delta^{j-1} & \text{for } j \leq \Delta \\ (4c + 1)\Delta^j & \text{for } j \geq \Delta \end{cases}$$

Furthermore, $\mathbf{F}_\Delta^o = 0$, that is, if $\delta(u, v) \leq \Delta^o$ then $\delta_H(u, v) \leq \mathbf{S}_\Delta^o$.

Proof. Taking the worst cases from Lemma 1 we have $\mathbf{F}_\Delta^j \leq 7\mathbf{F}_\Delta^{j-1} + 3\Delta^j$ and $\mathbf{S}_\Delta^j \leq \max\{\Delta\mathbf{S}_\Delta^{j-1}, (\Delta - 1)\mathbf{S}_\Delta^{j-1} + 4\mathbf{F}_\Delta^{j-1} + \Delta^{j-1} + 2\}$. One can easily verify by induction that $\mathbf{F}_\Delta^j \leq c\Delta^j$. To bound \mathbf{S}_Δ^j assume inductively that it is at most $\Delta^j + 4cj\Delta^{j-1} - 1$, and for $j \geq \Delta$, that it is at most $(4c + 1)\Delta^j - 1$; these

inequalities clearly hold for $j = 1$. First consider the case $j \leq \Delta$:

$$\begin{aligned}
S_\Delta^j &\leq \max\{\Delta S_\Delta^{j-1}, (\Delta - 1)S_\Delta^{j-1} + 4F_\Delta^{j-1} + \Delta^{j-1} + 2\} \\
&\leq \max\{\Delta^j + 4c(j-1)\Delta^{j-1} - \Delta, \\
&\quad (\Delta - 1)(\Delta^{j-1} + 4c(j-1)\Delta^{j-2} - 1) + 4c\Delta^{j-1} + \Delta^{j-1} + 2\} \\
&\leq \max\{\Delta^j + 4cj\Delta^{j-1} - 1, \\
&\quad \Delta^j + 4c(j-1)\Delta^{j-1} + 4c\Delta^{j-1} - (4c(j-1)\Delta^{j-2} + \Delta + 1)\} \\
&\leq \Delta^j + 4cj\Delta^{j-1} - 1
\end{aligned}$$

Notice that for $j = \Delta$ this bound is precisely $(4c+1)\Delta^j - 1$, which serves as our base case for the bounds on S_Δ^j for $j > \Delta$:

$$\begin{aligned}
S_\Delta^j &\leq \max\{\Delta S_\Delta^{j-1}, (\Delta - 1)S_\Delta^{j-1} + 4F_\Delta^{j-1} + \Delta^{j-1} + 2\} \\
&\leq \max\{(4c+1)\Delta^j - \Delta, (4c+1)(\Delta^j - \Delta^{j-1}) + 4c\Delta^{j-1} + \Delta^{j-1} - \Delta + 3\} \\
&\leq (4c+1)\Delta^j - 1
\end{aligned}$$

Lemma 2 states that in any spanner generated by some string in $\mathbf{A} \cdot \{\mathbf{A}, \mathbf{B}, \mathbf{C}\}^o$, the distortion is given by the function $f(d) = d + O(od^{1-1/o})$, provided that d is at least 8^o . As we will see in Theorem 1, o can be as large as $\log_{4/3} \log n$ which means that these spanners have weak guarantees for $d < 8^{\log_{4/3} \log n} < (\log n)^{7.23}$. See the full version [18] for spanners that better approximate polylogarithmic distances.

Theorem 1 illustrates some nice sparseness-distortion tradeoffs for spanners composed of schemes \mathbf{A} , \mathbf{B} , and \mathbf{C} . It only considers those generated by sequences \mathbf{ABBC}^{o-2} , which turns out to optimize sparseness without significantly affecting the distortion. (In other words, \mathbf{ABBBB} would be denser than \mathbf{ABBCC} and could only improve lower order terms in the distortion.)

Theorem 1. *The spanner generated by \mathbf{ABB} has $O(n^{6/5})$ edges and distortion function $f(d) = d + O(\sqrt{d})$. The spanner generated by \mathbf{ABBC}^{o-2} has $O(on^{1+\nu})$ edges, where $\nu = (\frac{3}{4})^{o-2} / (7 - 2(\frac{3}{4})^{o-2})$, and distortion $d + O(od^{1-1/o} + 8^o)$.*

Proof. Let H be the spanner defined by \mathbf{ABB} . H has on the order of $n/q_1 + n\sqrt{q_1/q_2} + n\sqrt{nq_2}$ edges, which is $O(n^{6/5})$ for $q_1 = n^{-1/5}$ and $q_2 = n^{-3/5}$. By Lemma 2, if $\delta(v, v') \leq \Delta^2$ then $\delta_H(v, v') \leq S_\Delta^2 = \Delta^2 + O(\Delta)$. (Recall that such a path cannot fail because every pair of vertices in V_2 is connected by a nearly shortest path.) In the general case let H be generated by $\mathbf{AB}^2\mathbf{C}^{o-2}$, for some $o \geq 3$. If v and v' are at distance at most $\Delta^o \geq 8^o$ then by Lemma 2 $\delta(v, v') \leq \min\{\Delta^o + O(o\Delta^{o-1}), O(\Delta^o)\}$. Thus, for any distance d (possibly less than 8^o) the distortion is $f(d) = d + O(od^{1-1/o} + 8^o)$. We now choose the sampling probabilities so as to optimize the size of H . They will be selected so that each of the levels zero through o contributes about the same number of edges, say $n^{1+\nu}$. Since the first three levels contribute $n/q_1 + n\sqrt{q_1/q_2} + n\sqrt{q_2/q_3}$ edges (scheme \mathbf{A} at level 0, \mathbf{B} at 1 and 2), it follows that $q_1 = n^{-\nu}$, $q_2 = n^{-3\nu}$,

and $q_3 = n^{-5\nu}$. Starting from the other end, level o (scheme **C**) contributes $n + n^{2.5}q_o^2$ implying $q_o = n^{-3/4+\nu/2}$. For $3 \leq j < o$, level j contributes on the order of $n + nq_j^2/q_{j+1}^{3/2}$ edges, implying $q_j = q_{j+1}^{3/4}n^{\nu/2}$. Assuming inductively that $q_{j+1} = n^{-(\frac{3}{4})^{o-j} + \nu(2 - \frac{3}{2}(\frac{3}{4})^{o-(j+1)})}$ (which holds for the base case $j+1 = o$), we have, for $3 \leq j < o$:

$$q_j = q_{j+1}^{3/4}n^{\nu/2} = n^{-(\frac{3}{4})^{o-j+1} + \frac{3}{4}\nu(2 - \frac{3}{2}(\frac{3}{4})^{o-(j+1)}) + \nu/2} = n^{-(\frac{3}{4})^{o-(j-1)} + \nu(2 - \frac{3}{2}(\frac{3}{4})^{o-j})}$$

The only sampling probability under two constraints is q_3 , which means that ν should be selected to satisfy $n^{-5\nu} = n^{-(\frac{3}{4})^{o-2} + \nu(2 - \frac{3}{2}(\frac{3}{4})^{o-3})}$. This equality holds for $\nu = (\frac{3}{4})^{o-2} / (7 - \frac{3}{2}(\frac{3}{4})^{o-3})$. The size of H is, therefore, on the order of $on^{1+\nu}$.

Let us briefly compare the size bounds obtained above to the spanners of Thorup and Zwick [22]. For distortions $d + O(\sqrt{d})$, $d + O(d^{2/3})$, and $d + O(d^{3/4})$ the spanners of Theorem 1 have sizes on the order of $n^{6/5}$, $n^{25/22}$, and $n^{103/94}$ in contrast to $n^{4/3}$, $n^{5/4}$, and $n^{6/5}$ obtained in [22]. The separation in density becomes sharper as o increases. For $o = \log_{4/3} \log n$, the size and distortion of our spanners is $O(n \log \log n)$ and $d + O(od^{1-1/o} + o^o)$, in contrast to [22], where the size and distortion are $O(on^{1+1/o})$ and $d + O(d^{1-1/o} + 2^o)$. In this case Theorem 1 gives a doubly exponential improvement in density. In some ways Theorem 1 is our strongest result. However, when the order o is large (close to $\log_{4/3} \log n$) and the distance being approximated very short, the spanners of Theorem 1 cannot guarantee good distortion. See the full version [18] for constructions that address these shortcomings.

4 The Connection Schemes

We only analyze schemes **A** and **C**. See [18] for a description of **B**, **D**, and **x**.

Connection scheme A. The subgraph returned by **A** at level i is, by definition, $\bigcup_{v \in V_i, u \in \bar{B}_i(v)} P(v, u)$, that is, a breadth first search tree from every $v \in V_i$ containing $p_{i+1}(v)$ and all vertices u closer to v than $p_{i+1}(v)$. The expected size of this subgraph is at most $\sum_{v \in V} \Pr[v \in V_i] \cdot \mathbb{E}[|\bar{B}_i(v)| - 1] \leq nq_i/q_{i+1}$.

Connection scheme C. To analyze scheme **C** we appeal to a lemma of Coppersmith and Elkin [8]. Let \mathcal{Q} be a set of shortest paths. We say that v is a *branching point* for two paths $P, P' \in \mathcal{Q}$ if P and P' intersect and v is an endpoint on the path $P \cap P'$. Notice that if P and P' have just one vertex in common it would be the unique endpoint on the edgeless path $P \cap P'$. Let $\text{br}(v)$ be the number of pairs $P, P' \in \mathcal{Q}$ for which v is a branching point, and let $\text{br}(\mathcal{Q}) = \sum_{v \in V} \text{br}(v)$.

Theorem 2. (*Coppersmith and Elkin*) *Let \mathcal{Q} be a set of shortest paths and $G(\mathcal{Q}) = \bigcup_{P \in \mathcal{Q}} P$. Then $|G(\mathcal{Q})| \leq n + O(\sqrt{n \text{br}(\mathcal{Q})})$.*

Proof. Let $\deg(v)$ be the degree of v in $G(\mathcal{Q})$. Notice that $\text{br}(v) \geq \binom{\deg(v)}{2}$. There must be at least $\lceil \deg(v)/2 \rceil$ paths in \mathcal{Q} that intersect v , no two of which use the same edges incident to v . Each pair of these paths contributes to $\text{br}(v)$. We can calculate $|G(\mathcal{Q})|$ as: $\frac{1}{2} \sum_v \deg(v) = n + \sum_{v: \deg(v) \geq 3} O(\sqrt{\text{br}(v)}) = n + O(\sqrt{n \text{br}(\mathcal{Q})})$. The last equality follows from the concavity of square root.

Theorem 3. *Let $\mathcal{Q} = \{P(v, u) : v \in V_i, u \in \bar{\mathcal{B}}_i^{1/3}(v)\}$. Then $\mathbb{E}[|G(\mathcal{Q})|] = n + O(nq_i^2/q_{i+1}^{3/2})$. If $\mathcal{Q}' = \{P(v, u) : (v, u) \in V_o \times V_o\}$ then $\mathbb{E}[|G(\mathcal{Q}')|] = n + O(n^{2.5}q_o^2)$.*

Proof. Let $v, w, v', w' \in V_i$, where $v' \in \bar{\mathcal{B}}_i^{1/3}(v)$ and $w' \in \bar{\mathcal{B}}_i^{1/3}(w)$. We first argue that if $P(v, v')$ and $P(w, w')$ intersect then $w, w' \in \mathcal{B}_i(v)$. For any vertex w , $\text{rad}_{i+1}(w) \leq \delta(w, v) + \text{rad}_{i+1}(v)$. Thus, if w lies outside $\mathcal{B}_i(v)$ then $\bar{\mathcal{B}}_i^{1/3}(w) \cap \bar{\mathcal{B}}_i^{1/3}(v)$ must be empty. Let v_a be the a th farthest vertex from $v = v_1$, breaking ties arbitrarily.

$$\begin{aligned} \mathbb{E}[\text{br}(\mathcal{Q})] &\leq 2 \sum_{v \in V, 1 < a < b < c} \Pr[\{v, v_a, v_b, v_c\} \subseteq V_i \wedge \{v_a, v_b, v_c\} \subseteq \mathcal{B}_i(v)] \\ &\leq \sum_{v \in V, c \geq 4} \Pr[|\mathcal{B}_i(v)| \geq c] \cdot c^2 \cdot q_i^4 \leq \sum_{v, c} (1 - q_{i+1})^c \cdot c^2 \cdot q_i^4 = O(nq_i^4/q_{i+1}^3) \end{aligned}$$

The second line follows since $v_c \in \mathcal{B}_i(v)$ if and only if $|\mathcal{B}_i(v)| \geq c$, and once v_c and $v = v_1$ are chosen there are $\binom{c-2}{2}$ ways to choose v_a and v_b . The last line follows since $(1 - q_{i+1})^c$ is bounded by a constant for $c < 1/q_{i+1}$ and geometrically decaying thereafter. Thus, $\mathbb{E}[|G(\mathcal{Q})|] = n + O(\sqrt{n \text{br}(\mathcal{Q})}) = n + O(nq_i^2/q_{i+1}^{3/2})$. Similarly, $\text{br}(\mathcal{Q}')$ is sharply concentrated around its mean—at most $(q_{i+1}n)^4$ —and $\mathbb{E}[|G(\mathcal{Q}')|] = n + O(\mathbb{E}[\sqrt{n \text{br}(\mathcal{Q}')}] = n + O(n^{2.5}q_o^2)$.

5 Conclusion

In this paper we have shown that nearly all the recent work on additive and low distortion spanners can be seen as merely instantiating a generic modular algorithm. The contribution of this work is not only a simpler way to look at spanners. On purely quantitative terms our constructions provide substantially better distortion than [11, 22] at any desired level of sparsity. Our constructions can also produce spanners with a linear number of edges. The last construction to achieve linearity was, quite surprisingly, Althöfer et al.'s [2] simple greedy algorithm.

Although the specific tradeoffs of our spanners could certainly be improved, the framework of this paper seems inherently incapable of generating arbitrarily sparse purely additive spanners. It is unclear whether a fundamentally new technique is required to find additive spanners or whether the path-buying algorithm of Baswana et al. [5] could be generalized for this purpose. In any case, proving or disproving the existence of additive spanners remains the chief open problem in this area.

Acknowledgment. Thanks to Michael Elkin, Mikkel Thorup, and Uri Zwick for inspiring much of this work.

References

1. D. Aingworth, C. Chekuri, P. Indyk, and R. Motwani. Fast estimation of diameter and shortest paths. *SIAM J. Comput.*, 28(4):1167–1181, 1999.
2. I. Althöfer, G. Das, D. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete and Computational Geometry*, 9:81–100, 1993.
3. B. Awerbuch. Complexity of network synchronization. *J. ACM*, 32:804–823, 1985.
4. S. Baswana and T. Kavitha. Faster algorithms for approximate distance oracles and all-pairs small stretch paths. In *FOCS 2006*.
5. S. Baswana, T. Kavitha, K. Mehlhorn, and S. Pettie. New constructions of (α, β) -spanners and purely additive spanners. In *SODA 2005*.
6. B. Bollobás, D. Coppersmith, and M. Elkin. Sparse subgraphs that preserve long distances and additive spanners. *SIAM J. Discr. Math.*, 9(4):1029–1055, 2006.
7. D. Coppersmith and M. Elkin. Sparse source-wise and pair-wise distance preservers. In *SODA 2005*.
8. D. Coppersmith and M. Elkin. Sparse source-wise and pair-wise preservers. *SIAM J. Discrete Math.*, to appear.
9. L. J. Cowen and C. G. Wagner. Compact roundtrip routing in directed networks. *J. Algor.*, 50(1):79–95, 2004.
10. D. Dor, S. Halperin, and U. Zwick. All-pairs almost shortest paths. *SIAM J. Comput.*, 29(5):1740–1759, 2000.
11. M. Elkin and D. Peleg. $(1 + \epsilon, \beta)$ -spanner constructions for general graphs. *SIAM J. Comput.*, 33(3):608–631, 2004.
12. M. Elkin and J. Zhang. Efficient algorithms for constructing $(1 + \epsilon, \beta)$ -spanners in the distributed and streaming models. In *PODC 2004*.
13. J. Fakcharoenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Syst. Sci.*, 69(3):485–497, 2004.
14. S. Halperin and U. Zwick. Unpublished result, 1996.
15. G. Narasimhan and M. Smid. *Geometric Spanner Networks*. 2007.
16. D. Peleg and A. A. Schaffer. Graph spanners. *J. Graph Theory*, 13:99–116, 1989.
17. D. Peleg and J. D. Ullman. An optimal synchronizer for the hypercube. *SIAM J. Comput.*, 18:740–747, 1989.
18. S. Pettie. Low distortion spanners. See <http://www.eecs.umich.edu/~pettie>.
19. L. Roditty, M. Thorup, and U. Zwick. Roundtrip spanners and roundtrip routing in directed graphs. In *SODA 2002*.
20. D. A. Spielman and S.-H. Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *STOC 2004*.
21. M. Thorup and U. Zwick. Compact routing schemes. In *SPAA 2001*.
22. M. Thorup and U. Zwick. Spanners and emulators with sublinear distance errors. In *SODA 2006*.
23. M. Thorup and U. Zwick. Approximate distance oracles. *J. ACM*, 52:1–24, 2005.
24. R. Wenger. Extremal graphs with no C^4 's, C^6 's, or C^{10} 's. *J. Combin. Theory Ser. B*, 52(1):113–116, 1991.
25. D. Woodruff. Lower bounds for additive spanners, emulators, and more. In *FOCS 2006*.